

CUBRID 2008 R4.1 Manual

Table of Contents

Introduction to Manual	1
Introduction to CUBRID	3
System Architecture	4
System Architecture	4
Database Volume Structure	5
Database Server	7
Broker	7
Interface Module	9
CUBRID Characteristics	10
Getting Started with CUBRID	13
Installing and Running	14
Installing and Running on Linux	14
Installing and Running on Windows	16
Configuring Environment Variable and Starting CUBRID	18
Configuring the Environment Variable	18
Language Setting	19
Starting the CUBRID Service	20
CSQL Interpreter	22
Starting the CSQL Interpreter	22
Executing the SQL with CSQL	23
Programming with JDBC	24
Setting up the JDBC Environment	24
JDBC Sample	25
Programming with PHP	29
Installing the PHP Module	29
PHP Sample	30
Programming with ODBC and ASP	32
Configuring the Environment of ODBC and ASP	32
ASP Sample	34
Programming with CCI	37
CCI Library	37
CCI Installation and Configuration	37

Printed Documentation

Using CCI _____	38
CCI Sample _____	39

CSQL Interpreter _____ 43

Introduction to the CSQL Interpreter _____	44
Executing CSQL _____	45
CSQL Execution Mode _____	45
Using CSQL (Syntax) _____	45
CSQL Startup Options _____	46
Session Commands _____	49

CUBRID SQL Guide _____ 57

Glossary _____	58
Comment _____	59
Identifier _____	60
Reserved Words _____	62
Data Types _____	66
Numeric Types _____	66
Date/Time Types _____	70
Bit Strings _____	77
Character Strings _____	78
BLOB/CLOB Data Types _____	84
Collection Types _____	90
Implicit Type Conversion _____	92
Table Definition _____	98
CREATE TABLE _____	98
ALTER TABLE _____	109
DROP TABLE _____	120
RENAME TABLE _____	121
Index Definition _____	122
CREATE INDEX _____	122
ALTER INDEX _____	122
DROP INDEX _____	123
VIEW Definition _____	124
CREATE VIEW _____	124
ALTER VIEW _____	126
DROP VIEW _____	128

RENAME VIEW	128
SERIAL	129
CREATE SERIAL	129
ALTER SERIAL	130
DROP SERIAL	131
Use SERIAL	131
Serial Function	132
Operators and Functions	134
Logical Operators	134
Comparison Operators	134
Arithmetic Operators	136
Set Operators	140
Containment Operators	144
BIT Functions and Operators	149
String Functions and Operators	151
Numeric and Operator Functions	175
Date/Time Functions and Operators	188
Data Type Conversion Functions and Operators	211
Aggregate Functions	223
Click Counter Functions	229
ROWNUM Functions	231
Information Functions	233
Encryption Function	238
Conditional Expressions and Functions	239
Conditional Expressions	245
Data Manipulation	253
SELECT	253
Outer Join	260
Subquery	263
Hierarchical Query	264
INSERT	271
UPDATE	274
REPLACE	275
DELETE	276
TRUNCATE	277
DO	277
PREPARED STATEMENT	278
SET	279
SHOW	281

Transaction and Lock	287
Overview	287
Database Transaction	287
Database Concurrency	290
Lock Protocol	291
Transaction Isolation Level	299
Transaction Termination and Restoration	311
Database User Authorization	313
Database User	313
Managing User	313
Granting Authorization	314
Revoking Authorization	315
User Authorization Management METHOD	315
Query Optimization	318
Updating Statistics	318
Checking Statistics Information	318
Using SQL Hint	319
Viewing Query Plan	320
Using Indexes	321
TRIGGER	329
CREATE TRIGGER	329
ALTER TRIGGER	334
DROP TRIGGER	335
RENAME TRIGGER	335
Deferred Condition and Action	335
Trigger on REPLACE and INSERT ... ON DUPLICATE KEY UPDATE	336
TRIGGER Debugging	337
TRIGGER Example	339
Java Stored Function/Procedure	341
Overview	341
Environment Configuration for Java Stored Function/Procedure	341
How to Write Java Stored Function/Procedure	342
Using Server-side Internal JDBC Driver	343
Connecting to Other Database	344
loadjava Utility	345
Loaded Java Class Publish	345
Java Stored Function/Procedure Call	347
Caution	349
METHOD	352

Overview	352
METHOD Type	352
Calling METHOD	352
Partitioning	354
What is Partitioning?	354
Range Partitioning	354
Hash Partitioning	357
List Partitioning	358
Partitioning Management	360
Class Inheritance	365
Overview	365
Class Attribute and Method	365
Order Rule for Inheritance	366
INHERIT Clause	366
ADD SUPERCLASS Clause	366
DROP SUPERCLASS Clause	367
Class Conflict Resolution	368
Overview	368
Resolution Specifier	368
Superclass Conflict	368
Subclass Conflict	369
Schema Invariant	370
Rule for Schema Changes	371
CUBRID System Catalog	374
Overview	374
System Catalog Classes	374
System Catalog Virtual Class	384
Catalog Class/Virtual Class Authorization	399
Consistency of Catalog Information	399
Querying on Catalog	399
Administrator's Guide	401
CUBRID Utilities	402
CUBRID Controls	404
How to Use CUBRID Utilities (Syntax)	404
CUBRID Services	404
Database Server	407
Broker	409

Printed Documentation

CUBRID Manager Server	420
Database Administration	422
How to Use the CUBRID Management Utilities (Syntax)	422
Database Users	422
databases.txt File	422
Creating Database	423
Adding Database Volume	428
Deleting Database	430
Renaming Database	430
Renaming Database Host	431
Copying/Moving Database	432
Registering Database	433
Checking Used Space	434
Compacting Used Space	435
Updating Statistics	436
Outputting Statistics Information of Server	437
Checking Lock Status	440
Checking Database Consistency	442
Killing Database Transactions	443
Checking the Query Plan Cache	444
Outputting Internal Database Information	445
Backup and Restore	446
Exporting and Importing	446
Dumping Parameters Used in Server/Client	446
Database Migration	447
Migrating Database	447
Unloading Database	447
Loading Database	450
How to Write Files to Load Database	454
Database Backup and Restore	457
Database Backup	457
Backup Strategy and Method	460
Managing Backup Files	462
Managing Archive Logs	462
Restoring Database	462
Restore Strategy and Procedure	465
Restoring Database to Different Server	466
CUBRID HA	468
Overview	468

CUBRID HA Concept _____	469
CUBRID HA Features _____	476
Quick Start _____	481
Environment Configuration _____	484
Running and Monitoring _____	490
Configuration _____	494
Constraints _____	503
Error Messages _____	504
Operation Scenario _____	508

Performance Tuning _____ 513

Database Server Configuration _____	514
Scope of Database Server Configuration _____	514
cubrid.conf Configuration File and Default Parameters _____	514
Connection-Related Parameters _____	517
Memory-Related Parameters _____	518
Disk-Related Parameters _____	519
Error Message-Related Parameters _____	521
Concurrency/Lock Parameters _____	522
Logging-Related Parameters _____	523
Transaction Processing-Related Parameters _____	526
Statement/Type-Related Parameters _____	526
Query Cache-Related Parameters _____	530
Utility-Related Parameters _____	530
HA-Related Parameters _____	531
Other Parameters _____	531
Changing Database Server Configuration _____	534
Broker Configuration _____	536
Broker Configuration File and Default Parameters _____	536
Common Parameters _____	537
Parameter by Broker _____	538

API Reference _____ 545

JDBC API _____	546
JDBC Programming _____	546
CUBRIDOID _____	556
CUBRIDPreparedStatement _____	562

Printed Documentation

CUBRIDResultSet	563
CUBRIDResultSetMetaData	564
CUBRIDStatement	565
ODBC API	567
ODBC Programming	567
OLE DB API	571
OLE DB Programming	571
PHP API	576
PHP Programming	576
cubrid_affected_rows	579
cubrid_bind	579
cubrid_client_encoding	582
cubrid_close	583
cubrid_close_prepare, cubrid_close_request	583
cubrid_col_get	584
cubrid_col_size	585
cubrid_column_names	585
cubrid_column_types	586
cubrid_commit	587
cubrid_connect	588
cubrid_connect_with_url	589
cubrid_current_oid	591
cubrid_data_seek	591
cubrid_db_name	592
cubrid_disconnect	593
cubrid_drop	594
cubrid_errno, cubrid_error_code	596
cubrid_error, cubrid_error_msg	596
cubrid_error_code_facility	597
cubrid_execute	598
cubrid_fetch	599
cubrid_fetch_array	600
cubrid_fetch_assoc	601
cubrid_fetch_field	602
cubrid_fetch_lengths	603
cubrid_fetch_object	604
cubrid_fetch_row	605
cubrid_field_flags	606
cubrid_field_len	606

cubrid_field_name	607
cubrid_field_seek	608
cubrid_field_table	609
cubrid_field_type	610
cubrid_free_result	610
cubrid_get	611
cubrid_get_autocommit	612
cubrid_get_charset	613
cubrid_get_class_name	614
cubrid_get_client_info	614
cubrid_get_db_parameter	615
cubrid_get_query_timeout	616
cubrid_get_server_info	617
cubrid_insert_id	618
cubrid_is_instance	618
cubrid_lob_close	619
cubrid_lob_export	620
cubrid_lob_get	620
cubrid_lob_send	621
cubrid_lob_size	622
cubrid_list_dbs	622
cubrid_lock_read	623
cubrid_lock_write	624
cubrid_move_cursor	625
cubrid_next_result	626
cubrid_num_cols/cubrid_num_fields	628
cubrid_num_rows	628
cubrid_pconnect	629
cubrid_pconnect_with_url	630
cubrid_ping	632
cubrid_prepare	632
cubrid_put	634
cubrid_query	635
cubrid_real_escape_string	636
cubrid_result	637
cubrid_rollback	638
cubrid_schema	639
cubrid_seq_drop	644
cubrid_seq_insert	645

Printed Documentation

cubrid_seq_put	646
cubrid_set_add	647
cubrid_set_autocommit	648
cubrid_set_db_parameter	649
cubrid_set_drop	650
cubrid_set_query_timeout	651
cubrid_unbuffered_query	651
cubrid_version	652
CCI API	654
CCI Overview	654
cci_bind_param	666
cci_bind_param_array	668
cci_bind_param_array_size	668
cci_blob_free	669
cci_blob_new	669
cci_blob_write	669
cci_blob_size	670
cci_blob_write	670
cci_clob_free	671
cci_clob_new	671
cci_clob_write	672
cci_clob_size	672
cci_clob_write	673
cci_close_req_handle	673
cci_col_get	673
cci_col_seq_drop	674
cci_col_seq_insert	675
cci_col_seq_put	675
cci_col_set_add	676
cci_col_set_drop	676
cci_col_size	677
cci_connect	677
cci_connect_with_url	678
cci_cursor	679
cci_cursor_update	680
cci_datasource_borrow	680
cci_datasource_create	681
cci_datasource_destroy	681
cci_datasource_release	682

cci_disconnect	682
cci_end_tran	683
cci_execute	684
cci_execute_array	685
cci_execute_batch	687
cci_execute_result	688
cci_fetch	689
cci_fetch_buffer_clear	689
cci_fetch_sensitive	690
cci_fetch_size	690
cci_get_autocommit	691
cci_get_bind_num	691
cci_get_class_num_objs	691
CCI_GET_COLLECTION_DOMAIN	692
cci_get_cur_oid	692
cci_get_data	692
cci_get_db_parameter	693
cci_get_db_version	694
cci_get_query_timeout	694
cci_get_result_info	695
CCI_GET_RESULT_INFO_ATTR_NAME	696
CCI_GET_RESULT_INFO_CLASS_NAME	696
CCI_GET_RESULT_INFO_IS_NON_NULL	696
CCI_GET_RESULT_INFO_NAME	697
CCI_GET_RESULT_INFO_PRECISION	697
CCI_GET_RESULT_INFO_SCALE	697
CCI_GET_RESULT_INFO_TYPE	698
CCI_IS_SET_TYPE, CCI_IS_MULTISSET_TYPE, CCI_IS_SEQUENCE_TYPE, CCI_IS_COLLECTION_TYPE	698
cci_is_updatable	698
cci_next_result	699
cci_oid	699
cci_oid_get	700
cci_oid_get_class_name	700
cci_oid_put	701
cci_oid_put2	701
cci_prepare	702
cci_prepare_and_execute	703
cci_property_create	704
cci_property_destroy	704

Printed Documentation

cci_property_get _____	705
cci_property_set _____	705
CCI_QUERY_RESULT_ERR_MSG _____	706
cci_query_result_free _____	706
CCI_QUERY_RESULT_RESULT _____	707
CCI_QUERY_RESULT_STMT_TYPE _____	707
cci_savepoint _____	708
cci_schema_info _____	708
cci_set_allocators _____	712
cci_set_autocommit _____	715
cci_set_db_parameter _____	715
cci_set_element_type _____	715
cci_set_free _____	716
cci_set_get _____	716
cci_set_isolation_level _____	717
cci_set_make _____	717
cci_set_max_row _____	717
cci_set_query_timeout _____	718
cci_set_size _____	719

Introduction to Manual

Manual Contents

The contents of the CUBRID Database Management System (CUBRID DBMS) product manual are as follows:

- [Introduction to CUBRID](#) : This chapter provides a description of the structure and characteristics of the CUBRID DBMS.
- [Getting Started with CUBRID](#) : The "Getting Started with CUBRID" provides users with a brief explanation on what to do when first starting CUBRID. The chapter contains information on new features added to CUBRID, on how to install and execute the system, and provides a simple guide on how to use the CSQL Interpreter and CUBRID Manager. The chapter also includes examples of how to write application programs using JDBC, PHP, ODBC, CCI, etc.
- [CSQL Interpreter](#) : CSQL is an application that allows you to use SQL statements through a command-driven interface. This chapter explains how to use the CSQL Interpreter and associated commands.
- [CUBRID SQL Guide](#) : This chapter describes SQL syntaxes such as data types, functions and operators, data retrieval or table manipulation. The chapter also provides SQL syntaxes used for indexes, triggers, partitioning, serial and user information changes, etc.
- [Administrator's Guide](#) : This chapter provides instructions on how to create, drop, back up, restore and migrate a database, and executing CUBRID HA. Also it includes instructions on how to use CUBRID utilities, which starts and stops the Server, Broker and CUBRID Manager servers, etc.
- [Performance Tuning](#) : The "Performance Tuning" chapter provides instructions on setting system parameters that may influence the performance. This chapter provides information on how to use the configuration file for the Server and Broker, and describes the meaning of each parameter.
- [API Reference](#) : This chapter provides information on JDBC API, ODBC API, OLE DB API, PHP API, and CCI API.

Manual Conventions

The following table provides conventions on definitions used in the CUBRID Database Management System product manual to identify "statements," "commands" and "reference within texts."

Convention	Description	Example
Italics	Italics type is used to show the variable names.	<i>persistent:</i> stringVariableName
Boldface	Boldface type is used for names such as the member function name, class name, constants, CUBRID keyword or names such as other required characters.	fetch () member function class odb_User
Constant Width	Constant Width type is used to show segments of code example or describes a command's execution and results.	csql database_name
UPPER-CASE	UPPER-CASE is used to show the CUBRID keyword (see Boldface).	SELECT
Single Quotes (')	Single quotes (') are used with braces and brackets, and shows the necessary sections of a syntax. Single quotes are also used to enclose strings.	{'const_list'}
Brackets ([])	Brackets ([]) indicate optional parameters or keywords.	[ONLY]

Underline(_)	Underline (_) indicates a default keyword if no keyword is specified.	[DISTINCT UNIQUE ALL]
Vertical bar()	Vertical bar () indicates that one or another option can be specified.	[COLUMN ATTRIBUTE]
Braces around parameters({ })	Braces around parameters indicate that one of those parameters must be specified in a statement syntax.	{2, 4, 6}
Braces with ellipsis({ } ...)	Braces before an ellipsis indicate that a parameter can be repeated.	{, class_name}...
Angle brackets(< >)	Angle brackets indicate a single key or a series of key strokes.	<Ctrl+n>

Introduction to CUBRID

This chapter explains the architecture and features of CUBRID. CUBRID is an object-relational database management system (DBMS) consisting of the Database Server, the Broker, and the CUBRID Manager. It is optimized for Internet data services, and provides various user-friendly features.

This chapter covers the following topics:

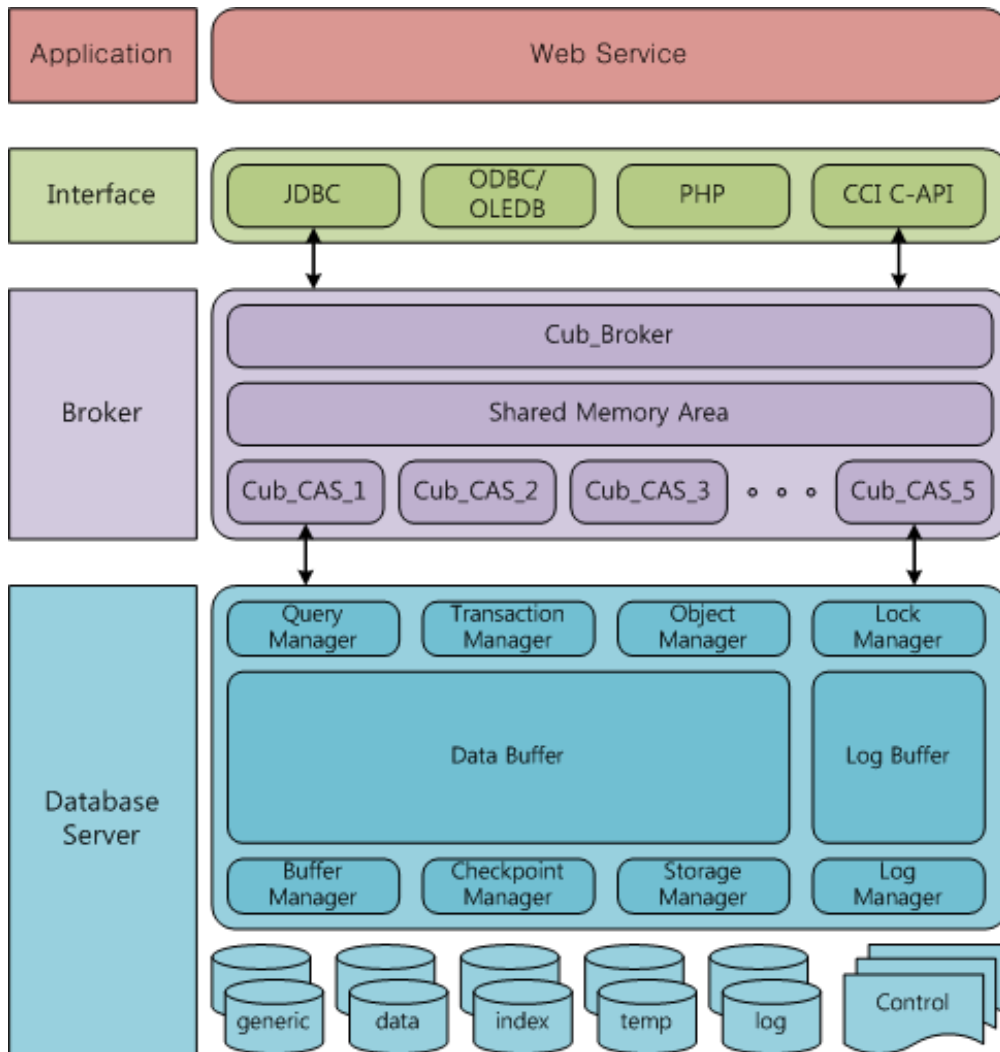
- System Architecture
- Features of CUBRID

System Architecture

System Architecture

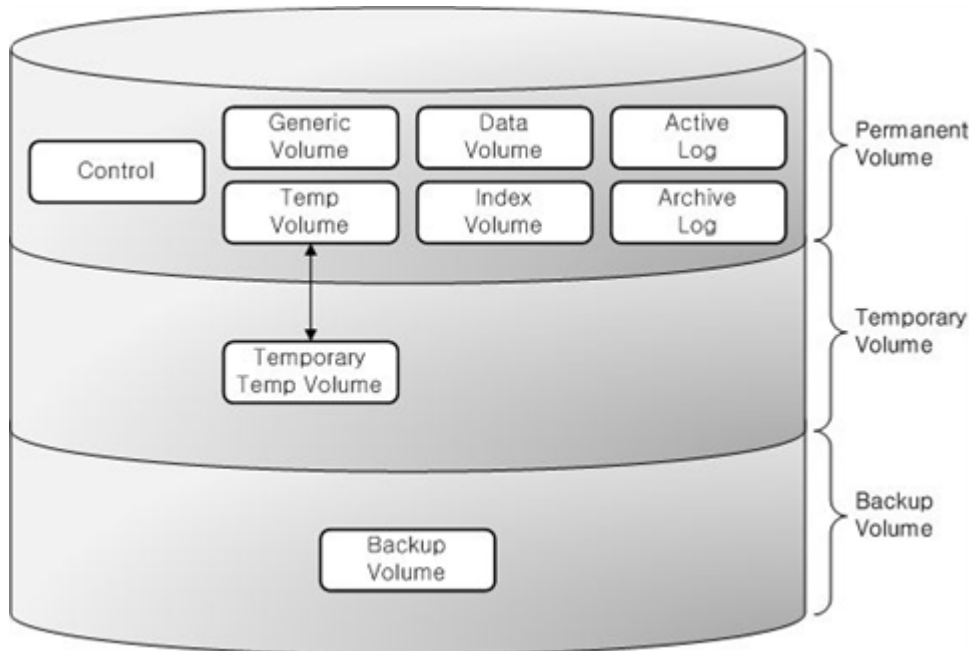
CUBRID is an object-relational database management system (DBMS) consisting of the Database Server, the Broker, and the CUBRID Manager.

- As the core component of the CUBRID Database Management System, the Database Server stores and manages data in multi-threaded client/server architecture. The Database Server processes the queries requested by users and manages objects in the database. The CUBRID Database Server provides seamless transactions using locking and logging methods even when multiple users use the database at once. It also supports database backup and restore for the operation.
- The Broker is a CUBRID-specific middleware that relays the communication between the Database Server and external applications. It provides functions including connection pooling, monitoring, and log tracing and analysis.
- The CUBRID Manager is a GUI tool that allows users to remotely manage the database and the Broker. It also provides the Query Editor, a convenient tool that allows users to execute SQL queries on the Database Server. See CUBRID Manager manual or online manual for more information on the CUBRID Manager.



Database Volume Structure

The following diagram illustrates the CUBRID database volume structure. As you can see, the database is divided into three volumes: permanent, temporary and backup. This chapter will examine each volume and its characteristics.



Permanent Volume

Permanent volume is a database volume that exists permanently once it is created. Its types include generic, data, temp, index, control, active log and archive log.

Generic Volume

For efficient management, the volume type to be added to the database can be specified as one of the followings: data, temp or index. If data usage is not specified, it is specified as a generic volume.

Data Volume

Data volume is a volume for storing data such as instances, tables and multimedia data.

Temp Volume

Temporary volume is a volume used temporarily for query processing and sorting. However, the temporary volume is not a volume where the storage is created and destroyed temporarily, but one of the permanent volumes with permanent spaces where the data is stored and destroyed temporarily. Therefore, the data in the temporary volume space gets initialized when CUBRID restarts without leaving any log info.

Index Volume

Index volume is a volume that holds the index information for fast query processing or integrity constraint checks.

Control File

The control file contains the volume, backup and log information in the database.

- **Volume Information** : The information that includes names, locations and internal volume identifiers of all the volumes in the database. When the database restarts, the CUBRID reads the volume information control file. It records a new entry to that file when a new database volume is added.
- **Backup Information** : Locations of all the backups for data, index, and generic volumes are recorded to a backup information control file. This control file is maintained where the log files are managed.

- **Log Information** : This information contains names of all active and archive logs. With the log information control file, you can verify the archive log information. The log information control file is created and managed at the same location as the log files.

Control files include the information about locations of database volumes, backups and logs. Since these files will be read when the database restarts, users must not modify them arbitrarily.

Active Log

Active log is a log that contains recent changes to the database. If a problem occurs, you can use active and archive logs to restore the database completely up to the point of the last commit before the occurrence of the fault.

Archive Log

Archive log is a volume to store logs continuously created after exhausting available active log space that contains recent changes. The archive log volume will be generated only after exhausting available active log volume space, just as the temporary temp volume will be generated after exhausting available permanent temp volume space. Temporary temp volume is automatically destroyed when server processes terminate but archive log is not automatically destroyed. Therefore, it is required to configure archive log to be deleted by system.

Note To get information on the conditions in which archive log can be deleted, see [Managing Archive Log](#).

Background Archive Log

Background archive log is a volume used in the background with log archiving temporarily before creating archive logs. It is created as the same volume size as active log and stored.

Temporary Volume

Temporary volume has the opposite meaning to the permanent volume. That is, the temporary volume is a storage created only when the accumulated data exceeds the space specified by the user as the permanent volume. The temporary volume is destroyed when the server process terminates. One of such volumes created or destroyed temporarily is the temporary temp volume.

Temporary Temp Volume

Temporary temp volume is a temporary volume created temporarily by the system after exhausting the space specified as the permanent temp volume, whereas the temporary volume belongs to the permanent volume with the permanent space specified. Therefore, adding permanent temp volume with an appropriate size can enhance performance based on operation environment; it is recommended for DBA to consider this case.

DBA should consider space where temporary temp volume can be created when creating a database. Once temporary temp volume is created, it is maintained until a database restarts and its size cannot be reduced. It is recommended to make temporary temp volume automatically delete by restarting a database if its size is too big.

The temporary temp volume is created to free up disk space needed for joining/sorting or index creation. Examples of such queries of creating temporary volume are: 1) SQL statements with a **GROUP BY** or **ORDER BY**, 2) SQL statements that contain coordinated subqueries, 3) join queries that perform sort-merge joins, and 4) a **CREATE INDEX** statement.

- **File name of the temporary temp volume** : The file name of the temporary temp volume of CUBRID has the format of *db_name_tnum*, where *db_name* is the database name and *num* is the volume identifier. The volume identifier is decremented by 1 from 32766.
- **Configuring the temporary temp volume size** : The number of temporary temp volumes to be created is determined by the system depending on the space size needed for processing transactions. However, users can limit the temporary temp volume size by configuring the **temp_file_max_size_in_pages** parameter value in the system parameter configuration file (**cubrid.conf**). If the **temp_file_max_size_in_pages** parameter value is configured to 0, the temporary temp volume will not be created even after exhausting the permanent temp volume.
- **Configuring storing location of the temporary temp volume** : By default, the temporary temp volume is created where the first database volume was created. However, you can specify a different directory to store the temporary temp volume by configuring the **temp_volume_path** parameter value.

- **Deleting the temporary temp volume** : The temporary temp volume exists temporarily only while the database is running. Therefore, you must not delete the temporary temp volume when running servers. The temporary temp volume is deleted if database servers are normally terminated while it is deleted when the servers restart if database servers are abnormally terminated.

Backup Volume

Backup volume is a database snapshot; based on such backup and log volumes, you can restore transactions to a certain point of time.

You can use the **cubrid backupdb** utility to copy all the data needed for database restore, or configure the **backup_volume_max_size_bytes** parameter value in the database configuration file (**cubrid.conf**) to adjust the backup volume partitioning size.

Database Server

Database Server Process

Each database has a single server process. The server process is the core component of the CUBRID Database Server, and handles a user's requests by directly accessing database and log files. The client process connects to the server process via TCP/IP communication. Each server process creates threads to handle requests by multiple client processes. System parameters can be configured for each database, that is, for each server process. The server process can connect to as many client processes as specified by the **max_clients** parameter value.

Master Process

The master process is a broker process that allows the client process to connect to and communicate with the server process. One master process runs for each host. (To be exact, one master process exists for each connection port number specified in the **cubrid.conf** system parameter file.) While the master process listens on the TCP/IP port specified, the client process connects to the master process through that port. The master process changes a socket to server port so that the server process can handle connection.

Execution Mode

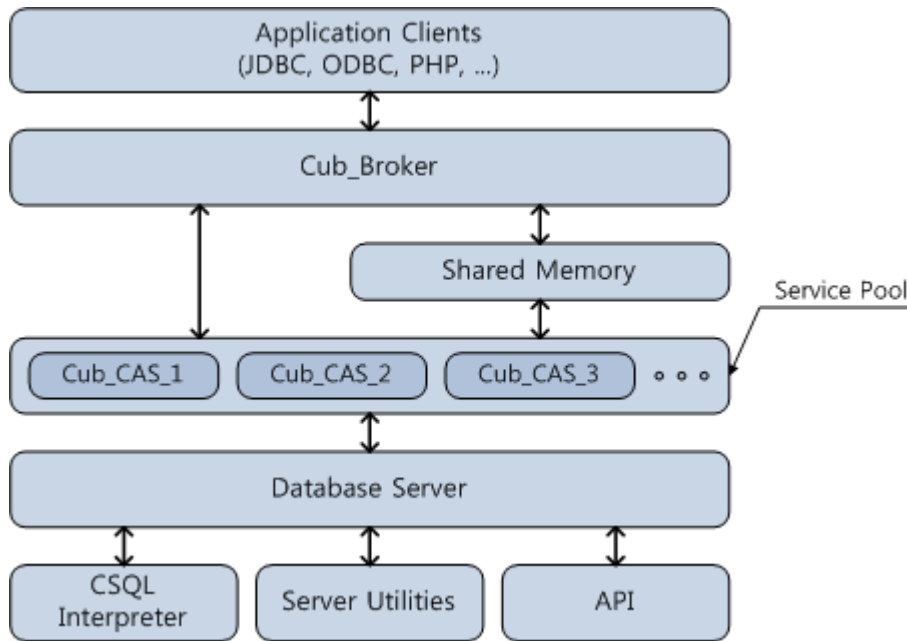
All CUBRID utilities except the server process have two modes: client/server mode and standalone mode.

- In client/server mode, an application accesses to the server process by operating itself as a client process.
- In standalone mode, an application accesses to the existing file of a server process.
- In the standalone mode, a process is shared between a client and a server, wherein a master process is not required and a database can be directly accessed.

For example, a database creation or a restore utility runs in the standalone mode so it can use the database exclusively by denying the access by multiple users. Another example is that the CSQL Interpreter can either connect to the server process in client/server mode or execute SQL statements by accessing the database in the standalone mode. Note that one database cannot be accessed simultaneously by a server process and a standalone program.

Broker

The Broker is a middleware that allows various application clients to connect to the Database Server. As shown below, the CUBRID system, which includes the Broker, has multi-layered architecture consisting of application clients, **cub_broker**, **cub_cas** and the Database Server.



Application Client

The interfaces that can be used in application clients include C-API (CCI, CUBRID Call Interface), ODBC, JDBC, PHP, Tcl/Tk, Python, and Ruby, OLEDB, and ADO.NET.

cub_cas

cub_cas (CUBRID Common Application Server) acts as a common application server used by all the application clients that request connections. **cub_cas** also acts as the Database Server's client and provides the connection to the Database Server upon the client's request. The number of **cub_cas**(s) running in the service pool can be specified in the configuration file, and this number is dynamically adjusted by **cub_broker**.

cub_cas is a program linked to the CUBRID Database Server's client library and functions as a client module in the server process. In the client module, tasks such as query parsing, optimization, execution plan creation are performed.

cub_broker

cub_broker relays the connection between the application client and the **cub_cas**. That is, when an application client requests access, the **cub_broker** checks the status of the **cub_cas** through the shared memory, and then delivers the request to an accessible **cub_cas**. It then returns the processing results of the request from the **cub_cas** to the application client.

The **cub_broker** also manages the server load by adjusting the number of **cub_cas**(s) in the service pool and monitors and manages the status of the **cub_cas**. If the **cub_broker** delivers the request to **cub_cas** but the connection to **cub_cas** 1 fails because of an abnormal termination, it sends an error message about the connection failure to the application client and restarts **cub_cas** 1. Restarted **cub_cas** 1 is now in a normal stand-by mode, and will be reconnected by a new request from a new application client.

Shared Memory

The status information of the **cub_cas** is stored in the shared memory, and the **cub_broker** refers to this information to relay the connection to the application client. With the status information stored in the shared memory, the system manager can identify which task the **cub_cas** is currently performing or which application client's request is currently being processed.

Interface Module

CUBRID provides various Application Programming Interfaces (APIs). The following APIs are supported by CUBRID. CUBRID also provides interfaces modules for each interface.

- JDBC : A standard API used to create database applications in Java. CUBRID provides the JDBC driver as an interface module.
- ODBC : A standard API used to create database applications on Windows. CUBRID provides the ODBC driver as an interface module.
- OLE DB : An API used to create COM-based database applications on Windows. CUBRID provides the OLE DB provider as an interface module.
- PHP : CUBRID provides a PHP interface module to create database applications in the PHP environment. The PHP module is based on the CCI library.
- CCI : CCI is a C language interface provided by CUBRID. The interface module is provided as a C library.

All interface modules access the Database Server through the Broker. The Broker is a middleware that allows various application clients to connect to the Database Server. When it receives a request from an interface module, it calls a native C API provided by the Database Server's client library.

You can find the latest information on interface modules, visit the Web site at <http://www.cubrid.org/wiki/apis>.

CUBRID Characteristics

Transaction Support

CUBRID supports the following features to completely ensure the atomicity, consistency, isolation and durability in transactions.

- Supporting commit, rollback, savepoint per transaction
- Ensuring transaction consistency in the event of system or database failure
- Ensuring transaction consistency between replications
- Supporting multiple granularity locking of databases, tables and records
- Resolving deadlocks automatically
- Supporting distributed transactions (two-phase commit)

Database Backup and Restore

A database backup is the process of copying CUBRID database volumes, control files and log files; a database restore is the process of restoring the database to a certain point in time using backup files, active logs and archive logs copied by the backup process. For a restore, there must be the same operating system and the same version of CUBRID installed as in the backup environment.

The backup methods which CUBRID supports include online, offline and incremental backups; the restore methods include restore using incremental backups as well as partial and full restore.

Table Partitioning

Partitioning is a method by which a table is divided into multiple independent logical units. Each logical unit is called a partition, and each partition is divided into a different physical space. This will lead performance improvement by only allowing access to the partition when retrieving records. CUBRID provides three partitioning methods:

- Range partitioning: Divides a table based on the range of a column value
- Hash partitioning: Divides a table based on the hash value of a column
- List partitioning: Divides a table based on the column value list

Supports a Variety of Index Functions

CUBRID supports the following index functions to utilize indices while executing a variety of conditional queries.

- Descending Index Scan: Descending Index Scan is available only with Ascending Index Scan, without creating separate reverse indexes.
- Covering Index: When the column of a **SELECT** list is included in the index, the requested data can be obtained with an index scan.
- **ORDER BY** Clause Optimization: If the required record sorting order is identical to the order of indices, no additional sorting is required (Skip **ORDER BY**).
- **GROUP BY** Clause Optimization: If all columns in the **GROUP BY** clause are included in the indices, they are available to use while executing queries. Therefore, no additional sorting is required (Skip **GROUP BY**).

HA

CUBRID provides High Availability (HA) to minimize system down time while continuing normal operation of server in the event of hardware, software, or network failure. The structure of CUBRID HA is shared-nothing. CUBRID monitors its system and status on a real time basis with the CUBRID Heartbeat and performs failover when failure occurs. It follows the two steps below to synchronize data from the master database server to slave database server.

- A transaction log multiplication step where the transaction log created in the database server is replicated in real time to another node
- A transaction log reflection step where data is applied to the slave database server through the analysis of the transaction log being replicated in real time

Java Stored Procedure

A stored procedure is a method to decrease the complexity of applications and to improve the reusability, security and performance through the separation of database logic and middleware logic. A stored procedure is written in Java (generic language), and provides Java stored procedures running on the Java Virtual Machine (JVM). To execute Java stored procedures in CUBRID, the following steps should be performed:

- Install and configure the Java Virtual Machine
- Create Java source files
- Compile the files and load Java resources
- Publish the loaded Java classes so they can be called from the database
- Call the Java stored procedures

Click Counter

In the Web, it is a common scenario to count and keep the number of clicks to the database in order to record retrieval history.

The above scenario is generally implemented by using the **SELECT** and **UPDATE** statements; **SELECT** retrieves the data and **UPDATE** increases the number of clicks for the retrieved queries.

This approach can cause significant performance degradation due to increased lock contention for **UPDATE** when a number of **SELECT** statements are executed against the same data.

To address this issue, CUBRID introduces the new concept of the click counter that will support optimized features in the Web in terms of usability and performance, and provides the **INCR** function and the **WITH INCREMENT FOR** statement.

Extending the Relational Data Model

Collection

For the relational data model, it is not allowed that a single column has multiple values. In CUBRID, however, you can create a column with several values. For this purpose, collection data types are provided in CUBRID. The collection data type is mainly divided into **SET**, **MULTISET** and **LIST**; the types are distinguished by duplicated availability and order.

- **SET** : A collection type that does not allow the duplication of elements. Elements are stored without duplication after being sorted regardless of their order of entry.
- **MULTISET** : A collection type that allows the duplication of elements. The order of entry is not considered.
- **LIST** : A collection type that allows the duplication of elements. Unlike with **SET** and **MULTISET**, the order of entry is maintained.

Inheritance

Inheritance is a concept to reuse columns and methods of a parent table in those of child tables. CUBRID supports reusability through inheritance. By using inheritance provided by CUBRID, you can create a parent table with some common columns and then create child tables inherited from the parent table with some unique columns added. In this way, you can create a database model which can minimize the number of columns.

Composition

In a relational database, the reference relationship between tables is defined as a foreign key. If the foreign key consists of multiple columns or the size of the key is significantly large, the performance of join operations between tables will be degraded. However, CUBRID allows the direct use of the physical address (OID) where the records of the referred table are located, so you can define the reference relationship between tables without using join operations.

That is, in an object-oriented database, you can create a composition relation where one record has a reference value to another by using the column displayed in the referred table as a domain (type), instead of referring to the primary key column from the referred table.

Getting Started with CUBRID

This chapter contains useful information on starting CUBRID such as how to install and run CUBRID; also it provides instructions on how to use the CSQL Interpreter. This chapter also includes examples on how to write application programs using JDBC, PHP, ODBC and CCI, etc.

This chapter covers the following topics:

- Installing and Running CUBRID
- Configuring Environment Variable and Starting CUBRID
- Using the CSQL Interpreter
- Writing Programs using JDBC
- Writing Programs using PHP
- Writing Programs using ODBC and ASP
- Writing Programs using CCI

CUBRID provides various tools for you to use easily through GUI. You can download these tools at the website below:

- CUBRID Query Browser : http://www.cubrid.org/wiki_tools/entry/cubrid-query-browser
- CUBRID Manager : http://www.cubrid.org/wiki_tools/entry/cubrid-manager
- CUBRID Migration Toolkit : http://www.cubrid.org/wiki_tools/entry/cubrid-migration-toolkit

You can download drives such as JDBC, PHP, ODBC, and OLE DB at the website below:

- CUBRID APIs Wiki : http://www.cubrid.org/wiki_apis

Installing and Running

Installing and Running on Linux

Details to Check when Installing

Check the following before installing CUBRID for Linux.

Category	Description
Operating system	Only supports glibc 2.3.4 or later. The glibc version can be checked as follows: rpm -q glibc
64-bit	Since version 2008 R2.0, CUBRID supports both 32-bit and 64-bit Linux. You can check the version as follows: % uname -a Linux host_name 2.6.18-53.1.14.el5xen #1 SMP Wed Mar 5 12:08:17 EST 2008 x86_64 x86_64 x86_64 GNU/Linux Make sure to install the CUBRID 32-bit version on 32-bit Linux and the CUBRID 64-bit version on 64-bit Linux. The followings are the libraries that should be added. Curses Library (rpm -q ncurses) gcrypt Library (rpm -q libgcrypt) stdc++ Library (rpm -q libstdc++)
Available free memory space	1 GB or more recommended.
Available free disk space	2 GB or more recommended (500 MB is required for initial installation and 1.5 GB is required for creating the default option database).
Required software	The CUBRID Manager and Java stored procedures require the Java Runtime Environment (JRE) version 1.6 or later.

Installing CUBRID

The installation program consists shell scripts that contain binary; thus it can be installed automatically. The following example shows how to install CUBRID with the "CUBRID-8.3.0.0312-linux.x86_64.sh" file on the Linux.

```
[cub_user@cubrid ~]$ sh CUBRID-8.3.1.0168-linux.x86_64.sh
Do you agree to the above license terms? (yes or no) : yes
Do you want to install this software(CUBRID) to the default(/home1/cub_user/CUBRID)
directory? (yes or no) [Default: yes] : yes
Install CUBRID to '/home1/cub_user/CUBRID' ...
In case a different version of the CUBRID product is being used in other machines, please
note that the CUBRID 2008 R3.1 servers are only compatible with the CUBRID 2008 R3.1
clients and vice versa.
Do you want to continue? (yes or no) [Default: yes] : yes
Copying old .cubrid.sh to .cubrid.sh.bak ...

CUBRID has been successfully installed.

demodb has been successfully created.

If you want to use CUBRID, run the following commands
% ./home1/cub_user/.cubrid.sh
% cubrid service start
```

As shown in the example above, after installing the downloaded file (CUBRID-8.3.1.0168-linux.x86_64.sh), the CUBRID related environment variables must be set in order to use the CUBRID database. Such setting has been made automatically when logging in the concerned terminal. Therefore there is no need to re-set after the first installation.

```
[cub_user@cubrid ~]$ ./home1/cub_user/.cubrid.sh
```

After the CUBRID Manager is installed, you can start the CUBRID Manager server and Broker as follows:

```
[cub_user@cubrid ~]$ cubrid service start
```

After starting the CUBRID service, if you wish to check whether the service was properly started, then check whether the cub_* processes have been started with grep (as shown below).

```
[cub_user@cubrid ~]$ ps -ef | grep cub
cub_user 15200 1 0 18:57 ? 00:00:00 cub_master
cub_user 15205 1 0 18:57 pts/17 00:00:00 cub_broker
cub_user 15210 1 0 18:57 pts/17 00:00:00 query editor cub cas 1
cub_user 15211 1 0 18:57 pts/17 00:00:00 query editor cub cas 2
cub_user 15212 1 0 18:57 pts/17 00:00:00 query editor cub cas 3
cub_user 15213 1 0 18:57 pts/17 00:00:00 query editor cub cas 4
cub_user 15214 1 0 18:57 pts/17 00:00:00 query_editor_cub_cas_5
cub_user 15217 1 0 18:57 pts/17 00:00:00 cub_broker
cub_user 15222 1 0 18:57 pts/17 00:00:00 broker1 cub cas 1
cub_user 15223 1 0 18:57 pts/17 00:00:00 broker1 cub cas 2
cub_user 15224 1 0 18:57 pts/17 00:00:00 broker1 cub cas 3
cub_user 15225 1 0 18:57 pts/17 00:00:00 broker1 cub cas 4
cub_user 15226 1 0 18:57 pts/17 00:00:00 broker1_cub_cas_5
cub_user 15229 1 0 18:57 ? 00:00:00 cub_auto start
cub_user 15232 1 0 18:57 ? 00:00:00 cub_js start
```

Installing CUBRID (rpm File)

You can install CUBRID by using rpm file that is created on CentOS5. The way of installing and uninstalling CUBRID is the same as that of using general rpm utility. While CUBRID is being installed, a new system group (cubrid) and a user account (cubrid) are created. After installation is complete, you should log in with a cubrid user account to start a CUBRID service.

```
$ rpm -Uvh CUBRID-8.3.1.0168-el5.x86_64.rpm
```

When rpm is executed, CUBRID is installed in the cubrid home directory (/opt/cubrid) and related configuration file (cubrid.[c]sh) is installed in the /etc/profile.d directory. Note that demodb is not automatically installed. Therefore, you must execute /opt/cubrid/demo/make_cubrid_demo.sh. When installation is complete, enter the code below to start CUBRID.

```
[cubrid@cubrid ~]$ cubrid service start
```

Note You must check RPM dependency when installing with RPM. If you ignore (--nodeps) dependency, it may not be executed.

Even if you remove RPM, user accounts and databases that are created after installing, you must remove it manually, if needed.

Installing CUBRID on Fedora/CentOS

To install CUBRID using the yum command, you should know where the CUBRID package is located. Choose appropriate location based on your operating system.

- http://www.cubrid.org/yum_repository

For example, if you are using Fedora 16, execute the command below. In the example, fc16 refers to Fedora 16.

```
$ rpm -i http://yumrepository.cubrid.org/cubrid_repo_settings/8.4.0/cubridrepo-8.4.0-1.fc16.noarch.rpm
```

If you are using CentOS 6.2, execute the command below. In this example, el6.2 refers to CentOS.

```
$ rpm -i http://yumrepository.cubrid.org/cubrid_repo_settings/8.4.0/cubridrepo-8.4.0-1.el6.2.noarch.rpm
```

You can install the CUBRID package you have desired based on the command you execute. To install the latest version, execute the command below.

```
$ yum install cubrid
```

To install the earlier version, you should include version information in the command.

```
$ yum install cubrid-8.3.1
```

After installation is complete, configure environment variables including installation path of CUBRID and then apply them to system.

Installing CUBRID on Ubuntu

To install CUBRID using the apt-get command on Ubuntu, add the CUBRID storage first and then update the apt index.

```
$ sudo add-apt-repository ppa:cubrid/cubrid
$ sudo apt-get update
```

To install the latest version, execute the command below.

```
$ sudo apt-get install cubrid
```

To install the earlier version, you should include version information in the command.

```
$ sudo apt-get install cubrid-8.3.1
```

After installation is complete, configure environment variables including installation path of CUBRID and then apply them to system.

Upgrading CUBRID

When you specify an installation directory where the previous version of CUBRID is already installed, a message which asks to overwrite files in the directory will appear. Entering **no** will stop the installation.

```
Directory '/home1/cub user/CUBRID' exist!
If a CUBRID service is running on this directory, it may be terminated abnormally.
And if you don't have right access permission on this directory(subdirectories or files),
install operation will be failed.
Overwrite anyway? (yes or no) [Default: no] : yes
```

Choose whether to overwrite the existing configuration files during the CUBRID installation. Entering **yes** will overwrite and back up them as extension .bak files.

```
The configuration file (.conf or .pass) already exists. Do you want to overwrite it? (yes
or no) : yes
```

Configuring Environment

You can modify the environment such as service ports etc. edit the parameters of a configuration file located in the **\$CUBRID/conf** directory. See [Environment Configuration](#) for more information.

Installing CUBRID Interfaces

You can see the latest information on interface modules such as JDBC, PHP, ODBC, and OLE DB and install them by downloading files from http://www.cubrid.org/wiki_apis.

Installing CUBRID Tools

You can see the latest information on tools such as CUBRID Manager and CUBRID Query Browser and install them by downloading files from http://www.cubrid.org/wiki_tools.

Installing and Running on Windows

Details to Check when Install

CUBRID 2008 R2.0 supports both 32-bit and 64-bit Windows. You can check the version by selecting [My Computer] > [System Properties]. Make sure to install the CUBRID 32-bit version on 32-bit Windows and the CUBRID 64-bit version on 64-bit Windows.

Category	Description
64-bit	Since version 2008 R2.0, CUBRID supports both 32-bit and 64-bit Windows. You can check the version by selecting [My Computer] > [System Properties]. Make sure to install the CUBRID 32-bit version on 32-bit Windows and the CUBRID 64-bit version on 64-bit Windows.
Available free memory space	1 GB or more recommended.

Available free disk space	2 GB or more recommended (500 MB is required for initial installation and 1.5 GB is required for creating the default option database).
Required software	The CUBRID Manager and Java stored procedures require the Java Runtime Environment (JRE) version 1.6 or later.

If CUBRID Service Tray does not automatically run upon system startup, you should check followings:

- Go to [Control Panel] > [Administrative Tools] > [Service] and verify whether Task Scheduler has started. If not, start Task Scheduler.
- Go to [Administrative Tools] > [Task Scheduler] and verify whether CUBRID Service Tray has been registered. If not, register CUBRID Service Tray.

Setup Type

- **Server and Driver Installation** : CUBRID Server, CSQL (a command line tool), interface drivers (OLEDB Provider, ODBC, JDBC, C API) are all installed.
- **Driver Installation** : The interface drivers (OLEDB Provider, ODBC, JDBC, C API) are only installed. You can select this type of installation if development or operation is performed by remote connection to the computer in which the CUBRID database server is installed.

Upgrading CUBRID

To install a new version of CUBRID in an environment in which a previous version has already been installed, select [CUBRID Service Tray] > [Exit] from the menu to stop currently running services, and then remove the previous version of CUBRID. Note that when you are prompted with "Do you want to delete all the existing version of databases and the configuration files?" you must select "No" to protect the existing databases.

For more information on migrating a database from a previous version to a new version, see [Migrating Database](#).

Configuring Environment

You can change configuration such as service ports to meet the user environment by changing the parameter values of following files which are located in the %CUBRID%\conf directory.

File	Description
cm.conf	CUBRID Manager's configuration file; the port number 8001 is configured by default. Two port numbers are required to use CUBRID; a configured number and the number added by 1 are used. For example, 8001 is configured for connection, the port number 8001 and 8002 are reserved.
cubrid.conf	Server configuration file is used to set the following: database memory, the number of threads due to the number of concurrent users, connection port between the Broker and Server, etc. See cubrid_broker.conf Configuration File and Default Parameters for details.
cubrid_broker.conf	Broker configuration file; the port is used by the broker that is operated. The file is used to set the number of CAS, SQL LOGs, etc. The ports shown in drivers such as JDBC are the concerned Broker's ports. See Parameter by Broker for details.

Installing CUBRID Interfaces

You can see the latest information on interface modules such as JDBC, PHP, ODBC, and OLE DB and install them by downloading files from http://www.cubrid.org/wiki_apis.

Installing CUBRID Tools

You can see the latest information on tools such as CUBRID Manager and CUBRID Query Browser and install them by downloading files from http://www.cubrid.org/wiki_tools.

Configuring Environment Variable and Starting CUBRID

Configuring the Environment Variable

The following environment variables need to be set in order to use the CUBRID. The necessary environment variables are automatically set when the CUBRID system is installed or can be changed, as needed, by the user.

CUBRID Environment Variables

- **CUBRID** : The default environment variable that designates the location where the CUBRID is installed. This variable must be set accurately since all programs included in the CUBRID system uses this environment variable as reference.
- **CUBRID_DATABASES** : The environment variable that designates the location of the database location information file. The CUBRID system stores and manages the absolute path of database volumes that are used in the **\$CUBRID_DATABASES/databases.txt** file. See [databases.txt file](#).
- **CUBRID_LANG** : The environment variable that designates the language that will be used in the CUBRID system. Currently, CUBRID provides English (**en_US**) and Korean (**ko_KR.euckr** and **ko_KR.utf8**). it is not a mandatory setting. Therefore, if the variable has not been set, then refer to the LANG environment variable or use **en_US**, which is the default value. See [Language Setting](#).

The above mentioned environment variables are set when the CUBRID is installed. However, the following commands can be used to verify the setting.

For Linux :

```
% printenv CUBRID
% printenv CUBRID DATABASES
% printenv CUBRID_LANG
```

In Windows :

```
C:\> set CUBRID
```

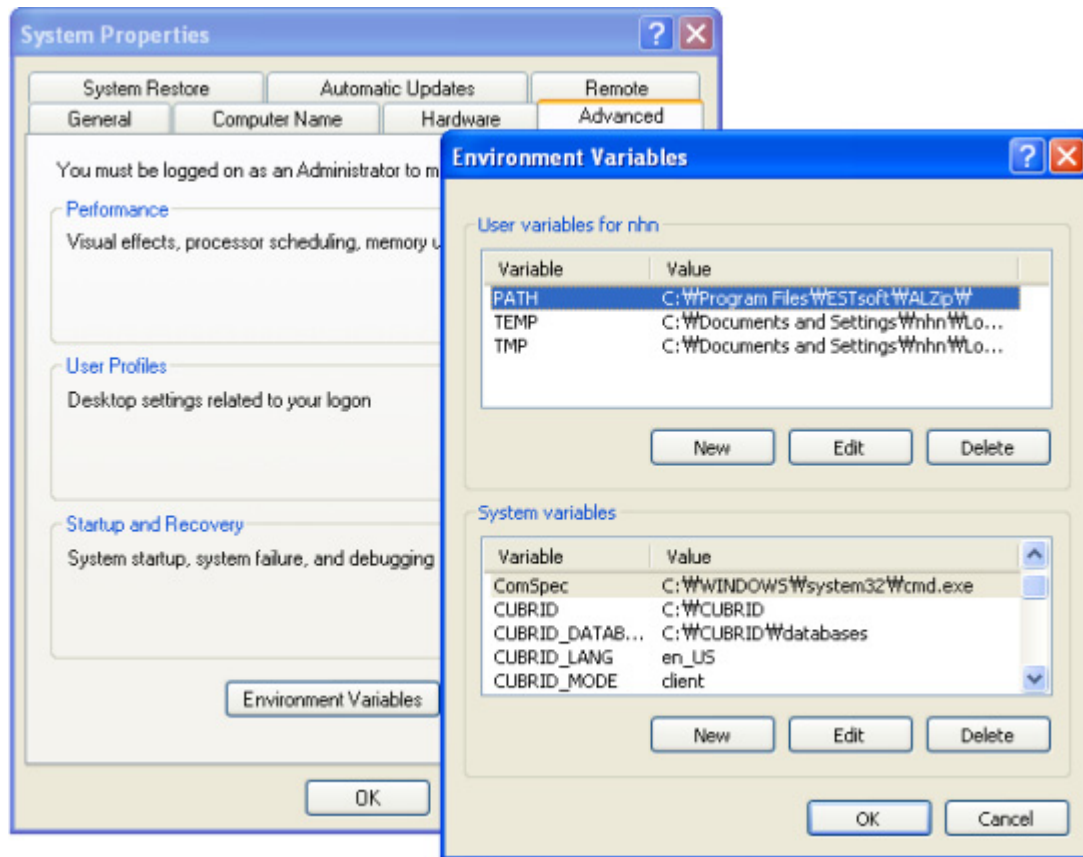
OS Environment and Java Environment Variables

- **PATH** : In the Linux environment, the directory **\$CUBRID/bin**, which includes a CUBRID system executable file, must be included in the PATH environment variable.
- **LD_LIBRARY_PATH** : In the Linux environment, **\$CUBRID/lib**, which is the CUBRID system's dynamic library file (**libjvm.so**), must be included in the **LD_LIBRARY_PATH** (or **SHLIB_PATH** or **LIBPATH**) environment variable.
- **Path** : In the Windows environment, the **%CUBRID%\bin**, which is a directory that contains CUBRID system's execution file, must be included in the **Path** environment variable.
- **JAVA_HOME** : To use the Java stored procedure in the CUBRID system, the Java Virtual Machine (JVM) version 1.6 or later must be installed, and the **JAVA_HOME** environment variable must designate the concerned directory. See the [Environment Configuration for Java Stored Functions/Procedures](#).

Configuring the Environment Variable

For Windows

If the CUBRID system has been installed on Windows, then the installation program automatically sets the necessary environment variable. Select [Systems Properties] in [My Computer] and select the [Advanced] tab. Click the [Environment Variable] button and check the setting in the [System Variable]. The settings can be changed by clicking on the [Edit] button. See the Windows help for more information on how to change the environment variable on Windows.



For Linux

If the CUBRID system has been installed on Linux, the installation program automatically creates the `.cubrid.sh` or `.cubrid.csh` file and makes configurations so that the files are automatically called from the installation account's shell log-in script. The following is the `.cubrid.sh` environment variable setting file that was created in an environment that uses `sh`, `bash`, etc.

```
CUBRID=/home1/cub user/CUBRID
CUBRID DATABASES=/home1/cub user/CUBRID/databases
CUBRID LANG=en_US
ld_lib_path=`printenv LD_LIBRARY_PATH`
if [ "$ld_lib_path" = "" ]
then
LD_LIBRARY_PATH=$CUBRID/lib
else
LD_LIBRARY_PATH=$CUBRID/lib:$LD_LIBRARY_PATH
fi
SHLIB_PATH=$LD_LIBRARY_PATH
LIBPATH=$LD_LIBRARY_PATH
PATH=$CUBRID/bin:$CUBRID/cubridmanager:$PATH
export CUBRID
export CUBRID DATABASES
export CUBRID_LANG
export LD_LIBRARY_PATH
export SHLIB_PATH
export LIBPATH
export PATH
```

Language Setting

The language that will be used in the CUBRID DBMS can be designated with the `CUBRID_LANG` environment variable. The following are values that can currently be set in the `CUBRID_LANG` environment variable.

- **en_US** : English (Default value)
- **ko_KR.euckr** : Korean EUC-KR encoding
- **ko_KR.utf8** : Korean utf-8 encoding

The language setting in the CUBRID system does not represent the character sets of data that is stored. In other words, even though the **CUBRID_LANG** is set to **ko_KR.utf8**, the data may not be changed to the concerned encoding. CUBRID's language setting will have an influence on the message printed from the program and will impact the date/time data type constant displayed throughout the use of the program.

If the **CUBRID_LANG** is not set, then the value of the LANG environment variable will be used. If the set value does not support the **CUBRID_LANG** or **LANG** value, then the action will be made as if the setting has been made to **en_US**, the default value.

Starting the CUBRID Service

Configure environment variables and language, and then start the CUBRID service. For more information on configuring environment variables and language, see [Registering Services](#) or [Starting and Stopping Services](#).

Shell Command

The following shell command can be used to start the CUBRID service and the demodb included in the installation package.

```
% cubrid service start
@ cubrid master start
++ cubrid master start: success
@ cubrid broker start
++ cubrid broker start: success
@ cubrid manager server start
++ cubrid manager server start: success

% cubrid server start demodb
@ cubrid server start: demodb

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.1

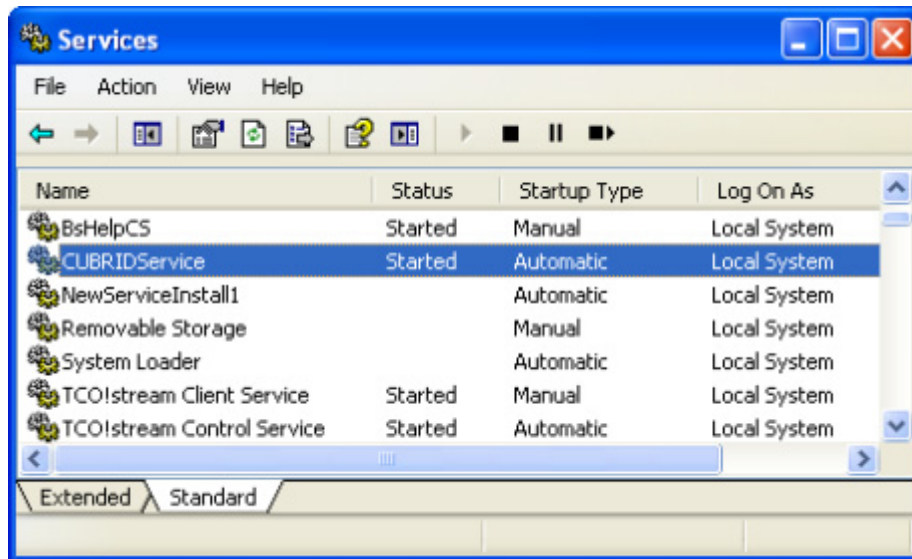
++ cubrid server start: success

@ cubrid server status
Server demodb (rel 8.4, pid 31322)
```

CUBRIDService or CUBRID Service Tray

On the Windows environment, you can start or stop a service as follows:

- Go to [Control Panel] > [Performance and Maintenance] > [Administrator Tools] > [Services] and select the CUBRIDService to start or stop the service.



- In the system tray, right-click the CUBRID Service Tray. To start CUBRID, select [Service Start]; to stop it, select [Service Stop]. Selecting [Service Start] or [Service Stop] menu would be like executing cubrid service start or cubrid service stop in a command prompt; this command runs or stops the processes configured in service parameters of cubrid.conf.
- If you click [Exit] while CUBRID is running, all the services and process in the server are stopped.

Note An administrator level (SYSTEM) authorization is required to start/stop CUBRID processes through the CUBRID Service tray; a login level user authorization is required to start/stop them with shell commands. If you cannot control the CUBRID processes on the Windows Vista or later version environment, select [Execute as an administrator (A)] in the [Start] > [All Programs] > [Accessories] > [Command Prompt]) or execute it by using the CUBRID Service Tray. When all processes of CUBRID Server stops, an icon on the CUBRID Service tray turns out gray.

Creating Databases

You can create databases by using the **cubrid createdb** utility and execute it where database volumes and log volumes are located. If you do not specify additional options such as **--db-volume-size** or **--log-volume-size**, 1.5 GB volume files are created by default (generic volume is set to 512 MB, active log is set to 512 MB, and background archive log is set to 512 MB).

```
% cd testdb
% cubrid createdb testdb

%ls -l
-rw----- 1 cubrid dbms 536870912 Jan 11 15:04 testdb
-rw----- 1 cubrid dbms 536870912 Jan 11 15:04 testdb_lgar t
-rw----- 1 cubrid dbms 536870912 Jan 11 15:04 testdb_lgat
-rw----- 1 cubrid dbms      176 Jan 11 15:04 testdb_lginf
-rw----- 1 cubrid dbms      183 Jan 11 15:04 testdb_vinf
```

In the above, testdb represents a generic volume file, testdb_lgar_t represents a background archive log file, testdb_lgat represents an active log file, testdb_lginf represents a log information file, and testdb_vinf represents a volume information file.

For details on volumes, see [Database Volume Structure](#). For details on creating volumes, see [Creating Database](#). It is recommended to classify and add volumes based on its purpose by using the **cubrid addvoldb** utility. For details, see [Adding Database Volume](#).

CSQL Interpreter

Starting the CSQL Interpreter

The CSQL Interpreter is a program used in CUBRID. The entered SQL statements and results can be stored in the file for later use. For more information [Introduction to the CSQL Interpreter](#) and [CSQL Execution Mode](#).

CUBRID offers the "CUBRID Manager" program, a convenient GUI program. All SQL can be executed and the results can be viewed from the CUBRID Manager's query editor. For more information, see CUBRID Manager manual or online manual.

In this section, we will provide information on using the CSQL Interpreter on Linux.

Starting the CSQL Interpreter

The CSQL program can be started in the shell as shown below.

```
% csql demodb
          CUBRID SQL Interpreter
Type ';'help' for help messages.
csql> ;help
=== <Help: Session Command Summary> ===
All session commands should be prefixed by ';' and only blanks/tabs
can precede the prefix. Capitalized characters represent the minimum
abbreviation that should be entered to execute the specified command.

;REAd    [<file-name>]      - read a file into command buffer.
;WRite   [<file-name>]      - (over)write command buffer into a file.
;APpend  [<file-name>]      - append command buffer into a file.
;PRINt   - print command buffer.
;SHELL   - invoke shell.
;CD      - change current working directory.
;EXit    - exit program.

;CLear   - clear command buffer.
;EDIT    - invoke system editor with command buffer.
;List    - display the content of command buffer.

;RUn     - execute sql in command buffer.
;Xrun    - execute sql in command buffer,
          and clears the command buffer.
;COMmit  - commit the current transaction.
;ROLLback - roll back the current transaction.
;AUTocommit [ON|OFF]      - enable/disable auto commit mode.
;REStart - restart database.

;SHELL_Cmd [shell-cmd]    - set default shell, editor, print and pager
;EDITOR_Cmd [editor-cmd]  - command to new one, or display the current
;PRINT_Cmd [print-cmd]    - one, respectively.

;DATE    - display the local time, date.
;DATAbase - display the name of database being accessed.
;SCHEMA class-name       - display schema information of a class.
;SYntax [sql-cmd-name]   - display syntax of a command.
;TRigger ['*' | trigger-name] - display trigger definition.
;Get system parameter    - get the value of a system parameter.
;SEt system parameter=value - set the value of a system parameter.
;PLan [simple|detail|off] - show query execution plan.
;Info <command>         - display internal information.
;Time [ON|OFF]          - enable/disable to display the query
                        execution time.

;HISTORYList            - display list of the executed queries.
;HISTORYRead <history num> - read entry on the history number into command buffer.
;HELP                  - display this help message.

csql>
```

Executing the SQL with CSQL

After the CSQL has been executed, you can enter the SQL into the CSQL prompt. Each SQL statement must end with a semicolon (;). Multiple SQL statements can be entered in a single line. You can find the simple usage of the session commands with the ;help command. For more information, see [Session Commands](#).

```
% csql demodb
CUBRID SQL Interpreter
Type `;help' for help messages.
csql> select * from olympic;
=== <Result of SELECT Command in Line 1> ===

      host year  host nation          host city          opening date  closing
date mascot              slogan              introduction
=====
2004  'Greece'          'Athens'          08/13/2004    08/29/2
004  'Athena Phevos'    'Welcome Home'    'In 2004 the Olympic Games re
turned to Greece, the home of both the ancient Olympics and the first modern Olympics.

<omitted>
25 rows selected.

Current transaction has been committed.

1 command(s) successfully processed.
csql> SELECT SUM(n) FROM (SELECT gold FROM participant WHERE nation code='KOR'
csql> UNION ALL SELECT silver FROM participant WHERE nation_code='JPN') AS t(n);

=== <Result of SELECT Command in Line 1> ===

      sum(n)
=====
          82

1 rows selected.

Current transaction has been committed.

1 command(s) successfully processed.
csql> ;exit
```

Programming with JDBC

Setting up the JDBC Environment

System Requirements

- JDK 1.6 or later
- CUBRID 2008 R1.0 or later
- CUBRID JDBC Driver 2008 R1.0 or later


Installing and Configuring Java Environment

You must already have Java installed and the JAVA_HOME environment variable set on your system. To install Java, download it from the Java homepage (<http://java.sun.com>). For more information, see [Environment Settings for Java Stored Functions/Procedures](#).


Configuring Environment Variables for Windows

After installing JAVA, double click [My Computer] and click [System Properties]. In the [Advanced] tab, click [Environment Variables]. The [Environment Variables] dialog will appear.

In the [System Variables], click [New]. Enter **JAVA_HOME** and Java installation path such as C:\Program Files\Java\jdk1.6.0_16 and then press [Enter].



Select "Path" and then click [Edit]. Add %JAVA_HOME%\bin to the variable and then click [OK].



You can configure **JAVA_HOME** and **PATH** in the shell.

```
set JAVA_HOME= C:\Program Files\Java\jdk1.6.0_16
set PATH=%PATH%;%JAVA_HOME%\bin
```

Configuring the Environment Variables for Linux

Specify the directory path where Java is installed (example : /usr/java/jdk1.6.0_16) in the **JAVA_HOME** environment variable, and add **\$JAVA_HOME/bin** to the **PATH** environment variable.

```
export JAVA_HOME=/usr/java/jdk1.6.0_16 //bash
export PATH=$JAVA_HOME/bin:$PATH //bash


setenv JAVA_HOME /usr/java/jdk1.6.0_16 //csh
```

```
set path = ($JAVA_HOME/bin $path) //csh
```

JDBC Driver Setting

To use the JDBC, set your **CLASSPATH** environment variable to the path where the CUBRID JDBC driver is located.

The CUBRID JDBC driver (**cuprid_jdbc.jar**) is located in jdbc directory which is subdirectory where CUBRID is installed.



Configuring the CLASSPATH Environment Variables for Windows

```
set CLASSPATH=%CUBRID%\jdbc\cuprid_jdbc.jar:.
```

Configuring the CLASSPATH Environment Variables for Linux

```
export CLASSPATH=$HOME/CUBRID/jdbc/cuprid_jdbc.jar:.
```

Note If a CUBRID JDBC driver has been installed in the same library directory (**\$JAVA_HOME/jre/lib/ext**) where the JRE is located, it may be loaded ahead of the server-side JDBC driver used by the Java stored procedure, causing it to malfunction. In a Java stored procedure environment, make sure not to install the generic CUBRID JDBC driver in the directory where the JRE is installed (**\$JAVA_HOME/jre/lib/ext**).

JDBC Sample

The following example shows how to connect to CUBRID by using the JDBC driver, and retrieve and insert data. To run the sample program, make sure that the database you are trying to connect to and the CUBRID Broker are running. In the sample, you will use the **demodb** database that is created automatically during the installation.

JDBC Driver Load

To connect to CUBRID, load the JDBC driver using the `forName()` method provided in the class. For more information, see "API Reference > JDBC API > JDBC Programming > CUBRID JDBC Driver."

```
Class.forName("cuprid.jdbc.driver.CUBRIDDriver");
```

How to Make the Connection to Database

When the JDBC driver is loaded, use the `getConnection()` method provided in the `DriverManager` to connect to the database. To create a `Connection` object, you must specify the url for describing the location of the database, database user name, password, etc. For more information, see the [Connection Configuration](#).

```
String url = "jdbc:cuprid:localhost:33000:demodb::";
String userid = "dba";
String password = "";
```

```
Connection conn = DriverManager.getConnection(url,userid,password);
```

Manipulating Database (Executing Queries and Processing the ResultSet)

To send a query statement to the connected database and execute it, create the **Statement**, **PreparedStatement**, and **CallableStatement** objects. When a statement object has been created, execute the query using the **executeQuery()** method or the **executeUpdate()** method for the statement object. The **next()** method can process the following row from the **ResultSet** that is returned from the **executeQuery()** method. For more information, see [CUBRID JDBC Driver](#).

Note If you execute commit after query execution, **ResultSet** is automatically closed. Therefore, you must not use **ResultSet** after commit. CUBRID is, in general, executed in auto-commit mode. If you does not want auto-commit mode, you must state **conn.setAutoCommit(false);** in the code.

Disconnecting from the Database

Each method can be disconnected from the database by executing the **close()** method.

Example 1

The following example shows how to create a table, execute a query with a prepared statement, and roll back the query. Modify the parameter value of the **getConnection()** method for practice.

```
import java.util.*;
import java.sql.*;

public class Basic {
    public static Connection connect() {
        Connection conn = null;
        try {
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            conn =
DriverManager.getConnection("jdbc:cubrid:localhost:33000:demodb:","dba","");
            conn.setAutoCommit (false) ;
        } catch ( Exception e ) {
            System.err.println("SQLException : " + e.getMessage());
        }
        return conn;
    }

    public static void printdata(ResultSet rs) {
        try {
            ResultSetMetaData rsmd = null;

            rsmd = rs.getMetaData();
            int numberOfColumn = rsmd.getColumnCount();

            while (rs.next ()) {
                for(int j=1; j<=numberOfColumn; j++ )
                    System.out.print(rs.getString(j) + " " );
                System.out.println("");
            }
        } catch ( Exception e ) {
            System.err.println("SQLException : " + e.getMessage());
        }
    }

    public static void main(String[] args) throws Exception {
        Connection conn = null;
        Statement stmt = null;
        ResultSet rs = null;
        PreparedStatement preStmt = null;

        try {
            conn = connect();

            stmt = conn.createStatement();
            stmt.executeUpdate("create class xoo ( a int, b int, c char(10))");

            preStmt = conn.prepareStatement("insert into xoo values(?,?,''100'')");
```



```

        preStmt.setInt (1, 1) ;
        preStmt.setInt (2, 1*10) ;
        int rst = preStmt.executeUpdate () ;

        rs = stmt.executeQuery("select a,b,c from xoo" );

        printdata(rs);

        conn.rollback();
        stmt.close();
        conn.close();
    } catch ( Exception e ) {
        conn.rollback();
        System.err.println("SQLException : " + e.getMessage());
    } finally {
        if ( conn != null ) conn.close();
    }
}
}
}

```

Example 2

The following example shows how to execute the **SELECT** statement by connecting to demodb provided by CUBRID during installation.

```

import java.sql.*;
public class SelectData {
    public static void main(String[] args) throws Exception {
        Connection conn = null;
        Statement stmt = null;
        ResultSet rs = null;
        try {
            // CUBRID0|| Connect
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            conn =
DriverManager.getConnection("jdbc:cubrid:localhost:33000:demodb::","dba","");
            String sql = "select name, players from event";
            stmt = conn.createStatement();
            rs = stmt.executeQuery(sql);
            while(rs.next()) {
                String name = rs.getString("name");
                String players = rs.getString("players");
                System.out.println("name ==> " + name);
                System.out.println("Number of players==> " + players);
                System.out.println("\n=====");
            }
            rs.close();
            stmt.close();
            conn.close();
        } catch ( SQLException e ) {
            System.err.println(e.getMessage());
        } catch ( Exception e ) {
            System.err.println(e.getMessage());
        } finally {
            if ( conn != null ) conn.close();
        }
    }
}
}

```

Example 3

The following example shows how to execute the **INSERT** statement by connecting to demodb provided by CUBRID during installation. You can delete or modify data the same way as you insert data. This means that you can reuse the code below by simply changing the query statements.

```

import java.sql.*;
public class insertData {
    public static void main(String[] args) throws Exception {
        Connection conn = null;
        Statement stmt = null;
        try {

```

```
        // CUBRIDOll Connect
        Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        conn =
DriverManager.getConnection("jdbc:cubrid:localhost:33000:demodb::", "dba", "");
        String sql = "insert into olympic(host_year, host_nation, host_city,
opening_date,          closing_date) values (2008, 'China', 'Beijing', to_date('08-08-
2008', 'mm-dd-yyyy'),          to date('08-24-2008', 'mm-dd-yyyy'))";
        stmt = conn.createStatement();
        stmt.executeUpdate(sql);
        System.out.println("Data is entered.");
        stmt.close();
    } catch ( SQLException e ) {
        System.err.println(e.getMessage());
    } catch ( Exception e ) {
        System.err.println(e.getMessage());
    } finally {
        if ( conn != null ) conn.close();
    }
}
}
```

Programming with PHP

Installing the PHP Module

Installing Required Software

You should install the software below, first:

- Apache Web server : Version 2.0 or later
- PHP : Version 5.2 or later
- CUBRID

Downloading PHP Modules

Download drives from ftp://ftp.cubrid.org/CUBRID_Drivers/PHP_Driver/. For latest information on PHP drives, visit Web site (http://www.cubrid.org/php_api_for_cubrid).

Installing for Windows

- Store the **php_cubrid.dll** file under PHP extension directory (default location : C:\Program Files\PHP\ext\).
- Configure the system configuration variable. The environment variable value of PHPRC should be **C:\Program Files\PHP** and **%PHPRC%** and **%PHPRC%\ext** should be added to the value of Path environment variable.
- Add the following line at the end of the **php.ini** file (default location : C:\Program Files\PHP\php.ini).

```
[PHP_CUBRID] extension=php_cubrid.dll
```

- Once configuration is complete, restart the Web server.

Installing for Linux

- Store the **cubrid.so** file under PHP extensions directory. The file is usually located under /usr/lib/php5/20090626 in PHP 5.3.3; however, the location depends on PHP versions).
- Add the following line at the end of the **php.ini** file (default location : /etc/php5/apache2/php.ini).

```
[CUBRID] extension=cubrid.so
```

- Once configuration is complete, restart the Web server.

Installing Package

Using the Installation Wizard for Windows

You can download CUBRID PHP API Installer from http://www.cubrid.org/php_install_wizard.

Installing by Using PEAR Package for Ubuntu Linux

First, you must set up phpize and PEAR package. Follow the steps below.

- Install Apache Web server and PHP. For information on installation of Apache Web server and PHP, visit the Web site (http://www.cubrid.org/cubrid_apache_php_ubuntu).
- Execute the following line to set up phpize.

```
sudo apt-get install php5-dev
```

- Execute the following line to install PEAR package.

```
sudo apt-get install php-pear
```

- Execute the pecl command in PEAR package so that the latest CUBRID PHP extension can be installed.

```
sudo pecl install cubrid
```

- To install the earlier version of PHP drivers, you must specify the desired version.

```
sudo pecl install cubrid-8.3.0.0005
```

- Add the following line at the end of the php.ini file (default location : /etc/php5/apache2/php.ini).

```
[CUBRID] extension=cubrid.so
```

- Once configuration is complete, restart the Web server.

Installing by Using PEAR Package for Linux Except for Version Ubuntu

First, you must install phpzie and PEAR package. Follow the steps below.

- Set up phpzie. The version php-dev must be PHP 5.2.x or PHP 5.3.x. If you are using the earlier version of PHP, update PHP before proceeding with installation.

```
yum install php-devel
```

- Download PEAR package.

```
wget http://pear.php.net/go-pear.phar
```

- Install PEAR package.

```
php go-pear.phar
```

- Install the latest versions of CUBRID-PHP Extension.

```
pecl install cubrid
```

- To install the earlier version of PHP drivers, you must specify the desired version.

```
sudo pecl install cubrid-8.3.0.0005
```

- Add the following line at the end of the php.ini file (default location : /etc/php5/apache2/php.ini).

```
[CUBRID] extension=cubrid.so
```

- Once configuration is complete, restart the Web server.

Note If you use the phpinfo() function to create test.php and verify that CUBRID configuration page is properly displayed after entering 오류! 하이퍼링크 참조가 잘못되었습니다.<IP address in which Web server is installed>/test.php, it means installation is successfully done.

PHP Sample

The following is a simple example that establishes a connection between PHP and CUBRID. This section will cover the most basic and notable features. Before running the sample program, a database and the Broker you are trying to connect must be running. This example uses the **demodb** database created during the installation.

Example of Data Retrieval

```
<html>
<head>
<meta http-equiv='content-type' content='text/html; charset=euc-kr'
</head>
<body>
<center>
<table border=2>
<?
// Set server information for CUBRID connection. host_ip is the IP address where the
CUBRID Broker is installed (localhost in this example), and host_port is the port number
of the CUBRID Broker. The port number is the default given during the installation. For
details, see "Administrator's Guide."
$host_ip = "localhost";
$host_port = 30000;
$db_name = "demodb";
// Connect to CUBRID Server. Do not make the actual connection, but only retain the
connection information. The reason for not making the actual connection is to handle
transaction more efficiently in the 3-tier architecture.
$cubrid_con = @cubrid_connect($host_ip, $host_port, $db_name);
if (!$cubrid_con) {
    echo "Database Connection Error";
    exit;
}
?>
<?>
```

```

$sql = "select sports, count(players) as players from event group by sports";
// Request the CUBRID Server for the results of the SQL statement. Now make the actual
connection to the CUBRID Server.
$result = cubrid_execute($cubrid_con, $sql);
if ($result) {
    // Get the column names from the result set created by the SQL query.
    $columns = cubrid_column_names($result);
    // Get the number of columns in the result set created by the SQL query.
    $num_fields = cubrid_num_cols($result);
    // List the column names of the result set on the screen.
    echo("<tr>");
    while (list($key, $colname) = each($columns)) {
        echo("<td align=center>$colname</td>");
    }
    echo("</tr>");
    // Get the results from the result set.
    while ($row = cubrid_fetch($result)) {
        echo("<tr>");
        for ($i = 0; $i < $num_fields; $i++) {
            echo("<td align=center>");
            echo($row[$i]);
            echo("</td>");
        }
        echo("</tr>");
    }
}
// The PHP module in the CUBRID runs in a 3-tier architecture. Even when calling SELECT
for transaction processing, it is processed as a part of the transaction. Therefore, the
transaction needs to be rolled back by calling commit or rollback even though SELECT was
called for smooth performance.
cubrid_commit($cubrid_con);
cubrid_disconnect($cubrid_con);
?>
</body></html>

```

Example of Data Insertion

```

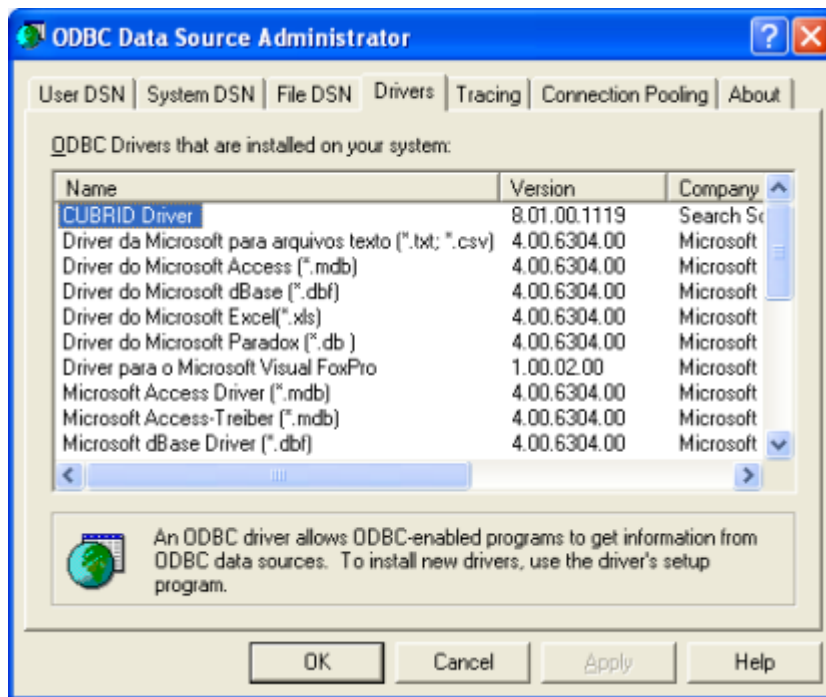
<html>
<head>
<meta http-equiv='content-type' content='text/html; charset=euc-kr'>
</head>
<body>
<center>
<table border=2>
<?
    $host_ip = "localhost";
    $host_port = 30000;
    $db_name = "demodb";
    $cubrid_con = @cubrid_connect($host_ip, $host_port, $db_name);
    if (!$cubrid_con) {
        echo "Database Connection Error";
        exit;
    }
?>
<?
    $sql = "insert into olympic (host year,host nation,host city,opening date,closing date)
values      (2008, 'China', 'Beijing', to date('08-08-2008','mm-dd-yyyy'),to date('08-24-
2008','mm-dd-yyyy'))      ";
    $result = cubrid_execute($cubrid_con, $sql);
    if ($result) {
        // Handled successfully, so commit.
        cubrid_commit($cubrid_con);
        echo("Inserted successfully ");
    } else {
        // Error occurred, so the error message is output and rollback is called.
        echo(cubrid_error_msg());
        cubrid_commit($cubrid_con);
    }
    cubrid_disconnect($cubrid_con);
?>
</body></html>

```

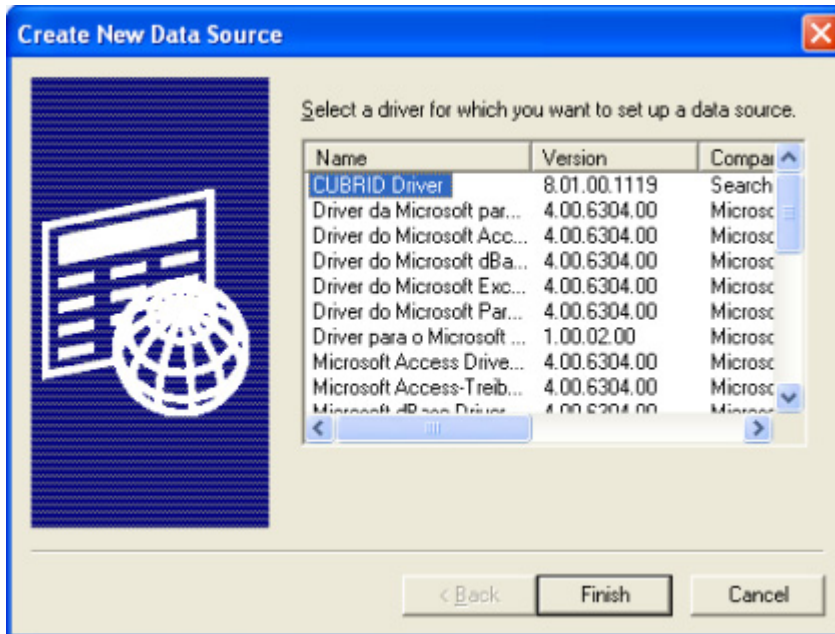
Programming with ODBC and ASP

Configuring the Environment of ODBC and ASP

CUBRID ODBC is compatible for version 3.52 ODBC and LEVEL2. Note that backward compatibility is not guaranteed for applications that are written with ODBC Spec 2.x. The CUBRID ODBC driver is automatically installed while CUBRID is installed. You can verify it from [Control Panel] > [Administrative Tools] > [Data Source (ODBC)] > [Drivers] tab.

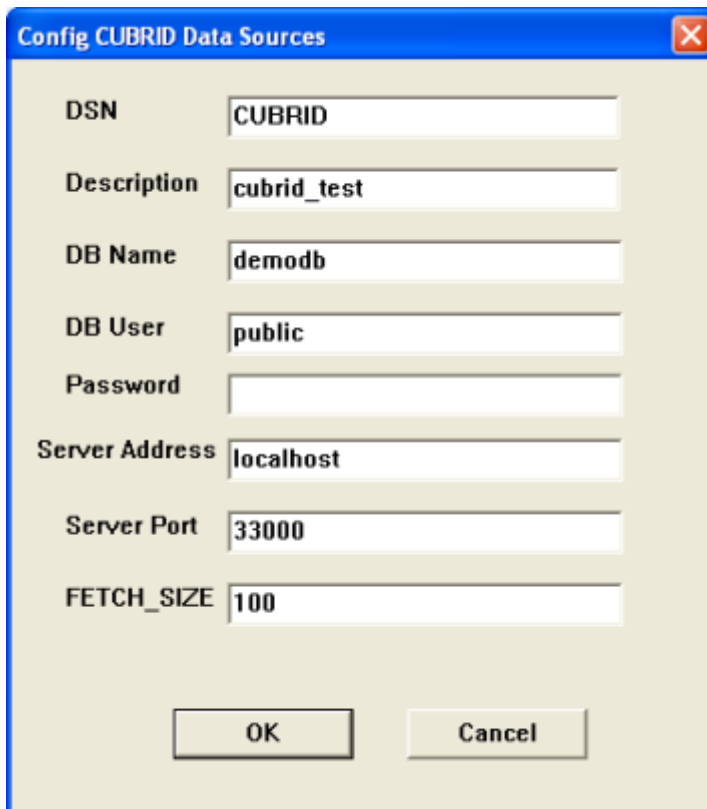


If the CUBRID ODBC driver is detected, set a DSN as a database where the application is trying to connect. To set up a DSN, click the [Add] button in the ODBC Data Source Administrator dialog box. Then, the following dialog box appears. Select "CUBRID Driver," and then click the [Finish] button.



When the following [Config CUBRID Data Sources] dialog box appears, enter the database name that you try to connect to in the [DB Name] field, the port number of the CUBRID Broker in the [Server Port] field, and then click [OK] button. You can verify the number in the **cubrid.broker.conf** file.

FETCH_SIZE refers to the number of records fetched from server whenever **cci_fetch()** function of CCI library is called; the CCI library is internally used by ODBC driver.



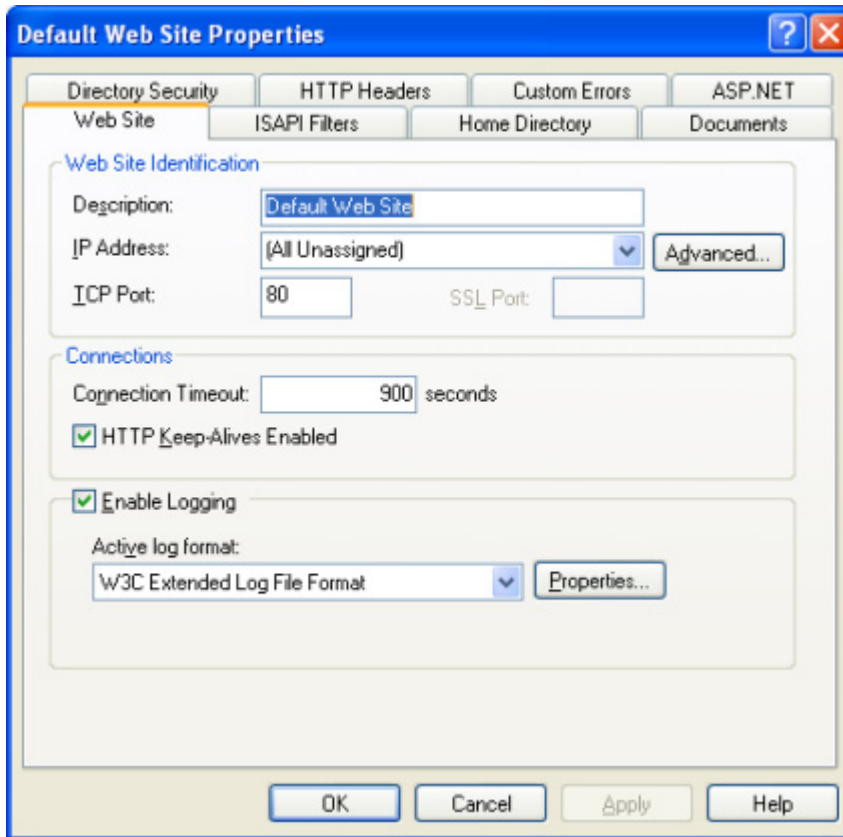
For more information on CUBRID ODBC driver, see "API Reference > ODBC API."

- [CUBRID ODBC Driver](#)

- [Using OIDs and Collections](#)
- [Supported Functions and Backward Compatibility](#)

ASP Sample

In the virtual directory where the ASP sample program runs, right-click "Default Web Site" and click [Properties].



The dialog box shown above will appear. Under the **Web Site Identification**, in the **IP Address** drop-down box, select "(All Unassigned)." This sets the IP address to localhost. If you want to run the sample program using a specific IP address, configure the directory with the IP address as a virtual directory and register the IP address in Properties.

The following example shows how to configure IP address as localhost.

Example

The following example shows how to create cubrid.asp and store it into a virtual directory.

```
<HTML>
  <HEAD>
    <meta http-equiv="Content-Type" content="text/html; charset=EUC-KR">
    <title>CUBRID Query Test Page</title>
  </HEAD>
  <BODY topmargin="0" leftmargin="0">

  <table border="0" width="748" cellspacing="0" cellpadding="0">
    <tr>
      <td width="200"></td>
      <td width="287">
        <p align="center"><font size="3" face="Times New Roman"><b><font
color="#FF0000">CUBRID</font>Query Test</b></font></td>
      <td width="200"></td>
    </tr>
  </table>
  <form action="cubrid.asp" method="post" >
  <table border="1" width="700" cellspacing="0" cellpadding="0" height="45">
```



```

<tr>
  <td width="113" valign="bottom" height="16" bgcolor="#DBD7BD"
bordercolorlight="#FFFFFF"><font size="2">SERVER IP</font></td>
  <td width="78" valign="bottom" height="16" bgcolor="#DBD7BD"
bordercolorlight="#FFFFFF"><font size="2">Broker PORT</font></td>
  <td width="148" valign="bottom" height="16" bgcolor="#DBD7BD"
bordercolorlight="#FFFFFF"><font size="2">DB NAME</font></td>
  <td width="113" valign="bottom" height="16" bgcolor="#DBD7BD"
bordercolorlight="#FFFFFF"><font size="2">DB USER</font></td>
  <td width="113" valign="bottom" height="16" bgcolor="#DBD7BD"
bordercolorlight="#FFFFFF"><font size="2">DB PASS</font></td>
  <td width="80" height="37" rowspan="4" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED">
    <p><input type="submit" value="Execute" name="B1" tabindex="7"></p></td>
</tr>
<tr>
  <td width="113" height="1" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED"><font
size="2"><input type="text" name="server ip" size="20" tabindex="1" maxlength="15"
value="<%=Request("server ip")%>"></font></td>
  <td width="78" height="1" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED"><font
size="2"><input type="text" name="cas_port" size="15" tabindex="2" maxlength="6"
value="<%=Request("cas_port")%>"></font></td>
  <td width="148" height="1" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED"><font
size="2"><input type="text" name="db name" size="20" tabindex="3" maxlength="20"
value="<%=Request("db name")%>"></font></td>
  <td width="113" height="1" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED"><font
size="2"><input type="text" name="db_user" size="15" tabindex="4"
value="<%=Request("db_user")%>"></font></td>
  <td width="113" height="1" bordercolorlight="#FFFFFF" bgcolor="#F5F5ED"><font
size="2"><input type="password" name="db pass" size="15" tabindex="5"
value="<%=Request("db pass")%>"></font></td>
</tr>
<tr>
  <td width="573" colspan="5" valign="bottom" height="18" bordercolorlight="#FFFFFF"
bgcolor="#DBD7BD"><font size="2">QUERY</font></td>
</tr>
<tr>
  <td width="573" colspan="5" height="25" bordercolorlight="#FFFFFF"
bgcolor="#F5F5ED"><textarea rows="3" name="query" cols="92"
tabindex="6"><%=Request("query")%></textarea></td>
</tr>
</table>
</form>
<hr>
</BODY>
</HTML>
<%
  ' Fetch the DSN and SQL statement.
  strIP = Request( "server ip" )
  strPort = Request( "cas_port" )
  strUser = Request( "db_user" )
  strPass = Request( "db_pass" )
  strName = Request( "db name" )
  strQuery = Request( "query" )

if strIP = "" then
Response.Write "Please enter the SERVER_IP"
  Response.End 'If no IP entered, end the page
end if
if strPort = "" then
  Response.Write "Please enter the port number"
  Response.End ' If no port entered, end the page
end if
if strUser = "" then
  Response.Write "Please enter the DB USER"
  Response.End ' If no DB User entered, end the page
end if
if strName = "" then
  Response.Write "Please enter the DB NAME "
  Response.End ' If no DB NAME entered, end the page
end if
if strQuery = "" then
  Response.Write "Please enter the query you want to check"
  Response.End ' If no Query entered, end the page

```

```

    end if
    ' Create the connection object
    strDsn = "driver={CUBRID Driver};server=" & strIP & ";port=" & strPort & ";uid=" &
strUser & ";pwd=" & strPass & ";db_name=" & strName & ";";
    ' Connect to DB
Set DBConn = Server.CreateObject("ADODB.Connection")
    DBConn.Open strDsn
    ' Execute SQL
Set rs = DBConn.Execute( strQuery )
    ' Show message depending on the SQL statement
    if InStr(Ucase(strQuery),"INSERT")>0 then
        Response.Write "The record has been added."
        Response.End
    end if

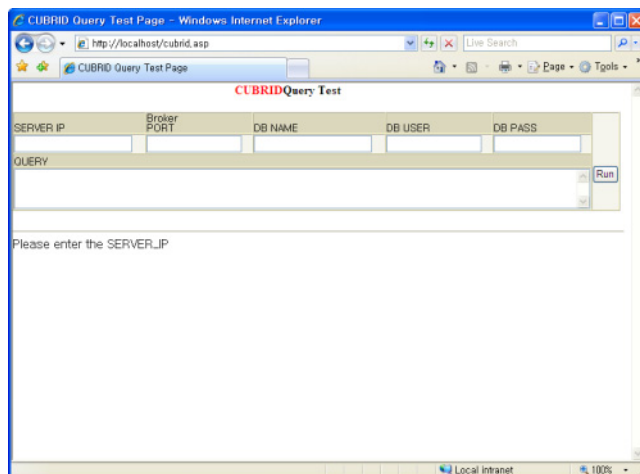
    if InStr(Ucase(strQuery),"DELETE")>0 then
        Response.Write "The record has been deleted."
        Response.End
    end if

    if InStr(Ucase(strQuery),"UPDATE")>0 then
        Response.Write "The record has been modified."
        Response.End
    end if
%>
<table>
<%
    ' Show the field name
    Response.Write "<tr bgColor=#f3f3f3>"
    For index = 0 to ( rs.fields.count-1 )
        Response.Write "<td><b>" & rs.fields(index).name & "</b></td>"
    Next
    Response.Write "</tr>"
    ' Show the field value
    Do While Not rs.EOF
        Response.Write "<tr bgColor=#f3f3f3>"
        For index = 0 to ( rs.fields.count-1 )
            Response.Write "<td>" & rs(index) & "</td>"
        Next
        Response.Write "</tr>"

        rs.MoveNext
    Loop
%>
<%
    set rs = nothing
%>
</table>

```


You can check the result of the sample program at <http://localhost/aSP/cubrid.asp>. When you execute the sample code above, you will get the following output. Enter appropriate values in each field, and then enter the query statement in the Query field. When you click [Run], the query result will be displayed at the lower portion of the page.



Programming with CCI

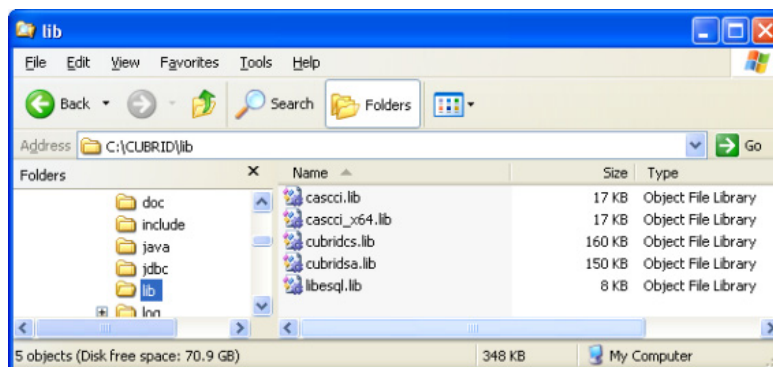
CCI Library

The CCI Library is a C language interface provided by CUBRID. CCI is connected to the application through the Broker, so you can manage it the same way as other interfaces such as JDBC, PHP and ODBC. In fact, CCI provides a foundation to implement PHP, ODBC, Python and, Ruby interfaces.



CCI Installation and Configuration

The CCI library is contained in the CUBRID installation package. The following figure shows where the files are located.



Operating System	Windows	UNIX/Linux
C header file	include/cas_cci.h	include/cas_cci.h
Static library	lib/cascci.lib	lib/libcascci.a
Dynamic library	lib/cascci.lib bin/cascci.dll	lib/libcascci.so

Using CCI

Basic Flow Diagram of the Application Using CCI

To use CUBRID, the following procedures are required for applications using the CCI libraries to execute queries: connection to CAS, query preparation, query execution, response handling, and disconnection. In each process, CCI communicates with the application using connection, query and response handles.

The following flowchart shows the process of the application using CCI and the functions used in each step. See CCI API in the API Reference for more information.

- Opening a database connection handle (related function : [cci_connect](#))
- Getting the request handle for a prepared statement (related function : [cci_prepare](#))
- Binding data to the prepared statement (related function : [cci_bind_param](#))
- Executing the prepared statement (related function : [cci_execute](#))
- Processing the execution result (related function : [cci_cursor](#), [cci_fetch](#), [cci_get_data](#), [cci_get_result_info](#))
- Closing the request handle (related function : [cci_close_req_handle](#))
- Closing a database connection handle (related function : [cci_disconnect](#))

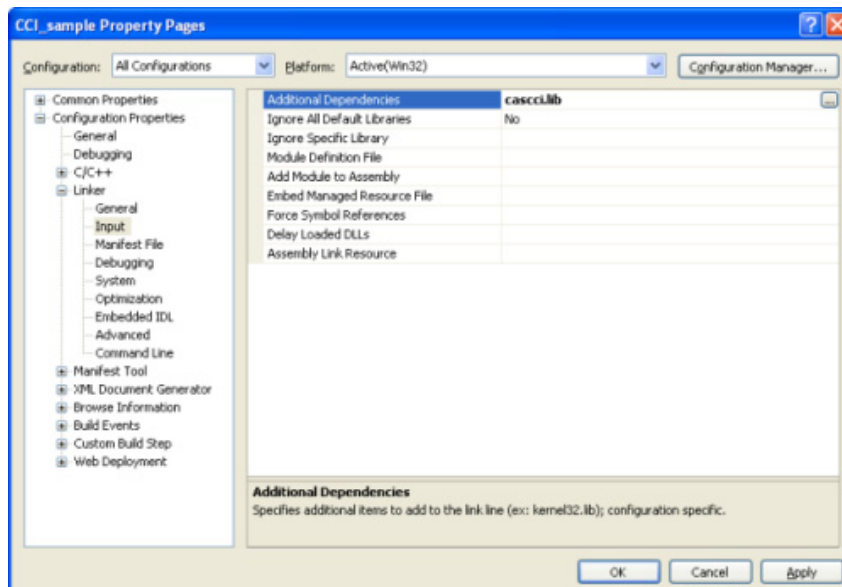
How to use

Once you have created the application using CCI, you should decide, according to its features, whether to execute CCI as a static link or dynamic link before you build it. Determine the library to use by referring to the table in the [CCI Installation and Configuration](#).

The following example shows Makefile to use the dynamic link library on UNIX/Linux:

```
CC=gcc
CFLAGS = -g -Wall -I. -I${CUBRID}/include
LDFLAGS = -L${CUBRID}/lib -lcascci -lnsl
TEST_OBJS = test.o
EXES = test
all: $(EXES)
test: $(TEST_OBJS)
$(CC) -o $@ $(TEST_OBJS) $(LDFLAGS)
```

The following is the settings for using the static library on Windows:



CCI Sample

Introduction

The sample program is to create a simple application using CCI through the connection to the **demodb** database deployed by default during the CUBRID installation. Follow the processes of connection to CAS, query preparation, query execution, response handling and disconnection in the sample. The sample is created in a way that uses dynamic links on Linux.

The following is schema information of the **olympic** table in the **demodb** database used in the sample.

```

csql> ;sc olympic

=== <Help: Schema of a Class> ===

<Class Name>

    olympic

<Attributes>

    host_year          INTEGER NOT NULL
    host_nation        CHARACTER VARYING(40) NOT NULL
    host_city           CHARACTER VARYING(20) NOT NULL
    opening_date        DATE NOT NULL
    closing_date        DATE NOT NULL
    mascot              CHARACTER VARYING(20)
    slogan              CHARACTER VARYING(40)
    introduction        CHARACTER VARYING(1500)

<Constraints>

    PRIMARY KEY pk_olympic_host_year ON olympic (host_year)

```

Preparation

Make sure that the **demodb** database and the Broker are running before you execute the sample program. You can start the **demodb** database and the Broker by executing the **cuprid** utility.

The following example shows how to run a database server and broker by executing the **cuprid** utility.

```

[tester@testdb ~]$ cuprid server start demodb
@ cuprid master start
++ cuprid master start: success
@ cuprid server start: demodb

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.0

++ cuprid server start: success
[tester@testdb ~]$ cuprid broker start
@ cuprid broker start
++ cuprid broker start: success

```

Build

With the program source and the Makefile ready, executing "make" will create an executable file called "test." If you use a static library, there is no need to deploy additional files and the execution will be faster. However, it increases the program size and memory usage. If you use a dynamic library, there will be some performance overhead, but the program size and memory usage can be optimized.

The following is a command line example. It builds the test program using the dynamic library instead of "make" on Linux.

```

cc -o test test.c -I$CUBRID/include -L$CUBRID/lib -lnsl -lcascci

```

Sample Code

```

#include <stdio.h>
#include <cas cci.h>
char *cci_client_name = "test";
int main (int argc, char *argv[])
{
    int con = 0, req = 0, col_count = 0, res, ind, i;
    T CCI_ERROR error;
    T CCI_COL_INFO *res_col_info;
    T CCI_SQLX_CMD cmd_type;
    char *buffer, db_ver[16];
    printf("Program started!\n");
    if ((con=cci_connect("localhost", 30000, "demodb", "PUBLIC", ""))<0) {
        printf( "%s(%d): cci connect fail\n", FILE , LINE );
        return -1;
    }

    if ((res=cci_get_db_version(con, db_ver, sizeof(db_ver)))<0) {
        printf( "%s(%d): cci get db version fail\n", FILE , LINE );
        goto handle_error;
    }
    printf("DB Version is %s\n",db_ver);
    if ((req=cci_prepare(con, "select * from event", 0,&error))<0) {
        printf( "%s(%d): cci prepare fail(%d)\n", FILE , LINE ,error.err code);
        goto handle_error;
    }
    printf("Prepare ok!(%d)\n",req);
    res_col_info = cci_get_result_info(req, &cmd_type, &col_count);
    if (!res_col_info) {
        printf( "%s(%d): cci get result info fail\n", FILE , LINE );
        goto handle_error;
    }

    printf("Result column information\n"
           "=====\n");
    for (i=1; i<=col_count; i++) {
        printf("name:%s type:%d(precision:%d scale:%d)\n",
              CCI_GET_RESULT_INFO_NAME(res_col_info, i),
              CCI_GET_RESULT_INFO_TYPE(res_col_info, i),
              CCI_GET_RESULT_INFO_PRECISION(res_col_info, i),
              CCI_GET_RESULT_INFO_SCALE(res_col_info, i));
    }
    printf("=====\n");
    if ((res=cci_execute(req, 0, 0, &error))<0) {
        printf( "%s(%d): cci execute fail(%d)\n", FILE , LINE ,error.err code);
        goto handle_error;
    }
    if ((res=cci_fetch_size(req, 100))<0) {
        printf( "%s(%d): cci fetch size fail\n", FILE , LINE );
        goto handle_error;
    }

    while (1) {
        res = cci_cursor(req, 1, CCI_CURSOR_CURRENT, &error);
        if (res == CCI_ER_NO_MORE_DATA) {
            printf("Query END!\n");
            break;
        }
        if (res<0) {
            printf( "%s(%d): cci cursor fail(%d)\n", FILE , LINE ,error.err code);
            goto handle_error;
        }

        if ((res=cci_fetch(req, &error))<0) {
            printf( "%s(%d): cci fetch fail(%d)\n", FILE , LINE ,error.err code);
            goto handle_error;
        }

        for (i=1; i<=col_count; i++) {
            if ((res=cci_get_data(req, i, CCI_A_TYPE_STR, &buffer, &ind))<0) {
                printf( "%s(%d): cci get data fail\n", FILE , LINE );
                goto handle_error;
            }
        }
    }
}

```

```
        }
        printf("%s \t|", buffer);
    }
    printf("\n");
}
if ((res=cci_close_req_handle(req))<0) {
    printf( "%s(%d): cci close req handle fail",  FILE  ,  LINE  );
    goto handle_error;
}
if ((res=cci_disconnect(con, &error)<0) {
    printf( "%s(%d): cci disconnect fail(%d)",  FILE  ,  LINE  ,error.err code);
    goto handle_error;
}
printf("Program ended!\n");
return 0;

handle_error:
if (req > 0)
    cci_close_req_handle(req);
if (con > 0)
    cci_disconnect(con, &error);
printf("Program failed!\n");
return -1;
}
```


CSQL Interpreter

To execute SQL statements in CUBRID, you need to use either a Graphical User Interface (GUI)-based CUBRID Manager or a console-based CSQL Interpreter.

CSQL is an application that allows users to use SQL statements through a command-driven interface. This section briefly explains how to use the CSQL Interpreter and associated commands.

- Introduction to the CSQL Interpreter
- Running CSQL
- Session Commands

Introduction to the CSQL Interpreter

A Tool for SQL

The CSQL Interpreter is an application installed with CUBRID that allows you to execute in an interactive or batch mode and viewing query results. The CSQL Interpreter has a command-line interface. With this, you can store SQL statements together with their results to a file for a later use.

The CSQL Interpreter provides the best and easiest way to use CUBRID. You can develop database applications with various APIs (e.g. JDBC, ODBC, PHP, CCI, etc.); you can use the CUBRID Manager, which is a management and query tool provided by CUBRID. With the CSQL Interpreter, users can create and retrieve data in a terminal-based environment.

The CSQL Interpreter directly connects to a CUBRID database and executes various tasks using SQL statements. Using the CSQL Interpreter, you can:

- Retrieve, update and delete data in a database by using SQL statements
- Execute external shell commands
- Save or print query results
- Create and execute SQL script files
- Select table schema
- Retrieve or modify parameters of the database server system
- Retrieve database information (e.g. schema, triggers, queued triggers, workspaces, locks, and statistics)

A Tool for DBA

A database administrator (**DBA**) performs administrative tasks by using various administrative utilities provided by CUBRID; a terminal-based interface of CSQL Interpreter is an environment where **DBA** executes administrative tasks.

It is also possible to run the CSQL Interpreter in a standalone mode. In this mode, the CSQL Interpreter directly accesses database files and executes commands including server process properties. That is, SQL statements can be executed to a database without running a separate database server process. The CSQL Interpreter is a powerful tool that allows you to use the database only with a **csql** utility, without any other applications such as the Database Server or the Brokers.

Executing CSQL

CSQL Execution Mode

Interactive Mode

With CSQL Interpreter, you can enter and execute SQL statements to handle schema and data in the database. Enter statements in a prompt that appears when running the **csql** utility. After executing the statements, the results are listed in the next line. This is called the interactive mode.

Batch Mode

You can store SQL statements in a file and execute them later to have the **csql** utility read the file. This is called the batch mode. For more information on the batch mode, see [CSQL Startup Options](#).

Standalone Mode

In the standalone mode, CSQL Interpreter directly accesses database files and executes commands including server process functions. That is, SQL statements can be sent and executed to a database without a separate database server process running for the task. Since the standalone mode allows only one user access at a given time, it is suitable for management tasks by Database Administrators (DBAs).

Client/Server Mode

CSQL Interpreter usually operates as a client process and accesses the server process.

Using CSQL (Syntax)

Connecting to Local Host

Description

Execute the CSQL Interpreter using the **csql** utility. You can set options as needed. To set the options, specify the name of the database to connect to as a parameter. The following is a **csql** utility statement to access the database on a local server:

Syntax

```
csql [ options ] database_name
```

Connecting to Remote Host

Description

The following is a **csql** utility statement to access the database on a remote host:

Syntax

```
csql [ options ] database_name@remote_host_name
```

Make sure that the following conditions are met before you run the CSQL Interpreter on a remote host.

- The CUBRID installed on the remote host must be the same version as the one on the local host.
- The port number used by the master process on the remote host must be identical to the one on the local host.
- You must access the remote host in a client/server mode using the **-C** option.

Example

The following example shows how to access the **demodb** database on the remote host with the IP address 192.168.1.3 and calls the **csql** utility.

```
csql -C demodb@192.168.1.3
```

CSQL Startup Options

To display the option list in the prompt, execute the **csql** utility without specifying the database name as follows:

```
% csql
interactive SQL utility, version 2008 R4.1
usage: csql [OPTION] database-name valid options:
  -S, --SA-mode           standalone mode execution
  -C, --CS-mode           client-server mode execution
  -u, --user=ARG          alternate user name
  -p, --password=ARG      password string, give "" for none
  -e, --error-continue   don't exit on statement error
  -i, --input-file=ARG    input-file-name
  -o, --output-file=ARG   output-file-name
  -s, --single-line       single line oriented execution
  -c, --command=ARG       CSQL-commands
  -l, --line-output       display each value in a line
  -r, --read-only         read-only mode
  --no-auto-commit        disable auto commit mode execution
  --no-pager              do not use pager
  --no-single-line        turn off single line oriented execution
```

Options

The following table lists the options that can be issued with the **csql** utility.

Option	Description
-S	Executes the csql utility in a standalone mode.
-C	Executes the csql utility in a client/server mode.
-u <i>user_name</i>	Specifies the user that tries to access the database. The default value is PUBLIC .
-p <i>password</i>	Specifies the password of the user that tries to access the database (if any).
-e	Continues the session even when an error occurs.
-i <i>input_file</i>	Executes the csql utility in a batch mode. The <i>input_file</i> parameter is the file name where SQL statements are stored.
-o <i>output_file</i>	Stores a result of the statement execution in the specified <i>output_file</i> without displaying it on the screen.
-s	It is used with the -i option and used to execute multiple SQL statements one by one in a file where they are stored consecutively. Use semicolons (;) to separate SQL statements.
-c "CSQL commands"	Executes SQL statements directly from the prompt. To use this option, enclose the SQL statement to execute in double quotes.
-l	Displays the query results in a line format instead of a column. By default, the results will be displayed in a column format.
-r	Connects to a database in read-only mode.
--no-auto-commit	Configures the auto-commit mode of the CSQL Interpreter to OFF.
--no-pager	Displays the results of the query performed by the CSQL Interpreter at once instead of page-by-page.
--no-single-line	Executes multiple SQL statements at once by using ;xr or ;r session command.

Executing in a standalone mode (-S)

The following example shows how to connect to a database in a standalone mode and execute the **csql** utility. If you want to use the database exclusively, use the **-S** option.

```
csql -S demodb
```

Executing in a client/server mode (-C)

The following example shows how to connect to a database in a client/server mode and execute the **csql** utility. In an environment where multiple clients connect to the database, use the **-C** option. Even when you connect to a database on a remote host in a client/server mode, the error log created during **csql** execution will be stored in the **cub_client.err** file on the local host.

```
csql -C demodb
```

Specifying the name of the input file to use in a batch mode (-i)

The following example shows how to specify the name of the input file that will be used in a batch mode with the **-i** option. In the **infile** file, more than one SQL statement are stored. Without the **-i** option specified, the CSQL Interpreter will run in an interactive mode.

```
csql -i infile demodb
```

Specifying the output file to store the execution results (-o)

The following example shows how to store the execution results to the specified file instead of displaying on the screen. It is useful to retrieve the results of the query performed by the CSQL Interpreter afterwards.

```
csql -o outfile demodb
```

Specifying the user name (-u)

The following example shows how to specify the name of the user that will connect to the specified database with the **-u** option. If the **-u** option is not specified, **PUBLIC** that has the lowest level of authorization will be specified as a user. If the user name is not valid, an error message is displayed and the **csql** utility is terminated. If there is a password for the user name you specify, you will be prompted to enter the password.

```
csql -u DBA demodb
```

Specifying the user password (-p)

The following example shows how to enter the password of the user specified with the **-p** option. Especially since there is no prompt to enter a password for the user you specify in a batch mode, you must enter the password using the **-p** option. When you enter an incorrect password, an error message is displayed and the **csql** utility is terminated.

```
csql -u DBA -p *** demodb
```

Executing SQL statements one by one (-s)

As an option used with the **-i** option, it executes multiple SQL statement one by one in a file with the **-s** option. This option is useful to allocate less memory for query execution and each SQL statement is separated by semicolons (;). If it is not specified, multiple SQL statements are retrieved and executed at once.

```
csql -s -i infile demodb
```

Executing SQL statements directly from the shell (-c)

The following example shows how to execute more than one SQL statement from the shell with the **-c** option. Multiple statements are separated by semicolons (;).

```
csql -c "select * from olympic;select * from stadium" demodb
```

Displaying the results in a line format (-l)

The following example shows how to display the execution results of the SQL statement in a line format with the **-l** option. The execution results will be output in a column format if the **-l** option is not specified.

```
csql -l demodb
```

Ignoring errors and keepgoing execution (-e)

The following example shows how to ignore errors and keep execution even though semantic or runtime errors occur with the **-e** option. However, if any SQL statements have syntax errors, query execution stops after errors occur despite specifying the **-e** option.

```
$ csql -e demodb
csql> SELECT * FROM aaa;SELECT * FROM athlete WHERE code=10000;
In line 1, column 10,
ERROR: Unknown class "aaa".

=== <Result of SELECT Command in Line 1> ===
      code name                gender          nation code
      event
=====
10000 'aaa'                    'M'          'NED'      'Rowing'

1 rows selected.
Current transaction has been committed.
1 command(s) successfully processed.
```

Connecting to a read-only database (-r)

You can connect to a read-only database with the **-r** option. Retrieving data is only allowed in a read-only database; creating databases and entering data are not allowed.

```
$ csql -r demodb
```

No auto-commit mode (--no-auto-commit)

The following example shows how to stop the auto-commit mode with the **--no-auto-commit** option. If you don't configure **--no-auto-commit** option, the CSQL Interpreter runs in an auto-commit mode by default, and the SQL statement is committed automatically at every execution. Executing the **;Autocommit** session command after starting the CSQL Interpreter will also have the same result.

```
csql --no-auto-commit demodb
```

Displaying all the execution results at once (--no-pager)

The following example shows how to display the execution results by the CSQL Interpreter at once instead of page-by-page with the **--no-pager** option. The results will be output page-by-page if **--no-pager** option is not specified.

```
csql --no-pager demodb
```

Executing all SQL statements at once (--no single-line)

The following example shows how to keep storing multiple SQL statements and execute them at once with the **;xr** or **;r** session command. If you do not specify this option, SQL statements are executed without **;xr** or **;r** session command.

```
csql --no-single-line demodb
```

Session Commands

In addition to SQL statements, CSQL Interpreter provides special commands allowing you to control the Interpreter. These commands are called session commands. All the session commands must start with a semicolon (;).

Session Commands

Enter the **;help** command to display a list of the session commands available in the CSQL Interpreter. Note that only the uppercase letters of each session command are required to make the CSQL Interpreter to recognize it. Session commands are not case-sensitive.

```
CUBRID SQL Interpreter
Type `;help' for help messages.
csql> ;help
=== <Help: Session Command Summary> ===
All session commands should be prefixed by `;' and only blanks/tabs
can precede the prefix. Capitalized characters represent the minimum
abbreviation that you need to enter to execute the specified command.
;REAd [<file-name>] - read a file into command buffer.
;WRite [<file-name>] - (over)write command buffer into a file.
;APpend [<file-name>] - append command buffer into a file.
;PRINt - print command buffer.
;SHELL - invoke shell.
;CD - change current working directory.
;EXit - exit program.
;CLear - clear command buffer.
;EDIT - invoke system editor with command buffer.
;LIst - display the content of command buffer.
;RUrun - execute sql in command buffer.
;XRrun - execute sql in command buffer, and clear the command
buffer.
;COMmit - commit the current transaction.
;ROllback - roll back the current transaction.
;AUtocommit [ON|OFF] - enable/disable auto commit mode.
;CHeckpoint - issue checkpoint.
;KILLtran - kill transaction.
;REStart - restart database.
;SHELL Cmd [shell-cmd] - set default shell, editor, print and pager
;EDITOR Cmd [editor-cmd] command to new one, or display the current
;PRINT_Cmd [print-cmd] one, respectively.
;PAger_cmd [pager-cmd]
;DATE - display the local time, date.
;DATAbase - display the name of database being accessed.
;SCHEMA class-name - display schema information of a class.
;SYntax [sql-cmd-name] - display syntax of a command.
;TRigger [*'|trigger-name] - display trigger definition.
;Get system parameter - get the value of a system parameter.
;SEt system parameter=value - set the value of a system parameter.
;PLan [simple/detail/off] - show query execution plan.
;INfo <command> - display internal information.
;TIme [ON|OFF] - enable/disable to display the query execution time.
;HISTORYList - display list of the executed queries.
;HISTORYRead <history_num> - read entry on the history number into command buffer.
;HELp - display this help message.
csql>
```

Options

Reading SQL statements from a file (;REAd)

The **;REAd** command reads the contents of a file into the buffer. This command is used to execute SQL commands stored in the specified file. To view the contents of the file loaded into the buffer, use the **;LIst** command.

```
csql> ;rea nation.sql
The file has been read into the command buffer.
csql> ;list
insert into "sport_event" ("event_code", "event_name", "gender_type", "num_player") values
```

```
(20001, 'Archery Individual', 'M', 1);
insert into "sport event" ("event code", "event name", "gender type", "num player") values
20002, 'Archery Individual', 'W', 1);
....
```

Saving SQL statements into a file (;Write)

The **;Write** command stores the contents of the command buffer into a file. This command is used to store SQL commands that you entered or modified in the CSQL Interpreter.

```
csql> ;w outfile
Command buffer has been stored.
```

Appending to a file (;Append)

This command appends the contents of the current command buffer to an **outfile** file.

```
csql> ;ap outfile
Command buffer has been stored.
```

Executing a shell command (;SHELL)

The **;SHELL** session command calls an external shell. Starts a new shell in the environment where the CSQL Interpreter is running. It returns to the CSQL Interpreter when the shell terminates. If the shell command to execute with the **;SHELL_Cmd** command has been specified, it starts the shell, executes the specified command, and returns to the CSQL Interpreter.

```
csql> ;shell
% ls -al
total 2088
drwxr-xr-x 16 DBA cubrid 4096 Jul 29 16:51 .
drwxr-xr-x 6 DBA cubrid 4096 Jul 29 16:17 ..
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 02:49 audit
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 16:17 bin
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 16:17 conf
drwxr-xr-x 4 DBA cubrid 4096 Jul 29 16:14 cubridmanager
% exit
csql>
```

Registering a shell command (;SHELL_Cmd)

The **;SHELL_Cmd** command registers a shell command to execute with the **SHELL** session command. As shown in the example below, enter the **;shell** command to execute the registered command.

```
csql> ;shell_c ls -la
csql> ;shell
total 2088
drwxr-xr-x 16 DBA cubrid 4096 Jul 29 16:51 .
drwxr-xr-x 6 DBA cubrid 4096 Jul 29 16:17 ..
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 02:49 audit
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 16:17 bin
drwxr-xr-x 2 DBA cubrid 4096 Jul 29 16:17 conf
drwxr-xr-x 4 DBA cubrid 4096 Jul 29 16:14 cubridmanager
csql>
```

Changing the current working directory (;CD)

This command changes the current working directory where the CSQL Interpreter is running to the specified directory. If you don't specify the path, the directory will be changed to the home directory.

```
csql> ;cd /home1/DBA/CUBRID
Current directory changed to /home1/DBA/CUBRID.
```

Exiting the CSQL Interpreter (;EXit)

This command exits the CSQL Interpreter.

```
csql> ;ex
```


Clearing the command buffer (;Clear)

The ;Clear session command clears the contents of the command buffer.

```
csql> ;cl
csql> ;list
```

Displaying the contents of the command buffer (;List)

The ;List session command lists the contents of the command buffer that have been entered or modified. The command buffer can be modified by ;READ or ;Edit command.

```
csql> ;l
```

Executing SQL statements (;Run)

This command executes SQL statements in the command buffer. Unlike the ;Xrun session command described below, the buffer will not be cleared even after the query execution.

```
csql> ;ru
```

Clearing the command buffer after executing the SQL statement (;Xrun)

This command executes SQL statements in the command buffer. The buffer will be cleared after the query execution.

```
csql> ;x
```

Committing transaction (;Commit)

This command commits the current transaction. You must enter a commit command explicitly if it is not in auto-commit mode. In auto-commit mode, transactions are automatically committed whenever SQL is executed.

```
csql> ;co
Current transaction has been committed.
```

Rolling back transaction (;Rollback)

This command rolls back the current transaction. Like a commit command (;Commit), it must enter a rollback command explicitly if it is not in auto-commit mode (OFF).

```
csql> ;ro
Current transaction has been rolled back.
```

Setting the auto-commit mode (;Autocommit)

This command sets auto-commit mode to ON or OFF. If any value is not specified, current configured value is applied by default. The default value is ON.

```
csql> ;au off
AUTOCOMMIT IS OFF
```

Checkpoint Execution (;Checkpoint)

This command executes the checkpoint within the CSQL session. This command can only be executed when a DBA group member, who is specified for the custom option (-u *user_name*), connects to the CSQL Interpreter in system administrator mode (--sysadm).

Checkpoint is an operation of flushing all dirty pages within the current data buffer to disks. You can also change the checkpoint interval using a command (;set *parameter_name* value) to set the parameter values in the CSQL session. You can see the examples of the parameter related to the checkpoint execution interval (**checkpoint_interval_in_mins** and **checkpoint_every_npages**). For more information, see [Logging-Related Parameters](#).

```
csql> ;ch
Checkpoint has been issued.
```

Transaction Monitoring Or Termination (;Killtran)

This command checks the transaction status information or terminates a specific transaction in the CSQL session. This command prints out the status information of all transactions on the screen if a parameter is omitted it terminates the

transaction if a specific transaction ID is specified for the parameter. It can only be executed when a DBA group member, who is specified for the custom option (`-u user_name`), connects to the CSQL Interpreter in system administrator mode (`--sysadm`).

```
csql> ;k
Tran index      User name      Host name      Process id     Program name
-----
1(+)            dba            myhost         664            cub_cas
2(+)            dba            myhost         6700           csql
3(+)            dba            myhost         2188           cub_cas
4(+)            dba            myhost         696            csql
5(+)            public         myhost         6944           csql

csql> ;k 3
The specified transaction has been killed.
```

Restarting database (;REStart)

A command that tries to reconnect to the target database in a CSQL session. Note that when you execute the CSQL Interpreter in CS (client/server) mode, it will be disconnected from the server. When the connection to the server is lost due to a HA failure and failover to another server occurs, this command is particularly useful in connecting to the switched server while maintaining the current session.

```
csql> ;res
The database has been restarted.
```

Displaying the current date (;DATE)

The `;DATE` command displays the current date and time in the CSQL Interpreter.

```
csql> ;date
Tue July 29 18:58:12 KST 2008
```

Displaying the database informatio (;DATAbase)

This command displays the database name and host name where the CSQL Interpreter is working. If the database is running, the HA mode (one of those followings: active, standby, or maintenance) will be displayed as well.

```
csql> ;data
demodb@localhost (active)
```

Displaying schema information of a class (;SCHEMA)

The `;SCHEMA` session command displays schema information of the specified table. The information includes the table name, its column name and constraints.

```
csql> ;sc event
=== <Help: Schema of a Class> ===
<Class Name>
  event
<Attributes>
  code          INTEGER NOT NULL
  sports        CHARACTER VARYING (50)
  name          CHARACTER VARYING (50)
  gender        CHARACTER (1)
  players       INTEGER
<Constraints>
  PRIMARY KEY pk_event_event_code ON event (code)
```

Displaying syntax (;SYntax)

This command displays the syntax of the SQL statement specified. If there is no specific syntax specified, all the syntaxes defined and their rules will be displayed.

```
csql> ;sy alter
=== <Help: Command Syntax> ===
<Name>
  ALTER
<Description>
  Change the definition of a class or virtual class.
```

```

<Syntax>
<alter> ::= ALTER [ <class type> ] <class name> <alter clause> ;
<class_type> ::= CLASS | TABLE | VCLASS | VIEW
<alter_clause> ::= ADD <alter_add> [ INHERIT <resolution_comma_list> ] |
                  DROP <alter_drop> [ INHERIT <resolution comma list> ] |
                  RENAME <alter_rename> [ INHERIT <resolution comma list> ] |
                  CHANGE <alter_change> |
                  INHERIT <resolution comma list>
>
<alter_add> ::= [ ATTRIBUTE | COLUMN ] <class_element_comma_list> |
               CLASS ATTRIBUTE <attribute_definition_comma_list> |
               FILE <file name comma list> |
               METHOD <method definition comma list> |
               QUERY <select statement> |
               SUPERCLASS <class name comma list>
.....

```

Displaying the trigger (;TRrigger)

This command searches and displays the trigger specified. If there is no trigger name specified, all the triggers defined will be displayed.

```

csql> ;tr
=== <Help: All Triggers> ===
    trig_delete_contents

```

Checking the parameter value(;Get)

You can check the parameter value currently set in the CSQL Interpreter using the ;Get session command. An error occurs if the parameter name specified is incorrect.

```

csql> ;g isolation_level
=== Get Param Input ===
isolation_level=4

```

Setting the parameter value (;SEt)

You can use the ;Set session command to set a specific parameter value. Note that changeable parameter values are only can be changed. To change the server parameter values, you must have DBA authorization. For information on list of changeable parameters, see [cubrid_broker.conf Configuration File and Default Parameters](#).

```

csql> ;se block_ddl_statement=1
=== Set Param Input ===
block_ddl_statement=1
-- Dynamically change the log_max_archives value in the csql accessed by dba account
csql>;se log_max_archives=5

```

Setting the view level of executing query plan (;PLan)

You can use the ;PLan session command to set the view level of executing query plan the level is composed of **simple**, **detail**, and **off**. Each command refers to the following:

- **off** : Not displaying the query execution plan
- **simple** : Displaying the query execution plan in simple version (OPT LEVEL=257)
- **detail** : Displaying the query execution plan in detailed version (OPT LEVEL=513)

Displaying information (;Info)

The ;Info session command allows you to check information such as schema, triggers, the working environment, locks and statistics.

```

csql> ;i lock
*** Lock Table Dump ***
Lock Escalation at = 100000, Run Deadlock interval = 1
Transaction (index 0, unknown, unknown@unknown|-1)
Isolation REPEATABLE CLASSES AND READ UNCOMMITTED INSTANCES
State TRAN ACTIVE
Timeout period -1
.....

```

Outputting statistics information of server processing (:Hist)

This command shows the statistics information of server processing. The information is collected after this command is entered. Therefore, the execution commands such as `;.dump_hist` or `;.x` must be entered to output the statistics information

This command is executable while the **communication_histogram** parameter in the **cubrid.conf** file is set to **yes**. You can also view this information by using the **cubrid statdump** utility. Following options are provided for this session command.

- **on** : Starts collecting statistics information for the current connection.
- **off** : Stops collecting statistics information of server.

This example shows the server statistics information for current connection. For information on specific items, see [Outputting Statistics Information of Server](#).

```

csql> ;.hist on
csql> ;.x
Histogram of client requests:
Name                                     Rcount   Sent size  Recv size , Server time
No server requests made

*** CLIENT EXECUTION STATISTICS ***
System CPU (sec)                         =         0
User CPU (sec)                           =         0
Elapsed (sec)                             =        20

*** SERVER EXECUTION STATISTICS ***
Num file creates                         =         0
Num file removes                         =         0
Num_file_ioreads                         =         0
Num_file_iowrites                        =         0
Num file iosynches                       =         0
Num data page fetches                    =         56
Num data page dirties                    =         14
Num data page ioreads                    =         0
Num_data_page_iowrites                   =         0
Num_data_page_victims                    =         0
Num data page iowrites for replacement =           0
Num log page ioreads                     =         0
Num log page iowrites                    =         0
Num_log_append_records                   =         0
Num_log_archives                         =         0
Num_log_checkpoints                      =         0
Num log wals                             =         0
Num page locks acquired                  =         2
Num object locks acquired                 =         2
Num_page_locks_converted                 =         0
Num_object_locks_converted               =         0
Num page locks re-requested              =         0
Num object locks re-requested            =         1
Num page locks waits                     =         0
Num_object_locks_waits                   =         0
Num_tran_commits                         =         1
Num_tran_rollbacks                       =         0
Num tran savepoints                      =         0
Num tran start topops                    =         3
Num tran end topops                      =         3
Num_tran_interrupts                      =         0
Num_btree_inserts                        =         0
Num_btree_deletes                        =         0
Num_btree_updates                        =         0
Num btree covered                        =         0
Num btree noncovered                     =         0
Num_btree_resumes                        =         0
Num_query_selects                        =         1
Num_query_inserts                        =         0
Num query deletes                        =         0
Num query updates                        =         0
Num_query_sscans                         =         1
Num_query_iscans                         =         0

```

```

Num query lscans          =          0
Num query setscans       =          0
Num_query_methscans      =          0
Num_query_nljoins        =          0
Num query mjoins         =          0
Num query objfetches     =          0
Num network requests     =          8
Num adaptive flush pages =          0
Num_adaptive_flush_log_pages =        0
Num_adaptive_flush_max_pages =        0

*** OTHER STATISTICS ***
Data page buffer hit ratio =      100.00
csql> ;h off

```

Displaying query execution time (;Time)

The **;Time** session command can be set to display the elapsed time to execute the query. It can be set to **ON** or **OFF**. The current setting is displayed if there is no value specified.

The **SELECT** query includes the time of outputting the fetched records. Therefore, to check the execution time of complete output of all records in the **SELECT** query, use the **--no-pager** option while executing the CSQC interpreter.

```

$ csql ?u dba --no-pager demodb
csql> ;ti ON
csql> ;ti
TIME IS ON

```

Displaying query history (;HISTORYList)

This command displays the list that contains previously executed commands (input) and their history numbers.

```

csql> ;historyl
----< 1 >----
select * from nation;
----< 2 >----
select * from athlete;

```

Reading input with the specified history number into the buffer (;HISTORYRead)

You can use **;HISTORYRead** session command to read input with history number in the **;HISTORYList** list into the command buffer. You can enter **;ru** or **;x** directly because it has the same effect as when you enter SQL statements directly.

```

csql> ;historyr 1

```

Calling the default editor (;EDIT)

This command calls the specified editor. The default editor is **vi** on Linux **Notepad** on Windows environment. Use **;EDITOR_Cmd** command to specify a different editor.

```

csql> ;edit

```

Specifying the editor (;EDITOR_Cmd)

This command specifies the editor to be used with **;EDIT** session command. As shown in the example below, you can specify other editor (ex: **emacs**) which is installed in the system.

```

csql> ;editor_c emacs
csql> ;edit

```


CUBRID SQL Guide

This chapter describes SQL syntax such as data types, functions and operators, data retrieval or table manipulation. You can also find SQL statements used for index, trigger, partition, serial and changing user information.

The main topics covered in this chapter are as follows:

- Glossary
- Comment
- Identifier
- Reserved words
- Data types
- Table definition
- Index definition
- VIEW definition
- SERIAL
- Operators and functions
- Data retrieval and manipulation
- Query optimization
- Triggers (TRIGGER)
- Java stored functions/procedures
- Methods
- Partitions
- Class inheritance
- Class conflict resolution
- CUBRID system catalog

Glossary

CUBRID is an object-relational database management system (ORDBMS), which supports object-oriented concepts such as inheritance. In this manual, relational database terminology is also used along with object-oriented terminology for better understanding. Object-oriented terminology such as class, instance and attribute is used to describe concepts including inheritance, and relational database terminology is mainly used to describe common SQL syntax.

The following table provides the summary:

Relational Database	CUBRID
table	class, table
column	attribute, column
record	instance, record
data type	domain, data type

Comment

The CSQL Interpreter is a SQL-style method; the SQL-style comment starts with the double dashes (--) and the comment line after the double dashes is regarded as comment. Additionally, it supports C++ style, which start with double slashes (//), and C-style, which starts and ends with '/*' and '*/' respectively.

The following are examples of comments supported in the CSQL Interpreter.

Example

- How to use --

```
-- This is a SQL-style comment.
```

- How to use //

```
This is a C++ style comment.
```

- How to use /* */

```
/* This is a C-style comment.*/  
/* This is an example to use two lines  
as comment by using the C-style. */
```

Identifier

Guidelines for Creating Identifiers

The guidelines for creating identifiers in the CSQL Interpreter are as follows:

- An identifier must begin with a letter it must not begin with a number or a symbol.
- It is not case-sensitive.
- CUBRID keywords are not allowed.

< identifier >

```
:: = < identifier_letter> [ { < other_identifier> } ]
```

< identifier_letter >

```
:: = < upper_case_letter>  
| < lower_case_letter>
```

< other_identifier >

```
:: = < identifier letter>  
| < digit>  
|  
| #
```

< digit >

```
:: = 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
```

< upper_case_letter >

```
:: = A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V  
| W | X | Y | Z
```

< lower_case_letter >

```
:: = a | b | c | d | e | f | g | h | i | j | k | l | m | n | o | p | q | r | s | t | u | v  
| w | x | y | z
```

Legal Identifiers

Beginning with a Letter

An identifier must begin with a letter. All other special characters except operator characters are allowed. The following are examples of legal identifiers.

```
a  
a_b  
ssn#  
this_is_an_example_#
```

Enclosing in Double Quotes, Square Brackets, or Backtick Symbol

Identifiers or a reserved keywords shown as below are not allowed' however, if they are enclosed in in double quotes, square brackets, or backtick symbol, they are allowed as an exception. Especially, the double quotations can be used as a symbol enclosing identifiers when the **ansi_quotes** parameter is set to **yes**. If this value is set to **no**, double quotations are used as a symbol enclosing character strings. The followings are examples of legal identifiers.

```
" select"  
" @lowcost"  
" low cost"  
" abc" " def"  
[position]
```

Illegal Identifiers

Beginning with special characters or numbers

An identifier starting with a special character or a number is not allowed. As an exception, an underline (`_`) and a sharp symbol (`#`) are allowed for the first character.

```
_a  
#ack  
%nums  
2fer  
88abs
```

An identifier containing a space

An identifier that has a space within characters is not allowed.

```
col1 t1
```

An identifier containing operator special characters

An identifier which contains operator special characters (`+`, `-`, `*`, `/`, `%`, `||`, `!`, `<`, `>`, `=`, `|`, `^`, `&`, `~`) is not allowed.

```
col+  
col~  
col& &
```

Reserved Words

The following keywords are previously reserved as a command, a function name or a type name in CUBRID. You are restricted to use these words for a class name, an attribute name, a variable name. Note than these reserved keywords can be used an identifier when they are enclosed in double quotes, square brackets, or backtick symbol (`).

ABSOLUTE	ACTION	ADD
ADD_MONTHS	AFTER	ALIAS
ALL	ALLOCATE	ALTER
AND	ANY	ARE
AS	ASC	ASSERTION
ASYNC	AT	ATTACH
ATTRIBUTE	AVG	
BEFORE	BETWEEN	BIGINT
BIT	BIT_LENGTH	BLOB
BOOLEAN	BOTH	BREADTH
BY		
CALL	CASCADE	CASCADED
CASE	CAST	CATALOG
CHANGE	CHAR	CHARACTER
CHECK	CLASS	CLASSES
CLOB	CLOSE	CLUSTER
COALESCE	COLLATE	COLLATION
COLUMN	COMMIT	COMPLETION
CONNECT	CONNECT_BY_ISCYCLE	CONNECT_BY_ISLEAF
CONNECT_BY_ROOT	CONNECTION	CONSTRAINT
CONSTRAINTS	CONTINUE	CONVERT
CORRESPONDING	COUNT	CREATE
CROSS	CURRENT	CURRENT_DATE
CURRENT_DATETIME	CURRENT_TIME	CURRENT_TIMESTAMP
CURRENT_USER	CURSOR	CYCLE
DATA	DATA_TYPE	DATABASE
DATE	DATETIME	DAY
DAY_HOUR	DAY_MILLISECOND	DAY_MINUTE
DAY_SECOND	DEALLOCATE	DEC
DECIMAL	DECLARE	DEFAULT
DEFERRABLE	DEFERRED	DELETE
DEPTH	DESC	DESCRIBE
DESCRIPTOR	DIAGNOSTICS	DICTIONARY
DIFFERENCE	DISCONNECT	DISTINCT

DISTINCTROW	DIV	DO
DOMAIN	DOUBLE	DUPLICATE
DROP		
EACH	ELSE	ELSEIF
END	EQUALS	ESCAPE
EVALUATE	EXCEPT	EXCEPTION
EXCLUDE	EXEC	EXECUTE
EXISTS	EXTERNAL	EXTRACT
FALSE	FETCH	FILE
FIRST	FLOAT	FOR
FOREIGN	FOUND	FROM
FULL	FUNCTION	
GENERAL	GET	GLOBAL
GO	GOTO	GRANT
GROUP		
HAVING	HOUR	HOUR_MILLISECOND
HOUR_MINUTE	HOUR_SECOND	
IDENTITY	IF	IGNORE
IMMEDIATE	IN	INDEX
INDICATOR	INHERIT	INITIALLY
INNER	INOUT	INPUT
INSERT	INT	INTEGER
INTERSECT	INTERSECTION	INTERVAL
INTO	IS	ISOLATION
JOIN		
KEY		
LANGUAGE	LAST	LDB
LEADING	LEAVE	LEFT
LESS	LEVEL	LIKE
LIMIT	LIST	LOCAL
LOCAL_TRANSACTION_ID	LOCALTIME	LOCALTIMESTAMP
LOOP	LOWER	
MATCH	MAX	METHOD
MILLISECOND	MIN	MINUTE
MINUTE_MILLISECOND	MINUTE_SECOND	MOD
MODIFY	MODULE	MONETARY
MONTH	MULTISET	MULTISET_OF
NA	NAMES	NATIONAL
NATURAL	NCHAR	NEXT
NO	NONE	NOT

NULL	NULLIF	NUMERIC
OBJECT	OCTET_LENGTH	OF
OFF	OID	ON
ONLY	OPEN	OPERATION
OPERATORS	OPTIMIZATION	OPTION
OR	ORDER	OTHERS
OUT	OUTER	OUTPUT
OVERLAPS		
PARAMETERS	PARTIAL	PENDANT
POSITION	PRECISION	PREORDER
PREPARE	PRESERVE	PRIMARY
PRIOR	PRIVATE	PRIVILEGES
PROCEDURE	PROTECTED	PROXY
QUERY		
READ	REAL	RECURSIVE
REF	REFERENCES	REFERENCING
REGISTER	RELATIVE	RENAME
REPLACE	RESIGNAL	RESTRICT
RETURN	RETURNS	REVOKE
RIGHT	ROLE	ROLLBACK
ROLLUP	ROUTINE	ROW
ROWNUM	ROWS	
SAVEPOINT	SCHEMA	SCOPE
SCROLL	SEARCH	SECOND
SECOND_MILLISECOND	SECTION	SELECT
SENSITIVE	SEQUENCE	SEQUENCE_OF
SERIALIZABLE	SESSION	SESSION_USER
SET	SET_OF	SETEQ
SHARED	SIBLINGS	SIGNAL
SIMILAR	SIZE	SMALLINT
SOME	SQL	SQLCODE
SQLERROR	SQLEXCEPTION	SQLSTATE
SQLWARNING	STATISTICS	STRING
STRUCTURE	SUBCLASS	SUBSET
SUBSETEQ	SUBSTRING	SUM
SUPERCLASS	SUPERSET	SUPERSETEQ
SYS_CONNECT_BY_PATH	SYS_DATE	SYS_DATETIME
SYS_TIME	SYS_TIMESTAMP	SYS_USER
SYSDATE	SYSDATETIME	SYSTEM_USER
SYSTIME		

TABLE	TEMPORARY	TEST
THEN	THERE	TIME
TIMESTAMP	TIMEZONE_HOUR	TIMEZONE_MINUTE
TO	TRAILING	TRANSACTION
TRANSLATE	TRANSLATION	TRIGGER
TRIM	TRUE	TRUNCATE
TYPE		
UNDER	UNION	UNIQUE
UNKNOWN	UPDATE	UPPER
USAGE	USE	USER
USING	UTIME	
VALUE	VALUES	VARCHAR
VARIABLE	VARYING	VCLASS
VIEW	VIRTUAL	VISIBLE
WAIT	WHEN	WHENEVER
WHERE	WHILE	WITH
WITHOUT	WORK	WRITE
XOR		
YEAR	YEAR_MONTH	
ZONE		

Data Types

Numeric Types

Definition and Characteristics

Definition

CUBRID supports the following numeric data types to store integers or real numbers.

Numeric Types Supported by CUBRID

Type	Bytes	Mix	Max	Exact/approx.
SHORT SMALLINT	2	-32,768	+32,767	exact numeric
INT INTEGER	4	-2,147,483,648	+2,147,483,647	exact numeric
BIGINT	8	-9,223,372,036,854,775,808	+9,223,372,036,854,775,807	exact numeric
NUMERIC DECIMAL	16	precision p : 1 scale s : 0	precision p : 38 scale s : 38	exact numeric
FLOAT REAL	4	-3.402823466E+38 (ANSI/IEEE 754-1985 standard)	+3.402823466E+38 (ANSI/IEEE 754-1985 standard)	approximate numeric floating point : 7
DOUBLE DOUBLE PRECISION	8	- 1.7976931348623157E+308 (ANSI/IEEE 754-1985 standard)	+1.7976931348623157E+308 (ANSI/IEEE 754-1985 standard)	approximate numeric floating point : 15
MONETARY	12	-3.402823466E+38	+3.402823466E+38	approximate numeric

Numeric data types are divided into exact and approximate types. Exact numeric data types (**SMALLINT**, **INT**, **BIGINT**, **NUMERIC**) are used for numbers whose values must be precise and consistent, such as the numbers used in financial accounting. Note that even when the literal values are equal, approximate numeric data types (**FLOAT**, **DOUBLE**, **MONETARY**) can be interpreted differently depending on the system.

CUBRID does not support the **UNSIGNED** type for numeric data types.

Characteristics

Precision and Scale

The precision of numeric data types is defined as the number of significant figures. This applies to both exact and approximate numeric data types.

The scale represents the number of digits following the decimal point. It is significant only in exact numeric data types. Attributes declared as exact numeric data types always have fixed precision and scale. **NUMERIC** (or **DECIMAL**) data type always has at least one-digit precision, and the scale should be between 0 and the precision declared. Scale cannot be greater than precision. For **INTEGER**, **SMALLINT**, or **BIGINT** data types, the scale is 0 (i.e. no digits following the decimal point), and the precision is fixed by the system.

Numeric Literals

Special signs can be used to input numeric values. The plus sign (+) and minus sign (-) are used to represent positive and negative numbers respectively. You can also use scientific notations. In addition, you can use currency signs specified in the system to represent currency values. The maximum precision that can be expressed by a numeric literal is 255.

Numeric Coercions

All numeric data type values can be compared with each other. To do this, automatic coercion to the common numeric data type is performed. For explicit coercion, use the **CAST** operator. When different data types are sorted or calculated in a numerical expression, the system performs automatic coercion. For example, when adding a **FLOAT** attribute value to an **INTEGER** attribute value, the system automatically coerces the **INTEGER** value to the most approximate **FLOAT** value before it performs the addition operation.

Caution Earlier version than CUBRID 2008 R2.0, the input constant value exceeds **INTEGER**, it is handled as **NUMERIC**. However, 2008 R2.0 or later versions, it is handled as **BIGINT**.

INT/INTEGER

Description

The **INTEGER** data type is used to represent integers. The value range is available is from -2,147,483,648 to +2,147,483,647. **SMALLINT** is used for small integers, and **BIGINT** is used for big integers.

INTEGER |
INT

Remark

- If a real number is entered for an **INT** type, the number is rounded to zero decimal place and the integer value is stored.
- **INTEGER** and **INT** are used interchangeably.

Example

```
If you specify 8934 as INTEGER, 8934 is stored.
If you specify 7823467 as INTEGER, 7823467 is stored.
If you specify 89.8 to an INTEGER, 90 is stored (all digits after the decimal point are rounded).
If you specify 3458901122 as INTEGER, an error occurs (if the allowable limit is exceeded).
```

SHORT/SMALLINT

Description

The **SMALLINT** data type is used to represent a small integer type. The value range is available is from -32,768 to +32,767.

SMALLINT |
SHORT

Remark

- If a real number is entered for an **SMALLINT** type, the number is rounded to zero decimal place and the integer value is stored.
- **SMALLINT** and **SHORT** are used interchangeably.

Example

```
If you specify 8934 as SMALLINT, 8934 is stored.
If you specify 34.5 as SMALLINT, 35 is stored (all digits after the decimal point are rounded).
If you specify 23467 as SMALLINT, 23467 is stored.
```

If you specify 89354 as SMALLINT, an error occurs (if the allowable limit is exceeded).

BIGINT

Description

The **BIGINT** data type is used to represent big integers. The value range is available from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.

BIGINT

Remark

- If a real number is entered for a **BIG** type, the number is rounded to zero decimal place and the integer value is stored.
- Based on the precision and the range of representation, the following order applies.

SMALLINT \subset **INTEGER** \subset **BIGINT** \subset **NUMERIC**

Example

If you specify 8934 as BIGINT, 8934 is stored.
 If you specify 89.1 as BIGINT, 89 is stored.
 If you specify 89.8 as BIGINT, 90 is stored (all digits after the decimal point are rounded).
 If you specify 3458901122 as BIGINT, 3458901122 is stored.

NUMERIC/DECIMAL

Description

NUMERIC or **DECIMAL** data types are used to represent fixed-point numbers. As an option, the total number of digits (precision) and the number of digits after the decimal point (scale) can be specified for definition. The minimum value for the precision p is 1. When the precision p is omitted, you cannot enter data whose integer part exceeds 15 digits because the default value is 15. If the scale s is omitted, an integer rounded to the first digit after the decimal point is returned because the default value is 0.

NUMERIC [(p [, s])]

Remark

- Precision must be equal to or greater than scale.
- Precision must be equal to or greater than the number of integer digits + scale.
- **NUMERIC**, **DECIMAL**, and **DEC** are used interchangeably.

Example

If you specify 12345.6789 as NUMERIC, 12346 is stored (it rounds to the first place after the decimal point since 0 is the default value of scale).
 If you specify 12345.6789 as NUMERIC(4), an error occurs (precision must be equal to or greater than the number of integer digits).
 If you declare NUMERIC(3,4), an error occurs (precision must be equal to or greater than the scale).
 If you specify 0.12345678 as NUMERIC(4,4), .1235 is stored (it rounds to the fifth place after the decimal point).
 If you specify -0.123456789 as NUMERIC(4,4), -.1235 is stored (it rounds to the fifth place after decimal point and then prefixes a minus (-) sign).

FLOAT/REAL

Description

The **FLOAT** (or **REAL**) data type represents floating point numbers.

The ranges of values that can be described as normalized values are from $-3.402823466E+38$ to $-1.175494351E-38$, 0, and from $+1.175494351E-38$ to $+3.402823466E+38$, whereas the values other than normalized values, which are closer to 0, are described as de-normalized values. It conforms to the ANSI/IEEE 754-1985 standard.

The minimum value for the precision p is 1 and the maximum value is 38. When the precision p is omitted or it is specified as seven or less, it is represented as single precision (in 7 significant figures). If the precision p is greater than 7 and equal to or less than 38, it is represented as double precision (in 15 significant figures) and it is converted into **DOUBLE** data type.

FLOAT data types must not be used if you want to store a precise value that exceeds the number of significant figures, as they only store the approximate value of any input value over 7 significant figures.

FLOAT [(p)]

Remark

- **FLOAT** is in 7 significant figures.
- Extra cautions are required when comparing data because the **FLOAT** type stores approximate numeric.
- **FLOAT** and **REAL** are used interchangeably.

Example

If you specify 16777217 as **FLOAT**, 16777216 is stored and 1.677722e+07 is displayed (if precision is omitted, 8-th digit is rounded up because it is represented as 7 significant figures).

If you specify 16777217 as **FLOAT(5)**, 16777216 is stored and 1.677722e+07 is displayed (if precision is in seven or less, 8-th digit is rounded up because it is represented as 7 significant figures).

If you specify 16777.217 as **FLOAT(5)**, 16777.216 is stored and 1.677722e+04 is displayed (if precision is in seven or less, 8-th digit is rounded up because it is represented as 7 significant figures).

If you specify 16777.217 as **FLOAT(10)**, 16777.217 is stored and 1.6777217000000000e+04 is displayed (if precision is greater than 7 and less than or equal to 38, zeroes are added because it is represented as 15 significant figures).

DOUBLE/DOUBLE PRECISION

Description

The **DOUBLE** data type is used to represent floating point numbers.

The ranges of values that can be described as normalized values are from $-1.7976931348623157E+308$ to $-2.2250738585072014E-308$, 0, and from $2.2250738585072014E-308$ to $1.7976931348623157E+308$, whereas the values other than normalized values, which are closer to 0, are described as de-normalized values. It conforms to the ANSI/IEEE 754-1985 standard.

The precision p is not specified. The data specified as this data type is represented as double precision (in 15 significant figures).

DOUBLE data types must not be used if you want to store a precise value that exceeds the number of significant figures, as they only store the approximate value of any input value over 15 significant figures.

DOUBLE

Remark

- **DOUBLE** is in 15 significant figures.
- Extra caution is required when comparing data because the **DOUBLE** type stores approximate numeric.
- **DOUBLE** and **DOUBLE PRECISION** are used interchangeably.

Example

If you specify 1234.56789 as **DOUBLE**, 1234.56789 is stored and 1.2345678900000000e+03 is displayed.

If you specify 9007199254740993 as DOUBLE, 9007199254740992 is stored and 9.007199254740992e+15 is displayed.

MONETARY

Description

The **MONETARY** data type is an approximate numeric data type. The range of valid value is the same as **DOUBLE**, which is represented to two decimal places; the value range can be different based on system. A comma is appended to every 1000th place.

MONETARY

Remark

You can use a dollar sign or a decimal point, but a comma is not allowed.

Example

If you specify 12345.67898934 as MONETARY, \$12,345.68 is stored (it is rounded to third decimal place).
If you specify 123456789 as MONETARY, \$123,456.789.00 is stored.

Date/Time Types

Definition and Characteristics

Definition

DATE-TIME data types are used to represent the date or time (or both together). CUBRID supports the following data types:

Date-Time Types Supported by CUBRID

Type	Min.	Max.	Note
DATE	0001-01-01	9999-12-31	As an exception, the TIMESTAMP '0000-00-00 00:00:00' value is allowed.
TIME	00:00:00	23:59:59	
TIMESTAMP	1970-01-01 00:00:01 (GMT) 1970-01-01 09:00:01 (KST)	2038-01-10 03:14:07 (GMT) 2038-01-19 12:14:07 (KST)	As an exception, the TIMESTAMP '0000-00-00 00:00:00' value is allowed.
DATETIME	0001-01-01 00:00:000	9999-12-31 23:59:599	As an exception, the DATETIME '0000-00-00 00:00:00' value is allowed.

Characteristics

Range and Resolution

- By default, the range of a time value is represented by the 24-hour system. Dates follow the Gregorian calendar. An error occurs if a value that does not meet these two constraints is entered as a date or time.
- The range of year in **DATE** is 0001 - 9999 AD.
- From the CUBRID 2008 R3.0 version, if time value is represented with two-digit numbers, a number from 00 to 69 is converted into a number from 2000 to 2069; a number from 70 to 99 is converted into a number from 1970 to 1999. In earlier than CUBRID 2008 R3.0 version, if time value is represented with two-digit numbers, a number from 01 to 99 is converted into a number from 0001 to 0099.
- The range of **TIMESTAMP** is between January 1, 1970 00:00:01 and January 19, 2038 03:14:07 (GMT). For KST (GMT+9), values from January 1, 1970 00:00:01 to January 19, 2038 12:14:07 can be stored.

- The results of date, time and timestamp operations may depend on the rounding mode. In these cases, for Time and Timestamp, the most approximate second is used as the minimum resolution; for Date, the most approximate date is used as the minimum resolution.

Coercions

The **Date-Time** types can be cast explicitly using the **CAST** operator only when they have the same field. For implicit coercion, see [Implicit Type Conversion](#). The following table shows types that allows explicit coercions. For implicit coercion, see [Arithmetic Operation and Type Casting of DATE/TIME Data Types](#).

Explicit Coercions

	TO			
FROM	DATE	TIME	DATETIME	TIMESTAMP
DATE	-	X	O	O
TIME	X	-	X	X
DATETIME	O	O	-	O
TIMESTAMP	O	O	O	-

Remarks

Every value of date and time can be 0 in **DATE**, **DATETIME**, or **TIMESTAMP** type. This is useful in terms that this value can be used if an index exists upon query execution of a column corresponding to the type.

- Some functions in which the **DATE**, **DATETIME**, or **TIMESTAMP** type is specified as an argument returns different value based on the **return_null_on_function_errors** system parameter if every argument value for date and time is 0. If **return_null_on_function_errors** is yes, **NULL** is returned; if no, an error is returned. The default value is **no**.
- The functions that return **DATE**, **DATETIME**, or **TIMESTAMP** type can return a value of 0 for date and time. However, these values cannot be stored in Date objects in Java applications. Therefore, it will be processed with one of the followings based on the configuration of zeroDateTimeBehaviorthe, connection URL property: being handled as an exception, returning **NULL**, or returning a minimum value (see "API Reference > JDBC API > JDBC Programming > Connection Configuration").

For details, see the description of each function.

DATE

Description

The **DATE** data type is used to represent the year (yyyy), month (mm) and day (dd). Supported range is "01/01/0001" to "12/31/9999." The year can be omitted. If it is, the year value of the current system is specified automatically.

Output and input formats are as follows:

```
'mm/dd[/yyyy]'
```

```
'[yyyy-]mm-dd'
```

Remark

- All fields must be entered as integer.
- The date value is displayed in the format of 'MM/DD/YYYY' in CSQL, and it is displayed in the format of 'YYYY-MM-DD' in JDBC application programs and the CUBRID Manager.
- The **TO_DATE()** function is used to convert a character string type into a **DATE** type.
- 0 is not allowed to input in year, month, and day; however, '0000-00-00', which every digit consisting of year, month, and day is 0, is allowed as an exception.

Example

```
DATE '2008-10-31' is stored as '10/31/2008'.
DATE '10/31' is stored as '10/31/2011'(if a value for year is omitted, the current year is automatically specified).
```

```
DATE '00-10-31' is stored as '10/31/2000'.
DATE '0000-10-31' is handled as an error (a year value should be at least 1).
DATE '70-10-31' is stored as '10/31/1970'.
DATE '0070-10-31' is stored as '10/31/0070'.
```

TIME

Description

The **TIME** data type is used to represent the hour (hh), minute (mm) and second (ss). Supported range is "00:00:00" to "23:59:59." Second can be omitted; if it is, 0 seconds is specified. Both 12-hour and 24-hour notations are allowed as an input format.

The input format of **TIME** is as follows:

```
'hh:mi [:ss] [am | pm]'
```

Remark

- All items must be entered as integer.
- AM/PM time notation is used to display time in the CSQL; while the 24-hour notation is used in the CUBRID Manager.
- AM/PM can be specified in the 24-hour notation. An error occurs if the time specified does not follow the AM/PM format.
- Every time value is stored in the 24-hour notation. **db_time_decode**, one of C API functions, is used to return a value in the 24-hour notation.
- The **TO_TIME()** function is used to return a character string type into a TIME type.

Example

```
TIME '00:00:00' is outputted as '12:00:00 AM'.
TIME '1:15' is regarded as '01:15:00 AM'.
TIME '13:15:45' is regarded as '01:15:45 PM'.
TIME '13:15:45 pm' is stored normally.
TIME '13:15:45 am' is an error (an input value does not match the AM/PM format).
```

TIMESTAMP

Description

The **TIMESTAMP** data type is used to represent a data value in which the date (year, month, date) and time (hour, minute, second) are combined. The range of representable value is between GMT 1970-01-01 00:00:01 and 2038-01-19 03:14:07. The **DATETIME** type can be used if the value is out of range or data in milliseconds is stored.

The input format of **TIMESTAMP** is as follows:

```
'hh:mi [:ss] [am|pm] mm/dd [/yyyy] '
'hh:mi [:ss] [am|pm] [yyyy-]mm-dd'

'mm/dd [/yyyy] hh:mi [:ss] [am|pm] '
'[yyyy-]mm-dd hh:mi [:ss] [am|pm]'
```

Remark

- All fields must be entered in integer format.
- If the year is omitted, the current year is specified by default. If the time value (hour/minute/second) is omitted, 12:00:00 AM is specified.
- You can store the timestamp value of the system in the **TIMESTAMP** type by using the **SYS_TIMESTAMP** (or **SYSTIMESTAMP**, **CURRENT_TIMESTAMP**) function. Note that the timestamp value is specified as a default value at the time of creating the table, not at the time of **INSERT** the data, if **SYS_TIMESTAMP** is specified as a **DEFAULT** value for a **TIMESTAMP** column when creating a table.
- The **TIMESTAMP()** or **TO_TIMESTAMP()** function is used to cast a character string type into a **TIMESTAMP** type.

- 0 is not allowed to input in year, month, and day; however, '0000-00-00 00:00:00', which every digit consisting of year, month, day, hour, minute, and second is 0, is allowed as an exception.

Example

```
TIMESTAMP '10/31' is outputted as '12:00:00 AM 10/31/2011' (if the value for year/time is
omitted, a default value is outputted ).
TIMESTAMP '10/31/2008' is outputted as '12:00:00 AM 10/31/2008' (if the value for time is
omitted, a default value is outputted ).
TIMESTAMP '13:15:45 10/31/2008' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '01:15:45 PM 2008-10-31' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '13:15:45 2008-10-31' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '10/31/2008 01:15:45 PM' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '10/31/2008 13:15:45' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '2008-10-31 01:15:45 PM' is outputted as '01:15:45 PM 10/31/2008'.
TIMESTAMP '2008-10-31 13:15:45' is outputted as '01:15:45 PM 10/31/2008'.
An error occurs on TIMESTAMP '2099-10-31 01:15:45 PM' (out of range to represent
TIMESTAMP).
```

DATETIME

Description

The **DATETIME** data type is used to represent a data value in which the data (year, month, date) and time (hour, minute, second) are combined. The range of representable value is between 0001-01-01 00:00:00.000 and 9999-12-31 23:59:59.999 (GMT).

The input format of **TIMESTAMP** is as follows:

```
'hh:mi [:ss[.msec]] [am|pm] mm/dd [/yyyy] '
'hh:mi [:ss[.msec]] [am|pm] [yyyy-]mm-dd'
'mm/dd[/yyyy] hh:mi[:ss[.ff]] [am|pm] '
'[yyyy-]mm-dd hh:mi[:ss[.ff]] [am|pm] '
```

Remark

- All fields must be entered as integer.
- If you year is omitted, the current year is specified by default. If the value (hour, minute/second) is omitted, 12:00:00.000 AM is specified.
- You can store the timestamp value of the system in the **DATETIME** type by using the **SYS_DATETIME** (or **SYSDATETIME**, **CURRENT_DATETIME**, **CURRENT_DATETIME()**, **NOW()**) function. Note that the timestamp value is specified as a default value at the time of creating the table, not at the time of **INSERT** the data, if **SYS_DATETIME** is specified as a **DEFAULT** value for a **DATETIME** column when creating a table.
- The **TO_DATETIME()** function is used to convert a string type into a **DATETIME** type.
- 0 is not allowed to input in year, month, and day; however, '0000-00-00 00:00:00', which every digit consisting of year, month, day, hour, minute, and second is 0, is allowed as an exception.

Example

```
DATETIME '10/31' is outputted as '12:00:00.000 AM 10/31/2011' (if the value for year/time
is omitted, a default value is outputted).
DATETIME '10/31/2008' is outputted as '12:00:00.000 AM 10/31/2008'.
DATETIME '13:15:45 10/31/2008' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '01:15:45 PM 2008-10-31' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '13:15:45 2008-10-31' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '10/31/2008 01:15:45 PM' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '10/31/2008 13:15:45' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '2008-10-31 01:15:45 PM' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '2008-10-31 13:15:45' is outputted as '01:15:45.000 PM 10/31/2008'.
DATETIME '2099-10-31 01:15:45 PM' is outputted as '01:15:45.000 PM 10/31/2099'.
```

Converting a String to Date/Time Type

Recommended Format for Strings in Date/Time Type

When you convert a string to Date/Time type, it is recommended to write the string in the following format:

- **DATE** Type

```
[[[ [Y]Y]YY]M]MDD
[[ [Y]Y]YY]-M]M-DD
MM/DD/YYYY
```

- **TIME** Type

```
HH[:MM[:SS]] ["am"|"pm"]
```

- **DATETIME** Type

```
YYYY-MM-DD HH:MM:SS[.msec] YY-MM-DD HH:MM:[SS[.msec]]
YY-MM-DD H
```

- **TIMESTAMP** Type

```
YYYY-MM-DD HH:MM:SS
YY-MM-DD HH:MM:[SS]
YY-MM-DD H
```

Available DATE String Format

```
[year sep] month sep day
```

- 2011-04-20
- 04-20

If a separator (*sep*) is a slash (/), strings are recognized in the following order:

```
month/day[/year]
```

- 04/20/2011
- 04/20

If you do not use a separator (*sep*), strings are recognized in the following format. It is allowed to use up to 4 digits for years and up to 2 digits for months. You must enter a 2-digit day.

```
[[[ [Y]Y]YY]M]MDD
```

- 20110420
- 110420
- 420

Available TIME String Format

```
[hour]:min[:[sec]] [.[msec]] [am|pm]
```

- 09:10:15.359 am
- 09:10:15
- 09:10
- :10

```
[[[[[ [Y]Y]Y]Y]M]MDD]HHMMSS[.[msec]] [am|pm]
```

- 20110420091015.359 am
- 0420091015

```
[H]HMMSS[.[msec]] [am|pm]
```

- 091015.359 am
- 91015

```
[M]MSS[.[msec]] [am|pm]
```

- 1015.359 am
- 1015

```
[S]S[.[msec]] [am|pm]
```

- 15.359 am
- 15

Note: The [H]H format was allowed in CUBRID 2008 R3.1 and the earlier versions. That is, the string '10' was converted to **TIME** '10:00:00' in the R3.1 and the earlier versions, and will be converted to **TIME** '00:00:10' in version R4.0 and later.

Available String Format in Time-Date

`[hour]:min[:sec[.msec]] [am|pm] sep [year-]month-day`

- 09:10:15.359 am 2011-04-20
- :10 04-20

`[hour]:min[:sec[.msec]] [am|pm] sep month/day[/[year]]`

- 09:10:15.359 am 04/20/2011
- :10 04/20

`hour[:min[:sec[.msec]]] [am|pm] sep [year-]month-day`

- 09:10:15.359 am 04-20
- 09 04-20

`hour[:min[:sec[.msec]]] [am|pm] sep month/day[/[year]]`

- 09:10:15.359 am 04/20
- 09 04/20

Available DATETIME String Format

`[year sep] month sep day [sep] [sep] hour [sep min[sep sec[.msec]]]`

- 04-20 09

`month/day[/year] [sep] hour [sep min [sep sec[.msec]]]`

- 04/20 09

`year sep month sep day sep hour [sep min[sep sec[.msec]]]`

- 2011-04-20 09

`month/day/year sep hour [sep min[sep sec [.msec]]]`

- 04/20/2011 09

`YYMMDDH (??? ? ?? ?? ?? ??)`

- 1104209

`YYMMDDHHMM[SS[.msec]]`

- 1104200910.359

`YYYYMMDDHHMMSS[.msec]`

- 201104200910.359

Rules

msec is a series of numbers representing milliseconds. The numbers after the fourth digit will be ignored.

sep represents the separator string allowed. The rules for the separator string are as follows:

- You should always use one colon (:) as a separator for the **TIME** separator.
- **DATE** and **DATETIME** strings can be represented as a series of numbers without the separator (*sep*), and non-alphanumeric characters can be used as separators. The **DATETIME** string can be divided into Time and Date with a space.
- Separators do not need to be identical in the input string.
- For the Time-Date string, you can only use colon (:) for a Time separator and hyphen (-) or slash (/) for a Date separator.
- For the **DATE** string, you can use colon (:) or other separators.

The following rules will be applied to the Date part in the string.

- You can omit the year as long as the syntax allows it.

- If you enter the year as two digits, it represents the range from 1970-2069. That is, if YY<70, it is treated as 2000+YY; if YY>=70, it is treated as 1900+YY. If you enter one, three or four digit numbers for the year, the numbers will be represented as they are.
- A space before and after a string and the string next to the space are ignored. The am/pm identifier for the **DATETIME** and **TIME** strings can be recognized as part of TIME value, but are not recognized as the am/pm identifier if non-space characters are added to it.

The **TIMESTAMP** type of CUBRID consists of **DATE** type and **TIME** type, and **DATETIME** type consists of **DATE** type and **TIME** type with milliseconds being added to them. Input strings can include Date (**DATE** string), Time (**TIME** string), or both (**DATETIME** strings). You can convert a string including a specific type of data to another type, and the following rules will be applied for the conversion.

- If you convert the **DATE** string to the **DATETIME** type, the time value will be '00:00:00.'
- If you convert the **TIME** string to the **DATETIME** type, colon (:) is recognized as a date separator, so that the **TIME** string can be recognized as a date string and the time value will be '00:00:00.'
- If you convert the **DATETIME** string to the **DATE** type, the time part will be ignored from the result but the time input value format should be valid.
- You can convert the **DATETIME** string to the **TIME** type, and you must follow the following rules.
- The DATE and TIME in the string must be divided by at least one blank.
- The date part of the result value is ignored but the date input value format should be valid.
- The year in the date part must be over 4 digits (available to start with 0) or the time part must include hours and minutes ([H]H:[M]M) at least. Otherwise the date part are recognized as the **TIME** type of the [MM]SS format, and the following string will be ignored.
- If the one of the units (year, month, date, hour, minute and second) of the **DATETIME** string is greater than 999999, it is not recognized as a number, so the string including the corresponding unit will be ignored. For example, in '2009-10-21 20:9943:10', an error occurs because the value in minutes is out of the range. However, if '2009-10-21 20:1000123:10' is entered, '2009' is recognized as the **TIME** type of the MMSS format, so that **TIME** '00:20:09' will be returned.
- If you convert the TIME-DATE string to the **TIME** type, the date part of the string is ignored but the date part format must be valid.
- All input strings including the time part allow [*msec*] on conversion, but only the **DATETIME** type can be maintained. If you convert this to a type such as **DATE**, **TIMESTAMP** or **TIME**, the *msec* value is discarded.
- All conversions in the **DATETIME**, **TIME** string allow English locale following after time value or am/pm identifier written in the current locale of a server.

Example

```
SELECT CAST('420' AS DATE);

    cast('420' as date)
=====
    04/20/2011

SELECT CAST('91015' AS TIME);

    cast('91015' as time)
=====
    09:10:15 AM

SELECT CAST('110420091035.359' AS DATETIME);

    cast('110420091035.359' as datetime)
=====
    09:10:35.359 AM 04/20/2011

SELECT CAST('110420091035.359' AS TIMESTAMP);

    cast('110420091035.359' as timestamp)
=====
    09:10:35 AM 04/20/2011
```

Bit Strings

Definition and Characteristics

Definition

A bit string is a sequence of bits (1's and 0's). Images (bitmaps) displayed on the computer screen can be stored as bit strings. CUBRID supports the following two types of bit strings:

- Fixed-length bit string (**BIT**)
- Variable-length bit string (**BIT VARYING**)

A bit string can be used as a method argument or an attribute domain. Bit string literals are represented in a binary or hexadecimal format. For binary format, append the string consisting of 0's and 1's to the letter **B** or append a value to the **0b** as shown example below.

```
B'1010'
0b1010
```

For hexadecimal format, append the string consisting of the numbers 0 - 9 and the letters A - F to the uppercase letter **X** or append a value to the **0x**. The following is hexadecimal representation of the same number that was represented above in binary format.

```
X'a'
0xA
```

The letters used in hexadecimal numbers are not case-sensitive. That is, X'4f' and X'4F' are considered as the same value.

Characteristics

Length

If a bit string is used in table attributes or method declarations, you must specify the maximum length. The maximum length for a bit string is 1,073,741,823 bits.

Bit String Coercion

Automatic coercion is performed between a fixed-length and a variable-length bit string for comparison. For explicit coercion, use the **CAST** operator.

BIT(*n*)

Description

Fixed-length binary or hexadecimal bit strings are represented as **BIT(*n*)**, where *n* is the maximum number of bits. If *n* is not specified, the length is set to 1.

Remark

- *n* must be a number greater than 0.
- If the length of the string exceeds *n*, it will be processed as an error.
- If a bit string smaller than *n* is stored, the remainder of the string is filled with 0s.

Example

```
CREATE TABLE bit_tbl(a1 BIT, a2 BIT(1), a3 BIT(8), a4 BIT VARYING);
INSERT INTO bit_tbl VALUES (B'1', B'1', B'1', B'1');
INSERT INTO bit_tbl VALUES (0b1, 0b1, 0b1, 0b1);
INSERT INTO bit_tbl(a3,a4) VALUES (B'1010', B'1010');
INSERT INTO bit_tbl(a3,a4) VALUES (0xaa, 0xaa);
SELECT * FROM bit_tbl;
      a1                a2                a3                a4
=====
```

X'8'	X'8'	X'80'	X'8'
X'8'	X'8'	X'80'	X'8'
NULL	NULL	X'a0'	X'a'
NULL	NULL	X'aa'	X'aa'

BIT VARYING(n)

Description

A variable-length bit string is represented as **BIT VARYING(n)**, where *n* is the maximum number of bits. If *n* is not specified, the length is set to 1,073,741,823 (maximum value).

Remark

- If the length of the string exceeds *n*, it will be processed as an error.
- The remainder of the string is not filled with 0s even if a bit string smaller than *n* is stored.
- *n* must be a number greater than 0.

Example

```
CREATE TABLE bitvar_tbl(a1 BIT VARYING, a2 BIT VARYING(8));
INSERT INTO bitvar_tbl VALUES (B'1', B'1');
INSERT INTO bitvar_tbl VALUES (0b1010, 0b1010);
INSERT INTO bitvar_tbl VALUES (0xaa, 0xaa);
INSERT INTO bitvar_tbl(a1) VALUES (0xaaa);
SELECT * FROM bitvar_tbl;
```

a1	a2
X'8'	X'8'
X'a'	X'a'
X'aa'	X'aa'
X'aaa'	NULL

```
INSERT INTO bitvar_tbl(a2) VALUES (0xaaa);
ERROR: Data overflow coercing X'aaa' to type bit varying.
```

Character Strings

Definition and Characteristics

Definition

CUBRID supports the following four types of character strings:

- Fixed-length character string: **CHAR(n)**
- Variable-length character string: **VARCHAR(n)**
- Fixed-length national character string: **NCHAR(n)**
- Variable-length national character string: **NCHAR VARYING(n)**

The followings are the rules that are applied when using the character string types.

- In general, single quotations are used to enclose character string. Double quotations may be used as well depending on the value of **ansi_quotes**, which is a parameter related to SQL statement. If the **ansi_quotes** value is set to **no**, character string enclosed by double quotations is handled as character string, not as an identifier. The default value is **yes**. For details, [Statement/Type-Related Parameters](#).
- If there are characters that can be considered to be blank (e.g. spaces, tabs, or line breaks) between two character strings, these two character strings are treated as one according to ANSI standard. For example, the following example shows that a line break exists between two character string.

```
'abc'
'def'
```

- The two strings above are considered identical to one string below.

```
'abcdef'
```

- If you want to include a single quote as part of a character string, enter two single quotes in a row. For example, the character string on the left is stored as the one on the right.

```
' 'abcde' 'fg hij'                      'abcde'fg hij
```

- The maximum size of the token for all the character strings is 16 KB.
- National character strings are used to store national (except English alphabet) character strings in a multilingual environment. Note that **N** (uppercase) should be followed by a single quote which encloses character strings.

```
N'Härder'
```

Characteristics

Length

For a **CHAR** or **VARCHAR** type, specify the length (bytes) of a character string for a **NCHAR** or **NCHAR VARYING** type, specify the number of character strings (number of characters).

When the length of the character string entered exceeds the length specified, the characters in excess of the specified length are truncated.

For a fixed-length character string type such as **CHAR** or **NCHAR**, the length is fixed at the declared length. Therefore, the right part (trailing space) of the character string is filled with space characters when the string is stored. For a variable-length character string type such as **VARCHAR** or **NCHAR VARYING**, only the entered character string is stored, and the space is not filled with space characters.

The maximum length of a **CHAR** or **VARCHAR** type to be specified is 1,073,741,823 the maximum length of a **NCHAR** or **NCHAR VARYING** type to be specified is 536,870,911. The maximum length that can be input or output in a CSQL statement is 8,192 KB.

Character Set, charset

A character set (charset) is a set in which rules are defined that relate to what kind of codes can be used for encoding when specified characters (symbols) are stored in the computer.

CUBRID supports the following character sets and you can specify them as the **CUBRID_LANG** environment variable. You can store data in other character sets (e.g. utf-8), but string function or **LIKE** search are not supported.

Character Set	CUBRID_LANG
8-bits ISO 8859-1 Latin	en_US
KSC 5601-1992 (EUC-KR)	ko_KR.euckr

Any characters from the above character sets can be included in a character string (the **NULL** character is represented as '\0').

Collating Character Sets

A collation is a set of rules used for comparing characters to search or sort values stored in the database when a certain character set is specified. Therefore, such rules are applied only to character string data types such as **CHAR** or **VARCHAR**. For a national character string type such as **NCAHR()** or **NCHAR VARYING()**, the sorting rules are determined according to the encoding algorithm of the specified character set.

Character String Coercion

Automatic coercion takes place between a fixed-length and a variable-length character string for the comparison of two characters, applicable only to characters that belong to the same character set. For example, when you extract a column value from a **CHAR(5)** data type and insert it into a column with a **CHAR(10)** data type, the data type is automatically coerced to **CHAR(10)**. If you want to coerce a character string explicitly, use the **CAST** operator (See [CAST Operator](#)).

CHAR(*n*)

Description

A fixed-length character string is represented as **CHAR(*n*)**, in which *n* is the number of bytes in an ASCII character string. For the English alphabet, each character takes up one byte. However, for Korean characters, note that the number of bytes taken up by each character differs depending on the character set of the data input environment (e.g. EUC-KR: 2 bytes, utf-8: 3 bytes). If *n* is not specified, the length is set to the default value 1.

When the length of a character string exceeds *n*, they are truncated. When character string which is shorter than *n* is stored, whitespace characters are used to fill up the trailing space.

CHAR(*n*) and **CHARACTER(*n*)** are used interchangeably.

Remark

- The **CHAR** data type is always based on the ISO 8859-1 Latin character set.
- *n* is an integer between 1 and 1,073,741,823 (1G).
- Empty quotes (' ') are used to represent a blank string. In this case, the return value of the **LENGTH** function is not 0, but is the fixed length defined in **CHAR(*n*)**. That is, if you enter a blank string into a column with **CHAR(10)**, the **LENGTH** is 10; if you enter a blank value into a **CHAR** with no length specified, the **LENGTH** is the default value 1.
- Space characters used as filling characters are considered to be smaller than any other characters, including special characters.

Example 1

```
If you specify 'pacesetter' as CHAR(12), 'pacesetter ' is stored (a 10-character string plus two whitespace characters).
If you specify 'pacesetter ' as CHAR(10), 'pacesetter' is stored (a 10-character string; two whitespace characters are truncated).
If you specify 'pacesetter' as CHAR(4), 'pace' is stored (truncated as the length of the character string is greater than 4).
If you specify 'p ' as CHAR, 'p' is stored (if n is not specified, the length is set to the default value 1).
```

Example 2

```
If you specify '큐브리드' as CHAR(10) in the EUC-KR encoding, it is processed normally.
If you specify '큐브리드' as CHAR(10) and use the CHAR_LENGTH() function in the EUC-KR encoding, 10 is stored.
If you specify '큐브리드' as CHAR(10) in the utf-8 encoding, the last character is broken (one Korean character takes up three bytes in the utf-8 encoding so it requires two more bytes).
If you specify '큐브리드' as CHAR(12) in the utf-8 encoding, it is processed normally.
```

VARCHAR(*n*)/CHAR VARYING(*n*)

Description

Variable-length character strings are represented as **VARCHAR(*n*)**, where *n* is the maximum number of ASCII character strings. Each English character takes up one byte. For Korean characters, note that the number of bytes taken up by each character differs depending on the character set of the data input environment (e.g. EUC-KR: 2 bytes, utf-8: 3 bytes). If *n* is not specified, the length is set to the maximum length of 1,073,741,823.

When the length of a character string exceeds *n*, they are truncated. When character string which is shorter than *n* is stored, whitespace characters are used to fill up the trailing space; for **VARCHAR(*n*)**, the length of string used are stored.

VARCHAR(*n*), **CHARACTER**, **VARYING(*n*)**, and **CHAR VARYING(*n*)** are used interchangeably.

Remark

- **STRING** is the same as the **VARCHAR** (maximum length).
- *n* is an integer between 1 and 1,073,741,823 (1G).
- Empty quotes (' ') are used to represent a blank string. In this case, the return value of the **LENGTH** function is not 0.

Example 1

If you specify 'pacesetter' as CHAR(4), 'pace' is stored (truncated as the length of the character string is greater than 4).
 If you specify 'pacesetter' as VARCHAR(12), 'pacesetter' is stored (a 10-character string).
 If you specify 'pacesetter ' as VARCHAR(12), 'pacesetter ' is stored (a 10-character string plus two whitespace characters).
 If you specify 'pacesetter ' as VARCHAR(10), 'pacesetter' is stored (a 10-character string; two whitespace characters are truncated).
 If you specify 'p ' as VARCHAR, 'p' is stored (if *n* is not specified, the default value 1,073,741,823 is used, and the trailing space is not filled with whitespace characters).

Example 2

If you specify '큐브리드' as VARCHAR(10) in the EUC-KR encoding, it is processed normally.
 If you specify '큐브리드' as CHAR(10) and then use CHAR_LENGTH() function in the EUC-KR encoding, 8 is stored.
 If you specify '큐브리드' as CHAR(10) in the utf-8 encoding, the last character is broken (one Korean character takes up three bytes in the utf-8 encoding so it requires two more bytes).
 If you specify '큐브리드' as VARCHAR(12) in the utf-8 encoding, it is processed normally.

STRING**Description**

STRING is a variable-length character string data type. **STRING** is the same as the [VARCHAR](#) with the length specified to the maximum value. That is, **STRING** and **VARCHAR**(1,073,741,823) have the same value.

NCHAR(n)**Description**

NCHAR(*n*) is used to store non-English character strings. It can be used only for character sets supported by CUBRID described above. *n* is the number of characters. If *n* is omitted, the length is specified as the default value 1. When the length of a character string exceeds *n*, they are truncated. When character string which is shorter than *n* is stored, whitespace characters are used to fill up the space.

To store a Korean character string as a national character string type, you must set the locale of the operating system to Korean, or set the value of the **CUBRID_LANG** environment variable to **ko_KR.euckr** before creating the table.

Remark

- *n* is an integer between 1 and 536,870,911.
- The number of national character sets that can be used in a single database is set to be one. For example, 8-bit ISO 8889-1 (Latin-1) and EUC code sets cannot be used simultaneously in the same database.
- An error occurs if a non-national character string (whether it is fixed-length or variable-length) is specified for an attribute declared as a national character string.
- Using two different character code sets at once also causes an error.

Example

If you specify '큐브리드' as NCHAR(5) in the EUC-KR encoding, it is processed normally.

If you specify '큐브리드' as NCHAR(5) and then use the CHAR_LENGTH() function in the EUC-KR encoding, 5 is stored.
 If you specify '큐브리드' as NCHAR(5) in the utf-8 encoding, an error occurs (utf-8 character set is not supported).

NCHAR VARYING(n)

Description

NCHAR VARYING(n) is a variable-length character string type. For details, see description and note of [NCHAR\(n\)](#). The difference is that the right part (trailing space) of the character string is not filled with whitespace characters, even when the number of strings is smaller than n.

NCHAR VARYING(n), NATIONAL CHAR VARYING(n), and NATIONAL CHARACTER VARYING(n) are used interchangeably.

Example

If you specify '큐브리드' as NCHAR VARYING(5) in the EUC-KR encoding, it is processed normally.
 If you specify '큐브리드' as NCHAR VARYING(5) and then use CHAR_LENGTH() function in the EUC-KR encoding, 4 is stored.
 If you specify '큐브리드' as HCHAR VARYING(5) in the utf-8 encoding, an error occurs (utf-8 character set is not supported).

Escape Special Characters

Description

CUBRID supports two kinds of methods to escape special characters. One is using quotes and the other is using backslash (\).

Escape with Quotes

If you set **no** for the system parameter **ansi_quotes** in the **cubrid.conf** file, you can use both double quotes (") and single quotes (') to wrap strings. The default value for the **ansi_quotes** parameter is **yes**, and you can use only single quotes to wrap the string. The numbers 2 and 3 below are applied only if you set for the **ansi_quotes** parameter to **no**.

- You should use two single quotes (') for the single quotes included in the strings wrapped in single quotes.
- You should use two double quotes (") for the double quotes included in the strings wrapped in double quotes.
- You don't need to escape the single quotes included in the string wrapped in double quotes.
- You don't need to escape the double quotes included in the string wrapped in single quotes.

Escape with Backslash

You can use escape using backslash (\) only if you set no for the system parameter **no_backslash_escapes** in the **cubrid.conf** file. The default value for the **no_backslash_escapes** parameter is **yes**. Depending on the input value, the following are the special characters.

- \ : Single quotes (')
- \" : Double quotes (")
- \n : Newline, linefeed character
- \r : Carriage return character
- \t : Tab character
- \\ : Backslash
- \% : Percent sign (%). For details, see the following description.
- _ : Underbar (_). For details, see the following description.

For all other escapes, the backslash will be ignored. For example, \"x\" is the same as entering only "x".

\% and _ are used in the pattern matching syntax such as **LIKE** to search percent signs and underbars and are used as a wildcard character if there is no backslash. Outside of the pattern matching syntax, "% and "_" are recognized as normal strings not wildcard characters. For details, see [LIKE Predicate](#).

Example 1

The following is the result of executing Escape if a value for the system parameter **ansi_quotes** in the **cubrid.conf** file is no, and a value for **no_backslash_escapes** is no.

```
SELECT STRCMP('single quotes test(')', 'single quotes test(\')');

      strcmp('single quotes test(')', 'single quotes test(\')')
=====
                                                                0

SELECT STRCMP("\a\b\c\d\e\f\g|h|i\j\k\l|m\n\o\p\q\r\s\t\u\v\w\x\y\z",
"a\b\c\defghijklm\nopq\r\s\tuvwxyz");

      strcmp('abcde fghijklm
s          uvwxyz', 'abcde fghijklm
s          uvwxyz')
=====
                                                                0

SELECT LENGTH('\');

      char length('\')
=====
                                                                1
```

Example 2

The following is the result of executing Escape if a value for the system parameter **ansi_quotes** in the **cubrid.conf** file is yes, and a value for **no_backslash_escapes** is yes.

```
SELECT STRCMP('single quotes test(')', 'single quotes test(\')');

In the command from line 2,
ERROR: unterminated string

In the command from line 2,
ERROR: syntax error, unexpected UNTERMINATED STRING

SELECT STRCMP("\a\b\c\d\e\f\g|h|i\j\k\l|m\n\o\p\q\r\s\t\u\v\w\x\y\z",
"a\b\c\defghijklm\nopq\r\s\tuvwxyz");

In line 1, column 18,
ERROR: [\a\b\c\d\e\f\g|h|i\j\k\l|m\n\o\p\q\r\s\t\u\v\w\x\y\z] is not defined.

In line 1, column 18,
ERROR: [a\b\c\defghijklm\nopq\r\s\tuvwxyz] is not defined.

SELECT LENGTH('\');

      char_length('\')
=====
                                                                2
```

Example 3

The following is the result of executing Escape if a value for the system parameter **ansi_quotes** in the **cubrid.conf** file is yes, and a value for **no_backslash_escapes** is no.

```
CREATE TABLE t1 (a varchar(200));
INSERT INTO t1 VALUES ('aaabbb'), ('aaa%');

SELECT a FROM t1 WHERE a LIKE 'aaa%' escape '\';

a
```

```
=====
'aaa%'
```

BLOB/CLOB Data Types

Definition and Characteristics

Definition

An External **LOB** type is data to process Large Object, such as text or images. When LOB-type data is created and inserted, it will be stored in a file to an external storage, and the location information of the relevant file (**LOB Locator**) will be stored in the CUBRID database. If the **LOB Locator** is deleted from the database, the relevant file that was stored in the external storage will be deleted as well. CUBRID supports the following two types of **LOB**:

- Binary Large Object (**BLOB**)
- Character Large Object (**CLOB**)

Related Terms

- **LOB** (Large Object) : Large-sized objects such as binaries or text.
- **FBO** (File Based Object) : An object that stores data of the database in an external file.
- **External LOB** : An object better known as FBO, which stores **LOB** data in a file into an external DB. It is supported by CUBRID. Internal **LOB** is an object that stores **LOB** data inside the DB.
- **External Storage** : An external storage to store LOB (example : POSIX file system).
- **LOB Locator** : The path name of a file stored in external storage.
- **LOB Data** : Details of a file in a specific location of LOB Locator.

File Names

When storing LOB data in external storage, the following naming convention will be applied:

```
{table_name}_{unique_name}
```

- *table_name* : It is inserted as a prefix and able to store the **LOB** data of many tables in one external storage.
- *unique_name* : The random name created by the DB server.

Default Storage

- **LOB** data is stored in the local file system of the DB server. LOB data is stored in the path specified in the **-lob-base-path option** value of **cubrid createdb**; if this value is omitted, the data will be stored in the [db-vol path]/lob path where the database volume will be created. For more details, see [Database Creation](#) and [Storage Creation and Management](#).
- If the relevant path is deleted despite a **LOB** data file path being registered in the database location file (**databases.txt**), please note that the utility that operates in database server (**cub_server**) and standalone will not function normally.

BLOB/CLOB

BLOB

- A type that stores binary data outside the database.
- The maximum length of **BLOB** data is the maximum file size creatable in an external storage.
- In SQL statements, the **BLOB** type expresses the input and output value in a bit array. That is, it is compatible with the **BIT(n)** and **BIT VARYING(n)** types, and only an explicit type change is allowed. If data lengths differ from one another, the maximum length is truncated to fit the smaller one.
- When converting the **BLOB** type value to a binary value, the length of the converted data cannot exceed 1GB. When converting binary data to the **BLOB** type, the size of the converted data cannot exceed the maximum file size provided by the **BLOB** storage.

CLOB

- A type that stores character string data outside the database.
- The maximum length of **CLOB** data is the maximum file size creatable in an external storage.
- In SQL statements, the CLOB type expresses the input and output value in a character string. That is, it is compatible with the **CHAR(n)**, **VARCHAR(n)**, **NCHAR(n)**, and **NCHAR VARYING(n)** types. However, only an explicit type change is allowed, and if data lengths are different from one another, the maximum length is truncated to fit to the smaller one.
- When converting the **CLOB** type value to a character string, the length of the converted data cannot exceed 1 GB. When converting a character string to the **CLOB** type, the size of the converted data cannot exceed the maximum file size provided by the **CLOB** storage.

Creating and Altering Columns

Description

BLOB/CLOB type columns can be created/added/deleted by using a **CREATE TABLE** statement or an **ALTER TABLE** statement.

Note

- You cannot create the index file for a **LOB** type column.
- You cannot define the **PRIMARY KEY**, **FOREIGN KEY**, **UNIQUE**, and **NOT NULL** constraints for a **LOB** type column. However, **SHARED** property cannot be defined and **DEFAULT** property can only be defined by the **NULL** value.
- **LOB** type column/data cannot be the element of collection type data.
- If you are deleting a record containing a **LOB** type column, all files located inside a **LOB** column value (Locator) and the external storage will be deleted. When a record containing a **LOB** type column is deleted in a basic key table, and a record of a foreign key table that refers to the foregoing details is deleted at once, all **LOB** files located in a **LOB** column value (Locator) and the external storage will be deleted. However, if the relevant table is deleted by using a **DROP TABLE** statement, or a **LOB** column is deleted by using an **ALTER TABLE...DROP** statement, only a **LOB** column value (**LOB** Locator) is deleted, and the **LOB** files inside the external storage which a **LOB** column refers to will not be deleted.

Example

```
-- creating a table and CLOB column
CREATE TABLE doc_t (doc_id VARCHAR(64) PRIMARY KEY, content CLOB);

-- an error occurs when UNIQUE constraint is defined on CLOB column
ALTER TABLE doc t ADD CONSTRAINT content unique UNIQUE(content);

-- an error occurs when creating an index on CLOB column
CREATE INDEX ON doc_t (content);

-- creating a table and BLOB column
CREATE TABLE image t (image id VARCHAR(36) PRIMARY KEY, doc id VARCHAR(64) NOT NULL, image
BLOB);

-- an error occurs when adding a BOLB column with NOT NULL constraint
ALTER TABLE image t ADD COLUMN thumbnail BLOB NOT NULL;

-- an error occurs when adding a BLOB column with DEFAULT attribute
ALTER TABLE image_t ADD COLUMN thumbnail2 BLOB DEFAULT BIT_TO_BLOB('X'010101');
```

Storing and Updating Columns

Description

In a **BLOB/CLOB** type column, each **BLOB/CLOB** type value is stored, and if binary or character string data is input, you must explicitly change the types by using each **BIT_TO_BLOB()**/**CHAR_TO_CLOB()** function.

If a value is input in a **LOB** column by using an **INSERT** statement, a file is created in an external storage internally and the relevant data is stored; the relevant file path (Locator) is stored in an actual column value.

If a record containing a **LOB** column uses a **DELETE** statement, a file to which the relevant **LOB** column refers will be deleted simultaneously. If a **LOB** column value is changed using an **UPDATE** statement, the column value will be changed following the operation below, according to whether a new value is **NULL** or not.

- If a **LOB** type column value is changed to a value that is not **NULL** : If a Locator that refers to an external file is already available in a **LOB** column, the relevant file will be deleted. A new file is created afterwards. After storing a value that is not **NULL**, a Locator for a new file will be stored in a **LOB** column value.
- If changing a **LOB** type column value to **NULL** : If a Locator that refers to an external file is already available in a **LOB** column, the relevant file will be deleted. And then **NULL** is stored in a **LOB** column value.

Example

```
-- inserting data after explicit type conversion into CLOB type column
INSERT INTO doc t (doc id, content) VALUES ('doc-1', CHAR TO CLOB('This is a Dog'));
INSERT INTO doc t (doc id, content) VALUES ('doc-2', CHAR TO CLOB('This is a Cat'));

-- inserting data after explicit type conversion into BLOB type column
INSERT INTO image t VALUES ('image-0', 'doc-0', BIT TO BLOB(X'000001'));
INSERT INTO image t VALUES ('image-1', 'doc-1', BIT TO BLOB(X'000010'));
INSERT INTO image t VALUES ('image-2', 'doc-2', BIT TO BLOB(X'000100'));

-- inserting data from a sub-query result
INSERT INTO image_t SELECT 'image-1010', 'doc-1010', image FROM image_t WHERE image_id =
'image-0';

-- updating CLOB column value to NULL
UPDATE doc t SET content = NULL WHERE doc id = 'doc-1';

-- updating CLOB column value
UPDATE doc_t SET content = CHAR_TO_CLOB('This is a Dog') WHERE doc_id = 'doc-1';

-- updating BLOB column value
UPDATE image t SET image = (SELECT image FROM image t WHERE image id = 'image-0') WHERE
image id = 'image-1';

-- deleting BLOB column value and its referencing files
DELETE FROM image_t WHERE image_id = 'image-1010';
```

Getting Column Values

Description

When you get a **LOB** type column, the data stored in a file to which the column refers will be displayed. You can execute an explicit type change by using **CAST** operator, **CLOB_TO_CHAR()** function, and **BLOB_TO_BIT()** function.

Remark

- If the query is executed in CSQL, a column value (Locator) will be displayed, instead of the data stored in a file. To display the data to which a **BLOB/CLOB** column refers, it must be changed to strings by using **CLOB_TO_CHAR()** function.
- To use the string process function, the strings need to be converted by using the **CLOB_TO_CHAR()** function.
- You cannot specify a **LOB** column in **GROUP BY** clause and **ORDER BY** clause.
- Comparison operators, relational operators, **IN**, **NOT IN** operators cannot be used to compare **LOB** columns. However, **IS NULL** expression can be used to compare whether it is a **LOB** column value (Locator) or **NULL**. This means that **TRUE** will be returned when a column value is **NULL**, and if a column value is **NULL**, there is no file to store **LOB** data.
- When a **LOB** column is created, and the file is deleted after data input, a **LOB** column value (Locator) will become a state that is referring to an invalid file. As such, using **CLOB_TO_CHAR()**, **BLOB_TO_BIT()**, **CLOB_LENGTH()**, and **BLOB_LENGTH()** functions on the columns that have mismatching **LOB** Locator and a **LOB** data file enables them to display **NULL**.

Example

```
-- displaying locator value when selecting CLOB and BLOB column in CSQL interpreter
```

```

SELECT doc t.doc id, content, image FROM doc t, image t WHERE doc t.doc id =
image t.doc id;

  doc_id          content          image
=====
'doc-1'          file:/home1/data1/ces 658/doc t.00001282208855807171 7329 file:/
home1/data1/ces 318/image t.00001282208855809474 7474
'doc-2'          file:/home1/data1/ces 180/doc t.00001282208854194135 5598 file:/
home1/data1/ces_519/image_t.00001282208854205773_1215

2 rows selected.

-- using string functions after coercing its type by CLOB TO CHAR( )
SELECT CLOB TO CHAR(content), SUBSTRING(CLOB TO CHAR(content), 10) FROM doc t;

  clob_to_char(content)  substring( clob_to_char(content) from 10)
=====
'This is a Dog'        ' Dog'
'This is a Cat'        ' Cat'

2 rows selected.

SELECT CLOB TO CHAR(content) FROM doc t WHERE CLOB TO CHAR(content) LIKE '%Dog%';

  clob to char(content)
=====
'This is a Dog'

SELECT CLOB TO CHAR(content) FROM doc t ORDER BY CLOB TO CHAR(content)

  clob to char(content)
=====
'This is a Cat'
'This is a Dog'

-- an error occurs when LOB column specified in WHERE/ORDER BY/GROUP BY clauses
SELECT * FROM doc t WHERE content LIKE 'This%';
SELECT * FROM doc t ORDER BY content;

```

Functions and Operators

CAST Operator

By using **CAST** operator, you can execute an explicit type change between **BLOB/CLOB** type and binary type/string type. For more details, see [CAST Operator](#).

Syntax

```

CAST (<bit_type_column_or_value> AS CLOB)
CAST (<bit_type_column_or_value> AS BLOB)
CAST (<char_type_column_or_value> AS BLOB)
CAST (<char_type_column_or_value> AS CLOB)

```

LOB Data Process and Type Change Functions

The next table shows the functions provided to process and change BLOB/CLOB types.

Function Expression	Description
CLOB_TO_CHAR (<lob_type_column>)	Changes number type, date/time type, and CLOB type to VARCHAR type.
BLOB_TO_BIT (<blob_type_column>)	Changes BLOB type to VARYING BIT type.
CHAR_TO_CLOB (<char_type_column_or_value>)	Changes text string type (CHAR , VARCHAR , NCHAR , NVARCHAR) to CLOB type.
BIT_TO_BLOB (<blob_type_column_or_value>)	Changes bit array type (BIT , VARYING BIT) to BLOB type.

CHAR_TO_BLOB (<i><char_type_column_or_value></i>)	Changes text string type (CHAR , VARCHAR , NCHAR , NVARCHAR) to BLOB type.
CLOB_FROM_FILE (<i><file_pathname></i>)	Reads file details from the file path of VARCHAR type and changes to CLOB type data. <i><file_pathname></i> is analyzed to a path of server which is operated by the DB client, such as CAS or CSQL. If a path is specified targeting this, the upper path will be the current work direction of the process. The statement that calls this function will not cache execution plans.
BLOB_FROM_FILE (<i><file_pathname></i>)	Reads file details from the file path of VARCHAR type, and changes to BLOB type data. The file path specified in is interpreted using the same method as the CLOB_FROM_FILE () function.
CLOB_LENGTH (<i><lob_column></i>)	Returns the length of LOB data stored in a CLOB file in bytes.
BLOB_LENGTH (<i><blob_column></i>)	Returns the length of LOB data stored in a BLOB file in bytes.
<i><blob_or_clob_column></i> IS NULL	Use an IS NULL expression to compare whether it is a LOB column value (Locator) or NULL ; returns TRUE if NULL .

Creating and Managing Storage

LOB File Path Specification

By default, the **LOB** data file is stored in the *<db-volume-path>/lob* directory where database volume is created. However, if the **--lob-base-path** option of **cubrid createdb** utility is used when creating the database, a **LOB** data file can be stored in the directory specified by option value. However, if there is no directory specified by option value, attempt to create a directory, and display an error message if it fails to create the directory. For more details, see the **--lob-base-path** option in [Creating Database](#).

```
# image_db volume is created in the current work directory, and a LOB data file will be
stored.
cubrid createdb image db

# LOB data file is stored in the "/home1/data1" path within a local file system.
cubrid createdb --lob-base-path="file:/home1/data1" image_db
```

Checking LOB File Store Directory

```
# You can check a directory where a LOB file will be stored by executing the cubrid
spacedb utility.
cubrid spacedb image_db

Space description for database 'image db' with pagesize 16.0K. (log pagesize: 16.0K)

Volid Purpose total size free size Vol Name
    0  GENERIC      512.0M    510.1M  /home1/data1/image_db
Space description for temporary volumes for database 'image db' with pagesize 16.0K.

Volid Purpose total size free size Vol Name
LOB space description file:/home1/data1
```

Changing or Expanding LOB File Store Directory

Secure disk space to create additional file storage, expand the **lob-base-path** of **databases.txt**, and change to the disk location. Restart the database server to apply the changes made to **databases.txt**. However, even if you change the **lob-base-path** of **databases.txt**, access to the **LOB** data stored in a previous storage is possible.

```
# You can change to a new directory from the lob-base-path of databases.txt file.
sh> cat $CUBRID DATABASES/databases.txt
#db-name      vol-path      db-host      log-path      lob-base-path
image_db      /home1/data1  localhost    /home1/data1  file:/home1/data2
```

Backing up and Recovering of LOB Files

While backup/recovery is not supported for **LOB** type columns, meta data (Locator) of the **LOB** type columns is supported with such service.

Copying Database with LOB Files

If you are copying a database by using the **cubrid copydb** utility, you must configure the **databases.txt** additionally, as the **LOB** file directory path will not be copied if the related option is not specified. For more details, see the **-B** and **--copy-lob-path** options in [Copying/Moving Database](#).

Supporting and Recovering Transactions

Description

Commit/rollback for **LOB** data changes are supported. That is, it ensures the validation of mapping between **LOB** Locator and actual **LOB** data within transactions, and it supports recovery during DB errors. This means that an error will be displayed in case of mapping errors between **LOB** Locator and **LOB** data due to the rollback of the relevant transactions, as the database is terminated during transactions. See the example below.

Example

```
;AUTOCOMMIT OFF

CREATE TABLE doc t (doc id VARCHAR(64) PRIMARY KEY, content CLOB);
INSERT INTO doc_t VALUES ('doc-10', CHAR_TO_CLOB('This is content'));
COMMIT;
UPDATE doc t SET content = CHAR TO CLOB('This is content 2') where doc id = 'doc-10';
ROLLBACK;
SELECT doc id, CLOB TO CHAR(content) FROM doc t WHERE doc id = 'doc-10';
  doc id  content
=====
'doc-10' 'This is content '

INSERT INTO doc t VALUES ('doc-11', CHAR_TO_CLOB ('This is content'));
COMMIT;
UPDATE doc_t SET content = CHAR_TO_CLOB('This is content 3') WHERE doc_id = 'doc-11';

-- system crash occurred and then restart server
SELECT doc id, CLOB TO CHAR(content) FROM doc t WHERE doc id = 'doc-11';

-- Error : LOB Locator references to the previous LOB data because only LOB Locator is
rollbacked.
```

Remark

- When selecting **LOB** data in an application through a driver such as JDBC, the driver can get **ResultSet** from DB server and fetch the record while changing the cursor location on **ResultSet**. That is, only Locator, the meta data of a **LOB** column, is stored at the time when **ResultSet** is imported, and **LOB** data that is referred by a File Locator will be fetched from the file Locator at the time when a record is fetched. Therefore, if **LOB** data is updated between two different points of time, there could be an error, as the mapping of **LOB** Locator and actual **LOB** data will be invalid.
- Since backup/recovery is supported only for meta data (Locator) of the **LOB** type columns, an error is likely to occur, as the mapping of **LOB** Locator and **LOB** data is invalid if recovery is performed based on a specific point of time.

- TO execute **INSERT** the **LOB** data into other device, LOB data referred by the meta data (Locator) of a **LOB** column must be read.
- In a CUBRID HA environment, the meta data (Locator) of a **LOB** column is replicated and data of a **LOB** type is not replicated. Therefore, if storage of a **LOB** type is located on the local machine, no tasks on the columns in a slave node or a master node after failover are allowed.

Caution Up to CUBRID 2008 R3.0, Large Objects are processed by using **glo** (Generalized Large Object) classes. However, the **glo** classes has been deprecated since the CUBRID 2008 R3.1. Instead of it, **LOB/CLOB** data type is supported. Therefore, both DB schema and application must be modified when upgrading CUBRID in an environment using the previous version of **glo** classes.

Collection Types

Definition and Characteristics

Definition

Allowing multiple data values to be stored in a single attribute is an extended feature of relational database. Elements of a collection are possible to have different domain each other. The domain can be one of the primitive data types or classes excluding virtual classes. For example, **SET** (INTEGER, tbl_1) can specify an integer or a set of row values of the user-defined class tbl_1 as a domain. When a domain list is not specified (e.g. **SET** ()), all data types are allowed as elements including user-defined classes.

The data of a collection-type column with at least two domain lists can be retrieved by using the **csql** utility or the C-API. It cannot be retrieved in CUBRID manager or CUBRID API (JDBC, ODBC, OLEDB, PHP, CCI).

Collection Types Supported by CUBRID

Type	Description	Definition	Input Data	Stored Data
SET	A union which does not allow duplicates	col_name SET VARCHAR(20) col_name SET (int, VARCHAR(20))	{'c','c','c','b','b','a'} {3,3,3,2,2,1,0,'c','c','c','b','b','a'}	{'a','b','c'} {0,1,2,3,'a','b','c'}
MULTISET	A union which allows duplicates	col_name MULTISET VARCHAR(20) col_name MULTISET (int, VARCHAR(20))	{'c','c','c','b','b','a'} {3,3,3,2,2,1,0,'c','c','c','b','b','a'}	{'a','b','b','c','c','c'} {0,1,2,2,3,3,3,'a','b','b','c','c','c'}
LIST SEQUENCE SEQUENCE	A union which allows duplicates and stores data in the order of input	col_name LIST VARCHAR(20) col_name LIST (int, VARCHAR(20))	{'c','c','c','b','b','a'} {3,3,3,2,2,1,0,'c','c','c','b','b','a'}	{'c','c','c','b','b','a'} {3,3,3,2,2,1,0,'c','c','c','b','b','a'}

As you see the table above, the value specified as a collection type can be inputted with braces ('{', '}') each value is separated with a comma (,).

Characteristics

Coercions

If the specified domains are identical, the collection types can be cast explicitly by using the **CAST** operator. The following table shows the collection types that allow explicit coercions.

Explicit Coercions

	TO			
FROM	SET	MULTISET	LIST	
SET	-	O	O	
MULTISET	O	-	X	
LIST	O	O	-	

SET**Description**

SET is a collection type in which each element has different values. Elements of a **SET** can have many different data types or even instances of different classes.

Example

```
CREATE TABLE set_tbl ( col_1 set(int, CHAR(1)));
INSERT INTO set_tbl VALUES ({3,3,3,2,2,1,0,'c','c','c','b','b','a'});
INSERT INTO set_tbl VALUES ({NULL});
INSERT INTO set_tbl VALUES ({});
SELECT * FROM set_tbl;
  col_1
=====
{0, 1, 2, 3, 'a', 'b', 'c'}
{NULL}
{' '}

SELECT CAST(col_1 AS MULTISET), CAST(col_1 AS LIST) FROM set_tbl;
  cast(col_1 as multiset)  cast(col_1 as sequence)
=====
{0, 1, 2, 3, 'a', 'b', 'c'} {0, 1, 2, 3, 'a', 'b', 'c'}
{NULL} {NULL}
{' '} {' '}

INSERT INTO set_tbl VALUES ({});
ERROR: Cannot coerce '' to type set.
```

MULTISET**Description**

MULTISET is a collection type in which duplicated elements are allowed. Elements of a **MULTISET** can have many different data types or even instances of different classes.

Example

```
CREATE TABLE multiset_tbl ( col_1 multiset(int, CHAR(1)));
INSERT INTO multiset_tbl VALUES ({3,3,3,2,2,1,0,'c','c','c','b','b','a'});
SELECT * FROM multiset_tbl;
  col_1
=====
{0, 1, 2, 2, 3, 3, 3, 'a', 'b', 'b', 'c', 'c', 'c'}

SELECT CAST(col_1 AS SET), CAST(col_1 AS LIST) FROM multiset_tbl;
  cast(col_1 as set)  cast(col_1 as sequence)
=====
{0, 1, 2, 3, 'a', 'b', 'c'} {3, 3, 3, 2, 2, 1, 0, 'c', 'c', 'c', 'b', 'b', 'a'}
{' '}
```

LIST/SEQUENCE

Description

LIST (=SEQUENCE) is a collection type in which the input order of elements is preserved, and duplications are allowed. Elements of a **LIST** can have many different data types or even instances of different classes.

Example

```
CREATE TABLE list_tbl ( col_1 list(int, CHAR(1)));
INSERT INTO list_tbl VALUES ({3,3,3,2,2,1,0,'c','c','c','b','b','a'});
SELECT * FROM list_tbl;
  col_1
=====
{3, 3, 3, 2, 2, 1, 0, 'c', 'c', 'c', 'b', 'b', 'a'}

SELECT CAST(col_1 AS SET), CAST(col_1 AS MULTISSET) FROM list_tbl;
  cast(col_1 as set)  cast(col_1 as multiset)
=====
{0, 1, 2, 3, 'a', 'b', 'c'} {0, 1, 2, 2, 3, 3, 3, 'a', 'b', 'b', 'c', 'c', 'c'
'}
```

Implicit Type Conversion

Rules

An implicit type conversion represents an automatic conversion of a type of expression to a corresponding type. **SET**, **MULTISET**, **LIST** and **SEQUENCE** should be converted explicitly.

If you convert the **DATETIME** and **TIMESTAMP** types to the **DATE** type or **TIME** type, data loss may occur. If you convert the **DATE** type to the **DATETIME** type or **TIMESTAMP** type, the time will be set to '12:00:00: AM.'

If you convert a string type or an exact numeric type to a floating-point numeric type, it may not be accurate. Because a string type and an exact type use a decimal precision to represent the value, but a floating-point numeric type uses a binary precision.

The implicit type conversion executed by CUBRID is as follows:

Implicit Type Conversion Table 1

From \ To	DATETIME	DATE	TIME	TIMESTAMP	DOUBLE	FLOAT	NUMERIC	BIGINT
DATETIME	-	O	O	O				
DATE	O	-		O				
TIME			-					
TIMESTAMP	O	O	O	-				
DOUBLE					-	O	O	O
FLOAT					O	-	O	O
NUMERIC					O	O	-	O
BIGINT					O	O	O	-
INT				O	O	O	O	O
SHORT					O	O	O	O
MONETARY					O	O	O	O
BIT								
VARBIT								
CHAR	O	O	O	O	O	O	O	O

VARCHAR	0	0	0	0	0	0	0	0
NCHAR	0	0	0	0	0	0	0	0
VARNCHAR	0	0	0	0	0	0	0	0

Implicit Type Conversion Table 2

From \ To	INT	SHORT	MONETARY	BIT	VARBIT	CHAR	VARCHAR	NCHAR	VARNCHAR
DATETIME					0	0	0	0	
DATE					0	0	0	0	
TIME					0	0	0	0	
TIMESTAMP					0	0	0	0	
DOUBLE	0	0	0		0	0	0	0	
FLOAT	0	0	0		0	0	0	0	
NUMERIC	0	0	0		0	0	0	0	
BIGINT	0	0	0		0	0	0	0	
INT	-	0	0		0	0	0	0	
SHORT	0	-	0		0	0	0	0	
MONETARY	0	0	-		0	0	0	0	
BIT				-	0	0	0	0	
VARBIT				0	-	0	0	0	0
CHAR	0	0	0	0	0	-	0	0	0
VARCHAR	0	0	0	0	0	0	-	0	0
NCHAR	0	0	0	0	0	0	0	-	0
VARNCHAR	0	0	0	0	0	0	0	0	-

INSERT and UPDATE

The type will be converted to the type of the column affected.

```
CREATE TABLE t(i INT);
INSERT INTO t VALUES ('123');

SELECT * FROM t;

      i
=====
     123
```

Function

If the parameter value entered in the function can be converted to the specified type, the parameter type will be converted. The strings are converted to numbers because the input parameter expected in the following function is a number.

```
SELECT MOD('123', '2');

      mod('123', '2')
=====
1.0000000000000000e+00
```

You can enter multiple type values in the function. If the type value not specified in the function is delivered, the type will be converted depending on the following priority order.

- Date/Time Type (**DATETIME** > **TIMESTAMP** > **DATE** > **TIME**)

- Approximate Numeric Type (**MONETARY** > **DOUBLE** > **FLOAT**)
- Exact Numeric Type (**NUMERIC** > **BIGINT** > **INT** > **SHORT**)
- String Type (**CHAR/NCHAR** > **VARCHAR/VARNCHAR**)

Comparison Operation

The following are the conversion rules according to an operand type of the comparison operator.

operand1 Type	operand2 Type	Conversion	Comparison
Numeric Type	Numeric Type	None	NUMERIC
	String Type	Converts operand2 to DOUBLE	NUMERIC
	Date/Time Type	None	N/a
String Type	Numeric Type	Converts operand1 to DOUBLE	NUMERIC
	String Type	None	String
	Date/Time Type	Converts operand1 to date/time type	Date/Time
Date/Time Type	Numeric Type	None	N/A
	String Type	Converts operand2 to date/time type	Date/Time
	Date/Time Type	Converts it to the type with higher priority	Date/Time

The following are the exceptions in the conversion rules for comparison operators:

- COLUMN <operator> value

operand1 Type	operand2 Type	Conversion	Comparison
String type	Numeric type	Converts operand2 to the string type	String
	Date/Time type	Converts operand2 to the string type	String

If operand2 is a set operator(**IS IN**, **IS NOT IN**, **= ALL**, **= ANY**, **< ALL**, **< ANY**, **<= ALL**, **<= ANY**, **>= ALL**, **>= ANY**), the exception above is not applied.

Numeric Type & String Type Operands

The string type operand will be converted to **DOUBLE**.

```
CREATE TABLE t(i INT, s STRING);
INSERT INTO t VALUES (1, '1'), (2, '2'), (3, '3'), (4, '4'), (12, '12');

SELECT i FROM t WHERE i < '11.3';

      i
=====
      1
      2
      3
      4

SELECT ('2' <= 11);

      ('2'<11)
=====
      1
```

String Type & Date/Time Type Operands

The string type operand will be converted to the date/time type.

```
SELECT ('2010-01-01' < date'2010-02-02');

      ('2010-01-01'<date '2010-02-02')
=====
      1
```

```

SELECT (date'2010-02-02' >= '2010-01-01');

      (date '2010-02-02'>='2010-01-01')
=====
1

```

String Type & Numeric Type Host Variable Operands

The numeric type host variable will be converted to the string type.

```

PREPARE s FROM 'SELECT s FROM t WHERE s < ?';
EXECUTE s USING 11;

      s
=====
      '1'

```

String Type & Numeric Type value Operands

The numeric type value will be converted to the string type.

```

SELECT s FROM t WHERE s > 11;

      s
=====
      '2'
      '3'
      '4'
      '12'

SELECT s FROM t WHERE s BETWEEN 11 AND 33;

      s
=====
      '2'
      '3'
      '12'

```

String Type Column & Date/Time Type Value Operands

The date/time type value will be converted to the string type.

```

SELECT s FROM t;

      s
=====
      '01/01/1998'
      '01/01/1999'
      '01/01/2000'

SELECT s FROM t WHERE s <= date'02/02/1998';

      s
=====
      '01/01/1998'
      '01/01/1999'
      '01/01/2000'

```

Range Operation

Numeric Type and String Type Operands

The string type operand will be converted to **DOUBLE**.

```

SELECT i FROM t WHERE i <= all {'11','12'};

      i
=====
      1
      2
      3
      4

```

String Type and Date/Time Type Operands

The string type operand will be converted to the date/time type.

```
SELECT s FROM t2;

      s
=====
'01/01/2000'
'01/01/1999'
'01/01/1998'

SELECT s FROM t2 WHERE s <= ALL {date'02/02/1998',date'01/01/2000'};

      s
=====
'01/01/1998'
```

If it is impossible to convert to the corresponding type, an error is returned.

Arithmetic Operation

Date/Time Type Operand

If the date/time type operands are given to '-' operator and the types are different from each other, it will be converted to the type with a higher priority. The following example shows that the operand data type on the left is converted from **DATE** to **DATETIME** so that the result of '-' operation of **DATETIME** can be outputted in milliseconds.

```
SELECT date'2002-01-01' - datetime'2001-02-02 12:00:00 am';

      date '2002-01-01'- datetime '2001-02-02 12:00:00 am'
=====
                                28771200000
```

Numeric Type Operand

If the numeric type operands are given and the types are different from each other, it will be converted to the type with the higher priority.

Date/Time Type & Numeric Type Operands

If the date/time type and the numeric type operands are given to '+' or '-' operator, the numeric type operand is converted to either **BIGINT**, **INT** or **SHORT**.

Date/Time Type & String Type Operands

If a date/time type and a string type are operands, only '+' and '-' operators are allowed. If the '+' operator is used, it will be applied according to the following rules.

- The string type will be converted to **BIGINT** with an interval value. The interval is the smallest unit for operands in the Date/Time type, and the interval for each type is as follows:
- **DATE** : Days
- **TIME**, **TIMESTAMP** : Seconds
- **DATETIME** : Milliseconds
- Floating-point numbers are rounded.
- The result type is the type of an date/time operand.

```
SELECT date'2002-01-01' + '10';

      date '2002-01-01'+ '10'
=====
01/11/2002
```

If the date/time type and a string type are operands and the '-' operator is used, they will be applied according to the following rules.

- If the date/time type operands are **DATE**, **DATETIME** and **TIMESTAMP**, the string will be converted to **DATETIME**; if the date/time operand is **TIME**, the string is converted to **TIME**.
- The result type is always **BIGINT**.

```
SELECT date'2002-01-01'-'2001-01-01';

      date '2002-01-01'-'2001-01-01'
=====
                        31536000000

-- this causes an error

SELECT date'2002-01-01'-'10';

In line 1, column 13,
ERROR: Cannot coerce '10' to type datetime.
```

Numeric Type & String Type Operands

If a numeric type and a string type are operands, they will be applied according to the following rules.

- Strings will be converted to **DOUBLE** when possible.
- The result type is **DOUBLE** or **MONETARY** and depends on the type of the numeric operand.

```
SELECT 4 + '5.2';

      4+'5.2'
=====
9.1999999999999999e+00
```

Unlike CUBRID 2008 R3.1 and the earlier versions, the string in the date/time format, that is, the string such as '2010-09-15' is not converted to the date/time type. You can use a literal (DATE'2010-09-15') with the date/time type for addition and subtraction operations.

```
SELECT '2002-01-01'+1;
      ERROR: Cannot coerce '2002-01-01' to type double.
SELECT DATE'2002-01-01'+1;
      date '2002-01-01'+1
=====
      01/02/2002
```

String Type Operand

If you multiply, divide or subtract both strings, the result returns a **DOUBLE** type value.

```
SELECT '3'*'2';

      '3'*'2'
=====
6.0000000000000000e+00
```

The '+' operator action depends on how to set the system parameter **plus_as_concat** in the **cubrid.conf** file. For details, see [Syntax/Type Related Parameter](#).

- If a value for **plus_as_concat** is yes, the concatenation of two strings will be returned.

```
SELECT '1'+ '1';

      '1'+ '1'
=====
      '11'
```

- If a value for **plus_as_concat** is no and two strings can be converted to numbers, the **DOUBLE** type value will be returned by adding the two numbers.

```
SELECT '1'+ '1';

      '1'+ '1'
=====
      2.0000000000000000e+00
```

If it is impossible to convert to the corresponding type, an error is returned.

Table Definition

CREATE TABLE

Table Definition

Description

To create a table, use the **CREATE TABLE** statement.

Syntax

```

CREATE {TABLE | CLASS} <table_name>
    [ <subclass definition> ]
    [ ( <column definition> [,<table constraint>]... ) ]
    [ AUTO_INCREMENT = initial_value ]
    [ CLASS ATTRIBUTE ( <column_definition_comma_list> ) ]
    [ METHOD <method_definition_comma_list> ]
    [ FILE <method_file_comma_list> ]
    [ INHERIT <resolution_comma_list> ]
    [ REUSE_OID ]

<column definition> ::=
column_name column_type [ [ <default_or_shared> ] | [ <column_constraint> ] ]...

<default or shared> ::=
{SHARED <value_specification> | DEFAULT <value_specification> } |
AUTO_INCREMENT [(seed, increment)]

<column_constraint> ::=
NOT NULL | UNIQUE | PRIMARY KEY | FOREIGN KEY <referential definition>

<table constraint> ::=
[ CONSTRAINT [ <constraint name> ] ] UNIQUE [ KEY | INDEX ]( column_name_comma_list ) |
[ { KEY | INDEX } [ <constraint_name> ]( column_name_comma_list ) |
[ PRIMARY KEY ( column_name_comma_list ) ] |
[ <referential constraint> ]

<referential constraint> ::=
FOREIGN KEY [ <foreign key name> ]( column_name_comma_list ) <referential definition>

<referential definition> ::=
REFERENCES [ referenced_table_name ] ( column_name_comma_list )
[ <referential triggered action> ... ]

<referential triggered action> ::=
{ ON UPDATE <referential_action> } |
{ ON DELETE <referential_action> } |
{ ON CACHE OBJECT cache_object_column_name }

<referential action> ::=
CASCADE | RESTRICT | NO ACTION | SET NULL

<subclass definition> ::=
{ UNDER | AS SUBCLASS OF } table_name_comma_list

<method definition> ::=
[ CLASS ] method_name
[ ( [ argument_type_comma_list ] ) ]
[ result_type ]
[ FUNCTION function_name ]

<resolution> ::=
[ CLASS ] { column_name | method_name } OF superclass_name
[ AS alias ]

```

- *table_name* : Specifies the name of the table to be created (maximum : 255 bytes).
- *column_name* : Specifies the name of the column to be created.

- *column_type* : Specifies the data type of the column.
- [**SHARED** *value* | **DEFAULT** *value*] : Specifies the initial value of the column (see [Column Definition](#) For details).
- *column_constraints* : Specifies the constraint of the column. Available constraints are **NOT NULL**, **UNIQUE**, **PRIMARY KEY** and **FOREIGN KEY** (see [Constraint Definition](#) For details).

Example

```
CREATE TABLE olympic (
  host_year      INT      NOT NULL PRIMARY KEY,
  host_nation    VARCHAR(40) NOT NULL,
  host_city      VARCHAR(20) NOT NULL,
  opening_date   DATE      NOT NULL,
  closing_date   DATE      NOT NULL,
  mascot         VARCHAR(20) ,
  slogan         VARCHAR(40) ,
  introduction    VARCHAR(1500)
);
```

Column Definition

A column is a set of data values of a particular simple type, one for each row of the table.

```
<column_definition> ::=
column_name column_type [ [ <default_or_shared> ] | [ <column_constraint> ] ]...

<default_or_shared> ::=
{ SHARED <value_specification> | DEFAULT <value_specification> } |
AUTO_INCREMENT [ (seed, increment) ]

<column_constraint> ::=
NOT NULL | UNIQUE | PRIMARY KEY | FOREIGN KEY <referential_definition>
```

Column Name

Description

How to create a column name, see [Identifier](#).

You can alter created column name by using [RENAME COLUMN](#) clause of the **ALTER TABLE**.

Example

The following example shows how to create the manager2 table that has the following two columns: full_name and age.

```
CREATE TABLE manager2 (full_name VARCHAR(40), age INT );
```

Caution

- The first character of a column name must be an alphabet. The maximum length is 255 characters.
- The column name must be unique in the table.

Setting the Column Initial Value (SHARED, DEFAULT)

Description

SHARED and **DEFAULT** are attributes related to the initial value of the column. You can change the value of **SHARED** and **DEFAULT** in the **ALTER TABLE** statement.

- **SHARED** : Column values are identical in all rows. If a value different from the initial value is **INSERT**ed, the column value is updated to a new one in every row.
- **DEFAULT** : The initial value set when the **DEFAULT** attribute was defined is stored even if the column value is not specified when a new row is inserted.

Note If you set `SYS_TIMESTAMP` as a **DEFAULT** value when creating a table, the **TIMESTAMP** value at the point of **CREATE TABLE**, not the point at which the data is **INSERT**ed, is specified by default. Therefore, you must specify the `SYS_TIMESTAMP` value for the **VALUES** of the **INSERT** statement when entering data.

Example

```
CREATE TABLE colval tbl
(id INT, name VARCHAR SHARED 'AAA', phone VARCHAR DEFAULT '000-0000');
INSERT INTO colval_tbl(id) VALUES (1),(2);
SELECT * FROM colval_tbl;
```

id	name	phone
1	'AAA'	'000-0000'
2	'AAA'	'000-0000'

```
--updating column values on every row
INSERT INTO colval tbl(id, name) VALUES (3,'BBB');
INSERT INTO colval_tbl(id) VALUES (4),(5);
SELECT * FROM colval_tbl;
```

id	name	phone
1	'BBB'	'000-0000'
2	'BBB'	'000-0000'
3	'BBB'	'000-0000'
4	'BBB'	'000-0000'
5	'BBB'	'000-0000'

```
--changing DEFAULT value in the ALTER TABLE statement
ALTER TABLE colval tbl CHANGE phone DEFAULT '111-1111'
INSERT INTO colval_tbl(id) VALUES (6);
SELECT * FROM colval_tbl;
```

id	name	phone
1	'BBB'	'000-0000'
2	'BBB'	'000-0000'
3	'BBB'	'000-0000'
4	'BBB'	'000-0000'
5	'BBB'	'000-0000'
6	'BBB'	'111-1111'

AUTO INCREMENT

Description

You can define the **AUTO_INCREMENT** attribute for the column to automatically give serial numbers to column values. This can be defined only for **SMALLINT**, **INTEGER**, **BIGINT(p,0)**, and **NUMERIC(p,0)** domains.

DEFAULT, **SHARED** and **AUTO_INCREMENT** cannot be defined for the same column. Make sure the value entered directly by the user and the value entered by the auto increment attribute do not conflict with each other.

You can change the initial value of **AUTO_INCREMENT** by using the **ALTER TABLE** statement. For details, see [AUTO_INCREMENT Statement](#) of **ALTER TABLE**.

Syntax

```
CREATE TABLE table_name (id int AUTO_INCREMENT[(seed, increment)]) |
CREATE TABLE table_name (id int AUTO_INCREMENT) AUTO_INCREMENT = seed;
```

- *seed* : The initial value from which the number starts. All integers (positive, negative, and zero) are allowed. The default is 1.
- *increment* : The increment value of each row. Only positive integers are allowed. The default value 1.

When you use the **CREATE TABLE table_name (id int AUTO_INCREMENT) AUTO_INCREMENT = seed;** statement, the constraints are as follows:

- You should define only one column with the **AUTO_INCREMENT** attribute.

- Don't use (*seed*, increment) and `AUTO_INCREMENT = seed` together.

Example

```
CREATE TABLE auto_tbl(id INT AUTO_INCREMENT, name VARCHAR);
INSERT INTO auto_tbl VALUES(NULL, 'AAA'),(NULL, 'BBB'),(NULL, 'CCC');
INSERT INTO auto_tbl(name) VALUES ('DDD'),('EEE');
SELECT * FROM auto_tbl;

=====
      id  name
=====
        1  'AAA'
        2  'BBB'
        3  'CCC'
        4  'DDD'
        5  'EEE'
```

```
CREATE TABLE tbl (id int AUTO INCREMENT, val string) AUTO INCREMENT = 3;
INSERT INTO tbl VALUES (NULL,'cubrid');

SELECT * FROM tbl;
      id  val
=====
        3  'cubrid'
```

```
CREATE TABLE t (id int AUTO INCREMENT, id2 int AUTO INCREMENT) AUTO INCREMENT = 5;
ERROR: To avoid ambiguity, the AUTO_INCREMENT table option requires the table to have
exactly one AUTO_INCREMENT column and no seed/increment specification.
```

```
CREATE TABLE t (i int AUTO INCREMENT(100, 2)) AUTO INCREMENT = 3;
ERROR: To avoid ambiguity, the AUTO INCREMENT table option requires the table to have
exactly one AUTO_INCREMENT column and no seed/increment specification.
```

Caution

- Even if a column has auto increment, the **UNIQUE** constraint is not satisfied.
- If **NULL** is specified in the column where auto increment is defined, the value of auto increment is stored.
- The initial value and the final value obtained by auto increment cannot exceed the minimum and maximum values allowed in the given domain.
- Because auto increment has no cycle, an error occurs when the maximum value of the type exceeds, and no rollback is executed. Therefore, you must delete and recreate the column in such cases.
- For example, if a table is created as below, the maximum value of A is 32767. Because an error occurs if the value exceeds 32767, you must make sure that the maximum value of the column A does not exceed the maximum value of the type when creating the initial table.

```
create table tbl(A smallint auto_increment, B char(5));
```

Constraint Definition

Description

You can define **NOT NULL**, **UNIQUE**, **PRIMARY KEY**, **FOREIGN KEY** as the constraints. You can also create an index by using **INDEX** or **KEY**.

```
<column_constraint> ::=
NOT NULL | UNIQUE | PRIMARY KEY | FOREIGN KEY <referential_definition>

<table_constraint> ::=
[ CONSTRAINT [ <constraint_name> ] ] UNIQUE [ KEY | INDEX ]( column_name comma list ) |
[ { KEY | INDEX } [ <constraint_name> ]( column_name_comma_list ) |
[ PRIMARY KEY ( column_name_comma_list ) ] |
[ <referential_constraint> ]

<referential_constraint> ::=
FOREIGN KEY ( column_name comma list ) <referential_definition>

<referential_definition> ::=
REFERENCES [ referenced_table_name ] ( column_name_comma_list )
[ <referential_triggered_action> ... ]
```

```

<referential triggered action> ::=
{ ON UPDATE <referential_action> } |
{ ON DELETE <referential_action> } |
{ ON CACHE OBJECT cache_object_column_name }

<referential action> ::=
CASCADE | RESTRICT | NO ACTION | SET NULL

```

NOT NULL Constraints

Description

A column for which the **NOT NULL** constraint has been defined must have a certain value that is not **NULL**. The **NOT NULL** constraint can be defined for all columns. An error occurs if you try to insert a **NULL** value into a column with the **NOT NULL** constraint by using the **INSERT** or **UPDATE** statement.

Example

```

CREATE TABLE const_tbl1(id INT NOT NULL, INDEX i_index(id ASC), phone VARCHAR);

CREATE TABLE const_tbl2(id INT NOT NULL PRIMARY KEY, phone VARCHAR);
INSERT INTO const_tbl2 (NULL,'000-0000');

In line 2, column 25,
ERROR: syntax error, unexpected Null

```

UNIQUE Constraint

Description

The **UNIQUE** constraint enforces a column to have a unique value. An error occurs if a new record that has the same value as the existing one is added by this constraint.

You can place a **UNIQUE** constraint on either a column or a set of columns. If the **UNIQUE** constraint is defined for multiple columns, the uniqueness is ensured not for each column, but the combination of multiple columns.

Example

If a **UNIQUE** constraint is defined on a set of columns, this ensures the uniqueness of the values in all the columns. As shown below, the second **INSERT** statement succeeds because the value of column a is the same, but the value of column b is unique. The third **INSERT** statement causes an error because the values of column a and b are the same as those in the first **INSERT** statement.

```

--UNIQUE constraint is defined on a single column only
CREATE TABLE const_tbl5(id INT UNIQUE, phone VARCHAR);
INSERT INTO const_tbl5(id) VALUES (NULL), (NULL);
INSERT INTO const_tbl5 VALUES (1, '000-0000');
SELECT * FROM const_tbl5;

      id  phone
=====
      NULL NULL
      NULL NULL
         1 '000-0000'

INSERT INTO const_tbl5 VALUES (1, '111-1111');

ERROR: Operation would have caused one or more unique constraint violations.

--UNIQUE constraint is defined on several columns
CREATE TABLE const_tbl6(id INT, phone VARCHAR, CONSTRAINT UNIQUE(id,phone));
INSERT INTO const_tbl6 VALUES (1,NULL), (2,NULL), (1,'000-0000'), (1,'111-1111');
SELECT * FROM const_tbl6;

```

id	phone
1	NULL
2	NULL
1	'000-0000'
1	'111-1111'

PRIMARY KEY Constraint

Description

A key in a table is a set of column(s) that uniquely identifies each row. A candidate key is a set of columns that uniquely identifies each row of the table. You can define one of such candidate keys a primary key. That is, the column defined as a primary key is uniquely identified in each row.

By default, the index created by defining the primary key is created in ascending order, and you can define the order by specifying **ASC** or **DESC** keyword next to the column.

Syntax

```
CREATE TABLE pk_tbl (a INT, b INT, PRIMARY KEY (a, b DESC));
```

Example

```
CREATE TABLE const_tbl7(
id INT NOT NULL,
phone VARCHAR,
CONSTRAINT pk id PRIMARY KEY(id));

--CONSTRAINT keyword
CREATE TABLE const_tbl8(
id INT NOT NULL PRIMARY KEY,
phone VARCHAR);

--primary key is defined on multiple columns
CREATE TABLE const_tbl8 (
host year    INT NOT NULL,
event code  INT NOT NULL,
athlete code INT NOT NULL,
medal       CHAR(1) NOT NULL,
score       VARCHAR(20),
unit        VARCHAR(5),
PRIMARY KEY(host year, event code, athlete code, medal)
);
```

FOREIGN KEY Constraint

Description

A foreign key is a column or a set of columns that references the primary key in other tables in order to maintain reference relationship. The foreign key and the referenced primary key must have the same data type. Consistency between two tables is maintained by the foreign key referencing the primary key, which is called referential integrity.

Syntax

```
[ CONSTRAINT <constraint_name> ]
FOREIGN KEY [ <foreign_key_name> ] ( column_name_comma_list )
REFERENCES [ referenced_table_name ] ( column_name_comma_list )
[ <referential_triggered action> ]

<referential_triggered action> :
ON UPDATE <referential action>
[ ON DELETE <referential action> [ ON CACHE OBJECT cache_object_column_name ] ]

<referential action> :
CASCADE | RESTRICT | NO ACTION | SET NULL
```

- *constraint_name* : Specifies the name of the table to be created.

- *foreign_key_name* : Specifies a name of the **FOREIGN KEY** constraint. You can skip the name specification. However, if you specify this value, *constraint_name* will be ignored, and the specified value will be used.
- *column_name* : Specifies the name of the column to be defined as a foreign key after the **FOREIGN KEY** keyword. The column number of foreign keys defined and primary keys must be same.
- *referenced_table_name* : Specifies the name of the table to be referenced.
- *column_name* : Specifies the name of the referred primary key column after the **FOREIGN KEY** keyword.
- *referential_triggered_action* : Specifies the trigger action that responds to a certain operation in order to maintain referential integrity. **ON UPDATE**, **ON DELETE** or **ON CACHE OBJECT** can be specified. Each action can be defined multiple times, and the definition order is not significant.
- **ON UPDATE** : Defines the action to be performed when attempting to update the primary key referenced by the foreign key. You can use either **NO ACTION**, **RESTRICT**, or **SET NULL** option. The default is **RESTRICT**.
- **ON DELETE** : Defines the action to be performed when attempting to delete the primary key referenced by the foreign key. You can use **NO ACTION**, **RESTRICT**, **CASCADE**, or **SET NULL** option. The default is **RESTRICT**.
- **ON CACHE OBJECT** : You can search an object using a direct object reference in object-oriented model. **ON CACHE OBJECT** option supports this feature in association with referential integrity (foreign key). **ON CACHE OBJECT** option adds an OID reference to a foreign key configuration. The OID is used as a **CACHE** point for the foreign key to the primary key table. Such OID is managed by the system internally; it cannot be changed by users. To define the **ON CACHE OBJECT** option, you must have defined a column whose domain is the table with a primary key and specified the column in the *cache_object_column_name*. The attribute defined with **ON CACHE OBJECT** can use the OID the same way as the one of the existing object type.
- *referential_action* : You can define an option that determines whether to maintain the value of the foreign key when the primary key value is deleted or updated.
- **CASCADE** : If the primary key is deleted, the foreign key is deleted as well. This option is supported only for the **ON DELETE** operation.
- **RESTRICT** : Prevents the value of the primary key from being deleted or updated, and rolls back any transaction that has been attempted.
- **SET NULL** : When a specific record is being deleted or updated, the column value of the foreign key is updated to **NULL**.
- **NO ACTION** : Its behavior is the same as that of the **RESTRICT** option.

Example

```
--creating two tables where one is referencing the other
CREATE TABLE a_tbl(
id INT NOT NULL DEFAULT 0 PRIMARY KEY,
phone VARCHAR(10));

CREATE TABLE b_tbl(
ID INT NOT NULL,
name VARCHAR(10) NOT NULL,
CONSTRAINT pk id PRIMARY KEY(id),
CONSTRAINT fk id FOREIGN KEY(id) REFERENCES a_tbl(id)
ON DELETE CASCADE ON UPDATE RESTRICT);

INSERT INTO a_tbl VALUES(1,'111-1111'), (2,'222-2222'), (3, '333-3333');
INSERT INTO b_tbl VALUES(1,'George'), (2,'Laura'), (3,'Max');
SELECT a.id, b.id, a.phone, b.name FROM a_tbl a, b_tbl b WHERE a.id=b.id;
```

id	id	phone	name
1	1	'111-1111'	'George'
2	2	'222-2222'	'Laura'
3	3	'333-3333'	'Max'

```
--when deleting primay key value, it cascades foreign key value
DELETE FROM a_tbl WHERE id=3;

1 rows affected.

SELECT a.id, b.id, a.phone, b.name FROM a_tbl a, b_tbl b WHERE a.id=b.id;
```

id	id	phone	name
1	1	'111-1111'	'George'
2	2	'222-2222'	'Laura'

```
=====
      1          1          '111-1111'      'George'
      2          2          '222-2222'      'Laura'

--when attempting to update primay key value, it restricts the operation
UPDATE  a tbl SET id = 10 WHERE phone = '111-1111';

In the command from line 1,

ERROR: Update/Delete operations are restricted by the foreign key 'fk_id'.

0 command(s) successfully processed.
```

Caution

- In a referential constraint, the name of the primary key table to be referenced and the corresponding column names are defined. If the list of column names are is not specified, the primary key of the primary key table is specified in the defined order.
- The number of primary keys in a referential constraint must be identical to that of foreign keys. The same column name cannot be used multiple times for the primary key in the referential constraint.
- The actions cascaded by reference constraints do not activate the trigger action.
- It is not recommended to use **referential_triggered_action** in the CUBRID HA environment. In the CUBRID HA environment, the trigger action is not supported. Therefore, if you use **referential_triggered_action**, the data between the master database and the slave database can be inconsistent. For details, see [CUBRID HA](#).

KEY or INDEX

Description

KEY and **INDEX** are used interchangeably. They create an index that uses the corresponding column as a key. You can specify the index name. If omitted, a name is assigned automatically.

Example

```
CREATE TABLE const_tbl3(id INT, phone VARCHAR, INDEX(id DESC, phone ASC));

CREATE TABLE const_tbl4(id INT, phone VARCHAR, KEY i_key(id DESC, phone ASC));
```

Column Option

Description

You can specify options such as **ASC** or **DESC** after the column name when defining **UNIQUE** or **INDEX** for a specific column. This keyword is specified to store the index value in ascending or descending order.

Syntax

```
column_name [ASC|DESC]
```

Example

```
CREATE TABLE const tbl(
id VARCHAR,
name VARCHAR,
CONSTRAINT UNIQUE INDEX(id DESC, name ASC)
);

INSERT INTO const tbl VALUES('1000', 'john'), ('1000','johnny'), ('1000', 'jone');
INSERT INTO const tbl VALUES('1001', 'johnny'), ('1001','john'), ('1001', 'jone');

SELECT * FROM const_tbl WHERE id > '100';
=====
      id      name
      1001    john
      1001    johnny
      1001     jone
```

1000	john
1000	johnny
1000	jone

Table Option (REUSE_OID)

Description

You can specify the **REUSE_OID** option when creating a table, so that OIDs that have been deleted due to the deletion of records (**DELETE**) can be reused when a new record is inserted (**INSERT**). Such a table is called an OID reusable or a non-referable table.

OID (Object Identifier) is an object identifier represented by physical location information such as the volume number, page number and slot number. By using such OIDs, CUBRID manages the reference relationships of objects and searches, stores or deletes them. When an OID is used, accessibility is improved because the object in the heap file can be directly accessed without referring to the table. However, the problem of decreased reusability of the storage occurs when there are many **DELETE/INSERT** operations because the object's OID is kept to maintain the reference relationship with the object even if it is deleted.

If you specify the **REUSE_OID** option when creating a table, the OID is also deleted when data in the table is deleted, so that another **INSERT**ed data can use it. OID reusable tables cannot be referred to by other tables, and OID values of the objects in the OID reusable tables cannot be viewed.

Example

```
--creating table with REUSE_OID option specified
CREATE TABLE reuse_tbl (a INT PRIMARY KEY) REUSE_OID;
INSERT INTO reuse_tbl VALUES (1);
INSERT INTO reuse_tbl VALUES (2);
INSERT INTO reuse_tbl VALUES (3);

--an error occurs when column type is a OID reusable table itself
CREATE TABLE tbl_1 ( a reuse_tbl);

ERROR: The class 'reuse_tbl' is marked as REUSE OID and is non-referable. Non-referable
classes can't be the domain of an attribute and their instances' OIDs cannot be returned.

--an error occurs when a table references a OID reusable table
CREATE TABLE tbl_2
(b int, FOREIGN KEY(b) REFERENCES reuse_tbl(a) ON CACHE OBJECT oid value);
INSERT INTO tbl_2(b) VALUES(1);
SELECT oid value.a FROM tbl_2;

ERROR: The class 'reuse_tbl' is marked as REUSE_OID and is non-referable. Non-referable
classes can't be the domain of an attribute and their instances' OIDs cannot be returned.
```

Caution

- OID reusable tables cannot be referred to by other tables.
- Updatable views cannot be created for OID reusable tables.
- OID reusable tables cannot be specified as class attribute domains of other tables.
- OID values of the objects in the OID reusable tables cannot be read.
- Instance methods cannot be called from OID reusable tables. Also, instance methods cannot be called if a subclass inherited from the class where the method is defined is defined as an OID reusable table.
- OID reusable tables are supported only by CUBRID 2008 R2.2 or above, and backward compatibility is not ensured. That is, the database in which the OID reusable table is located cannot be accessed from a lower version database.
- OID reusable tables can be managed as partitioned tables and can be replicated.

CREATE TABLE LIKE

Description

You can create a table that has the same schema as an existing table by using the **CREATE TABLE...LIKE** statement.

Column attribute, table constraint, and index are replicated from the existing table. An index name created from the existing table changes according to a new table name, but an index name defined by a user is replicated as it is. Therefore, you should be careful at a query statement that is supposed to use a specific index created by using the **USING INDEX**.

You cannot create the column definition because the **CREATE TABLE...LIKE** statement replicates the schema only.

Syntax

```
CREATE {TABLE | CLASS} <new_table_name> LIKE <old_table_name>
```

- *new_table_name* : A table name to be created.
- *old_table_name* : The name of the original table that already exists in the database. The following tables cannot be specified as original tables in the **CREATE TABLE...LIKE** statement.
 - Partition table
 - Table that contains an **AUTO_INCREMENT** column
 - Table that uses inheritance or methods

Example

```
CREATE TABLE a tbl(
id INT NOT NULL DEFAULT 0 PRIMARY KEY,
phone VARCHAR(10));
INSERT INTO a_tbl VALUES(1,'111-1111'), (2,'222-2222'), (3, '333-3333');
```

```
--creating an empty table with the same schema as a tbl
CREATE TABLE new_tbl LIKE a_tbl;
SELECT * FROM new_tbl;
```

There are no results.

```
;schema a tbl
```

```
=== <Help: Schema of a Class> ===
```

```
<Class Name>
```

```
  a_tbl
```

```
<Attributes>
```

```
  id                INTEGER DEFAULT 0 NOT NULL
  phone             CHARACTER VARYING(10)
```

```
<Constraints>
```

```
  PRIMARY KEY pk a_tbl id ON a_tbl (id)
```

Current transaction has been committed.

```
;schema new_tbl
```

```
=== <Help: Schema of a Class> ===
```

```
<Class Name>
```

```
  new_tbl
```

```
<Attributes>
```

```
  id                INTEGER DEFAULT 0 NOT NULL
```

```

phone                CHARACTER VARYING(10)

<Constraints>

PRIMARY KEY pk new tbl id ON new tbl (id)

Current transaction has been committed.

```

CREATE TABLE AS SELECT

Description

You can create a new table that contains the result records of the **SELECT** statement by using the **CREATE TABLE...AS SELECT** statement. You can define column and table constraints for the new table. The following rules are applied to reflect the result records of the **SELECT** statement.

- If *col_1* is defined in the new table and the same column *col_1* is specified in *select_statement*, the result record of the **SELECT** statement is stored as *col_1* value in the new table. Type casting is attempted if the column names are identical but the columns types are different.
- If *col_1* and *col_2* are defined in the new table, *col_1*, *col_2* and *col_3* are specified in the column list of the *select_statement* and there is a containment relationship between all of them, *col_1*, *col_2* and *col_3* are created in the new table and the result data of the **SELECT** statement is stored as values for all columns. Type casting is attempted if the column names are identical but the columns types are different.
- If columns *col_1* and *col_2* are defined in the new table and *col_1* and *col_3* are defined in the column list of *select_statement* without any containment relationship between them, *col_1*, *col_2* and *col_3* are created in the new table, the result data of the **SELECT** statement is stored only for *col_1* and *col_3* which are specified in *select_statement*, and **NULL** is stored as the value of *col_2*.
- Column aliases can be included in the column list of *select_statement*. In this case, new column alias is used as a new table column name. It is recommended to use an alias because invalid column name is created, if an alias does not exist when a function calling or an expression is used.
- The **REPLACE** option is valid only when the **UNIQUE** constraint is defined in a new table column (*col_1*). When duplicate values exist in the result record of *select_statement*, a **UNIQUE** value is stored for *col_1* if the **REPLACE** option has been defined, or an error message is displayed if the **REPLACE** option is omitted due to the violation of the **UNIQUE** constraint.

Syntax

```

CREATE {TABLE | CLASS} <table_name>
    [( <column_definition> [, <table_constraint>]... )]
    [REPLACE] AS <select_statement>

```

- *table_name* : A name of the table to be created.
- *column_definition* : Defines a column. If it is omitted, the column schema of **SELECT** statement is replicated; however, the constraint or the **AUTO_INCREMENT** attribute is not replicated.
- *table_constraint* : Defines table constraint.
- *select_statement* : A **SELECT** statement targeting a source table that already exists in the database.

Example

```

CREATE TABLE a_tbl(
id INT NOT NULL DEFAULT 0 PRIMARY KEY,
phone VARCHAR(10));
INSERT INTO a_tbl VALUES(1, '111-1111'), (2, '222-2222'), (3, '333-3333');

--creating a table without column definition
CREATE TABLE new_tbl1 AS SELECT * FROM a_tbl;
SELECT * FROM new_tbl1;

      id  phone
-----
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'

--all of column values are replicated from a_tbl

```

```

CREATE TABLE new_tbl2
(id INT NOT NULL AUTO INCREMENT PRIMARY KEY, phone VARCHAR) AS SELECT * FROM a_tbl;
SELECT * FROM new_tbl2;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'

--some of column values are replicated from a_tbl and the rest is NULL
CREATE TABLE new_tbl3
(id INT, name VARCHAR) AS SELECT id, phone FROM a_tbl;
SELECT * FROM new_tbl3

      name                id  phone
=====
      NULL                1  '111-1111'
      NULL                2  '222-2222'
      NULL                3  '333-3333'

--column alias in the select statement should be used in the column definition
CREATE TABLE new_tbl4
(id1 int, id2 int)AS SELECT t1.id id1, t2.id id2 FROM new_tbl1 t1, new_tbl2 t2;
SELECT * FROM new_tbl4;

      id1      id2
=====
      1         1
      1         2
      1         3
      2         1
      2         2
      2         3
      3         1
      3         2
      3         3

--REPLACE is used on the UNIQUE column
CREATE TABLE new_tbl5(id1 int UNIQUE) REPLACE AS SELECT * FROM new_tbl4;
SELECT * FROM new_tbl5;

      id1      id2
=====
      1         3
      2         3
      3         3

```

ALTER TABLE

Overview

Description

You can modify the structure of a table by using the **ALTER** statement. You can perform operations on the target table such as adding/deleting columns, creating/deleting indexes, and type casting existing columns as well as changing table names, column names and constraints. **TABLE** and **CLASS** are used interchangeably **VIEW** and **VCLASS**, and **COLUMN** and **ATTRIBUTE** as well.

You can also change the initial value of **AUTO_INCREMENT**.

Syntax

```

ALTER [ <class type> ] <table name> <alter clause> ;

<class_type> ::= TABLE | CLASS | VCLASS | VIEW

<alter_clause> ::= ADD <alter_add> [ INHERIT <resolution_comma_list> ] |
                ADD { KEY | INDEX } [index_name] (<index_col_name>) |

```

```

ALTER [ COLUMN ] column_name SET DEFAULT <value_specification> |
DROP <alter_drop> [ INHERIT <resolution_comma_list> ] |
DROP { KEY | INDEX } index_name |
DROP FOREIGN KEY constraint_name |
DROP PRIMARY KEY |
RENAME <alter_rename> [ INHERIT <resolution_comma_list> ] |
CHANGE <alter_change> |
INHERIT <resolution_comma_list>
AUTO_INCREMENT = <initial_value>

<alter_add> ::= [ ATTRIBUTE | COLUMN ] [( <class_element_comma_list> [ ] ) [ FIRST | AFTER
old column name ] ] |
CLASS ATTRIBUTE <column_definition_comma_list> |
CONSTRAINT <constraint_name> <column_constraint> ( column_name ) |
FILE <file_name_comma_list> |
METHOD <method_definition_comma_list> |
QUERY <select_statement> |
SUPERCLASS <class_name_comma_list>

<alter_change> ::= FILE <file_path_name> AS <file_path_name> |
METHOD <method_definition_comma_list> |
QUERY [ <unsigned_integer_literal> ] <select_statement> |
<column_name> DEFAULT <value_specification>

<alter_drop> ::= [ ATTRIBUTE | COLUMN | METHOD ]
<column_name_comma_list> |
FILE <file_name_comma_list> |
QUERY [ <unsigned_integer_literal> ] |
SUPERCLASS <class_name_comma_list> |
CONSTRAINT <constraint_name>

<alter_rename> ::= [ ATTRIBUTE | COLUMN | METHOD ]
<old_column_name> AS <new_column_name> |
FUNCTION OF <column_name> AS <function_name>
FILE <file_path_name> AS <file_path_name>

<resolution> ::= { column_name | method_name } OF <superclass_name>
[ AS alias ]

<class_element> ::= <column_definition> | <table_constraint>

<column_constraint> ::= UNIQUE [ KEY ] | PRIMARY KEY | FOREIGN KEY

<index_col_name> ::=
column_name [(length)] [ ASC | DESC ]

```

Caution

The table name can be changed only by the table owner, DBA and DBA members. The other users must be granted to change the name by the owner or DBA (see [Granting Authorization](#) For details on authorization).

ADD COLUMN Clause

Description

You can add a new column by using the **ADD COLUMN** clause. You can specify the location of the column to be added by using the **FIRST** or **AFTER** keyword.

If the newly added column has the **NOT NULL** constraint but no **DEFAULT** constraint, it will have the hard default when the database server configuration parameter, **add_column_update_hard_default** is set to yes. However, when the parameter is set to no, the column will have **NULL** even with the **NOT NULL** constraint.

If the newly added column has the **PRIMARY KEY** or **UNIQUE** constraints, an error will be returned when the database server configuration parameter **add_column_update_hard_default** is set to yes. When the parameter is set to no, all data will have **NULL**. The default value of **add_column_update_hard_default** is **no**.

For **add_column_update_hard_default** and the hard default, see [CHANGE Clause](#).

Syntax

```
ALTER [ TABLE | CLASS | VCLASS | VIEW ] table_name
ADD [ COLUMN | ATTRIBUTE ] [( <column_definition> ) ] [ FIRST | AFTER old_column_name ]

column_definition ::=
column_name column_type
    { [ NOT NULL | NULL ] |
      [ { SHARED <value_specification> | DEFAULT <value_specification> }
        | AUTO_INCREMENT [(seed, increment)] ] |
      [ UNIQUE [ KEY ] |
        [ PRIMARY KEY | FOREIGN KEY REFERENCES
          [ referenced_table_name ]( column_name_comma_list )
          [ <referential_triggered action> ... ]
        ]
      ] } ...

<referential_triggered action> ::=
{ ON UPDATE <referential_action> } |
{ ON DELETE <referential_action> } |
{ ON CACHE OBJECT cache_object_column_name }

<referential_action> ::=
CASCADE | RESTRICT | NO ACTION | SET NULL
```

- *table_name* : Specifies the name of a table that has a column to be added.
- *column_definition* : Specifies the name, data type, and constraints of a column to be added.
- **AFTER** *oid_column_name* : Specifies the name of an existing column before the column to be added.

Example

```
CREATE TABLE a tbl;
ALTER TABLE a tbl ADD COLUMN age INT DEFAULT 0 NOT NULL;
INSERT INTO a tbl (age) VALUES (20), (30), (40);
ALTER TABLE a tbl ADD COLUMN name VARCHAR FIRST;
ALTER TABLE a_tbl ADD COLUMN id INT NOT NULL AUTO_INCREMENT UNIQUE;
ALTER TABLE a_tbl ADD COLUMN phone VARCHAR(13) DEFAULT '000-0000-0000' AFTER name;

SELECT * FROM a tbl;
```

name	phone	age	id
NULL	'000-0000-0000'	20	NULL
NULL	'000-0000-0000'	30	NULL
NULL	'000-0000-0000'	40	NULL

```
--adding multiple columns
ALTER TABLE a_tbl ADD COLUMN (age1 int, age2 int, age3 int);
```

ADD CONSTRAINT Clause

Description

You can add a new constraint by using the **ADD CONSTRAINT** clause.

By default, the index created when you add **PRIMARY KEY** constraints is created in ascending order, and you can define the key sorting order by specifying the **ASC** or **DESC** keyword next to the column name.

Syntax

```
ALTER [ TABLE | CLASS | VCLASS | VIEW ] table_name
ADD CONSTRAINT < constraint_name > column_constraint ( column_name_comma_list )

column_constraint ::=
UNIQUE [ KEY ] |
PRIMARY KEY |
FOREIGN KEY [ <foreign key name> ] REFERENCES
[referenced_table_name]( column_name_comma_list )
[ <referential_triggered action> ... ]
```

```
<referential_triggered_action> ::=
{ ON UPDATE <referential_action> } |
{ ON DELETE <referential_action> } |
{ ON CACHE OBJECT cache_object_column_name }
```

```
<referential_action> ::=
CASCADE | RESTRICT | NO ACTION | SET NULL
```

- *table_name* : Specifies the name of a table that has a constraint to be added.
- *constraint_name* : Specifies the name of a constraint to be added, or it can be omitted. If omitted, a name is automatically assigned.
- *foreign_key_name* : Specifies a name of the **FOREIGN KEY** constraint. You can skip the name specification. However, if you specify this value, *constraint_name* will be ignored, and the specified value will be used.
- *column_constraint* : Defines a constraint for the specified column. For details, see [Constraint Definition](#).

Example

```
ALTER TABLE a tbl ADD CONSTRAINT PRIMARY KEY(id);
ALTER TABLE a tbl ADD CONSTRAINT PRIMARY KEY(id, no DESC);
ALTER TABLE a_tbl ADD CONSTRAINT UNIQUE u_key1(id);
```

ADD INDEX Clause

Description

You can define the index attributes for a specific column by using the **ADD INDEX** clause.

Syntax

```
ALTER [ TABLE | CLASS ] table_name ADD { KEY | INDEX } [index_name] (<index_col_name>)
```

```
<index_col_name> ::=
column_name [(length)] [ ASC | DESC ]
```

- *table_name* : Specifies the name of a table to be modified.
- *index_name* : Specifies the name of an index. If omitted, a name is automatically assigned.
- *index_col_name* : Specifies the column that has an index to be defined. **ASC** or **DESC** can be specified for a column option; *prefix_length* of an index key also can be specified for a column option.

Example

```
ALTER TABLE a tbl ADD INDEX (age ASC), ADD INDEX(phone DESC);
;schema a tbl
```

```
=== <Help: Schema of a Class> ===
```

```
<Class Name>
```

```
    a tbl
```

```
<Attributes>
```

```
    name                CHARACTER VARYING(1073741823) DEFAULT ''
    phone                CHARACTER VARYING(13) DEFAULT '111-1111'
    age                  INTEGER
    id                   INTEGER AUTO INCREMENT NOT NULL
```

```
<Constraints>
```

```
    UNIQUE u a tbl id ON a tbl (id)
    INDEX i a tbl age ON a tbl (age)
    INDEX i a tbl phone d ON a tbl (phone DESC)
```

```
Current transaction has been committed.
```

ALTER COLUMN ... SET DEFAULT Clause

Description

You can specify a new default value for a column that has no default value or modify the existing default value by using the **ALTER COLUMN ... SET DEFAULT**. You can use the **CHANGE** clause to change the default value of multiple columns with a single statement. For details, see the [CHANGE Clause](#).

Syntax

```
ALTER [ TABLE | CLASS ] table_name ALTER [COLUMN] column_name SET DEFAULT value
```

- *table_name* : Specifies the name of a table that has a column whose default value is to be modified.
- *column_name* : Specifies the name of a column whose default value is to be modified.
- *value* : Specifies a new default value.

Example

```
;schema a tbl
=== <Help: Schema of a Class> ===

<Class Name>

  a_tbl

<Attributes>

  name          CHARACTER VARYING(1073741823)
  phone         CHARACTER VARYING(13) DEFAULT '000-0000-0000'
  age           INTEGER
  id            INTEGER AUTO_INCREMENT NOT NULL

<Constraints>

  UNIQUE u_a_tbl_id ON a_tbl (id)

Current transaction has been committed.

ALTER TABLE a_tbl ALTER COLUMN name SET DEFAULT '';
ALTER TABLE a_tbl ALTER COLUMN phone SET DEFAULT '111-1111';

;schema a tbl
=== <Help: Schema of a Class> ===

<Class Name>

  a_tbl

<Attributes>

  name          CHARACTER VARYING(1073741823) DEFAULT ''
  phone         CHARACTER VARYING(13) DEFAULT '111-1111'
  age           INTEGER
  id            INTEGER AUTO_INCREMENT NOT NULL

<Constraints>

  UNIQUE u_a_tbl_id ON a_tbl (id)
```

AUTO_INCREMENT Clause

Description

The **AUTO_INCREMENT** clause can change the initial value of the increment value that is currently defined. However, there should be only one **AUTO_INCREMENT** column defined.

Syntax

```
ALTER TABLE table_name AUTO_INCREMENT = initial_value;
```

- *table_name* : Table name
- *initial_value* : Initial value to alter

Example

```
CREATE TABLE t (i int AUTO_INCREMENT);
ALTER TABLE t AUTO_INCREMENT = 5;

-- when 2 AUTO INCREMENT constraints are defined on one table, it returns error.
CREATE TABLE t (i int AUTO INCREMENT, j int AUTO INCREMENT);
ALTER TABLE t AUTO_INCREMENT = 5;

ERROR: To avoid ambiguity, the AUTO INCREMENT table option requires the table to have
exactly one AUTO_INCREMENT column and no seed/increment specification.
```

Caution

You must be careful not to violate constraints (such as a **PRIMARY KEY** or **UNIQUE**) when you alter the initial value of **AUTO_INCREMENT**.

CHANGE/MODIFY Clauses

Description

The **CHANGE** clause changes column names or changes the types and the attributes. If the existing column name and a new column name are the same, only the type and the attribute will be changed.

The **MODIFY** clause can modify the types and the attributes of columns but cannot change the names.

If you set the type and the attribute to apply to a new column with the **CHANGE** clause or the **MODIFY** clause, the attribute that is currently defined will not be passed to the attribute of the new column.

When you change data types using the **CHANGE** clause or the **MODIFY** clause, the data can be modified. For example, if you shorten the length of a column, the character string may be truncated.

Note ALTER TABLE <table_name> CHANGE <column_name> DEFAULT <default_value> syntax supported in CUBRID 2008 R3.1 or earlier version is no longer supported.

Syntax

```
ALTER TABLE tbl_name table_options;
```

table_options :

```
    table_option [, table_option]
```

table option :

```
    CHANGE [COLUMN | CLASS ATTRIBUTE] old_col_name new_col_name column_definition
           [FIRST | AFTER col_name]
    | MODIFY [COLUMN | CLASS ATTRIBUTE] col_name column_definition
           [FIRST | AFTER col_name]
```

- *tbl_name* : Specifies the name of the table including the column to change.
- *old_col_name* : Specifies the existing column name.
- *new_col_name* : Specifies the column name to change
- *column_definition* : Specifies the type and the attribute of the column to change.

- *col_name* : Specifies the column name to which the type and the attribute of the column to apply changes.

Example 1

```
CREATE TABLE t1 (a INTEGER);
ALTER TABLE t1 CHANGE a b INTEGER;

-- changing a column's constraint
ALTER TABLE t1 CHANGE a a INTEGER NOT NULL;
ALTER TABLE t1 MODIFY a INTEGER NOT NULL;

-- changing a column's type - "DEFAULT 1" constraint is removed.
CREATE TABLE t1 (col1 INT DEFAULT 1);
ALTER TABLE t1 MODIFY col1 BIGINT;

-- changing a column's type - "DEFAULT 1" constraint is kept.
CREATE TABLE t1 (col1 INT DEFAULT 1);
ALTER TABLE t1 MODIFY col1 BIGINT DEFAULT 1;
```

Example 2

```
-- changing the name and position of a column
CREATE TABLE t1(i1 int,i2 int);
INSERT INTO t1 VALUE (1,11),(2,22),(3,33);
SELECT * FROM t1 ORDER BY 1;
      i1      i2
=====
      1      11
      2      22
      3      33

ALTER TABLE t1 CHANGE i2 i0 INTEGER FIRST;
SELECT * FROM t1 ORDER BY 1;
      i0      i1
=====
      11      1
      22      2
      33      3
```

Example 3

```
-- adding NOT NULL constraint (strict)
-- alter table change type strict=yes

CREATE TABLE t1(i int);
INSERT INTO t1 values (11),(NULL),(22);

ALTER TABLE t1 change i i1 integer not null;

In the command from line 1,

ERROR: Cannot add NOT NULL constraint for attribute "i1": there are existing NULL values
for this attribute.
```

Example 4

```
-- adding NOT NULL constraint
-- alter table change type strict=no

CREATE TABLE t1(i int);
INSERT INTO t1 VALUES (11),(NULL),(22);

ALTER TABLE t1 CHANGE i i1 INTEGER NOT NULL;

SELECT * FROM t1;

      i1
=====
      22
      0
      11
```

Example 5

```
-- change the column's data type (no errors)

CREATE TABLE t1 (i1 int);
INSERT INTO t1 VALUES (1), (-2147483648), (2147483647);

ALTER TABLE t1 CHANGE i1 s1 CHAR(11);

SELECT * FROM t1;

  s1
=====
'2147483647 '
'-2147483648'
'1          '
```

Example 6

```
-- change the column's data type (errors), strict mode
-- alter_table_change_type_strict=yes

CREATE TABLE t1 (i1 int);
INSERT INTO t1 VALUES (1), (-2147483648), (2147483647);

ALTER TABLE t1 CHANGE i1 s1 CHAR(4);

In the command from line 1,

ERROR: ALTER TABLE .. CHANGE : changing to new domain : cast failed, current configuration
doesn't allow truncation or overflow.

-- change the column's data type (errors)
-- alter table change type strict=no

CREATE TABLE t1 (i1 INT);
INSERT INTO t1 VALUES (1), (-2147483648), (2147483647);

ALTER TABLE t1 CHANGE i1 s1 CHAR(4);

SELECT * FROM t1;

  s1
=====
'   '
'   '
'1  '
```

-- hard default values have been placed instead of signaling overflow

Syntax Operation According to Column Attributes

- **Type Change** : If the value of the system parameter `alter_table_change_type_strict` is set to no, then changing values to other types is allowed, but if it is set to yes then changing is not allowed. The default value of the parameter is **no**. You can change values to all types allowed by the **CAST** operator. Changing object types is allowed only by the upper classes (tables) of the objects.
- **NOT NULL**
- If the **NOT NULL** constraint is not specified, it will be removed from a new table even though it is present in the existing table.
- If the **NOT NULL** constraint is specified in the column to change, the result varies depending on the configuration of the system parameter, `alter_table_change_type_strict`.
- If `alter_table_change_type_strict` is set to yes, the column values will be checked. If **NULL** exists, an error will occur, and the change will not be executed.
- If the `alter_table_change_type_strict` is set to no, every existing **NULL** value will be changed to a hard default value of the type to change.
- **DEFAULT** : If the **DEFAULT** attribute is not specified in the column to change, it will be removed from a new table even though it is present in the existing table.

- **AUTO_INCREMENT** : If the **AUTO_INCREMENT** attribute is not specified in the column to change, it will be removed from a new table even though it is present in the existing table.
- **FOREIGN KEY** : You cannot change the column with the foreign key constraint that is referred to or refers to.
- Single Column **PRIMARY KEY**
- If the **PRIMARY KEY** constraint is specified in the column to change, a **PRIMARY KEY** is re-created only in which a **PRIMARY KEY** constraint exists in the existing column and the type is upgraded.
- If the **PRIMARY KEY** constraint is specified in the column to change but doesn't exist in the existing column, a **PRIMARY KEY** will be created.
- If a **PRIMARY KEY** constraint exists but is not specified in the column to change, the **PRIMARY KEY** will be maintained.
- Multicolumn **PRIMARY KEY** : If the **PRIMARY KEY** constraint is specified and the type is upgraded, a **PRIMARY KEY** will be re-created.
- Single Column **UNIQUE KEY**
- If the type is upgraded, a **UNIQUE KEY** will be re-created.
- If a **UNIQUE KEY** exists in the existing column and it is not specified in the column to change, it will be maintained.
- If a **UNIQUE KEY** exists in the existing column to change, it will be created.
- Multicolumn **UNIQUE KEY** : If the column type is changed, an index will be re-created.
- Column with a Non-unique Index : If the column type is changed, an index will be re-created.
- Partition Column: If a table is partitioned by a column, the column cannot be changed. Partitions cannot be added.
- Column with a Class Hierarchy : You can only change the tables that do not have a lower class. You cannot change the lower class that inherits from an upper class. You cannot change the inherited attributes.
- Trigger and View : You must redefine triggers and views directly because they are not changed according to the definition of the column to change.
- Column Sequence : You can change the sequence of columns.
- Name Change : You can change names as long as they do not conflict.

Syntax Operation According to the System Parameter, `alter_table_change_type_strict`

The `alter_table_change_type_strict` parameter determines whether the value conversion is allowed according to the type change. If the value is no, it can be changed when you change a column type or add a **NOT NULL** constraint. The default value is **no**.

When the value of the parameter, `alter_table_change_type_strict` is no, it will operate depending on the conditions as follows:

- Overflow Occurred while Converting Numbers or Character Strings to Numbers: The minimum value or the maximum value is specified according to the result type conditions, and the warning message will be recorded in the log for the record where overflow has occurred.
- If input values are numbers, their signs will be written to the log.
- If input values are character strings, the signs of the values converted to **DOUBLE** types will be written in the log.
- Character Strings to Convert to Shorter Ones: The record will be updated to the hard default value of the type that is defined and the warning message will be recorded in a log.
- Conversion Failure Due to Other Reasons : The record will be updated to the hard default value of the type that is defined and the warning message will be recorded in a log.

If the value of the `alter_table_change_type_strict` parameter is yes, an error message will be displayed and the changes will be rolled back.

The **ALTER CHANGE** statement checks the possibility of type conversion before updating a record but the type conversion of specific values may fail. For example, if the value format is not correct when you convert **VARCHAR** to **DATE**, the conversion may fail. In this case, the hard default value of the **DATE** type will be assigned.

The hard default value is a value that will be used when you add columns with the **ALTER TABLE** `ALTER TABLE` `ADD COLUMN` statement, add or change by converting types with the **ALTER TABLE** `ALTER TABLE` `CHANGE/MODIFY` statement. The operation will vary depending on the system parameter, `add_column_update_hard_default` in the **ADD COLUMN** statement.

Hard Default Value by Type

Type	Existence of Hard Default Value	Hard Default Value
INTEGER	Yes	0
FLOAT	Yes	0
DOUBLE	Yes	0
SMALLINT	Yes	0
DATE	Yes	date'01/01/0001'
TIME	Yes	time'00:00'
DATETIME	Yes	datetime'01/01/0001 00:00'
TIMESTAMP	Yes	timestamp'00:00:01 AM 01/01/1970' (GMT)
MONETARY	Yes	0
NUMERIC	Yes	0
CHAR	Yes	"
VARCHAR	Yes	"
NCHAR	Yes	N"
VARNCHAR	Yes	N"
SET	Yes	{}
MULTISET	Yes	{}
SEQUENCE	Yes	{}
BIGINT	Yes	0
BIT	Yes	
VARBIT	No	
OBJECT	No	
BLOB	No	
CLOB	No	
ELO	No	

RENAME COLUMN Clause**Description**

You can change the name of the column by using the **RENAME COLUMN** clause.

Syntax

```
ALTER [ TABLE | CLASS | VCLASS | VIEW ] table_name
RENAME [ COLUMN | ATTRIBUTE ] old_column_name { AS | TO } new_column_name
```

- *table_name* : Specifies the name of a table that has a column to be renamed.
- *old_column_name* : Specifies the name of a column.
- *new_column_name* : Specifies a new column name after the **AS** keyword.

Example

```
ALTER TABLE a_tbl RENAME COLUMN name AS name1;
```

DROP COLUMN Clause

Description

You can delete a column in a table by using the **DROP COLUMN** clause. You can specify multiple columns to delete simultaneously by separating them with commas (,).

Syntax

```
ALTER [ TABLE | CLASS | VCLASS | VIEW ] table_name
DROP [ COLUMN | ATTRIBUTE ] column_name, ...
```

- *table_name* : Specifies the name of a table that has a column to be deleted.
- *column_name* : Specifies the name of a column to be deleted. Multiple columns can be specified by separating them with commas (,).

Example

```
ALTER TABLE a_tbl DROP COLUMN age1,age2,age3;
```

DROP CONSTRAINT Clause

Description

You can drop the constraints pre-defined for the table, such as **UNIQUE**, **PRIMARY KEY** and **FOREIGN KEY** by using the **DROP CONSTRAINT** clause. In this case, you must specify a constraint name. You can check these names by using the CSQL command (`;<schema table_name>`).

Syntax

```
ALTER [ TABLE | CLASS ] table_name
DROP CONSTRAINT constraint_name
```

- *table_name* : Specifies the name of a table that has a constraint to be dropped.
- *constraint_name* : Specifies the name of a constraint to be dropped.

Example

```
ALTER TABLE a_tbl DROP CONSTRAINT pk_a_tbl_id;
ALTER TABLE a_tbl DROP CONSTRAINT fk_a_tbl_id;
ALTER TABLE a_tbl DROP CONSTRAINT u_a_tbl_id;
```

DROP INDEX Clause

Description

You can delete an index defined for a column by using the **DROP INDEX** clause.

Syntax

```
ALTER [ TABLE | CLASS ] table_name DROP INDEX index_name
```

- *table_name* : Specifies the name of a table that has an index attribute to be deleted.
- *index_name* : Specifies the name of an index to be deleted.

Example

```
ALTER TABLE a_tbl DROP INDEX i_a_tbl_age;
```

DROP PRIMARY KEY Clause

Description

You can delete a primary key constraint defined for a table by using the **DROP PRIMARY KEY** clause. You do have to specify the name of the primary key constraint because only one primary key can be defined by table.

Syntax

```
ALTER [ TABLE | CLASS ] table_name DROP PRIMARY KEY
```

- *table_name* : Specifies the name of a table that has a primary key constraint to be deleted.

Example

```
ALTER TABLE a_tbl DROP PRIMARY KEY;
```

DROP FOREIGN KEY Clause

Description

You can drop a foreign key constraint defined for a table using the **DROP FOREIGN KEY** clause.

Syntax

```
ALTER [ TABLE | CLASS ] table_name DROP FOREIGN KEY constraint_name
```

- *table_name* : Specifies the name of a table whose constraint is to be deleted.
- *constraint_name* : Specifies the name of foreign key constraint to be deleted.

Example

```
ALTER TABLE a_tbl DROP FOREIGN KEY fk_a_tbl_id;
```

DROP TABLE

Description

You can drop an existing table by the **DROP** statement. Multiple tables can be dropped by a single **DROP** statement. All rows of table are also dropped. If you use it together with the **IF EXISTS** statement, you can prevent errors from occurring and specify multiple tables in one statement.

Syntax

```
DROP [ TABLE | CLASS ] [ IF EXISTS ] <table specification comma list>
```

```
<table_specification_comma_list> ::=
<single_table_spec> | ( <table_specification_comma_list> )
```

```
<single_table_spec> ::=
|[ ONLY ] table name
| ALL table_name [ ( EXCEPT table_name, ... ) ]
```

- *table_name* : Specifies the name of the table to be dropped. You can delete multiple tables simultaneously by separating them with commas.
- If a super class name is specified after the **ONLY** keyword, only the super class, not the subclasses inheriting from it, is deleted. If a super class name is specified after the **ALL** keyword, the super classes as well as the subclasses inheriting from it are all deleted. You can specify the list of subclasses not to be deleted after the **EXCEPT** keyword.
- If subclasses that inherit from the super class specified after the **ALL** keyword are specified after the **EXCEPT** keyword, they are not deleted.

Example

```
DROP TABLE history ;
```

```
CREATE TABLE t (i INT);

-- DROP TABLE IF EXISTS
DROP TABLE IF EXISTS history, t;
2 command(s) successfully processed.
SELECT * FROM t; In line 1, column 10, ERROR: Unknown class "t".
```

RENAME TABLE

Description

You can change the name of a table by using the **RENAME TABLE** statement and specify a list of the table name to change the names of multiple tables.

Syntax

```
RENAME [ TABLE | CLASS | VIEW | VCLASS ] old_table_name { AS | TO } new_table_name [,
old_table_name { AS | TO } new_table_name, ... ]
```

- *old_table_name* : Specifies the old table name to be renamed.
- *new_table_name* : Specifies a new table name.

Example

```
RENAME TABLE a_tbl AS aa_tbl;

RENAME TABLE a_tbl TO aa_tbl, b_tbl TO bb_tbl;
```

Caution

The table name can be changed only by the table owner, **DBA** and **DBA** members. The other users must be granted to change the name by the owner or **DBA** (see [Granting Authorization](#) For details on authorization).

Index Definition

CREATE INDEX

Description

Use the **CREATE INDEX** statement to create an index in the specified table.

Syntax

```
CREATE [ REVERSE ] [ UNIQUE ] INDEX [ index_name ]
ON table_name ( column_name[(prefix_length)] [ASC | DESC] [ {, column_name[(prefix_length)]
[ASC | DESC]} ... ] ) [ ; ]
```

- **REVERSE** : Creates an index in the reverse order. A reverse index helps to increase sorting speed in descending order.
- **UNIQUE** : Creates an index with unique values.
- *index_name* : Specifies the name of the index to be created. The index name must be unique in the table. If omitted, a name is automatically assigned.
- *prefix_length* : When you specify an index for character- or bit string-type column, you can create an index by specifying the beginning part of the column name as a prefix. You can specify the length of the prefix in bytes in parentheses next to the column name. You cannot specify *prefix_length* in a multiple column index or a **UNIQUE** index. It is impossible to create an index by specifying *prefix_length* as a host variable. If you want to guarantee the query result order in the index in which *prefix_length* is specified, you must specify the **ORDER BY** clause.
- *table_name* : Specifies the name of the table where the index is to be created.
- *column_name* : Specifies the name of the column where the index is to be applied. To create a composite index, specify two or more column names.
- **ASC | DESC** : Specifies the sorting order of columns. In case of a **REVERSE** index, **ASC** is ignored and **DESC** is applied.

Example 1

The following example shows how to create a reverse index.

```
CREATE REVERSE INDEX gold_index ON participant(gold);
```

Example 2

The following example shows how to create a multiple column index.

```
CREATE INDEX name_nation_idx ON athlete(name, nation_code);
```

Example 3

The following example shows how to create a single column index. In this example, 1-byte long prefix is specified for the *nation_code* column when creating an index.

```
CREATE INDEX ON game(nation_code(1));
CREATE INDEX game_date_idx ON game(game_date);
```

ALTER INDEX

Description

Use the **ALTER INDEX** statement to rebuild an index. (That is, drop and rebuild an index.) There are the following two ways to specify an index to be rebuilt:

- Specifying it as the name of the index
- Specifying it as the name of the table or the column where the index is specified

Syntax

```
ALTER [ REVERSE ] [ UNIQUE ] INDEX index_name
[ON { ONLY } table_name ( column_name [ {, column_name } ...] ) ] REBUILD [ ; ]
```

```
ALTER [ REVERSE ] [ UNIQUE ] INDEX
ON { ONLY } table_name ( column_name [ {, column_name } ...] ) REBUILD [ ; ]
```

- **REVERSE** : Creates an index in the reverse order. A reverse index helps to increase sorting speed in descending order.
- **UNIQUE** : Creates an index with unique values.
- *index_name* : Specifies the name of the index to be altered. The index name must be unique in the table.
- *table_name* : Specifies the name of the table where the index is to be created.
- *column_name* : Specifies the name of the column where the index is to be applied. To create a multiple column index, specify two or more column names.

Example

The following are examples of many ways of re-creating indexes:

```
ALTER INDEX i_game_medal ON game(medal) REBUILD;
ALTER INDEX game_date_idx REBUILD;
```

DROP INDEX

Description

Use the **DROP INDEX** statement to drop an index. There are the following two ways to specify the index to be dropped:

- To specify the name of the index
- To specify the name of the table or the column where the index is specified

Syntax

```
DROP [ REVERSE ] [ UNIQUE ] INDEX index_name
[ON table_name ( column_name [ {, column_name } ...] ) ] [ ]
```

```
DROP [ REVERSE ] [ UNIQUE ] INDEX
ON table_name ( column_name [ {, column_name } ...] ) [ ]
```

- **REVERSE** : Specifies that the index to be dropped is a reverse index.
- **UNIQUE** : Specifies that the index to be dropped is a unique index.
- *index_name* : Specifies the name of the index to be dropped.
- *table_name* : Specifies the name of the table whose index is to be dropped.
- *column_name* : Specifies the name of the column whose index is to be dropped.

Example

The following are examples of many ways of dropping indexes:

```
DROP INDEX ON game(medal);
DROP INDEX game_date_idx;
DROP REVERSE INDEX gold_index ON participant(gold);
DROP INDEX name_nation_idx ON athlete(name, nation_code);
```

VIEW Definition

CREATE VIEW

Overview

Description

A view is a virtual table that does not exist physically. You can create a view by using an existing table or a query. **VIEW** and **VCLASS** are used interchangeably.

Use **CREATE VIEW** statement to create a view.

Syntax

```
CREATE [OR REPLACE] {VIEW | VCLASS} <view name>
    [ <subclass_definition> ]
    [ ( <view_column_def_comma_list> ) ]
    [ CLASS ATTRIBUTE
      ( <column definition comma list> ) ]
    [ METHOD <method_definition_comma_list> ]
    [ FILE <method_file_comma_list> ]
    [ INHERIT <resolution_comma_list> ]
    [ AS <select_statement> ]
    [ WITH CHECK OPTION ]

<view column definition> ::= <column definition> | <column name>

<column_definition> :
column name column type [ <default or shared> ] [ <column constraint list> ]

<default or shared> :
{SHARED [ <value specification> ] | DEFAULT <value specification> } |
AUTO_INCREMENT [ (seed, increment) ]

<column constraint> :
NOT NULL | UNIQUE | PRIMARY KEY | FOREIGN KEY REFERENCES...

<subclass definition> :
{ UNDER | AS SUBCLASS OF } table_name_comma_list

<method definition> :
[ CLASS ] method_name
[ ( [ argument type comma list ] ) ]
[ result type ]
[ FUNCTION function_name ]

<resolution> :
[ CLASS ] { column_name | method_name } OF superclass_name
[ AS alias ]
```

- **OR REPLACE** : If the keyword **OR REPLACE** is specified after **CREATE**, the existing view is replaced by a new one without displaying any error message, even when the *view_name* overlaps with the existing view name.
- *view_name* : Specify the name of a view to be created. It must be unique in a database.
- *view_column_definition*
- *column_name* : Defines the column of a view.
- *column_type* : Specifies the data type of a column.

AS select_statement : A valid **SELECT** statement must be specified. A view is created based on this.

WITH CHECK OPTION : If this option is specified, the update or insert operation is possible only when the condition specified in the **WHERE** clause of the *select_statement* is satisfied. Therefore, this option is used to disallow the update of a virtual table that violates the condition.

Example

```

CREATE TABLE a tbl(
id INT NOT NULL,
phone VARCHAR(10));
INSERT INTO a_tbl VALUES(1,'111-1111'), (2,'222-2222'), (3, '333-3333'), (4, NULL), (5,
NULL);

--creating a new view based on AS select statement from a tbl
CREATE VIEW b view AS SELECT * FROM a tbl WHERE phone IS NOT NULL WITH CHECK OPTION;
SELECT * FROM b_view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'

--WITH CHECK OPTION doesn't allow to update column value which violates WHERE clause
UPDATE b view SET phone=NULL;

In line 1, column 72,
ERROR: Check option exception on view b view.

--creating view which name is as same as existing view name
CREATE OR REPLACE VIEW b_view AS SELECT * FROM a_tbl ORDER BY id DESC;

--the existing view has been replaced as a new view by OR REPLACE keyword
SELECT * FROM b view;

      id  phone
=====
      5  NULL
      4  NULL
      3  '333-3333'
      2  '222-2222'
      1  '111-1111'

```

Condition for Creating Updatable VIEW

Description

To update data in a virtual table, it must be updatable because an option is needed to define data.

A virtual table is updatable if it satisfies the following conditions:

- The **FROM** clause must include only one table or updatable view. However, two tables included in parentheses as in **FROM** (class_x, class_y) can be updated because they represent one table.
- The **DISTINCT** or **UNIQUE** statement should not be included.
- The **GROUP BY... HAVING** statement should not be included.
- Aggregate functions such as **SUM()** or **AVG()** should not be included.
- The entire query must consist of queries that can be updated by **UNION ALL**, not by **UNION**. However, the table should exist only in one of the queries that constitute **UNION ALL**.
- If a row is inserted into a view created by using the **UNION ALL** statement, the system determines into which table the row will be inserted. This cannot be done by the user. To control this, the user must manually insert the row or create a separate view for insertion.

Even when all rules above are satisfied, each column of the a view may not be updatable. The following columns cannot be updated:

- Path expressions
- Numeric type column that includes an arithmetic operator

Even though the column defined in the view is updatable, a view can be updated only when an appropriate update authorization is granted on the table included in the **FROM** clause. Also there must be an access authorization to a

view. The way to grant an access authorization to a view is the same to grant an access authorization to a table. For details on granting authorization, see "[Granting Authorization](#)".

ALTER VIEW

ADD QUERY Clause

Description

You can add a new query to a query specification by using the **ADD QUERY** clause of the **ALTER VIEW** statement. 1 is assigned to the query defined when a virtual table was created, and 2 is assigned to the query added by the **ADD QUERY** clause.

Syntax

```
ALTER [ VIEW | VCLASS ] view_name
ADD QUERY select_statement
[ INHERIT resolution [ {, resolution }_ ] ]

resolution :
{ column_name | method_name } OF superclass_name [ AS alias ]
```

- *view_name* : Specifies the name of a view where the query to be added.
- *select_statement* : Specifies the query to be added.

Example

```
SELECT * FROM b view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
      4  NULL
      5  NULL

ALTER VIEW b view ADD QUERY SELECT * FROM a tbl WHERE id IN (1,2);
SELECT * FROM b view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
      4  NULL
      5  NULL
      1  '111-1111'
      2  '222-2222'
```

AS SELECT Clause

Description

You can change the **SELECT** query defined in the virtual table by using the **AS SELECT** clause in the **ALTER VIEW** statement. This function is working like the **CREATE OR REPLACE** statement. You can also change the query by specifying the query number 1 in the **CHANGE QUERY** clause of the **ALTER VIEW** statement.

Syntax

```
ALTER [ VIEW | VCLASS ] view_name AS select_statement
```

- *view_name* : Specifies the name of a view to be modified.

- *select_statement* : Specifies the new query statement to replace the **SELECT** statement defined when a view is created.

Example

```
ALTER VIEW b view AS SELECT * FROM a tbl WHERE phone IS NOT NULL;
SELECT * FROM b_view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
```

CHANGE QUERY Clause

Description

You can change the query defined in the query specification by using the **CHANGE QUERY** clause reserved word of the **ALTER VIEW** statement.

Syntax

```
ALTER [ VIEW | VCLASS ] view_name
      CHANGE QUERY [ integer ] select_statement [ ; ]
```

- *view_name* : Specifies the name of a view to be modified.
- *integer* : Specifies the number value of the query to be modified. The default value is 1.
- *select_statement* : Specifies the new query statement to replace the query whose query number is *integer*.

Example

```
--adding select statement which query number is 2 and 3 for each
ALTER VIEW b view ADD QUERY SELECT * FROM a tbl WHERE id IN (1,2);
ALTER VIEW b view ADD QUERY SELECT * FROM a tbl WHERE id = 3;
SELECT * FROM b_view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
      4  NULL
      5  NULL
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'

--altering view changing query number 2
ALTER VIEW b view CHANGE QUERY 2 SELECT * FROM a tbl WHERE phone IS NULL;
SELECT * FROM b view;

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
      4  NULL
      5  NULL
      4  NULL
      5  NULL
      3  '333-3333'
```

DROP QUERY Clause

Description

You can drop a query defined in the query specification by using the **DROP QUERY** of the **ALTER VIEW** statement.

Example

```
ALTER VIEW b view DROP QUERY 2,3;
SELECT * FROM b view;
```

```

      id  phone
=====
      1  '111-1111'
      2  '222-2222'
      3  '333-3333'
      4  NULL
      5  NULL
```

DROP VIEW**Description**

You can drop a view by using the **DROP VIEW** clause. The way to drop a view is the same as to drop a regular table.

Syntax

```
DROP [ VIEW | VCLASS ] view_name [ { ,view_name , ... } ]
```

- *view_name* : Specifies the name of a view to be dropped.

Example

```
DROP VIEW b_view;
```

RENAME VIEW**Description**

```
RENAME [ TABLE | CLASS | VIEW | VCLASS ] old_view_name AS new_view_name [ ; ]
```

- *old_view_name* : Specifies the name of a view to be modified.
- *new_view_name* : Specifies the new name of a view.

Example

The following example shows how to rename a view name to game_2004.

```
RENAME VIEW game_2004 AS info_2004;
```

SERIAL

CREATE SERIAL

Serial is an object that creates a unique sequence number, and has the following characteristics.

- The serial is useful in creating a unique sequence number in multi-user environment.
- Generated serial numbers are not related with table so, you can use the same serial in multiple tables.
- All users including **public** can create a serial object. Once it is created, all users can get the number by using **CURRENT_VALUE** and **NEXT_VALUE**.
- Only owner of a created serial object and **dba** can update or delete a serial object. If an owner is **public**, all users can update or delete it.

Description

You can create a serial object in the database by using the **CREATE SERIAL** statement.

Syntax

```
CREATE SERIAL serial_name
[ START WITH initial ]
[ INCREMENT BY interval ]
[ MINVALUE min | NOMINVALUE ]
[ MAXVALUE max | NOMAXVALUE ]
[ CACHE integer | NOCACHE ]
```

- *serial_identifier* : Specifies the name of the serial to be generated.
- **START WITH initial** : Specifies the initial value of serial with 38 digits or less. The default value of ascending serial is 1 and that of descending serial is -1.
- **INCREMENT BY interval** : Specifies the increment of the serial. You can specify any integer with 38 digits or less except zero at *interval*. The absolute value of the *interval* must be smaller than the difference between **MAXVALUE** and **MINVALUE**. If a negative number is specified, the serial is in descending order otherwise, it is in ascending order. The default value is 1.
- **MINVALUE** : Specifies the minimum value of the serial, with 38 digits or less. **MINVALUE** must be smaller than or equal to the initial value and smaller than the maximum value.
- **NOMINVALUE** : 1 is set automatically as a minimum value for the ascending serial $-(10)^{38}$ for the descending serial.
- **MAXVALUE** : Specifies the maximum number of the serial with 38 digits or less. **MAXVALUE** must be smaller than or equal to the initial value and greater than the minimum value.
- **NOMAXVALUE** : $(10)^{37}$ is set automatically as a maximum value for the ascending serial -1 for the descending serial.
- **CYCLE** : Specifies that the serial will be generated continuously after reaching the maximum or minimum value. When a serial in ascending order reaches the maximum value, the minimum value is created as the next value; when a serial in descending order reaches the minimum value, the maximum value is created as the next value.
- **NOCYCLE** : Specifies that the serial will not be generated any more after reaching the maximum or minimum value. The default value is **NOCYCLE**.
- **CACHE** : Stores as many serials as the number specified by "integer" in the cache to improve the performance of the serials and fetches a serial value when one is requested. If all cached values are used up, as many serials as "integer" are fetched again from the disk to the memory. If the database server stops accidentally, all cached serial values are deleted. For this reason, the serial values before and after the restart of the database server may be discontinuous. Because the transaction rollback does not affect the cached serial values, the request for the next serial will return the next value of the value used (or fetched) lastly when the transaction is rolled back. The "integer" after the **CACHE** keyword cannot be omitted. If the "integer" is equal to or smaller than 1, the serial cache is not applied.
- **NOCACHE** : Specifies that the serial cache is not used, and serial values are updated and retrieved from a disk upon every request.

Example 1

```
--creating serial with default values
CREATE SERIAL order no;

--creating serial within a specific range
CREATE SERIAL order_no START WITH 10000 INCREMENT BY 2 MAXVALUE 20000;
--creating serial with specifying the number of cached serial values
CREATE SERIAL order no START WITH 10000 INCREMENT BY 2 MAXVALUE 20000 CACHE 3;

--selecting serial information from the db serial class
SELECT * FROM db_serial;
```

name	current val	increment val	max val	min val	c
yclic	started	cached num	att name		
'order_no'	10006	2	20000	10000	
0	1	3	NULL		

Example 2

The following example shows how to create the athlete_idx table to store athlete codes and names and then create an instance by using the *order_no*. NEXT_VALUE increases the serial number and returns its value.

```
CREATE TABLE athlete_idx( code INT, name VARCHAR(40) );
CREATE SERIAL order_no START WITH 10000 INCREMENT BY 2 MAXVALUE 20000;
INSERT INTO athlete_idx VALUES (order_no.NEXT_VALUE, 'Park');
INSERT INTO athlete_idx VALUES (order no.NEXT VALUE, 'Kim');
INSERT INTO athlete_idx VALUES (order no.NEXT VALUE, 'Choo');
INSERT INTO athlete_idx VALUES (order no.CURRENT VALUE, 'Lee');
SELECT * FROM athlete_idx;
```

code	name
10000	'Park'
10002	'Kim'
10004	'Choo'
10004	'Lee'

ALTER SERIAL**Description**

With the **ALTER SERIAL** statement, you can update the increment of the serial value, set or delete its initial or minimum/maximum values, and set its cycle attribute.

Syntax

```
ALTER SERIAL serial identifier
[ INCREMENT BY interval ]
[ START WITH initial_value ]
[ MINVALUE min | NOMINVALUE ]
[ MAXVALUE max NOMAXVALUE ]
[ CACHE integer | NOCACHE ]
```

- *serial_identifier* : Specifies the name of the serial to be created.
- **INCREMENT BY** *interval* : Specifies the increment of the serial. For the *interval*, you can specify any integer with 38 digits or less except zero. The absolute value of the *interval* must be smaller than the difference between **MAXVALUE** and **MINVALUE**. If a negative number is specified, the serial is in descending order; otherwise, it is in ascending order. The default value is 1.
- **START WITH** *initial_value* : Changes the initial value of Serial.
- **MINVALUE** : Specifies the minimum value of the serial with 38 digits or less. **MINVALUE** must be smaller than or equal to the initial value and smaller than the maximum value. **NOMINVALUE** : 1 is set automatically as a minimum value for the ascending serial; $-(10)^{36}$ for the descending serial.
- **MAXVALUE** : Specifies the maximum number of the serial with 38 digits or less. **MAXVALUE** must be smaller than or equal to the initial value and greater than the minimum value.

- **NOMAXVALUE** : $(10)^{37}$ is set automatically as a maximum value for the ascending serial; -1 for the descending serial.
- **CYCLE** : Specifies that the serial will be generated continuously after reaching the maximum or minimum value. If the ascending serial reaches the maximum value, the minimum value is generated as the next value. If the descending serial reaches the minimum value, the maximum value is generated as the next value.
- **NOCYCLE** : Specifies that the serial will not be generated any more after reaching the maximum or minimum value. The default is **NOCYCLE**.
- **CACHE** : Stores as many serials as the number specified by integer in the cache to improve the performance of the serials and fetches a serial value when one is requested. The "integer" after the **CACHE** keyword cannot be omitted. If a number equal to or smaller than 1 is specified, the serial cache is not applied.
- **NOCACHE** : It does not use the serial cache feature. The serial value is updated every time and a new serial value is fetched from the disk upon each request.

Caution In CUBRID 2008 R1.x version, the serial value can be modified by updating the db_serial table, a system catalog. However, in CUBRID 2008 R2.0 version or above, the modification of the db_serial table is not allowed but use of the **ALTER SERIAL** statement is allowed. Therefore, if an **ALTER SERIAL** statement is included in the data exported (unloaddb) from CUBRID 2008 R2.0 or above, it is not allowed to import (loaddb) the data in CUBRID 2008 R1.x or below.

Example

```
--altering serial by changing start and incremental values
ALTER SERIAL order_no START WITH 100 INCREMENT BY 2;

--altering serial to operate in cache mode
ALTER SERIAL order_no CACHE 5;

--altering serial to operate in common mode
ALTER SERIAL order_no NOCACHE;
```

DROP SERIAL

Description

With the **DROP SERIAL** statement, you can drop a serial object from the database.

Syntax

```
DROP SERIAL serial_identifier
```

- *serial_identifier* : Specifies the name of the serial to be dropped.

Example

The following example shows how to drop the *order_no* serial.

```
DROP SERIAL order_no;
```

Use SERIAL

Description

You can access and update a serial by serial name and a reserved word pair.

Syntax

```
serial_identifier.CURRENT_VALUE
serial_identifier.NEXT_VALUE
```

- *serial_identifier*.**CURRENT_VALUE** : Returns the current serial value.
- *serial_identifier*.**NEXT_VALUE** : Increments the serial value and returns the result.

Example

The following example shows how to create a table `athlete_idx` where athlete numbers and names are stored and how to create the instances by using a serial `order_no`.

```
CREATE TABLE athlete_idx( code INT, name VARCHAR(40) );
INSERT INTO athlete_idx VALUES (order_no.NEXT_VALUE, 'Park');
INSERT INTO athlete_idx VALUES (order no.NEXT_VALUE, 'Kim');
INSERT INTO athlete_idx VALUES (order no.NEXT VALUE, 'Choo');
INSERT INTO athlete_idx VALUES (order no.NEXT VALUE, 'Lee');SELECT * FROM athlete idx;
```

code	name
10000	'Park'
10002	'Kim'
10004	'Choo'
10006	'Lee'

Caution

When you use a serial for the first time after creating it, `NEXT_VALUE` returns the initial value. Subsequently, the sum of the current value and the increment are returned.

Serial Function

Description

The **Serial** function consists of the `SERIAL_CURRENT_VALUE` and `SERIAL_NEXT_VALUE` functions.

The `SERIAL_CURRENT_VALUE` function returns the current serial value, which is the same value as `serial_name.current_value`.

This function returns as much added value as interval specified. The serial interval is determined by the value of a `CREATE SERIAL ... INCREMENT BY` statement. `SERIAL_NEXT_VALUE(serial_name, 1)` returns the same value as `serial_name.next_value`.

To get a large amount of serials at once, specify the desired number as an argument to call the `SERIAL_NEXT_VALUE` function only once; which has an advantage over calling repeatedly `serial_name.next_value` in terms of performance.

Assume that an application process is trying to get the number of `n` serials at once. To perform it, call `SERIAL_NEXT_VALUE(serial_name, N)` one time to store a return value and calculate a serial value between (a serial start value) and (the return value). (Serial value at the point of function call) is equal to the value of (return value) - (desired number of serials) * (serial interval).

For example, if you create a serial starting 101 and increasing by 1 and call `SERIAL_NEXT_VALUE(serial_name, 10)`, it returns 110. The start value at the point is $110 - (10 - 1) * 1 = 101$. Therefore, 10 serial values such as 101, 102, 103, ... 110 can be used by an application process. If `SERIAL_NEXT_VALUE(serial_name, 10)` is called in succession, 120 is returned; the start value at this point is $120 - (10 - 1) * 1 = 111$.

Syntax

```
SERIAL_CURRENT_VALUE(serial_name)
SERIAL_NEXT_VALUE(serial_name, number)
```

- `serial_name` : Serial name
- `number` : The number of serials to be obtained

Example

```
CREATE SERIAL order_no START WITH 10000 INCREMENT BY 2 MAXVALUE 20000;
SELECT SERIAL CURRENT VALUE(order no);
10000

-- At first, the first serial value starts with the initial serial value, 10000. So the
10'th serial value will be 10009.
```

```
SELECT SERIAL NEXT VALUE(order no, 10);  
10009  
  
SELECT SERIAL_NEXT_VALUE(order_no, 10);  
10019
```

Caution

If you create a serial and calls the **SERIAL_NEXT_VALUE** function for the first time, a value of (serial interval) * (desired number of serials - 1) added to the current value is returned. If you call the **SERIAL_NEXT_VALUE** function in succession, a value of (serial interval) * (desired number of serials) added to the current is returned (see the example above).

Operators and Functions

Logical Operators

Description

For logical operators, boolean expressions or expressions that evaluates to an integer value are specified as operands; **TRUE**, **FALSE** or **NULL** is returned as the result. If the **INTEGER** value is used, 0 is evaluated to **FALSE** and the other values are evaluated to **TRUE**. If a boolean value is used, 1 is evaluated to **TRUE** and 0 is evaluated to **FALSE**.

The following table shows the logic operators supported by CUBRID.

Logical Operators Supported by CUBRID

Logical Operator	Description	Condition
AND, &&	If all operands are TRUE , it returns TRUE .	a AND b
OR, 	If none of operands is NULL and one or more operand is TRUE , it returns TRUE . If pipes_as_concat is no that is a parameter related to SQL statements, a double pipe symbol can be used as OR operator.	a OR b
XOR	If none of operand is NULL and each of operand has a different value, it returns TRUE .	a XOR b
NOT, !	A unary operator. If a operand is FALSE , it returns TRUE . If it is TRUE , returns FALSE .	NOT a

True Table of Logical Operators

a	b	a AND b	a OR b	NOT a	a XOR b
TRUE	TRUE	TRUE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	TRUE
TRUE	NULL	NULL	TRUE	FALSE	NULL
FALSE	FALSE	FALSE	FALSE	TRUE	FALSE
FALSE	NULL	FALSE	NULL	TRUE	NULL

Remark

You should put the logical expressions in brackets in the **SELECT** list.

```
SELECT 1 = 1 FROM db root;
ERROR: syntax error, unexpected '='

SELECT (1 = 1) FROM db_root;
      (1=1)
=====
      1
```

Comparison Operators

Description

The comparison operators compare the operand on the left and on the right, and they return 1 or 0. The operands of comparison operations must have the same data type. Therefore, implicit type casting by the system or implicit type casting by the user is required.

The following table shows the comparison operators supported by CUBRID and their return values.

Comparison Operators Supported by CUBRID

Comparison Operator	Description	Predicate	Return Value
=	A general equal sign. It compares whether the values of the left and right operands are the same. Returns NULL if one or more operand is NULL .	1=>2 1=NULL	0 NULL
<=>	A NULL -safe equal sign. It compares whether the values of the left and right operands are the same including NULL . Returns 1 if both operands are NULL .	1<=>2 1<=> NULL	0 0 NULL
<>, !=	The value of left operand is not equal to that of right operand. If any operand value is NULL , NULL is returned.	1<>2	1
>	The value of left operand is greater than that of right operand. If any operand value is NULL , NULL is returned.	1>2	0
<	The value of left operand is less than that of right operand. If any operand value is NULL , NULL is returned.	1<2	1
>=	The value of left operand is greater than or equal to that of right operand. If any operand value is NULL , NULL is returned.	1>=2	0
<=	The value of left operand is less than or equal to that of right operand. If any operand value is NULL , NULL is returned.	1<=2	1
IS <i>boolean_value</i>	Compares whether the value of the left operand is the same as boolean value of the right. The boolean value may be TRUE , FALSE (or NULL).	1 IS FALSE	0
IS NOT <i>boolean_value</i>	Compares whether the value of the left operand is the same as boolean value of the right. The boolean value may be TRUE , FALSE (or NULL).	1 IS NOT FALSE	1

Syntax 1

```
expression comparison_operator expression
```

```
expression :
```

- bit string
- character string
- numeric value
- date-time value
- collection value
- **NULL**

```
comparison_operator :
```

```
=
| <=>
| <>
| !=
| >
| <
| >=
| <=
```

Syntax 2

```
expression IS [NOT] boolean_value
```

```
expression :
```

- bit string
- character string
- numeric value
- date-time value

- *collection value*
- **NULL**

```
boolean_value :
< UNKNOWN | NULL>
| TRUE
| FALSE
```

- *expression* : Declares an expression to be compared.
- *bit string* : A Boolean operation can be performed on bit strings, and all comparison operators can be used for comparison between bit strings. If you compare two expressions with different lengths, 0s are padded at the end of the shorter one.
- *character string* : When compared by a comparison operator, two character strings must have the same character sets. The comparison is determined by the collation sequence of the character code set. If you compare two character strings with different lengths, blanks are padded at the end of the shorter one before comparison so that they have the same length.
- *numeric value* : The Boolean operator can be performed for all numeric values and any types of comparison operator can be used. When two different numeric types are compared, the system implicitly performs type casting. For example, when an **INTEGER** value is compared with a **DECIMAL** value, the system first casts **INTEGER** to **DECIMAL** before it performs comparison. When you compare a **FLOAT** value, you must specify the range instead of an exact value because the processing of **FLOAT** is dependent on the system.
- *date-time value* : If two date-time values with the same type are compared, the order is determined in time order. That is, when comparing two date-time values, the earlier date is considered to be smaller than the later date. You cannot compare date-time values with different type by using a comparison operator; therefore, you must explicitly convert it. However, comparison operation can be performed between **DATE**, **TIMESTAMP**, and **DATETIME** because they are implicitly converted.
- *collection value* : When comparing two sequences each element of the two sequences is compared in the order that is specified at the time of sequence creation. Comparison between sets or multisets is overloaded by an appropriate operator. You can perform comparison operations on sets, multisets, lists or sequence sets by using a containment operator explained later in this chapter. For details, see [Containment Operators](#).
- **NULL** : The **NULL** value is not included in the value range of any data type. Therefore, comparison between **NULL** values is only allowed to determine if the given value is **NULL** or not. An implicit type cast does not take place when a **NULL** value is assigned to a different data type. For example, when an attribute of **INTEGER** type has a **NULL** and is compared with a floating point type, the **NULL** value is not coerced to **FLOAT** before comparison is made. A comparison operation on the **NULL** value does not return a result.

Example

```
EVALUATE (1 <> 0); -- 1 is displayed because it is TRUE.
EVALUATE (1 != 0); -- 1 is displayed because it is TRUE.
EVALUATE (0.01 = '0.01'); -- An error occurs because a numeric data type is compared with
a character string type.
EVALUATE (1 = NULL); -- NULL is displayed.
EVALUATE (1 <=> NULL); -- 0 is displayed because it is FALSE.
EVALUATE (1.000 = 1); -- 1 is displayed because it is TRUE.
EVALUATE ('cubrid' = 'CUBRID'); -- 0 is displayed because it is case sensitive.
EVALUATE ('cubrid' = 'cubrid'); -- 1 is displayed because it is TRUE.
EVALUATE (SYSTIMESTAMP = CAST(SYSDATETIME AS TIMESTAMP)); -- 1 is displayed after casting
the type explicitly and then performing comparison operator.
EVALUATE (SYSTIMESTAMP = SYSDATETIME)); 0 is displayed after casting the type implicitly
and then performing comparison operator.
EVALUATE (SYSTIMESTAMP <> NULL); -- NULL is returned without performing comparison
operator.
EVALUATE (SYSTIMESTAMP IS NOT NULL); -- 1 is returned because it is not NULL.
```

Arithmetic Operators

Arithmetic Operators

Description

For arithmetic operators, there are binary operators for addition, subtraction, multiplication, or division, and unary operators to represent whether the number is positive or negative. The unary operators to represent the numbers' positive/negative status have higher priority over the binary operators.

Arithmetic Operators Supported by CUBRID

Arithmetic Operator	Description	Operator	Return Value
+	Addition	1+2	3
-	Subtraction	1-2	-1
*	Multiplication	1*2	2
/	Division. Returns quotient.	1/2.0	0.500000000
DIV	Division. Returns quotient.	1 DIV 2	0
%, MOD	Division. Returns quotient. An operator must be an integer type, and it always returns integer. If an operand is real number, the MOD function can be used.	1 % 2 1 MOD 2	1

Syntax

```
expression mathematical_operator expression
```

```
expression :
```

- bit string
- character string
- numeric value
- date-time value
- collection value
- **NULL**

```
mathematical operator :
```

- set arithmetic operator
- arithmetic_operator

```
arithmetic_operator :
```

- +
- -
- *
- /, DIV
- %, MOD

```
set arithmetic operator :
```

- **UNION** (Union)
- **DIFFERENCE** (Difference)
- **INTERSECT** | **INTERSECTION** (Intersection)

- *expression* : Declares the mathematical operation to be calculated.
- *mathematical_operator* : A operator that performs an operation the arithmetic and the set operators are applicable.
- *set_arithmetic_operator* : A set arithmetic operator that performs operations such as union, difference and intersection on collection type operands.
- *arithmetic_operator* : An operator to perform the four fundamental arithmetic operations.

Arithmetic Operations and Type Casting of Numeric Data Types**Description**

All numeric data types can be used for arithmetic operations. The result type of the operation differs depending on the data types of the operands and the type of the operation. The following table shows the result data types of addition/subtraction/multiplication for each operand type.

Result Data Type by Operand Type

	INT	NUMERIC	FLOAT	DOUBLE	MONETARY
INT	INT (BIGINT)	NUMERIC	FLOAT	DOUBLE	MONETARY
NUMERIC	NUMERIC	NUMERIC (<i>p</i> and <i>s</i> are also converted)	DOUBLE	DOUBLE	MONETARY

FLOAT	FLOAT	DOUBLE	FLOAT	DOUBLE	MONETARY
DOUBLE	DOUBLE	DOUBLE	DOUBLE	DOUBLE	MONETARY
MONETARY	MONETARY	MONETARY	MONETARY	MONETARY	MONETARY

Note that the result type of the operation does not change if all operands are of the same data type but type casting occurs exceptionally in division operations. An error occurs when a denominator, i.e. a divisor, is 0.

If one of the operands is a **MONETARY** type, all operation results are cast to **MONETARY** type because a **MONETARY** type uses the same operation methods as the **DOUBLE** type.

The following table shows the total number of digits (*p*) and the number of digits after the decimal point (*s*) of the operation results when all operands are of the **NUMERIC** type.

Result of NUMERIC Type Operation

Operation	Maximum Precision	Maximum Scale
$N(p1, s1) + N(p2, s2)$	$\max(p1-s1, p2-s2) + \max(s1, s2) + 1$	$\max(s1, s2)$
$N(p1, s1) - N(p2, s2)$	$\max(p1-s1, p2-s2) + \max(s1, s2)$	$\max(s1, s2)$
$N(p1, s1) * N(p2, s2)$	$p1+p2+1$	$s1+s2$
$N(p1, s1) / N(p2, s2)$	Let $Pt = p1 + \max(s1, s2) + s2 - s1$ when $s2 > 0$ and $Pt = p1$ in other cases; $St = s1$ when $s1 > s2$ and $s2$ in other cases; the number of decimal places is $\min(9-St, 38-Pt) + St$ when $St < 9$ and St in other cases.	

Example

```
--int * int
SELECT 123*123;
      123*123
=====
      15129

-- int * int returns overflow error
SELECT (1234567890123*1234567890123);

ERROR: Data overflow on data type bigint.

-- int * numeric returns numeric type
SELECT (1234567890123*CAST(1234567890123 AS NUMERIC(15,2)));
      (1234567890123* cast(1234567890123 as numeric(15,2)))
=====
      1524157875322755800955129.00

-- int * float returns float type
SELECT (1234567890123*CAST(1234567890123 AS FLOAT));
      (1234567890123* cast(1234567890123 as float))
=====
                        1.524158e+024

-- int * double returns double type
SELECT (1234567890123*CAST(1234567890123 AS DOUBLE));
      (1234567890123* cast(1234567890123 as double))
=====
                        1.524157875322756e+024

-- numeric * numeric returns numeric type
SELECT (CAST(1234567890123 AS NUMERIC(15,2))*CAST(1234567890123 AS NUMERIC(15,2)));
      ( cast(1234567890123 as numeric(15,2))* cast(1234567890123 as numeric(15,2)))
=====
      1524157875322755800955129.0000

-- numeric * float returns double type
SELECT (CAST(1234567890123 AS NUMERIC(15,2))*CAST(1234567890123 AS FLOAT));
      ( cast(1234567890123 as numeric(15,2))* cast(1234567890123 as float))
=====
                        1.524157954716582e+024
```



```

-- numeric * double returns double type
SELECT (CAST(1234567890123 AS NUMERIC(15,2))*CAST(1234567890123 AS DOUBLE));
( cast(1234567890123 as numeric(15,2))* cast(1234567890123 as double))
=====
1.524157875322756e+024

-- float * float returns float type
SELECT (CAST(1234567890123 AS FLOAT)*CAST(1234567890123 AS FLOAT));
( cast(1234567890123 as float)* cast(1234567890123 as float))
=====
1.524158e+024

-- float * double returns float type
SELECT (CAST(1234567890123 AS FLOAT)*CAST(1234567890123 AS DOUBLE));
( cast(1234567890123 as float)* cast(1234567890123 as double))
=====
1.524157954716582e+024

-- double * double returns float type
SELECT (CAST(1234567890123 AS DOUBLE)*CAST(1234567890123 AS DOUBLE));
( cast(1234567890123 as double)* cast(1234567890123 as double))
=====
1.524157875322756e+024

-- int / int returns int type without type conversion or rounding
SELECT 100100/100000;
100100/100000
=====
1

-- int / int returns int type without type conversion or rounding
SELECT 100100/200200;
100100/200200
=====
0

-- int / zero returns error
SELECT 100100/(100100-100100);
ERROR: Attempt to divide by zero.

```

Arithmetic Operations and Type Casting of DATE/TIME Data Types

Description

If all operands are date/time type, only a subtraction operation is allowed and its return value is **BIGINT**. Note that the unit of the operation differs depending on the types of the operands. Both addition and subtraction operations are allowed in case of date/time and integer types. In this case, operation units and return values are date/time data type.

The following table shows operations allowed for each operand type, and their result types.

Allowable Operation and Result Data Type by Operand Type

	TIME (in seconds)	DATE (in day)	TIMESTAMP (in seconds)	DATETIME (in milliseconds)	INT
TIME	A subtraction is allowed. BIGINT	X	X	X	An addition and a subtraction are allowed. TIME
DATE	X	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	An addition and a subtraction are allowed. DATE
TIMESTAMP	X	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	An addition and a subtraction are allowed. TIMESTAMP

DATETIME	X	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	A subtraction is allowed. BIGINT	An addition and a subtraction are allowed. DATETIME
INT	An addition and a subtraction are allowed. TIME	An addition and a subtraction are allowed. DATE	An addition and a subtraction are allowed. TIMESTAMP	An addition and a subtraction are allowed. DATETIME	All operations are allowed.

Remark

If any of the date/time arguments contains **NULL**, **NULL** is returned.

Example

```
-- initial systimestamp value
SELECT SYSDATETIME;
      SYSDATETIME
=====
      07:09:52.115 PM 01/14/2010

-- time type + 10(seconds) returns time type
SELECT (CAST (SYSDATETIME AS TIME) + 10);
      ( cast( SYS DATETIME  as time)+10)
=====
      07:10:02 PM

-- date type + 10 (days) returns date type
SELECT (CAST (SYSDATETIME AS DATE) + 10);
      ( cast( SYS DATETIME  as date)+10)
=====
      01/24/2010

-- timestamp type + 10(seconds) returns timestamp type
SELECT (CAST (SYSDATETIME AS TIMESTAMP) + 10);
      ( cast( SYS DATETIME  as timestamp)+10)
=====
      07:10:02 PM 01/14/2010

-- systimestamp type + 10(milliseconds) returns systimestamp type
SELECT (SYSDATETIME + 10);
      ( SYS DATETIME +10)
=====
      07:09:52.125 PM 01/14/2010

SELECT DATETIME '09/01/2009 03:30:30.001 pm'- TIMESTAMP '08/31/2009 03:30:30 pm';
      datetime '09/01/2009 03:30:30.001 pm'-timestamp '08/31/2009 03:30:30 pm'
=====
      86400001

SELECT TIMESTAMP '09/01/2009 03:30:30 pm'- TIMESTAMP '08/31/2009 03:30:30 pm';
      timestamp '09/01/2009 03:30:30 pm'-timestamp '08/31/2009 03:30:30 pm'
=====
      86400
```

Set Operators

Set Arithmetic Operators

Set Arithmetic Operators

To evaluate set operations such as union, difference or intersection for **SET**, **MULTISET** or **LIST (SEQUENCE)** types, you can use +, - or * operators respectively.

The following table shows a summary of how to use these operators.

Result Data Type by Operand Type

	SET	MULTISET	LIST (=SEQUENCE)
SET	+, -, *: SET	+, -, *: MULTISET	+, -, *: MULTISET
MULTISET	+, -, *: MULTISET	+, -, *: MULTISET	+, -, *: MULTISET
LIST (=SEQUENCE)	+: MULTISET -: MULTISET *: MULTISET	+: MULTISET -: MULTISET *: MULTISET	+: LIST -: MULTISET *: MULTISET

Syntax

```
value_expression set_arithmetic_operator value_expression
```

value expression :

- collection value
- **NULL**

set_arithmetic_operator :

- **+** (union)
- **-** (difference)
- ***** (intersection)

Example

```
SELECT ((CAST ({3,3,3,2,2,1} AS SET))+ (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as set))+ ( cast({4, 3, 3, 2} as multiset)))
=====
{1, 2, 2, 3, 3, 3, 4}

SELECT ((CAST ({3,3,3,2,2,1} AS MULTISET))+ (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as multiset))+ ( cast({4, 3, 3, 2} as multiset)))
=====
{1, 2, 2, 2, 3, 3, 3, 3, 3, 4}

SELECT ((CAST ({3,3,3,2,2,1} AS LIST))+ (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as sequence))+ ( cast({4, 3, 3, 2} as multiset)))
=====
{1, 2, 2, 2, 3, 3, 3, 3, 3, 4}

SELECT ((CAST ({3,3,3,2,2,1} AS SET))- (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as set))- ( cast({4, 3, 3, 2} as multiset)))
=====
{1}

SELECT ((CAST ({3,3,3,2,2,1} AS MULTISET))- (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as multiset))- ( cast({4, 3, 3, 2} as multiset)))
=====
{1, 2, 3}

SELECT ((CAST ({3,3,3,2,2,1} AS LIST))- (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as sequence))- ( cast({4, 3, 3, 2} as multiset)))
=====
{1, 2, 3}

SELECT ((CAST ({3,3,3,2,2,1} AS SET))* (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as set))* ( cast({4, 3, 3, 2} as multiset)))
=====
{2, 3}

SELECT ((CAST ({3,3,3,2,2,1} AS MULTISET))* (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as multiset))* ( cast({4, 3, 3, 2} as multiset)))
=====
{2, 3, 3}

SELECT ((CAST ({3,3,3,2,2,1} AS LIST))* (CAST ({4,3,3,2} AS MULTISET)));
(( cast({3, 3, 3, 2, 2, 1} as sequence))* ( cast({4, 3, 3, 2} as multiset)))
=====
{2, 3, 3}
```

Assigning Collection Value to Variable

For a collection value to be assigned to a variable, the outer query must return a single row as the result. The following example shows how to assign a collection value to a variable. The outer query must return only a single row as follows:

```
SELECT SET (SELECT name
FROM people
WHERE ssn in {'1234', '5678'})
TO : "names"
FROM TABLE people;
```

Statement Set Operators

Description

Statement set operators are used to get union, difference or intersection on the result of more than one query statement specified as an operand. Note that the data types of the data to be retrieved from the target tables of the two query statements must be identical or implicitly castable.

The following table shows statement set operators supported by CUBRID and their examples.

Statement Set Operators Supported by CUBRID

Statement Set Operator	Description	Note
UNION	Union Duplicates are not allowed.	Outputs all instance results containing duplicates with UNION ALL
DIFFERENCE	Difference Duplicates are not allowed.	Same as the EXCEPT operator Outputs all instance results containing duplicates with DIFFERENCE ALL
INTERSECTION	Intersection Duplicates are not allowed.	Same as the INTERSECTION operator Outputs all instance results containing duplicates with INTERSECTION ALL

Syntax

```
query term statement set operator[qualifier] query term
[ {statement set operator[qualifier] query term} ];
```

```
query term :
• query_specification
• subquery
```

```
qualifier :
• DISTINCT or DISTINCTROW (A returned instance is a distinct value.)
• UNIQUE (A returned instance is a unique value.)
• ALL (All instances are returned. Duplicates are allowed.)
```

```
statement set operator :
• UNION (union)
• DIFFERENCE (difference)
• INTERSECTION | INTERSECT (intersection)
```

Example

```
CREATE TABLE nojoin_tbl_1 (ID INT, Name VARCHAR(32));

INSERT INTO nojoin_tbl_1 VALUES (1, 'Kim');
INSERT INTO nojoin_tbl_1 VALUES (2, 'Moy');
INSERT INTO nojoin_tbl_1 VALUES (3, 'Jonas');
INSERT INTO nojoin_tbl_1 VALUES (4, 'Smith');
INSERT INTO nojoin_tbl_1 VALUES (5, 'Kim');
INSERT INTO nojoin_tbl_1 VALUES (6, 'Smith');
INSERT INTO nojoin_tbl_1 VALUES (7, 'Brown');

CREATE TABLE nojoin_tbl_2 (id INT, Name VARCHAR(32));
```

```

INSERT INTO nojoin tbl 2 VALUES (5,'Kim');
INSERT INTO nojoin tbl 2 VALUES (6,'Smith');
INSERT INTO nojoin_tbl_2 VALUES (7,'Brown');
INSERT INTO nojoin_tbl_2 VALUES (8,'Lin');
INSERT INTO nojoin tbl 2 VALUES (9,'Edwin');
INSERT INTO nojoin tbl 2 VALUES (10,'Edwin');

--Using UNION to get only distinct rows
SELECT id, name FROM nojoin_tbl_1
UNION
SELECT id,name FROM nojoin tbl 2;

      id  name
=====
      1  'Kim'
      2  'Moy'
      3  'Jonas'
      4  'Smith'
      5  'Kim'
      6  'Smith'
      7  'Brown'
      8  'Lin'
      9  'Edwin'
     10  'Edwin'

--Using UNION ALL not eliminating duplicate selected rows
SELECT id, name FROM nojoin_tbl_1
UNION ALL
SELECT id,name FROM nojoin tbl 2;

      id  name
=====
      1  'Kim'
      2  'Moy'
      3  'Jonas'
      4  'Smith'
      5  'Kim'
      6  'Smith'
      7  'Brown'
      5  'Kim'
      6  'Smith'
      7  'Brown'
      8  'Lin'
      9  'Edwin'
     10  'Edwin'

--Using DIFFERENCE to get only rows returned by the first query but not by the second
SELECT id, name FROM nojoin tbl 1
DIFFERENCE
SELECT id,name FROM nojoin_tbl_2;

      id  name
=====
      1  'Kim'
      2  'Moy'
      3  'Jonas'
      4  'Smith'

--Using INTERSECTION to get only those rows returned by both queries
SELECT id, name FROM nojoin tbl 1
INTERSECT
SELECT id,name FROM nojoin_tbl_2;

      id  name
=====
      5  'Kim'
      6  'Smith'
      7  'Brown'

```

Containment Operators

Containment Operators

Description

Containment operators are used to check the containment relationship by performing comparison operation on operands of the set data type. Set data types or subqueries can be specified as operands. The operation returns TRUE or FALSE if there is a containment relationship between the two operands of identical/different/subset/proper subset.

The description and return values about the containment operators supported by CUBRID are as follows:

Containment Operators Supported by CUBRID

Containment Operator	Description	Predicates	Return Value
A SETEQ B	A = B Elements in A and B are same each other.	{1,2} SETEQ {1,2,2}	0
A SETNEQ B	A ≠ B Elements in A and B are not same each other.	{1,2} SETNEQ {1,2,3}	1
A SUPERSET B	A ⊃ B B is a proper subset of A.	{1,2} SUPERSET {1,2,3}	0
A SUBSET B	A ⊂ B A is a proper subset of B.	{1,2} SUBSET {1,2,3}	1
A SUPERSETEQ B	A ⊇ B B is a subset of A.	{1,2} SUPERSETEQ {1,2,3}	0
A SUBSETEQ B	A ⊆ B A is a subset of B.	{1,2} SUBSETEQ {1,2,3}	1

The following table shows than possibility of operation by operand and type conversion if a containment operator is used.

Possibility of Operation by Operand

	SET	MULTISET	LIST(=SEQUENCE)
SET	Operation possible	Operation possible	Operation possible
MULTISET	Operation possible	Operation possible	Operation possible (LIST is converted into MULTISET)
LIST(=SEQUENCE)	Operation possible	Operation possible (LIST is converted into MULTISET)	Some operation possible (SETEQ, SETNEQ) Error occurs for the rest of operators.

Syntax

```
collection operand containment operator collection operand
```

collection operand:

- set
- multiset
- sequence (or list)
- subquery
- **NULL**

containment_operator:

- SETEQ
- SETNEQ
- SUPERSET
- SUBSET
- SUPERSETEQ
- SUBSETEQ

- *collection_operand* : This expression that can be specified as an operand is a single SET-valued attribute, an arithmetic expression containing a SET operator or a SET value enclosed in braces. If the type is not specified, the SET value enclosed in braces is treated as a **LIST** type by default.

Subqueries can be specified as operands. If a column which is not a **SET** type is queried, a SET data type keyword is required for the subquery (e.g. **SET(subquery)**). The column retrieved by a subquery must return a single set so that it can be compared with the set of the other operands.

If the element type is an object, the OIDs, not its contents, are compared. For example, two objects with different OIDs are considered to be different even though they have the same attribute values.

- **NULL** : Any of operands to be compared is **NULL**, **NULL** is returned.

Example

```
--empty set is a subset of any set
EVALUATE ({} SUBSETEQ (CAST ({3,1,2} AS SET)));
      Result
=====
          1

--operation between set type and null returns null
EVALUATE ((CAST ({3,1,2} AS SET)) SUBSETEQ NULL);
      Result
=====
          NULL

--{1,2,3} seteq {1,2,3} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SETEQ (CAST ({1,2,3,3} AS SET)));
      Result
=====
          1

--{1,2,3} seteq {1,2,3,3} returns false
EVALUATE ((CAST ({3,1,2} AS SET)) SETEQ (CAST ({1,2,3,3} AS MULTISSET)));
      Result
=====
          0

--{1,2,3} setneq {1,2,3,3} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SETNEQ (CAST ({1,2,3,3} AS MULTISSET)));
      Result
=====
          1

--{1,2,3} subseteq {1,2,3,4} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SUBSETEQ (CAST ({1,2,4,4,3} AS SET)));
      Result
=====
          1

--{1,2,3} subseteq {1,2,3,4,4} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SUBSETEQ (CAST ({1,2,4,4,3} AS MULTISSET)));
      Result
=====
          1

--{1,2,3} subseteq {1,2,4,4,3} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SUBSETEQ (CAST ({1,2,4,4,3} AS LIST)));
      Result
=====
          0

--{1,2,3} subseteq {1,2,3,4,4} returns true
EVALUATE ((CAST ({3,1,2} AS SET)) SUBSETEQ (CAST ({1,2,3,4,4} AS LIST)));
      Result
```

```

=====
1
--{3,1,2} seteq {3,1,2} returns true
EVALUATE ((CAST ({3,1,2} AS LIST)) SETEQ (CAST ({3,1,2} AS LIST)));
Result
=====
1
--error occurs because LIST subseteq LIST is not supported
EVALUATE ((CAST ({3,1,2} AS LIST)) SUBSETEQ (CAST ({3,1,2} AS LIST)));
Result
=====
error

```

SETEQ Operator

Description

The **SETEQ** operator returns **TRUE** if the first operand is the same as the second one. It can perform comparison operator for all collection data type.

Syntax

```
collection_operand SETEQ collection_operand
```

Example

```

--creating a table with SET type address column and LIST type zip code column

CREATE TABLE contain_tbl (id int primary key, name char(10), address SET varchar(20),
zip code LIST int);
INSERT INTO contain_tbl VALUES(1, 'Kim', {'country', 'state'},{1, 2, 3});
INSERT INTO contain_tbl VALUES(2, 'Moy', {'country', 'state'},{3, 2, 1});
INSERT INTO contain_tbl VALUES(3, 'Jones', {'country', 'state', 'city'},{1,2,3,4});
INSERT INTO contain_tbl VALUES(4, 'Smith', {'country', 'state', 'city',
'street'},{1,2,3,4});
INSERT INTO contain_tbl VALUES(5, 'Kim', {'country', 'state', 'city', 'street'},{1,2,3,4});
INSERT INTO contain_tbl VALUES(6, 'Smith', {'country', 'state', 'city',
'street'},{1,2,3,5});
INSERT INTO contain_tbl VALUES(7, 'Brown', {'country', 'state', 'city', 'street'},{});

--selecting rows when two collection_operands are same in the WEHRE clause
SELECT id, name, address, zip code FROM contain_tbl WHERE address SETEQ {'country','state',
'city'};

=====
id name address zip code
=====
3 'Jones ' {'city', 'country', 'state'} {1, 2, 3, 4}

1 row selected.

--selecting rows when two collection operands are same in the WEHRE clause
SELECT id, name, address, zip_code FROM contain_tbl WHERE zip_code SETEQ {1,2,3};

=====
id name address zip code
=====
1 'Kim ' {'country', 'state'} {1, 2, 3}

1 rows selected.

```

SETNEQ Operator

Description

The **SETNEQ** operator returns **TRUE(1)** if a first operand is different from a second operand. A comparable operation can be performed for all collection data types.

Syntax

```
collection_operand SETNEQ collection_operand
```


Example

```
--selecting rows when two collection operands are not same in the WEHRE clause
SELECT id, name, address, zip code FROM contain_tbl WHERE address SETNEQ
{'country','state','city'};
=====
      id name          address                                zip_code
=====
      1 'Kim           ' {'country','state'} {1, 2, 3}
      2 'Moy           ' {'country','state'} {3, 2, 1}
      4 'Smith         ' {'city','country','state','street'} {1, 2, 3, 4}
      5 'Kim           ' {'city','country','state','street'} {1, 2, 3, 4}
      6 'Smith         ' {'city','country','state','street'} {1, 2, 3, 5}
      7 'Brown        ' {'city','country','state','street'} {}

6 rows selected.

--selecting rows when two collection operands are not same in the WEHRE clause
SELECT id, name, address, zip_code FROM contain_tbl WHERE zip_code SETNEQ {1,2,3};
=====
      id name          address                                zip_code
=====
      2 'Moy           ' {'country','state'} {3, 2, 1}
      3 'Jones         ' {'city','country','state'} {1, 2, 3, 4}
      4 'Smith         ' {'city','country','state','street'} {1, 2, 3, 4}
      5 'Kim           ' {'city','country','state','street'} {1, 2, 3, 4}
      6 'Smith         ' {'city','country','state','street'} {1, 2, 3, 5}
      7 'Brown        ' {'city','country','state','street'} {}
```

SUPERSET Operator

Description

The **SUPERSET** operator returns **TRUE(1)** when a second operand is a proper subset of a first operand; that is, the first one is larger than the second one. If two operands are identical, **FALSE(0)** is returned. Note that **SUPERSET** is not supported if all operands are **LIST** type.

Syntax

```
collection_operand SUPERSET collection_operand
```

Example

```
--selecting rows when the first operand is a superset of the second operand and they are
not same
SELECT id, name, address, zip_code FROM contain_tbl WHERE address SUPERSET
{'country','state','city'};
=====
      id name          address                                zip code
=====
      4 'Smith         ' {'city','country','state','street'} {1, 2, 3, 4}
      5 'Kim           ' {'city','country','state','street'} {1, 2, 3, 4}
      6 'Smith         ' {'city','country','state','street'} {1, 2, 3, 5}
      7 'Brown        ' {'city','country','state','street'} {}

--SUPERSET operator cannot be used for comparison between LIST and LIST type values
SELECT id, name, address, zip code FROM contain_tbl WHERE zip code SUPERSET {1,2,3};

ERROR: ' superset ' operator is not defined on types sequence and sequence.

--Comparing operands with a SUPERSET operator after casting LIST type as SET type
SELECT id, name, address, zip code FROM contain_tbl WHERE zip code SUPERSET (CAST ({1,2,3}
AS SET));
=====
      id name          address                                zip_code
=====
      3 'Jones         ' {'city','country','state'} {1, 2, 3, 4}
      4 'Smith         ' {'city','country','state','street'} {1, 2, 3, 4}
      5 'Kim           ' {'city','country','state','street'} {1, 2, 3, 4}
      6 'Smith         ' {'city','country','state','street'} {1, 2, 3, 5}
```

SUPERSETEQ Operator

Description

The **SUPERSETEQ** operator returns **TRUE(1)** when a second operand is a subset of a first operand; that is, the first one is identical to or larger than the second one. Note that **SUPERSETEQ** is not supported if an operand is **LIST** type.

Syntax

```
collection_operand SUPERSETEQ collection_operand
```

Example

```
--selecting rows when the first operand is a superset of the second operand
SELECT id, name, address, zip code FROM contain tbl WHERE address SUPERSETEQ
{'country','state','city'};

=====
id  name      address                                     zip_code
=====
3   'Jones    {'city', 'country', 'state'} {1, 2, 3, 4}
4   'Smith    {'city', 'country', 'state', 'street'} {1, 2, 3, 4}
5   'Kim      {'city', 'country', 'state', 'street'} {1, 2, 3, 4}
6   'Smith    {'city', 'country', 'state', 'street'} {1, 2, 3, 5}
7   'Brown    {'city', 'country', 'state', 'street'} {}

--SUPERSETEQ operator cannot be used for comparison between LIST and LIST type values
SELECT id, name, address, zip code FROM contain tbl WHERE zip code SUPERSETEQ {1,2,3};

ERROR: ' superseteq ' operator is not defined on types sequence and sequence.

--Comparing operands with a SUPERSETEQ operator after casting LIST type as SET type
SELECT id, name, address, zip code FROM contain tbl WHERE zip code SUPERSETEQ (CAST
({1,2,3} AS SET));

=====
id  name      address                                     zip_code
=====
1   'Kim      {'country', 'state'} {1, 2, 3}
3   'Jones    {'city', 'country', 'state'} {1, 2, 3, 4}
4   'Smith    {'city', 'country', 'state', 'street'} {1, 2, 3, 4}
5   'Kim      {'city', 'country', 'state', 'street'} {1, 2, 3, 4}
6   'Smith    {'city', 'country', 'state', 'street'} {1, 2, 3, 5}
```

SUBSET Operator

Description

The **SUBSET** operator returns **TRUE(1)** if the second operand contains all elements of the first operand. If the first and the second collection have the same elements, **FALSE(0)** is returned. Note that both operands are the **LIST** type, the **SUBSET** operation is not supported.

Syntax

```
collection_operand SUBSET collection_operand
```

Example

```
--selecting rows when the first operand is a subset of the second operand and they are not
same
SELECT id, name, address, zip code FROM contain tbl WHERE address SUBSET
{'country','state','city'};

=====
id  name      address                                     zip_code
=====
1   'Kim      {'country', 'state'} {1, 2, 3}
2   'Moy      {'country', 'state'} {3, 2, 1}

--SUBSET operator cannot be used for comparison between LIST and LIST type values
SELECT id, name, address, zip_code FROM contain_tbl WHERE zip_code SUBSET {1,2,3};

ERROR: ' subset ' operator is not defined on types sequence and sequence.
```

```
--Comparing operands with a SUBSET operator after casting LIST type as SET type
SELECT id, name, address, zip code FROM contain tbl WHERE zip code SUBSET (CAST ({1,2,3}
AS SET));
```

id	name	address	zip_code
7	'Brown'	{'city', 'country', 'state', 'street'}	{}

SUBSETEQ Operator

Description

The **SUBSETEQ** operator returns **TRUE(1)** when a first operand is a subset of a second operand; that is, the second one is identical to or larger than the first one. Note that **SUBSETEQ** is not supported if an operand is **LIST** type.

Syntax

```
collection_operand SUBSETEQ collection_operand
```

Example

```
--selecting rows when the first operand is a subset of the second operand
SELECT id, name, address, zip code FROM contain tbl WHERE address SUBSETEQ
{'country','state','city'};
```

id	name	address	zip_code
1	'Kim'	{'country', 'state'}	{1, 2, 3}
2	'Moy'	{'country', 'state'}	{3, 2, 1}
3	'Jones'	{'city', 'country', 'state'}	{1, 2, 3, 4}

```
--SUBSETEQ operator cannot be used for comparison between LIST and LIST type values
SELECT id, name, address, zip code FROM contain tbl WHERE zip code SUBSETEQ {1,2,3};

ERROR: ' subseteq ' operator is not defined on types sequence and sequence.

--Comparing operands with a SUBSETEQ operator after casting LIST type as SET type
SELECT id, name, address, zip_code FROM contain_tbl WHERE zip_code SUBSETEQ (CAST ({1,2,3}
AS SET));
```

id	name	address	zip code
1	'Kim'	{'country', 'state'}	{1, 2, 3}
7	'Brown'	{'city', 'country', 'state', 'street'}	{}

BIT Functions and Operators

Bitwise Operator

A **Bitwise** operator performs operations in bits, and can be used in arithmetic operations. An integer type is specified as the operand and the **BIT** type cannot be specified. An integer of **BIGINT** type (64-bit integer) is returned as the result of the operation. If one or more operand is **NULL**, **NULL** is returned.

The following table shows the bitwise operators supported by CUBRID.

The bitwise operators supported by CUBRID

Bitwise operator	Description	Expression	Return Value
&	Performs AND operation in bits and returns a BIGINT integer.	17 & 3	1
	Performs OR operation in bits and returns a BIGINT integer.	17 3	19
^	Performs XOR operation in bits and returns a BIGINT integer.	17 ^ 3	18
~	A unary operator. It performs complementary operation that reverses (INVERT) the bit order of the operand and returns a BIGINT integer.	~17	-18

<<	Performs the operation of moving bits of the left operand as far to the left as the value of the right operand, and returns a BIGINT integer.	17 << 3	136
>>	Performs the operation of moving bits of the left operand as far to the right as the value of the right operand, and returns a BIGINT integer.	17 >> 3	2

BIT_AND Function

Description

An aggregate function. It performs **AND** operations in bits on every bit of *expr*. The return value is a **BIGINT** type. If there is no row that satisfies the expression, **NULL** is returned.

Syntax

BIT_AND (*expr*)

- expr*: An expression of integer type

Example

```
CREATE TABLE bit_tbl(id int);
INSERT INTO bit_tbl VALUES (1), (2), (3), (4), (5);
SELECT 1&3&5, BIT_AND(id) FROM bit_tbl WHERE id in(1,3,5);
=====
1                               bit and(id)
1                               1
```

BIT_OR Function

Description

An aggregate function. It performs **OR** operations in bits on every bit of *expr*. The return value is a **BIGINT** type. If there is no row that satisfies the expression, **NULL** is returned.

Syntax

BIT_OR (*expr*)

- expr*: An expression of integer type

Example

```
SELECT 1|3|5, BIT_OR(id) FROM bit_tbl WHERE id in(1,3,5);
=====
1|3|5                               bit_or(id)
7                                   7
```

BIT_XOR Function

Description

An aggregate function. It performs **XOR** operations in bits on every bit of *expr*. The return value is a **BIGINT** type. If there is no row that satisfies the expression, **NULL** is returned.

Syntax

BIT_XOR (*expr*)

- expr*: An expression of integer type

Example

```
SELECT 1^2^3, BIT_XOR(id) FROM bit_tbl WHERE id in(1,3,5);
```

1^3^5	bit xor(id)
7	7

BIT_COUNT Function

Description

The **BIT_COUNT** function returns the number of bits of *expr* that have been set to 1; it is not an aggregate function. The return value is a **BIGINT** type.

Syntax

```
BIT_COUNT (expr)
```

- *expr* : An expression of integer type

Example

```
SELECT BIT_COUNT(id) FROM bit_tbl WHERE id in(1,3,5);
  bit_count(id)
=====
             1
             2
             2
```

String Functions and Operators

Concatenation Operator

Description

A concatenation operator gets a character string or bit string data type as an operand and returns a concatenated string. The plus sign (+) and double pipe symbol (||) are provided as concatenation operators for character string data. If **NULL** is specified as an operand, a **NULL** value is returned.

If **pipes_as_concat** that is a parameter related to SQL statement is set to **no** (default value: yes), a double pipe (||) symbol is interpreted as an **OR** operator. If **plus_as_concat** is set to no (default value: yes), a plus (+) symbol is interpreted as a plus (+) operator. In such case, It is recommended to concatenate strings or bit strings, by using the **CONCAT** function.

Syntax

```
concat_operand1 + concat_operand1
concat_operand2 || concat_operand2
concat_operand1 :
• bit string
• NULL

concat_operand2 :
• bit string
• character string
• NULL
```

- *concat_operand1* : Left string after concatenation. String or bit string can be specified.
- *concat_operand2* : Right string after concatenation. String or bit string can be specified.

Example

```
SELECT 'CUBRID' || ',' + '2008';
'CUBRID','+2008'
=====
'CUBRID,2008'

SELECT 'cubrid' || ',' || B'0010' || B'0000' || B'0000' || B'1000';;
'cubrid','|B'0010'|B'0000'|B'0000'|B'1000'
```

```

=====
'cubrid,2008'

SELECT ((EXTRACT(YEAR FROM SYS_TIMESTAMP)) || (EXTRACT(MONTH FROM SYS_TIMESTAMP)));
(( extract(year from SYS_TIMESTAMP ) || ( extract(month from SYS_TIMESTAMP )))
=====
'200812'

SELECT 'CUBRID' || ',' + NULL;
'CUBRID' || ',' + null
=====
NULL

```

ASCII Function

Description

The **ASCII** function returns the ASCII code of the most left character in numeric value. If an input string is **NULL**, **NULL** is returned.

This function supports single-byte character sets only. If a numeric value is entered, it is converted into character string and then the ASCII code of the most left character is returned.

Syntax

```
ASCII( str )
```

- *str*: An input string

Example

```

SELECT ASCII('5');
53
SELECT ASCII('ab');
97

```

BIN Function

Description

The **BIN** function converts a **BIGINT** type number into binary string. If an input string is **NULL**, **NULL** is returned.

Syntax

```
BIN( n )
```

- *n*: A **BIGINT** type number

Example

```

SELECT BIN(12);
'1100'

```

BIT_LENGTH Function

Description

The **BIT_LENGTH** function returns the length (bits) of a character string or bit string as an integer value. The return value of the **BIT_LENGTH** function may depend on the character set, because for the character string, the number of bytes taken up by a single character is different depending on the character set of the data input environment (e.g., EUC-KR: 2*8 bits). For details about character sets supported by CUBRID, see [Definition and Characteristics](#).

Syntax

```
BIT_LENGTH( string )
```

```
string :
```

- bit string
- character string
- **NULL**

- *string* : Specifies the character string or bit string whose number of bits is to be calculated. If this value is **NULL**, **NULL** is returned.

Example

```

SELECT BIT LENGTH('');
      bit length('')
=====
              0

SELECT BIT LENGTH('CUBRID');
      bit length('CUBRID')
=====
              48

SELECT BIT_LENGTH('큐브리드');
      bit_length('큐브리드')
=====
              64

SELECT BIT LENGTH(B'010101010');
      bit length(B'010101010')
=====
              9

CREATE TABLE bit_length_tbl (char 1 CHAR, char 2 CHAR(5), varchar 1 VARCHAR, bit var 1 BIT
VARYING);
INSERT INTO bit_length_tbl VALUES(' ', ' ', ' ', B''); --Length of empty string
INSERT INTO bit_length_tbl VALUES('a', 'a', 'a', B'010101010'); --English character
INSERT INTO bit_length_tbl VALUES(NULL, '큐', '큐', B'010101010'); --Korean character and
NULL
INSERT INTO bit_length_tbl VALUES(' ', ' 큐', ' 큐', B'010101010'); --Korean character and
space

SELECT BIT LENGTH(char 1), BIT LENGTH(char 2), BIT LENGTH(varchar 1), BIT LENGTH(bit var 1)
FROM bit_length_tbl;

bit_length(char_1)  bit_length(char_2)      bit_length(varchar_1)  bit_length(bit_var_1)
=====
8                  40                      0                      0
8                  40                      8                      9
NULL              40                      16                     9
8                  40                      24                     9

```

CHAR_LENGTH/CHARACTER_LENGTH/LENGTHB/LENGTH Functions

Description

CHAR_LENGTH, **LENGTHB**, and **LENGTH** are used interchangeably.

They return the length of a character string (byte) as an integer. The return value may be different depending on the character set (e.g., EUC-KR: 2 bites).

For details about the character sets supported by CUBRID, see [Definition and Characteristics](#).

Syntax

```

CHAR_LENGTH( string )
CHARACTER_LENGTH( string )
LENGTHB( string )
LENGTH( string )

```

string :

- character string
- **NULL**

- *string* : Specifies the character string whose number of characters is to be calculated. If the character string is NULL, NULL is returned.

Remark

- The length of each space character that is included in a character string is one byte.
- For multi-byte strings, the length of a single character is calculated as 2 or 3 bytes depending on the character set of the data input environment.
- The length of empty quotes (") to represent a space character is 0. Note that in a CHAR(*n*) type, the length of a space character is *n*, and it is specified as 1 if *n* is omitted.

Example

```
--character set is euc-kr for Korean characters
SELECT LENGTH(' ');
char length(' ')
=====
0

SELECT LENGTH('CUBRID');
char length('CUBRID')
=====
6

SELECT LENGTH('큐브리드');
char length('큐브리드')
=====
8

CREATE TABLE length_tbl (char_1 CHAR, char_2 CHAR(5), varchar_1 VARCHAR, varchar_2
VARCHAR);
INSERT INTO length_tbl VALUES(' ', ' ', ' ', ' '); --Length of empty string
INSERT INTO length_tbl VALUES('a', 'a', 'a', 'a'); --English character
INSERT INTO length_tbl VALUES(NULL, '큐', '큐', '큐'); --Korean character and NULL
INSERT INTO length_tbl VALUES(' ', ' 큐', ' 큐', ' 큐'); --Korean character and space

SELECT LENGTH(char 1), LENGTH(char 2), LENGTH(varchar 1), LENGTH(varchar 2) FROM
length_tbl;

char_length(char_1) char_length(char_2) char_length(varchar_1) char_length(varchar_2)
=====
1 5 0 0
1 5 1 1
NULL 5 2 2
1 5 3 3
```

CHR Function

Description

The CHR function returns a character that corresponds to the return value of the expression specified as an argument. It returns 0 if it exceeds range of character code.

Syntax

```
CHR( number_operand )
```

- *number_operand* : Specifies an expression that returns a numeric value.

Example

```
SELECT CHR(68) || CHR(68-2);
chr(68) || chr(68-2)
=====
'DB'
```


CONCAT Function

Description

The **CONCAT** function has at least one argument specified for it and returns a string as a result of concatenating all argument values. The number of parameters that can be specified is unlimited. Automatic type casting takes place if a non-string type is specified as the argument. If any of the arguments is specified as **NULL**, **NULL** is returned.

If you want to insert separators between strings specified as arguments for concatenation, use the [CONCAT_WS Function](#).

Syntax

```
CONCAT( string1, string2 [,string3 [, ... [, stringN]...]])
```

string :

- character string
- **NULL**

Example

```
SELECT CONCAT('CUBRID', '2008', 'R3.0');
      concat('CUBRID', '2008', 'R3.0')
=====
'CUBRID2008R3.0'

--it returns null when null is specified for one of parameters
SELECT CONCAT('CUBRID', '2008', 'R3.0', NULL);
      concat('CUBRID', '2008', 'R3.0', null)
=====
      NULL

--it converts number types and then returns concatenated strings
SELECT CONCAT(2008, 3.0);
      concat(2008, 3.0)
=====
      '20083.0'
```

CONCAT_WS Function

Description

The **CONCAT_WS** function has at least two arguments specified for it. The function uses the first argument value as the separator and returns the result.

```
CONCAT_WS( string1, string2 [,string3 [, ... [, stringN]...]])
```

string :

- character string
- **NULL**

Example

```
SELECT CONCAT_WS(' ', 'CUBRID', '2008', 'R3.0');
      concat_ws(' ', 'CUBRID', '2008', 'R3.0')
=====
      'CUBRID 2008 R3.0'

--it returns strings even if null is specified for one of parameters
SELECT CONCAT_WS(' ', 'CUBRID', '2008', NULL, 'R3.0');
      concat_ws(' ', 'CUBRID', '2008', null, 'R3.0')
=====
      'CUBRID 2008 R3.0'

--it converts number types and then returns concatenated strings with separator
SELECT CONCAT_WS(' ', 2008, 3.0);
      concat_ws(' ', 2008, 3.0)
=====
```

'2008 3.0'

ELT Function

Description

If N is 1, the **ELT** function returns *string1* and if N is 2, it returns *string2*. The return value is a **VARCHAR** type. You can add conditional expressions as needed.

The maximum length of the character string is 33,554,432 and if this length is exceeded, **NULL** will be returned.

If N is 0 or a negative number, an empty string will be returned. If N is greater than the number of this input character string, **NULL** will be returned as it is out of range. If N is a type that cannot be converted to an integer, an error will be returned.

Syntax

```
ELT( $N$ , string1, string2, ... )
```

```
string :
• character string
• NULL
```

Example

```
SELECT ELT(3,'string1','string2','string3');
      elt(3, 'string1', 'string2', 'string3')
=====
      'string3'

SELECT ELT('3','1/1/1','23:00:00','2001-03-04');
      elt('3', '1/1/1', '23:00:00', '2001-03-04')
=====
      '2001-03-04'

SELECT ELT(-1, 'string1','string2','string3');
      elt(-1, 'string1','string2','string3')
=====
      NULL

SELECT ELT(4,'string1','string2','string3');
      elt(4, 'string1', 'string2', 'string3')
=====
      NULL

SELECT ELT(3.2,'string1','string2','string3');
      elt(3.2, 'string1', 'string2', 'string3')
=====
      'string3'

SELECT ELT('a','string1','string2','string3');
ERROR: Cannot coerce value of domain "character" to domain "bigint".
```

FIELD Function

Description

The **FIELD** function returns the location index value (position) of a string of *string1*, *string2*. The function returns 0 if it does not have a parameter value which is the same as *search_string*. It returns 0 if *search_string* is **NULL** because it cannot perform the comparison operation with the other arguments.

If all arguments specified for **FIELD** function are of string type, string comparison operation is performed: if all of them are of number type, numeric comparison operation is performed. If the type of one argument is different from that of another, a comparison operation is performed by casting each argument to the type of the first argument. If type casting fails during the comparison operation with each argument, the function considers the result of the comparison operation as **FALSE** and resumes the other operations.

Syntax

```
FIELD( search_string, string1 [,string2 [, ... [, stringN]...]])
```

string :

- character string
- NULL

Example

```
SELECT FIELD('abc', 'a', 'ab', 'abc', 'abcd', 'abcde');
      field('abc', 'a', 'ab', 'abc', 'abcd', 'abcde')
=====
                                     3

--it returns 0 when no same string is found in the list
SELECT FIELD('abc', 'a', 'ab', NULL);
      field('abc', 'a', 'ab', null)
=====
                                     0

--it returns 0 when null is specified in the first parameter
SELECT FIELD(NULL, 'a', 'ab', NULL);
      field(null, 'a', 'ab', null)
=====
                                     0

SELECT FIELD('123', 1, 12, 123.0, 1234, 12345);
      field('123', 1, 12, 123.0, 1234, 12345)
=====
                                     0

SELECT FIELD(123, 1, 12, '123.0', 1234, 12345);
      field(123, 1, 12, '123.0', 1234, 12345)
=====
                                     3
```

FIND_IN_SET Function

Description

The **FIND_IN_SET** function looks for the string *str* in the string list *strlist* and returns a position of *str* if it exists. A string list is a string composed of substrings separated by a comma (,).

If *str* is not in *strlist* or *strlist* is an empty string, 0 is returned. If either argument is **NULL**, **NULL** is returned. This function does not work properly if *str* contains a comma (,).

Syntax

```
FIND_IN_SET(str, strlist)
```

- *str* : A string to be searched
- *strlist* : A group of strings separated by a comma

Example

```
SELECT FIND IN SET('b', 'a,b,c,d');
2
```

INSERT Function

Description

The **INSERT** function inserts a partial character string as long as the length from the specific location of the input character string. The return value is a **VARCHAR** type.

The maximum length of the character string is 33,554,432 and if this length is exceeded, **NULL** will be returned.

Syntax

```
INSERT( str, pos, len, string )
```

- *str* : Input character string
- *pos* : *str* location. Starts from 1. If *pos* is smaller than 1 or greater than the length of *string* + 1, the *string* will not be inserted and the *str* will be returned instead.
- *len* : Length of *string* to insert *pos* of *str*. If *len* exceeds the length of the partial character string, insert as many values as *string* in the *pos* of the *str*. If *len* is a negative number, *str* will be the end of the character string.
- *string* : Partial character string to insert to *str*

Example

```
SELECT INSERT('cubrid',2,2,'dbsql');
insert('cubrid', 2, 2, 'dbsql')
=====
'cdbsqlrid'

SELECT INSERT('cubrid',0,3,'db');
insert('cubrid', 0, 3, 'db')
=====
'cubrid'

SELECT INSERT('cubrid',-3,3,'db');
insert('cubrid', -3, 3, 'db')
=====
'cubrid'

SELECT INSERT('cubrid',3,100,'db');
insert('cubrid', 3, 100, 'db')
=====
'cudb'

SELECT INSERT('cubrid',7,100,'db');
insert('cubrid', 7, 100, 'db')
=====
'cubriddb'

SELECT INSERT('cubrid',3,-1,'db');
insert('cubrid', 3, -1, 'db')
=====
'cudb'
```

INSTR Function

Description

The **INSTR** function, similarly to the **POSITION**, returns the position of a *substring* within *string*; the position. For the **INSTR** function, you can specify the starting position of the search for *substring* to make it possible to search for duplicate *substring*.

Note that the function calculates the starting position and the length of the character string in bytes, not in characters. For a multi-byte character set, the number of bite representing onc character is different, so the return value may not be the same.

Syntax

```
INSTR( string , substring [, position] )
```

```
string , substring :
• character string
• NULL
position :
• INT
• NULL
```

- *string* : Specifies the input character string.
- *substring* : Specifies the character string whose position is to be returned.

- *position* : Optional. Represents the position of a *string* where the search begins. If omitted, the default value 1 is applied. The first position of the *string* is specified as 1. If the value is negative, the system counts backward from the end of the *string*.

Example

```
--character set is euc-kr for Korean characters
--it returns position of the first 'b'
SELECT INSTR ('12345abcdeabcde','b');
instr('12345abcdeabcde', 'b', 1)
=====
7

-- it returns position of the first 'ㄴ' on double byte charset
SELECT INSTR ('12345가나다라마가나다라마', 'ㄴ' );
instr('12345가나다라마가나다라마', 'ㄴ', 1)
=====
8

-- it returns position of the second 'ㄴ' on double byte charset
SELECT INSTR ('12345가나다라마가나다라마', 'ㄴ', 16 );
instr('12345가나다라마가나다라마', 'ㄴ', 16)
=====
18

--it returns position of the 'b' searching from the 8th position
SELECT INSTR ('12345abcdeabcde','b', 8);
instr('12345abcdeabcde', 'b', 8)
=====
12

--it returns position of the 'b' searching backwardly from the end
SELECT INSTR ('12345abcdeabcde','b', -1);
instr('12345abcdeabcde', 'b', -1)
=====
12

--it returns position of the 'b' searching backwardly from a specified position
SELECT INSTR ('12345abcdeabcde','b', -8);
instr('12345abcdeabcde', 'b', -8)
=====
7
```

LCASE/LOWER Functions

Description

The **LCASE** and **LOWER** functions are used interchangeably. They convert uppercase characters included in string to lowercase characters. Note that these functions may not work properly in character sets which are not supported by CUBRID. For details on the character sets supported by CUBRID, see [Definition and Characteristics](#).

Syntax

```
LCASE ( string )
LOWER ( string )

string :
• character string
• NULL
```

- *string* : Specifies the string in which uppercase characters are to be converted to lowercase. If the value is **NULL**, **NULL** is returned.

Example

```
SELECT LOWER('');
lower('')
```

```

=====
' '
SELECT LOWER(NULL);
  lower(null)
=====
NULL
SELECT LOWER('Cubrid');
  lower('Cubrid')
=====
'cubrid'

```

LEFT Function

Description

The **LEFT** function returns a length number of characters from the leftmost of *string*. If any of the arguments is **NULL**, **NULL** is returned. If a value greater than the *length* of the *string* or a negative number is specified for a length, the entire string is returned.

To extract a length number of characters from the rightmost of the string, use the [RIGHT Function](#).

Syntax

```

LEFT( string , length )

string :
• character string
• NULL

length :
• INT
• NULL

```

Example

```

SELECT LEFT('CUBRID', 3);
  left('CUBRID', 3)
=====
' CUB'

SELECT LEFT('CUBRID', 10);
  left('CUBRID', 10)
=====
'CUBRID'

```

LOCATE Function

Description

The **LOCATE** function returns the location index value of a *substring* within a character string. The third argument *position* can be omitted. If this argument is specified, the function searches for *substring* from the given position and returns the location index value of the first occurrence. If the *substring* cannot be found within the string, 0 is returned.

The **LOCATE** function is working like the [POSITION Function](#), but you cannot use **LOCATE** for bit strings.

Syntax

```

LOCATE ( substring, string [, position] )

string :
• character string
• NULL

```

Example

```

--it returns 1 when substring is empty space

```

```

SELECT LOCATE ('', '12345abcdeabcde');
locate('', '12345abcdeabcde')
=====
1

--it returns position of the first 'abc'
SELECT LOCATE ('abc', '12345abcdeabcde');
locate('abc', '12345abcdeabcde')
=====
6

--it returns position of the second 'abc'
SELECT LOCATE ('abc', '12345abcdeabcde', 8);
locate('abc', '12345abcdeabcde', 8)
=====
11

--it returns 0 when no substring found in the string
SELECT LOCATE ('ABC', '12345abcdeabcde');
locate('ABC', '12345abcdeabcde')
=====
0

```

LPAD Function

Description

The **LPAD** function pads the left side of a string with a specific set of characters.

Syntax

```
LPAD( char1, n, [, char2 ] )
```

char1 :

- *character string*
- *string valued column*
- **NULL**

n :

- *integer*
- **NULL**

char2 :

- *character string*
- **NULL**

- *char1* : Specifies the string to pad characters to. If *n* is smaller than the length of *char1*, padding is not performed, and *char1* is truncated to length *n* and then returned. A single character is processed as 2 or 3 bytes in multi-byte character set environment. If *char1* is truncated up to the first byte representing a character according to a value of *n*, the last byte is removed and a space character (1 byte) is added to the left because the last character cannot be represented normally. If the value is **NULL**, **NULL** is returned.
- *n* : Specifies the total length of *char1* in bytes. Note that the number and the length of the character strings may be different in multi-byte character set environment. If the value is **NULL**, **NULL** is returned.
- *char2* : Specifies the string to pad to the left until the length of *char1* reaches *n*. If it is not specified, empty characters (' ') are used as a default. If the value is **NULL**, **NULL** is returned.

Example

```

--character set is euc-kr for Korean characters

--it returns only 3 characters if not enough length is specified
SELECT LPAD ('CUBRID', 3, '?');
lpad('CUBRID', 3, '?')
=====
'CUB'

--on multi-byte charset, it returns the first character only with a left padded space
SELECT LPAD ('큐브리드', 3, '?');
lpad('큐브리드', 3, '?')

```

```

=====
' 큐'

--padding spaces on the left till char_length is 10
SELECT LPAD ('CUBRID', 10);
lpad('CUBRID', 10)
=====
' CUBRID'

--padding specific characters on the left till char_length is 10
SELECT LPAD ('CUBRID', 10, '?');
lpad('CUBRID', 10, '?')
=====
'????CUBRID'

--padding specific characters on the left till char_length is 10
SELECT LPAD ('큐브리드', 10, '?');
lpad('큐브리드', 10, '?')
=====
'??큐브리드'

--padding 4 characters on the left
SELECT LPAD ('큐브리드', LENGTH('큐브리드')+4, '?');
lpad('큐브리드', char_length('큐브리드')+4, '?')
=====
'????큐브리드'

```

LTRIM Function

Description

The **LTRIM** function removes all specified characters from the left-hand side of a string.

Syntax

```
LTRIM( string [, trim_string])
```

string :

- *character string*
- *string valued column*
- **NULL**

trim_string :

- *character string*
- **NULL**

- *string* : Enters a string or string-type column to trim. If this value is **NULL**, **NULL** is returned.
- *trim_string* : You can specify a specific string to be removed in the left side of *string*. If it is not specified, empty characters (' ') is automatically specified so that the empty characters in the left side are removed.

Example

```

--trimming spaces on the left
SELECT LTRIM ('    Olympic    ');
ltrim('    Olympic    ')
=====
'Olympic    '

--If NULL is specified, it returns NULL
SELECT LTRIM ('iiiiiiOlympiciiiiii', NULL);
ltrim('iiiiiiOlympiciiiiii', null)
=====
NULL

-- trimming specific strings on the left
SELECT LTRIM ('iiiiiiOlympiciiiiii', 'i');
ltrim('iiiiiiOlympiciiiiii', 'i')

```



```
=====
'Olympiciiii'
```

MID Function

Description

The **MID** function extracts a string with the length of *substring_length* from a *position* within the *string* and then returns it. If a negative number is specified as a *position* value, the *position* is calculated in a reverse direction from the end of the *string*. **substring_length** cannot be omitted. If a negative value is specified, the function considers this as 0 and returns an empty string.

The **MID** function is working like the [SUBSTR Function](#), but there are differences in that it cannot be used for bit strings, that the *substring_length* argument must be specified, and that it returns an empty string if a negative number is specified for *substring_length*.

Syntax

```
string :
• character string
• NULL

position :
• integer
• NULL

substring_length :
• integer
• NULL
```

- *string* : Specifies an input character string. If this value is **NULL**, **NULL** is returned.
- *position* : Specifies the starting position from which the string is to be extracted. The position of the first character is 1. It is considered to be 1 even if it is specified as 0. If the input value is **NULL**, **NULL** is returned.
- *substring_length* : Specifies the length of the string to be extracted. If 0 or a negative number is specified, an empty string is returned; if **NULL** is specified, **NULL** is returned.

Example

```
CREATE TABLE mid_tbl(a VARCHAR);
INSERT INTO mid_tbl VALUES('12345abcdeabcde');

--it returns empty string when substring length is 0
SELECT MID(a, 6, 0), SUBSTR(a, 6, 0), SUBSTRING(a, 6, 0) FROM mid_tbl;
mid(a, 6, 0)          substr(a, 6, 0)          substring(a from 6 for 0)
=====
''                   ''                   ''

--it returns 4-length substrings counting from the 6th position
SELECT MID(a, 6, 4), SUBSTR(a, 6, 4), SUBSTRING(a, 6, 4) FROM mid_tbl;
mid(a, 6, 4)          substr(a, 6, 4)          substring(a from 6 for 4)
=====
'abcd'               'abcd'                 'abcd'

--it returns a empty string when substring length < 0
SELECT MID(a, 6, -4), SUBSTR(a, 6, -4), SUBSTRING(a, 6, -4) FROM mid_tbl;
mid(a, 6, -4)          substr(a, 6, -4)          substring(a from 6 for -4)
=====
''                   NULL                   'abcdeabcde'

--it returns 4-length substrings at 6th position counting backward from the end
SELECT MID(a, -6, 4), SUBSTR(a, -6, 4), SUBSTRING(a, -6, 4) FROM mid_tbl;
mid(a, -6, 4)          substr(a, -6, 4)          substring(a from -6 for 4)
=====
'eabc'               'eabc'                 '1234'
```

OCTET_LENGTH Function

Description

The **OCTET_LENGTH** function returns the length (byte) of a character string or bit string as an integer. Therefore, it returns 1 (byte) if the length of the bit string is 8 bits, but 2 (bytes) if the length is 9 bits.

Syntax

```
OCTET_LENGTH ( string )
```

string :

- *bit string*
- *character string*
- **NULL**

- *string* : Specifies the character or bit string whose length is to be returned in bytes. If the value is **NULL**, **NULL** is returned.

Example

```
--character set is euc-kr for Korean characters

SELECT OCTET_LENGTH('');
   octet length('')
=====
                0

SELECT OCTET_LENGTH('CUBRID');
   octet length('CUBRID')
=====
                6

SELECT OCTET_LENGTH('큐브리드');
   octet_length('큐브리드')
=====
                8

SELECT OCTET_LENGTH(B'010101010');
   octet length(B'010101010')
=====
                2

CREATE TABLE octet_length_tbl (char 1 CHAR, char 2 CHAR(5), varchar 1 VARCHAR, bit var 1
BIT VARYING);
INSERT INTO octet_length_tbl VALUES(' ', ' ', ' ', B''); --Length of empty string
INSERT INTO octet_length_tbl VALUES('a', 'a', 'a', B'010101010'); --English character
INSERT INTO octet_length_tbl VALUES(NULL, '큐', '큐', B'010101010'); --Korean character and
NULL
INSERT INTO octet_length_tbl VALUES(' ', ' 큐', ' 큐', B'010101010'); --Korean character
and space

SELECT OCTET_LENGTH(char 1), OCTET_LENGTH(char 2), OCTET_LENGTH(varchar 1),
OCTET_LENGTH(bit var 1) FROM octet_length_tbl;
octet length(char 1) octet length(char 2) octet length(varchar 1) octet length(bit var 1)
=====
1                5                0                0
1                5                1                2
NULL            5                2                2
1                5                3                2
```

POSITION Function

Description

The **POSITION** function returns the position of a character string corresponding to *substring* within a character string corresponding to *string*. Note that it returns the position in bytes, not in characters. Therefore, the return values may differ because the number of bytes representing a single character is different in multi-byte character sets.

An expression that returns a character string or a bit string can be specified as an argument of this function. The return value is an integer greater than or equal to 0. This function returns the position value in bytes for a character string, and in bits for a bit string.

The **POSITION** function is occasionally used in combination with other functions. For example, if you want to extract a certain string from another string, you can use the result of the **POSITION** function as an input to the **SUBSTRING** function.

Syntax

```
POSITION ( substring IN string )
```

```
substring :
• bit string
• character string
• NULL
```

- *substring* : Specifies the character string whose position is to be returned. If the value is an empty character, 1 is returned. If the value is **NULL**, **NULL** is returned.

Example

```
--character set is euc-kr for Korean characters

--it returns 1 when substring is empty space
SELECT POSITION ('' IN '12345abcdeabcde');
      position('' in '12345abcdeabcde')
=====
                          1

--it returns position of the first 'b'
SELECT POSITION ('b' IN '12345abcdeabcde');
      position('b' in '12345abcdeabcde')
=====
                          7

-- it returns position of the first '나' on double byte charset
SELECT POSITION ('나' IN '12345가나다라마가나다라마');
      position('나' in '12345가나다라마가나다라마')
=====
                          8

--it returns 0 when no substring found in the string
SELECT POSITION ('f' IN '12345abcdeabcde');
      position('f' in '12345abcdeabcde')
=====
                          0

SELECT POSITION (B'1' IN B'000011110000');
      position(B'1' in B'000011110000')
=====
                          5
```

REPEAT Function

Description

The **REPEAT** function returns the character string with a length equal to the number of repeated input character strings. The return value is a **VARCHAR** type. The maximum length of the character string is 33,554,432 and if it this length is exceeded, **NULL** will be returned. If one of the parameters is **NULL**, **NULL** will be returned.

Syntax

```
REPEAT( string, count )
```

- *string* : Character string
- *count* : Repeat count. If you enter 0 or a negative number, an empty string will be returned and if you enter a non-numeric data type, an error will be returned.

Example

```
SELECT REPEAT('cubrid',3);
      repeat('cubrid', 3)
=====
      'cubridcubridcubrid'

SELECT REPEAT('cubrid',32000000);
      repeat('cubrid', 32000000)
=====
      NULL

SELECT REPEAT('cubrid',-1);
      repeat('cubrid', -1)
=====
      ''

SELECT REPEAT('cubrid','a');
ERROR: Cannot coerce value of domain "character" to domain "integer".
```

REPLACE Function

Description

The **REPLACE** function searches for a character string, *search_string*, within a given character string, *string*, and replaces it with a character string, *replacement_string*. If the string to be replaced, *replacement_string* is omitted, all *search_strings* retrieved from *string* are removed. If **NULL** is specified as an argument, **NULL** is returned.

Syntax

```
REPLACE( string, search string [, replacement string ] )
```

```
string :
• character string
• NULL
```

```
search string :
• character string
• NULL
```

```
replacement string :
• character string
• NULL
```

- *string* : Specifies the original string. If the value is **NULL**, **NULL** is returned.
- *search_string* : Specifies the string to be searched. If the value is **NULL**, **NULL** is returned.
- *replacement_string* : Specifies the string to replace the *search_string*. If this value is omitted, *string* is returned with the *search_string* removed. If the value is **NULL**, **NULL** is returned.

Example

```
--it returns NULL when an argument is specified with NULL value
SELECT REPLACE('12345abcdeabcde','abcde',NULL);
replace('12345abcdeabcde', 'abcde', null)
=====
      NULL

--not only the first substring but all substrings into 'ABCDE' are replaced
SELECT REPLACE('12345abcdeabcde','abcde','ABCDE');
replace('12345abcdeabcde', 'abcde', 'ABCDE')
=====
      '12345ABCDEABCDE'

--it removes all of substrings when replace string is omitted
SELECT REPLACE('12345abcdeabcde','abcde');
replace('12345abcdeabcde', 'abcde')
=====
      '12345'
```

REVERSE Function

Description

The **REVERSE** function returns *string* converted in the reverse order.

Syntax

```
REVERSE( string )
```

string :

- character string
- **NULL**

- *string* : Specifies an input character string. If the value is an empty string, empty value is returned. If the value is **NULL**, **NULL** is returned.

Example

```
SELECT REVERSE('CUBRID');
reverse('CUBRID')
=====
      'DIRBUC'
```

RIGHT Function

Description

The **RIGHT** function returns a *length* number of characters from the rightmost of a *string*. If any of the arguments is **NULL**, **NULL** is returned. If a value greater than the length of the *string* or a negative number is specified for a *length*, the entire string is returned.

To extract a length number of characters from the leftmost of the string, use the [LEFT Function](#).

Syntax

```
RIGHT( string , length )
```

string :

- character string
- **NULL**

length :

- *INT*
- **NULL**

Example

```
SELECT RIGHT('CUBRID', 3);
```

```

right('CUBRID', 3)
=====
'RID'

SELECT RIGHT ('CUBRID', 10);
right('CUBRID', 10)
=====
'CUBRID'

```

RPAD Function

Description

The **RPAD** function pads the right side of a string with a specific set of characters.

Syntax

```
RPAD( char1, n, [, char2 ] )
```

char1 :

- *character string*
- *string valued column*
- **NULL**

n :

- *integer*
- **NULL**

char2 :

- *character string*
- **NULL**

- *char1* : Specifies the string to pad characters to. If *n* is smaller than the length of *char1*, padding is not performed, and *char1* is truncated to length *n* and then returned. A single character is processed as 2 or 3 bytes in multi-byte character set environment. If *char1* is truncated up to the first byte representing a character according to a value of *n*, the last byte is removed and an empty character (1 byte) is added to the left because the last character cannot be represented normally. If the value is **NULL**, **NULL** is specified.
- *n* : Specifies the total length of *char1* in bytes. Note that the number and the length of the character strings may be different in multi-byte character set environment. If the value is **NULL**, **NULL** is specified.
- *char2* : Specifies the string to pad to the right until the length of *char1* reaches *n*. If it is not specified, empty characters (' ') are used as a default. If the value is **NULL**, **NULL** is returned.

Example

```

--character set is euc-kr for Korean characters

--it returns only 3 characters if not enough length is specified
SELECT RPAD ('CUBRID', 3, '?');
rpad('CUBRID', 3, '?')
=====
'CUB'

--on multi-byte charset, it returns the first character only with a right-padded space
SELECT RPAD ('큐브리드', 3, '?');
rpad('큐브리드', 3, '?')
=====
'큐 '

--padding spaces on the right till char length is 10
SELECT RPAD ('CUBRID', 10);
rpad('CUBRID', 10)
=====
'CUBRID   '

--padding specific characters on the right till char length is 10
SELECT RPAD ('CUBRID', 10, '?');
rpad('CUBRID', 10, '?')
=====

```

```

'CUBRID????'

--padding specific characters on the right till char_length is 10
SELECT RPAD ('큐브리드', 10, '?');
      rpad('큐브리드', 10, '?')
=====
      '큐브리드??'

--padding 4 characters on the right
SELECT RPAD ('큐브리드', LENGTH('큐브리드')+4, '?');
      rpad(' ', char_length('')+4, '?')
=====
      '큐브리드????'

```

RTRIM Function

Description

The **RTRIM** function removes specified characters from the right-hand side of a string.

Syntax

```
RTRIM( string [, trim_string])
```

string :

- character string
- string valued column
- **NULL**

trim string :

- character string
- **NULL**

- *string* : Enters a string or string-type column to trim. If this value is **NULL**, **NULL** is returned.
- *trim_string* : You can specify a specific string to be removed in the right side of *string*. If it is not specified, empty characters (' ') is automatically specified so that the empty characters in the right side are removed.

Example

```

SELECT RTRIM ('      Olympic      ');
      rtrim('      Olympic      ')
=====
      '      Olympic'

--If NULL is specified, it returns NULL
SELECT RTRIM ('iiiiOlympiciiii', NULL);
      rtrim('iiiiOlympiciiii', null)
=====
      NULL

-- trimming specific strings on the right
SELECT RTRIM ('iiiiOlympiciiii', 'i');
      rtrim('iiiiOlympiciiii', 'i')
=====
      'iiiiOlympic'

```

SPACE Function

Description

The **SPACE** function returns as many empty strings as the number specified. The return value is a **VARCHAR** type.

Syntax

```
SPACE (N)
```

- *N*: Space count. It cannot be greater than the value specified in the system parameter, **string_max_size_bytes** (default 1048576). If it exceeds the specified value, **NULL** will be returned. The maximum value is 33,554,432; if this length is exceeded, **NULL** will be returned. If you enter 0 or a negative number, an empty string will be returned; if you enter a type that can't be converted to a numeric value, an error will be returned.

Example

```

SELECT SPACE(8);
      space(8)
=====
      '      '

SELECT LENGTH(space(1048576));
      char_length( space(1048576))
=====
                                1048576

SELECT LENGTH(space(1048577));
      char_length( space(1048577))
=====
                                NULL

-- string max size bytes=33554432
SELECT LENGTH(space('33554432'));
      char_length( space('33554432'))
=====
                                33554432

SELECT SPACE('aaa');
ERROR: Cannot coerce value of domain "character" to domain "bigint".

```

STRCMP Function

Description

The **STRCMP** function compares two strings, *string1* and *string2*, and returns 0 if they are identical, 1 if *string1* is greater, or -1 if *string1* is smaller. If any of the parameters is **NULL**, **NULL** is returned.

Syntax

```

STRCMP( string1 , string2 )

string :
• character string
• NULL

```

Example

```

SELECT STRCMP('abc', 'abc');
=====
                                0
SELECT STRCMP ('acc', 'abc');
=====
                                1

--STRCMP works case-insensitively
SELECT STRCMP ('ABC','abc');
=====
                                0

```


SUBSTR Function

Description

The **SUBSTR** function extracts a character string with the length of *substring_length* from a position, *position*, within character string, *string*, and then returns it. If a negative number is specified as a *position* value, the position is calculated in a reverse direction from the end of the string. If *substring_length* is omitted, character strings between the given position, *position*, and the end of the string are extracted, and then returned.

Note that it returns the starting position and the length of character string in bytes, not in characters. Therefore, in a multi-byte character set, you must specify the parameter in consideration of the number of bytes representing a single character.

Syntax

```
SUBSTR( string, position [, substring_length])
```

string :

- character string
- bit string
- **NULL**

position :

- integer
- **NULL**

substring_length :

- integer

- *string* : Specifies the input character string. If the input value is **NULL**, **NULL** is returned.
- *position* : Specifies the position from where the string is to be extracted in bytes. Even though the position of the first character is specified as 1 or a negative number, it is considered as 1. If a value greater than the string length or **NULL** is specified, **NULL** is returned.
- *substring_length* : Specifies the length of the string to be extracted in bytes. If this argument is omitted, character strings between the given position, *position*, and the end of them are extracted. **NULL** cannot be specified as an argument value of this function. If 0 is specified, an empty string is returned; if a negative value is specified, **NULL** is returned.

Example

```
--character set is euc-kr for Korean characters

--it returns empty string when substring_length is 0
SELECT SUBSTR('12345abcdeabcde',6, 0);
  substr('12345abcdeabcde', 6, 0)
=====
  ''

--it returns 4-length substrings counting from the position
SELECT SUBSTR('12345abcdeabcde', 6, 4), SUBSTR('12345abcdeabcde', -6, 4);
  substr('12345abcdeabcde', 6, 4)  substr('12345abcdeabcde', -6, 4)
=====
  'abcd'                          'eabc'

--it returns substrings counting from the position to the end
SELECT SUBSTR('12345abcdeabcde', 6), SUBSTR('12345abcdeabcde', -6);
  substr('12345abcdeabcde', 6)  substr('12345abcdeabcde', -6)
=====
  'abcdeabcde'                  'eabcde'

-- it returns 4-length substrings counting from 16th position on double byte charset
SELECT SUBSTR ('12345가나다라마가나다라마', 16 , 4);
  substr('12345가나다라마가나다라마', 16 , 4)
=====
  '가나'
```

SUBSTRING Function

Description

The **SUBSTRING** function, operating like **SUBSTR**, extracts a character string having the length of *substring_length* from a position, *position*, within character string, *string*, and returns it.

If a negative number is specified to the *position* value, the **SUBSTRING** function calculates the position from the beginning of the string. And **SUBSTR** function calculates the position from the end of the string. If a negative number is specified to the *substring_length* value, the **SUBSTRING** function handles the argument is omitted, but the **SUBSTR** function returns **NULL**.

Syntax

```
SUBSTRING( string, position [, substring_length])
SUBSTRING( string FROM position [FOR substring_length] )
```

string :

- bit string
- character string
- **NULL**

position :

- integer
- **NULL**

substring length :

- integer

- *string* : Specifies the input character string. If the input value is **NULL**, **NULL** is returned.
- *position* : Specifies the position from where the string is to be extracted in bytes. Even though the position of the first character is specified as 1 or a negative number, it is considered as 1. If a value greater than the string length is specified, an empty string is returned. If **NULL**, **NULL** is returned.
- *substring_length* : Specifies the length of the string to be extracted in bytes. If this argument is omitted, character strings between the given position, *position*, and the end of them are extracted. **NULL** cannot be specified as an argument value of this function. If 0 is specified, an empty string is returned; if a negative value is specified, **NULL** is returned.

Example

```
SELECT SUBSTRING('12345abcdeabcde', -6 ,4), SUBSTR('12345abcdeabcde', -6 ,4);
=====
'1234'                'eabc'

SELECT SUBSTRING('12345abcdeabcde', 16), SUBSTR('12345abcdeabcde', 16);
=====
''                    NULL

SELECT SUBSTRING('12345abcdeabcde', 6, -4), SUBSTR('12345abcdeabcde', 6, -4);
=====
'abcdeabcde'        NULL
```

SUBSTRING_INDEX Function

Description

The **SUBSTRING_INDEX** function counts the separators included in the partial character string and will return the partial character string before *count*th. The return value is a **VARCHAR** type.

Syntax

```
SUBSTRING_INDEX (string, delim, count)
```

- *string* : Input character string. The maximum length is 33,554,432 and if this length is exceeded, **NULL** will be returned.

- *delim* : Delimiter. It is case-sensitive.
- *count* : Delimiter occurrence count. If you enter a positive number, it counts the character string from the left and if you enter a negative number, it counts it from the right. If it is 0, an empty string will be returned. If the type cannot be converted, an error will be returned.

Example

```

SELECT SUBSTRING_INDEX('www.cubrid.org','.', '2');
substring_index('www.cubrid.org', '.', '2')
=====
'www.cubrid'

SELECT SUBSTRING_INDEX('www.cubrid.org','.', '2.3');
substring_index('www.cubrid.org', '.', '2.3')
=====
'www.cubrid'

SELECT SUBSTRING_INDEX('www.cubrid.org', ':', '2.3');
substring_index('www.cubrid.org', ':', '2.3')
=====
'www.cubrid.org'

SELECT SUBSTRING_INDEX('www.cubrid.org', 'cubrid', 1);
substring_index('www.cubrid.org', 'cubrid', 1)
=====
'www.'

SELECT SUBSTRING_INDEX('www.cubrid.org', '.', 100);
substring_index('www.cubrid.org', '.', 100)
=====
'www.cubrid.org'

```

TRANSLATE Function

Description

The **TRANSLATE** function replaces a character into the character specified in *to_substring* if the character exists in the specified *string*. Correspondence relationship is determined based on the order of characters specified in *from_substring* and *to_substring*. Any characters in *from_substring* that do not have one on one relationship to *to_substring* are all removed. This function is working like the **REPLACE** function but the argument of *to_substring* cannot be omitted in this function.

Syntax

```
TRANSLATE( string, from_substring, to_substring )
```

string :

- *character string*
- **NULL**

from_substring :

- *character string*
- **NULL**

to_substring :

- *character string*
- **NULL**

- *string* : Specifies the original string. If the value is **NULL**, **NULL** is returned.
- *from_substring* : Specifies the string to be retrieved. If the value is **NULL**, **NULL** is returned.
- *to_substring* : Specifies the character string in the *from_substring* to be replaced. It cannot be omitted. If the value is **NULL**, **NULL** is returned.

Example

```

--it returns NULL when an argument is specified with NULL value
SELECT TRANSLATE('12345abcdeabcde', 'abcde', NULL);
translate('12345abcdeabcde', 'abcde', null)

```

```

=====
NULL

--it translates 'a','b','c','d','e' into '1', '2', '3', '4', '5' respectively
SELECT TRANSLATE('12345abcdeabcde', 'abcde', '12345');
  translate('12345abcdeabcde', 'abcde', '12345')
=====
'123451234512345'

--it translates 'a','b','c' into '1', '2', '3' respectively and removes 'd's and 'e's
SELECT TRANSLATE('12345abcdeabcde','abcde', '123');
  translate('12345abcdeabcde', 'abcde', '123')
=====
'12345123123'

--it removes 'a's,'b's,'c's,'d's, and 'e's in the string
SELECT TRANSLATE('12345abcdeabcde','abcde', '');
  translate('12345abcdeabcde', 'abcde', '')
=====
'12345'

--it only translates 'a','b','c' into '3', '4', '5' respectively
SELECT TRANSLATE('12345abcdeabcde','ABabc', '12345');
  translate('12345abcdeabcde', 'ABabc', '12345')
=====
'12345345de345de'

```

TRIM Function

Description

```
TRIM ( [ [ LEADING | TRAILING | BOTH ] [ trim_string ] FROM ] string )
```

trim_string :

- character string
- **NULL**

string :

- character string literal
- string valued column
- **NULL**

- *trim_string* : Specifies a specific string to be removed that is in front of or at the back of the target string. If it is not specified, an empty character (') is automatically specified so that spaces in front of or at the back of the target string are removed.
- *string* : Enters a string or string-type column to trim. If this value is **NULL**, **NULL** is returned.
- [**LEADING** | **TRAILING** | **BOTH**] : You can specify an option to trim a specified string that is in a certain position of the target string. If it is **LEADING**, trimming is performed in front of a character string if it is **TRAILING**, trimming is performed at the back of a character string if it is **BOTH**, trimming is performed in front and at the back of a character string. If the option is not specified, **BOTH** is specified by default.
- The character string of *trim_string* and *string* should have the same character set.

Example

```

--trimming NULL returns NULL
SELECT TRIM (NULL);
  trim(both from null)
=====
NULL

--trimming spaces on both leading and trailing parts
SELECT TRIM ('   Olympic   ');
  trim(both from '   Olympic   ')
=====
'Olympic'

--trimming specific strings on both leading and trailing parts
SELECT TRIM ('i' FROM 'iiiiOlympiciiii');
  trim(both 'i' from 'iiiiOlympiciiii')
=====

```

```

'Olympic'

--trimming specific strings on the leading part
SELECT TRIM (LEADING 'i' FROM 'iiiiOlympiciiii');
trim(leading 'i' from 'iiiiOlympiciiii')
=====
'Olympiciiii'

--trimming specific strings on the trailing part
SELECT TRIM (TRAILING 'i' FROM 'iiiiOlympiciiii');
trim(trailing 'i' from 'iiiiOlympiciiii')
=====
'iiiiOlympic'

```

UCASE/UPPER Functions

Description

The **UCASE** and **UPPER** functions convert lowercase characters that are included in a character string to uppercase characters. Note that the **UPPER** function may not work properly in character sets that are not supported by CUBRID. For details about the character sets supported by CUBRID, see [Definition and Characteristics](#).

Syntax

```

UCASE ( string )
UPPER ( string )

```

string :

- *character string*
- **NULL**

- *string* : Specifies the string in which lowercase characters are to be converted to uppercase. If the value is **NULL**, **NULL** is returned.

Example

```

SELECT UPPER('');
upper('')
=====
''

SELECT UPPER(NULL);
upper(null)
=====
NULL

SELECT UPPER('Cubrid');
upper('Cubrid')
=====
'CUBRID'

```

Numeric and Operator Functions

ABS Function

Description

The **ABS** function returns the absolute value of a given number. The data type of the return value is the same as that of the argument.

Syntax

```

ABS( number_operand )

```

- *number_operand* : An operator which returns a numeric value

Example

```
--it returns the absolute value of the argument
SELECT ABS(12.3), ABS(-12.3), ABS(-12.3000), ABS(0.0);
   abs(12.3)          abs(-12.3)          abs(-12.3000)       abs(0.0)
=====
   12.3              12.3              12.3000            .0
```

ACOS Function

Description

The **ACOS** function returns an arc cosine value of the argument. That is, it returns a value whose cosine is *x* in radian. The return value is a **DOUBLE** type. *x* must be a value between -1 and 1, inclusive. Otherwise, **NULL** is returned.

Syntax

```
ACOS ( x )
```

- *x* : An expression that returns a numeric value.

Example

```
SELECT ACOS(1), ACOS(0), ACOS(-1);
   acos(1)          acos(0)          acos(-1)
=====
  0.000000000000000e+00  1.570796326794897e+00  3.141592653589793e+00
```

ASIN Function

Description

The **ASIN** function returns an arc sine value of the argument. That is, it returns a value whose sine is *x* in radian. The return value is a **DOUBLE** type. *x* must be a value between -1 and 1, inclusive. Otherwise, **NULL** is returned.

Syntax

```
ASIN ( x )
```

- *x* : An expression that returns a numeric value.

Example

```
SELECT ASIN(1), ASIN(0), ASIN(-1);
   asin(1)          asin(0)          asin(-1)
=====
  1.570796326794897e+00  0.000000000000000e+00  -1.570796326794897e+00
```

ATAN Function

Description

The **ATAN** function returns a value whose tangent is *x* in radian. The argument *y* can be omitted. If *y* is specified, the function calculates the arc tangent value of *y/x*. The return value is a **DOUBLE** type.

Syntax

```
ATAN ( [y,] x )
```

- *x, y* : An expression that returns a numeric value.

Example

```
SELECT ATAN(1), ATAN(-1), ATAN(1,-1);
           atan(1)          atan(-1)          atan2(1, -1)
=====
```

```
7.853981633974483e-01 -7.853981633974483e-01 2.356194490192345e+00
```

ATAN2 Function

Description

The **ATAN2** function returns the arc tangent value of y/x in radian. This function is working like the [ATAN Function](#). Arguments x and y must be specified. The return value is a **DOUBLE** type.

Syntax

```
ATAN2 ( y, x )
```

- x, y : An expression that returns a numeric value.

Example

```
SELECT ATAN2 (1,1), ATAN2 (-1,-1), ATAN2 (Pi (), 0);

atan2 (1, 1)          atan2 (-1, -1)          atan2 ( pi (), 0)
=====
7.853981633974483e-01  -2.356194490192345e+00  1.570796326794897e+00
```

CEIL Function

Description

The **CEIL** function returns the smallest integer that is not less than its argument. The return value is determined based on the valid number of digits that are specified as the *number_operand* argument.

Syntax

```
CEIL( number_operand )
```

- *number_operand*: An expression that returns a numeric value.

Example

```
SELECT CEIL (34567.34567), CEIL (-34567.34567);
ceil (34567.34567)      ceil (-34567.34567)
=====
34568.00000           -34567.00000

SELECT CEIL (34567.1), CEIL (-34567.1);
ceil (34567.1)         ceil (-34567.1)
=====
34568.0               -34567.0
```

CONV Function

Description

The **CONV** function converts numbers between different number bases. This function returns a string representation of a converted number.

The minimum value is 2 and the maximum value is 36. If *to_base* (representing the base to be returned) is negative, *number* is regarded as a signed number. Otherwise, it regarded as a unsigned number.

Syntax

```
CONV (number, from_base, to_base)
```

- *number*: An input number
- *from_base*: The base of an input number
- *to_base*: The base of an returned value

Example

```
SELECT CONV('f',16,2);
'1111'
SELECT CONV('6H',20,8);
'211'
SELECT CONV(-30,10,-20);
```

COS Function

Description

The **COS** function returns a cosine value of the argument. The argument *x* must be a radian value. The return value is a **DOUBLE** type.

Syntax

```
COS ( x )
```

- *x* : An expression that returns a numeric value.

Example

```
SELECT COS(pi()/6), COS(pi()/3), COS(pi());
cos( pi()/6)          cos( pi()/3)          cos( pi())
=====
8.660254037844387e-01  5.000000000000001e-01  -1.000000000000000e+00
```

COT Function

Description

The **COT** function returns the cotangent value of the argument *x*. That is, it returns a value whose tangent is *x* in radian. The return value is a **DOUBLE** type.

Syntax

```
COT ( x )
```

- *x* : An expression that returns a numeric value.

Example

```
SELECT COT(1), COT(-1), COT(0);
cot(1)          cot(-1)          cot(0)
=====
6.420926159343306e-01  -6.420926159343306e-01  NULL
```

DEGREES Function

Description

The **DEGREES** function returns the argument *x* specified in radian converted to a degree value. The return value is a **DOUBLE** type.

Syntax

```
DEGREES ( x )
```

- *x* : An expression that returns a numeric value.

Example

```
SELECT DEGREES(pi()/6), DEGREES(pi()/3), DEGREES(pi());
degrees( pi()/6)          degrees( pi()/3)          degrees( pi())
=====
3.000000000000000e+01  5.999999999999999e+01  1.800000000000000e+02
```


DRANDOM/DRAND Functions

Description

The **DRANDOM/DRAND** function returns a random double-precision floating point value in the range of between 0.0 and 1.0. A *seed* argument that is **INTEGER** type can be specified. It rounds up real numbers and an error is returned when it exceeds the range of **INTEGER**.

The **DRAND** function performs the operation only once to produce only one random number regardless of the number of rows where the operation is output, but the **DRANDOM** function performs the operation every time the statement is repeated to produce a different random value for each row. Therefore, to output rows in a random order, you must use the **DRANDOM** function in the **ORDER BY** clause.

To obtain a random integer value, use the [RANDOM/RAND Functions](#).

Syntax

```
DRANDOM( [seed] )
DRAND ( [seed] )
```

Example

```
SELECT DRAND(), DRAND(1), DRAND(1.4);
=====
                drand()                drand(1)                drand(1.4)
=====
2.849646518006921e-001    4.163034446537495e-002    4.163034446537495e-002

SELECT * FROM rand_tbl;
=====
   id  name
=====
    1  'a'
    2  'b'
    3  'c'
    4  'd'
    5  'e'
    6  'f'
    7  'g'
    8  'h'
    9  'i'
   10  'j'

--drandom() returns random values on every row
SELECT DRAND(), DRANDOM() FROM rand_tbl;
=====
   drand()                drandom()
=====
7.638782921842098e-001    1.018707846308786e-001
7.638782921842098e-001    3.191320535905026e-001
7.638782921842098e-001    3.461714529862361e-001
7.638782921842098e-001    6.791894283883175e-001
7.638782921842098e-001    4.533829767754143e-001
7.638782921842098e-001    1.714224677266762e-001
7.638782921842098e-001    1.698049867244484e-001
7.638782921842098e-001    4.507583849604786e-002
7.638782921842098e-001    5.279091769157994e-001
7.638782921842098e-001    7.021088290047914e-001

--selecting rows in random order
SELECT * FROM rand_tbl ORDER BY DRANDOM();
=====
   id  name
=====
    6  'f'
    2  'b'
    7  'g'
    8  'h'
    1  'a'
    4  'd'
   10  'j'
    9  'i'
    5  'e'
    3  'c'
```

EXP Function

Description

The **EXP** function returns e^x (the base of natural logarithm) raised to a power.

Syntax

EXP (*x*)

- *x* : An operator which returns a numeric value

Example

```
SELECT EXP(1), EXP(0);
   exp(1)                exp(0)
=====
2.718281828459045e+000  1.000000000000000e+000

SELECT EXP(-1), EXP(2.00);
   exp(-1)                exp(2.00)
=====
3.678794411714423e-001  7.389056098930650e+000
```

FLOOR Function

Description

The **FLOOR** function returns the largest integer that is not greater than its argument. The data type of the return value is the same as that of the argument.

Syntax

FLOOR(*number_operand*)

- *number_operand* : An operator which returns a numeric value

Example

```
--it returns the largest integer less than or equal to the arguments
SELECT FLOOR(34567.34567), FLOOR(-34567.34567);
   floor(34567.34567)    floor(-34567.34567)
=====
34567.00000             -34568.00000

SELECT FLOOR(34567), FLOOR(-34567);
   floor(34567)         floor(-34567)
=====
34567                  -34567
```

FORMAT Function

Description

The **FORMAT** function displays the number *x* by using commas as thousands delimiters, so that its format becomes '#,###,###.#####' and performs rounding after the decimal point to express as many as *dec* digits after it. The return value is a string type.

Syntax

FORMAT (*x* , *dec*)

- *x* , *dec* : An expression that returns a numeric value.

Example

```
SELECT FORMAT(12000.123456,3), FORMAT(12000.123456,0);
   format(12000.123456, 3)    format(12000.123456, 0)
```

```
=====
'12,000.123'      '12,000'
```

GREATEST Function

Description

The **GREATEST** function compares more than one expression specified as parameters and returns the greatest value. If only one expression has been specified, the expression is returned because there is no expression to be compared with.

Therefore, more than one expression that are specified as parameters must be of the type that can be compared with each other. If the types of the specified parameters are identical, so are the types of the return values; if they are different, the type of the return value becomes a convertible common data type.

That is, the **GREATEST** function compares the values of column 1, column 2 and column 3 in the same row and returns the greatest value while the **MAX** function compares the values of column in all result rows and returns the greatest value.

Syntax

```
GREATEST( expression [, expression]* )
```

- *expression* : Specifies more than one expression. Their types must be comparable each other. One of the arguments is **NULL**, **NULL** is returned.

Example

The following example shows how to retrieve the number of every medals and the highest number that Korea won (demodb).

```
SELECT gold, silver , bronze, GREATEST (gold, silver, bronze) FROM participant
WHERE nation code = 'KOR';
      gold      silver      bronze  greatest(gold, silver, bronze)
=====
          9          12          9           12
          8          10          10          10
          7          15           5           15
          12           5          12           12
          12          10          11           12
```

HEX Function

Description

The **HEX** function returns a decimal string if a hexadecimal string is specified as an argument; it returns a hexadecimal string if a decimal string is specified as an argument. If a number is specified as an argument, it returns a value like **CONV(num, 10, 16)**.

Syntax

```
HEX(str)
HEX(num)
```

- *str* : A hexadecimal string
- *num* : A decimal string

Example

```
SELECT HEX('ab'), HEX(128), CONV(HEX(128), 16, 10);
hex('ab')          hex(128)          conv(hex(128), 16, 10)
=====
'6162'              '80'              '128'
```

LEAST Function

Description

The **LEAST** function compares more than one expression specified as parameters and returns the smallest value. If only one expression has been specified, the expression is returned because there is no expression to be compared with.

Therefore, more than one expression that are specified as parameters must be of the type that can be compared with each other. If the types of the specified parameters are identical, so are the types of the return values; if they are different, the type of the return value becomes a convertible common data type.

That is, the **LEAST** function compares the values of column 1, column 2 and column 3 in the same row and returns the smallest value while the **MIN** function compares the values of column in all result rows and returns the smallest value.

Syntax

```
LEAST ( expression [, expression]* )
```

- *expression* : Specifies more than one expression. Their types must be comparable each other. One of the arguments is **NULL**, **NULL** is returned.

Example

The following example shows how to retrieve the number of every medals and the lowest number that Korea won (demodb).

```
SELECT gold, silver, bronze, LEAST(gold, silver, bronze) FROM participant
WHERE nation_code = 'KOR';
```

gold	silver	bronze	least(gold, silver, bronze)
9	12	9	9
8	10	10	8
7	15	5	5
12	5	12	5
12	10	11	10

LN Function

Description

The **LN** function returns the natural log value (base = e) of an antilogarithm *x*. The return value is a **DOUBLE** type. If the antilogarithm is 0 or a negative number, an error is returned.

Syntax

```
LN ( x )
```

- *x* : An expression that returns a positive value.

Example

```
SELECT ln(1), ln(2.72);
```

ln(1)	ln(2.72)
0.0000000000000000e+00	1.000631880307906e+00

LOG2 Function

Description

The **LOG2** function returns a log value whose antilogarithm is *x* and base is 2. The return value is a **DOUBLE** type. If the antilogarithm is 0 or a negative number, an error is returned.

Syntax

```
LOG2 ( x )
```

- x : An expression that returns a positive number.

Example

```
SELECT log2(1), log2(8);
      log2(1)                log2(8)
=====
0.0000000000000000e+00    3.0000000000000000e+00
```

LOG10 Function

Description

The **LOG10** function returns the common log value of an antilogarithm x . The return value is a **DOUBLE** type. If the antilogarithm is 0 or a negative number, an error is returned.

Syntax

```
LOG10 ( x )
```

- x : An expression that returns a positive number.

Example

```
SELECT log10(1), log10(1000);
      log10(1)                log10(1000)
=====
0.0000000000000000e+00    3.0000000000000000e+00
```

MOD Function

Description

The **MOD** function returns the remainder of the first parameter m divided by the second parameter n . If n is 0, m is returned without the division operation being performed.

Note that if the dividend, the parameter m of the **MOD** function, is a negative number, the function operates differently from a typical operation (classical modulus) method.

Result of MOD

m	n	MOD(m, n)	Classical Modulus $m-n*\text{FLOOR}(m/n)$
11	4	3	3
11	-4	3	-1
-11	4	-3	1
-11	-4	-3	-3
11	0	11	Divided by 0 error

Syntax

```
MOD ( m, n )
```

- m : Represents a dividend. It is an expression that returns a numeric value.
- n : Represents a divisor. It is an expression that returns a numeric value.

Example

```
--it returns the remainder of m divided by n
SELECT MOD(11, 4), MOD(11, -4), MOD(-11, 4), MOD(-11, -4), MOD(11, 0);
      mod(11, 4)  mod(11, -4)  mod(-11, 4)  mod(-11, -4)  mod(11, 0)
=====
3                3            -3            -3            11
```

```

SELECT MOD(11.0, 4), MOD(11.000, 4), MOD(11, 4.0), MOD(11, 4.000);
=====
mod(11.0, 4)          mod(11.000, 4)          mod(11, 4.0)          mod(11, 4.000)
=====
3.0                  3.000                  3.0                  3.000
    
```

PI Function

Description

The **PI** function returns the π value of type **DOUBLE**.

Syntax

```
PI ( )
```

Example

```

SELECT PI ( ), PI ( ) / 2;
=====
pi ( )          pi ( ) / 2
=====
3.141592653589793e+00    1.570796326794897e+00
    
```

POW/POWER Functions

Description

The **POW** function returns x to the power of y . **POW** and **POWER** are used interchangeably. The return value is a **DOUBLE** type.

Syntax

```
POW ( x, y )
POWER ( x, y )
```

- x : It represents the base. It is an expression that returns a numeric value. An expression that returns a numeric value.
- y : It represents the exponent. An expression that returns a numeric value. If the base is a negative number, an integer must be specified as the exponent.

Example

```

SELECT POWER(2, 5), POWER(-2, 5), POWER(0, 0), POWER(1,0);
=====
power(2, 5)          power(-2, 5)          power(0, 0)          power(1, 0)
=====
3.2000000000000000e+01    -
3.2000000000000000e+01    1.0000000000000000e+00    1.0000000000000000e+00

--it returns an error when the negative base is powered by a non-int exponent
SELECT POWER(-2, -5.1), POWER(-2, -5.1);
=====
ERROR
    
```

RADIANS Function

Description

The **RADIANS** function returns the argument x specified in degrees converted to a radian value. The return value is a **DOUBLE** type.

Syntax

```
RADIANS ( x )
```

- x : An expression that returns a numeric value.

Example

```
SELECT RADIANS(90), RADIANS(180), RADIANS(360);
```

radians (90)	radians (180)	radians (360)
1.570796326794897e+00	3.141592653589793e+00	6.283185307179586e+00

RANDOM/RAND Functions

Description

The **RANDOM/RAND** function returns any integer value between $0^{2^{31}}$ and a *seed* argument that is **INTEGER** type can be specified. It rounds up real numbers and an error is returned when it exceeds the range of **INTEGER**.

The **RAND** function performs the operation only once to produce only one random number regardless of the number of rows where the operation is output, but the **RANDOM** function performs the operation every time the statement is repeated to produce a different random value for each row. Therefore, to output rows in a random order, you must use the **RANDOM** function.

To obtain a random real number, use the [DRANDOM/DRAND Functions](#).

Syntax

```
RANDOM( [seed] )
RAND( [seed] )
```

Example

```
SELECT RAND(), RAND(1), RAND(1.4);
      rand()      rand(1)      rand(1.4)
=====
1526981144      89400484      89400484

--creating a new table
SELECT * FROM rand_tbl;
      id  name
=====
1      'a'
2      'b'
3      'c'
4      'd'
5      'e'
6      'f'
7      'g'
8      'h'
9      'i'
10     'j'

--random() returns random values on every row
SELECT RAND(),RANDOM() FROM rand_tbl;
      rand()      random()
=====
2078876566      1753698891
2078876566      1508854032
2078876566      625052132
2078876566      279624236
2078876566      1449981446
2078876566      1360529082
2078876566      1563510619
2078876566      1598680194
2078876566      1160177096
2078876566      2075234419

--selecting rows in random order
SELECT * FROM rand_tbl ORDER BY RANDOM();
      id  name
=====
6      'f'
1      'a'
5      'e'
4      'd'
2      'b'
```

```

7 'g'
10 'j'
9 'i'
3 'c'
8 'h'

```

ROUND Function

Description

The **ROUND** function returns the specified argument, *number_operand*, rounded to the number of places after the decimal point specified by the *integer*. If the *integer* argument is a negative number, it rounds to a place before the decimal point, that is, at the integer part.

Syntax

```
ROUND( number_operand, integer )
```

- *number_operand* : An expression that returns a numeric value
- *integer* : Specifies the place to round to. If a positive integer *n* is specified, the number is represented to the *n*th place after the decimal point; if a negative integer *n* is specified, the number is rounded to the *n*th place before the decimal point.
- The return value has the same type as the *number_operand*.

Example

```

--it rounds a number to one decimal point when the second argument is omitted
SELECT ROUND(34567.34567), ROUND(-34567.34567);
   round(34567.34567, 0)   round(-34567.34567, 0)
=====
   34567.00000           -34567.00000

--it rounds a number to three decimal point
SELECT ROUND(34567.34567, 3), ROUND(-34567.34567, 3) FROM db root;
   round(34567.34567, 3)   round(-34567.34567, 3)
=====
   34567.34600           -34567.34600

--it rounds a number three digit to the left of the decimal point
SELECT ROUND(34567.34567, -3), ROUND(-34567.34567, -3);
   round(34567.34567, -3)   round(-34567.34567, -3)
=====
   35000.00000           -35000.00000

```

SIGN Function

Description

The **SIGN** function returns the sign of a given number. It returns 1 for a positive value, -1 for a negative value, and 0 for zero.

Syntax

```
SIGN(number_operand)
```

- *number_operand* : An operator which returns a numeric value

Example

```

--it returns the sign of the argument
SELECT SIGN(12.3), SIGN(-12.3), SIGN(0);
   sign(12.3)   sign(-12.3)   sign(0)
=====
           1           -1           0

```


SIN Function

Description

The **SIN** function returns a sine value of the parameter. The argument x must be a radian value. The return value is a **DOUBLE** type.

Syntax

```
SIN (  $x$  )
```

- x : An expression that returns a numeric value.

Example

```
SELECT SIN(pi()/6), SIN(pi()/3), SIN(pi());
      sin( pi()/6)          sin( pi()/3)          sin( pi())
=====
4.9999999999999999e-01    8.660254037844386e-01    1.224646799147353e-16
```

SQRT Function

Description

The **SQRT** function returns the square root of x as a **DOUBLE** type.

Syntax

```
SQRT (  $x$  )
```

- x : An expression that returns a numeric value. An error is returned if this value is a negative number.

Example

```
SELECT SQRT(4), SQRT(16.0);
      sqrt(4)          sqrt(16.0)
=====
2.0000000000000000e+00    4.0000000000000000e+00
```

TAN Function

Description

The **TAN** function returns a tangent value of the argument. The argument x must be a radian value. The return value is a **DOUBLE** type.

Syntax

```
TAN (  $x$  )
```

- x : An expression that returns a numeric value.

Example

```
SELECT TAN(pi()/6), TAN(pi()/3), TAN(pi()/4);
      tan( pi()/6)          tan( pi()/3)          tan( pi()/4)
=====
5.773502691896257e-01    1.732050807568877e+00    9.999999999999999e-01
```

TRUNC/TRUNCATE Functions

Description

The **TRUNC** and **TRUNCATE** function truncates the numbers of the specified argument x to the right of the *dec* position. If the *dec* argument is a negative number, it displays 0s to the *dec*-th position left to the decimal point. Note that the *dec* argument of the **TRUNC** function can be omitted, but that of the **TRUNCATE** function cannot be omitted.

If the *dec* argument is a negative number, it displays 0s to the *dec*-th position left to the decimal point. The number of digits of the return value to be represented follows the argument *x*.

Syntax

```
TRUNC ( x [, dec] )
TRUNCATE ( x, dec )
```

- *x* : An expression that returns a numeric value.
- *dec* : The place to be truncated is specified. If a positive integer *n* is specified, the number is represented to the *n*-th place after the decimal point; if a negative integer *n* is specified, the number is truncated to the *n*-th place before the decimal point. It truncates to the first place after the decimal point if the *dec* argument is 0 or omitted. Note that the *dec* argument cannot be omitted in the **TRUNCATE** function.

Example

```
--it returns a number truncated to 0 places
SELECT TRUNC(34567.34567), TRUNCATE(34567.34567, 0);
      trunc(34567.34567, 0)      trunc(34567.34567, 0)
=====
34567.00000                    34567.00000

--it returns a number truncated to three decimal places
SELECT TRUNC(34567.34567, 3), TRUNC(-34567.34567, 3);
      trunc(34567.34567, 3)      trunc(-34567.34567, 3)
=====
34567.34500                    -34567.34500

--it returns a number truncated to three digits left of the decimal point
SELECT TRUNC(34567.34567, -3), TRUNC(-34567.34567, -3);
      trunc(34567.34567, -3)      trunc(-34567.34567, -3)
=====
34000.00000                    -34000.00000
```

Date/Time Functions and Operators

ADDDATE/DATE_ADD Functions

Description

The **ADDDATE** function performs an addition or subtraction operation on a specific **DATE** value; **ADDDATE** and **DATE_ADD** are used interchangeably. The return value is a **DATE** or **DATETIME** type. The **DATETIME** type is returned in the following cases.

- The first argument is a **DATETIME** or **TIMESTAMP** type
- The first argument is a **DATE** type and the unit of **INTERVAL** value specified is less than the unit of day

Therefore, to return value of **DATETIME** type, you should convert the value of first argument by using the **CAST** function. Even though the date resulting from the operation exceeds the last day of the month, the function returns a valid **DATE** value considering the last date of the month.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

If the calculated value is between '0000-00-00 00:00:00' and '0001-01-01 00:00:00', a value having 0 for all arguments is returned in **DATE** or **DATETIME** type. Note that operation in JDBC program is determined by the configuration of zeroDateTimeBehavior, connection URL property (see "API Reference > JDBC API > JDBC Programming > Connection Configuration").

Syntax

```
ADDDATE (date, INTERVAL expr unit)
DATE_ADD (date, INTERVAL expr unit)
ADDDATE (date, days)
```

- *date* : It is a **DATE**, **TIMETIME**, or **TIMESTAMP** expression that represents the start date. If an invalid **DATE** value such as '2006-07-00' is specified, an error is returned.
- *expr* : It represents the interval value to be added to the start date. If a negative number is specified next to the **INTERVAL** keyword, the interval value is subtracted from the start date.
- *unit* : It represents the unit of the interval value specified in the *expr* expression. See the following table to specify the format for the interpretation of the interval value. If the value of *expr* is less than the number requested in the *unit*, it is specified from the smallest unit. For example, if it is HOUR_SECOND, three values such as 'HOURS:MINUTES:SECONDS' are required. In the case, if only two values such as "1:1" are given, it is regarded as 'MINUTES:SECONDS'.

expr value for unit

Unit Value	expr Value
MILLISECOND	MILLISECONDS
SECOND	SECONDS
MINUTE	MINUTES
HOUR	HOURS
DAY	DAYS
WEEK	WEEKS
MONTH	MONTHS
QUARTER	QUARTERS
YEAR	YEARS
SECOND_MILLISECOND	'SECONDS.MILLISECONDS'
MINUTE_MILLISECOND	'MINUTES:SECONDS.MILLISECONDS'
MINUTE_SECOND	'MINUTES:SECONDS'
HOUR_MILLISECOND	'HOURS:MINUTES:SECONDS.MILLISECONDS'
HOUR_SECOND	'HOURS:MINUTES:SECONDS'
HOUR_MINUTE	'HOURS:MINUTES'
DAY_MILLISECOND	'DAYS HOURS:MINUTES:SECONDS.MILLISECONDS'
DAY_SECOND	'DAYS HOURS:MINUTES:SECONDS'
DAY_MINUTE	'DAYS HOURS:MINUTES'
DAY_HOUR	'DAYS HOURS'
YEAR_MONTH	'YEARS-MONTHS'

Example

```

SELECT SYSDATE, ADDDATE(SYSDATE,INTERVAL 24 HOUR), ADDDATE(SYSDATE, 1);

  SYS_DATE      date_add( SYS_DATE , INTERVAL 24 HOUR)  adddate( SYS_DATE , 1)
=====
  03/30/2010  12:00:00.000 AM 03/31/2010              03/31/2010

--it subtracts days when argument < 0
SELECT SYSDATE, ADDDATE(SYSDATE,INTERVAL -24 HOUR), ADDDATE(SYSDATE, -1);

  SYS DATE      date add( SYS DATE , INTERVAL -24 HOUR)  adddate( SYS DATE , -1)
=====
  03/30/2010  12:00:00.000 AM 03/29/2010              03/29/2010

--when expr is not fully specified for unit
select sys datetime, adddate(sys datetime, interval '1:20' HOUR SECOND);

  SYS DATETIME                                date add( SYS DATETIME , INTERVAL '1:20' HOUR SECOND)
=====

```

```

06:18:24.149 PM 06/28/2010      06:19:44.149 PM 06/28/2010

SELECT ADDDATE('0000-00-00', 1 );

ERROR: Conversion error in date format.

-- In JDBC program, java error or depends on 'zeroDateTimeBehavior'
SELECT ADDDATE('0001-01-01 00:00:00', -1);

adddate('0001-01-01 00:00:00', -1)
=====
'12:00:00.000 AM 00/00/0000'
    
```

ADDTIME Function

Description

The **ADDTIME** function adds or subtracts a value of specific time.

The first argument is **DATE**, **DATETIME**, **TIMESTAMP**, or **TIME** type and the second argument is **TIME**, **DATETIME**, or **TIMESTAMP** type. Time should be include in the second argument, and the date of the second argument is ignored. The return type for each argument type is follows:

First Argument Type	Second Argument Type	Return Type	Note
TIME	TIME, DATETIME, TIMESTAMP	TIME	The result value must be equal to or less than 24 hours.
DATE	TIME, DATETIME, TIMESTAMP	DATETIME	
DATETIME	TIME, DATETIME, TIMESTAMP	DATETIME	
date/time string	TIME, DATETIME, TIMESTAMP or time string	VARCHAR	The result string includes time.

Syntax

```
ADDTIME ( expr1, expr2)
```

- *expr1* : **DATE**, **DATETIME**, **TIME** or **TIMESTAMP** type
- *expr2* : **DATETIME**, **TIMESTAMP**, **TIME** type or date/time string

Example

```

SELECT ADDTIME(datetime'2007-12-31 23:59:59', time'1:1:2');
      addtime(datetime '2007-12-31 23:59:59', time '1:1:2')
=====
01:01:01.000 AM 01/01/2008

SELECT ADDTIME(time'01:00:00', time'02:00:01');
      addtime(time '01:00:00', time '02:00:01')
=====
03:00:01 AM
    
```

ADD_MONTHS Function

The **ADD_MONTHS** function adds a *month* value to the expression *date_argument* of **DATE** type, and it returns a **DATE** type value. If the day (*dd*) of the value specified as an argument exists within the month of the result value of the operation, it returns the given day (*dd*); otherwise returns the last day of the given month (*dd*). If the result value of the operation exceeds the expression range of the **DATE** type, it returns an error.

Syntax

```
ADD_MONTHS ( date_argument , month )
```

- date_argument* :
- *date*
 - **NULL**

month :

- *integer*
- **NULL**

- *date_argument* : Specifies an expression of **DATE** type. To specify a **TIMESTAMP** or **DATETIME** value, an explicit casting to **DATE** type is required. If the value is **NULL**, **NULL** is returned.
- *month* : Specifies the number of the months to be added to the *date_argument*. Both positive and negative values can be specified. If the given value is not an integer type, conversion to an integer type by an implicit casting (rounding to the first place after the decimal point) is performed. If the value is **NULL**, **NULL** is returned.

Example

```
--it returns DATE type value by adding month to the first argument
SELECT ADD_MONTHS (DATE '2008-12-25', 5), ADD_MONTHS (DATE '2008-12-25', -5);
  add_months(date '2008-12-25', 5)    add_months(date '2008-12-25', -5)
=====
  05/25/2009                          07/25/2008

SELECT ADD_MONTHS (DATE '2008-12-31', 5.5), ADD_MONTHS (DATE '2008-12-31', -5.5);
  add_months(date '2008-12-31', 5.5)  add_months(date '2008-12-31', -5.5)
=====
  06/30/2009                          06/30/2008

SELECT ADD_MONTHS (CAST (SYS_DATETIME AS DATE), 5), ADD_MONTHS (CAST (SYS_TIMESTAMP AS DATE),
5);
  add_months( cast( SYS_DATETIME as date), 5)    add_months( cast( SYS_TIMESTAMP as date),
5)
=====
  07/03/2010                                  07/03/2010
```

CURDATE/CURRENT_DATE/CURRENT_DATE()/SYS_DATE/SYSDATE

Description

The **CURDATE()**, **CURRENT_DATE**, **CURRENT_DATE**, **SYS_DATE**, and **SYSDATE** are used interchangeably, and they return the current date as the **DATE** type (*MM/DD/YYYY* or *YYYY-MM-DD*). The unit is day.

If every argument value of date is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```
CURDATE ()
CURRENT_DATE ()
CURRENT_DATE
SYS_DATE
SYSDATE
```

Example

```
--it returns the current date in DATE type
SELECT CURDATE(), CURRENT DATE(), CURRENT DATE, SYS DATE, SYSDATE;

  SYS DATE    SYS DATE    SYS DATE    SYS DATE    SYS DATE
=====
  04/01/2010  04/01/2010  04/01/2010  04/01/2010  04/01/2010

--it returns the date 60 days added to the current date
SELECT CURDATE()+60;

  SYS DATE +60
=====
  05/31/2010

SELECT TO DAYS('0000-00-00');
ERROR: Conversion error in date format.
```

CURRENT_DATETIME/CURRENT_DATETIME()/NOW()/SYS_DATETIME/SYSDATETIME

Description

CURRENT_DATETIME, **CURRENT_DATETIME()**, **NOW()**, **SYS_DATETIME**, and **SYSDATETIME** are used interchangeably, and they return the current date and time in **DATETIME** type. The unit is millisecond.

Syntax

```
CURRENT_DATETIME
CURRENT_DATETIME ()
NOW ()
SYS_DATETIME
SYSDATETIME
```

Example

```
--it returns the current date and time in DATETIME type
SELECT NOW(), SYS DATETIME;

      SYS DATETIME                SYS DATETIME
=====
04:08:09.829 PM 02/04/2010      04:08:09.829 PM 02/04/2010

--it returns the timestamp value 1 hour added to the current sys datetime value
SELECT TO CHAR(SYSDATETIME+3600*1000, 'YYYY-MM-DD HH:MI');
to char( SYS DATETIME +3600*1000, 'YYYY-MM-DD HH:MI', 'en US')
=====
'2010-02-04 04:08'
```

CURTIME()/CURRENT_TIME/CURRENT_TIME()/SYS_TIME/SYSTIME

Description

The **CURTIME()**, **CURRENT_TIME**, **CURRENT_TIME()**, **SYS_TIME**, and **SYSTIME** are used interchangeably, and they return the current time as **TIME** type (*HH:MI:SS*). The unit is second.

Syntax

```
CURTIME ()
CURRENT_TIME
CURRENT_TIME ()
SYS_TIME
SYSTIME
```

Example

```
--it returns the current time in TIME type
SELECT CURTIME(), CURRENT_TIME(), CURRENT_TIME, SYS_TIME, SYSTIME;
      SYS TIME      SYS TIME      SYS TIME      SYS TIME      SYS TIME
=====
04:37:34 PM 04:37:34 PM 04:37:34 PM 04:37:34 PM 04:37:34 PM

--it returns the time value 1 hour added to the current sys_time
SELECT CURTIME()+3600;
      SYS TIME +3600
=====
05:37:34 PM
```

CURRENT_TIMESTAMP/CURRENT_TIMESTAMP()/SYS_TIMESTAMP/SYSTIMESTAMP/LOCALTIME/LOCALTIME()/LOCALTIMESTAMP/LOCALTIMESTAMP()

Description

The **CURRENT_TIMESTAMP**, **CURRENT_TIMESTAMP()**, **SYS_TIMESTAMP**, **SYSTIMESTAMP**, **LOCALTIME**, **LOCALTIME()**, **LOCALTIMESTAMP**, and **LOCALTIMESTAMP()** are used interchangeably, and they return the current date and time as **TIMESTAMP** type. The unit is second.

If you define **DEFAULT** value for column initial value and specify the initial value to **SYS_DATETIME**, the default value is specified to the timestamp at the time of creating a table, not inserting a table. Note that the default value is not specified in case of **INSERT**. Therefore, you must specify **SYS_DATETIME** in the **VALUES** of **INSERT** statement upon inserting data.

Syntax

```
CURRENT_TIMESTAMP
CURRENT_TIMESTAMP ()
SYS_TIMESTAMP
SYSTIMESTAMP
LOCALTIME
LOCALTIME ()
LOCALTIMESTAMP
LOCALTIMESTAMP ()
```

Example

```
--it returns the current date and time in TIMESTAMP type
SELECT LOCALTIME, SYS_TIMESTAMP;
SYS_TIMESTAMP          SYS_TIMESTAMP
=====
07:00:48 PM 04/01/2010    07:00:48 PM 04/01/2010

--it returns the timestamp value 1 hour added to the current sys timestamp value
SELECT CURRENT_TIMESTAMP()+3600;
SYS_TIMESTAMP +3600
=====
08:02:42 PM 04/01/2010
```

DATE Function

Description

The **DATE** function extracts the date part from specified argument, and returns it as *MM/DD/YYYY* format string. Arguments that can be specified are **DATE**, **TIMESTAMP** and **DATETIME** types. The return value is a **VARCHAR** type.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date and time, string where 0 is specified for every date value is returned.

Syntax

```
DATE (date)
```

- *date* : The **DATE**, **TIMESTAMP** or **DATETIME** can be specified.

Example

```
SELECT DATE('2010-02-27 15:10:23');
date('2010-02-27 15:10:23')
=====
'02/27/2010'

SELECT DATE(NOW());
date( SYS_DATETIME )
=====
'04/01/2010'

SELECT DATE('0000-00-00 00:00:00');
date('0000-00-00 00:00:00')
=====
'00/00/0000'
```

DATEDIFF Function

Description

The **DATEDIFF** function returns the difference between two arguments as an integer representing the number of days. Arguments that can be specified are **DATE**, **TIMESTAMP** and **DATETIME** types and its return value is only **INTEGER** type.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```
DATEDIFF (date1, date2)
```

- *date1, date2*: The **DATE**, **TIMESTAMP** or **DATETIME** type or date/time format string can be specified. If an invalid string is specified, an error is returned.

Example

```
SELECT DATEDIFF('2010-2-28 23:59:59', '2010-03-02');
datediff('2010-2-28 23:59:59', '2010-03-02')
=====
-2

SELECT DATEDIFF('0000-00-00 00:00:00', '2010-2-28 23:59:59');
ERROR: Conversion error in date format.
```

DATE_SUB()/SUBDATE() Functions

Description

The **DATE_SUB** and **SUBDATE()** are used interchangeably, and they perform an addition or subtraction operation on a specific **DATE** value. The value is returned in **DATE** or **DATETIME** type. If the date resulting from the operation exceeds the last day of the month, the function returns a valid **DATE** value considering the last date of the month.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

If the calculated value is between '0000-00-00 00:00:00' and '0001-01-01 00:00:00', a value having 0 for all arguments is returned in **DATE** or **DATETIME** type. Note that operation in JDBC program is determined by the configuration of zeroDateTimeBehavior, connection URL property (see "API Reference > JDBC API > JDBC Programming > Connection Configuration").

Syntax

```
DATE_SUB (date, INTERVAL expr unit)
SUBDATE (date, INTERVAL expr unit)
SUBDATE (date, days)
```

- *date*: It is a **DATE** or **TIMESTAMP** expression that represents the start date. If an invalid **DATE** value such as '2006-07-00' is specified, **NULL** is returned.
- *expr*: It represents the interval value to be subtracted from the start date. If a negative number is specified next to the **INTERVAL** keyword, the interval value is added to the start date.
- *unit*: It represents the unit of the interval value specified in the *expr* expression. To check the *expr* argument for the unit value, see the table of [ADDDATE/DATE_ADD Functions](#).

Example

```
SELECT SYSDATE, SUBDATE(SYSDATE, INTERVAL 24 HOUR), SUBDATE(SYSDATE, 1);
SYS DATE      date sub( SYS DATE , INTERVAL 24 HOUR)  subdate( SYS DATE , 1)
=====
03/30/2010  12:00:00.000 AM 03/29/2010                03/29/2010

--it adds days when argument < 0
```



```

SELECT SYSDATE, SUBDATE(SYSDATE,INTERVAL -24 HOUR), SUBDATE(SYSDATE, -1);
  SYS DATE      date sub( SYS DATE , INTERVAL -24 HOUR)      subdate( SYS DATE , -1)
=====
  03/30/2010  12:00:00.000 AM 03/31/2010                      03/31/2010

SELECT SUBDATE('0000-00-00 00:00:00', -50);
ERROR: Conversion error in date format.

SELECT SUBDATE('0001-01-01 00:00:00', 10);
  subdate('0001-01-01 00:00:00', 10)
=====
'12:00:00.000 AM 00/00/0000'

```

DAY/DAYOFMONTH Functions

Description

The **DAY** or **DAYOFMONTH** function returns day in the range of 1 to 31 from the specified parameter. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date, 0 is returned as an exception.

Syntax

```

DAY (date)
DAYOFMONTH (date)

```

- *date* : Date

Example

```

SELECT DAYOFMONTH('2010-09-09');
  dayofmonth('2010-09-09')
=====
  9

SELECT DAY('2010-09-09 19:49:29');
  day('2010-09-09 19:49:29')
=====
  9

SELECT DAYOFMONTH('01:02:03');
ERROR: Conversion error in date format.

SELECT DAYOFMONTH('0000-00-00 00:00:00');
  dayofmonth('0000-00-00 00:00:00')
=====
  0

```

DAYOFWEEK Function

Description

The **DAYOFWEEK** function returns a day in the range of 1 to 7 (1: Sunday, 2: Monday, ..., 7: Saturday) from the specified parameters. The day index is same as the ODBC standards. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```

DAYOFWEEK (date)

```

- *date* : Date

Example

```

SELECT DAYOFWEEK('2010-09-09');
      dayofweek('2010-09-09')
=====
                          5

SELECT DAYOFWEEK('2010-09-09 19:49:29');
      dayofweek('2010-09-09 19:49:29')
=====
                          5

SELECT DAYOFWEEK('10:28:00');
ERROR: Conversion error in date format.

SELECT DAYOFWEEK('0000-00-00');
ERROR: Conversion error in date format.

```

DAYOFYEAR Function**Description**

The **DAYOFYEAR** function returns the day of a year in the range of 1 to 366. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** types; the value is returned in **INTEGER** type.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```
DAYOFYEAR (date)
```

- *date* : Date

Example

```

SELECT DAYOFYEAR('2010-09-09');
      dayofyear('2010-09-09')
=====
                          252

SELECT DAYOFYEAR('2010-09-09 19:49:29');
      dayofyear('2010-09-09 19:49:29')
=====
                          252

SELECT DAYOFYEAR('10:28:00');
ERROR: Conversion error in date format.

SELECT DAYOFYEAR('0000-00-00');
ERROR: Conversion error in date format.

```

EXTRACT Operator**Description**

The **EXTRACT** operator extracts the values from *date-time_argument* and then converts the value type into **INTEGER**.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date and time, 0 is returned as an exception.

Syntax

```
EXTRACT ( field FROM date-time_argument )
```

- field* :
- **YEAR**
 - **MONTH**

- **DAY**
- **HOURL**
- **MINUTE**
- **SECOND**
- **MILLISECOND**

date-time argument :

- *expression*

- *field* : Specifies a value to be extracted from date-time expression.
- *date-time argument* : An expression that returns a value of date-time. This expression must be one of **TIME**, **DATE**, **TIMESTAMP**, or **DATETIME** types. If the value is **NULL**, **NULL** is returned.

Example

```

SELECT EXTRACT(MONTH FROM DATETIME '2008-12-25 10:30:20.123' );
  extract(month from datetime '2008-12-25 10:30:20.123')
=====
                                                12

SELECT EXTRACT(HOUR FROM DATETIME '2008-12-25 10:30:20.123' );
  extract(hour from datetime '2008-12-25 10:30:20.123')
=====
                                                10

SELECT EXTRACT(MILLISECOND FROM DATETIME '2008-12-25 10:30:20.123' );
  extract(millisecond from datetime '2008-12-25 10:30:20.123')
=====
                                                123

SELECT EXTRACT(MONTH FROM '0000-00-00 00:00:00');
  extract(month from '0000-00-00 00:00:00')
=====
                                                0

```

FROM_DAYS Function

Description

The **FROM_DAYS** function returns a date value in **DATE** type if **INTEGER** type is inputted as an argument.

It is not recommended to use the **FROM_DAYS** function for dates prior to the year 1582 because the function does not take dates prior to the introduction of the Gregorian Calendar into account.

If a value in the range of 0 to 3,652,424 can be inputted as an argument. If a value in the range of 0 to 365 is inputted, 0 is returned. 3,652,424, which is the maximum value, means the last day of year 9999.

Syntax

FROM_DAYS (*N*)

- *N* : Integer in the range of 0 to 3,652,424

Example

```

SELECT FROM_DAYS(719528);
  from days(719528)
=====
  01/01/1970

SELECT FROM_DAYS('366');
  from days('366')
=====
  01/03/0001

SELECT FROM_DAYS(3652424);
  from days(3652424)
=====
  12/31/9999

```

```

SELECT FROM DAYS(3652425);
ERROR: Conversion error in date format.

SELECT FROM_DAYS(-1);
ERROR: Conversion error in date format.

SELECT FROM DAYS(0);
      from days(0)
=====
      00/00/0000

```

FROM_UNIXTIME Function

Description

The **FROM_UNIXTIME** function returns the date and time in the format of 'YYYY-MM-DD HH:MM:SS.' You can specify **INTEGER** type that corresponds to the UNIX timestamp; the value is returned in **VARCHAR** type and is displayed in the current time zone.

It displays the result according to the format that you specified, and the time *format* format follows the Date/Time Format 2 table of [DATE_FORMAT Function](#).

The relationship is not one of one-to-one correspondence between **TIMESTAMP** and UNIX timestamp so if you use **UNIX_TIMESTAMP** or **FROM_UNIXTIME** function, partial value could be lost. For details, see [UNIX_TIMESTAMP Function](#).

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date and time, string where 0 is specified for every date and time value is returned. Note that operation in JDBC program is determined by the configuration of zeroDateTimeBehavior, connection URL property (see "API Reference > JDBC API > JDBC Programming > Connection Configuration").

Syntax

```
FROM_UNIXTIME( unix_timestamp [, format] )
```

- *unix_timestamp* : Positive integer
- *format* : Time format. Follows the date/time format of the [DATE_FORMAT Function](#).

Example

```

SELECT FROM_UNIXTIME(1234567890);
      from unixtime(1234567890)
=====
      01:31:30 AM 02/14/2009

SELECT FROM_UNIXTIME('1000000000');
      from_unixtime('1000000000')
=====
      04:46:40 AM 09/09/2001

SELECT FROM_UNIXTIME(1234567890, '%M %Y %W');
      from_unixtime(1234567890, '%M %Y %W')
=====
      'February 2009 Saturday'

SELECT FROM_UNIXTIME('1234567890', '%M %Y %W');
      from_unixtime('1234567890', '%M %Y %W')
=====
      'February 2009 Saturday'

SELECT FROM_UNIXTIME(-1);
ERROR: Conversion error in timestamp format.
Download in other formats:

SELECT FROM_UNIXTIME(-1);
ERROR: Conversion error in timestamp format.
Download in other formats:

```

```
SELECT FROM UNIXTIME(0);
      from unixtime(0)
=====
12:00:00 AM 00/00/0000
```

LAST_DAY Function

Description

The **LAST_DAY** function returns the last day of the given month as **DATE** type.

If every argument value of date and time is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```
LAST_DAY ( date_argument )
```

date_argument :

- **date**
- **NULL**

- *date_argument* : Specifies an expression of **DATE** type. To specify a **TIMESTAMP** or **DATETIME** value, explicit casting to **DATE** is required. If the value is **NULL**, **NULL** is returned.

Example

```
--it returns last day of the month in DATE type
SELECT LAST_DAY (DATE '1980-02-01'), LAST_DAY (DATE '2010-02-01');
      last day(date '1980-02-01')      last day(date '2010-02-01')
=====
02/28/1980                          02/28/2010

--it returns last day of the month when explicitly casted to DATE type
SELECT LAST_DAY (CAST (SYS_TIMESTAMP AS DATE)), LAST_DAY (CAST (SYS_DATETIME AS DATE));
      last day( cast( SYS_TIMESTAMP as date))      last day( cast( SYS_DATETIME as date))
=====
02/28/2010                          02/28/2010

SELECT LAST_DAY ('0000-00-00');
ERROR: Conversion error in date format.
```

MAKEDATE Function

Description

The **MAKEDATE** function returns a date from the specified parameter. You can specify an **INTEGER** type corresponding to the day of the year in the range of 1 to 9999 as an argument; the value in the range of 1/1/1 to 12/31/9999 is returned in **DATE** type. If the day of the year has passed the corresponding year, it will become the next year. For example, **MAKEDATE**(1999, 366) will return 2000-01-01.

However, if you input a value in the range of 0 to 69 as the year, it will be processed as the year 2000-2069, if it is a value in the range of 70 to 99, it will be processed as the year 1970-1999.

If every value specified in *year* and *dayofyear* is 0, the return value is determined by the **return_null_on_function_errors** system parameter; if it is set to yes, then **NULL** is returned; if it is set to no, an error is returned. The default value is **no**.

Syntax

```
MAKEDATE (year, dayofyear)
```

- *year* : Year in the range of 1 to 9999
- *dayofyear* : If you input a value in the range of 0 to 99 in the argument, it is handled as an exception; *dayofyear* must be equal to or less than 3,615,902 and the return value of **MAKEDATE**(100, 3615902) is 9999/12/31.

Example

```

SELECT MAKEDATE (2010,277);
    makedate(2010, 277)
=====
    10/04/2010

SELECT MAKEDATE (10,277);
    makedate(10, 277)
=====
    10/04/2010

SELECT MAKEDATE (70,277);
    makedate(70, 277)
=====
    10/04/1970

SELECT MAKEDATE (100,3615902);
    makedate(100, 3615902)
=====
    12/31/9999

SELECT MAKEDATE ('9999','365');
    makedate('9999', '365')
=====
    12/31/9999

SELECT MAKEDATE (9999,366);
ERROR: Conversion error in date format.

SELECT MAKEDATE (0,0);
ERROR: Conversion error in date format.

```

MAKETIME Function**Description**

The **MAKETIME** function returns the hour from specified argument in the AM/PM format. You can specify the **INTEGER** types corresponding hours, minutes and seconds as arguments; the value is returned in **DATETIME**.

Syntax

```
MAKETIME (hour, min, sec)
```

- *hour* : Integers representing the hours in the range of 0 to 23
- *min* : Integers representing the minutes in the range of 0 to 59
- *sec* : Integers representing the seconds in the range of 0 to 59

Example

```

SELECT MAKETIME (13,34,4);
    maketime(13, 34, 4)
=====
    01:34:04 PM

SELECT MAKETIME ('1','34','4');
    maketime('1', '34', '4')
=====
    01:34:04 AM

SELECT MAKETIME (24,0,0);
ERROR: Conversion error in time format.

```

MINUTE Function

Description

The **MINUTE** function returns the minutes in the range of 0 to 59 from specified argument. You can specify the **TIME**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

Syntax

```
MINUTE (time)
```

- *time* : Time

Example

```
SELECT MINUTE('12:34:56');
      minute('12:34:56')
=====
                          34

SELECT MINUTE('2010-01-01 12:34:56');
      minute('2010-01-01 12:34:56')
=====
                          34

SELECT MINUTE('2010-01-01 12:34:56.7890');
      minute('2010-01-01 12:34:56.7890')
=====
                          34

SELECT MINUTE('2010-01-01');

      In the command from line 1,
      ERROR: Conversion error in time format.
```

MONTH Function

Description

The **MONTH** function returns the month in the range of 1 to 12 from specified argument. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date, 0 is returned as an exception.

Syntax

```
MONTH (date)
```

- *date* : Date

Example

```
SELECT MONTH('2010-01-02');
      month('2010-01-02')
=====
                          1

SELECT MONTH('2010-01-02 12:34:56');
      month('2010-01-02 12:34:56')
=====
                          1

SELECT MONTH('2010-01-02 12:34:56.7890');
      month('2010-01-02 12:34:56.7890')
=====
                          1

SELECT MONTH ('12:34:56');
```

```
ERROR: Conversion error in date format.
```

```
SELECT MONTH('0000-00-00');
       month('0000-00-00')
=====
0
```

MONTHS_BETWEEN Function

Description

The **MONTHS_BETWEEN** function returns the difference between the given **DATE** value. The return value is **DOUBLE** type. An integer value is returned if the two dates specified as arguments are identical or are the last day of the given month; otherwise, a value obtained by dividing the day difference by 31 is returned.

Syntax

```
MONTHS_BETWEEN(date_argument, date_argument)
```

date_argument :

- *date*
- **NULL**

- *date_argument* : Specifies an expression of **DATE** type. To specify a **TIMESTAMP** or **DATETIME** value, explicit casting to **DATE** is required. If the value is **NULL**, **NULL** is returned.

Example

```
--it returns the negative months when the first argument is the previous date
SELECT MONTHS_BETWEEN(DATE '2008-12-31', DATE '2010-6-30');
       months between(date '2008-12-31', date '2010-6-30')
=====
-1.8000000000000000e+001

--it returns integer values when each date is the last dat of the month
SELECT MONTHS_BETWEEN(DATE '2010-6-30', DATE '2008-12-31');
       months between(date '2010-6-30', date '2008-12-31')
=====
1.8000000000000000e+001

--it returns months between two arguments when explicitly casted to DATE type
SELECT MONTHS_BETWEEN(CAST (SYS_TIMESTAMP AS DATE), DATE '2008-12-25');
       months between( cast( SYS_TIMESTAMP as date), date '2008-12-25')
=====
1.332258064516129e+001

--it returns months between two arguments when explicitly casted to DATE type
SELECT MONTHS_BETWEEN(CAST (SYS_DATETIME AS DATE), DATE '2008-12-25');
       months between( cast( SYS DATETIME as date), date '2008-12-25')
=====
1.332258064516129e+001
```

QUARTER Function

Description

The **QUARTER** function returns the quarter in the range of 1 to 4 from specified argument. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

Syntax

```
QUARTER(date)
```

- *date* : Date

Example

```
SELECT QUARTER('2010-05-05');
       quarter('2010-05-05')
```



```

=====
                2
SELECT QUARTER('2010-05-05 12:34:56');
       quarter('2010-05-05 12:34:56')
=====
                2

SELECT QUARTER('2010-05-05 12:34:56.7890');
       quarter('2010-05-05 12:34:56.7890')
=====
                2

SELECT QUARTER('12:34:56');
ERROR: Conversion error in date format.

```

SEC_TO_TIME Function

Description

The **SEC_TO_TIME** function returns the time including hours, minutes and seconds from specified argument. You can specify the **INTEGER** type in the range of 0 to 86,399; the value is returned in **TIME** type.

Syntax

```
SEC_TO_TIME(second)
```

- *second* : Seconds in the range of 0 to 86,399

Example

```

SELECT SEC_TO_TIME(82800);
       sec to time(82800)
=====
       11:00:00 PM

SELECT SEC_TO_TIME('82800.3');
       sec to time('82800.3')
=====
       11:00:00 PM

SELECT SEC_TO_TIME(86399)
       sec_to_time(86399)
=====
       11:59:59 PM

SELECT SEC TO TIME(86400);
ERROR: Conversion error in time format.

```

SECOND Function

Description

The **SECOND** function returns the seconds in the range of 0 to 59 from specified argument. You can specify the **TIME**, **TIMESTAMP**, or **DATETIME**; the value is returned in **INTEGER** type.

If the function fails, **NULL** is returned when the database server configuration parameter **return_null_on_function_errors** is set to yes. When the parameter is set to no, the function outputs error message. The default value of **return_null_on_function_errors** is no.

Syntax

```
SECOND(time)
```

- *time* : Time

Example

```

SELECT SECOND('12:34:56');
      second('12:34:56')
=====
                          56

SELECT SECOND('2010-01-01 12:34:56');
      second('2010-01-01 12:34:56')
=====
                          56

SELECT SECOND('2010-01-01 12:34:56.7890');
      second('2010-01-01 12:34:56.7890')
=====
                          56

SELECT SECOND ('2010-01-01');
ERROR: Conversion error in time format.

```

STR_TO_DATE Function**Description**

The **STR_TO_DATE** function converts the given character string to a date/time value by interpreting it according to the specified format and operates in the opposite way to the [DATE_FORMAT Function](#). The return value is determined by the date/time part included in the character string and it is one of the **DATETIME**, **DATE**, and **TIME** types.

If the *string* includes an invalid date/time value or the character string cannot be interpreted by applying the format specifier specified in the *format*, an error will be returned.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date and time, the value of **DATE** or **DATETIME** type that has 0 for every date and time value is returned as an exception. Note that operation in JDBC program is determined by the configuration of zeroDateTimeBehavior, connection URL property (see [Connection Configuration](#) of JDBC API).

Syntax

```
STR_TO_DATE(string, format)
```

- *string* : All character string types can be specified.
- *format* : Specifies the format to interpret the character string. You should use character strings including % for the format specifiers. See the table, date/time format 2 of [DATE_FORMAT Function](#).

Example

```

SELECT STR_TO_DATE('01,5,2013','%d,%m,%Y');
      str_to_date('01,5,2013', '%d,%m,%Y')
=====
      05/01/2013

SELECT STR_TO_DATE('May 1, 2013','%M %d,%Y');
      str_to_date('May 1, 2013', '%M %d,%Y')
=====
      05/01/2013

SELECT STR_TO_DATE('13:30:17','%h:%i');
      str_to_date('13:30:17', '%h:%i')
=====
      01:30:00 PM

SELECT STR_TO_DATE('09:30:17 PM','%r');
      str_to_date('09:30:17 PM', '%r')
=====
      09:30:17 PM

SELECT STR_TO_DATE('0,0,0000','%d,%m,%Y');

```

```
str to date('0,0,0000', '%d,%m,%Y')
=====
00/00/0000
```

TIME Function

Description

The **TIME** function extracts the time part from specified argument and returns it in the 'HH:MM:SS' format. You can specify the **TIME**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **VARCHAR** type.

Syntax

```
TIME(time)
```

- *time*: Time

Example

```
SELECT TIME('12:34:56');
      time('12:34:56')
=====
      '12:34:56.000'

SELECT TIME('2010-01-01 12:34:56');
      time('2010-01-01 12:34:56')
=====
      '12:34:56.000'

SELECT TIME(datetime'2010-01-01 12:34:56');
      time(datetime '2010-01-01 12:34:56')
=====
      '12:34:56.000'

SELECT TIME('invalid string');

      ERROR: ERROR: Conversion error in time format.
```

TIME_TO_SEC Function

Description

The **TIME_TO_SEC** function returns the seconds in the range of 0 to 86,399 from specified argument. You can specify the **TIME**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

Syntax

```
TIME_TO_SEC(time)
```

- *time*: Time

Example

```
SELECT TIME_TO_SEC('23:00:00');
      time to sec('23:00:00')
=====
              82800

SELECT TIME_TO_SEC('2010-10-04 23:00:00');
      time_to_sec('2010-10-04 23:00:00')
=====
              82800

SELECT TIME TO SEC('2010-10-04 23:00:00.1234');
      time_to_sec('2010-10-04 23:00:00.1234')
=====
              82800

SELECT TIME_TO_SEC('2010-01-01');
```

```
ERROR: Conversion error in time format.
```

TIMEDIFF Function

Description

The **TIMEDIFF** function returns the time difference between the two specified time arguments.

You can enter a date/time type, the **TIME**, **DATE**, **TIMESTAMP**, or **DATETIME** type and the data types of the two arguments must be identical. The **TIME** will be returned and the time difference between the two arguments must be in the range of 00:00:00 -23:59:59. If it exceeds the range, an error will be returned.

Syntax

```
TIMEDIFF(expr1, expr2)
```

- *expr1*, *expr2* : Time. The data types of the two arguments must be identical.

Example

```
SELECT TIMEDIFF(time '17:18:19', time '12:05:52');
      timediff(time '17:18:19', time '12:05:52')
=====
      05:12:27 AM

SELECT TIMEDIFF('17:18:19', '12:05:52');
      timediff('17:18:19', '12:05:52')
=====
      05:12:27 AM

SELECT TIMEDIFF('2010-01-01 06:53:45', '2010-01-01 03:04:05');
      timediff('2010-01-01 06:53:45', '2010-01-01 03:04:05')
=====
      03:49:40 AM

SELECT TIMEDIFF('2010-01-02 06:53:45', '2010-01-01 03:04:05');
ERROR: ERROR: Conversion error in time format.
```

TIMESTAMP Function

Description

The **TIMESTAMP** function converts a **DATE** or **TIMESTAMP** type expression to **DATETIME** type.

If the **DATE** format string ('YYYY-MM-DD' or 'MM/DD/YYYY') or **TIMESTAMP** format string ('YYYY-MM-DD HH:MI:SS' or 'HH:MI:SS MM/DD/YYYY') is specified as the first argument, the function returns it as **DATETIME**.

If the **TIME** format string ('HH:MI:SS') is specified as the second, the function adds it to the first argument and returns the result as a **DATETIME** type. If the second argument is not specified, **12:00:00.000 AM** is specified by default.

Syntax

```
TIMESTAMP(date [, time])
```

- *date* : The following format strings can be specified : 'YYYY-MM-DD', 'MM/DD/YYYY', 'YYYY-MM-DD HH:MI:SS', and 'HH:MI:SS MM/DD/YYYY'
- *time* : The following format string can be specified : 'HH:MI:SS'

Example

```
SELECT TIMESTAMP('2009-12-31'), TIMESTAMP('2009-12-31', '12:00:00');
      timestamp('2009-12-31')          timestamp('2009-12-31', '12:00:00')
=====
      12:00:00.000 AM 12/31/2009      12:00:00.000 PM 12/31/2009
```

```

SELECT TIMESTAMP('2010-12-31 12:00:00','12:00:00');
timestamp('2010-12-31 12:00:00', '12:00:00')
=====
12:00:00.000 AM 01/01/2011

SELECT TIMESTAMP('13:10:30 12/25/2008');
timestamp('13:10:30 12/25/2008')
=====
01:10:30.000 PM 12/25/2008

```

TO_DAYS Function

Description

The **TO_DAYS** function returns the number of days after year 0 in the range of 366 to 3652424 from specified argument. You can specify **DATE** type; the value is returned in **INTEGER** type.

It is not recommended to use the **TO_DAYS** function for dates prior to the year 1582, as the function does not take dates prior to the introduction of the Gregorian Calendar into account.

Syntax

```
TO_DAYS (date)
```

- *date* : Date

Example

```

SELECT TO_DAYS('2010-10-04');
to days('2010-10-04')
=====
734414

SELECT TO_DAYS('2010-10-04 12:34:56');
to days('2010-10-04 12:34:56')
=====
734414

SELECT TO_DAYS('2010-10-04 12:34:56.7890');
to days('2010-10-04 12:34:56.7890')
=====
734414

SELECT TO_DAYS('1-1-1');
to days('1-1-1')
=====
366

SELECT TO_DAYS('9999-12-31');
to_days('9999-12-31')
=====
3652424

SELECT TO_DAYS ('12:34:56');
ERROR: Conversion error in date format.

```

UNIX_TIMESTAMP Function

Description

The arguments of the **UNIX_TIMESTAMP** function can be omitted. If they are omitted, the function returns the interval between '1970-01-01 00:00:00' UTC and the current system date/time in seconds as **INTEGER** type. If the date argument is specified, the function returns the interval between '1970-01-01 00:00:00' UTC and the specified date/time in seconds.

0 is not allowed in the argument value corresponding to year, month, and day; however, if 0 is inputted in every argument value corresponding to date and time, 0 is returned as an exception.

Syntax

```
UNIX_TIMESTAMP ( [date] )
```

- *date* : **DATE** type or **TIMESTAMP** type, **DATE** format string ('YYYY-MM-DD' or 'MM/DD/YYYY'), **TIMESTAMP** format string ('YYYY-MM-DD HH:MI:SS' or 'HH:MI:SS MM/DD/YYYY') or 'YYYYMMDD' string can be specified.

Example

```
SELECT UNIX_TIMESTAMP('1970-01-02'), UNIX_TIMESTAMP();
       unix_timestamp('1970-01-02')    unix_timestamp()
=====
                          54000          1270196737

SELECT UNIX_TIMESTAMP ('0000-00-00 00:00:00');
       unix_timestamp('0000-00-00 00:00:00')
=====
                                          0
```

UTC_DATE Function

Description

The **UTC_DATE** function returns the UTC date in 'YYYY-MM-DD' format.

Syntax

```
UTC_DATE ()
```

Example

```
SELECT UTC_DATE ();
       utc_date()
=====
       01/12/2011
```

UTC_TIME Function

Description

The **UTC_TIME** function returns the UTC time in 'HH:MM:SS' format.

Syntax

```
UTC_TIME ()
```

Example

```
SELECT UTC_TIME ();
       utc_time()
=====
       10:35:52 AM
```

WEEK Function

Description

The **WEEK** function returns the week in the range of 0 to 53 from specified argument. You can specify the **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

You can omit the second argument, *mode* and must input a value in the range of 0 to 7. You can set that a week starts from Sunday or Monday and the range of the return value is from 0 to 53 or 1 to 53 with this value. If you omit the *mode*, the system parameter, **default_week_format** value will be used. The *mode* value means as follows:

mode	Start Day of the Week	Range	The First Week of the Year
0	Sunday	0~53	The first week that Sunday is included in the year
1	Monday	0~53	The first week that more than three days are included in the year
2	Sunday	1~53	The first week in the year that includes a Sunday
3	Monday	1~53	The first week in the year that includes more than three days
4	Sunday	0~53	The first week in the year that includes more than three days
5	Monday	0~53	The first week in the year that includes Monday
6	Sunday	1~53	The first week in the year that includes more than three days
7	Monday	1~53	The first week in the year that includes Monday

If the *mode* value is one of 0, 1, 4 or 5, and the date corresponds to the last week of the previous year, the **WEEK** function will return 0. The purpose is to see what nth of the year the week is so it returns 0 for the 52th week of the year 1999.

```
SELECT YEAR('2000-01-01'), WEEK('2000-01-01',0);
      year('2000-01-01')    week('2000-01-01', 0)
=====
                2000                0
```

To see what n-th the week is based on the year including the start day of the week, use 0, 2, 5 or 7 as the *mode* value.

```
SELECT WEEK('2000-01-01',2);
      week('2000-01-01', 2)
=====
                52
```

Syntax

```
WEEK(date[, mode])
```

- *date* : Date
- *mode* : Value in the range of 0 to 7

Example

```
SELECT WEEK('2010-04-05');
      week('2010-04-05', 0)
=====
                14

SELECT WEEK('2010-04-05 12:34:56',2);
      week('2010-04-05 12:34:56',2)
=====
                14

SELECT WEEK('2010-04-05 12:34:56.7890',4);
      week('2010-04-05 12:34:56.7890',4)
=====
                14

SELECT WEEK ('12:34:56');

ERROR: Conversion error in date format.

SELECT WEEK('2010-04-05',8);

ERROR: Conversion error in date format.
```

WEEKDAY Function

Description

The **WEEKDAY** function returns the day of week in the range of 0 to 6 (0: Sunday, 1: Monday, ..., 6: Saturday) from the specified parameter. The day of week index is same as the ODBC standards. You can specify **DATE**, **TIMESTAMP**, **DATETIME** types as parameters and an **INTEGER** type will be returned.

Syntax

WEEKDAY (*date*)

- date* : Date

Example

```
SELECT WEEKDAY('2010-09-09');
      weekday('2010-09-09')
=====
                          3

SELECT WEEKDAY('2010-09-09 13:16:00');
      weekday('2010-09-09 13:16:00')
=====
                          3

SELECT WEEKDAY('10:28:00');
ERROR: Conversion error in date format.
```

YEAR Function

Description

The **YEAR** function returns the year in the range of 1 to 9,999 from the specified parameter. You can specify **DATE**, **TIMESTAMP**, or **DATETIME** type; the value is returned in **INTEGER** type.

Syntax

YEAR (*date*)

- date* : Date

Example

```
SELECT YEAR('2010-10-04');
      year('2010-10-04')
=====
                2010

SELECT YEAR('2010-10-04 12:34:56');
      year('2010-10-04 12:34:56')
=====
                2010

SELECT YEAR('2010-10-04 12:34:56.7890');
      year('2010-10-04 12:34:56.7890')
=====
                2010

SELECT YEAR('12:34:56');
ERROR: Conversion error in date format.
```


Data Type Conversion Functions and Operators

CAST Operator

Description

The **CAST** operator can be used to explicitly cast one data type to another in the **SELECT** statement. A query list or a value expression in the **WHERE** clause can be cast to another data type.

Depending on the situation, data type can be automatically converted without using the **CAST** operator. For details, see [Implicit Type Conversion](#).

See [Converting the String of Date/Time Data Type into Date/Time Type](#) regarding to convert the string of date/time type into date/time type.

The following table shows a summary of explicit type conversions (casts) using the **CAST** operator in CUBRID.

From \ To	EN	AN	VC	FC	VB	FB	BLOB	CLOB	D	T	UT	DT	S	MS	SQ
EN	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No
AN	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No
VC	Yes	Yes	Yes*	Yes*	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
FC	Yes	Yes	Yes*	Yes*	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
VB	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No
FB	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No
BLOB	No	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No
CLOB	No	No	Yes	Yes	Yes	Yes	No	Yes	No	No	No	No	No	No	No
D	No	No	Yes	Yes	No	No	No	No	Yes	No	Yes	Yes	No	No	No
T	No	No	Yes	Yes	No	No	No	No	No	Yes	No	No	No	No	No
UT	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No
DT	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No
S	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
MS	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
SQ	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes

* The **CAST** operation is allowed only when the value expression and the data type to be cast have the same character code set.

Data Type Key

- **EN** : Exact numeric data type (**INTEGER**, **SMALLINT**, **BIGINT**, **NUMERIC**, **DECIMAL**)
- **AN** : Approximate numeric data type (**FLOAT/REAL**, **DOUBLE PRECISION**, **MONETARY**)
- **VC** : Variable-length character string (**VARCHAR(n)**, **NCHAR VARYING(n)**)
- **FC** : Fixed-length character string (**CHAR(n)**, **NCHAR(n)**)
- **VB** : Variable-length bit string (**BIT VARYING(n)**)
- **FB** : Fixed-length bit string (**BIT(n)**)
- **BLOB** : Binary data that is stored outside DB
- **CLOB** : String data that is stored inside DB
- **D** : Date (**DATE**)
- **T** : Time (**TIME**)
- **UT** : Timestamp (**TIMESTAMP**)
- **S** : Set (**SET**)

- **MS** : Multiset (**MULTISET**)
- **SQ** : Sequence set (**LIST, SEQUENCE**)

Syntax

CAST (*cast operand* **AS** *cast target*)

cast_operand :

- *value expression*
- **NULL**

cast_target :

- *data type*

- *cast_operand* : Declares the value to cast to a different data type.
- *cast_target* : Specifies the type to cast to.

Example

The following example shows how to explicitly cast and return a **VARCHAR** record in kg unit to a **FLOAT**.

```
--operation after casting character as INT type returns 2
SELECT (1+CAST ('1' AS INT));
      (1+ cast('1' as integer))
=====
              2

--cannot cast the string which is out of range as SMALLINT
SELECT (1+CAST('1234567890' AS SMALLINT));

ERROR: Cannot coerce value of domain "character" to domain "smallint".
--operation after casting returns 1+1234567890
SELECT (1+CAST('1234567890' AS INT));
      (1+ cast('1234567890' as integer))
=====
              1234567891

--'1234.567890' is casted to 1235 after rounding up
SELECT (1+CAST('1234.567890' AS INT));
      (1+ cast('1234.567890' as integer))
=====
              1236

--'1234.567890' is casted to string containing only first 5 letters.
SELECT (CAST('1234.567890' AS CHAR(5)));
      ( cast('1234.567890' as char(5)))
=====
      '1234.'

--numeric type can be casted to CHAR type only when enough length is specified
SELECT (CAST(1234.567890 AS CHAR(5)));

ERROR: Cannot coerce value of domain "numeric" to domain "character".
--numeric type can be casted to CHAR type only when enough length is specified
SELECT (CAST(1234.567890 AS CHAR(11)));
      ( cast(1234.567890 as char(11)))
=====
      '1234.567890'

--numeric type can be casted to CHAR type only when enough length is specified
SELECT (CAST(1234.567890 AS VARCHAR));
      ( cast(1234.567890 as varchar))
=====
      '1234.567890'

--string can be casted to time/date types only when its literal is correctly specified
SELECT (CAST('2008-12-25 10:30:20' AS TIMESTAMP));
      ( cast('2008-12-25 10:30:20' as timestamp))
=====
      10:30:20 AM 12/25/2008

SELECT (CAST('10:30:20' AS TIME));
```

```

( cast('10:30:20' as time))
=====
10:30:20 AM

--string can be casted to TIME type when its literal is same as TIME's.
SELECT (CAST('2008-12-25 10:30:20' AS TIME));
( cast('2008-12-25 10:30:20' as time))
=====
10:30:20 AM

--string can be casted to TIME type after specifying its type of the string
SELECT (CAST(TIMESTAMP'2008-12-25 10:30:20' AS TIME));
( cast(timestamp '2008-12-25 10:30:20' as time))
=====
10:30:20 AM

SELECT CAST('abcde' AS BLOB);
cast('abcde' as blob)
=====
file:/home1/user1/db/tdb/lob/ces_743/ces_temp.00001283232024309172_1342

SELECT CAST(B'11010000' as varchar(10));
cast(B'11010000' as varchar(10))
=====
'd0'

SELECT CAST('1A' AS BLOB);
cast('1A' as bit(16))
=====
x'1a00'

```

Remark

- **CAST** is allowed only between data types having the same character set.
- If you cast an approximate data type to integer type, the number is rounded to zero decimal places.
- If you cast a numeric data type to string character type, it should be longer than the length of significant figures + decimal point. An error occurs otherwise.
- If you cast a character string type A to a character string type B, B should be longer than the A. The end of character string is truncated otherwise.
- If you cast a character string type A to a date-time date type B, it is converted only when literal of A and B type match one another. An error occurs otherwise.
- You must explicitly do type casting for numeric data stored in a character string so that an arithmetic operation can be performed.

DATE_FORMAT Function

Description

The **DATE_FORMAT** function converts the value of strings with **DATE** format ('YYYY-MM-DD' or 'MM/DD/YYYY') or that of date/time data type (**DATE**, **TIMESTAMP**, **DATETIME**) to specified date/time format and then return the value with the **VARCHAR** data type.

Syntax

```
DATE_FORMAT (date, format)
```

- *date* : A value of strings with the **DATE** format ('YYYY-MM-DD' or 'MM/DD/YYYY') or that of date/time data type (**DATE**, **TIMESTAMP**, **DATETIME**) can be specified .
- *format* : Specifies the output format. Use a string that contains '%' as a specifier. See the following table to specify the format. Date/Time formats described in the following [Date/Time Format 2](#) table are used in **DATE_FORMAT** function, and [TIME_FORMAT Function](#), and [STR_TO_DATE Function](#).

Default Date/Time Format

Date/Time Type	Default Output Format
DATE	'MM/DD/YYYY'

TIME	'HH:MI:SS AM'
TIMESTAMP	'HH:MI:SS AM MM/DD/YYYY'
DATETIME	'HH:MI:SS.FF AM MM/DD/YYYY'
Date/Time Format 2	
format	Value
Meaning	
%a	Weekday, English abbreviation (Sun, ..., Sat)
%b	Month, English abbreviation (Jan, ..., Dec)
%c	Month (1, ..., 12)
%D	Day of the month, English ordinal number (1st, 2nd, 3rd, ...)
%d	Day of the month, two-digit number (01, ..., 31)
%e	Day of the month (1, ..., 31)
%f	Microseconds, three-digit number (000, ..., 999)
%H	Hour, 24-hour based, number with at least two--digit (00, ..., 23, ..., 100, ...)
%h	Hour, 12-hour based two-digit number (01, ..., 12)
%I	Hour, 12-hour based two-digit number (01, ..., 12)
%i	Minutes, two-digit number (00, ..., 59)
%j	Day of year, three-digit number (001, ..., 366)
%k	Hour, 24-hour based, number with at least one-digit (0, ..., 23, ..., 100, ...)
%l	Hour, 12-hour based (1, ..., 12)
%M	Month, English string (January, ..., December)
%m	Month, two-digit number (01, ..., 12)
%p	AM or PM
%r	Time, 12-hour based, hour:minute:second (hh:mm:ss AM or hh:mm:ss PM)
%S	Seconds, two-digit number (00, ..., 59)
%s	Seconds, two-digit number (00, ..., 59)
%T	Time, 24-hour based, hour:minute:second (hh:mm:ss)
%U	Week, two-digit number, week number of the year with Sunday being the first day Week (00, ..., 53)
%u	Week, two-digit number, week number of the year with Monday being the first day (00, ..., 53)
%V	Week, two-digit number, week number of the year with Sunday being the first day Week (00, ..., 53) (Available to use in combination with %X)
%v	Week, two-digit number, week number of the year with Monday being the first day (00, ..., 53) (Available to use in combination with %X)
%W	Weekday, English string (Sunday, ..., Saturday)
%w	Day of the week, number index (0=Sunday, ..., 6=Saturday)
%X	Year, four-digit number calculated as the week number with Sunday being the first day of the week (0000, ..., 9999) (Available to use in combination with %V)
%x	Year, four-digit number calculated as the week number with Monday being the first day of the week (0000, ..., 9999)

	(Available to use in combination with %V)
%Y	Year, four-digit number (0001, ..., 9999)
%y	Year, two-digit number (00, 01, ..., 99)
%%	Output the special character "%" as a string
%x	Output an arbitrary character x as a string out of English letters that are not used as format specifiers.

Example

```

SELECT DATE_FORMAT('2009-10-04 22:23:00', '%W %M %Y');
date_format('2009-10-04 22:23:00', '%W %M %Y')
=====
'Sunday October 2009'

SELECT DATE_FORMAT('2007-10-04 22:23:00', '%H:%i:%s');
date_format('2007-10-04 22:23:00', '%H:%i:%s')
=====
'22:23:00'

SELECT DATE_FORMAT('1900-10-04 22:23:00', '%D %y %a %d %m %b %j');
date_format('1900-10-04 22:23:00', '%D %y %a %d %m %b %j')
=====
'4th 00 Thu 04 10 Oct 277'

SELECT DATE_FORMAT('1999-01-01', '%X %V');
date_format('1999-01-01', '%X %V')
=====
'1998 52'

```

TIME_FORMAT Function

Description

The **TIME_FORMAT** function converts the value of strings with **TIME** format (*HH:MI:SS*) or that of date/time data type (**DATE**, **TIMESTAMP**, **DATETIME**) to specified date/time format and then return the value with the **VARCHAR** data type.

Syntax

```
TIME_FORMAT(time, format)
```

- time* : A value of string with **TIME** (*HH:MI:SS*) or that of date/time data type (**DATE**, **TIMESTAMP**, **DATETIME**) can be specified.
- format* : Specifies the output format. Use a string that contains ‘%’ as a specifier. See the table of the [Date/Time Format 2](#) table. If un-related format specifier is used, the English letters themselves are displayed.

Example

```

SELECT TIME_FORMAT('22:23:00', '%H %i %s');
time_format('22:23:00', '%H %i %s')
=====
'22 23 00'

SELECT TIME_FORMAT('23:59:00', '%H %h %i %s %f');
time_format('23:59:00', '%H %h %i %s %f')
=====
'23 11 59 00 000'

SELECT SYSTIME, TIME_FORMAT(SYSTIME, '%T');
SYS TIME      time_format( SYS TIME , '%T')
=====
08:46:53 PM  '20:46:53'

```

TO_CHAR Function (date_time)

Description

The **TO_CHAR** function converts the value of strings with **TIME** format (*HH:MI:SS*) or that of date/time type (**TIME**, **TIMESTAMP**, **DATETIME**) by [Date/Time Format 1](#) and then return the value with the **VARCHAR** data type. If a format argument is not specified, it converts the value based by default format. If a format which is not corresponding to the given value, an error is returned.

Syntax

```
TO_CHAR( date_time [, format [, date_lang_string_literal ] ] )
```

date_time :

- *date*
- *time*
- *timestamp*
- *datetime*
- **NULL**

format :

- character strings (see [Date/Time Format 1](#))
- **NULL**

date_lang_string_literal : (see [date lang string literal](#))

- 'en_US'
- 'ko_KR'

- *date_time* : Specifies an expression that returns date-time type string. If the value is **NULL**, **NULL** is returned.
- *format* : Specifies a format of return value. If a format is not specified, the default format is used. If the value is **NULL**, **NULL** is returned.
- *date_lang_string_literal* : Specifies a language applied to a return value (see [date lang string literal](#)). The default value is 'en_US'. You can modify the value by specifying the **CUBRID_DATE_LANG** environment variable.

Default Date/Time Format

Date/Time Type	Default Output Format
DATE	'MM/DD/YYYY'
TIME	'HH:MI:SS AM'
TIMESTAMP	'HH:MI:SS AM MM/DD/YYYY'
DATETIME	'HH:MI:SS.FF AM MM/DD/YYYY'

Date/Time Format 1

Format Element	Description
CC	Century
YYYY, YY	Year with 4 numbers, Year with 2 numbers
Q	Quarter (1, 2, 3, 4; January - March = 1)
MM	Month (01-12; January = 01) Note : MI represents the minute of hour.
MONTH	Month in characters
MON	Abbreviated month name
DD	Day (1 - 31)
DAY	Day of the week in characters
DY	Abbreviated day of the week
D or d	Day of the week in numbers (1 - 7)

AM or PM	AM/PM
A.M. or P.M.	AM/PM with periods
HH or HH12	Hour (1 -12)
HH24	Hour (0 - 23)
MI	Minute (0 - 59)
SS	Second (0 - 59)
FF	Millisecond (0-999)
- / , . ; : "text"	Punctuation and quotation marks are represented as they are in the result

Example of date_lang_string_literal

Format Element	Date_lang_string_literal	
	'en_US'	'ko_KR'
MONTH	JANUARY	1 월
MON	JAN	1
DAY	MONDAY	월요일
DY	MON	월
Month	January	1 월
Mon	Jan	1
Day	Monday	월요일
Dy	Mon	월
month	january	1 월
mon	jan	1
day	monday	월요일
Dy	mon	월
AM	AM	오전
Am	Am	오전
am	am	오전
A.M.	A.M.	오전
A.m.	A.m.	오전
a.m.	a.m.	오전
PM	AM	오전
Pm	Am	오전
pm	am	오전
P.M.	A.M.	오전

P.m.	A.m.	오전
p.m.	a.m	오전

The Number of Digits Format

Format Element	Number of Digits
MONTH (Month, month)	9 (ko_KR : 4)
MON (Mon, mon)	3 (ko_KR : 2)
DAY (Day, day)	9 (ko_KR : 6)
DY (Dy, dy)	3 (ko_KR : 2)
HH12, HH24	2
"text"	The length of the text
Other formats	Same as the length of the format

Example

```
--creating a table having date/time type columns
CREATE TABLE datetime_tbl(a TIME, b DATE, c TIMESTAMP, d DATETIME);
INSERT INTO datetime_tbl VALUES(SYSTIME, SYSDATE, SYSTIMESTAMP, SYSDATETIME);

--selecting a VARCHAR type string from the data in the specified format
SELECT TO_CHAR(b, 'DD, DY , MON, YYYY') FROM datetime_tbl;
to_char(b, 'DD, DY , MON, YYYY', 'en_US')
=====
'04, THU , FEB, 2010'

SELECT TO_CHAR(c, 'HH24:MI, DD, MONTH, YYYY') FROM datetime_tbl;
to char(c, 'HH24:MI, DD, MONTH, YYYY', 'en US')
=====
'16:50, 04, FEBRUARY , 2010'

SELECT TO_CHAR(c, 'HH24:MI:FF, DD, MONTH, YYYY') FROM datetime_tbl;
ERROR: Invalid format.

SELECT TO_CHAR(d, 'HH12:MI:SS:FF pm, YYYY-MM-DD-DAY') FROM datetime_tbl;
to char(d, 'HH12:MI:SS:FF pm, YYYY-MM-DD-DAY', 'en US')
=====
'04:50:11:624 pm, 2010-02-04-THURSDAY '

SELECT TO_CHAR(TIMESTAMP'2009-10-04 22:23:00', 'Day Month yyyy');
to char(timestamp '2009-10-04 22:23:00', 'Day Month yyyy', 'en US')
=====
'Sunday October 2009'
```

TO_CHAR Function (number)

Description

The **TO_CHAR** function converts a [Number Format](#) or numeric data type to a character string according to the number format and returns it. The type of the return value is **VARCHAR**. If the number format has not been specified as an argument, all significant figures are converted to a character string according to the default format.

Syntax

```
TO_CHAR(number argument[, format argument ])
```

number_argument :

- numeric(decimal)
- integer
- smallint
- bigint

- *float* (*real*)
- *double*
- **NULL**

format argument :

- character strings (see [Number Format](#))
- **NULL**

- *number_argument* : Specifies an expression that returns numeric data type string. If the input value is **NULL**, **NULL** is returned. If the input value is character type, the character itself is returned.
- *format_argument* : Specifies a format of return value. If format is not specified, all significant figures are returned as character string by default. If the value is **NULL**, **NULL** is returned.

Number Format

Format Element	Example	Description
9	9999	The number of 9's represents the number of significant figures to be returned. If the number of significant figures specified in the format is not sufficient, only the decimal part is rounded. If it is less than the number of digits in an integer, # is outputted. If the number of significant figures specified in the format is sufficient, the part preceding the integer part is filled with space characters and the decimal part is filled with 0.
0	0999	If the number of significant figures specified in the format is sufficient, the part preceding the integer part is filled with 0, not space characters before the value is returned.
S	S9999	Outputs the negative/positive sign in the specified position. These signs can be used only at the beginning of character string.
C	C9999	Returns the ISO currency code at the specified position.
,(comma)	9,999	Returns a comma (",") at the specified position. Multiple commas are allowed in the format.
.(percimal point)	9.999	Outputs the decimal point (".") that distinguishes the integer and the decimal part at a specified position. Only one decimal point is allowed in the format.
EEEE	9.99EEEE	Returns a scientific notation number.

Example

```
--selecting a string casted from a number in the specified format
SELECT TO_CHAR(12345,'S999999'), TO_CHAR(12345,'S099999');

=====
' +12345'          '+012345'

SELECT TO_CHAR(1234567,'C9,999,999,999');
=====
'    $1,234,567'

SELECT TO_CHAR(123.4567,'99'), TO_CHAR(123.4567,'999.99999'),
TO_CHAR(123.4567,'99999.999');
to_char(123.4567,'99','en_US')  to_char(123.4567,'999.99999',
'en_US')  to_char(123.4567,'99999.999','en_US')
=====
'###'              '123.45670'          ' 123.457'

SELECT TO_CHAR(1.234567,'99.999EEEE'), TO_CHAR(1.234567E-4);
to_char(1.234567,'99.999EEEE','en_US')  to_char(1.234567E-4)
=====
'1.235E+00'        '0.0001234567'
```

TO_DATE Function

Description

The **TO_DATE** function interprets a character string based on the date format given as an argument, converts it to a **DATE** type value, and returns it. For the format, see [TO_CHAR Function \(date_time\)](#). If a format is not specified, the "MM/DD/YYYY" format is applied by default.

Syntax

```
TO_DATE(string_argument[,format_argument[,date_lang_string_literal]])
```

string_argument :

- character strings
- **NULL**

format_argument :

- character strings (see [Date/Time Format 1](#))
- **NULL**

date_lang_string_literal : (see [date lang string literal](#))

- 'en_US'
- 'ko_KR'

- *string_argument* : Specifies an expression that returns character string. If the value is **NULL**, **NULL** is returned.
- *format_argument* : Specifies a format of return value to be converted as **DATE** type. See the "Default Date-Time Format" table of [TO_CHAR Function \(date_time\)](#). If the value is **NULL**, **NULL** is returned.
- *date_lang_string_literal* : Specifies the language for the input value to be applied. You can modify the value by using the **CUBRID_DATE_LANG** environment.

Example

```
--selecting a date type value casted from a string in the specified format
```

```
SELECT TO_DATE('12/25/2008');
to date('12/25/2008')
=====
12/25/2008

SELECT TO_DATE('25/12/2008', 'DD/MM/YYYY');
to date('25/12/2008', 'DD/MM/YYYY', 'en_US')
=====
12/25/2008

SELECT TO_DATE('081225', 'YMMDD');
to date('081225', 'YMMDD', 'en_US')
=====
12/25/2008

SELECT TO_DATE('2008-12-25', 'YYYY-MM-DD');
to date('2008-12-25', 'YYYY-MM-DD', 'en_US')
=====
12/25/2008
```

TO_DATETIME Function

Description

The **TO_DATETIME** function interprets a character string based on the date-time format given as an argument, converts it to a **DATETIME** type value, and returns it. For the format, see [TO_CHAR Function \(date_time\)](#). If format is not specified, the "HH:MI:SS.FF [am|pm] MM/DD/YYYY" format is applied by default.

Syntax

```
TO_DATETIME(string_argument[,format_argument[,date_lang_string_literal]])
```

string_argument :

- *character strings*
- **NULL**

format_argument :

- *character strings* (see the table [Date/Time Format 1](#))
- **NULL**

date_lang_string_literal : (see the table [Example of date lang string literal](#))

- 'en_US'
- 'ko_KR'

- *string_argument* : Specifies an expression that returns character string. If the value is **NULL**, **NULL** is returned.
- *format_argument* : Specifies a format of return value to be converted as **DATETIME** type. See the "Default Date-Time Format" table of [TO_CHAR Function \(date_time\)](#). If the value is **NULL**, **NULL** is returned.
- *date_lang_string_literal* : Specifies the language for the input value to be applied. You can modify the value by using the **CUBRID_DATE_LANG** environment.

Example

```
--selecting a datetime type value casted from a string in the specified format

SELECT TO_DATETIME('13:10:30 12/25/2008');
to_datetime('13:10:30 12/25/2008')
=====
01:10:30.000 PM 12/25/2008

SELECT TO_DATETIME('08-Dec-25 13:10:30.999', 'YY-Mon-DD HH24:MI:SS.FF');
to_datetime('08-Dec-25 13:10:30.999', 'YY-Mon-DD HH24:MI:SS.FF', 'en_US')
=====
01:10:30.999 PM 12/25/2008

SELECT TO_DATETIME('DATE: 12-25-2008 TIME: 13:10:30.999', '"DATE:" MM-DD-YYYY "TIME:"
HH24:MI:SS.FF');
to_datetime('DATE: 12-25-2008 TIME: 13:10:30.999', '"DATE:" MM-DD-YYYY "TIME:"
HH24:MI:SS.FF', 'en_US')
=====
01:10:30.999 PM 12/25/2008
```

TO_NUMBER Function

Description

The **TO_NUMBER** function interprets a character string based on the number format given as an argument, converts it to a **NUMERIC** type value, and returns it. If the number format is not specified, returns all significant figures that are included in the character string as **NUMERIC** type numbers by default.

Syntax

```
TO_NUMBER(string_argument [, format_argument ])
```

string_argument :

- *character strings*
- **NULL**

format_argument :

- *character strings*
- **NULL**

- *string_argument* : Specifies an expression that returns character string. If the value is **NULL**, **NULL** is returned.
- *format_argument* : Specifies a format of return value to be converted as **NUMBER** type. See the "Number Format" table of [TO_CHAR Function \(number\)](#). If the value is **NULL**, an error is returned.

Example

```
--selecting a number casted from a string in the specified format
SELECT TO_NUMBER('-1234');
to_number('-1234')
=====
-1234
```

```

SELECT TO_NUMBER('12345', '999999');
to_number('12345', '999999')
=====
12345

SELECT TO_NUMBER('$12,345.67', 'C99,999.999');
to_number('$12,345.67', 'C99,999.999')
=====
12345.670

SELECT TO_NUMBER('12345.67', '99999.999');
to_number('12345.67', '99999.999')
=====
12345.670

```

TO_TIME Function

Description

The **TO_TIME** function interprets a character string based on the time format given as an argument, converts it to a **TIME** type value, and returns it. For the format, see [TO_CHAR Function \(date_time\)](#). If a format is not specified, the "HH:MI:SS" format is applied by default.

Syntax

```

TO_TIME(string argument[,format argument [,date lang string literal]]):

string_argument :
• character strings
• NULL

format_argument :
• character strings (see Date/Time Format 1)
• NULL

date lang string literal : (see date lang string literal)
• 'en_US'
• 'ko_KR'

```

- *string_argument* : Specifies an expression that returns character string. If the value is **NULL**, **NULL** is returned.
- *format_argument* : Specifies a format of return value to be converted as **TIME** type. See the "Default Date-Time Format" table of [TO_CHAR Function \(date_time\)](#). If the value is **NULL**, **NULL** is returned.
- *date_lang_string_literal* : Specifies the language for the input value to be applied. You can modify the value by using the **CUBRID_DATE_LANG** environment.

Example

```

--selecting a time type value casted from a string in the specified format

SELECT TO TIME ('13:10:30');
to time('13:10:30')=====
01:10:30 PM

SELECT TO TIME('HOUR: 13 MINUTE: 10 SECOND: 30', '"HOUR:" HH24 "MINUTE:" MI "SECOND:" SS');
to time('HOUR: 13 MINUTE: 10 SECOND: 30', '"HOUR:" HH24 "MINUTE:" MI "SECOND:" SS',
'en_US')=====
01:10:30 PM

SELECT TO TIME ('13:10:30', 'HH24:MI:SS');
to time('13:10:30', 'HH24:MI:SS', 'en_US')
=====
01:10:30 PM

SELECT TO_TIME ('13:10:30', 'HH12:MI:SS');

```

```
ERROR: Conversion error in date format.
```

TO_TIMESTAMP Function

Description

The **TO_TIMESTAMP** function interprets a character string based on the time format given as an argument, converts it to a **TIMESTAMP** type value, and returns it. For the format, see [TO_CHAR Function \(date_time\)](#). If a format is not specified, the "HH:MI[:SS] [am|pm] MM/DD/YYYY" format is applied by default.

Syntax

```
TO_TIMESTAMP(string argument[, format argument[, date lang string literal]])
```

string argument :

- character strings
- **NULL**

format argument :

- character strings (see [Date/Time Format 1](#) table)
- **NULL**

date lang string literal : (see [date lang string literal](#) table)

- 'en_US'
- 'ko_KR'

- *string_argument* : Specifies an expression that returns character string. If the value is **NULL**, **NULL** is returned.
- *format_argument* : Specifies a format of return value to be converted as **TIMESTAMP** type. See the "Default Date-Time Format" table of [TO_CHAR Function \(date_time\)](#). If the value is **NULL**, **NULL** is returned.
- *date_lang_string_literal* : Specifies the language for the input value to be applied. You can modify the value by using the **CUBRID_DATE_LANG** environment.

Example

```
--selecting a timestamp type value casted from a string in the specified format
```

```
SELECT TO_TIMESTAMP('13:10:30 12/25/2008');
to_timestamp('13:10:30 12/25/2008')
```

```
=====
01:10:30 PM 12/25/2008
```

```
SELECT TO_TIMESTAMP('08-Dec-25 13:10:30', 'YY-Mon-DD HH24:MI:SS');
to_timestamp('08-Dec-25 13:10:30', 'YY-Mon-DD HH24:MI:SS', 'en_US')
```

```
=====
01:10:30 PM 12/25/2008
```

```
SELECT TO_TIMESTAMP('YEAR: 2008 DATE: 12-25 TIME: 13:10:30', '"YEAR:" YYYY "DATE:" MM-DD
"TIME:" HH24:MI:SS');
to_timestamp('YEAR: 2008 DATE: 12-25 TIME: 13:10:30', '"YEAR:" YYYY "DATE:" MM-DD "TIME:"
HH24:MI:SS', 'en_US')
```

```
=====
01:10:30 PM 12/25/2008
```

Aggregate Functions

AVG Function

Description

The **AVG** function calculates the arithmetic average of the value of an expression representing all rows. Only one *expression* is specified as a parameter. You can get the average without duplicates by using the **DISTINCT** or **UNIQUE** keyword in front of the expression or the average of all values by omitting the keyword or by using **ALL**.

Syntax

```
AVG ( [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

- *expression* : Specifies an expression that returns a numeric value. A collection expression cannot be specified.
- **ALL** : Calculates an average value for all data (default).
- **DISTINCT** or **UNIQUE** : Calculates an average value without duplicates.

Example

The following example shows how to retrieve the average number of gold medals that Korea won in Olympics (demodb).

```
SELECT AVG(gold)
FROM participant
WHERE nation_code = 'KOR';
```

Result value : 9

COUNT Function

Description

The **COUNT** function returns the number of rows returned by a query. If an asterisk (*) is specified, the number of all rows satisfying the condition (including the rows with the **NULL** value) is returned. If the **DISTINCT** or **UNIQUE** keyword is specified in front of the expression, only the number of rows that have a unique value (excluding the rows with the **NULL** value) is returned after duplicates have been removed. Therefore, the value returned is always an integer and **NULL** is never returned.

A column that has collection type and object domain (user-defined class or multimedia class) can also be specified in the *expression*.

Syntax

```
COUNT ( * | [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

- *expression* : Specifies an expression.
- **ALL** : Gets the number of rows given in the *expression* (default).
- **DISTINCT** or **UNIQUE** : Gets the number of rows without duplicates.

Example

The following example shows how to retrieve the number of Olympic Games that have a mascot (demodb).

```
SELECT COUNT(*)
FROM olympic
WHERE mascot IS NOT NULL;
```

Result value : 9

GROUP_CONCAT Function

Description

The **GROUP_CONCAT** function connects the values that are not **NULL** in the group and returns the character string in the **VARCHAR** type. If there are no rows of query result or there are only **NULL** values, **NULL** will be returned.

The maximum size of the return value follows the configuration of the system parameter, **group_concat_max_len**. The default is **1024** bytes, the minimum value is 4 bytes and the maximum value is 33,554,432 bytes. If it exceeds the maximum value, **NULL** will be returned.

To remove the duplicate values, use the **DISTINCT** clause. The default separator for the group result values is comma (.). To represent the separator explicitly, add the character string to use as a separator in the **SEPARATOR** clause and after that. If you want to remove separators, enter empty strings after the **SEPARATOR** clause.

If the non-character string type is passed to the result character string, an error will be returned.

To use the **GROUP_CONCAT** function, you must meet the following conditions.

- Only one expression (or a column) is allowed for an input parameter.
- Sorting with **ORDER BY** is available only in the the expression used as a parameter.
- The character string used as a separator allows not only character string type but also allows other types.

Syntax

```
GROUP_CONCAT([DISTINCT] {col | expression}
             [ORDER BY {col | unsigned_int} [ASC | DESC]]
             [SEPARATOR str_val])
```

- *expression* : Operation returning numerical values or character strings
- *str_val* : Character string to use as a separator
- **DISTINCT** : Removes duplicate values from the result.
- **ORDER BY** : Specifies the order of result values.
- **SEPARATOR** : Specifies the separator to divide the result values. If it is omitted, the default character, comma (,) will be used as a separator.

Example

```
SELECT GROUP_CONCAT(s name) FROM code;
      group_concat(s name)
=====
      'X,W,M,B,S,G'

SELECT GROUP_CONCAT(s name ORDER BY s name SEPARATOR ':') from code;
      group_concat(s name order by s name separator ':')
=====
      'B:G:M:S:W:X'

CREATE TABLE t(i int);
INSERT INTO t VALUES (4), (2), (3), (6), (1), (5);

SELECT GROUP_CONCAT(i*2+1 ORDER BY 1 SEPARATOR '') FROM t;
      group_concat(i*2+1 order by 1 separator '')
=====
      '35791113'
```

MAX Function

Description

The **MAX** function gets the greatest value of expressions of all rows. Only one *expression* is specified.

For expressions that return character strings, the string that appears later in alphabetical order becomes the maximum value; for those that return numbers, the greatest value becomes the maximum value.

Syntax

```
MAX ( [ [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

- *expression* : Specifies an expression that returns a numeric or string value. A collection expression cannot be specified.
- **ALL** : Gets the maximum value for all data (default).
- **DISTINCT** or **UNIQUE** : Gets the maximum value without duplicates.

Example

The following example shows how to retrieve the maximum number of gold medals that Korea won in the Olympics (demodb).

```
SELECT MAX(gold) FROM participant WHERE nation code = 'KOR';
      max(gold)
=====
      12
```

MIN Function

Description

The **MIN** function gets the smallest value of expressions of all rows. Only one *expression* is specified.

For expressions that return character strings, the string that appears earlier in alphabetical order becomes the minimum value; for those that return numbers, the smallest value becomes the minimum value.

Syntax

```
MIN ( [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

- *expression* : Specifies an expression that returns a numeric or string value. A collection expression cannot be specified.
- **ALL** : Gets the minimum value for all data (default).
- **DISTINCT** or **UNIQUE** : Gets the maximum value without duplicates.

Example

The following example shows how to retrieve the minimum number of gold medals that Korea won in the Olympics (demodb).

```
SELECT MIN(gold) FROM participant WHERE nation code = 'KOR';
      min(gold)
=====
              7
```

STDDEV/STDDEV_POP Functions

Description

The **STDDEV** function returns a standard deviation of the expression values of all rows. Only one *expression* is specified as a parameter. You can get the standard deviation without duplicates by inserting the **DISTINCT** or **UNIQUE** keyword in front of the expression, or get the standard deviation of all values by omitting the keyword or by using **ALL**.

The return value may be different from the actual evaluation value because it follows the type of the expression specified as a parameter.

Syntaxs

```
STDDEV( [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL] expression )
```

- *expression* : Specifies an expression that returns a numeric value.
- **ALL** : Calculates the standard deviation for all data (default).
- **DISTINCT** or **UNIQUE** : Calculates the standard deviation without duplicates.

Example

```
CREATE TABLE test_table (d DOUBLE);
INSERT INTO test_table VALUES(78), (63.65), (230.54), (32), (17.2), (195.7689), (57.57);

SELECT STDDEV POP(d) FROM test table;
      stddev_pop(d)
=====
      7.672456168942171e+01

SELECT STDDEV POP(POWER(d,2)+d*2+1) FROM test table;
      stddev_pop( power(d, 2)+d*2+1)
=====
      1.995964904967644e+04

TRUNCATE TABLE test table;
SELECT STDDEV POP(d) FROM test table;
      stddev_pop(d)
```



```
=====
NULL
```

STDDEV_SAMP Function

Description

The **STDDEV_SAMP** function calculates the sample standard deviation. Only one *expression* is specified as a parameter. If the **DISTINCT** or **UNIQUE** keyword is included, it calculates the sample standard deviation after deleting the duplicates; if the keyword is omitted or is **ALL**, it calculates the sample standard deviation for all values.

The return value is the same as the square root of the [VAR_SAMP Function](#) return value and it is a **DOUBLE** type. If there are no rows that can be used for calculating a result, **NULL** will be returned.

The following are the formulas applied to the function.

$$\text{STDDEV_SAMP} = [\{ 1 / (N-1) \} * \text{SUM}(\{ x_i - \text{mean}(x) \}^2)]^{1/2}$$

- **SUM** : Sum
- **mean** : Average

Syntax

```
STDDEV_SAMP ( [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

- *expression* : An expression that returns a numeric value
- **ALL** : Used to calculate the standard deviation for all values. It is the default value.
- **DISTINCT** or **UNIQUE** : Used used to calculate the standard deviation for the unique values without duplicates.

Example

```
CREATE TABLE test_table (d DOUBLE);
INSERT INTO test table VALUES(78), (63.65), (230.54), (32), (17.2), (195.7689), (57.57);

SELECT STDDEV SAMP(d) FROM test table;
          stddev_samp(d)
=====
          8.287199825135663e+01

SELECT STDDEV SAMP(POWER(d,2)+d*2+1) FROM test table;
          stddev_samp( power(d, 2)+d*2+1)
=====
          2.155888498702931e+04

TRUNCATE TABLE test table;
SELECT STDDEV SAMP(d) FROM test table;
          stddev_samp(d)
=====
          NULL
```

SUM Function

Description

The **SUM** function returns the sum of expressions of all rows. Only one *expression* is specified as a parameter. You can get the sum without duplicates by inserting the **DISTINCT** or **UNIQUE** keyword in front of the expression, or get the sum of all values by omitting the keyword or by using **ALL**.

Syntax

```
SUM ( [ { DISTINCT | DISTINCTROW } | UNIQUE | ALL ] expression )
```

You can specify a single-value expression as an input for **SUM** function.

- *expression* : Specifies an expression that returns a numeric value.
- **ALL** : Gets the sum for all data (default).
- **DISTINCT** or **UNIQUE** : Gets the sum of unique values without duplicates

Example

The following example shows how to retrieve the top 10 countries and the total number of gold medals based on the sum of gold medals won in the Olympics (demodb).

```
SELECT nation_code, SUM(gold) FROM participant GROUP BY nation_code
ORDER BY SUM(gold) DESC
FOR ORDERBY NUM() BETWEEN 1 AND 10 ;
```

=== <Result of SELECT Command in Line 1> ===

nation_code	sum(gold)
'USA'	190
'CHN'	97
'RUS'	85
'GER'	79
'URS'	55
'FRA'	53
'AUS'	52
'ITA'	48
'KOR'	48
'EUN'	45

VAR_POP/VARIANCE Functions

Description

The **VARIANCE** function returns a variance of expression values of all rows. Only one *expression* is specified as a parameter. You can get the variance without duplicates by using the **DISTINCT** or **UNIQUE** keyword in front of the expression or the variance of all values by omitting the keyword or by using **ALL**.

The return value may be different from the actual evaluation value because it follows the type of the expression specified as a parameter.

The following is a formula that is applied to the function.

$$\frac{[\text{SUM}(x^2) - \frac{(\text{SUM}(x))^2}{n}]}{(n-1)}$$

Syntax

```
VARIANCE( [DISTINCT | UNIQUE | ALL] expression )
```

- *expression* : Specifies an expression that returns a numeric value.
- **ALL** : Gets the variance for all values (default).
- **DISTINCT** or **UNIQUE** : Gets the variance of unique values without duplicates.

Example

```
CREATE TABLE test table (d double);
INSERT INTO test table VALUES(78), (63.65), (230.54), (32), (17.2), (195.7689), (57.57);
SELECT VAR POP(d) FROM test table;
           var pop(d)
=====
           5.886658366433878e+003

SELECT VAR POP(POWER(d,2)+d*2+1) FROM test table;
           var pop( power(d, 2)+d*2+1)
=====
           3.983875901862495e+008

TRUNCATE TABLE test table;
SELECT VAR POP(d) FROM test table;
           var pop(d)
=====
           NULL
```

VAR_SAMP Function

Description

The **VAR_SAMP** function returns the sample variance. The denominator is the number of all rows - 1. Only one *expression* is specified as a parameter. If the **DISTINCT** or **UNIQUE** keyword is included, it calculates the sample variance after deleting the duplicates, and if the keyword is omitted or is **ALL**, it calculates the sample variance for all values.

The return value is a **DOUBLE** type. If there are no rows that can be used for calculating a result, **NULL** will be returned.

The following are the formulas applied to the function.

$$\text{VAR_POP} = \{ 1 / (N-1) \} * \text{SUM}(\{ x_i - \text{AVG}(x) \}^2)$$

- **SUM** : Sum
- **AVG** : Average

```
VAR_SAMP( [ DISTINCT | UNIQUE | ALL] expression )
```

- *expression* : Specifies one expression to return the numeric.
- **ALL** : Is used to calculate the sample variance of unique values without duplicates. It is the default value.
- **DISTINCT** or **UNIQUE** : Is used to calculate the sample variance for the unique values without duplicates.

Example

```
CREATE TABLE test table (d double);
INSERT INTO test table VALUES(78), (63.65), (230.54), (32), (17.2), (195.7689), (57.57);
SELECT VAR SAMP(d) FROM test table;
           var_samp(d)
=====
           6.867768094172856e+03

SELECT VAR SAMP(POWER(d,2)+d*2+1) FROM test table;
           var_samp( power(d, 2)+d*2+1)
=====
           4.647855218839577e+08

TRUNCATE TABLE test table;
SELECT VAR SAMP(d) FROM test table;
           var_samp(d)
=====
           NULL
```

Click Counter Functions

INCR/DECR Functions

Description

The **INCR** function increments the column's value given as a parameter for a [SELECT](#) statement by 1. The **DECR** function decrements the value of the column by 1.

Syntax

```
SELECT [ qualifier ] select_expression
[ { TO | INTO } variable [ {, variable }...; ]
...;
select_expression :
*





```

The **INCR** and **DECR** functions are called "click counters" and can be effectively used to increase the number of post views for a Bulletin Board System (BBS) type of web service. In a scenario where you want to [SELECT](#) a post and immediately increase the number of views by 1 using an **UPDATE** statement, you can view the post and increment the number at once by using the **INCR** function in a single **SELECT** statement.

The **INCR** function increments the column value specified as an argument. Only integer type numbers can be used as arguments. If the value is **NULL**, the **INCR** function returns the **NULL**. That is, a value must be valid in order to be incremented by the **INCR** function. The **DECR** function decrements the column value specified as a parameter.

If an **INCR** function is specified in the [SELECT](#) statement, the **COUNTER** value is incremented by 1 and the query result is displayed with the values before the increment. Furthermore, the **INCR** function does not increment the value of the tuple affected by the query process but rather the one affected by the final result.

Remark

- The **INCR/DECR** function executes independent of user-defined transactions and is applied automatically to the database by the top operation internally used in the system, apart from the transaction's **COMMIT/ROLLBACK**.
- When multiple **INCR/DECR** functions are specified in a single [SELECT](#) statement, the failure of any of the **INCR/DECR** functions leads to the failure of all of them.
- The **INCR/DECR** functions apply only to top-level [SELECT](#) statements. **SUB SELECT** statements such as **INSERT ... SELECT ...** statement and **UPDATE** table **SET** col = **SELECT ...** statement are not supported. The following example shows where the **INCR** function is not allowed.

```
SELECT b.content, INCR(b.read_count) FROM (SELECT * FROM board WHERE id = 1) AS b
```

- If the **SELECT** statement with **INCR/DECR** function(s) returns more than one row as a result, it is treated as an error. The final result must have only one row to be considered valid.
- The **INCR/DECR** function can be used only in numerical domains. Applicable domains are limited to integer data types such as **SMALLINT** and **INTEGER**. They cannot be used in other domains.
- When the **INCR** function is called, the value to be returned will be the current value, while the value to be stored will be the current value + 1. Execute the following statement to select the value to be stored as the result :

```
SELECT content, INCR(read_count) + 1 FROM board WHERE id = 1;
```

- If the defined maximum value of the domain is exceeded, the **INCR** function initializes the column value to 0. Likewise, the column value is also initialized to 0 when the **DECR** function applies to the minimum value.
- Data inconsistency can occur because the **INCR/DECR** functions are executed regardless of **UPDATE** trigger. The following example shows the database inconsistency in that situation.

```
CREATE TRIGGER event_tr BEFORE UPDATE ON event EXECUTE REJECT;
SELECT INCR(players) FROM event WHERE gender='M';
```

- The **INCR/DECR** functions returns an error in the write-protected broker mode such as slave mode of HA configuration, **CSQL Interpreter (csql -r)** of read-only, **Read Only**, **Slave Only** or **Preferred Host Read Only** mode.

Example

Suppose that the following three rows of data are inserted into the 'board' table.

```
CREATE TABLE board (
id INT, title VARCHAR(100), content VARCHAR(4000), read count INT );
INSERT INTO board VALUES (1, 'aaa', 'text...', 0);
INSERT INTO board VALUES (2, 'bbb', 'text...', 0);
INSERT INTO board VALUES (3, 'ccc', 'text...', 0);
```

The following example shows how to increment the value of the 'read_count' column in a data whose 'id' value is 1 by using the **INCR** function.

```
SELECT content, INCR(read count) FROM board WHERE id = 1;
content          read_count
=====
'text...'                0
```

In the example, the column value becomes read_count + 1 as a result of the **INCR** function in the **SELECT** statement. You can check the result using the following **SELECT** statement.

```
SELECT content, read count FROM board WHERE id = 1;
content          read_count
```

```
=====
'text...' 1
```

ROWNUM Functions

ROWNUM/INST_NUM()

Description

The **ROWNUM** function returns the number representing the order of the records that will be generated by the query result. The first result record is assigned 1, and the second result record is assigned 2.

ROWNUM and **INST_NUM()** can be used in the **SELECT** statement, and **GROUPBY_NUM()** can be used in the **SELECT** statement with **GROUP BY** clauses. The **ROWNUM** function can be used to limit the number of result records of the query in several ways. For example, it can be used to search only the first 10 records or to return even or odd number records.

The **ROWNUM** function has a result value as an integer, and can be used wherever an expression is valid such as the **SELECT** or **WHERE** clause. However, it is not allowed to compare the result of the **ROWNUM** function with the attribute or the correlated subquery.

Syntax

```
INST_NUM()
ROWNUM
```

Remark

- The **ROWNUM** function specified in the **WHERE** clause works the same as the **INST_NUM()** function. Whereas **INST_NUM()** is a scalar function, **GROUPBY_NUM()** is a kind of an aggregate function. In a **SELECT** statement with a **GROUP BY** clause, **GROUPBY_NUM()** must be used instead of **INST_NUM()**.
- The **ROWNUM** function belongs to each **SELECT** statement. That is, if a **ROWNUM** function is used in a subquery, it returns the sequence of the subquery result while it is being executed. Internally, the result of the **ROWNUM** function is generated right before the searched record is written to the query result set. At this moment, the counter value that generates the serial number of the result set records increases.
- If an **ORDER BY** clause is included in the **SELECT** statement, the value of the **ROWNUM** function specified in the **WHERE** clause is generated before sorting for the **ORDER BY** clause. If a **GROUP BY** clause is included in the **SELECT** statement, the value of the **GROUPBY_NUM()** function specified in the **HAVING** clause is calculated after the query results are grouped. After the sorting process is completed using the **ORDER BY** clause, you need to use the **ORDERBY_NUM()** function in the **ORDER BY** clause in order to get a sequence of the result records.
- The **ROWNUM** function can also be used in SQL statements such as **INSERT**, **DELETE** and **UPDATE** in addition to the **SELECT** statement. For example, as in the query **INSERT INTO table_name SELECT ... FROM ... WHERE ...**, you can search for part of the row from one table and then insert it into another by using the **ROWNUM** function in the **WHERE** clause.

Example

The following example shows how to retrieve country names ranked first to fourth based on the number of gold medals in the 1988 Olympics (demodb).

```
--Limiting 4 rows using ROWNUM in the WHERE condition
SELECT * FROM
(SELECT nation_code FROM participant WHERE host_year = 1988
 ORDER BY gold DESC) AS T
WHERE ROWNUM <5;
 nation code
=====
'URS'
'GDR'
'USA'
'KOR'

--Limiting 4 rows using FOR ORDERBY_NUM()
```

```

SELECT ROWNUM, nation code FROM participant WHERE host year = 1988
ORDER BY gold DESC
FOR ORDERBY_NUM() < 5;
      rownum  nation_code
=====
          156  'URS'
          155  'GDR'
          154  'USA'
          153  'KOR'

--Unexpected results : ROWNUM operated before ORDER BY
SELECT ROWNUM, nation code FROM participant
WHERE host year = 1988 AND ROWNUM < 5
ORDER BY gold DESC;
      rownum  nation_code
=====
          1   'ZIM'
          2   'ZAM'
          3   'ZAI'
          4   'YMD'

```

GROUPBY_NUM() Function

Description

The **GROUPBY_NUM()** function is used with the **ROWNUM()** or **INST_NUM()** function to limit the number of result rows. The difference is that the **GROUPBY_NUM()** function is combined after the **GROUP BY ... HAVING** clause to give order to a result that has been already sorted. In addition, while the **INST_NUM()** function is a scalar function, the **GROUPBY_NUM()** function is kind of an aggregate function.

That is, when retrieving only some of the result rows by using **ROWNUM** in a condition clause of the **SELECT** statement that includes the **GROUP BY** clause, **ROWNUM** is applied first and then group sorting by **GROUP BY** is performed. On the other hand, when retrieving only some of the result rows by using the **GROUPBY_NUM()** function, **ROWNUM** is applied to the result of group sorting by **GROUP BY**.

Syntax

```
GROUPBY_NUM()
```

Example

The following example shows how to retrieve the fastest record in the previous five Olympic Games from the history table (demodb).

```

--Group-ordering first and then limiting rows using GROUPBY NUM()
SELECT host year, MIN(score) FROM history
GROUP BY host_year HAVING GROUPBY_NUM() BETWEEN 1 AND 5;
      host_year  min(score)
=====
          1968  '8.9'
          1980  '01:53.0'
          1984  '13:06.0'
          1988  '01:58.0'
          1992  '02:07.0'

--Limiting rows first and then Group-ordering using ROWNUM
SELECT host year, MIN(score) FROM history
WHERE ROWNUM BETWEEN 1 AND 5 GROUP BY host_year;
      host_year  min(score)
=====
          2000  '03:41.0'
          2004  '01:45.0'

```

ORDERBY_NUM() Function

Description

The **ORDERBY_NUM()** function is used with the **ROWNUM()** or **INST_NUM()** function to limit the number of result rows. The difference is that the **ORDERBY_NUM()** function is combined after the **ORDER BY** clause to give order to a result that has been already sorted.

That is, when retrieving only some of the result rows by using **ROWNUM** in a condition clause of the **SELECT** statement that includes the **ORDER BY** clause, **ROWNUM** is applied first and then group sorting by **ORDER BY** is performed. On the other hand, when retrieving only some of the result rows by using the **ORDER_NUM()** function, **ROWNUM** is applied to the result of sorting by **ORDER BY**.

Syntax

```
FOR ORDERBY_NUM()
```

Example

The following example shows how to retrieve athlete names ranked 3rd to 5th and their records in the history table (demodb).

```
--Ordering first and then limiting rows using FOR ORDERBY_NUM()
SELECT athlete, score FROM history
ORDER BY score FOR ORDERBY NUM() BETWEEN 3 AND 5;
  athlete          score
=====
  'Luo Xuejuan'    '01:07.0'
  'Rodal Vebjorn'  '01:43.0'
  'Thorpe Ian'     '01:45.0'

--Limiting rows first and then Ordering using ROWNUM
SELECT athlete, score FROM history
WHERE ROWNUM BETWEEN 3 AND 5 ORDER BY score;
  athlete          score
=====
  'Thorpe Ian'     '01:45.0'
  'Thorpe Ian'     '03:41.0'
  'Hackett Grant'  '14:43.0'
```

Information Functions

CURRENT_USER/USER

Description

CURRENT_USER and **USER** are used interchangeably. They return the user name that is currently logged in to the database as a string.

USER() and **SYSTEM_USER()** are used interchangeably. They return the user name with a host name.

Syntax

```
CURRENT_USER
USER
```

Example

```
--selecting the current user on the session
SELECT USER;
  CURRENT_USER
=====
  'PUBLIC'

SELECT USER(), CURRENT_USER;
  user()          CURRENT_USER
=====
  'PUBLIC'
```

```

=====
'PUBLIC@cubs006.cub' 'PUBLIC'

--selecting all users of the current database from the system table
SELECT name, id, password FROM db user;
name                id password
=====
'DBA'                NULL NULL
'PUBLIC'            NULL NULL
'SELECT_ONLY_USER'  NULL db_password
'ALMOST DBA USER'   NULL db_password
'SELECT_ONLY_USER2' NULL NULL

```

DATABASE/SCHEMA

Description

The **DATABASE** and **SCHEMA** functions are used interchangeably. They return the name of currently-connected database as a **VARCHAR** type.

Syntax

```

DATABASE ()
SCHEMA ()

```

Example

```

SELECT DATABASE(), SCHEMA();
database()      schema()
=====
'demodb'        'demodb'

```

DEFAULT Function

Description

The **DEFAULT** function returns a default value defined for a column. If a default value has not been specified for the given column, **NULL** or an error is returned. If any of constraints is not defined or the **UNIQUE** constraint is defined for the column where a default value is not defined, **NULL** is returned. If **NOT NULL** or **PRIMARY KEY** constraint is defined, an error is returned.

Syntax

```

DEFAULT (column_name)

```

Example

```

CREATE TABLE info_tbl(id INT DEFAULT 0, name VARCHAR)
INSERT INTO info_tbl VALUES (1,'a'), (2,'b'), (NULL,'c');

3 rows affected.

SELECT id, DEFAULT(id) FROM info_tbl;
id      default(id)
=====
1              0
2              0
NULL           0

```

INDEX_CARDINALITY Function

Description

The **INDEX_CARDINALITY** function returns the index cardinality in a table. The index cardinality is the number of unique values defining the index. The index cardinality can be applied even to the partial key of the multiple column

index and displays the number of the unique value for the partial key by specifying the column location with the third parameter.

The return value is 0 or a positive integer and if any of the input parameters is **NULL**, **NULL** is returned. If tables or indexes that are input parameters are not found, or *key_pos* is out of range, **NULL** is returned.

For the table and the index names which are the first and the second input parameters, they cannot be passed as **NCHAR** or **VARNCHAR** types.

Syntax

```
INDEX_CARDINALITY (table, index, key_pos)
```

- *table* : Table name
- *index* : Index name that exists in the *table*
- *key_pos* : Partial key location. It *key_pos* starts from 0 and has a range that is smaller than the number of columns consisting of keys; that is, the *key_pos* of the first column is 0. For the single column index, it is 0. It can be one of the following types.
 - Character string that can be converted to a numeric type. **NCHAR** and **VARNCHAR** are not supported.
 - Numeric type that can be converted to an integer type. The **FLOAT** or the **DOUBLE** types will be the value converted by the **ROUND** function.

Example

```
CREATE TABLE t1( i1 INTEGER ,
i2 INTEGER not null,
i3 INTEGER unique,
s1 VARCHAR(10),
s2 VARCHAR(10),
s3 VARCHAR(10) UNIQUE);

CREATE INDEX i t1 i1 ON t1(i1 DESC);
CREATE INDEX i t1 s1 ON t1(s1(7));
CREATE INDEX i t1 i1 s1 on t1(i1,s1);
CREATE UNIQUE INDEX i_t1_i2_s2 ON t1(i2,s2);

INSERT INTO t1 VALUES (1,1,1,'abc','abc','abc');
INSERT INTO t1 VALUES (2,2,2,'zabc','zabc','zabc');
INSERT INTO t1 VALUES (2,3,3,'+abc','+abc','+abc');

SELECT INDEX_CARDINALITY('t1','i_t1_i1_s1',0);
       index_cardinality('t1', 'i_t1_i1_s1', 0)
=====
                               2

SELECT INDEX_CARDINALITY('t1','i t1 i1 s1',1);
       index_cardinality('t1', 'i_t1_i1_s1', 1)
=====
                               3

SELECT INDEX_CARDINALITY('t1','i t1 i1 s1',2);
       index_cardinality('t1', 'i t1 i1 s1', 2)
=====
                               NULL

SELECT INDEX_CARDINALITY('t123','i t1 i1 s1',1);
       index_cardinality('t123', 'i t1 i1 s1', 1)
=====
                               NULL
```

LAST_INSERT_ID Function

Description

The **LAST_INSERT_ID** function returns the value created at the end of the **AUTO_INCREMENT** column of all tables.

If no values are **INSERT**ed successfully, the last successful value will be maintained, and the SQL statement on execution does not affect the **LAST_INSERT_ID()** value. If you enter multiple rows with one **INSERT** statement, the **LAST_INSERT_ID()** will return the input row value entered at the end. If the execution result of the previous SQL statement returns an error, the **LAST_INSERT_ID()** value is not defined, and the rollback cannot recover the **LAST_INSERT_ID()** value as the previous transaction value.

You cannot check the **LAST_INSERT_ID()** value used in the trigger, outside trigger.

The created ID is maintained independently for the connection of each client.

Syntax

LAST_INSERT_ID()

Example

```
CREATE TABLE ss (id INT AUTO_INCREMENT NOT NULL PRIMARY KEY, text VARCHAR(32));
INSERT into ss VALUES (NULL, 'cubrid');
SELECT LAST_INSERT_ID();

    last insert id()
=====
                    1

INSERT INTO ss VALUES (NULL, 'database'), (NULL, 'manager');
SELECT LAST_INSERT_ID();

    last_insert_id()
=====
                    3
```

Caution

If you insert multiple rows with a single **INSERT** statement, **LAST_INSERT_ID()** returns the first **AUTO_INCREMENT** value.

```
CREATE TABLE tbl (id INT AUTO_INCREMENT);
INSERT INTO tbl values (500), (NULL), (NULL);
SELECT LAST_INSERT_ID();

    last insert id()
=====
                    1

INSERT INTO tbl values (500), (NULL), (NULL);
SELECT LAST_INSERT_ID();

    last_insert_id()
=====
                    3

SELECT * FROM tbl;

                    id
=====
                    500
                     1
                     2
                    500
                     3
                     4
```

LIST_DBS Function

Description

The **LIST_DBS** function outputs the list of all databases in the CUBRID database server, separated by blanks.

Syntax

```
LIST_DBS()
```

Example

```
SELECT LIST_DBS();
      dbs
=====
'testdb demodb'
```

ROW_COUNT Function**Description**

The **ROW_COUNT** function returns the number of rows updated (**UPDATE**, **INSERT**, **DELETE**) by the previous statement. Note that the **ROW_COUNT** function execution area at the SQL level is limited to the client session in which the SQL was created. If this function is called after executing SQL with the **;run** or **;xrun** command, it returns -1.

Syntax

```
ROW_COUNT()
```

Example

```
CREATE TABLE rc (i int);
INSERT INTO rc VALUES (1), (2), (3), (4), (5), (6), (7);
SELECT ROW_COUNT();
      row_count()
=====
              7

UPDATE rc SET i = 0 WHERE i > 3;
SELECT ROW_COUNT();
      row_count()
=====
              4

DELETE FROM rc WHERE i = 0;
SELECT ROW_COUNT();
      row_count()
=====
              4
```

USER/SYSTEM_USER Functions**Description**

The **USER** function and the **SYSTEM_USER** function are identical and they return the user name together with the host name.

The [CURRENT_USER, USER](#) with a similar feature returns the user names who has logged on to the current database as character strings.

Syntax

```
USER()
SYSTEM_USER()
```

Example

```
--selecting the current user on the session
SELECT USER;
      CURRENT_USER
=====
      'PUBLIC'

SELECT USER(), CURRENT_USER;
```

```

user()                CURRENT USER
=====
'PUBLIC@cubs006.cub' 'PUBLIC'

--selecting all users of the current database from the system table
SELECT name, id, password FROM db user;
name                id password
=====
'DBA'                NULL NULL
'PUBLIC'            NULL NULL
'SELECT ONLY USER'  NULL db password
'ALMOST DBA USER'  NULL db password
'SELECT_ONLY_USER2' NULL NULL
    
```

VERSION Function

Description

The **VERSION** function returns the version character string representing the CUBRID server version.

Syntax

```
VERSION( )
```

Example

```

SELECT VERSION();
version()
=====
'8.3.1.2015'
    
```

Encryption Function

MD5 Function

Description

The **MD5** function returns the MD5 128-bit checksum for the input character string. The result value is displayed as a character string that is expressed in 32 hexadecimal, which you can use to create hash keys, for example.

The return value is a **VARCHAR(32)** type and if an input parameter is **NULL**, **NULL** will be returned.

Syntax

```
MD5(string)
```

- *string* : Input string. If a value that is not a **VARCHAR** type is entered, it will be converted to **VARCHAR**.

Example

```

SELECT MD5('cubrid');
md5('cubrid')
=====
'685c62385ce717a04f909047d0a55a16'

SELECT MD5(255);
md5(255)
=====
'fe131d7f5a6b38b23cc967316c13dae2'
SELECT MD5('01/01/2010');
md5('01/01/2010')
=====
'4a2f373c30426a1b8e9cf002ef0d4a58'

SELECT MD5(CAST('2010-01-01' as DATE));
md5(cast('2010-01-01' as date))
=====
    
```

```
'4a2f373c30426a1b8e9cf002ef0d4a58'
```

Conditional Expressions and Functions

CASE

Description

The **CASE** expression uses the SQL statement to perform an **IF ... THEN** statement. When a result of comparison expression specified in a **WHEN** clause is true, a value specified in **THEN** value is returned. A value specified in an **ELSE** clause is returned otherwise. If no **ELSE** clause exists, **NULL** is returned.

Syntax

```
CASE control_expression simple_when_list
[ else_clause ]
END

CASE searched_when_list
[ else_clause ]
END

simple_when :
WHEN expression THEN result

searched_when :
WHEN search_condition THEN result

else_clause :
ELSE result

result :
expression | NULL
```

The **CASE** expression must end with the **END** keyword. A *control_expression* argument and an *expression* argument in *simple_when* expression should be comparable data types. The data types of *result* specified in the **THEN ... ELSE** statement should all same, or they can be convertible to common data type.

The data type for a value returned by the **CASE** expression is determined based on the following rules.

- If data types for result specified in the **THEN** statement are all same, a value with the data type is returned.
- If data types can be convertible to common data type even though they are not all same, a value with the data type is returned.
- If any of values for *result* is a variable length string, a value data type is a variable length string. If values for *result* are all a fixed length string, the longest character string or bit string is returned.
- If any of values for result is an approximate numeric data type, a value with a numeric data type is returned. The number of digits after the decimal point is determined to display all significant figures.

Example

```
--creating a table
CREATE TABLE case_tbl( a INT);
INSERT INTO case_tbl VALUES (1);
INSERT INTO case_tbl VALUES (2);
INSERT INTO case_tbl VALUES (3);
INSERT INTO case_tbl VALUES (NULL);

--case operation with a search when clause
SELECT a,
       CASE WHEN a=1 THEN 'one'
            WHEN a=2 THEN 'two'
            ELSE 'other'
       END
FROM case_tbl;
      a case when a=1 then 'one' when a=2 then 'two' else 'other' end
=====
      1 'one'
```

```

        2 'two'
        3 'other'
        NULL 'other'

--case operation with a simple when clause
SELECT a,
       CASE a WHEN 1 THEN 'one'
             WHEN 2 THEN 'two'
             ELSE 'other'
       END
FROM case_tbl;
      a case a when 1 then 'one' when 2 then 'two' else 'other' end
=====
      1 'one'
      2 'two'
      3 'other'
      NULL 'other'

--result types are converted to a single type containing all of significant figures
SELECT a,
       CASE WHEN a=1 THEN 1
            WHEN a=2 THEN 1.2345
            ELSE 1.234567890
       END
FROM case_tbl;
      a case when a=1 then 1 when a=2 then 1.2345 else 1.234567890 end
=====
      1 1.000000000
      2 1.234500000
      3 1.234567890
      NULL 1.234567890

--an error occurs when result types are not convertible
SELECT a,
       CASE WHEN a=1 THEN 'one'
            WHEN a=2 THEN 'two'
            ELSE 1.2345
       END
FROM case_tbl;
ERROR: Cannot coerce 'one' to type double.

```

COALESCE Function

Description

The **COALESCE** function has more than one expression as an argument. If a first argument is non-**NULL**, the corresponding value is returned if it is **NULL**, a second argument is returned. If all expressions which have an argument are **NULL**, **NULL** is returned. Therefore, this function is generally used to replace **NULL** with other default value.

Operation is performed by converting the type of every argument into that with the highest priority. If there is an argument whose type cannot be converted, the type of every argument is converted into a **VARCHAR** type. The following list shows priority of conversion based on input argument type.

- **CHAR < VARCHAR**
- **NCHAR < NCHAR VARYING**
- **BIT < VARBIT**
- **SHORT < INT < BIGINT < NUMERIC < FLOAT < DOUBLE**
- **DATE < TIMESTAMP < DATETIME**

For example, if a type of a is **INT**, b, **BIGINT**, c, **SHORT**, and d, **FLOAT**, then **COALESCE**(a, b, c, d) returns a **FLOAT** type. If a type of a is **INTEGER**, b, **DOUBLE**, c, **FLOAT**, and d, **TIMESTAMP**, then **COALESCE**(a, b, c, d) returns a **VARCHAR** type.

Syntax

```
COALESCE(expression [, ...])
```

```
result :
expression | NULL
```

COALESCE(*a*, *b*) works the same as the **CASE** statement as follows:

```
CASE WHEN a IS NOT NULL
THEN a
ELSE b
END
```

Example

```
SELECT * FROM case_tbl;
      a
=====
      1
      2
      3
      NULL

--substituting a default value 10.0000 for NULL valuse
SELECT a, COALESCE(a, 10.0000) FROM case_tbl;
      a coalesce(a, 10.0000)
=====
      1 1.0000
      2 2.0000
      3 3.0000
      NULL 10.0000
```

DECODE Function

Description

As well as a **CASE** expression, the **DECODE** function performs the same functionality as the **IF ... THEN ... ELSE** statement. It compares the *expression* argument with *search* argument, and returns the *result* corresponding to *search* that has the same value. It returns *default* if there is no *search* with the same value, and returns **NULL** if *default* is omitted. An expression argument and a search argument to be comparable should be same or convertible each other. The number of digits after the decimal point is determined to display all significant figures including valid number of all *result*.

Syntax

```
DECODE( expression, search, result [, search, result]* [, default] )
```

```
result :
result | default | NULL
```

DECODE(*a*, *b*, *c*, *d*, *e*) has the same meaning as the **CASE** statement below.

```
CASE WHEN a = b THEN c
WHEN a = c THEN d
ELSE e
END
```

Example

```
SELECT * FROM case_tbl;
      a
=====
      1
      2
      3
      NULL

--Using DECODE function to compare expression and search values one by one
SELECT a, DECODE(a, 1, 'one', 2, 'two', 'other') FROM case_tbl;
      a decode(a, 1, 'one', 2, 'two', 'other')
=====
      1 'one'
      2 'two'
```

```

        3 'other'
        NULL 'other'

--result types are converted to a single type containing all of significant figures
SELECT a, DECODE(a, 1, 1, 2, 1.2345, 1.234567890) FROM case_tbl;
        a decode(a, 1, 1, 2, 1.2345, 1.234567890)
=====
        1 1.000000000
        2 1.234500000
        3 1.234567890
        NULL 1.234567890

--an error occurs when result types are not convertible
SELECT a, DECODE(a, 1, 'one', 2, 'two', 1.2345) FROM case_tbl;
ERROR: Cannot coerce 'one' to type double.

```

IF Function

Description

The **IF** function returns *expression2* if the value of the arithmetic expression specified as the first parameter is **TRUE**, or *expression3* if the value is **FALSE** or **NULL**. *expression2* and *expression3* which are returned as the result must be the same or of a convertible common type. If one is explicitly **NULL**, the result of the function follows the type of the non-**NULL** parameter.

Syntax

```
IF( expression1, expression2, expression3 )
```

```
result :
expression2 | expression3
```

IF(*a*, *b*, *c*) has the same meaning as the **CASE** statement in the following example:

```
CASE WHEN a IS TRUE THEN b
ELSE c
END
```

Example

```

SELECT * FROM case_tbl;
        a
=====
        1
        2
        3
        NULL

--IF function returns the second expression when the first is TRUE
SELECT a, IF(a=1, 'one', 'other') FROM case_tbl;
        a if(a=1, 'one', 'other')
=====
        1 'one'
        2 'other'
        3 'other'
        NULL 'other'

--If function in WHERE clause
SELECT * FROM case_tbl WHERE IF(a=1, 1, 2) = 1;
        a
=====
        1

```


IFNULL, NVL Function

Description

The **IFNULL** function is working like the **NVL** function; however, only the **NVL** function supports set data type as well. The **IFNULL** function (which has two arguments) returns *expr1* if the value of the first expression is not **NULL** or returns *expr2*, otherwise.

Operation is performed by converting the type of every argument into that with the highest priority. If there is an argument whose type cannot be converted, the type of every argument is converted into a **VARCHAR** type. The following list shows priority of conversion based on input argument type.

- **CHAR < VARCHAR**
- **NCHAR < NCHAR VARYING**
- **BIT < VARBIT**
- **SHORT < INT < BIGINT < NUMERIC < FLOAT < DOUBLE**
- **DATE < TIMESTAMP < DATETIME**

For example, if a type of a is **INT** and b is **BIGINT**, then **IFNULL(a, b)** returns a **BIGINT** type. If a type of a is **INTEGER** and b is **TIMESTAMP**, then **IFNULL(a, b)** returns a **VARCHAR** type.

Syntax

```
IFNULL( expr1, expr2 )
NVL( expr1, expr2 )
```

```
result :
expr1 | expr2
```

IFNULL(a, b) or **NVL(a, b)** has the same meaning as the **CASE** statement below.

```
CASE WHEN a IS NULL THEN b
ELSE a
END
```

Example

```
SELECT * FROM case_tbl;
      a
-----
      1
      2
      3
      NULL

--returning a specific value when a is NULL
SELECT a, NVL(a, 10.0000) FROM case_tbl;
      a      nvl(a, 10.0000)
-----
      1      1.0000
      2      2.0000
      3      3.0000
      NULL   10.0000

--IFNULL can be used instead of NVL and return values are converted to the string type
SELECT a, IFNULL(a, 'UNKNOWN') FROM case_tbl;
      a      ifnull(a, 'UNKNOWN')
-----
      1      '1'
      2      '2'
      3      '3'
      NULL   'UNKNOWN'
```

NULLIF Function

Description

The **NULLIF** function returns **NULL** if the two expressions specified as the parameters are identical, and returns the first parameter value otherwise.

Syntax

```
NULLIF(expr1, expr2)
```

```
result :
expr1 | NULL
```

NULLIF(*a*, *b*) is the same of the **CASE** statement.

```
CASE
WHEN a = b THEN NULL
ELSE a
END
```

Example

```
SELECT * FROM case_tbl;
      a
=====
      1
      2
      3
      NULL

--returning NULL value when a is 1
SELECT a, NULLIF(a, 1) FROM case_tbl;
      a  nullif(a, 1)
=====
      1          NULL
      2           2
      3           3
      NULL        NULL

--returning NULL value when arguments are same
SELECT NULLIF (1, 1.000) FROM db root;
      nullif(1, 1.000)
=====
      NULL

--returning the first value when arguments are not same
SELECT NULLIF ('A', 'a') FROM db root;
      nullif('A', 'a')
=====
      'A'
```

NVL2 Function

Description

Three parameters are specified for the **NVL2** function. The second expression (*expr2*) is returned if the first expression (*expr1*) is not **NULL**; the third expression (*expr3*) is returned if it is **NULL**.

Operation is performed by converting the type of every argument into that with the highest priority. If there is an argument whose type cannot be converted, the type of every argument is converted into a **VARCHAR** type. The following list shows priority of conversion based on input argument type.

- **CHAR < VARCHAR**
- **NCHAR < NCHAR VARYING**
- **BIT < VARBIT**
- **SHORT < INT < BIGINT < NUMERIC < FLOAT < DOUBLE**
- **DATE < TIMESTAMP < DATETIME**

For example, if a type of a is **INT**, b, **BIGINT**, and c, **SHORT**, then **NVL2**(a, b, c) returns a **BIGINT** type. If a type of a is **INTEGER**, b, **DOUBLE**, and c, **TIMESTAMP**, then **NVL2**(a, b, c) returns a **VARCHAR** type.

Syntax

```
NVL2( expr1, expr2, expr3 )

result :
expr2 | expr3
```

Example

```
SELECT * FROM case_tbl;
      a
=====
      1
      2
      3
      NULL

--returning a specific value of INT type
SELECT a, NVL2(a, a+1, 10.5678) FROM case_tbl;
      a  nv12(a, a+1, 10.5678)
=====
      1                2
      2                3
      3                4
      NULL             11
```

Conditional Expressions

Basic Conditional Expressions

A conditional expression is an expression that is included in the **WHERE** clause of the **SELECT**, **UPDATE** and **DELETE** statements, and in the **HAVING** clause of the **SELECT** statement. There are simple comparison, **ANY/SOME/ALL**, **BETWEEN**, **EXISTS**, **IN/NOT IN**, **LIKE** and **IS NULL** conditional expressions, depending on the kinds of the operators combined.

A simple comparison conditional expression compares two comparable data values. Expressions or subqueries are specified as operands, and the conditional expression always returns **NULL** if one of the operands is **NULL**. The following table shows operators that can be used in the simple comparison conditional expressions. For details, see [Comparison Operator](#).

Operators for Conditional Expressions

Comparison Operator	Description	Conditional Expression	Return Value
=	A value of left operand is the same as that of right operand.	1=2	0
<>, !=	A value of left operand is not the same as that of right operand.	1<>2	1
>	A value of left operand is greater than that of right operand.	1>2	0
<	A value of left operand is less than that of right operand.	1<2	1
>=	A value of left operand is equal to or greater than that of right operand.	1>=2	0
<=	A value of left operand is equal to or less than that of right operand.	1<=2	1

ANY/SOME/ALL Conditional Expressions

Description

Group conditional expressions that include quantifiers such as **ANY/SOME/ALL** perform comparison operation on one data value and on some or all values included in the list. A conditional expression that includes **ANY** or **SOME** returns **TRUE** if the value of the data on the left satisfies simple comparison with at least one of the values in the list specified as an operand on the right. A group conditional expression that includes **ALL** returns **TRUE** if the value of the data on the left satisfies simple comparison with all values in the list on the right.

When a comparison operation is performed on **NULL** in a group conditional expression that includes **ANY** or **SOME**, **UNKNOWN** or **TRUE** is returned as the result; when a comparison operation is performed on **NULL** in a group conditional expression that includes **ALL**, **UNKNOWN** or **FALSE** is returned.

Syntax

```
expression comp_op SOME expression
expression comp_op ANY expression
expression comp_op ALL expression
```

- *comp_op* : A comparison operator >, = or <= can be used.
- *expression* (left) : A single-value column, path expression, constant value or arithmetic function that produces a single value can be used.
- *expression* (right) : A column name, path expression, list (set) of constant values or subquery can be used. A list is a set represented within braces ({}). If a subquery is used, *expression* (left) and comparison operation on all results of the subquery execution is performed.

Example

```
--creating a table

CREATE TABLE condition_tbl (id int primary key, name char(10), dept name VARCHAR, salary
INT);
INSERT INTO condition_tbl VALUES(1, 'Kim', 'devel', 4000000);
INSERT INTO condition_tbl VALUES(2, 'Moy', 'sales', 3000000);
INSERT INTO condition_tbl VALUES(3, 'Jones', 'sales', 5400000);
INSERT INTO condition_tbl VALUES(4, 'Smith', 'devel', 5500000);
INSERT INTO condition_tbl VALUES(5, 'Kim', 'account', 3800000);
INSERT INTO condition_tbl VALUES(6, 'Smith', 'devel', 2400000);
INSERT INTO condition_tbl VALUES(7, 'Brown', 'account', NULL);

--selecting rows where department is sales or devel
SELECT * FROM condition_tbl WHERE dept_name = ANY{'devel','sales'};
=====
      id name      dept name      salary
=====
      1 'Kim      '      'devel'      4000000
      2 'Moy      '      'sales'      3000000
      3 'Jones    '      'sales'      5400000
      4 'Smith    '      'devel'      5500000
      6 'Smith    '      'devel'      2400000

--selecting rows comparing NULL value in the ALL group conditions
SELECT * FROM condition_tbl WHERE salary > ALL{3000000, 4000000, NULL};
There are no results.

--selecting rows comparing NULL value in the ANY group conditions
SELECT * FROM condition_tbl WHERE salary > ANY{3000000, 4000000, NULL};
=====
      id name      dept name      salary
=====
      1 'Kim      '      'devel'      4000000
      3 'Jones    '      'sales'      5400000
      4 'Smith    '      'devel'      5500000
      5 'Kim      '      'account'    3800000

--selecting rows where salary*0.9 is less than those salary in devel department
SELECT * FROM condition_tbl WHERE (
(0.9 * salary) < ALL (SELECT salary FROM condition_tbl
WHERE dept_name = 'devel')
```

```
);
```

id	name	dept name	salary
6	'Smith'	'devel'	2400000

BETWEEN Conditional Expression

Description

The **BETWEEN** conditional expression makes a comparison to determine whether the data value on the left exists between two data values specified on the right. It returns **TRUE** even when the data value on the left is the same as a boundary value of the comparison target range. If **NOT** comes before the **BETWEEN** keyword, the result of a **NOT** operation on the result of the **BETWEEN** operation is returned.

i **BETWEEN** g **AND** m and the compound condition $i \geq g$ **AND** $i \leq m$ have the same effect.

Syntax

```
expression [ NOT ] BETWEEN expression AND expression
```

- expression* : A column name, path expression, constant value, arithmetic expression or aggregate function can be used. For a character string expression, the conditions are evaluated in alphabetical order. If **NULL** is specified for at least one of the expressions, the **BETWEEN** predicate returns **UNKNOWN** as the result.

Example

```
--selecting rows where 3000000 <= salary <= 4000000
SELECT * FROM condition_tbl WHERE salary BETWEEN 3000000 AND 4000000;
SELECT * FROM condition_tbl WHERE (salary >= 3000000) AND (salary <= 4000000);
```

id	name	dept name	salary
1	'Kim'	'devel'	4000000
2	'Moy'	'sales'	3000000
5	'Kim'	'account'	3800000

```
--selecting rows where salary < 3000000 or salary > 4000000
SELECT * FROM condition_tbl WHERE salary NOT BETWEEN 3000000 AND 4000000;
```

id	name	dept name	salary
3	'Jones'	'sales'	5400000
4	'Smith'	'devel'	5500000
6	'Smith'	'devel'	2400000

```
--selecting rows where name starts from A to E
SELECT * FROM condition_tbl WHERE name BETWEEN 'A' AND 'E';
```

id	name	dept name	salary
7	'Brown'	'account'	NULL

EXISTS Conditional Expression

Description

The **EXISTS** conditional expression returns **TRUE** if one or more results of the execution of the subquery specified on the right exist, and returns **FALSE** if the result of the operation is an empty set.

Syntax

```
EXISTS expression
```

- expression* : Specifies a subquery and compares to determine whether the result of the subquery execution exists. If the subquery does not produce any result, the result of the conditional expression is **FALSE**.

Example

```
--selecting rows using EXISTS and subquery
SELECT 'raise' FROM db root WHERE EXISTS(
SELECT * FROM condition_tbl WHERE salary < 2500000);
```

```
'raise'
=====
'raise'

--selecting rows using NOT EXISTS and subquery
SELECT 'raise' FROM db root WHERE NOT EXISTS(
SELECT * FROM condition tbl WHERE salary < 2500000);
There are no results.
```

IN Conditional Expression

Description

The **IN** conditional expression compares to determine whether the single data value on the left is included in the list specified on the right. That is, the predicate returns **TRUE** if the single data value on the left is an element of the expression specified on the right. If **NOT** comes before the **IN** keyword, the result of a **NOT** operation on the result of the **IN** operation is returned.

Syntax

```
expression [ NOT ] IN expression
```

- *expression* (left) : A single-value column, path expression, constant value or arithmetic function that produces a single value can be used.
- *expression* (right) : A column name, path *expression*, list (set) of constant values or subquery can be used. A list is a set represented within parentheses (()) or braces ({}). If a subquery is used, comparison with *expression*(left) is performed for all results of the subquery execution.

Example

```
--selecting rows where department is sales or devel
SELECT * FROM condition tbl WHERE dept name IN {'devel','sales'};
SELECT * FROM condition tbl WHERE dept name = ANY{'devel','sales'};
=====
      id name      dept_name      salary
=====
      1 'Kim      '      'devel'      4000000
      2 'Moy      '      'sales'      3000000
      3 'Jones    '      'sales'      5400000
      4 'Smith    '      'devel'      5500000
      6 'Smith    '      'devel'      2400000

--selecting rows where department is neither sales nor devel
SELECT * FROM condition tbl WHERE dept name NOT IN {'devel','sales'};
=====
      id name      dept name      salary
=====
      5 'Kim      '      'account'    3800000
      7 'Brown    '      'account'    NULL
```

IS NULL Conditional Expression

Description

The **IS NULL** conditional expression compares to determine whether the expression specified on the left is **NULL**, and if it is **NULL**, returns **TRUE** and it can be used in the conditional expression. If **NOT** comes before the **NULL** keyword, the result of a **NOT** operation on the result of the **IS NULL** operation is returned.

Syntax

```
expression IS [ NOT ] NULL
```

- *expression* : A single-value column, path expression, constant value or arithmetic function that produces a single value can be used.

Example

```
SELECT * FROM condition_tbl WHERE salary IS NULL;
=====
      id name      dept_name      salary
```

```

=====
          7 'Brown      '          'account'          NULL
--selecting rows where salary is NOT NULL
SELECT * FROM condition_tbl WHERE salary IS NOT NULL;
=====
          id name          dept name          salary
=====
          1 'Kim          '          'devel'          4000000
          2 'Moy          '          'sales'          3000000
          3 'Jones         '          'sales'          5400000
          4 'Smith         '          'devel'          5500000
          5 'Kim          '          'account'        3800000
          6 'Smith         '          'devel'          2400000
--simple comparison operation returns NULL when operand is NULL
SELECT * FROM condition_tbl WHERE salary = NULL;
There are no results.

```

ISNULL Function

Description

The **ISNULL** function performs a comparison to determine if the result of the expression specified as an argument is **NULL**. The function returns 1 if it is **NULL** or 0 otherwise. You can check if a certain value is **NULL**. This function is working like the **ISNULL** expression.

Syntax

```
ISNULL(expression)
```

- *expression* : An arithmetic function that has a single-value column, path expression, constant value is specified.

Example

```

--Using ISNULL function to select rows with NULL value
SELECT * FROM condition_tbl WHERE ISNULL(salary);
=====
          id name          dept_name          salary
=====
          7 'Brown      '          'account'          NULL

```

LIKE Conditional Expression

Description

The **LIKE** conditional expression compares patterns between character string data, and returns **TRUE** if a character string whose pattern matches the search word is found. Pattern comparison target domains are **CHAR**, **VARCHAR** and **STRING**. The **LIKE** search cannot be performed on an **NCHAR** or **BIT** type. If **NOT** comes before the **LIKE** keyword, the result of a **NOT** operation on the result of the **LIKE** operation is returned.

A wild card string corresponding to any character or character string can be included in the search word on the right of the **LIKE** operator. % (percent) and _ (underscore) can be used. .% corresponds to any character string whose length is 0 or greater, and _ corresponds to one character. An escape character is a character that is used to search for a wild card character itself, and can be specified by the user as another character (**NULL**, alphabet, or number_ whose length is 1. See below for an example of using a character string that includes wild card or escape characters.

Syntax

```
expression [ NOT ] LIKE expression [ ESCAPE char ]
```

- *expression* (left) : Specify the data type column of the character string. Pattern comparison, which is case-sensitive, starts from the first character of the column.
- *expression* (right) : Enter the search word. A character string with a length of 0 or greater is required. Wild card characters (% or _) can be included as the pattern of the search word. The length of the character string is 0 or greater.

- **ESCAPE char** : NULL, alphabet, or number is allowed for *char*. If the string pattern of the search word includes "_" or "%" itself, an ESCAPE character must be specified. For example, if you want to search for the character string "10%" after specifying backslash (\) as the ESCAPE character, you must specify "10\%" for the expression (right). If you want to search for the character string "C:\", you can specify "C:\\\" for the expression (right).

Remark

The **LIKE** conditional expression is case sensitive. To disable case sensitive, use the [REGEXP/RLIKE Conditional Expression](#).

LIKE search may not work properly for data entered in multi-byte character set environment such as utf-8. This is because byte units for string comparison operation depends on the character sets. You can get normal results by adding a parameter(**single_byte_compare=yes**) to the **cubrid.conf** file that enables string comparison in a single-byte units, and restarting the DB.

For details about character sets supported in CUBRID, see [Definition and Characteristics](#). For details about the **single_byte_compare** parameter, see [Other Parameters](#).

Whether to detect the escape characters of the **LIKE** conditional expression is determined depending on the configuration of **no_backslash_escapes** and **require_like_escape_character** in the **cubrid.conf** file. For details, see [Statement/Type-Related Parameters](#).

Example

```
--selection rows where name contains lower case 's', not upper case
SELECT * FROM condition tbl WHERE name LIKE '%s%';
=====
      id name      dept name      salary
=====
      3 'Jones      '      'sales'      5400000

--selection rows where second letter is 'O' or 'o'
SELECT * FROM condition tbl WHERE UPPER(name) LIKE ' O%';
=====
      id name      dept name      salary
=====
      2 'Moy      '      'sales'      3000000
      3 'Jones      '      'sales'      5400000

--selection rows where name is 3 characters
SELECT * FROM condition tbl WHERE name LIKE '   ';
=====
      id name      dept name      salary
=====
      1 'Kim      '      'devel'      4000000
      2 'Moy      '      'sales'      3000000
      5 'Kim      '      'account'    3800000
```

REGEXP/RLIKE Conditional Expression

Description

The **REGEXP** and **RLIKE** conditional expressions are used interchangeably; A regular expressions is a powerful way to specify a pattern for a complex search. CUBRID uses Henry Spencer's implementation of regular expressions, which conforms the POSIX 1003.2 standards. The details on regular expressions are not described in this page. For more information on regular expressions, see Henry Spencer's [regex\(7\)](#).

The following list describes basic characteristics of regular expressions.

- "." matches any single character(including new-line and carriage-return).
- "[...]" matches one of characters within square brackets. For example, "[abc]" matches "a", "b", or "c". To represent a range of characters, use a dash (-). "[a-z]" matches any alphabet letter whereas "[0-9]" matches any single number.
- "*" matches 0 or more instances of the thing preceding it. For example, "abc*" matches "xab", "xabc", "xabcc", and "xabcxabc" etc. "[0-9][0-9]*" matches any numbers, and "."* matches every string.
- To match special characters such as "\n", "\t", "\r", and "\"", some must be escaped with the backslash (\) by specifying the value of **no_backslash_escapes** (default: yes) to **no**. For details on **no_backslash_escapes**, see [Escape Special Characters](#).

The difference between **REGEXP** and **LIKE** are as follows:

- The **LIKE** operator succeeds only if the pattern matches the entire value.
- The **REGEXP** operator succeeds if the pattern matches anywhere in the value. To match the entire value, you should use "^" at the beginning and "\$" at the end.
- The **LIKE** operator is case sensitive, but patterns of regular expressions in **REGEXP** is not case sensitive. To enable case sensitive, you should use **REGEXP BINARY** statement.

In the syntax below, if *expr* matches *pat*, 1 is returned; otherwise, 0 is returned. If either *expr* or *pat* is **NULL**, **NULL** is returned.

The second syntax has the same meaning as the third syntax, which both syntaxes are using **NOT**.

Syntax

```
expr REGEXP|RLIKE [BINARY] pat
expr NOT REGEXP|RLIKE pat
NOT (expr REGEXP|RLIKE pat)
```

- *expr* : Column or input expression
- *pat* : Pattern used in regular expressions; not case sensitive

Example

```
-- When REGEXP is used in SELECT list, enclosing this with parentheses is required. But
-- used in WHERE clause, no need parentheses.
-- case insensitive, except when used with BINARY.
SELECT name FROM athlete where name REGEXP '^[a-d]';
name
=====
'Dziouba Irina'
'Dzieciol Iwona'
'Dzamalutdinov Kamil'
'Crucq Maurits'
'Crosta Daniele'
'Bukovec Brigita'
'Bukic Perica'
'Abdullayev Namik'

-- \n : match a special character, when no backslash escapes=no
SELECT ('new\nline' REGEXP 'new
line');
('new
line' regexp 'new
line')
=====
1

-- ^ : match the beginning of a string
SELECT ('cubrid dbms' REGEXP '^cub');
('cubrid dbms' regexp '^cub')
=====
1

-- $ : match the end of a string
SELECT ('this is cubrid dbms' REGEXP 'dbms$');
('this is cubrid dbms' regexp 'dbms$')
=====
1

--.: match any character
SELECT ('cubrid dbms' REGEXP '^c.*$');
('cubrid dbms' regexp '^c.*$')
=====
1

-- a+ : match any sequence of one or more a characters. case insensitive.
SELECT ('Aaaapricot' REGEXP '^A+pricot');
('Aaaapricot' regexp '^A+pricot')
=====
1
```

```

-- a? : match either zero or one a character.
SELECT ('Apricot' REGEXP '^Aa?pricot');
('Apricot' regexp '^Aa?pricot')
=====
1
SELECT ('Aapricot' REGEXP '^Aa?pricot');
('Aapricot' regexp '^Aa?pricot')
=====
1

SELECT ('Aaapricot' REGEXP '^Aa?pricot');
('Aaapricot' regexp '^Aa?pricot')
=====
0

-- (cub)* : match zero or more instances of the sequence abc.
SELECT ('cubcub' REGEXP '^ (cub)*$');
('cubcub' regexp '^ (cub)*$')
=====
1

-- [a-dX], [^a-dX] : matches any character that is (or is not, if ^ is used) either a, b,
c, d or X.
SELECT ('aXbc' REGEXP '^[a-dXYZ]+');
('aXbc' regexp '^[a-dXYZ]+')
=====
1

SELECT ('strike' REGEXP '^[^a-dXYZ]+$');
('strike' regexp '^[^a-dXYZ]+$')
=====
1

```

Remark

The following shows RegEx-Spencer's license, which is library used to implement the **REGEXP** conditional expression.

```

Copyright 1992, 1993, 1994 Henry Spencer. All rights reserved.
This software is not subject to any license of the American Telephone
and Telegraph Company or of the Regents of the University of California.

```

Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it, subject to the following restrictions:

1. The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
2. The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
4. This notice may not be removed or altered.

Data Manipulation

SELECT

Overview

Description

The **SELECT** statement specifies columns that you want to retrieve from a table.

Syntax

```

SELECT [ <qualifier> ] <select_expressions>
  [ { TO | INTO } <variable_comma_list> ]
  [ FROM <extended_table_specification_comma_list> ]
  [ WHERE <search_condition> ]
  [ GROUP BY { <col_name | expr> [ ASC | DESC ],... [ WITH ROLLUP ] ]
  [ HAVING <search_condition> ]
  [ ORDER BY { <col_name | expr> [ ASC | DESC ],... [ FOR <orderby_for_condition> ] ]
  [ LIMIT [ <offset,> ] <row_count> ]
  [ USING INDEX { <index_name [,> <index_name,>... ] | NONE } ]

<qualifier> ::= ALL | DISTINCT | DISTINCTROW | UNIQUE

<select_expressions> ::= * | <expression_comma_list> | *, <expression_comma_list>

<extended_table_specification_comma_list> ::=
<table_specification> [ { , <table_specification> | <join_table_specification> }... ]

<table_specification> ::=
  <single_table_spec> [ <correlation> ] [ WITH ( <lock_hint> ) ] |
  <metaclass_specification> [ <correlation> ] |
  <subquery> <correlation> |
  TABLE ( <expression> ) <correlation>

<correlation> ::= [ AS ] <identifier> [ ( <identifier_comma_list> ) ]

<single_table_spec> ::= [ ONLY ] <table_name> |
  ALL <table_name> [ EXCEPT <table_name> ]

<metaclass_specification> ::= CLASS <class_name>

<join_table_specification> ::=
[ INNER | [ LEFT | RIGHT [ OUTER ] ] JOIN <table_specification> ON <search_condition>

lock_hint :
READ UNCOMMITTED

<orderby_for_condition> ::=
ORDERBY_NUM( ) { BETWEEN <int> AND <int> } |
  { { = | <= | < | > | >= } <int> } |
  IN ( <int> , ... )

```

- *qualifier* : A qualifier. It can be omitted. When omitted, it is set to **ALL**.
- **ALL** : Retrieves all records of the table.
- **DISTINCT** : Retrieves only records with unique values without allowing duplicates. **DISTINCT** and **DISTINCTROW** are used interchangeably.
- **UNIQUE** : Like **DISTINCT**, retrieves only records with unique values without allowing duplicates.
- *select_expression* :
- * : By using **SELECT *** statement, you can retrieve all the columns from the table specified in the **FROM** clause.
- *expression_comma_list* : *expression* can be a path expression, variable or table name. All general expressions including arithmetic operations can also be used. Use a comma (,) to separate each expression in the list. You can specify aliases by using the **AS** keyword for columns or expressions to be queried. Specified aliases are used as column names in **GROUP BY**, **HAVING**, **ORDER BY** and **FOR** clauses. The position index of a column

is assigned based on the order in which the column was specified. The starting value is 1.

As **AVG**, **COUNT**, **MAX**, **MIN**, or **SUM**, an aggregate function that manipulates the retrieved data can also be used in the *expression*. As the aggregate function returns only one result, you cannot specify a general column which has not been grouped by an aggregate function in the **SELECT** column list.

- *table_name*. * : Specifying the table name and using * has the same effect as specifying all columns for the given table.
- *variable* : The data retrieved by the *select_expression* can be stored in more than one variables.
- [:]*identifier* : By using the :identifier after **TO** (or **INTO**), you can store the data to be retrieved in the ':identifier' variable.

Example 1

The following example shows how to retrieve host countries of the Olympic Games without any duplicates. This example is performed on the olympic table of demodb.

The **DISTINCT** or **UNIQUE** keyword allows only unique values in the query result set. For example, when there are multiple **olympic** records whose **host_nation** values are 'Greece', you can use such keywords to display only one value in the query result.

```
SELECT DISTINCT host_nation FROM olympic;
host_nation
=====
'Australia'
'Belgium'
'Canada'
'Finland'
'France'
...
```

Example 2

The following example shows how to define an alias to a column to be queried and sort the result record by using the column alias in the **ORDER BY** clause. At this time, the number of the result records is limited to 5 by using the **LIMIT** clause and **FOR ORDERBY_NUM()**.

```
SELECT host_year as col1, host_nation as col2 FROM olympic ORDER BY col2 LIMIT 5;
col1 col2
=====
2000 'Australia'
1956 'Australia'
1920 'Belgium'
1976 'Canada'
1948 'England'

SELECT CONCAT(host_nation, ', ', host_city) AS host_place FROM olympic
ORDER BY host_place FOR ORDERBY_NUM() BETWEEN 1 AND 5;
host_place
=====
'Australia, Melbourne'
'Australia, Sydney'
'Belgium, Antwerp'
'Canada, Montreal'
'England, London'
```

FROM Clause

General

Description

The **FROM** clause specifies the table in which data is to be retrieved in the query. If no table is referenced, the **FROM** clause can be omitted. Retrieval paths are as follows:

- Single table
- Subquery

- Derived table

Syntax

```

SELECT [ <qualifier> ] <select_expressions>
          [ FROM <table_specification> [ {, <table_specification>
| <join table specification> }... ] ]

<select_expressions> ::= * | <expression_comma_list> | *, <expression_comma_list>

<table_specification> ::=
<single_table_spec> [ <correlation> ] [ WITH (lock_hint) ] |
<metaclass_specification> [ <correlation> ] |
<subquery> <correlation> |
TABLE ( <expression> ) <correlation>

<correlation> ::= [ AS ] <identifier> [ ( <identifier_comma_list> ) ]

<single table spec> ::= [ ONLY ] <table name> |
ALL <table_name> [ EXCEPT <table_name> ]

<metaclass_specification> ::= CLASS <class_name>

lock_hint ::=
READ UNCOMMITTED

```

- *select_expressions* : One or more columns or expressions to query is specified. Use * to query all columns in the table. You can also specify an alias for a column or an expression to be queried by using the AS keyword. This keyword can be used in **GROUP BY**, **HAVING**, **ORDER BY** and **FOR** clauses. The position index of the column is given according to the order in which the column was specified. The starting value is 1.
- *table_specification* : At least one table name is specified after the **FROM** clause. Subqueries and derived tables can also be used in the **FROM** clause. For details on subquery derived tables, see [Subquery Derived Table](#).
- *lock_hint* : You can set **READ UNCOMMITTED** for the table isolation level. **READ UNCOMMITTED** is a level where dirty reads are allowed; see [Transaction Isolation Level](#) For details on the CUBRID transaction isolation level.

Example

```

--FROM clause can be omitted in the statement
SELECT 1+1 AS sum value;
  sum_value
=====
          2

--db root can be used as a dummy table
SELECT 1+1 AS sum value FROM db root;
  sum_value
=====
          2

SELECT CONCAT('CUBRID', '2008' , 'R3.0') AS db version;
  db version
=====
'CUBRID2008R3.0'

```

Derived Table

In the query statement, subqueries can be used in the table specification of the **FROM** clause. Such subqueries create derived tables where subquery results are treated as tables. A correlation specification must be used when a subquery that creates a derived table is used.

Derived tables are also used to access the individual element of an attribute that has a set value. In this case, an element of the set value is created as an instance in the derived table.

Subquery Derived Table

Description

Each instance in the derived table is created from the result of the subquery in the **FROM** clause. A derived table created from a subquery can have any number of columns and records.

Syntax

```
FROM (subquery) [ AS ] derived_table_name [( column_name [ {, column_name }_ ] )]
```

- The number of *column_name* and the number of columns created by the *subquery* must be identical.

Example 1

The following example shows how to retrieve the sum of the number of gold medals won by Korea and that of silver medals won by Japan. This example shows a way of getting an intermediate result of the subquery and processing it as a single result, by using a derived table. The query returns the sum of the **gold** values whose **nation_code** is 'KOR' and the **silver** values whose **nation_code** column is 'JPN'.

```
SELECT SUM(n) FROM (SELECT gold FROM participant WHERE nation code='KOR'
UNION ALL SELECT silver FROM participant WHERE nation code='JPN') AS t(n);
      sum(n)
=====
          82
```

Example 2

Subquery derived tables can be useful when combined with outer queries. For example, a derived table can be used in the **FROM** clause of the subquery used in the **WHERE** clause.

The following example shows **nation_code**, **host_year** and **gold** fields of the instances whose number of gold medals is greater than average sum of the number of silver and bronze medals when one or more silver or bronze medals were won. In this example, the query (the outer **SELECT** clause) and the subquery (the inner **SELECT** clause) share the **nation_code** attribute.

```
SELECT nation code, host year, gold
FROM participant p
WHERE gold > ( SELECT AVG(s)
              FROM ( SELECT silver + bronze
                    FROM participant
                    WHERE nation code = p.nation code
                    AND silver > 0
                    AND bronze > 0
                    ) AS t(s));
nation_code      host_year      gold
=====
'JPN'            2004           16
'CHN'            2004           32
'DEN'            1996            4
'ESP'            1992           13
```

WHERE Clause

Description

In a query, a column can be processed based on conditions. The **WHERE** clause specifies a search condition for data.

Syntax

```
WHERE search_condition
```

```
search condition :
• comparison predicate
• between predicate
• exists predicate
• in predicate
```

- *null predicate*
- *like predicate*
- *quantified predicate*
- *set_predicate*

The **WHERE** clause specifies a condition that determines the data to be retrieved by *search_condition* or a query. Only data for which the condition is true is retrieved for the query results. (**NULL** value is not retrieved for the query results because it is evaluated as unknown value.)

- *search_condition* : It is described in detail in the following sections.
- [Basic Conditional Expression](#)
- [BETWEEN Conditional Expression](#)
- [EXISTS Conditional Expression](#)
- [IN Conditional Expression](#)
- [IS NULL Conditional Expression](#)
- [LIKE Conditional Expression](#)
- [ANY/SOME/aLL Conditional Expressions](#)

The logical operator **AND** or **OR** can be used for multiple conditions. If **AND** is specified, all conditions must be true. If **OR** is specified, only one needs to be true. If the keyword **NOT** is preceded by a condition, the meaning of the condition is reserved. The following table shows the order in which logical operators are evaluated.

Priority	Operator	Function
1	()	Logical expressions in parentheses are evaluated first.
2	NOT	Negates the result of the logical expression.
3	AND	All conditions in the logical expression must be true.
4	OR	One of the conditions in the logical expression must be true.

GROUP BY ... HAVING Clause

Description

The **GROUP BY** clause is used to group the result retrieved by the **SELECT** statement based on a specific column. This clause is used to sort by group or to get the aggregation by group using the aggregation function. Herein, a group consists of records that have the same value for the column specified in the **GROUP BY** clause.

You can also set a condition for group selection by including the **HAVING** clause after the **GROUP BY** clause. That is, only groups satisfying the condition specified by the **HAVING** clause are queried out of all groups that are grouped by the **GROUP BY** clause.

By SQL standard, you cannot specify a column (hidden column) not defined in the **GROUP BY** clause to the **SELECT** column list. However, by using extended CUBRID grammars, you can specify the hidden column to the **SELECT** column list. If you do not use the extended CUBRID grammars, the **only_full_group_by** parameter should be set to **yes**. For details, see [Statement/Type-Related Parameters](#).

Syntax

```
SELECT ...
GROUP BY { col_name | expr | position } [ ASC | DESC ],...
        [ WITH ROLLUP ] [ ORDER BY NULL ] [ HAVING <search_condition> ]
```

- *col_name | expr | position* : Specify one or more column names, expressions, aliases or column location. Items are separated by commas. Columns are sorted on this basis.
- [**ASC** | **DESC**] : Specify the **ASC** or **DESC** sorting option after the columns specified in the **GROUP BY** clause. If the sorting option is not specified, the default value is **ASC**.
- *search_condition* : Specify the search condition in the **HAVING** clause. In the **HAVING** clause you can refer to the hidden columns not specified in the **GROUP BY** clause as well as to columns and aliases specified in the **GROUP BY** clause and columns used in aggregate functions.

- **WITH ROLLUP** : If you specify the **WITH ROLLUP** modifier in the **GROUP BY** clause, the aggregate information of the result value of each GROUPED BY column is displayed for each group, and the total of all result rows is displayed at the last row.
- **ORDER BY NULL** : You can avoid the sorting overhead caused by **GROUP BY** by specifying the **ORDER BY NULL** modifier in the GROUP BY clause.

Example

```
--creating a new table
CREATE TABLE sales_tbl
(dept no int, name VARCHAR(20) PRIMARY KEY, sales month int, sales amount int DEFAULT 100);
INSERT INTO sales_tbl VALUES
(201, 'George' , 1, 450),
(201, 'Laura' , 2, 500),
(301, 'Max' , 4, 300),
(501, 'Stephan', 4, DEFAULT),
(501, 'Chang' , 5, 150),
(501, 'Sue' , 6, 150),
(NULL, 'Yoka' ,4, NULL);

--selecting rows grouped by dept no with ORDER BY NULL modifier
SELECT dept no, avg(sales amount) FROM sales_tbl
GROUP BY dept_no ORDER BY NULL;
=====
dept_no avg(sales_amount)
=====
NULL NULL
201 475
301 300
501 133

--conditions in WHERE clause operate first before GROUP BY
SELECT dept no, avg(sales amount) FROM sales_tbl
WHERE sales amount > 100 GROUP BY dept no;
=====
dept no avg(sales amount)
=====
201 475
301 300
501 150

--conditions in HAVING clause operate last after GROUP BY
SELECT dept_no, avg(sales_amount) FROM sales_tbl
WHERE sales_amount > 100 GROUP BY dept_no HAVING avg(sales_amount) > 200;
=====
dept no avg(sales amount)
=====
201 475
301 300

--selecting and sorting rows with using column alias
SELECT dept no AS a1, avg(sales amount) AS a2 FROM sales_tbl
WHERE sales amount > 200 GROUP BY a1 HAVING a2 > 200 ORDER BY a2;
=====
a1 a2
=====
301 300
201 475

--selecting rows grouped by dept no with WITH ROLLUP modifier
SELECT dept no AS a1, name AS a2, avg(sales amount) AS a3 FROM sales_tbl
WHERE sales_amount > 100 GROUP BY a1,a2 WITH ROLLUP;
=====
a1 a2 a3
=====
201 'George' 450
201 'Laura' 500
201 NULL 475
301 'Max' 300
301 NULL 300
501 'Chang' 150
501 'Sue' 150
501 NULL 150
NULL NULL 310
```


ORDER BY Clause

Description

The **ORDER BY** clause sorts the query result set in ascending or descending order. If you do not specify a sorting option such as **ASC** or **DESC**, the result set in ascending order by default. If you do not specify the **ORDER BY** clause, the order of records to be queried may vary depending on query.

Syntax

```
SELECT ...
ORDER BY {col name | expr | position} [ASC | DESC],...
        [ FOR <orderby for condition> ] ]

<orderby_for_condition> ::=
ORDERBY_NUM() { BETWEEN int AND int } |
{ { = | =< | < | > | >= } int } |
IN ( int, ...)
```

- *col_name | expr | position* : Specify an column name, expression, alias, or column location. One or more column names, expressions or aliases can be specified. Items are separated by commas. A column that is not specified in the list of **SELECT** columns can be specified.
- [**ASC** | **DESC**] : **ASC** means sorting in ascending order, and **DESC** is sorting in descending order. If the sorting option is not specified, the default value is **ASC**.

Example

```
--selecting rows sorted by ORDER BY clause
SELECT * FROM sales tbl ORDER BY dept no DESC, name ASC;
      dept no  name                sales month  sales amount
=====
          501  'Chang'                        5             150
          501  'Stephan'                     4             100
          501  'Sue'                          6             150
          301  'Max'                          4             300
          201  'George'                       1             450
          201  'Laura'                         2             500
          NULL  'Yoka'                          4             NULL

--sorting reversely and limiting result rows by LIMIT clause
SELECT dept no AS a1, avg(sales amount) AS a2 FROM sales tbl
GROUP BY a1 ORDER BY a2 DESC LIMIT 0,3;
      a1      a2
=====
          201      475
          301      300
          501      133

--sorting reversely and limiting result rows by FOR clause
SELECT dept_no AS a1, avg(sales_amount) AS a2 FROM sales_tbl
GROUP BY a1 ORDER BY a2 DESC FOR ORDERBY NUM() BETWEEN 1 AND 3;
      a1      a2
=====
          201      475
          301      300
          501      133
```

LIMIT Clause

Description

The **LIMIT** clause can be used to limit the number of records displayed. It takes one or two arguments. You can specify a very big integer for *row_count* to output to the last row, starting from a specific row.

The **LIMIT** clause can be used as a prepared statement. In this case, the bind parameter (?) can be used instead of an argument.

INST_NUM() and **ROWNUM** cannot be included in the **WHERE** clause in a query that contains the **LIMIT** clause. Also, **LIMIT** cannot be used together with **FOR ORDERBY_NUM()** or **HAVING GROUPBY_NUM()**.

Syntax

```
LIMIT [offset,] row_count
```

- *offset* : Specify the offset value of the starting row to be output. The offset value of the starting row of the result set is 0; it can be omitted and the default value is 0.
- *row_count* : Specify the number of records to be output. You can specify an integer greater than 0.

Example

```
--LIMIT clause can be used in prepared statement
PREPARE STMT FROM 'SELECT * FROM sales_tbl LIMIT ?, ?';
EXECUTE STMT USING 0, 10;

--selecting rows with LIMIT clause
SELECT * FROM sales tbl WHERE sales amount > 100 LIMIT 5;
    dept_no  name                sales_month  sales_amount
=====
         201  'George'                1             450
         201  'Laura'                 2             500
         301  'Max'                   4             300
         501  'Chang'                 5             150
         501  'Sue'                   6             150

--LIMIT clause can be used in subquery
SELECT t1.* FROM
(SELECT * FROM sales tbl AS t2 WHERE sales amount > 100 LIMIT 5) AS t1 LIMIT 1,3;
    dept_no  name                sales_month  sales_amount
=====
         201  'Laura'                 2             500
         301  'Max'                   4             300
         501  'Chang'                 5             150
```

Outer Join

Description

A join is a query that combines the rows of two or more tables or virtual tables (views). In a join query, a condition that compares the columns that are common in two or more tables is called a join condition. Rows are retrieved from each joined table, and are combined only when they satisfy the specified join condition.

A join query using an equality operator (=) is called an equi-join, and one without any join condition is called a cartesian product. Meanwhile, joining a single table is called a self join. In a self join, table **ALIAS** is used to distinguish columns, because the same table is used twice in the **FROM** clause.

A join that outputs only rows that satisfy the join condition from a joined table is called an inner or a simple join, whereas a join that outputs both rows that satisfy and do not satisfy the join condition from a joined table is called an outer join. An outer join is divided into a left outer join which outputs all rows of the left table as the result, a right outer join which outputs all rows of the right table as the result and a full outer join which outputs all rows of both tables. If there is no column value that corresponds to a table on one side in the result of an outer join query, all rows are returned as **NULL**.

Syntax

```
FROM table_specification [{, table_specification | join_table_specification}...]

table_specification :
correlation ]
CLASS table_name [ correlation ]
subquery correlation
TABLE (expression) correlation

join_table_specification :
```

```
[ INNER | {LEFT | RIGHT} [ OUTER ] ] JOIN table_specification
join_condition

join_condition :
ON search_condition
```

- *oin_table_specification*
- { **LEFT** | **RIGHT** } [**OUTER**] **JOIN** : **LEFT** is used for a left outer join query, and **RIGHT** is for a right outer join query.

CUBRID does not support full outer joins. Path expressions that include subqueries and sub-columns cannot be used in the join conditions of an outer join.

Join conditions of an outer join are specified in a different way from those of an inner join. In an inner join, join conditions are expressed in the **WHERE** clause; in an outer join, they appear after the **ON** keyword in the **FROM** clause. Other retrieval conditions can be used in the **WHERE** or **ON** clause, but the retrieval result depends on whether the condition is used in the **WHERE** or **ON** clause.

The table execution order is fixed according to the order specified in the **FROM** clause. Therefore, when using an outer join, you should create a query statement in consideration of the table order. It is recommended to use standard statements using { **LEFT** | **RIGHT** } [**OUTER**] **JOIN**, because using an Oracle-style join query statements by specifying an outer join operator (+) in the **WHERE** clause, even if possible, might lead the execution result or plan in an unwanted direction.

Example 1

The following example shows how to retrieve the years and host countries of the Olympic Games since 1950 where a world record has been set. The following query retrieves instances whose values of the **host_year** column in the **history** table are greater than 1950.

```
SELECT DISTINCT h.host_year, o.host_nation FROM history h, olympic o
WHERE h.host_year=o.host_year AND o.host_year>1950;
  host_year  host_nation
=====
          1968  'Mexico'
          1980  'U.S.S.R.'
          1984  'United States of America'
          1988  'Korea'
          1992  'Spain'
          1996  'United States of America'
          2000  'Australia'
          2004  'Greece'
```

Example 2

The following example shows how to retrieve the years and host countries of the Olympic Games since 1950 where a world record has been set, but including the Olympic Games where any world records haven't been set in the result. This example can be expressed in the following right outer join query. In this example, all instances whose values of the **host_year** column in the **history** table are not greater than 1950 are also retrieved. All instances of **host_nation** are included because this is a right outer join. **host_year** that does not have a value is represented as **NULL**.

```
SELECT DISTINCT h.host_year, o.host_nation
FROM history h RIGHT OUTER JOIN olympic o ON h.host_year=o.host_year WHERE
o.host_year>1950;
  host_year  host_nation
=====
          NULL  'Australia'
          NULL  'Canada'
          NULL  'Finland'
          NULL  'Germany'
          NULL  'Italy'
          NULL  'Japan'
          1968  'Mexico'
          1980  'U.S.S.R.'
          1984  'United States of America'
          1988  'Korea'
          1992  'Spain'
```

```

1996 'United States of America'
2000 'Australia'
2004 'Greece'

```

Example 3

A right outer join query can be converted to a left outer join query by switching the position of two tables in the **FROM** clause. The right outer join query in the previous example can be expressed as a left outer join query as follows:

```

SELECT DISTINCT h.host year, o.host nation
FROM olympic o LEFT OUTER JOIN history h ON h.host_year=o.host_year WHERE o.host_year>1950;
  host_year  host_nation
=====
      NULL   'Australia'
      NULL   'Canada'
      NULL   'Finland'
      NULL   'Germany'
      NULL   'Italy'
      NULL   'Japan'
      1968   'Mexico'
      1980   'U.S.S.R.'
      1984   'United States of America'
      1988   'Korea'
      1992   'Spain'
      1996   'United States of America'
      2000   'Australia'
      2004   'Greece'
14 rows selected.

```

In this example, **h.host_year=o.host_year** is an outer join condition, and **o.host_year > 1950** is a search condition. If the search condition is used not in the **WHERE** clause but in the **ON** clause, the meaning and the result will be different. The following query also includes instances whose values of **o.host_year** are not greater than 1950.

```

SELECT DISTINCT h.host_year, o.host_nation
FROM olympic o LEFT OUTER JOIN history h ON h.host year=o.host year AND
o.host year>1950;

=== <Result of SELECT Command in Line 3> ===
  host_year  host_nation
=====
      NULL   'Australia'
      NULL   'Belgium'
      NULL   'Canada'
...
      1996   'United States of America'
      2000   'Australia'
      2004   'Greece'

```

Example 4

Outer joins can also be represented by using (+) in the **WHERE** clause. The above example is a query that has the same meaning as the example using the **LEFT OUTER JOIN**. The (+) syntax is not ISO/aNSI standard, so it can lead to ambiguous situations. It is recommended to use the standard syntax **LEFT OUTER JOIN** (or **RIGHT OUTER JOIN**) if possible.

```

SELECT DISTINCT h.host year, o.host nation FROM history h, olympic o
WHERE o.host year=h.host year(+) AND o.host year>1950;
  host_year  host_nation
=====
      NULL   'Australia'
      NULL   'Canada'
      NULL   'Finland'
      NULL   'Germany'
      NULL   'Italy'
      NULL   'Japan'
      1968   'Mexico'
      1980   'U.S.S.R.'
      1984   'United States of America'
      1988   'Korea'
      1992   'Spain'
      1996   'United States of America'

```

```
2000 'Australia'
2004 'Greece'
```

Subquery

A subquery can be used wherever expressions such as **SELECT** or **WHERE** clause can be used. If the subquery is represented as an expression, it must return a single column; otherwise it can return multiple rows. Subqueries can be divided into single-row subquery and multiple-row subquery depending on how they are used.

Single-Row Subquery

Description

A single-row subquery outputs a row that has a single column. If no row is returned by the subquery, the subquery expression has a **NULL** value. If the subquery is supposed to return more than one rows, an error occurs.

Example

The following example shows how to retrieve the history table as well as the host country where a new world record has been set. This example shows a single-row subquery used as an expression. In this example, the subquery returns **host_nation** values for the rows whose values of the **host_year** column in the **olympic** table are the same as those of the **host_year** column in the **history** table. If there are no values that meet the condition, the result of the subquery is **NULL**.

```
SELECT h.host year, (SELECT host nation FROM olympic o WHERE o.host year=h.host year),
h.event code, h.score, h.unit from history h;
      host_year (SELECT host_nation FROM olympic o WHERE
o.host_year=h.host_year)      event_code      score      unit
=====
==
      2004 'Greece'      20283 '07:53.0'      'time'
      2004 'Greece'      20283 '07:53.0'      'time'
      2004 'Greece'      20281 '03:57.0'      'time'
      2004 'Greece'      20281 '03:57.0'      'time'
      2004 'Greece'      20281 '03:57.0'      'time'
      2004 'Greece'      20281 '03:57.0'      'time'
      2004 'Greece'      20326 '210'      'kg'
      2000 'Australia'      20328 '225'      'kg'
      2004 'Greece'      20331 '237.5'      'kg'
...

```

Multiple-Row Subquery

Description

The multiple-row subquery returns one or more rows that contain the specified column. The result of the multiple-row subquery can be used to create a set, a multiset or a list/sequence set using an appropriate keyword (**SET**, **MULTISET**, **LIST** or **SEQUENCE**).

Example

The following example shows how to retrieve countries and their capital cities from the nation table, and returning lists of host countries and host cities of the Olympic Games. In this example, the subquery result is used to create a list from the values of the **host_city** column in the **olympic** table. This query returns **name** and **capital** value for **nation** table, as well as a set that contains **host_city** values of the **olympic** table with **host_nation** value. If the **name** value is an empty set in the query result, it is excluded. If there is no **olympic** table that has the same value as the **name**, an empty set is returned.

```
SELECT name, capital, list(SELECT host_city FROM olympic WHERE host_nation = name) FROM
nation;
      name      capital      sequence((SELECT host city FROM olympic
WHERE host nation=name))
=====

```

```
'Somalia'           'Mogadishu'           {}
'Sri Lanka'         'Sri Jayewardenepura Kotte' {}
'Sao Tome & Principe' 'Sao Tome'           {}
...
'U.S.S.R.'          'Moscow'              {'Moscow'}
'Uruguay'           'Montevideo'         {}
'United States of America' 'Washington.D.C'     {'Atlanta ', 'St. Louis', 'Los Angeles', 'Los Angeles'}
'Uzbekistan'        'Tashkent'           {}
'Vanuatu'           'Port Vila'          {}
```

Such multiple-row subquery expressions can be used anywhere a set value expression is allowed. However, they cannot be used where a set constant value is required as in the **DEFAULT** specification in the class attribute definition.

If the **ORDER BY** clause is not used explicitly in the subquery, the order of the multiple-row query result is not set. Therefore, the order of the multiple-row subquery result that creates a sequence set must be specified by using the **ORDER BY** clause.

Hierarchical Query

START WITH ... CONNECT BY Clause

Description

This clause is used to obtain a set of data organized in a hierarchy. The **START WITH ... CONNECT BY** clause is used in combination with the **SELECT** clause in the following form.

Syntax

```
SELECT column_list
FROM table_joins | tables
[WHERE join_conditions and/or filtering_conditions]
[START WITH condition]
CONNECT BY [NOCYCLE] condition
```

START WITH Clause

The **START WITH** clause will filter the rows from which the hierarchy will start. The rows that satisfy the **START WITH** condition will be the root nodes of the hierarchy. If **START WITH** is omitted, then all the rows will be considered as root nodes.

Note If **START WITH** clause is omitted or the rows that satisfy the **START WITH** condition does not exist, all of rows in the table are considered as root nodes; which means that hierarchy relationship of sub rows which belong each root is searched. Therefore, some of results can be duplicate.

CONNECT BY [NOCYCLE] or PRIOR Clause

- **PRIOR** : The **CONNECT BY** condition is tested for a pair of rows. If it evaluates to true, the two rows satisfy the parent-child relationship of the hierarchy. We need to specify the columns that are used from the parent row and the columns that are used from the child row. We can use the **PRIOR** operator when applied to a column, which will refer to the value of the parent row for that column. If **PRIOR** is not used for a column, the value in the child row is used.
- **NOCYCLE** : In some cases, the resulting rows of the table joins may contain cycles, depending on the **CONNECT BY** condition. Because cycles cause an infinite loop in the result tree construction, CUBRID detects them and either returns an error doesn't expand the branches beyond the point where a cycle is found (if the **NOCYCLE** keyword is specified).
This keyword may be specified after the **CONNECT BY** keywords. It makes CUBRID run a statement even if the processed data contains cycles.
If a **CONNECT BY** statement causes a cycle at runtime and the **NOCYCLE** keyword is not specified, CUBRID will return an error and the statement will be canceled. When specifying the **NOCYCLE** keyword, if CUBRID detects a cycle while processing a hierarchy node, it will set the **CONNECT_BY_ISCYCLE** attribute for that node to the value of 1 and it will stop further expansion of that branch.

Example

For the following samples, you will need the following structures:

tree Table

ID	MgrID	Name	BirthYear
1	NULL	Kim	1963
2	NULL	Moy	1958
3	1	Jonas	1976
4	1	Smith	1974
5	2	Verma	1973
6	2	Foster	1972
7	6	Brown	1981

tree_cycle table

ID	MgrID	Name
1	NULL	Kim
2	11	Moy
3	1	Jonas
4	1	Smith
5	3	Verma
6	3	Foster
7	4	Brown
8	4	Lin
9	2	Edwin
10	9	Audrey
11	10	Stone

```
-- Creating tree table and then inserting data
CREATE TABLE tree(ID INT, MgrID INT, Name VARCHAR(32), BirthYear INT);

INSERT INTO tree VALUES (1,NULL,'Kim', 1963);
INSERT INTO tree VALUES (2,NULL,'Moy', 1958);
INSERT INTO tree VALUES (3,1,'Jonas', 1976);
INSERT INTO tree VALUES (4,1,'Smith', 1974);
INSERT INTO tree VALUES (5,2,'Verma', 1973);
INSERT INTO tree VALUES (6,2,'Foster', 1972);
INSERT INTO tree VALUES (7,6,'Brown', 1981);

-- Creating tree cycle table and then inserting data
CREATE TABLE tree cycle(ID INT, MgrID INT, Name VARCHAR(32));

INSERT INTO tree cycle VALUES (1,NULL,'Kim');
INSERT INTO tree cycle VALUES (2,11,'Moy');
INSERT INTO tree cycle VALUES (3,1,'Jonas');
INSERT INTO tree cycle VALUES (4,1,'Smith');
INSERT INTO tree cycle VALUES (5,3,'Verma');
INSERT INTO tree cycle VALUES (6,3,'Foster');
INSERT INTO tree cycle VALUES (7,4,'Brown');
INSERT INTO tree cycle VALUES (8,4,'Lin');
INSERT INTO tree cycle VALUES (9,2,'Edwin');
INSERT INTO tree cycle VALUES (10,9,'Audrey');
INSERT INTO tree cycle VALUES (11,10,'Stone');

-- Executing a hierarchy query with CONNECT BY clause
SELECT id, mgrid, name
```

```

FROM tree
CONNECT BY PRIOR id=mgrid
ORDER BY id;

id  mgrid      name
=====
1   null       Kim
2   null       Moy
3   1          Jonas
3   1          Jonas
4   1          Smith
4   1          Smith
5   2          Verma
5   2          Verma
6   2          Foster
6   2          Foster
7   6          Brown
7   6          Brown
7   6          Brown

-- Executing a hierarchy query with START WITH clause
SELECT id, mgrid, name
FROM tree
START WITH mgrid IS NULL
CONNECT BY prior id=mgrid
ORDER BY id;

id  mgrid      name
=====
1   null       Kim
2   null       Moy
3   1          Jonas
4   1          Smith
5   2          Verma
6   2          Foster
7   6          Brown

```

Hierarchical Query for Table Join

Join Conditions

The table joins are evaluated first using the join conditions, if any. The conditions found in the **WHERE** clause are classified as join conditions or filtering conditions. All the conditions in the **FROM** clause are classified as join conditions. Only the join conditions are evaluated; the filtering conditions are kept for later evaluation. We recommend placing all join conditions in the **FROM** clause only so that conditions that are intended for joins are not mistakenly classified as filtering conditions.

Query Results

The resulting rows of the table joins are filtered according to the **START WITH** condition to obtain the root nodes for the hierarchy. If no **START WITH** condition is specified, then all the rows resulting from the table joins will be considered as root nodes.

After the root nodes are obtained, CUBRID will select the child rows for the root nodes. These are all nodes from the table joins that respect the **CONNECT BY** condition. This step will be repeated for the child nodes to determine their child nodes and so on until no more child nodes can be added.

In addition, CUBRID evaluates the **CONNECT BY** clause first and all the rows of the resulting hierarchy tree by using the filtering condition in the **WHERE** clause.

Example

The example illustrates how joins can be used in **CONNECT BY** queries. The joins are evaluated before the **CONNECT BY** condition and the join result will be the starting table on which the two clauses (**START WITH** clause and **CONNECT BY** clause).

```
-- Creating tree2 table and then inserting data
```



```

CREATE TABLE tree2(id int, treeid int, job varchar(32));

INSERT INTO tree2 VALUES(1,1,'Partner');
INSERT INTO tree2 VALUES(2,2,'Partner');
INSERT INTO tree2 VALUES(3,3,'Developer');
INSERT INTO tree2 VALUES(4,4,'Developer');
INSERT INTO tree2 VALUES(5,5,'Sales Exec. ');
INSERT INTO tree2 VALUES(6,6,'Sales Exec. ');
INSERT INTO tree2 VALUES(7,7,'Assistant');
INSERT INTO tree2 VALUES(8,null,'Secretary');

-- Executing a hierarchical query onto table joins
SELECT t.id,t.name,t2.job,level
  FROM tree t
        inner join tree2 t2 on t.id=t2.treeid
 START WITH t.mgrid is null
 CONNECT BY prior t.id=t.mgrid
 ORDER BY t.id;

```

id	name	job	level
1	Kim	Partner	1
2	Moy	Partner	1
3	Jonas	Developer	2
4	Smith	Developer	2
5	Verma	Sales Exec.	2
6	Foster	Sales Exec.	2
7	Brown	Assistant	3

Pseudo Columns Available in Hierarchical Queries

LEVEL

LEVEL is a pseudo column representing depth of hierarchical queries. The **LEVEL** of root node is 1 and the **LEVEL** of its child node is 2.

The **LEVEL** (pseudo column) can be used in the **WHERE** clause, **ORDER BY** clause, and **GROUP BY ... HAVING** clause of the **SELECT** statement. And it can also be used in the statement using aggregate functions.

The following example shows how to retrieve the **LEVEL** value to view level of node.

```

-- Viewing LEVEL value
SELECT id, mgrid, name, LEVEL
  FROM tree
 WHERE LEVEL=2
 START WITH mgrid IS NULL
 CONNECT BY PRIOR id=mgrid
 ORDER BY id;

```

id	mgrid	name	level
3	1	Jonas	2
4	1	Smith	2
5	2	Verma	2
6	2	Foster	2

CONNECT_BY_ISLEAF

This pseudo-column indicates whether a hierarchical node is a leaf node or not. If the value for a row is 1, then the associated node is a leaf node.; otherwise, it will have the value 0 indicating that the node has children.

In this example, the **CONNECT_BY_ISLEAF** shows that the rows with the IDs 3, 4, 5 and 7 have no children.

```

-- Executing a hierarchical query with CONNECT_BY_ISLEAF
SELECT id, mgrid, name, CONNECT BY ISLEAF
  FROM tree
 START WITH mgrid IS NULL
 CONNECT BY PRIOR id=mgrid
 ORDER BY id;

```

id	mgrid	name	connect_by_isleaf
----	-------	------	-------------------

```
=====
```

1	null	Kim	0
2	null	Moy	0
3	1	Jonas	1
4	1	Smith	1
5	2	Verma	1
6	2	Foster	0
7	6	Brown	1

CONNECT_BY_ISCYCLE

This pseudo-column indicates that a cycle was detected while processing the node, meaning that a child was also found to be an ancestor. A value of 1 for a row means a cycle was detected; the pseudo-column's value is 0, otherwise.

The **CONNECT_BY_ISCYCLE** pseudo-column may be used in the **SELECT** list, **WHERE** clause, **ORDER BY** clause, **GROUP BY** and **HAVING** clauses and also in aggregate functions (when the **GROUP BY** class exists in the statement).

Note This pseudo-column is available only when the **NOCYCLE** keyword is used in the statement.

The following example shows how to execute a hierarchical query with **CONNECT_BY_ISCYCLE** operator.

```
-- --Executing a hierarchical query with CONNECT BY ISCYCLE
SELECT id, mgrid, name, CONNECT BY ISCYCLE
FROM tree_cycle
START WITH name in ('Kim', 'Moy')
CONNECT BY NOCYCLE PRIOR id=mgrid
ORDER BY id;
```

id	mgrid	name	connect by iscycle
1	null	Kim	0
2	11	Moy	0
3	1	Jonas	0
4	1	Smith	0
5	3	Verma	0
6	3	Foster	0
7	4	Brown	0
8	4	Lin	0
9	2	Edwin	0
10	9	Audrey	0
11	10	Stone	1

Operators Available When Using the CONNECT BY Clause

CONNECT_BY_ROOT Operator

This operator can be applied to columns and it returns the parent row or root row values for that column. This operator may be used in the **SELECT** list, **Where** clause and **ORDER BY** clause. When using the **CONNECT BY** clause some column operators become available.

The following example shows how to execute a hierarchical query with **CONNECT_BY_ROOT** operator.

```
-- Executing a hierarchical query with CONNECT_BY_ROOT operator
SELECT id, mgrid, name, CONNECT_BY_ROOT id
FROM tree
START WITH mgrid IS NULL
CONNECT BY PRIOR id=mgrid
ORDER BY id;
```

id	mgrid	name	connect_by_root id
1	null	Kim	1
2	null	Moy	2
3	1	Jonas	1
4	1	Smith	1
5	2	Verma	2
6	2	Foster	2
7	6	Brown	2

PRIOR Operator

This operator may be applied to a column; it will return the parent node value for that column. For a root node, the operator will return the **NULL** value if it is applied to a column. This operator may be used in the **SELECT** list, **WHERE** clause, **ORDER BY** clause and also in the **CONNECT BY** clause.

The following example shows how to execute a hierarchical query with **PRIOR** operator.

```
-- Executing a hierarchical query with PRIOR operator
SELECT id, mgrid, name, PRIOR id as "prior id"
   FROM tree
  START WITH mgrid IS NULL
 CONNECT BY PRIOR id=mgrid
  ORDER BY id;
```

id	mgrid	name	prior_id
1	null	Kim	null
2	null	Moy	null
3	1	Jonas	1
4	1	Smith	1
5	2	Verma	2
6	2	Foster	2
7	6	Brown	6

Functions Available When Using the CONNECT BY Clause

Description

The **SYS_CONNECT_BY_PATH** function returns the branch of the node in the hierarchy. It returns a string that represents the concatenation of all the values obtained by evaluating the scalar expression for all the parents of a row, including that row, separated by the separator character, ordered ascending by level.

This function may be used in the **SELECT** list, **WHERE** clause and **ORDER BY** clause.

Syntax

```
SYS_CONNECT_BY_PATH (column_name, separator_char)
```

Example

The following example shows how to execute a hierarchical query with **SYS_CONNECT_BY_PATH** function.

```
-- Executing a hierarchical query with SYS CONNECT BY PATH function
SELECT id, mgrid, name, SYS CONNECT BY PATH(name, '/') as [hierarchy]
   FROM tree
  START WITH mgrid IS NULL
 CONNECT BY PRIOR id=mgrid
  ORDER BY id;
```

id	mgrid	name	hierarchy
1	null	Kim	/Kim
2	null	Moy	/Moy
3	1	Jonas	/Kim/Jonas
4	1	Smith	/Kim/Smith
5	2	Verma	/Moy/Verma
6	2	Foster	/Moy/Foster
7	6	Brown	/Moy/Foster/Brown

Ordering Data with the Hierarchical Query

Description

The **ORDER SIBLINGS BY** clause will cause the ordering of the rows while preserving the hierarchy ordering so that the child nodes with the same parent will be stored according to the column list.

Syntax

```
ORDER SIBLINGS BY col_1 [ASC|DESC] [, col_2 [ASC|DESC] [...[, col_n [ASC|DESC]]...]]
```

Example 1

The following example shows how to display information about seniors and subordinates in a company in the order of birth year.

The result with hierarchical query shows parent and child nodes in a row according to the column list specified in **ORDER SIBLINGS BY** statement by default. Sibling nodes that share the same parent node have outputted in a specified order.

```
-- Outputting a parent node and its child nodes, which sibling nodes that share the same
parent are sorted in the order of birth year.
SELECT id, mgrid, name, birthyear, level
FROM tree
START WITH mgrid IS NULL
CONNECT BY PRIOR id=mgrid
ORDER SIBLINGS BY birthyear;
```

id	mgrid	name	birthyear	level
2	NULL	'Moy'	1958	1
6	2	'Foster'	1972	2
7	6	'Brown'	1981	3
5	2	'Verma'	1973	2
1	NULL	'Kim'	1963	1
4	1	'Smith'	1974	2
3	1	'Jonas'	1976	2

Example 2

The following example shows how to display information about seniors and subordinates in a company in the order of joining. For the same level, the employee ID numbers are assigned in the order of joining. id indicates employee ID numbers (parent and child nodes) and mgrid indicates the employee ID numbers of their seniors.

```
-- Outputting siblings in a row
SELECT id, mgrid, name, LEVEL
FROM tree
START WITH mgrid IS NULL
CONNECT BY PRIOR id=mgrid
ORDER SIBLINGS BY id;
```

id	mgrid	name	level
1	null	Kim	1
3	1	Jonas	2
4	1	Smith	2
2	null	Moy	1
5	2	Verma	2
6	2	Foster	2
7	6	Brown	3

Scenario of Using Hierarchical Query

First of all let's start by giving a rough SQL translation of the **SELECT** statement with a **CONNECT BY** clause. For this we can consider that we have a table that contains a recurrent reference. We can consider that table to have two columns named ID and ParentID; ID is the primary key for the table and ParentID is a foreign-key to the same table. Naturally, the root nodes will have a ParentID value of **NULL**.

Now let us consider the fact that we want to get the full rows and a column with the level of the row in the hierarchy tree. For this we can write something similar to by querying with **UNION ALL**.

```
SELECT L1.ID, L1.ParentID, ..., 1 AS [Level]
FROM tree_table AS L1
WHERE L1.ParentID IS NULL
UNION ALL
SELECT L2.ID, L2.ParentID, ..., 2 AS [Level]
```

```

FROM tree table AS L1
  INNER JOIN tree table AS L2 ON L1.ID=L2.ParentID
WHERE L1.ParentID IS NULL
UNION ALL
SELECT L3.ID, L3.ParentID, ..., 3 AS [Level]
FROM tree table AS L1
  INNER JOIN tree table AS L2 ON L1.ID=L2.ParentID
  INNER JOIN tree table AS L3 ON L2.ID=L3.ParentID
WHERE L1.ParentID IS NULL
UNION ALL ...

```

The problem with our approach is that we do not know how many levels we have. This could be rewritten in a stored procedure with a cycle until no new rows are retrieved, but we will have to check the tree for cycles at every step. Using a **SELECT** statement with a **CONNECT BY** clause we can rewrite this as follows.

This query will return the full hierarchy with the level of each row in the hierarchy.

```

SELECT ID, ParentID, ..., Level
FROM tree_table
START WITH ParentID IS NULL
CONNECT BY ParentID=PRIOR ID

```

If we want to avoid the potential error caused by cycles we can write it as follows:

```

SELECT ID, ParentID, ..., Level
FROM tree table
START WITH ParentID IS NULL
CONNECT BY NOCYCLE ParentID=PRIOR ID

```

Performance of Hierarchical Query

Although this form is shorter and clearer, please keep in mind that it has its limitations regarding speed. If the result of the query contains all the rows of the table, the **CONNECT BY** form might be slower as it has to do additional processing (such as cycle detection, pseudo-column bookkeeping and others). However, if the result of the query only contains a part of the table rows, the **CONNECT BY** form might be faster.

For example, if we have a table with 20,000 records and we want to retrieve a sub-tree of roughly 1,000 records, a **SELECT** statement with a **START WITH ... CONNECT BY** clause will run up to 30% faster than an equivalent **UNION ALL** with **SELECT** statements.

INSERT

Overview

Description

You can insert a new record into a table in a database by using the **INSERT** statement. CUBRID supports **INSERT...VALUES**, **INSERT...SET** and **INSERT...SELECT** statements.

INSERT...VALUES and **INSERT...SET** statements are used to insert a new record based on the value that is explicitly specified while the **INSERT...SELECT** statement is used to insert query result records obtained from different tables. Use the **INSERT VALUES** or **INSERT...SELECT** statement to insert multiple rows by using the single **INSERT** statement.

Syntax

```

<INSERT ... VALUES statement>
INSERT [INTO] table_name [(column_name, ...)]
  {VALUES | VALUE}({expr | DEFAULT}, ...)[, ({expr | DEFAULT}, ...),...]
  [ON DUPLICATE KEY UPDATE column_name = expr, ... ]
INSERT [INTO] table_name DEFAULT [ VALUES ]
INSERT [INTO] table_name VALUES ()

<INSERT ... SET statement>
INSERT [INTO] table_name
  SET column name = {expr | DEFAULT}[, column name = {expr | DEFAULT},...]
  [ON DUPLICATE KEY UPDATE column_name = expr, ... ]

```

```
<INSERT ... SELECT statement>
INSERT [INTO] table_name [(column_name, ...)]
SELECT ...
[ON DUPLICATE KEY UPDATE column_name = expr, ... ]
```

- *table_name* : Specify the name of the target table into which you want to insert a new record.
- *column_name* : Specify the name of the column into which you want to insert the value. If you omit to specify the column name, it is considered that all columns defined in the table have been specified. Therefore, you must specify the values for all columns next to the **VALUES** keyword. If you do not specify all the columns defined in the table, a **DEFAULT** value is assigned to the non-specified columns; if the **DEFAULT** value is not defined, a **NULL** value is assigned.
- *expr* | **DEFAULT** : Specify values that correspond to the columns next to the **VALUES** keyword. Expressions or the **DEFAULT** keyword can be specified as a value. At this time, the order and number of the specified column list must correspond to the column value list. The column value list for a single record is described in parentheses.
- **DEFAULT** : You can use the **DEFAULT** keyword to specify a default value as the column value. If you specify **DEFAULT** in the column value list next to the **VALUES** keyword, a default value column is stored for the given column: if you specify **DEFAULT** before the **VALUES** keyword, default values are stored for all columns in the table. **NULL** is stored for the column whose default value has not been defined.
- **ON DUPLICATE KEY UPDATE** : In case constraints are violated because a duplicated value for a column where **PRIMARY KEY** or **UNIQUE** attribute is defined is inserted, the value that makes constraints violated is changed into a specific value by performing the action specified in the **ON DUPLICATE KEY UPDATE** statement.

Example

```
CREATE TABLE a_tbl1(
id INT UNIQUE,
name VARCHAR,
phone VARCHAR DEFAULT '000-0000');

--insert default values with DEFAULT keyword before VALUES
INSERT INTO a_tbl1 DEFAULT VALUES;

--insert multiple rows
INSERT INTO a_tbl1 VALUES (1,'aaa', DEFAULT),(2,'bbb', DEFAULT);

--insert a single row specifying column values for all
INSERT INTO a_tbl1 VALUES (3,'ccc', '333-3333');

--insert two rows specifying column values for only
INSERT INTO a_tbl1(id) VALUES (4), (5);

--insert a single row with SET clauses
INSERT INTO a_tbl1 SET id=6, name='eee';
INSERT INTO a_tbl1 SET id=7, phone='777-7777';

SELECT * FROM a_tbl1;
=====
      id  name                phone
=====
      NULL NULL                '000-0000'
        1  'aaa'                '000-0000'
        2  'bbb'                '000-0000'
        3  'ccc'                '333-3333'
        4  NULL                '000-0000'
        5  NULL                '000-0000'
        6  'eee'                '000-0000'
        7  NULL                '777-7777'
```

INSERT ... SELECT Statement

Description

If you use the **SELECT** query in the **INSERT** statement, you can insert query results obtained from at least one table. The **SELECT** statement can be used in place of the **VALUES** keyword, or be included as a subquery in the column value list next to **VALUES**. If you specify the **SELECT** statement in place of the **VALUES** keyword, you can insert

multiple query result records into the column of the table at once. However, there should be only one query result record if the **SELECT** statement is specified in the column value list.

In this way, you can extract data from another table that satisfies a certain retrieval condition, and insert it into the target table by combining the **SELECT** statement with the **INSERT** statement.

Syntax

```
INSERT [INTO] table_name [(column_name, ...)]
  SELECT...
  [ON DUPLICATE KEY UPDATE column_name = expr, ... ]
```

Example

```
--creating an empty table which schema replicated from a tbl1
CREATE TABLE a_tbl2 LIKE a_tbl1;

--inserting multiple rows from SELECT query results
INSERT INTO a_tbl2 SELECT * FROM a_tbl1 WHERE id IS NOT NULL;

--inserting column value with SELECT subquery specified in the value list
INSERT INTO a_tbl2 VALUES(8, SELECT name FROM a_tbl1 WHERE name <'bbb', DEFAULT);

SELECT * FROM a_tbl2;
=====
      id  name                phone
=====
      1  'aaa'                '000-0000'
      2  'bbb'                '000-0000'
      3  'ccc'                '333-3333'
      4  NULL                 '000-0000'
      5  NULL                 '000-0000'
      6  'eee'                '000-0000'
      7  NULL                 '777-7777'
      8  'aaa'                '000-0000'
```

ON DUPLICATE KEY UPDATE Statement

Description

In a situation in which a duplicate value is inserted into a column for which the **UNIQUE** index or the **PRIMARY KEY** constraint has been set, you can update to a new value without outputting the error by specifying the **ON DUPLICATE KEY UPDATE** clause in the **INSERT** statement.

However, the **ON DUPLICATE KEY UPDATE** clause cannot be used in a table in which a trigger for **INSERT** or **UPDATE** has been activated, or in a nested **INSERT** statement.

Syntax

```
<INSERT ... VALUES statement>
<INSERT ... SET statement>
<INSERT ... SELECT statement>
  INSERT ...
  [ON DUPLICATE KEY UPDATE column_name = expr, ... ]
```

- *column_name* = *expr* : Specifies the name of the column whose value you want to change next to **ON DUPLICATE KEY UPDATE** and a new column value by using the equal sign.

Example

```
--creating a new table having the same schema as a tbl1
CREATE TABLE a_tbl3 LIKE a_tbl1;
INSERT INTO a_tbl3 SELECT * FROM a_tbl1 WHERE id IS NOT NULL and name IS NOT NULL;
SELECT * FROM a_tbl3;
=====
      id  name                phone
=====
      1  'aaa'                '000-0000'
      2  'bbb'                '000-0000'
      3  'ccc'                '333-3333'
      6  'eee'                '000-0000'
```

```
--insert duplicated value violating UNIQUE constraint
INSERT INTO a_tbl3 VALUES(2, 'bbb', '222-2222');

ERROR: Operation would have caused one or more unique constraint violations.

--insert duplicated value with specifying ON DUPLICATED KEY UPDATE clause
INSERT INTO a_tbl3 VALUES(2, 'bbb', '222-2222')
ON DUPLICATE KEY UPDATE phone = '222-2222';

SELECT * FROM a_tbl3 WHERE id=2;
      id  name          phone
=====
      2  'bbb'            '222-2222'
```

UPDATE

Description

You can update the column value of a record stored in the target table to a new one by using the **UPDATE** statement. Specify the name of the column to update and a new value in the **SET** clause, and specify the condition to be used to extract the record to be updated in the [WHERE Clause](#). You can also specify the number of records to be updated in the **LIMIT** clause. You can use the update with the **ORDER BY** clause if you want to maintain the execution order or lock order of triggers.

Syntax

```
UPDATE table_name SET column_name = {expr | DEFAULT} [, column_name = {expr | DEFAULT}...]
[WHERE search_condition]
[ORDER BY {col_name | expr}]
[LIMIT row_count]
```

- *table_name* : Specify the name of the table to be updated.
- *column_name* : Specify the columns to be updated.
- *expr* | **DEFAULT** : Specify a new value for the column, and specify an expression or the **DEFAULT** keyword as the value. You can also specify the **SELECT** query, which returns a single result record.
- *search_condition* : You can update the column value only for the record that satisfies the condition by specifying one in the [WHERE Clause](#).
- *col_name* | *expr* : Specifies a column used as a basis for the update order.
- *row_count* : Specify the number of records to be updated after the [LIMIT Clause](#). An integer greater than 0 can be specified.

Remark

One column can be updated only once in the same **UPDATE** statement.

Example

```
--creating a new table having all records copied from a tbl1
CREATE TABLE a_tbl5 AS SELECT * FROM a_tbl1;
SELECT * FROM a_tbl5 WHERE name IS NULL;
      id  name          phone
=====
      NULL NULL            '000-0000'
      4  NULL            '000-0000'
      5  NULL            '000-0000'
      7  NULL            '777-7777'

UPDATE a_tbl5 SET name='yyy', phone='999-9999' WHERE name IS NULL LIMIT 3;
SELECT * FROM a_tbl5;
      id  name          phone
=====
      NULL 'yyy'          '999-9999'
      1  'aaa'          '000-0000'
      2  'bbb'          '000-0000'
      3  'ccc'          '333-3333'
```



```

4 'yyy'          '999-9999'
5 'yyy'          '999-9999'
6 'eee'          '000-0000'
7 NULL          '777-7777'

-- using triggers, that the order in which the rows are updated is modified by the ORDER
BY clause.

CREATE TABLE t (i INT,d INT);
CREATE TRIGGER trigger1 BEFORE UPDATE ON t IF new.i < 10 EXECUTE PRINT 'trigger1 executed';
CREATE TRIGGER trigger2 BEFORE UPDATE ON t IF new.i > 10 EXECUTE PRINT 'trigger2 executed';
INSERT INTO t VALUES (15,1), (8,0), (11,2), (16,1), (6,0), (1311,3), (3,0);
UPDATE t SET i = i + 1 WHERE 1 = 1;

trigger2 executed
trigger1 executed
trigger2 executed
trigger2 executed
trigger1 executed
trigger2 executed
trigger1 executed

TRUNCATE TABLE t;
INSERT INTO t VALUES (15,1), (8,0), (11,2), (16,1), (6,0), (1311,3), (3,0);
UPDATE t SET i = i + 1 WHERE 1 = 1 ORDER BY i;

trigger1 executed
trigger1 executed
trigger1 executed
trigger2 executed
trigger2 executed
trigger2 executed
trigger2 executed

```

REPLACE

Description

The **REPLACE** statement is working like **INSERT**, but the difference is that it inserts a new record after deleting the existing record without displaying the error when a duplicate value is inserted into a column for which **PRIMARY KEY** and **UNIQUE** constraints have defined. You must have both **INSERT** and **DELETE** authorization to use the **REPLACE** statement, because it performs insertion or insertion after deletion operations.

The **REPLACE** statement determines whether a new record causes the duplication of **PRIMARY KEY** or **UNIQUE** index column values. Therefore, for performance reasons, it is recommended to use the **INSERT** statement for a table for which a **PRIMARY KEY** or **UNIQUE** index has not been defined. The **REPLACE** statement is an extension of the SQL standard. See the following regarding the use of this statement.

- The **REPLACE** statement cannot contain subqueries.
- The **REPLACE** statement cannot be used for tables for which an **INSERT** or **DELETE** trigger has been set.
- An assignment statement such as **SET col_name = col_name + 1** is not valid. Change such a statement to **SET col_name = DEFAULT(col_name) + 1**. Here, a non-NULL default value should be set for the *col_name* column.

Syntax

```

<REPLACE ... VALUES statement>
REPLACE [INTO] table_name [(column_name, ...)]
    {VALUES | VALUE}{(expr | DEFAULT), ...}[, ((expr | DEFAULT), ...),...]

<REPLACE ... SET statement>
REPLACE [INTO] table_name
    SET column_name = {expr | DEFAULT}[, column_name = {expr | DEFAULT},...]

<REPLACE ... SELECT statement>
REPLACE [INTO] table_name [(column_name, ...)]
    SELECT...

```

- *table_name* : Specify the name of the target table into which you want to insert a new record.

- *column_name* : Specify the name of the column into which you want to insert the value. If you omit to specify the column name, it is considered that all columns defined in the table have been specified. Therefore, you must specify the value for the column next to **VALUES**. If you do not specify all the columns defined in the table, a **DEFAULT** value is assigned to the non-specified columns; if the **DEFAULT** value is not defined, a NULL value is assigned.
- *expr* | **DEFAULT** : Specify values that correspond to the columns after **VALUES**. Expressions or the **DEFAULT** keyword can be specified as a value. At this time, the order and number of the specified column list must correspond to the column value list. The column value list for a single record is described in parentheses.

Example

```
--creating a new table having the same schema as a_tbl1
CREATE TABLE a_tbl4 LIKE a_tbl1;
INSERT INTO a_tbl4 SELECT * FROM a_tbl1 WHERE id IS NOT NULL and name IS NOT NULL;
SELECT * FROM a_tbl4;
=====
      id  name                phone
=====
      1  'aaa'                '000-0000'
      2  'bbb'                '000-0000'
      3  'ccc'                '333-3333'
      6  'eee'                '000-0000'

--insert duplicated value violating UNIQUE constraint
REPLACE INTO a_tbl4 VALUES(1, 'aaa', '111-1111'),(2, 'bbb', '222-2222');
REPLACE INTO a_tbl4 SET id=6, name='fff', phone=DEFAULT;

SELECT * FROM a_tbl4;
=====
      id  name                phone
=====
      3  'ccc'                '333-3333'
      1  'aaa'                '111-1111'
      2  'bbb'                '222-2222'
      6  'fff'                '000-0000'
```

DELETE

Description

You can delete records in the table by using the **DELETE** statement. You can specify delete conditions by combining the statement with the [WHERE Clause](#). If you want to limit the number of records to be deleted, you can do so by specifying the number of records to be deleted after the [LIMIT Clause](#). In this case, only **row_count** records are deleted even when the number of records satisfying the [WHERE Clause](#) exceeds **row_count**.

Syntax

```
DELETE FROM <table specification> [ WHERE <search condition> ] [LIMIT row count]

<table specification> ::= <table hierarchy> | ( <table hierarchy comma list > )

<table_hierarchy> ::= [ ONLY ] <table_name> |
                    ALL <table_name> [ EXCEPT <table_specification> ]
```

- *table_name* : Specifies the name of the table that contains the data to be deleted.
- *search_condition* : Delete only the data that meets the *search_condition* by using the [WHERE Clause](#). If it is not specified, all the data in the table will be deleted.
- *row_count* : Specify the number of records to be deleted after the [LIMIT Clause](#). An integer greater than 0 can be specified.

Example

```
CREATE TABLE a_tbl(
id INT NOT NULL,
phone VARCHAR(10));
INSERT INTO a_tbl VALUES(1,'111-1111'), (2,'222-2222'), (3, '333-3333'), (4, NULL), (5,
NULL);

DELETE FROM a_tbl WHERE phone IS NULL LIMIT 1;
```

```
--delete one record only from a tbl
SELECT * FROM a_tbl;
      id  phone
=====
          1  '111-1111'
          2  '222-2222'
          3  '333-3333'
          5  NULL

--delete all records from a_tbl
DELETE FROM a_tbl;
```

TRUNCATE

Description

You can delete all records in the specified table by using the **TRUNCATE** statement.

This statement internally delete first all indexes and constraints defined in a table and then deletes all records. Therefore, it performs the job faster than using the **DELETE FROM *table_name*** statement without a **WHERE** clause.

If the **PRIMARY KEY** constraint is defined in the table and this is referred by one or more **FOREIGN KEY**, it follows the **FOREIGN KEY ACTION**. If the **ON DELETE** action of **FOREIGN KEY** is **RESTRICT** or **NO_ACTION**, the **TRUNCATE** statement returns an error. If it is **CASCADE**, it deletes **FOREIGN KEY**. The **TRUNCATE** statement initializes the **AUTO INCREMENT** column of the table. Therefore, if data is inserted, the **AUTO INCREMENT** column value increases from the initial value.

Note To execute the **TRUNCATE** statement, the authorization of **ALTER**, **INDEX**, and **DELETE** is required on the table. For granting authorization, see [Granting Authorization](#).

Syntax

```
TRUNCATE [ TABLE ] <table_name>
```

- *table_name* : Specify the name of the table that contains the data to be deleted.

Example

```
CREATE TABLE a_tbl(A INT AUTO INCREMENT(3,10) PRIMARY KEY);
INSERT INTO a_tbl VALUES (NULL), (NULL), (NULL);
SELECT * FROM a_tbl;
      a
=====
          3
         13
         23

--AUTO_INCREMENT column value increases from the initial value after truncating the table
TRUNCATE TABLE a_tbl;
INSERT INTO a_tbl VALUES (NULL);
SELECT * FROM a_tbl;
      a
=====
          3
```

DO

Description

The **DO** statement executes the specified expression, but does not return the result. In general, the execution speed of the **DO** statement is higher than that of the **SELECT** expression statement, because the database server does not return the operation result or errors.

Syntax

```
DO expression
```

- *expression* : Specify an expression.

PREPARED STATEMENT

Overview

In general, the prepared statement is executed through the interface functions of JDBC, PHP, or ODBC; it can also be executed in the SQL level. The following SQL statements are provided for execution of prepared statement.

- Prepare the SQL statement to execute.

```
PREPARE stmt_name FROM preparable_stmt
```

- Execute the prepared statement.

```
EXECUTE stmt_name [USING value [, value] ...]
```

- Drop the prepared statement.

```
{DEALLOCATE | DROP} PREPARE stmt_name
```

PREPARE Statement

Description

The **PREPARE** statement prepares the query specified in *preparable_stmt* of the **FROM** clause and assigns the name to be used later when the SQL statement is referenced to *stmt_name*. See [EXECUTE Statement](#) for example.

```
PREPARE stmt_name FROM preparable_stmt
```

- *stmt_name* : The prepared statement is specified. If an SQL statement with the same *stmt_name* exists in the given client session, clear the existing prepared statement and prepare a new SQL statement. If the **PREPARE** statement is not executed properly due to an error in the given SQL statement, it is processed as if the *stmt_name* assigned to the SQL statement does not exist.
- *preparable_stmt* : You must use only one SQL statement. Multiple SQL statements cannot be specified. You can use a question mark (?) as a bind parameter in the *preparable_stmt* statement and it should not be enclosed with quotes.

Caution

The **PREPARE** statement is started by connecting an application to a server and will be maintained until the application terminates the connection. The connection maintained during this period is called a session. You can set the session time with the **session_state_timeout** parameter of **cubrid.conf**; the default value is **21600** seconds (=6 hours).

The data managed by the session includes the **PREPARE** statement, user-defined variables, the last ID inserted (**LAST_INSERT_ID**), and the number of rows affected by the statement (**ROW_COUNT**) that you execute at the end.

EXECUTE Statement

Description

The **EXECUTE** statement executes the prepared statement. You can bind the data value after the **USING** clause if a bind parameter (?) is included in the prepared statement. You cannot specify user-defined variables like an attribute in the **USING** clause. An value such as literal and an input parameter only can be specified.

Syntax

```
EXECUTE stmt_name [USING value [, value] ...]
```

- *stmt_name* : The name given to the prepared statement to be executed is specified. An error message is displayed if the *stmt_name* is not valid, or if the prepared statement does not exist.
- *value* : The data to bind is specified if there is a bind parameter in the prepared statement. The number and the order of the data must correspond to that of the bind parameter. If it does not, an error message is displayed.

Example

```

PREPARE st FROM 'SELECT 1 + ?';
EXECUTE st USING 4;
1+ ?:0
=====
5

SET @a=3;
EXECUTE st USING @a;
1+ ?:0
=====
4

PREPARE st FROM 'SELECT ? + ?';
EXECUTE st USING 1,3;
?:0 + ?:1
=====
4

PREPARE st FROM 'SELECT ? + ?';
EXECUTE st USING 'a','b';
?:0 + ?:1
=====
'ab'

PREPARE st FROM 'SELECT FLOOR(?)';
EXECUTE st USING '3.2';
floor( ?:0 )
=====
3.0000000000000000e+000

PREPARE st FROM 'SELECT FLOOR(?)';
EXECUTE st USING 3.2;
floor( ?:0 )
=====
3.0

```

DEALLOCATE PREPARE/DROP PREPARE Statements

Description

The **DEALLOCATE PREPARE** and **DROP PREPARE** statements are used interchangeably and they clear the prepared statement. All prepared statements are cleared automatically by the server when the client session is terminated even if the **DEALLOCATE PREPARE** or **DROP PREPARE** statement is not executed.

Syntax

```
{DEALLOCATE | DROP} PREPARE stmt_name
```

- *stmt_name* : The name given to the prepared statement to be cleared is specified. An error message is displayed if the *stmt_name* is not valid, or if the prepared statement does not exist.

Example

```
DEALLOCATE PREPARE stmt1;
```

SET

Description

The **SET** statement is the syntax that specifies user-defined variables and the method that you can use to store values.

You can create user-defined variables in two ways. One is to use the **SET** statement and the other is to use the assignment statement of user-defined variables within SQL statements. You can delete the user-defined variables that you defined with the **DEALLOCATE** or the **DROP** statements.

The user-defined variables are also called session variables as they are used for maintaining connections within one application. The user-defined variables are used within the part of a connection session, and the user-defined variables defined by an application cannot be accessed by other applications. When an application terminates connections, all variables will be removed automatically. The user-defined variables are limited to twenty per connection session for an application. If you already have twenty user-defined variables and want to define a new user-defined variable, you must remove some variables with the **DROP VARIABLE** statement.

You can use user-defined variables in most SQL statements. If you define user-defined variables and refer to them in one statement, the sequence is not guaranteed. That is, if you refer to the variables specified in the **SELECT** list of the **HAVING**, **GROUP BY** or **ORDER BY** clause, you may not get the values in the sequence you expect. You cannot also use user-defined variables as identifiers, such as column names or table names within SQL statements

The user-defined variables are not case-sensitive. The user-defined variable type can be one of the **SHORT**, **INTEGER**, **BIGINT**, **FLOAT**, **DOUBLE**, **NUMERIC**, **CHAR**, **VARCHAR**, **NCHAR**, **VARNCHAR**, **BIT** and **BIT VARYING**. Other types will be converted to the **VARCHAR** type.

```
SET @v1 = 1, @v2=CAST(1 AS BIGINT), @v3 = '123', @v4 = DATE'2010-01-01';

SELECT typeof(@v1), typeof(@v2), typeof(@v3), typeof(@v4);

  typeof(@v1)      typeof(@v2)      typeof(@v3)      typeof(@v4)
=====
'integer'         'bigint'         'character var'  'character varying (10)'
```

The user-defined variables can be changed when you define values.

```
SELECT @v := 1, typeof(@v1), @v1:='1', typeof(@v1);

  @v := 1      typeof(@v1)      @v1 := '1'      typeof(@v1)
=====
1             'integer'         '1'             'character (1)'
```

Syntax

```
<set_statement>
: <set statement>, <udf assignment>
| SET <udv_assignment>
;

<udv_assignment>
: @<name> = <expression>
| @<name> := <expression>
;

{DEALLOCATE|DROP} VARIABLE <variable name list>
<variable_name_list>
: <variable_name_list> ',' @<name>
```

- You must define the variable names with alphanumeric characters and underscores (_).
- When you define the variables within SQL statements, you should use the ':=' operator.

Example

The following example shows how to define the variable 'a' and assign a value 1 to it.

```
SET @a = 1;

SELECT @a;

  @a
=====
1
```

The following example shows how to count the number of rows in the **SELECT** statement by using the user-defined variable.

```
CREATE TABLE t (i INTEGER);
INSERT INTO t(i) VALUES (2), (4), (6), (8);

SET @a = 0;
```

```
SELECT @a := @a+1 AS row no, i FROM t;

  row_no          i
  =====
  1              2
  2              4
  3              6
  4              8
4 ROWS selected.
```

The following example shows how to use the user-defined variable as the input of bind parameter specified in the prepared statement.

```
SET @a:=3;

PREPARE stmt FROM 'SELECT i FROM t WHERE i < ?';
EXECUTE stmt USING @a;

      i
  =====
      2
```

The following example shows how to declare the user-defined variable by using the ':=' operator.

```
SELECT @a := 1, @user defined variable := 'user defined variable';
UPDATE t SET i = (@var := 1);
```

The following example shows how to delete the user-defined variable 'a' and 'user_defined_variable'.

```
DEALLOCATE VARIABLE @a, @user defined variable;
DROP VARIABLE @a, @user_defined_variable;
```

Caution

The user-defined variables that are defined by the **SET** statement are started by connecting an application to a server and will be maintained until the application terminates the connection. The connection maintained during this period is called a session. When an application terminates the connection or when there are no requests for a certain period of time, the session will expire, and the user-defined variables will be deleted as a result. You can set the session time with the **session_state_timeout** parameter of **cubrid.conf**; the default value is **21600** seconds (=6 hours).

The data managed by the session includes **PREPARE** statements, the user-defined variables, the last ID inserted (**LAST_INSERT_ID**) and the number of rows affected by the statement that you execute at the end (**ROW_COUNT**).

SHOW

SHOW TABLES Statement

Description

Displays the list of all the table names within a database. The name of the result column will be **tables_in_<database name>** and it will have one column. If you use the **LIKE** clause, you can search the table names matching this and if you use the **WHERE** clause, you can search table names with more general terms. **SHOW FULL TABLES** displays the second column, **table_type** together. The table must have the value, **BASE TABLE** and the view has the value, **VIEW**.

Syntax

```
SHOW [FULL] TABLES [LIKE 'pattern' | WHERE expr]
```

Example

The following is the result of executing the query with the demodb.

```
SHOW TABLES;
  Tables in demodb
  =====
```

```

'athlete'
'code'
'event'
'game'
'history'
'nation'
'olympic'
'participant'
'record'
'stadium'

SHOW FULL TABLES;
Tables in demodb      Table type
=====
'athlete'             'BASE TABLE'
'code'                'BASE TABLE'
'event'               'BASE TABLE'
'game'                'BASE TABLE'
'history'             'BASE TABLE'
'nation'              'BASE TABLE'
'olympic'             'BASE TABLE'
'participant'         'BASE TABLE'
'record'              'BASE TABLE'
'stadium'             'BASE TABLE'

SHOW FULL TABLES LIKE '%c%';
Tables_in_demodb      Table_type
=====
'code'                'BASE TABLE'
'olympic'             'BASE TABLE'
'participant'         'BASE TABLE'
'record'              'BASE TABLE'

SHOW FULL TABLES WHERE table type = 'BASE TABLE' and TABLES IN demodb LIKE
'%co%'; Tables in demodb      Table type
=====
'code'                'BASE TABLE'
'record'              'BASE TABLE'

```

SHOW COLUMN Statement

Description

Displays the column information of a table. You can use the **LIKE** clause to search the column names matching it. If you use the **WHERE** clause, you can search column names with more general terms like, "General Considerations for All **SHOW** Statements." If you use the **FULL** keyword, the additional information of a column will be displayed as follows:

- Field : Column name
- Type : Column data type
- Null : If you can store **NULL**, the value is YES and if not, it is NO
- Key : Whether a column has an index or not. If there is more than one key value in the given column of a table, this displays only the one that appears first in the order of PRI, UNI and MUL.
- If the key is a space, the column doesn't have an index, it is not the first column in the multiple column index or the index is non-unique.
- If the value is PRI, it is a primary key or the primary key of multiple columns.
- If the value is UNI, it is a unique index. (The unique index allows multiple NULL values but you can also set a NOT NULL constraint.)
- If the value is MUL, it is the first column of the non-unique index that allows the given value to be displayed in the column several times. If the column composes a composite unique index, the value will be MUL. The combination of column values can be unique but the value of each column can appear several times.
- Default : Default value defined in the column
- Extra : Additional information available about the given column. **AUTO_INCREMENT** The column attribute must have the auto_increment value.

SHOW FIELDS is the same command as **SHOW COLUMNS**.

The **DESCRIBE**(abbreviated **DESC**) statement and the **EXPLAIN** statement provide similar information to **SHOW COLUMNS**.

Syntax

```
SHOW COLUMNS {FROM | IN} tbl_name [LIKE 'pattern' | WHERE expr]
```

Example

The following is the result of a query for the demodb.

```
SHOW COLUMNS FROM athlete;
  Field      Type      Null      Key
Default      Extra
=====
'code'      'INTEGER' 'NO'      'PRI'
NULL        'auto_increment'
'name'      'STRING(40)' 'NO'      ''
NULL        ''
'gender'    'CHAR(1)' 'YES'     ''
NULL        ''
'nation_code' 'CHAR(3)' 'YES'     ''
NULL        ''
'event'     'STRING(30)' 'YES'     ''
NULL        ''

SHOW COLUMNS FROM athlete LIKE '%c%';
  Field      Type      Null      Key
Default      Extra
=====
'code'      'INTEGER' 'NO'      'PRI'
NULL        'auto_increment'
'nation code' 'CHAR(3)' 'YES'     ''
NULL        ''

SHOW COLUMNS FROM athlete WHERE "type" = 'INTEGER' and "key"='PRI' AND
extra='auto_increment';
  Field      Type      Null      Key
Default      Extra
=====
'code'      'INTEGER' 'NO'      'PRI'
NULL        'auto_increment'
```

SHOW INDEX Statement

Description

The **SHOW INDEX** statement displays the index information. The query must have the following columns:

- Table : Table Name
- Non_unique
- 0 : Duplicate data are not allowed
- 1 : Duplicate data are allowed
- Key_name : Index name
- Seq_in_index : Serial number of the column in the index. Starts from 1.
- Column_name : Column name
- Collation :Method of sorting columns in the index. 'A' means ascending and **NULL** means not sorted.
- Cardinality : The numerical value of measuring the unique values in the index. Higher cardinality increases the opportunity of using an index. This value is updated every time **SHOW INDEX** is executed.
- Sub_part : The number of bytes of the indexed characters if the columns are indexed partially. **NULL** if all columns are indexed.
- Packed : Shows how keys are packed. If they are not packed, it will be **NULL**.

- Null : YES if a column can include NULL, NO if not.
- Index_type : Index to be used (currently, only the BTREE is supported.)

Syntax

```
SHOW {INDEX | INDEXES | KEYS } {FROM | IN} tbl_name
```

Example

The following is the result of a query for the demodb.

```
SHOW INDEX IN athlete;
Table      Non unique  Key name      Seq in index  Column name    Collation      Cardin
ality     Sub part  Packed      Null      Index type
=====
'athlete'   0          'pk_athlete_code'  1           'code'         'A'            6677
          NULL      NULL        'NO'      'BTREE'

CREATE TABLE t1( i1 INTEGER , i2 INTEGER NOT NULL, i3 INTEGER UNIQUE, s1 VARCHAR(10), s2
VARCHAR(10), s3 VARCHAR(10) UNIQUE);

CREATE INDEX i_t1_i1 ON t1(i1 desc);
CREATE INDEX i_t1_s1 ON t1(s1(7));
CREATE INDEX i_t1_i1_s1 ON t1(i1,s1);
CREATE UNIQUE INDEX i_t1_i2_s2 ON t1(i2,s2);

SHOW INDEXES FROM t1;
Table      Non unique  Key_name      Seq_in_index  Column_name    Collation      Cardinality
          Sub part  Packed      Null      Index type
=====
't1'       0          'i_t1_i2_s2'  1           'i2'           'A'            0
          NULL      NULL        'NO'      'BTREE'
't1'       0          'i_t1_i2_s2'  2           's2'           'A'            0
          NULL      NULL        'YES'     'BTREE'
't1'       0          'u_t1_i3'     1           'i3'           'A'            0
          NULL      NULL        'YES'     'BTREE'
't1'       0          'u_t1_s3'     1           's3'           'A'            0
          NULL      NULL        'YES'     'BTREE'
't1'       1          'i_t1_i1'     1           'i1'           NULL           0
          NULL      NULL        'YES'     'BTREE'
't1'       1          'i_t1_i1_s1'  1           'i1'           'A'            0
          NULL      NULL        'YES'     'BTREE'
't1'       1          'i_t1_i1_s1'  2           's1'           'A'            0
          NULL      NULL        'YES'     'BTREE'
't1'       1          'i_t1_s1'     1           's1'           'A'            0
          7          NULL        'YES'     'BTREE'
```

SHOW GRANTS Statement

Description

The **SHOW GRANT** statement displays the permissions associated with the database user accounts.

Syntax

```
SHOW GRANTS FOR 'user'
```

Example

```
CREATE TABLE testgrant (id int);
CREATE USER user1;
GRANT INSERT,SELECT ON testgrant TO user1;

SHOW GRANTS FOR user1;
Grants for USER1
=====
'GRANT INSERT, SELECT ON testgrant TO USER1'
```

SHOW CREATE VIEW Statement

Description

The **SHOW CREATE VIEW** statement outputs the corresponding **CREATE VIEW** statement if view name is specified.

Syntax

```
SHOW CREATE VIEW view_name
```

Example

The following example shows the result of executing query on demodb.

```
SHOW CREATE VIEW "db_class";
View          Create View
=====
'db class'    'SELECT c.class name, CAST(c.owner.name AS VARCHAR(255)), CASE
c.class type WHEN 0 THEN 'CLASS' WHEN 1 THEN 'VCLASS' ELSE
              'UNKNOW' END, CASE WHEN MOD(c.is_system_class, 2) = 1 THEN 'YES' ELSE
'NO' END, CASE WHEN c.sub_classes IS NULL THEN 'NO'
              ELSE NVL((SELECT 'YES' FROM db partition p WHERE p.class of = c and
p.pname IS NULL), 'NO') END, CASE WHEN
              MOD(c.is_system_class / 8, 2) = 1 THEN 'YES' ELSE 'NO' END FROM
db class c WHERE CURRENT USER = 'DBA' OR {c.owner.name}
              SUBSETEQ ( SELECT SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name}),
SET{}) FROM db_user u, TABLE(groups) AS t(g) WHERE
              u.name = CURRENT_USER) OR {c} SUBSETEQ ( SELECT
SUM(SET{au.class of}) FROM db auth au WHERE {au.grantee.name} SUBSETEQ
              ( SELECT SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name}), SET{}) FROM
db user u, TABLE(groups) AS t(g) WHERE u.name =
              CURRENT_USER) AND au.auth_type = 'SELECT')
```

SHOW EXEC STATISTICS Statement

Description

The **SHOW EXEC STATISTICS** statement outputs statistics information of executing query.

- To start collecting **@collect_exec_stats** statistics information, configure the value of session variable (**@collect_exec_stats**) to 1; to stop, configure it to 0.
- It outputs the result of collecting statistics information.
- The **SHOW EXEC STATISTICS** statement outputs four part of data page statistics information; **data_page_fetches**, **data_page_dirties**, **data_page_ioreads**, and **data_page_iowrites**. The result columns consist of variable column (name of statistics name) and value column (value of statistics value). Once the **SHOW EXEC STATISTICS** statement is executed, the statistics information which has been accumulated is initialized.
- The **SHOW EXEC STATISTICS ALL** statement outputs all items of statistics information.

For details, see [Outputting Statistics Information of Server](#).

Syntax

```
SHOW EXEC STATISTICS [ ALL ]
```

Example

The following example shows the result of executing query on demodb.

```
-- set session variable @collect_exec_stats as 1 to start collecting the statistical
information.
SET @collect_exec_stats = 1;
SELECT * FROM db class;
...

-- print the statistical information of the data pages.
SHOW EXEC STATISTICS;
variable value
```

```

=====
'data_page_fetches' 332
'data_page_dirties' 85
'data_page_ioreads' 18
'data_page_iowrites' 28

SELECT * FROM db index;
...

-- print all of the statistical information.
SHOW EXEC STATISTICS ALL;

variable value
=====
'file_creates' 0
'file_removes' 0
'file_ioreads' 6
'file_iowrites' 0
'file_iosynches' 0
'data_page_fetches' 548
'data_page_dirties' 34
'data_page_ioreads' 6
'data_page_iowrites' 0
'data_page_victims' 0
'data_page_iowrites_for_replacement' 0
'log_page_ioreads' 0
'log_page_iowrites' 0
'log_append_records' 0
'log_checkpoints' 0
'log_wals' 0
'page_locks_acquired' 13
'object_locks_acquired' 9
'page_locks_converted' 0
'object_locks_converted' 0
'page_locks_re-requested' 0
'object_locks_re-requested' 8
'page_locks_waits' 0
'object_locks_waits' 0
'tran_commits' 3
'tran_rollbacks' 0
'tran_savepoints' 0
'tran_start_topops' 6
'tran_end_topops' 6
'tran_interrupts' 0
'btree_inserts' 0
'btree_deletes' 0
'btree_updates' 0
'btree_covered' 0
'btree_noncovered' 2
'btree_resumes' 0
'query_selects' 4
'query_inserts' 0
'query_deletes' 0
'query_updates' 0
'query_sscans' 2
'query_iscans' 4
'query_lscans' 0
'query_setscans' 2
'query_methscans' 0
'query_nljoins' 2
'query_mjoins' 0
'query_objfetches' 0
'network_requests' 88
'adaptive_flush_pages' 0
'adaptive_flush_log_pages' 0
'adaptive_flush_max_pages' 0
'network_requests' 88
'adaptive_flush_pages' 0
'adaptive_flush_log_pages' 0
'adaptive_flush_max_pages' 0

```

Transaction and Lock

Overview

This chapter covers issues relating to concurrency and restore, as well as how to commit or roll back transactions.

In multi-user environment, controlling access and update is essential to protect database integrity and ensure that a user's transaction will have accurate and consistent data. Without appropriate control, data could be updated incorrectly in the wrong order.

To control parallel operations on the same data, data must be locked during transaction, and unacceptable access to the data by another transaction must be blocked until the end of the transaction. In addition, any updates to a certain class must not be seen by other users before they are committed. If updates are not committed, all queries entered after the last commit or rollback of the update can be invalidated.

All examples introduced here were executed by csql. Outputs in the examples are displayed in *italics*.

Database Transaction

Overview

A database transaction groups CUBRID queries into a unit of consistency (for ensuring valid results in multi-user environment) and restore (for making the results of committed transactions permanent and ensuring that the aborted transactions are canceled in the database despite any failure, such as system failure). A transaction is a collection of one or more queries that access and update the database.

CUBRID allows multiple users to access the database simultaneously and manages accesses and updates to prevent inconsistency of the database. For example, if data is updated by one user, the changes made by this transaction are not seen to other users or the database until the updates are committed. This principle is important because the transaction can be rolled back without being committed.

You can delay permanent updates to the database until you are confident of the transaction result. Also, you can remove (**ROLLBACK**) all updates in the database if an unsatisfactory result or failure occurs in the application or computer system during the transaction. The end of the transaction is determined by the **COMMIT WORK** or **ROLLBACK WORK** statement. The **COMMIT WORK** statement makes all updates permanent while the **ROLLBACK WORK** statement cancels all updates entered in the transaction. For details, see the [Transaction Commit](#) and [Transaction Rollback](#) sections.

Transaction Commit

Description

Updates that occurred in the database are not permanently stored until the **COMMIT WORK** statement is executed. "Permanently stored" means that storing the updates in the disk is completed; The **WORK** keyword can be omitted. In addition, other users of the database cannot see the updates until they are permanently applied. For example, when a new row is inserted into a class, only the user who inserted the row can access it until the database transaction is committed. (If the **UNCOMMITTED INSTANCES** isolation level is used, other users can see inconsistent uncommitted updates.)

All locks obtained by the transaction are released after the transaction is committed.

Syntax

```
COMMIT [ WORK ]
```

Example

The database transaction in the following example consists of three **UPDATE** statements and changes three column values of seats from the stadium. To compare the results, check the current values and names before the update is made. Since, by default, csql runs in an autocommit mode, the following example is executed after setting the autocommit mode to off.

```
;autocommit off
AUTOCOMMIT IS OFF
SELECT name, seats
FROM stadium WHERE code IN (30138, 30139, 30140);
name                seats
=====
'Athens Olympic Tennis Centre'      3200
'Goudi Olympic Hall'                5000
'Vouliagmeni Olympic Centre'       3400
```

Let each **UPDATE** statement have the current seats of each stadium. To verify whether the command is correctly executed, you can retrieve the columns related to the seats table.

```
UPDATE stadium
SET seats = seats + 1000
WHERE code IN (30138, 30139, 30140);

SELECT name, seats FROM stadium WHERE code in (30138, 30139, 30140);
name                seats
=====
'Athens Olympic Tennis Centre'      4200
'Goudi Olympic Hall'                6000
'Vouliagmeni Olympic Centre'       4400
```

If the update is properly done, the changes can be semi-permanently fixed. In this time, use the **COMMIT WORK** as below:

```
COMMIT WORK;
```

Note In CUBRID, an auto-commit mode is set by default for transaction management.

An auto-commit mode is a mode that commits or rolls back all SQL statements. The transaction is committed automatically if the SQL is executed successfully, or is rolled back automatically if an error occurs. Such auto commit modes are supported in any interfaces.

In CCI, PHP, ODBC and OLE DB interfaces, you can configure auto-commit mode by using **CCI_DEFAULT_AUTOCOMMIT** upon startup of an application. If configuration on broker parameter is omitted, the default value is set to **ON**. To change auto-commit mode, use the following functions by interface: **cci_set_autocommit()** for CCI interface and **cubrid_set_autocommit()** for PHP interface.

For session command (**;**Autocommit****) which enables auto-commit configuration in CSQL Interpreter, see [Session Commands](#).

Transaction Rollback

Description

The **ROLLBACK WORK** statement removes all updates to the database since the last transaction. The **WORK** keyword can be omitted. By using this statement, you can cancel incorrect or unnecessary updates before they are permanently applied to the database. All locks obtained during the transaction are released.

Syntax

```
ROLLBACK [ WORK ]
```

Example

The following example shows two commands that modify the definition and the row of the same table.

```
ALTER TABLE code DROP s_name;
```

```
INSERT INTO code (s_name, f_name) VALUES ('D','Diamond');
```

```
ERROR: s_name is not defined.
```

The **INSERT** statement fails because the `s_name` column has been dropped in the definition of `code`. The data intended to be entered to the `code` table is correct, but the `s_name` column is wrongly removed. At this point, you can use the **ROLLBACK WORK** statement to restore the original definition of the `code` table.

```
ROLLBACK WORK;
```

Later, remove the `s_name` column by entering the **ALTER TABLE** again and modify the **INSERT** statement. The **INSERT** command must be entered again because the transaction has been aborted. If the database update has been done as intended, commit the transaction to make the changes permanent.

```
ALTER TABLE code drop s_name;
INSERT INTO code (f_name) VALUES ('Diamond');
COMMIT WORK;
```

Savepoint and Partial Rollback

Description

A savepoint is established during the transaction so that database changes made by the transaction are rolled back to the specified savepoint. Such operation is called a partial rollback. In a partial rollback, database operations (insert, update, delete, etc.) after the savepoint are rolled back, and transaction operations before it are not rolled back. The transaction can proceed with other operations after the partial rollback is executed. Or the transaction can be terminated with the **COMMIT WORK** or **ROLLBACK WORK** statement. Note that the savepoint does not commit the changes made by the transaction.

A savepoint can be created at a certain point of the transaction, and multiple savepoints can be used for a certain point. If a partial rollback is executed to a savepoint before the specified savepoint or the transaction is terminated with the **COMMIT WORK** or **ROLLBACK WORK** statement, the specified savepoint is removed. The partial rollback after the specified savepoint can be performed multiple times.

Savepoints are useful because intermediate steps can be created and named to control long and complicated utilities. For example, if you use a savepoint during the update operation, you don't need to perform all statements again when you made a mistake.

Syntax 1

```
SAVEPOINT mark
mark:
    a SQL identifier
    a host variable (starting with :)
```

If you make *mark* all the same value when you specify multiple savepoints in a single transaction, only the latest savepoint appears in the partial rollback. The previous savepoints remain hidden until the rollback to the latest savepoint is performed and then appears when the latest savepoint disappears after being used.

Syntax 2

```
ROLLBACK [ WORK ] [ TO [ SAVEPOINT ] mark ] [ ]
mark:
    a SQL identifier
    a host variable (starting with :)
```

Previously, the **ROLLBACK WORK** statement canceled all database changes added since the latest transaction. The **ROLLBACK WORK** statement is also used for the partial rollback that rolls back the transaction changes after the specified savepoint.

If *mark* value is not given, the transaction terminates canceling all changes including all savepoints created in the transaction. If *mark* value is given, changes after the specified savepoint are canceled and the ones before it are remained.

Example

The following example shows how to roll back part of the transaction.

First, set savepoints SP1 and SP2.

```
CREATE TABLE athlete2 (name VARCHAR(40), gender CHAR(1), nation code CHAR(3), event
VARCHAR(30));
INSERT INTO athlete2(name, gender, nation code, event)
VALUES ('Lim Kye-Sook', 'W', 'KOR', 'Hockey');
SAVEPOINT SP1;

SELECT * from athlete2;
INSERT INTO athlete2(name, gender, nation code, event)
VALUES ('Lim Jin-Suk', 'M', 'KOR', 'Handball');

SELECT * FROM athlete2;
SAVEPOINT SP2;

RENAME TABLE athlete2 AS sportsman;
SELECT * FROM sportsman;
ROLLBACK WORK TO SP2;
```

In the example above, the name change of the athlete2 table is rolled back by the partial rollback. The following example shows how to execute the query with the original name and examining the result.

```
SELECT * FROM athlete2;
DELETE FROM athlete2 WHERE name = 'Lim Jin-Suk';
SELECT * FROM athlete2;
ROLLBACK WORK TO SP2;
```

In the example above, deleting 'Lim Jin-Suk' is discarded by rollback work to SP2 command.

The following example shows how to roll back to SP1.

```
SELECT * FROM athlete2;
ROLLBACK WORK TO SP1;
SELECT * FROM athlete2;
COMMIT WORK;
```

Database Concurrency

If there are multiple users with read and write authorization to a database, possibility exists that more than one user will access the database simultaneously. Controlling access and update in multi-user environment is essential to protect database integrity and ensure that users and transactions should have accurate and consistent data. Without appropriate control, data could be updated incorrectly in the wrong order.

Like most commercial database systems, CUBRID adopts serializability, an element that is essential to maintaining data concurrency within the database. Serializability ensures no interference between transactions when multiple transactions are executed at once. It is guaranteed more with the higher isolation level. This principle is based on the assumption that database consistency is guaranteed as long as transaction is executed automatically. This will be covered in the [Lock Protocol](#) section in detail.

The transaction must ensure database concurrency, and each transaction must guarantee appropriate results. When multiple transactions are being executed at once, an event in transaction T1 should not affect an event in transaction T2. This means isolation. Transaction isolation level is the degree to which a transaction is separated from all other concurrent transactions. The higher isolation level means the lower interference from other transactions. The lower isolation level means the higher the concurrency. A database determines whether which lock is applied to tables and records based on these isolation levels. Therefore, can control the level of consistency and concurrency specific to a service by setting appropriate isolation level.

You can set an isolation level by using the [SET TRANSACTION ISOLATION LEVEL](#) statement or system parameters provided by CUBRID. For details, see [Concurrency/Lock Parameters](#).

The read operations that allow interference between transactions with isolation levels are as follows:

- **Dirty read** : A transaction T2 can read D' before a transaction T1 updates data D to D' and commits it.

- **Non-repeatable read** : A transaction T1 can read other value, if a transaction T2 updates data while data is retrieved in the transaction T2 multiple times.
- **Phantom read** : A transaction T1 can read E, if a transaction T2 inserts new record E while data is retrieved in the transaction T1 multiple times.

The default value of CUBRID isolation level is [REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES](#) (3).

Isolation Levels Provided by CUBRID

CUBRID Isolation Level (isolation_level)	Other DBMS Isolation Level (isolation_level)	DIRTY READ	UNREPEATABLE READ	PHANTOM READ	Schema Changes of the Table Being Retrieved
SERIALIZABLE (6)	SERIALIZABLE (4)	N	N	N	N
REPEATABLE READ CLASS with REPEATABLE READ INSTANCES (5)	REPEATABLE READ (3)	N	N	Y	N
REPEATABLE READ CLASS with READ COMMITTED INSTANCES (4)	READ COMMITTED (2)	N	Y	Y	N
REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES (3)	READ UNCOMMITTED (1)	Y	Y	Y	N
READ COMMITTED CLASS with READ COMMITTED INSTANCES (2)		N	Y	Y	Y
READ COMMITTED CLASS with READ UNCOMMITTED INSTANCES (1)		Y	Y	Y	Y

Lock Protocol

Overview

In the two-phase locking protocol used by CUBRID, a transaction obtains a shared lock before it reads an object, and an exclusive lock before it updates the object so that conflicting operations are not executed simultaneously.

If transaction T1 requires a lock, CUBRID checks if the requested lock conflicts with the existing one. If it does, transaction T1 enters a standby state and delays the lock. If another transaction T2 releases the lock, transaction T1 resumes and obtains it. Once the lock is released, the transaction do not require any more new locks.

Granularity Locking

CUBRID uses a granularity locking protocol to decrease the number of locks. In the granularity locking protocol, a database can be modeled as a hierarchy of lockable units: bigger locks have more granular locks.

For example, suppose that a database consists of multiple tables and each table consists of multiple instances. If the database is locked, all tables and instances are implicitly considered to be locked. A lock on a big unit results in less overhead, because only one lock needs to be managed. However, it leads to decreased concurrency because almost all concurrent transactions conflict with each other. The finer the granularity, the better the concurrency; it causes more overhead because more locks need to be managed. CUBRID selects a locking granularity level based on the operation being executed. For example, if a transaction retrieves all instances of a table, the entire tables will be locked, rather than each instance. If the transaction accesses a few instances of the table, the instances are locked individually.

If the locking granularities overlap, effects of a finer granularity are propagated in order to prevent conflicts. That is, if a shared lock is required on an instance of a table, an intention shared lock will be set on the table. If an exclusive lock is required on an instance of a table, an intention exclusive lock will be set on the table. An intention shared lock on a table means that a shared lock can be set on an instance of the table. An intention exclusive lock on a table means that a shared/exclusive lock can be set on an instance of the table. That is, if an intention shared lock on a table is allowed in one transaction, another transaction cannot obtain an exclusive lock on the table (for example, to add a new column). However, the second transaction may obtain a shared lock on the table. If an intention exclusive lock on the table is allowed in one transaction, another transaction cannot obtain a shared lock on the table (for example, a query on an instance of the tables cannot be executed because it is being changed).

A mechanism called lock escalation is used to limit the number of locks being managed. If a transaction has more than a certain number of locks (a number which can be changed by the `lock_escalation` system parameter), the system begins to require locks at the next higher level of granularity. This escalates the locks to a coarser level of granularity. CUBRID performs lock escalation when no transactions have a higher level of granularity in order to avoid a deadlock caused by lock conversion.

Lock Mode Types And Compatibility

CUBRID determines the lock mode depending on the type of operation to be performed by the transaction, and determines whether or not to share the lock depending on the mode of the lock preoccupied by another transaction. Such decisions concerning the lock are made by the system automatically. Manual assignment by the user is not allowed. To check the lock information of CUBRID, use the `cubrid lockdb db_name` command. For details, see [Checking Lock Status](#).

- **Shared lock (shared lock, S_LOCK)** : This lock is obtained before the read operation is executed on the object. It can be obtained by multiple transactions for the same object.
Transaction T1 obtains the shared lock first before it performs the read operation on a certain object X, and releases it immediately after it completes the operation even before transaction T1 is committed. Here, transaction T2 and T3 can perform the read operation on X concurrently, but not the update operation.
- **Exclusive lock (exclusive lock, X_LOCK)** : This lock is obtained before the update operation is executed on the object. It can only be obtained by one transaction.
Transaction T1 obtains the exclusive lock first before it performs the update operation on a certain object X, and does not release it until transaction T1 is committed even after the update operation is completed. Therefore, transaction T2 and T3 cannot perform the read operation as well on X before transaction T1 releases the exclusive lock.
- **Update lock (update lock, U_LOCK)** : This lock is obtained when the read operation is executed in the expression before the update operation is performed.
For example, when an UPDATE statement combined with a **WHERE** clause is executed, execute the operation by obtaining the update lock for each tuple and the exclusive lock only for the result tuples that satisfy the condition when performing index search or full scan search in the **WHERE** clause. The update lock is converted to an exclusive lock when the actual update operation is performed. It can be called a quasi-exclusive lock because it does not allow the read lock on the same object for another transaction.
- **Intention lock (intention lock)** : A lock that is set inherently in a higher-level object than X to protect the lock on the object X of a certain level.
For example, when a shared lock is requested for a certain tuple, prevent a situation from occurring in which the table is locked by another transaction by setting the intention shared lock as well on the table at the higher level in hierarchy. Therefore, the intention lock is not set on tuples at the lowest level, but is set on higher-level objects. The types of intention locks are as follows:
- **Intention shared lock (intention shared lock, IS_LOCK)** : If the intention shared lock is set on the table, which is the higher-level object, as the result of the shared lock set on a certain tuple, another transaction cannot perform operations such as changing the schema of the table (e.g. adding a column or changing the table name) or updating all tuples. However updating some tuples or viewing all tuples is allowed.
- **Intention exclusive lock (intention exclusive lock, IX_LOCK)** : If the intention exclusive lock is set on the table, which is the higher-level object, as the result of the exclusive lock set on a certain tuple, another transaction cannot perform operations such as changing the schema of the table, updating or viewing all tuples. However updating some tuples is allowed.
- **Shared with intent exclusive (shared with intent exclusive, SIX_LOCK)** : This lock is set on the higher-level object inherently to protect the shared lock set on all objects at the lower hierarchical level and the intention exclusive lock on some object at the lower hierarchical level.

Once the shared intention exclusive lock is set on a table, another transaction cannot change the schema of the table, update all/some tuples or view all tuples. However, viewing some tuples is allowed.

The following table briefly shows the lock compatibility between the locks described below. Compatibility means that the lock requester can obtain a lock while the lock holder is keeping the lock obtained for the object X. N/a means 'not applicable'.

Lock Compatibility

	Lock Holder							
	NULL_LOCK	IS_LOCK	S_LOCK	IX_LOCK	SIX_LOCK	U_LOCK	X_LOCK	
Lock Requester (lock requester)	NULL_LOCK	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
	IS_LOCK	TRUE	TRUE	TRUE	TRUE	TRUE	N/A	FALSE
	S_LOCK	TRUE	TRUE	TRUE	FALSE	FALSE	FALSE	FALSE
	IX_LOCK	TRUE	TRUE	FALSE	TRUE	FALSE	N/A	FALSE
	SIX_LOCK	TRUE	TRUE	FALSE	FALSE	FALSE	N/A	FALSE
	U_LOCK	TRUE	N/A	TRUE	N/A	N/A	FALSE	FALSE
	X_LOCK	TRUE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

- **NULL_LOCK** : No lock

Example

session 1

```

;autocommit off
AUTOCOMMIT IS OFF
set transaction isolation level 4;
Isolation level set to:
REPEATABLE READ SCHEMA, READ COMMITTED
INSTANCES.
    
```

session 2

```

;autocommit off
AUTOCOMMIT IS OFF
set transaction isolation level 4;
Isolation level set to:
REPEATABLE READ SCHEMA, READ
COMMITTED INSTANCES.

/*
C:\CUBRID> cubrid lockdb demodb

*** Lock Table Dump ***
...

Object Lock Table:
    Current number of
objects which are locked      = 0
    Maximum number of
objects which can be locked =
10000
...
*/
    
```

```

SELECT nation_code, gold FROM participant
WHERE nation_code='USA';
    nation_code      gold
=====
'USA'                36
'USA'                37
'USA'                44
'USA'                37
'USA'                36

/*
C:\CUBRID> cubrid lockdb demodb
    
```

```

*** Lock Table Dump ***
...

Object type: Root class.
LOCK HOLDERS:
  Tran_index = 2, Granted_mode
= IS_LOCK, Count = 1, Nsubgranules = 1

Object type: Class = participant.
LOCK HOLDERS:
  Tran_index = 2, Granted_mode
= IS_LOCK, Count = 2, Nsubgranules = 0
*/

UPDATE participant SET gold =
11 WHERE nation_code = 'USA' ;

SELECT nation_code, gold FROM participant
WHERE nation_code='USA';

/* no results until transaction 2 releases a
lock

C:\CUBRID> cubrid lockdb demodb
*** Lock Table Dump ***
...

Object type: Instance of class
( 0| 551| 7) = participant.
LOCK HOLDERS:
  Tran index = 3, Granted mode
= X_LOCK, Count = 2

...

Object type: Root class.
LOCK HOLDERS:
  Tran index = 3, Granted mode
= IX_LOCK, Count = 1, Nsubgranules = 3
NON_2PL_RELEASED:
  Tran_index = 2, Non_2_phase_lock
= IS_LOCK

...

Object type: Class = participant.
LOCK HOLDERS:
  Tran_index = 3, Granted_mode
= IX_LOCK, Count = 3, Nsubgranules = 5
  Tran_index = 2, Granted_mode
= IS_LOCK, Count = 2, Nsubgranules = 0
*/

COMMIT;

Current transaction has been
committed.

nation_code      gold
=====
'USA'            11
'USA'            11
'USA'            11
'USA'            11
'USA'            11

```

```

/*
C:\CUBRID> cubrid lockdb demodb
...

Object type: Root class.
LOCK HOLDERS:
    Tran_index = 2, Granted_mode
= IS_LOCK, Count = 1, Nsubgranules = 1

Object type: Class = participant.
LOCK HOLDERS:
    Tran_index = 2, Granted_mode
= IS_LOCK, Count = 3, Nsubgranules = 0
...
*/

COMMIT;

Current transaction has been committed.

/*
C:\CUBRID> cubrid lockdb demodb
...

Object Lock Table:
    Current number of objects which are
locked = 0
    Maximum number of objects which can
be locked = 10000
*/

```

Transaction Deadlock

A deadlock is a state in which two or more transactions wait at once for another transaction's lock to be released. CUBRID resolves the problem by rolling back one of the transactions because transactions in a deadlock state will hinder the work of another transaction. The transaction to be rolled back is usually the transaction which has made the least updates; it is usually the one that started more recently. As soon as a transaction is rolled back, the lock held by the transaction is released and other transactions in a deadlock are permitted to proceed.

It is impossible to predict such deadlocks, but it is recommended that you reduce the range to which lock is applied by setting the index, shortening the transaction, or setting the transaction isolation level as low in order to decrease such occurrences.

Note that if you configure the value of **error_log_level**, which indicates the severity level, to NOTIFICATION, information on lock is stored in error log file of server upon deadlock occurrences.

In the following error log file, (1) indicates a table name which causes deadlock state and (2) indicates an index name.

```

demodb 20111102 1811.err
...
OID = -532| 520| 1
(1) Object type: Index key of class ( 0| 417| 7) = tbl.
BTID = 0| 123| 530
(2) Index Name : i tbl coll
Total mode of holders = NS LOCK, Total mode of waiters = NULL LOCK.
Num holders= 1, Num blocked-holders= 0, Num waiters= 0
LOCK HOLDERS:
Tran index = 2, Granted mode = NS LOCK, Count = 1
...

```

Example

session 1

```
;autocommit off
AUTOCOMMIT IS OFF
set transaction isolation level 6;
Isolation level set to:
SERIALIZABLE
```

```
CREATE TABLE lock_tbl(host_year
integer, nation_code char(3));
INSERT INTO lock_tbl VALUES (2004,
'KOR');
INSERT INTO lock tbl VALUES (2004,
'USA');
INSERT INTO lock tbl VALUES (2004,
'GER');
INSERT INTO lock_tbl VALUES (2008,
'GER');
COMMIT;

SELECT * FROM lock_tbl;
  host_year  nation_code
=====
          2004  'KOR'
          2004  'USA'
          2004  'GER'
          2008  'GER'
```

session 2

```
;autocommit off
AUTOCOMMIT IS OFF
set transaction isolation level 6;
Isolation level set to:
SERIALIZABLE
```

```
SELECT * FROM lock_tbl;
  host_year  nation_code
=====
          2004  'KOR'
          2004  'USA'
          2004  'GER'
          2008  'GER'
```

```
DELETE FROM lock_tbl WHERE
host_year=2008;

/* no result until transaction 2
releases a lock

C:\CUBRID> cubrid lockdb demodb
*** Lock Table Dump ***
...

Object type: Class = lock_tbl.
LOCK HOLDERS:
  Tran_index = 2, Granted mode
= S LOCK, Count = 2, Nsubgranules
= 0

BLOCKED LOCK HOLDERS:
  Tran_index = 1, Granted mode
= S LOCK, Count = 3, Nsubgranules
= 0
  Blocked_mode = SIX_LOCK
  Start waiting at = Fri Feb 12
14:22:58 2010
  Wait_for_nsecs = -1

*/
```

```
INSERT INTO lock_tbl VALUES (2004,
'AUS');
```

```
ERROR: Your transaction (index 1,
dba@ 090205|4760) has been
unilaterally aborted by the system.

/*
System rolled back the transaction 1
to resolve a deadlock

C:\CUBRID> cubrid lockdb demodb
*** Lock Table Dump ***

Object type: Class = lock_tbl.

LOCK HOLDERS:
    Tran index =    2, Granted mode =
SIX LOCK, Count =    3, Nsubgranules
=    0
*/
```

Transaction Timeout

CUBRID provides the lock timeout feature, which sets the waiting time for the lock until the transaction lock setting is allowed.

If the lock is allowed within the lock timeout, CUBRID rolls back the transaction and outputs an error message when the timeout has passed. If a transaction deadlock occurs within the lock timeout, CUBRID rolls back the transaction whose waiting time is closest to the timeout.

Setting the Lock Timeout

Description

The system parameter `lock_timeout_in_secs` in the `$CUBRID/conf/cubrid.conf` file or the `SET TRANSACTION` statement sets the timeout (in seconds) during which the application will wait for the lock and rolls back the transaction and outputs an error message when the specified time has passed. The default value of the `lock_timeout_in_secs` parameter is `-1`, which means the application will wait indefinitely until the transaction lock is allowed. Therefore, the user can change this value depending on the transaction pattern of the application. If the lock timeout value has been set to `0`, an error message will be displayed as soon as a lock occurs.

Syntax

```
SET TRANSACTION LOCK TIMEOUT timeout_spec [ ; ]
timeout_spec:
- INFINITE
- OFF
- unsigned integer
- variable
```

- **INFINITE** : Wait indefinitely until the transaction lock is allowed. Has the same effect as setting the system parameter `lock_timeout_in_secsto -1`.
- **OFF** : Do not wait for the lock, but roll back the transaction and display an error message. Has the same effect as setting the system parameter `lock_timeout_in_secsto 0`.
- *unsigned integer* : Set in seconds. Wait for the transaction lock for the specified time period.
- *variable* : A variable can be specified. Wait for the transaction lock for the value stored by the variable.

Example 1

```
vi $CUBRID/conf/cubrid.conf
...
lock_timeout_in_secs = 10
```

...

Example 2

```
SET TRANSACTION LOCK TIMEOUT 10;
```

Checking the Lock Timeout

Description

You can check the lock timeout set for the current application by using the **GET TRANSACTION** statement, or store this value in a variable.

Syntax

```
GET TRANSACTION LOCK TIMEOUT [ { INTO | TO } variable ] [ ; ]
```

Example

```
GET TRANSACTION LOCK TIMEOUT;
      Result
=====
1.000000e+001
```

Checking and Handling Lock Timeout Error Message

The following message is displayed if lock timeout occurs in a transaction that has been waiting for another transaction's lock to be released. To configuration level information to be displayed in details, see the description of **lock_timeout_message_type** in [Concurrency/Lock Parameters](#).

```
ERROR: Your transaction (index 3, cub user@cubs006.cub|15668) timed out waiting
on X LOCK lock on instance 0|636|34 of class participant. You are waiting for
user(s) to finish.
```

- Your transaction(index 3 ...): This means that the index of the transaction that was rolled back due to timeout while waiting for the lock is 3. The transaction index is a number that is sequentially assigned when the client connects to the database server. You can also check this number by executing the **cubrid lockdb** utility.
- (...cub_user@cubs006.cub|15668): cub_user is the login ID of the client and the part after @ is the name of the host where the client was running. The part after| is the process ID (PID) of the client.
- X_LOCK: This means the exclusive lock set on the object to perform data update. For details, see [Lock Mode Types And Compatibility](#).
- Instance 0|636|34 of class participant: This means that **X_LOCK** has been set on a specific row in the table named participant and the OID (unique ID assigned to the given object) of the row is 0|636|34.

That is, the above lock error message can be interpreted as meaning that "Because another client is holding **X_LOCK** on a specific row in the participant table, transaction 3 which running on the host cubs006.cub waited for the lock and was rolled back as the timeout has passed."

If you want to check the lock information of the transaction specified in the error message, you can do so by using the **cubrid lockdb** utility to search for the OID value (ex: 0|636|34) of a specific row where the **X_LOCK** is set currently to find the transaction ID currently holding the lock, the client program name and the process ID (PID). For details, see [Checking Lock Status](#). You can also check the transaction lock information in the CUBRID Manager.

You can organize the transactions by checking uncommitted queries through the SQL log after checking the transaction lock information in the manner described above. For information on checking the SQL log, see [Broker Log](#).

Also, you can forcefully stop problematic transactions by using the **cubrid killtran** utility. For details, see [Killing Transactions](#).

Transaction Isolation Level

Overview

The transaction isolation level is determined based on how much interference occurs. The more isolation means the less interference from other transactions and more serializable. The less isolation means the more interference from other transactions and higher level of concurrency. You can control the level of consistency and concurrency specific to a service by setting appropriate isolation level.

Note A transaction can be restored in all supported isolation levels because updates are not committed before the end of the transaction.

SET TRANSACTION ISOLATION LEVEL

Description

You can set the level of transaction isolation by using `isolation_level` and the `SET TRANSACTION` statement in the `$CUBRID/conf/cubrid.conf`. The level of `REPEATABLE READ CLASS` and `READ UNCOMMITTED INSTANCES` are set by default, which indicates the level 3 through level 1 to 6. For details, see [Database Concurrency](#).

Syntax

```
SET TRANSACTION ISOLATION LEVEL isolation_level_spec [ ; ]
isolation_level_spec:
_ SERIALIZABLE
_ CURSOR STABILITY
_ isolation_level [ { CLASS | SCHEMA } [ , isolation_level INSTANCES ] ]
_ isolation_level [ INSTANCES [ , isolation_level { CLASS | SCHEMA } ] ]
_ variable
isolation_level:
_ REPEATABLE READ
_ READ COMMITTED
_ READ UNCOMMITTED
```

Example 1

```
vi $CUBRID/conf/cubrid.conf
...
isolation_level = 1
...

or

isolation_level = "TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE"
```

Example 2

```
SET TRANSACTION ISOLATION LEVEL 1;
or
SET TRANSACTION ISOLATION LEVEL READ COMMITTED CLASS, READ UNCOMMITTED INSTANCES;
```

The following table shows the isolation levels from 1 to 6. It consists of table schema (row) and isolation level. For the unsupported isolation level, see [Unsupported Combination of Isolation Level](#).

Levels of Isolation Supported by CUBRID

Name	Description
SERIALIZABLE (6)	In this isolation level, problems concerning concurrency (e.g. dirty read, non-repeatable read, phantom read, etc.) do not occur.
REPEATABLE READ CLASS with REPEATABLE READ INSTANCES (5)	Another transaction T2 cannot update the schema of table A while transaction T1 is viewing table A. Transaction T1 may experience phantom read for the record R that was inserted by another transaction T2 when it is repeatedly retrieving a specific record.

REPEATABLE READ CLASS with READ COMMITTED INSTANCES (or CURSOR STABILITY) (4)	Another transaction T2 cannot update the schema of table A while transaction T1 is viewing table A. Transaction T1 may experience R read (non-repeatable read) that was updated and committed by another transaction T2 when it is repeatedly retrieving the record R.
REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES (3)	Default isolation level. Another transaction T2 cannot update the schema of table A while transaction T1 is viewing table A. Transaction T1 may experience R' read (dirty read) for the record that was updated but not committed by another transaction T2.
READ COMMITTED CLASS with READ COMMITTED INSTANCES (2)	Transaction T1 may experience A' read (non-repeatable read) for the table that was updated and committed by another transaction T2 while it is viewing table A repeatedly. Transaction T1 may experience R' read (non-repeatable read) for the record that was updated and committed by another transaction T2 while it is retrieving the record R repeatedly.
READ COMMITTED CLASS with READ UNCOMMITTED INSTANCES (1)	Transaction T1 may experience A' read (non-repeatable read) for the table that was updated and committed by another transaction T2 while it is repeatedly viewing table A. Transaction T1 may experience R' read (dirty read) for the record that was updated but not committed by another transaction T2.

If the transaction level is changed in an application while a transaction is executed, the new level is applied to the rest of the transaction being executed. Therefore, some object locks that have already been obtained may be released during the transaction while the new isolation level is applied. For this reason, it is recommended that the transaction isolation level be modified when the transaction starts (after commit, rollback or system restart) because an isolation level which has already been set does not apply to the entire transaction, but can be changed during the transaction.

GET TRANSACTION ISOLATION LEVEL

Description

You can assign the current isolation level to *variable* by using the **GET TRANSACTION** statement. The following is a statement that verifies the isolation level. *variable*.

Syntax

```
GET TRANSACTION ISOLATION LEVEL [ { INTO | TO } variable ] [ ; ]
```

Example

```
GET TRANSACTION ISOLATION LEVEL;
      Result
=====
READ COMMITTED SCHEMA, READ UNCOMMITTED INSTANCES
```

SERIALIZABLE

The highest isolation level (6). Problems concerning concurrency (e.g. dirty read, non-repeatable read, phantom read, etc.) do not occur.

The following are the rules of this isolation level:

- Transaction T1 cannot read or modify the record being updated by another transaction T2.
- Transaction T1 cannot read or modify the record being viewed by another transaction T2.
- Another transaction T2 cannot insert a new record into table A while transaction T1 is retrieving the records of table A.

This isolation level uses a two-phase locking protocol for shared and exclusive lock: the lock is held until the transaction ends even after the operation has been executed.

Example

The following example shows that another transaction cannot access the table or record while one transaction is reading or updating the object when the transaction level of the concurrent transactions is **SERIALIZABLE**.

session 1	session 2
<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 6; Isolation level set to: SERIALIZABLE</pre>	<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 6; Isolation level set to: SERIALIZABLE</pre>
<pre>--creating a table CREATE TABLE isol6_tbl(host_year integer, nation_code char(3)); INSERT INTO isol6_tbl VALUES (2008, 'AUS'); COMMIT;</pre>	<pre>--selecting records from the table SELECT * FROM isol6 tbl WHERE nation_code = 'AUS'; host_year nation_code ===== 2008 'AUS'</pre>
<pre>INSERT INTO isol6 tbl VALUES (2004, 'AUS'); /* unable to insert a row until the tran 2 committed */</pre>	<pre>COMMIT;</pre>
<pre>COMMIT;</pre>	<pre>SELECT * FROM isol6_tbl WHERE nation_code = 'AUS'; /* unable to select rows until tran 1 committed */ host_year nation_code ===== 2008 'AUS' 2004 'AUS'</pre>
<pre>DELETE FROM isol6_tbl WHERE nation_code = 'AUS' and host_year=2008; /* unable to delete rows until tran 2 committed */</pre>	<pre>COMMIT;</pre>
<pre>COMMIT;</pre>	<pre>SELECT * FROM isol6 tbl WHERE nation_code = 'AUS'; /* unable to select rows until tran 1 committed */ host_year nation_code ===== 2004 'AUS'</pre>
<pre>ALTER TABLE isol6_tbl ADD COLUMN gold INT; /* unable to alter the table schema</pre>	<pre>/* repeatable read is ensured while tran_1 is altering table schema */ SELECT * FROM isol6_tbl WHERE</pre>

until tran 2 committed */	<pre>nation_code = 'AUS'; host_year nation_code ===== 2004 'AUS'</pre>
COMMIT;	COMMIT;
COMMIT;	<pre>SELECT * FROM isol6 tbl WHERE nation_code = 'AUS'; /* unable to access the table until tran_1 committed */</pre>
COMMIT;	<pre>host_year nation_code gold ===== 2004 'AUS' NULL</pre>

REPEATABLE READ CLASS with REPEATABLE READ INSTANCES

A relatively high isolation level (5). A dirty or non-repeatable read does not occur, but a phantom read may.

The following are the rules of this isolation level:

- Transaction T1 cannot read or modify the record being updated by another transaction T2.
- Transaction T1 cannot read or modify the record being viewed by another transaction T2.
- Another transaction T2 can insert a new record into table A while transaction T1 is retrieving records of table A. However, transaction T1 and T2 cannot set the lock on the same record.

This isolation level uses a two-phase locking protocol.

Example

The following example shows that phantom read may occur because another transaction can add a new record while one transaction is performing the object read when the transaction level of the concurrent transactions is **REPEATABLE READ CLASS** with **REPEATABLE READ INSTANCES**.

session 1	session 2
<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 5; Isolation level set to: REPEATABLE READ SCHEMA, REPEATABLE READ INSTANCES.</pre>	<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 5; Isolation level set to: REPEATABLE READ SCHEMA, REPEATABLE READ INSTANCES.</pre>
<pre>--creating a table CREATE TABLE isol5_tbl(host_year integer, nation_code char(3)); CREATE UNIQUE INDEX on isol5_tbl(nation_code, host_year); INSERT INTO isol5 tbl VALUES (2008, 'AUS'); INSERT INTO isol5 tbl VALUES (2004, 'AUS'); COMMIT;</pre>	<pre>--selecting records from the table SELECT * FROM isol5 tbl WHERE nation_code='AUS'; host_year nation_code ===== 2004 'AUS' 2008 'AUS'</pre>

<pre>INSERT INTO isol5_tbl VALUES (2004, 'KOR'); INSERT INTO isol5_tbl VALUES (2000, 'AUS'); /* able to insert new rows only when locks are not conflicted */</pre>	
	<pre>SELECT * FROM isol5_tbl WHERE nation_code='AUS'; /* phantom read may occur when tran 1 committed */</pre>
<pre>COMMIT;</pre>	<pre>host_year nation_code ===== 2000 'AUS' 2004 'AUS' 2008 'AUS'</pre>
<pre>DELETE FROM isol5_tbl WHERE nation_code = 'AUS' and host_year=2008; /* unable to delete rows until tran 2 committed */</pre>	
<pre>COMMIT;</pre>	
	<pre>SELECT * FROM isol5_tbl WHERE nation_code = 'AUS'; /* unable to select rows until tran 1 committed */</pre>
<pre>COMMIT;</pre>	<pre>host_year nation_code ===== 2000 'AUS' 2004 'AUS'</pre>
<pre>ALTER TABLE isol5_tbl ADD COLUMN gold INT; /* unable to alter the table schema until tran 2 committed */</pre>	
	<pre>/* repeatable read is ensured while tran_1 is altering table schema */ SELECT * FROM isol5_tbl WHERE nation_code = 'AUS'; host_year nation_code ===== 2000 'AUS' 2004 'AUS'</pre>
<pre>COMMIT;</pre>	
	<pre>SELECT * FROM isol5_tbl WHERE nation_code = 'AUS'; /* unable to access the table until tran_1 committed */</pre>
<pre>COMMIT;</pre>	<pre>host_year nation_code gold ===== 2000 'AUS' NULL 2004 'AUS' NULL</pre>

REPEATABLE READ CLASS with READ COMMITTED INSTANCES

A relatively low isolation level (4). A dirty read does not occur, but non-repeatable or phantom read may. That is, transaction T1 can read another value because insert or update by transaction T2 is allowed while transaction T1 is repeatedly retrieving one object.

The following are the rules of this isolation level:

- Transaction T1 cannot read the record being updated by another transaction T2.
- Transaction T1 can update/insert record to the table being viewed by another transaction T2.
- Transaction T1 cannot change the schema of the table being viewed by another transaction T2.

This isolation level uses a two-phase locking protocol for an exclusive lock. A shared lock on a row is released immediately after it is read; however, an intention lock on a table is released when a transaction terminates to ensure repeatable read on the schema.

Example

The following example shows that a phantom or non-repeatable read may occur because another transaction can add or update a record while one transaction is performing the object read but repeatable read for the table schema update is ensured when the transaction level of the concurrent transactions is **REPEATABLE READ CLASS with READ COMMITTED INSTANCES**.

session 1	session 2
<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 4; Isolation level set to: REPEATABLE READ SCHEMA, READ COMMITTED INSTANCES.</pre>	<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 4; Isolation level set to: REPEATABLE READ SCHEMA, READ COMMITTED INSTANCES.</pre>
<pre>--creating a table CREATE TABLE isol4 tbl(host year integer, nation_code char(3)); INSERT INTO isol4_tbl VALUES (2008, 'AUS'); COMMIT;</pre>	<pre>--selecting records from the table SELECT * FROM isol4_tbl; host_year nation_code ===== 2008 'AUS'</pre>
<pre>INSERT INTO isol4_tbl VALUES (2004, 'AUS'); INSERT INTO isol4 tbl VALUES (2000, 'NED'); /* able to insert new rows even if tran 2 uncommitted */</pre>	<pre>SELECT * FROM isol4_tbl; /* phantom read may occur when tran 1 committed */</pre>
<pre>COMMIT;</pre>	<pre> host_year nation_code ===== 2008 'AUS'</pre>

	<pre> 2004 'AUS' 2000 'NED' </pre>
<pre> INSERT INTO isol4_tbl VALUES (1994, 'FRA'); </pre>	
	<pre> SELECT * FROM isol4_tbl; /* unrepeatable read may occur when tran 1 committed */ </pre>
<pre> DELETE FROM isol4_tbl WHERE nation_code = 'AUS' and host_year=2008; /* able to delete rows while tran 2 is selecting rows*/ </pre>	
<pre> COMMIT; </pre>	<pre> host_year nation_code ===== 2004 'AUS' 2000 'NED' 1994 'FRA' </pre>
<pre> ALTER TABLE isol4_tbl ADD COLUMN gold INT; /* unable to alter the table schema until tran 2 committed */ </pre>	
	<pre> /* repeatable read is ensured while tran_1 is altering table schema */ SELECT * FROM isol4_tbl; host_year nation_code ===== 2004 'AUS' 2000 'NED' 1994 'FRA' </pre>
	<pre> COMMIT; </pre>
	<pre> SELECT * FROM isol4_tbl; /* unable to access the table until tran_1 committed */ </pre>
<pre> COMMIT; </pre>	<pre> host_year nation_code gold ===== 2004 'AUS' NULL 2000 'NED' NULL 1994 'FRA' NULL </pre>

REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES

The default isolation of CUBRID (3). The concurrency level is high. A dirty, non-repeatable or phantom read may occur for the tuple, but repeatable read is ensured for the table. That is, transaction T2 can read an object while transaction T1 is updating one.

The following are the rules of this isolation level:

- Transaction T1 can read the record being updated by another transaction T2.
- Transaction T1 can update/insert record to the table being viewed by another transaction T2.
- Transaction T1 cannot change the schema of the table being viewed by another transaction T2.

This isolation level uses a two-phase locking protocol for an exclusive and update lock. However, the shared lock on the tuple is released immediately after it is retrieved. The intention lock on the table is released when the transaction ends to ensure repeatable reads.

Example

The following example shows that another transaction can read dirty data uncommitted by one transaction but repeatable reads are ensured for table schema update when the transaction level of the concurrent transactions is **REPEATABLE READ CLASS** with **READ UNCOMMITTED INSTANCES**.

session 1	session 2
<pre> ;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 3; Isolation level set to: REPEATABLE READ SCHEMA, READ UNCOMMITTED INSTANCES. </pre>	<pre> ;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 3; Isolation level set to: REPEATABLE READ SCHEMA, READ UNCOMMITTED INSTANCES. </pre>
<pre> --creating a table CREATE TABLE isol3_tbl(host_year integer, nation_code char(3)); CREATE UNIQUE INDEX on isol3_tbl(nation_code, host_year); INSERT INTO isol3_tbl VALUES (2008, 'AUS'); COMMIT; </pre>	
	<pre> --selecting records from the table SELECT * FROM isol3_tbl; host_year nation_code ===== 2008 'AUS' </pre>
<pre> INSERT INTO isol3_tbl VALUES (2004, 'AUS'); INSERT INTO isol3_tbl VALUES (2000, 'NED'); /* able to insert new rows even if tran 2 uncommitted */ </pre>	
	<pre> SELECT * FROM isol3_tbl; host_year nation_code ===== 2008 'AUS' 2004 'AUS' 2000 'NED' /* dirty read may occur so that tran 2 can select new rows uncommitted by tran_1 */ </pre>
<pre> ROLLBACK; </pre>	
	<pre> SELECT * FROM isol3_tbl; host_year nation_code ===== 2008 'AUS' /* unrepeatable read may occur so that selected results are different */ </pre>
<pre> INSERT INTO isol3_tbl VALUES (1994, 'FRA'); </pre>	

<pre>DELETE FROM isol3_tbl WHERE nation_code = 'AUS' and host_year=2008; /* able to delete rows even if tran 2 uncommitted */</pre>	<pre>SELECT * FROM isol3_tbl; host_year nation_code ===== 1994 'FRA'</pre>
<pre>ALTER TABLE isol3_tbl ADD COLUMN gold INT; /* unable to alter the table schema until tran 2 committed */</pre>	<pre>/* repeatable read is ensured while tran_1 is altering table schema */ SELECT * FROM isol3_tbl; host_year nation_code ===== 1994 'FRA'</pre>
	<pre>COMMIT;</pre>
	<pre>SELECT * FROM isol3_tbl;</pre>
<pre>COMMIT;</pre>	<pre>host_year nation_code gold ===== 1994 'FRA' NULL</pre>

Note CUBRID flushes dirty data (or dirty instances) in the client buffers to the database (server) such as the following situations. For details, see [How to Handle Dirty Instances](#).

READ COMMITTED CLASS with READ COMMITTED INSTANCES

A relatively low isolation level (2). A dirty read does not occur, but non-repeatable or phantom read may occur. That is, this level is similar to **REPEATABLE READ CLASS with READ COMMITTED INSTANCES**(level 4) described above, but works differently for table schema. Non-repeatable read due to a table schema update may occur because another transaction T2 can change the schema of the table being viewed by the transaction T1.

The following are the rules of this isolation level:

- Transaction T1 cannot read the record being updated by another transaction T2.
- Transaction T1 can update/insert a record to the table being viewed by another transaction T2.
- Transaction T1 can change the schema of the table being viewed by another transaction T2.

This isolation level uses a two-phase locking protocol for an exclusive lock. However, non-repeatable read may occur because the shared lock on the tuple is released immediately after it is retrieved and the intention lock on the table is released immediately as well.

Example

The following example shows that phantom or non-repeatable read for the record as well as for the table schema may occur because another transaction can add or update a new record while one transaction is performing the object read when the transaction level of the concurrent transactions is **READ COMMITTED CLASS with READ COMMITTED INSTANCES**.

session 1	session 2
<pre>;autocommit off AUTOCOMMIT IS OFF</pre>	<pre>;autocommit off AUTOCOMMIT IS OFF</pre>

<pre>SET TRANSACTION ISOLATION LEVEL 2; Isolation level set to: READ COMMITTED SCHEMA, READ COMMITTED INSTANCES.</pre>	<pre>SET TRANSACTION ISOLATION LEVEL 2; Isolation level set to: READ COMMITTED SCHEMA, READ COMMITTED INSTANCES.</pre>
<pre>--creating a table CREATE TABLE isol2_tbl(host_year integer, nation_code char(3)); CREATE UNIQUE INDEX on isol2_tbl(nation_code, host_year); INSERT INTO isol2 tbl VALUES (2008, 'AUS'); COMMIT;</pre>	
	<pre>--selecting records from the table SELECT * FROM isol2_tbl; host_year nation_code ===== 2008 'AUS'</pre>
<pre>INSERT INTO isol2 tbl VALUES (2004, 'AUS'); INSERT INTO isol2_tbl VALUES (2000, 'NED'); /* able to insert new rows even if tran 2 uncommitted */</pre>	
<pre>COMMIT;</pre>	<pre>SELECT * FROM isol2_tbl; /* phantom read may occur when tran 1 committed */</pre>
	<pre> host_year nation_code ===== 2008 'AUS' 2004 'AUS' 2000 'NED'</pre>
<pre>INSERT INTO isol2_tbl VALUES (1994, 'FRA');</pre>	
	<pre>SELECT * FROM isol2_tbl; /* unrepeatable read may occur when tran 1 committed */</pre>
<pre>DELETE FROM isol2_tbl WHERE nation_code = 'AUS' and host_year=2008; /* able to delete rows even if tran 2 uncommitted */</pre>	
<pre>COMMIT;</pre>	<pre> host_year nation_code ===== 2004 'AUS' 2000 'NED' 1994 'FRA'</pre>
<pre>ALTER TABLE isol2_tbl ADD COLUMN gold INT; /* able to alter the table schema even if tran 2 is uncommitted yet*/</pre>	

	/* unrepeatable read may occur so that result shows different schema */
	SELECT * FROM isol2_tbl;
COMMIT;	host_year nation_code gold =====
	2004 'AUS' NULL
	2000 'NED' NULL
	1994 'FRA' NULL

READ COMMITTED CLASS with READ UNCOMMITTED INSTANCES

The lowest isolation level (1). The concurrency level is the highest. A dirty, non-repeatable or phantom read may occur for the tuple and a non-repeatable read may occur for the table as well. Similar to **REPEATABLE READ CLASS** with **READ UNCOMMITTED INSTANCES**(level 3) described above, but works differently for the table schema. That is, non-repeatable read due to table schema update may occur because another transaction T2 can change the schema of the table being viewed by the transaction T1.

The following are the rules of this isolation level:

- Transaction T1 can read the record being updated by another transaction T2.
- Transaction T1 can update/insert record to the table being viewed by another transaction T2.
- Transaction T1 can change the schema of the table being viewed by another transaction T2.

This isolation level uses a two-phase locking protocol for an exclusive and update lock. However, the shared lock on the tuple is released immediately after it is retrieved. The intention lock on the table is released immediately after the retrieval as well.

Example

session 1	session 2
<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 1; Isolation level set to: READ COMMITTED SCHEMA, READ UNCOMMITTED INSTANCES.</pre>	<pre>;autocommit off AUTOCOMMIT IS OFF SET TRANSACTION ISOLATION LEVEL 1; Isolation level set to: READ COMMITTED SCHEMA, READ UNCOMMITTED INSTANCES.</pre>
<pre>--creating a table CREATE TABLE isol1_tbl(host_year integer, nation_code char(3)); CREATE UNIQUE INDEX on isol1_tbl(nation_code, host_year); INSERT INTO isol1_tbl VALUES (2008, 'AUS'); COMMIT;</pre>	<pre>--selecting records from the table SELECT * FROM isol1_tbl; host_year nation_code ===== 2008 'AUS'</pre>
<pre>INSERT INTO isol1_tbl VALUES (2004, 'AUS'); INSERT INTO isol1_tbl VALUES (2000, 'NED');</pre>	

<pre>/* able to insert new rows even if tran 2 uncommitted */</pre>	<pre>SELECT * FROM isoll_tbl; host_year nation_code ===== 2008 'AUS' 2004 'AUS' 2000 'NED' /* dirty read may occur so that tran 2 can select new rows uncommitted by tran_1 */</pre>
<pre>ROLLBACK;</pre>	<pre>SELECT * FROM isoll_tbl; host_year nation_code ===== 2008 'AUS' /* unrepeatable read may occur so that selected results are different */</pre>
<pre>INSERT INTO isoll_tbl VALUES (1994, 'FRA'); DELETE FROM isoll_tbl WHERE nation_code = 'AUS' and host_year=2008; /* able to delete rows while tran 2 is selecting rows*/</pre>	<pre>SELECT * FROM isoll_tbl; host_year nation_code ===== 1994 'FRA'</pre>
<pre>ALTER TABLE isoll_tbl ADD COLUMN gold INT; /* able to alter the table schema even if tran 2 is uncommitted yet*/</pre>	<pre>/* unrepeatable read may occur so that result shows different schema */ SELECT * FROM isoll_tbl;</pre>
<pre>COMMIT;</pre>	<pre>host_year nation_code gold ===== 1994 'FRA' NULL</pre>

UPDATE INCONSISTENCY

In this isolation level, uncommitted updates may be lost, which makes a transaction unrestorable (cannot be rolled back) because the data are committed before the end of the transaction. CUBRID does not support this isolation level because this can cause the updates made by the user to be lost. However, if this isolation level is specified, CUBRID provides an appropriate level to the user application.

The following are the rules of this isolation level:

- A transaction does not overwrite an object being modified by another transaction.

Note A transaction can be restored in all supported isolation levels because updates are not committed before the end of the transaction.

Combination of Unsupported Isolation Level

You can set customized isolation levels by using the **SET TRANSACTION ISOLATION LEVEL** statement. However, combinations of isolation levels below are not supported. If they are used, a system error message is shown up and an isolation level closest to the one specified is chosen.

The following are unsupported isolation levels. If table schema is changed while data is selected, unrepeatable read occurs; therefore, the combinations below are not supported.

- **READ COMMITTED CLASS** with **REPEATABLE READ INSTANCES**
- **READ UNCOMMITTED CLASS** with **REPEATABLE READ INSTANCES**

Neither are isolation levels below supported because updating a row by a transaction is not allowed while table schema is changed by other transaction.

- **READ UNCOMMITTED CLASS** with **READ COMMITTED INSTANCES**
- **READ UNCOMMITTED CLASS** with **READ UNCOMMITTED INSTANCES**

How to Handle Dirty Instance

CUBRID flushes dirty data (or dirty instances) in the client buffers to the database (server) such as the following situations. In additions to those, there can be more situations where flushes can be performed.

- Dirty data can be flushed to server when a transaction is committed.
- Some of dirty data can be flushed to server when a lot of data is loaded into the client buffers.
- Dirty data of table A can be flushed to server when the schema of table A is updated.
- Dirty data of table A can be flushed to server when the table A is retrieved (**SELECT**)
- Some of dirty data can be flushed to server when a server function is called.

Transaction Termination and Restoration

Overview

The restore process in CUBRID makes it possible that the database is not affected even if a software or hardware error occurs. In CUBRID, all read and update commands that are made during a transaction must be atomic. This means that either all of the transaction's commands are committed to the database or none are. The concept of atomicity is extended to the set of operations that consists of a transaction. The transaction must either commit so that all effects are permanently applied to the database or roll back so that all effects are removed. To ensure transaction atomicity, CUBRID applies the effects of the committed transaction again every time an error occurs without the updates of the transaction being written to the disk. CUBRID also removes the effects of partially committed transactions in the database every time the site fails (some transactions may have not committed or applications may have requested to cancel transactions). This restore feature eases the burden for the applications of maintaining the database consistency depending on the system error. The restore process used in CUBRID is based on the undo/redo logging mechanism.

CUBRID provides an automatic restore method to maintain the transaction atomicity when a hardware or software error occurs. You do not have to take the responsibility for restore since CUBRID's restore feature always returns the database to a consistent state even when an application or computer system error occurs. For this purpose, CUBRID automatically rolls back part of committed transactions when the application fails or the user requests explicitly. For example, a system error that occurred during the execution of the **COMMIT WORK** statement must be stopped if the transaction has not committed yet (it cannot be confirmed that the user's operation has been committed). Automatic stop prevents errors causing undesired changes to the database by canceling uncommitted updates.

Restarting Database

CUBRID uses log volumes/files and database backups to restore committed or uncommitted transactions when a system or media (disk) error occurs. Logs are also used to support the user-specified rollback. A log consists of a collection of sequential files created by CUBRID. The most recent log is called the active log, and the rest are called archive logs. A log file refers to both the active log and archive logs.

All updates of the database are written to the log. Actually, two copies of the updates are logged. The first one is called a before image and used to restore data during execution of the user-specified **ROLLBACK WORK** statement or during media or system errors. The second copy is an after image and used to re-apply the updates when a media or system error occurs.

When the active log is full, CUBRID copies it to an archive log to store in the disk. The archive log is needed to restore the database when a system failure occurs.

Normal Termination or Error

CUBRID restores the database if it restarts due to a normal termination or a device error. The restore process re-applies the committed changes that have not been applied to the database and removes the uncommitted changes stored in the database. The general operation of the database resumes after the restore is completed. This restore process does not use any archive logs or database backup.

In a client/server environment, the database can restart by using server utilities.

Media Error

The user's intervention is somewhat needed to restart the database after a media error occurs. The first step is to restore the database by installing a backup of a known good state. In CUBRID, the most recent log file (the one after the last backup) must be installed. This specific log (archive or active) is applied to a backup copy of the database. As with normal termination, the database can restart after restoration is committed.

Note To minimize the possibility of losing database updates, it is recommended to create a snapshot and store it in the backup media before it is deleted from the disk. The DBA can backup and restore the database by using the **cubrid backupdb** and **cubrid restoredb** utilities. For details on these utilities, see [Database Backup](#).

Database User Authorization

Database User

CUBRID has two types of users by default: **DBA** and **PUBLIC**.

- All users have authorization granted to the **PUBLIC** user. All users of the database are automatically the members of **PUBLIC**. Granting authorization to the **PUBLIC** means granting it all users.
- The **DBA** user has the authorization of the database administrator. The **DBA** automatically becomes the member of all users and groups. That is, the **DBA** is granted the access for all tables. Therefore, there is no need to grant authorization explicitly to the **DBA** and **DBA** members. Each database user has a unique name. The database administrator can create multiple users simultaneously using the **cubrid createdb** utility (see [How to Use the Database Management Utilities](#) for details). A database user cannot have a member who already has the same authorization. If authorization is granted to a user, all members of the user is automatically granted the same authorization.

Managing User

Description

DBA and **DBA** members can create, drop and alter users by using SQL statements.

Syntax

```
CREATE USER user_name
[ PASSWORD password ]
[ GROUPS user_name [ {, user_name } ... ] ]
[ MEMBERS user_name [ {, user_name } ... ] ] ;
DROP USER user_name;
ALTER USER user_name PASSWORD password;
```

- *user_name* : Specifies the user name to create, delete or change.
- *password* : Specifies the user password to create or change.

Example 1

The following example shows how to create a user (Fred), change a password, and delete the user.

```
CREATE USER Fred;
ALTER USER Fred PASSWORD '1234';
DROP USER Fred;
```

Example 2

The following example shows how to create a user and add member to the user. By the following statement, company becomes a group that has engineering, marketing and design as its members. marketing becomes a group with members smith and jones, design becomes a group with a member smith, and engineering becomes a group with a member brown.

```
CREATE USER company;
CREATE USER engineering GROUPS company;
CREATE USER marketing GROUPS company;
CREATE USER design GROUPS company;
CREATE USER smith GROUPS design, marketing;
CREATE USER jones GROUPS marketing;
CREATE USER brown GROUPS engineering;
```

Example 3

The following example shows how to create the same groups as above but use the **MEMBERS** keyword instead of **GROUPS**.

```
CREATE USER smith;
CREATE USER brown;
CREATE USER jones;
```

```
CREATE USER engineering MEMBERS brown;
CREATE USER marketing MEMBERS smith, jones;
CREATE USER design MEMBERS smith;
CREATE USER company MEMBERS engineering, marketing, design;
```

Granting Authorization

Description

In CUBRID, the smallest grant unit of authorization is a table. You must grant appropriate authorization to other users (groups) before allowing them to access the table you created.

You don't need to grant authorization individually because the members of the granted group have the same authorization. The access to the (virtual) table created by a **PUBLIC** user is allowed to all other users. You can grant access authorization to a user by using the **GRANT** statement.

Syntax

```
GRANT operation [ { ,operation }_ ] ON table_name [ { ,table_name }_ ]
TO user [ { ,user }_ ] [ WITH GRANT OPTION ] [ ; ]
```

- *operation* : Indicates an operation that can be used when granting authorization. The following table shows operations.
- **SELECT** : Allows to read the table definitions and retrieve records. The most general type of permissions.
- **INSERT** : Allows to create records in the table.
- **UPDATE** : Allows to modify the records already existing in the table.
- **DELETE** : Allows to delete records in the table.
- **ALTER** : Allows to modify the table definition, rename or delete the table.
- **INDEX** : Allows to call table methods or instance methods.
- **EXECUTE** : Allows to call table methods or instance methods.
- **ALL PRIVILEGES** : Includes all permissions described above.
- *table_name* : Specifies the name of the table or virtual table to be granted.
- *user* : Specifies the name of the user (group) to be granted. Enter the login name of the database user or **PUBLIC**, a system-defined user. If **PUBLIC** is specified, all database users are granted with the permission.
- **WITH GRANT OPTION** : **WITH GRANT OPTION** allows the grantee of authorization to grant that same authorization to another user.

Example 1

The following example shows how to grant the **SELECT** authorization for the olympic table to Fred (including his members).

```
GRANT SELECT ON olympic TO Fred;
```

Example 2

The following example shows how to grant the **SELECT**, **INSERT**, **UPDATE** and **DELETE** authorization on the nation and athlete tables to Jeniffer and Daniel (including their members).

```
GRANT SELECT, INSERT, UPDATE, DELETE ON nation, athlete TO Jeniffer, Daniel;
```

Example 3

The following example shows how to grant every authorization on the game and event tables to all users.

```
GRANT ALL PRIVILEGES ON game, event TO public;
```

Example 4

The following example shows how to grant retrieving authorization on the record and history tables to ROSS. Using **WITH GRANT OPTION** allows ROSS to grant retrieving to another users. Ross can grant authorization to others within her authorization.


```
GRANT SELECT ON record, history TO Ross WITH GRANT OPTION;
```

Caution

- The grantor of authorization must be the owner of all tables listed before the grant operation or have **WITH GRANT OPTION** specified.
- Before granting **SELECT**, **UPDATE**, **DELETE** and **INSERT** authorization for a virtual table, the owner of the virtual table must have **SELECT** and **GRANT** authorization for all the tables included in the query specification. The **DBA** user and the members of the **DBA** group are automatically granted all authorization for all tables.
- To execute the **TRUNCATE** statement, the **ALTER**, **INDEX**, and **DELETE** authorization is **required**.

Revoking Authorization

Description

You can revoke authorization using the **REVOKE** statement. The authorization granted to a user can be revoked anytime. If more than one authorization are granted to a user, all or part of the authorization can be revoked. In addition, if authorization on multiple tables is granted to more than one user using one **GRANT** statement, the authorization can be selectively revoked for specific users and tables.

If the authorization (**WITH GRANT OPTION**) is revoked from the grantor, the authorization granted to the grantee by that grantor is also revoked.

Syntax

```
REVOKE operation [ { , operation }_ ] ON table_name [ { , class_name }_ ]
FROM user [ { , user }_ ] [ ; ]
```

- *operation* : Indicates an operation that can be used when granting authorization (see **Syntax** in [Granting Authorization](#) for details).
- *table_name* : Specifies the name of the table or virtual table to be granted.
- *user* : Specifies the name of the user (group) to be granted.

Example 1

The following example shows how to grant **SELECT**, **INSERT**, **UPDATE**, and **DELETE** authorization to Fred and John so that they can perform on the nation and athlete tables.

```
GRANT SELECT, INSERT, UPDATE, DELETE ON nation, athlete TO Fred, John;
```

Example 2

The following example shows how to execute the **REVOKE** statement; this allows John to have only **SELECT** authorization. If John has granted authorization to another user, the user is also allowed to execute **SELECT** only.

```
REVOKE INSERT, UPDATE, DELETE ON nation, athlete FROM John;
```

Example 3

The following example shows how to execute the **REVOKE** statement revoking all authorization that has granted to Fred. Fred is not allowed to execute any operations on the nation and athlete tables once this statement is executed.

```
REVOKE ALL PRIVILEGES ON nation, athlete FROM Fred;
```

User Authorization Management METHOD

Description

The database administrator (**DBA**) can check and modify user authorization by calling authorization-related methods defined in **db_user** where information about database user is stored, and **db_authorization** (the system authorization class). The administrator can specify **db_user** or **db_authorization** depending on the method to be called, and store the return value of a method to a variable. In addition, some methods can be called only by **DBA** or members of **DBA** group.

Syntax

```
CALL method_definition ON CLASS auth_class [ TO variable ] [ ; ]
CALL method_definition ON variable [ ; ]
```

login() method

As a class method of **db_user** class, this method is used to change the users who are currently connected to the database. The name and password of a new user to connect are given as parameters, and they must be string type. If there is no password, a blank character ("") can be used as the parameter. **DBA** and **DBA** members can call the **login()** method without a password.

```
-- Connect as DBA user who has no password
CALL login ('dba', '') ON CLASS db user;
-- Connect as a user_1 whose password is cubrid
CALL login ('user_1', 'cubrid') ON CLASS db_user;
```

add_user() method

As a class method of **db_user** class, this method is used to add a new user. The name and password of a new user to add are given as parameters, and they must be string type. At this time, the new user name should not duplicate any user name already registered in a database. The **add_user()** can be called only by **DBA** or members of **DBA** group.

```
-- Add user 2 who has no password
CALL add_user ('user_3', '') ON CLASS db_user;
-- Add user_3 who has no password, and store the return value of a method into an admin
variable
CALL add_user ('user_2', '') ON CLASS db_user to admin;
```

drop_user() method

As a class method of **db_user** class, this method is used to drop an existing user. Only the user name to be dropped is given as a parameter, and it must be a string type. However, the owner of a class cannot be dropped thus **DBA** needs to specify a new owner of the class before dropping the user. The **drop_user()** method can be also called only by **DBA** or members of **DBA**.

```
-- Delete user 2
CALL drop_user ('user_2') ON CLASS db_user;
```

find_user() method

As a class method of **db_user** class, this method is used to find a user who is given as a parameter. The name of a user to be found is given as a parameter, and the return value of the method is stored into a variable that follows 'to'. The stored value can be used in a next query execution.

```
-- Find user 2 and store it into a variable called 'admin'
CALL find_user ('user_2') ON CLASS db_user to admin;
```

set_password() method

This method is an instance method that can call each user instance, and it is used to change a user's password. The new password of a specified user is given as a parameter. General users other than **DBA** and **DBA** group members can only change their own passwords.

```
-- Add user 4 and store it into a variable called user common
CALL add user ('user 4', '') ON CLASS db user to user common;
-- Change the password of user 4 to 'abcdef'
CALL set_password('abcdef') on user_common;
```

change_owner() method

As a class method of **db_authorizations** class, this method is used to change the owner of a class. The name of a class for which you want to change the owner, and the name of a new owner are given as parameters. At this time, the class and owner that are specified as a parameter must exist in a database. Otherwise, an error occurs. **change_owner()** can be called only by **DBA** or members of **DBA** group.

```
-- Change the owner of table_1 to user_4
```

```
CALL change_owner ('table_1', 'user_4') ON CLASS db_authorizations;
```

Example

The following example shows a **CALL** statement that calls the `find_user` method defined in the system table `db_user`. It is called to determine whether the database user entered as the `find_user` exists. The first statement calls the table method defined in the `db_user` class. The name (`db_user` in this case) is stored in `x` if the user is registered in the database. Otherwise, **NULL** is stored.

The second statement outputs the value stored in the variable `x`. In this query statement, the `DB_ROOT` is a system class that can have only one record. It can be used to output the value of `sys_date` or other registered variables. For this purpose, the `DB_ROOT` can be replaced by another table having only one record.

```
CALL find_user('dba') ON CLASS db_user to x;
Result
=====
db user

SELECT x FROM db_root;
x
=====
db user
```

With `find_user`, you can determine if the user exists in the database depending on whether the return value is **NULL** or not.

Query Optimization

Updating Statistics

Description

With the **UPDATE STATISTICS ON** statement, you can generate internal statistics used by the query processor. Such statistics allow the database system to perform query optimization more efficiently.

Syntax

```
UPDATE STATISTICS ON { table_spec [ {, table_spec } ] | ALL CLASSES | CATALOG CLASSES }
[ ; ]
table_spec :
single_table_spec
( single_table_spec [ {, single_table_spec } ] )
single_table_spec :
[ ONLY ] table_name
| ALL table_name [ ( EXCEPT table_name ) ]
```

- **ALL CLASSES** : If the **ALL CLASSES** keyword is specified, the statistics on all the tables existing in the database are updated.

Checking Statistics Information

Description

You can check the statistics Information with the session command of the CSQL Interpreter.

Syntax

```
csql> ;info stats <table_name>
```

- *table_name* : Table name to check the statistics Information

Example

The following example shows how to display the statistics Information of the t1 table in the CSQL Interpreter.

```
CREATE TABLE t1 (code INT);
INSERT INTO t1 VALUES (1), (2), (3), (4), (5);
CREATE INDEX ON t1(code);
UPDATE STATISTICS ON t1;
;info stats t1
CLASS STATISTICS
*****
Class name: t1 Timestamp: Mon Mar 14 16:26:40 2011
Total pages in class heap: 1
Total objects: 5
Number of attributes: 1
Attribute: code
  id: 0
  Type: DB TYPE INTEGER
  Minimum value: 1
  Maximum value: 5
  B+tree statistics:
    BTID: { 0 , 1049 }
    Cardinality: 5 (5) , Total pages: 2 , Leaf pages: 1 , Height: 2
```

Using SQL Hint

Description

Using hints can affect the performance of query execution. you can allow the query optimizer to create more efficient execution plan by referring the SQL HINT. The SQL HINTs related tale join, index, and statistics information are provided by CUBRID.

Syntax

```
CREATE /**+ NO_STATS */ [TABLE | CLASS] ...;
ALTER /**+ NO_STATS */ [TABLE | CLASS] ...;

CREATE /**+ NO_STATS */ INDEX ...;
ALTER /**+ NO_STATS */ INDEX ...;
DROP /**+ NO_STATS */ INDEX ...;

SELECT /**+ hint [ { hint } ... ] */
SELECT --+ hint [ { hint } ... ]
SELECT //+ hint [ { hint } ... ]

hint :
USE_NL[(spec-name[ {, spec-name}...])]
USE_IDX[(spec-name[ {, spec-name}...])]
USE_MERGE[(spec-name[ {, spec-name}...])]
ORDERED
USE_DESC_IDX
NO_DESC_IDX
NO_COVERING_IDX
```

SQL hints are specified by using plus signs and comments. CUBRID interprets this comment as a list of hints separated by blanks. The hint comment must appear after the **SELECT**, **CREATE**, or **ALTER** keyword, and the comment must begin with a plus sign (+), following the comment delimiter.

- *hint* : The following hints can be specified.
- **USE_NL** : Related to a table join, the query optimizer creates a nested loop join execution plan with this hint.
- **USE_MERGE** : Related to a table join, the query optimizer creates a sort merge join execution plan with this hint.
- **ORDERED** : Related to a table join, the query optimizer create a join execution plan with this hint, based on the order of tables specified in the FROM clause. The left table in the FROM clause becomes the outer table; the right one becomes the inner table.
- **USE_IDX** : Related to a index, the query optimizer creates a index join execution plan corresponding to a specified table with this hint.
- **USE_DESC_IDX** : This is a hint for the scan in descending index.
- **NO_DESC_IDX** : This is a hint not to use the descending index. For details, see [Index Scan in Descending Order](#).
- **NO_COVERING_IDX** : This is a hint not to use the covering index. For details, see [Covering Index](#).
- **NO_STATS** : Related to statistics information, the query optimizer does not update statistics information. Query performance for the corresponding queries can be improved; however, query plan is not optimized because the information is not updated.
- *spec_name* : If the *spec_name* is specified together with **USE_NL**, **USE_IDX** or **USE_MERGE**, the specified join method applies only to the *spec_name*. If **USE_NL** and **USE_MERGE** are specified together, the given hint is ignored. In some cases, the query optimizer cannot create a query execution plan based on the given hint. For example, if **USE_NL** is specified for a right outer join, the query is converted to a left outer join internally, and the join order may not be guaranteed.

Example 1

The following example shows how to retrieve the years when Sim Kwon Ho won medals and the types of medals. Here, a nested loop join execution plan needs to be created which has the **athlete** table as an outer table and the **game** table as an inner table. It can be expressed by the following query. The query optimizer creates a nested loop join execution plan that has the **game** table as an outer table and the **athlete** table as an inner table.

```
SELECT /**+ USE_NL ORDERED */ a.name, b.host_year, b.medal
FROM athlete a, game b WHERE a.name = 'Sim Kwon Ho' AND a.code = b.athlete_code;
      name                host_year  medal
```

```
=====
'Sim Kwon Ho'          2000  'G'
'Sim Kwon Ho'          1996  'G'
2 rows selected.
```

Example 2

The following example shows how to retrieve query execution time with **NO_STAT** hint to improve the functionality of drop partitioned table (before_2008); any data is not stored in the table. Assuming that there are more than 1 million data in the **participant2** table. The execution time in the example depends on system performance and database configuration.

```
-- Not using NO_STATS hint
ALTER TABLE participant2 DROP partition before_2008;
SQL statement execution time: 31.684550 sec
Current transaction has been committed.
1 command(s) successfully processed.

-- Using NO_STATS hint
ALTER /*+ NO_STATS */ TABLE participant2 DROP partition before_2008;
SQL statement execution time: 0.025773 sec
Current transaction has been committed.
1 command(s) successfully processed.
```

Viewing Query Plan

Description

To view a query plan for a CUBRID SQL query, change the value of the optimization level by using the **SET OPTIMIZATION** statement. You can get the current optimization level value by using the **GET OPTIMIZATION** statement.

The CUBRID query optimizer determines whether to perform query optimization and output the query plan by referencing the optimization level value set by the user. The query plan is displayed as standard output; the following explanations are based on the assumption that the plan is used in a terminal-based program such as the CSQL Interpreter. In the CSQL query editor, you can view execution plan by executing the **plan** command. See [Session Commands](#). For information on how to view a query plan, see the CUBRID Manager.

Syntax

```
SET OPTIMIZATION LEVEL opt-level [;]
GET OPTIMIZATION LEVEL [ { TO | INTO } variable ] [;]
```

- *opt-level*: A value that specifies the optimization level. It has the following meanings.
- 0 : Does not perform query optimization. The query is executed using the simplest query plan. This value is used only for debugging.
- 1 : Create a query plan by performing query optimization and executes the query. This is a default value used in CUBRID, and does not have to be changed in most cases.
- 2 : Creates a query plan by performing query optimization. However, the query itself is not executed. In general, this value is not used; it is used together with the following values to be set for viewing query plans.
- 257 : Performs query optimization and outputs the created query plan. This value works for displaying the query plan by internally interpreting the value as 256+1 related with the value 1.
- 258 : Performs query optimization and outputs the created query plan. The difference from the value 257 is that the query is not executed. That is, this value works for displaying the query plan by internally interpreting the value as 256+2 related with the value 2. This setting is useful to examine the query plan but not to intend to see the query results.
- 513 : Performs query optimization and outputs the detailed query plan. This value works for displaying more detailed query plan than the value 257 by internally interpreting the value as 512+1.
- 514 : Performs query optimization and outputs the detailed query plan. However, the query is not executed. This value works for displaying more detailed query plan than the value 258 by internally interpreting the value as 512+2.

Example

The following example shows how to view query plan by using the example retrieving year when Sim Kwon Ho won medal and metal type.

```

GET OPTIMIZATION LEVEL
      Result
=====
                        1

SET OPTIMIZATION LEVEL 258;

SELECT a.name, b.host_year, b.medal
FROM athlete a, game b WHERE a.name = 'Sim Kwon Ho' AND a.code = b.athlete code
Query plan:
  Nested loops
    Sequential scan(game b)
    Index scan(athlete a, pk_athlete_code, a.code=b.athlete_code)
There are no results.
0 rows selected.

```

Using Indexes

USING INDEX Clause

Description

The **USING INDEX** clause forces a sequential scan or an index scan to be used or an index for better performance to be included. The **USING INDEX** clause must be specified after the **WHERE** clause of **SELECT**, **DELETE**, or **UPDATE** statement.

If you specify the list of index names in the **USING INDEX** clause, the query optimizer calculates the query execution cost and makes the most optimized execution plan comparing the cost between the index scan and the sequential scan which are specified (CUBRID performs the query optimization based on the cost to select the execution plan).

You can use the **USING INDEX** clause in the order that you want without using **ORDER BY**. If you do an index scan in CUBRID, the result will be created in the order of being stored in the index and you can **USING INDEX** to get the query result in the specific index order when one table has multiple indexes.

Syntax

```

SELECT . . . FROM . . . WHERE . . .
[USING INDEX { NONE | index_spec [ {, index_spec } ...] } ] [ ; ]
DELETE FROM . . . WHERE . . .
[USING INDEX { NONE | index_spec [ {, index_spec } ...] } ] [ ; ]
UPDATE . . . SET . . . WHERE . . .
[USING INDEX { NONE | index_spec [ {, index_spec } ...] } ] [ ; ]
index_spec :
[table_name.]index_name [(+)]

```

- **NONE** : If **NONE** is specified, a sequential scan is selected.
- **(+)** : If **(+)** is specified after the index name, an index scan using the specified index is selected.

Example

The following example shows how to create an index based on the table creation statement of the **athlete** table.

```

CREATE TABLE athlete (
  code          SMALLINT      NOT NULL PRIMARY KEY,
  name          VARCHAR(40)  NOT NULL,
  gender        CHAR(1)      ,
  nation code   CHAR(3)      ,
  event         VARCHAR(30)
);
CREATE UNIQUE INDEX athlete_idx ON athlete(code, nation_code);
CREATE INDEX char_idx ON athlete(gender, nation_code);

```

For the following query, the query optimizer can choose an index scan that uses the **athlete_idx** index.

```
SELECT * FROM athlete WHERE gender='M' AND nation_code='USA';
```

USING INDEX char is the same meaning as **USING INDEX**.

If the index scan cost is less than the sequential scan cost, an index scan is performed.

```
SELECT * FROM athlete WHERE gender='M' AND nation_code='USA'
USING INDEX char_idx;
```

To forcefully specify an index scan that uses the **char_idx** index, place (+) after the index name.

```
SELECT * FROM athlete WHERE gender='M' AND nation_code='USA'
USING INDEX char_idx(+);
```

To allow a sequential scan to be selected, specify **NONE** in the **USING INDEX** clause as follows:

```
SELECT * FROM athlete WHERE gender='M' AND nation code='USA'
USING INDEX NONE;
```

If more than one indexes were specified in the **USING INDEX** clause as shown below, the query optimizer chooses an appropriate one from the specified indexes.

```
SELECT * FROM athlete WHERE gender='M' AND nation code='USA'
USING INDEX char_idx, athlete_idx;
```

If you execute queries for multiple tables, you can specify to perform an index scan on one table by using a special index, and a sequential scan on other tables. These queries have the following form.

```
SELECT ... FROM tab1, tab2 WHERE ... USING INDEX tab1.idx1, tab2.NONE;
```

If you execute a query including the **USING INDEX** clause, the query optimizer considers all indexes available of the corresponding table for the tables not specified indexes. For example, if the table tab1 has indices idx1 and idx2, and the table tab2 has indices idx3, idx4 and idx5, specify the index for only tab1 and if you do not specify tab2 index, the query optimizer works considering tab2 index.

```
SELECT ... FROM tab1, tab2 WHERE ... USING INDEX tab1.idx1;
```

- Select the best query plan by comparing the sequential scan and index scan of table tab1.
- Select the most optimized query plan by comparing the sequential scan on the table tab2 and the index scan on idx3, idx4 and idx5.

To perform an index scan on the table tab2 and a sequential scan on the table tab1, specify tab1.NONE so as not to perform an index scan on the tab1 table.

```
SELECT * from tab1,tab2 WHERE tab1.id > 2 and tab2.id < 3 USING index i_tab2_id, tab1.NONE;
```

Index Scan in Descending Order

Description

When a query is executed by sorting in descending order as follows, it usually creates a reverse index.

```
SELECT * FROM tab [WHERE ...] ORDER BY a DESC
```

However, if you create an ascending index and an descending index in the same column, the possibility of deadlock increases. In order to decrease the possibility of such case, CUBRID supports the descending scan without the separate descending index creation. Users can use the **USE_DESC_IDX** hint to specify the use of the descending scan. If the hint is not specified, the following three query executions should be considered, provided that the columns listed in the **ORDER BY** clause can use the index.

- Sequential scan + Sort in descending order
- Scan in general ascending order + sort in descending
- Scan in descending order that does not require a separate scan

Although the **USE_DESC_IDX** hint is omitted for the scan in descending order, the query optimizer decides the last execution plan of the three listed for an optimal plan.

Note The **USE_DESC_IDX** hint is not supported for the join query.

Example

```
CREATE TABLE di (i INT);
CREATE INDEX on di (i);
INSERT INTO di VALUES (5), (3), (1), (4), (3), (5), (2), (5);
```

The following example shows how to execute queries by using the **USE_DESC_IDX** hint.

```
-- We now run the following query, using the 'use desc idx' SQL hint:

SELECT /*+ USE_DESC_IDX */ * FROM di WHERE i > 0 LIMIT 3;

Query plan:
  Index scan(di di, i di i, (di.i range (0 gt inf max) and inst num() range (min inf le 3))
  (covers) (desc index))

      i
=====
      5
      5
      5
```

Even though the example below is the same as that above, the output result may be different because it cannot be scanned in descending order; which is caused by not using the **USE_DESC_IDX** hint.

```
-- The same query, without the hint, will have a different output, since descending scan
is not used.

SELECT * FROM di WHERE i > 0 LIMIT 3;

Query plan:
  Index scan(di di, i_di_i, (di.i range (0 gt_inf max) and inst_num() range (min inf_le 3))
  (covers))

      i
=====
      1
      2
      3
```

The following example shows how to sort in descending order by using **ORDER BY DESC**; the example below is the same as that above. There is no **USE_DESC_IDX** hint in the following example; however it is scanned in descending order and the result is the same as the example 1.

```
-- We also run the same query , this time asking that the results are displayed in
descending order. However, no hint will be given. Since the

-- ORDER BY...DESC clause is present, CUBRID will use descending scan, even if the hint
is was not given, thus avoiding to sort the records.

SELECT * FROM di WHERE i > 0 ORDER BY i DESC LIMIT 3;

Query plan:
  Index scan(di di, i_di_i, (di.i range (0 gt_inf max)) (covers) (desc_index))

      i
=====
      5
      5
      5
```

Covering Index

Description

The covering index is the index including the data of all columns in the **SELECT** list and the **WHERE**, **HAVING**, **GROUP BY**, and **ORDER BY** clauses.

You only need to scan the index pages, as the covering index contains all the data necessary for executing a query, and it also reduces the I/O costs as it is not necessary to scan the data storage any further. To increase data search speed, you

can consider creating a covering index but you should be aware that the **INSERT** and the **DELETE** processes may be slowed down due to the increase in index size.

The rules about the applicability of the covering index are as follows:

- If the covering index is applicable, you should use the CUBRID query optimizer first.
- For the join query, if the index includes columns of the table in the **SELECT** list, use this index.
- You cannot use the covering index if an index cannot be used.

Example

```
CREATE TABLE t (col1 INT, col2 INT, col3 INT);
CREATE INDEX ON t (col1,col2,col3);
INSERT INTO t VALUES (1,2,3), (4,5,6), (10,8,9);
```

The following example shows that the index is used as a covering index because columns of both **SELECT** and **WHERE** condition exist within the index.

```
csql>;plan simple
SELECT * FROM t WHERE col1 < 6;

Query plan:
  Index scan(t t, i t col1 col2 col3, [(t.col1 range (min inf lt t.col3))]) (covers)
=====
      col1          col2          col3
-----
          1            2            3
          4            5            6
```

Caution

If the covering index is applied when you get the values from the **VARCHAR** type column, the empty strings that follow will be truncated. If the covering index is applied to the execution of query optimization, the resulting query value will be retrieved. This is because the value will be stored in the index with the empty string being truncated.

If you don't want this, use the **NO_COVERING_IDX** hint, which does not use the covering index function. If you use the hint, you can get the result value from the data area rather than from the index area.

The following is a detailed example of the above situation. First, create a table with columns in **VARCHAR** types, and then **INSERT** the value with the same start character string value but the number of empty characters. Next, create an index in the column.

```
CREATE TABLE tab(c VARCHAR(32));
INSERT INTO tab VALUES('abcd'),('abcd  '),('abcd ');
CREATE INDEX ON tab(c);
```

If you must use the index (the covering index applied), the query result is as follows:

```
csql>;plan simple
SELECT * FROM tab where c='abcd  ' USING INDEX i tab c(+);

Query plan:
  Index scan(tab tab, i tab c, (tab.c='abcd  ')) (covers)

c
=====
'abcd'
'abcd'
'abcd'
```

The following is the query result when you don't use the index.

```
SELECT * FROM tab WHERE c='abcd  ' USING INDEX tab.NONE;

Query plan:
  Sequential scan(tab tab)

c
=====
'abcd'
'abcd  '
'abcd  '
```

As you can see in the above comparison result, the value in the **VARCHAR** type retrieved from the index will appear with the following empty string truncated when the covering index has been applied.

Optimizing ORDER BY Statement

Description

The index including all columns in the **ORDER BY** clause is referred to as the ordered index. In general, for an ordered index, the columns in the **ORDER BY** clause should be located at the front of the index.

```
SELECT * FROM tab WHERE col1 > 0 ORDER BY col1, col2
```

- The index consisting of tab(col1, col2) is an ordered index.
- The index consisting of tab(col1, col2, col3) is also an ordered index. This is because the col3, which is not referred by the ORDER BY clause comes after col1 and col2.
- The index consisting of tab(col1) is not an ordered index.
- You can use the index consisting of tab(col3, col1, col2) or tab(col1, col3, col2) for optimization. This is because col3 is not located at the back of the columns in the ORDER BY clause.

Although the columns composing an index do not exist in the **ORDER BY** clause, you can use an ordered index if the column condition is a constant.

```
SELECT * FROM tab WHERE col2=val ORDER BY col1, col3;
```

If the index consisting of tab(col1, col2, col3) exists and the index consisting of tab(col1, col2) do not exist when executing the above query, the query optimizer uses the index consisting of tab(col1, col2, col3) as an ordered index.

You can get the result in the requested order when you execute an index scan, so you don't need to sort rows.

If you can use the sorted index and the covering index, use the latter first. If you use the covering index, you don't need to retrieve additional data, because the data result requested is included in the index page, and you won't need to sort the result if you are satisfied with the index order.

If the query doesn't include any conditions and uses an ordered index, the ordered index will be used under the condition that the first column meets the **NOT NULL** condition.

Example

```
CREATE TABLE tab (i INT, j INT, k INT);
CREATE INDEX on tab (j,k);
INSERT INTO tab VALUES (1,2,3), (6,4,2), (3,4,1), (5,2,1), (1,5,5), (2,6,6), (3,5,4);
```

The following example shows that indexes consisting of tab(j,k) become sorted indexes and no separate sorting process is required because **GROUP BY** is executed by j and k columns.

```
SELECT i,j,k FROM tab WHERE j > 0 ORDER BY j,k;
```

```
-- the selection from the query plan dump shows that the ordering index i_tab_j_k was
used and sorting was not necessary
```

```
-- (/* --> skip ORDER BY */)
```

```
Query plan:
```

```
iscan
```

```
  class: tab node[0]
  index: i_tab_j_k term[0]
  sort:  2 asc, 3 asc
  cost:  fixed 0(0.0/0.0) var 1(0.0/1.0) card 0
```

```
Query stmt:
```

```
select tab.i, tab.j, tab.k from tab tab where ((tab.j> ? :0 )) order by 2, 3
```

```
/* ---> skip ORDER BY */
```

i	j	k
5	2	1
1	2	3
3	4	1
6	4	2
3	5	4
1	5	5
2	6	6

The following example shows that j and k columns execute **ORDER BY** and the index including all columns are selected so that indexes consisting of tab(j,k) are used as covering indexes; no separate process is required because the value is selected from the indexes themselves.

```
SELECT /*+ RECOMPILE */ j,k FROM tab WHERE j > 0 ORDER BY j,k;

-- in this case the index i tab j k is a covering index and also respects the orderind
index property.
-- Therefore, it is used as a covering index and sorting is not performed.

Query plan:
iscan
  class: tab node[0]
  index: i_tab_j_k term[0] (covers)
  sort: 1 asc, 2 asc
  cost: fixed 0(0.0/0.0) var 1(0.0/1.0) card 0

Query stmt: select tab.j, tab.k from tab tab where ((tab.j> ?:0 )) order by 1, 2
/* ----> skip ORDER BY */
```

j	k
2	1
2	3
4	1
4	2
5	4
5	5
6	6

The following example shows that i column exists, **ORDER BY** is executed by j and k columns, and columns that perform **SELECT** are i, j, and k. Therefore, indexes consisting of tab(i,j,k) are used as covering indexes; separate sorting process is required for **ORDER BY** j, k even though the value is selected from the indexes themselves.

```
CREATE INDEX ON tab (i,j,k);
SELECT /*+ RECOMPILE */ i,j,k FROM tab WHERE i > 0 ORDER BY j,k;

-- since an index on (i,j,k) is now available, it will be used as covering index. However,
sorting the results according to
-- the ORDER BY clause is needed.

Query plan:
temp(order by)
  subplan: iscan
    class: tab node[0]
    index: i tab i j k term[0] (covers)
    sort: 1 asc, 2 asc, 3 asc
    cost: fixed 0(0.0/0.0) var 1(0.0/1.0) card 1
  sort: 2 asc, 3 asc
  cost: fixed 6(5.0/1.0) var 1(0.0/1.0) card 1

Query stmt: select tab.i, tab.j, tab.k from tab tab where ((tab.i> ?:0 )) order by 2, 3
```

i	j	k
5	2	1
1	2	3
3	4	1
6	4	2
3	5	4
1	5	5
2	6	6

GROUP BY Clause Optimization

Description

GROUP BY clause optimization works on the premise that if all columns in the **GROUP BY** clause are included in an index, you can use the index upon executing a query, so you don't execute a separate sorting job. The columns in the **GROUP BY** clause must exist in front side of the column forming the index.

```
SELECT * FROM tab WHERE col1 > 0 GROUP BY col1,col2
```

- You can use the index consisting of tab(col1, col2) for optimization.
- The index consisting of tab(col1, col2, col3) can be used because col3 no referred by **GROUP BY** comes after col1 and col2.
- You cannot use the index consisting of tab(col1) for optimization.
- You also cannot use the index consisting of tab(col3, col1, col2) or tab(col1, col3, col2), because col3 is not located at the back of the column in the **GROUP BY** clause.

You can use the index if the column condition is a constant although the column consisting of the index doesn't exist in the **GROUP BY** clause.

```
SELECT * FROM tab WHERE col2=val GROUP BY col1,col3
```

If there is any index that consists of tab(col1, col2, col3) in the above example, use the index for optimizing **GROUP BY**.

Row sorting by **GROUP BY** is not required, because you can get the result as the requested order on the index scan.

If the index consisting of the **GROUP BY** column and the first column of the index is **NOT NULL**, even though there is no **WHERE** clause, the **GROUP BY** optimization will be applied.

GROUP BY optimization is applied only when **MIN()** or **MAX()** are used in an aggregate function, and to use the two aggregate functions together, an identical column must be used.

```
CREATE INDEX ON T(a, b, c);
SELECT a, MIN(b), c, MAX(b) FROM T WHERE a > 18 GROUP BY a, b;
```

Example

```
CREATE TABLE tab (i INT, j INT, k INT);
CREATE INDEX ON tab (j,k);
INSERT INTO tab VALUES (1,2,3), (6,4,2), (3,4,1), (5,2,1), (1,5,5), (2,6,6), (3,5,4);
```

The following example shows that indexes consisting of tab(j,k) are used and no separate sorting process is required because **GROUP BY** is executed by j and k columns.

```
SELECT i,j,k FROM tab WHERE j > 0 GROUP BY j,k;

-- the selection from the query plan dump shows that the index i tab j k was used and
-- sorting was not necessary
-- (/* ---> skip GROUP BY */)

Query plan:
iscan
  class: tab node[0]
  index: i tab j k term[0]
  sort: 2 asc, 3 asc
  cost: fixed 0(0.0/0.0) var 1(0.0/1.0) card 0

Query stmt:
select tab.i, tab.j, tab.k from tab tab where ((tab.j> ?:0 )) group by tab.j, tab.k
/* ---> skip GROUP BY */
  i           j           k
  5           2           1
  1           2           3
  3           4           1
  6           4           2
  3           5           4
  1           5           5
  2           6           6
```

The following example shows that an index consisting of tab(j,k) is used and no separate sorting process is required while **GROUP BY** is executed by j and k columns, no condition exists for j, and j column has **NOT NULL** attribute.

```
ALTER TABLE tab CHANGE COLUMN j j INT NOT NULL;
SELECT * FROM tab GROUP BY j,k;

-- the selection from the query plan dump shows that the index i tab j k was used (since
-- j has the NOT NULL constraint )
-- and sorting was not necessary (/* ---> skip GROUP BY *)
```

```
Query plan:
iscan
  class: tab node[0]
  index: i_tab_j_k
  sort: 2 asc, 3 asc
  cost: fixed 0(0.0/0.0) var 1(0.0/1.0) card 0

Query stmt: select tab.i, tab.j, tab.k from tab tab group by tab.j, tab.k
/* ---> skip GROUP BY */
=== <Result of SELECT Command in Line 1> ===
      i          j          k
=====
      5          2          1
      1          2          3
      3          4          1
      6          4          2
      3          5          4
      1          5          5
      2          6          6
```

TRIGGER

CREATE TRIGGER

Guideline for TRIGGER Definition

Trigger definition provides various and powerful functionalities. Before creating a trigger, you must consider the following:

- **Does the trigger condition expression cause unexpected results (side effect)?**

You must use the SQL statements within an expectable range.

- **Does the trigger action change the table given as its event target?**

While this type of design is not forbidden in the trigger definition, it must be carefully applied, because a trigger can be created that falls into an infinite loop. When the trigger action modifies the event target table, the same trigger can be called again. If a trigger occurs in a statement that contains a **WHERE** clause, there is no side effect in the table affected by the **WHERE** clause.

- **Does the trigger cause unnecessary overhead?**

If the desired action can be expressed more effectively in the source, implement it directly in the source.

- **Is the trigger executed recursively?**

If the trigger action calls a trigger and this trigger calls the previous trigger again, a recursive loop is created in the database. If a recursive loop is created, the trigger may not be executed correctly, or the current session must be forced to terminate to break the ongoing infinite loop.

- **Is the trigger definition unique?**

A trigger defined in the same table or the one started in the same action becomes the cause of an unrecoverable error. A trigger in the same table must have a different trigger event. In addition, trigger priority must be explicitly and unambiguously defined.

TRIGGER Definition

Description

A trigger is created by defining a trigger target, condition and action to be performed in the **CREATE TRIGGER** statement. A trigger is a database object that performs a defined action when a specific event occurs in the target table.

Syntax

```
CREATE TRIGGER trigger_name
[ STATUS { ACTIVE | INACTIVE } ]
[ PRIORITYkey ]
event_time event_type[ event_target ]
[ IFcondition ]
EXECUTE [ AFTER | DEFERRED ] action [ ; ]
```

event time:

- BEFORE
- AFTER
- DEFERRED

event type:

- INSERT
- STATEMENT INSERT
- UPDATE
- STATEMENT UPDATE
- DELETE
- STATEMENT DELETE
- ROLLBACK
- COMMIT

event_target:

```

• ONtable_name
• ONtable name [ (column name) ]

condition:
• expression

action:
• REJECT
• INVALIDATE TRANSACTION
• PRINT message_string
• INSERT statement
• UPDATE statement
• DELETE statement

```

- *trigger_name* : Specifies the name of the trigger to be defined.
- [**STATUS** { **ACTIVE** | **INACTIVE** }] : Defines the state of the trigger (if not defined, the default value is **ACTIVE**).
- If **ACTIVE** state is specified, the trigger is executed every time the corresponding event occurs.
- If **INACTIVE** state is specified, the trigger is not executed even when the corresponding event occurs. The state of the trigger can be modified. For details, see [Altering TRIGGER Definition](#) section.
- [**PRIORITY** *key*] : Specifies a trigger priority if multiple triggers are called for an event. *key* must be a floating point value that is not negative. If the priority is not defined, the lowest priority 0 is assigned. Triggers having the same priority are executed in a random order. The priority of triggers can be modified. For details, see [Altering TRIGGER Definition](#) section.
- *event_time* : Specifies the point of time when the conditions and actions are executed. **BEFORE**, **AFTER** or **DEFERRED** can be specified. For details, see the [Event Time](#) section.
- *event_type* : Trigger types are divided into a user trigger and a table trigger. For details, see the [TRIGGER Event Type](#) section.
- *event_target* : An event target is used to specify the target for the trigger to be called. For details, see the [TRIGGER Event Target](#) section.
- *condition* : Specifies the trigger condition. For details, see the [TRIGGER Condition](#) section.
- *action* : Specifies the trigger action. For details, see the [TRIGGER Action](#) section.

Example

The following example shows how to create a trigger that rejects the update if the number of medals won is smaller than 0 when an instance of the participant table is updated.

As shown below, the update is rejected if you try to change the number of gold medals that Korea won in the 2004 Olympic Games to a negative number.

```

CREATE TRIGGER medal_trigger
BEFORE UPDATE ON participant
IF new.gold < 0 OR new.silver < 0 OR new.bronze < 0
EXECUTE REJECT;

UPDATE participant SET gold = -5 WHERE nation_code = 'KOR'
AND host_year = 2004;

ERROR: The operation has been rejected by trigger "medal_trigger".

```

Event Time

Description

Specifies the point of time when trigger conditions and actions are executed. The types of event time are **BEFORE**, **AFTER** and **DEFERRED**.

- **BEFORE** : Checks the condition before the event is processed.
- **AFTER** : Checks the condition after the event is processed.
- **DEFERRED** : Checks the condition at the end of the transaction for the event. If you specify **DEFERRED**, you cannot use **COMMIT** or **ROLLBACK** as the event type.

Trigger Type

User Trigger

- A trigger relevant to a specific user of the database is called a user trigger.
- A user trigger has no event target and is executed only by the owner of the trigger (the user who created the trigger). Event types that define a user trigger are **COMMIT** and **ROLLBACK**. Table Trigger
- A trigger that has a table as the event target is called a table trigger (class trigger).
- A table trigger can be seen by all users who have the **SELECT** authorization on a target table.
- Event types that define a table trigger are instance and statement events.

TRIGGER Event Type

Description

- Instance events : An event type whose unit of operation is an instance. The types of instance events are as follows:
 - **INSERT**
 - **UPDATE**
 - **DELETE**
- Statement events : If you define a statement event as an event type, the trigger is called only once when the trigger starts even when there are multiple objects (instances) affected by the given statement (event). The types of statement events are as follows:
 - **STATEMENT INSERT**
 - **STATEMENT UPDATE**
 - **STATEMENT DELETE**
- Other events : **COMMIT** and **ROLLBACK** cannot be applied to individual instances.
 - **COMMIT**
 - **ROLLBACK**

Example 1

The following example shows how to use an instance event. The example trigger is called by each instance affected by the database update. For example, if the score values of five instances in the history table are modified, the trigger is called five times. If you want the trigger to be called only once, before the first instance of the score column is updated, use the **STATEMENT UPDATE** type as in example 2.

```
CREATE TRIGGER example
...
BEFORE UPDATE ON history(score)
...
```

Example 2

The following example shows how to use a statement event. If you define a statement event, the trigger is called only once before the first instance gets updated even when there are multiple instances affected by the update.

```
CREATE TRIGGER example
...
BEFORE STATEMENT UPDATE ON history(score)
...
```

Caution

- You must specify the event target when you define an instance or statement event as the event type.
- **COMMIT** and **ROLLBACK** cannot have an event target.

TRIGGER Event Target

Description

An event target specifies the target for the trigger to be called. The target of a trigger event can be specified as a table or column name. If a column name is specified, the trigger is called only when the specified column is affected by the event. If a column is not specified, the trigger is called when any column of the table is affected. Only **UPDATE** and **STATEMENT UPDATE** events can specify a column as the event target.

Example

The following example shows how to specify the score column of the history table as the event target of the example trigger.

```
CREATE TRIGGER example
...
BEFORE UPDATE ON history(score)
...
```

Combination of Event Type and Target

Description

A database event calling triggers is identified by the trigger event type and event target in a trigger definition. The following table shows the trigger event type and target combinations, along with the meaning of the CUBRID database event that the trigger event represents.

Event Type	Event Target	Corresponding Database Activity
UPDATE	Table	Trigger is called when the UPDATE statement for a table is executed.
INSERT	Table	Trigger is called when the INSERT statement for a table is executed.
DELETE	Table	Trigger is called when the DELETE statement for a table is executed.
COMMIT	None	Trigger is called when database transaction is committed.
ROLLBACK	None	Trigger is called when database transaction is rolled back.

TRIGGER Condition

Description

You can specify whether a trigger action is to be performed by defining a condition when defining the trigger.

- If a trigger condition is specified, it can be written as an independent compound expression that evaluates to true or false. In this case, the expression can contain arithmetic and logical operators allowed in the **WHERE** clause of the **SELECT** statement. The trigger action is performed if the condition is true; if it is false, action is ignored.
- If a trigger condition is omitted, the trigger becomes an unconditional trigger, which refers to that the trigger action is performed whenever it is called.

Example 1

The following example shows how to use a correlation name in an expression within a condition. If the event type is **INSERT**, **UPDATE** or **DELETE**, the expression in the condition can reference the correlation names **obj**, **new** or **old** to access a specific column. This example prefixes **obj** to the column name in the trigger condition to show that the example trigger tests the condition based on the current value of the record column.

```
CREATE TRIGGER example
.....
IF obj.record * 1.20 < 500
.....
```

Example 2

The following example shows how to use the **SELECT** statement in an expression within a condition. The trigger in this example uses the **SELECT** statement that contains an aggregate function **COUNT(*)** to compare the value with a constant. The **SELECT** statement must be enclosed in parentheses and must be placed at the end of the expression.

```
CREATE TRIGGER example
.....
IF 1000 > (SELECT COUNT( * ) FROM participant)
.....
```

Caution

The expression given in the trigger condition may cause side effects on the database if a method is called while the condition is performed. A trigger condition must be constructed to avoid unexpected side effects in the database.

Correlation Name

You can access the column values defined in the target table by using a correlation name in the trigger definition. A correlation name is the instance that is actually affected by the database operation calling the trigger. A correlation name can also be specified in a trigger condition or action.

The types of correlation names are **new**, **old** and **obj**. These correlation names can be used only in instance triggers that have an **INSERT**, **UPDATE** or **DELETE** event.

As shown in the table below, the use of correlation names is further restricted by the event time defined for the trigger condition.

	BEFORE	AFTER or DEFERRED
INSERT	new	obj
UPDATE	obj new	obj old (AFTER)
DELETE	obj	N/A

Correlation Name Representative Attribute Value

obj	Refers to the current attribute value of an instance. This can be used to access attribute values before an instance is updated or deleted. It is also used to access attribute values after an instance has been updated or inserted.
new	Refers to the attribute value proposed by an insert or update operation. The new value can be accessed only before the instance is actually inserted or updated.
old	Refers to the attribute value that existed prior to the completion of an update operation. This value is maintained only while the trigger is being performed. Once the trigger is completed, the old values get lost.

TRIGGER Action

Description

A trigger action describes what to be performed if the trigger condition is true or omitted. If a specific point of time (**AFTER** or **DEFERRED**) is not given in the action clause, the action is executed at once as the trigger event.

The following is a list of actions that can be used for trigger definitions.

- **REJECT** : **REJECT** discards the operation that initiated the trigger and keeps the former state of the database, if the condition is not true. Once the operation is performed, **REJECT** is allowed only when the action time is **BEFORE** because the operation cannot be rejected. Therefore, you must not use **REJECT** if the action time is **AFTER** or **DEFERRED**.

- **INVALIDATE TRANSACTION** : **INVALIDATE TRANSACTION** allows the event operation that called the trigger, but does not allow the transaction that contains the commit to be executed. You must cancel the transaction by using the **ROLLBACK** statement if it is not valid. Such action is used to protect the database from having invalid data after a data-changing event happens.
- **PRINT** : **PRINT** outputs trigger actions on the terminal screen in text messages, and can be used during developments or tests. The results of event operations are not rejected or discarded.
- **INSERT** : **INSERT** inserts one or more new instances to the table.
- **UPDATE** : **UPDATE** updates one or more column values in the table.
- **DELETE** : **DELETE** deletes one or more instances from the table.

Example

The following example shows how to define an action when a trigger is created. The medal_trig trigger defines **REJECT** in its action. **REJECT** can be specified only when the action time is **BEFORE**.

```
CREATE TRIGGER medal_trig
BEFORE UPDATE ON participant
IF new.gold < 0 OR new.silver < 0 OR new.bronze < 0
EXECUTE REJECT;
```

Caution

- Trigger may fall into an infinite loop when you use **INSERT** in an action of a trigger where an **INSERT** event is defined.
- If a trigger where an **UPDATE** event is defined runs on a partitioned table, you must be careful because the defined partition can be broken or unintended malfunction may occur. To prevent such situation, CUBRID outputs an error so that the **UPDATE** causing changes to the running partition is not executed. Trigger may fall into an infinite loop when you use **UPDATE** in an action of a trigger where an **UPDATE** event is defined.

ALTER TRIGGER

Description

In the trigger definition, **STATUS** and **PRIORITY** options can be changed by using the **ALTER** statement. If you need to alter other parts of the trigger (event targets or conditional expressions), you must delete and then re-create the trigger.

Syntax

```
ALTER TRIGGER trigger_name trigger_option [ ; ]
trigger_option :
• STATUS { ACTIVE | INACTIVE }
• PRIORITY key
```

- *trigger_name* : Specifies the name of the trigger to be changed.
- *trigger_option* :
- **STATUS** { **ACTIVE** | **INACTIVE** } : Changes the status of the trigger.
- **PRIORITY** *key* : Changes the priority.

Example

The following example shows how to create the medal_trig trigger and then change its state to **INACTIVE** and its priority to 0.7.

```
CREATE TRIGGER medal_trig
STATUS ACTIVE
BEFORE PDATE ON participant
IF new.gold < 0 OR new.silver < 0 OR new.bronze < 0
EXECUTE REJECT;
ALTER TRIGGER medal_trig STATUS INACTIVE;
ALTER TRIGGER medal_trig PRIORITY 0.7;
```

Caution

- Only one option can be specified in a single **ALTER TRIGGER** statement.

- To change a table trigger, you must be the trigger owner or granted the **ALTER** authorization on the table where the trigger belongs.
- A user trigger can only be changed by its owner. For details on these options, see the [CREATE TRIGGER \(Syntax\)](#) section. The key specified together with the **PRIORITY** option must be a non-negative floating point value.

DROP TRIGGER

Description

You can drop a trigger by using the **DROP TRIGGER** statement.

Syntax

```
DROP TRIGGER trigger_name [ ; ]
```

- *trigger_name* : Specifies the name of the trigger to be dropped.

Example

The following example shows how to drop the medal_trig trigger.

```
DROP TRIGGER medal_trig;
```

Caution

- A user trigger (i.e. the trigger event is **COMMIT** or **ROLLBACK**) can be seen and dropped only by the owner.
- Only one trigger can be dropped by a single **DROP TRIGGER** statement. A table trigger can be dropped by a user who has an **ALTER** authorization on the table.

RENAME TRIGGER

Description

You can change a trigger name by using the **TRIGGER** reserved word in the **RENAME** statement.

Syntax

```
RENAME TRIGGER old_trigger_name AS new_trigger_name [ ; ]
```

- *old_trigger_name* : Specifies the current name of the trigger.
- *new_trigger_name* : Specifies the name of the trigger to be modified.

Example

```
RENAME TRIGGER medal_trigger AS medal_trig;
```

Caution

- A trigger name must be unique among all trigger names. The name of a trigger can be the same as the table name in the database.
- To rename a table trigger, you must be the trigger owner or granted the **ALTER** authorization on the table where the trigger belongs. A user trigger can only be renamed by its user.

Deferred Condition and Action

Definition

A deferred trigger action and condition can be executed later or canceled. These triggers include a **DEFERRED** time option in the event time or action clause. If the **DEFERRED** option is specified in the event time and the time is omitted before the action, the action is deferred automatically.

Executing Deferred Condition and Action

Description

Executes the deferred condition or action of a trigger immediately.

Syntax

```
EXECUTE DEFERRED TRIGGER trigger_identifier [ ; ]
trigger_identifier :
• trigger_name
• ALL TRIGGERS
```

- *trigger_identifier* :
- *trigger_name* : Executes the deferred action of the trigger when a trigger name is specified.
- **ALL TRIGGERS** : Executes all currently deferred actions.

Dropping Deferred Condition and Action

Description

Drops the deferred condition and action of a trigger.

Syntax

```
DROP DEFERRED TRIGGER trigger_identifier [ ; ]
trigger_option :
• trigger_name
• ALL TRIGGERS
```

- *trigger_option* :
- *trigger_name* : Cancels the deferred action of the trigger when a trigger name is specified.
- **ALL TRIGGERS** : Cancels currently deferred actions.

Granting TRIGGER Authorization

Description

Trigger authorization is not granted explicitly. Authorization on the table trigger is automatically granted to the user if the authorization is granted on the event target table described in the trigger definition. In other words, triggers that have table targets (**INSERT**, **UPDATE**, etc.) are seen by all users. User triggers (**COMMIT** and **ROLLBACK**) are seen only by the user who defined the triggers. All authorizations are automatically granted to the trigger owner.

Caution

- To define a table trigger, you must have an **ALTER** authorization on the table.
- To define a user trigger, the database must be accessed by a valid user.

Trigger on REPLACE and INSERT ... ON DUPLICATE KEY UPDATE

Deferred Actions

Description

When the **REPLACE** statement and **INSERT ... ON DUPLICATE KEY UPDATE** statements are executed, the trigger is executed in CUBRID, while **DELETE**, **UPDATE**, **INSERT** jobs occur internally. The following table shows the order in which the trigger is executed in CUBRID depending on the event that occurred when the **REPLACE** or **INSERT ... ON DUPLICATE KEY UPDATE** statement is executed. Both the **REPLACE** statement and **INSERT ... ON DUPLICATE KEY UPDATE** statement do not execute triggers in the inherited class (table).

Execution Sequence of Triggers in the REPLACE and the INSERT ... ON DUPLICATE KEY UPDATE statements

Event	Execution Sequence of Triggers
REPLACE When a record is deleted and new one is inserted	BEFORE DELETE > AFTER DELETE > BEFORE INSERT > AFTER INSERT
INSERT ... ON DUPLICATE KEY UPDATE When a record is updated	BEFORE UPDATE > AFTER UPDATE
REPLACE, INSERT ... ON DUPLICATE KEY UPDATE Only when a record is inserted	BEFORE INSERT > AFTER INSERT

Example

The following is an example in which the trigger inserts records to the trigger table if **INSERT ... ON DUPLICATE KEY UPDATE** and **REPLACE** are executed in the `with_trigger` table.

```
CREATE TABLE with_trigger (id INT UNIQUE);
INSERT INTO with_trigger VALUES (11);

CREATE TABLE trigger_actions (val INT);

CREATE TRIGGER trig_1 BEFORE INSERT ON with_trigger EXECUTE INSERT INTO trigger_actions
VALUES (1);
CREATE TRIGGER trig_2 BEFORE UPDATE ON with_trigger EXECUTE INSERT INTO trigger_actions
VALUES (2);
CREATE TRIGGER trig_3 BEFORE DELETE ON with_trigger EXECUTE INSERT INTO trigger_actions
VALUES (3);

INSERT INTO with_trigger VALUES (11) ON DUPLICATE KEY UPDATE id=22;

SELECT * FROM trigger_actions;
      va
=====
      2

REPLACE INTO with_trigger VALUES (22);

SELECT * FROM trigger_actions;
      va
=====
      2
      3
      1
```

TRIGGER Debugging

Definition and Example

Description

Once a trigger is defined, it is recommended to check whether it is running as intended. Sometimes the trigger takes more time than expected in processing. This means that it is adding too much overhead to the system or has fallen into a recursive loop. This section explains several ways to debug the trigger.

Example

The following example shows a trigger that was defined to fall into a recursive loop when it is called. A loop trigger is somewhat artificial in its purpose; it can be used as an example of debugging trigger.

```
CREATE TRIGGER loop_tgr
BEFORE UPDATE ON participant(gold)
IF new.gold > 0
```

```
EXECUTE UPDATE participant
  SET gold = new.gold - 1
  WHERE nation_code = obj.nation_code AND host_year = obj.host_year;
```

Viewing TRIGGER Execution Log

Description

You can view the execution log of the trigger from a terminal by using the **SET TRIGGER TRACE** statement.

Syntax

```
SET TRIGGER TRACE switch [ ; ]
switch:
• ON
• OFF
```

- *switch* :
- **ON** : Executes **TRACE** until the switch is set to **OFF** or the current database session terminates.
- **OFF** : Stops the **TRACE**.

Example

The following example shows how to execute the **TRACE** and the loop trigger to view the trigger execution logs. To identify the trace for each condition and action executed when the trigger is called, a message is displayed on the terminal. The following message appears 15 times because the loop trigger is executed until the gold value becomes 0.

```
SET TRIGGER TRACE ON;

UPDATE participant SET gold =15 WHERE nation code = 'KOR' AND host year = 1988;
TRACE: Evaluating condition for trigger "loop".
TRACE: Executing action for trigger "loop".
```

Limiting Nested TRIGGER

Description

With the **MAXIMUM DEPTH** keyword of the **SET TRIGGER** statement, you can limit the number of triggers to be initiated at each step. By doing so, you can prevent a recursively called trigger from falling into an infinite loop.

Syntax

```
SET TRIGGER [ MAXIMUM ] DEPTH count [ ; ]
count:
• unsigned_integer_literal
```

- *unsigned_integer_literal* : A positive integer value that specifies the number of times that a trigger can recursively start another trigger or itself. If the number of triggers reaches the maximum depth, the database request stops(aborts) and the transaction is marked as invalid. The specified **DEPTH** applies to all other triggers except the current session. The maximum value is 32.

Example

The following example shows how to configure the maximum number of times of recursive trigger calling to 10. This applies to all triggers that start subsequently. In this example, the gold column value is updated to 15, so the trigger is called 16 times in total. This exceeds the currently set maximum depth and the following error message occurs.

```
SET TRIGGER MAXIMUM DEPTH 10;
UPDATE participant SET gold = 15 WHERE nation code = 'KOR' AND host year = 1988;

ERROR: Maximum trigger depth 10 exceeded at trigger "loop_tgr".
```


TRIGGER Example

Description

This section covers trigger definitions in the demo database. The triggers created in the demodb database are not complex, but use most of the features available in CUBRID. If you want to maintain the original state of the demodb database when testing such triggers, you must perform a rollback after changes are made to the data.

Triggers created by the user in the own database can be as powerful as applications created by the user.

Example 1

The following trigger created in the participant table rejects an update to the medal column (gold, silver, bronze) if a given value is smaller than 0. The evaluation time must be **BEFORE** because a correlation name new is used in the trigger condition. Although not described, the action time of this trigger is also **BEFORE**.

```
CREATE TRIGGER medal trigger
BEFORE UPDATE ON participant
IF new.gold < 0 OR new.silver < 0 OR new.bronze < 0
EXECUTE REJECT;
```

The medal_trigger trigger starts when the number of gold medals of the country whose nation code is 'BLA' is updated. Since a negative value is not permitted for the number of gold medals as shown above, this update is not allowed.

```
UPDATE participant
SET gold = -10
WHERE nation_code = 'BLA';
```

Example 2

The following trigger has the same condition as the one above except that **STATUS INACTIVE** is added. If the STATUS statement is omitted, the default value is **ACTIVE**. You can change the status to **INACTIVE** by using the **ALTER TRIGGER** statement.

You can specify whether or not to execute the trigger depending on the **STATUS** value.

```
CREATE TRIGGER medal trig
STATUS ACTIVE
BEFORE UPDATE ON participant
IF new.gold < 0 OR new.silver < 0 OR new.bronze < 0
EXECUTE REJECT;

ALTER TRIGGER medal trig
STATUS INACTIVE;
```

Example 3

The following trigger shows how integrity constraint is enforced when a transaction is committed. This example is different from the previous ones, in that one trigger can have specific conditions for multiple tables.

```
CREATE TRIGGER check null first
BEFORE COMMIT
IF 0 < (SELECT count(*) FROM athlete WHERE gender IS NULL)
OR 0 < (SELECT count(*) FROM game WHERE nation_code IS NULL)
EXECUTE REJECT;
```

Example 4

The following trigger delays the update integrity constraint check for the record table until the transaction is committed. Since the **DEFERRED** keyword is given as the event time, the trigger is not executed at the time.

```
CREATE TRIGGER deferred_check_on_record
DEFERRED UPDATE ON record
IF obj.score = '100'
EXECUTE INVALIDATE TRANSACTION;
```

Once completed, the update in the record table can be confirmed at the last point (commit or rollback) of the current transaction. The correlation name **old** cannot be used in the conditional clause of the trigger where **DEFERRED UPDATE** is used. Therefore, you cannot create a trigger as the following.

```
CREATE CLASS foo (n int);
CREATE TRIGGER foo trigger
  DEFERRED UPDATE ON foo
  IF old.n = 100
  EXECUTE PRINT 'foo_trigger';
```

If you try to create a trigger as shown above, an error message is displayed and the trigger fails.

```
ERROR: Error compiling condition for 'foo_trigger' : old.n is not defined
```

The correlation name **old** can be used only with **AFTER**.

Java Stored Function/Procedure

Overview

Stored functions and procedures are used to implement complicated program logic that is not possible with SQL. They allow users to manipulate data more easily. Stored functions/procedures are blocks of code that have a flow of commands for data manipulation and are easy to manipulate and administer.

CUBRID supports to develop stored functions and procedures in Java. Java stored functions/procedures are executed on the JVM (Java Virtual Machine) hosted by CUBRID.

You can call Java stored functions/procedures from SQL statements or from Java applications using JDBC.

The advantages of using Java stored functions/procedures are as follows:

- **Productivity and usability** : Java stored functions/procedures, once created, can be reused anytime. They can be called from SQL statements or from Java applications using JDBC.
- **Excellent interoperability and portability** : Java stored functions/procedures use the Java Virtual Machine. Therefore, they can be used on any system where the Java Virtual Machine is available.

Environment Configuration for Java Stored Function/Procedure

To use Java-stored functions/procedures in CUBRID, you must have JRE (Java Runtime Environment) 1.6 or better installed in the environment where the CUBRID server is installed. You can download JRE from the Developer Resources for Java Technology (<http://java.sun.com>).

If the `java_stored_procedure` parameter in the CUBRID configuration file (`cubrid.conf`) is set to `yes`, CUBRID 64-bit needs a 64-bit Java Runtime Environment, and CUBRID 32-bit needs a 32-bit Java Runtime Environment. For example, when you run CUBRID 64-bit in the system in which a 32-bit JAVA Runtime Environment is installed, the following error may occur.

```
% cubrid server start demodb

This may take a long time depending on the amount of recovery works to do.
WARNING: Java VM library is not found :
/usr/java/jdk1.6.0_15/jre/lib/amd64/server/libjvm.so: cannot open shared object file: No
such file or directory.
Consequently, calling java stored procedure is not allowed
```

Execute the following command to check the JRE version if you have it already installed in the system.

```
% java -version
Java(TM) SE Runtime Environment (build 1.6.0_05-b13)
Java HotSpot(TM) 64-Bit Server VM (build 10.0-b19, mixed mode)
```

Windows Environment

For Windows, CUBRID loads the `jvm.dll` file to run the Java Virtual Machine. CUBRID first locates the `jvm.dll` file from the `PATH` environment variable and then loads it. If it cannot find the file, it uses the Java runtime information registered in the system registry.

You can configure the `JAVA_HOME` environment variable and add the directory in which the Java executable file is located to `Path`, by executing the command as follows: For information on configuring environment variables using GUI, see [Setting up the JDBC Environment](#).

- An example of installing 64 Bit JDK 1.6 and configuring the environment variables

```
% set JAVA_HOME=C:\jdk1.6.0
% set PATH=%PATH%;%JAVA_HOME%\jre\bin\server
```

- An example of installing 32 Bit JDK 1.6 and configuring the environment variables

```
% set JAVA_HOME=C:\jdk1.6.0
% set PATH=%PATH%;%JAVA_HOME%\jre\bin\client
```

To use other vendor's implementation instead of Sun's Java Virtual Machine, add the path of the **jvm.dll** file to the **PATH** variable during the installation.

Linux/UNIX Environment

For Linux/UNIX environment, CUBRID loads the **libjvm.so** file to run the Java Virtual Machine. CUBRID first locates the **libjvm.so** file from the **LD_LIBRARY_PATH** environment variable and then loads it. If it cannot find the file, it uses the **JAVA_HOME** environment variable. For Linux, glibc version 2.3.4 or higher is supported. The following example shows how to configure the Linux environment variable (e.g., **.profile**, **.cshrc**, **.bashrc**, **.bash_profile**, etc.).

- An example of installing 64 Bit JDK 1.6 and configuring the environment variables in a bash shell

```
% JAVA_HOME=/usr/java/jdk1.6.0_10
%
LD_LIBRARY_PATH=$JAVA_HOME/jre/lib/amd64:$JAVA_HOME/jre/lib/amd64/server:$LD_LIBRARY_PATH
% export JAVA_HOME
% export LD_LIBRARY_PATH
```

- An example of installing 32 Bit JDK 1.6 and configuring the environment variables in a bash shell

```
% JAVA_HOME=/usr/java/jdk1.6.0_10
%
LD_LIBRARY_PATH=$JAVA_HOME/jre/lib/i386:$JAVA_HOME/jre/lib/i386/client:$LD_LIBRARY_PATH
% export JAVA_HOME
% export LD_LIBRARY_PATH
```

- An example of installing 64 Bit JDK 1.6 and configuring the environment variables in a csh

```
% setenv JAVA_HOME /usr/java/jdk1.6.0_10
% setenv LD_LIBRARY_PATH
$JAVA_HOME/jre/lib/amd64:$JAVA_HOME/jre/lib/amd64/server:$LD_LIBRARY_PATH
% set path=($path $JAVA_HOME/bin .)
```

- An example of installing 32 Bit JDK 1.6 and configuring the environment variables in a csh shell

```
% setenv JAVA_HOME /usr/java/jdk1.6.0_10
% setenv LD_LIBRARY_PATH
$JAVA_HOME/jre/lib/i386:$JAVA_HOME/jre/lib/i386/client:$LD_LIBRARY_PATH
% set path=($path $JAVA_HOME/bin .)
```

To use other vendor's implementation instead of Sun's Java Virtual Machine, add the path of the JVM (**libjvm.so**) to the library path during the installation.

The path of the **libjvm.so** file can be different depending on the platform. For example, the path is the **\$JAVA_HOME/jre/lib/sparc** directory in a SUN Sparc machine.

How to Write Java Stored Function/Procedure

Steps to write a Java stored function/procedure are as follows:

- [Check the cubrid.conf file](#)
- [Write and compile the Java source code](#)
- [Load the compiled Java class into CUBRID](#)
- [Publish the loaded Java class](#)
- [Call the Java stored function/procedure](#)

Check the cubrid.conf file

By default, the **java_stored_procedure** is set to **no** in the **cubrid.conf** file. To use a Java stored function/procedure, this value must be changed to **yes**. For details on this value, see [Other Parameters](#) in Database Server Configuration.

Write and compile the Java source code

Compile the SpCubrid.java file as follows:

```
public class SpCubrid{
    public static String HelloCubrid() {
```

```

        return "Hello, Cubrid !!";
    }
    public static int SpInt(int i) {
        return i + 1;
    }
    public static void outTest(String[] o) {
        o[0] = "Hello, CUBRID";
    }
}
%javac SpCubrid.java

```

Here, the Java class method must be public static.

Load the compiled Java class into CUBRID

Load the compiled Java class into CUBRID.

```
% loadjava demodb
```

Publish the loaded Java class

Create a CUBRID stored function and publish the Java class as shown below.

```

csql> create function hello() return string
as language java
name 'SpCubrid.HelloCubrid()' return java.lang.String';

```

Call the Java stored function/procedure

Call the published Java stored function as follows:

```

csql> call hello() into :Hello;
Result
=====
'Hello, Cubrid !!'

```

Using Server-side Internal JDBC Driver

To access the database from a Java stored function/procedure, you must use the server-side JDBC driver. As Java stored functions/procedures are executed within the database, there is no need to make the connection to the server-side JDBC driver again. To acquire a connection to the database using the server-side JDBC driver, you can either use **"jdbc:default:connection:"** as the URL for JDBC connection, or call the **getDefaultConnection()** method of the **cubrid.jdbc.driver.CUBRIDDriver** class.

```

Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
Connection conn = DriverManager.getConnection("jdbc:default:connection:");

```

or

```
cubrid.jdbc.driver.CUBRIDDriver.getDefaultConnection();
```

If you connect to the database using the JDBC driver as shown above, the transaction in the Java stored function/procedure is ignored. That is, database operations executed in the Java stored function/procedure belong to the transaction that called the Java stored function/procedure. In the following example, **conn.commit()** method of the **Athlete** class is ignored.

```

import java.sql.*;
public class Athlete{
    public static void Athlete(String name, String gender, String nation code, String
event) throws SQLException{
        String sql="INSERT INTO ATHLETE (NAME, GENDER, NATION CODE, EVENT)" + "VALUES
(?, ?, ?, ?)";
        try{
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            Connection conn = DriverManager.getConnection("jdbc:default:connection:");
            PreparedStatement pstmt = conn.prepareStatement(sql);

            pstmt.setString(1, name);
            pstmt.setString(2, gender);

```

```

        pstmt.setString(3, nation code);
        pstmt.setString(4, event);
        pstmt.executeUpdate();

        pstmt.close();
        conn.commit();
        conn.close();
    } catch (Exception e) {
        System.err.println(e.getMessage());
    }
}
}

```

Connecting to Other Database

You can connect to another outside database instead of the currently connected one even when the server-side JDBC driver is being used. Acquiring a connection to an outside database is not different from a generic JDBC connection. For details, see JDBC API.

If you connect to other databases, the connection to the CUBRID database does not terminate automatically even when the execution of the Java method ends. Therefore, the connection must be explicitly closed so that the result of transaction operations such as **COMMIT** or **ROLLBACK** will be reflected in the database. That is, a separate transaction will be performed because the database that called the Java stored function/procedure is different from the one where the actual connection is made.

```

import java.sql.*;
public class SelectData {
    public static void SearchSubway(String[] args) throws Exception {
        Connection conn = null;
        Statement stmt = null;
        ResultSet rs = null;
        try {
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            conn =
            DriverManager.getConnection("jdbc:CUBRID:localhost:33000:demodb:::", "", "");
            String sql = "select line_id, line from line";
            stmt = conn.createStatement();
            rs = stmt.executeQuery(sql);
            while(rs.next()) {
                int host_year = rs.getString("host_year");
                String host_nation = rs.getString("host_nation");
                System.out.println("Host Year ==> " + host_year);
                System.out.println(" Host Nation==> " + host_nation);
                System.out.println("\n=====\n");
            }
            rs.close();
            stmt.close();
            conn.close();
        } catch ( SQLException e ) {
            System.err.println(e.getMessage());
        } catch ( Exception e ) {
            System.err.println(e.getMessage());
        } finally {
            if ( conn != null ) conn.close();
        }
    }
}

```

When the Java stored function/procedure being executed should run only on JVM located in the database server, you can check where it is running by calling `System.getProperty("cubrid.server.version")` from the Java program source. The result value is the database version if it is called from the database; otherwise, it is **NULL**.

loadjava Utility

Description

To load a compiled Java or JAR (Java Archive) file into CUBRID, use the **loadjava** utility. If you load a Java *.class or *.jar file using the **loadjava** utility, the file is moved to the specified database path.

Syntax

```
loadjava <option> database-name java-class-file
```

- *database-name* : The name of the database where the Java file is to be loaded.
- *java-class-file* : The name of the Java class or jar file to be loaded.
- *<option>* :
- **-y** : Automatically overwrites a class file with the same name, if any. The default value is **no**. If you load the file without specifying the **-y** option, you will be prompted to ask if you want to overwrite the class file with the same name (if any).

Loaded Java Class Publish

Overview

In CUBRID, it is required to publish Java classes to call Java methods from SQL statements or Java applications. You must publish Java classes by using call specifications because it is not known how a function in a class will be called by SQL statements or Java applications when Java classes are loaded.

Call Specifications

To use a Java stored function/procedure in CUBRID, you must write call specifications. With call specifications, Java function names, parameter types, return values and their types can be accessed by SQL statements or Java applications. To write call specifications, use **CREATE FUNCTION** or **CREATE PROCEDURE** statement. Java stored function/procedure names are not case sensitive. The maximum number of characters a Java stored function/procedure can have is 256. The maximum number of parameters a Java stored function/procedure can have is 64.

Syntax

```
CREATE {PROCEDURE procedure name[(param[, param]...)] | FUNCTION function name[(param[, param]...)] RETURN sql_type }
{IS | AS} LANGUAGE JAVA
NAME 'method_fullname (java_type_fullname[, java_type_fullname]... [return java_type_fullname]';

parameter name [IN|OUT|IN OUT|INOUT] sql_type
(default IN)
```

If the parameter of a Java stored function/procedure is set to **OUT**, it will be passed as a one-dimensional array whose length is 1. Therefore, a Java method must store its value to pass in the first space of the array.

Example

```
CREATE FUNCTION Hello() RETURN VARCHAR
AS LANGUAGE JAVA
NAME 'SpCubrid.HelloCubrid() return java.lang.String';
CREATE FUNCTION Sp_int(i int) RETURN int
AS LANGUAGE JAVA
NAME 'SpCubrid.SpInt(int) return int';

CREATE PROCEDURE Phone Info(name varchar, phoneno varchar)
AS LANGUAGE JAVA
NAME 'PhoneNumber.Phone(java.lang.String, java.lang.String)';
```

When a Java stored function/procedure is published, it is not checked whether the return definition of the Java stored function/procedure coincides with the one in the declaration of the Java file. Therefore, the Java stored

function/procedure follows the *sql_type* return definition provided at the time of registration. The return definition in the declaration is significant only as user-defined information.

Data Type Mapping

In call specifications, the data types SQL must correspond to the data types of Java parameter and return value. The following table shows SQL/Java data types allowed in CUBRID.

Data Type Mapping

SQL Type	Java Type
CHAR, VARCHAR	java.lang.String, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.lang.Byte, java.lang.Short, java.lang.Integer, java.lang.Long, java.lang.Float, java.lang.Double, java.math.BigDecimal, byte, short, int, long, float, double
NUMERIC, SHORT, INT, FLOAT, DOUBEL, CURRENCY	java.lang.Byte, java.lang.Short, java.lang.Integer, java.lang.Long, java.lang.Float, java.lang.Double, java.math.BigDecimal, java.lang.String, byte, short, int, long, float, double
DATE, TIME, TIMESTAMP	java.sql.Date, java.sql.Time, java.sql.Timestamp, java.lang.String
SET, MULTISSET, SEQUENCE	java.lang.Object[], java primitive type array, java.lang.Integer[] ...
OBJECT	cubrid.sql.CUBRIDOID
CURSOR	cubrid.jdbc.driver.CUBRIDResultSet

Checking the Published Java Stored Function/Procedure Information

You can check the information on the published Java stored function/procedure. The **db_stored_procedure** system virtual table provides virtual table and the **db_stored_procedure_args** system virtual table. The **db_stored_procedure** system virtual table provides the information on stored names and types, return types, number of parameters, Java class specifications, and the owner. The **db_stored_procedure_args** system virtual table provides the information on parameters used in the stored function/procedure.

```

SELECT * from db_stored_procedure;
sp_name      sp_type      return_type      arg_count
sp_name      sp_type      return_type      arg_count      lang
target
owner
=====
'hello'      'FUNCTION'    'STRING'        0              'JAVA' 'SpCu
brid.HelloCubrid() return java.lang.String' 'DBA'

'sp_int'     'FUNCTION'    'INTEGER'       1              'JAVA' 'SpCu
brid.SpInt(int) return int' 'DBA'

'athlete add' 'PROCEDURE'   'void'          4              'JAVA' 'Athl
ete.Athlete(java.lang.String, java.lang.String, java.lang.String,
java.lang.String)' 'DBA'
SELECT * from db_stored_procedure_args;
sp name      index of arg name data type      mode
=====
'sp_int'     0 'i'           'INTEGER'      'IN'
'athlete add' 0 'name'        'STRING'       'IN'
'athlete_add' 1 'gender'      'STRING'       'IN'
'athlete_add' 2 'nation_code' 'STRING'       'IN'
'athlete_add' 3 'event'       'STRING'       'IN'
    
```

Deleting Java Stored Functions/Procedures

You can delete published Java stored functions/procedures in CUBRID. To delete a Java function/procedure, use the **DROP FUNCTION *function_name*** or **DROP PROCEDURE *procedure_name*** statement. Also, you can delete

multiple Java stored functions/procedures at a time with several *function_names* or *procedure_names* separated by a comma (,).

A Java stored function/procedure can be deleted only by the user who published it or by DBA members. For example, if a **PUBLIC** user published the 'sp_int' Java stored function, only the **PUBLIC** or **DBA** members can delete it.

```
drop function hello[, sp int]
drop procedure Athlete_Add
```

Java Stored Function/Procedure Call

Using CALL Statement

You can call the Java stored functions/procedures by using a **CALL** statement, from SQL statements or Java applications.

The following shows how to call them by using the **CALL** statement. The name of the Java stored function/procedure called from a **CALL** statement is not case sensitive.

Syntax

```
CALL {procedure_name ([param[, param]...)] | function_name ([param[, param]...)]
INTO :host_variable
param {literal | :host_variable}
```

Example

```
call Hello() into :HELLO;
call Sp_int(3) into :i;
call phone_info('Tom','016-111-1111');
```

In CUBRID, the Java functions/procedures are called by using the same **CALL** statement. Therefore, the **CALL** statement is processed as follows:

- It is processed as a method if there is a target class in the **CALL** statement.
- If there is no target class in the **CALL** statement, it is checked whether a Java stored function/procedure is executed or not; a Java stored function/procedure will be executed if one exists.
- If no Java stored function/procedure exists in step 2 above, it is checked whether a method is executed or not; a method will be executed if one with the same name exists.

The following error occurs if you call a Java stored function/procedure that does not exist.

```
CALL deposit()
ERROR: Stored procedure/function 'deposit' is not exist.

CALL deposit('Tom', 3000000)
ERROR: Methods require an object as their target.
```

If there is no argument in the **CALL** statement, a message "ERROR: Stored procedure/function 'deposit' is not exist." appears because it can be distinguished from a method. However, if there is an argument in the **CALL** statement, a message "ERROR: Methods require an object as their target." appears because it cannot be distinguished from a method.

If the **CALL** statement is nested within another **CALL** statement calling a Java stored function/procedure, or if a subquery is used in calling the Java function/procedure, the **CALL** statement is not executed.

```
call phone_info('Tom', call sp_int(999));
call phone_info((select * from Phone where id='Tom'));
```

If an exception occurs during the execution of a Java stored function/procedure, the exception is logged and stored in the *dbname*_java.log file. To display the exception on the screen, change a handler value of the **\$CUBRID/java/logging.properties** file to "java.lang.logging.ConsoleHandler." Then, the exception details are displayed on the screen.

Calling from SQL Statement

You can call a Java stored function from a SQL statement as shown below.

```
select Hello() from db root;
select sp_int(99) from db_root;
```

You can use a host variable for the IN/OUT data type when you call a Java stored function/procedure as follows:

```
SELECT 'Hi' INTO :out_data FROM db root;
CALL test_out(:out_data);
SELECT :out_data FROM db_root;
```

The first clause calls a Java stored procedure in out mode by using a parameter variable; the second is a query clause retrieving the assigned host variable `out_data`.

Calling from Java Application

To call a Java stored function/procedure from a Java application, use a **CallableStatement** object.

Create a phone class in the CUBRID database.

```
CREATE TABLE phone(
  name varchar(20),
  phoneno varchar(20)
)
```

Compile the following `PhoneNumber.java` file, load the Java class file into CUBRID, and publish it.

```
import java.sql.*;
import java.io.*;
public class PhoneNumber{    public static void Phone(String name, String phoneno) throws
Exception{
    String sql="INSERT INTO PHONE(NAME, PHONENO)+" "VALUES (?, ?)";
    try{
        Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        Connection conn = DriverManager.getConnection("jdbc:default:connection:");
        PreparedStatement pstmt = conn.prepareStatement(sql);

        pstmt.setString(1, name);
        pstmt.setString(2, phoneno);
        pstmt.executeUpdate();
        pstmt.close();
        conn.commit();
        conn.close();
    } catch (SQLException e) {
        System.err.println(e.getMessage());
    }
}
}
create PROCEDURE phone_info(name varchar, phoneno varchar)
as language java
name 'PhoneNumber.Phone(java.lang.String, java.lang.String)';
```

Create and run the following Java application.

```
import java.sql.*;
public class StoredJDBC{
    public static void main(){
        Connection conn = null;
        Statement stmt= null;
        int result;
        int i;
        try{
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            conn =
DriverManager.getConnection("jdbc:CUBRID:localhost:33000:demodb:::", "", "");
            CallableStatement cs;
            cs = conn.prepareCall("call PHONE INFO(?, ?)");
            cs.setString(1, "Jane");
            cs.setString(2, "010-1111-1111");
            cs.executeUpdate();
            conn.commit();
            cs.close();
            conn.close();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

```
}}
```

Retrieve the phone class after executing the program above; the following result would be displayed.

```
SELECT * from phone;
name          phoneno
=====
'Jane'        '010-111-1111'
```

Caution

Returning Value of Java Stored Function/Procedure and Precision Type on IN/OUT

To limit the return value of Java stored function/procedure and precision type on IN/OUT, CUBRID processes as follows:

Checks the `sql_type` of the Java stored function/procedure.

Passes the value returned by Java to the database with only the type converted if necessary, ignoring the number of digits defined during creating the Java stored function/procedure. In principle, the user manipulates the passed data directly in the database.

Take a look at the following `typestring()` Java stored function.

```
public class JavaSP1{
    public static String typestring(){
        String temp = " ";
        for(int i=0 i< 1 i++)
            temp = temp + "1234567890";
        return temp;
    }
}

create function typestring() return char(5)
as language java
name 'JavaSP1.typestring() return java.lang.String';

Call typestring()

Result
=====
' 1234567890'
```

Returning `java.sql.ResultSet` in Java Stored Procedure

In CUBRID, you must use `CURSOR` as the data type when you declare a Java stored function/procedure that returns a `java.sql.ResultSet`.

```
create function rset() return cursor
as language java
name 'JavaSP2.TResultSet() return java.sql.ResultSet'
```

Before the Java file returns `java.sql.ResultSet`, it is required to cast to the `CUBRIDResultSet` class and then to call the `setReturnable()` method.

```
public static class JavaSP2 {
public static ResultSet TResultSet(){
    try{
        Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        Connection conn = DriverManager.getConnection("jdbc:default:connection:");
        ((CUBRIDConnection)con).setCharset("euc kr");
        String sql = "select * from station";
        Statement stmt=con.createStatement();
        ResultSet rs = stmt.executeQuery(sql);
        ((CUBRIDResultSet)rs).setReturnable();
        return rs;
    } catch (Exception e) {
        e.printStackTrace();
    }
    return null;
}
}
```

```
}

```

In the calling block, you must set the OUT argument with **Types.JAVA_OBJECT**, get the argument to the **getObject()** function, and then cast it to the **java.sql.ResultSet** type before you use it. In addition, the **java.sql.ResultSet** is only available to use in **CallableStatement** of JDBC.

```
import java.sql.*;
public class TestResultSet{
    public static void main(String[] args) {
        Connection conn = null;
        Statement stmt= null;
        int result;
        int i;

        try{
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
            conn = DriverManager.getConnection("jdbc:CUBRID:localhost:33000:demodb:::", "", "");

            CallableStatement cstmt = con.prepareCall("=?=CALL rset()");
            cstmt.registerOutParameter(1, Types.JAVA_OBJECT);
            cstmt.execute();
            ResultSet rs = (ResultSet) cstmt.getObject(1);
            while(rs.next()) {
                System.out.println(rs.getString(1));
            }
            rs.close();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

You cannot use the **ResultSet** as an input argument. If you pass it to an IN argument, an error occurs. An error also occurs when calling a function that returns **ResultSet** in a non-Java environment.

IN/OUT of Set Type in Java Stored Function/Procedure

If the set type of the Java stored function/procedure in CUBRID is IN OUT, the value of the argument changed in Java must be applied to IN OUT. When the set type is passed to the OUT argument, it must be passed as a two-dimensional array.

```
Create procedure setoid(x in out set, z object)
as language java name
'SetOIDTest.SetOID(cubrid.sql.CUBRIDOID[][], cubrid.sql.CUBRIDOID';

public static void SetOID(cubrid.sql.CUBRID[][] set, cubrid.sql.CUBRIDOID aoid){
    Connection conn=null;
    Statement stmt=null;
    String ret="";
    Vector v = new Vector();
    cubrid.sql.CUBRIDOID[] set1 = set[0];
    try {
        if(set1!=null) {
            int len = set1.length;
            int i = 0;
            for (i=0 i< len i++)
                v.add(set1[i]);
        }
        v.add(aoid);
        set[0]=(cubrid.sql.CUBRIDOID[]) v.toArray(new cubrid.sql.CUBRIDOID[]{});
    } catch(Exception e) {
        e.printStackTrace();
        System.err.println("SQLException:"+e.getMessage());
    }
}
```

Using OID in Java Stored Function/Procedure

In case of using the OID type value for IN/OUT in CUBRID, use the value passed from the server.

```
create procedure tOID(i inout object, q string)
as language java
```

```
name 'OIDtest.tOID(cubrid.sql.CUBRIDOID[], java.lang.String)';

public static void tOID(CUBRIDOID[] oid, String query)
{
    Connection conn=null;
    Statement stmt=null;
    String ret="";

    try {
        Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        conn=DriverManager.getConnection("jdbc:default:connection:");

        conn.setAutoCommit(false);
        stmt = conn.createStatement();
        ResultSet rs = stmt.executeQuery(query);
        System.out.println("query:"+ query);
        while(rs.next()) {
            oid[0]=(CUBRIDOID)rs.getObject(1);
            System.out.println("oid:"+oid[0].getTableName());
        }
        stmt.close();
        conn.close();
    } catch (SQLException e) {
        e.printStackTrace();
        System.err.println("SQLException:"+e.getMessage());
    } catch (Exception e) {
        e.printStackTrace();
        system.err.println("Exception:"+ e.getMessage());
    }
}
```

METHOD

Overview

This chapter describes methods (software routines) that extend or customize the features of the CUBRID database system.

The methods are written in C and called by the **CALL** or **EVALUATE** statement. A method program is loaded and linked with the application currently running by the dynamic loader when the method is called. The return value created as a result of the method execution is passed to the caller.

This chapter describes the following topics:

- Method Types
- Calling a Method

METHOD Type

The CSQL language supports the following two types of methods: class and instance methods.

- The **class method** is a method called by a class object. It is usually used to create a new class instance or to initialize it. It is also used to access or update class attributes.
- The **instance method** is a method called by a class instance. It is used more often than the class method because most operations are executed in the instance. For example, an instance method can be written to calculate or update the instance attribute. This method can be called from any instance of the class in which the method is defined or of the subclass that inherits the method.

The method inheritance rules are similar to those of the attribute inheritance. The subclass inherits classes and instance methods from the super class. The subclass has only the name of a class or instance method definition inherited from the super class.

The rules for resolving method name conflicts are same as those for attribute name conflicts. For details about attribute/method inheritance conflicts, see [Overview](#) in Class Conflict Resolution.

Calling METHOD

Overview

Methods are executed by the **CALL** or **EVALUATE** statement, and their results are returned the same way as the query results.

These statements are also used to call a method from a query. (The **CALL** or **EVALUATE** keyword is omitted.)

CALL Statement

Description

In CUBRID, the **CALL** statement is used to call a method defined in the database. Both table and record methods can be called by the **CALL** statement.

Syntax

```
CALL method_call [ ; ]
method_call :
• method_name ( [ arg_value [ {, arg_value }_ ] ] ) ON call_target [ to_variable ]
• method_name ( call_target [, arg_value [ {, arg_value }_ ] ] ) [ to_variable ]
arg_value :
• any CSQL expression
call_target :
• an object-valued expression
```

```
to variable :
• INTO variable
• TO variable
```

- The *method_name* is either the method name defined in the table or the system-defined method name provided with CUBRID. A method requires one or more parameters. If there is no parameter for the method, a set of blank parentheses must be used.
- *call_target* can use an object-valued expression that contains a class name, a variable, another method call (which returns an object). To call a class method for a class object, you must place the **CLASS** keyword before the *call_target*. In this case, the table name must be the name of the class where the table method is defined. To call a record method, you must specify the expression representing the record object. You can optionally store the value returned by the table or record method in the *to_variable*. This returned variable value can be used in the **CALL** statement just like the *call_target* or *arg_value* parameter.
- Calling nested methods is possible when other *method_call* is the *call_target* of the method or given as one of the *arg_value* parameters.

EVALUATE Statement

Description

The **EVALUATE** statement is also used to call a method defined in the database.

In the **EVALUATE** statement, a method call is a *term* in an expression. If the method returns a constant value, another constant (or a method returning a constant) can also be a term in an expression. Both class and instance methods can be called by the **EVALUATE** statement.

Syntax

```
EVALUATE expression [ ; ]

expression:
• [ + | - ] term [ { + | - | * | / } term ]

term:
• method call

method_call :
• method_name ( call_target [, arg_value [ {, arg_value }_ ] ] ) [ to_variable ]
  method_name ( [ arg value [ {, arg value } ] ] )
  ON call_target [ to_variable ]

arg value :
• literal
• variable
• expression

call target :
• CLASS class_name
• variable
• expression
• method_call

to variable :
• INTO variable
• TO variable
```

In the **EVALUATE** statement, the target argument for the specified method is represented in the parentheses following the *method_name*. The target can be the first field in the list, followed by method arguments. If the method executed is a class method, the **CLASS** keyword must precede the target class as the first field in the list. If only the method arguments are included in the parentheses, the *call_target* should be in the **ON** clause.

The **EVALUATE** statement also supports nested method calls by allowing one method call to be expressed as the target or the argument of another method. In these types of expressions, the result of the inner method is used to determine that of the outer method.

Partitioning

What is Partitioning?

Partitioning is a method by which a table is divided into multiple independent logical units. Each logical unit used in partitioning is called a partition. Partitioning can enhance manageability, performance and availability. Some advantages of partitioning are as follows:

- Improved management of large capacity tables
- Improved performance by narrowing the range of access when retrieving data
- Improved performance and decreased physical loads by distributing disk I/O
- Decreased possibility of data corruption and improved availability by partitioning a table into multiple chunks
- Optimized storage cost

Three types of partitioning methods are supported by CUBRID: range partitioning, hash partitioning, and list partitioning.

The maximum number of partitions cannot exceed 1,024. Each partition of a table is created as its subtable. The subtables created by the partitioning process cannot be altered or deleted by users. The name of the subtable is stored in the system table in a 'class_name__p__partition_name' format. Database users can check the partitioning information in the db_class and db_partition virtual tables. They can also check the information by using the ;sc <table name> command in the CUBRID Manager or the CSQL Interpreter.

Range Partitioning

Range Partitioning Definition

Description

You can define a range partition by using the **PARTITION BY RANGE** clause.

Syntax

```
CREATE TABLE (
...
)
PARTITION BY RANGE ( <partition_expression> ) (
PARTITION <partition_name> VALUES LESS THAN ( <range_value> ),
PARTITION <partition_name> VALUES LESS THAN ( <range_value> ) ),
... )
)
```

- *partition_expression* : Specifies the partition expression. The expression can be specified by the name of the column to be partitioned or by a function. For details of the data types and functions available, see [Data Types Available for Partition Expression](#).
- *partition_name* : Specifies the partition name.
- *range_value* : Specifies the partition-by value.

Example 1

The following example shows how to create the participant2 table with the participating countries, and insert data that partitions the years into before and after the 2000 Olympic Games. When inserting data, the countries that participated in the 1988 and 1996 Olympic Games are stored in before_2000; the rest of them are stored in before_2008.

```
CREATE TABLE participant2 (host_year INT, nation CHAR(3), gold INT, silver INT, bronze INT)
PARTITION BY RANGE (host_year)
(PARTITION before_2000 VALUES LESS THAN (2000),
PARTITION before_2008 VALUES LESS THAN (2008) );

INSERT INTO participant2 VALUES (1988, 'NZL', 3, 2, 8);
```



```
INSERT INTO participant2 VALUES (1988, 'CAN', 3, 2, 5);
INSERT INTO participant2 VALUES (1996, 'KOR', 7, 15, 5);
INSERT INTO participant2 VALUES (2000, 'RUS', 32, 28, 28);
INSERT INTO participant2 VALUES (2004, 'JPN', 16, 9, 12);
```

Example 2

As shown below, the partition key value in a range partition is **NULL**, the data are stored in the first partition.

```
INSERT INTO participant2 VALUES (NULL, 'AAA', 0, 0, 0);
```

Caution

- The maximum number of partitions possible for a given table is 1024.
- If the partition key value is **NULL**, the data is stored in the first partition (see Example 2).

Range Partitioning Redefinition

Description

You can redefine a partition by using the **REORGANIZE PARTITION** clause of the **ALTER** statement. By redefining partitions, you can combine multiple partitions into one or divide one into multiple.

Syntax

```
ALTER {TABLE | CLASS} <table name>
REORGANIZE PARTITION
<alter partition name comma list>
INTO ( <partition definition comma list> )

partition definition comma list:
PARTITION <partition name> VALUES LESS THAN ( <range_value> ),...
```

- *table_name* : Specifies the name of the table to be redefined.
- *alter partition name comma list* : Specifies the partition to be redefined. Multiple partitions are separated by commas (,).
- *partition definition comma list* : Specifies the redefined partitions. Multiple partitions are separated by commas (,).

Example 1

The following example shows how to perform repartitioning the before_2000 partition into the before_1996 and before_2000 partitions.

```
CREATE TABLE participant2 ( host year INT, nation CHAR(3), gold INT, silver INT, bronze
INT)
PARTITION BY RANGE (host year)
( PARTITION before 2000 VALUES LESS THAN (2000),
PARTITION before_2008 VALUES LESS THAN (2008) );

ALTER TABLE participant2 REORGANIZE PARTITION before 2000 INTO (
PARTITION before 1996 VALUES LESS THAN (1996),
PARTITION before 2000 VALUES LESS THAN (2000)
);
```

Example 2

The following example shows how to combine two partitions redefined in Example 1 back into a single before_2000 partition.

```
ALTER TABLE participant2 REORGANIZE PARTITION before 1996, before 2000 INTO
(PARTITION before_2000 VALUES LESS THAN (2000) );
```

Caution

- When redefining a range or list partition, duplicate ranges or values are not allowed.
- The **REORGANIZE PARTITION** clause cannot be used to change the partition table type. For example, a range partition cannot be changed to a hash partition, or vice versa.

- The maximum number of partitions cannot exceed 1,024. There must be at least one partition remaining after deleting partitions. In a range-partitioned table, only adjacent partitions can be redefined.

Adding Range Partitioning

Description

You can add range partitions by using the **ADD PARTITION** clause of the **ALTER** statement.

Syntax

```
ALTER {TABLE | CLASS} <table_name>
ADD PARTITION <partition_definitions comma list>
partition definition comma list:
PARTITION <partition_name> VALUES LESS THAN ( <range_value> ),...
```

- *table_name* : Specifies the name of the table to which partitions are added.
- *partition definition comma list* : Specifies the partitions to be added. Multiple partitions are separated by commas (,).

Example

Currently, the partition before the 2008 Olympic Games is defined in the participant2 table. The following example shows how to add the before_2012 and before_2016 partitions; the former will store the information about the 2012 Olympic Games and the latter will store the information about the 2016 Olympic Games.

```
ALTER TABLE participant2 ADD PARTITION (
PARTITION before_2012 VALUES LESS THAN (2012),
PARTITION before_2016 VALUES LESS THAN MAXVALUE );
```

Caution

- When a range partition is added, only the partition by value greater than the existing partition value can be added. Therefore, as shown in the above example, if the maximum value is specified by **MAXVALUE**, no more partitions can be added (you can add partitions by changing the **MAXVALUE** value by redefining the partition).
- To add the partition by value smaller than the existing partition value, use the redefining partitions (see [Range Partitioning Redefinition](#)).

Dropping Range Partitioning

Description

You can drop a partition by using the **DROP PARTITION** clause of the **ALTER** statement.

Syntax

```
ALTER {TABLE | CLASS} <table_name>
DROP PARTITION <partition_name>
```

- *table_name* : Specifies the name of the partitioned table.
- *partition_name* : Specifies the name of the partition to be dropped.

Example

The following example shows how to drop the before_2000 partition in the participant2 table.

```
ALTER TABLE participant2 DROP PARTITION before_2000;
```

Caution

- When dropping a partitioned table, all stored data in the partition are also dropped.
- If you want to change the partitioning of a table without losing data, use the **ALTER TABLE...REORGANIZE PARTITION** statement (see [Range Partitioning Redefinition](#)).
- The number of rows deleted is not returned when a partition is dropped. If you want to delete the data, but want to maintain the table and partitions, use the **DELETE** statement.

Hash Partitioning

Hash Partitioning Definition

Description

You can define a hash partition by using the **PARTITION BY HASH** clause.

Syntax

```
CREATE TABLE (
...
)
( PARTITION BY HASH ( <partition_expression> )
PARTITIONS ( <number_of_partitions> )
)
```

- *partition_expression* : Specifies a partition expression. The expression can be specified by the name of the column to be partitioned or by a function.
- *number_of_partitions* : Specifies the number of partitions.

Example 1

The following example shows how to create the nation2 table with country codes and country names, and define 4 hash partitions based on code values. Only the number of partitions, not the name, is defined in hash partitioning; names such as p0 and p1 are assigned automatically.

```
CREATE TABLE nation2
( code CHAR(3),
name VARCHAR(50) )
PARTITION BY HASH ( code) PARTITIONS 4;
```

Example 2

The following example shows how to insert data to the hash partition created in the example 1. When a value is inserted into a hash partition, the partition to store the data is determined by the hash value of the partition key. If the partition key value is **NULL**, the data is stored in the first partition.

```
INSERT INTO nation2 VALUES ('KOR','Korea');
INSERT INTO nation2 VALUES ('USA','USA United States of America');
INSERT INTO nation2 VALUES ('FRA','France');
INSERT INTO nation2 VALUES ('DEN','Denmark');
INSERT INTO nation2 VALUES ('CHN','China');
INSERT INTO nation2 VALUES (NULL,'AAA');
```

Caution

The maximum number of partitions cannot exceed 1024.

Hash Partitioning Redefinition

Description

You can redefine a partition by using the **COALESCE PARTITION** clause of the **ALTER** statement. Instances are preserved if the hash partition is redefined.

Syntax

```
ALTER {TABLE | CLASS} <table name>
COALESCE PARTITION <unsigned integer>
```

- *table_name* : Specifies the name of the table to be redefined.
- *unsigned integer* : Specifies the number of partitions to be deleted.

Example

The following example shows how to decrease the number of partitions in the nation2 table from 4 to 2.

```
ALTER TABLE nation2 COALESCE PARTITION 2;
```

Caution

- Decreasing the number of partitions is only available.
- To increase the number of partitions, use the **ALTER TABLE ... ADD PARTITION** statement as in range partitioning (see [Adding Range Partitioning](#) For details).
- There must be at least one partition remaining after redefining partitions.

List Partitioning

List Partitioning Definition

Description

You can define a list partition by using the **PARTITION BY LIST** statement.

Syntax

```
CREATE TABLE (
...
)
PARTITION BY LIST ( <partition_expression> ) (
PARTITION <partition_name> VALUES IN ( <partition_value_list> ),
PARTITION <partition_name> VALUES IN ( <partition_value_list>
,
...
);
```

- *partition_expression* : Specifies a partition expression. The expression can be specified by the name of the column to be partitioned or by a function. For details on the data types and functions available, see [Data Types Available for Partition Expression](#).
- *partition_name* : Specifies the partition name.
- *partition_value_list* : Specifies the list of the partition by values.

Example 1

The following example shows how to create the athlete2 table with athlete names and sport events, and define list partitions based on event values.

```
CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball')
);
```

Example 2

The following example shows how to insert data to the list partition created in the example 1. In the last query of the example 2, if you insert an argument that has not been specified in the partition expression of the example 1, data inserting fails.

```
INSERT INTO athlete2 VALUES ('Hwang Young-Cho', 'Athletics');
INSERT INTO athlete2 VALUES ('Lee Seung-Yuop', 'Baseball');
INSERT INTO athlete2 VALUES ('Moon Dae-Sung', 'Taekwondo');
INSERT INTO athlete2 VALUES ('Cho In-Chul', 'Judo');
INSERT INTO athlete2 VALUES ('Hong Kil-Dong', 'Volleyball');
```

Example 3

The following example shows in which an error occurs with no data inserted when the partition key value is **NULL**. To define a partition where a **NULL** value can be inserted, define one that has a list including a **NULL** value as in the event3 partition as below.

```
INSERT INTO athlete2 VALUES ('Hong Kil-Dong', 'NULL');

CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics ' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball', NULL)
);
```

Caution

The maximum number of partitions cannot exceed 1,024.

List Partitioning Redefinition

Description

You can redefine a partition by using the **REORGANIZE PARTITION** clause of the **ALTER** statement. By redefining partitions, you can combine multiple partitions into one or divide one into multiple.

Syntax

```
ALTER {TABLE | CLASS} <table_name>
REORGANIZEPARTITION
<alter partition name comma list>
INTO ( <partition definition comma list> )
partition definition comma list:
PARTITION <partition_name> VALUES IN ( <partition_value_list>),...
```

- *table_name* : Specifies the name of the table to be redefined.
- *alter partition name comma list* : Specifies the partition to be redefined. Multiple partitions are separated by commas (.).
- *partition definition comma list* : Specifies the redefined partitions. Multiple partitions are separated by commas (,).

Example 1

The following example shows how to create the athlete2 table partitioned by the list of sport events, and redefine the event2 partition to be divided into event2_1 (Judo) and event2_2 (Taekwondo, Boxing).

```
CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics ' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball')
);

ALTER TABLE athlete2 REORGANIZE PARTITION event2 INTO
(PARTITION event2_1 VALUES IN ('Judo'),
PARTITION event2_2 VALUES IN ( 'Taekwondo','Boxing'));
```

Example 2

The following example shows how to combine the event2_1 and event2_2 partitions divided in Example 1 back into a single event2 partition.

```
ALTER TABLE athlete2 REORGANIZE PARTITION event2_1, event2_2 INTO
(PARTITION event2 VALUES IN('Judo','Taekwondo','Boxing'));
```

Dropping List Partitioning

Description

You can drop a partition by using the **DROP PARTITION** clause of the **ALTER** statement.

Syntax

```
ALTER {TABLE | CLASS} <table_name>
DROP PARTITION <partition_name>
```

- *table_name* : Specifies the name of the partitioned table.
- *partition_name* : Specifies the name of the partition to be dropped.

Example

The following example shows how to create the athlete2 table partitioned by the list of sport events and drop the event3 partition.

```
CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball')
);

ALTER TABLE athlete2 DROP PARTITION event3;
```

Partitioning Management

Retrieving and Manipulating Data in Partitioning

Description

When retrieving data, the **SELECT** statement can be used not only for partitioned tables but also for each partition.

Example

The following example shows how to create the athlete2 table to be partitioned by the list of sport events, insert data, and retrieve the event1 and event2 partitions.

```
CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball')
);

INSERT INTO athlete2 VALUES ('Hwang Young-Cho', 'Athletics');
INSERT INTO athlete2 VALUES ('Lee Seung-Yuop', 'Baseball');
INSERT INTO athlete2 VALUES ('Moon Dae-Sung', 'Taekwondo');
INSERT INTO athlete2 VALUES ('Cho In-Chul', 'Judo');
SELECT * from athlete2 p event1;
name                event
=====
'Hwang Young-Cho'   'Athletics'

SELECT * from athlete2 p event2;
name                event
=====
'Moon Dae-Sung'    'Taekwondo'
'Cho In-Chul'      'Judo'
```

Caution

Data manipulation such as insert, update and delete for each partition of the partitioned table is not allowed.

Moving Data by Changing Partitioning Key Value

Description

If a partition key value is changed, the changed instance can be moved to another partition by the partition expression.

Example

The following example shows how to move the instance to another partition by changing the partition key value. If you change the sport event information of Hwang Young-Cho in the event1 partition from Athletics to Football, the instance is moved to the event3 partition.

```
CREATE TABLE athlete2( name VARCHAR(40), event VARCHAR(30) )
PARTITION BY LIST (event) (
PARTITION event1 VALUES IN ('Swimming', 'Athletics ' ),
PARTITION event2 VALUES IN ('Judo', 'Taekwondo','Boxing'),
PARTITION event3 VALUES IN ('Football', 'Basketball', 'Baseball')
);

INSERT INTO athlete2 VALUES ('Hwang Young-Cho', 'Athletics');
INSERT INTO athlete2 VALUES ('Lee Seung-Yuop', 'Baseball');

  name                event
=====
'Hwang Young-Cho'    'Athletics'

UPDATE athlete2 SET event = 'Football' WHERE name = 'Hwang Young-Cho';

SELECT * FROM athlete2 p event3;
  name                event
=====
'Lee Seung-Yuop'    'Baseball'
'Hwang Young-Cho'    'Football'
```

Caution

Be aware that when moving data between partitions by changing a partition key value, it can cause performance degradation due to internal deletions and insertions.

Altering Regular Table into Partitioning Table

Description

To alter a regular table into a partitioned one, use the **ALTER TABLE** statement. Three partitioning methods can be used with the **ALTER TABLE** statement. The data in the existing table are moved to and stored in each partition according to the partition definition.

Syntax

```
ALTER {TABLE | CLASS} table_name
PARTITION BY {RANGE | HASH | LIST } ( <partition_expression> )
( PARTITION partition_name VALUES LESS THAN { MAXVALUE | ( <partition_value_option> ) }
| PARTITION partition_name VALUES IN ( <partition_value_option list> ) ]
| PARTITION <UNSIGNED_INTEGER> )

<partition_expression>
expression_
<partition value option>
literal_
```

- *table_name* : Specifies the name of the table to be altered.
- *partition_expression* : Specifies a partition expression. The expression can be specified by the name of the column to be partitioned or by a function. For details on the data types and functions available, see [Data Types Available for Partition Expressions](#).
- *partition_name* : Specifies the name of the partition.
- *partition_value_option* : Specifies the value or the value list on which the partition is based.

Example

The following are examples of altering the record table into a range, list and hash table respectively.

```
ALTER TABLE record PARTITION BY RANGE (host_year)
( PARTITION before_1996 VALUES LESS THAN (1996),
  PARTITION after_1996 VALUES LESS THAN MAXVALUE);

ALTER TABLE record PARTITION BY list (unit)
( PARTITION time_record VALUES IN ('Time'),
  PARTITION kg_record VALUES IN ('kg'),
  PARTITION meter_record VALUES IN ('Meter'),
  PARTITION score_record VALUES IN ('Score') );

ALTER TABLE record
PARTITION BY HASH (score) PARTITIONS 4;
```

Caution

If there is data that does not satisfy the partition condition, partitions cannot be defined.

Altering Partitioning Table into Regular Table

Description

To alter an existing partitioned table into a regular one, use the **ALTER TABLE** statement.

Syntax

```
ALTER {TABLE | CLASS} <table name>
REMOVE PARTITIONING
```

- *table_name* : Specifies the name of the table to be altered.

Example

The following example shows how to alter the partitioned table of name "nation2" into a regular one.

```
ALTER TABLE nation2 REMOVE PARTITIONING;
```

Partition Pruning

Description

Partition pruning is an optimization, limiting the scope of your query according to the criteria you have specified. It is the skipping of unnecessary data partitions in a query. By doing this, you can greatly reduce the amount of data output from the disk and time spent on processing data as well as improve query performance and resource availability.

Example 1

The following example shows how to create the olympic2 table to be partitioned based on the year the Olympic Games were held, and retrieve the countries that participated in the Olympic Games since the 2000 Sydney Olympic Games.

In the **WHERE** clause, partition pruning takes place when equality or range comparison is performed between a partition key and a constant value. In this example, the before_1996 partition that has a smaller year value than 2000 is not scanned.

```
CREATE TABLE olympic2
( opening_date DATE, host_nation VARCHAR(40))
PARTITION BY RANGE ( EXTRACT (YEAR FROM opening_date) )
( PARTITION before_1996 VALUES LESS THAN (1996),
  PARTITION before_MAX VALUES LESS THAN MAXVALUE );

SELECT opening_date, host_nation FROM olympic2 WHERE EXTRACT ( YEAR FROM (opening_date))
>= 2000;
```


Example 2

The following example shows how to retrieve the method of getting the effects of partition pruning by retrieving data with a specific partition when partition pruning does not occur. In the first query, partition pruning does not occur because the value compared is not in the same format as that of the partition expression.

Therefore, you can use the same effect of partition pruning by specifying the appropriate partition as shown in the second query.

```
SELECT host_nation FROM olympic2 WHERE opening_date >= '2000 - 01 - 01';
SELECT host_nation FROM olympic2__p__before_max WHERE opening_date >= '2000 - 01 - 01';
```

Example 3

The following example shows how to specify the search condition to make a partition pruning in the hash partitioned table, called the manager table. For hash partitioning, partition pruning occurs only when equality comparison is performed between a partition key and a constant value in the **WHERE** clause.

```
CREATE TABLE manager (
code INT,
name VARCHAR(50))
PARTITION BY HASH (code) PARTITIONS 4;
SELECT * FROM manager WHERE code = 10053;
```

Caution

The partition expression and the value compared must be in the same format.

Data Types Available for Partitioning Expression**Description**

The following table shows data types of the column that can or cannot be used as a partition key.

Data Types Available	Data Types Unavailable
CHAR	FLOAT
VARCHAR	REAL
NCHAR	DOUBLE
VARNCHAR	BIT
INTEGER	BIT VARYING
SMALLINT	NUMERIC OR DECIMAL
DATE	MONETARY
TIME	SET
TIMESTAMP	LIST OR SEQUENCE
	MULTISET
	OBJECT

The following operator functions can be used in partition expressions to be applied to partition keys.

- Number Operations
+, -, *, /, MOD, STRCAT, FLOOR, CEIL, POWER, ROUND, ABS, TRUNC
- String Operations
POSITION, SUBSTRING, OCTEC_LENGTH, BIT_LENGTH, CHAR_LENGTH, LOWER, UPPER, TRIM, LTRIM, RTRIM, LPAD, RPAD, REPLACE, TRANSLATE
- Date Operations
ADD_MONTH, LAST_DAY, MONTH_BETWEEN, SYS_DATE, SYS_TIME, SYS_TIMESTAMP, TO_DATE, TO_NUMBER, TO_TIME, TO_TIMESTAMP, TO_CHAR
- Others
EXTRACT, CAST

Creating VIEW with Partitioning Table

Description

You can define a virtual table by using each partition of a partitioned table. Retrieving data from the virtual table created is possible, but data insert, delete and update operations are not allowed.

Example

The following example shows how to create the participant2 table partitioned based on the participating year, and create and retrieve a virtual table with the participant2__p__before_2000 partition.

```
CREATE TABLE participant2 (host_year INT, nation CHAR(3), gold INT, silver INT, bronze INT)
PARTITION BY RANGE (host year)
( PARTITION before 2000 VALUES LESS THAN (2000),
  PARTITION before 2008 VALUES LESS THAN (2008) );

INSERT INTO participant2 VALUES (1988, 'NZL', 3, 2, 8);
INSERT INTO participant2 VALUES (1988, 'CAN', 3, 2, 5);
INSERT INTO participant2 VALUES (1996, 'KOR', 7, 15, 5);
INSERT INTO participant2 VALUES (2000, 'RUS', 32, 28, 28);
INSERT INTO participant2 VALUES (2004, 'JPN', 16, 9, 12);

CREATE VIEW v_2000 AS
SELECT * FROM participant2__p__before_2000
WHERE host year = 1988;
```

host year	nation	gold	silver	bronze
1988	'NZL'	3	2	8
1988	'CAN'	3	2	5

Updating Statistics on Partitioning Tables

You can update statistics on the database by using the **cubrid optimizedb** utility or the SQL statement called **UPDATE STATISTICS ON CLASSES**. You can also use the **ANALYZE PARTITION** statement for partitioned tables.

The following example shows the **ANALYZE PARTITION** statement.

```
ALTER TABLE t1 ANALYZE PARTITION p3;
```

Class Inheritance

Overview

Description

To explain the concept of inheritance, a table is represented as a class and a column is represented as an attribute.

Classes in CUBRID database can have class hierarchy. Attributes and methods can be inherited through such hierarchy.

As shown in the previous section, you can create a Manager class by inheriting attributes from an Employee class. The Manager class is called the **subclass** of the Employee class, and the Employee class is called the **super class** of the Manager class. Inheritance can simplify class creation by reusing the existing class hierarchy.

CUBRID allows multiple inheritance, which means that a class can inherit attributes and methods from more than one super class. However, inheritance can cause conflicts when an attribute or method of the super class is added or deleted.

Such conflict occurs in multiple inheritance if there are attributes or methods with the same name in different super classes. For example, if it is likely that a class inherits attributes of the same name and type from more than one super class, you must specify the attributes to be inherited. In such a case, if the inherited super class is deleted, a new attribute of the same name and type must be inherited from another super class. In most cases, the database system resolves such problems automatically. However, if you don't like the way that the system resolves a problem, you can resolve it manually by using the INHERIT clause.

When attributes are inherited from more than one super class, it is possible that their names are to be the same, while their domains are different. For example, two super classes may have the same attribute, whose domain is a class. In this case, a subclass automatically inherits attributes with more specialized (a lower in the class hierarchy) domains. If such conflict occurs between basic data types (e.g. STRING or INTEGER) provided by the system, inheritance fails.

Conflicts during inheritance and their resolutions will be covered in the [Resolving Class Conflicts](#) section.

Caution

The following cautions must be observed during inheritance:

- The class name must be unique in the database. A class can be created as a subclass of one or more super class names in the database optionally. An error occurs if you create a class that inherits another class that does not exist.
- The name of a method/attribute must be unique within a class. The name cannot contain spaces, and cannot be a reserved keyword of CUBRID. Alphabets as well as '_', '#', '%' are allowed in the class name, but the first character cannot be '_'. A class name cannot exceed 255 English letters. Class names are not case-sensitive. A class name will be stored in the system after being converted to lowercase characters.

Note A super class name can begin with the user name so that the owner of the class can be easily identified.

Class Attribute and Method

You can create class attributes to store the aggregate property of all instances in the class. When you define a **CLASS** attribute or method, you must precede the attribute or method name with the keyword **CLASS**. Because a class attribute is associated with the class itself, not with an instances of the class, it has only one value. For example, a class attribute can be used to store the average value determined by a class method or the timestamp when the class was created. A class method is executed on the class object itself. It can be used to calculate the aggregate value for the instances of the class.

When a subclass inherits a super class, each class has a separate storage space for class attributes, so that two classes may have different values of class attribute. Therefore, the subclass does not change even when the attributes of the super class are changed.

The name of a class attribute can be the same as that of an instance attribute of the same class. Likewise, the name of a class method can be the same as that of an instance method of the same class.

Order Rule for Inheritance

The following rules apply to inheritance. The term class is generally used to describe the inheritance relationship between classes and virtual classes in the database.

- For an object without a super class, attributes are defined in the same order as in the **CREATE** statement (an ANSI standard).
- If there is one super class, locally created attributes are placed after the super class attributes. The order of the attributes inherited from the super class follows the one defined during the super class definition. For multiple inheritance, the order of the super class attributes is determined by the order of the super classes specified during the class definition.
- If more than one super class inherits the same class, the attribute that exists in both super classes is inherited to the subclass only once. At this time, if a conflict occurs, the attribute of the first super class is inherited.
- If a name conflict occurs in more than one super class, you can inherit only the ones you want from the super class attributes by using the **INHERIT** clause in order to resolve the conflict.
- If the name of the super class attribute is changed by the alias option of the **INHERIT** clause, its position is maintained.

INHERIT Clause

Description

When a class is created as a subclass, the class inherits all attributes and methods of the super class. A name conflict that occurs during inheritance can be handled by either a system or a user. To resolve the name conflict directly, add the **INHERIT** clause to the **CREATE CLASS** statement.

Syntax

```
CREATE CLASS
.
.
.
INHERIT resolution [ {, resolution }_ ]

resolution :
{ column_name | method_name } OF super_class [ AS alias ]
```

For the *attr_mthd_name* in the **INHERIT** clause, specify the name of the attribute or method of the super class to inherit. With the **ALIAS** clause, you can resolve a name conflict that occurs in multiple inheritance statements by inheriting a new name.

ADD SUPERCLASS Clause

Description

To extend class inheritance, add a super class to a class. A relationship between two classes is created when a super class is added to an existing class. Adding a super class does not mean adding a new class.

Syntax

```
ALTER CLASS
.
.
.
ADD SUPERCLASS [ user_name.]class_name [ {, [ user_name.]class_name }_ ]
[ INHERIT resolution [ {, resolution }_ ] ] [ ; ]
resolution:
{ column_name | method_name } OF superclass_name [ AS alias ]
```

For the first *class_name*, specify the name of the class where a super class is to be added. Attributes and methods of the super class can be inherited by using the syntax above.

Name conflicts can occur when adding a new super class. If a name conflict cannot be resolved by the database system, attributes or methods to inherit from the super class can be specified by using the **INHERIT** clause. You can use aliases to inherit all attributes or methods that cause the conflict. For details on super class name conflicts, see the [Resolving Class Conflict](#) section.

Example

The following example shows how to create the `female_event` class by inheriting the `event` class included in `demodb`.

```
CREATE CLASS female_event UNDER event;
```

DROP SUPERCLASS Clause

Description

Deleting a super class from a class means removing the relationship between two classes. If a super class is deleted from a class, it changes inheritance relationship of the classes as well as of all their subclasses.

Syntax

```
ALTER CLASS
.
.
.
DROP SUPERCLASS class_name [ { , class_name } ]
[ INHERIT resolution [ { , resolution } ] ] [ ; ]

resolution:
{ column_name | method_name } OF superclass_name [ AS alias ]
```

For the first `class_name`, specify the name of the class to be modified. For the second `class_name`, specify the name of the super class to be deleted. If a name conflict occurs after deleting a super class, see the [Resolving Class Conflict](#) section for the resolution.

Example 1

The following example shows how to inherit the `female_event` class from the `event` class.

```
CREATE CLASS female_event UNDER event
```

Example 2

The following example shows how to delete the super class `event` from the `female_event` class. Attributes that the `female_event` class inherited from the `event` class no longer exist.

```
ALTER CLASS female event
DROP SUPERCLASS event;
```

Class Conflict Resolution

Overview

If you modify the schema of the database, conflicts can occur between attributes or methods of inheritance classes. Most conflicts are resolved automatically by CUBRID otherwise, you must resolve the conflict manually. Therefore, you need to examine the possibility of conflicts before modifying the schema.

Two types of conflicts can cause damage to the database schema. One is conflict with a subclass when the subclass schema is modified. The other is conflict with a super class when the super class is modified. The following are operations that may cause conflicts between classes.

- Adding an attribute
- Deleting an attribute
- Adding a super class
- Deleting a super class
- Deleting a class

If a conflict occurs as the result of the above operations, CUBRID applies a basic resolution to the subclass where the conflict occurred. Therefore, the database schema can always maintain consistent state.

Resolution Specifier

Description

Conflicts between the existing classes or attributes, and inheritance conflicts can occur if the database schema is modified. If the system fails to resolve a conflict automatically or if you don't like the way the system resolved the problem, you can suggest how to resolve the conflict by using the **INHERIT** clause of the **ALTER** statement (often referred as resolution specifier).

When the system resolves the conflict automatically, basically, the existing inheritance is maintained (if any). If the previous resolution becomes invalid when the schema is modified, the system will arbitrarily select another one. Therefore, you must avoid excessive reuse of attributes or methods in the schema design stage because the way the system will resolve the conflict cannot always be predictable.

What will be discussed concerning conflicts is applied commonly to both attributes and methods.

Syntax

```
ALTER [ class_type ] class_name alter_clause
[ INHERIT resolution [ { , resolution }_ ] ] [ ]
resolution:
{ column_name | method_name } OF super_class_name [ AS alias ]
```

Superclass Conflict

Adding a super class

The **INHERIT** clause of the **ALTER CLASS** statement is optional, but must be used when a conflict occurs due to class changes. You can specify more than one resolutions after the **INHERIT** clause.

super_class_name specifies the name of the super class that has the new attribute or method to inherit when a conflict occurs. *attr_mthd_name* specifies the name of the attribute or method to inherit. You can use the **alias** clause when you need to change the name of the attribute or method to inherit.

The following example shows how to create the `soccer_stadium` class by inheriting the `event` and `stadium` classes in the olympic database of `demodb`. Because both `event` and `stadium` classes have the `name` and `code` attributes, you must specify the attributes to inherit using the **INHERIT** clause.

```
CREATE CLASS soccer_stadium UNDER event, stadium
INHERIT name OF stadium, code OF stadium
```

When the two super classes (event and stadium) have the name attribute, if the soccer_stadium class needs to inherit both attributes, it can inherit the name unchanged from the stadium class and the name changed from the event class by using the **alias** clause of the **INHERIT**.

The following example shows in which the name attribute of the stadium class is inherited as it is, and that of the event class is inherited as the 'purpose' alias.

```
ALTER CLASS soccer_stadium
INHERIT name OF event AS purpose
```

Deleting a super class

A name conflict may occur again if a super class that explicitly inherited an attribute or method is dropped by using the **INHERIT**. In this case, you must specify the attribute or method to be explicitly inherited when dropping the super class.

The following example shows how to create the seoul_1988_soccer class by inheriting game, participant and stadium classes from demodb, and delete the participant class from the super class. Because nation_code and host_year are explicitly inherited from the participant class, you must resolve their name conflicts before deleting it from the super class. However, host_year does not need to be specified explicitly because it exists only in the game class.

```
CREATE CLASS seoul_1988_soccer UNDER game, participant, stadium
INHERIT nation_code OF participant, host_year OF participant
ALTER CLASS seoul_1988_soccer
DROP super class participant
INHERIT nation_code OF stadium
```

Compatible Domains

When an attribute conflict occurs among two or more super classes, the statement resolving the conflict is not possible only if all attributes have compatible domains.

For example, the class that inherits a super class with the phone attribute of integer type cannot have another super class with the phone attribute of string type. If the types of the phone attributes of the two super classes are both String or Integer, you can add a new super class by resolving the conflict with the **INHERIT** clause.

Compatibility is checked when inheriting an attribute with the same name, but with the different domain. In this case, the attribute that has a lower class in the class inheritance hierarchy as the domain is automatically inherited. If the domains of the attributes to inherit are compatible, the conflict must be resolved in the class where an inheritance relationship is defined.

Subclass Conflict

Any changes in a class will be automatically propagated to all subclasses. If a problem occurs in the subclass due to the changes, CUBRID resolves the corresponding subclass conflict and then displays a message saying that the conflict has been resolved automatically by the system.

Subclass conflicts can occur due to operations such as adding a super class, or creating/deleting a method or an attribute. Any changes in a class will affect all subclasses. Since changes are automatically propagated, harmless changes can even cause side effects in subclasses.

Adding Attributes and Methods

The simplest subclass conflict occurs when an attribute is added. A subclass conflict occurs if an attribute added to a super class has the same name as one already inherited by another super class. In such cases, CUBRID will automatically resolve the problem. That is, the added attribute will not be inherited to all subclasses that have already inherited the attribute with the same name.

The following example shows how to add an attribute to the event class. The super classes of the soccer_stadium class are the event and the stadium classes, and the nation_code attribute already exists in the stadium class. Therefore, a

conflict occurs in the soccer_stadium class if the nation_code attribute is added to the event class. However, CUBRID resolves this conflict automatically.

```
ALTER CLASS event
ADD ATTRIBUTE nation_code CHAR(3)
```

If the event class is dropped from the soccer_stadium super class, the cost attribute of the stadium class will be inherited automatically.

Dropping Attributes and Methods

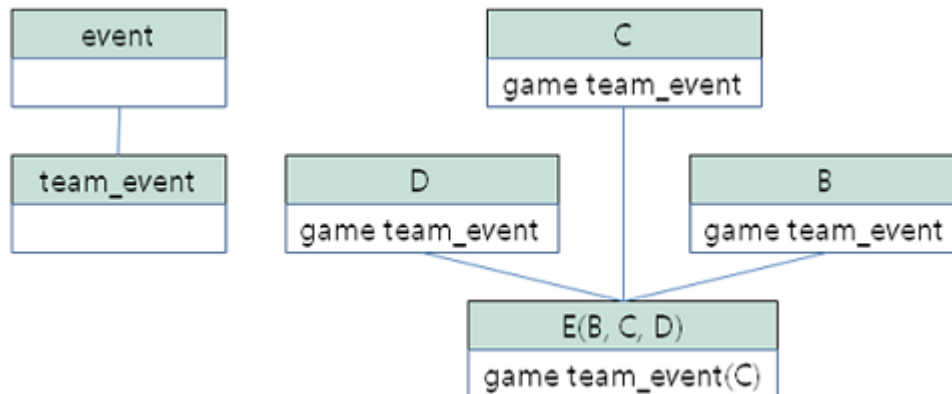
When an attribute is dropped from a class, any resolution specifiers which refer to the attribute by using the **INHERIT** clause are also removed. If a conflict occurs due to the deletion of an attribute, the system will determine a new inheritance hierarchy. If you don't like the inheritance hierarchy determined by the system, you can determine it by using the **INHERIT** clause of the **ALTER** statement. The following example shows such conflict.

Suppose there is a subclass that inherits attributes from three different super classes. If a name conflict occurs in all super classes and the explicitly inherited attribute is dropped, one of the remaining two attributes will be inherited automatically to resolve the problem.

The following example shows subclass conflict. Classes B, C and D are super classes of class E, and have an attribute whose name is team and the domain is team_event. Class E was created with the place attribute inherited from class C as follows:

```
create class E under B, C, D
inherit place of C
```

In this case, the inheritance hierarchy is as follows:



Suppose that you decide to delete class C from the super class. This drop will require changes to the inheritance hierarchy. Because the domains of the remaining classes B and D with the game attribute are at the same level, the system will randomly choose to inherit from one of the two classes. If you don't want the system to make a random selection, you can specify the class to inherit from by using the **INHERIT** clause when you change the class.

```
ALTER CLASS E
INHERIT game OF D
ALTER CLASS C
DROP game
```

Note If the domain of one game attribute in one super class is event and that of another super class is team_event, team_event is more specific than event because team_event is the descendant of event. Therefore, a super class that has the team_event attribute as a domain will be inherited; a user cannot forcefully inherit a super class that has the event attribute as a domain.

Schema Invariant

Invariants of a database schema are a property of the schema that must be preserved consistently (before and after the schema change). There are four types of invariants: invariants of class hierarchy, name, inheritance and consistency.

- **Invariant of class hierarchy** has a single root and defines a class hierarchy as a Directed Acyclic Graph (DAG) where all connected classes have a single direction. That is, all classes except the root have one or more super classes, and cannot become their own super classes. The root of DAG is "object," a system-defined class.
- **Invariant of name** means that all classes in the class hierarchy and all attributes in a class must have unique names. That is, attempts to create classes with the same name or to create attributes or methods with the same name in a single class are not allowed.
Invariant of name is redefined by the 'rename' qualifier. The 'rename' qualifier allows the name of an attribute or method to be changed.
- **Invariant of inheritance** means that a class must inherit all attributes and methods from all super classes. This invariant can be distinguished with three qualifiers: source, conflict and domain. The names of inherited attributes and methods can be modified. For default or shared value attributes, the default or shared value can be modified. Invariant of inheritance means that such changes will be propagated to all classes that inherit these attributes and methods.
- A **source qualifier** means that if class C inherits subclasses of class S, only one of the subclass attributes (methods) inherited from class S can be inherited to class C. That is, if an attribute (method) defined in class S is inherited by other classes, it is in effect a single attribute (method), even though it exists in many subclasses. Therefore, if a class multiply inherits from classes that have attributes (methods) of the same source, only one appearance of the attribute (method) is inherited.
- A **conflict qualifier** means that if class C inherits from two or more classes that have attributes (methods) with the same name but of different sources, it can inherit more than one class. To inherit attributes (methods) with the same name, you must change their names so as not to violate the invariant of name.
- A **domain qualifier** means that a domain of an inherited attribute can be converted to the domain's subclass.
- **Invariant of consistency** means that the database schema must always follow the invariants of a schema and all rules ([Rules for Schema Changes](#)) except when it is being changed.

Rule for Schema Changes

The Invariants of a Schema section has described the characteristics of schema that must be preserved all the time. There are some methods for changing schemas, and all these methods must be able to preserve the invariants of a schema. For example, suppose that in a class which has a single super class, the relationship with the super class is to be removed. If the relationship with the super class is removed, the class becomes a direct subclass of the object class, or the removal attempt will be rejected if the user specified that the class should have at least one super class. To have some rules for selecting one of the methods for changing schemas, even though such selection seems arbitrary, will be definitely useful to users and database designers.

The following three types of rules apply: conflict-resolution rules, domain-change rule and class-hierarchy rule.

Seven conflict-resolution rules reinforce the invariant of inheritance. Most schema change rules are needed because of name conflicts. A domain-change rule reinforces a domain resolution of the invariant of inheritance. A class-hierarchy rule reinforces the invariant of class hierarchy.

Conflict-Resolution Rules

- **Rule 1** : If an attribute (method) name of class C and an attribute name of the super class S conflict with each other (that is, their names are same), the attribute of class C is used. The attribute of S is not inherited.

If a class has one or more super classes, three aspects of the attribute (method) of each super class must be considered to determine whether the attributes are semantically equal and which attribute to inherit. The three aspects of the attribute (method) are the name, domain and source. The following table shows eight combinations of these three aspects that can happen with two super classes. In Case 1 (two different super classes have attributes with the same name, domain and source), only one of the two subclasses should be inherited because two attributes are identical. In Case 8 (two different super classes have attributes with different names, domains and sources), both classes should be inherited because two attributes are totally different ones.

Case	Name	Domain	Source
1	Same	Same	Same
2	Same	Same	Different
3	Same	Different	Same

4	Same	Different	Different
5	Different	Same	Same
6	Different	Same	Different
7	Different	Different	Same
8	Different	Different	Different

Five cases (1, 5, 6, 7, 8) out of eight have clear meaning. Invariant of inheritance is a guideline for resolving conflicts in such cases. In other cases (2, 3, 4), it is very difficult to resolve conflicts automatically. Rules 2 and 3 can be resolutions for these conflicts.

- **Rule 2** : When two or more super classes have attributes (methods) with different sources but the same name and domain, one or more attributes (methods) can be inherited if the conflict-resolution statement is used. If the conflict-resolution statement is not used, the system will select and inherit one of the two attributes.

This rule is a guideline for resolving conflicts of Case 2 in the table above.

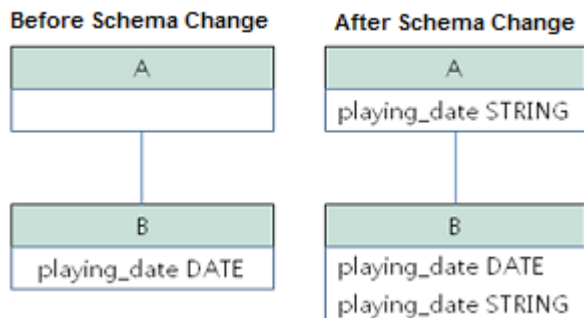
- **Rule 3** : If two or more super classes have attributes with different sources and domains but the same name, attributes (methods) with more detailed (lower in the inheritance hierarchy) domains are inherited. If there is no inheritance relationship between domains, schema change is not allowed.

This rule is a guideline for resolving conflicts of Case 3 and 4. If Case 3 and 4 conflict with each other, Case 3 has the priority.

- **Rule 4** : The user can make any changes except the ones in Case 3 and 4. In addition, the resolution of subclass conflicts cannot cause changes in the super class.

The philosophy of Rule 4 is that "an inheritance is a privilege that subclass has obtained from a super class, so changes in a subclass cannot affect the super class." Rule 4 means that the name of the attribute (method) included in the super class cannot be changed to resolve conflicts between class C and super classes. Rule 4 has an exception in cases where the schema change causes conflicts in Case 3 and 4.

- For example, suppose that class A is the super class of class B, and class B has the playing_date attribute of DATE type. If an attribute of STRING type named playing_date is added to class A, it conflicts with the playing_date attribute in class B. This is what happens in Case 4. The precise way to resolve such conflict is for the user to specify that class B must inherit the playing_date attribute of class A. If a method refers to the attribute, the user of class B needs to modify the method properly so that the appropriate playing_date attribute will be referenced. Schema change of class A is not allowed because the schema falls into an inconsistent state if the user of class B does not describe an explicit statement to resolve the conflict occurring from the schema change.



- **Rule 5** : If a conflict occurs due to a schema change of the super class, the original resolution is maintained as long as the change does not violate the rules. However, if the original resolution becomes invalid due to the schema change, the system will apply another resolution.

Rule 5 is for cases where a conflict is caused to a conflict-free class or where the original resolution becomes invalid.

This is the case where the name or domain of an attribute (method) is modified or a super class is deleted when the attribute (method) is added to the super class or the one inherited from the super class is deleted. The philosophy of Rule 5 coincides with that of Rule 4. That is, the user can change the class freely without considering what effects the subclass that inherits from the given class will have on the inherited attribute (method).

When you change the schema of class C, if you decide to inherit an attribute of the class due to an earlier conflict with another class, this may cause attribute (method) loss of class C. Instead, you must inherit one of the attributes (methods) that caused conflicts earlier.

The schema change of the super class can cause a conflict between the attribute (method) of the super class and the (locally declared or inherited) attribute (method) of class C. In this case, the system resolves the conflict automatically by applying Rule 2 or 3 and may inform the user.

Rule 5 cannot be applied to cases where a new conflict occurs due to the addition or deletion of the relationship with the super class. The addition/deletion of a super class must be limited to within the class. That is, the user must provide an explicit resolution.

- **Rule 6** : Changes of attributes or methods are propagated only to subclasses without conflicts.

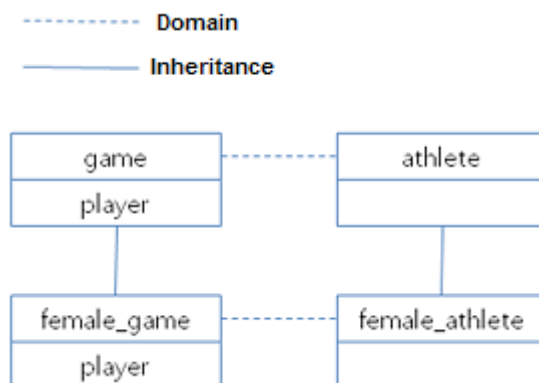
This rule limits the application of Rule 5 and the invariant of inheritance. Conflicts can be detected and resolved by applying Rule 2 and 3.

- **Rule 7** : Class C can be dropped even when an attribute of class R uses class C as a domain. In this case, the domain of the attribute that uses class C as a domain can be changed to object.

Domain-Change Rules

- **Rule 8** : If the domain of an attribute of class C is changed from D to a super class of D, the new domain is less generic than the corresponding domain in the super class from which class C inherited the attribute. The following example explains the principle of this rule.

Suppose that in the database there are the game class with the player attribute and the female_game class which inherits game. The domain of the player attribute of the game class is the athlete class, but the domain of the player attribute of the female_game class is changed to female_athlete which is a subclass of athlete. The following diagram shows such relationship. The domain of the player attribute of the female_game class can be changed back to athlete, which is the super class of female_athlete.



Class-Hierarchy Rules

- **Rule 9** : A class without a super class becomes a direct subclass of object. The class-hierarchy rule defines characteristics of classes without super classes. If you create a class without a super class, object becomes the super class. If you delete the super class S, which is a unique super class of class C, class C becomes a direct subclass of object.

CUBRID System Catalog

Overview

You can easily get various schema information from the SQL statement by using the system catalog virtual class. For example, you can get the following schema information by using the catalog virtual class.

```

-- Classes that refer to the 'b user' class
SELECT class_name
FROM db_attribute
WHERE domain class name = 'db user';

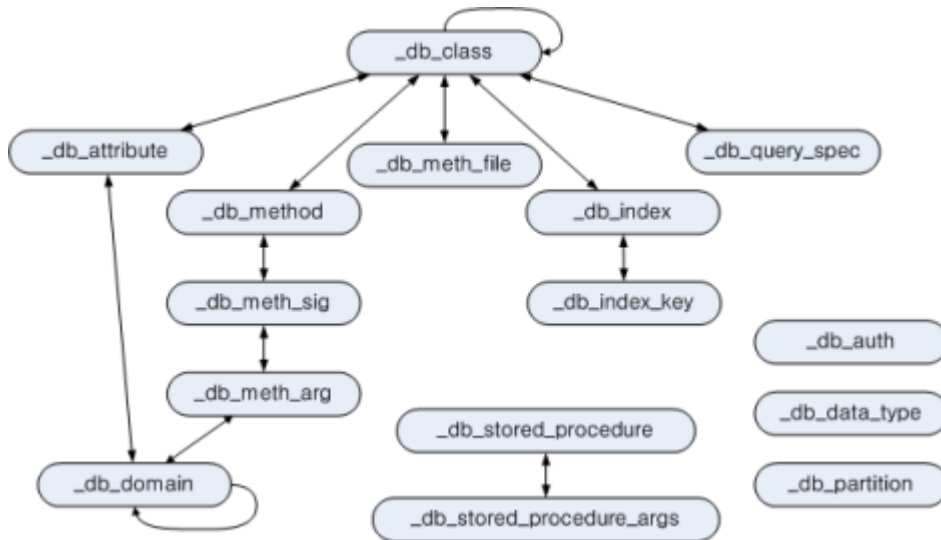
-- The number of classes that the current user can access
SELECT COUNT(*)
FROM db_class;

-- Attribute of the 'db user' class
SELECT attr name, data type
FROM db attribute
WHERE class_name = 'db user';
    
```

System Catalog Classes

System Catalog Classes

To define a catalog virtual class, define a catalog class first. The figure below shows catalog classes to be added and their relationships. The arrows represent the reference relationship between classes, and the classes that start with an underline () are catalog classes.



Added catalog classes represent information about all classes, attributes and methods in the database. Catalog classes are made up of class composition hierarchy and designed to have OIDs of catalog class instances for cross reference.

_db_class

Represents class information. An index for class_name is created.

Attribute Name	Data Type	Description
class_of	object	A class object. Represents a meta information object for the class stored in the system.
class_name	VARCHAR(255)	Class name

class_type	INTEGER	0 for a class, and 1 for a virtual class
is_system_class	INTEGER	0 for a user-defined class, and 1 for a system class
owner	db_user	Class owner
inst_attr_count	INTEGER	The number of instance attributes
class_attr_count	INTEGER	The number of class attributes
shard_attr_count	INTEGER	The number of shared attributes
inst_meth_count	INTEGER	The number of instance methods
class_meth_count	INTEGER	The number of class methods
sub_classes	SEQUENCE OF _db_class	Class one level down
super_classes	SEQUENCE OF _db_class	Class one level up
inst_attrs	SEQUENCE OF _db_attribute	Instance attribute
class_attrs	SEQUENCE OF _db_attribute	Class attribute
shared_attrs	SEQUENCE OF _db_attribute	Shared attribute
inst_meths	SEQUENCE OF _db_method	Instance method
class_meths	SEQUENCE OF _db_method	Class method
meth_files	SEQUENCE OF _db_methfile	File path in which the function for the method is located
query_specs	SEQUENCE OF _db_queryspec	SQL definition statement for a virtual class
indexes	SEQUENCE OF _db_index	Index created in the class

Example

The following example shows how to retrieve all subclasses under the class owned by user 'PUBLIC' (for the child class female_event in the result, see the example in [Adding a super class](#)).

```
SELECT class name, SEQUENCE(SELECT class name FROM db class s WHERE s IN c.sub_classes)
FROM _db_class c
WHERE c.owner.name = 'PUBLIC' AND c.sub_classes IS NOT NULL;
class name          sequence((select class name from db class s where s in
c.sub_classes))
=====
'event'             {'female_event'}
```

Note All examples of system catalog classes have been written in the csql utility. In this example, **--no-auto-commit** (inactive mode of auto-commit) and **-u** (specifying user DBA) options are used.

```
% csql --no-auto-commit -u dba demodb
```

_db_attribute

Represents attribute information. Indexes for class_of and attr_name are created.

Attribute Name	Data Type	Description
class_of	_db_class	Class to which the attribute belongs

attr_name	VARCHAR(255)	Attribute name
attr_type	INTEGER	Type defined for the attribute. 0 for an instance attribute, 1 for a class attribute, and 2 for a shared attribute.
from_class_of	_db_class	If the attribute is inherited, the super class in which the attribute is defined is specified. Otherwise, NULL is specified.
from_attr_name	VARCHAR(255)	Inherited attribute. If an attribute name has changed to resolve a name conflict, the original name define in the super class is specified. Otherwise, NULL is specified.
def_order	INTEGER	Order of attributes in the class. Begins with 0. If the attribute is inherited, the order is the one defined in the super class. For example, if class y inherits attribute a from class x and a was first defined in x, def_order becomes 0.
data_type	INTEGER	Data type of the attribute. One of the values specified in the "Data Types Supported by CUBRID" table below.
default_value	VARCHAR(255)	Default value. Stores as a character string regardless of data types. If there is no default value, NULL . If the default value is NULL , NULL is used. If the data type is an object, 'volume id page id slot id' is used. If the data type is a set, '{element 1, element 2, ...' is used.
domains	SEQUENCE OF _db_domain	Domain information of the data type
is_nullable	INTEGER	0 if a not null constraint is configured, and 1 otherwise.

Data Types Supported by CUBRID

Value	Meaning	Value	Meaning
1	INTEGER	13	MONETARY
2	FLOAT	18	SHORT
3	DOUBLE	20	OID
4	STRING	22	NUMERIC
5	OBJECT	23	BIT
6	SET	24	VARBIT
7	MULTISET	25	CHAR
8	SEQUENCE	26	NCHAR
9	ELO	27	VARNCHAR
10	TIME	31	BIGINT
11	TIMESTAMP	32	DATETIME
12	DATE	33	BLOB
		34	CLOB

Character Sets Supported by CUBRID

Value	Meaning
0	US English - ASCII encoding
3	Latin 1 - ISO 8859 encoding
4	KSC 5601 1990 - EUC encoding

Example

The following example shows how to retrieve user classes (from_class_of.is_system_class = 0) among the ones owned by user 'PUBLIC':

```
SELECT class_of.class_name, attr_name
FROM _db_attribute
WHERE class_of.owner.name = 'PUBLIC' AND FROM class_of.is_system_class = 0
ORDER BY 1, def_order;
class_of.class_name  attr_name
=====
'female_event'      'code'
'female_event'      'sports'
'female_event'      'name'
'female_event'      'gender'
'female_event'      'players'
```

_db_domain

Represents domain information. An index for object_of is created.

Attribute Name	Data Type	Description
object_of	object	Attribute that refers to the domain, which can be a method parameter or domain
data_type	INTEGER	Data type of the domain (a value in the "Value" column of the "Data Types Supported by CUBRID" table in _db_attribute)
prec	INTEGER	Precision of the data type. 0 is used if the precision is not specified.
scale	INTEGER	Scale of the data type. 0 is used if the scale is not specified.
class_of	_db_class	Domain class if the data type is an object, NULL otherwise.
code_set	INTEGER	Character set (value of table "character sets supported by CUBRID" in _db_attribute) if it is character data type. 0 otherwise.
set_domains	SEQUENCE OF _db_domain	Domain information about the data type of collection element if it is collection data type. NULL otherwise.

_db_method

Represents method information. Indexes for class_of and meth_name are created.

Attribute Name	Data Type	Description
class_of	_db_class	Class to which the method belongs
meth_type	INTEGER	Type of the method defined in the class. 0 for an instance method, and 1 for a class method.
from_class_of	_db_class	If the method is inherited, the super class in which it is defined is used otherwise NULL
from_meth_name	VARCHAR(255)	If the method is inherited and its name is changed to resolve a name conflict, the original name defined in the super class is used otherwise NULL
meth_name	VARCHAR(255)	Method name
signatures	SEQUENCE OF _db_meth_sig	C function executed when the method is called

Example

The following example shows how to retrieve class methods of the class with a class method (c.class_meth_count > 0), among classes owned by user 'DBA':

```

SELECT class name, SEQUENCE(SELECT meth name
                             FROM db method m
                             WHERE m in c.class_meths)
FROM _db_class c
WHERE c.owner.name = 'DBA' AND c.class meth count > 0
ORDER BY 1;
class name          sequence((select meth name from db method m where m in
c.class meths))
=====
'db_serial'         {'change_serial_owner'}
'db authorizations' {'add_user', 'drop_user', 'find_user', 'print authorizations',
'info', 'change owner', 'change trigg
r owner', 'get owner'}
'db authorization' {'check authorization'}
'db_user'           {'add_user', 'drop_user', 'find_user', 'login'}
'db_root'           {'add_user', 'drop_user', 'find_user', 'print authorizations',
'info', 'change owner', 'change trigg
r owner', 'get_owner', 'change_sp_owner'}

```

_db_meth_sig

Represents configuration information of a C function on the method. An index for meth_of is created.

Attribute Name	Data Type	Description
meth_of	_db_method	Method for the function information
arg_count	INTEGER	The number of input arguments of the function
func_name	VARCHAR(255)	Function name
return_value	SEQUENCE OF _db_meth_arg	Return value of the function
arguments	SEQUENCE OF _db_meth_arg	Input arguments of the function

_db_meth_arg

Represents method argument information. An index for meth_sig_of is created.

Attribute Name	Data Type	Description
meth_sig_of	_db_meth_sig	Information of the function to which the argument belongs
data_type	INTEGER	Data type of the argument (a value in the "Value" column of the "Data Types Supported by CUBRID" in db attribute)
index_of	INTEGER	Order of the argument listed in the function definition. Begins with 0 if it is a return value, and 1 if it is an input argument.
domains	SEQUENCE OF _db_domain	Domain of the argument

_db_meth_file

Represents information of a file in which a function is defined. An index for class_of is created.

Attribute Name	Data Type	Description
class_of	_db_class	Class to which the method file information belongs
from_class_of	_db_class	If the file information is inherited, the super class in which it is defined is used otherwise, NULL
path_name	VARCHAR(255)	File path in which the method is located

db_query_spec

Represents the SQL statement of a virtual class. An index for class_of is created.

Attribute Name	Data Type	Description
class_of	_db_class	Class information of the virtual class
spec	VARCHAR(4096)	SQL definition statement of the virtual class

db_index

Represents index information. An index for class_of is created.

Attribute Name	Data Type	Description
class_of	_db_class	Class to which to index belongs
index_name	varchar(255)	Index name
is_unique	INTEGER	1 if the index is unique, and 0 otherwise.
key_count	INTEGER	The number of attributes that comprise the key
key_attrs	SEQUENCE OF _db_index_key	Attributes that comprise the key
is_reverse	INTEGER	1 for a reverse index, and 0 otherwise.
is_primary_key	INTEGER	1 for a primary key, and 0 otherwise.
is_foreign_key	INTEGER	1 for a foreign key, and 0 otherwise.

Example

The following example shows how to retrieve names of indexes that belong to the class.

```

SELECT class_of.class name, index name
FROM _db_index
ORDER BY 1;
  class_of.class name  index name
=====
' db attribute'      'i db attribute class of attr name'
' db auth'          'i db auth grantee'
' _db_class'        'i _db_class_class_name'
' _db_domain'       'i _db_domain_object_of'
' db index'         'i db index class of'
' db index key'     'i db index key index of'
' db meth arg'      'i db meth arg meth sig of'
' _db_meth_file'    'i _db_meth_file_class_of'
' _db_meth_sig'     'i _db_meth_sig_meth_of'
' db method'        'i db method class of meth name'
' db partition'     'i db partition class of pname'
' db query spec'    'i db query spec class of'
' db stored procedure' 'u db stored procedure sp name'
' _db_stored_procedure_args' 'i _db_stored_procedure_args_sp_name'
'athlete'          'pk_athlete_code'
'db serial'        'pk_db_serial_name'
'db user'          'i db user name'
'event'           'pk_event_code'
'game'            'pk_game_host_year_event_code_athlete_code'
'game'            'fk_game_event_code'
'game'            'fk_game_athlete_code'
'history'         'pk_history_event_code_athlete'
'nation'          'pk_nation_code'
'olympic'         'pk_olympic_host_year'
'participant'     'pk_participant_host_year_nation_code'
'participant'     'fk_participant_host_year'
'participant'     'fk_participant_nation_code'
'record'          'pk_record_host_year_event_code_athlete_code_medal'
'stadium'         'pk_stadium_code'

```

db_index_key

Represents key information on an index. An index for index_of is created.

Attribute Name	Data Type	Description
index_of	_db_index	Index to which the key attribute belongs
key_attr_name	VARCHAR(255)	Name of the attribute that comprises the key
key_order	INTEGER	Order of the attribute in the key. Begins with 0.
asc_desc	INTEGER	1 if the order of attribute values is descending, and 0 otherwise.
key_prefix_length	INTEGER	Length of prefix to be used as a key

Example

The following example shows how to retrieve the names of index that belongs to the class.

```
SELECT class_of.class_name, SEQUENCE(SELECT key_attr_name
                                     FROM _db_index_key k
                                     WHERE k in i.key_attrs)
FROM db index i
WHERE key count >= 2;
class_of.class_name  sequence((select key_attr_name from db index key k where k in
i.key_attrs))
=====
' db partition'      {'class of', 'pname'}
' db method'         {'class of', 'meth name'}
' db attribute'      {'class of', 'attr name'}
'participant'       {'host_year', 'nation code'}
'game'               {'host_year', 'event_code', 'athlete_code'}
'record'             {'host_year', 'event_code', 'athlete_code', 'medal'}
'history'            {'event_code', 'athlete'}
```

db_auth

Represents user authorization information of the class. An index for the grantee is created.

Attribute Name	Data Type	Description
grantor	db_user	Authorization grantor
grantee	db_user	Authorization grantee
class_of	_db_class	Class object to which authorization is to be granted
auth_type	VARCHAR(7)	Type name of the authorization granted
is_grantable	INTEGER	1 if authorization for the class can be granted to other users, and 0 otherwise.

Authorization types supported by CUBRID are as follows:

- **SELECT**
- **INSERT**
- **UPDATE**
- **DELETE**
- **ALTER**
- **INDEX**
- **EXECUTE**

Example

The following example shows how to retrieve authorization information defined in the class 'db_trig'.

```
SELECT grantor.name, grantee.name, auth type
FROM db auth
WHERE class_of.class_name = 'db_trig';
```

grantor.name	grantee.name	auth type
'DBA'	'PUBLIC'	'SELECT'

_db_data_type

Represents the data type supported by CUBRID (see the "Data Types Supported by CUBRID" table in [_db_attribute](#)).

Attribute Name	Data Type	Description
type_id	INTEGER	Data type identifier. Corresponds to the "Value" column in the "Data Types Supported by CUBRID" table.
type_name	VARCHAR(9)	Data type name. Corresponds to the "Meaning" column in the "Data Types Supported by CUBRID" table.

Example

The following example shows how to retrieve attributes and type names of the 'event' class.

```
SELECT a.attr name, t.type name
FROM _db_attribute a join _db_data_type t ON a.data_type = t.type_id
WHERE class_of.class_name = 'event'
ORDER BY a.def order;
attr name          type name
=====
'code'             'INTEGER'
'sports'           'STRING'
'name'             'STRING'
'gender'           'CHAR'
'players'          'INTEGER'
```

_db_partition

Represents partition information. Indexes for class_of and pname are created.

Attribute Name	Data Type	Description
class_of	_db_class	OID of the parent class
pname	VARCHAR(255)	Parent - NULL
pptype	INTEGER	0 - HASH 1 - RANGE 2 - LIST
pexpr	VARCHAR(255)	Parent only
pvalues	SEQUENCE OF	Parent - Column name, Hash size RANGE - MIN/MAX value : - Infinite MIN/MAX is stored as NULL LIST - value list

_db_stored_procedure

Represents Java stored procedure information. An index for sp_name is created.

Attribute Name	Data Type	Description
sp_name	VARCHAR(255)	Stored procedure name
sp_type	INTEGER	Stored procedure type (function or procedure)
return_type	INTEGER	Return value type
arg_count	INTEGER	The number of arguments
args	SEQUENCE OF _db_stored_procedure_args	Argument list

lang	INTEGER	Implementation language (currently, Java)
target	VARCHAR(4096)	Name of the Java method to be executed
owner	db_user	Owner

_db_stored_procedure_args

Represents Java stored procedure argument information. An index for sp_name is created.

Attribute Name	Data Type	Description
sp_name	VARCHAR(255)	Stored procedure name
index_of	INTEGER	Order of the arguments
arg_name	VARCHAR(255)	Argument name
data_type	INTEGER	Data type of the argument
mode	INTEGER	Mode (IN, OUT, INOUT)

db_user

Attribute Name	Data Type	Description
name	VARCHAR(1073741823)	User name
id	INTEGER	User identifier
password	db_password	User password. Not displayed to the user.
direct_groups	SET OF db_user	Groups to which the user belongs directly
groups	SET OF db_user	Groups to which the user belongs directly or indirectly
authorization	db_authorization	Information of the authorization owned by the user
triggers	SEQUENCE OF object	Triggers that occur due to user actions

Function Names

- **set_password()**
- **set_password_encoded()**
- **add_member()**
- **drop_member()**
- **print_authorizations()**
- **add_user()**
- **drop_user()**
- **find_user()**
- **login()**

db_authorization

Attribute Name	Data Type	Description
owner	db_user	User information
grants	SEQUENCE OF object	Sequence of {object for which the user has authorization, authorization grantor of the object, authorization type}

Method Name

- **check_authorization**(varchar(255), integer)

db_trigger

Attribute Name	Data Type	Description
owner	db_user	Trigger owner
name	VARCHAR(1073741823)	Trigger name
status	INTEGER	1 for INACTIVE, and 2 for ACTIVE. The default value is 2.
priority	DOUBLE	Execution priority between triggers. The default value is 0.
event	INTEGER	0 is set for UPDATE, 1 for UPDATE STATEMENT, 2 for DELETE, 3 for DELETE STATEMENT, 4 for INSERT, 5 for INSERT STATEMENT, 8 for COMMIT, and 9 for ROLLBACK.
target_class	object	Class object for the trigger target class
target_attribute	VARCHAR(1073741823)	Trigger target attribute name. If the target attribute is not specified, NULL is used.
target_class_attribute	INTEGER	If the target attribute is an instance attribute, 0 is used. If it is a class attribute, 1 is used. The default value is 0.
condition_type	INTEGER	If a condition exist, 1; otherwise NULL .
condition	VARCHAR(1073741823)	Action condition specified in the IF statement
condition_time	INTEGER	1 for BEFORE, 2 for AFTER, and 3 for DEFERRED if a condition exists; NULL , otherwise.
action_type	INTEGER	1 for one of INSERT, UPDATE, DELETE, CALL and EVALUATE, 2 for REJECT, 3 for INVALIDATE_TRANSACTION, and 4 for PRINT.
action_definition	VARCHAR(1073741823)	Execution statement to be triggered
action_time	INTEGER	1 for BEFORE, 2 for AFTER, and 3 for DEFERRED.

db_ha_apply_info

A table that stores the progress status every time the **applylogdb** utility applies replication logs. This table is updated at every point the **applylogdb** utility commits, and the accumulative count of operations are stored in the *_counter column. The meaning of each column is as follows:

Column Name	Column Type	Meaning
db_name	VARCHAR(255)	Name of the database stored in the log
db_creation_time	DATETIME	Creation time of the source database for the log to be applied
copied_log_path	VARCHAR(4096)	Path to the log file to be applied
page_id	INTEGER	Page of the replication log committed in the slave database
offset	INTEGER	Offset of the replication log committed in the slave database
log_record_time	DATETIME	Timestamp included in replication log committed in the slave database, i.e. the creation time of the log
last_access_time	DATETIME	Time when applylogdb was committed in the slave database
status	INTEGER	Progress status (0: IDLE, 1: BUSY)
insert_counter	BIGINT	Number of times that applylogdb was inserted
update_counter	BIGINT	Number of times that applylogdb was updated

delete_counter	BIGINT	Number of times that applylogdb was deleted
schema_counter	BIGINT	Number of times that applylogdb changed the schema
commit_counter	BIGINT	Number of times that applylogdb was committed
fail_counter	BIGINT	Number of times that applylogdb failed to be inserted/updated/deleted/committed and to change the schema
required_page_id	INTEGER	Minimum pageid that applylogdb can read
start_time	DATETIME	Time when the applylogdb process accessed the slave database

System Catalog Virtual Class

System Catalog Virtual Class

General users can only see information of classes for which they have authorization through system catalog virtual classes.

This section explains which information each system catalog virtual class represents, and virtual class definition statements.

DB_CLASS

Represents information of classes for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
class_name	VARCHAR(255)	Class name
owner_name	VARCHAR(255)	Name of class owner
class_type	VARCHAR(6)	'CLASS' for a class, and 'VCLASS' for a virtual class
is_system_class	VARCHAR(3)	'YES' for a system class, and 'NO' otherwise.
partitioned	VARCHAR(3)	'YES' for a partitioned group class, and 'NO' otherwise.
is_reuse_oid_class	VARCHAR(3)	'YES' for a REUSE_OID class, and 'NO' otherwise.

Definition

```
CREATE VCLASS db_class (class_name, owner_name, class_type, is_system_class, partitioned,
is_reuse_oid_class)
AS
SELECT c.class name, CAST(c.owner.name AS VARCHAR(255)),
CASE c.class type WHEN 0 THEN 'CLASS' WHEN 1 THEN 'VCLASS' ELSE 'UNKNOW' END,
CASE WHEN MOD(c.is_system_class, 2) = 1 THEN 'YES' ELSE 'NO' END,
CASE WHEN c.sub classes IS NULL THEN 'NO' ELSE NVL((SELECT 'YES' FROM db partition p
WHERE p.class of = c and p.pname IS NULL), 'NO') END,
CASE WHEN MOD(c.is system class / 8, 2) = 1 THEN 'YES' ELSE 'NO' END
FROM db class c
WHERE CURRENT_USER = 'DBA' OR
(c.owner.name) SUBSETEQ (
SELECT SET{CURRENT USER} + COALESCE(SUM(SET{t.g.name}), SET{})
FROM db user u, TABLE(groups) AS t(g)
WHERE u.name = CURRENT USER) OR
(c) SUBSETEQ (
SELECT SUM(SET{au.class_of})
FROM _db_auth au
WHERE {au.grantee.name} SUBSETEQ(
SELECT SET{CURRENT USER} + COALESCE(SUM(SET{t.g.name}), SET{})
FROM db user u, TABLE(groups) AS t(g)
WHERE u.name = CURRENT_USER) AND au.auth_type = 'SELECT');
```

Example 1

The following example shows how to retrieve classes owned by the current user.

```
SELECT class_name
FROM db_class
WHERE owner name = CURRENT USER;
  class name
=====
'stadium'
'code'
'nation'
'event'
'athlete'
'participant'
'olympic'
'game'
'record'
'history'
'female_event'
```

Note All examples of system catalog classes have been written in the `csql` utility. In this example, the user option is omitted (if omitted, the default user is **PUBLIC**). If not otherwise specified, **--no-auto-commit** (No auto-commit mode) and **-u** (Specify the user **dba**) options are used.

```
% csql --no-auto-commit -u dba demo
```

Example 2

The following example shows how to retrieve virtual classes that can be accessed by the current user.

```
SELECT class_name
FROM db class
WHERE class type = 'VCLASS';
  class name
=====
'db_stored_procedure_args'
'db_stored_procedure'
'db_partition'
'db_trig'
'db_auth'
'db_index_key'
'db_index'
'db_meth_file'
'db_meth_arg_setdomain_elm'
'db_meth_arg'
'db_method'
'db_attr_setdomain_elm'
'db_attribute'
'db_vclass'
'db_direct_super_class'
'db_class'
```

The following example shows how to retrieve system classes that can be accessed by the current user user (**PUBLIC** user).

```
SELECT class_name
FROM db_class
WHERE is_system_class = 'YES' AND class_type = 'CLASS'
ORDER BY 1;
  class name
=====
'db_authorization'
'db_authorizations'
'db_root'
'db_serial'
'db_user'
```

DB_DIRECT_SUPER_CLASS

Represents the names of super classes (if any) of the class for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
class_name	VARCHAR(255)	Class name
super_class_name	VARCHAR(255)	super class name

Definition

```
CREATE VCLASS db direct super class (class name, super class name)
AS
SELECT c.class_name, s.class_name
FROM _db_class c, TABLE(c.super_classes) AS t(s)
WHERE (CURRENT USER = 'DBA' OR
       {c.owner.name} subseteq (
           SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}),
set{}))
           from db_user u, table(groups) as t(g)
           where u.name = CURRENT USER ) OR
       {c} subseteq (
SELECT sum(set{au.class of})
           FROM db_auth au
           WHERE {au.grantee.name} subseteq (
           SELECT set{CURRENT_USER} +
coalesce(sum(set{t.g.name}), set{}))
           from db user u, table(groups) as t(g)
           where u.name = CURRENT USER ) AND
           au.auth type =
'SELECT'))
```

Example

The following example shows how to retrieve super classes of the 'female_event' class (see [ADD SUPERCLASS Clause](#)).

```
SELECT super class name
FROM db direct super class
WHERE class_name = 'female_event';
  super_class_name
=====
'event'
```

The following example shows how to retrieve super classes of the class owned by the current user (**PUBLIC** user).

```
SELECT c.class name, s.super class name
FROM db class c, db direct super class s
WHERE c.class name = s.class name AND c.owner name = user
ORDER BY 1;
  class_name          super_class_name
=====
'female_event'      'event'
```

DB_VCLASS

Represents SQL definition statements of virtual classes for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
vclass_name	VARCHAR(255)	Virtual class name
vclass_def	VARCHAR(4096)	SQL definition statement of the virtual class

Definition

```
CREATE VCLASS db vclass (vclass name, vclass def)
AS
SELECT q.class_of.class_name, q.spec
```



```

FROM db query spec q
WHERE CURRENT_USER = 'DBA' OR
      {q.class_of.owner.name} subseteq (
        SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
        from db user u, table(groups) as t(g)
        where u.name = CURRENT_USER ) OR
      {q.class_of} subseteq (
SELECT sum(set{au.class_of})
      FROM _db_auth au
      WHERE {au.grantee.name} subseteq (
        SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
        from db user u, table(groups) as t(g)
        where u.name = CURRENT_USER ) AND
        au.auth_type = 'SELECT');

```

Example

The following example shows how to retrieve SQL definition statements of the 'db_class' virtual class.

```

SELECT vclass def
FROM db vclass
WHERE vclass name = 'db class';
'SELECT c.class_name, CAST(c.owner.name AS VARCHAR(255)), CASE c.class_type WHEN 0 THEN
'CLASS' WHEN 1 THEN 'VCLASS' WHEN 2 THEN 'PROXY' ELSE 'UNKNOW' END, CASE WHEN
MOD(c.is system class, 2) = 1 THEN 'YES' ELSE 'NO' END, CASE WHEN c.sub_classes IS NULL
THEN 'NO' ELSE NVL((SELECT 'YES' FROM db partition p WHERE p.class_of = c and p.pname IS
NULL), 'NO') END FROM db class c WHERE CURRENT_USER = 'DBA' OR {c.owner.name} SUBSETEQ
( SELECT SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name}), SET{}) FROM db user u,
TABLE(groups) AS t(g) WHERE u.name = CURRENT_USER) OR {c} SUBSETEQ ( SELECT
SUM(SET{au.class_of}) FROM _db_auth au WHERE {au.grantee.name} SUBSETEQ ( SELECT
SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name}), SET{}) FROM db user u, TABLE(groups) AS
t(g) WHERE u.name = CURRENT_USER) AND au.auth_type = 'SELECT')'

```

DB_ATTRIBUTE

Represents the attribute information of a class for which the current user has access authorization in the database.

Attribute Name	Data Type	Description
attr_name	VARCHAR(255)	Attribute name
class_name	VARCHAR(255)	Name of the class to which the attribute belongs
attr_type	VARCHAR(8)	'INSTANCE' for an instance attribute, 'CLASS' for a class attribute, and 'SHARED' for a shared attribute.
def_order	INTEGER	Order of attributes in the class. Begins with 0. If the attribute is inherited, the order is the one defined in the super class.
from_class_name	VARCHAR(255)	If the attribute is inherited, the super class in which it is defined is used. Otherwise, NULL
from_attr_name	VARCHAR(255)	If the attribute is inherited and its name is changed to resolve a name conflict, the original name defined in the super class is used. Otherwise, NULL
data_type	VARCHAR(9)	Data type of the attribute (one in the "Meaning" column of the "Data Types Supported by CUBRID" table in db_attribute)
prec	INTEGER	Precision of the data type. 0 is used if the precision is not specified.
scale	INTEGER	Scale of the data type. 0 is used if the scale is not specified.
code_set	INTEGER	Character set (value of table "character sets supported by CUBRID" in db_attribute) if it is string type. 0 otherwise.
domain_class_name	VARCHAR(255)	Domain class name if the data type is an object. NULL otherwise.
default_value	VARCHAR(255)	Saved as a character string by default, regardless of data types. If no default value is specified, NULL is stored if a default value is

		NULL, it is displayed as 'NULL'. An object data type is represented as 'volume id page id slot id' while a set data type is represented as '{element 1, element 2, ...}'.
--	--	---

is_nullable	VARCHAR(3)	'NO' if a not null constraint is set, and 'YES' otherwise.
-------------	------------	--

Definition

```
CREATE VCLASS db_attribute (
attr_name, class_name, attr_type, def_order, from_class_name, from_attr_name, data_type,
prec, scale, code_set, domain_class_name, default_value, is_nullable)
AS
SELECT a.attr name, c.class name,
CASE WHEN a.attr type = 0 THEN 'INSTANCE'
WHEN a.attr_type = 1 THEN 'CLASS'
ELSE 'SHARED' END,
a.def_order, a.from_class_of.class_name, a.from_attr_name, t.type_name,
d.prec, d.scale, d.code set, d.class of.class name, a.default value,
CASE WHEN a.is nullable = 0 THEN 'YES' ELSE 'NO' END
FROM db class c, db attribute a, db domain d, db data type t
WHERE a.class_of = c AND d.object_of = a AND d.data_type = t.type_id AND
(CURRENT_USER = 'DBA' OR
{c.owner.name} subseteq (
SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}),
set{}))
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR
{c} subseteq (
SELECT sum(set{au.class of})
FROM db auth au
WHERE {au.grantee.name} subseteq (
SELECT set{CURRENT_USER} +
coalesce(sum(set{t.g.name}), set{}))
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) AND
au.auth type =
'SELECT'))
```

Example 1

The following example shows how to retrieve attributes and data types of the 'event' class.

```
SELECT attr name, data type, domain class name
FROM db attribute
WHERE class name = 'event'
ORDER BY def_order;
attr_name          data_type          domain_class_name
=====
'code'             'INTEGER'         NULL
'sports'          'STRING'          NULL
'name'            'STRING'          NULL
'gender'          'CHAR'            NULL
'players'         'INTEGER'         NULL
```

Example 2

The following example shows how to retrieve attributes of the 'female_event' class and its super class.

```
SELECT attr name, from class name
FROM db attribute
WHERE class_name = 'female_event'
ORDER BY def_order;
attr name          from class name
=====
'code'            'event'
'sports'          'event'
'name'            'event'
'gender'          'event'
'players'         'event'
```

Example 3

The following example shows how to retrieve classes whose attribute names are similar to 'name,' among the ones owned by the current user. (The user is **PUBLIC**.)

```
SELECT a.class_name, a.attr_name
FROM db_class c join db_attribute a ON c.class_name = a.class_name
WHERE c.owner name = CURRENT USER AND attr name like '%name%'
ORDER BY 1;
class name          attr name
=====
'athlete'          'name'
'code'             'f_name'
'code'             's_name'
'event'            'name'
'female event'    'name'
'nation'           'name'
'stadium'          'name'
```

DB_ATTR_SETDOMAIN_ELM

Among attributes of the class to which the current user has access authorization in the database, if an attribute's data type is a set (set, multiset, sequence), this macro represents the data type of the element of the set.

Attribute Name	Data Type	Description
attr_name	VARCHAR(255)	Attribute name
class_name	VARCHAR(255)	Name of the class to which the attribute belongs
attr_type	VARCHAR(8)	'INSTANCE' for an instance attribute, 'CLASS' for a class attribute, and 'SHARED' for a shared attribute.
data_type	VARCHAR(9)	Data type of the element
prec	INTEGER	Precision of the data type of the element
scale	INTEGER	Scale of the data type of the element
code_set	INTEGER	Character set if the data type of the element is a character
domain_class_name	VARCHAR(255)	Domain class name if the data type of the element is an object

Definition

```
CREATE VCLASS db attr setdomain elm (
attr name, class name, attr type,data type, prec, scale, code set, domain class name)
AS
SELECT a.attr_name, c.class_name,
CASE WHEN a.attr type = 0 THEN 'INSTANCE'
WHEN a.attr type = 1 THEN 'CLASS'
ELSE 'SHARED' END,
et.type_name, e.prec, e.scale, e.code_set, e.class_of.class_name
FROM db_class c, db_attribute a, db_domain d,
TABLE(d.set domains) AS t(e), db data type et
WHERE a.class of = c AND d.object of = a AND e.data type = et.type id AND
(CURRENT USER = 'DBA' OR
{c.owner.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name})), set{}}
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR
{c} subseteq (
SELECT sum(set{au.class of})
FROM db_auth au
WHERE {au.grantee.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name})), set{}}
from db user u, table(groups) as t(g)
where u.name = CURRENT_USER ) AND
au.auth_type = 'SELECT'));
```

If the set_attr attribute of class D is of a SET (A, B, C) type, the following three records exist.

Attr_name	Class_name	Attr_type	Data_type	Prec	Scale	Code_set	Domain_class_name
'set_attr'	'D'	'INSTANCE'	'SET'	0	0	0	'A'
'set_attr'	'D'	'INSTANCE'	'SET'	0	0	0	'B'
'set_attr'	'D'	'INSTANCE'	'SET'	0	0	0	'C'

Example

The following example shows how to retrieve set type attributes and data types of the 'city' class (the city table defined in [Containment Operators](#) is created).

```
SELECT attr name, attr type, data type, domain class name
FROM db_attr_setdomain_elm
WHERE class_name = 'city';
attr name          attr type          data type          domain class name
=====
'sports'           'INSTANCE'          'STRING'          NULL
```

DB_METHOD

Represents method information of a class for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
meth_name	VARCHAR(255)	Method name
class_name	VARCHAR(255)	Name of the class to which the method belongs
meth_type	VARCHAR(8)	'INSTANCE' for an instance method, and 'CLASS' for a class method.
from_class_name	VARCHAR(255)	If the method is inherited, the super class in which it is defined is used otherwise NULL
from_meth_name	VARCHAR(255)	If the method is inherited and its name is changed to resolve a name conflict, the original name defined in the super class is used otherwise NULL
func_name	VARCHAR(255)	Name of the C function for the method

Definition

```
CREATE VCLASS db method (
meth name, class name, meth type, from class name, from meth name, func name)
AS
SELECT m.meth name, m.class of.class name,
CASE WHEN m.meth type = 0 THEN 'INSTANCE' ELSE 'CLASS' END,
m.from class of.class name, m.from meth name, s.func name
FROM db method m, db meth sig s
WHERE s.meth_of = m AND
(CURRENT_USER = 'DBA' OR
{m.class of.owner.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}),
set{))
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR
{m.class of} subseteq (
SELECT sum(set{au.class of})
FROM db auth au
WHERE {au.grantee.name} subseteq (
SELECT set{CURRENT_USER} +
coalesce(sum(set{t.g.name}), set{))
from db user u, table(groups) as t(g)
where u.name = CURRENT_USER ) AND
au.auth type =
'SELECT'))
```

Example

The following example shows how to retrieve methods of the 'db_user' class.

```
SELECT meth_name, meth_type, func_name
FROM db_method
WHERE class name = 'db user'
ORDER BY meth type, meth name;
  meth name          meth type          func name
=====
'add_user'          'CLASS'            'au_add_user_method'
'drop_user'         'CLASS'            'au_drop_user_method'
'find user'         'CLASS'            'au find user method'
'login'             'CLASS'            'au login method'
'add member'        'INSTANCE'         'au add member method'
'drop member'       'INSTANCE'         'au drop member method'
'print_authorizations' 'INSTANCE'         'au_describe_user_method'
'set_password'      'INSTANCE'         'au_set_password_method'
'set_password encoded' 'INSTANCE'         'au set password encoded method'
'set_password_encoded_shal' 'INSTANCE'         'au_set_password_encoded_shal_method'
```

DB_METH_ARG

Represents the input/output argument information of the method of the class for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
meth_name	VARCHAR(255)	Method name
class_name	VARCHAR(255)	Name of the class to which the method belongs
meth_type	VARCHAR(8)	'INSTANCE' for an instance method, and 'CLASS' for a class method.
index_of	INTEGER	Order in which arguments are listed in the function definition. Begins with 0 if it is a return value, and 1 if it is an input argument.
data_type	VARCHAR(9)	Data type of the argument
prec	INTEGER	Precision of the argument
scale	INTEGER	Scale of the argument
code_set	INTEGER	Character set if the data type of the argument is a character.
domain_class_name	VARCHAR(255)	Domain class name if the data type of the argument is an object.

Definition

```
CREATE VCLASS db meth arg (
meth name, class name, meth type,
index_of, data_type, prec, scale, code_set, domain_class_name)
AS
SELECT s.meth of.meth name, s.meth of.class of.class name,
CASE WHEN s.meth of.meth type = 0 THEN 'INSTANCE' ELSE 'CLASS' END,
a.index of, t.type name, d.prec, d.scale, d.code set,
d.class of.class name
FROM _db_meth_sig s, _db_meth_arg a, _db_domain d, _db_data_type t
WHERE a.meth_sig_of = s AND d.object_of = a AND d.data_type = t.type_id AND
(CURRENT USER = 'DBA' OR
{s.meth of.class of.owner.name} subseteq (
SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}), set{})
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR
{s.meth of.class of} subseteq (
SELECT sum(set{au.class of})
FROM db_auth au
WHERE {au.grantee.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
from db_user u, table(groups) as t(g)
```

```
where u.name = CURRENT USER ) AND
      au.auth_type = 'SELECT'));
```

Example

The following example shows how to retrieve input arguments of the method of the 'db_user' class.

```
SELECT meth name, data type, prec
FROM db meth arg
WHERE class name = 'db user';
=====
meth name          data type          prec
=====
'append_data'     'STRING'          1073741823
```

DB_METH_ARG_SETDOMAIN_ELM

If the data type of the input/output argument of the method of the class is a set, for which the current user has access authorization in the database, this macro represents the data type of the element of the set.

Attribute Name	Data Type	Description
meth_name	VARCHAR(255)	Method name
class_name	VARCHAR(255)	Name of the class to which the method belongs
meth_type	VARCHAR(8)	'INSTANCE' for an instance method, and 'CLASS' for a class method.
index_of	INTEGER	Order of arguments listed in the function definition. Begins with 0 if it is a return value, and 1 if it is an input argument.
data_type	VARCHAR(9)	Data type of the element
prec	INTEGER	Precision of the element
scale	INTEGER	Scale of the element
code_set	INTEGER	Character set if the data type of the element is a character
domain_class_name	VARCHAR(255)	Domain class name if the data type of the element is an object

Definition

```
CREATE VCLASS db meth arg setdomain elm(
meth name, class name, meth type,
index_of, data_type, prec, scale, code_set, domain_class_name)
AS
SELECT s.meth of.meth name, s.meth of.class of.class name,
      CASE WHEN s.meth of.meth type = 0 THEN 'INSTANCE' ELSE 'CLASS' END,
      a.index of, et.type name, e.prec, e.scale, e.code set,
      e.class of.class name
FROM _db_meth_sig s, _db_meth_arg a, _db_domain d,
     TABLE(d.set domains) AS t(e), db data type et
WHERE a.meth sig of = s AND d.object of = a AND e.data type = et.type id AND
      (CURRENT USER = 'DBA' OR
       {s.meth of.class of.owner.name} subseteq (
         SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
         from db_user u, table(groups) as t(g)
         where u.name = CURRENT_USER ) OR
       {s.meth of.class of} subseteq (
         SELECT sum(set{au.class of})
         FROM db auth au
         WHERE {au.grantee.name} subseteq (
           SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
           from db user u, table(groups) as t(g)
           where u.name = CURRENT_USER ) AND
           au.auth_type = 'SELECT'));
```

DB_METH_FILE

Represents information of a file in which the method of the class for which the current user has access authorization in the database is defined.

Attribute Name	Data Type	Description
class_name	VARCHAR(255)	Name of the class to which the method file belongs
path_name	VARCHAR(255)	File path in which the C function is defined
from_class_name	VARCHAR(255)	Name of the super class in which the method file is defined if the method is inherited, and otherwise NULL

Definition

```
CREATE VCLASS db meth file (class name, path name, from class name)
AS
SELECT f.class of.class name, f.path name, f.from class of.class name
FROM _db_meth_file f
WHERE (CURRENT_USER = 'DBA' OR
       {f.class of.owner.name} subseteq (
           SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}),
set{}))
           from db_user u, table(groups) as t(g)
           where u.name = CURRENT_USER ) OR
       {f.class of} subseteq (
SELECT sum(set{au.class of})
           FROM db_auth au
           WHERE {au.grantee.name} subseteq (
               SELECT set{CURRENT_USER} +
coalesce(sum(set{t.g.name}), set{}))
               from db user u, table(groups) as t(g)
               where u.name = CURRENT_USER ) AND
               au.auth type =
'SELECT'))
```

DB_INDEX

Represents information of indexes created for the class for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
index_name	VARCHAR(255)	Index name
is_unique	VARCHAR(3)	'YES' for a unique index, and 'NO' otherwise.
is_reverse	VARCHAR(3)	'YES' for a reversed index, and 'NO' otherwise.
class_name	VARCHAR(255)	Name of the class to which the index belongs
key_count	INTEGER	The number of attributes that comprise the key
is_primary_key	VARCHAR(3)	'YES' for a primary key, and 'NO' otherwise.
is_foreign_key	VARCHAR(3)	'YES' for a foreign key, and 'NO' otherwise.

Definition

```
CREATE VCLASS db index (index name, is unique, is reverse, class name, key count,
is primary key, is foreign key)
AS
SELECT i.index_name, CASE WHEN i.is unique = 0 THEN 'NO' ELSE 'YES' END,
CASE WHEN i.is_reverse = 0 THEN 'NO' ELSE 'YES' END, i.class_of.class_name, i.key_count,
CASE WHEN i.is primary key = 0 THEN 'NO' ELSE 'YES' END, CASE WHEN i.is foreign key = 0
THEN 'NO' ELSE 'YES' END
FROM db_index i
WHERE (CURRENT_USER = 'DBA' OR
       {i.class_of.owner.name} subseteq (
           SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{}))
           from db_user u, table(groups) as t(g)
```

```

        where u.name = CURRENT USER ) OR
        {i.class of} subseteq (
SELECT sum(set{au.class_of})
        FROM _db_auth au
        WHERE {au.grantee.name} subseteq (
                SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}), set{})
                from db user u, table(groups) as t(g)
                where u.name = CURRENT USER ) AND
                au.auth_type = 'SELECT');

```

Example

The following example shows how to retrieve index information of the class.

```

SELECT class name, index name, is unique
FROM db index
ORDER BY 1;

```

class_name	index_name	is_unique
'athlete'	'pk athlete code'	'YES'
'city'	'pk city city name'	'YES'
'db_serial'	'pk_db_serial_name'	'YES'
'db_user'	'i_db_user_name'	'NO'
'event'	'pk event code'	'YES'
'female event'	'pk event code'	'YES'
'game'	'pk game host year event code athlete code'	'YES'
'game'	'fk game event code'	'NO'
'game'	'fk_game_athlete_code'	'NO'
'history'	'pk_history_event_code_athlete'	'YES'
'nation'	'pk nation code'	'YES'
'olympic'	'pk olympic host year'	'YES'
'participant'	'pk participant host year nation code'	'YES'
'participant'	'fk participant host year'	'NO'
'participant'	'fk_participant_nation_code'	'NO'
'record'	'pk_record_host_year_event_code_athlete_code_medal'	'YES'
'stadium'	'pk_stadium_code'	'YES'

DB_INDEX_KEY

Represents the key information of indexes created for the class for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
index_name	VARCHAR(255)	Index name
class_name	VARCHAR(255)	Name of the class to which the index belongs
key_attr_name	VARCHAR(255)	Name of attributes that comprise the key
key_order	INTEGER	Order of attributes in the key. Begins with 0.
asc_desc	VARCHAR(4)	'DESC' if the order of attribute values is descending, and 'ASC' otherwise.
key_prefix_length	INTEGER	Length of prefix to be used as a key

Definition

```

CREATE VCLASS db index key (index name, class name, key attr name, key order,
key prefix length)
AS
SELECT k.index_of.index_name, k.index_of.class_of.class_name, k.key_attr_name, k.key_order
CASE k.asc_desc
WHEN 0 THEN 'ASC'
WHEN 1 THEN 'DESC' ELSE 'UNKN' END,
k.key prefix length
FROM db index key k
WHERE (CURRENT USER = 'DBA' OR
{k.index_of.class_of.owner.name}
subseteq (

```



```

SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}), set{})
from db user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR {k.index_of.class_of}
subsetq (
  SELECT sum(set{au.class of})
  FROM db auth au
  WHERE {au.grantee.name} subsetq (
    SELECT set{CURRENT USER} + coalesce(sum(set{t.g.name}), set{})
    from db_user u, table(groups) as t(g)
    where u.name = CURRENT_USER ) AND
  au.auth_type = 'SELECT');

```

Example

The following example shows how to retrieve index key information of the class.

```

SELECT class name, key attr name, index name
FROM db_index_key
ORDER BY class_name, key_order;
'athlete'          'code'              'pk athlete code'
'city'             'city name'         'pk city city name'
'db serial'        'name'              'pk db serial name'
'db_user'          'name'              'i_db_user_name'
'event'            'code'              'pk_event_code'
'female event'    'code'              'pk_event_code'
'game'             'host_year'         'pk game host_year event code athlete code'
'game'            'event code'        'fk game event code'
'game'            'athlete code'      'fk game athlete code'
...

```

DB_AUTH

Represents authorization information of classes for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
grantor_name	VARCHAR(255)	Name of the user who grants authorization
grantee_name	VARCHAR(255)	Name of the user who is granted authorization
class_name	VARCHAR(255)	Name of the class for which authorization is to be granted
auth_type	VARCHAR(7)	Name of the authorization type granted
is_grantable	VARCHAR(3)	'YES' if authorization for the class can be granted to other users, and 'NO' otherwise.

Definition

```

CREATE VCLASS db_auth (grantor name, grantee name, class name, auth type, is grantable )
AS
SELECT CAST(a.grantor.name AS VARCHAR(255)),
       CAST(a.grantee.name AS VARCHAR(255)),
       a.class_of.class name, a.auth type,
       CASE WHEN a.is grantable = 0 THEN 'NO' ELSE 'YES' END
FROM db_auth a
WHERE (CURRENT_USER = 'DBA' OR
      {a.class_of.owner.name} subsetq (
        SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
        from db user u, table(groups) as t(g)
        where u.name = CURRENT_USER ) OR
      {a.class_of} subsetq (
SELECT sum(set{au.class of})
  FROM db_auth au
  WHERE {au.grantee.name} subsetq (
    SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
    from db user u, table(groups) as t(g)
    where u.name = CURRENT_USER ) AND
    au.auth_type = 'SELECT'));

```

Example

The following example shows how to retrieve authorization information of the classes whose names begin with 'db_a'.

```
SELECT class_name, auth_type, grantor_name
FROM db_auth
WHERE class name like 'db a%'
ORDER BY 1;
=====
class name          auth type          grantor name
=====
'db_attr_setdomain_elm' 'SELECT'          'DBA'
'db_attribute'       'SELECT'          'DBA'
'db_auth'            'SELECT'          'DBA'
'db_authorization'   'EXECUTE'         'DBA'
'db_authorization'   'SELECT'          'DBA'
'db_authorizations' 'EXECUTE'         'DBA'
'db_authorizations' 'SELECT'          'DBA'
```

DB_TRIG

Represents information of a trigger that has the class for which the current user has access authorization to a database, or its attribute as the target.

Attribute Name	Data Type	Description
trigger_name	VARCHAR(255)	Trigger name
target_class_name	VARCHAR(255)	Target class
target_attr_name	VARCHAR(255)	Target attribute. If not specified in the trigger, NULL
target_attr_type	VARCHAR(8)	Target attribute type. If specified, 'INSTANCE' is used for an instance attribute, and 'CLASS' is used for a class attribute.
action_type	INTEGER	1 for one of INSERT, UPDATE, DELETE, CALL and EVALUATE, 2 for REJECT, 3 for INVALIDATE_TRANSACTION, and 4 for PRINT.
action_time	INTEGER	1 for BEFORE, 2 for AFTER, and 3 for DEFERRED.

Example

The following example shows how to display information of the trigger that has the class for which the current user has access authorization, or its attribute as the target.

```
CREATE VCLASS db trig (
trigger_name, target_class_name, target_attr_name, target_attr_type, action_type,
action_time)
AS
SELECT CAST(t.name AS VARCHAR(255)), c.class name,
CAST(t.target attribute AS VARCHAR(255)),
CASE WHEN t.target_class_attribute = 0 THEN 'INSTANCE' ELSE 'CLASS' END,
t.action_type, t.action_time
FROM db class c, db trigger t
WHERE t.target class = c.class of AND
(CURRENT USER = 'DBA' OR
{c.owner.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
from db_user u, table(groups) as t(g)
where u.name = CURRENT_USER ) OR
{c} subseteq (
SELECT sum(set{au.class of})
FROM db_auth au
WHERE {au.grantee.name} subseteq (
SELECT set{CURRENT_USER} + coalesce(sum(set{t.g.name}), set{})
from db user u, table(groups) as t(g)
where u.name = CURRENT_USER ) AND
au.auth_type = 'SELECT'));
```

DB_PARTITION

Represents information of partitioned classes for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
class_name	VARCHAR(255)	Class name
partition_name	VARCHAR(255)	Partition name
partition_class_name	VARCHAR(255)	Partitioned class name
partition_type	VARCHAR(32)	Partition type (HASH, RANGE, LIST)
partition_expr	VARCHAR(255)	Partition expression
partition_values	SEQUENCE OF	RANGE - MIN/MAX value - For infinite MIN/MAX, NULL LIST - value list

Definition

```
CREATE VCLASS db partition
(sp name, sp type, return type, arg count, lang, target, owner)
AS
SELECT p.class_of.class_name AS class_name, p.pname AS partition_name,
       p.class_of.class_name || '__p__' || p.pname AS partition_class_name,
       CASE WHEN p.ptype = 0 THEN 'HASH'
            WHEN p.ptype = 1 THEN 'RANGE'
            ELSE 'LIST' END AS partition_type,
       TRIM(SUBSTRING( pi.pexpr FROM 8 FOR (POSITION(' FROM ' IN pi.pexpr)-8))) AS
       partition_expression,
       p.pvalues AS partition values
FROM db partition p,
     ( select * from db partition sp
  where sp.class of = p.class of AND sp.pname is null) pi
WHERE p.pname is not null AND
     ( CURRENT_USER = 'DBA'
   OR
     {p.class of.owner.name} SUBSETEQ
     ( SELECT SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name})), SET{}
       FROM db user u, TABLE(groups) AS t(g)
       WHERE u.name = CURRENT_USER
     )
   OR
     {p.class of} SUBSETEQ
     ( SELECT SUM(SET{au.class of})
       FROM db auth au
       WHERE {au.grantee.name} SUBSETEQ
         ( SELECT SET{CURRENT_USER} + COALESCE(SUM(SET{t.g.name})), SET{}
           FROM db user u, TABLE(groups) AS t(g)
           WHERE u.name = CURRENT_USER) AND
           au.auth type = 'SELECT'
         )
     )
 )
 )
```

Example

The following example shows how to retrieve the partition information currently configured for the participant2 class (see examples in [Defining Range Partitions](#)).

```
SELECT * FROM db partition WHERE class name = 'participant2';
 class_name      partition_name      partition_class_name      partition_type
 partition_expr      partition_values
=====
'participant2'      'before 2000'      'participant2 p before 2000' 'RANGE'
'host year'         {NULL, 2000}
'participant2'      'before 2008'      'participant2__p__before_2008' 'RANGE'
'host_year'         {2000, 2008}
```

DB_STORED_PROCEDURE

Represents information of Java stored procedure for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
sp_name	VARCHAR(255)	Stored procedure name
sp_type	VARCHAR(16)	Stored procedure type (function or procedure)
return_type	VARCHAR(16)	Return value type
arg_count	INTEGER	The number of arguments
lang	VARCHAR(16)	Implementing language (currently, Java)
target	VARCHAR(4096)	Name of the Java method to be executed
owner	VARCHAR(256)	Owner

Definition

```
CREATE VCLASS db_stored_procedure
(sp_name, sp_type, return_type, arg_count, lang, target, owner)
AS
SELECT sp.sp name,
       CASE sp.sp type   WHEN 1 THEN 'PROCEDURE'
                        ELSE 'FUNCTION' END,
       CASE WHEN sp.return_type = 0 THEN 'void'
            WHEN sp.return_type = 28 THEN 'CURSOR'
       ELSE ( SELECT dt.type name
              FROM db data type dt
              WHERE sp.return type = dt.type id) END,
       sp.arg_count,
       CASE sp.lang     WHEN 1 THEN 'JAVA'
                        ELSE '' END, sp.target, sp.owner.name
FROM _db_stored_procedure sp
```

Example

The following example shows how to retrieve Java stored procedures owned by the current user.

```
SELECT sp_name, target from db_stored_procedure WHERE sp_type = 'FUNCTION' AND owner =
CURRENT_USER
 sp_name          target
=====
'hello'          'SpCubrid.HelloCubrid() return java.lang.String'
'sp_int'         'SpCubrid.SpInt(int) return int'
```

DB_STORED_PROCEDURE_ARGS

Represents argument information of Java stored procedure for which the current user has access authorization to a database.

Attribute Name	Data Type	Description
sp_name	VARCHAR(255)	Stored procedure name
index_of	INTEGER	Order of the arguments
arg_name	VARCHAR(256)	Argument name
data_type	VARCHAR(16)	Data type of the argument
mode	VARCHAR(6)	Mode (IN, OUT, INOUT)

Definition

```
CREATE VCLASS db_stored_procedure_args (sp_name, index_of, arg_name, data_type, mode)
AS
SELECT sp.sp name, sp.index of, sp.arg name,
       CASE sp.data_type WHEN 28 THEN 'CURSOR'
```

```

ELSE ( SELECT dt.type name FROM db data type dt
      WHERE sp.data type = dt.type id) END,
CASE WHEN sp.mode = 1 THEN 'IN' WHEN sp.mode = 2 THEN 'OUT'
ELSE 'INOUT' END
FROM db stored procedure args sp
ORDER BY sp.sp_name, sp.index_of ;

```

Example

The following example shows how to retrieve arguments the 'phone_info' Java stored procedure in the order of the arguments.

```

SELECT index_of, arg_name, data_type, mode FROM db_stored_procedure_args
WHERE sp name = 'phone info'
ORDER BY index_of

```

index_of	arg_name	data_type	mode
0	'name'	'STRING'	'IN'
1	'phoneno'	'STRING'	'IN'

Catalog Class/Virtual Class Authorization

Catalog classes are created to be owned by **dba**. However, **dba** can only execute **SELECT** operations. If **dba** executes operations such as **UPDATE/DELETE**, an authorization failure error occurs. General users cannot execute queries on system catalog classes.

Although catalog virtual classes are created to be owned by **dba**, all users can perform the **SELECT** statement on catalog virtual classes. Of course, **UPDATE/DELETE** operations on catalog virtual classes are not allowed.

Updating catalog classes/virtual classes is automatically performed by the system when users execute a DDL statement that creates/modifies/deletes a class/attribute/index/user/authorization.

Consistency of Catalog Information

Catalog information is represented by the instance of a catalog class/virtual class. If such information is accessed at the **READ UNCOMMITTED INSTANCES (TRAN_REP_CLASS_UNCOMMIT_INSTANCE** or **TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE)** isolation level, incorrect values (values being changed) can be read. Therefore, to get correct catalog information, you must use the **SELECT** query on the catalog class/virtual class at the **READ COMMITTED INSTANCES** isolation level or higher.

Querying on Catalog

To query on catalog classes, you must convert identifiers such as class, virtual class, attribute, trigger, method and index names to lowercases, and create them. Therefore, you must use lowercases when querying on catalog classes.

```

CREATE TABLE Foo(name varchar(255));
SELECT class name, partitioned FROM db class WHERE class name = 'Foo';

```

There are no results.

```

SELECT class_name, partitioned FROM db_class WHERE class_name = 'foo';
class_name  partitioned
=====
'foo'      'NO'

```


Administrator's Guide

The "Administrator's Guide" provides the database administrators (**DBA**) with details on how to operate the CUBRID system. The guide includes instructions on the following: database management tasks (creating and deleting databases, adding volume, etc.), migration tasks (moving database to a different location or making changes so that it fits the system's version), and making back-ups and rollbacks of the database in case of failures.

It also includes instructions on how to use the CUBRID utilities, which starts and stops various processes of the CUBRID server, the broker and manager server.

This chapter contains the following:

- How to use CUBRID utilities
- How to control the CUBRID (service, database server, broker, manager server)
- How to use the database administrative utilities
- Database migration
- Database backup and restore
- CUBRID HA

CUBRID Utilities

The CUBRID utilities provide features that can be used to comprehensively manage the CUBRID service. The CUBRID utilities are divided into the service management utility, which is used to manage the CUBRID service process, and the database management utility, which is used to manage the database.

The service management utilities are as follows:

- Service utility : Operates and manages the master process.
- cubrid service
- Server utility : Operates and manages the server process.
- cubrid server
- Broker utility : Operates and manages the broker process and application server (CAS) process.
- cubrid broker
- Manager utility : Operates and manages the manager server process.
- cubrid manager
- HA utility : Operates and manages the HA related process.
- cubrid heartbeat

See [Registering Services](#) for details.

The database management utilities are as follows:

- Creating database, adding volume, and deleting database
- cubrid createdb
- cubrid addvoldb
- cubrid deletedb
- Renaming database, altering host, copying database, and installing database
- cubrid renamedb
- cubrid alterdbhost
- cubrid copydb
- cubrid installdb
- Checking and compacting database space
- cubrid spacedb
- cubrid compactdb
- Checking query plan and optimizing database
- cubrid plandump
- cubrid optimizedb
- cubrid statdump
- Checking database lock, killing transaction, and checking consistency
- cubrid lockdb
- cubrid killtran
- cubrid checkdb
- Diagnosing database and dumping parameter
- cubrid diagdb
- cubrid paramdump
- Loading and unloading database
- cubrid loaddb
- cubrid unloaddb
- Backing up and restoring database
- cubrid backupdb
- cubrid restoredb

- HA utilities
- cubrid changemode
- cubrid copylogdb
- cubrid applylogdb
- cubrid applyinfo

See [How to Use the CUBRID Management Utilities \(Syntax\)](#) for details.

The following result is displayed if cubrid is entered in a prompt.

```
% cubrid

cubrid utility, version 2008 R4.1
usage: cubrid <utility-name> [args]
Type 'cubrid <utility-name>' for help on a specific utility.

Available service's utilities:
  service
  server
  broker
  manager
  heartbeat

Available administrator's utilities:
  addvoldb
  alterdbhost
  backupdb
  checkdb
  compactdb
  copydb
  createdb
  deletedb
  diagdb
  installdb
  killtran
  loaddb
  lockdb
  optimizedb
  plandump
  renamedb
  restoredb
  spacedb
  unloaddb
  paramdump
  statdump
  changemode
  copylogdb
  applylogdb
  applyinfo

cubrid is a tool for DBMS.
For additional information, see http://www.cubrid.com
```

Remark

If you want to control the service by using cubrid utility on Windows Vista or later, it is recommended that you run the command prompt with an administrator account.

If you use cubrid utility without an administrator account, the result message is not displayed even though you can run it through the User Account Control (UAC) dialog.

To run the command prompt on Windows Vista or later with an administrator account, right-click [Start] > [All Programs] > [Accessories] > [Command Prompt] and select [Run as Administrator]. In the dialog verifying authorization, click [Yes]; then the command prompt runs with an administrator account.

CUBRID Controls

How to Use CUBRID Utilities (Syntax)

How to use CUBRID utilities (syntax) is as follows:

Controlling CUBRID Service

The following **cubrid** utility syntax shows how to control services registered in the configuration file. One of the followings can be specified in *command*: **start**, **stop**, **restart**, or **status**. No additional option or argument is required.

```
 cubrid service  command
command : { start | stop | restart | status }
```

Controlling Database Server

The following **cubrid** utility syntax shows how to control database server process. One of the followings can be specified in *command*: **start**, **stop**, **restart**, or **status**. Database name must be specified as an argument of commands (except **status**).

```
 cubrid server  command [<database_name>]
command : { start | stop | restart | status }
```

Controlling Broker

The following **cubrid** utility syntax shows how to control CUBRID broker process. One of the followings can be specified in *command*: **start**, **stop**, **restart**, or **status**; in addition to these commands, **on** (to start a specific broker) and **off** (to stop a specific broker) are provided.

```
 cubrid broker  command
command : { start | stop | restart | status [<broker_name>] | on <broker_name> | off
<broker_name> | reset <broker_name> | acl {status|reload} <broker_name> }
```

Controlling CUBRID Manager Server

To use the CUBRID Manager, the Manager server must be running where database server is running. The following **cubrid** utility syntax shows how to control the CUBRID Manager process. One of the followings can be specified in *command*: **start**, **stop**, or **status**.

```
 cubrid manager  command
command : { start | stop | status }
```

Controlling CUBRID HA

The following **cubrid heartbeat** utility syntax shows how to use CUBRID HA. One of the followings can be specified in *command*: **start**, **stop**, **reload**, **deact**, or **act**. For details, see [Utilities of cubrid heartbeat](#).

```
 cubrid heartbeat  command
command : { start | stop | reload | deact | act }
```

CUBRID Services

Registering Services

You can register one or more of database servers, CUBRID brokers, CUBRID Manager(s) or CUBRID HAs as CUBRID service in the configuration file (**cubrid.conf**). If you do not register any service, only master process is registered by default. It is convenient for you to view status of all related processes at a glance or start and stop the processes at once with the **cubrid service** utility once it is registered as CUBRID service. For details on CUBRID HA configuration, see [Utilities of cubrid service](#).

The following example shows how to register database server and broker as service in the **cubrid.conf** file and enable databases (demodb and testdb) to start automatically at once when CUBRID server starts running.

```
# cubrid.conf
...

[service]

# The list of processes to be started automatically by 'cubrid service start' command
# Any combinations are available with server, broker, manager and heartbeat.
service=server,broker

# The list of database servers in all by 'cubrid service start' command.
# This property is effective only when the above 'service' property contains 'server'
keyword.
server=demodb,testdb
```

Starting and Stopping Services

Starting Services

In Linux environment, you can enter the code below to start CUBRID after installation. If no server is registered in the configuration file, only master process runs by default. In the Windows environment, the code below is normally executed only in case that a user with system permission has logged in. An administrator or general user can start or stop the CUBRID server by clicking its icon on the taskbar tray.

```
% cubrid service start
@ cubrid master start
++ cubrid master start: success
```

The following message is returned if master process is already running.

```
% cubrid service start
@ cubrid master start
++ cubrid master is running.
```

The following message is returned if master process fails to run. The example shows that service fails to start due to conflicts of the **cubrid_port_id** parameter value specified in the **cubrid.conf** file. In a such case, you can resolve the problem by changing the port. If it fails to start even though no port is occupied by process, delete /tmp/CUBRID1523 file and then restart the process.

```
% cubrid service start
@ cubrid master start
cub master: '/tmp/CUBRID1523' file for UNIX domain socket exist.... Operation not
permitted
++ cubrid master start: fail
```

After registering service as explained in [Registering Services](#), enter the code below to start the service. You can verify that database server process and broker as well as registered demodb and testdb are starting at once.

```
% cubrid service start
@ cubrid master start
++ cubrid master start: success
@ cubrid server start: demodb
```

This may take a long time depending on the amount of restore works to do.

```
CUBRID 2008 R4.1.....
```

```
++ cubrid server start: success
@ cubrid server start: testdb
```

This may take a long time depending on the amount of recovery works to do.

```
CUBRID 2008 R4.1.....
```

```
++ cubrid server start: success
@ cubrid broker start
++ cubrid broker start: success
```

Stopping Services

Enter code below to stop CUBRID service. If no services are registered by a user, only master process stops and then restarts.

```
% cubrid service stop
@ cubrid master stop
++ cubrid master stop: success
```

Enter code below to stop registered CUBRID service. You can verify that server process, broker process, and master process as well as demodb and testdb stop at once.

```
% cubrid service stop
@ cubrid server stop: demodb
Server demodb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
@ cubrid server stop: testdb
Server testdb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
@ cubrid broker stop
++ cubrid broker stop: success
@ cubrid master stop
++ cubrid master stop: success
```

Restarting Services

Enter code below to restart CUBRID service. If no services are registered by a user, only master process stops and then restarts.

```
% cubrid service restart
@ cubrid master stop
++ cubrid master stop: success
@ cubrid master start
++ cubrid master start: success
```

Enter code below to restart registered CUBRID service. You can verify that server process, broker process, and master process as well as demodb and testdb stop and then restart at once.

```
% cubrid service restart
@ cubrid server stop: demodb
Server demodb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
@ cubrid server stop: testdb
Server testdb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
@ cubrid broker stop
++ cubrid broker stop: success
@ cubrid master stop
++ cubrid master stop: success
@ cubrid master start
++ cubrid master start: success
@ cubrid server start: demodb
```

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.1.....

```
++ cubrid server start: success
@ cubrid server start: testdb
```

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.1.....

```
++ cubrid server start: success
@ cubrid broker start
++ cubrid broker start: success
```

Managing Service Status

The following example shows how to check the status of master process and database server registered.

```
% $ cubrid service status
@ cubrid master status
++ cubrid master is running.
@ cubrid server status

Server testdb (rel 8.4, pid 31059)
Server demodb (rel 8.4, pid 30950)

@ cubrid broker status
% query editor - cub cas [15464,40000]
/home1/cubrid1/CUBRID/log/broker//query_editor.access
/home1/cubrid1/CUBRID/log/broker//query_editor.err
JOB QUEUE:0, AUTO ADD APPL SERVER:ON, SQL LOG MODE:ALL:100000
LONG TRANSACTION TIME:60.00, LONG QUERY TIME:60.00, SESSION TIMEOUT:300
KEEP CONNECTION:AUTO, ACCESS MODE:RW
-----
ID   PID   QPS   LQS PSIZE STATUS
-----
 1 15465    0    0 48032 IDLE
 2 15466    0    0 48036 IDLE
 3 15467    0    0 48036 IDLE
 4 15468    0    0 48036 IDLE
 5 15469    0    0 48032 IDLE

@ cubrid manager server status
++ cubrid manager server is not running.
```

The following message is returned if master process has stopped.

```
% cubrid service status
@ cubrid master status
++ cubrid master is not running.
```

Database Server

Starting and Stopping Database Server

Starting Database Server

The following example shows how to run demodb server.

```
% cubrid server start demodb
@ cubrid server start: demodb

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.1

++ cubrid server start: success
```

If you start demodb server while master process has stopped, master process automatically runs at first and then a specified database server runs.

```
% cubrid server start demodb
@ cubrid master start
++ cubrid master start: success
@ cubrid server start: demodb

This may take a long time depending on the amount of recovery works to do.

CUBRID 2008 R4.1

++ cubrid server start: success
```

The following message is returned while demodb server is running.

```
% cubrid server start demodb
@ cubrid server start: demodb
```

```
++ cubrid server 'demodb' is running.
```

cubrid server start runs cub_server process of a specific database regardless of HA mode configuration. To run database in HA environment, you should use **cubrid heartbeat start**.

Stopping Database Server

The following example shows how to stop demodb server.

```
% cubrid server stop demodb
@ cubrid server stop: demodb
Server demodb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
```

The following message is returned while demodb server has stopped.

```
% cubrid server stop demodb
@ cubrid server stop: demodb
++ cubrid server 'demodb' is not running.
```

cubrid server stop stops cub_server process of a specific database regardless of HA mode configuration. Be careful not to restart the database server or occur failover. To stop database in HA environment, you should use **cubrid heartbeat stop**.

Restarting Database Server

The following example shows how to restart demodb server. demodb server that has already run stops and the server restarts.

```
% cubrid server restart demodb
@ cubrid server stop: demodb
Server demodb notified of shutdown.
This may take several minutes. Please wait.
++ cubrid server stop: success
@ cubrid server start: demodb
```

This may take a long time depending on the amount of recovery works to do.

```
CUBRID 2008 R4.1
```

```
++ cubrid server start: success
```

Checking Database Server Status

The following example shows how to check the status of a database server. Names of currently running database servers are displayed.

```
% cubrid server status
@ cubrid server status
Server testdb (rel 8.4, pid 24465)
Server demodb (rel 8.4, pid 24342)
```

The following example shows the message when master process has stopped.

```
% cubrid server status
@ cubrid server status
++ cubrid master is not running.
```

Limiting Database Server Access

Description

To limit brokers and the CSQL Interpreter connecting to the database server, configure the parameter value of **access_ip_control** in the **cubrid.conf** file to yes and enter the path of a file in which the list of IP addresses allowed to access the **access_ip_control_file** parameter value is written. You should enter the absolute file path. If you enter the relative path, the system will search the file under the **\$CUBRID/conf** directory on Linux and under the **%CUBRID%\conf** directory on Windows.

The following example shows how to configure the **cubrid.conf** file.

```
# cubrid.conf
access_ip_control=yes
access_ip_control_file="/home1/cubrid1/CUBRID/db.access"
```

The following example shows the format of the **access_ip_control_file** file.

```
[@<db_name>]
<ip_addr>
...
```

- **<db_name>** : The name of a database in which access is allowed
- **<ip_addr>** : The IP address allowed to access a database. Using an asterisk (*) at the last digit means that all IP addresses are allowed. Several lines of **<ip_addr>** can be added in the next line of the name of a database.

To configure several databases, it is possible to specify additional **[@<db_name>]** and **<ip_addr>**.

Accessing any IP address except localhost is blocked by server if **access_ip_control** is configured to yes but **ip_control_file** is not configured. A server will not run if analyzing **access_ip_control_file** fails caused by incorrect format.

The following example shows **access_ip_control_file**.

```
[@dbname1]
10.10.10.10
10.156.*

[@dbname2]
*

[@dbname3]
192.168.1.15
```

The example above shows that dbname1 database allows the access of IP addresses starting with 10.156; dbname2 database allows the access of every IP address; dbname3 database allows the access of an IP address, 192.168.1.15, only.

For the database which has already been running, you can modify a configuration file or you can check the currently applied status by using the following commands.

Syntax

To change the contents of **access_ip_control_file** and apply it to server, use the following command.

```
cubrid server acl reload <database_name>
```

- **database_name** : Database name

To display the IP configuration of a sever which is currently running, use the following command.

```
cubrid server acl status <database_name>
```

- **database_name** : Database name

Database Server Log

The following log is created in the file of a server error log if an IP address that is not allowed to access is used.

```
Time: 10/29/10 17:32:42.360 - ERROR *** ERROR CODE = -1022, Tran = 0, CLIENT =
(unknown):(unknown)(-1), EID = 2
Address(10.24.18.66) is not authorized.
```

Note For details on how to limit an access to the broker server, see [Limiting Broker Server Access](#).

Broker

Starting and Stopping Broker

Enter the code below to start the Broker.

```
% cubrid broker start
```

```
@ cubrid broker start
++ cubrid broker start: success
```

The following message is returned if the Broker is already running.

```
% cubrid broker start
@ cubrid broker start
++ cubrid broker is running.
```

Enter the code below to stop the Broker.

```
% cubrid broker stop
@ cubrid broker stop
++ cubrid broker stop: success
```

The following message is returned if the Broker has stopped.

```
% cubrid broker stop
@ cubrid broker stop
++ cubrid broker is not running.
```

Checking Broker Status

Description

The **cubrid broker status** utility allows you to check the broker status such as number of completed jobs and the number of standby jobs.

Syntax

The following syntax shows how to check the CUBRID status; specifying *<expr>* as an argument indicates status monitoring of a specific broker and omitting an argument indicates status monitoring of all brokers registered in the Broker configuration file (**cubrid_broker.conf**).

```
cubrid broker status options [<expr>]
options : [ -b | -f [-l secs] | -q | -t | -s secs ]
```

Options

The following table shows options available with the **cubrid broker status** utility.

Option	Description
<i>expr</i>	Displays the status information of a specific broker whose name includes <i><expr></i> . If it is not specified, status information of all brokers are displayed.
-b	Displays the status information of a broker but does not display information on application server (CAS).
-f [<i>-l secs</i>]	Displays information of DB and host accessed by broker. If it is used with the -b option, additional information on CAS is displayed. The <i>-l secs</i> option is used to specify accumulation period (unit : sec.) when displaying the number of application servers whose client status is Waiting or Busy. If it is omitted, the default value (1 second) is specified.
-q	Displays standby jobs in the job queue.
-t	Displays results in tty mode on the screen. The output can be redirected and used as a file.
-s secs	Regularly displays the status of broker based on specified period. It returns to a command prompt if p is entered.
-f	Displays information of DB and host accessed by broker.

Example

If you do not specify an option or argument to check the status of all brokers, the following result is displayed.

```
% cubrid broker status
@ cubrid broker status
```



```
% query editor - cub cas [28433,40820] /home/CUBRID/log/broker/query_editor.access
/home/CUBRID/
JOB_QUEUE:0, AUTO_ADD_APPL_SERVER:ON, SQL_LOG_MODE:ALL:100000, SLOW_LOG:ON
LONG_TRANSACTION_TIME:60, LONG_QUERY_TIME:60, SESSION_TIMEOUT:300
KEEP_CONNECTION:AUTO, ACCESS_MODE:RW, MAX_QUERY_TIMEOUT:0
```

ID	PID	QPS	LQS	PSIZE	STATUS
1	28434	0	0	50144	IDLE
2	28435	0	0	50144	IDLE
3	28436	0	0	50144	IDLE
4	28437	0	0	50140	IDLE
5	28438	0	0	50144	IDLE

```
% broker1 - cub_cas [28443,40821] /home/CUBRID/log/broker/broker1.access /home/CUBRID/
JOB_QUEUE:0, AUTO_ADD_APPL_SERVER:ON, SQL_LOG_MODE:ALL:100000, SLOW_LOG:ON
LONG_TRANSACTION_TIME:60, LONG_QUERY_TIME:60, SESSION_TIMEOUT:300
KEEP_CONNECTION:AUTO, ACCESS_MODE:RW, MAX_QUERY_TIMEOUT:0
```

ID	PID	QPS	LQS	PSIZE	STATUS
1	28444	0	0	50144	IDLE
2	28445	0	0	50140	IDLE
3	28446	0	0	50144	IDLE
4	28447	0	0	50144	IDLE
5	28448	0	0	50144	IDLE

- % query_editor : Broker name
- cub_cas : Type of the CUBRID application server
- [28433, 40820] : Broker process ID and connection port number of the Broker
- /home/CUBRID/log/broker/query_editor.access : Path of the access log file of query_editor
- JOB_QUEUE : The number of standby jobs in the job queue
- AUTO_ADD_APPL_SERVER : The value of the AUTO_ADD_APPL_SERVER parameter in **cubrid_broker.conf** is ON, which enables the application server to be added automatically.
- SQL_LOG_MODE : The value of the SQL_LOG parameter in the **cubrid_broker.conf** file is ALL, which enables logs for all SQLs to be stored.
- SLOW_LOG : The value of the SQL_LOG parameter in the **cubrid_broker.conf** file is ON, which enables long-duration queries or queries where an error occurred to be recorded in the SLOW SQL LOG file.
- LONG_TRANSACTION_TIME : Execution time of transactions determined by long-duration transaction. It is regarded as long-duration transaction if transaction execution time exceeds 60 seconds.
- LONG_QUERY_TIME : Execution time of queries determined by long-duration query. It is regarded as long-duration query if query execution time exceeds 60 seconds.
- SESSION_TIMEOUT : The timeout value specified to disconnect application server (CAS) sessions in idle state (which any commit or rollback happens) after the transaction has started. If it exceeds specified time in this state, connection between application client and server is closed. The value of SESSION_TIMEOUT parameter in the **cubrid_broker.conf** file is 300 seconds.
- KEEP_CONNECTION : The value of KEEP_CONNECTION parameter in the **cubrid_broker.conf** file is AUTO, which enables an application client to be connected to its application server (CAS) automatically.
- ACCESS_MODE : The Broker action mode; both manipulation and looking up database are allowed in RW mode.
- MAX_QUERY_TIMEOUT : Timeout value of query execution. If it exceeds specified time, the executed query is rolled back. No time limits if the value is 0.
- ID : Serial number of the application server (CAS) within the Broker
- PID : Application server (CAS) process ID within the Broker
- QPS : The number of queries processed per second
- LQS : The number of long-duration queries processed per second
- PSIZE : Size of the application server process
- STATUS : The current status of the application server (BUSY, IDLE, CLIENT_WAIT, CLOSE_WAIT)

To check the status of broker, enter the code below.

```
% cubrid broker status -b
@ cubrid broker status
NAME          PID  PORT  AS  JQ      REQ  TPS  QPS  LONG-T  LONG-Q  ERR-Q
```

```
=====
* query editor 4094 30000 5 0 0 0 0 0/60 0/60 0
* broker1 4104 33000 5 0 0 0 0 0/60 0/60 0
=====
```

- NAME : Broker name
- PID : Process ID of the Broker
- PORT : Port number of the Broker
- AS : The number of application servers
- JQ : The number of standby jobs in the job queue
- REQ : The number of client requests processed by the Broker
- TPS : The number of transactions processed per second (calculated only when the option is configured to "-b -s <sec>")
- QPS : The number of queries processed per second (calculated only when the option is configured to "-b -s <sec>")
- LONG-T : The number of transactions which exceed LONG_TRANSACTION_TIME; the value of the LONG_TRANSACTION_TIME parameter
- LONG-Q : The number of queries which exceed LONG_QUERY_TIME; the value of the LONG_QUERY_TIME parameter
- ERR-Q : The number of queries with errors found

Enter code below to check the status of broker whose name includes broker1 with the **-q** option and job status of a specific broker in the job queue. If you do not specify broker1 as an argument, list of jobs in the job queue for all brokers is displayed.

```
% cubrid broker status -q broker1
@ cubrid broker status
% broker1 - cub cas [28443,40821] /home/CUBRID/log/broker/broker1.access /home/CUBRID/
JOB QUEUE:0, AUTO ADD APPL SERVER:ON, SQL LOG MODE:ALL:100000, SLOW LOG:ON
LONG TRANSACTION TIME:60, LONG QUERY TIME:60, SESSION TIMEOUT:300
KEEP CONNECTION:AUTO, ACCESS MODE:RW, MAX QUERY TIMEOUT:0
-----
ID   PID   QPS   LQS  PSIZE STATUS
-----
1 28444   0     0 50144 IDLE
2 28445   0     0 50140 IDLE
3 28446   0     0 50144 IDLE
4 28447   0     0 50144 IDLE
5 28448   0     0 50144 IDLE
```

Enter code below to input the monitoring interval of broker whose name includes broker1 with the **-s** option and monitor broker status regularly. If you do not specify broker1 as an argument, monitoring status for all brokers is performed regularly. It returns to a command prompt if q is not entered.

```
% cubrid broker status -s 5 broker1
% broker1 - cub cas [28443,40821] /home/CUBRID/log/broker/broker1.access /home/CUBRID/
JOB QUEUE:0, AUTO ADD APPL SERVER:ON, SQL LOG MODE:ALL:100000, SLOW LOG:ON
LONG TRANSACTION TIME:60, LONG QUERY TIME:60, SESSION TIMEOUT:300
KEEP_CONNECTION:AUTO, ACCESS_MODE:RW, MAX_QUERY_TIMEOUT:0
-----
ID   PID   QPS   LQS  PSIZE STATUS
-----
1 28444   0     0 50144 IDLE
2 28445   0     0 50140 IDLE
3 28446   0     0 50144 IDLE
4 28447   0     0 50144 IDLE
5 28448   0     0 50144 IDLE
```

Display information of TPS and QPS to a file with the **-t** option. To cancel the process, press <CTRL+C> to stop program.

```
% cubrid broker status -b -t -s 1 > log_file
```

Enter code below to regularly monitor status of all brokers including TPS and QPS with the **-b** and **-s** options.

```
% cubrid broker status -b -s 1
NAME          PID  PORT  AS  JQ    REQ  TPS  QPS  LONG-T  LONG-Q  ERR-Q
=====
* query_editor 28433 40820 5 0     0    0    0    0/60    0/60    0
* broker1      28443 40821 5 0     0    0    0    0/60    0/60    0
```

Enter code below to view information of server/database accessed by broker, access time, the IP addresses accessed to CAS with the **-f** option.

```
$ cubrid broker status -f broker1
@ cubrid broker status
% broker1 - cub cas [28443,40821] /home/CUBRID/log/broker/broker1.access /home/CUBRID/
JOB QUEUE:0, AUTO ADD APPL SERVER:ON, SQL LOG MODE:ALL:100000, SLOW LOG:ON
LONG TRANSACTION TIME:60, LONG QUERY TIME:60, SESSION TIMEOUT:300
KEEP_CONNECTION:AUTO, ACCESS_MODE:RW, MAX_QUERY_TIMEOUT:0
-----
ID   PID   QPS   LQS  PSIZE STATUS          LAST ACCESS TIME          DB      HOST   LAST
CONNECT TIME          CLIENT IP    SQL LOG MODE    TRANSACTION STIME # CONNECT # RESTART
-----
1 26946   0     0 51168 IDLE             2011/11/16 16:23:42    demodb localhost 2011/11/16
16:23:40          10.0.1.101          NONE 2011/11/16 16:23:42          0         0
2 26947   0     0 51172 IDLE             2011/11/16 16:23:34          -         -
          0         0         0.0.0.0          -         -
3 26948   0     0 51172 IDLE             2011/11/16 16:23:34          -         -
          0         0         0.0.0.0          -         -
4 26949   0     0 51172 IDLE             2011/11/16 16:23:34          -         -
          0         0         0.0.0.0          -         -
5 26950   0     0 51172 IDLE             2011/11/16 16:23:34          -         -
          0         0         0.0.0.0          -         -
          0         0
```

Meaning of every column in code above is as follows:

- LAST ACCESS TIME : Last time when an application client accesses an application server (CAS)
- DB : Name of a database which an application server (CAS) accesses most recently
- HOST : Name of a which an application server (CAS) accesses most recently
- LAST CONNECT TIME: Most recent time when an application server (CAS) accesses a database
- CLIENT IP : IP of an application clients currently being connected to an application server (CAS). If no application client is connected, 0.0.0.0 is displayed.
- SQL_LOG_MODE : SQL logging mode of an application server (CAS). If the mode is same as the mode configured in Broker, "-" is displayed.
- TRANSACTION STIME : Transaction start time
- # CONNECT : The number of connections that an application client accesses to an application server (CAS) after starting Broker
- # RESTART : The number of connection that an application server (CAS) is re-running after starting Broker

Enter code below to display information on AS (T W B Ns-W Ns-B) and CANCELED with the **-b** and **-f** options.

```
// The -f option is added upon execution of Broker status information. Configuring Ns-W
and Ns-B are displayed as long as N seconds by using the -l.
% cubrid broker status -b -f -l 2
@ cubrid broker status
NAME          PID   PSIZE PORT AS(T W B 2s-W 2s-B) JQ REQ TPS QPS LONG-T LONG-Q ERR-Q
CANCELED ACCESS MODE SQL LOG
=====
query_editor 16784 56700 38000 5 0 0 0 0 0 0 0 0/60.0
0/60.0 0 0 RW ALL
```

Meaning of every column in code above is as follows:

- AS(T) : Total number of CASs being executed
- AS(W) : The number of CASs in the status of Waiting
- AS(B) : The number of CASs in the status of Busy
- AS(Ns-W) : The number of CASs that the client belongs to has been waited for N seconds.
- AS(Ns-B) : The number of CASs that the client belongs to has been Busy for N seconds.
- CANCELED : The number of queries have canceled by user interruption since Broker is started (if it is used with the **-l N** option, it specifies the number of accumulations for N seconds.).

Limiting Broker Server Access

Description

To limit the client applications accessing the broker, set to **ON** for the **ACCESS_CONTROL** parameter in the **cubrid_broker.conf** file, and enter a name of the file in which the users and the list of databases and IP addresses allowed to access the **ACCESS_CONTROL_FILE** parameter value are written. The default value of the **ACCESS_CONTROL** broker parameter is **OFF**. The **ACCESS_CONTROL** and **ACCESS_CONTROL_FILE** parameters must be written under [broker] in which common parameters are specified.

The format of **ACCESS_CONTROL_FILE** is as follows:

```
[%<broker_name>]
<db_name>:<db_user>:<ip_list_file>
...
```

- <broker_name> : A broker name. It is the one of broker names specified in **cubrid_broker.conf**.
- <db_name> : A database name. If it is specified as *, all databases are allowed to access the broker server.
- <db_user> : A database user ID. If it is specified as *, all database user IDs are allowed to access the broker server.
- <ip_list_file> : Names of files in which the list of accessible IPs are stored. Several files such as ip_list_file1, ip_list_file2, ... can be specified by using a comma (,).

[%<broker_name>] and <db_name>:<db_user>:<ip_list_file> can be specified separately for each broker.

The format of the ip_list_file is as follows:

```
<ip_addr>
...
```

- <ip_addr> : An IP address that is allowed to access the server. If the last digit of the address is specified as *, all IP addresses in that range are allowed to access the broker server.

If a value for **ACCESS_CONTROL** is set to **ON** and a value for **ACCESS_CONTROL_FILE** is not specified, the broker will only allow the access requests from the localhost. If the analysis of **ACCESS_CONTROL_FILE** and ip_list_file fails while a broker is running, the broker will only allow the access requests from the localhost.

If the analysis of **ACCESS_CONTROL_FILE** and ip_list_file fails while a broker is running, the broker will not run.

```
# cubrid broker.conf
[broker]
MASTER SHM ID          =30001
ADMIN LOG FILE         =log/broker/cubrid broker.log
ACCESS_CONTROL         =ON
ACCESS_CONTROL_FILE    =/home1/cubrid/access_file.txt
[%QUERY_EDITOR]
SERVICE               =ON
BROKER PORT            =38000
.....
```

The following example shows the content of **ACCESS_CONTROL_FILE**. The * symbol represents everything, and you can use it when you want to specify database names, database user IDs and IPs in the IP list file which are allowed to access the broker server.

```
[%QUERY_EDITOR]
dbname1:dbuser1:READIP.txt
dbname1:dbuser2:WRITEIP1.txt,WRITEIP2.txt
*:dba:READIP.txt
*:dba:WRITEIP1.txt
*:dba:WRITEIP2.txt

[%BROKER2]
dbname:dbuser:iplist2.txt

[%BROKER3]
dbname:dbuser:iplist2.txt

[%BROKER4]
dbname:dbuser:iplist2.txt
```

The brokers specified above are **QUERY_EDITOR**, **BROKER2**, **BROKER3** and **BROKER4**.

The QUERY_EDITOR broker only allows the following application access requests.

- When a user logging into dbname1 with a dbuser1 account connects from IPs registered in READIP.txt
- When a user logging into dbname1 with a dbuser2 account connects from IPs registered in WRITEIP1.txt and WRITEIP2.txt

When a user logging into every database with a dba account connects from IPs registered in READIP.txt, WRITEIP1.txt, and WRITEIP2.txt

The following example shows how to specify the IPs allowed in ip_list_file.

```
192.168.1.25
192.168.*
10.*
*
```

The descriptions for the IPs specified in the example above are as follows:

- The first line setting allows an access from 192.168.1.25.
- The second line setting allows an access from all IPs starting with 192.168.
- The third line setting allows an access from all IPs starting with 10.
- The fourth line setting allows an access from all IPs.

For the broker which has already been running, you can modify the configuration file or check the currently applied status of configuration by using the following commands.

Syntax

To configure databases, database user IDs and IPs allowed to access the broker and then apply the modified configuration to the server, use the following command.

```
cubrid broker acl reload [<BR_NAME>]
```

- *BR_NAME* : A broker name. If you specify this value, you can apply the changes only to specified brokers. If you omit it, you can apply the changes to all brokers.

To display the databases, database user IDs and IPs that are allowed to access the broker in running on the screen, use the following command.

```
cubrid broker acl status [<BR_NAME>]
```

- *BR_NAME* : A broker name. If you specify the value, you can display the specified broker configuration. If you omit it, you can display all broker configurations.

Broker Logs

If you try to access brokers through IP addresses that are not allowed, the following logs will be created.

- ACCESS_LOG

```
1 192.10.10.10 - - 1288340944.198 1288340944.198 2010/10/29 17:29:04 ~ 2010/10/29
17:29:04 14942 - -1 db1 dba : rejected
```

- SQL LOG

```
10/29 10:28:57.591 (0) CLIENT IP 192.10.10.10 10/29 10:28:57.592 (0) connect db db1
user dba url jdbc:cubrid:192.10.10.10:30000:db1:: - rejected
```

Note For details on how to limit an access to the database server, see [Limiting Database Server Access](#).

Managing a Specific Broker

Enter the code below to run broker1 only. Note that broker1 should have been configured in the shared memory.

```
% cubrid broker on broker1
```

The following message is returned if broker1 is not configured in the shared memory.

```
% cubrid broker on broker1
Cannot open shared memory
```

Enter the code below to stop *broker1* only. Note that service pool of *broker1* is also removed.

```
% cubrid broker off broker1
```

Enter the code below to restart broker1.

```
% cubrid broker restart broker1
```

The broker reset feature enables broker application servers (CAS) to disconnect the existing connection and reconnect when the servers are connected to unwanted databases due to failover and etc in HA. For example, once Read Only broker is connected to active servers, it is not automatically connected to standby servers although standby servers are available. Connecting to standby servers is allowed only with the **cubrid broker reset** command.

Enter the code below to reset broker1.

```
% cubrid broker reset broker1
```

Dynamically Changing Broker Parameters

Description

You can configure the parameters related to running the Broker in the Broker configuration file (**cubrid_broker.conf**/cubrid_broker.conf). For details, see [Parameter by Broker](#) in the "Performance Management Guide." You can also modify some broker parameters temporarily while the Broker is running by using the `broker_changer` utility. The following broker parameters can be modified dynamically.

- ACCESS_MODE
- ACCESS_LOG
- APPL_SERVER_MAX_SIZE
- KEEP_CONNECTION
- LOG_BACKUP
- LONG_QUERY_TIME
- LONG_TRANSACTION_TIME
- MAX_QUERY_TIMEOUT
- SLOW_LOG
- SQL_LOG
- SQL_LOG_MAX_SIZE
- STATEMENT_POOLING
- TIME_TO_KILL

Syntax

The syntax for the **broker_changer** utility, which is used to change broker parameters while the Broker is running, is as follows. Enter the name of the currently running Broker for the *broker_name*. The *parameters* can be used only for dynamically modifiable parameters. The *value* must be specified based on the parameter to be modified. You can specify CAS identifier (*cas_id*) to apply the changes to the specific CAS. *cas_id* is an ID to be output by **cubrid broker status** command.

```
broker_changer broker_name [cas_id] parameters value
```

Example 1

Enter the following to configure the `SQL_LOG` parameter to ON so that SQL logs can be written to the currently running Broker. Such dynamic parameter change is effective only while the Broker is running.

```
% broker changer query editor sql log on
OK
```

Example 2

Enter the following to change Broker's `ACCESS_MODE` to **Read Only** and automatically reset the Broker in HA environment.

```
% broker_changer broker_m access_mode ro
```

OK

Note If you want to control the service using Cubrid utilities on Windows Vista or the later versions of Windows, you are recommended to open the command prompt window as an administrator. For details, see the notes of [CUBRID Utilities](#).

Broker Logs

There are three types of logs that relate to starting the Broker: access, error and SQL logs. Each log can be found in the log directory under the installation directory. You can change the directory where these logs are to be stored through LOG_DIR and ERROR_LOG_DIR parameters of the Broker configuration file (cubrid_broker.conf).

Checking the Access Log

The access log file records information on the application client and is stored with the name of *broker_name.access*. If the LOG_BACKUP parameter is configured to ON in the Broker configuration file, when the Broker stops properly, the access log file is stored with the date and time that the Broker has stopped. For example, if broker1 stopped at 12:27 P.M. on June 17, 2008, an access file named broker1.access.20080617.1227 is generated in the **log/broker** directory. The following example shows an access log.

The following example and description show an access log file created in the log directory:

```
1 192.168.1.203 - - 972523031.298 972523032.058 2008/06/17 12:27:46~2008/06/17 12:27:47
7118 - -1
2 192.168.1.203 - - 972523052.778 972523052.815 2008/06/17 12:27:47~2008/06/17 12:27:47
7119 ERR 1025
1 192.168.1.203 - - 972523052.778 972523052.815 2008/06/17 12:27:49~2008/06/17 12:27:49
7118 - -1
```

- 1 : ID assigned to the application server of the Broker
- 192.168.1.203 : IP address of the application client
- 972523031.298 : UNIX timestamp value when the client's request processing started
- 2008/06/17 12:27:46 : Time when the client's request processing started
- 972523032.058 : Unix timestamp value when the client's request processing finished
- 2008/06/17 12:27:47 : Time when the client's request processing finished
- 7118 : Process ID of the application server
- -1 : No error occurred during the request processing
- ERR 1025 : Error occurred during the request processing. Error information exists in offset=1025 of the error log file

Checking the Error Log

The error log file records information on errors that occurred during the client's request processing and is stored with the name of *broker_name_app_server_num.err*.

The following example and description show an error log:

```
Time: 02/04/09 13:45:17.687 - SYNTAX ERROR *** ERROR CODE = -493, Tran = 1, EID = 38
Syntax: Unknown class "unknown_tbl". select * from unknown_tbl
```

- Time : 02/04/09 13:45:17.687 : Time when the error occurred
- - SYNTAX ERROR : Type of error (e.g. SYNTAX ERROR, ERROR, etc.)
- *** ERROR CODE = -493 : Error code
- Tran = 1 : Transaction ID. -1 indicates that no transaction ID is assigned.
- EID = 38 : Error ID. This ID is used to find the SQL log related to the server or client logs when an error occurs during SQL statement processing.
- Syntax... : Error message (An ellipsis (...) indicates omission.)

Managing the SQL Log

The SQL log file records SQL statements requested by the application client and is stored with the name of *broker_name_app_server_num.sql.log*. The SQL log is generated in the log/broker/sql_log directory when the

SQL_LOG parameter is set to ON. Note that the size of the SQL log file to be generated cannot exceed the value set for the SQL_LOG_MAX_SIZE parameter. CUBRID offers the **broker_log_top**, **broker_log_converter**, and **broker_log_runner** utilities to manage SQL logs. Each utility should be executed in a directory where the corresponding SQL log exists.

The following examples and descriptions show SQL log files:

```
02/04 13:45:17.687 (38) prepare 0 insert into unique_tbl values (1)
02/04 13:45:17.687 (38) prepare srv_h_id 1
02/04 13:45:17.687 (38) execute srv_h_id 1 insert into unique_tbl values (1)
02/04 13:45:17.687 (38) execute error:-670 tuple 0 time 0.000, EID = 39
02/04 13:45:17.687 (0) auto rollback
02/04 13:45:17.687 (0) auto rollback 0
*** 0.000
02/04 13:45:17.687 (39) prepare 0 select * from unique_tbl
02/04 13:45:17.687 (39) prepare srv_h_id 1 (PC)
02/04 13:45:17.687 (39) execute srv_h_id 1 select * from unique_tbl
02/04 13:45:17.687 (39) execute 0 tuple 1 time 0.000
02/04 13:45:17.687 (0) auto commit
02/04 13:45:17.687 (0) auto_commit 0
*** 0.000
```

- 02/04 13:45:17.687 : Time when the application sent the request
- (39) : Sequence number of the SQL statement group. If prepared statement pooling is used, it is uniquely assigned to each SQL statement in the file.
- prepare 0 : Whether or not it is a prepared statement
- prepare srv_h_id 1 : Prepares the SQL statement as srv_h_id 1.
- (PC) : It is displayed if the data in the plan cache is used.
- SELECT... : SQL statement to be executed. (An ellipsis (...) indicates omission.) For statement pooling, the binding variable of the WHERE clause is represented as a question mark (?).
- Execute 0 tuple 1 time 0.000 : One row is executed. The time spent is 0.000 seconds.
- auto_commit/auto_rollback : Automatically committed or rolled back. The second auto_commit/auto_rollback is an error code. 0 indicates that the transaction has been completed without an error.

The **broker_log_top** utility analyses the SQL logs which are generated for a specific period. As a result, the information of SQL statements and time execution are displayed in files by order of the longest execution time; the results of SQL statements are stored in **log.top.q** and those of execution time are stored in **log.top.res**, respectively.

The **broker_log_top** utility is useful to analyse the long query. The syntax is as follows:

```
broker_log_top [options] sql_log_file_list
options : {-t | -F from_date | -T to_date}
```

The result is displayed in transaction unit if the **-t** option is specified.

SQL statements which are used for a specific period time can be analyzed by using the **-F** and **-T** options. The input format is MM/DD[hh:mm[:ss[:msec]]], and the part enclosed by [] can be omitted. If you omit the value, it is regarded as that 01 is input for DD, and 0 is input for hh, mm, ss and msec.

```
-- Set the search range to milliseconds
broker_log_top -F "01/19 15:00:25.000" -T "01/19 15:15:25.180" log1.log

-- The part where the time format is omitted is set to 0 by default. This means that -F
"01/19 00:00:00.000" -T "01/20 00:00:00.000" is input.
broker_log_top -F "01/19" -T "01/20" log1.log
```

All logs are displayed by SQL statement if any option is not specified.

The following logs are the results of executing the **broker_log_top** utility; logs are generated from Nov. 11th to Nov. 12th, and it is displayed in the order of the longest execution of SQL statements. Each month and day are separated by a slash (/) when specifying period. Note that "*.sql.log" is not recognized so the SQL logs should be separated by a white space on Windows.

```
--Execution broker log top on Linux
% broker_log_top -F "11/11" -T "11/12" -t *.sql.log

query_editor_1.sql.log
query_editor_2.sql.log
```



```

query_editor_3.sql.log
query_editor_4.sql.log
query_editor_5.sql.log

--Executing broker log top on Windows
% broker log top -F "11/11" -T "11/12" -t query_editor_1.sql.log query_editor_2.sql.log
query_editor_3.sql.log query_editor_4.sql.log query_editor_5.sql.log

```

The `log.top.q` and `log.top.res` files are generated in the same directory where the analyzed logs are stored when executing the example above; In the `log.top.q` file, you can view each SQL statement, and its line number. In the `log.top.res`, you can the minimum, maximum and avg. time, and the number of execution queries for each SQL statement.

```

--log.top.q file
[Q1]-----
broker1 6.sql.log:137734
11/11 18:17:59.396 (27754) execute all srv_h_id 34 select a.int col, b.var col from
dml v view 6 a, dml v view 6 b, dml v view 6 c , dml v view 6 d, dml v view 6 e where
a.int col=b.int col and b.int col=c.int col and c.int col=d.int col and
d.int_col=e.int_col order by 1,2;
11/11 18:18:58.378 (27754) execute_all 0 tuple 497664 time 58.982
.
.
[Q4]-----
broker1 100.sql.log:142068
11/11 18:12:38.387 (27268) execute_all srv_h_id 798 drop table list_test;
11/11 18:13:08.856 (27268) execute_all 0 tuple 0 time 30.469

-- log.top.res file contents
max          min          avg          cnt(err)
-----
[Q1]          58.982       30.371       44.676       2 (0)
[Q2]          49.556       24.023       32.688       6 (0)
[Q3]          35.548       25.650       30.599       2 (0)
[Q4]          30.469       0.001        0.103      1050 (0)

```

To store SQL logs created in `log/broker/sql_log` under the installation directory to a separate file, the **broker_log_converter** utility is executed. The syntax of the **broker_log_converter** utility is as follows. The example shows how to store queries in the `query_editor_1.sql.log` file to the `query_convert.in` file.

```
broker_log_converter SQL_log_file output_file
```

The following example shows how to convert the query in the `query_editor_1.sql.log` file into the `query_convert.in` file.

```
% broker_log_converter query_editor_1.sql.log query_convert.in
```

To re-execute queries stored in the query file which has been created by the **broker_log_converter** utility, the **broker_log_runner** utility is executed. The syntax of the **broker_log_runner** utility is as follows: The example shows how to re-executes queries store in the `query_convert.in` of `demodb`. It is assumed that the IP address of the Broker is 192.168.1.10 and its port number is 30,000.

```
broker_log_runner
broker_log_runner options input_file
options : -I broker_ip -P broker_port -d dbname [-u dbuser [-p dbpasswd ]] [-t
num_thread] [-r repeat_count] [-Q] [-o result_file]
```

broker_log_runner Utility Options

Option	Description
-I <i>broker_ip</i>	IP address or host name of the CUBRID Broker
-P <i>broker_port</i>	Port number of the CUBRID Broker
-d <i>dbname</i>	Name of the database against which queries are to be executed
-u <i>dbuser</i>	Database user name (default value : public)
-p <i>dbpasswd</i>	Database password
-t <i>numthread</i>	The number of threads (default value : 1)
-r <i>repeat_count</i>	The number of times that the query is to be executed (default value : 1)

-Q Stores the query plan in *result_file*.
Name of the file where execution results are to be stored

-o result_file Name of the file where execution results are to be stored

```
% broker log runner -I 192.168.1.10 -P 30000 -d demodb -t 2 query convert.in
broker ip = 192.168.1.10
broker_port = 30000
num_thread = 2
repeat = 1
dbname = demodb
dbuser = public
dbpasswd =
exec_time : 0.001
exec_time : 0.000
0.000500 0.000500 -

% broker log runner -I 192.168.1.10 -P 30000 -d demodb -o result -Q query convert.in
...
%cat result.0
----- query -----
SELECT * FROM athlete where code=10099;

cci execute:1
----- query plan -----
Join graph segments (f indicates final):
seg[0]: [0]
seg[1]: code[0] (f)
seg[2]: name[0] (f)
seg[3]: gender[0] (f)
seg[4]: nation_code[0] (f)
seg[5]: event[0] (f)
Join graph nodes:
node[0]: athlete athlete(6677/107) (sargs 0)
Join graph terms:
term[0]: (athlete.code=10099) (sel 0.000149768) (sarg term) (not-join eligible) (indexable
code[0]) (loc 0)

Query plan:

iscan
  class: athlete node[0]
  index: pk_athlete_code term[0]
  cost: fixed 0(0.0/0.0) var 0(0.0/0.0) card 1

Query stmt:

select athlete.code, athlete.[name], athlete.gender, athlete.nation code, athlete.event
from athlete athlete where (athlete.code= ?:0 )

----- query result -----
10099|Andersson Magnus|M|SWE|Handball|
-- 1 rows -----
```

CUBRID Manager Server

Starting and Stopping CUBRID Manager

Starting the CUBRID Manager

The following example shows how to start the CUBRID Manager server.

```
% cubrid manager start
```

The following message is returned if the CUBRID Manager server is already running.

```
% cubrid manager start
@ cubrid manager server start
```

```
++ cubrid manager server is running.
```

Stopping the CUBRID Manager

The following example shows how to stop the CUBRID Manager server.

```
% cubrid manager stop  
@ cubrid manager server stop  
++ cubrid manager server stop: success
```

CUBRID Manager Server Log

The logs of CUBRID Manager server are stored in the log/manager directory under the installation directory. There are four types of log files depending on server process of CUBRID Manager.

- cub_auto.access.log : Access log of a client that has successfully logged into and out of the CUBRID Manager server
- cub_auto.error.log : Access log of a client that failed to log into or out of the CUBRID Manager Server
- cub_js.access.log : Job log processed by the CUBRID Manager server
- cub_js.error.log : Error log that occurred while the CUBRID Manager server has been processing jobs

Database Administration

How to Use the CUBRID Management Utilities (Syntax)

The following shows how to use the CUBRID management utilities.

```

cubrid utility_name
utility_name :
createdb [option] <database_name> --- Creating a database
deletedb [option] <database_name> --- Deleting a database
installdb [option] <database_name> --- Installing a database
renamedb [option] <source-database-name> <target-database-name> --- Renaming a database
copydb [option] <source-database-name> <target-database-name> --- Copying a database
backupdb [option] <database-name> --- Backing up a database
restoredb [option] <database-name> --- Restoring a database
addvol [option] <database-name> --- Adding a database volume file
spacedb [option] <database-name> --- Displaying details of database space
lockdb [option] <database-name> --- Displaying details of database lock
killtran [option] <database-name> --- Removing transactions
optimizedb [option] <database-name> --- Updating database statistics
statdump [option] <database-name> --- Outputting statistic information of database
server execution
compactdb [option] <database-name> --- Optimizing space by freeing unused space
diagdb [option] <database-name> --- Displaying internal information
checkdb [option] <database-name> --- Checking database consistency
alterdbhost [option] <database-name> --- Altering database host
plandump [option] <database-name> --- Displaying details of the query plan
loaddb [option] <database-name> --- Loading data and schema
unloaddb [option] <database-name> --- Unloading data and schema
paramdump [option] <database-name> --- Checking out the parameter values configured in
a database
changemode [option] <database-name> --- Displaying or changing the server HA mode
copylogdb [option] <database-name> --- Multiplating transaction logs to configure HA
applylogdb [option] <database-name> --- Reading and applying replication logs from
transaction logs to configure HA

```

Database Users

A CUBRID database user can have members with the same authorization. If authorization **A** is granted to a user, the same authorization is also granted to all members belonging to the user. A database user and its members are called a "group."

CUBRID provides **DBA** and **PUBLIC** users by default.

- **DBA** can access every object in the database, that is, it has authorization at the highest level. Only **DBA** has sufficient authorization to add, alter and delete the database users.
- All users including **DBA** are members of **PUBLIC**. Therefore, all database users have the authorization granted to **PUBLIC**. For example, if authorization **B** is added to **PUBLIC** group, all database members will automatically have the **B** authorization.

databases.txt File

Description

CUBRID stores information on the locations of all existing databases in the **databases.txt** file. This file is called the "database location file." A database location file is used when CUBRID executes utilities for creating, renaming, deleting or replicating databases; it is also used when CUBRID runs each database. By default, this file is located in the **databases** directory under the installation directory. The directory is located through the environment variable **CUBRID_DATABASES**.

Syntax

```
db_name db_directory server_host logfile_directory
```

The format of each line of a database location file is the same as defined by the above syntax; it contains information on the database name, database path, server host and the path to the log files. The following example shows how to check the contents of a database location file.

```
% more databases.txt
dist testdb /home1/user/CUBRID/bin d85007 /home1/user/CUBRID/bin
dist demodb /home1/user/CUBRID/bin d85007 /home1/user/CUBRID/bin
testdb /home1/user/CUBRID/databases/testdb d85007 /home1/user/CUBRID/databases/testdb
demodb /home1/user/CUBRID/databases/demodb d85007 /home1/user/CUBRID/databases/demodb
```

By default, the database location file is stored in the **databases** directory under the installation directory. You can change the default directory by modifying the value of the **CUBRID_DATABASES** environment variable. The path to the database location file must be valid so that the **cuprid** utility for database management can access the file properly. You must enter the directory path correctly and check if you have write permission on the file. The following example shows how to check the value configured in the **CUBRID_DATABASES** environment variable.

```
% set | grep CUBRID DATABASES
CUBRID_DATABASES=/home1/user/CUBRID/databases
```

An error occurs if an invalid directory path is set in the **CUBRID_DATABASES** environment variable. If the directory path is valid but the database location file does not exist, a new location information file is created. If the **CUBRID_DATABASES** environment variable has not been configured at all, CUBRID retrieves the location information file in the current working directory.

Creating Database

Description

The **cuprid createdb** utility creates databases and initializes them with the built-in CUBRID system tables. It can also define initial users to be authorized in the database and specify the locations of the logs and databases. In general, the **cuprid createdb** utility is used only by DBA.

Syntax

```
cuprid createdb options database_name
options :
[--db-volume-size=size] [--db-page-size=size] [--log-volume-size=size] [--log-page-size=size]
[--comment=comment] [{-F |--file-path=path] [{-L |--log-path=path] [{-B |--lob-base-path=path]
[--server-name=host] [-r|--replace] [--more-volume-file=file] [--user-definition-file=file]
[--csql-initialization-file=file] [{-o |--output-file=file] [-v|--verbose]
```

- **cuprid** : An integrated utility for the CUBRID service and database management.
- **createdb** : A command used to create a new database.
- *options* : A short option starts with a single dash (-) while a full name option starts with a double dash (--).
- *database_name* : Specifies a unique name for the database to be created, without including the path name to the directory where the database will be created. If the specified database name is the same as that of an existing database name, CUBRID halts creation of the database to protect existing files.

Options

The following table shows options available with the **cuprid createdb** utility (options are case sensitive).

Option	Description
--db-volume-size	Specifies the size of the database volume that will be created first in bytes. Default value : A value of <i>db_volume_size</i> , the system parameter
--db-page-size	Specifies the database page size in bytes. Default value : 16K
--log-volume-size	Specifies the log volume size in bytes.
--log-page-size	Specifies the page size of log volume in bytes. Default value : Database page size

--comment	Adds information on the database to be created in the form of a comment.
-F --file-path	Specifies the directory path where the database will be created. Default value : Current working directory
-L --log-path	Specifies the directory path where log files will be stored. Default value : A directory path specified with the -F option
-B --lob-base-path	Specifies the directory path where LOB data files will be stored. Default value : <location of database volumns created>/lob directory
--server-name	Specifies the name of the server host to connect to. Default value : localhost
-r --replace	Allows overwriting if the name of the database to be created is the same as that of an existing database. Default value : Deactivated
--more-volume-file	Specifies the file that includes the specifications for creating an additional volume of the database.
--user-definition-file	Specifies the file that includes user definitions.
--csql-initialization-file	Specifies the file for csql initialization.
-o --output-file	Specifies the file where output messages concerning database creation are stored.
-v --verbose	Displays detailed messages to the screen concerning database creation. Default value : Deactivated

Size of the first database volume (--db-volume-size)

The **--db-volume-size** option specifies the size of the database volume that will be created first. The default value is the value of the system parameter **db_volume_size**, and the minimum value is 20M. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB), and terabytes (TB) respectively. If you omit the unit, bytes will be applied.

The following example shows how to create a database named testdb and assign 512 MB to its first volume.

```
cubrid createdb --db-volume-size=512M testdb
```

Database page size (--db-page-size)

The **--db-page-size** option specifies the size of the database page; the minimum value is 4K and the maximum value is **16K** (default). K stands for kilobytes (KB).

The value of page size is one of the followings: 4K, 8K, or 16K. If a value between 4K and 16K is specified, system rounds up the number. If a value greater than 16K or less than 4K, the specified number is used.

The following example shows how to create a database named testdb and configure its page size 16K.

```
cubrid createdb --db-page-size=16K testdb
```

Log volume size (--log-volume-size)

The **--log-volume-size** option specifies the size of the database log volume. The default value is the same as database volume size, and the minimum value is 20M. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB), and terabytes (TB) respectively. If you omit the unit, bytes will be applied.

The following example shows how to create a database named testdb and assign 256 MB to its log volume.

```
cubrid createdb --log-volume-size=256M testdb
```

Log page size (--log-page-size)

The **--log-page-size** option specifies the size of the log volume page. The default value is the same as data page size. The minimum value is 4K and the maximum value is 16K. K stands for kilobytes (KB).

The value of page size is one of the followings: 4K, 8K, or 16K. If a value between 4K and 16K is specified, system rounds up the number. If a value greater than 16K or less than 4K, the specified number is used.

The following example shows how to create a database named testdb and configure its log volume page size 8K.

```
cubrid createdb --log-page-size=8K testdb
```

Comment (--comment)

The **--comment** option specifies a comment to be included in the database volume header. If the character string contains spaces, the comment must be enclosed in double quotes.

The following example shows how to create a database named testdb and add a comment to the database volume.

```
cubrid createdb --comment "a new database for study" testdb
```

Database directory path (-F)

The **-F** option specifies an absolute path to a directory where the new database will be created. If the **-F** option is not specified, the new database is created in the current working directory.

The following example shows how to create a database named testdb in the directory /dbtemp/new_db.

```
cubrid createdb -F "/dbtemp/new_db/" testdb
```

Log file directory path (-L)

The **-L** option specifies an absolute path to the directory where database log files are created. If the **-L** option is not specified, log files are created in the directory specified by the **-F** option. If neither **-F** nor **-L** option is specified, database log files are created in the current working directory.

The following example shows how to create a database named testdb in the directory /dbtemp/newdb and log files in the directory /dbtemp/db_log.

```
cubrid createdb -F "/dbtemp/new_db/" -L "/dbtemp/db_log/" testdb
```

LOB data file directory path (-B)

The **--lob-base-path** option specifies a directory where LOB data files are stored when BLOB/CLOB data is used. If the **--lob-base-path** option is not specified, LOB data files are store in *<location of database volumns created>*/lob directory.

The following example shows how to create a database named testdb in the working directory and specify /home/data1 of local file system as a location of LOB data files.

```
cubrid createdb --lob-base-path "file:/home1/data1" testdb
```

Server host name (--server-name)

The **--server-name** option enables the server of a specific database to run in the specified host when CUBRID client/server is used. The information of a host specified is stored in the **databases.txt** file. If this option is not specified, the current localhost is specified by default.

The following example shows how to create a database named testdb and register it on the host aa_host.

```
cubrid createdb --server-name aa_host testdb
```

Overwriting (-r)

The **-r** option creates a new database and overwrites an existing database if one with the same name exists. If the **-r** option is not specified, database creation is halted.

The following example shows how to create a new database named testdb and overwrite the existing database with the same name.

```
cubrid createdb -r testdb
```

Adding a database volume (--more-volume-file)

The **--more-volume-file** option creates an additional volume based on the specification contained in the file specified by the option. The volume is created in the same directory where the database is created. Instead of using this option, you can add a volume by using the **cubrid addvoldb** utility.

The following example shows how to create a database named testdb as well as an additional volume based on the specification stored in the **vol_info.txt** file.

```
cubrid createdb --more-volume-file vol_info.txt testdb
```

The following is a specification of the additional volume contained in the **vol_info.txt** file. The specification of each volume must be written on a single line.

```
#XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
# NAME volname COMMENTS volcmnts PURPOSE volpurp NPAGES volnpgs
NAME data v1 COMMENTS "data information volume" PURPOSE data NPAGES 1000
NAME data v2 COMMENTS "data information volume" PURPOSE data NPAGES 1000
NAME data v3 PURPOSE data NPAGES 1000
NAME index v1 COMMENTS "index information volume" PURPOSE index NPAGES 500
NAME temp_v1 COMMENTS "temporary information volume" PURPOSE temp NPAGES 500
NAME generic_v1 COMMENTS "generic information volume" PURPOSE generic NPAGES 500
#XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

As shown in the example, the specification of each volume consists followings.

```
NAME volname COMMENTS volcmnts PURPOSE volpurp NPAGES volnpgs
```

- *volname* : The name of the volume to be created. It must follow the UNIX file name conventions and be a simple name not including the directory path. The specification of a volume name can be omitted. If it is, the "database name to be created by the system_volume identifier" becomes the volume name.
- *volcmnts* : Comment to be written in the volume header. It contains information on the additional volume to be created. The specification of the comment on a volume can also be omitted.
- *volpurp* : It must be one of the following types: **data**, **index**, **temp**, or **generic** based on the purpose of storing volumes. The specification of the purpose of a volume can be omitted in which case the default value is **generic**.
- *volnpgs* : The number of pages of the additional volume to be created. The specification of the number of pages of the volume cannot be omitted; it must be specified.

User information file (--user-definition-file)

The **--user-definition-file** option adds users who have access to the database to be created. It adds a user based on the specification contained in the user information file specified by the parameter. Instead of using the **--user-definition-file** option, you can add a user by using the **CREATE USER** statement (for details, see [Managing USER](#)).

The following example shows how to create a database named testdb and add users to testdb based on the user information defined in the **user_info.txt** file.

```
cubrid createdb --user-definition-file user_info.txt testdb
```

The syntax of a user information file is as follows:

```
USER user_name [ groups clause | members clause ] |
groups clause:
[ GROUPS group_name [ { group_name }... ] ]
members_clause:
[ MEMBERS member_name [ { member_name... } ] ]
```

- The *user_name* is the name of the user who has access to the database. It must not include spaces.
- The **GROUPS** clause is optional. The *group_name* is the upper level group that contains the *user_name*. Here, the *group_name* can be multiply specified and must be defined as **USER** in advance.
- The **MEMBERS** clause is optional. The *member_name* is the name of the lower level member that belongs to the *user_name*. Here, the *member_name* can be multiply specified and must be defined as **USER** in advance.

Comments can be used in a user information file. A comment line must begin with a consecutive hyphen lines (--). Blank lines are ignored.

The following example shows a user information in which grandeur and sonata are included in sedan group, tuscan is included in suv group, and i30 is included in hatchback group. The name of the user information file is **user_info.txt**.

```
--
-- Example 1 of a user information file
```



```
--
USER sedan
USER suv
USER hatchback
USER grandeur GROUPS sedan
USER sonata GROUPS sedan
USER tuscan GROUPS suv
USER i30 GROUPS hatchback
```

The following example shows a file that has the same user relationship information as the file above. The difference is that the **MEMBERS** statement is used in the file below.

```
--
-- Example 2 of a user information file
--
USER grandeur
USER sonata
USER tuscan
USER i30
USER sedan MEMBERS sonata grandeur
USER suv MEMBERS tuscan
USER hatchback MEMBERS i30
```

File where CSQL statements are stored (--csql-initialization-file)

The **--csql-initialization-file** option executes an SQL statement on the database to be created by using the CSQL Interpreter. A schema can be created based on the SQL statement contained in the file specified by the parameter.

The following example shows how to create a database named testdb and execute the SQL statement defined in table_schema.sql through the CSQL Interpreter.

```
cubrid createdb --csql-initialization-file table_schema.sql testdb
```

Storing output messages to a file (-o)

The **-o** option stores messages related to the database creation to the file given as a parameter. The file is created in the same directory where the database was created. If the **-o** option is not specified, messages are displayed on the console screen. The **-o** option allows you to use information on the creation of a certain database by storing messages, generated during the database creation, to a specified file.

The following example shows how to create a database named testdb and store the output of the utility to the **db_output** file instead of displaying it on the console screen.

```
cubrid createdb -o db_output testdb
```

Verbose output (-v)

The **-v** option displays all information on the database creation operation onto the screen. Like the **-o** option, this option is useful in checking information related to the creation of a specific database. Therefore, if you specify the **-v** option together with the **-o** option, you can store the output messages in the file given as a parameter; the messages contain the operation information about the **cubrid createdb** utility and database creation process.

The following example shows how to create a database named testdb and display detailed information on the operation onto the screen.

```
cubrid createdb -v testdb
```

Remark

temp_file_max_size_in_pages is a parameter that configures the maximum number of pages assigned to store the temporary temp volume - used for complicated queries or storing arrays - on the disk.

While the default value is **-1**, the temporary temp volume may be increased up to the amount of extra space on the disk specified by the **temp_volume_path** parameter. If the value is 0, the temporary temp volume cannot be created. In this case, the permanent temp volume should be added by using the [cubrid addvoldb](#) utility.

For the efficient management of the volume, it is recommended to add a volume for each usage. By using the [cubrid spacedb](#) utility, you can check the remaining space of each volume. By using the [cubrid addvoldb](#) utility, you can add

more volumes as needed while managing the database. When adding a volume while managing the database, you are advised to do so when there is less system load. Once the assigned volume for a usage is completely in use, a generic volume will be created, so it is suggested to add extra volume for a usage that is expected to require more space.

Next, we will look at how to add volumes for **data**, **index**, and **temp** by creating the database and separating the volume usage.

```
cubrid createdb --db-volume-size=512M --log-volume-size=256M cubriddb
cubrid addvoldb -p data -n cubriddb_DATA01 --db-volume-size=512M cubriddb
cubrid addvoldb -p data -n cubriddb_DATA02 --db-volume-size=512M cubriddb
cubrid addvoldb -p index -n cubriddb_INDEX01 cubriddb --db-volume-size=512M cubriddb
cubrid addvoldb -p temp -n cubriddb_TEMP01 cubriddb --db-volume-size=512M cubriddb
```

Adding Database Volume

Description

Adds database volume.

Syntax

```
cubrid addvoldb options database_name
options :
[--db-volume-size=size] [-n |--volume-name=name] [-F |--file-path=path] [--comment=comment] [-p |--purpose] [-S |--SA-mode |-C |--CS-mode]
```

- **cubrid** : An integrated utility for CUBRID service and database management.
- **addvoldb** : A command that adds a specified number of pages of the new volume to a specified database.
- *options* : A short option starts with a single dash (-) while a full name option starts with a double dash (--).
- *database_name* : Specifies the name of the database to which a volume is to be added without including the path name to the directory where the database is to be created.

Options

The following table shows options available with the **cubrid addvoldb** utility.

Option	Description
--db-volume-size	Specifies the database volume size in bytes. Default value : A value of <i>db_volume_size</i> , the system parameter
-n --volume-name	Specifies the name of the database volume to be added. Default value : A value in the format of <i>dbname_number</i> , configured by the system
-F --file-path	Specifies the directory path where the database volume to be added will be created. Default value : A value of <i>volume_extension_path</i> , the system parameter
--comment	Inserts a comment about the database volume to be added.
-p --purpose	Specifies the purpose of the database volume to be added. Default value : Generic volume
-S --SA-mode	Adds the database volume in standalone mode.
-C --CS-mode	Adds the database volume in client/server mode.

Size of the extended volume (**--db-volume-size**)

--db-volume-size is an option that specifies the size of the volume to be added to a specified database. If the **--db-volume-size** option is omitted, the value of the system parameter **db_volume_size** is used by default. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB), and terabytes (TB) respectively. If you omit the unit, bytes will be applied.

The following example shows how to add a volume for which 256 MB are assigned to the testdb database.

```
cubrid addvoldb -p data --db-volume-size=256M testdb
```

Name of the extended volume (-n)

-n is an option that specifies the name of the volume to be added to a specified database. The volume name must follow the file name protocol of the operating system and be a simple one without including the directory path or spaces. If the **-n** option is omitted, the name of the volume to be added is configured by the system automatically as "database name_volume identifier." For example, if the database name is testdb, the volume name testdb_x001 is automatically configured.

The following example shows how to add a volume for which 256 MB are assigned to the testdb database in standalone mode. The volume name testdb_v1 will be created.

```
cubrid addvoldb -S -n testdb_v1 --db-volume-size=256M testdb
```

Path of the extended volume (-F)

The **-F** option is used to specify the directory path where the volume to be added will be stored. If the **-F** option is omitted, the value of the system parameter **volume_extension_path** is used by default.

The following example shows how to add a volume for which 256 MB are assigned to the testdb database in standalone mode. The added volume is created in the /dbtemp/addvol directory. Because the **-n** option is not specified for the volume name, the volume name testdb_x001 will be created.

```
cubrid addvoldb -S -F /dbtemp/addvol/ --db-volume-size=256M testdb
```

Comment about the added volume (--comment)

The **--comment** option is used to facilitate to retrieve information on the added volume by adding such information in the form of comments. It is recommended that the contents of a comment include the name of **DBA** who adds the volume, or the purpose of adding the volume. The comment must be enclosed in double quotes.

The following example shows how to add a volume for which 256 MB are assigned to the testdb database in standalone mode and inserts a comment about the volume.

```
cubrid addvoldb -S --comment "data volume added_cheolsoo kim" --db-volume-size=256M testdb
```

Purpose of the volume (-p)

The **-p** option is used to specify the purpose of the volume to be added. The reason for specifying the purpose of the volume is to improve the I/O performance by storing volumes separately on different disk drives according to their purpose. Parameter values that can be used for the **-p** option are **data**, **index**, **temp** and **generic**. The default value is **generic**. For the purpose of each volume, see "[Database Volume Structure](#)."

The following example shows how to add a volume for which 256 MB are assigned to the testdb database in standalone mode.

```
cubrid addvoldb -S -p index --db-volume-size=256M testdb
```

Standalone mode (-S)

The **-S** option is used to access the database in standalone mode without running the server process. This option has no parameter. If the **-S** option is not specified, the system assumes to be in client/server mode.

```
cubrid addvoldb -S --db-volume-size=256M testdb
```

Client/server mode (-C)

The **-C** option is used to access the database in client/server mode by running the server and the client separately. There is no parameter. Even when the **-C** option is not specified, the system assumes to be in client/server mode by default.

```
cubrid addvoldb -C --db-volume-size=256M testdb
```

Example

The following example shows how to create a database, classify volume usage, and add volumes such as **data**, **index**, and **temp**.

```

cubrid createdb --db-volume-size=512M --log-volume-size=256M cubriddb
cubrid addvoldb -p data -n cubriddb_DATA01 --db-volume-size=512M cubriddb
cubrid addvoldb -p data -n cubriddb_DATA02 --db-volume-size=512M cubriddb
cubrid addvoldb -p index -n cubriddb_INDEX01 cubriddb --db-volume-size=512M cubriddb
cubrid addvoldb -p temp -n cubriddb_TEMP01 cubriddb --db-volume-size=512M cubriddb

```

Deleting Database

Description

The **cubrid deletedb** utility is used to delete a database. You must use the **cubrid deletedb** utility to delete a database, instead of using the file deletion commands of the operating system; a database consists of a few interdependent files. The **cubrid deletedb** utility also deletes the information on the database from the database location file (**databases.txt**). The **cubrid deletedb** utility must be run offline, that is, in standalone mode when nobody is using the database.

Syntax

```

cubrid deletedb options database_name
options : [{"-o|--output-file=} file] [{"-d|--delete-backup"}]

```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **deletedb** : A command to delete a database, its related data, logs and all backup files. It can be executed successfully only when the database is in a stopped state.
- *options* : **-o** and **-d** options are provided.
- *database_name* : Specifies the name of the database to be deleted without including the path name.

Options

Storing output messages (-o or --output-file)

The following example shows how to delete testdb and write messages to the file specified by using the **-o** option.

```

cubrid deletedb -o deleted_db.out testdb

```

The **cubrid deletedb** utility also deletes the database information contained in the database location file (**databases.txt**). The following message is returned if you enter a utility that tries to delete a non-existing database.

```

cubrid deletedb testdb
Database "testdb" is unknown, or the file "databases.txt" cannot be accessed.

```

Deleting backup files simultaneously (-d or --delete-backup)

The following example shows how to delete testdb and its backup volumes and backup information files simultaneously by using the **-d** option. If the **-d** option is not specified, backup volume and backup information files are not deleted.

```

cubrid deletedb -d testdb

```

Renaming Database

Description

The **cubrid renamedb** utility renames a database. The names of information volumes, log volumes and control files are also renamed to conform to the new database one.

The **cubrid alterdbhost** utility configures or changes the host name of the specified database. It changes the host name configuration in the **databases.txt** file.

Syntax

```

cubrid renamedb options src_database_name dest_database_name
options : [{"-E | --extended-volumn-path=} path] [{"-i | --control-file=} file] [{"-d | --delete-backup"}]

```

- **cubrid** : An integrated utility for the CUBRID service and database management.

- **renamedb** : A command that changes the existing name of a database to a new one. It executes successfully only when the database is in a stopped state. The names of related information volumes, log volumes and control files are also changed to new ones accordingly.
- *options* : The **-E**, **-i** and **-d** options are supported. For details about each option, see its description and the examples.
- *src_database_name* : The name of the existing database to be renamed. The path name to the directory where the database is to be created must not be included.
- *dest_database_name* : The new name of the database. It must not be the same as that of an existing database. The path name to the directory where the database is to be created must not be included.

Options

Saving the renamed extended volume to a new directory (-E or --extended-volume-path)

The following example shows how to rename an extended volume created in a specific directory path (e.g. /dbtemp/addvol/) with a **-E** option, and then moves the volume to a new directory. The **-E** option is used to specify a new directory path (e.g. /dbtemp/newaddvols/) where the renamed extended volume will be moved. If the **-E** option is not specified, the extended volume is only renamed in the existing path without being moved. If a directory path outside the disk partition of the existing database volume or an invalid one is specified, the rename operation is not executed. This option cannot be used together with the **-i** option.

```
cubrid renamedb -E /dbtemp/newaddvols/ testdb testdb_1
```

Specifying the input file where the directory information is stored (-i or --control-file)

The following example shows how to specify an input file in which directory information is stored to change all database name of volumes or files and assign different directory at once. To perform this work, the **-i** option is used. The **-i** option cannot be used together with the **-E** option.

```
cubrid renamedb -i rename_path testdb testdb_1
```

The followings are the syntax and example of a file that contains the name of each volume, the current directory path and the directory path where renamed volumes will be stored.

```
valid source_fullvolname dest_fullvolname
```

- *valid* : An integer that is used to identify each volume. It can be checked in the database volume control file (database_name_vinf).
- *source_fullvolname* : The current directory path to each volume.
- *dest_fullvolname* : The target directory path where renamed volumes will be moved. If the target directory path is invalid, the database rename operation is not executed.

```
-5 /home1/user/testdb_vinf /home1/CUBRID/databases/testdb_1_vinf
-4 /home1/user/testdb_lginf /home1/CUBRID/databases/testdb_1_lginf
-3 /home1/user/testdb_bkvinf /home1/CUBRID/databases/testdb_1_bkvinf
-2 /home1/user/testdb_lgat /home1/CUBRID/databases/testdb_1_lgat
0 /home1/user/testdb /home1/CUBRID/databases/testdb_1
1 /home1/user/backup/testdb_x001/home1/CUBRID/databases/backup/testdb_1_x001
```

Deleting and renaming backup files simultaneously (-d or --delete-backup)

By using the **-d** option, the following example shows how to rename the testdb database and at once forcefully delete all backup volumes and backup information files that are in the same location as testdb. Note that you cannot use the backup files with the old names once the database is renamed. If the **-d** option is not specified, backup volumes and backup information files are not deleted.

```
cubrid renamedb -d testdb testdb_1
```

Renaming Database Host

Decription

The **cubrid alterdbhost** utility sets or changes the host name of the specified database. It changes the host name set in the **databases.txt** file.

Syntax

```
cubrid alterdbhost [option] database_name  
option : [ {-h | --host=} host_name ]
```

- **cubrid**: An integrated utility for the CUBRID service and database management
- **alterdbhost**: A command used to change the host name of the current database
- *option*: Specifies the host name to be changed after **-h** or **--host=**. When this option is omitted, specifies the host name to localhost.

Copying/Moving Database

Description

The **cubrid copydb** utility copy or move a database to another location. As arguments, source and target name of database must be given. A target database name must be different from a source database name. When the target name argument is specified, the location of target database name is registered in the **databases.txt** file. The **cubrid copydb** utility can be executed only offline (that is, state of a source database stop).

Syntax

```
cubrid copydb [options] src-database-name dest-database-name  
options : [--server-name=}host] [{"-F |--file-path=}database_path] [{"-L |--log-path=}log path] [{"-B |--lob-base-path=}lob file path] [{"-E |--extended-volume-path=}path] [{"-i |--control-file=}FILE] [{"-r|--replace}] [{"-d|--delete-source}] [--copy-lob-path]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **copydb** : A command that copy or move a database from one to another location.
- *options* : For details about each option, see its description and the examples. If options are omitted, a target database is copied into the same directory of a source database.
- *src-database-name* : The names of source and target databases to be copied or moved.
- *dest-database-name* : A new (target) database name.

Options

Registering a host name (--server-name)

The following example shows how to specify a host name of new database. The host name is registered in the **databases.txt** file. If this option is omitted, a local host is registered.

```
cubrid copydb --server-name=cub_server1 demodb new_demodb
```

Storing a new database volume in a specific directory (-F or --file-path)

The following example shows how to specify a specific directory path where a new database volume is stored with an **-F** option. It represents specifying an absolute path. If the specified directory does not exist, an error is displayed. If this option is omitted, a new database volume is created in the current working directory. And this information is specified in **vol-path** of the **databases.txt** file.

```
cubrid copydb -F /home/usr/CUBRID/databases demodb new_demodb
```

Storing a new database log volume in a specific directory (-L or --log-path)

The following example shows how to specify a specific directory path where a new database volume is stored with an **-L** option. It represents specifying an absolute path. If the specified directory does not exist, an error is displayed. If this option is omitted, a new database volume is created in the current working directory. And this information is specified in **log-path** of the **databases.txt** file.

```
cubrid copydb -L /home/usr/CUBRID/databases/logs demodb new_demodb
```

Storing a new database extended volume in a specific directory (-E or --extended-volume-path)

The following example shows how to specify a specific directory path where a new database extended volume is stored with an **-E**. If this option is omitted, a new database extended volume is created in the location of a new database volume or in the registered path of controlling file. The **-i** option cannot be used with this option.

```
cubrid copydb -E home/usr/CUBRID/databases/extvols demodb new_demodb
```

Specifying an input file where directory path information is stored (-i or --control file)

The following example shows how to specify an input file where a new directory path information and a source volume are stored to copy or move multiple volumes into a different directory, respectively. The **-i** option cannot be used with the **-E** option. An input file named `copy_path` is specified in the example below.

```
cubrid copydb -i copy_path demodb new_demodb
```

The following is an example of input file that contains each volume name, current directory path, and new directory and volume names.

```
# valid source fullvolname dest fullvolname
0 /usr/databases/demodb /drive1/usr/databases/new_demodb
1 /usr/databases/demodb data1 /drive1/usr/databases/new_demodb new_data1
2 /usr/databases/ext/demodb index1 /drive2/usr/databases/new_demodb new_index1
3 /usr/databases/ext/demodb index2 /drive2/usr/databases/new_demodb new_index2
```

- *valid* : An integer that is used to identify each volume. It can be checked in the database volume control file (**database_name_vinf**).
- *source_fullvolname* : The current directory path to each source database volume.
- *dest_fullvolname* : The target directory path where new volumes will be stored. You should specify a valid path.

Overwriting if same database exists (-r or --replace)

If the **-r** option is specified, a new database name overwrites the existing database name if it is identical, instead of outputting an error.

```
cubrid copydb -r -F /home/usr/CUBRID/databases demodb new_demodb
```

Deleting a source database if it is copied (-d or --delete-source)

If the **-d** option is specified, a source database is deleted after the database is copied. This execution brings the same result as executing **cubrid deletedb** utility after copying a database. Note that if a source database contains LOB data, LOB file directory path of a source database is copied into a new database and it is registered in the **lob-base-path** of the **databases.txt** file.

```
cubrid copydb -d -opyhome/usr/CUBRID/databases demodb new_demodb
```

Copying LOB file directory (--copy-lob-path)

If the **--copy-lob-path** option is specified, a new directory path for LOB files is created and a source database is copied into a new directory path. If this option is omitted, the directory path is not created. Therefore, the **lob-base-path** of the **databases.txt** file should be modified separately. This option cannot be used with the **-B** option.

```
cubrid copydb --copy-lob-path demodb new_demodb
```

Copying LOB file directory simultaneously with specifying it (-B or --lob-base-path)

If the **-B** option is specified, a specified directory is specified as for LOB files of a new database and a source database is copied. This option cannot be used with the **--copy-lob-path** option.

```
cubrid copydb -B /home/usr/CUBRID/databases/new_lob demodb new_demodb
```

Registering Database

Description

The **cubrid installdb** utility is used to register the information of a newly installed database to **databases.txt**, which stores database location information. The execution of this utility does not affect the operation of the database to be registered.

Syntax

```
cubrid installdb options database_name
options : [--server-name=host] [-F | --file-path=database_path] [ {-L | --log-path=log_path }
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **installdb** : A command that registers the information of a moved or copied database to **databases.txt**.
- *options* : **--server-name**, **-F**, **-L** options are available. For details on each option, see the option description and example. If no option is used with a command, the command must be executed in the directory where the corresponding database exists.
- *database_name* : The name of database to be registered to **databases.txt**.

Options

Registering a host name (--server-name)

The following example shows how to register the server host information of a database to **databases.txt** with a specific host name. If this option is not specified, the current host information is registered.

```
cubrid installdb --server-name=cub_server1 testdb
```

Registering the directory path of a database volume (-F or --file-path)

The following example shows how to register the directory path of a database volume to **databases.txt** by using the **-F** option. If this option is not specified, the path of a current directory is registered as default.

```
cubrid installdb -F /home/cubrid/CUBRID/databases/testdb testdb
```

Registering the directory path of a database log volume (-L or --log-path)

The following example shows how to register the directory path of a database log volume to **databases.txt** by using the **-L** option. If this option is not specified, the directory path of a volume is registered.

```
cubrid installdb -L /home/cubrid/CUBRID/databases/logs/testdb testdb
```

Checking Used Space

Description

The **cubrid spacedb** utility is used to check how much space of database volumes is being used. It shows a brief description of all permanent data volumes in the database. Information returned by the **cubrid spacedb** utility includes the ID, name, purpose and total/free space of each volume. You can also check the total number of volumes and used/unused database pages.

Syntax

```
cubrid spacedb options database name
options : [-o|--output-file=file] [-S|--SA-mode|-C|--CS-mode] [--size-unit=PAGE|M|G|T|H]
[-s|--summarize]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **spacedb** : A command that checks the space in the database. It executes successfully only when the database is in a stopped state.
- *options* : The **-o**, **-S**, **-C** **--size-unit**, and **-s** options are supported. For details about each option, refer to its description and the examples.
- *database_name* : The name of the database whose space is to be checked. The path-name to the directory where the database is to be created must not be included.

Options

Storing output messages to a file (-o)

The following syntax shows how to store the result of checking the space information of testdb to a file named *db_output*.

```
cubrid spacedb -o db_output testdb
```

Executing in stand-alone mode (-S or --SA-mode)

The **-S** option is used to access a database in standalone, which means it works without processing server; it does not have an argument. If **-S** is not specified, the system recognizes that a database is running in client/server mode.

```
cubrid spacedb --SA-mode testdb
```

Executing in client/server mode (-C or --CS-mode)

The **-C** option is used to access a database in client/server mode, which means it works in client/server process respectively; it does not have an argument. If **-C** is not specified, the system recognize that a database is running in client/server mode by default.

```
cubrid spacedb --CS-mode testdb
```

Outputting in specified size unit (--size-unit)

The **--size-unit** option is used to specify the size unit of the space information of the database to be one of PAGE, M(MB), G(GB), T(TB), H(print-friendly). The default value is **H**. If you set the value to H, the unit is automatically determined as follows: M if 1 MB = DB size < 1024 MB, G if 1 GB = DB size < 1024 GB.

```
cubrid spacedb --size_unit=M testdb
cubrid spacedb --size_unit=H testdb
```

Outputs total pages, used pages, free pages by volume usage (-s or --summarize)

Aggregates total_pages, used_pages and free_pages by DATA, INDEX, GENERIC, TEMP and TEMP TEMP, and outputs it.

```
cubrid spacedb -s testdb
```

Compacting Used Space

Description

The **cubrid compactdb** utility is used to secure unused space of the database volume. In case the database server is not running (offline), you can perform the job in stand-alone mode. In case the database server is running, you can perform it in client-server mode.

The **cubrid compactdb** utility secures the space being taken by OIDs of deleted objects and by class changes. When an object is deleted, the space taken by its OID is not immediately freed because there might be other objects that refer to the deleted one. Reference to the object deleted during compacting is displayed as **NULL**, which means this can be reused by OIDs.

Syntax

```
cubrid compactdb options database name [class name], class name2,...]
options : [-v | --verbose] [-S|--SA-mode | -C | --CS-mode]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **compactdb** : A command that compacts the space of the database so that OIDs assigned to deleted data can be reused.
- **options** : The **-v**, **-S** and **-C** options are supported. Options (**-I**, **-i**, **-c**, **-d**, **-p**) that is applied in client/server mode only.
- **database_name** : The name of the database whose space is to be compacted. The path name to the directory where the database is to be created must not be included.
- **class_name_list** : You can specify the list of tables names that you want to compact space after a database name; the **-i** option cannot be used together. It is used in client/server mode only.

Options

Displaying detailed messages during execution (-v)

You can output messages that shows which class is currently being compacted and how many instances have been processed for the class by using the `-v` option.

```
cubrid compactdb -v testdb
```

Executing in stand-alone mode (-S or --SA mode)

The `-S` option is specified to compact used space in stand-alone mode while database server is not running; no argument is specified. If the `-S` option is not specified, system recognizes that the job is executed in client/server mode.

```
cubrid compactdb --SA-mode testdb
```

Executing in client/server mode (C or --CS mode)

The `-C` option is specified to compact used space in client/server mode while database server is running; no argument is specified. Even though this option is omitted, system recognizes that the job is executed in client/server mode. The following options can be used in client/server mode only.

- `-i, --input-class-file=FILE` : You can specify an input file name that contains the table name with this option. Write one table name in a single line; invalid table name is ignored. Note that you cannot specify the list of the table names after a database name in case of you use this option.
- `-p, --pages-committed-once=NUMBER` : You can specify the number of maximum pages that can be committed once with this option. The default value is 10, the minimum value is 1, and the maximum value is 10. The less option value is specified, the more concurrency is enhanced because the value for class/instance lock is small; however, it causes slowdown on operation, and vice versa.
- `-d, --delete-old-repr` : You can delete an existing table representation from catalog with this option.
- `-I, --Instance-lock-timeout` : You can specify a value of instance lock timeout with this option. The default value is 2 (seconds), the minimum value is 1, and the maximum value is 10. The less option value is specified, the more operation speeds up. However, the number of instances that can be processed becomes smaller, and vice versa.
- `-c, --class-lock-timeout` : You can specify a value of instance lock timeout with this option. The default value is 10 (seconds), the minimum value is 1, and the maximum value is 10. The less option value is specified, the more operation speeds up. However, the number of tables that can be processed becomes smaller, and vice versa.

```
cubrid compactdb --CS-mode -p 10 testdb tbl1, tbl2, tbl5
```

Updating Statistics

Description

Updates statistical information such as the number of objects, the number of pages to access, and the distribution of attribute values.

Syntax

```
cubrid optimizedb options database_name
options : [{-n|--class-name=} name]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **optimizedb** : Updates the statistics information, which is used for cost-based query optimization of the database. If the option is specified, only the information of the specified class is updated.
- *options* : The `-n` option is supported.
- *database_name* : The name of the database whose cost-based query optimization statistics are to be updated.

Options

Updating the query statistics of the target database

The following example shows how to update the query statistics information of all classes in the database.

```
cubrid optimizedb testdb
```

Updating the query statistics of a specific class in the database (-n or --class-name)

The following example shows how to update the query statistics information of the given class by using the **-n** option.

```
cubrid optimizedb -n event_table testdb
```

Outputting Statistics Information of Server

Description

The cubrid statdump utility checks statistics information processed by the CUBRID database server. The statistics information mainly consists of the followings: File I/O, Page buffer, Logs, Transactions, Concurrency/Lock, Index, and Network request

Note that you must specify the parameter **communication_histogram** to **yes** in the **cubrid.conf** before executing the utility. You can also check statistics information of server with session commands (**;****h on**) in the CSQL.

Syntax

```
cubrid statdump options database name
options : [{"-o |--output-file=}file_name] [{"-i |--interval=}secs] [{"-c|--cumulative}] [{"-s |--substr=}sub_string]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **installdb** : A command that dumps the statistics information on the database server execution.
- *options* : **--o**, **-i**, **-c**, and **-s** options are available.
- *database_name* : The name of database which has the statistics data to be dumped.

Options

Outputting statistics information periodically (-i or --interval)

```
cubrid statdump -i 5 testdb
```

```
Thu April 07 23:10:08 KST 2011
```

```
*** SERVER EXECUTION STATISTICS ***
Num file creates                =          0
Num file removes                =          0
Num_file_ioreads                =          0
Num_file_iowrites               =          0
Num file iosynches              =          0
Num data page fetches           =          0
Num data page dirties           =          0
Num_data_page_ioreads           =          0
Num_data_page_iowrites          =          0
Num_data_page_victims           =          0
Num_data_page_iowrites for replacement =          0
Num log page ioreads            =          0
Num log page iowrites           =          0
Num log append records          =          0
Num_log_archives                =          0
Num_log_checkpoints             =          0
Num log wals                    =          0
Num page locks acquired         =          0
Num object locks acquired       =          0
Num_page_locks_converted        =          0
Num_object_locks_converted      =          0
Num page locks re-requested     =          0
Num object locks re-requested   =          0
Num page locks waits            =          0
Num object locks waits          =          0
Num_tran_commits                =          0
Num_tran_rollbacks              =          0
Num tran savepoints             =          0
Num tran start topops           =          0
Num tran end topops             =          0
Num_tran_interrupts            =          0
```

```

Num btree inserts          =          0
Num btree deletes         =          0
Num btree_updates         =          0
Num btree_covered        =          0
Num btree_noncovered      =          0
Num btree_resumes        =          0
Num query_selects        =          0
Num query_inserts        =          0
Num_query_deletes        =          0
Num_query_updates        =          0
Num_query_sscans         =          0
Num_query_iscans         =          0
Num_query_lscans         =          0
Num_query_setscans       =          0
Num_query_methscans      =          0
Num_query_nljoins        =          0
Num_query_mjoins         =          0
Num query_objfetches     =          0
Num network_requests     =          1
Num_adaptive_flush_pages =          0
Num_adaptive_flush_log_pages =          0
Num adaptive_flush_max_pages =          900

*** OTHER STATISTICS ***
Data_page_buffer_hit_ratio =          0.00
    
```

Category of Statistics Information

Category	Item	Description
File I/O	Num_file_removes	The number of files removed
	Num_file_creates	The number of files created
	Num_file_ioreads	The number of files read
	Num_file_iowrites	The number of files stored
	Num_file_iosynches	The number of file synchronization
Page buffer	Num_data_page_fetches	The number of pages fetched
	Num_data_page_dirties	The number of duty pages
	Num_data_page_ioreads	The number of pages read
	Num_data_page_iowrites	The number of pages stored
	Num_data_page_victims	The number specifying the victim data to be flushed from the data page to the disk
	Num_data_page_iowrites_for_replacement	The number of the written data pages specified as victim
	Num_adaptive_flush_pages	The number of data pages flushed from the data buffer to the disk
Logs	Num_adaptive_flush_log_pages	The number of log pages flushed from the log buffer to the disk
	Num_adaptive_flush_max_pages	The maximum number of pages allowed to flush from data and the log buffer to the disk
	Num_log_page_ioreads	The number of log pages read
	Num_log_page_iowrites	The number of log pages stored
	Num_log_append_records	The number of log records appended

	Num_log_archives	The number of logs archived
	Num_log_checkpoints	The number of checkpoints
	Num_log_wals	Not used
Transactions	Num_tran_commits	The number of commits
	Num_tran_rollbacks	The number of rollbacks
	Num_tran_savepoints	The number of savepoints
	Num_tran_start_topops	The number of top operations started
	Num_tran_end_topops	The number of top operations stopped
	Num_tran_interrupts	The number of interruptions
Concurrency/lock	Num_page_locks_acquired	The number of locked pages acquired
	Num_object_locks_acquired	The number of locked objects acquired
	Num_page_locks_converted	The number of locked pages converted
	Num_object_locks_converted	The number of locked objects converted
	Num_page_locks_re-requested	The number of locked pages requested
	Num_object_locks_re-requested	The number of locked objects requested
	Num_page_locks_waits	The number of locked pages waited
	Num_object_locks_waits	The number of locked objects waited
Index	Num_btree_inserts	The number of nodes inserted
	Num_btree_deletes	The number of nodes deleted
	Num_btree_updates	The number of nodes updated
	Num_btree_covered	The number of cases in which an index includes all data upon query execution
	Num_btree_noncovered	The number of cases in which an index includes some or no data upon query execution
	Num_btree_resumes	The exceeding number of index scan specified in index_scan_oid_buffer_pages
Query Related	Num_query_selects	The number of SELECT query execution
	Num_query_inserts	The number of INSERT query execution
	Num_query_deletes	The number of DELETE query execution
	Num_query_updates	The number of UPDATE query execution
	Num_query_sscans	The number of sequential scans (full scan)
	Num_query_iscans	The number of index scans
	Num_query_lscans	The number of LIST scans
	Num_query_setscans	The number of SET scans
	Num_query_methscans	The number of METHOD scans
	Num_query_nljoins	The number of nested loop joins
	Num_query_mjoins	The number of parallel joins
	Num_query_objfetches	The number of fetch objects

Network request related	Num_network_requests	The number of network requested
	Data_page_buffer_hit_ratio	Hit Ratio of page buffers (Num_data_page_fetches - Num_data_page_ioreads)*100 / Num_data_page_fetches

Saving statistics information to a file (-o or --output-file)

The `-o` options is used to store statistics information of server processing for the database to a specified file.

```
cubrid statdump -o statdump.log testdb
```

Displays the accumulated operation statistics information (-c or --cumulative)

You can display the accumulated operation statistics information of the target database server by using the `-c` option. By combining this with the `-i` option, you can check the operation statistics information at a specified interval.

```
cubrid statdump -i 5 -c testdb
```

Displays statistics that includes specified string (-s or --substr)

You can display statistics about items of which name include the specified string by using `-s` option.

The following example shows how to display statistics about items of which name include "data".

```
cubrid statdump -s data testdb

*** SERVER EXECUTION STATISTICS ***
Num data page fetches           =          135
Num data page dirties           =           0
Num data page ioreads           =           0
Num_data_page_iowrites          =           0
Num_data_page_victims           =           0
Num data page iowrites for replacement =           0

*** OTHER STATISTICS ***
Data_page_buffer_hit_ratio      =        100.00
```

Note Each status information consists of 64-bit INTEGER data and the corresponding statistics information can be lost if the accumulated value exceeds the limit.

Checking Lock Status

Description

The `cubrid lockdb` utility is used to check the information on the lock being used by the current transaction in the database.

Syntax

```
cubrid lockdb options database name
options : [{-o|--output-file=} file ]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **lockdb** : A command used to check the information on the lock being used by the current transaction in the database.
- *options* : The `-o` option is supported.
- *database_name* : The name of the database where lock information of the current transaction is to be checked.

Options

Displaying lock information on a screen

The following example shows how to display lock information of the testdb database on a screen without any option.

```
cubrid lockdb testdb
```

Displaying lock information to the specified file (-o)

The following example shows how to display lock information of the testdb database as a output.txt by using the **-o** option.

```
cubrid lockdb -o output.txt testdb
```

Output Contents

The output contents of **cubrid lockdb** are divided into three logical sections:

- Server lock settings
- Clients that are accessing the database
- The contents of an object lock table

Server lock settings

The first section of the output of **cubrid lockdb** is the database lock settings.

```
*** Lock Table Dump ***
Lock Escalation at = 100000, Run Deadlock interval = 0
```

The lock escalation level is 100,000 records, and the interval to detect deadlock is set to 0 seconds (For a description of the related system parameters, **lock_escalation** and **deadlock_detection_interval**, see [Concurrency/Lock Parameters](#)).

Clients that are accessing the database

The second section of the output of **cubrid lockdb** includes information on all clients that are connected to the database. This includes the transaction index, program name, user ID, host name, process ID, isolation level and lock timeout settings of each client.

```
Transaction (index 1, csql, dba@cubriddb|12854)
Isolation READ COMMITTED CLASSES AND READ UNCOMMITTED INSTANCES
Timeout_period -1
```

Here, the transaction index is 1, the program name is csql, the user ID is dba, the host name is cubriddb, the client process identifier is 12854, the isolation level is READ COMMITTED CLASSES AND READ UNCOMMITTED INSTANCES, and the lock timeout is unlimited.

A client for which transaction index is 0 is the internal system transaction. It can obtain the lock at a specific time, such as the processing of a checkpoint by a database. In most cases, however, this transaction will not obtain any locks.

Because **cubrid lockdb** utility accesses the database to obtain the lock information, the **cubrid lockdb** is an independent client and will be output as such.

Object lock table

The third section of the output of the **cubrid lockdb** includes the contents of the object lock table. It shows which client has the lock for which object in which mode, and which client is waiting for which object in which mode. The first part of the result of the object lock table shows how many objects are locked.

```
Object lock Table:
Current number of objects which are locked = 2001
```

cubrid lockdb outputs the OID, object type and table name of each object that obtained lock. In addition, it outputs the number of transactions that hold lock for the object (Num holders), the number of transactions (Num blocked-holders) that hold lock but are blocked since it could not convert the lock to the upper lock (e.g., conversion from U_LOCK to

X_LOCK), and the number of different transactions that are waiting for the lock of the object (Num waiters). It also outputs the list of client transactions that hold lock, blocked client transactions and waiting client transactions.

The example below shows an object in which the object type is an instance of a class, or record that will be blocked, because the OID(2| 50| 1) object that has S_LOCK for transaction 1 and S_LOCK for transaction 2 cannot be converted into X_LOCK. It also shows that transaction 3 is blocked because transaction 2 is waiting for X_LOCK even when transaction 3 is waiting for S_LOCK.

```
OID = 2| 50| 1
Object type: instance of class ( 0| 62| 5) = athlete
Num holders = 1, Num blocked-holders= 1, Num waiters = 1
LOCK HOLDERS :
  Tran_index = 2, Granted mode = S LOCK, Count = 1
BLOCKED LOCK HOLDERS :
  Tran_index = 1, Granted_mode = U_LOCK, Count = 3
  Blocked mode = X LOCK
  Start waiting at = Fri May 3 14:44:31 2002
  Wait for_nsecs = -1
LOCK WAITERS :
  Tran_index = 3, Blocked_mode = S_LOCK
  Start_waiting_at = Fri May 3 14:45:14 2002
  Wait_for_nsecs = -1
```

It outputs the lock information on the index of the table when the object type is the Index key of class (index key).

```
OID = -662| 572|-32512
Object type: Index key of class ( 0| 319| 10) = athlete.
Index name: pk athlete code
Total mode of holders = NX LOCK, Total mode of waiters = NULL LOCK.
Num holders= 1, Num blocked-holders= 0, Num waiters= 0
LOCK HOLDERS:
  Tran_index = 1, Granted_mode = NX_LOCK, Count = 1
```

Granted_mode refers to the mode of the obtained lock, and Blocked_mode refers to the mode of the blocked lock.

Starting_waiting_at refers to the time at which the lock was requested, and Wait_for_nsecs refers to the waiting time of the lock. The value of Wait_for_nsecs is determined by lock_timeout_in_secs, a system parameter.

When the object type is a class (table), Nsubgranules is displayed, which is the sum of the record locks and the key locks obtained by a specific transaction in the table.

```
OID = 0| 62| 5
Object type: Class = athlete
Num holders = 2, Num blocked-holders= 0, Num waiters= 0
LOCK HOLDERS:
Tran_index = 3, Granted mode = IS LOCK, Count = 2, Nsubgranules = 0
Tran_index = 1, Granted_mode = IX_LOCK, Count = 3, Nsubgranules = 1
Tran_index = 2, Granted_mode = IS_LOCK, Count = 2, Nsubgranules = 1
```

Checking Database Consistency

Description

The **cuprid checkdb** utility is used to check the consistency of a database. You can use **cuprid checkdb** to identify data structures that are different from indexes by checking the internal physical consistency of the data and log volumes. If the **cuprid checkdb** utility reveals any inconsistencies, you must try automatic repair by using the **--repair** option.

Syntax

```
cuprid checkdb options database name [class name1 class name2 ...]
options : [-S|--SA-mode | -C|--CS-mode] [-r | --repair] | [-i table_list.txt|--input-
class-file]
```

- **cuprid** : An integrated utility for CUBRID service and database management.
- **checkdb** : A utility that checks the data consistency of a specific database.
- *options* : **-S**, **-C**, **-r**, and **-i** options are supported.
- *database_name* : The name of the database whose consistency status will be either checked or restored.


```

table_list.txt : A file name to store the list of the tables for consistency check or
recovery
class_name1 class_name2 : List the table names for consistency check or recovery

```

Options

Checking the database consistency in standalone mode (-S or --SA-mode)

The **-S** option is used to access a database in standalone, which means it works without processing server; it does not have an argument. If **-S** is not specified, the system recognizes that a database is running in client/server mode.

```
cubrid checkdb -S testdb
```

Checking the database consistency in client/server mode (-C or --CS-mode)

The **-C** option is used to access a database in client/server mode, which means it works in client/server process respectively; it does not have an argument. If **-C** is not specified, the system recognize that a database is running in client/server mode by default.

```
cubrid checkdb -C testdb
```

Restoring in case of a database consistency problem (-r or --repair)

The **-r** option is used to restore an issue if a consistency error occurs in a database.

```
cubrid checkdb -r testdb
```

Specifying a table in which consistency is checked or restored (-i, --input-class-file or table name)

You can specify a table in which consistency is check or restored by specifying the **-i table_list.txt** option or listing the table names after a database name. In this way, you can limit the target to be restored and both ways can be used. If a specific target is not specified, entire database will be a target of consistency check or restoration.

```

cubrid checkdb testdb tbl1 tbl2
cubrid checkdb -r testdb tbl1 tbl2
cubrid checkdb -r -i tbl_list.txt testdb tbl1 tbl2

```

Empty string, tab, carriage return and comma are separators among table names in the table list file specified by **-i** option. The following example shows the table list file; from t1 to t10, it is recognized as a table for consistency check or restoration.

```

t1 t2 t3,t4 t5
t6, t7 t8 t9

t10

```

Killing Database Transactions

Description

The **cubrid killtran** is used to check transactions or abort specific transaction. Only a DBA can execute this utility.

Syntax

```

cubrid killtran options database name
options :
[{-i|--kill-transaction-index=}index] [--kill-user-name=id] [--kill-host-name=host] [--
kill-program-name=program_name] [{-p|--dba-password=}password] [-d|--display-information]
[-f|--force]

```

- **cubrid** : An integrated utility for the CUBRID service and database management
- **killtran** : A utility that manages transactions for a specified database
- **options** : Some options refer to killing specified transactions; others refer to outputting active transactions. If no option is specified, **-d** is specified by default so all transactions are displayed on the screen. **-p** A value followed by the **-p** option is a password of the **DBA**, and should be entered in the prompt.
- **database_name** : The name of database whose transactions are to be killed

Options

Displaying all transactions (no option)

```
cubrid killtran testdb
```

Tran index	User name	Host name	Process id	Program name
1(+)	dba	myhost	664	cub_cas
2(+)	dba	myhost	6700	csql
3(+)	dba	myhost	2188	cub_cas
4(+)	dba	myhost	696	csql
5(+)	public	myhost	6944	csql

Killing transactions in a specified index (-i or --kill-transaction-index)

```
cubrid killtran -i 1 testdb
```

Ready to kill the following transactions:

Tran index	User name	Host name	Process id	Program name
1(+)	dba	myhost	4760	csql

Do you wish to proceed ? (Y/N)y

Killing transaction associated with transaction index 1

Displaying all transactions (-d or --display)

```
cubrid killtran -d testdb
```

Tran index	User name	Host name	Process id	Program name
2(+)	dba	myhost	6700	csql
3(+)	dba	myhost	2188	cub_cas
4(+)	dba	myhost	696	csql
5(+)	public	myhost	6944	csql

Killing transactions for a specified OS user ID (--kill-user-name)

```
cubrid killtran --kill-user-name=os_user_id testdb
```

Killing transactions of a specified client host (--kill-host-name)

```
cubrid killtran --kill-host-name=myhost testdb
```

Killing transactions for a specified program (--kill-program-name)

```
cubrid killtran --kill-program-name=cub_cas testdb
```

Omitting a prompt to check transactions to be stopped (-f or --force)

```
cubrid killtran -f -i 1 testdb
```

Checking the Query Plan Cache

Description

The **cubrid plandump** utility is used to display information on the query plans stored (cached) on the server.

Syntax

```
cubrid plandump options database name  
options : [-d|--drop] [{-o|--output-file=} file]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **plandump** : A utility that displays the query plans stored in the current cache of a specific database.
- *options* : The **-d** and **-o** options are supported.

- *database_name* : The name of the database where the query plans are to be checked or dropped from its sever cache.

Options

Checking the query plans stored in the cache

```
cubrid plandump testdb
```

Dropping the query plans stored in the cache (-d or --drop)

```
cubrid plandump -d testdb
```

Saving the results of the query plans stored in the cache to a file (-o or --output)

```
cubrid plandump -o output.txt testdb
```

Outputting Internal Database Information

Description

You can check various pieces of internal information on the database with the **cubrid diagdb** utility. Information provided by **cubrid diagdb** is helpful in diagnosing the current status of the database or figuring out a problem.

Syntax

```
cubrid diagdb options database_name  
options : [{-d | --dump-type=} type]
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **diagdb** : A command that is used to check the current storage state of the database by outputting the information contained in the binary file managed by CUBRID in text format. It normally executes only when the database is in a stopped state. You can check the whole database or the file table, file size, heap size, class name or disk bitmap selectively by using the provided option.
- *options* : The **-d** option is provided.
- *database_name* : The name of the database of which internal information is to be diagnosed.

Options

Specifying the output range (-d or --dump-type)

The following example shows how to display the information of all files in the testdb database. If any option is not specified, the default value of 1 is used.

```
cubrid diagdb -d 1 myhost testdb
```

The utility has 9 types of -d options as follows:

Type	Description
-1	Displays all database information.
1	Displays file table information.
2	Displays file capacity information.
3	Displays heap capacity information.
4	Displays index capacity information.
5	Displays class name information.
6	Displays disk bitmap information.
7	Displays catalog information.
8	Displays log information.
9	Displays hip information.

Backup and Restore

DBA must perform regular backups of the database so that it can be restored successfully to a state at a certain point in time in case of system failure. For details, see [Database Backup](#).

Exporting and Importing

To use a newer version of CUBRID database, the existing version must be migrated to a new one. For this purpose, you can use "Export to a ASCII text file" and "Import from a ASCII text file" features provided by CUBRID. For details on export and import, see [Migrating Database](#).

Dumping Parameters Used in Server/Client

Description

The **cubrid paramdump** utility outputs parameter information used in the server/client process.

Syntax

```
 cubrid paramdump  options database_name
options : [{-o|--output-file=}filename] [{-b|--both}] [{-S|--SA-mode}] [{-C|--CS-mode}]
```

- **cubrid** : An integrated utility for the CUBRID service and database management
- **paramdump** : A utility that outputs parameter information used in the server/client process
- *options* : A short name option starts with a single dash (-) while a full name option starts with a double dash (--). **-o**, **-b**, **-S** and **-C** options are provided.
- *database_name* : The name of the database in which parameter information is to be displayed

Options

Storing output information to a file (-o)

The **-o** option is used to store information of the parameters used in the server/client process of the database into a specified file. The file is created in the current directory. If the **-o** option is not specified, the message is displayed on a console screen.

```
cubrid paramdump -o db_output testdb
```

Displaying information of the server/client parameters (-b)

The **-b** option is used to display parameter information used in server/client process on a console screen. If the **-b** option is not specified, only server-side information is displayed.

```
cubrid paramdump -b testdb
```

Displaying parameter information of the server process in standalone mode (-S or --SA-mode)

```
cubrid paramdump -S testdb
```

Displaying parameter information of the server process in client/server mode (-C or --CS-mode)

```
cubrid paramdump -C testdb
```

Database Migration

Migrating Database

To use a newer version of CUBRID database, you might migrate an existing data to a new one. For this purpose, you can use the "Export to a ASCII text file" and "Import from a ASCII text file" features provided by CUBRID. The following section explains migration steps using the **cubrid unloaddb** and **cubrid loaddb** utilities.

Recommended scenario and procedures

The following is an explanation of a migration scenario that can be applied while the existing version of CUBRID is running. For database migration, the **cubrid unloaddb** and **cubrid loaddb** utilities are used. For details, see [Unloading Database](#) and [Loading Database](#).

1. Back up the existing database

Back up the existing version of the database by using the **cubrid backupdb** utility. The purpose of this step is to safeguard against failures that might occur during the database unload/load operations. For details on the database backup, see [Database Backup](#).

2. Unload the existing database

Unload the database created for the existing version of CUBRID by using the **cubrid unloaddb** utility. For details on the database unload, see [Unloading Database](#).

3. Storing the existing CUBRIDG configuration files

Save configurations files such as **cubrid.conf**, **cubrid_broker.conf** and **cm.conf** located in the **CUBRID/conf** directory. The purpose of this step is to conveniently apply parameter values for the existing CUBRID database environment to the new one.

4. Install a new version of CUBRID

Once backing up and unloading of the data created by the existing version of CUBRID have been completed, delete the existing version of CUBRID and its databases and then install the new version of CUBRID. For details on installing CUBRID, see [Installing and Running on Linux](#) in "Getting Started."

5. Configure the new CUBRID

You can configure the new version of CUBRID by referring to configuration files of the existing database stored in the step 3, "**Save configuration files of the existing version of CUBRID.**" For details on configuration, see [Installing and Running on Windows](#) in "Getting Started."

6. Load the new database

Create a database by using the **cubrid createdb** utility and then use the **cubrid loaddb** utility to load into the new database the data which had previously been unloaded. For details on creating a database, see [Creating Database](#) in "Administrator's Guide." For details on database loading, see [Loading Database](#).

7. Back up the new database

Once the data has been successfully loaded into the new database, back up the database created for the new version of CUBRID by using the **cubrid backupdb** utility. The reason for this step is because you cannot restore the data backed up in the existing version of CUBRID when using the new version. For details on backing up the database, see [Database Backup](#).

Unloading Database

Description

The purposes of loading/unloading databases are as follows:

- To rebuild databases by volume reconfiguration
- To migrate database in different system environments
- To migrate database in different versions

Syntax

```

cubrid unloaddb [ options ] database_name
[ options ]
-i | -O | -s | -d | -v | -S | -C |
--input-class-file | --output-path | --schema-only | --data-only | --verbose | --SA-mode |
--CS-mode | --include-reference | --input-class-only | --lo-count | --estimated-size | --
cached-pages | --output-prefix | --hash-file | --datafile-per-class

```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **unloaddb** : A utility that creates ASCII files from a database. It is used together with the **cubrid loaddb** utility for replacing system, upgrading product version or reorganizing database volumes. It can be used both in standalone and client/server modes. Data can be unloaded even when the database is running.
- *options* : A short option starts with a single dash (-) while a full name option starts with a double dash (--). Note that options are case sensitive.
- *database_name* : Specifies the name of the database to be unloaded.

Return values

Return values of **cubrid unloaddb** utility are as follows:

- 0 : Success
- Non-zero : Failure

Created Files

- Schema file (*database-name_schema*) : A file that contains information on the schema defined in the database.
- Object file (*database-name_objects*) : A file that contains information on the records in the database.
- Index file (*database-name_indexes*) : A file that contains information on the indexes defined in the database.
- Trigger file (*database-name_trigger*) : A file that contains information on the triggers defined in the database. If you don't want triggers to be running while loading the data, load the trigger definitions after the data loading has completed.

The schema, object, index, and trigger files are created in the same directory.

Options

The following table shows options available with the **cubrid unloaddb** utility (options are case sensitive).

Option	Description
-i --input-class-file	Unloads the database class into the input file specified in an argument.
-O --output-path	Specifies the directory in which to create schema and object files. If the option is not specified, files are created in the current directory.
-s --schema-only	Creates only the schema file, not the data file.
-d --data-only	Creates only the data file, not the schema file.
-v --verbose	Displays detailed information on the database being unloaded.
-S --SA-mode	Unloads the database in standalone mode.
-C --CS-mode	Unloads the database in client/server mode.
--include-reference	Unloads the object reference as well when the specified database class is unloaded with the -i option.
--input-class-only	Is used with the -i option. Creates only the schema files which are related to tables included in the input file.

<code>--lo-count</code>	Specifies the number of large object (LO) data files to be created in a single directory. Default value : 0
<code>--estimated-size</code>	Specifies the number of records expected.
<code>--cached-pages</code>	Configures the number of object tables to be cached in the memory. Default value : 100
<code>--output-prefix</code>	Specifies the prefix for schema and object file names.
<code>--hash-file</code>	Specifies the name of the hash file.
<code>--datafile-per-class</code>	Generates a data file per each table.

Input file with the list of tables to be unloaded (-i or --input-class-file)

The following example shows an input file (table_list.txt).

```
table_1
table_2
..
table_n
```

The `-i` option specifies the input file where the list of tables to be unloaded is stored so that only specified part of the database can be unloaded.

```
cubrid unloaddb -i table_list.txt demodb
```

The `-i` option can be used together with the `--input-class-only` option that creates the schema file related to only those tables included in the input file.

```
cubrid unloaddb --input-class-only -i table_list.txt demodb
```

The `-i` option can be used together with the `--include-reference` option that creates the object reference as well.

```
cubrid unloaddb --include-reference -i table_list.txt demodb
```

Specifying the directory where files created will be stored (-O or --output-path)

The `-O` option specifies the directory where the output files generated by the unload operation is stored. If the `-O` option is not specified, output files are created in the current working directory.

```
cubrid unloaddb -O ./CUBRID/Databases/demodb demodb
```

If the specified directory does not exist, the following error message will be displayed.

```
unloaddb: No such file or directory.
```

Creating the schema file only (-s or --schema-only)

The `-s` option specifies that only the schema file will be created from amongst all the output files which can be created by the unload operation.

```
cubrid unloaddb -s demodb
```

Creating the data file only (-d or -data-only)

The `-d` option specifies that only the data file will be created from amongst all of the output files which can be created by the unload operation.

```
cubrid unloaddb -d demodb
```

Creates data files by table (--datafile-per-class)

`--datafile-per-class` is the option specifying that the output file generated through unload operation creates a data file per each table. The file name is generated as `<Database Name>_<Table Name>_objects` for each table. However, all column values in object types are unloaded as NULL and `%id class_name class_id` part is not written in the unloaded file (see [How to Write a File to Load Database](#)).

```
cubrid unloaddb -d demodb
```

Displaying the unload status information (-v or --verbose)

The `-v` option displays detailed information on the database tables and records being unloaded while the unload operation is under way.

```
cubrid unloaddb -v demodb
```

Standalone mode (-S or --SA-mode)

The `-S` option performs the unload operation by accessing the database in standalone mode.

```
cubrid unloaddb -S demodb
```

Client/server mode (-C or --CS-mode)

The `-C` option performs the unload operation by accessing the database in client/server mode.

```
cubrid unloaddb -C demodb
```

Number of estimated records (--estimated-size)

The `--estimated-size` option allows you to assign hash memory to store records of the database to be unloaded. If the `--estimated-size` option is not specified, the number of records of the database is determined based on recent statistics information. This option can be used if the recent statistics information has not been updated or if a large amount of hash memory needs to be assigned. Therefore, if the number given as the argument for the option is too small, the unload performance deteriorates due to hash conflicts.

```
cubrid unloaddb --estimated-size 1000 demodb
```

Number of pages to be cached (--cached-pages)

The `--cached-pages` option specifies the number of pages of tables to be cached in the memory. Each page is 4,096 bytes. The administrator can configure the number of pages taking into account the memory size and speed. If this option is not specified, the default value is 100 pages.

```
cubrid unloaddb --cached-pages 500 demodb
```

Specifying the prefix for the name of the file to be created (--output-prefix)

The `--output-prefix` option specifies the prefix for the names of schema and object files created by the unload operation. Once the example is executed, the schema file name becomes `abcd_schema` and the object file name becomes `abcd_objects`. If the `--output-prefix` option is not specified, the name of the database to be unloaded is used as the prefix.

```
cubrid unloaddb --output-prefix abcd demodb
```

Loading Database

Description

You can load a database by using the `cubrid loaddb` utility in the following situations:

- Migrating previous version of CUBRID database to new version of CUBRID database
- Migrating a third-party DBMS database to a CUBRID database
- Inserting massive amount of data faster than using the `INSERT` statement

In general, the `cubrid loaddb` utility uses files (schema definition, object input, and index definition files) created by the `cubrid unloaddb` utility .

Syntax

```
cubrid loaddb [ options ] database name
[ options ]
-u | -p | -l | -v | -c | -s | -i | -d |
--user | --password | --load-only | --verbose | --periodic-commit | --schema-file | --
index-file | --data-file | --data-file-check-only | --estimated-size | --no-oid | --no-
statistics | --ignore-class-file | --error-control-file | --no-logging
```


- **cubrid** : An integrated utility for the CUBRID service and database management.
- **loaddb** : A utility loads files which is generated by the unload operation and then creates a new database. It is also used to enter mass data into a database faster than ever by loading the input file written by a user. Database loading is performed in standalone mode with **DBA** authorization.
- *options* : A short name option starts with a single dash (-) while a full name option starts with a double dash (--). The options are case sensitive.
- *database_name* : Specifies the name of the database to be created.

Return Value

Return values of **cubrid loaddb** utility are as follows:

- 0 : Success
- Non-zero : Failure

Input Files

- Schema file (*database-name_schema*) : A file generated by the unload operation; it contains schema information defined in the database.
- Object file (*database-name_objects*) : A file created by an unload operation. It contains information on the records in the database.
- Index file (*database-name_indexes*) : A file created by an unload operation. It contains information on the indexes defined in the database.
- Trigger file (*database-name_trigger*) : A file created by an unload operation. It contains information on the triggers defined in the database.
- User-defined object file (*user_defined_object_file*) : A file in table format written by the user to enter mass data. (For details, see [How to Write Files to Load Database.](#))

Options

The following table shows options available with the **cubrid loaddb** utility (options are case sensitive).

Option	Description
-u --user	Enters the database user's account. The default value is PUBLIC .
-p --password	Enters the database user's password.
-l --load-only	Skips checking statements and types included in the object file and loads records.
-v --verbose	Displays detailed information on the data loading status on the screen.
-c --periodic-commit	Commits the transaction whenever a specified number of records has been entered.
-s --schema-file	Specifies the schema file created by the unload operation and performs schema loading.
-i --index-file	Specifies the index file created by the unload operation and loads indexes.
-d --data-file	Specifies the data file created by the unload operation and loads records.
--data-file-check-only	Performs checking only for statements and types included in the data file, but does not load records.
--estimated-size	Specifies the number of records expected.
--no-oid	Ignores the OID reference relationship included in the data file and loads records.

--no-statistics	Loads records without updating database statistics information.
--ignore-class-file	Specifies the ignoring classes.
--error-control-file	Specifies the file that describes how to handle specific errors occurring during data loading.
--no-logging	Can load data quickly during execution because transaction logs are stored; however, it has risk, which data cannot be recovered in case of error occurred. Thus, you should read the messages in the Remarks section below in this page carefully.

Entering a user account (-u or --user)

The **-u** option specifies the user account of a database where records are loaded. If the option is not specified, the default value is **PUBLIC**.

```
cubrid loaddb -u admin -d demodb_objects newdb
```

Entering the password (-p or --password)

The **-p** option specifies the password of a database user who will load records. If the option is not specified, you will be prompted to enter the password.

```
cubrid loaddb -p admin -d demodb_objects newdb
```

Loading records without checking syntax (-l or --load-only)

The **-l** option loads data directly without checking the syntax for the data to be loaded. The following example shows how to load data included in demodb_objects to newdb.

If the **-l** option is used, loading speed increases because data is loaded without checking the syntax included in demodb_objects, but an error might occur.

```
cubrid loaddb -l -d demodb_objects newdb
```

Displaying the loading status information (-v or --verbose)

The following example shows how to display detailed information on the tables and records of the database being loaded while the database loading operation is performed. You can check the detailed information such as the progress level, the class being loaded and the number of records entered by using the **-v** option.

```
cubrid loaddb -v -d demodb_objects newdb
```

Configuring the commit interval (-c or --periodic-commit)

The following command performs commit regularly every time 100 records are entered into the newdb by using the **-c** option. If the **-c** option is not specified, all records included in demodb_objects are loaded to newdb before the transaction is committed. If the **-c** option is used together with the **-s** or **-i** option, commit is performed regularly every time 100 DDL statements are loaded. The recommended commit interval varies depending on the data to be loaded. It is recommended that the parameter of the **-c** option be configured to 50 for schema loading, 1,000 for record loading, and 1 for index loading.

```
cubrid loaddb -c 100 -d demodb_objects newdb
```

Schema loading (-s or --schema-file)

The following statement loads the schema information defined in demodb into the newly created newdb database. demodb_schema is a file created by the unload operation and contains the schema information of the unloaded database. You can load the actual records after loading the schema information first by using the **-s** option.

```
cubrid loaddb -u dba -s demodb_schema newdb
```

```
Start schema loading.
Total 86 statements executed.
Schema loading from demodb_schema finished.
Statistics for Catalog classes have been updated.
```

The following statement loads the triggers defined in demodb into the newly created newdb database. demodb_trigger is a file created by the unload operation and contains the trigger information of the unloaded database. It is recommended to load the schema information after loading the records.

```
cubrid loaddb -u dba -s demodb_trigger newdb
```

Index loading (-i or --index-file)

The following command loads the index information defined in demodb into the newly created newdb database. demo_indexes is a file created by the unload operation and contains the index information of the unloaded database. You can create indexes after loading records by using the **-i** option together with the **-d** option.

```
cubrid loaddb -u dba -i demodb_indexes newdb
```

Data loading (-d or -data-file)

The following command loads the record information into newdb by specifying the data file or the user-defined object file with the **-d** option. demodb_objects is either an object file created by the unload operation or a user-defined object file written by the user for mass data loading.

```
cubrid loaddb -u dba -d demodb_objects newdb
```

Checking the syntax for the data to be loaded only (--data-file-check-only)

The following is a command that checks the statements for the data contained in demodb_objects by using the **--data-file-check-only** option. Therefore, the execution of the command below does not load records.

```
cubrid loaddb --data-file-check-only -d demodb_objects newdb
```

Number of expected records (--estimated-size)

The **--estimated-size** option can be used to improve loading performance when the number of records to be unloaded exceeds the default value of 5,000. That is, you can improve the load performance by assigning large hash memory for record storage with this option.

```
cubrid loaddb --estimated-size 8000 -d demodb_objects newdb
```

Loading records while ignoring the reference relationship (--no-oid)

The following is a command that loads records into newdb ignoring the OIDs in demodb_objects.

```
cubrid loaddb --no-oid -d demodb_objects newdb
```

Loading records without updating statistics information (--no-statistics)

The following is a command that does not update the statistics information of newdb after loading demodb_objects. It is useful especially when small data is loaded to a relatively big database; you can improve the load performance by using this command.

```
cubrid loaddb --no-statistics -d demodb_objects newdb
```

Specifying the ignoring classes (--ignore-class-file)

You can specify a file that lists classes to be ignored during loading records. All records of classes except ones specified in the file will be loaded.

```
cubrid loaddb --ignore-class-file=skip_class_list -d demodb_objects newdb
```

Specifying the error information file (--error-control-file)

This option specifies the file describing how to handle specific errors occurring during database loading.

```
cubrid loaddb --error-control-file=error_test -d demodb_objects newdb
```

Remark

The **--no-logging** option enables to load data file quickly when **loaddb** is executed by not storing transaction logs; however, it has risk, which data cannot be recovered in case of errors occurred such as incorrect file format or system

failure. In this case, you must rebuild database to solve the problem. Thus, in general, it is not recommended to use this option exception of building a new database which does not require data recovery.

How to Write Files to Load Database

You can add mass data to the database more rapidly by writing the object input file used in the **cubrid loaddb** utility. An object input file is a text file in simple table form that consists of comments and command/data lines.

Comment

In CUBRID, a comment is represented by two hyphens (--).

```
-- This is a comment!
```

Command Line

A command line begins with a percent character (%) and consists of **%class** and **%id** commands; the former defines classes, and the latter defines aliases and identifiers used for class identification.

Assigning an Identifier to a Class

You can assign an identifier to class reference relationships by using the **%id** command.

Syntax

```
%id class_name class_id
class_name:
    identifier
class_id:
    integer
```

The *class_name* specified by the **%id** command is the class name defined in the database, and *class_id* is the numeric identifier which is assigned for object reference.

Example 1

```
%id employee 2
%id office 22
%id project 23
%id phone 24
```

Specifying the Class and Attribute

You can specify the classes (tables) and attributes (columns) upon loading data by using the **%class** command. The data line should be written based on the order of attributes specified.

Syntax

```
%class class_name ( attr_name [ { attr_name } _ ]
```

The schema must be pre-defined in the database to be loaded.

The *class_name* specified by the **%class** command is the class name defined in the database and the *attr_name* is the name of the attribute defined.

Example 2

The following example shows how to specify a class and three attributes by using the **%class** command to enter data into a class named employee. Three pieces of data should be entered on the data lines after the **%class** command. For this, see [Example 3](#) in the "Configuring Reference Relationship" section.

```
%class employee (name age department)
```

Data Line

A data line comes after the **%class** command line. Data loaded must have the same type as the class attributes specified by the **%class** command. The data loading operation stops if these two types are different.

Data for each attribute must be separated by at least one space and be basically written as a single line. However, if the data to be loaded takes more than one line, you should specify the plus sign (+) at the end of the first data line to enter data continuously on the following line. Note that no space is allowed between the last character of the data and the plus sign.

Loading an Instance

As shown below, you can load an instance that has the same type as the specified class attribute. Each piece of data is separated by at least one space.

Example 1

```
%class employee (name)
'jordan'
'james'
'garnett'
'malone'
```

Assigning an Instance Number

You can assign a number to a given instance at the beginning of the data line. An instance number is a unique positive number in the specified class. Spaces are not allowed between the number and the colon (:). Assigning an instance number is used to configure the reference relationship for later.

Example 2

```
%class employee (name)
1: 'jordan'
2: 'james'
3: 'garnett'
4: 'malone'
```

Configuring Reference Relationship

You can configure the object reference relationship by specifying the reference class after an "at sign (@)" and the instance number after the "vertical line (|)".

Syntax

```
@class_ref | instance_no
class_ref:
  class_name
  class_id
```

Specify a class name or a class id after the @ sign, and an instance number after a vertical line (|). Spaces are not allowed before and after a vertical line (|).

Example 3

The following example shows how to load class instances into the paycheck class. The name attribute references an instance of the employee class. As in the last line, data is loaded as **NULL** if you configure the reference relationship by using an instance number not specified earlier.

```
%class paycheck(name department salary)
@employee|1 'planning' 8000000
@employee|2 'planning' 6000000
@employee|3 'sales' 5000000
@employee|4 'development' 4000000
@employee|5 'development' 5000000
```

Example 4

Since the id 21 was assigned to the employee class by using the **%id** command in the [Assigning an Identifier to a Class](#) section, Example 3 can be written as follows:

```
%class paycheck(name department salary)
@21|1 'planning' 8000000
@21|2 'planning' 6000000
@21|3 'sales' 5000000
@21|4 'development' 4000000
@21|5 'development' 5000000
```

Database Backup and Restore

Database Backup

A database backup is the procedure of storing CUBRID database volumes, control files and log files, and it is executed by using the **cubrid backupdb** utility or the CUBRID Manager. **DBA** must regularly back up the database so that the database can be properly restored in the case of storage media or file errors. The restore environment must have the same operating system and the same version of CUBRID as the backup environment. For such a reason, you must perform a backup in a new environment immediately after migrating a database to a new version.

To recover all database pages, control files and the database to the state at the time of backup, the **cubrid backupdb** utility copies all necessary log records.

Syntax

```
cubrid backupdb [ options ] database_name
[ options ]
-D | -r | -l | -o | -S | -C | -t | -z | -e |
--destination-path | --remove-archive | --level | --output-file | --SA-mode | --CS-mode |
--thread-count | --compress | --except-active-log | --no-check
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **backupdb** : A utility for database backup. Performs an online, offline, compressed or parallel backup depending on the option used. This utility can only be executed by a user who has the backup authorization (e.g. **DBA**).
- *options* : A short option starts with a single dash (-) while a full name option starts with a double dash (--). Options are case sensitive.
- *database_name* : Specifies the name of the database to be backed up.

Return Value

- 0 : Success
- Non-zero : Failure

Options

The following table shows options available with the **cubrid backupdb** utility (options are case sensitive).

Option	Description
-D --destination-path	Specifies the directory path name or device name where backup volumes are to be created. The default value is the location of log-path specified in the database location file (databases.txt) which was generated upon database creation.
-r --remove-archive	Removes unnecessary archive logs after the backup is complete.
-l --level	Specifies the backup level to 0, 1 or 2. The default value is a full backup (0).
-o --output-file	Specifies the name of the file where progress information is to be displayed.
-S --SA-mode	Performs a backup in standalone mode.
-C --CS-mode	Performs a backup in client/server mode.
-t --thread-count	Specifies the maximum number of threads allowed for a parallel backup. The default value is the number of CPUs in the system.

-z --compress	Performs a compressed backup.
-e --except-active-log	Specifies that active log volumes are not included in the backup.
--sleep-msecs	Specifies the interval of idle time after reading 1 MB of data from a backup file. The default value is 0 in milliseconds.
--no-check	Does not perform a consistency check on a database before making a backup.

Performing backup by specifying the directory in which backup files are to be stored (-D or --destination-path)

The following example shows how to use the **-D** option to store backup files in the specified directory. The backup file directory must be specified before performing this job. If the **-D** option is not specified, backup files are stored in the directory specified in the **databases.txt** file which stores database location information.

```
cubrid backupdb -D /home/cubrid/backup demodb
```

The following example shows how to store backup files in the current directory by using the **-D** option. If you enter a period (.) following the **-D** option as an argument, the current directory is specified.

```
cubrid backupdb -D . demodb
```

Removing archive logs after backup (-r or --remove-archive)

Writes an active log to a new archive log file when the active log is full. If a backup is performed in such a situation and backup volumes are created, backup logs created before the backup will not be used in subsequent backups. The **-r** option is used to remove archive log files that will not be used any more in subsequent backups after the current one is complete.

The **-r** option only removes unnecessary archive log files that were created before backup, and does not have any impact on backup; however, if an administrator removes the archive log file after a backup, it may become impossible to restore everything. For this reason, archive logs should be removed only after careful consideration.

If you perform an incremental backup (backup level 1 or 2) with the **-r** option, there is the risk that normal recovery of the database will be impossible later on. Therefore, it is recommended that the **-r** option only be used when a full backup is performed.

```
cubrid backupdb -r demodb
```

The **-r** option does not affect the restore because it removes only unnecessary archive logs before the backup, but full restore may not be possible if the administrator removes archive logs created after the backup as well; when you remove archive logs, you must check if those logs would be required in any subsequent restore.

If you perform an incremental backup (backup level 1 or 2) with the **-r** option, there is the risk that normal recovery of the database will be impossible later on. Therefore, it is recommended that the **-r** option only be used when a full backup is performed.

Performing a backup with the backup level specified (-l or --level)

The following example shows how to execute an incremental backup of the level specified by using the **-l** option. If the **-l** option is not specified, a full backup is performed. For details on backup levels, see [Incremental Backup](#).

```
cubrid backupdb -l 1 demodb
```

Storing backup progress information in the specified file (-o or --output-file)

The following example shows how to write the progress of the database backup to the `info_backup` file by using the **-o** option.

```
cubrid backupdb -o info_backup demodb
```

The following example shows the contents of the `info_backup` file. You can check the information on the number of threads, compression method, backup start time, the number of permanent volumes, backup progress and backup end time.


```
[ Database(demodb) Full Backup start ]
- num-threads: 1
- compression method: NONE
- backup start time: Mon Jul 21 16:51:51 2008
- number of permanent volumes: 1
- backup progress status
-----
volume name          | # of pages | backup progress status | done
-----
demodb_vinf          |          1 | #####                 | done
demodb                |        25000 | #####                 | done
demodb_lginf         |          1 | #####                 | done
demodb_lgat          |        25000 | #####                 | done
-----
# backup end time: Mon Jul 21 16:51:53 2008
[Database(demodb) Full Backup end]
```

Performing backup in standalone mode (-S or --SA-mode)

The following example shows how to perform backup in standalone mode (that is, backup offline) by using the **-S** option. If the **-S** option is not specified, the backup is performed in client/server mode.

```
cubrid backupdb -S demodb
```

Performing backup in client/server mode (-C or --CS-mode)

The following example shows how to perform backup in client/server mode by using the **-C** option and the demodb database is backed up online. If the **-C** option is not specified, a backup is performed in client/server mode.

```
cubrid backupdb -C demodb
```

Parallel backup (-t or --thread-count)

The following example shows how to execute parallel backup with the number of threads specified by the administrator by using the **-t** option. Even when the argument of the **-t** option is not specified, a parallel backup is performed by automatically assigning as many threads as CPUs in the system.

```
cubrid backupdb -t 4 demodb
```

Compressed backup (-z or --compress)

The following example shows how to compress the database and stores it in the backup file by using the **-z** option. The size of the backup file and the time required for backup can be reduced by using the **-z** option.

```
cubrid backupdb -z demodb
```

Enabling to exclude active log volumes (-e or --except-active-log)

The following example shows how to execute backup excluding active logs of the database by using the **-e** option. You can reduce the time required for backup by using the **-e** option. However, extra caution is required because active logs needed for completing a restore to the state of a certain point from the backup point are not included in the backup file, which may lead to an unsuccessful restore.

```
cubrid backupdb -e demodb
```

Adjusting the interval of idle time during a backup (--sleep-msecs)

The **--sleep-msecs** option allows you to specify the interval of idle time during the database backup. The default value is 0 in milliseconds. The system becomes idle for the specified amount of time whenever it reads 1 MB of data from a file. This option is used to reduce the performance degradation of an active server during a live backup. The idle time will prevent excessive disk I/O operations.

```
cubrid backupdb --sleep-msecs=5 demodb
```

Disabling database consistency check (--no-check)

The following example shows how to execute backup without checking the consistency of the database by using the **--no-check** option.

```
cubrid backupdb --no-check demodb
```

Backup Strategy and Method

The following must be considered before performing a backup:

- **Selecting the data to be backed up**
- Determine whether it is valid data worth being preserved.
- Determine whether to back up the entire database or only part of it.
- Check whether there are other files to be backed up along with the database.
- **Choosing a backup method**
- Choose the backup method from one of incremental and online backups. Also, specify whether to use compression backup, parallel backup, and mode.
- Prepare backup tools and devices available.
- **Determining backup time**
- Identify the time when the least usage in the database occur.
- Check the size of the archive logs.
- Check the number of clients using the database to be backed up.

Online Backup

An online backup (or a hot backup) is a method of backing up a currently running database. It provides a snapshot of the database image at a certain point in time. Because the backup target is a currently running database, it is likely that uncommitted data will be stored and the backup may affect the operation of other databases.

To perform an online backup, use the **cubrid backupdb -C** command.

Offline Backup

An offline backup (or a cold backup) is a method of backing up a stopped database. It provides a snapshot of the database image at a certain point in time.

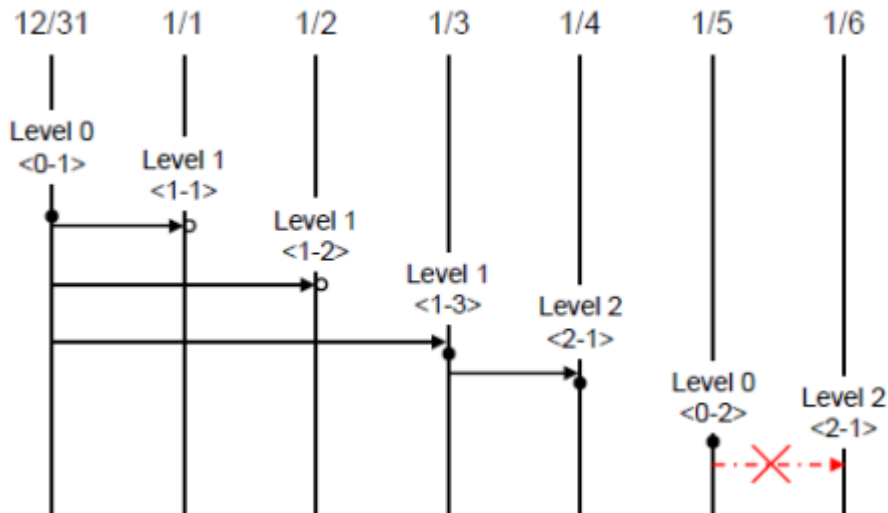
To perform an offline backup, use the **cubrid backupdb -S** command.

Incremental Backup

An incremental backup, which is dependent upon a full backup, is a method of only backing up data that have changed since the last backup. This type of backup has an advantage of requiring less volume and time than a full backup. CUBRID supports backup levels 0, 1 and 2. A higher level backup can be performed sequentially only after a lower level backup is complete.

To perform an incremental backup, use the **cubrid backupdb -I <level>** command.

The following example shows incremental backup. Let's example backup levels in details.



- **Full backup (backup level 0)** : Backup level 0 is a full backup that includes all database pages.

The level of a backup which is attempted first on the database naturally becomes a 0 level. **DBA** must perform full backups regularly to prepare for restore situations. In the example, full backups were performed on December 31st and January 5th.

- **First incremental backup (backup level 1)** : Backup level 1 is an incremental backup that only stores changes since the level 0 full backup, and is called a "first incremental backup."

Note that the first incremental backups are attempted sequentially such as <1-1>, <1-2> and <1-3> in the example, but they are always performed based on the level 0 full backup.

Suppose that backup files are created in the same directory. If the first incremental backup <1-1> is performed on January 1st and then the first incremental backup <1-2> is attempted again on January 2nd, the incremental backup file created in <1-1> is overwritten. The final incremental backup file is created on January 3rd because the first incremental backup is performed again on that day.

Since there can be a possibility that the database needs to be restored the state of January 1st or January 2nd, it is recommended for **DBA** to store the incremental backup files <1-1> and <1-2> separately in storage media before overwriting with the final incremental file.

- **Second incremental backup (backup level 2)** : Backup level 2 is an incremental backup that only stores data that have changed since the first incremental backup, and is called a "second incremental backup."

A second incremental backup can be performed only after the first incremental backup. Therefore, the second incremental backup attempted on January fourth succeeds; the one attempted on January sixth fails.

Backup files created for backup levels 0, 1 and 2 may all be required for database restore. To restore the database to its state on January fourth, for example, you need the second incremental backup generated at <2-1>, the first incremental backup file generated at <1-3>, and the full backup file generated at <0-1>. That is, for a full restore, backup files from the most recent incremental backup file to the earliest created full backup file are required.

Compress Backup

A compress backup is a method of backing up the database by compressing it. This type of backup reduces disk I/O costs and stores disk space because it requires less backup volume.

To perform a compress backup, use the **cubrid backupdb -z|--compress** command.

Parallel Backup Mode

A parallel or multi-thread backup is a method of performing as many backups as the number of threads specified. In this way, it reduces backup time significantly. Basically, threads are given as many as the number of CPUs in the system.

To perform a parallel backup, use the **cubrid backupdb -t|--thread-count** command.

Managing Backup Files

One or more backup files can be created in sequence based on the size of the database to be backed up. A unit number is given sequentially (000, 001-0xx) to the extension of each backup file based in the order of creation.

Managing Disk Capacity during the Backup

During the backup process, if there is not enough space on the disk to store the backup files, a message saying that the backup cannot continue appears on the screen. This message contains the name and path of the database to be backed up, the backup file name, the unit number of backup files and the backup level. To continue the backup process, the administrator can choose one of the following options:

- Option 0 : An administrator enters 0 to discontinue the backup.
- Option 1 : An administrator inserts a new disk into the current device and enters 1 to continue the backup.
- Option 2 : An administrator changes the device or the path to the directory where backup files are stored and enters 2 to continue the backup.

```
*****
Backup destination is full, a new destination is required to continue:
Database Name: /local1/testing/demodb
  Volume Name: /dev/rst1
    Unit Num: 1
  Backup Level: 0 (FULL LEVEL)
Enter one of the following options:
Type
- 0 to quit.
- 1 to continue after the volume is mounted/loaded. (retry)
- 2 to continue after changing the volume's directory or device.
*****
```

Managing Archive Logs

You must not delete archive logs by using the file deletion command such as `rm` or `del` by yourself; the archive logs should be deleted by system configuration or CUBRID backup utility. In the following three cases, archive logs can be deleted.

- In an HA environment, configure the value of **force_remove_log_archives** to **no** and specify the number in the value of **log_max_archives** (logs are deleted after replication applied).
- In non-HA environment, configure the value of **force_remove_log_archives** to **yes** and specify the number in the value of **log_max_archives**.
- Use **cubrid backupdb -r**; note that it should not be used in an HA environment.

If you want to delete logs as much as possible while operating a database, configure the value of **log_max_archives** to 0 or as small as possible and configure the value of **force_remove_log_archives** to **yes**. Note that in an HA environment, if the value of **force_remove_log_archives** is **yes**, archive logs that have not replicated in a slave node are deleted, which can cause replication errors. Therefore, it is recommended that you configure it to **no**. Although the value of **force_remove_log_archives** is set to **no**, files that are complete for replication can be deleted by HA management process.

Restoring Database

A database restore is the procedure of restoring the database to its state at a certain point in time by using the backup files, active logs and archive logs which have been created in an environment of the same CUBRID version. To perform a database restore, use the **cubrid restoredb** utility or the CUBRID Manager.

The **cubrid restoredb** utility (`restoradb.exe` on Windows) restores the database from the database backup by using the information written to all the active and archive logs since the execution of the last backup.

Syntax

```
cubrid restoredb [ options ] database_name
[ options ]
```

```
-d | -B | -l | -p | -o | -u |
--up-to-date | --backup-file-path | --level | --partial-recovery | --output-file | --use-
database-location-path | --list
```

- **cubrid** : An integrated utility for the CUBRID service and database management.
- **restoredb** : A command for restoration of the specified database. For a successful restoration, you must prepare backup files, active log files and archive log files. This command can be performed only in standalone mode.
- *options* : A short name option starts with a single dash (-) while a full name option starts with a double dash (--). This option is case sensitive.
- *database_name* : Specifies the name of the database to be restored.

Return Value

- 0 : Success
- Non-zero : Failure

Options

The following table shows options available with the **cubrid restoredb** utility (options are case sensitive).

Option	Description
-d --up-to-date	Directly sets the time to backup the database or specifies the backuptime keyword.
-B --backup-file-path	Specifies the directory pathname or device name where backup files are to be located.
-l --level	Sets the restoration level to 0, 1 or 2. The default value is full restoration (0).
-p --partial-recovery	Performs a partial restoration.
-o --output-file	Specifies the name of the file where restoration information is to be displayed.
-u --use-database-location-path	Restores the database to the path specified in the database location file (databases.txt).
--list	Displays information on backup volumes of the database on the screen.

Performing restoration by specifying a specific point (-d or --up-to-date)

The following syntax shows how to restore demodb. If no option is specified, demodb is restored to the point of the last commit by default. If no active/archive log files are required to restore to the point of the last commit, the database is restored only to the point of the last backup.

```
cubrid restoredb demodb
```

demodb can be restored to the given point by using the **-d** option and the command which specifies the date and time of the restoration. The user can specify the restoration point manually in the dd-mm-yyyy:hh:mm:ss (e.g. 14-10-2008:14:10:00) format. If no active log/archive log files are required to restore to the point specified, the database is restored only to the point of the last backup.

```
cubrid restoredb -d 14-10-2008:14:10:00 demodb
```

The following command specifies the restoration point by using the **-d** option and the **backuptime** keyword and restores demodb to the point of the last backup.

```
cubrid restoredb -d backuptime demodb
```

Performing restoration by specifying the directory path to a backup file (-B or --backup-file-path)

You can specify the directory where backup files are to be located by using the **-B** option. If this option is not specified, the system retrieves the backup information file (**dbname_bkvinf**) generated upon a database backup; the backup information file is located in the **log-path** directory specified in the database location information file (**databases.txt**).

And then it searches the backup files in the directory path specified in the backup information file. However, if the backup information file has been damaged or the location information of the backup files has been deleted, the system will not be able to find the backup files. Therefore, the administrator must manually specify the directory where the backup files are located by using the **-B** option.

```
cubrid restoredb -B /home/cubrid/backup demodb
```

If the backup files of demodb is in the current directory, the administrator can specify the directory where the backup files are located by using the **-B** option.

```
cubrid restoredb -B . demodb
```

Performing restoration by specifying backup level (-l or --level)

You can perform restoration by specifying the backup level of the database to 0, 1, or 2. For details on backup levels, see [Increment Backup](#).

```
cubrid restoredb -l 1 demodb
```

Performing partial restoration (-p or --partial-recovery)

The following syntax shows how to perform partial restoration without requesting for the user's response by using the **-p** option. If active or archive logs written after the backup point are not complete, by default the system displays a request message informing that log files are needed and prompting the user to enter an execution option. The partial restoration can be performed directly without such a request message by using the **-p** option. Therefore, if the **-p** option is used when performing restoration, data is always restored to the point of the last backup.

```
cubrid restoredb -p demodb
```

When the **-p** option is not specified, the message requesting the user to select the execution option is as follows:

```
*****
Log Archive /home/cubrid/test/log/demodb lgar002
is needed to continue normal execution.
  Type
  - 0 to quit.
  - 1 to continue without present archive. (Partial recovery)
  - 2 to continue after the archive is mounted/loaded.
  - 3 to continue after changing location/name of archive.
*****
```

- Option 0 : Stops restoring
- Option 1 : Performing partial restoration without log files.
- Option 2 : Performing restoration after locating a log to the current device.
- Option 3 : Resuming restoration after changing the location of a log

Storing restore progress information in the specified file (-o or --output-file)

The following syntax shows how to write the restoration progress of a database to the info_restore file by using the **-o** option.

```
cubrid restoredb -o info_restore demodb
```

Restoring data to the directory specified in the database location file (-u or --use-database-location-path)

The following syntax shows how to restore a database to the path specified in the database location file (**databases.txt**) by using the **-u** option. The **-u** option is useful when you perform a backup on server A and store the backup file on server B.

```
cubrid restoredb -u demodb
```

Checking the backup information of a database (--list)

The following syntax shows how to display information on backup files of a database by using the **--list** option; restoration procedure is not performed with this command.

```
cubrid restoredb --list demodb
```

The following example shows how to display backup information by using the `--list` option. You can specify the path to which backup files of the database are originally stored as well as backup levels.

```
*** BACKUP HEADER INFORMATION ***
Database Name: /local1/testing/demodb
DB Creation Time: Mon Oct 1 17:27:40 2008
Pagesize: 4096
Backup Level: 1 (INCREMENTAL LEVEL 1)
Start_lsa: 513|3688
Last_lsa: 513|3688
Backup Time: Mon Oct 1 17:32:50 2008
Backup Unit Num: 0
Release: 8.1.0
Disk Version: 8
Backup Pagesize: 4096
Zip Method: 0 (NONE)
Zip Level: 0 (NONE)
Previous Backup level: 0 Time: Mon Oct 1 17:31:40 2008
(start lsa was -1|-1)
Database Volume name: /local1/testing/demodb_vinf
Volume Identifier: -5, Size: 308 bytes (1 pages)
Database Volume name: /local1/testing/demodb
Volume Identifier: 0, Size: 2048000 bytes (500 pages)
Database Volume name: /local1/testing/demodb_lginf
Volume Identifier: -4, Size: 165 bytes (1 pages)
Database Volume name: /local1/testing/demodb_bkvinf
Volume Identifier: -3, Size: 132 bytes (1 pages)
```

With the backup information displayed by using the `--list` option, you can check that backup files have been created at the backup level 1 as well as the point where the full backup of backup level 0 has been performed. Therefore, to restore the database in the example, you must prepare backup files for backup levels 0 and 1.

Restore Strategy and Procedure

You must consider the followings before restoring databases.

- **Preparing backup files**
- Identify the directory where the backup and log files are to be stored.
- If the database has been incrementally backed up, check whether an appropriate backup file for each backup level exists.
- Check whether the backed-up CUBRID database and the CUBRID database to be backed up are the same version.
- **Choosing restore method**
- Determine whether to perform a partial or full restore.
- Determine whether or not to perform a restore using incremental backup files.
- Prepare restore tools and devices available.
- **Determining restore point**
- Identify the point in time when the database server was terminated.
- Identify the point in time when the last backup was performed before database failure.
- Identify the point in time when the last commit was made before database failure.

Database Restore Procedure

The following procedure shows how to perform backup and restoration described in the order of time.

- Performs a full backup of demodb which stopped running at 2008/8/14 04:30.
- Performs the first incremental backup of demodb running at 2008/8/14 10:00.
- Performs the first incremental backup of demodb running at 2008/8/14 15:00. Overwrites the first incremental backup file in step 2.
- A system failure occurs at 2008/8/14 15:30, and the system administrator prepares the restore of demodb. Sets the restore time as 15:25, which is the time when the last commit was made before database failure
- The system administrator prepares the full backup file created in Step 1 and the first incremental backup file created in Step 3, restores the demodb database up to the point of 15:00, and then prepares the active and archive logs to restore the database up to the point of 15:25.

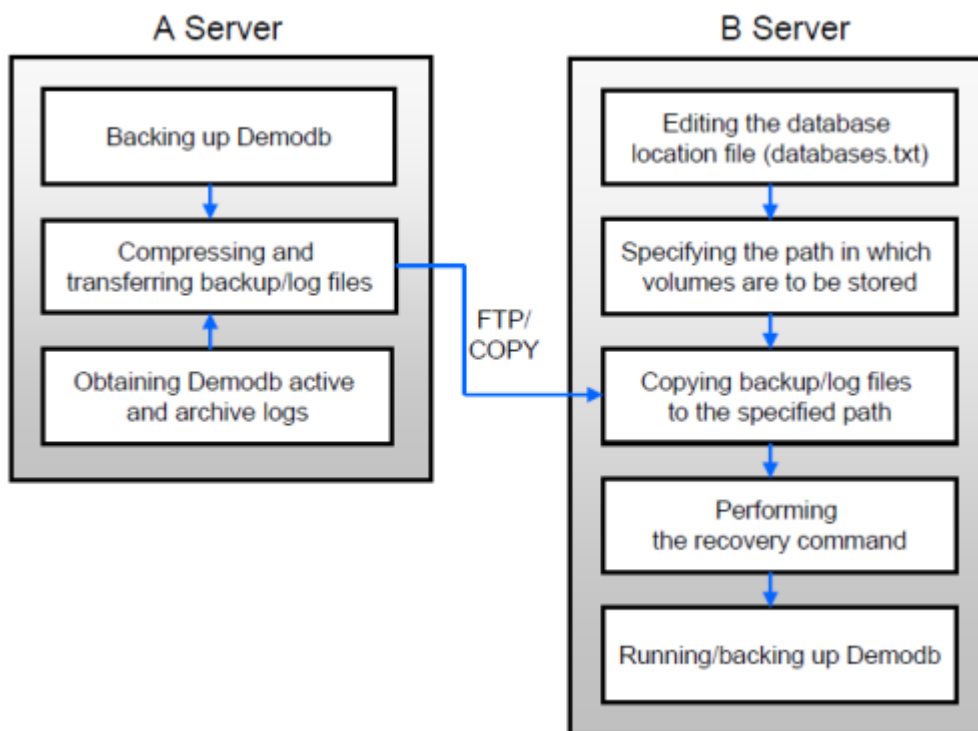
Time	Command	Description
2008/8/14 04:25	<code>cubrid server stop demodb</code>	Shuts down demodb.
2008/8/14 04:30	<code>cubrid backupdb -S -D /home/backup -l 0 demodb</code>	Performs a full backup of demodb in offline mode and creates backup files in the specified directory.
2008/8/14 05:00	<code>cubrid server start demodb</code>	Starts demodb.
2008/8/14 10:00	<code>cubrid backupdb -C -D /home/backup -l 1 demodb</code>	Performs the first incremental backup of demodb online and creates backup files in the specified directory.
2008/8/14 15:00	<code>cubrid backupdb -C -D /home/backup -l 1 demodb</code>	Performs the first incremental backup of demodb online and creates backup files in the specified directory. Overwrites the first incremental backup file created at 10:00.
2008/8/14 15:30		A system failure occurs.
2008/8/14 15:40	<code>cubrid restoredb -l 1 -d 08/14/2008:15:25:00 demodb</code>	Restores demodb based on the full backup file, first incremental backup file, active logs and archive logs. The database is restored to the point of 15:25 by the full and first incremental backup files, the active and archive logs.

Restoring Database to Different Server

The following shows how to back up demodb on server A and restore it on server B with the backed up files.

Backup and Restore Environments

Suppose that demodb is backed up in the `/home/cubrid/db/demodb` directory on server A and restored into `/home/cubrid/data/demodb` on server B.



1. Backing up on server A

Back up demodb on server A. If a backup has been performed earlier, you can perform an incremental backup for data only that have changed since the last backup. The directory where the backup files are created, if not specified in the **-D** option, is created by default in the location where the log volume is stored. The following is a backup command with recommended options. For details on the options, see [Database Backup](#).

```
cubrid backupdb -z -t demodb
```

2. Editing the database location file on Server B

Unlike a general scenario where a backup and restore are performed on the same server, in a scenario where backup files are restored using a different server, you need to add the location information on database restore in the database location file (**databases.txt**) on server B. In the diagram above, it is supposed that **demodb** is restored in the **/home/cubrid/data/demodb** directory on server B (hostname: pmlinux); edit the location information file accordingly and create the directory on server B.

Put the database location information in one single line. Separate each item with a space. The line should be written in [database name]. [data volume path] [host name] [log volume path] format; that is, write the location information of **demodb** as follows:

```
demodb /home/cubrid/data/demodb pmlinux /home/cubrid/data/demodb
```

3. Transferring backup/log files to server B

For a restore, you must prepare a backup file (e.g. demodb_bk0v000) and a backup information file (e.g. demodb_bkvinf) of the database to be backed up. To restore the entire data up to the point of the last commit, you must prepare an active log (e.g. demodb_lgat) and an archive log (e.g. demodb_lgar000). Then, transfer the backup information, active log, and archive log files created on server A to server B. That is, the backup information, active log and archive log files must be located in a directory (e.g. /home/cubrid/temp) on server B.

4. Restoring the database on server B

Perform database restore by calling the **cubrid restoredb** utility from the directory into which the backup, backup information, active log and archive log files which were transferred to server B had been stored. With the **-u** option, demodb is restored in the directory path from the **databases.txt** file.

```
cubrid restoredb -u demodb
```

To call the **cubrid restoredb** utility from a different path, specify the directory path to the backup file by using the **-B** option as follows:

```
cubrid restoredb -u -B /home/cubrid/temp demodb
```

5. Backing up the restored database on server B

Once the restore of the target database is complete, run the database to check if it has been properly restored. For stable management of the restored database, it is recommended to restore the database again on the server B environment.

CUBRID HA

Overview

CUBRID HA

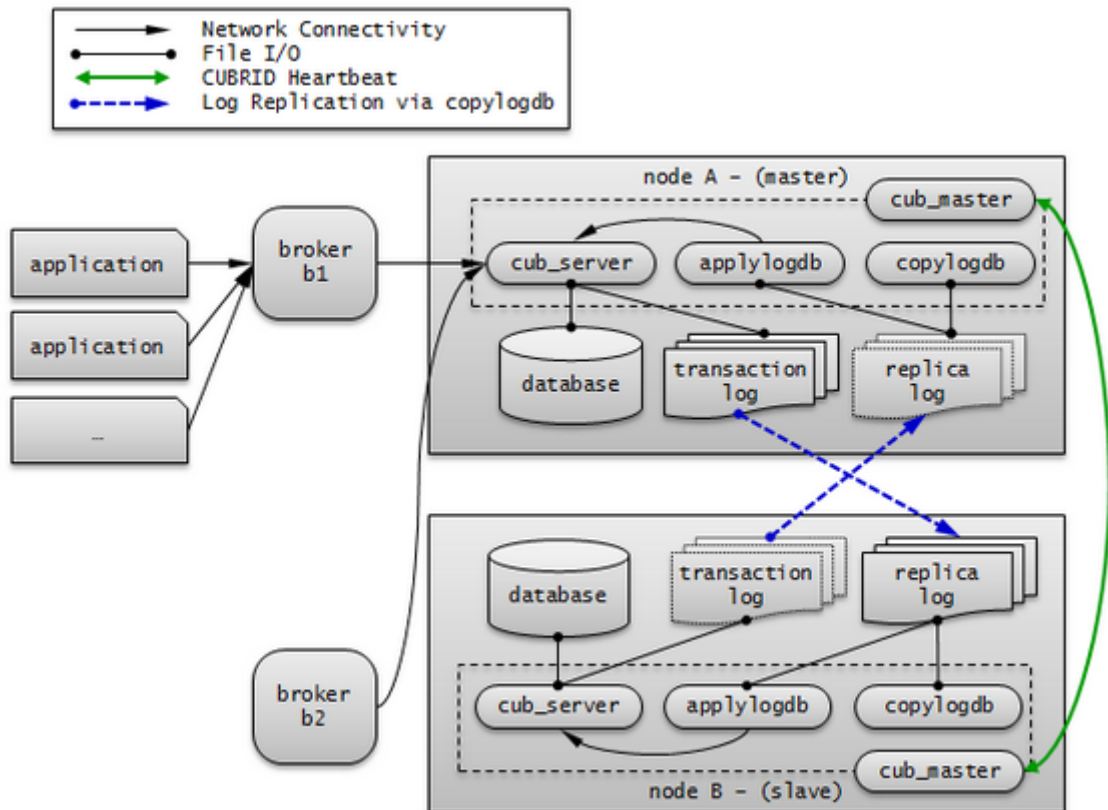
High Availability (HA) refers to a feature to provide uninterrupted service in the event of hardware, software, or network failure. This ability is a critical element in the network computing area where services should be provided 24/7. An HA system consists of more than two server systems, each of which provides uninterrupted services, even when a failure occurs in one of them.

CUBRID HA is an implementation of High Availability. CUBRID HA ensures database synchronization among multiple servers when providing service. When an unexpected failure occurs in the system which is operating services, this feature minimizes the service down time by allowing the other system to carry out the service automatically.

CUBRID HA is in a shared-nothing structure. To synchronize data from an active server to a standby server, CUBRID HA executes the following two steps.

- Transaction log multiplexing: Replicates the transaction logs created by an active server to another node in real time.
- Transaction log reflection: Analyzes replicated transaction logs in real time and reflects the data to a standby server.

CUBRID HA executes the steps described above in order to always maintain data synchronization between an active server and a standby server. For this reason, if an active server is not working properly because of a failure occurring in the master node that had been providing service, the standby server of the slave node provides service instead of the failed server. CUBRID HA monitors the status of the system and CUBRID in real time. It uses heartbeat messages to execute an automatic failover when a failure occurs.



CUBRID HA Concept

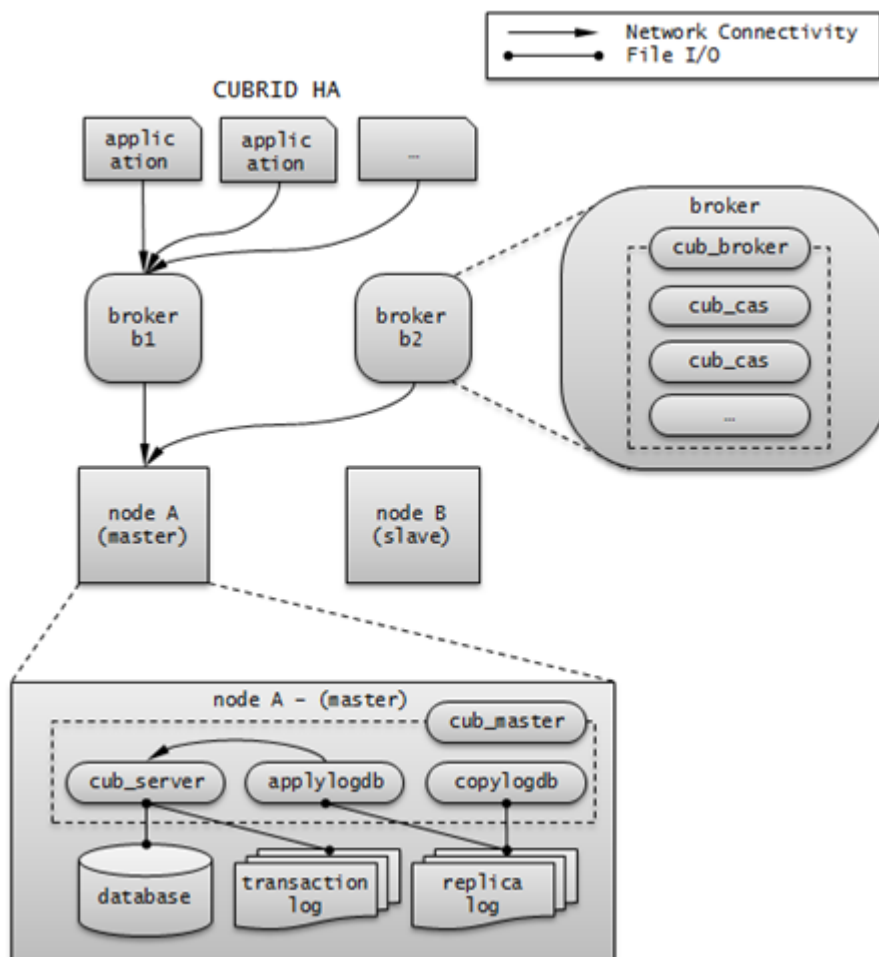
Groups and Nodes

A node is a logical unit that makes up CUBRID HA. It can become one of the following nodes according to its status: master node, slave node, or replica node.

- **Master node** : A node to be replicated. It provides all services which are read, write, etc. using an active server.
- **Slave node** : A node that has the same information as a master node. Changes made in the master node are automatically reflected to the slave node. It provides the read service using a standby server, and a failover will occur when the master node fails.
- **Replica node** : A node that has the same information as a master node. Changes made in the master node are automatically reflected to the replica node. It provides the read service using a standby server, and no failover will occur when the master node fails.

The CUBRID HA group consists of the nodes described above. You can configure the members of this group by using the `ha_node_list` and `ha_replica_list` in the `cubrid.conf` file. Nodes in a group have the same information. They exchange status checking messages periodically and a failover will occur when the master node fails.

A node includes the master process (`cub_master`), the database server process (`cub_server`), the replication log copy process (`copylogdb`), the replication log reflection process (`applylogdb`), etc.





Processes



A CUBRID HA node consists of one master process (`cub_master`), one or more database server processes (`cub_server`), one or more replication log copy processes (`copylogdb`), and one or more replication log reflection processes

(applylogdb). When a database is configured, database server processes, replication log copy processes, and replication log reflection processes will start. Because copy and reflection of a replication log are executed by different processes, the delay in replicating reflections does not affect the transaction that is being executed.

- **Master process (cub_master)** : Exchanges heartbeat messages to control the internal management processes of CUBRID HA.
- **Database server process (cub_server)** : Provides services such as read or write to the user. For details, see [Server](#).
- **Replication log copy process (copylogdb)** : Copies all transaction logs in a group. When the replication log copy process requests a transaction log from the database server process of the target node, the database server process returns the corresponding log. The location of copied transaction logs can be configured in the **REPL_LOG_HOME** of **cubrid-ha**. Use [cubrid applyinfo](#) utility to verify the information of copied replication logs. The replication log copy process has following three modes: SYNC, SEMISYNC, and ASYNC. You can configure it with the **LW_SYNC_MODE** of **cubrid-ha**. For details on these modes, see [Multiplexing Logs](#).



- **Replication log reflection process (applylogdb)** : Reflects the log that has been copied by the replication log copy process to a node. The information of reflected replications is stored in the internal catalog (db_ha_apply_info). You can use the [cubrid applyinfo](#) utility to verify this information.

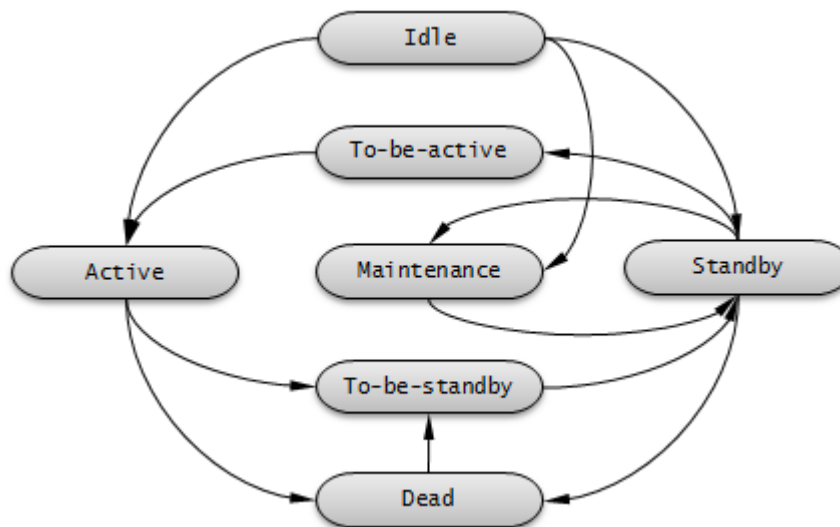


Servers

Here, the word "server" is a logical representation of database server processes. Depending on its status, a server can be either an active server or a standby server.

- **Active server** : A server that belongs to a master node; the status is active. An active server provides all services, including read, write, etc. to the user.
- **Standby server**: A standby server that belongs to a non-master node; the status is standby. A standby server provides only the read service to the user.

The server status changes based on the status of the node. You can use the [cubrid changemode](#) utility to verify server status. The maintenance mode exists for operational convenience and you can change it by using the **cubrid changemode** utility.

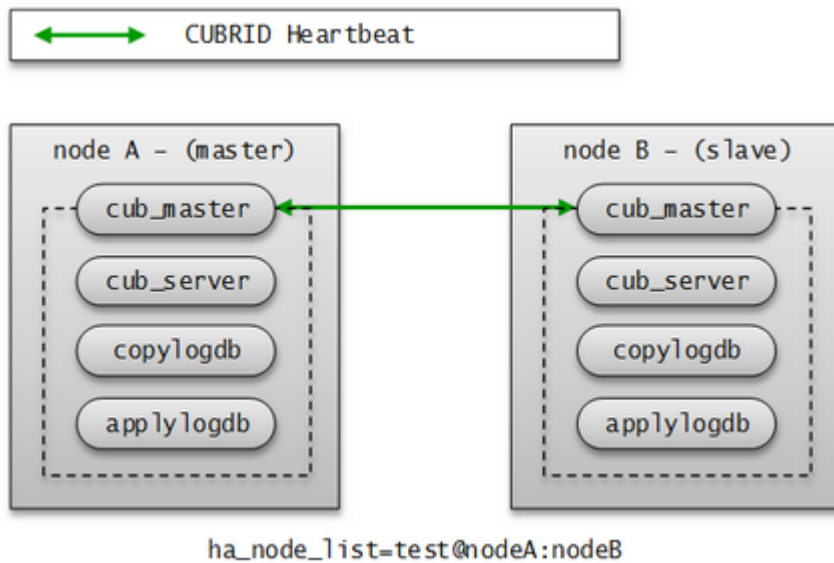


- **active** : The status of servers that run on a master node is usually active. In this status, all services including read, write, etc. are provided.
- **standby** : The status of servers that run on a slave node or a replica node is standby. In this status, only the read service is provided.
- **maintenanc** : The status of servers can be manually changed for operational convenience is maintenance. In this status, only a csq1 can access and no service is provided to the user.
- **to-be-active** : The status in which a standby server will become active for reasons such as failover, etc. is to-be-active. In this status, servers prepare to become active by reflecting transaction logs from the existing master node to its own server.
- **Other** : This status internally used.

heartbeat Message

As a core element to provide HA, it is a message exchanged among master, slave, and replica nodes to monitor the status of other nodes. A master process periodically exchanges heartbeat messages with all other master processes in the group. A heartbeat message is exchanged through the UDP port configured in the **ha_port_id** parameter of **cubrid.conf**. The exchange interval of heartbeat messages is determined by an internally configured value.


When the master node fails, a failover occurs to a slave node.



failover and failback

A failover means that the highest priority slave node automatically becomes a new master node when the original master node fails to provide services due to a failure. A master process calculates scores for all nodes in the CUBRID HA group based on the collected information, promotes slave nodes to master modes when it is necessary, and then notifies the management process of the changes it has made.

A failback means that the previously failed master node automatically becomes a master node back after the failure node is restored. The CUBRID HA does not currently support this functionality.



If a heartbeat message fails to deliver, a failover will occur. For this reason, servers with unstable connection may experience failover even though no actual failures occur. To prevent a failover from occurring in the situation described above, configure `ha_ping_ports`. Configuring `ha_ping_ports` will send a ping message to a node specified in `ha_ping_ports` in order to verify whether the network is stable or not when a heartbeat message fails to deliver. For details on configuring `ha_ping_ports`, see [cubrid_ha.conf](#).

Broker Mode

A broker can access a server with one of the following modes: **Read Write**, **Read Only**, **Slave Only**, or **Preferred Host Read Only**. This configuration value is determined by a user.

A broker finds and connects to a suitable server by trying to establish a connection in the order of server connections; this is, if it fails to establish a connection, it tries another connection to the next server defined until it reaches the last server. If no connection is made even after trying all servers, the broker fails to connect to a server.

For details on how to configure broker mode, see [cubrid_broker.conf](#).

Read Write

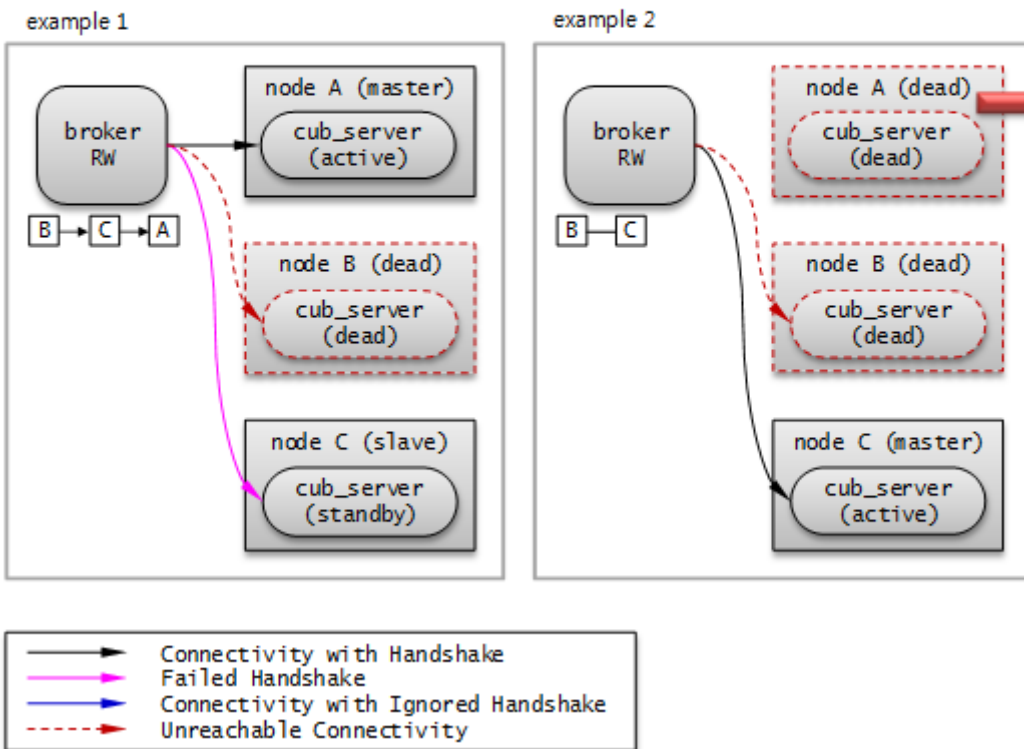
A broker that provides read and write services. This broker is usually connected to an active server. If no active servers exist, this broker will be connected to a standby server. For this reason, a Read Write broker can be temporarily connected to a standby server.

When the broker temporarily establishes a connection to a standby server, it will disconnect itself from the standby server at the end of every transaction so that it can attempt to find an active server at the beginning of the next transaction. When it is connected to the standby server, only read service is available. Any write requests will result in a server error.

The order of server connection is described below:

- The broker tries to establish a connection to an existing server connected (if exists). The active status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence. The active status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence and connects to the first available host.

databases.txt			
#db-name	vol-path	db-host	log-path
tdb	/home/db	node B:node C:node A	/home/db/log



Read Only

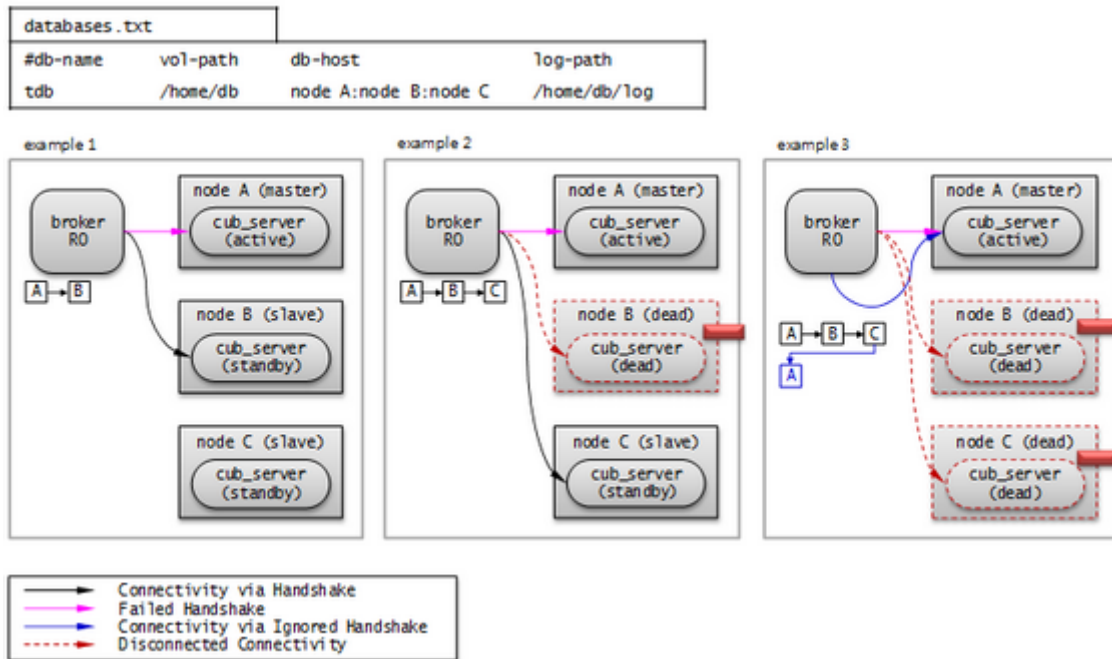
A broker that provides the read service. This broker is connected to a standby server if possible. For this reason, a Read Only broker can be connected to an active server temporarily.

Once it establishes a connection with an active server, it will maintain that connection even if a standby server exists. To disconnect from the active server and reconnect to a standby server, you should execute the **cubrid_broker reset**

command. An error will occur when a Read Only broker receives a write request; therefore, only the read service will be available even if it is connected to an active server.

The order of server connection is described below:

- The broker tries to establish a connection to an existing server connected (if exists). The standby status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence. The standby status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence and connects to the first available host.

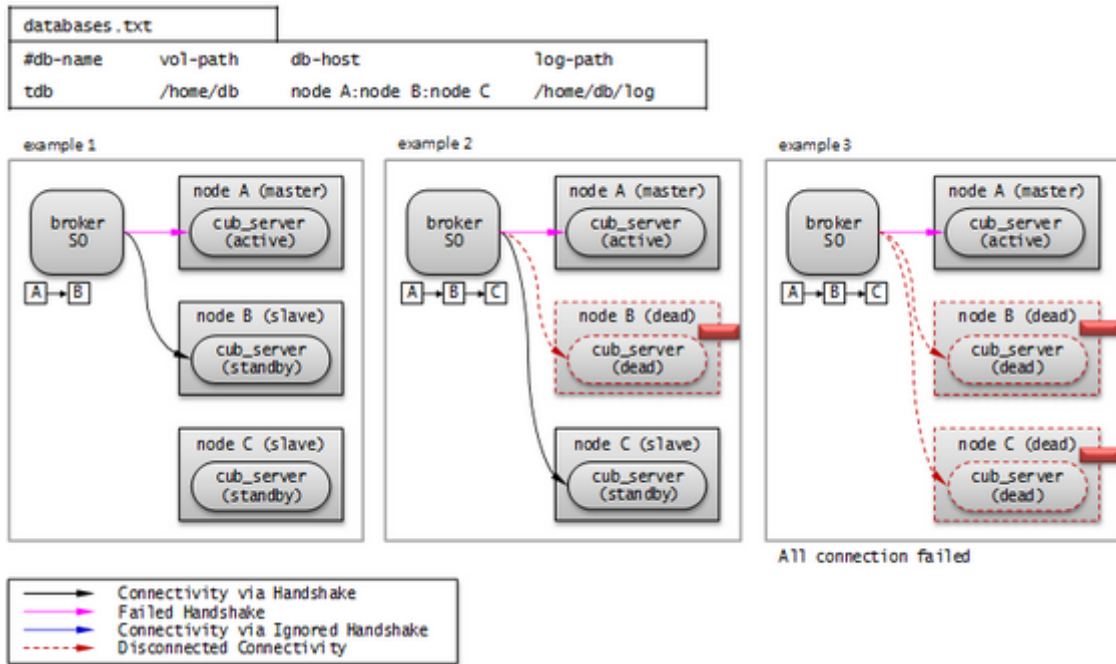


Slave Only

A broker that provides the read service. This broker can only be connected to a standby server. If no standby server exists, no service will be provided.

The order of server connection is described below:

- The broker tries to establish a connection to an existing server connected (if exists). The standby status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence. The standby status of the server means the connection is complete.

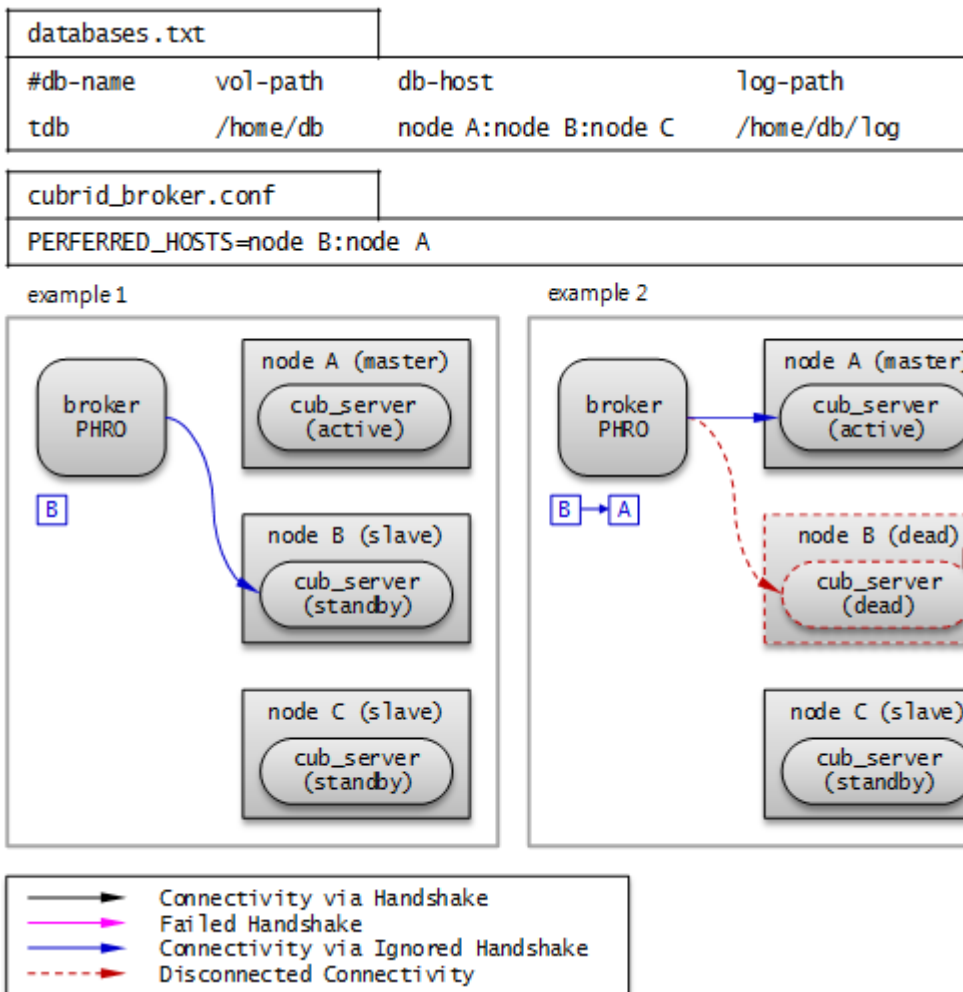


Preferred Host Read Only

A broker that provides the read service. This works in the same manner as the Read Only broker except its server connection order and server selecting criteria. The server connection order and server selecting criteria can be configured in **PREFERRED_HOSTS**. For details on configuring these, see [cubrid_broker.conf](#).

The order of server connection is described below:

- The broker tries to establish a connection to the hosts specified in **PREFERRED_HOSTS** in a sequence and connects to the first available host.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence. The standby status of the server means the connection is complete.
- The broker tries to establish a connection to the hosts specified in the **databases.txt** file in a sequence and connects to the first available host.



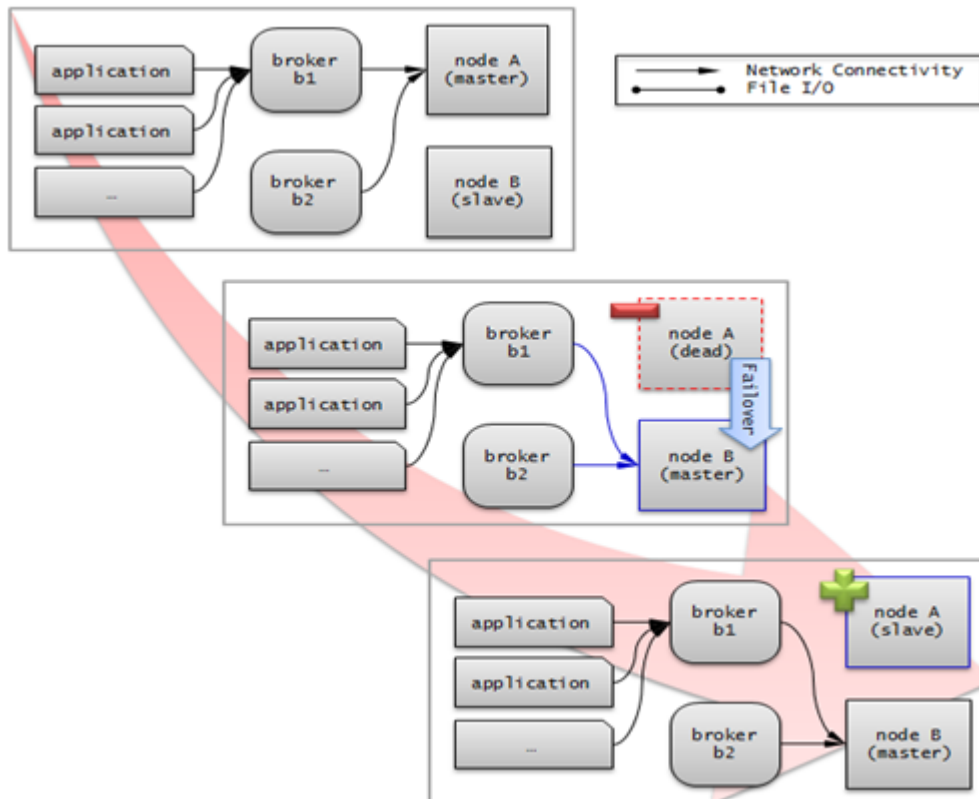
CUBRID HA Features

Duplexing Servers

Duplexing servers is building a system by configuring duplicate hardware equipment to provide CUBRID HA. This method will prevent any interruptions in a server in case of occurring a hardware failure.

Server failover

A broker defines server connection order and connects to a server according to the defined order. If the connected server fails, the broker connects to the server with the next highest priority. This requires no processing in the application side. The actions taken when the broker connects to another server may differ according to the current mode of the broker. For details on the server connection order and configuring broker mode, see [cubrid_broker.conf](#).



Server failback

CUBRID HA does not automatically support server failback. Therefore, to manually apply failback, restore the master node that has been abnormally terminated and run it as a slave node, terminate the node that has become the master from the slave due to failover, and finally, change the role of each node again.

For example, when nodeA is the master and nodeB is the slave, nodeB becomes the master and nodeA becomes the slave after a failover. After terminating nodeB (**cubrid heartbeat stop**) check (**cubrid heartbeat status**) whether the status of nodeA has become active. Start (**cubrid heartbeat start**) nodeB and it will become the slave.

Duplexing Brokers

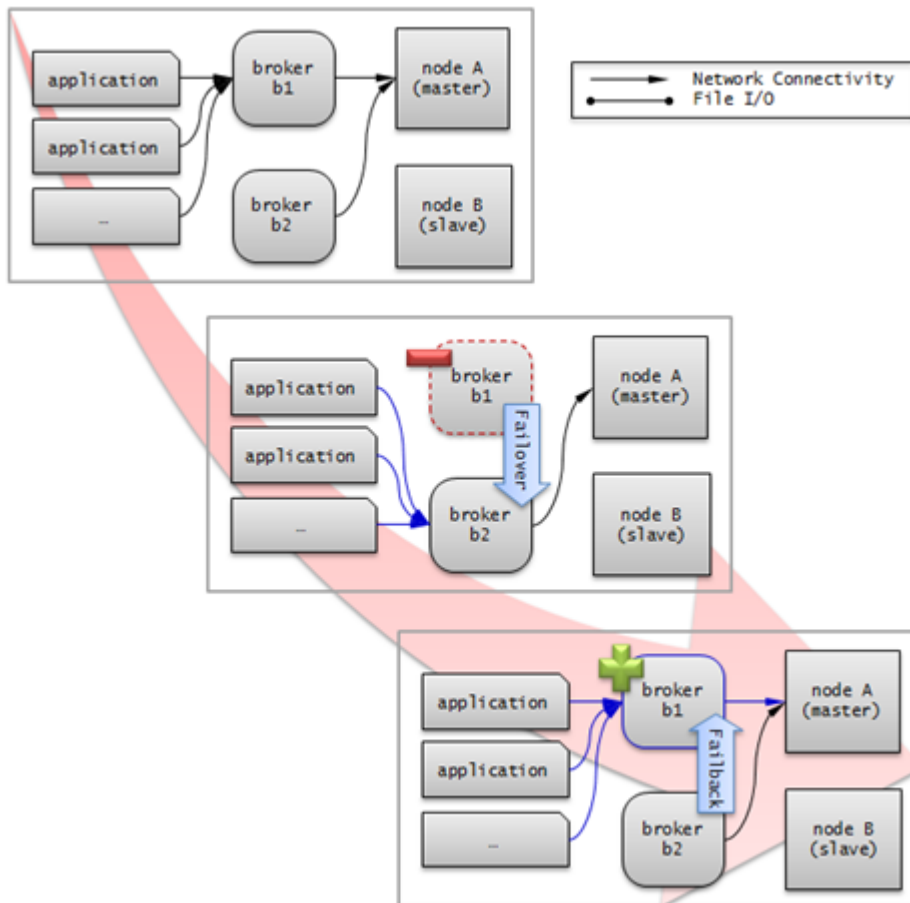
As a 3-tier DBMS, CUBRID has middleware called the broker which relays applications and database servers. To provide HA, the broker also requires duplicate hardware equipment. This method will prevent any interruptions in a broker in case of occurring a hardware failure.

The configuration of broker redundancy is not determined by the configuration of server redundancy; it can be user-defined. In addition, it can be separated by piece of individual equipment.

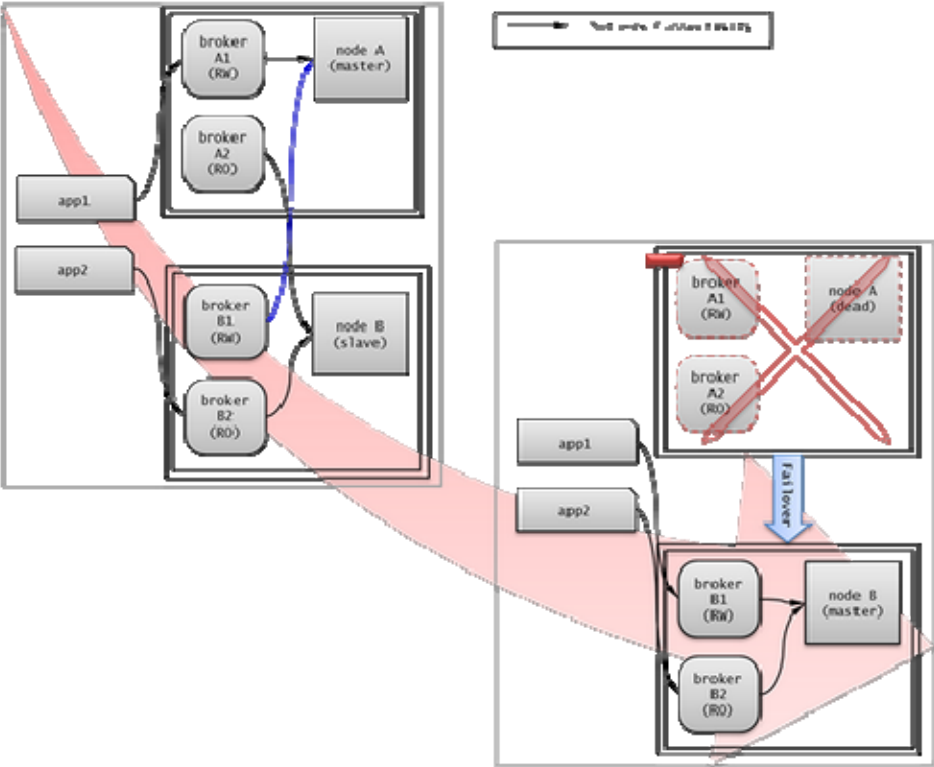
To use the failover and failback functionalities of a broker, the **althosts** attribute must be added to the connection URL of the JDBC, CCI, or PHP. For a description of this, see JDBC Configuration, CCI Configuration and PHP Configuration.

To set a broker, configure the **cubrid_broker.conf** file. To set the order of failovers of a database server, configure the **databases.txt** file. For more information, see Broker Configuration.

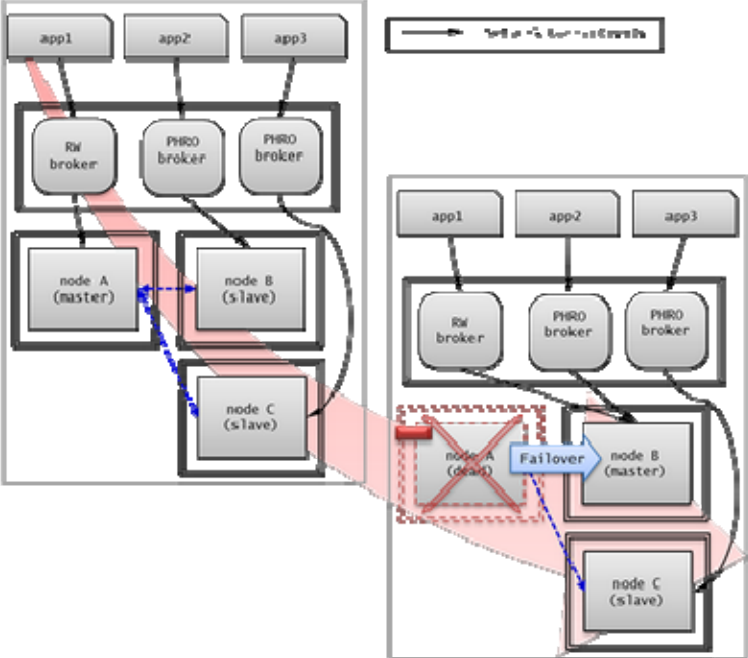
The following is an example in which two Read Write (RW) brokers are configured. When the first connection broker of the application URL is set to broker B1 and the second connection broker to broker B2, the application connects to broker B2 when it cannot connect to broker B1. When broker B1 becomes available again, the application reconnects to broker B1.



The following is an example in which the Read Write (RW) broker and the Read-Only (RO) broker are configured in each piece of equipment of the master node and the slave node. First, the app1 and the app2 URL connect to broker A1 (RW) and broker B2 (RO), respectively. The second connection (althosts) is made to broker A2 (RW) and broker B1 (RO). When equipment that includes nodeA fails, app1 and the app2 connect to the broker that includes nodeB.



The following is an example of a configuration in which broker equipment includes one Read Write broker (master node) and two Preferred Host Read Only brokers (slave nodes). The Preferred Host Read Only brokers are connected to nodeB and nodeC to distribute read load.



Broker failover

Broker failover is not automatically failed over by the settings of system parameters. It is available in the JDBC, CCI, and PHP applications only when broker hosts are configured in the **althosts** of the connection URL. Applications connect to the broker with the highest priority. When the connected broker fails, the application connects to the broker

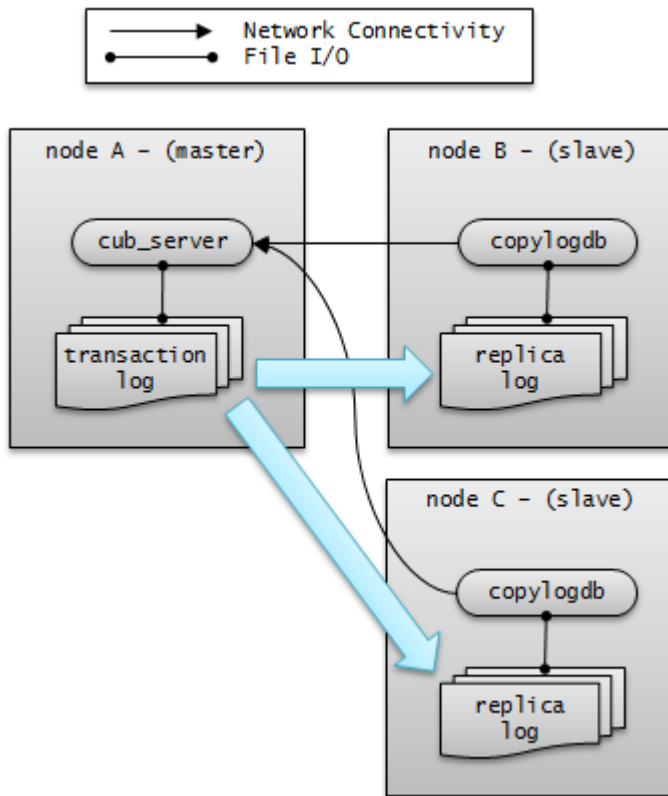
with the next highest priority. Configuring the **althosts** of the connection URL is the only necessary action, and it is processed in the JDBC, CCI, and PHP drivers.

Broker failback

If the failed broker is recovered after a failover, the connection to the existing broker is terminated and a new connection is established with the recovered broker which has the highest priority. This requires no processing in the application side as it is processed within the JDBC, CCI, and PHP drivers. Execution time of failback depends on the value configured in JDBC connection URL. For details, see [JDBC Configuration](#).

Log Multiplexing

CUBRID HA keeps every node in the CUBRID HA group with the identical structure by copying and reflecting transaction logs to all nodes included in the CUBRID HA group. As the log copy structure of CUBRID HA is a mutual copy between the master and the slave nodes, it has a disadvantage of increasing the size of a log volume. However, it has an advantage of flexibility in terms of configuration and failure handling, comparing to the chain-type copy structure.



The transaction log copy modes include **SYNC**, **SEMISYNC**, and **ASYN**C. This value can be configured by the user in [cubrid_ha.conf](#) file.

SYNC Mode

When transactions are committed, the created transaction logs are copied to the slave node and stored as a file. The transaction commit is complete after receiving a notice on its success. Although the time it takes to execute commit in this mode may be longer than that in other modes, this is the safest method because the copied transaction logs are always guaranteed to be reflected to the standby server even if a failover occurs.

SEMISYNC Mode

When transactions are committed, the created transaction logs are copied to the slave node and stored as a file according to the internally optimized interval. The transaction commit is complete after receiving a notice of its success. The committed transactions in this mode are guaranteed to be reflected to the slave node sometime in the future.

Because SEMISYNC mode does not always store replication logs as a file, the execution time of commit can decrease, comparing to the SYNC mode. However, data synchronization between nodes may be delayed because replication logs are not reflected until it is stored as a file.

ASYNC Mode

When transactions are committed, commit is complete without verifying the transfer of transaction logs to a slave node. Therefore, it is not guaranteed that committed transactions are reflected to a slave node in a master node side.

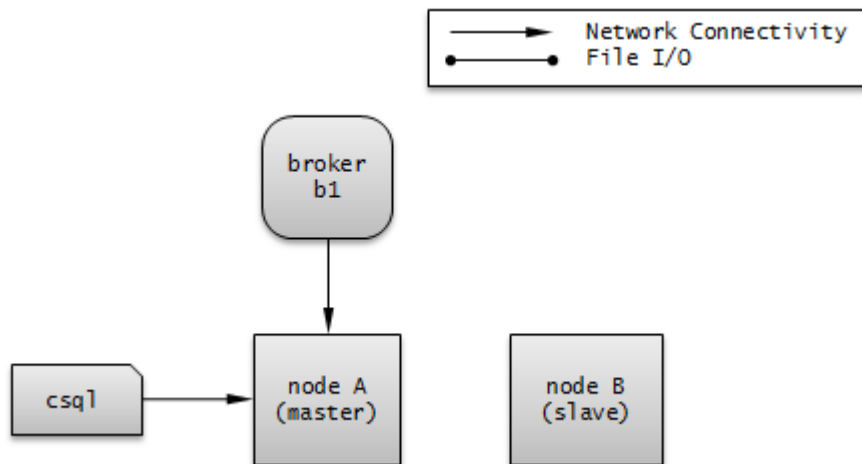
Although ASYNC mode provides a better performance as it has almost no delay when executing commit, there may be data inconsistency in its nodes.

Quick Start

Preparation

Structure Diagram

The diagram below aims to help users who are new to CUBRID HA, by explaining a simple procedure of the CUBRID HA configuration.



Specifications

Linux and CUBRID version 2008 R2.2 or higher must be installed on the equipment to be used as the master and the slave nodes.

Specifications of Configuring the CUBRID HA Equipment

	CUBRID Version	OS
For master nodes	CUBRID 2008 R2.2 or higher	Linux
For slave nodes	CUBRID 2008 R2.2 or higher	Linux

Note This document describes the HA configuration in CUBRID 2008 R4.1 Patch 2 or higher versions. Note that the previous versions have different settings. For example, `cubrid_ha.conf` is only available in CUBRID 2008 R4.0 or higher. `ha_make_slavedb.sh` describes CUBRID 2008 R4.1 Patch 2 or higher.

Creating Databases and Configuring Servers

Creating Databases

Create databases to be included in CUBRID HA at each node of the CUBRID HA in the same manner. Modify the options for database creation as needed.

```
[master]$ cd $CUBRID_DATABASES
[master]$ mkdir testdb
[master]$ cd testdb
[master]$ mkdir log
[master]$ cubrid createdb -L ./log testdb
Creating database with 5000 pages.

CUBRID 2008 R4.1

[master]$
```

cubrid.conf

Ensure **ha_mode** of `$CUBRID/conf/cubrid.conf` in every CUBRID HA node has the same value. Especially, take caution when configuring the **log_max_archives** and **force_remove_log_archives** parameters (logging parameters) and the **ha_mode** parameter (HA parameter).

```
# Service parameters
[service]
service=server,broker,manager

# Server parameters
server=testdb
data_buffer_size=512M
log_buffer_size=4M
sort_buffer_size=2M
max_clients=100
cubrid_port_id=1523
db_volume_size=512M
log_volume_size=512M

# Adds when configuring HA (Logging parameters)
log_max_archives=100
force_remove_log_archives=no

# Adds when configuring HA (HA mode)
ha_mode=on
```

Configuring cubrid_ha.conf

Ensure **ha_port_id**, **ha_node_list**, **ha_db_list** of `$CUBRID/conf/cubrid_ha.conf` in every CUBRID HA node has the same value.

```
[common]
ha_port_id=12345
ha_node_list=cubrid@nodeA:nodeB
ha_db_list=testdb
ha_copy_sync_mode=sync:sync
ha_apply_max_mem_size=500
```

Starting and Verifying CUBRID HA

Starting CUBRID HA

Execute the **cubrid heartbeat start** at each node in the CUBRID HA group. Note that the node executing **cubrid heartbeat start** first will become a master node.

- Master node

```
[master]$ cubrid heartbeat start
```

- Slave node


```
[slave]$ cubrid heartbeat start
```

Verifying CUBRID HA Status

Execute **cubrid heartbeat status** at each node in the CUBRID HA group to verify its configuration status.

```
[master]$ cubrid heartbeat status
@cubrid heartbeat list
HA-Node Info (current master-node-name, state master)
  Node slave-node-name (priority 2, state slave)
  Node master-node-name (priority 1, state master)
HA-Process Info (master 9289, state master)
  Applylogdb testdb@localhost:/home1/cubrid1/DB/testdb_slave.cub (pid 9423, state
registered)
  Copylogdb testdb@slave-node-name:/home1/cubrid1/DB/testdb_slave.cub (pid 9418, state
registered)
  Server testdb (pid 9306, state registered and active)

[master]$
```

Use the **cubrid changemode** utility at each node in the CUBRID HA group to verify the status of the server.

- Master node

```
[master]$ cubrid changemode testdb@localhost
The server 'testdb@localhost''s current HA running mode is active.
```

- Slave node

```
[slave]$ cubrid changemode testdb@localhost
The server 'testdb@localhost''s current HA running mode is standby.
```

Verifying the CUBRID HA Operation

Ensure the change has been correctly reflected to the standby server of the slave node after writing the active server of the master node. Note that a primary key must exist when creating a table.

- Master node

```
[master]$ csq1 -u dba demodb@localhost -c "create table abc(a int, b int, c int,
primary key(a));"
[master]$ csq1 -u dba demodb@localhost -c "insert into abc values (1,1,1);"
[master]$
```

- Slave node

```
[slave]$ csq1 -u dba testdb@localhost -l -c "select * from abc;"
=== <Result of SELECT Command in Line 1> ===
<00001> a: 1
        b: 1
        c: 1
[slave]$
```

Configuring and Starting Broker, and Verifying the Broker Status

Configuring the Broker

To provide normal service during a database failover, it is necessary to configure an available node in the **db-host** of **databases.txt**. And **ACCESS_MODE** in the **cubrid_broker.conf** file must be specified; if it is omitted, the default value is configured to Read Write mode. If you want to divide into a separate device, you must configure **cubrid_broker.conf** and **databases.txt** in the device.

- databases.txt

```
#db-name          vol-path          db-host          log-path          lob-base-path
demodbtestdb     /home1/cubrid1/DBCUBRID/demodbtestdb  nodeA:nodeB     /home1/c
ubrid1/DBCUBRID/demodbtestdb/log file:/home1/cubrid1/CUBRID/testdb/lob
```

- cubrid_broker.conf

```
[%testdb RWbroker]
SERVICE          =ON
BROKER PORT       =33000
MIN_NUM_APPL_SERVER =5
MAX_NUM_APPL_SERVER =40
```

```

APPL_SERVER_SHM_ID =33000
LOG_DIR            =log/broker/sql_log
ERROR_LOG_DIR     =log/broker/error_log
SQL_LOG           =ON
TIME_TO_KILL      =120
SESSION_TIMEOUT   =300
KEEP_CONNECTION   =AUTO
CCI_DEFAULT_AUTOCOMMIT =ON

# broker mode parameter
ACCESS_MODE       =RW

```

Starting Broker and Verifying its Status

A broker is used to access applications such as JDBC, CCI or PHP. Therefore, to simply test server redundancy, execute the CSQL interpreter that is directly connected to the server processes, without having to start a broker. To start a broker, execute **cubrid broker start**. To stop it, execute **cubrid broker stop**.

The following example shows how to execute a broker from the master node

```

[master]$ cubrid broker start
@ cubrid broker start
++ cubrid broker start: success
[master]$ cubrid broker status
@ cubrid broker status
% testdb RWbroker - cub cas [9531,33000] /home1/cubrid1/CUBRID/log/broker//testdb.access
/home1/cubrid1/CUBRID/log/broker//testdb.err
JOB_QUEUE:0, AUTO_ADD_APPL_SERVER:ON, SQL_LOG_MODE:ALL:100000
LONG_TRANSACTION_TIME:60.00, LONG_QUERY_TIME:60.00, SESSION_TIMEOUT:300
KEEP_CONNECTION:AUTO, ACCESS_MODE:RW
-----
ID   PID   QPS   LQS  PSIZE STATUS
-----
1   9532   0     0  48120  IDLE

```

Configuring Applications

Specifies the host name and port for an application to connect in the connection URL. The **althosts** attribute defines the broker where the next connection will be made when the connection to a broker fails. The following is an example of a JDBC program. For more information on CCI and PHP, see [CCI Configuration](#) and [PHP Configuration](#).

```

Connection connection =
DriverManager.getConnection("jdbc:CUBRID:nodeA broker:33000:testdb::?charset=utf-
8&althosts=nodeB_broker:33000", "dba", "");

```

Environment Configuration

cubrid.conf

The **cubrid.conf** file that has general information on configuring CUBRID is located in the **\$CUBRID/conf** directory. This section explains the **cubrid.conf** parameters used by CUBRID HA.

ha_mode

A parameter used to configure whether to use CUBRID HA. The default value is **off**.

- **off** : CUBRID HA is not used.
- **on** : CUBRID HA is used. Failover is supported for its node.
- **replica** : CUBRID HA is used. Failover is not supported for its node.

The **ha_mode** parameter can be re-configured in the **[@<database>]** section; however, only **off** can be entered in the case. An error is displayed if a value other than **off** is entered in the **[@<database>]** section.

If **ha_mode** is **on**, the CUBRID HA values are configured by reading **cubrid_ha.conf**.

This parameter cannot be modified dynamically. To modify the value of this parameter, you must restart the corresponding node.

log_max_archives

This parameter configures the minimum number of archive log files to be archived. The minimum value is 0 and the default is `INT_MAX`. The performance of the parameter is affected by `force_remove_log_archives`.

The existing archive log files to which the activated transaction refers or the archive log files of the master node not reflected to the slave node in HA environment will not be deleted. For details, see the following `force_remove_log_archives`.

For details about `log_max_archives`, see [Logging-Related Parameters](#).

force_remove_log_archives

It is recommended to always maintain the archive logs needed to perform the HA related process by setting a value for `force_remove_log_archives` to no in order to implement the HA environment by setting a value for `ha_mode` to on.

If you set for `force_remove_log_archives` to yes, the archive log files which will be used in the HA related process can be deleted, and this may lead to an inconsistency between replicated databases. If you want to maintain free disk space even with this risk, set for `force_remove_log_archives` to yes.

For details about `force_remove_log_archives`, see [Logging-Related Parameters](#).

max_clients

A parameter used to specify the maximum number of clients to be connected to a database server simultaneously. The default is `50`.

Because the replication log copy process and the replication log reflection process are started by default if CUBRID HA is used, you must configure the value to twice the number of all nodes in the CUBRID HA group, except the corresponding node. Furthermore, you must consider the case in which a client that was connected to another node at the time of failover attempts to connect to that node.

The Parameters That Must Have the Same Value for All Nodes

- **log_buffer_size**: The size of a log buffer. This must be same for all nodes, as it affects the protocol between `copylogdb` that duplicate the server and logs.
- **log_volume_size**: The size of a log volume. In CUBRID HA, the format and contents of a transaction log are the same as that of the replica log. Therefore, the parameter must be same for all nodes. If each node creates its own DB, the `cubrid createdb` options (`--db-volume-size`, `--db-page-size`, `--log-volume-size`, `--log-page-size`, etc.) must be the same.
- **cubrid_port_id**: The TCP port number for creating a server connection. It must be same for all nodes in order to connect `copylogdb` that duplicate the server and logs.
- **HA-related parameters**: HA parameters included in `cubrid_ha.conf` must be identical by default. However, the following parameters can be set differently according to the node.
 - The `ha_mode` parameter in replica node
 - The `ha_copy_sync_mode` parameter
 - The `ha_ping_hosts` parameter

Example

The following example shows how to configure `cubrid.conf`. Please take caution when configuring `log_max_archives` and `force_remove_log_archives` (logging-related parameters), and `ha_mode` (an HA-related parameter).

```
max_clients=200
# Service Parameters
[service]
service=server,broker,manager

# Server Parameters
server=testdb
data_buffer_size=512M
log_buffer_size=4M
sort_buffer_size=2M
```

```

max_clients=200
cubrid port id=1523
db_volume_size=512M
log_volume_size=512M

# Adds when configuring HA (Logging parameters)
log_max_archives=100
force_remove_log_archives=no

# Adds when configuring HA (HA mode)
ha_mode=on
log_max_archives=100

```

cubrid_ha.conf

The **cubrid_ha.conf** file that has generation information on CUBRID HA is located in the **\$CUBRID/conf** directory.

ha_node_list

A parameter used to specify the group name to be used in the CUBRID HA group and the host name of member nodes in which failover is supported. The group name is separated by **@**. The name before **@** is for the group, and the names after **@** are for host names of member nodes. A colon (**:**) is used to separate individual host names. The default is **localhost@localhost**.

The host name of the member nodes specified in this parameter cannot be replaced with the IP. When a host name is used, the name must be registered in **/etc/hosts**. A node in which the **ha_mode** value is set to **on** must be specified in **ha_node_list**. The value of the **ha_node_list** of all nodes in the CUBRID HA group must be identical. When a failover occurs, a node becomes a master node in the order specified in the parameter.

This parameter can be modified dynamically. If you modify the value of this parameter, you must execute [cubrid heartbeat reload](#) to apply the changes.

ha_replica_list

A parameter used to specify the group name to be used in the CUBRID HA group and the host name of member nodes in which failover is not supported. The group name is separated by **@**. The name before **@** is for the group, and the names after **@** are for host names of member nodes. A colon (**:**) is used to separate individual host names. The default is **NULL**.

The group name must be identical to the name specified in **ha_node_list**. The host names of member nodes and the host names of nodes specified in this parameter must be registered in **/etc/hosts**. A node in which the **ha_mode** value is set to **replica** must be specified in **ha_replica_list**. The **ha_node_list** values of all nodes in the CUBRID HA group must be identical.

This parameter can be modified dynamically. If you modify the value of this parameter, you must execute [cubrid heartbeat reload](#) to apply the changes.

ha_port_id

A parameter used to specify the UDP port number; the UDP port is used to detect failure when exchanging heartbeat messages. The default is **59,901**.

If a firewall exists in the service environment, the firewall must be configured to allow the configured port to pass through it.

ha_ping_hosts

A parameter used to specify the host which verifies whether or not a failover occurs due to unstable network when a failover has started in a slave node. The default is **NULL**.

The host name of the member nodes specified in this parameter can be replaced with the IP. When a host name is used, the name must be registered in **/etc/hosts**.

Configuring this parameter can prevent split-brain, a phenomenon in which two master nodes simultaneously exist as a result of the slave node erroneously detecting an abnormal termination of the master node due to unstable network status and then promoting itself as the new master. When specifying multiple hosts, separate each host with a colon (:).

ha_copy_sync_mode

A parameter used to specify the mode of storing the transaction log copy. The default is **SYNC**.

The value can be one of the followings: **SYNC**, **SEMISYNC**, or **ASYNC**. The number of values must be the same as the number of nodes specified in **ha_node_list**. They must be ordered by the specified value. You can specify multiple nodes by using a colon (:). The replica node is always working in **ASYNC** mode regardless of this value.

For details, see [Multiplexing Logs](#).

ha_copy_log_base

A parameter used to specify the location of storing the transaction log copy. The default is **\$CUBRID_DATABASES**.

For details, see [Multiplexing Logs](#).

ha_db_list

A parameter used to specify the name of the database that will run in CUBRID HA mode. The default is **NULL**. You can specify multiple databases by using a comma (,).

ha_apply_max_mem_size

A parameter used to specify the value of maximum memory that the replication log reflection process of CUBRID HA can use. The default is **500** (unit: MB). When the value is larger than the size allowed by the system, memory allocation fails and the HA replication reflection process may malfunction. For this reason, you must check whether or not the memory resource can handle the specified value before setting it.

ha_applylogdb_ignore_error_list

A parameter used to specify to ignore any error that occurs in the replication process of CUBRID HA. When specifying errors to be ignored, separate each error with a comma (,). This value has a high priority. Therefore, when this value is the same as the value of the **ha_applylogdb_retry_error_list** parameter or the error code of "List of Retry Errors," the values of the **ha_applylogdb_retry_error_list** parameter or the error code of "List of Retry Errors" are ignored and the tasks that cause the error are not retried. For "List of Retry Errors," see the description of **ha_applylogdb_retry_error_list** below.

ha_applylogdb_retry_error_list

A parameter used to specify repeatedly retry a task that caused an error in the replication log reflection process of CUBRID HA until the task succeeds. When specifying errors to be retried, separate each error with a comma (,). The following table shows the default "List of Retry Errors." If these values exist in **ha_applylogdb_ignore_error_list**, the error will be overridden.

List of Retry Errors

Error	Error Code
ER_LK_UNILATERALLY_ABORTED	-72
ER_LK_OBJECT_TIMEOUT_SIMPLE_MSG	-73
ER_LK_OBJECT_TIMEOUT_CLASS_MSG	-74
ER_LK_OBJECT_TIMEOUT_CLASSOF_MSG	-75
ER_LK_PAGE_TIMEOUT	-76
ER_PAGE_LATCH_TIMEDOUT	-836
ER_PAGE_LATCH_ABORTED	-859

ER_LK_OBJECT_DL_TIMEOUT_SIMPLE_MSG	-966
ER_LK_OBJECT_DL_TIMEOUT_CLASS_MSG	-967
ER_LK_OBJECT_DL_TIMEOUT_CLASSOF_MSG	-968
ER_LK_DEADLOCK_CYCLE_DETECTED	-1021

Example

The following example shows how to configure **cubrid_ha.conf**.

```
[common]
ha_node_list=cubrid@masterdb.cub:slavedb.cub
ha_db_list=testdb
ha_copy_sync_mode=sync:sync
ha_apply_max_mem_size=500
```

Remark

The following example shows how to configure the value of `/etc/hosts` (a host name of a member node: `masterdb.cub`, IP: `192.168.0.1`).

```
127.0.0.1 localhost.localdomain localhost
192.168.0.1 masterdb.cub
```

cubrid_broker.conf

The **cubrid_broker.conf** file that has general information on configuring CUBRID broker is located in the **SCUBRID/conf** directory. This section explains the parameters of **cubrid_broker.conf** that are used by CUBRID HA.

ACCESS_MODE

A parameter used to specify the mode of a broker. The default is **RW**.

Its value can be one of the followings: **RW** (Read Write), **RO** (Read Only), **SO** (Slave Only), or **PHRO** (Preferred Host Read Only). For details, see [Broker Mode](#).

PREFERRED_HOSTS

A parameter used only when the **ACCESS_MODE** parameter value is **PHRO**. The default value is **NULL**.

You can specify multiple nodes by using a colon (:). First, it tries to connect to host in the following order: host specified in the **PREFERRED_HOSTS** parameter first and host specified in **SCUBRID_DATABASES/databases.txt** second.

Example

The following example shows how to configure **cubrid_broker.conf**.

```
[%PHRO broker]
SERVICE                =ON
BROKER_PORT              =33000
MIN_NUM_APPL_SERVER     =5
MAX_NUM_APPL_SERVER     =40
APPL_SERVER_SHM_ID      =33000
LOG_DIR                  =log/broker/sql log
ERROR_LOG_DIR            =log/broker/error log
SQL_LOG                  =ON
TIME_TO_KILL             =120
SESSION_TIMEOUT         =300
KEEP_CONNECTION          =AUTO
CCI_DEFAULT_AUTOCOMMIT  =ON

# Broker mode setting parameter
ACCESS_MODE              =PHRO
PREFERRED_HOSTS          =nodeA:nodeB:nodeC
```

databases.txt

The **databases.txt** file that has information on servers to be connected by a broker and their order is located in the **\$CUBRID_DATABASES** (if not specified, **\$CUBRID/databases**) directory; the information can be configured by using **db_hosts**. You can specify multiple nodes by using a colon (:).

The following example shows how to configure **databases.txt**.

```
#db-name      vol-path      db-host      log-path      lob-base-path
testdb01      /home/cubrid/DB/testdb
masterdb.cub:slavedb.cub  /home/cubrid/DB/testdb01/log  file:/home/cubrid/DB/testdb/lob
testdb02      /home/cubrid/DB/testdb02
masterdb.cub:slavedb.cub  /home/cubrid/DB/testdb02/log  file:/home/cubrid/DB/testdb02/lob
```

JDBC Configuration

To use CUBRID HA in JDBC, you must specify the connection information of another broker (**secondary_broker**) to be connected when a failure occurs in broker (**primary_broker**). The attribute configured for CUBRID HA is **althosts** which represents information of more than one broker nodes to be connected. For details, see "API Reference > JDBC API > JDBC Programming > Connection Configuration."

The following example shows how to configure JDBC:

```
Connection connection =
DriverManager.getConnection("jdbc:CUBRID:primary_broker:33000:testdb::?charset=utf-8&althosts=secondary_broker:33000", "dba", "");
```

CCI Configuration

To use CUBRID HA in CCI, you must use the **cci_connect_with_url** function which additionally allows specifying connection information in connection URL; the connection information is used when a failure occurs in broker. The attribute configured for CUBRID HA is **althosts** which represents information of one or more broker nodes to be connected.

The following example shows how to configure CCI.

```
con = cci_connect_with_url
("cci:CUBRID:primary_broker:33000:testdb::?althosts=secondary_broker:33000", "dba", NULL);
if (con < 0)
{
    printf ("cannot connect to database\n");
    return 1;
}
```

PHP Configuration

To use the functions of CUBRID HA in PHP, connect it to the broker by using **cubrid_connect_with_url**, which is used to specify the connection information of the failover broker in the connection URL. The attribute specified for CUBRID HA is **althosts**, the information on one or more broker nodes to connect when a failover occurs.

The following example shows how to configure PHP.

```
<?php
$con = cubrid_connect_with_url
("cci:CUBRID:primary_broker:33000:testdb::?althosts=secondary_broker:33000", "dba", NULL);
if ($con < 0)
{
    printf ("cannot connect to database\n");
    return 1;
}
?>
```

Running and Monitoring

Utilities of cubrid heartbeat

start

This utility is used to start all components of CUBRID HA in the node (database server process, replication log copy process, replication log reflection process).

Note that a master node or a slave node is determined based on the execution order of **cubrid heartbeat start**.

How to execute the command is as shown below.

```
$ cubrid heartbeat start
$
```

cubrid server start only starts cub_server process of the database, regardless of HA mode configuration. If you want to start all HA related processes, you can execute **cubrid heartbeat start**.

stop

This utility is used to stop all components of CUBRID HA in the node (database server process, replication log copy process, replication log reflection process). The node that executes this command stops and a failover occurs to the next slave node according to the CUBRID HA configuration.

How to use this utility is as shown below.

```
$ cubrid heartbeat stop
$
```

cubrid server stop only starts cub_server process of the database, regardless of HA mode configuration. The database does not restart, and failover does not occur. If you want to stop all HA related processes, you can execute **cubrid heartbeat stop**.

reload

This utility is used to retrieve the CUBRID HA information again, and it starts or stops the CUBRID HA components according to new CUBRID HA configuration. You can modify the information of **ha_node_list** and **ha_replica_list**. If an error occurs during the command execution, the node will stop.

How to use this utility is as shown below.

```
$ cubrid heartbeat reload
$
```

deact

This utility is used to exclude the node from the CUBRID HA group. A node in which **deact** is executed will be excluded from the CUBRID HA group and the CUBRID HA components will stop. The status of this node is displayed as **unknown** when you verify it by using **cubrid heartbeat status**. You can include the node to the CUBRID HA group back by executing **act**.

It is recommended that this command be used only when it is unavoidable.

How to use this utility is as shown below.

```
$ cubrid heartbeat deact
$
```

deregister

This utility is used to terminate **applylogdb** or **copylogdb**, the CUBRID HA configuration processes. Specifies the process to be terminated when executing **deregister** by using the process ID. It is used to pause the replication of the HA log (**copylogdb**) or the reflection of replication logs (**applylogdb**) for reconfiguring the master node. To restart HA functionalities, manually execute **cubrid copylogdb** or **cubrid applylogdb**.

It is recommended that this command only be used when it is unavoidable.

The following example shows how to use it.

```
$ cubrid heartbeat deregister <process-id>
$
```

The following is an example of re-executing the **copylogdb** after executing **deregister**. The **-L** option is used to specify the location where the copy of a transaction log will be stored. The **-m** option is used to save the replica of the transaction log; it behaves like the **ha_copy_sync_mode** parameter of **cubrid_ha.conf**. Specify the option values with the values set in the **cubrid_ha.conf** file.

```
$ cubrid copylogdb -L /home/cubrid/DB/testdb01_masterdb.cub -m async testdb
$
```

The following is an example of re-executing the **applylogdb** after executing **deregister**. The **-L** option is used to specify the location where the saved transaction log will be read. The **--max_mem-size** is the maximum memory size that the **applylogdb** will use. This behaves like the **ha_apply_max_mem_size** parameter in the **cubrid_ha.conf** file. Specify the option values with the values set in the **cubrid_ha.conf** file.

```
$ cubrid applylogdb -L /home/cubrid/DB/testdb01_masterdb.cub --max-mem-size=500 testdb
$
```

act

This utility is used to includes nodes back in the CUBRID HA group, and it starts the CUBRID HA components.

It is recommended that this command only be used when it is unavoidable.

How to use this utility is as shown below.

```
$ cubrid heartbeat act
$
```

status

This utility is used to output the information of CUBRID HA group and CUBRID HA components.

How to use this utility is as shown below.

```
$ cubrid heartbeat status
@ cubrid heartbeat list

HA-Node Info (current slavedb.cub, state slave)
  Node slavedb.cub (priority 2, state slave)
  Node masterdb.cub (priority 1, state master)

HA-Process Info (master 2143, state slave)
  Applylogdb testdb01@localhost:/home/cubrid/DB/testdb01_slavedb.cub (pid 2510, state
registered)
  Copylogdb testdb01@masterdb.cub:/home/cubrid/DB/testdb01_masterdb.cub (pid 2505, state
registered)
  Server testdb01 (pid 2393, state registered)

$
```

Utilities of cubrid service

If you register heartbeat to CUBRID service, you can use the utilities of **cubrid service** to start, stop or check all the related processes at once. The processes specified by **service** parameter in **[service]** section in **cubrid.conf** file are registered to CUBRID service. If this parameter includes **heartbeat**, you can start/stop all the service processes and the HA related processes by using **cubrid service start/stop** command.

How to configure **cubrid.conf** file is shown below.

```
# cubrid.conf
...
[service]
...
service=broker,heartbeat
```

```
...
[common]
...
ha_mode=on
```

cubrid applyinfo

Description

This utility is used to monitor the replication status of CUBRID HA.

Syntax

```
cubrid applyinfo [option] <database-name>
```

- *database-name* : Specifies the name of a server to monitor. A node name is not included.

Options

Option	Default	Description
-r	none	Configures the name of a target node in which transaction logs are copied. Using this option will output the information of active logs (Active Info.) of a target node.
-a		Outputs the information of replication reflection of a node executing cubrid applyinfo. The -L option is required to use this option.
-L	none	Configures the location of transaction logs copied from the other node. Using this option will output the information of transaction logs copied (Copied Active Info.) from the other node.
-p	0	Outputs the information of a specific page in the copied logs. This is available only when the -L option is enabled.
-v		Outputs detailed information.

Example

```
$ cubrid applyinfo -L /home/cubrid/DB/tdb01 masterdb.cub -r master node name -a tdb01

*** Applied Info. ***
Committed page           : 1913 | 2904
Insert count             : 645
Update count             : 0
Delete count             : 0
Schema count             : 60
Commit count             : 15
Fail count               : 0

*** Copied Active Info. ***
DB name                  : testdb01
DB creation time         : 11:28:00.000 AM 12/17/2010 (1292552880)
EOF LSA                  : 1913 | 2976
Append LSA               : 1913 | 2976
HA server state          : active

*** Active Info. ***
DB name                  : testdb01
DB creation time         : 11:28:00.000 AM 12/17/2010 (1292552880)
EOF LSA                  : 1913 | 2976
Append LSA               : 1913 | 2976
HA server state          : active
$
```

- Applied Info.
- Committed page : The information of committed pageid and offset of a transaction reflected last through replication log reflection process. This information is internally used; which means that replication reflection will be delayed if a big difference exists between the EOF LSA value of "Copied Active Info." and this value.
- Insert Count : The number of Insert queries reflected through replication log reflection process.

- Update Count : The number of Update queries reflected through replication log reflection process.
- Delete Count : The number of Delete queries reflected through replication log reflection process.
- Schema Count : The number of DDL statements reflected through replication log reflection process.
- Commit Count : The number of commits reflected through replication log reflection process.
- Fail Count : The number of DML and DDL statements in which log reflection through replication log reflection process fails.
- Copied Active Info.
- DB name : Database name of a database server process to which replication log copy process copies replication logs
- DB creation time : The creation time of a database copied through replication log copy process
- EOF LSA : The last information of pageid and offset of a database server process replication log copied through replication log copy process. There will be a delay in replication log copy process as much as difference with the EOF LSA value of "Active Info." and with the Append LSA value of "Copied Active Info."
- Append LSA : The last information of pageid and offset of a log received from the database server process through replication log copy process. This value can be less than or equal to EOF LSA. There will be a delay in replication log copy process as much as difference between the EOF LSA value of "Copied Active Info." and this value.
- HA server state : Status of a database server process which replication log copy process receives replication logs from. For details on status, see [Server](#).
- Active Info.
- DB name : Database name of a database server process of a node that is configured in the **-r** option.
- DB creation time : Database creation time of a node that is configured in the **-r** option.
- EOF LSA : The last information of pageid and offset of a database server process replication log of a node that is configured in the **-r** option. There will be a delay in replication log copy process as much as difference between the EOF LSA value of "Copied Active Info." and this value.
- Append LSA : The last information of pageid and offset of a replication log which is written in a database server process of a node that is configured in the **-r** option.
- HA server state : Status of a database server process of a node that is configured in the **-r** option.

cubrid changemode

Description

This utility is used to check and change the server status of CUBRID HA.

Syntax

```
cubrid changemode [option] <database-name>
```

- *database-name* : Specifies the name of a server to monitor and the node name; separate them by using @.

Options

Option	Default	Description
-m	none	Changes the server status. You can enter one of the followings: standby , maintenance , or active .
-f		Configures whether or not to forcibly change the server status. This option must be configured if you want to change the server status from to-be-active to active. If it is not configured, the status will not be changed to active. Forcibly change may cause replication inconsistency; so it is not recommended.
-t	5 (in seconds)	Configures the waiting time for the normal completion of the transaction that is being processed when the node status switches from standby to maintenance . If the transaction is still in progress beyond the configured time, it will be forced to terminate and switch to maintenance status; if all transactions have completed normally within the configured time, it will switch to maintenance status

immediately.

Status Changeable Map

This table shows changeable modes depending on current status. However, replication inconsistency may occur if the status of the current server is changed from to-be-active to active. Therefore, it is recommended that only a user who is familiar with this condition uses this option.

		Future Status		
		active	standby	maintenance
Current Status	standby	X	O	O
	to-be-standby	X	X	X
	active	O	X	X
	to-be-active	O*	X	X
	maintenance	X	O	O

* When the server status is to-be-active, forcibly change may cause replication inconsistency. It is not recommended if you are not skilled enough.

Example

The following example shows how to switch the testdb01 server status in the localhost node to **maintenance**. The waiting time for all transactions in progress to complete normally is 5 seconds, which is the default value for the **-t** option. If all transactions are complete within this time limit, the status will be switched immediately. However, if there are transactions still being processed after this time limit, they will be rolled back before changing the status.

```
$ cubrid changemode -m maintenance testdb01@localhost
The server 'testdb01@localhost''s current HA running mode is maintenance.
```

The following example shows how to look up the testdb01 server status in the localhost node.

```
$ cubrid changemode testdb01@localhost
The server 'testdb01@localhost''s current HA running mode is active.
```

Monitoring CUBRID Manager HA

CUBRID Manager is a dedicated CUBRID database management tool that provides the CUBRID database management and query features in a GUI environment. CUBRID Manager provides the HA dashboard, which shows the relationship diagram for the CUBRID HA group and server status. For details, see CUBRID Manager manual.

Configuration

Overview

There are four possible structures for CUBRID HA: HA basic structure, multiple-slave node structure, load balancing structure, and multiple-standby server structure. In the table below, M stands for a master node, S for a slave node, and R for a replica node.

Structure	Node structure (M:S:R)	Characteristic
Basic Structure	1:1:0	The most basic structure of CUBRID HA consists of one master node and one slave node and provides availability which is a unique feature of CUBRID HA.
Multiple-Slave Node Structure	1:N:0	This is a structure in which availability is increased by several slave nodes. However, note that there may be a situation in which data is inconsistent in the CUBRID HA group when multiple failures occur.


Load Balancing Structure	1:1:N	Several replica nodes are added in the basic structure. Read service load can be distributed, and the HA load is reduced, comparing to a multiple-slave node structure.
Note that replica nodes do not failover.		
Multiple-Standby Server Structure	1:1:0	Basically, this structure is the same as the basic structure. However, several slave nodes are installed on a single physical server.

HA Basic Structure

The most basic structure of CUBRID HA consists of one master node and one slave node.

The default configuration is one master node and one slave node. To distribute the write load, a multi-slave node or load-distributed configuration is recommended. In addition, to access a specific node such as a slave node or replica node in read-only mode, configure a Read Only broker or the Preferred Host Read Only broker. For details about broker configuration, see [Duplexing Brokers](#).

An Example of Node Configuration



You can configure each node in the basic structure of HA as shown below:

- **nodeA** (master node)
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**:

```
ha port id=12345
ha node list=cubrid@nodeA:nodeB
ha_db_list=testdb
```

- **nodeB** (slave node): Configure this node in the same manner as nodeA.

For the **databases.txt** file of a broker node, it is necessary to configure the list of hosts configured as HA in **db-host** according to their priority. The following example shows the **databases.txt** file.

```
#db-name    vol-path          db-host          log-path         lob-base-path
testdb     /home/cubrid/DB/testdb1  nodeA:nodeB     /home/cubrid/DB/testdb/log
file:/home/cubrid/DB/testdb/lob
```

The **cubrid_broker.conf** file can be set in a variety of ways according to configuration of the broker. It can also be configured as separate equipment with the **databases.txt** file.

The example below shows that the RW broker is set in each node, and nodeA and nodeB have the same value.

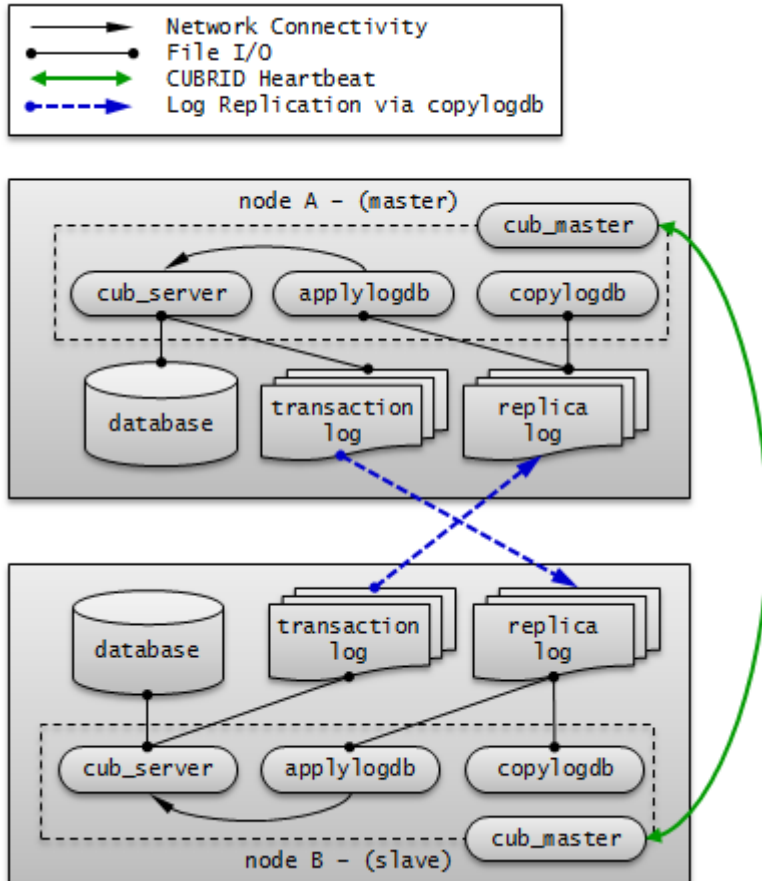
```
[%RW broker]
...
# Broker mode setting parameter
ACCESS_MODE      =RW
```

Application Connection Setup

See [JDBC Configuration](#), [CCI Configuration](#), and [PHP Configuration](#) in Environment Configuration.

Remark

The path of a transaction log in these configurations is as follows:



Multiple-Slave Node Structure

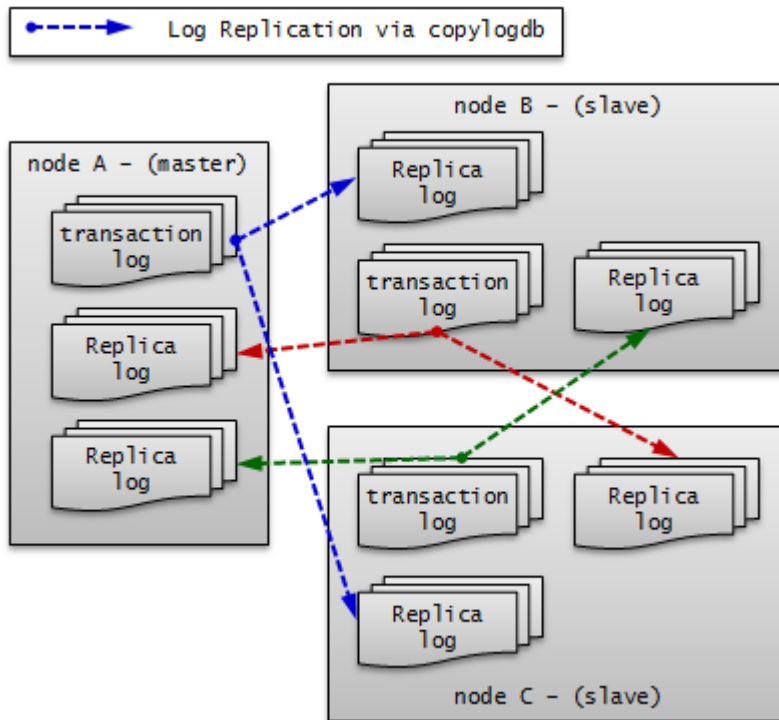
In multiple-slave node structure, there is one master node and several slave nodes to improve the service availability of CUBRID.

Because replication log copy process and replication log reflection process are running at all nodes in the CUBRID HA group, a load of copying replication log occurs. Therefore, all nodes in the CUBRID HA group have high network and disk usage.

Because there are many nodes with HA enabled, read and write services never fail as long as a single node is alive.

In the multiple-slave node structure, the node becoming a master node when failover occurs is determined by the order specified in `ha_node_list`. If the value of `ha_node_list` is `node1:node2:node3` and the master node is `node1`, `node2` will become a new master node when `node1` fails.

An Example of Node Configuration



You can configure each node in the basic structure of HA as shown below:

- **node A** (master node)
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**:

```
ha_port_id=12345
ha_node_list=cubrid@nodeA:nodeB:nodeC
ha_db_list=testdb
```

- **node B** (slave node): Configure this node in the same manner as nodeA.
- **node C** (slave node): Configure this node in the same manner as nodeA.

You must enter the list of hosts configured in HA in order of priority in the **databases.txt** file of a broker node. The following is an example of the **databases.txt** file.

```
#db-name    vol-path          db-host          log-path        lob-base-path
testdb     /home/cubrid/DB/testdb1  nodeA:nodeB:nodeC  /home/cubrid/DB/testdb/log
file:/home/cubrid/DB/testdb/lob
```

The **cubrid_broker.conf** file can be set in a variety of ways according to configuration of the broker. It can also be configured as separate equipment with the **databases.txt** file.

In this example, the RW broker is configured in nodeA, nodeB, and nodeC.

The following is an example of the **databases.txt** file in nodeA, nodeB, and nodeC.

```
[%RW broker]
...
# Broker mode setting parameter
ACCESS_MODE          =RW
```

Application Connection Setup

Connect the application to access to the broker of nodeA, nodeB, or nodeC.

```
Connection connection =
DriverManager.getConnection("jdbc:CUBRID:nodeA:33000:testdb:::charset=utf-
8&althosts=nodeB:33000,nodeC:33000", "dba", "");
```

For details, see [JDBC Configuration](#), [CCI Configuration](#), and [PHP Configuration](#) in Environment Configuration.

Caution

The data in the CUBRID HA group may lose integrity when there are multiple failures in this structure and the example is shown below.

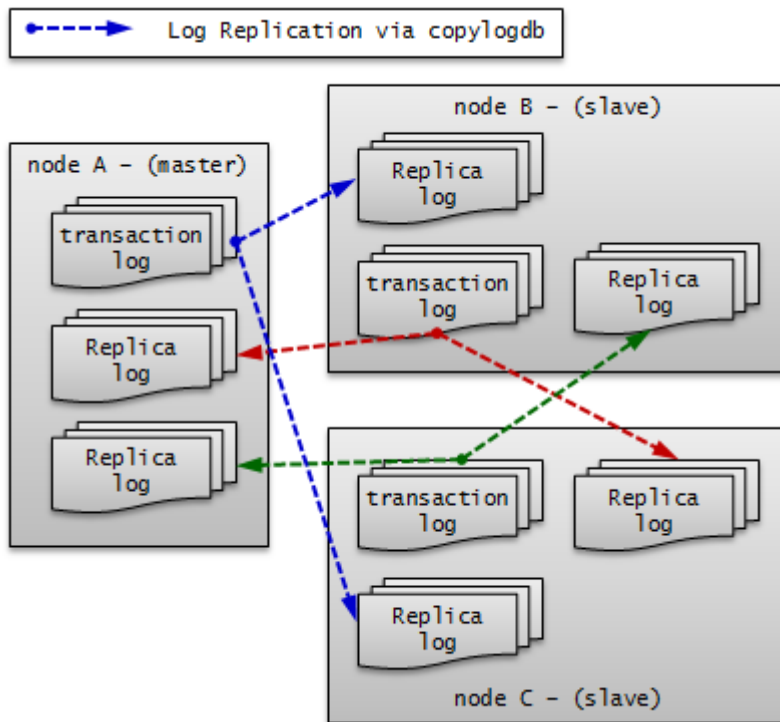
- In a situation where a failover occurs in the first slave node while replication in the second slave node is being delayed due to restart
- In a situation where a failover re-occurs before replication reflection of a new master node is not complete due to frequent failover

In addition, if the mode of replication log copy process is ASYNC, the data in the CUBRID HA group may lose integrity.

If the data in the CUBRID HA group loses integrity for any of the reasons above, you can fix it by using [Rebuilding Replication](#).

Remark

The path of a transaction log in these configurations is as follows:



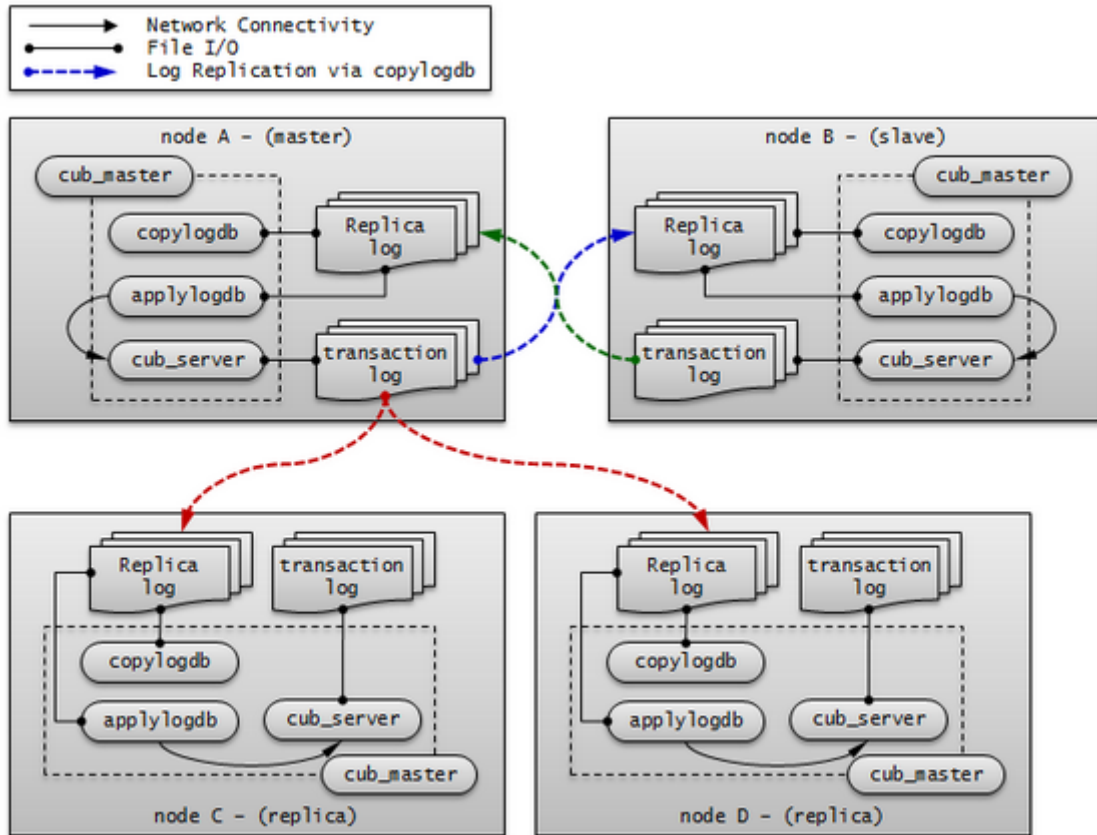
Load Balancing Structure

The load balancing structure increases the availability of the CUBRID service by placing several nodes in the HA configuration (one master node and one slave node) and distributes read-load.

Because the replica nodes receive replication logs from the nodes in the HA configuration and maintain the same data, and because the nodes in the HA configuration do not receive replication logs from the replica nodes, its network and disk usage rate is lower than that of the multiple-slave structure.

Because replica nodes are not included in the HA structure, they provide read service without failover, even when all other nodes in the HA structure fail.

An Example of Node Configuration



You can configure each node in load balancing structure as shown below:

- **node A** (master node)
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**:

```
ha_port_id=12345
ha_node_list=cubrid@nodeA:nodeB
ha_replica_list=cubrid@nodeC:nodeD
ha_db_list=testdb
```

- **node B** (slave node): Configure this node in the same manner as nodeA.
- **node C** (replica node)
- Configure the **ha_mode** of the **cubrid.conf** file to **replica**.

```
ha_mode=replica
```

- You can configure the **cubrid_ha.conf** file in the same manner as nodeA.
- **node D** (replica node): Configure this node in the same manner as nodeC.

You must enter the list of DB server hosts in the order so that each broker can be connected appropriate HA or load balancing server in the **databases.txt** file of a broker node.

The following is an example of the **databases.txt** file in nodeA and nodeB.

```
#db-name    vol-path          db-host          log-path          lob-base-path
testdb     /home/cubrid/DB/testdb1  nodeA:nodeB     /home/cubrid/DB/testdb/log
file:/home/cubrid/CUBRID/testdb/lob
```

The following is an example of the **databases.txt** file in nodeC.

```
#db-name    vol-path          db-host          log-path          lob-base-path
```

```
testdb /home/cubrid/DB/testdb nodeC /home/cubrid/DB/testdb/log file:/home/cubrid/CUBRID/testdb/lob
```

The following is an example the **databases.txt** nodeD.

```
#db-name vol-path db-host log-path lob-base-path
testdb /home/cubrid/DB/testdb nodeD /home/cubrid/DB/testdb/log
file:/home/cubrid/CUBRID/testdb/lob
```

The **cubrid_broker.conf** can be set in a variety of ways according to configuration of the broker. It can also be configured as separate equipment with the **databases.txt** file.

In this example, the RW broker is configured in nodeA and nodeB and the PHRO broker is configured in nodeC and nodeD.

The following is an example of **cubrid_broker.conf** in nodeA and nodeB.

```
[%RW_broker]
...
# Broker mode setting parameter
ACCESS_MODE =RW
```

The following is an example **cubrid_broker.conf** in nodeC.

```
[%PHRO_broker]
...
# Broker mode setting parameter
ACCESS MODE =PHRO
PREFERRED_HOSTS =nodeC:nodeD
```

The following is an example **cubrid_broker.conf** in nodeD.

```
[%PHRO broker]
...
# Broker mode setting parameter
ACCESS MODE =PHRO
PREFERRED_HOSTS =nodeD:nodeC
```

Application Connection Setup

Connect the application to access in read/write mode to the broker of nodeA or nodeB. The following is an example of a JDBC application.

```
Connection connection =
DriverManager.getConnection("jdbc:CUBRID:nodeA:33000:testdb::?charset=utf-8&althosts=nodeB:33000", "dba", "");
```

Connect the application to access in read-only mode to the broker of nodeC or nodeD. The following is an example of a JDBC application.

```
Connection connection =
DriverManager.getConnection("jdbc:CUBRID:nodeC:33000:testdb::?charset=utf-8&althosts=nodeD:33000", "dba", "");
```

For details, see [JDBC Configuration](#), [CCI Configuration](#), and [PHP Configuration](#) in Environment Configuration.

Caution

The data in the CUBRID HA group may lose integrity when there are multiple failures in this structure.

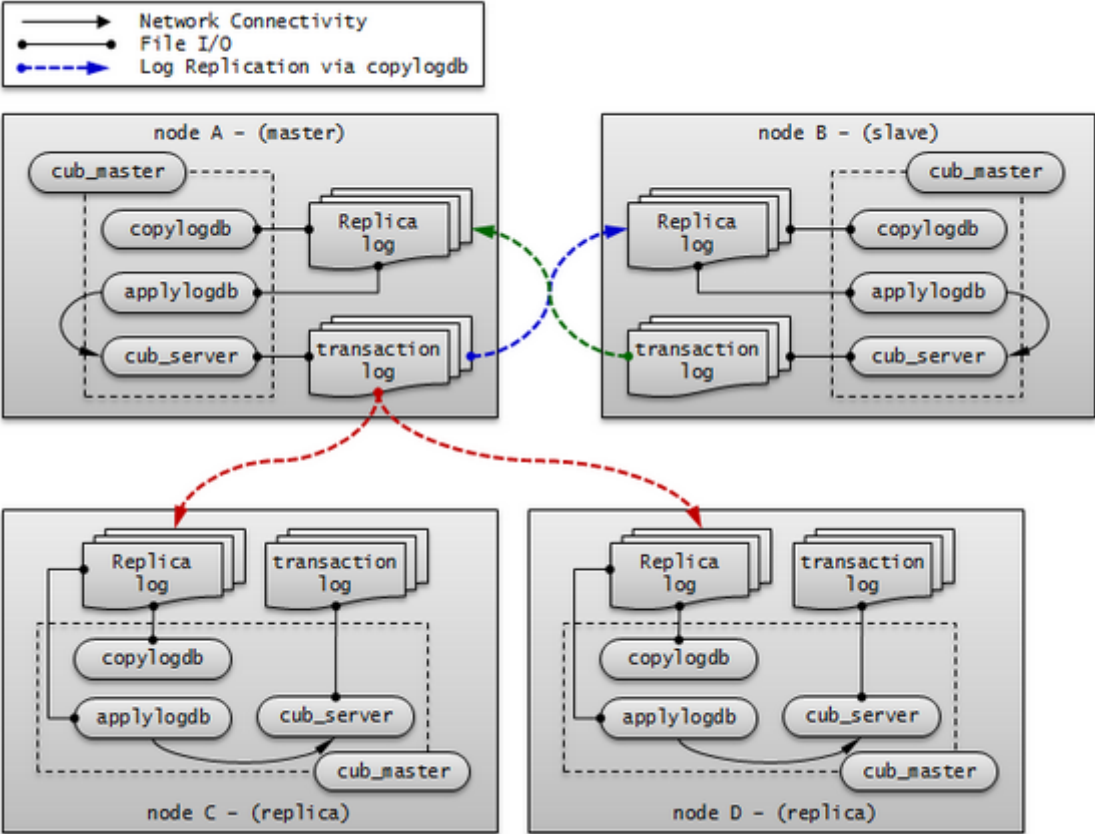
- In a situation where a failover occurs in the first slave node while replication in the second slave node is being delayed due to restart
- In a situation where a failover re-occurs before replication reflection of a new master node is not complete due to frequent failover

In addition, if the mode of replication log copy process is ASYNC, the data in the CUBRID HA group may lose integrity.

If the data in the CUBRID HA group loses integrity for any of the reasons above, you can fix it by using [Rebuilding Replication](#).

Remark

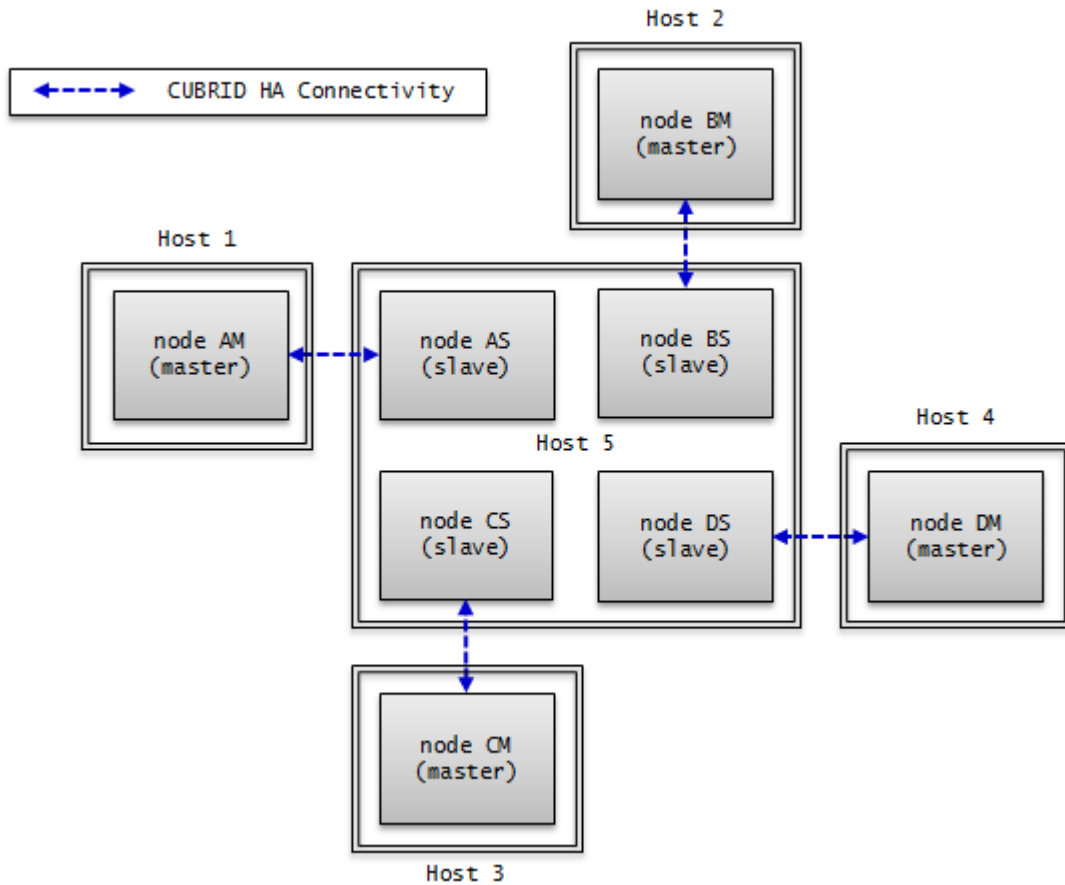
The path of a transaction log in these configurations is as follows:



Multiple-Standby Server Structure

Although its node structure has a single master node and a single slave node, many slave nodes from different services are physically configured in a single server.

This structure is for very small services in which the read load of slave nodes are light. It is strictly for the availability of the CUBRID service. For this reason, when a master node that failed after a failover has been restored, the load must be moved back to the original master node to minimize the load of the server with multiple-slave nodes.



An Example of Node Configuration

You can configure each node in the basic structure of HA as shown below:

- **node AM, node AS:** Configure node AM and node AS in the same manner.
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**.

```
ha port id=10000
ha_node_list=cubridA@Host1:Host5
ha_db_list=testdbA1,testdbA2
```

- **node BM, node BS:** Configure node BM and node BS in the same manner.
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**.

```
ha port id=10001
ha_node_list=cubridB@Host2:Host5
ha_db_list=testdbB1,testdbB2
```

- **node CM, node CS:** Configure node CM and node CS in the same manner.
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**.

```
ha port id=10002
ha_node_list=cubridC@Host3:Host5
ha_db_list=testdbC1,testdbC2
```

- **node DM, node DS:** Configure node DM and node DS in the same manner.
- Configure the **ha_mode** of the **cubrid.conf** file to **on**.

```
ha_mode=on
```

- The following example shows how to configure **cubrid_ha.conf**.

```
ha port id=10003
ha node list=cubridD@Host4:Host5
ha_db_list=testdbD1,testdbD2
```

Constraints

Supported Platform

Currently, CUBRID HA is supported only by Linux. All nodes within a CUBRID HA group must be configured with the same platform.

Table Primary Key

CUBRID HA synchronizes data among nodes with the following method (as known as transaction log shipping): First, it replicates the primary key-based replication logs generated from the server of a master node to a slave node. Second, it reflects the replication logs.

If data of the specific table within a CUBRID HA group is not synchronized, you should check whether the appropriate primary key has specified for the table.

Table Trigger and Java Stored Procedure

Using triggers and java stored procedures in CUBRID HA can cause duplicate execution because triggers and java stored procedures executed in a master node will be executed in a slave node again. This may cause data inconsistency among nodes within a CUBRID HA group.

It is not recommended to use triggers and java stored procedures in CUBRID HA.

Method and CUBRID Manager

CUBRID HA synchronizes data among nodes within a CUBRID HA group based on replication logs. Therefore, using method that does not generate replication logs or configuring **NOT NULL** through CUBRID Manager may cause data inconsistency among nodes within a CUBRID HA group.

Stand-Alone Mode

The replication logs are not generated as for tasks performed in stand-alone mode. For this reason, data inconsistency among nodes in a CUBRID HA group may occur when performing tasks in stand-alone mode.

Serial Cache

To enhance performance, a serial cache does not access Heap and does not generate replication logs when retrieving or updating serial information. Therefore, if you use a serial cache, the current values of serial caches will be inconsistent among the nodes in a CUBRID HA group.

cubrid backupdb -r

This command is used to back up a specified database. If the **-r** option is used, logs that are not required for recovery will be deleted. This deletion may result in data inconsistency among nodes in a HA group. Therefore, you must not use the **-r** option.

INCR/DECR Function

If you use **INCR/DECR** (click counter functions) in a slave node of HA configuration, an error is returned.

LOB (BLOB/CLOB) Type

In a CUBRID HA environment, the meta data (Locator) of a **LOB** column is replicated and data of a **LOB** type is not replicated. Therefore, if storage of a **LOB** type is located on the local machine, no tasks on the columns in a slave node or a master node after failover are allowed.

Error Messages

Replication Log Copy Process

The error messages from the replication log copy process are stored in `$CUBRID/log/db-name@remote-node-name_copylogdb.err`. The severity of error messages found in the replication log copy process are as follows: fatal, error, and notification. The default severity is error. Therefore, to store notification error messages, it is necessary to change the value of `error_log_level` in `cubrid.conf`. For details, see [Error Message-Related Parameters](#).

Initialization Error Message

The error messages that may be stored in the initialization stage of replication log copy process are as follows:

Error Code	Error Message	Severity	Description	Solution
10	Unable to mount disk volume ?.	error	Opening a replication log file has failed.	Check if there is a replication log. For the location of replicated logs, see Default Environment Configuration .
78	Internal error: an I/O error occurred while reading logical log page ? (physical page ?) of ?.	fatal	Reading a replication log has failed.	Check the replication log by using the cubrid applyinfo utility.
81	Internal error: logical log page ? may be corrupted.	fatal	A replication log page error, in which the replication log copy process has been copied from the connected database server process.	Check the error log of the database server process to which the replication log copy process is connected. This error log can be found in <code>\$CUBRID/log/server</code> .
1039	log writer: log writer started. mode: ?	error	The replication log copy process has been successfully initialized and started.	No action is required, because this error message has been stored to display the starting information of the replication log copy process. Error messages that have been displayed after the start of the replication log copy process and before this error message is displayed can be ignored, as they can be displayed in normal conditions.

Replication Log Request and Reception Error Messages

The replication log copy process requests a replication log from the connected database server, and receives the corresponding replication log. The error messages that may occur during this process are as follows:

Error Code	Error Message	Severity	Description	Solution
89	Log ? does not belong to the given database.	error	The previously replicated log and the log to be replicated do not match.	Check the information of the database server/host to which the replication log copy process is connected. If you need to change the

				database server/host information, reinitialize it by deleting the existing replication log, and then restarting.
186	Error receiving data from server.	error	Incorrect information has been received from the database server to which the replication log copy process is connected.	This is recovered internally.
199	Server no longer responding.	error	The connection to the database server has been terminated.	This is recovered internally.

Replication Log Writing Error Messages

The replication log copy process copies the replication log (ha_copy_base) that was received from the connected database server process to the location specified in **cubrid_ha.conf** . The error messages that may occur during this process are as follows:

Error Code	Error Message	Severity	Description	Solution
10	Unable to mount disk volume ?.	error	Opening a replication log file has failed.	Check if the replication log exists.
79	Internal error: an I/O error occurred while writing logical log page ? (physical page ?) of ?.	fatal	Writing a replication log has failed.	This is recovered internally.
80	Insufficient space in operating system device when writing logical log page ? (physical page ?) of ?. Could not write more than ? bytes.	fatal	Writing a replication log has failed due to insufficient file system space.	Check if there is sufficient space left in the disk partition.

Replication Log Archive Error Messages

The replication log copy process periodically archives the replication log that has been received from the connected database server process. The error messages that may occur during this process are as follows:

Error Code	Error Message	Severity	Description	Solution
78	Internal error: an I/O error occurred while reading logical log page ? (physical page ?) of ?.	fatal	Reading a replication log has failed during archiving.	Check the replication log by using the cubrid applyinfo utility.
79	Internal error: an I/O error occurred while writing logical log page ? (physical page ?) of ?.	fatal	Writing the archive log has failed.	This is recovered internally.
81	Internal error: logical log page ? may be corrupted.	fatal	A replication log error has been found during archiving.	Check the replication log by using the cubrid applyinfo utility.
98	Unable to create archive log ? to archive pages from ? to ?.	fatal	Creating the archive log file has failed.	Check if there is sufficient space left in the disk partition.
974	Archive log ? is created to archive pages from ? to ?.	notification	Archive log file information	No action is required, because this error message is recorded for the log information of the newly

created archive.

Stop and Restart Error Message

The error messages that may occur at the beginning and the end of the replication log copy process are as follows:

Error Code	Error Message	Severity	Description	Solution
1037	log writer: log writer shut itself down by signal.	error	The copylogdb process has been terminated by a specific signal.	This is recovered internally.

Replication Log Reflection Process

The error messages from the replication log reflection process are stored in `$CUBRID/log/db-name@local-node-name_applylogdb_db-name_remote-node-name.err`. The severity of error messages found in the replication log reflection process can be as follows: fatal, error, and notification. The default severity is error. Therefore, to store notification error messages, it is necessary to change the value of `error_log_level` in `cubrid.conf`. For details, see [Error Message-Related Parameters](#).

Initialization Error Message

The error messages that may be stored in the initialization stage of the replication log reflection process are as follows:

Error Code	Error Message	Severity	Description	Solution
10	Unable to mount disk volume ?	error	An applylogdb that is trying to reflect the same replica log is already running.	Check if there is a applylogdb process that is trying to reflect the same replication log.
1038	log applier: log applier started. required LSA: ???. last committed LSA: ???.	error	Starts normally once the initialization of applylogdb is successful.	No action is required, because this error message has been stored to display the starting information of the replication log reflection process.

Log Analysis Error Message

The replication log reflection process reads, analyzes, and reflects the replication logs that have been copied by the replication log copy process. The error messages that may occur during a replication log analysis are as follows:

Error Code	Error Message	Severity	Description	Solution
13	An I/O error occurred while reading page ? of volume ?.	error	Reading the log page to be reflected has failed.	Check the replication log by using the cubrid applyinfo utility.
17	Internal error: fetching deallocated pageid ? of volume ?.	fatal	Trying to read a log page that does not exist in the replication log	Check the replication log by using the cubrid applyinfo utility.
81	Internal error: logical log page ? may be corrupted.	fatal	There is an inconsistency between an old log under replication reflection and the current log, or there is a replication log record error.	Check the replication log by using the cubrid applyinfo utility.
82	Unable to mount log disk volume/file ?.	error	No replication log file	Check if there is a replication log. Check the

				replication log by using the cubrid applyinfo utility.
97	Internal error: unable to find log page ? in log archives.	error	The log page does not exist in the replication log.	Check the replication log by using the cubrid applyinfo utility.
897	Decompression failed.	error	Decompressing the log record has failed.	Check the replication log by using the cubrid applyinfo utility.
1028	log applier: unexpected EOF record in archive log. LSA: ??.	error	An incorrect log record exists in the archive log.	Check the replication log by using the cubrid applyinfo utility.
1029	log applier: invalid replication log page/offset. page HDR: ??, final: ??, append LSA: ??, EOF LSA: ??, ha file status: ?, is end-of-log: ?.	error	An incorrect log record has been included.	Check the replication log by using the cubrid applyinfo utility.
1030	log applier: invalid replication record. LSA: ??, forw LSA: ??, backw LSA: ??, Trid: ?, prev tran LSA: ??, type: ?.	error	A log record header error	Check the replication log by using the cubrid applyinfo utility.

Replication Log Reflection Error Message

The replication log reflection process reads, analyzes, and reflects the replication logs that have been copied by the replication log copy process. The error messages that may occur when reflecting a replication log analysis are as follows:

Error Code	Error Message	Severity	Description	Solution
72	Your transaction (index ?, ?@??) has been unilaterally aborted by the system.	error	Replication reflection has been failed due to deadlock, etc.	This is recovered internally.
111	Your transaction has been aborted by the system due to server failure or mode change.	error	Replication reflection is failed because the database server process for replication reflection has been terminated, or its mode has been changed.	This is recovered internally.
191	Cannot connect to server ? on ?.	error	The connection to the database server process for replication reflection has been terminated.	This is recovered internally.
195	Server communications error: ?.	error	The connection to the database server process for replication reflection has been terminated.	This is recovered internally.
224	A database has not been restarted.	error	The connection to the database server process for replication reflection has	This is recovered internally.

			been terminated.	
1027	log applier: failed to change apply state from ? to ?.	error	The status change of replication reflection has been failed.	This is recovered internally.
1031	log applier: failed to apply schema replication log. class: ?, schema: ?, internal error: ?.	error	SCHEMA replication reflection has been failed.	Check the consistency of the replica. If it is inconsistent, reconfigure the HA replication.
1032	log applier: failed to apply insert replication log. class: ?, key: ?, internal error: ?.	error	INSERT replication reflection has been failed.	Check the consistency of the replica. If it is inconsistent, reconfigure the HA replication.
1033	log applier: failed to apply update replication log. class: ?, key: ?, internal error: ?.	error	UPDATE replication reflection has been failed.	Check the consistency of the replica. If it is inconsistent, reconfigure the HA replication.
1034	log applier: failed to apply delete replication log. class: ?, key: ?, internal error: ?.	error	DELETE replication reflection has been failed.	Check the consistency of the replica. If it is inconsistent, reconfigure the HA replication.
1040	HA generic: ?.	notification	Change the last record of the archive log or replication reflection status.	No action is required, because this error message is stored to provide general information.

Stop and Restart Error Message

The error messages that may occur at the beginning and the end of the replication log reflection process are as follows:

Error Code	Error Message	Severity	Description	Solution
1035	log applier: mem size(? MB) of log applier is greater than max mem size (? MB) or has been grow more than 2 times (? MB). required LSA: ? ? last committed LSA: ? ?.	error	The replication log reflection process has been restarted due to reaching the maximum memory size limit.	This is recovered internally.
1036	log applier: log applier shut itself down by signal.	error	The replication log reflection process has been terminated by a specified signal.	This is recovered internally.

Operation Scenario

Rebuilding Replication

Replication rebuilding is required in CUBRID HA when data in the CUBRID HA group is inconsistent because of multiple failures in multiple-slave node structure, or because of a generic error. Rebuilding replications in CUBRID HA is perform done through a script. With the **cubrid applyinfo** utility, you can check the replication progress; however replication inconsistency is not detected. If you want to determine whether replication is inconsistent, you must examine data of the master and slave nodes yourself.

For rebuilding replications, the following environment must be the same in the slave, master, and replica nodes.

- CUBRID version
- Environmental variable (**\$CUBRID**, **\$CUBRID_DATABASES**, **\$LD_LIBRARY_PATH**, **\$PATH**)
- The paths of database volume, log, and replication
- Username and password of the Linux server
- HA-related parameters except for **ha_mode** and **ha_copy_sync_mode**, **ha_ping_hosts**

ha_make_slavedb.sh Script

To rebuild replications, use the **ha_make_slavedb.sh** script. This script is located in **\$CUBRID/share/script/ha**. Before rebuilding replications, the following items must be configured for the environment of the user. This script is supported since the version 2008 R2.2 Patch 9 and its configuration is different from 2008 R4.1 Patch 2 or earlier. This document describes it in CUBRID 2008 R4.1 Patch 2 or higher.

- **target_host**: The host name of the source node for rebuilding replication. It should be registered in **/etc/hosts**. A slave node can be replicated as the master node or the replica node. A replica node can be replicated and rebuilt as another replica node.
- **repl_log_home**: Specifies the home directory of the replication log of the master node. It is usually the same as **\$CUBRID_DATABASES**.

The following are optional items:

- **db_name**: Specifies the name of the database to be replicated. If not specified, the first name that appears in **ha_db_list** in **\$CUBRID/conf/cubrid_ha.conf** is used.
- **backup_dest_path**: Specifies the path in which the backup volume is created when executing **backupdb** in source node for rebuilding replication.
- **backup_option**: Specifies necessary options when executing **backupdb** in source node in which replication will be rebuilt.
- **restore_option**: Specifies necessary options when executing **restoredb** in slave node in which replication will be rebuilt.
- **scp_option**: Specifies the **scp** option which enables backup of source node in which replication is rebuilt to copy into the slave node. The default option is **-l 131072**, which does not impose a overload on network (limits the transfer rate to 16 M).

Once the script has been configured, execute the **ha_make_slavedb.sh** script in slave node in which replication will be rebuilt. When the script is executed, rebuilding replication happens in a number of phases. To move to the next stage, the user must enter an appropriate value. The following are the descriptions of available values.

- **yes**: Keeps going.
- **no**: Does not move forward with any stages from now on.
- **skip**: Skips to the next stage. This input value is used to ignore a stage that has not necessarily been executed when retrying the script after it has failed.


Constraints

- Remote ssh connection must be available when using the script because it executes connection commands in the remote node by using expect and ssh.
- **Online backup of rebuilding replication node**: Existing backup of the replica or slave nodes cannot be used for rebuilding replication. Therefore, you must use the online backup of the master node that is automatically created by the script.
- **Error while executing the rebuilding replication script**: The rebuilding replication script is not automatically rolled back to its previous stage even when an error occurs during the execution. This is because the slave node cannot provide normal service before rebuilding replication script is executed. To return to the phase before rebuilding replication script is executed, you must back up the existing replication logs and **db_ha_apply_info** information which is internal catalog of the master and slave nodes before building replication is executed.

Caution

To replicate, you must copy the physical image of the database volume in the target node to the database of the node to be replicated. However, **cubrid unloaddb** backs up only logical images so replication using **cubrid unloaddb** and **cubrid loaddb** is unavailable. Because **cubrid backupdb** backs up physical images, replication is possible by using this utility. The **ha_make_slavedb.sh** script performs replication by using **cubrid backupdb**.

Example



- The host name in master node: master
- The host name in slave node: slave
- The host name in replica node1: replica1
- The host name in replica node2: replica2

In the HA-configured server, when a slave node fails and must be replicated using a replica node, the **target_host** of **ha_make_slavedb.sh** must be changed to replica1 or replica2. If the value of **REPL_LOG_HOME** is not **SCUBRID_DATABASES**, set the value of **repl_log_home** as well.

```
[slave]$ cd $CUBRID/share/script/ha
[slave]$ vi ha make slavedb.sh
target host="replica2"

# if REPL LOG HOME != $CUBRID DATABASES then
repl_log_home=$USER_SPECIFIC_REPL_LOG_HOME
```

The following example shows how to execute a script after storing changes.

```
[slave]$ ./ha_make_slavedb.sh
```

Operation Scenario during Read/Write Service

The operation scenario written in this page is not affected by read/write services. Therefore, its impact on the services caused by CUBRID operation is very limited. There can be two types of operation scenarios in which failover occurs or it does not occur.

When Failover Does Not Occur

You can perform the following operations without stopping and restarting nodes in CUBRID HA group.

General Operation	Scenario	Consideration
Online Backup	Operation task is performed at each master node and slave node each during operation.	Note that there may be a delay in the transaction of master node due to the operation task.
Schema change (excluding basic key change), index change, authorization change	When an operation task occurs at a master node, it is automatically replicated to a slave node.	Because replication log is copied and reflected to a slave node after an operation task is completed in a master node, operation task time is doubled. Changing schema must be processed without any failover. Index change and authority change other than the schema change can be performed by stopping each node and executing standalone mode (ex: the -S

		option of the csql utility) when the operation time is important.
Add volume	Operation task is performed at each DB regardless of HA structure.	Note that there may be a delay in the transaction of master node due to the operation task. If operation task time is an issue, operation task can be performed by stopping each node and executing standalone mode (ex: the -S of the cubrid addvoldb utility).
Failure node server replacement	It can be replaced without restarting the CUBRID HA group when a failure occurs.	The failure node must be registered in the ha_node_list of CUBRID HA group, and the node name must not be changed during replacement.
Failure broker server replacement	It can be replaced without restarting the broker when a failure occurs.	The connection to a broker replaced at a client can be made by rctime which is configured in URL string.
DB server expansion	You can execute cubrid heartbeat reload after configuration change (ha_node_list , ha_replica_list) without restarting the previously configured CUBRID HA group.	Note that all the management processes of a node are stopped when cubrid heartbeat reload fails.
Broker server expansion	Run additional brokers without restarting existing brokers.	Modify the URL string to connect to a broker where a client is added.

When Failover Occurs

You must stop nodes in CUBRID HA group and complete operation before performing the following operations.

General Operation	Scenario	Consideration
DB server configuration change	A node whose configuration is changed is restarted when the configuration in cubrid.conf is changed.	
Change broker configuration, add broker, and delete broker	A broker whose configuration is changed is restarted when the configuration in cubrid_broker.conf is changed.	
DBMS version patch	Restart nodes and brokers in HA group after version patch.	Version patch means there is no change in the internal protocol, volume, and log of CUBRID.

Operation Scenario during Read Service

The operation scenario written in this page is only applied to read service. It is required to allow read service only or dynamically change mode configuration of broker to Read Only. There can be two types of operation scenarios in which failover occurs or it does not occur.

When Failover Does Not Occur

You can perform the following operations without stopping and restarting nodes in CUBRID HA group.

General Operation	Scenario	Consideration
Schema change (primary key change)	When an operation task is performed at the master node, it is automatically reflected to the slave node.	In order to change the primary key, the existing key must be deleted and a new one added. For this reason, replication reflection may not occur due to the HA internal structure which reflects primary key-based replication logs. Therefore, operation

		tasks must be performed during the read service.
Schema change (excluding basic key change), index change, authorization change	When an operation task is performed at the master node, it is automatically reflected to the slave node.	Because replication log is copied and reflected to a slave node after an operation task is completed in a master node, operation task time is doubled. Changing schema must be processed without any failover. Index change and authority change other than the schema change can be performed by stopping each node and executing standalone mode (ex: the span class="nkeyword">-S option of csql) when the operation time is important.

When Failover Occurs

You must stop nodes in CUBRID HA group and complete operation before performing the following operations.

General Operation	Scenario	Consideration
DBMS version upgrade	Restart each node and broker in the CUBRID HA group after they are upgraded.	A version upgrade means that there have been changes in the internal protocol, volume, or log of CUBRID. Because there are two different versions of the protocols, volumes, and logs of a broker and server during an upgrade, an operation task must be performed to make sure that each client and broker (before/after upgrade) are connected to the corresponding counterpart in the same version.
Massive data processing (INSERT / UPDATE / DELETE)	Stop the node that must be changed, perform an operation task, and then execute the node.	This processes massive data that cannot be segmented.

Operation Scenario after Service Stop

You must stop all nodes in CUBRID HA group before performing the following operation.

General Operation	Scenario	Consideration
Changing the host name and IP of a DB server	Stop all nodes in the CUBRID HA group, and restart them after the operation task.	When a host name has been changed, change the databases.txt file of each broker and reset the broker connection with cubrid broker reset .

Performance Tuning

This chapter provides information about configuring system parameters that can affect the system performance. System parameters determine overall performance and operation of the system. This chapter explains how to use configuration files for database server and broker as well as a description of each parameter. For CUBRID Manager server configuration, see CUBRID Manager Manual.

This chapter covers the following topics:

- Configuring the Database server
- Configuring the Broker

Database Server Configuration

Scope of Database Server Configuration

CUBRID consists of the Database Server, the Broker and the CUBRID Manager. Each component has its configuration file. The system parameter configuration file for the Database Server is **cubrid.conf** located in the **\$CUBRID/conf** directory. System parameters configured in **cubrid.conf** affect overall performance and operation of the database system. Therefore, it is very important to understand the Database Server configuration.

The CUBRID Database Server has a client/server architecture. To be more specific, it is divided into a Database Server process linked to the server library and a Broker process linked to the client library. The server process manages the database storage structure and provides concurrency and transaction functionalities. The client process prepares for query execution and manages object/schema.

System parameters for the database server, which can be set in the **cubrid.conf** file, are classified into a client parameter, a server parameter and a client/server parameter according to the range to which they are applied. A client parameter is only applied to client processes such as the broker. A server parameter affects the behaviors of the server processes. A client/server parameter must be applied to both the server and the client.

Location of cubrid.conf File and How It Works

- A Database Server process refers only to the **\$CUBRID/conf/cubrid.conf** file. Database-specific configurations are distinguished by sections in the **cubrid.conf** file.
- A client process (i) refers to the **\$CUBRID/conf/cubrid.conf** file and then (ii) additionally refers to the **cubrid.conf** file in the current directory (**\$PWD**). The configuration of the file in the current directory (**\$PWD/cubrid.conf**) overwrites that of the **\$CUBRID/conf/cubrid.conf** file. That is, if the same parameter configuration exists in **\$PWD/cubrid.conf** and in **\$CUBRID/conf/cubrid.conf**, the configuration in **\$PWD/cubrid.conf** has the priority.

cubrid.conf Configuration File and Default Parameters

CUBRID consists of the Database Server, the Broker and the CUBRID Manager. The name of the configuration file for each component is as follows. These files are all located in the **\$CUBRID/conf** directory.

- Database Server configuration file : **cubrid.conf**
- Broker configuration file : **cubrid_broker.conf**
- CUBRID Manager server configuration file : **cm.conf**

cubrid.conf is a configuration file that sets system parameters for the CUBRID Database Server and determines overall performance and operation of the database system. In the **cubrid.conf** file, some important parameters needed for system installation are provided, having their default values.

Database Server System Parameters

The following are Database Server system parameters that can be used in the **cubrid.conf** configuration file. For the scope of **client** and **server parameters**, see [Scope of Database Server Configuration](#).

You can change the parameters that are capable of dynamically changing the setting value through the **SET SYSTEM PARAMETERS** statement or a session command of the CSQL Interpreter, **;set** dynamically. If you are a DBA, you can change parameters regardless of the applied classification. However, if you are not a DBA, you can only change client parameters.

Category	Parameter Name	Applied	Type	Default Value	Dynamicity
Connection	cubrid_port_id	client parameter	int	1523	
	db_hosts	client parameter	string	NULL	available

	max_clients	server parameter	int	100	
Memory	data_buffer_size	server parameter	int	512M	
	index_scan_oid_buffer_size	server parameter	int	64K	
	sort_buffer_size	server parameter	int	2M	
	temp_file_memory_size_in_pages	server parameter	int	4	
	thread_stack_size	server parameter	int	1048576	
Disk	db_volume_size	server parameter	int	512M	
	dont_reuse_heap_file	server parameter	bool	no	
	temp_file_max_size_in_pages	server parameter	int	-1	
	temp_volume_path	server parameter	string	NULL	
	unfill_factor	server parameter	float	0.1	
	volume_extension_path	server parameter	string	NULL	
	log_volume_size	server parameter	int	512M	
Error message	call_stack_dump_activation_list	client/server parameter	string	NULL	available
	call_stack_dump_deactivation_list	client/server parameter	string	NULL	available
	call_stack_dump_on_error	client/server parameter	bool	no	available
	error_log	client/server parameter	string	cub_client.err, cub_server.err	
	error_log_level	client/server parameter	string	SYNTAX	available
	error_log_warning	client/server parameter	bool	no	available
	error_log_size	client/server parameter	int	8000000	available
Concurrency/Lock	deadlock_detection_interval_in_secs	server parameter	float	1.0	available
	isolation_level	client parameter	int	3	available
	lock_escalation	server parameter	int	100000	
	lock_timeout_in_secs	client parameter	int	-1	available
	lock_timeout_message_type	server parameter	int	0	available
Logging	adaptive_flush_control	server parameter	bool	yes	available
	background_archiving	server parameter	bool	yes	available
	checkpoint_every_npages	server parameter	int	10000	
	checkpoint_interval_in_mins	server parameter	int	720	available
	force_remove_log_archives	server parameter	bool	yes	available
	log_buffer_size	server parameter	int	2M	
	log_max_archives	server parameter	int	INT_MAX	available
	max_flush_pages_per_second	server parameter	int	10000	available
	page_flush_interval_in_msecs	server parameter	int	0	available
sync_on_nflush	server parameter	int	200	available	
Transaction handling	async_commit	server parameter	bool	no	
	group_commit_interval_in_msecs	server parameter	int	0	available
Syntax/Type	add_column_update_hard_default	client parameter	bool	no	available
	alter_table_change_type_strict	client/server parameter	bool	no	available

	ansi_quotes	client parameter	bool	yes	
	block_ddl_statement	client parameter	bool	no	available
	block_nowhere_statement	client parameter	bool	no	available
	compat_numeric_division_scale	client/server parameter	bool	no	available
	default_week_format	client/server parameter	int	0	available
	group_concat_max_len	server parameter	int	1024	available
	intl_mbs_support	client parameter	bool	no	
	no_backslash_escapes	client parameter	bool	yes	
	only_full_group_by	client parameter	bool	no	available
	oracle_style_empty_string	client parameter	bool	no	
	pipes_as_concat	client parameter	bool	yes	
	plus_as_concat	client parameter	bool	yes	
	require_like_escape_character	client parameter	bool	no	
	return_null_on_function_errors	client/server parameter	bool	no	available
Query cache	max_plan_cache_entries	client/server parameter	int	1000	
Utility	backup_volume_max_size_bytes	server parameter	int	-1	
	communication_histogram	client parameter	bool	no	available
	compactdb_page_reclaim_only	server parameter	int	0	
	csql_history_num	client parameter	int	50	available
HA	ha_mode	server parameter	string	off	
Others	access_ip_control	server parameter	bool	no	
	access_ip_control_file	server parameter	string		
	auto_restart_server	server parameter	bool	yes	available
	index_scan_in_oid_order	client parameter	bool	no	available
	index_unfill_factor	server parameter	float	0.05	
	insert_execution_mode	client parameter	int	1	
	java_stored_procedure	server parameter	bool	no	
	multi_range_optimization_limit	server parameter	int	100	available
	pthread_scope_process	server parameter	bool	yes	
	server	server parameter	string		
	service	server parameter	string		
	session_state_timeout	server parameter	int	21600	
	single_byte_compare	server parameter	bool	no	
	use_orderby_sort_limit	server parameter	bool	yes	available

Section by Parameter

Parameters specified in **cubrid.conf** have the following three sections:

- Used when the CUBRID service starts : [service] section
- Applied commonly to all databases : [common] section

- Applied individually to each database : [`@<database>`] section

Where `<database>` is the name of the database to which each parameter applies. If a parameter configured in `[common]` is the same as the one configured in [`@<database>`], the one configured in [`@<database>`] is applied.

Default Parameters

cubrid.conf, a default database configuration file created during the CUBRID installation, includes some default Database Server parameters that must be changed. You can change the value of a parameter that is not included as a default parameter by manually adding or editing one.

The following is the content of the **cubrid.conf** file.

```
# Copyright (C) 2008 Search Solution Corporation. All rights reserved by Search Solution.
#
# $Id$
#
# cubrid.conf#

# For complete information on parameters, see the CUBRID
# Database Administration Guide chapter on System Parameters

# Service section - a section for 'cubrid service' command
[service]

# The list of processes to be started automatically by 'cubrid service start' command
# Any combinations are available with server, broker and manager.
service=server,broker,manager

# The list of database servers in all by 'cubrid service start' command.
# This property is effective only when the above 'service' property contains 'server'
keyword.
#server=server, broker, manager

# Common section - properties for all databases
# This section will be applied before other database specific sections.
[common]

# Read the manual for detailed description of system parameters
# Manual > Performance Tuning > Database Server Configuration > Default Parameters

# Size of data buffer are using K, M, G, T unit
data buffer size=512M

# Size of log buffer are using K, M, G, T unit
log buffer size=4M

# Size of sort buffer are using K, M, G, T unit
# The sort buffer should be allocated per thread.
# So, the max size of the sort buffer is sort_buffer_size * max_clients.
sort buffer size=2M

# The maximum number of concurrent client connections the server will accept.
# This value also means the total # of concurrent transactions.
max_clients=100

# TCP port id for the CUBRID programs (used by all clients).
cubrid_port_id=1523
```

Connection-Related Parameters

The following are parameters related to the Database Server. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min	Max
cubrid_port_id	int	1523	1	
db_hosts	string	NULL		
max_clients	int	100	10	1024

cubrid_port_id

cubrid_port_id is a parameter that configures the port to be used by the master process. The default value is **1,523**. If the port 1,523 is already being used on the server where CUBRID is installed or it is blocked by a firewall, an error message, which means the master server is not connected because the master process cannot be running properly, is displayed. If such port conflict occurs, the administrator must change the value of **cubrid_port_id** considering the server environment.

db_hosts

db_hosts is a parameter that specifies a list of Database Server hosts to which clients can connect, and the connection order. The server host list consists of more than one server host names, and host names are separated by spaces or colons (:). Duplicate or non-existent names are ignored.

The following example shows the values of the **db_hosts** parameter. In this example, connections are attempted in the order of **host1 > host2 > host3**.

```
db_hosts="hosts1:hosts2:hosts3"
```

To connect to the server, the client first tries to connect to the specified server host referring to the database location file (**databases.txt**). If the connection fails, the client then tries to connect to the first one of the secondarily specified server hosts by referring to the value of the **db_hosts** parameter in the database configuration file (**cubrid.conf**).

max_clients

max_clients is a parameter that configures the maximum number of clients (usually Broker application processes (CAS)) which allow concurrent connections to the database server. The **max_clients** parameter refers to the number of concurrent transactions. The default value is **100**.

To grantee performance while increasing the number of concurrent users in CUBRID environment, you need to make the appropriate value of the **max_clients** (**cubrid.conf**) parameter and the [MAX_NUM_APPL_SERVER](#) (**cubrid_broker.conf**) parameter. That is, you are required to configure the number of concurrent connections allowed by databases with the **max_clients** parameter. You should also configure the number of concurrent connections allowed by brokers with the **MAX_NUM_APPL_SERVER** parameter.

For example, in the **cubrid_broker.conf** file, two node of a broker where the **MAX_NUM_APPL_SERVER** value of [%query_editor] is 50 and the **MAX_NUM_APPL_SERVER** value of [%BROKER1] is 50 is trying to connect one database server, the concurrent connections (**max_clients** value) allowed by the database server can be configured as follows:

- (the maximum number of 100 by each node of a broker) * (two node of a broker) + (10 spare for database server connections of internal CUBRID process such as database server connection of CSQL Interpreter or HA log replication process) = 210

Especially, in HA environment, the value must be greater than the sum specified in **MAX_NUM_APPL_SERVER** of every broker node which connects to the same database.

Memory-Related Parameters

The following are parameters related to the memory used by the Database Server or client. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min	Max
data_buffer_size	int	512M	16M	
index_scan_oid_buffer_size	int	64K	1K	256K
sort_buffer_size	int	2M	64K	
temp_file_memory_size_in_pages	int	4	0	20
thread_stacksize	int	1048576	65536	

data_buffer_size

data_buffer_size is a parameter that configures the size of data buffer to be cached in the memory by the Database Server. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB), and terabytes (TB) respectively. If you omit the unit, bytes will be applied. The default value is 512M, and the minimum value is 16M.

The greater the value of the **data_buffer_size** parameter, the more data pages to be cached in the buffer, thus providing the advantage of decreased disk I/O cost. However, if this parameter is too large, the buffer pool can be swapped out by the operating system because the system memory is excessively occupied. It is recommended to configure the **data_buffer_size** parameter in a way the required memory size is less than two-thirds of the system memory size.

- Required memory size = data buffer size (**data_buffer_size**)

index_scan_oid_buffer_size

index_scan_oid_buffer_size is a parameter that configures the size of buffer where the OID list is to be temporarily stored during the index scan. You can set units as K, M, G and T, which stand for KB (kilobytes), MB (megabytes), GB (gigabytes) and TB (terabytes), respectively. If you omit the unit, bytes will be applied. The default value is 2M, and the minimum value is 64K.

The size of the OID buffer tends to vary in proportion to the value of the **index_scan_oid_buffer_size** parameter and the page size set when the database was created. In addition, the bigger the size of such OID buffer, the more the index scan cost. You can set the value of the **index_scan_oid_buffer_size** by considering these factors.

sort_buffer_size

sort_buffer_size is a parameter that configures the size of buffer to be used when sorting. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB), and terabytes (TB) respectively. If you omit the unit, bytes will be applied. The default value is 2M, and the minimum value is 64K.

The server assigns one sort buffer for each client request, and releases the assigned buffer memory when sorting is complete.

temp_file_memory_size_in_pages

temp_file_memory_size_in_pages is a parameter that configures the number of buffer pages to cache temporary result of a query. The default value is 4 and the maximum value is 20.

- Required memory size = the number of temporary memory buffer pages (**temp_file_memory_size_in_pages** * **page size**)
- The number of temporary memory buffer pages = the value of the **temp_file_memory_size_in_pages** parameter
- Page size = the value of the page size specified by the **-s** option of the **cuprid createdb** utility during the database creation

thread_stacksize

thread_stacksize is a parameter that configures the stack size of a thread. The default value is **1048576** bytes. The value of the **thread_stacksize** parameter must not exceed the stack size allowed by the operating system.

Disk-Related Parameters

The following are disk-related parameters for defining database volumes and storing files. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min.	Max.
db_volume_size	int	512M	20M	20G
dont_reuse_heap_file	bool	no		

log_volume_size	int	512M	20M	4G
temp_file_max_size_in_pages	int	-1		
temp_volume_path	string	NULL		
unfill_factor	float	0.1	0.0	0.3
volume_extension_path	string	NULL		

db_volume_size

A parameter used to specify the following values. The default value is **512M**.

- The default database volume size when **cubrid createdb** and **cubrid addvoldb** utility is used without **--db-volume-size** option.
- The default size of generic volume that is added automatically when database volume is full.

dont_reuse_heap_file

A parameter used to specify whether or not heap files, which are deleted when deleting the table (DROP TABLE), are to be reused when creating a new table (CREATE TABLE). If this parameter is set to 0, the deleted heap files can be reused; if it is set to 1, the deleted heap files are not used when creating a new table. The default value is 0.

log_volume_size

A parameter used to specify the default size of log volume file when cubrid createdb utility is used without --log-volume-size option. You can set units as K, M, G and T, which stand for kilobytes (KB), megabytes (MB), gigabytes (GB) and terabytes (TB) respectively. If you omit the unit, bytes will be applied. The default value is 512M.

temp_file_max_size_in_pages

A parameter used to specify the maximum number of pages to store temporary volumes in the disk, which are used for the execution of complex queries or sorting; the default value is **-1**. If this parameter is configured to the default value, unlimited number of temporary temp volumes are created and stored in the directory specified by the **temp_volume_path** parameter. If it is configured to 0, the administrator must create permanent temp volumes manually by using the **cubrid addvoldb** utility because temporary temp volumes are not created automatically.

temp_volume_path

A parameter used to specify the directory in which to create temporary temp volumes used for the execution of complex queries or sorting. The default value is the volume location configured during the database creation.

unfill_factor

A parameter used to specify the rate of disk space to be allocated in a heap page for data updates. The default value is **0.1**. That is, the rate of free space is configured to 10%. In principle, data in the table is inserted in physical order. However, if the size of the data increases due to updates and there is not enough space for storage in the given page, performance may degrade because updated data must be relocated to another page. To prevent such a problem, you can configure the rate of space for a heap page by using the **unfill_factor** parameter. The allowable maximum value is 0.3 (30%). In a database where data updates rarely occur, you can configure this parameter to 0.0 so that space will not be allocated in a heap page for data updates. If the value of the **unfill_factor** parameter is negative or greater than the maximum value, the default value (**0.1**) is used.

volume_extension_path

A parameter used to specify the directory where automatically extended volumes are to be created. The default value is the volume location configured during the database creation.

Error Message-Related Parameters

The following are parameters related to processing error messages recorded by CUBRID. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value
call_stack_dump_activation_list	string	NULL
call_stack_dump_deactivation_list	string	NULL
call_stack_dump_on_error	bool	no
error_log	string	cub_client.err, cub_server.err
error_log_level	string	SYNTAX
error_log_warning	bool	no
error_log_size	int	8000000

call_stack_dump_activation_list

call_stack_dump_activation_list is a parameter that specifies a certain error number for which a call stack is to be dumped as an exception even when you configure that a call stack will not be dumped for any errors. Therefore, the **call_stack_dump_activation_list** parameter is effective only when **call_stack_dump_on_error=no**. The following example shows how to configure the parameter so that call stacks will not be dumped for any errors, except the ones whose numbers are -115 and -116.

```
call_stack_dump_on_error= no
call_stack_dump_activation_list=-115,-116
```

call_stack_dump_deactivation_list

call_stack_dump_deactivation_list is a parameter that specifies a certain error number for which a call stack is not to be dumped when you configure that a call stack will be dumped for any errors. Therefore, the **call_stack_dump_deactivation_list** parameter is effective only when **call_stack_dump_on_error=yes**. The following example shows how to configure the parameter so that call stacks will be dumped for any errors, except the ones whose numbers are -115 and -116.

```
call_stack_dump_on_error= yes
call_stack_dump_deactivation_list=-115,-116
```

call_stack_dump_on_error

call_stack_dump_on_error is a parameter that determines whether or not to dump a call stack when an error occurs in the Database Server. If this parameter is configured to no, a call stack for any errors is not dumped. If it is configured to yes, a call stack for all errors is dumped. The default value is **no**.

error_log

error_log is a server/client parameter that specifies the name of the error log file when an error occurs in the database server. The name of the error log file must be in the form of `<database_name>_<date>_<time>.err`. However, the naming rule of the error log file does not apply to errors for which the system cannot find the Database Server information. Therefore, error logs are recorded in the **cubrid.err** file. The error log file **cubrid.err** is stored in the **\$CUBRID/log/server** directory.

error_log_level

error_log_level is a server parameter that specifies a error message to be stored based on severity. There are five different levels which ranges from **NOTIFICATION** (lowest level), **WARNING**, **SYNTAX**, **ERROR**, and **SYNTAX** (highest level). An error message with **SYNTAX**, **ERROR**, and **FATAL** levels are stored in the log file if severity of error is **SYNTAX**, default value.

error_log_warning

The server parameter **error_log_warning** specifies whether or not error messages with a severity level of **WARNING** are to be displayed. Its default value is no. Therefore, only error messages with levels other than **WARNING** will be stored even when it is set to **error_log_level = NOTIFICATION**. For this reason, you must set **error_log_warning = yes** to store **WARNING** messages to an error log file.

error_log_size

error_log_size is a parameter that specifies the maximum number of lines per an error log file. The default value is **8,000,000**. If it reaches up the specified number, the `<database_name>_<date>_<time>.err.bak` file is created.

Concurrency/Lock Parameters

The following are parameters related to concurrency control and locks of the Database Server. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min	Max
deadlock_detection_interval_in_secs	float	1.0	0.1	
isolation_level	int	3	1	6
lock_escalation	int	100000	5	
lock_timeout_in_secs	int	-1	-1	
lock_timeout_message_type	int	0	0	2

deadlock_detection_interval_in_secs

deadlock_detection_interval_in_secs is a parameter that configures the interval (in seconds) in which deadlocks are detected for stopped transactions. If a deadlock occurs, CUBRID resolves the problem by rolling back one of the transactions. The default value is 1 second and the minimum value is 0.1 second. This value is rounded up by 0.1 sec. unit. For example, if an input value is 0.12 seconds, the value is rounded up to 0.2 seconds. Note that deadlocks cannot be detected if the detection interval is too long.

isolation_level

isolation_level is a parameter that configures the isolation level of a transaction. The higher the isolation level, the less concurrency and the less interruption by other concurrent transactions. The **isolation_level** parameter can be configured to an integer value from 1 to 6, which represent isolation levels, or character strings. The default value is **TRAN_REP_CLASS_UNCOMMIT_INSTANCE**. For details about each isolation level and parameter values, see [Setting Isolation Level](#) and the following table.

Isolation Level	isolation_level Parameter Value
SERIALIZABLE	"TRAN_SERIALIZABLE" or 6
REPEATABLE READ CLASS with REPEATABLE READ INSTANCES	"TRAN_REP_CLASS_REP_INSTANCE" or "TRAN_REP_READ" or 5
REPEATABLE READ CLASS with READ COMMITTED INSTANCES(or CURSOR STABILITY)	"TRAN_REP_CLASS_COMMIT_INSTANCE" or "TRAN_READ_COMMITTED" or "TRAN_CURSOR_STABILITY" or 4
REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES	"TRAN_REP_CLASS_UNCOMMIT_INSTANCE" or "TRAN_READ_UNCOMMITTED" or 3
READ COMMITTED CLASS with READ COMMITTED	"TRAN_COMMIT_CLASS_COMMIT_INSTANCE" or 2

INSTANCES

READ COMMITTED CLASS "TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE" or 1
with READ UNCOMMITTED
INSTANCES

- **TRAN_SERIALIZABLE** : This isolation level ensures the highest level of consistency. For details, see [SERIALIZABLE](#).
- **TRAN_REP_CLASS_REP_INSTANCE** : This isolation level can occur phantom read. For details, see [REPEATABLE READ CLASS with REPEATABLE READ INSTANCES](#).
- **TRAN_REP_CLASS_COMMIT_INSTANCE** : This isolation level can occur unrepeatable read. For details, see [REPEATABLE READ CLASS with READ COMMITTED INSTANCES](#).
- **TRAN_REP_CLASS_UNCOMMIT_INSTANCE** : This isolation level can occur dirty read. For details, see [REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES](#).
- **TRAN_COMMIT_CLASS_COMMIT_INSTANCE** : This isolation level can occur unrepeatable read. It allows modification of table schema by current transactions while data is being retrieved. For details, see [READ COMMITTED CLASS with READ COMMITTED INSTANCES](#).
- **TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE** : This isolation level can occur dirty read. It allows modification of table schema by current transactions while data is being retrieved. For details, see [READ COMMITTED CLASS with READ UNCOMMITTED INSTANCES](#).

lock_escalation

lock_escalation is a parameter that specifies the maximum number of locks permitted before row level locking is extended to table level locking. The default value is **100,000**. If the value of the **lock_escalation** parameter is small, the overhead by memory lock management is small as well; however, the concurrency decreases. On the other hand, if the configured value is large, the overhead is large as well; however, the concurrency increases.

lock_timeout_in_secs

lock_timeout_in_secs is a client parameter that configures the lock waiting time. If the lock is not permitted within the specified time period, the given transaction is canceled, and an error message is returned. If the parameter is configured to **-1**, which is the default value, the waiting time is infinite until the lock is permitted. If it is configured to 0, there is no waiting for locks.

lock_timeout_message_type

lock_timeout_message_type is a parameter that configures the level of information that is to be included in the message returned when a lock timeout occurs. If the parameter is configured to **0**, which is the default value, the information about lock ownership is not included in the message. If it is configured to 1, single lock ownership information is included. If it is configured to 2, all information about lock ownership is included.

- If **lock_timeout_message_type = 0**

```
ERROR: Your transaction (index 3, cub user@cdfs006.cub|15668) timed out waiting
on X_LOCK lock on instance 0|636|34 of class participant. You are waiting for
user(s) to finish.
```

- If **lock_timeout_message_type = 1**

```
ERROR : Your transaction (index 3, cub user@cdfs006.cub|15668) timed out waiting
on X LOCK lock on instance 0|636|34 of class participant. You are waiting for user(s)
cub_user@cdfs006.cub|15615 to finish.
```

- If **lock_timeout_message_type = 2**

```
ERROR: Your transaction (index 3, cub user@cdfs006.cub|15668) timed out waiting
on X LOCK lock on instance 0|636|34 of class participant. You are waiting for user(s)
cub_user@cdfs006.cub|15615, cub_user@cdfs006.cub|15596 to finish.
```

Logging-Related Parameters

The following are parameters related to logs used for database backup and restore. The types and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min.	Max.
adaptive_flush_control	bool	yes		
background_archiving	bool	yes		
checkpoint_every_npages	int	10000	10	
checkpoint_interval_in_mins	int	720	1	
force_remove_log_archives	bool	yes		
log_buffer_size	int	2 MB	192 KB	
log_max_archives	int	INT_MAX	0	
max_flush_pages_per_second	int	10000	1	INT_MAX
page_flush_interval_in_msecs	int	0	-1	
sync_on_nflush	int	200	1	INT_MAX

adaptive_flush_control

The **adaptive_flush_control** parameter automatically adjusts the flush capacity at every 50 ms depending on the current status of the flushing operation. Its default value is **yes**. That is, this capacity is increased if a large number of **INSERT** or **UPDATE** operations are concentrated at a certain point of time and the number of flushed pages reaches the **max_flush_pages_per_second** parameter value; and is decreased otherwise. In the same way, you can distribute the I/O load by adjusting the flush capacity on a regular basis depending on the workload.

background_archiving

The **background_archiving** parameter generates a temporary archive log periodically at a specific time. It is useful when balancing disk I/O load which has been caused by archiving logs. The default is **yes**.

checkpoint_every_npages

The **checkpoint_every_npages** parameter configures checkpoint interval by log page. The default value is **10,000**.

You can distribute disk I/O overload at the checkpoint by specifying lower number in the **checkpoint_every_npages** parameter, especially in the environment where **INSERT/UPDATE** are heavily loaded at a specific time.

Checkpoint is a job to record every modified page in data buffers to database volumes (disk) at a specific point. It can restore data back to the latest checkpoint if database failure occurs. It is important to choose efficient checkpoint interval because large increase of log files stored in a disk may affect database operation, causing unnecessary disk I/O.

The **checkpoint_interval_in_mins** and **checkpoint_every_npages** parameters are related to setting checkpoint cycle. The checkpoint is periodically executed whenever the time specified in **checkpoint_interval_in_mins** parameter has elapsed or the number of log pages specified in **checkpoint_every_npages** parameter has reached.

checkpoint_interval_in_mins

The **checkpoint_interval_in_mins** parameter configures execution period of checkpoint in minutes. The default value is **720**.

force_remove_log_archives

The **force_remove_log_archives** parameter configures whether to allow the deletion of the files other than the recent log archive files of which the number is specified by **log_max_archives**. The default value is **yes**.

If the value is set to yes, the files will be deleted other than the recent log archive files for which the number is specified by **log_max_archives**. If it is set to no, the log archive files will not be deleted. Exceptionally, if **ha_mode** is set to on, the files other than the log archive files required for the HA related process and the recent log archive files of which the number is specified by **log_max_archives** will be deleted.

If you want to build a CUBRID HA environment, see [Configuration](#).

log_buffer_size

The **log_buffer_size** parameter configures the size of log buffer to be cached in the memory. There are four types of unit available: K, M, G, and T; K stands for kilobytes (KB), M stands for megabytes (MB), G stands for gigabytes (GB), and T stands for terabytes (TB). If unit is omitted, byte-unit is applied and the default value is **2M**.

If the value of the **log_buffer_size** parameter is large, performance can be improved (due to the decrease in disk I/O) in an environment where transactions are long and numerous. It is recommended to configure an appropriate value considering the memory size and operations of the system where CUBRID is installed.

- Required memory size = the size of log buffer (**log_buffer_size**)

log_max_archives

The **log_max_archives** parameter configures the maximum number of archive log files. The minimum value is 0 and default value is **INT_MAX**. Its operations can differ depending on the configuration of **force_remove_log_archives**. For example, when **log_max_archives** is 3 and **force_remove_log_archives** is **yes** in the cubrid.conf file, the most recent three archive log files are recorded and when a fourth archiving log file is generated, the oldest archive log file is automatically deleted; the information about the deleted archive logs are recorded in the *_**lginf** file.

However, if an active transaction still refers to an existing archive log file, the archive log file will not be deleted. That is, if a transaction starts at the point that the first archive log file is generated, and it is still active until the fifth archive log is generated, the first archive log file cannot be deleted.

For how to set up the CUBRID HA environment, see [cubrid.conf](#).

max_flush_pages_per_second

The **max_flush_pages_per_second** parameter configures the maximum flush capacity when the flushing operation is performed from a buffer to a disk. Its default value is **10,000**. That is, you can prevent concentration of I/O load at a certain point of time by configuring this parameter to control the maximum flush capacity per second.

If a large number of **INSERT** or **UPDATE** operations are concentrated at a certain point of time, and the flush capacity reaches the maximum capacity set by this parameter, only log pages are flushed to the disk, and data pages are no longer flushed. Therefore, you must set an appropriate value for this parameter considering the workload of the service environment.

page_flush_interval_in_msecs

The **page_flush_interval_in_msecs** parameter configures the interval in milliseconds (msec.) at which dirty pages in a data buffer are flushed to a disk. Its default value is **0**. When the minimum value is set to -1, it work as that is set to 0. This is a parameter that is related to I/O load and buffer concurrency. For this reason, you must set its value in consideration of the workload of the service environment.

sync_on_nflush

The **sync_on_nflush** parameter configures the interval in pages between after data and log pages are flushed from buffer and before they are synchronized with FILE I/O of operating system. Its default value is **200**. That is, the CUBRID Server performs synchronization with the FILE I/O of the operating system whenever 200 pages have been flushed. This is also a parameter related to I/O load.

Transaction Processing-Related Parameters

The following are parameters for improving transaction commit performance. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min	Max
async_commit	bool	no		
group_commit_interval_in_msecs	int	0		0

async_commit

async_commit is a parameter that activates the asynchronous commit functionality. If the parameter is configured to **no**, which is the default value, the asynchronous commit is not performed; if it is configured to **yes**, the asynchronous commit is executed. The asynchronous commit is a functionality that improves commit performance by completing the commit for the client before commit logs are flushed on the disk and having the log flush thread (LFT) perform log flushing in the background. Note that already committed transactions cannot be restored if a failure occurs on the Database Server before log flushing is performed.

group_commit_interval_in_msecs

group_commit_interval_in_msecs is a parameter that configures the interval (in milliseconds), at which the group commit is to be performed. If the parameter is configured to **0**, which is the default value, the group commit is not performed. The group commit is a functionality that improves commit performance by combining multiple commits that occurred in the specified time period into a group so that commit logs are flushed on the disk at once.

Statement/Type-Related Parameters

The following are parameters related to SQL statements and data types supported by CUBRID. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value
add_column_update_hard_default	bool	no
alter_table_change_type_strict	bool	no
ansi_quotes	bool	yes
block_ddl_statement	bool	no
block_nowhere_statement	bool	no
compat_numeric_division_scale	bool	no
default_week_format	int	0
group_concat_max_len	int	1024
intl_mbs_support	bool	no
no_backslash_escapes	bool	yes
only_full_group_by	bool	no
oracle_style_empty_string	bool	no
pipes_as_concat	bool	yes
plus_as_concat	bool	yes
require_like_escape_character	bool	no
return_null_on_function_errors	bool	no

add_column_update_hard_default

The **add_column_update_hard_default** parameter is used to determine whether or not to provide the `hard_default` value as the input value for a column when you add a new column to the **ALTER TABLE ... ADD COLUMN** clause.

If a value for this parameter is set to `yes`, enter a new input value of a column as a hard default value when you have **NOT NULL** constraints but no **DEFAULT** constraints. If the parameter value is set to `no`, enter **NULL**, even if **NOT NULL** constraints exist. If a value for this parameter is set to `yes` and there is no hard default value for the column type to add, an error message will be displayed and a roll-back occurs. For the hard default for each type, see the [CHANGE Clause](#) of the **ALTER TABLE** statement.

```
-- add column update hard default=no

CREATE TABLE tbl (i INT);
INSERT INTO tbl VALUES (1), (2);
ALTER TABLE tbl ADD COLUMN j INT NOT NULL;

SELECT * FROM TBL;

           i           j
=====
           2           NULL
           1           NULL

-- add column update hard default=yes

CREATE TABLE tbl (i int);
INSERT INTO tbl VALUES (1), (2);
ALTER TABLE tbl ADD COLUMN j INT NOT NULL;

SELECT * FROM tbl;

           i           j
=====
           2             0
           1             0
```

alter_table_change_type_strict

The **alter_table_change_type_strict** parameter is used to determine whether or not to allow the conversion of column values according to the type change, and the default value is **no**. If a value for this parameter is set to `no`, the value may be changed when you change the column types or when you add **NOT NULL** constraints; if it is set to `yes`, the value is not changed. For details, see **CHANGE Clause** in the [CHANGE/MODIFY Clause](#).

ansi_quotes

ansi_quotes is a parameter that enclose symbols and character string to handle identifiers. The default value is **yes**. If this parameter value is set to **yes**, double quotations are handled as identifier symbols and single quotations are handled as character string symbols. If it is set to **no**, both double and single quotations are handled as character string symbols.

block_ddl_statement

block_ddl_statement is a parameter that restricts the execution of DDL (Data Definition Language) statements by the client. If the parameter is configured to `no`, the given client is allowed to execute DDL statements. If it is configured to `yes`, the client is not permitted to execute DDL statements. The default value is **no**.

block_nowhere_statement

block_nowhere_statement is a parameter that restricts the execution of **UPDATE/DELETE** statements without a condition clause (**WHERE**) by the client. If the parameter is configured to `no`, the given client is allowed to execute **UPDATE/DELETE** statements without a condition clause. If it is configured to `yes`, the client is not permitted to execute **UPDATE/DELETE** statements without a condition clause. The default value is **no**.

compact_numeric_division_scale

compact_numeric_division_scale is a parameter that configures the scale to be displayed in the result (quotient) of a division operation. If the parameter is configured to **no**, the scale of the quotient is 9 if it is configured to **yes**, the scale is determined by that of the operand. The default value is **no**.

default_week_format

The **default_week_format** parameter is used to set the default value for the *mode* attribute of the WEEK function. The default value is 0. For details, see [WEEK Function](#).

group_concat_max_len

The **group_concat_max_len** parameter is used to limit the return value size of the **GROUP_CONCAT** function. The default value is 1024 bytes, the minimum value is 4 bytes, and the maximum value is 33,554,432 bytes. If the return value of the **GROUP_CONCAT** function exceeds the limitation, **NULL** will be returned.

intl_mbs_support

intl_mbs_support is a parameter that specifies whether or not to support multi-byte character set. If the parameter is configured to **no**, a multi-byte character set is not allowed if it is configured to **yes**, a multi-byte character set is allowed. To improve performance, it is recommended to configure the **intl_mbs_support** parameter to **no** and use alphabets for table and column names because operation cost for supporting multi-byte character set is high.

no_backslash_escapes

The **no_backslash_escapes** is used to determine whether or not to use backslash (\) as an escape character, and the default value is **yes**. If a value for this parameter is set to **no**, backslash (\) will be used as an escape character; if it is set to **yes**, backslash (\) will be used as a normal character. For details, see [Escape Special Characters](#).

only_full_group_by

only_full_group_by is a parameter that specifies whether extended syntax about using **GROUP BY** statement is used or not.

If this parameter value is set to **no**, an extended syntax is applied thus, a column that is not specified in the **GROUP BY** statement can be specified in the **SELECT** column list. If it is set to **yes**, a column that is only specified in the **GROUP BY** statement can be the **SELECT** column list.

The default value is **no**. Therefore, specify the **only_full_group_by** parameter value to **yes** to execute queries by SQL standards. Because the extended syntax is not applied in this case, an error below is displayed.

```
ERROR: Attributes exposed in aggregate queries must also appear in the group by clause.
```

oracle_style_empty_string

oracle_style_empty_string is a parameter that improves compatibility with other DBMS (Database Management Systems) and specifies whether or not to process empty strings as **NULL** as in Oracle DBMS. If the **oracle_style_empty_string** parameter is configured to **no**, the character string is processed as a valid string if it is configured to **yes**, the empty string is processed as **NULL**.

pipes_as_concat

pipes_as_concat is a parameter about using a double pipe symbol. The default value is **yes**. If this parameter value is set to **yes**, a double pipe symbol is handled as a concatenation operator if **no**, it is handled as the **OR** operator.

plus_as_concat

The **plus_as_concat** parameter is a parameter for the use of the + operator, and the default value is yes. If a value for this parameter is set to yes, the + operator will be interpreted as a concatenation operator; if it is set to no, the operator will be interpreted as a numeric operator.

```
-- plus as concat = yes
SELECT '1'+ '1';
      '1'+ '1'
=====
      '11'  SELECT '1'+ 'a';
      '1'+ 'a'
=====
      '1a'
```

```
-- plus_as_concat = no
SELECT '1'+ '1';
      '1'+ '1'
=====
2.0000000000000000e+000

SELECT '1'+ 'a';

ERROR: Cannot coerce 'a' to type double.
```

require_like_escape_characterintl_mbs_support

The **require_like_escape_character** parameter is used to determine whether or not to use an ESCAPE character in the **LIKE** clause, and the default value is no. If a value for this parameter is set to yes and a value for **no_backslash_escapes** is set to no, backslash (\) will be used as an ESCAPE character in the strings of the LIKE clause, otherwise you should specify an ESCAPE character by using the **LIKE... ESCAPE** clause. For details, see [LIKE Predicate](#).

return_null_on_function_errors

The **return_null_on_function_errors** parameter is used to define actions when errors occur in some SQL functions, and the default value is no. If a value for this parameter is set to yes, **NULL** is returned; if it is set to no, an error is returned when the error occurs in functions, and the related message is displayed.

The following SQL functions are affected by this system parameter.

- ADDTIME
- DATEDIFF
- DAY
- DAYOFMONTH
- DAYOFWEEK
- DAYOFYEAR
- FROM_DAYS
- FROM_UNIXTIME
- HOUR
- LAST_DAY
- MAKEDATE
- MAKETIME
- MINUTE
- MONTH
- QUARTER
- SEC_TO_TIME
- SECOND
- TIME

- TIME_TO_SEC
- TIMEDIFF
- TO_DAYS
- WEEK
- WEEKDAY
- YEAR

```
-- return null on function errors=no
SELECT HOUR('2010-01-01');

ERROR: Conversion error in time format.

-- return null on function errors=yes
SELECT HOUR('2010-01-01');

    hour('2010-01-01')
=====
NULL
```

Query Cache-Related Parameters

The following are parameters related to the query cache functionality that provides execution results cached for the same **SELECT** statement. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min	Max
max_plan_cache_entries	int	1,000		

max_plan_cache_entries

max_plan_cache_entries is a parameter that configures the maximum number of query plans to be cached in the memory. If the **max_plan_cache_entries** parameter is configured to -1 or 0, generated query plans are not stored in the memory cache; if it is configured to an integer value equal to or greater than 1, a specified number of query plans are cached in the memory. Also, the value of this parameter must be configured to an integer value equal to or greater than 1 to use the query cache functionality that caches the results of the same query.

The following example shows how to cache up to 1,000 queries.

```
max_plan_cache_entries=1000
```

Utility-Related Parameters

The following are parameters related to utilities used in CUBRID. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min.	Max.
backup_volume_max_size_bytes	int	-1	1024*32	
communication_histogram	bool	no		
compactdb_page_reclaim_only	int	0		
csql_history_num	int	50	1	200

backup_volume_max_size_bytes

backup_volume_max_size_bytes is a parameter that configures the size of the backup volume file created by the **cubrid backupdb** utility in byte units. If the parameter is configured to -1, which is the default value, the created backup volume is not partitioned; otherwise, the backup volume is partitioned as much as it is specified size.

there is no limit to the size of the backup volume to be created. If it is not configured, the size of the backup volume is allowed up to the size limit of the storage media.

communication_histogram

communication_histogram is a parameter related to the **cubrid statdump** utility. It is related to [Session Commands ;h](#) of the CSQL Interpreter and the default value is **no**. For details, see [Outputting Statistics Information of Server](#).

compactdb_page_reclaim_only

compactdb_page_reclaim_only is a parameter related to the **compactdb** utility, which compacts the storage of already deleted objects to reuse OIDs of the already assigned storage. Storage optimization with the **compactdb** utility can be divided into three steps. The optimization steps can be selected through the **compactdb_page_reclaim_only** parameter. If the parameter is configured to **0**, which is the default value, step 1, 2 and 3 are all performed, so the storage is optimized in data, table and file units. If it is configured to 1, step 1 is skipped to have the storage optimized in table and file units. If it is configured to 2, steps 1 and 2 are skipped to have the storage optimized only in file units.

- Step 1 : Optimizes the storage only in data units.
- Step 2 : Optimizes the storage in table units.
- Step 3 : Optimizes the storage in file (heap file) units.

csql_history_num

csql_history_num is a parameter related to the CSQL Interpreter, and configures the number of SQL statements to be stored in the history of the CSQL Interpreter. The default value is **50**.

HA-Related Parameters

The following are HA-related parameters. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value
ha_mode	string	off

ha_mode

The **ha_mode** parameter is used to set CUBRID HA, and the default value is **off**.

- off : CUBRID HA is not used.
- on : CUBRID HA is used using the configured node as a node for failover.
- replica : CUBRID HA is used without using the configured node as a node for failover.

To use the CUBRID HA feature, you should set HA related parameters in the **cubrid_ha.conf** file in addition to the **ha_mode** parameter. For details, see [CUBRID HA](#).

Other Parameters

The following are other parameters. The type and value range for each parameter are as follows:

Parameter Name	Type	Default Value	Min.	Max.
access_ip_control	bool	no		
access_ip_control_file	string			
auto_restart_server	bool	yes		
index_scan_in_oid_order	bool	no		
index_unfill_factor	float	0.05	0	0.5

insert_execution_mode	int	1	1	7
java_stored_procedure	bool	no		
multi_range_optimization_limit	int	100	0	10000
pthread_scope_process	bool	yes		
server	string			
service	string			
session_state_timeout	int	21600 (6 hours)	60 (1 minute)	31536000 (1 year)
single_byte_compare	bool	no		
use_orderby_sort_limit	bool	yes		

access_ip_control

A parameter used to specify whether to use feature limiting the IP addresses that allow server access. The default value is **no**. For details, see Database [Server Access Limitation](#).

access_ip_control_file

A parameter used to specify a file name in which the list of IPs allowed by servers is stored. If **access_ip_control** value is set to **yes**, database server allows the list of IP addresses only stored in the file specified by this parameter. For details, [Limiting Database Server Access](#).

auto_restart_server

A parameter used to specify whether to restart the process when it stops due to fatal errors being occurred in database server process. If **auto_restart_server** value is set to **yes**, the server process automatically restarts when it has stopped due to errors; it does not restart in case it stops by following normal process (by using **STOP** command).

index_scan_in_oid_order

A parameter used to specify the result data to be retrieved in OID order after the index scan. If the parameter is set to **no**, which is the default value, results are retrieved in data order; if it is set to **yes**, they are retrieved in OID order.

index_unfill_factor

If there is no free space because index pages are full when the **INSERT** or **UPDATE** operation is executed after the first index is created, the split of index page nodes occurs. This substantially affects the performance by increasing the operation time. **index_unfill_factor** is a parameter that specifies the percent of free space defined for each index page node when an index is created. The **index_unfill_factor** value is applied only when an index is created for the first time. The percent of free space defined for the page is not maintained dynamically. Its value ranges between 0 and 0.35. The default value is **0.05**.

If an index is created without any free space for the index page node (**index_unfill_factor** is set to 0), the split of index page nodes occurs every time an additional insertion is made. This may degrade the performance.

If the value of **index_unfill_factor** is large, a large amount of free space is available when an index is created. Therefore, better performance can be obtained because the split of index nodes does not occur for a relatively long period of time until the free space for the nodes is filled after the first index is created.

If this value is small, the amount of free space for the nodes is small when an index is created. Therefore, it is likely that the index nodes are split by **INSERT** or **UPDATE** because the free space for the index nodes is filled in a short period of time.

insert_execution_mode

insert_execution_mode has execution modes ranging from 1 to 7. Queries are usually executed on the server according to the query plan created by the client, but this parameter is used to directly insert queries on the server side. A selected execution mode is executed directly on the server, and other execution modes are executed on the client. This parameter can be used to perform an INSERT operation to the server in an environment in which dirty reading of INSERTed data is required, or in which the memory capacity of the client is limited.

The following are three types of **INSERT** statements for execution modes. This parameter can be set through a combination of integer values corresponding to each execution mode.

- **INSERT_SELECT** : When using the **SELECT** statement in the **INSERT** statement.

```
INSERT INTO code2(s_name, f_name) SELECT s_name, f_name from code;
```

- **INSERT_VALUES** : The common **INSERT** statement.

```
INSERT INTO code2(s_name, f_name) VALUES ('S', 'Silver');
```

- **INSERT_DEFAULT** : When inserting the default value because a column with the default value is omitted in the **INSERT** statement.

```
CREATE TABLE code2(s_name char(1) DEFAULT ' ', f_name varchar(40));
INSERT INTO code2(f_name) DEFAULT VALUES;
```

- **INSERT_REPLACE** : For example, when the **REPLACE** statement is executed, the corresponding integer value is 8.

```
CREATE TABLE code2(s_name char(1) NOT NULL UNIQUE, f_name varchar(40));
REPLACE INTO code2 VALUES ('S', 'Silver');
```

- **INSERT_ON_DUP_KEY_UPDATE** : In addition, when the **ON DUPLICATE KEY UPDATE** clause is specified in the **INSERT** statement, the corresponding integer value is 16.

```
CREATE TABLE code2(s_name char(1) NOT NULL UNIQUE, f_name varchar(40));
INSERT INTO code2 VALUES ('S', 'Silver') ON DUPLICATE KEY UPDATE f_name='Silver';
```

The sum of the execution mode values above is the execution mode to be configured.

- Example 1 : If you want to execute **INSERT_SELECT** and **INSERT_VALUES** on the server, the **insert_execution_mode** is 3. (1 + 2 = 3)
- Example 2 : If you want to execute **INSERT_SELECT**, **INSERT_DEFAULT**, **INSERT_REPLACE**, an **INSERT_ON_DUP_KEY_UPDATE** on the server, the **insert_execution_mode** is 29(1+4+8+16=29).

java_stored_procedure

A parameter used to specify whether to use Java stored procedures by running the Java Virtual Machine (JVM). If the parameter is set to **no**, which is the default value, JVM is not executed; if it is set to **yes**, JVM is executed so you can use Java stored procedures. Therefore, configure the parameter to yes if you plan to use Java stored procedures.

multi_range_optimization_limit

If the number of rows specified by the **LIMIT** clause in the query, which has multiple ranges (col IN (?, ?, ...,?)) and is available to use an index, is within the number specified by the **multi_range_optimization_limit** parameter, the optimization for the way of index sorting will be performed. The default value is 100.

For example, if a value for this parameter is set to 50, **LIMIT 10** means that it is within the value specified by this parameter, so that the values that meet the conditions will be sorted to produce the result. If **LIMIT** is 60, it means that it exceeds the parameter configuration value, so that it gets and sorts out all values that meet the conditions.

Depending on the setting value, the differences are made between collecting the result with on-the-fly sorting of the intermediate values and sorting the result values after collecting them, and the bigger value could make more unfavorable performance.

pthread_scope_process

A parameter used to specify the contention scope of threads. It only applies to AIX systems. If the parameter is set to **no**, the contention scope becomes **PTHREAD_SCOPE_SYSTEM**; if it is set to **yes**, it becomes **PTHREAD_SCOPE_PROCESS**. The default value is **yes**.

server

A parameter used to register database server process that starts automatically when the CUBRID service starts.

service

A parameter used to register process that starts automatically when the CUBRID service starts. There are four types of processes: **server**, **broker**, **manager**, and **heartbeat**. Three processes are usually registered as in **service=server,broker,manager**.

- If the parameter is set to **server**, the database process specified by the **@server** parameter gets started.
- If the parameter is set to **broker**, the Broker process gets started.
- If the parameter is set to **manager**, the manager process gets started.
- If the parameter is set to **heartbeat**, the HA-related process gets started.

session_state_timeout

A parameter used to define how long the CUBRID session data will be kept. The session data will be deleted when the driver terminates the connection or the session time is expired. The session time will expire if a client does not have any requests until a value specified in **session_state_timeout**.

The default value is **21600** seconds (6 hours).

The following are CUBRID session data.

- Custom variables defined by **SET**.
- **PREPARE** statements.
- **LAST_INSERT_ID**
- Number of records affected by the last executed statement(**ROW_COUNT**)

Custom variables defined by **SET** and **PREPARE** statements can be deleted by **DROP/DEALLOCATE** statements before session timeout.

single_byte_compare

A parameter used to specify whether or not to compare strings in single byte units. If the parameter is set to **no**, which is the default value, strings are compared in two byte units; if it is set to **yes**, they are compared in single byte units. That is, you can retrieve/compare strings on data stored as UTF-8.

use_orderby_sort_limit

A parameter used to specify whether to keep the intermediate result of sorting and merging process in the statement including the **ORDER BY ... LIMIT row_count** clause as many as **row_count**. If it is set to **yes**, you can decrease unnecessary comparing and merging processes because as many as intermediate results will be kept as the value of **row_count**. The default value is **yes**.

Changing Database Server Configuration

Editing the Configuration File

You can add/delete parameters or change parameter values by manually editing the system parameter configuration file (**cubrid.conf**) in the **\$CUBRID/conf** directory.

The following parameter syntax rules are applied when configuring parameters in the configuration file:

- Parameter names are not case-sensitive.
- The name and value of a parameter must be entered in the same line.
- An equal sign (=) can be used to configure the parameter value. Spaces are allowed before and after the equal sign.
- If the value of a parameter is a character string, enter the character string without quotes. However, use quotes if spaces are included in the character string.

Using SQL Statements

Description

You can configure a parameter value by using SQL statements in the CSQL Interpreter or CUBRID Manager's Query Editor. Note that you cannot change every parameter. For updatable parameters, see [cubrid.conf Configuration File and Default Parameters](#).

Syntax

```
SET SYSTEM PARAMETERS 'parameter_name=value [{; name=value}...]'
```

parameter_name is the name of a client parameter whose value is editable. In this syntax, *value* is the value of the given parameter. You can change multiple parameter values by separating them with semicolons (;). You must take caution when you apply changes of parameter values.

Example

The following example shows how to retrieve the result of an index scan in OID order and configure the number of queries to be stored in the history of the CSQL Interpreter to 70.

```
SET SYSTEM PARAMETERS 'index_scan_in_oid_order=1; csql_history_num=70'
```

Using Session Commands of the CSQL Interpreter

Description

You can configure system parameter values by using session commands (;SEt) in the CSQL Interpreter. Note that you cannot change every parameter. For updatable parameters, see [cubrid.conf Configuration File and Default Parameters](#).

Example

The following example shows how to configure the **block_ddl_statement** parameter to 1 so that execution of DDL statements is not allowed.

```
csql> ;se block ddl statement=1
=== Set Param Input ===
block_ddl_statement=1
```

Broker Configuration

Broker Configuration File and Default Parameters

Broker System Parameters

The following table shows the Broker parameters available in the Broker configuration file (**cupbrid_broker.conf**). For details, see [Common Parameters](#) and [Parameter by Broker](#).

Type	Parameter Name	Type	Default Value
Common	ACCESS_CONTROL	bool	no
	ACCESS_CONTROL_FILE	string	
	ADMIN_LOG_FILE	string	log/broker/cubrid_broker.log
	MASTER_SHM_ID	int	30001
Broker	ACCESS_LIST	string	-
	ACCESS_LOG	string	ON
	ACCESS_MODE	string	RW
	APPL_SERVER	string	CAS
	APPL_SERVER_MAX_SIZE	int	32-bit Windows : 40 64-bit Windows : 80 Linux : 0
	APPL_SERVER_MAX_SIZE_HARD_LIMIT	int	1024
	APPL_SERVER_PORT	int	BROKER_PORT+1
	APPL_SERVER_SHM_ID	int	30000
	BROKER_PORT	int	30000 (max. : 65535)
	CCI_DEFAULT_AUTOCOMMIT	string	ON
	CCI_PCONNECT	string	OFF
	ERROR_LOG_DIR	string	log/broker/error_log
	KEEP_CONNECTION	string	AUTO
	LOG_BACKUP	string	OFF
	LOG_DIR	string	log/broker/sql_log
	LONG_QUERY_TIME	int	60
	LONG_TRANSACTION_TIME	int	40
	MAX_NUM_APPL_SERVER	int	40
	MAX_PREPARED_STMT_COUNT	int	2000 (min. : 1)
	MAX_QUERY_TIMEOUT	int	0 (max. : 86400 (seconds))
	MAX_STRING_LENGTH	int	-1
	MIN_NUM_APPL_SERVER	int	5
	PREFERRED_HOSTS	string	-
	SELECT_AUTO_COMMIT	string	OFF
	SERVICE	string	ON
	SESSION_TIMEOUT	int	300

	SLOW_LOG	string ON
	SLOW_LOG_DIR	string log/broker/sql_log
	SOURCE_ENV	string cubrid.env
	SQL_LOG	string ON
	SQL_LOG_MAX_SIZE	int 100000
	STATEMENT_POOLING	string ON
TIME_TO_KILL	int	120

Default Parameters

The **cubrid_broker.conf** file, a default Broker configuration file created during CUBRID installation, includes some parameters that must be modified by default. If you want to modify the values of parameters that are not included in the configuration file by default, you can add or modify one yourself.

The following is the content of the **cubrid_broker.conf** file provided by default.

```
[broker]
MASTER SHM ID           =30001
ADMIN LOG FILE          =log/broker/cubrid_broker.log

[%query_editor]
SERVICE                =ON
BROKER PORT             =30000
MIN_NUM_APPL_SERVER    =5
MAX_NUM_APPL_SERVER    =40
APPL_SERVER_SHM_ID     =30000
LOG_DIR                 =log/broker/sql_log
ERROR_LOG_DIR          =log/broker/error_log
SQL_LOG                 =ON
TIME_TO_KILL           =120
SESSION_TIMEOUT        =300
KEEP_CONNECTION        =AUTO

[%BROKER1]
SERVICE                =ON
BROKER PORT             =33000
MIN_NUM_APPL_SERVER    =5
MAX_NUM_APPL_SERVER    =40
APPL_SERVER_SHM_ID     =33000
LOG_DIR                 =log/broker/sql_log
ERROR_LOG_DIR          =log/broker/error_log
SQL_LOG                 =ON
TIME_TO_KILL           =120
SESSION_TIMEOUT        =300
KEEP_CONNECTION        =AUTO
```

Broker Configuration File Related Environment Variables

You can specify the location of broker configuration file (**cubrid_broker.conf**) file by using the **CUBRID_BROKER_CONF_FILE** variable. The variable is used when executing several brokers with different configuration.

Common Parameters

The following are parameters commonly applied to entire Brokers; it is written under [broker] section.

ACCESS_CONTROL

A parameter used to limit applications which are trying to connect a broker. The default value is **OFF**. For details, see [Broker Server Access Limitation](#).

ACCESS_CONTROL_FILE

A parameter used to specify the name of a file in which a database name, a database user ID, and the list of IPs are stored. For details, see [Broker Server Access Limitation](#).

ADMIN_LOG_FILE

A parameter used to specify a file in which time of running CUBRID Broker is stored. The default value is a **log/broker/cubrid_broker.log** file.

MASTER_SHM_ID

A parameter used to specify the identifier of shared memory which is used to manage the CUBRID Broker. Its value must be unique in the system. The default value is **30001**.

Parameter by Broker

The following describes parameters to configure the environment variables of Brokers; each parameter is located under [%broker_name].

ACCESS_LIST

A parameter used to specify the name of a file where the list of IP addresses of an application which allows access to the CUBRID Broker is stored. To allow access by IP addresses access 210.192.33.* and 210.194.34.*, store them to a file (ip_lists.txt) and then assign the file name with the value of this parameter.

ACCESS_LOG

A parameter used to specify whether to store the access log of Broker. The default value is **ON**. The name of the access log file for the Broker is *broker_name_id.access* and the file is stored under **\$CUBRID/log/broker** directory.

ACCESS_MODE

A parameter used to specify default mode of Broker. The default value is **RW**. For details, see [cubrid_broker.conf](#) of "Administrator's Guide".

APPL_SERVER

A parameter used to specify types of application servers generated and managed by the CUBRID Broker. The default value is **CAS**.

APPL_SERVER_MAX_SIZE

A parameter used to specify the maximum size of the process memory usage handled by application servers (CAS); the unit is MB.

Specifying this parameter makes transactions terminate (commit or rollback) only when it is executed by a user. In contrast to this, specifying **APPL_SERVER_MAX_SIZE_HARD_LIMIT** makes transactions forcibly terminate (rollback) and restart application servers (CAS).

Note that the default values of Windows and Linux from each other.

For 32-bit Windows, the default value is **40 MB**; for 64-bit Windows, it is **80 MB**. At the time when current process size exceeds the value of **APPL_SERVER_MAX_SIZE**, broker restarts the corresponding application server. For Linux, the default value is **0**; an application server restarts in the following conditions.

- The value is zero or negative : At the point when current process size becomes twice as large as initial memory
- The value is positive : At the point when it exceeds the value specified in **APPL_SERVER_MAX_SIZE**

Note Be careful not to make the value too small because application servers may restart frequently and unexpectedly. In general, the value of `APPL_SERVER_MAX_SIZE_HARD_LIMIT` is greater than that of `APPL_SERVER_MAX_SIZE`. For details, see description of `APPL_SERVER_MAX_SIZE_HARD_LIMIT`.

APPL_SERVER_MAX_SIZE_HARD_LIMIT

A parameter used to specify the maximum size of process memory usage handled by application servers (CAS); the unit is MB and default value is **1024** MB.

Specifying this parameter makes transactions being processed forcibly terminate (rollback) and restart application servers (CAS). In contrast to this, specifying `APPL_SERVER_MAX_SIZE` makes transactions terminate only when it is executed by a user.

Note Be careful not to make the value too small because application servers may restart frequently and unexpectedly. When restarting application servers (CAS), `APPL_SERVER_MAX_SIZE` is specified to wait for normal termination of transactions although memory usage increases; `APPL_SERVER_MAX_SIZE_HARD_LIMIT` is specified to forcibly terminate transactions if memory usage exceeds the maximum value allowed. Therefore, in general, the value of `APPL_SERVER_MAX_SIZE_HARD_LIMIT` is greater than that of `APPL_SERVER_MAX_SIZE`.

APPL_SERVER_PORT

A parameters used to specify the connection port of application server (CAS) that communicates with an application; it is only supported for Windows. The default value is determined by adding plus 1 to the `BROKER_PORT` parameter value. The maximum number of application servers is specified in the `MAX_NUM_APPL_SERVER` parameter in `cubrid_broker_conf`; therefore, the maximum number of connection ports is also determined by the value of `MAX_NUM_APPL_SERVER` parameter.

On the Windows system, if firewall exists between an application and the CUBRID Broker, the communication port specified in `BROKER_PORT` and `APPL_SERVER_PORT` must be open.

APPL_SERVER_SHM_ID

A parameter used to specify the ID of shared memory used by application servers (CAS); the value must be unique within system. The default value is the same as the port value of Broker.

BROKER_PORT

A parameter used to specify the port number of the Broker; the value must be unique and smaller than 65,535. The default port value of `query_editor` Broker is 30,000 and the port value of `broker1` is 33,000.

CCI_DEFAULT_AUTOCOMMIT

A parameter used to specify whether to use automatic commit of applications developed by CCI APIs. The default value is **ON**. This parameter affects applications developed by CCI APIs or an applications using interfaces (PHP, ODBC, and OLE DB) developed by CCI; it does not affect the applications developed by JDBC.

CCI_PCONNECT

A parameter used to specify whether or not to use the CCI connection pooling. The default value is **OFF**. This parameter affects applications developed by CCI APIs or an applications using interfaces (PHP, ODBC, and OLE DB) developed by CCI; it does not affect the applications developed by JDBC.

ERROR_LOG_DIR

A parameter used to specify a default directory in which error logs about Broker is stored. The default value is `log/broker/error_log`. The log file name for Broker error is `broker_name_id.err`.

KEEP_CONNECTION

A parameter used to specify the way of connection between application servers (CAS) and application clients; it is set to one of the followings: **ON**, **OFF** or **AUTO**. If this value is **OFF**, clients are connected to servers in transaction unit; for **ON**, it is connected in connection unit. If it is **AUTO** and the number of servers is more than that of clients, transaction unit is used; in the reverse case, connection unit is used. The default value is **AUTO**.

LOG_BACKUP

A parameter used to specify whether to back up access and error log files of the Broker when CUBRID stops. The default value is set to **OFF**. An access log file (*broker_name access*) in the **\$CUBRID/log/broker** directory is deleted when CUBRID stops. If the value is set to **ON**, an access log file is stored (backed up) as *broker_name access* *yyyymmdd.hhmm* when CUBRID stops.

LOG_DIR

A parameter used to specify the directory where SQL logs are stored. The default value is **log/broker/sql_log**. The file name of the SQL logs is *broker_name_id.sql.log*.

ERROR_LOG_DIR

A parameter used to specify the directory where error logs for the Broker are stored. The default value is **log/broker/error_log**. The name of the error log file for the Broker is *broker_name_id.err*.

LONG_QUERY_TIME

A parameter used to specify execution time of query which is evaluated as long-duration query. The default value is **60** (seconds) and can be value in msec. with a decimal separator. For example, the value should be configured into 0.5 to configure 500 msec. Note that a parameter value is configured to 0, it is not evaluated as a long-duration query.

LONG_TRANSACTION_TIME

A parameter used to specify execution time of query which is evaluated as long-duration transaction. The default value is **60** (seconds) and can be value in msec. with a decimal separator. For example, the value should be configured into 0.5 to configure 500 msec. Note that a parameter is configured to 0, it is not evaluated as a long-duration transaction.

MAX_NUM_APPL_SERVER

A parameter used to specify the maximum number of simultaneous connections allowed. The default value is **40**.

In an environment where connection pool is maintained by using middleware such as DBCP or WAS, the value of **MAX_NUM_APPL_SERVER** parameter and the number of connection pools should be same.

MAX_PREPARED_STMT_COUNT

A parameter used to limit the number of prepared statements by user (application) access. The default value is **2,000** and the minimum value is 1. The problem in which prepared statement exceeding allowed memory is mistakenly generated by system can be prohibited by making users specify the parameter value.

MAX_QUERY_TIMEOUT

A parameter used to specify timeout value of query execution. When time exceeds a value specified in this parameter after starting query execution, the query being executed stops and rolls back.

The default value is **0** (seconds) and it means infinite wait. The value range is available from 8 to 86,400 seconds (one day). The smallest value (except 0) between the **MAX_QUERY_TIMEOUT** value and query timeout value of an application is applied if query timeout is configured in an application.

Note See the **cci_connect_with_url** and **cci_set_query_timeout** functions to configure query timeout of CCI applications. For configuring query timeout of JDBC applications, see the **setQueryTimeout** method.

MAX_STRING_LENGTH

A parameter used to specify the maximum string length for bit, varbit, char, varchar, nchar, nchar varying data types. If the value is **-1**, which is the default value, the length defined in the database is used. If the value is **100**, the value acts like 100 being applied even when a certain attribute is defined as varchar(1000).

MIN_NUM_APPL_SERVER

A parameter used to specify the minimum number of application servers (**CAS**) even if any request to connect the Broker has not been made. The default value is **5**.

PREFERRED_HOSTS

A parameter mandatorily configured if Broker mode is set to PHRO. The default value is **NULL**. FOR details, see [cubrid_broker.conf](#) of "Administrator's Guide."

MAX_NUM_APPL_SERVER

A parameter used to specify the maximum number of application servers (**CAS**). The default value is **40**. In environment in which connection pool is maintained by using a middle ware such as WAS, you must specify the value of **MAX_NUM_APPL_SERVER** parameter as same as that of connection pool.

SELECT_AUTO_COMMIT

A parameter used to specify auto-commit mode for **SELECT** statements in **CCI** or **PHP**. The default value is **OFF**. Note that auto-commit is performed only at the point at which the result set for all n query statements is fetched from the server when there are n prepared statements. An example is as follows. For details, see "API Reference > CCI API > cci_end_tran."

```

SELECT 1 prepare
SELECT 1 execute // AUTO COMMIT 0

SELECT_1 prepare
SELECT_2 prepare
SELECT 1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed
SELECT 2 execute // AUTO COMMIT 0

SELECT_1 prepare
SELECT_1 execute // AUTO COMMIT 0
INSERT_1 prepare
INSERT 1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed

INSERT 1 prepare
INSERT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed
SELECT_1 prepare
SELECT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed

SELECT 1 prepare
INSERT 1 prepare
SELECT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed
INSERT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed

INSERT 1 prepare
SELECT 1 prepare
INSERT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed
SELECT_1 execute // AUTO COMMIT X -> An EXPLICIT COMMIT needed

```

SERVICE

A parameter used to specify whether to run Broker. It can be either **ON** or **OFF**. The default value is **ON**. The Broker can run only when this value is configured to **ON**.

SESSION_TIMEOUT

A parameter used to specify timeout value for the session of Broker. The default value is **300** (seconds). If there is no response to the job request for the specified time period, session will be terminated. If a value exceeds the value specified in this parameter without any action taken after starting transaction, the connections is terminated.

SLOW_LOG

A parameter used to specify whether to log. The default value is **ON**. If the value is **ON**, long transaction query which exceeds the time specified in **LONG_QUERY_TIME** or query where an error occurred is stored in the **SLOW SQL** log file. The name of file created is *broker_name_id.slow.log* and it is located under **SLOW_LOG_DIR**.

SLOW_LOG_DIR

A parameter used to specify the location of directory where the log file is generated. The default value is **log/broker/sql_log**.

SOURCE_ENV

A parameter used to specify the file to independently configure operating system environment variables for each broker. The extension of the file must be **env**. All parameters specified in **cubrid.conf** can also be configured by environment variables. For example, the **lock_timeout_in_secs** parameter in **cubrid.conf** can also be configured by the **CUBRID_LOCK_TIMEOUT_IN_SECS** environment variable. As another example, to block execution of DDL statements on broker1, you can configure **CUBRID_BLOCK_DDL_STATEMENT 1** in the file specified by **SOURCE_ENV**.

An environment variable, if exists, has priority over **cubrid.conf**. The default value is **cubrid.env**.

STATEMENT_POOLING

A parameter used to specify whether to use statement pooling. The default value is **ON**.

When transaction is committed or rolled back, CUBRID closes all the prepared statement handles that exist in the client session. However, if the parameter is set to **STATEMENT_POOLING=ON**, the prepared statement handles remain in the pool, so that the handles can be reused. Therefore, you must maintain the default setting (**STATEMENT_POOLING=ON**) in general applications that reuse prepared statements or in environments in which a library such as DBCP, in which the statement pooling is implemented, is applied.

When the parameter value is configured to **STATEMENT_POOLING=OFF** and the prepared statement is executed after the transaction is committed or terminated, the following message is displayed.

```
Caused by: cubrid.jdbc.driver.CUBRIDException: Attempt to access a closed Statement.
```

SQL_LOG

A parameter used to specify whether to leave logs for SQL statements processed by the application server (CAS) when an application server handles requests from a client. The default value is **ON**. When this parameter is configured to **ON**, all logs are stored. Log file name becomes *broker_name_id.sql.log*. The file is created in the **log/broker/sql_log** directory under the installation directory. The parameter values are as follows:

- **OFF** : Does not leave any logs
- **ERROR** : Leaves logs for queries which occur an error. only queries where an error occurs
- **NOTICE** : Leaves logs for the long-duration execution queries which exceeds the configured time/transaction, or leaves logs for queries which occur an error
- **TIMEOUT** : Leaves logs for the long-duration execution queries which exceeds the configured time/transaction
- **ON/ALL** : Leaves all logs

SQL_LOG_MAX_SIZE

A parameter used to specify the maximum size of the SQL log file. The default value is **100,000** (KB). If the size of the SQL log file, which is created when the **SQL_LOG** parameter is configured to **ON**, reaches the value configured by the parameter, *broker_name_id*. **sql.log.bak** is created.

STATEMENT_POOLING

A parameter used to specify whether to use statement pool feature. The default value is **ON**.

CUBRID closes all handles of prepared statement in the corresponding client sessions when transaction commit or rollback is made. If the value of **STATEMENT_POOLING** is set to **ON**, the handles are reusable because they are maintained in the pool. Therefore, in an environment where libraries, such as general applications reusing prepared statement or DBCP where statement pooling is implemented, are applied, the default configuration (**ON**) should be maintained.

If the prepared statement is executed after transaction commit or termination while **STATEMENT_POOLING** is set to **OFF**, the following message will be displayed.

```
Caused by: cubrid.jdbc.driver.CUBRIDException: Attempt to access a closed Statement.
```

TIME_TO_KILL

A parameter used to specify the time to remove application servers (CAS) in idle state among application servers added dynamically. The default value is **120** (seconds). An idle state is one in which the server is not involved in any jobs. If this state continues exceeding the value specified in **TIME_TO_KILL**, the application server (CAS) is removed.

The value configured in this parameter affects only application server added dynamically, so it applies only when the **AUTO_ADD_APPL_SERVER** parameter is configured to **ON**. Note that times to add or remove the application servers (CAS) will be increased more if the **TIME_TO_KILL** value is so small.

API Reference

This chapter covers the following APIs:

- JDBC API
- ODBC API
- OLE DB API
- PHP API
- CCI API

JDBC API

JDBC Programming

CUBRID JDBC Driver

The CUBRID JDBC driver (**cubrid_jdbc.jar**) enables the system to make a connection to the CUBRID database in an application written in Java. The driver is located in the "location of CUBRID installed/jdbc" directory.

The CUBRID JDBC driver has been developed based on the JDBC 2.0 specification and provides compilation output generated in JDK version 1.6.

Checking the CUBRID JDBC Driver Version

You can check the JDBC driver version as follows:

```
% jar -tf cubrid_jdbc.jar
META-INF/ META-INF/MANIFEST.MF
cubrid/ cubrid/jdbc/
cubrid/jdbc/driver/
cubrid/jdbc/jci/
cubrid/sql/
CUBRID-JDBC-8.1.4.1032
cubrid/jdbc/driver/CUBRIDBlob.class
...
```

Registering the CUBRID JDBC Driver

Use the **Class.forName** (*driver-class-name*) command to register the JDBC driver. The following example shows how to load the **cubrid.jdbc.driver.CUBRIDDriver** class to register the CUBRID JDBC driver.

```
import java.sql.*;
import cubrid.jdbc.driver.*;

public class LoadDriver {
    public static void main(String[] Args) {
        try {
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        } catch (Exception e) {
            System.err.println("Unable to load driver.");
            e.printStackTrace();
        }
    }
    ...
}
```

CUBRID JDBC Interface

The following table shows the JDBC standard and extended interfaces supported by CUBRID JDBC. Note that some methods are not supported even though they are specified in the JDBC 2.0 specification.

Supported Interface by CUBRID

JDBC Standard Interface	JDBC Extended Interface	Supported
java.sql.Blob	java.sql.CUBRIDConnection	Supported.
java.sql.CallableStatement	java.sql.CUBRIDPreparedStatement	
java.sql.Clob	java.sql.CUBRIDResultSet	
java.sql.Connection	java.sql.CUBRIDResultSetMetaData	
java.sql.DatabaseMetaData	CUBRIDOID	
java.sql.Driver		
java.sql.PreparedStatement		
java.sql.ResultSet		
java.sql.ResultSetMetaData		
java.sql.Statement	java.sql.CUBRIDStatement	The getGeneratedKeys() method of JDBC 3.0 is supported.

java.sql.DriverManager	Supported
Java.sql.SQLException	Java.sql.CUBRIDException
Supported	
java.sql.Array	Not Supported
java.sql.ParameterMetaData	
java.sql.Ref	
java.sql.Savepoint	
java.sql.SQLData	
java.sql.SQLInput	
java.sql.Struct	

Connection Configuration

The **DriverManager** is a basic interface for JDBC driver management and performs functions such as selecting a database driver and creating a new database connection. If the CUBRID JDBC driver is registered, database connection is made by calling the **DriverManager.getConnection** (*db-url*, *user-id*, and *password*) method. The **getConnection** method returns the **Connection** object, which is used for query and command executions and transaction commit or rollback. The parameter *db-url*, which is for connection configuration, is as follows:

```
jdbc:cubrid:<host>:<port>:<db-name>:[user-id]:[password]:[?<property> [& <property>]]

<host> ::=
hostname | ip address

<property> ::=
althosts=<alternative_hosts> | rctime=<second> | connectTimeout=<second> |
queryTimeout=<second> | charset=<character_set> | zeroDateTimeBehavior=<behavior_type> |
logFile=<file_name> | logOnException=<bool_type> |
logSlowQueries=<bool_type>&slowQueryThresholdMillis=<millisecond>

<alternative_hosts> ::=
<standby_broker1_host>:<port> [, <standby_broker2_host>:<port>]
<behavior_type> ::= exception | round | convertToNull
<bool_type> ::= true | false
```

- *<host>* : IP address or host name where the CUBRID Broker is running
- *<port>* : Broker port number (default : 33,000)
- *<db-name>* : The name of the database to connect
- *[user-id]* : The user that will be connected to the database. There are two users in the database by default: **dba** and **public**. If you enter an empty string (" "), you will connect to the database as a **public** user.
- *[password]* : If there is no password set for the user, enter an empty string (" ").
- **althosts** : One or more host IP of standby broker and connection port to be failed over in HA environment
- **rctime** : Interval time (in seconds) to fail over an active server during system failure. For more information, see the example in "Administrator's Guide > CUBRID HA > Environment Configuration > JDBC Configuration."
- **connectTimeout** : Configures timeout value for database connection in seconds (default value : 0). The **DriverManger.setLoginTimeout()** method can also be used to configure timeout value; however, if the value is configured in connection URL, configuration value specified as a method is ignored.
- **queryTimeout** : Configures timeout value for query execution in seconds (default value : 0, infinite). This value can be changed by the **DriverManger.setQueryTimeout** method.
- **charset** : Character set (charset) of database to be connected
- **zeroDateTimeBehavior** : JDBC does not allow a value having 0 for both date and time regardless of date and time of the **java.sql.Date** type. This property (*zeroDateTimeBehavior*) is used to determine how to handle this value in case that it should be displayed. The default operation is **exception**. The each operation by configuration value is as follows:
 - **exception** : Default operation. It is handled as a SQLException exception.
 - **round** : Returns the minimum value.
 - **convertToNull** : Returns **NULL**.

For information on the value having 0 for both date and time, see "CUBRID SQL Guide > Data Types > Date/Time Types > Definition and Characteristics."

- **logFile** : Name of a log file for debugging (default value : cubrid_jdbc.log)
- **logOnException** : Whether exception logging for debugging exists (default value : false)
- **logSlowQueries** : Whether slow query logging for debugging exists (default value : false)
- **slowQueryThresholdMillis** : Time out of a slow query for debugging (default value : 60,000 milliseconds)

Example 1

```
--connection URL string when user name and password omitted
URL=jdbc:CUBRID:192.168.0.1:33000:db1:::

--connection URL string when zeroDateTimeBehavior property specified
URL=jdbc:CUBRID:127.0.0.1:31000:db1:::?zeroDateTimeBehavior=convertToNull

--connection URL string when charset property specified
URL=jdbc:CUBRID:192.168.0.1:33000:db1::?charset=utf-8

--connection URL string when queryTimeout and charset property specified
URL=jdbc:CUBRID:127.0.0.1:31000:db1::?queryTimeout=1&charset=utf-8

--connection URL string when a property(alhosts) specified for HA
URL=jdbc:CUBRID:192.168.0.1:33000:db1::?alhosts=192.168.0.2:33000,192.168.0.3:33000

--connection URL string when properties(alhosts,rctime) specified for HA
URL=jdbc:CUBRID:192.168.0.1:33000:db1::?alhosts=192.168.0.2:33000,192.168.0.3:33000&rcti
me=600

--connection URL string when properties(alhosts,rctime, charset) specified for HA
URL=jdbc:CUBRID:192.168.0.1:33000:db1::?alhosts=192.168.0.2:33000,192.168.0.3:33000&rcti
me=600&charset=utf-8
```

Example 2

```
String url = "jdbc:cubrid:192.168.0.1:33000:demodb:::";
String userid = "";
String password = "";

try {
    Connection conn =
        DriverManager.getConnection(url,userid,password);
    // Do something with the Connection

    ...

} catch (SQLException e) {
    System.out.println("SQLException:" + e.getMessage());
    System.out.println("SQLState: " + e.getSQLState());
}
...

```

Remark

Because a colon (:) and a question mark (?) are used as a separator in URL string, it is not allowed to include them for password of URL string. To use them, you must specify a user name (*user-id*) and a password (*password*) as a separate parameter in the **getConnection** method.

Note The rollback method, which requests the transaction rollback, exits when the server completes the rollback job.

Verifying Foreign Key Information

Description

You can verify foreign key information by using **getImportedKeys**, **getExportedKeys**, and **getCrossReference** methods provided by **DatabaseMetaData** interface. Usage and examples of each method are as follows:

Syntax

```
getImportedKeys(String catalog, String schema, String table)
getExportedKeys(String catalog, String schema, String table)
getCrossReference(String parentCatalog, String parentSchema, String parentTable, String
foreignCatalog, String foreignSchema, String foreignTable)
```

- **getImportedKeys method** : A method that retrieves the information of primary key columns which are referred by foreign key columns in a given table. The results are sorted by **PKTABLE_NAME** and **KEY_SEQ**.
- **getExportedKeys method** : A method that retrieves the information of all foreign key columns which refer to primary key columns in a given table. The results are sorted by **FKTABLE_NAME** and **KEY_SEQ**.
- **getCrossReference method** : A method that retroeves the information of primary key columns which are referred by foreign key columns in a given table. The results are sorted by **PKTABLE_NAME** and **KEY_SEQ**.

Return Value

When the methods above are called, the following ResultSet, consisting of 14 columns, is returned.

Name	Type	Note
PKTABLE_CAT	String	Always null
PKTABLE_SCHEM	String	Always null
PKTABLE_NAME	String	Table name of primary key
PKCOLUMN_NAME	String	Table name of primary key
FKTABLE_CAT	String	Always null
FKTABLE_SCHEM	String	Always null
FKTABLE_NAME	String	Table name of foreign key
FKCOLUMN_NAME	String	Column name of foreign key
KEY_SEQ	short	Sequence of foreign or primary keys (starting from 1)
UPDATE_RULE	short	A corresponding value to referring action defined as to foreign keys when primary keys are updated Cascade=0, Restrict=2, No action=3, Set null=4
DELETE_RULE	short	A corresponding value to referring action defined as to foreign keys when primary keys are deleted Cascade=0, Restrict=2, No action=3, Set null=4
FK_NAME	String	Foreign key name
PK_NAME	String	Primary key name
DEFERRABILITY	short	Always 6(DatabaseMetaData.importedKeyInitiallyImmediate)

Example

```
ResultSet rs = null;

    DatabaseMetaData dbmd = conn.getMetaData();

    System.out.println("\n==== Test getImportedKeys");
    System.out.println("====");
    rs = dbmd.getImportedKeys(null, null, "pk table");
    Test.printFkInfo(rs);
    rs.close();

    System.out.println("\n==== Test getExportedKeys");
    System.out.println("====");
    rs = dbmd.getExportedKeys(null, null, "fk table");
    Test.printFkInfo(rs);
    rs.close();

    System.out.println("\n==== Test getCrossReference");
    System.out.println("====");
```

```

        rs = dbmd.getCrossReference(null, null, "pk table", null, null,
"fk table");
        Test.printFkInfo(rs);
        rs.close();

```

Using OIDs and Collections

In addition to the methods defined in the JDBC specification, the CUBRID JDBC driver provides methods that handle OIDs and collections (set, multiset and sequence).

To use these methods, you must import **cubrid.sql.***; in addition to the CUBRID JDBC driver classes which are imported by default. In addition, to get the results, you must convert **ResultSet** to **CUBRIDResultSet** first. (**ResultSet** is provided by the standard JDBC API, by default.)

```

import cubrid.jdbc.driver.* ;
import cubrid.sql.* ;
...
CUBRIDResultSet urs = (CUBRIDResultSet) stmt.executeQuery(
"SELECT city FROM location");

```

Caution AUTO COMMIT does not work even though it is configured to **TRUE** if CUBRID extended APIs are used. Therefore, you must manually commit open connections. The CUBRID extended APIs are methods that handle OIDs and collections.

Using OIDs

You must follow the following rules to use OIDs.

- To use **CUBRIDOID**, you should import **cubrid.sql.***; (a)
- You can retrieve an OID by specifying a class name in the **SELECT** statement. The name can be used together with other attributes. (b)
- The **ResultSet** of a query must be **CUBRIDResultSet**. (c)
- The method that retrieves the OID from the **CUBRIDResultSet** is **getOID()**. (d)
- To retrieve a value from an OID, use the **getValues()** method. Its result is **ResultSet**. (e)
- To substitute a value for an OID, use the **setValues()** method. (f)
- When you use the extended APIs, you must always perform **commit()** to make connection. (g)

```

import java.sql.*;
import cubrid.sql.*; //a
import cubrid.jdbc.driver.*;

/*
CREATE TABLE oid test(
    id INTEGER,
    name VARCHAR(10),
    age INTEGER
);

INSERT INTO oid test VALUES(1, 'Laura', 32);
INSERT INTO oid_test VALUES(2, 'Daniel', 39);
INSERT INTO oid_test VALUES(3, 'Stephen', 38);
*/

class OID Sample
{
    public static void main (String args [])
    {
        // Making a connection
        String url= "jdbc:cubrid:localhost:33000:demodb::";
        String user = "dba";
        String passwd = "";

        // SQL statement to get OID values
        String sql = "SELECT oid test from oid test"; //b
        // columns of the table
        String[] attr = { "id", "name", "age" } ;

```

```

// Declaring variables for Connection and Statement
Connection con = null;
Statement stmt = null;
CUBRIDResultSet rs = null;
ResultSetMetaData rsmd = null;

try {
    Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
} catch (ClassNotFoundException e) {
    throw new IllegalStateException("Unable to load Cubrid driver", e);
}

try {
    con = DriverManager.getConnection(url, user, passwd);
    stmt = con.createStatement();
    rs = (CUBRIDResultSet)stmt.executeQuery(sql); //c
    rsmd = rs.getMetaData();

    // Printing columns
    int numOfColumn = rsmd.getColumnCount();
    for (int i = 1; i <= numOfColumn; i++ ) {
        String columnName = rsmd.getColumnName(i);
        String jdbcType = rsmd.getColumnTypeName(i);
        System.out.print(columnName );
        System.out.print("(" + jdbcType + ")");
        System.out.print(" | ");
    }
    System.out.print("\n");
    // Printing rows
    CUBRIDResultSet rsoid = null;
    int k = 1;

    while (rs.next()) {
        CUBRIDOID oid = rs.getOID(1); //d
        System.out.print("OID");
        System.out.print(" | ");
        rsoid = (CUBRIDResultSet)oid.getValues(attr); //e

        while (rsoid.next()) {
            for( int j=1; j <= attr.length; j++ ) {
                System.out.print(rsoid.getObject(j));
                System.out.print(" | ");
            }
        }
        System.out.print("\n");

        // New values of the first row
        Object[] value = { 4, "Yu-ri", 19 };
        if (k == 1) oid.setValues(attr, value); //f

        k = 0;
    }
    con.commit(); //g

} catch(CUBRIDException e) {
    e.printStackTrace();

} catch(SQLException ex) {
    ex.printStackTrace();

} finally {
    if(rs != null) try { rs.close(); } catch(SQLException e) {}
    if(stmt != null) try { stmt.close(); } catch(SQLException e) {}
    if(con != null) try { con.close(); } catch(SQLException e) {}
}
}
}

```

Using Collections

The line marked by 'a' in the example 1 below is where data of a collection type (SET, MULTISSET, LIST) is fetched from the **CUBRIDResultSet**. The results are returned as array format. Note that this function is supported only when data types of elements defined in the collection type are same.

Example 1

```
import java.sql.*;
import java.lang.*;
import cubrid.sql.*;
import cubrid.jdbc.driver.*;

// create class collection test(
// setttest set(integer),
// multisettest multiset(integer),
// listttest list(Integer)
// );
//

// insert into collection test values({1,2,3},{1,2,3},{1,2,3});
// insert into collection_test values({2,3,4},{2,3,4},{2,3,4});
// insert into collection_test values({3,4,5},{3,4,5},{3,4,5});

class Collection Sample
{
    public static void main (String args [])
    {
        String url= "jdbc:cubrid:127.0.0.1:33000:demodb::";
        String user = "";
        String passwd = "";
        String sql = "select setttest,multisettest,listttest from collection test";
        try {
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        } catch(Exception e){
            e.printStackTrace();
        }
        try {
            Connection con = DriverManager.getConnection(url,user,passwd);
            Statement stmt = con.createStatement();
            CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);
            CUBRIDResultSetMetaData rsmd = (CUBRIDResultSetMetaData) rs.getMeta Data();
            int numbOfColumn = rsmd.getColumnCount();
            while (rs.next ()) {
                for (int j=1; j<=numbOfColumn; j++ ) {
                    Object[] reset = (Object[]) rs.getCollection(j); //a
                    for (int m=0 ; m < reset.length ; m++)
                        System.out.print(reset[m] +",");
                    System.out.print(" | ");
                }
                System.out.print("\n");
            }
            rs.close();
            stmt.close();
            con.close();
        } catch(SQLException e) {
            e.printStackTrace();
        }
    }
}
```

Example 2

```
import java.sql.*;
import java.io.*;
import java.lang.*;
import cubrid.sql.*;
import cubrid.jdbc.driver.*;

// create class collection test(
// setttest set(integer),
// multisettest multiset(integer),
// listttest list(Integer)
```

```

// );
//
// insert into collection_test values({1,2,3},{1,2,3},{1,2,3});
// insert into collection_test values({2,3,4},{2,3,4},{2,3,4});
// insert into collection_test values({3,4,5},{3,4,5},{3,4,5});

class SetOP Sample
{
    public static void main (String args [])
    {
        String url = "jdbc:cubrid:127.0.0.1:33000:demodb:::" ;
        String user = "";
        String passwd = "";
        String sql = "select collection test from collection test";
        try {
            Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
        } catch(Exception e){
            e.printStackTrace();
        }
        try {
            CUBRIDConnection con =(CUBRIDConnection)
            DriverManager.getConnection(url,user,passwd);
            Statement stmt = con.createStatement();
            CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);
            while (rs.next ()) {
                CUBRIDOID oid = rs.getOID(1);
                oid.addToSet ("settest",new Integer(10));
                oid.addToSet ("multisettest",new Integer(20));
                oid.addToSequence ("listtest",1,new Integer(30));
                oid.addToSequence ("listtest",100,new Integer(100));
                oid.putIntoSequence ("listtest",99,new Integer(99));
                oid.removeFromSet ("settest",new Integer(1));
                oid.removeFromSet ("multisettest",new Integer(2));
                oid.removeFromSequence ("listtest",99);
                oid.removeFromSequence ("listtest",1);
            }
            con.commit();
            rs.close();
            stmt.close();
            con.close();
        } catch(SQLException e) {
            e.printStackTrace();
        }
    }
}

```

Getting Auto-Increment Column Values

Auto-increment Feature

The auto-increment feature (**AUTO_INCREMENT**) is a column-related feature that increments the numeric value of each row. For more information, see "CUBRID SQL Guide > Table Definition > CREATE TABLE > Column Definition" in Creating Tables. This feature can be defined only for numeric domains (**SMALLINT**, **INTEGER**, **DECIMAL**(*p*, 0), **NUMERIC**(*p*, 0)).

The auto-increment feature is recognized as an automatically created key in a JDBC program. To retrieve the key, you need to specify the time to insert a row from which the automatically created key value is to be retrieved. To perform it, you must set the flag by calling **Connection.prepareStatement** and **Statement.execute**. In this case, the command to be executed should be the **INSERT** statement or **INSERT** within **SELECT** statement. For other commands, the JDBC driver ignores the flag-setting parameter.

Steps

- Use one of the followings to indicate whether or not to return a key created automatically. The following method forms are used for tables of the database server that supports the auto-increment columns. Each method form can be applied only to a single-row **INSERT** statement.
- Create a **PreparedStatement** object by referring to the followings:
Connection.prepareStatement(*sql statement*, **Statement.RETURN_GENERATED_KEYS**);

- To insert a row using the **Statement.execute** method, use one of the forms of the **Statement.execute** method by referring to the followings:
Statement.execute(sql statement, Statement.RETURN_GENERATED_KEYS);
- Retrieve a **ResultSet** object that contains an automatically created key value by calling the **PreparedStatement.getGeneratedKeys** or **Statement.getGeneratedKeys** method. Note that the data type of the automatically created key in **ResultSet** is **DECIMAL** regardless of the data type of the given domain.

Example

The following example shows how to create a table with the auto-increment feature, enter data into the table so that automatically created key values are entered into auto-increment columns, and check whether the key values are successfully retrieved by using the **Statement.getGeneratedKeys()** method. Each step is explained in the comments for commands that correspond to the steps above.

```
import java.sql.*;
import java.math.*;
import cubrid.jdbc.driver.*;

Connection con;
Statement stmt;
ResultSet rs;
java.math.BigDecimal idColVar;
...
stmt = con.createStatement(); // Create a Statement object

stmt.executeUpdate(
"CREATE TABLE EMP_PHONE (EMPNO CHAR(6), PHONENO CHAR(4), "
+ "IDENTCOL INTEGER AUTO_INCREMENT)"); // Create table with identity column

stmt.execute(
"INSERT INTO EMP_PHONE (EMPNO, PHONENO) "
+ "VALUES ('000010', '5555')", // Insert a row <Step 1>
Statement.RETURN_GENERATED_KEYS); // Indicate you want automatically

rs = stmt.getGeneratedKeys(); // generated keys
// Retrieve the automatically <Step 2>
// generated key value in a ResultSet.
// Only one row is returned.
// Create ResultSet for query

while (rs.next()) {
    java.math.BigDecimal idColVar = rs.getBigDecimal(1);
    // Get automatically generated key
    // value
    System.out.println("automatically generated key value = " + idColVar);
}
rs.close(); // Close ResultSet
stmt.close(); // Close Statement
```

Using BLOB/CLOB

The interfaces that handle **LOB** data in JDBC is implemented based on JDBC 4.0 specification. The constraints of interfaces are as follows:

- It supports sequential writes only when creating the objects of **BLOB** or **CLOB**. Writing to arbitrary locations is not supported.
- You cannot change the data of **BLOB** or **CLOB** by calling methods of **BLOB** or **CLOB** object which are received from **ResultSet**.
- It does not support **Blob.truncate**, **Clob.truncate**, **Blob.position**, and **Clob.position**.
- You cannot bind **LOB** data by calling **PreparedStatement.setAsciiStream**, **PreparedStatement.setBinaryStream**, and **PreparedStatement.setCharacterStream** methods of **BLOB/CLOB** type columns.
- To use **BLOB/CLOB** types in an environment where JDBC 4.0 specification is not supported such as JDB version 1.5 or earlier, you must convert a conn object to **CUBRIDConnection**, explicitly. See the example below.

```
// JDK 1.6 or later
import java.sql.*;
Connection conn = DriverManager.getConnection(url, id, passwd);
```



```

Blob blob = conn.createBlob();
...
// JDK 1.5 or earlier
import java.sql.*;
import cubrid.jdbc.driver.*;

Connection conn = DriverManager.getConnection(url, id, passwd);
Blob blob = ((CUBRIDConnection)conn).createBlob();
...

```

Saving LOB Data

The way to bind **LOB** type data is as follows:

- Create **java.sql.Blob** or **java.sql.Clob** object and store the file contents in the object. Use, then, **setBlob()** or **setClob()** of **PreparedStatement** (example 1).
- Perform query and get **java.sql.Blob** or **java.sql.Clob** object from the **ResultSet** object. Bind, then, the object in **PreparedStatement** (example 2).

Example 1

```

Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
Connection conn = DriverManager.getConnection("jdbc:cubrid:localhost:33000:image_db::",
"", "");
PreparedStatement pstmt1 = conn.prepareStatement("INSERT INTO doc(image id, doc id, image)
VALUES (?, ?, ?)");
pstmt1.setString(1, "image-21");
pstmt1.setString(2, "doc-21");

//Creating an empty file in the file system
Blob bImage = conn.createBlob();
byte[] bArray = new byte[256];
...

//Inserting data into the external file. Position is start with 1.
bImage.setBytes(1, bArray);
//Appending data into the external file
bImage.setBytes(257, bArray);
...
pstmt1.setBlob(3, bImage);
pstmt1.executeUpdate();
...

```

Example 2

```

Class.forName("cubrid.jdbc.driver.CUBRIDDriver");
Connection conn = DriverManager.getConnection("jdbc:cubrid:localhost:33000:image_db::",
"", "");
conn.setAutoCommit(false);
PreparedStatement pstmt1 = conn.prepareStatement("SELECT image FROM doc WHERE image id = ?
");
pstmt1.setString(1, "image-21");
ResultSet rs = pstmt1.executeQuery();

while (rs.next())
{
Blob bImage = rs.getBlob(1);
PreparedStatement pstmt2 = conn.prepareStatement("INSERT INTO doc(image_id, doc_id, image)
VALUES (?, ?, ?)");
pstmt2.setString(1, "image-22")
pstmt2.setString(2, "doc-22")
pstmt2.setBlob(3, bImage);
pstmt2.executeUpdate();
pstmt2.close();
}
pstmt1.close();
conn.commit();
conn.setAutoCommit(true);
conn.close();
...

```

Getting LOB Data

The way to get **LOB** type data is as follows:

- Get data directly from **ResultSet** by using **getBytes()** or **getString()** method (example 1).
- Get the **java.sql.Clob** object from **ResultSet** by calling **getBlob()** or **getClob()** method and then get data by using **getBytes()** or **getSubString()** method for this object (example 2).

Example 1

```
Connection conn = DriverManager.getConnection ("jdbc:cubrid:localhost:33000:image db::",
"", "");

// Getting data directly from ResultSet
PreparedStatement pstmt1 = conn.prepareStatement("SELECT content FROM doc t WHERE doc id = ?");
pstmt1.setString(1, "doc-10");
ResultSet rs = pstmt1.executeQuery();
while (rs.next())
{
    String sContent = rs.getString(1);
    System.out.println("doc.content= "+sContent.);
}
}
```

Example 2

```
Connection conn = DriverManager.getConnection ("jdbc:cubrid:localhost:33000:image db::",
"", "");

//Getting Blob data from ResultSet and getting data from the Blob object
PreparedStatement pstmt2 = conn.prepareStatement("SELECT image FROM image_t WHERE image_id = ?");
pstmt2.setString(1,"image-20");
ResultSet rs = pstmt2.executeQuery();
while (rs.next())
{
    Blob bImage = rs.getBlob(1);
    Bytes[] bArray = bImage.getBytes(1, (int)bImage.length());
}
}
```

Note If a string longer than defined size in a column is inserted(INSERT) or updated(UPDATE), the string will be truncated.

CUBRIDOID

Overview

A **CUBRIDOID** class contains the following methods to process OIDs.

Return Type	Method Name
void	addToSequence(String attrName, int index, Object value)
void	addToSet(String attrName, Object value)
static CUBRIDOID	getNewInstance(CUBRIDConnection con, String oidStr)
String	getOidString()
String	getTableName()
ResultSet	getValues(String[] attrNames)
Boolean	isInstance()
void	putIntoSequence(String attrName, int index, Object value)
void	remove()
void	removeFromSequence(String attrName, int index)

void	removeFromSet(String attrName, Object value)
void	setReadLock()
void	setValues(String[] attrNames, Object[] values)
void	setWriteLock()

addToSequence

Description

Inserts the value specified in *value* into the attribute named *attrName* and associated with **SEQUENCE** constraints on the **CUBRIDOID** instance, specifically in front of the *index*-th element in the **SEQUENCE** attribute.

Syntax

```
void addToSequence(String attrName, int index, Object value)
```

Example

```
//create class foo(c list of int )
//insert into foo values({3})

String sql = "select foo from foo" ;

Connection con = DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);          // get OID
    oid.addToSequence("c",1, new Integer(22)); // c: {3}-> {22,3}
}
```

addToSet

Description

Inserts the value specified in *value* into the attribute named *attrName* and associated with **SET** or **MULTISET** constraints on the **CUBRIDOID** instance.

Syntax

```
void addToSet(String attrName, Object value)
```

Example

```
//create class foo(a set of int, b multiset of int )
//insert into foo values({1},{2})
String sql = "select foo from foo" ;

Connection con = DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);
while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);          // get OID
    oid.addToSet("a",new Integer(11));    // a : {1} -> {1,11}
    oid.addToSet("b",new Integer(13));    // b : {2} -> {2, 13}
}
```

getNewInstance

Description

Converts an OID string to a **CUBRIDOID** object, and then returns the **CUBRIDOID** object.

Syntax

```
static CUBRIDOID getNewInstance(CUBRIDConnection con, String oidStr)
```

Return Value

- **CUBRIDOID** object

Example

```
String sql = "select foo from foo" ;

CUBRIDConnection con = (CUBRIDConnection)
    DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();

CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID realoid = rs.getOID(1); // get OID (CUBRIDOID)
    // CUBRIDOID -> OID string
    String stringoid = realoid.getOidString();
    // OID string -> CUBRIDOID
    realoid = CUBRIDOID.getNewInstance(con, stringoid);
}
```

getOidString

Description

Converts a **CUBRIDOID** object to an OID string, and then returns the string.

Syntax

```
String getOidString()
```

Return Value

- Character string

Example

```
String sql = "select foo from foo" ;

CUBRIDConnection con = (CUBRIDConnection)
    DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();
CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID realoid = rs.getOID(1); // get OID
    // CUBRIDOID -> OID string
    String stringoid = realoid.getOidString();
    // OID string -> CUBRIDOID
    realoid = CUBRIDOID.getNewInstance(con, stringoid);
}
```

getTableName

Description

Returns the table name of the instance corresponding to the **CUBRIDOID** object.

Syntax

```
String getTableName()
```

Return Value

- A table name of an instance that corresponds to **CUBRIDOID**

Example

```
String sql = "select foo from foo" ;
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);
while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);

    String tablename = oid.getTableName();
    System.out.println(tablename );
}
```

getValues

Description

Returns the **ResultSet** which contains values of the requested attribute.

Syntax

```
ResultSet getValues(String[] attrNames)
```

Return Value

- **ResultSet**

Example

```
// create class foo ( a string, b int )
// insert into foo values('CUBRID', 2001)

String sql = "select foo from foo";
String[] attr = { "a", "b" }; // class's column name list
CUBRIDResultSet rs= (CUBRIDResultSet) stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    ResultSet rsoid = oid.getValues(attr);
}
```

isInstance

Description

Returns true if the instance corresponding to the **CUBRIDOID** exists. If otherwise, it returns false.

Syntax

```
Boolean isInstance()
```

Return Value

- TRUE : An instance that corresponds to **CUBRIDOID** exists.
- FALSE : An instance that corresponds to **CUBRIDOID** does not exist.

Example

```
String sql = "select foo from foo" ;

CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    System.out.print("isInstance : " + oid.isInstance()); // true
    oid.remove(); // remove the object in the oid
    System.out.print("After remove, isInstance : "
        + oid.isInstance()); // false
}
```

putIntoSequence**Description**

Modifies the *index*-th value in the attribute associated with the **SEQUENCE** constraint on the **CUBRIDOID** instance as the value specified in *value*.

Syntax

```
void putIntoSequence(String attrName, int index, Object value)
```

Example

```
//create class foo(c list of int )
//insert into foo values({1})

String sql = "select foo from foo" ;

Connection con = DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1); // get OID
    oid.putIntoSequence("c",1, new Integer(10)); // c:{1}->{10}
}
```

remove**Description**

Removes the instance corresponding to the **CUBRIDOID**.

Syntax

```
void remove()
```

Example

```
String sql = "select foo from foo" ;

CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    System.out.print("isInstance : " + oid.isInstance()); // true
    oid.remove(); // remove the object in the oid
    System.out.print(" After remove .isInstance : " +
```

```
oid.isInstance()); // false
}
```

removeFromSequence

Description

Removes the *index*-th value from the attribute associated with the **SEQUENCE** constraint on the **CUBRIDOID** instance.

Syntax

```
void removeFromSequence(String attrName, int index)
```

Example

```
//create class foo(c list of int )
//insert into foo values(1,3)

String sql = "select foo from foo" ;

Connection con = DriverManager.getConnection(url,user,passwd);

Statement stmt = con.createStatement();
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);      // get OID
    oid.removeFromSequence("c",1);     // c: {1,3} -> {3}
}
```

removeFromSet

Description

Removes the corresponding value specified in *value* from the attribute associated with the **SET** constraint on the **CUBRIDOID** instance. If the corresponding value is more than one, the very value found for the first time becomes removed.

Syntax

```
void removeFromSet(String attrName, Object value)
```

Example

```
//create class foo(a set of int, b multiset of int )
//insert into foo values({1,11},{2,13})

String sql = "select foo from foo"
Connection con = DriverManager.getConnection(url,user,passwd)

Statement stmt = con.createStatement()
CUBRIDResultSet rs= (CUBRIDResultSet) stmt.executeQuery(sql)

while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1)      // get OID
    oid.removeFromSet("a",new Integer(11)) // a: {1,11} -> {1}
    oid.removeFromSet("a",new Integer(13)) // b: {2,13} -> {2}
}
```

setReadLock

Description

Sets read lock on the instance corresponding to the **CUBRIDOID**.

Syntax

```
void setReadLock()
```

Example

```
String sql = "select foo from foo" ;
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);
while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    oid.setReadLock();
}
```

setValues**Description**

Replaces the value specified in the *attrNames* with the value specified in the *values*.

Syntax

```
void setValues(String[] attrNames, Object[] values)
```

Example

```
// create class foo ( a string, b int )
String sql = "select foo from foo";
String[] attr = { "a", "b" }; // a list of attribute names
String[] values = {"CUBRID", new Integer(2001)};
CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);
while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    oid.setValues(attr, values );
}
```

setWriteLock**Description**

Sets write lock on the instance corresponding to the **CUBRIDOID**.

Syntax

```
void setWriteLock()
```

Example

```
String sql = "select foo from foo" ;
CUBRIDResultSet rs = (CUBRIDResultSet) stmt.executeQuery(sql);
while (rs.next ()) {
    CUBRIDOID oid = rs.getOID(1);
    oid.setWriteLock();
}
```

CUBRIDPreparedStatement**Overview**

The **CUBRIDPreparedStatement** class extends the standard **PreparedStatement** and contains the following additional methods.

Return Type	Method Name
CUBRIDOID	executeInsert()
void	setCollection(int index, Object[] array)
void	setOID(int index, CUBRIDOID oid)

executeInsert

Description

Executes an **INSERT** statement within the **CUBRIDPreparedStatement** object and returns the **CUBRIDOID** corresponding to the inserted object.

Syntax

```
CUBRIDOID executeInsert()
```

Return Value

- A **CUBRIDOID** that corresponds to the inserted object

Example

```
String sql = "insert into testtable(a) values(?)";

CUBRIDPreparedStatement pstmt = (CUBRIDPreparedStatement)
    con.prepareStatement(sql);
pstmt.setString(1, "CUBRID");
CUBRIDOID oid = pstmt.executeInsert();
```

setCollection

Description

Configures the *index*-th parameter in the prepared statement as a collection corresponding to *array*. CUBRID has three types of collections: Set, Multiset and Sequence.

Syntax

```
void setCollection(int index, Object[] array)
```

Example

```
String[] strs = { "abc", "def"};

psmt.setCollection(1, strs);
```

setOID

Description

Configures the *index*-th parameter in the prepared statement as the **CUBRIDOID** specified in *oid*.

Syntax

```
void setOID(int index, CUBRIDOID oid)
```

CUBRIDResultSet

Overview

The **CUBRIDResultSet** class is extended from the standard **ResultSet** class and has the following additional methods.

Return Type	Method Name
Object	getCollection(int attrIndex)
Object	getCollection(String attrName)
CUBRIDOID	getOid()
CUBRIDOID	getOid(int attrIndex)
CUBRIDOID	getOid(String attrName)

getCollection

Description

Returns the index specified in *attrIndex* or the attribute value specified in *attrName*. The returned object can be converted to an array such as `String[]`.

Syntax

```
Object getCollection(int attrIndex)
Object getCollection(String attrName)
```

Return Value

- An index specified by *attrIndex* or a value of the column that corresponds to the column name specified by *attrName*

getOID

Description

Returns the index specified in *attrIndex* or the attribute value specified in *attrName* to **CUBRIDOID**, thus it returns the **CUBRIDOID**.

If *attrIndex* or *attrName* is not specified, **CUBRIDOID** of the current row of **ResultSet** is returned. This is valid only when **ResultSet** is **TYPE_SCROLL_SENSITIVE** or **CONCUR_UPDATABLE**.

Syntax

```
CUBRIDOID getOID(int attrIndex)
CUBRIDOID getOID(String attrName)
CUBRIDOID getOID()
```

Return Value

- CUBRIDOID**

CUBRIDResultSetMetaData

Overview

The **CUBRIDResultSetMetaData** class is extended from the standard **ResultSetMetaData** and has the following additional methods.

Return Type	Method Name
int	getElementType(int columnIndex)
String	getElementTypeName(int columnIndex)

getElementType

Description

Returns a type of the COLLECTION element as *int* defined in the `java.sql.Types`. If a domain of the *columnIndex*-th attribute is not COLLECTION such as **SET**, **MULTISET**, or **SEQUENCE**, **SQLException** occurs in the end.

Syntax

```
int getElementType(int columnIndex)
```

Return Value

- Collection element type (int)

getElementTypeName

Description

Returns the name of the type in the COLLECTION elements. If a domain of the *columnIndex*-th attribute is not COLLECTION such as **SET**, **MULTISET**, or **SEQUENCE**, **SQLException** occurs in the end.

Syntax

```
String getElementTypeName(int columnIndex)
```

Return Value

- Collection element's type name

Example

```
// The following schema is used in this example.
//
// create class foo(
//   a set(int),
//   b multiset(int),
//   c sequence(int)
// );

String sql = "select * from foo" ;
Connection con = DriverManager.getConnection(url,user,passwd) ;

Statement stmt = con.createStatement();

CUBRIDResultSet rs = (CUBRIDResultSet)stmt.executeQuery(sql);
CUBRIDResultSetMetaData rsmd = (CUBRIDResultSetMetaData)
    rs.getMetaData();

int numberOfColumn = rsmd.getColumnCount();
for (int i=1; i <= numberOfColumn; i++ ) {
    System.out.println(rsmd.getElementType(i) );
    System.out.println(rsmd.getElementTypeName(i) );
}
```

CUBRIDStatement

Overview

The **CUBRIDStatement** class is extended from the standard **Statement** class and has the following additional methods.

Return Type	Method Name
CUBRIDOID	executeInsert(String insertStmt)

executeInsert

Description

Returns the **CUBRIDOID** corresponding to a new tuple (row) inserted by the SQL statement, *insertStmt*.

Syntax

```
CUBRIDOID executeInsert(String insertStmt)
```

Return Value

- **CUBRIDOID** of the added row

Example

```
String sql = "insert into testable(a) values (1) "  
CUBRIDStatement stmt = (CUBRIDStatement) con.createStatement();  
CUBRIDOID oid = stmt.executeInsert(sql);
```

ODBC API

ODBC Programming

CUBRID ODBC Driver

Description

The CUBRID ODBC driver supports ODBC version 3.52, ODBC core, and some of Level 1 and Level 2 APIs. Since it has been developed based on ODBC Spec 3.x, backward compatibility is not completely ensured for programs written using ODBC Spec 2.x. Only 32 bit are supported. In the 64-bit Windows environment, you can check the ODBC driver with CUBRID 32 bit by executing "C:\Windows\SysWOW64\odbcad32.exe."

For more information on configuring CUBRID ODBC driver, see "Getting Started with CUBRID > Programming with ODBC and ASP > Configuring the Environment of ODBC and ASP."

Data Type Mapping of CUBRID and ODBC

The following table shows the data mapping relationship between data types of ODBC and those supported by CUBRID.

CUBRID Data Type	ODBC Data Type
Char	SQL_CHAR
Varchar	SQL_VARCHAR
String	SQL_LONGVARCHAR
Nchar	SQL_CHAR
Varnchar	SQL_VARCHAR
Bit	SQL_BINARY
varying bit	SQL_VARBINARY
Numeric	SQL_NUMERIC
Int	SQL_INTEGER
Short	SQL_SMALLINT
Float	SQL_FLOAT
Double	SQL_DOUBLE
Bigint	SQL_BIGINT
Date	SQL_TYPE_DATE
Time	SQL_TYPE_TIME
Timestamp	SQL_TYPE_TIMESTAMP
Datetime	SQL_TYPE_TIMESTAMP
Monetary	SQL_DOUBLE
Oid	SQL_CHAR(32)
set, multiset, sequence	SQL_VARCHAR(MAX_STRING_LENGTH)

Configuring Connection Strings

When you are programming CUBRID ODBC, you can write connection strings as follows:

Item	Example	Description
Driver	CUBRID Driver	Driver name
UID	PUBLIC	User ID
PWD	xxx	Password
FETCH_SIZE	100	Fetch size
PORT	33000	Broker port number
SERVER	127.0.0.1	IP address or host name of a CUBRID Broker server
DB_NAME	demodb	Database name
DESCRIPTION	cubrid_test	Description
CHARSET	utf-8	Character set

The following example shows how to use connecting strings above.

```
"DRIVER=CUBRID
Driver;UID=PUBLIC;PWD=xxx;FETCH_SIZE=100;PORT=33000;SERVER=127.0.0.1;DB_NAME=demodb;DESCRIPTION=cubrid_test;CHARSET=utf-8"
```

Remark

Because a semi-colon is used as a separator in URL string, it is not allowed to include a semi-colon for password (PWD) when specifying it in connection string.

Supported Functions and Backward Compatibility

Information on supported functions by CUBRID ODBC, versions, compatibility with ODBC Spec is as follows:

API	Version Introduced	Standards Compliance	Support
SQLAllocHandle	3.0	ISO 92	YES
SQLBindCol	1.0	ISO 92	YES
SQLBindParameter	2.0	ODBC	YES
SQLBrowseConnect	1.0	ODBC	NO
SQLBulkOperations	3.0	ODBC	YES
SQLCancel	1.0	ISO 92	YES
SQLCloseCursor	3.0	ISO 92	YES
SQLColAttribute	3.0	ISO 92	YES
SQLColumnPrivileges	1.0	ODBC	NO
SQLColumns	1.0	X/Open	YES
SQLConnect	1.0	ISO 92	YES
SQLCopyDesc	3.0	ISO 92	YES
SQLDescribeCol	1.0	ISO 92	YES
SQLDescribeParam	1.0	ISO 92	NO
SQLDisconnect	1.0	ISO 92	YES
SQLDriverConnect	1.0	ODBC	YES
SQLEndTran	3.0	ISO 92	YES

SQLExecDirect	1.0	ISO 92	YES
SQLExecute	1.0	ISO 92	YES
SQLFetch	1.0	ISO 92	YES
SQLFetchScroll	3.0	ISO 92	YES
SQLForeignKeys	1.0	ODBC	YES (2008 R3.1 or later)
SQLFreeHandle	3.0	ISO 92	YES
SQLFreeStmt	1.0	ISO 92	YES
SQLGetConnectAttr	3.0	ISO 92	YES
SQLGetCursorName	1.0	ISO 92	YES
SQLGetData	1.0	ISO 92	YES
SQLGetDescField	3.0	ISO 92	YES
SQLGetDescRec	3.0	ISO 92	YES
SQLGetDiagField	3.0	ISO 92	YES
SQLGetDiagRec	3.0	ISO 92	YES
SQLGetEnvAttr	3.0	ISO 92	YES
SQLGetFunctions	1.0	ISO 92	YES
SQLGetInfo	1.0	ISO 92	YES
SQLGetStmtAttr	3.0	ISO 92	YES
SQLGetTypeInfo	1.0	ISO 92	YES
SQLMoreResults	1.0	ODBC	YES
SQLNativeSql	1.0	ODBC	YES
SQLNumParams	1.0	ISO 92	YES
SQLNumResultCols	1.0	ISO 92	YES
SQLParamData	1.0	ISO 92	YES
SQLPrepare	1.0	ISO 92	YES
SQLPrimaryKeys	1.0	ODBC	YES (2008 R3.1 or later)
SQLProcedureColumns	1.0	ODBC	YES (2008 R3.1 or later)
SQLProcedures	1.0	ODBC	YES (2008 R3.1 or later)
SQLPutData	1.0	ISO 92	YES
SQLRowCount	1.0	ISO 92	YES
SQLSetConnectAttr	3.0	ISO 92	YES
SQLSetCursorName	1.0	ISO 92	YES
SQLSetDescField	3.0	ISO 92	YES
SQLSetDescRec	3.0	ISO 92	YES
SQLSetEnvAttr	3.0	ISO 92	NO
SQLSetPos	1.0	ODBC	YES
SQLSetStmtAttr	3.0	ISO 92	YES
SQLSpecialColumns	1.0	X/Open	YES
SQLStatistics	1.0	ISO 92	YES
SQLTablePrivileges	1.0	ODBC	YES (2008 R3.1 or later)

SQLTables	1.0	X/Open	YES
-----------	-----	--------	-----

Some functions for which backward compatibility is not supported must be converted into appropriate ones by using the mapping table below.

ODBC 2.x Function	ODBC 3.x Function
SQLAllocConnect	SQLAllocHandle
SQLAllocEnv	SQLAllocHandle
SQLAllocStmt	SQLAllocHandle
SQLBindParam	SQLBindParameter
SQLColAttributes	SQLColAttribute
SQLError	SQLGetDiagRec
SQLFreeConnect	SQLFreeHandle
SQLFreeEnv	SQLFreeHandle
SQLFreeStmt with SQL_DROP	SQLFreeHandle
SQLGetConnectOption	SQLGetConnectAttr
SQLGetStmtOption	SQLGetStmtAttr
SQLParamOptions	SQLSetStmtAttr
SQLSetConnectOption	SQLSetConnectAttr
SQLSetParam	SQLBindParameter
SQLSetScrollOption	SQLSetStmtAttr
SQLSetStmtOption	SQLSetStmtAttr
SQLTransact	SQLEndTran

Using OIDs and Collections

ODBC is designed for relational DBMSs. Therefore, CUBRID ODBC does not support some object-oriented features such as CUBRID OIDs and collections. It is because CUBRID is an object-relational DBMS that integrates relational and object-oriented data models.

Using OIDs

Because the CUBRID ODBC driver considers an OID as a string (char(32)), the **INSERT**, **UPDATE** and **DELETE** statements containing OIDs can be used as follows. The OID string should be used with single quotes ('). The domain of the member attribute in the following example is the same as the OID.

```
insert into foo(member) values('@12|34|56')
delete from foo where member = '@12|34|56'
update foo set age = age + 1 where member = '@12|34|56'
```

Using Collections

Collection types : **SET**, **MULTISET** and **SEQUENCE** are supported. The CUBRID ODBC driver considers a collection as a string (longvarchar). You can obtain a collection by separating each element in the **SELECT** statement using commas in braces as with "{value_1, value_2, ...value_n}."

Note If a string longer than defined size in a column is inserted (**INSERT**) or updated (**UPDATE**), the string will be truncated.

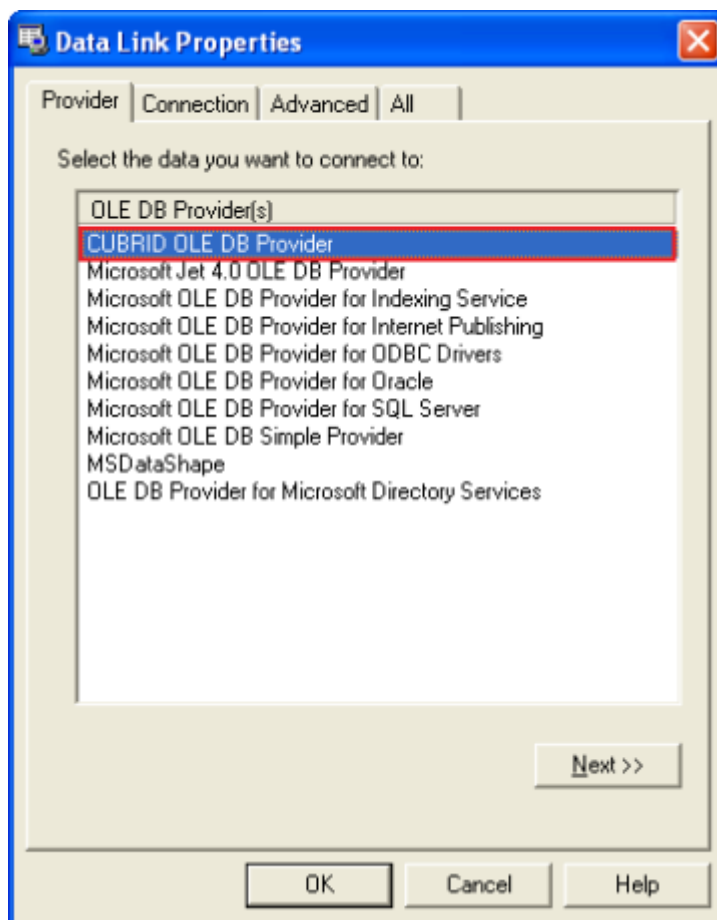
OLE DB API

OLE DB Programming

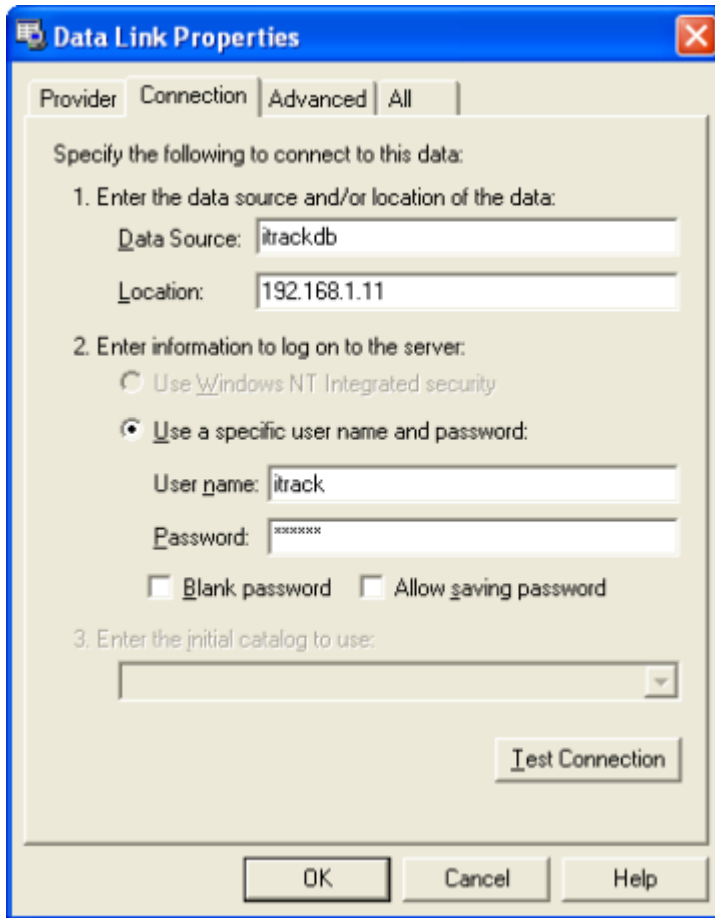
Using Data Link Property Dialog Box

In the [Data Link Properties] dialog box, you can check and configure various OLE DB providers provided by the current Windows operating system.

If you have properly installed the CUBRID OLE DB Provider for Windows, 'CUBRID OLE DB Provider' is displayed in the provider list of the [Data Link Properties] dialog box, as shown below.

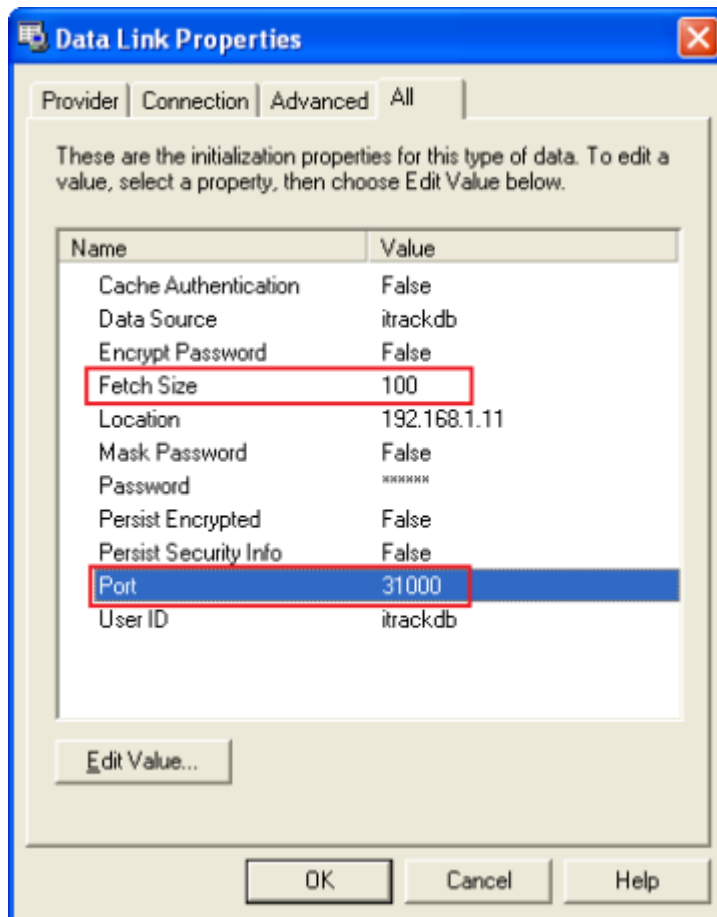


If you click the [Next] button after selecting 'CUBRID OLE DB Provider', the [Connection] tab appears as shown below. Set the desired link properties in the [Connection] tab.



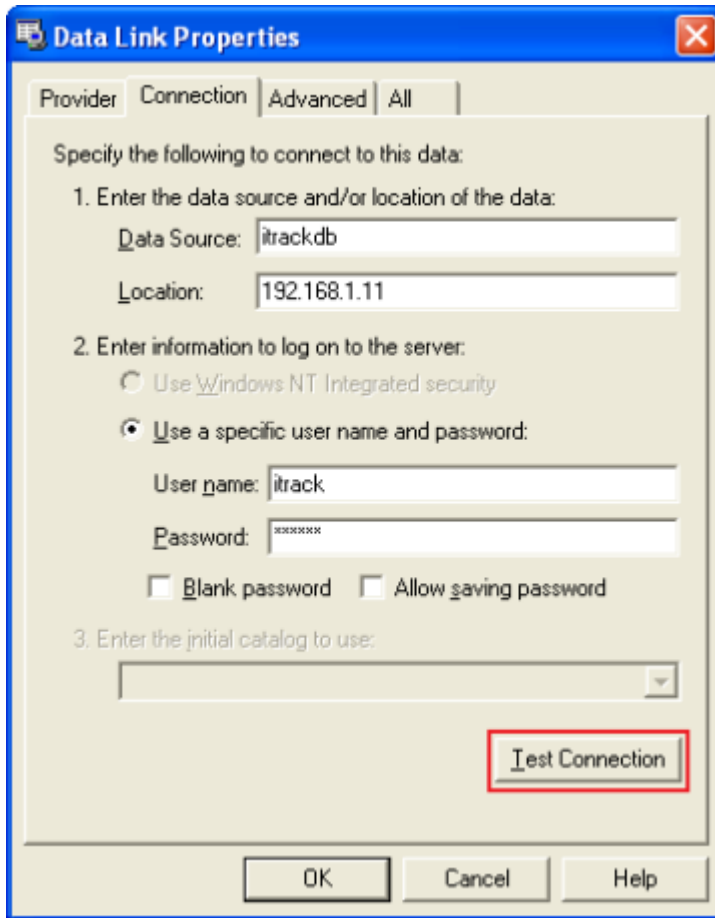
- **Data source** : Enter the name of the CUBRID database.
- **Location** : Enter the IP address or host name of the server where the CUBRID Broker is running.
- **User name** : Enter the name of the user who will log on to the database server.
- **Password** : Enter the password to be used for the database server logon.

Select all connection properties and then click the [All] tab.



To check every value currently configured, click the [All] tab; to edit the value, double-click the item you want. When the [Edit Property Value] dialog box appears, enter the desired value and then click [OK]. The figure above shows an example that configures the [Port] to "31000," and [Fetch Size] to "100."

You can check whether the connection is working properly by clicking the [Test Connection] button in the [Connection] tab after completing all configurations.



Configuring Connection String

When you program the CUBRID OLE DB Provider using ADO (ActiveX Data Object) or ADO.net, write the connection string as follows:

Item	Example	Description
Provider	CUBRIDProvider	Provider name
Data source	demodb	Database name
Location	127.0.0.1	The IP address of the CUBRID Broker Server
User ID	PUBLIC	User ID
Password	xxx	Password
Port	33000	Broker port number
Fetch Size	100	Fetch size

A connection string using the above example is as follows:

```
"Provider=CUBRIDProvider;Data Source=demodb;Location=127.0.0.1;User ID=PUBLIC;Password=xxx;Port= 33000;Fetch Size=100"
```

Remark

Because a semi-colon is used as a separator in URL string, it is not allowed to include a semi-colon for password (Password) when specifying it in connection string.

Note If INSERT/UPDATE for a longer string defined in a column is executed, the string will be truncated.

Multi-Thread Programming in .NET Environment

To develop programs by using the CUBRID OLE DB Provider in the Microsoft .NET, you should consider the followings:

If you develop multi-thread programs by using ADO.NET in the management environment, you need to change the value of the ApartmentState attribute of the Thread object to a ApartmentState.STA value because the CUBRID OLE DB Provider supports only Single Threaded Apartment (STA) attributes.

Without any change of given values, the default value of the attribute in the Thread object returns Unknown value, thereby causing abnormal process or errors during multi-threads programming.

Caution All OLE DB objects are COM objects. Currently, the CUBRID OLE DB Provider supports only the apartment threading model among COM threading models. It does not support the free threading model. This applies to not only the .NET but all multi-threaded environment.

Note If a string longer than defined size in a column is inserted(INSERT) or updated(UPDATE), the string will be truncated.

PHP API

PHP Programming

General Features

Connection

Connecting to a database : The first step of a database application is to use the [cubrid_connect\(\)](#) or [cubrid_connect_with_url\(\)](#) function which provide a database connection. Once the [cubrid_connect\(\)](#) or [cubrid_connect_with_url\(\)](#) function is executed successfully, you can use any functions available in the database. It is very important to call the [cubrid_disconnect\(\)](#) function before the application is terminated completely. The [cubrid_disconnect\(\)](#) function terminates the current transaction as well as the connection handle and all request handles created by the [cubrid_connect\(\)](#) function.

Transactions and auto-commit

CUBRID PHP supports both transaction and auto-commit mode. Auto-commit mode means that every query that you run has its own implicit transaction. You can use the [cubrid_get_autocommit\(\)](#) function to get the status of current connection auto-commit mode, and use the [cubrid_set_autocommit\(\)](#) function to enable/disable auto-commit mode of current connection. When [cubrid_set_autocommit\(\)](#) function is called, concurrent transactions are committed regardless of the auto-commit mode.

You can configure the default value of auto-commit mode by using `CCI_DEFAULT_AUTOCOMMIT` (broker parameter) upon startup of an application. If configuration on broker parameter is omitted, the default value is **ON**. You can also use [cubrid_connect_with_url\(\)](#) function to set the autocommit-mode when you establish the database connection. For example:

```
$con = cubrid_connect_with_url("cci:CUBRID:localhost:33000:demodb:dba::?autocommit=true");
```

If you need a transaction, you must use the [cubrid_set_autocommit\(\)](#) function to disable the auto-commit mode.

The [cubrid_commit\(\)](#) or [cubrid_rollback\(\)](#) function is used to commit or roll back a transaction. The [cubrid_disconnect\(\)](#) function terminates the transaction and rolls back uncommitted ones.

Processing Queries

The following are basic steps of query execution.

- Creating a connection handle
- Creating a request handle for an SQL query request
- Fetching the result
- Terminating the request handle

```
$con = cubrid_connect("192.168.0.10", 33000, "demodb");
if($con) {
    $req = cubrid_execute($con, "select * from code");
    if($req) {
        while ($row = cubrid_fetch($req)) {
            echo $row["s name"];
            echo $row["f name"];
        }
        cubrid_close_request($req);
    }
    cubrid_disconnect($con);
}
```

Column types and names of the query result

The [cubrid_column_types\(\)](#) function is used to get an array containing column types, and the [cubrid_column_names\(\)](#) function is used to get an array containing column names.

```

$req = cubrid_execute($con, "select host year, host city from olympic");
if($req) {
    $col_types = cubrid_column_types($req);
    $col_names = cubrid_column_names($req);

    while (list($key, $col_type) = each($col_types)) {
        echo $col_type;
    }
    while (list($key, $col_name) = each($col_names)) {
        echo $col_name;
    }
    cubrid_close_request($req);
}

```

Adjusting the cursor

You can configure the position of the query result. The [cubrid_move_cursor\(\)](#) function is used to move the cursor to a certain position from one of three points: the beginning of the query result, the current cursor position and the end of the query result.

```

$req = cubrid_execute($con, "select host year, host city from olympic order by host year");
if($req) {
    cubrid_move_cursor($req, 20, CUBRID_CURSOR_CURRENT)
    while ($row = cubrid_fetch($req, CUBRID_ASSOC)) {
        echo $row["host year"]." ";
        echo $row["host city"]."\n"; }
}

```

Result array types

One of the following three types of arrays is used in the result of the [cubrid_fetch\(\)](#) function. The type of the array can be determined when the [cubrid_fetch\(\)](#) function is called. The associative array uses character string indexes. The numeric array uses numeric order indexes. The last array type includes both associative and numeric arrays.

- Numeric array

```

while (list($id, $name) = cubrid_fetch($req, CUBRID_NUM)) {
    echo $id;
    echo $name;
}

```

- Associative array

```

while ($row = cubrid_fetch($req, CUBRID_ASSOC)) {
    echo $row["id"];
    echo $row["name"];
}

```

Catalog Operation

Information about the database schema such as classes, virtual classes, attributes, functions, triggers and constraints can be obtained by calling the [cubrid_schema\(\)](#) function. The return value of the [cubrid_schema\(\)](#) function is a two-dimensional array.

```

$pk = cubrid_schema($con, CUBRID_SCH_PRIMARY_KEY, "game");
if ($pk) {
    print_r($pk);
}

$fk = cubrid_schema($con, CUBRID_SCH_IMPORTED_KEYS, "game");
if ($fk) {
    print_r($fk);
}

```

Processing Errors

When an error occurs, most PHP interface functions display the error message and return false or -1. Each error message, error code or error facility code can be checked by using the [cubrid_error_msg\(\)](#), [cubrid_error_code\(\)](#), and [cubrid_error_code_facility\(\)](#) functions.

The return value of the [cubrid_error_code_facility\(\)](#) function is one of **CUBRID_FACILITY_DBMS** (DBMS error), **CUBRID_FACILITY_CAS** (CAS server error), **CUBRID_FACILITY_CCI** (CCI error) and **CUBRID_FACILITY_CLIENT** (PHP module error).

CUBRID Features

Using OIDs

With a query that can update the **CUBRID_INCLUDE_OID** option in the [cubrid_execute\(\)](#) function, you can get the OID value of the current row updated by the executing [cubrid_current_oid\(\)](#).

```
$req = cubrid execute($con, "select * from person where id = 1", CUBRID INCLUDE OID);
if ($req) {
    while ($row = cubrid fetch($req)) {
        echo cubrid current oid($req);
        echo $row["id"];
        echo $row["name"];
    }
    cubrid close request($req);
}
```

You can get all attributes, the specified attribute or an attribute of an instance by using the OID.

If you don't specify any attribute in the [cubrid_get\(\)](#) function, the values of all attributes are returned (a). If you specify an attribute as an array data type, an associative array containing the values of the specified attribute is returned (b). If you specify an attribute as a character string array, the value of the attribute is returned (c).

```
$attrarray = cubrid get($con, $oid); // (a)
$attrarray = cubrid get($con, $oid, array("id", "name")); // (b)
$attrarray = cubrid_get($con, $oid, "id"); // (c)
```

You can also update an attribute value of an instance by using the OID. To update a single attribute value, specify the attribute name as a character string type and its value (a). To set multiple attribute values, specify an associative array containing the attribute names and values (b).

```
$cubrid_put ($con, $oid, "id", 1); // (a)
$cubrid_put ($con, $oid, array("id"=>1, "name"=>"Tomas")); // (b)
```

Using Collections

- Collection data types can be used by using either PHP array data types or PHP functions that support array data types. The following example shows how to fetch the query result with the [cubrid_fetch\(\)](#) function.

```
$row = cubrid fetch ($req);
$col = $row["customer"];
while (list ($key, $cust) = each ($col)) {
    echo $cust;
}
```

- You can also get values of collection attributes. The following example shows how to get collection attribute values with the [cubrid_col_get\(\)](#) function.

```
$tels = cubrid col get ($con, $oid, "tels");
while (list ($key, $tel) = each ($tels)) {
    echo $tel."\n";
}
```

- You can directly update collection type values with [cubrid_set_add\(\)](#) and [cubrid_set_drop\(\)](#) functions.

```
$tels = cubrid col get ($con, $oid, "tels");
while (list ($key, $tel) = each ($tels)) {
    $res = cubrid set drop ($con, $oid, "tel", $tel);
}
cubrid_commit ($con);
```

Note A string is truncated if longer string specified in the column is inserted or updated (**INSERT/UPDATE**).

cubrid_affected_rows

Description

The **cubrid_affected_rows** function gets the number of rows that have been affected by the SQL statements (**INSERT**, **DELETE**, and **UPDATE**).

Syntax

```
int cubrid_affected_rows([ resource $req_identifier ])
```

- *req_identifier* : Request identifier. If the request identifier is not specified, the last request is assumed.

Return Value

- Success : Returns the number of rows affected by the SQL statement.
- When last SQL statement is not **INSERT**, **UPDATE** or **DELETE** : -1
- When request identifier is not specified and there is no last request : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE cubrid test");
cubrid_execute($conn, "CREATE TABLE cubrid test (t varchar)");

for ($i = 0; $i < 5; $i++) {
    cubrid_execute($conn, "INSERT INTO cubrid_test(t) VALUES('cubrid_test')");
}

cubrid_execute($conn, "DELETE FROM cubrid test");

$affected_num = cubrid_affected_rows();
var_dump($affected_num);

cubrid_disconnect($conn);
?>
```

The above example will output:

```
int(5)
```

See Also

- [cubrid_execute](#)

cubrid_bind

Description

The **cubrid_bind** function substitutes a value for a variable of the [cubrid_prepare\(\)](#) with parameters, a various types in PHP and corresponding types in SQL. If *bind_value_type* is not given, string will be the default. The following table shows the types of substitute values:

Support	Bind type	Corresponding SQL type
Supported	STRING	CHAR, VARCHAR
	NCHAR	NCHAR, NVARCHAR
	BIT	BIT, VARBIT
	NUMERIC or NUMBER	SHORT, INT, NUMERIC
	FLOAT	FLOAT

	DOUBLE	DOUBLE
	TIME	TIME
	DATE	DATE
	TIMESTAMP	TIMESTAMP
	OBJECT	OBJECT
	BLOB	BLOB
	CLOB	CLOB
	NULL	NULL
Not supported	SET	SET
	MULTISET	MULTISET
	SEQUENCE	SEQUENCE

Syntax

```
bool cubrid_bind(resource $req_identifier, mixed $bind_param, mixed $bind_value[, string $bind_value_type])
```

- *req_identifier* : Request identifier as a result of [cubrid_prepare\(\)](#)
- *bind_param* : Parameter identifier. For a prepared statement using named placeholders, this will be a parameter name of the form :name (Note that the name can only contain digit, alphabet, and underscore, and it cannot begin with digit. The name length cannot be longer than 32). For a prepared statement using question mark placeholders, this will be the 1-indexed position of the parameter.
- *bind_value* : Actual value to be bound
- *bind_value_type* : Type of the value to be bound. It can be omitted by default. If it is omitted, the type is automatically cast to an appropriate one. However, **NCHAR**, **BLOB/CLOB** and **BIT** types must be passed as arguments.

Note If data is bound to **BLOB/CLOB**, CUBRID maps the data into PHP stream, which is a unified approach to handle files and sockets in PHP extension. If the actual value to be bound is not stream, CUBRID returns it as string; this string includes the **BLOB/CLOB** type meta data (Locator) which is a file path and name recorded in an external storage.

Return Value

- Success : TRUE
- Failure : FALSE

Example 1

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$result = cubrid_execute($conn, "SELECT code FROM event WHERE sports='Basketball' and
gender='M'");
$row = cubrid_fetch_array($result, CUBRID_ASSOC);
$event_code = $row["code"];

cubrid_close_request($result);

$game_req = cubrid_prepare($conn, "SELECT athlete_code FROM game WHERE host_year=1992 and
event_code=? and nation_code='USA'");
cubrid_bind($game_req, 1, $event_code, "number");
cubrid_execute($game_req);

printf("--- Dream Team (1992 United States men's Olympic basketball team) ---\n");
while ($athlete_code = cubrid_fetch_array($game_req, CUBRID_NUM)) {
    $athlete_req = cubrid_prepare($conn, "SELECT name FROM athlete WHERE code=? AND
nation_code='USA' AND event='Basketball' AND gender='M'");
    cubrid_bind($athlete_req, 1, $athlete_code[0], "number");
    cubrid_execute($athlete_req);
    $row = cubrid_fetch_assoc($athlete_req);
```

```

    printf("%s\n", $row["name"]);
}
cubrid_close_request($game_req);
cubrid_close_request($athlete_req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```

--- Dream Team (1992 United States men's Olympic basketball team) ---
Stockton John
Robinson David
Pippen Scottie
Mullin C.
Malone Karl
Laettner C.
Jordan Michael
Johnson Earvin
Ewing Patrick
Drexler Clyde
Bird Larry
Barkley Charles

```

Example 2

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$sql stmt = <<<EOD
SELECT s.name FROM stadium s, game g
WHERE s.code = g.stadium_code AND g.medal = :medal_type
GROUP BY g.stadium code ORDER BY count(medal) DESC LIMIT 1;
EOD;

$req = cubrid_prepare($conn, $sql stmt);
printf("%-30s %s\n", "Medal Type", "Stadium where most medals were ever won");

cubrid_bind($req, ":medal_type", "G");
cubrid_execute($req);
$row = cubrid_fetch_assoc($req);
printf("%-30s %s\n", "Gold", $row["name"]);

cubrid_bind($req, ":medal_type", "S");
cubrid_execute($req);
$row = cubrid_fetch_assoc($req);
printf("%-30s %s\n", "Silver", $row["name"]);

cubrid_bind($req, ":medal_type", "B");
cubrid_execute($req);
$row = cubrid_fetch_assoc($req);
printf("%-30s %s\n", "Bronze", $row["name"]);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

Medal Type	Stadium where most medals were ever won
Gold	Olympic Aquatic Centre
Silver	Olympic Aquatic Centre
Bronze	Sydney Convention and Exhibition Centre

Example 3

```

<?php
$con = cubrid_connect("localhost", 33000, "foo");
if ($con) {
    $sql = "INSERT INTO php_cubrid_lob_test(doc_content) VALUES(?)";
    $req = cubrid_prepare($con, $sql);

    $fp = fopen("book.txt", "rb");

```

```

cubrid bind($req, 1, $fp, "blob");
cubrid execute($req);
}
?>

```

Example 4

```

<?php
$con = cubrid connect("localhost", 33000, "foo");
if ($con) {
    $sql = "INSERT INTO php_cubrid_lob_test(image) VALUES(?)";
    $req = cubrid prepare($con, $sql);

    cubrid_bind($req, 1, "cubrid_logo.png", "blob");
    cubrid execute($req);
}
?>

```

See Also

- [cubrid_execute](#)
- [cubrid_prepare](#)

cubrid_client_encoding

Description

The `cubrid_client_encoding` function returns the charset configured in the currently connected CUBRID in string.

The [cubrid_get_charset](#) and `cubrid_client_encoding` functions are used interchangeably exception that an input argument can be omitted in the `cubrid_client_encoding` function. If the input argument is omitted, a connection identifier created most recently is handled as an input argument.

Syntax

```
string cubrid_client_encoding ([ resource $conn_identifier ])
```

- *conn_identifier*: The CUBRID connection. If the connection identifier is not specified, the last connection opened is assumed.

Return Value

- Success: A string that represents the CUBRID connection charset
- Failure: FALSE

Example

```

<?php
$con = cubrid_connect("localhost", 33000, "demodb");
if (!$con)
{
    die('Could not connect.');
```

```

}

printf("CUBRID current charset: %s\n", cubrid_client_encoding($con));
?>

```

See Also

- [cubrid_get_charset](#)

cubrid_close

Description

The **cubrid_close** function stops transactions currently being executed, disconnect connection from server, and close a connection handle. If any request handles not being closed yet exists, all of them will be closed. The **cubrid_close** and [cubrid_disconnect](#) functions are used interchangeably exception that an input argument can be omitted in the **cubrid_close** function. If the input argument is omitted, a connection identifier created most recently is handled as an input argument.

Syntax

```
bool cubrid_close ([resource $con_identifier])
```

- *con_identifier* : Connection identifier. If the connection identifier is not specified, the last connection opened is assumed.

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
$con = cubrid_connect("192.168.0.10", 33000, "demodb");
if ($con) {
    echo "connected successfully";
    $req = cubrid_execute( $con, "insert into person values(1,'James')");
    if ($req) {
        cubrid_close_request($req);
        cubrid_commit($con);
    } else {
        cubrid_rollback($con);
    }
    cubrid_close($con);
}
```

See Also

- [cubrid_connect](#)

cubrid_close_prepare, cubrid_close_request

Description

The **cubrid_close_prepare** and **cubrid_close_request** functions close the request handle given to the *req_identifier* parameter and release the memory area related to the handle.

Syntax

```
int cubrid_close_prepare (resource $req_identifier)
int cubrid_close_request (resource $req_identifier)
```

- *req_identifier* : Request identifier

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = cubrid_prepare ($conn, "SELECT * FROM olympic WHERE host_year=?");
```

```

$host_year = 2004;
cubrid_bind($req, 1, $host_year, "number");
cubrid_execute($req);

printf("%-9s %-11s %-9s %-12s %-12s %-15s %-15s\n",
       "host year", "host nation", "host city", "opening date",
       "closing date", "mascot", "slogan");

while ($row = cubrid_fetch_assoc($req)) {
    printf("%-9s %-11s %-9s %-12s %-12s %-15s %-15s\n",
           $row["host year"], $row["host nation"], $row["host city"],
           $row["opening date"], $row["closing date"], $row["mascot"], $row["slogan"]);
}

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

host year	host nation	host city	opening date	closing date	mascot	slogan
2004	Greece	Athens	2004-8-13	2004-8-29	Athena Phevos	Welcome Home

See Also

- [cubrid_execute](#)

cubrid_col_get

Description

The `cubrid_col_get` function gets the elements of the given collection type (set, multiset, sequence) attribute in the form of an array.

Syntax

```
array cubrid_col_get (resource $conn_identifier, string $oid, string $attr_name)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the attribute to be read from the instance

Return Value

- Success : An array that contains the desired elements (0 : default numeric array)
- Failure : FALSE. If an error occurs, a warning message is displayed to distinguish it from a collection without attributes or NULL. You can check the error with [cubrid_error_code\(\)](#).

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

```

```
$size = cubrid_col_size($conn, $oid, "b");
var_dump($size);
```

```
cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(3) {
  [0]=>
    string(1) "1"
  [1]=>
    string(1) "2"
  [2]=>
    string(1) "3"
}
int(3)
```

cubrid_col_size

Description

The **cubrid_col_size** function gets the number of elements of a collection type (set, multiset, sequence) attribute.

Syntax

```
int cubrid_col_size(resource $conn_identifier, string $oid, string $attr_name)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the desired attribute of the instance

Return Value

- Success : The number of elements
- Failure : FALSE

Example

```
$elem_count = cubrid_col_size($con, $oid, "tel");
echo "$oid (tel) has $elem_count elements\n";
```

cubrid_column_names

Description

The **cubrid_column_names** function gets column names in the query results by using *req_identifier*.

Syntax

```
array cubrid_column_names(resource $req_identifier)
```

- *req_identifier* : Request identifier

Return Value

- Success : An array that contains the column names
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
```

```

$result = cubrid execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$column_names = cubrid_column_names($result);
$column_types = cubrid_column_types($result);

printf("%-30s %-30s %-15s\n", "Column Names", "Column Types", "Column Maxlen");
for($i = 0, $size = count($column_names); $i < $size; $i++) {
    $column_len = cubrid_field_len($result, $i);
    printf("%-30s %-30s %-15s\n", $column_names[$i], $column_types[$i], $column_len);
}

cubrid disconnect($conn);
?>

```

The above example will output:

Column Names	Column Types	Column Maxlen
host year	integer	11
event_code	integer	11
athlete_code	integer	11
stadium code	integer	11
nation code	char(3)	3
medal	char(1)	1
game_date	date	10

See Also

- [cubrid_execute](#)
- [cubrid_prepare](#)
- [cubrid_column_types](#)

cubrid_column_types

Description

The `cubrid_column_types` function gets column types in the query results by using `req_identifier`.

Syntax

```
array cubrid_column_types (resource $req_identifier)
```

- `req_identifier`: Request identifier

Return Value

- Success: An array that contains the column types
- Failure: FALSE

Example

```

<?php
$conn = cubrid connect("localhost", 33000, "demodb");
$result = cubrid execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$column_names = cubrid_column_names($result);
$column_types = cubrid_column_types($result);

printf("%-30s %-30s %-15s\n", "Column Names", "Column Types", "Column Maxlen");
for($i = 0, $size = count($column_names); $i < $size; $i++) {
    $column_len = cubrid_field_len($result, $i);
    printf("%-30s %-30s %-15s\n", $column_names[$i], $column_types[$i], $column_len);
}

cubrid disconnect($conn);
?>

```

The above example will output:

Column Names	Column Types	Column Maxlen
host_year	integer	11
event_code	integer	11
athlete code	integer	11
stadium code	integer	11
nation code	char(3)	3
medal	char(1)	1
game_date	date	10

See Also

- [cubrid_execute](#)
- [cubrid_prepare](#)
- [cubrid_column_names](#)

cubrid_commit

Description

The **cubrid_commit** function commits on the transaction pointed by *conn_identifier*, currently in progress. Connection to the server is closed after the **cubrid_commit**() function is called; the connection handle is still valid, though.

You can configure the default value of auto-commit mode by using **CCI_DEFAULT_AUTOCOMMIT** (broker parameter) upon startup of an application. If configuration on broker parameter is omitted, the default value is **ON**; use the [cubrid_set_autocommit](#)() function to change auto-commit mode within an application. If auto-commit mode is **OFF**, you must explicitly commit or roll back transaction by using [cubrid_commit](#)() or [cubrid_rollback](#)() function.

Syntax

```
bool cubrid_commit (resource $conn_identifier)
```

- *conn_identifier* : Connection identifier

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect("127.0.0.1", 33000, "demodb");

@cubrid execute($conn, "DROP TABLE publishers");

$sql = <<EOD
CREATE TABLE publishers(
pub id CHAR(3),
pub name VARCHAR(20),
city VARCHAR(15),
state CHAR(2),
country VARCHAR(15)
)
EOD;

if (!cubrid_execute($conn, $sql)) {
    printf("Error facility: %d\nError code: %d\nError msg: %s\n",
cubrid_error_code_facility(), cubrid_error_code(), cubrid_error_msg());

    cubrid_disconnect($conn);
    exit;
}

$req = cubrid_prepare($conn, "INSERT INTO publishers VALUES(?, ?, ?, ?, ?)");
$id_list = array("P01", "P02", "P03", "P04");
```

```

$name_list = array("Abatis Publishers", "Core Dump Books", "Schadenfreude Press",
"Tenterhooks Press");
$city_list = array("New York", "San Francisco", "Hamburg", "Berkeley");
$state_list = array("NY", "CA", NULL, "CA");
$country_list = array("USA", "USA", "Germany", "USA");

for ($i = 0, $size = count($id_list); $i < $size; $i++) {
    cubrid_bind($req, 1, $id_list[$i]);
    cubrid_bind($req, 2, $name_list[$i]);
    cubrid_bind($req, 3, $city_list[$i]);
    cubrid_bind($req, 4, $state_list[$i]);
    cubrid_bind($req, 5, $country_list[$i]);

    if (!$ret = cubrid_execute($req)) {
        break;
    }
}

if (!$ret) {
    cubrid_rollback($conn);
} else {
    cubrid_commit($conn);

    $req = cubrid_execute($conn, "SELECT * FROM publishers");
    while ($result = cubrid_fetch_assoc($req)) {
        printf("%-3s %-20s %-15s %-3s %-15s\n",
            $result["pub_id"], $result["pub_name"], $result["city"], $result["state"],
            $result["country"]);
    }
}

cubrid_disconnect($conn);
?>

```

The above example will output:

P01	Abatis Publishers	New York	NY	USA
P02	Core Dump Books	San Francisco	CA	USA
P03	Schadenfreude Press	Hamburg		Germany
P04	Tenterhooks Press	Berkeley	CA	USA

See Also

- [cubrid_rollback](#)

cubrid_connect

Description

The **cubrid_connect** function configures the connection environment with the server by using the given information such as the server address, port number, database name, user name and password. If the user name and password are not set, **PUBLIC** is used as default.

Syntax

```
resource cubrid_connect (string $host, int $port, string $dbname[, string $userid[, string $passwd[, bool $new_link]])
```

- *host* : IP address and host name of the Broker Server
- *port* : Port number of the Broker Server
- *dbname* : Database name
- *userid* : Database user name
- *passwd* : Database user password
- *new_link* : Whether to re-use an existing connection if an environment is identical in one HTTP request. If the value is **true**, a new connection identifier is created; if **false** and a connection with the same address, port number, database name, and user name exists, the existing connection is used. Note that this option is only valid within HTTP request; when HTTP request is closed, all connections are closed as well.

Return Value

- Success : Connection handle
- Failure : FALSE

Example

```
<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid version());

printf("\n");

$conn = cubrid connect("localhost", 33000, "demodb");

if (!$conn) {
    die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg());
}

$db_params = cubrid get db parameter($conn);

while (list($param name, $param value) = each($db params)) {
    printf("%-30s %s\n", $param_name, $param_value);
}

printf("\n");

$server_info = cubrid get server info($conn);
$client_info = cubrid_get_client_info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid_get_charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid_disconnect($conn);
?>
```

The above example will output:

```
CUBRID PHP Version:          8.3.1.0005

PARAM_ISOLATION_LEVEL      3
LOCK_TIMEOUT                -1
MAX_STRING_LENGTH          1073741823
PARAM_AUTO_COMMIT          0

Server Info:                8.3.1.0173
Client Info:                8.3.1

CUBRID Charset:            iso8859-1
```

See Also

- [cubrid_disconnect](#)

cubrid_connect_with_url

Description

The **cubrid_connect_with_url** function tries to connect a database by using connection information passed with a *url* string argument. If CUBRID HA is enabled in PHP, you must specify connection information of the active server and connection information of the standby server, which is used for failover when failure occurs, in the *url* string argument of this function.

Syntax

```
resource cubrid_connect_with_url (string $conn_url [, string $db_user,
string $db_password[, bool $new link]])

<conn_url> ::= cci:cubrid:<host>:<db_name>:<db_user>:<db_password>:[?<properties>]
<properties> ::= <property> [&<property>]
<property> ::= login_timeout=<milli_sec>
<property> ::= query_timeout=<milli_sec>
<property> ::= disconnect_on_query_timeout=true|false
<property> ::= autocommit=<autocommit mode>
<property> ::= althosts=<alternative_hosts> [ &rcptime=<time>]
<alternative_hosts> ::= <host>:<port> [,<host>:<port>]
<host> := HOSTNAME | IP_ADDR
<time> := SECOND
```

- *conn_url* : A character string that contains server connection information
- *host* : A host name or IP address of the master database
- *db_name* : A name of the database
- *db_user* : (IN) A name of the database user
- *db_password* : A database user password
- **login_timeout** : Timeout value (unit : msec.) upon login to a database. If time exceeds the value specified in this parameter, an error is returned. The default value is 0, which means infinite wait.
- **query_timeout** : Timeout value (unit : msec.). It configures timeout value based on query request. If timeout occurs, a cancellation message of query request, which has been sent to a server, is returned. The return value of a function which executes query is different based on configuration of **disconnect_on_query_timeout** and the request can succeed even though a cancellation message is sent to a server.
- **disconnect_on_query_timeout** : Configures whether to return an error immediately if timeout on the query request happens. The default value is **false**. If timeout on query request happens and this value is **true**, an error is returned after a cancellation message is sent and a socket is disconnected. In this case, a user should close the connection handle to a database explicitly by using the **cubrid_disconnect** function. If the value is **false**, a cancellation message is sent and it will wait until a response on query request is made.
- **autocommit=true/false** : The database connection auto commit mode.
- **althosts=standby_broker1_host, standby_broker2_host, . . .** : Specifies the broker information of the standby server, which is used for failover when it is impossible to connect to the active server. You can specify multiple brokers for failover, and the connection to the brokers is attempted in the order listed in **althosts**.
- **rcptime** : An interval between the attempts to connect to the active broker in which failure occurred. After a failure occurs, the system connects to the broker specified by **althosts** (failover), terminates the transaction, and then attempts to connect to the active broker of the master database at every **rcptime**. The default value is 600 seconds.
- *db_user* : A name of the database user
- *db_passwd* : A database user password
- *new_link* : Whether to re-use an existing connection if an environment is identical in one HTTP request. If the value is **true**, a new connection identifier is created; if **false** and a connection with the same address, port number, database name, and user name exists, the existing connection is used. Note that this option is only valid within HTTP request; when HTTP request is closed, all connections are closed as well.

Return Value

- Success : Connection identifier
- Failure : FALSE

Example

```
<?php
$con = cubrid_connect_with_url("cci:CUBRID:localhost:33000:demodb:dba::?autocommit=true");
?>
```

Remark

Because a colon (:) and a question mark (?) is used as a separator in URL string, it is not allowed to include them for password of URL string. To use them, you must specify a user name (*db_user*) and a password (*db_password*) as a separate parameter.

cubrid_current_oid

Description

The **cubrid_current_oid** function gets the OID of the current cursor location from the query result. To use **cubrid_current_oid()**, the query executed must be an updatable query, and the **CUBRID_INCLUDE_OID** option must be included during the query execution.

Syntax

```
string cubrid_current_oid (resource $req_identifier)
```

- *req_identifier* : Request identifier

Return Value

- Success : OID of the current cursor position
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect("localhost", 33000, "demodb");

$req = cubrid execute($conn, "SELECT * FROM code", CUBRID INCLUDE OID);
$oid = cubrid current oid($req);
$res = cubrid_get($conn, $oid);

print_r($res);

cubrid disconnect($conn);
?>
```

The above example will output:

```
Array
(
    [s_name] => X
    [f_name] => Mixed
)
```

See Also

- [cubrid_execute](#)

cubrid_data_seek

Description

The **cubrid_data_seek** function moves the internal row pointer of the CUBRID result associated with the specified result identifier to point to the specified *row_number*. The next call to a CUBRID fetch function, such as [cubrid_fetch_assoc\(\)](#), would return that row.

Syntax

```
bool cubrid_data_seek (resource $req_identifier, int $row_number)
```

- *req_identifier* : Result identifier
- *row_number* : The desired row number of the new result pointer

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("127.0.0.1", 33000, "demodb");

$req = cubrid_execute($conn, "SELECT * FROM code");
cubrid_data_seek($req, 0);

$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_data_seek($req, 2);
$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_data_seek($req, 4);
$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(2) {
  [0]=>
  string(1) "X"
  [1]=>
  string(5) "Mixed"
}
array(2) {
  [0]=>
  string(1) "M"
  [1]=>
  string(3) "Man"
}
array(2) {
  [0]=>
  string(1) "S"
  [1]=>
  string(6) "Silver"
}
```

cubrid_db_name

Description

The `cubrid_db_name` function gets db name from results of [cubrid_list_dbs\(\)](#).

Syntax

```
string cubrid_db_name(resource $result, int $index)
```

- *result* : The result from a call to `cubrid_list_dbs`
- *index* : The index into the result set

Return Value

- Success : database name
- Failure : The index into the result set

Example

```
<?php
error_reporting(E_ALL);

$conn = cubrid_connect('dbhost', 33000, 'demodb');
$db_list = cubrid_list_dbs($conn);
```

```

$i = 0;
$cnt = cubrid_num_rows($db_list);
while ($i < $cnt) {
    echo cubrid_db_name($db_list, $i) . "\n";
    $i++;
}
?>

```

See Also

- [cubrid_list_dbs](#)

cubrid_disconnect

Description

The **cubrid_disconnect** function is used to stop transactions currently being executed, disconnect connection from server, and close a connection handle.

If any request handles not being closed yet exists, all of them will be closed. The **cubrid_disconnect** and [cubrid_close](#) functions are used interchangeably exception that an input argument can be omitted in the **cubrid_close** function. If the input argument is omitted, a connection identifier created most recently is handled as an input argument.

Syntax

```
bool cubrid_disconnect (resource $conn_identifier)
```

- *conn_identifier* : Connection identifier

Return Value

- Success : TRUE
- Failure : FALSE

Example

```

<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid_version());

printf("\n");

$conn = cubrid_connect("localhost", 33000, "demodb");

if (!$conn) {
    die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg());
}

$db_params = cubrid_get_db_parameter($conn);

while (list($param_name, $param_value) = each($db_params)) {
    printf("%-30s %s\n", $param_name, $param_value);
}

printf("\n");

$server_info = cubrid_get_server_info($conn);
$client_info = cubrid_get_client_info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid_get_charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid_disconnect($conn);
?>

```

The above example will output:

```
CUBRID PHP Version:      8.3.1.0005

PARAM ISOLATION LEVEL   3
LOCK TIMEOUT            -1
MAX STRING LENGTH      1073741823
PARAM_AUTO_COMMIT      0

Server Info:           8.3.1.0173
Client Info:          8.3.1

CUBRID Charset:       iso8859-1
```

See Also

- [cubrid_connect](#)

cubrid_drop

Description

The **cubrid_drop** function drops the desired instance from the database by using the OID.

Syntax

```
bool cubrid_drop (resource $conn_identifier, string $oid)
```

- *conn_identifier*: Connection identifier
- *oid*: OID of the instance to be deleted

Return Value

- Success: TRUE
- Failure: FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES (1, {1,2,3}, {11,22,33,333},
'a')");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES (2, {4,5,7}, {44,55,66,666},
'b')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

printf("--- Before Drop: ---\n");
$attr = cubrid_get($conn, $oid);
var_dump($attr);

if (cubrid_drop($conn, $oid)) {
    cubrid_commit($conn);
} else {
    cubrid_rollback($conn);
}

cubrid_close_request($req);

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);
```



```

cubrid move cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid current oid($req);

printf("\n--- After Drop: ---\n");
$attr = cubrid get($conn, $oid);
var_dump($attr);

cubrid close request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```

--- Before Drop: ---
array(4) {
  ["a"]=>
    string(1) "1"
  ["b"]=>
    array(3) {
      [0]=>
        string(1) "1"
      [1]=>
        string(1) "2"
      [2]=>
        string(1) "3"
    }
  ["c"]=>
    array(4) {
      [0]=>
        string(2) "11"
      [1]=>
        string(2) "22"
      [2]=>
        string(2) "33"
      [3]=>
        string(3) "333"
    }
  ["d"]=>
    string(10) "a          "
}

--- After Drop: ---
array(4) {
  ["a"]=>
    string(1) "2"
  ["b"]=>
    array(3) {
      [0]=>
        string(1) "4"
      [1]=>
        string(1) "5"
      [2]=>
        string(1) "7"
    }
  ["c"]=>
    array(4) {
      [0]=>
        string(2) "44"
      [1]=>
        string(2) "55"
      [2]=>
        string(2) "66"
      [3]=>
        string(3) "666"
    }
  ["d"]=>
    string(10) "b          "
}

```

See Also

- [cubrid_is_instance](#)

cubrid_errno, cubrid_error_code

Description

The `cubrid_errno` and `cubrid_error_code` functions get the code of the error that occurred during the API execution. Usually, the error message can be fetched when the API returns **FALSE**.

Syntax

```
int cubrid_errno ()
int cubrid_error_code ()
```

Return Value

- Error code

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = cubrid_prepare($conn , "SELECT * FROM code WHERE s_name=?");

$req = @cubrid_execute($req);
if (!$req) {
    printf("Error facility: %d\nError code: %d\nError msg: %s\n",
        cubrid_error_code_facility(), cubrid_error_code(), cubrid_error_msg());

    cubrid_disconnect($conn);
    exit;
}
?>
```

The above example will output:

```
Error facility: 4
Error code: -2015
Error msg: Some parameter not binded
```

See Also

- [cubrid_error_code_facility](#)
- [cubrid_error_msg](#)

cubrid_error, cubrid_error_msg

Description

The `cubrid_error` and `cubrid_error_msg` functions get the error message that occurred during the API execution. Usually, the error message can be fetched when the API returns **FALSE**.

Syntax

```
string cubrid_error ()
string cubrid_error_msg ()
```

Return Value

- Occurred error message

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

if (!@cubrid_schema($conn, 100000)) {
    printf("Error facility: %d\nError code: %d\nError msg: %s\n",
```

```

        cubrid_error_code_facility(), cubrid_error_code(), cubrid_error_msg());

    cubrid_disconnect($conn);
    exit;
}
?>

```

The above example will output:

```

Error facility: 2
Error code: -1015
Error msg: Invalid T_CCI_SCH_TYPE value

```

See Also

- [cubrid_error_code](#)
- [cubrid_error_code_facility](#)

cubrid_error_code_facility

Description

The **cubrid_error_code_facility** function gets a facility code (level at which the error occurred) from the code of the error that occurred during the API execution. Usually, the error code can be fetched when the API returns **FALSE**.

Syntax

```
int cubrid_error_code_facility ()
```

Return Value

- Facility code of the occurred error code :
CUBRID_FACILITY_DBMS, CUBRID_FACILITY_CAS,
CUBRID_FACILITY_CCI, CUBRID_FACILITY_CLIENT

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = @cubrid_execute($conn, "SELECT * FROM unknown");
if (!$req) {
    printf("Error facility: %d\nError code: %d\nError msg: %s\n",
        cubrid_error_code_facility(), cubrid_error_code(), cubrid_error_msg());

    cubrid_disconnect($conn);
    exit;
}
?>

```

The above example will output:

```

Error facility: 1
Error code: -493
Error msg: Syntax: syntax error, unexpected UNKNOWN

```

See Also

- [cubrid_error_code](#)
- [cubrid_error_msg](#)

cubrid_execute

Description

The **cubrid_execute** function executes a given SQL statement. It executes a query by using *conn_identifier* and SQL and then returns the request identifier created. This is an appropriate way to simply execute a query when parameter binding is not necessary.

This function is also used when executing **Prepared Statement** with [cubrid_prepare](#) and [cubrid_bind](#). In this case, required parameters are *req_identifier* and *option*.

The *option* parameter is used to determine whether to get OID after query execution and whether to execute the query in synchronous or asynchronous mode. **CUBRID_INCLUDE_OID** and **CUBRID_ASYNC** (or **CUBRID_EXEC_QUERY_ALL** if you want to execute multiple SQL statements) can be specified by using a bitwise OR operator (|). If not specified, neither of them is selected.

If the flag **CUBRID_EXEC_QUERY_ALL** is set, a synchronous mode (*sync_mode*) is used to retrieve query results and in such case the following rules are applied.

- The return value is the result of the first query.
- If an error occurs in any query, the execution is processed as a failure.
- For a query composed of in a query composed of q1 q2 q3 if an error occurs in q2 after q1 succeeds the execution, the result of q1 remains valid. That is, the previous successful query executions are not rolled back when an error occurs.
- If a query is executed successfully, the result of the second query can be obtained using [cubrid_next_result\(\)](#).

If *req_identifier* is the first argument for the execution of [cubrid_prepare\(\)](#), only **CUBRID_ASYNC** or **CUBRID_EXEC_QUERY_ALL** can be used as an option.

Syntax

```
resource cubrid_execute (resource $conn_identifier, string $SQL [, int $option])
```

- *conn_identifier* : Connection identifier
- *SQL* : SQL statement to be executed
- *option* : Query execution option - **CUBRID_INCLUDE_OID**, **CUBRID_ASYNC**, **CUBRID_EXEC_QUERY_ALL**

```
bool cubrid_execute (resource &req_identifier[, int $option])
```

- *req_identifier* : request identifier
- *option* : Query execution option - **CUBRID_ASYNC**, **CUBRID_EXEC_QUERY_ALL**

Return Value

- Success
- Request identifier : When process is successful and first parameter is *conn_identifier*
- TRUE : When process is successful and first argument is *req_identifier*
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$result = cubrid_execute($conn, "SELECT code FROM event WHERE name='100m Butterfly' and
gender='M'", CUBRID_ASYNC);
$row = cubrid_fetch_array($result, CUBRID_ASSOC);
$event_code = $row["code"];

cubrid_close_request($result);

$history_req = cubrid_prepare($conn, "SELECT * FROM history WHERE event code=?");
cubrid_bind($history_req, 1, $event_code, "number");
cubrid_execute($history_req);
```

```

printf("%-20s %-9s %-10s %-5s\n", "athlete", "host year", "score", "unit");
while ($row = cubrid_fetch_array($history_req, CUBRID_ASSOC)) {
    printf("%-20s %-9s %-10s %-5s\n",
        $row["athlete"], $row["host year"], $row["score"], $row["unit"]);
}

cubrid_close_request($history_req);

cubrid_disconnect($conn);
?>

```

The above example will output:

athlete	host_year	score	unit
Phelps Michael	2004	51.25	time

See Also

- [cubrid_close_request](#)
- [cubrid_commit](#)
- [cubrid_rollback](#)
- [cubrid_prepare](#)
- [cubrid_bind](#)

cubrid_fetch

Description

The **cubrid_fetch** function gets one row from the query result. After getting the result, a cursor automatically moves to the next row.

Syntax

```
mixed cubrid_fetch (resource &result [, int &type])
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *type* : Type of the result array to be fetched. CUBRID_NUM, CUBRID_ASSOC, CUBRID_BOTH, CUBRID_OBJECT

Return Value

- Success : Result array or object.

It is determined by the *type* parameter. If the *type* parameter is omitted, **CUBRID_BOTH** is used. If you want to get the query result as an object data type, column names must comply with identifier name rules allowed in PHP. For example, a column name "count(*)" cannot be fetched and used as an object type.

The following are different result types depending on *type*.

- CUBRID_NUM : Numeric array (0-default)
- CUBRID_ASSOC : Associative array
- CUBRID_BOTH : Numeric and associative arrays (default value)
- CUBRID_OBJECT : An object that has the attribute whose name is the same as the column name of the query result
- After getting the last row : FALSE
- Failure : FALSE

Example

```

<?php
$conn = cubrid_connect("localhost", 33088, "demodb");
$req = cubrid_execute($conn, "SELECT * FROM stadium WHERE nation code='GRE' AND seats > 10000");

```

```

printf("%-40s %-10s %-6s %-20s\n", "name", "area", "seats", "address");
while ($row = cubrid_fetch($req)) {
    printf("%-40s %-10s %-6s %-20s\n",
        $row["name"], $row["area"], $row["seats"], $row["address"]);
}

cubrid_close_request($req);

cubrid_disconnect($conn);
?>

```

The above example will output:

name	area	seats	address
Panathinaiko Stadium	86300.00	50000	Athens, Greece
Olympic Stadium	54700.00	13000	Athens, Greece
Olympic Indoor Hall	34100.00	18800	Athens, Greece
Olympic Hall	52400.00	21000	Athens, Greece
Olympic Aquatic Centre	42500.00	11500	Athens, Greece
Markopoulo Olympic Equestrian Centre	64000.00	15000	Markopoulo, Athens, Greece
Faliro Coastal Zone Olympic Complex	34650.00	12171	Faliro, Athens, Greece
Athens Olympic Stadium	120400.00	71030	Maroussi, Athens, Greece
Ano Liossia	34000.00	12000	Ano Liosia, Athens, Greece

See Also

- [cubrid_execute](#)

cubrid_fetch_array

Description

The **cubrid_fetch_array** function gets a single row from the query result and returns an array. A cursor automatically moves to the next row after getting the result.

Syntax

```
mixed cubrid_fetch_array (resource $result[, int $type])
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *type* : Type of the result array to be fetched. CUBRID_NUM, CUBRID_ASSOC, CUBRID_BOTH

Return Value

- Success : Returns an array of strings that corresponds to the fetched row, when process is successful.

The type of returned array depends on how type is defined. By using **CUBRID_BOTH** (default), you'll get an array with both associative and numeric indices, and you can decide which data type to use by setting the type argument.

The *type* variable can be set to one of the following values:

- CUBRID_NUM : Numeric array (0-based)
- CUBRID_ASSOC : Associative array
- CUBRID_BOTH : Numeric and associative arrays (default value)
- After getting the last row : FALSE
- Failure : FALSE

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = cubrid_execute($conn, "SELECT name,area,seats,address FROM stadium WHERE
nation_code='GRE' AND seats > 10000");

printf("%-40s %-10s %-6s %-20s\n", "name", "area", "seats", "address");
while ($row = cubrid_fetch_array($req, CUBRID_NUM)) {
    printf("%-40s %-10s %-6s %-20s\n", $row[0], $row[1], $row[2], $row[3]);
}

```

```
cubrid close request($req);
```

```
cubrid_disconnect($conn);
?>
```

The above example will output:

name	area	seats	address
Panathinaiko Stadium	86300.00	50000	Athens, Greece
Olympic Stadium	54700.00	13000	Athens, Greece
Olympic Indoor Hall	34100.00	18800	Athens, Greece
Olympic Hall	52400.00	21000	Athens, Greece
Olympic Aquatic Centre	42500.00	11500	Athens, Greece
Markopoulo Olympic Equestrian Centre	64000.00	15000	Markopoulo, Athens, Greece
Faliro Coastal Zone Olympic Complex	34650.00	12171	Faliro, Athens, Greece
Athens Olympic Stadium	120400.00	71030	Maroussi, Athens, Greece
Ano Liossia	34000.00	12000	Ano Liosia, Athens, Greece

See Also

- [cubrid_execute](#)

cubrid_fetch_assoc

Description

The **cubrid_fetch_assoc** function returns an associative array corresponding to the fetched row and moves forward the internal data pointer. If any record exists, **FALSE** is returned.

Syntax

```
array cubrid_fetch_assoc ( resource $result )
```

- *result* : Result that comes from a call to the [cubrid_execute](#) function

Return Value

- Success : Associative arrays
- After getting the last row : FALSE
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect("localhost", 33000, "demodb");
$req = cubrid execute($conn, "SELECT name,area,seats,address FROM stadium WHERE
nation_code='GRE' AND seats > 10000");
```

```
printf("%-40s %-10s %-6s %-20s\n", "name", "area", "seats", "address");
while ($row = cubrid fetch assoc($req)) {
    printf("%-40s %-10s %-6s %-20s\n",
        $row["name"], $row["area"], $row["seats"], $row["address"]);
}
```

```
cubrid close request($req);
```

```
cubrid_disconnect($conn);
?>
```

The above example will output:

name	area	seats	address
Panathinaiko Stadium	86300.00	50000	Athens, Greece
Olympic Stadium	54700.00	13000	Athens, Greece
Olympic Indoor Hall	34100.00	18800	Athens, Greece
Olympic Hall	52400.00	21000	Athens, Greece
Olympic Aquatic Centre	42500.00	11500	Athens, Greece
Markopoulo Olympic Equestrian Centre	64000.00	15000	Markopoulo, Athens, Greece

Faliro Coastal Zone Olympic Complex	34650.00	12171	Faliro, Athens, Greece
Athens Olympic Stadium	120400.00	71030	Maroussi, Athens, Greece
Ano Liossia	34000.00	12000	Ano Liosia, Athens, Greece

cubrid_fetch_field

Description

The **cubrid_fetch_field** function returns an object containing field information. The function can be used to obtain information about fields in the provided query result. The characteristics of the object are:

- `name` : Column name
- `table` : Name of the table where the column belongs
- `def` : Default value of the column
- `max_length` : Maximum length of the column
- `not_null` : 1 if the column cannot be **NULL**
- `unique_key` : 1 if the column is a unique key
- `multiple_key` : 1 if the column is a non-unique key
- `numeric` : 1 if the column is numeric
- `type` : The type of the column

Syntax

```
object cubrid_fetch_field ( resource $result [, int $field_offset= 0 ] )
```

- `result` : Result that comes from a call to [cubrid_execute\(\)](#)
- `field_offset` : The numerical field offset. If the field offset is not specified, the next field that was not yet retrieved by this function is retrieved. The `field_offset` starts at 0.

Return Value

- Success: Associative arrays
- After getting the last row : FALSE
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect("localhost", 33000, "demodb");
$req = cubrid execute($conn, "SELECT event code,athlete code,nation code,game date FROM
game WHERE host year=1988 and event code=20001;");

var_dump(cubrid_fetch_row($req));

cubrid field seek($req, 1);
$field = cubrid fetch field($req);

printf("\n--- Field Properties ---\n");
printf("%-30s %s\n", "name:", $field->name);
printf("%-30s %s\n", "table:", $field->table);
printf("%-30s \"%s\"\n", "default value:", $field->def);
printf("%-30s %d\n", "max lenght:", $field->max length);
printf("%-30s %d\n", "not null:", $field->not null);
printf("%-30s %d\n", "unique key:", $field->unique key);
printf("%-30s %d\n", "multiple key:", $field->multiple_key);
printf("%-30s %d\n", "numeric:", $field->numeric);
printf("%-30s %s\n", "type:", $field->type);

cubrid close request($req);

cubrid_disconnect($conn);
?>
```

The above example will output:


```

array(4) {
  [0]=>
    string(5) "20001"
  [1]=>
    string(5) "16681"
  [2]=>
    string(3) "KOR"
  [3]=>
    string(9) "1988-9-30"
}

--- Field Properties ---
name:           athlete code
table:          game
default value:  ""
max length:    5
not null:      1
unique key:    1
multiple key:  0
numeric:       1
type:          integer

```

cubrid_fetch_lengths

Description

The **cubrid_fetch_lengths** function returns the length of each field value in the last row fetched in array. If it fails, **FALSE** is returned.

Syntax

```
array cubrid_fetch_lengths (resource $result)
```

- *result* : The result handle that is being evaluated. This result comes from a call to [cubrid_execute\(\)](#).

Note If field data type is **BLOB/CLOB**, you should get its length by using [cubrid_lob_size\(\)](#).

Return Value

- Success : Array in which the length of each field value in the last row fetched is stored.
- After getting the last row : **FALSE**
- Failure : **FALSE**

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$result = cubrid_execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$row = cubrid_fetch_row($result);
print_r($row);

$lens = cubrid_fetch_lengths($result);
print_r($lens);

cubrid_disconnect($conn);
?>
The above example will output:

Array
(
    [0] => 2004
    [1] => 20085
    [2] => 15118
    [3] => 30134
    [4] => AUS
    [5] => G
    [6] => 2004-8-20

```

```

)
Array
(
    [0] => 4
    [1] => 5
    [2] => 5
    [3] => 5
    [4] => 3
    [5] => 1
    [6] => 9
)

```

cubrid_fetch_object

Description

As an object that has attributes corresponding to the specified record fetched, the **cubrid_fetch_object** function returns the return set of current row and moves forward an internal data pointer. The object returned at the execution of this function has attributes corresponding to the filed name of the record.

Syntax

```
object cubrid_fetch_object (resource $result[, string $class_name[, array $params]])
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *class_name* : The name of the class to instantiate. If not specified, a **stdClass** (stdClass is PHP's generic empty class that's used when casting other types to objects) object is returned.
- *params* : An optional array of parameters to pass to the constructor for class_name objects

Return Value

- Success: Object
- After getting the last row : FALSE
- Failure : FALSE

Example

```

<?php
$conn = cubrid connect("127.0.0.1", 33000, "demodb", "PUBLIC", "");
$res = cubrid_execute($conn, "SELECT * FROM code");

var dump(cubrid fetch object($res));

class demodb code {

public $s_name = null;
public $f_name = null;

public function toString() {
    var dump($this);
}
}

var dump(cubrid fetch object($res, "demodb code"));

class demodb_code_construct extends demodb_code {

public function    construct($s, $f) {
    $this->s name = $s;
    $this->f name = $f;
}
}

var dump(cubrid fetch object($res, 'demodb code construct', array('s name', 'f name')));
var dump(cubrid fetch object($res));

```

```

cubrid close request($res);
cubrid disconnect($conn);
?>

Output:
object(stdClass)#1 (2) {
  ["s_name"]=>
  string(1) "X"
  ["f_name"]=>
  string(5) "Mixed"
}
object(demodb code)#1 (2) {
  ["s_name"]=>
  string(1) "W"
  ["f_name"]=>
  string(5) "Woman"
}
object(demodb code construct)#1 (2) {
  ["s_name"]=>
  string(6) "s_name"
  ["f_name"]=>
  string(6) "f name"
}
object(stdClass)#1 (2) {
  ["s_name"]=>
  string(1) "B"
  ["f_name"]=>
  string(6) "Bronze"
}

```

cubrid_fetch_row

Description

The **cubrid_fetch_row** function returns a numerical array that corresponds to the fetched row and moves the internal data pointer ahead, or **FALSE** if there are no more rows. The value of each field is stored in array.

Syntax

```
array cubrid_fetch_row (resource $result)
```

- result* : Result that comes from a call to [cubrid_execute\(\)](#)

Return Value

- Success: Numeric array
- After getting the last row : **FALSE**
- Failure : **FALSE**

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = cubrid_execute($conn, "SELECT name,area,seats,address FROM stadium WHERE
nation code='GRE' AND seats > 10000");

printf("%-40s %-10s %-6s %-20s\n", "name", "area", "seats", "address");
while ($row = cubrid_fetch_row($req)) {
  printf("%-40s %-10s %-6s %-20s\n", $row[0], $row[1], $row[2], $row[3]);
}

cubrid_close_request($req);

cubrid_disconnect($conn);
?>

```

The above example will output:

```

name                               area      seats  address

```

Panathinaiko Stadium	86300.00	50000	Athens, Greece
Olympic Stadium	54700.00	13000	Athens, Greece
Olympic Indoor Hall	34100.00	18800	Athens, Greece
Olympic Hall	52400.00	21000	Athens, Greece
Olympic Aquatic Centre	42500.00	11500	Athens, Greece
Markopoulo Olympic Equestrian Centre	64000.00	15000	Markopoulo, Athens, Greece
Faliro Coastal Zone Olympic Complex	34650.00	12171	Faliro, Athens, Greece
Athens Olympic Stadium	120400.00	71030	Maroussi, Athens, Greece
Ano Liossia	34000.00	12000	Ano Liosia, Athens, Greece

cubrid_field_flags

Description

The **cubrid_field_flags** function returns the field flags of the specified field. The flags are reported as a single word per flag separated by a single space, so that you can split the return value using **explode()**.

Syntax

```
string cubrid_field_flags ( resource $result , int $field_offset )
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *field_offset* : The field_offset starting 0.

Return Value

- Success: A string with flags
- Invalid *field_offset* value : FALSE
- SQL statement excluding SELECT : -1

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$result = cubrid_execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$col num = cubrid_num_cols($result);

printf("%-30s %s\n", "Field Name", "Field Flags");
for($i = 0; $i < $col num; $i++) {
    printf("%-30s %s\n", cubrid_field_name($result, $i), cubrid_field_flags($result, $i));
}

cubrid_disconnect($conn);
?>
```

The above example will output:

Field Name	Field Flags
host_year	not_null primary_key unique_key
event_code	not_null primary_key unique_key foreign_key
athlete code	not null primary key unique key foreign key
stadium code	not null
nation code	
medal	
game_date	

cubrid_field_len

Description

The **cubrid_field_len** function returns the length of the specified field on success, or **FALSE** on failure.

Syntax

```
string cubrid_field_len ( resource $result , int $field_offset )
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *field_offset* : The numerical field offset. The field_offset starts at 0. If field_offset does not exist, an error occurs.

Return Value

- Success: Maximum length
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$result = cubrid_execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$column_names = cubrid_column_names($result);
$column_types = cubrid_column_types($result);

printf("%-30s %-30s %-15s\n", "Column Names", "Column Types", "Column Maxlen");
for($i = 0, $size = count($column_names); $i < $size; $i++) {
    $column_len = cubrid_field_len($result, $i);
    printf("%-30s %-30s %-15s\n", $column_names[$i], $column_types[$i], $column_len);
}

cubrid_disconnect($conn);
?>
```

The above example will output:

Column Names	Column Types	Column Maxlen
host year	integer	11
event code	integer	11
athlete_code	integer	11
stadium_code	integer	11
nation code	char(3)	3
medal	char(1)	1
game_date	date	10

cubrid_field_name

Description

The **cubrid_field_name** function returns the name of the specified field index on success, or **FALSE** on failure.

Syntax

```
string cubrid_field_name ( resource $result , int $field_offset )
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *field_offset* : The numerical field offset. The field_offset starts at 0. If field_offset does not exist, an error occurs.

Return Value

- Success: Name of specified field index
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$result = cubrid_execute($conn, "SELECT * FROM game WHERE host year=2004 AND
nation code='AUS' AND medal='G'");

$col_num = cubrid_num_cols($result);
```

```
printf("%-30s %s\n", "Field Name", "Field Flags");
for($i = 0; $i < $col_num; $i++) {
    printf("%-30s %s\n", cubrid_field_name($result, $i), cubrid_field_flags($result, $i));
}

cubrid_disconnect($conn);
?>
```

The above example will output:

Field Name	Field Flags
host year	not null primary key unique key
event code	not null primary key unique key foreign key
athlete_code	not_null primary_key unique_key foreign_key
stadium_code	not_null
nation code	
medal	
game_date	

cubrid_field_seek

Description

The `cubrid_field_seek` function sets a field offset value to be used in `cubrid_fetch_field()` function. If the `cubrid_fetch_field()` function that does not include a field offset is called, the field offset specified in this function is returned.

Syntax

```
bool cubrid_field_seek ( resource $result , int $field_offset )
```

- *result* : Result that comes from a call to `cubrid_execute()`
- *field_offset* : The numerical field offset. The `field_offset` starts at 0. If `field_offset` does not exist, an error occurs.

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$req = cubrid_execute($conn, "SELECT event code,athlete code,nation code,game date FROM
game WHERE host year=1988 and event code=20001;");

var_dump(cubrid_fetch_row($req));

cubrid_field_seek($req, 1);
$field = cubrid_fetch_field($req);

printf("\n--- Field Properties ---\n");
printf("%-30s %s\n", "name:", $field->name);
printf("%-30s %s\n", "table:", $field->table);
printf("%-30s \"%s\"\n", "default value:", $field->def);
printf("%-30s %d\n", "max length:", $field->max_length);
printf("%-30s %d\n", "not null:", $field->not_null);
printf("%-30s %d\n", "unique key:", $field->unique_key);
printf("%-30s %d\n", "multiple key:", $field->multiple_key);
printf("%-30s %d\n", "numeric:", $field->numeric);
printf("%-30s %s\n", "type:", $field->type);

cubrid_close_request($req);

cubrid_disconnect($conn);
?>
```

The above example will output:

```

array(4) {
  [0]=>
  string(5) "20001"
  [1]=>
  string(5) "16681"
  [2]=>
  string(3) "KOR"
  [3]=>
  string(9) "1988-9-30"
}

--- Field Properties ---
name:                athlete code
table:               game
default value:      ""
max lenght:         5
not null:           1
unique key:         1
multiple key:       0
numeric:            1
type:               integer

```

cubrid_field_table

Description

The `cubrid_field_table` function returns the name of the table that the specified field is in.

Syntax

```
string cubrid_field_table (resource $result , int $field_offset)
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *field_offset* : The numerical field offset. The `field_offset` starts at 0. If `field_offset` does not exist, an error occurs.

Return Value

- Success : Name of the table of the specified field
- Invalid *field_offset* value : FALSE
- SQL sentence is not **SELECT** : -1

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");
$result = cubrid_execute($conn, "SELECT * FROM code");

$col num = cubrid_num_cols($result);

printf("%-15s %-15s %s\n", "Field Table", "Field Name", "Field Type");
for($i = 0; $i < $col num; $i++) {
    printf("%-15s %-15s %s\n",
        cubrid_field_table($result, $i), cubrid_field_name($result, $i),
        cubrid_field_type($result, $i));
}

cubrid_disconnect($conn);
?>

```

The above example will output:

Field Table	Field Name	Field Type
code	s name	char(1)
code	f_name	varchar(6)

cubrid_field_type

Description

The **cubrid_field_type** function returns the type of specified field and the field type returned is one of the types (example : "int", "float", or "string") that are supported by CUBRID.

Syntax

```
string cubrid_field_type (resource $result , int $field_offset)
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *field_offset* : The numerical field offset. The field_offset starts at 0. If field_offset does not exist, an error occurs.

Result Values

- Success : Type of the column
- Invalid *field_offset* value : FALSE
- SQL sentence is not **SELECT** : -1

Example

```
<?php
$conn = cubrid connect("localhost", 33000, "demodb");
$result = cubrid execute($conn, "SELECT * FROM code");

$col_num = cubrid_num_cols($result);

printf("%-15s %-15s %s\n", "Field Table", "Field Name", "Field Type");
for($i = 0; $i < $col_num; $i++) {
    printf("%-15s %-15s %s\n",
        cubrid_field_table($result, $i), cubrid_field_name($result, $i),
        cubrid_field_type($result, $i));
}

cubrid disconnect($conn);
?>
```

The above example will output:

Field Table	Field Name	Field Type
code	s name	char(1)
code	f name	varchar(6)

Output:

```
Your 'func' table has 4 fields and 1 record(s)
The table has the following fields:
string name 64 not null primary key binary
int ret 1 not null
string dl 128 not_null
string type 9 not_null enum
```

cubrid_free_result

Description

The **cubrid_free_result** function frees the memory occupied by the result data.

Note The **cubrid_free_result()** function can only free the client fetch buffer now, and if you want free all memory occupied by the result data, use function [cubrid_close_request\(\)](#).

Syntax

```
bool cubrid_free_result (resource $result)
```

- *result* : Result that comes from a call to [cubrid_execute](#)

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$req = cubrid_execute($conn, "SELECT * FROM history WHERE host_year=2004 ORDER BY
event code");
$row = cubrid_fetch_assoc($req);
var_dump($row);

cubrid_free_result($req);
cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(5) {
  ["event code"]=>
  string(5) "20005"
  ["athlete"]=>
  string(12) "Hayes Joanna"
  ["host_year"]=>
  string(4) "2004"
  ["score"]=>
  string(5) "12.37"
  ["unit"]=>
  string(4) "time"
}
```

cubrid_get

Description

The **cubrid_get** function gets a desired attribute of an instance by using OID. You can get a single attribute by using a character string type for the *attr* argument, or multiple attributes by using an array type.

Syntax

```
mixed cubrid_get (resource $conn_identifier, string $oid[, mixed $attr])
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the instance whose value you want to get
- *attr* : Name of the attribute whose value you want to get

Return Value

A character string is returned if a character string type is set for the *attr* argument; an associative array is returned if an array type (0 - default numeric array) is set. If the *attr* argument is omitted, all attributes of the instance are returned as an associative array.

- Success : Content of the attribute(s) requested
- Failure : FALSE. If an error occurs, a warning message is displayed to distinguish it from an empty character string or NULL. You can check the error with [cubrid_error_code\(\)](#).

Example

```
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO_INCREMENT, b set(int), c list(int), d
char(10))");
```

```

cubrid execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(2, {4,5,7}, {44,55,66,666},
'b')");

$req = cubrid execute($conn, "SELECT * FROM foo", CUBRID INCLUDE OID);

cubrid move cursor($req, 1, CUBRID CURSOR FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid get($conn, $oid, "b");
var_dump($attr);

$attr = cubrid get($conn, $oid);
var_dump($attr);

cubrid close request($req);
cubrid disconnect($conn);
?>
The above example will output:

```

```

string(9) "{1, 2, 3}"
array(4) {
  ["a"]=>
  string(1) "1"
  ["b"]=>
  array(3) {
    [0]=>
    string(1) "1"
    [1]=>
    string(1) "2"
    [2]=>
    string(1) "3"
  }
  ["c"]=>
  array(4) {
    [0]=>
    string(2) "11"
    [1]=>
    string(2) "22"
    [2]=>
    string(2) "33"
    [3]=>
    string(3) "333"
  }
  ["d"]=>
  string(10) "a          "
}

```

See Also

- [cubrid_put](#)

cubrid_get_autocommit

Description

The **cubrid_get_autocommit** function returns the result of auto-commit mode of database connection.

Syntax

```
bool cubrid_get_autocommit (resource $conn_identifier)
```

- *conn_identifier* : Connection identifier

Return Value

- Auto-commit mode ON : TRUE
- Auto-commit mode OFF : FALSE

See Also

- [cubrid_set_autocommit](#)

cubrid_get_charset

Description

The `cubrid_get_charset` function returns the charset configured in the currently connected CUBRID in string.

The `cubrid_get_charset` and [cubrid_client_encoding](#) functions are used interchangeably exception that an input argument can be omitted in the `cubrid_client_encoding` function. If the input argument is omitted, a connection identifier created most recently is handled as an input argument.

Syntax

```
string cubrid_get_charset (resource $req_identifier)
```

- `req_identifier` : Request identifier

Return Value

- Success : A string that represents the CUBRID connection charset
- Failure : FALSE

Example

```
<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid version());

printf("\n");

$conn = cubrid_connect("localhost", 33088, "demodb");

if (!$conn) {
    die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg());
}

$db_params = cubrid_get_db_parameter($conn);

while (list($param name, $param value) = each($db_params)) {
    printf("%-30s %s\n", $param name, $param value);
}

printf("\n");

$server_info = cubrid_get_server_info($conn);
$client_info = cubrid_get_client_info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid_get_charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid_disconnect($conn);

?>
```

The above example will output:

```
CUBRID PHP Version:      8.3.1.0005

PARAM_ISOLATION_LEVEL   3
LOCK_TIMEOUT            -1
MAX_STRING_LENGTH       1073741823
```

```
PARAM AUTO COMMIT          0
Server Info:                8.3.1.0173
Client Info:                8.3.1
CUBRID Charset:            iso8859-1
```

cubrid_get_class_name

Description

The `cubrid_get_class_name` function gets a class name from the OID.

Syntax

```
mixed cubrid_get_class_name (resource $conn_identifier, string $oid)
```

- `conn_identifier` : Connection identifier
- `oid` : OID of an instance, for which you want to check whether it exists

Return Value

- Success : Class name
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect ("localhost", 33088, "demodb");

$req = cubrid_execute ($conn, "SELECT * FROM code", CUBRID_INCLUDE_OID);
$oid = cubrid_current_oid ($req);
$class name = cubrid_get_class_name ($conn, $oid);

print_r ($class name);

cubrid_disconnect ($conn);
?>
```

The above example will output:
code

See Also

- [cubrid_is_instance](#)
- [cubrid_drop](#)

cubrid_get_client_info

Description

The `cubrid_get_client_info` function returns the library version of the current cci in string.

```
string cubrid_get_client_info ( void )
```

Return Value

- Success : A string that represents the client library version
- Failure : FALSE

Example

```
<?php
printf ("% -30s %s\n", "CUBRID PHP Version:", cubrid_version ());
printf ("\n");
```

```

$conn = cubrid connect("localhost", 33088, "demodb");

if (!$conn) {
    die('Connect Error ('. cubrid error code() .)'. cubrid error msg());
}

$db_params = cubrid get db parameter($conn);

while (list($param_name, $param_value) = each($db_params)) {
    printf("%-30s %s\n", $param_name, $param_value);
}

printf("\n");

$server_info = cubrid_get_server_info($conn);
$client_info = cubrid get client info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid get charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);
cubrid_disconnect($conn);

?>

```

The above example will output:

```

CUBRID PHP Version:          8.3.1.0005

PARAM ISOLATION LEVEL       3
LOCK TIMEOUT                 -1
MAX STRING LENGTH           1073741823
PARAM_AUTO_COMMIT           0

Server Info:                 8.3.1.0173
Client Info:                 8.3.1

CUBRID Charset:             iso8859-1

```

cubrid_get_db_parameter

Description

The `cubrid_get_db_parameter` function returns the CUBRID system parameters. It returns the CUBRID system parameters or it returns FALSE on failure. It returns an associative array with the values for the following parameters:

- **PARAM_ISOLATION_LEVEL** : In CUBRID PHP, you can set the level of transaction isolation by using [cubrid_set_db_parameter\(\)](#) function, `isolation_level` in the `$CUBRID/conf/cubrid.conf` or the **SET TRANSACTION** statement. For levels of isolation supported by CUBRID, see "CUBRID SQL Guide > Transaction and Lock > Transaction Isolation Level > SET TRANSACTION ISOLATION LEVEL."
- **PARAM_LOCK_TIMEOUT** : CUBRID provides the lock timeout feature, which sets the waiting time for the lock until the transaction lock setting is allowed. You can set lock timeout by using [cubrid_set_db_parameter\(\)](#) function, parameter `lock_timeout_in_secs` in the `$CUBRID/conf/cubrid.conf` file or the **SET TRANSACTION** statement (in seconds). The default value of the `lock_timeout_in_secs` parameter is -1, which means the application client will wait indefinitely until the transaction lock is allowed.
- **PARAM_MAX_STRING_LENGTH** : The maximum string length of a parameter
- **PARAM_AUTO_COMMIT** : In CUBRID PHP, an auto-commit mode is enabled by default for transaction management. If you want to start a transaction, you should set auto-commit mode to off by using the [cubrid_set_autocommit\(\)](#) function. And auto commit modes can be applied only for SELECT statements by setting broker parameters.

Syntax

```
array cubrid_get_db_parameter ( $req )
```

Return Value

- Success : An associative array that includes the value of CUBRID system parameter
- Failure : FALSE

Example

```
<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid version());

printf("\n");

$conn = cubrid connect("localhost", 33088, "demodb");

if (!$conn) {
    die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg());
}

$db_params = cubrid get db parameter($conn);

while (list($param_name, $param_value) = each($db_params)) {
    printf("%-30s %s\n", $param_name, $param_value);
}

printf("\n");

$server_info = cubrid_get_server_info($conn);
$client_info = cubrid_get_client_info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);
printf("\n");

$charset = cubrid_get_charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid disconnect($conn);

?>
```

The above example will output:

```
CUBRID PHP Version:          8.3.1.0005

PARAM ISOLATION LEVEL       3
LOCK TIMEOUT                 -1
MAX STRING LENGTH           1073741823
PARAM AUTO COMMIT           0

Server Info:                 8.3.1.0173
Client Info:                 8.3.1

CUBRID Charset:             iso8859-1
```

See Also

- [cubrid_set_db_parameter](#)

cubrid_get_query_timeout**Description**

The **cubrid_get_query_timeout** function returns timeout value configured for query execution.

Syntax

```
int cubrid_get_query_timeout(resource $conn_identifier)
```

- *conn_identifier* : Connection identifier

Return Value

- Success : Timeout value configured in the current request handle (unit : msec.)
- Failure : FALSE

cubrid_get_server_info

Description

The `cubrid_get_server_info` function returns the server version of CUBRID being currently connected in string.

Syntax

```
string cubrid_get_server_info ( void )
```

Return Value

- Success : A string that represents the CUBRID server version
- Failure : FALSE

Example

```
<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid version());

printf("\n");

$conn = cubrid_connect("localhost", 33088, "demodb");

if (!$conn) {
    die('Connect Error ('. cubrid error code() .')' . cubrid error msg());
}

$db_params = cubrid_get_db_parameter($conn);

while (list($param name, $param value) = each($db params)) {
    printf("%-30s %s\n", $param name, $param value);
}

printf("\n");

$server info = cubrid get server info($conn);
$client info = cubrid get client info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid get charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid disconnect($conn);

?>
```

The above example will output:

```
CUBRID PHP Version:          8.3.1.0005

PARAM_ISOLATION_LEVEL       3
LOCK_TIMEOUT                 -1
```

MAX STRING LENGTH	1073741823
PARAM AUTO COMMIT	0
Server Info:	8.3.1.0173
Client Info:	8.3.1
CUBRID Charset:	iso8859-1

cubrid_insert_id

Description

The **cubrid_insert_id** function retrieves the ID generated for the **AUTO_INCREMENT** columns which is updated by the previous **INSERT** query and returns every **AUTO_INCREMENT** column and its value in array. If the value of **AUTO_INCREMENT** is not generated in the previous query, 0 is returned; if CUBRID connection fails, **FALSE** is returned.

Note If more than one **AUTO_INCREMENT** column in a single table exists, you should not use this function.

Syntax

```
array cubrid_insert_id (string $class_name [, resource $conn_identifier])
```

- *class_name* : The name of the class (table) that was used in the last **INSERT** statement for which the auto increment values are retrieved.
- *connection_identifier* : Connection identifier previously obtained from a call to [cubrid_connect\(\)](#)

Return Value

- Success : Associative array that has every **AUTO_INCREMENT** column and its value
- If the previous query does not generate a new row : 0
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid execute($conn, "DROP TABLE cubrid test");
cubrid execute($conn, "CREATE TABLE cubrid test (d int AUTO INCREMENT(1, 2), t varchar)");

for ($i = 0; $i < 10; $i++) {
    cubrid execute($conn, "INSERT INTO cubrid test(t) VALUES('cubrid test')");
}

$id_list = cubrid_insert_id("cubrid test");
var_dump($id_list);

cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(1) {
  ["d"]=>
  int(19)
}
```

cubrid_is_instance

Description

The **cubrid_is_instance** function checks whether an instance referred to by the OID exists in the database.

Syntax

```
int cubrid_is_instance (resource $conn_identifier, string $oid)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of an instance, for which you want to check whether it exists

Return Value

- Instance : 1
- No instance : 0
- Error : -1

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$sql = <<<EOD
SELECT host year, medal, game date
FROM game
WHERE athlete code IN
      (SELECT code FROM athlete WHERE name='Thorpe Ian');
EOD;

$req = cubrid_execute($conn, $sql, CUBRID_INCLUDE_OID);
$oid = cubrid_current_oid($req);

$res = cubrid_is_instance ($conn, $oid);
if ($res == 1) {
    echo "Instance pointed by $oid exists.\n";
} else if ($res == 0){
    echo "Instance pointed by $oid doesn't exist.\n";
} else {
    echo "error\n";
}

cubrid_disconnect($conn);
?>
```

The above example will output:

```
Instance pointed by @0|0|0 doesn't exist.
```

See Also

- [cubrid_drop](#)
- [cubrid_get_class_name](#)

cubrid_lob_close

Description

The **cubrid_lob_close** function closes the external storage file of **BLOB/CLOB** type returned by [cubrid_lob_get\(\)](#).

Syntax

```
bool cubrid_lob_close (array $lob_identifier_array)
```

- *lob_identifier_array* : LOB identifier array returned from [cubrid_lob_get\(\)](#)

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$lobs = cubrid_lob_get($con, "SELECT doc content FROM doc WHERE doc id=5");
cubrid_lob_export($conn, $lobs[0], "doc 5.txt");
cubrid_lob_close($lobs);
?>
```

See Also

- [cubrid_lob_export](#)
- [cubrid_lob_get](#)
- [cubrid_lob_send](#)
- [cubrid_lob_size](#)

cubrid_lob_export

Description

The **cubrid_lob_export** function exports **BLOB/CLOB** type data to a file.

Syntax

```
bool cubrid_lob_export(resource $conn_identifier, resource $lob_identifier, string
$path_name )
```

- *conn_identifier* : Connection identifier
- *lob_identifier* : LOB identifier
- *path_name* : Path name of file

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$lobs = cubrid_lob_get($con, "SELECT doc content FROM doc WHERE doc id=5");
cubrid_lob_export($conn, $lobs[0], "doc 5.txt");
cubrid_lob_close($lobs);
?>
```

See Also

- [cubrid_lob_close](#)
- [cubrid_lob_get](#)
- [cubrid_lob_send](#)
- [cubrid_lob_size](#)

cubrid_lob_get

Description

The **cubrid_lob_get** function executes SQL statement and returns the value of every **BLOB/CLOB** type in resource array. Be sure that the SQL should have only one column and the type must be **BLOB/CLOB**.

Remember to use [cubrid_lob_close\(\)](#) to release the **BLOB** or **CLOB** if you don't need it any more.

Syntax

```
array cubrid_lob_get (resource $conn_identifier, string $SQL)
```

- *conn_identifier* : Connection identifier
- *SQL* : SQL statement to be executed.

Return Value

- Success : An array of LOB resources
- Failure : FALSE

Example

```
<?php
$llobs = cubrid_lob_get($con, "SELECT doc_content FROM doc WHERE doc_id=5");
cubrid_lob_export($conn, $llobs[0], "doc_5.txt");
cubrid_lob_close($llobs);
?>
```

See Also

- [cubrid_lob_close](#)
- [cubrid_lob_export](#)
- [cubrid_lob_send](#)
- [cubrid_lob_size](#)

cubrid_lob_send

Description

The **cubrid_lob_send** function reads **BLOB/CLOB** data and passes it straight through to the browser. To use this function, you should get **BLOB/CLOB** information by using the [cubrid_lob_get\(\)](#) function.

Syntax

```
bool cubrid_lob_send(resource $conn_identifier, resource $lob_identifier )
```

- *conn_identifier* : Connection identifier
- *lob_identifier* : LOB identifier

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$llobs = cubrid_lob_get($con, "SELECT image FROM person WHERE id=1");
Header("Content-type: image/jpeg");
cubrid_lob_send($conn, $llobs[0]);
cubrid_lob_close($llobs);
?>
```

See Also

- [cubrid_lob_close](#)
- [cubrid_lob_export](#)
- [cubrid_lob_get](#)
- [cubrid_lob_size](#)

cubrid_lob_size

Description

The **cubrid_lob_size** function returns the size of **BLOB/CLOB** data. The maximum size of **BLOB/CLOB** data is the same as the maximum file size in an external storage. LOB type size of CUBRID PHP is 64-bit integer but 64-bit integer cannot be returned in PHP. Therefore, string is returned instead.

Syntax

```
bool cubrid_lob_size(resource $lob_identifier )
```

- *lob_identifier* : LOB identifier

Return Value

- Success : LOB data size string
- Failure : FALSE

Example

```
<?php
$lobs = cubrid_lob_get($conn, "SELECT doc content FROM doc WHERE doc id=5");
echo "Doc size:".cubrid_lob_size($lobs[0]);
cubrid_lob_export($conn, $lobs[0], "doc 5.txt");
cubrid_lob_close($lobs);
?>
```

See Also

- [cubrid_lob_close](#)
- [cubrid_lob_export](#)
- [cubrid_lob_get](#)
- [cubrid_lob_send](#)

cubrid_list_dbs

Description

The **cubrid_list_dbs** function returns the name of every database existing in the server in array.

Syntax

```
array cubrid_list_dbs (resource $conn_identifier)
```

- *conn_identifier* : Connection identifier previously obtained from a call to [cubrid_connect\(\)](#)

Return Value

- Success : Numeric array of every CUBRID database
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33088, "demodb");

$db_list = cubrid_list_dbs($conn);
var_dump($db_list);

cubrid_disconnect($conn);
?>
The above example will output:
array(1) {
```

```
[0]=>
string(6) "demodb"
}
```

See Also

- [cubrid_db_name](#)

cubrid_lock_read

Description

The **cubrid_lock_read** function configures read lock on the desired instance by using the OID.

Syntax

```
bool cubrid_lock_read (resource $conn_identifier, string $oid)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of an instance on which you want to configure lock

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid connect("localhost", 33088, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");
cubrid execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(2, {4,5,7}, {44,55,66,666},
'b')");

$req = cubrid execute($conn, "SELECT * FROM foo", CUBRID INCLUDE OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

cubrid lock read($conn, $oid);

$attr = cubrid get($conn, $oid, "b");
var_dump($attr);

$attr = cubrid get($conn, $oid);
var_dump($attr); cubrid close request($req);
cubrid disconnect($conn);
?>
```

The above example will output:

```
string(9) "{1, 2, 3}"
array(4) {
  ["a"]=>
  string(1) "1"
  ["b"]=>
  array(3) {
    [0]=>
    string(1) "1"
    [1]=>
    string(1) "2"
    [2]=>
    string(1) "3"
  }
}
```

```

["c"]=>
array(4) {
  [0]=>
  string(2) "11"
  [1]=>
  string(2) "22"
  [2]=>
  string(2) "33"
  [3]=>
  string(3) "333"
}
["d"]=>
string(10) "a      "
}

```

See Also

- [cubrid_lock_write](#)

cubrid_lock_write

Description

The `cubrid_lock_write` function configures write lock on the desired instance by using the OID.

Syntax

```
bool cubrid_lock_write (resource $conn_identifier, string $oid)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of an instance on which you want to configure lock

Return Value

- Success : TRUE
- Failure : FALSE

Example

```

<?php
$conn = cubrid connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(2, {4,5,7}, {44,55,66,666},
'b')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID INCLUDE OID);

cubrid_move_cursor($req, 1, CUBRID CURSOR FIRST);
$oid = cubrid_current_oid($req);

cubrid_lock_write($conn, $oid);

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_put($conn, $oid, "b", array(2, 4, 8));

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```
array(3) {
  [0]=>
  string(1) "1"
  [1]=>
  string(1) "2"
  [2]=>
  string(1) "3"
}
array(3) {
  [0]=>
  string(1) "2"
  [1]=>
  string(1) "4"
  [2]=>
  string(1) "8"
}
```

See Also

- [cubrid_lock_read](#)

cubrid_move_cursor

Description

The **cubrid_move_cursor** function moves the current cursor position of *req_identifier* to the distance configured by the offset argument in the direction in the origin argument. For origin, the first position in the result (**CUBRID_CURSOR_FIRST**), the current position in the result (**CUBRID_CURSOR_CURRENT**) and the last position in the result (**CUBRID_CURSOR_LAST**) can be used. If origin is not specified, **CUBRID_CURSOR_CURRENT** is used by default.

If the amount of cursor movement exceeds the range of the result, the cursor moves to a position next to the end of the result range. For example, if the cursor moves to the position 20 when the size of the result is 10, it moves to the 11th position and returns **CUBRID_NO_MORE_DATA**.

Syntax

```
int cubrid_move_cursor (resource $req_identifier, int $offset[, int $origin])
```

- *req_identifier* : Request identifier
- *offset* : The number of positions to which the cursor is to be moved
- *origin* : Origin of the cursor movement
CUBRID_CURSOR_FIRST,
CUBRID_CURSOR_CURRENT,
CUBRID_CURSOR_LAST

Return Value

- Success : CUBRID_CURSOR_SUCCESS
- No data : CUBRID_NO_MORE_DATA
- Failure : CUBRID_CURSOR_ERROR

Example

```
<?php
$conn = cubrid_connect("127.0.0.1", 33000, "demodb");

$req = cubrid_execute($conn, "SELECT * FROM code");
cubrid_move_cursor($req, 1, CUBRID_CURSOR_LAST);

$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
```

```

$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_CURRENT);
$result = cubrid_fetch_row($req);
var_dump($result);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```

array(2) {
  [0]=>
  string(1) "G"
  [1]=>
  string(4) "Gold"
}
array(2) {
  [0]=>
  string(1) "X"
  [1]=>
  string(5) "Mixed"
}
array(2) {
  [0]=>
  string(1) "M"
  [1]=>
  string(3) "Man"
}

```

See Also

- [cubrid_execute](#)

cubrid_next_result

Description

The **cubrid_next_result** function returns the result of next query when the [cubrid_execute](#) function is executed by specifying the CUBRID_EXEC_QUERY_ALL flag. If the next query succeeds, a database is updated to the current query information.

Syntax

```
bool cubrid_next_result (resource $result)
```

- *result* : Result of calling [cubrid_execute\(\)](#)

Return Value

- Success : TRUE
- No next result : FALSE
- Failure : FALSE

Example

```

<?php
$conn = cubrid_connect($host, $port, $db, $user, $passwd);

$sql_stmt = "SELECT * FROM code; SELECT * FROM history WHERE host_year=2004 AND
event code=20281";
$res = cubrid_execute($conn, $sql_stmt, CUBRID_EXEC_QUERY_ALL);

get_result_info($res);

cubrid_next_result($res);

```



```

get result info($res);

function get_result_info($req)
{
    printf("\n----- get result info ----- \n");

    $row num = cubrid num rows($req);
    $col_num = cubrid_num_cols($req);

    $column name list = cubrid column names($req);
    $column type list = cubrid column types($req);

    $column last name = cubrid field name($req, $col num - 1);
    $column_last_table = cubrid_field_table($req, $col_num - 1);

    $column last type = cubrid field type($req, $col num - 1);
    $column last len = cubrid field len($req, $col num - 1);

    $column_1_flags = cubrid_field_flags($req, 1);

    printf("%-30s %d\n", "Row count:", $row num);
    printf("%-30s %d\n", "Column count:", $col num);
    printf("\n");

    printf("%-30s %-30s %-15s\n", "Column Names", "Column Types", "Column Len");
    printf("-----
\n");
    $size = count($column name list);
    for($i = 0; $i < $size; $i++) {
        $column len = cubrid field len($req, $i);
        printf("%-30s %-30s %-15s\n", $column_name_list[$i], $column_type_list[$i],
    $column_len);
    }
    printf("\n\n");

    printf("%-30s %s\n", "Last Column Name:", $column last name);
    printf("%-30s %s\n", "Last Column Table:", $column_last_table);
    printf("%-30s %s\n", "Last Column Type:", $column_last_type);
    printf("%-30s %d\n", "Last Column Len:", $column last len);
    printf("%-30s %s\n", "Second Column Flags:", $column 1 flags);

    printf("\n\n");
}
?>

```

The above example will output:

```

----- get_result_info -----
Row count:          6
Column count:      2

Column Names          Column Types          Column Len
-----
s_name                char(1)                1
f_name                varchar(6)             6

Last Column Name:    f name
Last Column Table:  code
Last Column Type:   varchar(6)
Last Column Len:    6
Second Column Flags:

----- get result info -----
Row count:          4
Column count:      5

Column Names          Column Types          Column Len
-----

```

event code	integer	11
athlete	varchar(40)	40
host_year	integer	11
score	varchar(10)	10
unit	varchar(5)	5
Last Column Name:	unit	
Last Column Table:	history	
Last Column Type:	varchar(5)	
Last Column Len:	5	
Second Column Flags:	not_null primary_key unique_key	

See Also

- [cubrid_execute](#)

cubrid_num_cols/cubrid_num_fields

Description

The `cubrid_num_cols()` and `cubrid_num_fields()` functions are used interchangeably and they get the number of columns in the query result. This is available only with the **SELECT** statement.

Syntax

```
int cubrid_num_cols (resource $req_identifier)
int cubrid_num_fields (resource $req_identifier)
```

- *req_identifier* : Request identifier

Return Value

- Success : The number of columns
- Error : -1

Example

```
$req = cubrid_execute ($con, "select * from member");
if ($req) {
    $rows_count = cubrid_num_rows ($req);
    $cols_count = cubrid_num_cols ($req);
    echo "result set rows count : $rows\n";
    echo "result set columns count : $cols\n";
    cubrid_close_request ($req);
}
```

See Also

- [cubrid_execute](#)
- [cubrid_num_rows](#)

cubrid_num_rows

Description

The `cubrid_num_rows` function returns the number of rows in the query result. This is available only with the **SELECT** statement. Use [cubrid_affected_rows\(\)](#) if you want to know the results of **INSERT**, **UPDATE** and **DELETE** queries. `cubrid_num_rows()` can be used only with synchronous queries. It returns 0 if the query is asynchronous.

Syntax

```
int cubrid_num_rows (resource $req_identifier)
```

- *req_identifier* : Request identifier

Return Value

- Success : The number of rows
- Asynchronous query : 0
- Failure : -1

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$req = cubrid_execute($conn, "SELECT * FROM code");

$row_num = cubrid_num_rows($req);
$col_num = cubrid_num_cols($req);

printf("Row Num: %d\nColumn Num: %d\n", $row_num, $col_num);

cubrid_disconnect($conn);
?>
```

The above example will output:

```
Row Num: 6
Column Num: 2
```

See Also

- [cubrid_execute](#)
- [cubrid_num_cols](#)
- [cubrid_affected_rows](#)

cubrid_pconnect

Description

The **cubrid_pconnect** function configures the persistent connection to a database server.

The **cubrid_pconnect** function is working similar to the [cubrid_connect](#) function but there are two differences as follows:

- When the **cubrid_pconnect** function is executed, it looks for a persistent link with the same host, port number, database name, and user name and returns an existing connection identifier if it exists.
- The connection made by the **cubrid_pconnect** function is maintained although the [cubrid_close](#) or [cubrid_disconnect](#) function is called.

Syntax

```
resource cubrid_pconnect( string $host, int $port, string $dbname[, string $userid[,
string $passwd]] )
```

- *host* : IP address and host name of the Broker Server
- *port* : Port number of the Broker Server (**BROKER_PORT** defined in **\$CUBRID/conf/cubrid_broker.conf**)
- *dbname* : Database name
- *userid* : Database user name
- *passwd* : Database user password

Return Value

- Success : Connection identifier
- Failure : FALSE

Example

```

<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid version());

printf("\n");

$conn = cubrid_pconnect("localhost", 33000, "demodb");

if (!$conn) { die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg()); }

$db_params = cubrid_get_db_parameter($conn);

while (list($param name, $param value) = each($db params)) { printf("%-30s %s\n",
$param name, $param value); }

printf("\n");

$server_info = cubrid_get_server_info($conn);
$client_info = cubrid_get_client_info();

printf("%-30s %s\n", "Server Info:", $server_info);
printf("%-30s %s\n", "Client Info:", $client_info);

printf("\n");

$charset = cubrid_get_charset($conn);

printf("%-30s %s\n", "CUBRID Charset:", $charset);

cubrid_disconnect($conn);
?>

```

The above example will output:

```

CUBRID PHP Version: 8.4.1.0001

PARAM ISOLATION LEVEL 3
LOCK TIMEOUT 0
MAX STRING LENGTH 1073741823
PARAM AUTO COMMIT 0

```

```

Server Info: 8.4.1.0508
Client Info: 8.4.1
CUBRID Charset: iso8859-1

```

cubrid_pconnect_with_url

Description

The `cubrid_pconnect_with_url` function configures the persistent connection to a database server.

The `cubrid_pconnect_with_url` function is working similar to the [cubrid_connect_with_url](#) function but there are two differences as follows:

- When `cubrid_pconnect_with_url` function is executed, it looks for a persistent link with the same host, port number, database name, and user name and returns an existing identifier if it exists.
- The connection made by the `cubrid_pconnect_with_url` function is maintained although the [cubrid_close](#) or [cubrid_disconnect](#) function is called.

Syntax

```

resource cubrid_pconnect_with_url( string $conn_url[, string $userid[, string $passwd]] )

<conn url> ::= [cc:|]CUBRID:<host>:<db name>:<db user>:<db password>:[?<properties>]
<properties> ::= <property> [&<property>]
<property> ::= autocommit=<autocommit_mode>
<property> ::= althosts=<alternative_hosts> [ &rcptime=<time>]
<property> ::= login_timeout=<milli_sec>
<property> ::= query_timeout=<milli_sec>
<property> ::= disconnect_on_query_timeout=true|false

```

```

<alternative_hosts> ::= <host>:<port>[, <host>:<port>]
<host> := HOSTNAME | IP_ADDR

<time> := SECOND
<milli_sec> := MILLISECOND

```

- *conn_url* : A character string that contains server connection information
 - *host* : IP address and host name of the Broker Server
 - *db_name* : Database name
 - *db_user* : Database user name
 - *db_password* : Database user password
 - **autocommit=true|false** : Whether to configure auto-commit mode upon database connection
 - **althosts=standby_broker1_host, standby_broker2_host, ...** : Specifies the broker information of the standby server, which is used for failover when it is impossible to connect to the active server. You can specify multiple brokers for failover, and the connection to the brokers is attempted in the order listed in **althosts**.
- *host* : A host name or IP address of the master database
- *port* : Port number of a broker server (**BROKER_PORT** defined in **SCUBRID/conf/cubrid_broker.conf**)
- **rctime** : An interval between the attempts to connect to the active broker in which failure occurred. After a failure occurs, the system connects to the broker specified by **althosts** (failover), terminates the transaction, and then attempts to connect to the active broker of the master database at every **rctime**. The default value is **600** seconds.
 - **login_timeout** : Timeout value (unit : msec.) upon login to a database. If time exceeds the value specified in this parameter, a **CCI_ER_LOGIN_TIMEOUT** error is returned. The default value is 0, which means infinite wait.
 - **query_timeout** : Timeout value (unit : msec.). It configures timeout value based on query request. If timeout occurs, a cancellation message of query request, which has been sent to a server, is returned. The return value of a function which executes query is different based on configuration of **disconnect_on_query_timeout** and the request can succeed even though a cancellation message is sent to a server.
 - **disconnect_on_query_timeout** : Configures whether to return an error immediately if timeout on the query request happens. The default value is **false**. If timeout on query request happens and this value is **true**, a **CCI_ER_QUERY_TIMEOUT** error is returned after a cancellation message is sent and a socket is disconnected. In this case, a user should close the connection handle to a database explicitly by using the **cubrid_disconnect** function. If the value is **false**, a cancellation message is sent and it will wait until a response on query request is made.
- *userid* : Database user name
- *passwd* : Database user password

Return Value

- Success : Connection identifier
- Failure : FALSE

Example

```

Example #1 cubrid_pconnect_with_url() url without properties example
<?php
$conn url = "CUBRID:127.0.0.1:33000:demodb:dba:123456:?autocommit=off"
$con = cubrid_pconnect with url ($conn url);

if ($con) {
echo "connected successfully";
$req =cubrid_execute($con, "insert into person values(1,'James')");

if ($req) { cubrid close request ($req); cubrid commit ($con); } else { cubrid rollback
($con); }
cubrid_disconnect ($con);
}
?>

Example #2 cubrid_pconnect with url() url with properties example
<?php

```

```

$conn url =
"CUBRID:127.0.0.1:33000:demodb:dba:123456:?autocommit=off&althost=10.34.63.132:33088&rctime=100"
$con = cubrid_pconnect_with_url ($conn_url);

if ($con) {
echo "connected successfully";
$req =cubrid_execute($con, "insert into person values(1,'James')");

if ($req) { cubrid_close_request ($req); cubrid_commit ($con); } } else { cubrid_rollback ($con); }
cubrid_disconnect ($con);
}
?>

```

cubrid_ping

Description

The **cubrid_ping** function pings a server connection or reconnection if there is no connection.

Syntax

```
bool cubrid_ping ([resource $conn_identifier ])
```

- *conn_identifier* : Connection identifier. If the connection identifier is not specified, the last connection is assumed.

Return Value

- If the connection to the database server is working : TRUE
- Otherwise : FALSE

Example

```

<?php
set time limit(0);

$conn = cubrid_connect('localhost', 33000, 'demodb');

/* Assuming this query will take a long time */
$result = cubrid_query($sql);
if (!$result) {
    echo 'Query #1 failed, exiting.';
    exit;
}

/* Make sure the connection is still alive, if not, try to reconnect */
if (!cubrid_ping($conn)) {
    echo 'Lost connection, exiting after query #1';
    exit;
}
cubrid_free_result($result);

/* So the connection is still alive, let's run another query */
$result2 = cubrid_query($sql2);
?>

```

cubrid_prepare

Description

The **cubrid_prepare** function is an API that represents a pre-compiled SQL statement on the given connection handle. The SQL statement is pre-compiled and then included in **cubrid_prepare()**. This method can be used to efficiently execute the statement multiple times or to effectively process Long Data. You can use only a single statement and a parameter can insert a question mark (?) into appropriate position in the SQL statement. You can also add a parameter

to the position in the **VALUES** clause of the **INSERT** statement or in the **WHERE** clause of the SQL statement, for which the value is to be substituted. Substituting a value for a question mark (?) can be performed only by [cubrid_bind](#).

Syntax

```
resource cubrid_prepare (resource $conn_identifier, string $prepare_stmt [, int $option])
```

- *conn_identifier* : Connection handle
- *prepare_stmt* : A prepare query
- *option* : OID return option - CUBRID_INCLUDE_OID

Return Value

- Success : Request handle
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$sql = <<<EOD
SELECT g.event code, e.name
FROM game g
JOIN event e ON g.event_code=e.code
WHERE host year = ? AND event code NOT IN (SELECT event code FROM game WHERE host year=?)
GROUP BY event code;
EOD;

$req = cubrid_prepare($conn, $sql);

cubrid_bind($req, 1, 2004);
cubrid_bind($req, 2, 2000);
cubrid_execute($req);

$row_num = cubrid_num_rows($req);
printf("There are %d event that exists in 2004 olympic but not in 2000. For example:\n\n",
$row_num);

printf("%-15s %s\n", "Event code", "Event name");
printf("-----\n");

$row = cubrid_fetch_assoc($req);
printf("%-15d %s\n", $row["event code"], $row["name"]);
$row = cubrid_fetch_assoc($req);
printf("%-15d %s\n", $row["event code"], $row["name"]);

cubrid_disconnect($conn);
?>
```

The above example will output:

There are 27 event that exists in 2004 olympic but not in 2000. For example:

Event code	Event name

20063	+91kg
20070	64kg

See Also

- [cubrid_execute](#)
- [cubrid_bind](#)

cubrid_put

Description

The **cubrid_put** function changes attribute values of an instance by using the given OID. You can update single attribute by using string data type to set *attr*. In such case, you can use integer, floating-point, or character string data type for the value argument. To change multiple attributes simultaneously, pass value argument in the form of associative array data type without specifying the *attr* argument.

Syntax

```
int cubrid_put (resource $conn_identifier, string $oid[, string $attr], mixed $value)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the instance whose value you want to change
- *attr* : Name of the attribute whose value you want to change
- *value* : Value of the attribute you want to change

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO_INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(2, {4,5,7}, {44,55,66,666},
'b')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_put($conn, $oid, "b", array(2, 4, 8));

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(3) {
  [0]=>
  string(1) "1"
  [1]=>
  string(1) "2"
  [2]=>
  string(1) "3"
}
array(3) {
  [0]=>
  string(1) "2"
  [1]=>
  string(1) "4"
```



```
[2]=>
string(1) "8"
}
```

See Also

- [cubrid_get](#)
- [cubrid_set_add](#)
- [cubrid_set_drop](#)
- [cubrid_seq_insert](#)
- [cubrid_seq_drop](#)
- [cubrid_seq_put](#)

cubrid_query

Description

The **cubrid_query** function sends a single query to a database server associated with *conn_identifier*. It cannot send multiple queries.

When you call [cubrid_num_rows\(\)](#) to get information on the number of rows returned with the **SELECT** statement or call [cubrid_affected_rows\(\)](#) to get information on the number of rows affected by **DELETE**, **INSERT**, **REPLACE**, or **UPDATE** statement, use a result identifier returned by the **cubrid_query** function.

Syntax

```
resource cubrid_query (string $query[, resource $conn_identifier])
```

- *query* : SQL query statement
- *conn_identifier* : Specifies the CUBRID connection. If it is not specified, the last connection is used.

Return Value

- Success : Request identifier
- No access permission to the table : FALSE
- Failure : FALSE

Example

```
<?php
// This could be supplied by a user, for example
$firstname = 'fred';
$lastname = 'fox';

$conn = cubrid_connect('localhost', 33000, 'foo');

// Formulate Query
// This is the best way to perform an SQL query
// For more examples, see cubrid_real_escape_string()
$query = sprintf("SELECT firstname, lastname, address, age FROM friends WHERE
firstname='%s' AND lastname='%s'",
cubrid_real_escape_string($firstname),
cubrid_real_escape_string($lastname));

// Perform Query
$result = cubrid_query($query);

// Check result
// This shows the actual query sent to CUBRID, and the error. Useful for debugging.
if (!$result) {
    $message = 'Invalid query: ' . cubrid_error() . "\n";
    $message .= 'Whole query: ' . $query;
    die($message);
}

// Use result
```

```
// Attempting to print $result won't allow access to information in the resource
// One of the cubrid result functions must be used
// See also cubrid_result(), cubrid_fetch_array(), cubrid_fetch_row(), etc.
while ($row = cubrid_fetch_assoc($result)) {
    echo $row['firstname'];
    echo $row['lastname'];
    echo $row['address'];
    echo $row['age'];
}

// Free the resources associated with the result set
// This is done automatically at the end of the script
cubrid_free_result($result);
?>
```

See Also

- [cubrid_unbuffered_query](#)

cubrid_real_escape_string

Description

The **cubrid_real_escape_string** function returns the escaped string version of the given string. Follow two escape sequence methods should be supported. On the quoted escape sequence, a string quoted with " or ' can be applied when system parameter *ansi_quotes* is set to "yes". If this option is set to "no", only a string quoted with ' can be applied. The default value is "no".

Quoted escape sequence:

- A ' inside a string quoted with ' may be written as ''
- A " inside a string quoted with " may be written as "" (applied when *ansi_quotes*=yes)
- A ' inside a string quoted with " needs no special treatment and need not be doubled or escaped. (applied when *ansi_quotes*=yes)
- In the same way, " inside a string quoted with ' needs no special treatment.

Backslash escape sequence: This sequence is on by system parameter *no_backslash_escapes*.

The following characters can be escaped by backslash: \, \", \n, \r, \t, \\\, \% _.

If this option is set to "no", backslash escaping will work. The default value is "yes".

Syntax

```
string cubrid_real_escape_string (string $unescape_string [, resource $link_identifier ] )
```

- *unescape_string* : The string that is to be escaped.
- *conn_identifier* : The CUBRID connection. If the link identifier is not specified, the last link opened by [cubrid_connect\(\)](#) is assumed.

Return Value

- Success : Escaped characters
- Failure : FALSE

Example

```
< ?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$unescape_str = ' !"#%&\'()*+,-
./0123456789:;<=>?@ABCDEFGHIJKLMNopqrstuvwxyz[\]^ `abcdefghijklmnopqrstuvwxy{|}~';
$escaped_str = cubrid_real_escape_string($unescape_str);

$len = strlen($unescape_str);

@cubrid_execute($conn, "DROP TABLE cubrid_test");
```

```

cubrid execute($conn, "CREATE TABLE cubrid test (t char($len))");
cubrid execute($conn, "INSERT INTO cubrid test (t) VALUES('$escaped str')");

$req = cubrid_execute($conn, "SELECT * FROM cubrid_test");
$row = cubrid_fetch_assoc($req);

var_dump($row);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

The above example will output:

array(1) {
  ["t"]=>
  string(95) " !"#%&'()*+,-
./0123456789:;<=>?@ABCDEFGHIJKLMNopQRSTUVWXYZ[\]^ `abcdefghijklmnopqrstuvwxy{|}~"
}

```

cubrid_result

Description

The **cubrid_result** function retrieves the contents of one cell from a CUBRID result set on success, or **FALSE** on failure.

When working on large result sets, you should consider using one of the functions that fetch an entire row. As these functions return the contents of multiple cells in one function call, they're MUCH quicker than **cubrid_result()**. Also, note that specifying a numeric offset for the field argument is much quicker than specifying a fieldname or tablename.fieldname argument.

Syntax

```
string cubrid_result ( resource $result , int $row [, mixed $field= 0 ] )
```

- *result* : Result that comes from a call to [cubrid_execute\(\)](#)
- *row* : The row number from the result that's being retrieved. Row numbers start at 0.
- *field* : The name or offset of the field being retrieved. It can be the field's offset, the field's name, or the field's table dot field name (tablename.fieldname). If the column name has been aliased ('select foo as bar from...'), use the alias instead of the column name. If undefined, the first field is retrieved.

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$req = cubrid_execute($conn, "SELECT * FROM code");

$result = cubrid_result($req, 0);
var_dump($result);

$result = cubrid_result($req, 0, 1);
var_dump($result);

$result = cubrid_result($req, 5, "f name");
var_dump($result);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```

string(1) "X"
string(5) "Mixed"
string(4) "Gold"

```

cubrid_rollback

Description

The **cubrid_rollback** function rolls back the transaction being executed in the connection referred by the *conn_identifier*. The connection with the server is terminated after the **cubrid_rollback** method is called, but the connection identifier remains valid.

Syntax

```
bool cubrid_rollback (resource $conn_identifier)
```

- *conn_identifier* : Connection identifier

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("127.0.0.1", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE publishers");

$sql = <<EOD
CREATE TABLE publishers(
pub id CHAR(3),
pub name VARCHAR(20),
city VARCHAR(15),
state CHAR(2),
country VARCHAR(15)
)
EOD;

if (!cubrid_execute($conn, $sql)) {
    printf("Error facility: %d\nError code: %d\nError msg: %s\n",
cubrid_error_code_facility(), cubrid_error_code(), cubrid_error_msg());

    cubrid_disconnect($conn);
    exit;
}

$req = cubrid_prepare($conn, "INSERT INTO publishers VALUES(?, ?, ?, ?, ?)");

$id_list = array("P01", "P02", "P03", "P04");
$name_list = array("Abatis Publishers", "Core Dump Books", "Schadenfreude Press",
"Tenterhooks Press");
$city_list = array("New York", "San Francisco", "Hamburg", "Berkeley");
$state_list = array("NY", "CA", NULL, "CA");
$country_list = array("USA", "USA", "Germany", "USA");

for ($i = 0, $size = count($id_list); $i < $size; $i++) {
    cubrid_bind($req, 1, $id_list[$i]);
    cubrid_bind($req, 2, $name_list[$i]);
    cubrid_bind($req, 3, $city_list[$i]);
    cubrid_bind($req, 4, $state_list[$i]);
    cubrid_bind($req, 5, $country_list[$i]);

    if (!$ret = cubrid_execute($req)) {
        break;
    }
}

if (!$ret) {
    cubrid_rollback($conn);
} else {
    cubrid_commit($conn);
}
```

```

$req = cubrid execute($conn, "SELECT * FROM publishers");
while ($result = cubrid fetch assoc($req)) {
    printf("%-3s %-20s %-15s %-3s %-15s\n",
        $result["pub_id"], $result["pub_name"], $result["city"], $result["state"],
        $result["country"]);
    }
}

cubrid_disconnect($conn);
?>

```

The above example will output:

```

P01 Abatis Publishers      New York      NY  USA
P02 Core Dump Books       San Francisco CA  USA
P03 Schadenfreude Press   Hamburg              Germany
P04 Tenterhooks Press     Berkeley      CA  USA

```

See Also

- [cubrid_commit](#)
- [cubrid_disconnect](#)

cubrid_schema

Description

The **cubrid_schema** function gets specific schema information of a database. You should specify *class_name* to get information related to a specific class, and *attr_name* to get information related to a specific attribute (currently, only used with CUBRID_SCH_ATTR_PRIVILEGE).

Syntax

```
array cubrid_schema (resource $conn_identifier, int $schema_type[, string $class_name[, string $attr_name]])
```

- *conn_identifier* : Connection identifier
- *schema_type* : Type of schema you want to get
- *class_name* : Class from which schema is to be obtained
- *attr_name* : Attribute from which schema is to be obtained

Return Value

- Success : Array in which schema information is contained
- Failure : FALSE

The result of the cubrid_schema() function is returned as a two-dimensional array(column (associative array) * row (numeric array)). The following table shows types of schema and the column structure of the result array to be returned based on the schema type.

Schema	Column Number	Column Name	Value
CUBRID_SCH_CLASS	1	NAME	0 : System class 1 : vclass 2 : class
	2	TYPE	
CUBRID_SCH_VCLASS	1	NAME	1 : vclass
	2	TYPE	
CUBRID_SCH_QUERY_SPEC	1	QUERY_SPEC	
CUBRID_SCH_ATTRIBUTE	1	ATTR_NAME	

	2	DOMAIN	
	3	SCALE	
	4	PRECISION	
	5	INDEXED	1 : indexed
	6	NON NULL	1 : non null
	7	SHARED	1 : shared
	8	UNIQUE	1 : unique
	9	DEFAULT	
	10	ATTR_ORDER	1 : base
	11	CLASS_NAME	
	12	SOURCE_CLASS	
CUBRID_SCH_CLASS_ATTRIBUTE	1	ATTR_NAME	
	2	DOMAIN	
	3	SCALE	
	4	PRECISION	
	5	INDEXED	1 : indexed
	6	NON NULL	1 : non null
	7	SHARED	1 : shared
	8	UNIQUE	1 : unique
	9	DEFAULT	
	10	ATTR_ORDER	1 : base
	11	CLASS_NAME	
	12	SOURCE_CLASS	
CUBRID_SCH_METHOD	1	NAME	
	2	RET_DOMAIN	
	3	ARG_DOMAIN	
CUBRID_SCH_METHOD_FILE	1	METHOD_FILE	
CUBRID_SCH_SUPERCLASS	1	CLASS_NAME	
	2	TYPE	
CUBRID_SCH_SUBCLASS	1	CLASS_NAME	
	2	TYPE	
CUBRID_SCH_CONSTRAINT	1	TYPE	0 : unique 1 : index
	2	NAME	
	3	ATTR_NAME	
CUBRID_SCH_TRIGGER	1	NAME	
	2	STATUS	
	3	EVENT	
	4	TARGET_CLASS	
	5	TARGET_ATTR	

	6	ACTION_TIME	
	7	ACTION	
	8	PRIORITY	
	9	CONDITION_TIME	
	10	CONDITION	
CUBRID_SCH_CLASS_PRIVILEGE	1	CLASS_NAME	
	2	PREVILEGE	
	3	GRANTABLE	
CUBRID_SCH_ATTR_PRIVILEGE	1	ATTR_NAME	
	2	PREVILEGE	
	3	GRANTABLE	
CUBRID_SCH_PRIMARY_KEY	1	ATTR_NAME	
	2	KEY_SEQ	1 : base
	3	KEY_NAME	
	4	KEY_NAME	
CUBRID_SCH_IMPORTED_KEYS	1	PKTABLE_NAME	
	2	PKCOLUMN_NAME	
	3	FKTABLE_NAME	
	4	FKCOLUMN_NAME	
	5	KEY_SEQ	
	6	UPDATE_ACTION	0 : cascade 1 : restrict 2 : no action 3 : set null
	7	DELETE_ACTION	0 : cascade 1 : restrict 2 : no action 3 : set null
	8	FK_NAME	
	9	PK_NAME	
CUBRID_SCH_EXPORTED_KEYS	1	PKTABLE_NAME	
	2	PKCOLUMN_NAME	
	3	FKTABLE_NAME	
	4	FKCOLUMN_NAME	
	5	KEY_SEQ	
	6	UPDATE_ACTION	0 : cascade 1 : restrict 2 : no action 3 : set null
	7	DELETE_ACTION	0 : cascade 1 : restrict 2 : no

			action 3 : set null
	8	FK_NAME	
	9	PK_NAME	
CUBRID_SCH_CROSS_REFERENCE	1	PKTABLE_NAME	
	2	PKCOLUMN_NAME	
	3	FKTABLE_NAME	
	4	FKCOLUMN_NAME	
	5	KEY_SEQ	
	6	UPDATE_ACTION	0 : cascade 1 : restrict 2 : no action 3 : set null
	7	DELETE_ACTION	0 : cascade 1 : restrict 2 : no action 3 : set null
	8	FK_NAME	
	9	PK_NAME	

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

printf("\n--- Primary Key ---\n");
$pk = cubrid_schema($conn, CUBRID_SCH_PRIMARY_KEY, "game");
var_dump($pk);

printf("\n--- Foreign Keys ---\n");
$fk = cubrid_schema($conn, CUBRID_SCH_IMPORTED_KEYS, "game");
var_dump($fk);

printf("\n--- Column Attribute ---\n");
$attr = cubrid_schema($conn, CUBRID_SCH_ATTRIBUTE, "stadium", "area");
var_dump($attr);

cubrid_disconnect($conn);
?>
The above example will output:

--- Primary Key ---
array(3) {
  [0]=>
  array(4) {
    ["CLASS_NAME"]=>
    string(4) "game"
    ["ATTR_NAME"]=>
    string(12) "athlete code"
    ["KEY_SEQ"]=>
    string(1) "3"
    ["KEY_NAME"]=>
    string(41) "pk game host year event code athlete code"
  }
  [1]=>
  array(4) {
    ["CLASS_NAME"]=>
    string(4) "game"
    ["ATTR_NAME"]=>

```



```

    string(10) "event code"
    ["KEY_SEQ"]=>
    string(1) "2"
    ["KEY_NAME"]=>
    string(41) "pk game host year event code athlete code"
  }
  [2]=>
  array(4) {
    ["CLASS_NAME"]=>
    string(4) "game"
    ["ATTR_NAME"]=>
    string(9) "host year"
    ["KEY_SEQ"]=>
    string(1) "1"
    ["KEY_NAME"]=>
    string(41) "pk_game_host_year_event_code_athlete_code"
  }
}

--- Foreign Keys ---
array(2) {
  [0]=>
  array(9) {
    ["PKTABLE_NAME"]=>
    string(7) "athlete"
    ["PKCOLUMN_NAME"]=>
    string(4) "code"
    ["FKTABLE_NAME"]=>
    string(4) "game"
    ["FKCOLUMN_NAME"]=>
    string(12) "athlete code"
    ["KEY_SEQ"]=>
    string(1) "1"
    ["UPDATE_RULE"]=>
    string(1) "1"
    ["DELETE_RULE"]=>
    string(1) "1"
    ["FK_NAME"]=>
    string(20) "fk_game_athlete_code"
    ["PK_NAME"]=>
    string(15) "pk athlete code"
  }
  [1]=>
  array(9) {
    ["PKTABLE_NAME"]=>
    string(5) "event"
    ["PKCOLUMN_NAME"]=>
    string(4) "code"
    ["FKTABLE_NAME"]=>
    string(4) "game"
    ["FKCOLUMN_NAME"]=>
    string(10) "event code"
    ["KEY_SEQ"]=>
    string(1) "1"
    ["UPDATE_RULE"]=>
    string(1) "1"
    ["DELETE_RULE"]=>
    string(1) "1"
    ["FK_NAME"]=>
    string(18) "fk game event code"
    ["PK_NAME"]=>
    string(13) "pk_event_code"
  }
}

--- Column Attribute ---
array(1) {
  [0]=>
  array(13) {
    ["ATTR_NAME"]=>
    string(4) "area"
    ["DOMAIN"]=>
    string(1) "7"
  }
}

```

```

["SCALE"]=>
string(1) "2"
["PRECISION"]=>
string(2) "10"
["INDEXED"]=>
string(1) "0"
["NON NULL"]=>
string(1) "0"
["SHARED"]=>
string(1) "0"
["UNIQUE"]=>
string(1) "0"
["DEFAULT"]=>
NULL
["ATTR_ORDER"]=>
string(1) "4"
["CLASS NAME"]=>
string(7) "stadium"
["SOURCE CLASS"]=>
string(7) "stadium"
["IS_KEY"]=>
string(1) "0"
}

```

cubrid_seq_drop

Description

The **cubrid_seq_drop** function drops elements from the given **SEQUENCE** type attribute in the database.

Syntax

```
bool cubrid_seq_drop(resource $conn_identifier, string $oid, string $attr_name, int $index)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the desired attribute of the instance
- *index* : Index of the element to be dropped. The default value is 1.

Return Value

- Success : TRUE
- Failure : FALSE

Example

```

<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid execute($conn, "DROP TABLE foo");
cubrid execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c sequence(int),
d char(10))");
cubrid execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_seq_drop($conn, $oid, "c", 4);

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_close_request($req);

```

```
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(4) {
  [0]=>
  string(2) "11"
  [1]=>
  string(2) "22"
  [2]=>
  string(2) "33"
  [3]=>
  string(3) "333"
}
array(3) {
  [0]=>
  string(2) "11"
  [1]=>
  string(2) "22"
  [2]=>
  string(2) "33"
}
```

See Also

- [cubrid_seq_insert](#)
- [cubrid_seq_put](#)

cubrid_seq_insert

Description

The **cubrid_seq_insert** function inserts an element into a specific position of a **SEQUENCE** type attribute.

Syntax

```
bool cubrid_seq_insert (resource $conn_identifier, string $oid, string $attr_name, int $index, string $seq_element)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the desired attribute of the instance
- *index* : Position into which the new element is to be inserted (default value: 1)
- *seq_string* : Content of the element to be inserted

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c sequence(int), d char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333}, 'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);
```

```

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_seq_insert($conn, $oid, "c", 5, "44");

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>

```

The above example will output:

```

array(4) {
  [0]=>
    string(2) "11"
  [1]=>
    string(2) "22"
  [2]=>
    string(2) "33"
  [3]=>
    string(3) "333"
}
array(5) {
  [0]=>
    string(2) "11"
  [1]=>
    string(2) "22"
  [2]=>
    string(2) "33"
  [3]=>
    string(3) "333"
  [4]=>
    string(2) "44"
}

```

See Also

- [cubrid_seq_drop](#)
- [cubrid_seq_put](#)

cubrid_seq_put

Description

The **cubrid_seq_put** function changes the content of an element of the given SEQUENCE type attribute.

Syntax

```
bool cubrid_seq_put (resource $conn_identifier, string $oid, string $attr_name, int index, string $seq_element)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the desired attribute of the instance
- *index* : Index of the element to be changed (default value: 1)
- *seq_element* : Content of the element to be changed

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO_INCREMENT, b set(int), c sequence(int),
d char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_seq_put($conn, $oid, "c", 1, "111");

$attr = cubrid_col_get($conn, $oid, "c");
var_dump($attr);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(4) {
  [0]=>
  string(2) "11"
  [1]=>
  string(2) "22"
  [2]=>
  string(2) "33"
  [3]=>
  string(3) "333"
}
array(4) {
  [0]=>
  string(3) "111"
  [1]=>
  string(2) "22"
  [2]=>
  string(2) "33"
  [3]=>
  string(3) "333"
}
```

See Also

- [cubrid_seq_insert](#)
- [cubrid_seq_drop](#)

cubrid_set_add

Description

The **cubrid_set_add** function inserts an element to the given SET type (set, multiset) attribute.

Syntax

```
bool cubrid_set_add (resource $conn_identifier, string $oid, string $attr_name, string
$set_element)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance

- *attr_name* : Name of the desired attribute of the instance
- *set_string* : Content of the element to be inserted

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

@cubrid execute($conn, "DROP TABLE foo");
cubrid execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid_move_cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid_current_oid($req);

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_set_add($conn, $oid, "b", "4");

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid_close_request($req);
cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(3) {
  [0]=>
  string(1) "1"
  [1]=>
  string(1) "2"
  [2]=>
  string(1) "3"
}
array(4) {
  [0]=>
  string(1) "1"
  [1]=>
  string(1) "2"
  [2]=>
  string(1) "3"
  [3]=>
  string(1) "4"
}
```

See Also

- [cubrid_set_drop](#)

cubrid_set_autocommit

Description

The **cubrid_set_autocommit** function sets the status of CUBRID database connection auto-commit mode of the current database connection. This function just turns on/off the auto-commit mode. When this function is called, concurrent transactions are committed regardless of the auto-commit mode.

CCI_DEFAULT_AUTOCOMMIT, broker parameter configured in the **cubrid_broker.conf** file, determines whether it is in auto-commit mode upon program startup.

Syntax

```
bool cubrid_set_autocommit(resource $conn_identifier, int $mode)
```

- *conn_identifier* : Connection identifier
- *mode* : Whether to turn on auto-commit or not. It should be **CUBRID_AUTOCOMMIT_FALSE** or **CUBRID_AUTOCOMMIT_TRUE**.

Return Value

- Success : TRUE
- Failure : FALSE

See Also

- [cubrid_get_autocommit](#)

cubrid_set_db_parameter

Description

The **cubrid_set_db_parameter** function sets the CUBRID system parameters. It can set the following CUBRID system parameters:

- **CUBRID_PARAM_ISOLATION_LEVEL** : Transaction isolation level. For more information, see "CUBRID SQL Guide > Transaction and Lock > Transaction Isolation Level > SET TRANSACTION ISOLATION LEVEL."
- **CUBRID_PARAM_LOCK_TIMEOUT** : Lock timeout. Time when transaction lock is held.

Syntax

```
bool cubrid_set_db_parameter ( resource $conn_identifier, int $param_type, int $param_value)
```

- *conn_identifier* : Connection identifier
- *param_type* : System parameter type
- *param_value* : System parameter value

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
<?php
$conn = cubrid_connect("localhost", 33000, "demodb");

$params = cubrid_get_db_parameter($conn);
var_dump($params);

cubrid_set_autocommit($conn, CUBRID_AUTOCOMMIT_TRUE);
cubrid_set_db_parameter($conn, CUBRID_PARAM_ISOLATION_LEVEL, 2);

$params_new = cubrid_get_db_parameter($conn);
var_dump($params_new);

cubrid_disconnect($conn);
?>
```

The above example will output:

```
array(4) {
  ["PARAM_ISOLATION_LEVEL"]=>
```

```

int (3)
["PARAM_LOCK_TIMEOUT"]=>
int (-1)
["PARAM_MAX_STRING_LENGTH"]=>
int (1073741823)
["PARAM_AUTO_COMMIT"]=>
int (0)
}
array (4) {
["PARAM_ISOLATION_LEVEL"]=>
int (2)
["PARAM_LOCK_TIMEOUT"]=>
int (-1)
["PARAM_MAX_STRING_LENGTH"]=>
int (1073741823)
["PARAM_AUTO_COMMIT"]=>
int (1)
}

```

See Also

- [cubrid_get_db_parameter](#)

cubrid_set_drop

Description

The `cubrid_set_drop` function drops an element from the given SET type (set, multiset) attribute.

Syntax

```
bool cubrid_set_drop (resource $conn_identifier, string $oid, string $attr_name, string $set_element)
```

- *conn_identifier* : Connection identifier
- *oid* : OID of the desired instance
- *attr_name* : Name of the desired attribute of the instance
- *set_element* : Content of the element to be dropped.

Return Value

- Success : TRUE
- Failure : FALSE

Example

```

<?php
$conn = cubrid connect("localhost", 33000, "demodb");

@cubrid_execute($conn, "DROP TABLE foo");
cubrid_execute($conn, "CREATE TABLE foo(a int AUTO INCREMENT, b set(int), c list(int), d
char(10))");
cubrid_execute($conn, "INSERT INTO foo(a, b, c, d) VALUES(1, {1,2,3}, {11,22,33,333},
'a')");

$req = cubrid_execute($conn, "SELECT * FROM foo", CUBRID_INCLUDE_OID);

cubrid move cursor($req, 1, CUBRID_CURSOR_FIRST);
$oid = cubrid current oid($req);

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

cubrid set drop($conn, $oid, "b", "1");

$attr = cubrid_col_get($conn, $oid, "b");
var_dump($attr);

```



```
cubrid close request($req);
cubrid disconnect($conn);
?>
```

The above example will output:

```
array(3) {
  [0]=>
  string(1) "1"
  [1]=>
  string(1) "2"
  [2]=>
  string(1) "3"
}
array(2) {
  [0]=>
  string(1) "2"
  [1]=>
  string(1) "3"
}
```

See Also

- [cubrid_set_add](#)

cubrid_set_query_timeout

Description

The **cubrid_set_query_timeout** function configures timeout value for query execution.

The timeout value configured by **cci_set_query_timeout** affects the [cubrid_prepare](#) and [cubrid_execute](#) function. When timeout occurs in the function and if the **disconnect_on_query_timeout** value configured in [cubrid_connect_with_url](#) connection URL is yes, it returns a **CUBRID_ER_QUERY_TIMEOUT** error.

The **cubrid_prepare** and **cubrid_execute** functions return an error in case that **login_timeout** is configured in the connection URL, which is an argument of the **cubrid_connect_with_url** function; this means that login timeout happens between application client and server (CAS) during re-connection.

It is going through the process of re-connection between application client and server (CAS) when an application restarts or it is re-scheduled. Re-scheduling is a process that an application server chooses an application client, and starts and stops connection in the unit of transaction. If **KEEP_CONNECTION**, Broker parameter, is OFF, it always happens; if AUTO, it can happen depending on its situation. For details, see **KEEP_CONNECTION** description in the "Performance Tuning > Broker Configuration > Parameter by Broker" on the manual.

Syntax

```
bool cubrid_set_query_timeout(resource $conn_identifier, int $timeout);
```

- *conn_identifier* : Connection identifier
- *timeout* : Timeout(timeout); unit is millisecond (msec).

Return Value

- Success : Configured value before changes
- Failure : FALSE

cubrid_unbuffered_query

Description

The **cubrid_unbuffered_query** function sends a specified single query to the server without fetching the result row or buffering as [cubrid_execute\(\)](#) does. This keeps considerable amount of memory occupied by SQL statements that

produce large result sets and enables work on the result set immediately after the first row has been retrieved; in other words, it does not need to wait until execution of SQL statement is complete.

The optional argument *link_identifier* must be specified in multiple connection environment.

The benefits `cubrid_unbuffered_query()` come at a cost; however, you cannot use `cubrid_num_rows()` and `cubrid_data_seek()` on a result set returned from `cubrid_unbuffered_query()`.

Syntax

```
resource cubrid_unbuffered_query ( string $query [, int $conn_identifier ] )
```

- *query* : A SQL query
- *link_identifier* : The CUBRID connection. If the connection identifier is not specified, the last link opened by `cubrid_connect()` is used.

Return Value

- Success : TRUE
- Failure : FALSE

Example

```
< ?php
$result = cubrid_unbuffered_query("INSERT INTO mytable (product) values ('kossu')", $link);
if (!$result) {
    echo 'Could not run query: ' . cubrid_error_msg();
    exit;
}
printf("Last inserted record has id %d\n", cubrid_insert_id());
?>
```

cubrid_version

Description

The `cubrid_version` function checks the version information of the CUBRID PHP module.

Syntax

```
string cubrid_version ()
```

Return Value

- n version information (e.g. "1.2.0")

Example

```
<?php
printf("%-30s %s\n", "CUBRID PHP Version:", cubrid_version());

printf("\n");

$conn = cubrid_connect("localhost", 33088, "demodb");

if (!$conn) {
    die('Connect Error (' . cubrid_error_code() . ') ' . cubrid_error_msg());
}

$db_params = cubrid_get_db_parameter($conn);

while (list($param_name, $param_value) = each($db_params)) {
    printf("%-30s %s\n", $param_name, $param_value);
}

printf("\n");

$server_info = cubrid_get_server_info($conn);
```

```
$client_info = cubrid_get_client_info();  
  
printf("%-30s %s\n", "Server Info:", $server_info);  
printf("%-30s %s\n", "Client Info:", $client_info);  
  
printf("\n");  
  
$charset = cubrid_get_charset($conn);  
  
printf("%-30s %s\n", "CUBRID Charset:", $charset);  
  
cubrid_disconnect($conn);  
  
?>
```

The above example will output:

```
CUBRID PHP Version:      8.3.1.0005  
  
PARAM_ISOLATION_LEVEL   3  
LOCK_TIMEOUT            -1  
MAX_STRING_LENGTH       1073741823  
PARAM_AUTO_COMMIT       0  
  
Server Info:            8.3.1.0173  
Client Info:            8.3.1  
  
CUBRID Charset:        iso8859-1
```

See Also

- [cubrid_error_code](#)
- [cubrid_error_code_facility](#)

CCI API

CCI Overview

Overview

The CCI (C Client Interface) is an interface that exists between the CUBRID broker and the application client, through which a C-based application client can access the CUBRID database server using a broker. This interface is also used as an infrastructure for making tools that utilize CAS (e.g. PHP and ODBC). The CUBRID broker delivers the query received from an application client to the broker, and transfers the execution result to the client.

A header file and library files are required to use CCI.

- Header file : `$CUBRID/include/cas_cci.h`
- Library file
- `$CUBRID/lib/libcascci.so` (Windows : `cascci.dll`)
- `$CUBRID/lib/libcascci.a` (Windows : `cascci.lib`)

Writing CCI Application Program

The basic steps used for writing programs are as follows, and a step for binding the data to a variable is added to use the prepared statement. The steps are implemented in example codes 1 and 2.

You can configure the default value of auto-commit mode by using `CCI_DEFAULT_AUTOCOMMIT` (broker parameter) upon startup of an application. If configuration on broker parameter is omitted, the default value is **ON**; use the `cci_set_autocommit()` function to change auto-commit mode within an application. If auto-commit mode is **OFF**, you must explicitly commit or roll back transaction by using the `cci_end_tran()` function.

- Opening a database connection handle (related functions : [cci_connect\(\)](#), [cci_connect_with_url\(\)](#))
- Preparing an SQL statement (related function : [cci_prepare\(\)](#))
- Binding data to a prepared SQL statement (related function : [cci_bind_param\(\)](#))
- Executing a prepared SQL statement (related function : [cci_execute\(\)](#))
- Processing the execution result (related functions : [cci_cursor\(\)](#), [cci_fetch\(\)](#), [cci_get_data\(\)](#), [cci_get_result_info\(\)](#))
- Closing the request handle (related function : [cci_close_req_handle\(\)](#))
- Closing a database connection handle (related function : [cci_disconnect\(\)](#))
- Using database connection pool (related functions : [cci_property_create\(\)](#), [cci_property_destroy\(\)](#), [cci_property_set\(\)](#), [cci_datasource_create\(\)](#), [cci_datasource_destroy\(\)](#), [cci_datasource_borrow\(\)](#), [cci_datasource_release\(\)](#))

Example 1

```
//Example to execute a simple query
#include <stdio.h>
#include "cas_cci.h"
#define BUFSIZE (1024)

int
main (void)
{
    int con = 0, req = 0, col count = 0, i, ind;
    int error;
    char *data;
    T_CCI_ERROR cci_error;
    T_CCI_COL_INFO *col_info;
    T_CCI_SQLX_CMD cmd type;
    char *query = "select * from code";

    //getting a connection handle for a connection with a server
    con = cci_connect ("localhost", 33000, "demodb", "dba", "");
    if (con < 0)
```

```

    {
        printf ("cannot connect to database\n");
        return 1;
    }

//preparing the SQL statement
req = cci_prepare (con, query, 0, &cci_error);
if (req < 0)
    {
        printf ("prepare error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle_error;
    }

//getting column information when the prepared statement is the SELECT query
col_info = cci_get_result_info (req, &cmd_type, &col_count);
if (col_info == NULL)
    {
        printf ("get result info error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle_error;
    }

//Executing the prepared SQL statement
error = cci_execute (req, 0, 0, &cci_error);
if (error < 0)
    {
        printf ("execute error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle_error;
    }
while (1)
    {

//Moving the cursor to access a specific tuple of results
error = cci_cursor (req, 1, CCI_CURSOR_CURRENT, &cci_error);
if (error == CCI_ER_NO_MORE_DATA)
    {
        break;
    }
if (error < 0)
    {
        printf ("cursor error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle_error;
    }

//Fetching the query result into a client buffer
error = cci_fetch (req, &cci_error);
if (error < 0)
    {
        printf ("fetch error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle_error;
    }
for (i = 1; i <= col_count; i++)
    {

//Getting data from the fetched result
error = cci_get_data (req, i, CCI_A_TYPE_STR, &data, &ind);
if (error < 0)
    {
        printf ("get data error: %d, %d\n", error, i);
        goto handle_error;
    }
    printf ("%s\t|", data);
    }
    printf ("\n");
    }

//Closing the request handle
error = cci_close_req_handle (req);
if (error < 0)

```

```

    {
        printf ("close req handle error: %d, %s\n", cci_error.err code,
                cci_error.err_msg);
        goto handle_error;
    }

//Disconnecting with the server
error = cci disconnect (con, &cci_error);
if (error < 0)
    {
        printf ("error: %d, %s\n", cci_error.err code, cci_error.err msg);
        goto handle_error;
    }

return 0;

handle error:
if (req > 0)
    cci close req handle (req);
if (con > 0)
    cci_disconnect (con, &cci_error);

return 1;
}

```

Example 2

```

//Example to execute a query with a bind variable

char *query = "select * from nation where name = ?";
char namebuf[128];

//getting a connection handle for a connection with a server
con = cci connect ("localhost", 33000, "demodb", "dba", "");
if (con < 0)
    {
        printf ("cannot connect to database ");
        return 1;
    }

//preparing the SQL statement
req = cci_prepare (con, query, 0, &cci_error);
if (req < 0)
    {
        printf ("prepare error: %d, %s ", cci_error.err code,
                cci_error.err_msg);
        goto handle_error;
    }

//Binding date into a value
strcpy (namebuf, "Korea");
error =
    cci_bind_param (req, 1, CCI_A_TYPE_STR, &namebuf, CCI_U_TYPE_STRING,
                    CCI_BIND_PTR);

if (error < 0)
    {
        printf ("bind param error: %d ", error);
        goto handle_error;
    }
}

```

Example 3

```

#include <stdio.h>
#include "cas cci.h"

//Example to use connection/statement pool in CCI
int main ()
{
    T_CCI_PROPERTIES *ps = NULL;
    T_CCI_DATASOURCE *ds = NULL;
    T_CCI_ERROR err;
    T_CCI_CONN cons[20];
    int rc = 1, i;
}

```

```

ps = cci_property_create ();
if (ps == NULL)
{
    fprintf (stderr, "Could not create T CCI PROPERTIES.\n");
    rc = 0;
    goto cci_pool_end;
}

cci_property_set (ps, "user", "dba");
cci_property_set (ps, "url", "cci:cubrid:localhost:33000:demodb::");
cci_property_set (ps, "pool size", "10");
cci_property_set (ps, "max wait", "1200");
cci_property_set (ps, "pool prepared statement", "true");
cci_property_set (ps, "default_autocommit", "false");
cci_property_set (ps, "default_isolation", "TRAN_REP_CLASS_UNCOMMIT_INSTANCE");
cci_property_set (ps, "default lock timeout", "10");
cci_property_set (ps, "login timeout", "300000");
cci_property_set (ps, "query timeout", "3000");

ds = cci_datasource_create (ps, &err);
if (ds == NULL)
{
    fprintf (stderr, "Could not create T CCI DATASOURCE.\n");
    fprintf (stderr, "E[%d,%s]\n", err.err code, err.err msg);
    rc = 0;
    goto cci_pool_end;
}

for (i = 0; i < 3; i++)
{
    cons[i] = cci_datasource_borrow (ds, &err);
    if (cons[i] < 0)
    {
        fprintf (stderr,
            "Could not borrow a connection from the data source.\n");
        fprintf (stderr, "E[%d,%s]\n", err.err code, err.err msg);
        continue;
    }
    // put working code here.
    cci_work (cons[i]);
}

sleep (1);

for (i = 0; i < 3; i++)
{
    if (cons[i] < 0)
    {
        continue;
    }
    cci_datasource_release (ds, cons[i], &err);
}
cci_pool_end:
cci_property_destroy (ps);
cci_datasource_destroy (ds);

return 0;
}

// working code
int cci_work (T_CCI_CONN con)
{
    T_CCI_ERROR err;
    char sql[4096];
    int req, res, error, ind;
    int data;

    cci_set_autocommit (con, CCI_AUTOCOMMIT_TRUE);
    cci_set_lock_timeout (con, 100, &err);
    cci_set_isolation_level (con, TRAN_REP_CLASS_COMMIT_INSTANCE, &err);

    error = 0;

```

```

snprintf (sql, 4096, "SELECT host year FROM record WHERE athlete code=11744");
req = cci prepare (con, sql, 0, &err);
if (req < 0)
{
    printf ("prepare error: %d, %s\n", err.err code, err.err msg);
    return error;
}

res = cci_execute (req, 0, 0, &err);
if (res < 0)
{
    printf ("execute error: %d, %s\n", err.err code, err.err msg);
    goto cci work end;
}

while (1)
{
    error = cci cursor (req, 1, CCI_CURSOR_CURRENT, &err);
    if (error == CCI_ER_NO_MORE_DATA)
    {
        break;
    }
    if (error < 0)
    {
        printf ("cursor error: %d, %s\n", err.err code, err.err msg);
        goto cci_work_end;
    }

    error = cci fetch (req, &err);
    if (error < 0)
    {
        printf ("fetch error: %d, %s\n", err.err_code, err.err_msg);
        goto cci_work_end;
    }

    error = cci get data (req, 1, CCI_A_TYPE_INT, &data, &ind);
    if (error < 0)
    {
        printf ("get data error: %d\n", error);
        goto cci work end;
    }
    printf ("%d\n", data);
}

error = 1;
cci work end:
    cci close req handle (req);
    return error;
}

```

Using BLOB/CLOB with CCI

Storing LOB Data

You can create LOB data file and bind the data by using the following functions in CCI applications.

- Creating LOB data file (related function : [cci_blob_new\(\)](#), [cci_blob_write\(\)](#))
- Binding LOB data (related function : [cci_bind_param\(\)](#))
- Freeing memory of LOB structure (related function : [cci_blob_free\(\)](#))

Example 1

```

int con = 0; /* connection handle */
int req = 0; /* request handle */
int res;
int n executed;
int i;
T_CCI_ERROR error;
T_CCI_BLOB blob = NULL;
char data[1024] = "bulabula";

```



```

con = cci connect ("localhost", 33000, "tdb", "PUBLIC", "");
if (con < 0) {
goto handle_error;
}
req = cci prepare (con, "insert into doc (doc id, content) values (?,?)", 0, &error);
if (req < 0)
{
goto handle error;
}

res = cci bind param (req, 1 /* binding index*/, CCI A TYPE STR, "doc-10", &ind,
CCI U TYPE STRING);

/* Creating an empty LOB data file
res = cci_blob_new (con, &blob, &error);
res = cci_blob_write (con, blob, 0 /* start position */, 1024 /* length */, data, &error);

/* Binding BLOB data */
res = cci_bind_param (req, 2 /* binding index*/, CCI A TYPE BLOB, (void *)blob,
CCI_U_TYPE_BLOB, CCI_BIND_PTR);

n executed = cci execute (req, 0, 0, &error);
if (n executed < 0)
{
goto handle error
}

/* Memory free */
cci_blob_free(blob);
return 0;

handle_error:
if (blob != NULL)
{
cci_blob_free(blob);
}
if (req > 0)
{
cci_close_req_handle (req);
}
if (con > 0)
{
cci_disconnect(con, &error);
}
return -1;

```

Getting LOB Data

Description

You can get LOB data by using the following functions in CCI applications. Note that if you enter data in LOB type column, the actual LOB data is stored externally and Locator value referring to the file is stored in LOB type column itself. Therefore, you must call the `cci_blob_read()` function (not the `cci_get_data()` function) to get LOB data stored in the file.

- Getting LOB type column meta data (Locator) (related function : [cci_get_data\(\)](#))
- Getting LOB data (related function : [cci_blob_read\(\)](#))
- Freeing memory of LOB structure (related function : [cci_blob_free\(\)](#))

Example

```

int con = 0; /* connection handle */
int req = 0; /* request handle */
int ind; /* NULL indicator, 0 if not NULL, -1 if NULL*/
int res;
int i;
T CCI ERROR error;
T CCI BLOB blob;
char buffer[1024];

```

```

con = cci connect ("localhost", 33000, "image db", "PUBLIC", "");
if (con < 0)
{
    goto handle_error;
}
req = cci prepare (con, "select content from doc t", 0 /*flag*/, &error);
if (req < 0)
{
    goto handle_error;
}
res = cci execute (req, 0/*flag*/, 0/*max col size*/, &error);
res = cci fetch size (req, 100 /* fetch size */);

while (1) {
    res = cci_cursor (req, 1/* offset */, CCI_CURSOR_CURRENT/* cursor position */, &error);
    if (res == CCI_ER_NO_MORE_DATA)
    {
        break;
    }
    res = cci_fetch (req, &error);

    /* Fetching CLOB Locator */
    res = cci_get_data (req, 1 /* colume index */, CCI_A_TYPE_BLOB,
        (void *)&blob /* BLOB handle */, &ind /* NULL indicator */);

    /* Fetching CLOB data */
    res = cci_blob_read (con, blob, 0 /* start position */, 1024 /* length */, buffer, &error);
    printf ("content = %s\n", buffer);
}
/* Memory free */
cci_blob_free(blob);
res=cci_close_req_handle(req);
res = cci_disconnect (con, &error);
return 0;

handle error:
if (req > 0)
{
    cci_close_req_handle (req);
}
if (con > 0)
{
    cci_disconnect(con, &error);
}
return -1;

```

CCI Error Code and Message

The following table shows the error codes and their messages of CCI.

Error Code	Error Message	Note
CCI_ER_ALLOC_CON_HANDLE	"Cannot allocate connection handle"	
CCI_ER_ATYPE	"Invalid T_CCI_A_TYPE value"	
CCI_ER_BIND_ARRAY_SIZE	"Array binding size is not specified"	
CCI_ER_BIND_INDEX	"Parameter index is out of range"	Index that binds data is not valid.
CCI_ER_COLUMN_INDEX	"Column index is out of range"	
CCI_ER_COMMUNICATION	"Cannot communicate with server"	
CCI_ER_CON_HANDLE	"Invalid connection handle"	

CCI_ER_CONNECT	"Cannot connect to CUBRID CAS"	Fails to connect the CAS when trying connection to the server.
CCI_ER_DELETED_TUPLE	"Current row was deleted"	
CCI_ER_FILE	"Cannot open file"	Fails to open/read/write a file.
CCI_ER_HOSTNAME	"Unknown host name"	
CCI_ER_INVALID_CURSOR_POS	"Invalid cursor position"	
CCI_ER_INVALID_URL	"Invalid url string"	
CCI_ER_ISOLATION_LEVEL	"Unknown transaction isolation level"	
CCI_ER_NO_MORE_DATA	"Invalid cursor position"	
CCI_ER_NO_MORE_MEMORY	"Memory allocation error"	Insufficient memory
CCI_ER_OBJECT	"Invalid oid string"	
CCI_ER_OID_CMD	"Invalid T_CCI_OID_CMD value"	
CCI_ER_TRAN_TYPE	"Unknown transaction type"	
CCI_ER_PARAM_NAME	"Invalid T_CCI_DB_PARAM value"	
CCI_ER_REQ_HANDLE	"Cannot allocate request handle"	
CCI_ER_SAVEPOINT_CMD	"Invalid T_CCI_SAVEPOINT_CMD value"	Invalid T_CCI_SAVEPOINT_CMD value is used as an argument of cci_savepoint() function.
CCI_ER_SET_INDEX	"Invalid set index"	Invalid index is specified when an set element in the T_SET is retrieved.
CCI_ER_STRING_PARAM	"Invalid string argument"	string parameter is NULL or an empty string.
CCI_ER_THREAD_RUNNING	"Thread is running"	The thread is still executed when cci_execute() is executed with CCI_EXEC_THREAD flagged and check the result of thread execution through cci_get_thread_result().
CCI_ER_TRAN_TYPE	"Unknown transaction type"	Connection to the server has succeeded, connection to a database fails.
CCI_ER_TYPE_CONVERSION	"Type conversion error"	Cannot convert the given value into an actual data type.
CCI_ER_DBMS CAS_ER_DBMS	"CUBRID DBMS Error"	Fails to database connection.
CAS_ER_COLLECTION_DOMAIN	"Heterogeneous set is not supported"	Not supported set type.
CAS_ER_COMMUNICATION	"Cannot receive data from client"	
CAS_ER_DB_VALUE	"Cannot make DB_VALUE"	

CAS_ER_DBSERVER_DISCONNECTED	"Cannot communicate with DB Server"	
CAS_ER_FREE_SERVER	"Cannot process the request. Cannot assign CAS. Try again later"	
CAS_ER_INVALID_CALL_STMT	"Illegal CALL statement"	
CAS_ER_NO_MORE_DATA	"Invalid cursor position"	
CAS_ER_NO_MORE_MEMORY	"Memory allocation error"	
CAS_ER_NO_MORE_RESULT_SET	"No More Result"	
CAS_ER_NOT_AUTHORIZED_CLIENT	"Authorization error"	Access is denied.
CAS_ER_NOT_COLLECTION	"The attribute domain must be the set type"	No set type.
CAS_ER_NUM_BIND	"Invalid parameter binding value argument"	The number of data to be bound is not matched with the number of delivered data.
CAS_ER_OBJECT	"Invalid oid"	
CAS_ER_OPEN_FILE	"Cannot open file"	
CAS_ER_PARAM_NAME	"Invalid T_CCI_DB_PARAM value"	Invalid get_db_parameter and , set_db_parameter parameter name.
CAS_ER_QUERY_CANCEL	"Cannot cancel the query"	
CAS_ER_UNKNOWN_U_TYPE	"Invalid T_CCI_U_TYPE value"	
CAS_ER_TYPE_CONVERSION	"Type conversion error"	
CAS_ER_SCHEMA_TYPE	"Invalid T_CCI_SCH_TYPE value"	
CAS_ER_STMT_POOLING	"Invalid plan"	
CAS_ER_TRAN_TYPE	"Invalid transaction type argument"	
CAS_ER_TYPE_CONVERSION	"Type conversion error"	
CAS_ER_UNKNOWN_U_TYPE	"Invalid T_CCI_U_TYPE value"	
CAS_ER_VERSION	"Version mismatch"	Invalid Server and Client version.

C Type Definition

Name	Type Member	Description
T_CCI_ERROR	struct char err_msg[1024] int err_code	Representation of database error info
T_CCI_BIT	struct int size char *buf	Representation of bit type
T_CCI_DATE	struct short yr short mon short day short hh	Representation of timestamp, date, time type

		short mm	
		short ss	
		short ms	
T_CCI_SET	void*		Representation of set type
T_CCI_COL_INFO	struct	T_CCI_U_TYPE type char is_non_null short scale int precision char *col_name char *real_attr char *class_name	Representation of column information for the SELECT statement
T_CCI_QUERY_RESULT	struct	int result_count int stmt_type char *err_msg char oid[32]	Results of batch execution
T_CCI_PARAM_INFO	struct	T_CCI_PARAM_MODE mode T_CCI_U_TYPE type short scale int precision	Representation of input parameter info
T_CCI_U_TYPE	enum	CCI_U_TYPE_UNKNOWN CCI_U_TYPE_NULL CCI_U_TYPE_CHAR CCI_U_TYPE_STRING CCI_U_TYPE_NCHAR CCI_U_TYPE_VARNCHAR CCI_U_TYPE_BIT CCI_U_TYPE_VARBIT CCI_U_TYPE_NUMERIC CCI_U_TYPE_INT CCI_U_TYPE_SHORT CCI_U_TYPE_MONETARY CCI_U_TYPE_FLOAT CCI_U_TYPE_DOUBLE CCI_U_TYPE_DATE CCI_U_TYPE_TIME CCI_U_TYPE_TIMESTAMP CCI_U_TYPE_SET CCI_U_TYPE_MULTISSET CCI_U_TYPE_SEQUENCE	Database type info

		CCI_U_TYPE_OBJECT	
		CCI_U_TYPE_BIGINT	
		CCI_U_TYPE_DATETIME	
T_CCI_A_TYPE	enum	CCI_A_TYPE_STR	Representation of type info used in API
		CCI_A_TYPE_INT	
		CCI_A_TYPE_FLOAT	
		CCI_A_TYPE_DOUBLE	
		CCI_A_TYPE_BIT	
		CCI_A_TYPE_DATE	
		CCI_A_TYPE_SET	
		CCI_A_TYPE_BIGINT	
		CCI_TYPE_BLOB	
		CCI_TYPE_CLOB	
T_CCI_DB_PARAM	enum	CCI_PARAM_ISOLATION_LEVEL	System parameter name
		CCI_PARAM_LOCK_TIMEOUT	
		CCI_PARAM_MAX_STRING_LENGTH	
		CCI_PARAM_AUTO_COMMIT	
T_CCI_SCH_TYPE	enum	CCI_SCH_CLASS	
		CCI_SCH_VCLASS	
		CCI_SCH_QUERY_SPEC	
		CCI_SCH_ATTRIBUTE	
		CCI_SCH_CLASS_ATTRIBUTE	
		CCI_SCH_METHOD	
		CCI_SCH_CLASS_METHOD	
		CCI_SCH_METHOD_FILE	
		CCI_SCH_SUPERCLASS	
		CCI_SCH_SUBCLASS	
		CCI_SCH_CONSTRIT	
		CCI_SCH_TRIGGER	
		CCI_SCH_CLASS_PRIVILEGE	
		CCI_SCH_ATTR_PRIVILEGE	
		CCI_SCH_DIRECT_SUPER_CLASS	
		CCI_SCH_PRIMARY_KEY	
		CCI_SCH_IMPORTED_KEYS	
		CCI_SCH_EXPORTED_KEYS	
		CCI_SCH_CROSS_REFERENCE	
T_CCI_CUBRID_STMT	enum	CUBRID_STMT_ALTER_CLASS	
		CUBRID_STMT_ALTER_SERIAL	
		CUBRID_STMT_COMMIT_WORK	
		CUBRID_STMT_REGISTER_DATABASE	

CUBRID_STMT_CREATE_CLASS

CUBRID_STMT_CREATE_INDEX

CUBRID_STMT_CREATE_TRIGGER

CUBRID_STMT_CREATE_SERIAL

CUBRID_STMT_DROP_DATABASE

CUBRID_STMT_DROP_CLASS

CUBRID_STMT_DROP_INDEX

CUBRID_STMT_DROP_LABEL

CUBRID_STMT_DROP_TRIGGER

CUBRID_STMT_DROP_SERIAL

CUBRID_STMT_EVALUATE

CUBRID_STMT_RENAME_CLASS

CUBRID_STMT_ROLLBACK_WORK

CUBRID_STMT_GRANT

CUBRID_STMT_REVOKE

CUBRID_STMT_STATISTICS

CUBRID_STMT_INSERT

CUBRID_STMT_SELECT

CUBRID_STMT_UPDATE

CUBRID_STMT_DELETE

CUBRID_STMT_CALL

CUBRID_STMT_GET_ISO_LVL

CUBRID_STMT_GET_TIMEOUT

CUBRID_STMT_GET_OPT_LVL

CUBRID_STMT_SET_OPT_LVL

CUBRID_STMT_SCOPE

CUBRID_STMT_GET_TRIGGER

CUBRID_STMT_SET_TRIGGER

CUBRID_STMT_SAVEPOINT

CUBRID_STMT_PREPARE

CUBRID_STMT_ATTACH

CUBRID_STMT_USE

CUBRID_STMT_REMOVE_TRIGGER

CUBRID_STMT_RENAME_TRIGGER

CUBRID_STMT_ON_LDB

CUBRID_STMT_GET_LDB

CUBRID_STMT_SET_LDB

CUBRID_STMT_GET_STATS

CUBRID_STMT_CREATE_USER

CUBRID_STMT_DROP_USER

		CUBRID_STMT_ALTER_USER
T_CCI_CURSOR_POS	enum	CCI_CURSOR_FIRST CCI_CURSOR_CURRENT CCI_CURSOR_LAST
T_CCI_TRAN_ISOLATION	enum	TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE TRAN_COMMIT_CLASS_COMMIT_INSTANCE TRAN_REP_CLASS_UNCOMMIT_INSTANCE TRAN_REP_CLASS_COMMIT_INSTANCE TRAN_REP_CLASS_REP_INSTANCE TRAN_SERIALIZABLE
T_CCI_PARAM_MODE	enum	CCI_PARAM_MODE_UNKNOWN CCI_PARAM_MODE_IN CCI_PARAM_MODE_OUT CCI_PARAM_MODE_INOUT

Note If a string longer than defined size in a column is inserted (INSERT) or updated (UPDATE), the string will be truncated.

cci_bind_param

Description

The `cci_bind_param` function binds data in the `bind` variable of prepared statement. Converts `value` of the given `a_type` to an actual binding type and stores it. Subsequently, whenever `cci_execute()` is called, the stored data is sent to the server. If `cci_bind_param()` is called multiple times for the same `index`, the last set value is configured.

If `NULL` is bound to the database, there can be two scenarios.

- `value` is a `NULL` pointer.
- `u_type` is `CCI_U_TYPE_NULL`.

If `CCI_BIND_PTR` is configured for `flag`, the pointer of `value` variable is copied (shallow copy), but no value is copied. If it is not configured for `flag`, the value of `value` variable is copied (deep copy) by allocating memory. If multiple columns are bound by using the same memory buffer, `CCI_BIND_PTR` must not be configured for the `flag`.

`T_CCI_A_TYPE` is a C language type that is used in CCI applications for data binding, and consists of primitive types such as `int` and `float` and user-defined types defined by CCI such as `T_CCI_BIT` and `T_CCI_DATE`. The identifier for each type is defined as shown in the table below.

a_type	value Type
CCI_A_TYPE_STR	char*
CCI_A_TYPE_INT	int*
CCI_A_TYPE_FLOAT	float*
CCI_A_TYPE_DOUBLE	double*
CCI_A_TYPE_BIT	T_CCI_BIT*
CCI_A_TYPE_SET	T_CCI_SET*
CCI_A_TYPE_DATE	T_CCI_DATE*
CCI_A_TYPE_BIGINT	int64_t* (For Windows : __int64*)

CCI_A_TYPE_BLOB	T_CCI_BLOB
CCI_A_TYPE_CLOB	T_CCI_CLOB

T_CCI_U_TYPE is a column type of database to convert binding data by the *value* argument. For the definition of the identifier for each type, see the table below.

The reason these types are used is to inform information converting convert A-type data which C language can interpret into U-type data which database can interpret. There are various A-type data that are allowed by U-type data. For example, **CCI_U_TYPE_INT** can receive **CCI_A_TYPE_STR** as A-type data including **CCI_A_TYPE_INT**. For information on type conversion, see "CUBRID SQL Guide > Data Types > Implicit Type Conversion > Rules."

Both **T_CCI_A_TYPE** and **T_CCI_U_TYPE** enum(s) are defined in the **cas_cci.h** file. The definition of each identifier is described in the table below.

u_type	Corresponding a_type (default)
CCI_U_TYPE_CHAR	CCI_A_TYPE_STR
CCI_U_TYPE_STRING	CCI_A_TYPE_STR
CCI_U_TYPE_NCHAR	CCI_A_TYPE_STR
CCI_U_TYPE_VARCHAR	CCI_A_TYPE_STR
CCI_U_TYPE_BIT	CCI_A_TYPE_BIT
CCI_U_TYPE_VARBIT	CCI_A_TYPE_BIT
CCI_U_TYPE_NUMERIC	CCI_A_TYPE_STR
CCI_U_TYPE_INT	CCI_A_TYPE_INT
CCI_U_TYPE_SHORT	CCI_A_TYPE_INT
CCI_U_TYPE_MONETARY	CCI_A_TYPE_DOUBLE
CCI_U_TYPE_FLOAT	CCI_A_TYPE_FLOAT
CCI_U_TYPE_DOUBLE	CCI_A_TYPE_DOUBLE
CCI_U_TYPE_DATE	CCI_A_TYPE_DATE
CCI_U_TYPE_TIME	CCI_A_TYPE_DATE
CCI_U_TYPE_TIMESTAMP	CCI_A_TYPE_DATE
CCI_U_TYPE_OBJECT	CCI_A_TYPE_STR
CCI_U_TYPE_BIGINT	CCI_A_TYPE_BIGINT
CCI_U_TYPE_DATETIME	CCI_A_TYPE_DATE

Syntax

```
int cci_bind_param(int req_handle, int index, T_CCI_A_TYPE a_type, void *value,
T_CCI_U_TYPE u_type, char flag)
```

- *req_handle* : (IN) Request handle of a prepared statement
- *index* : (IN) Location of binding! it starts with 1.
- *a_type* : (IN) Data type of *value*
- *value* : (IN) Data value to bind
- *u_type* : (IN) Data type to be applied to the database
- *flag* : (IN) bind_flag (**CCI_BIND_PTR**)

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_TYPE_CONVERSION
- CCI_ER_BIND_INDEX
- CCI_ER_ATYPE
- CCI_ER_NO_MORE_MEMORY

cci_bind_param_array**Description**

The `cci_bind_param_array` function binds a parameter array for a prepared `req_handle`. Subsequently, whenever [cci_execute_array\(\)](#) occurs, data is sent to the server by the stored `value` pointer. If `cci_bind_param_array()` is called multiple times for the same `index`, the last configured value is used. If `NULL` is bound to the data, a non-zero value is configured to `null_ind`.

If `value` is a `NULL` pointer, or `u_type` is `CCI_U_TYPE_NULL`, all data are bound to `NULL` and the data buffer used by `value` cannot be reused.

For the data type of `value` for `a_type`, see the [cci_bind_param\(\)](#) function description.

Syntax

```
int cci_bind_param_array(int req_handle, int index, T_CCI_A_TYPE a_type, void *value, int *null_ind, T_CCI_U_TYPE u_type)
```

- `req_handle` : (IN) Request handle of a prepared SQL statement
- `index` : (IN) Binding location
- `a_type` : (IN) Data type of `value`
- `value` : (IN) Data value to be bound
- `null_ind` : (IN) `NULL` indicator array (0 : not `NULL`, 1 : `NULL`)
- `u_type` : (IN) Data type to be applied to the database.

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_TYPE_CONVERSION
- CCI_ER_BIND_INDEX
- CCI_ER_ATYPE
- CCI_ER_BIND_ARRAY_SIZE

cci_bind_param_array_size**Description**

The `cci_bind_param_array_size` function determines the size of the array to be used in [cci_bind_param_array\(\)](#). `cci_bind_param_array_size()` must be called first before [cci_bind_param_array\(\)](#) is used.

Syntax

```
int cci_bind_param_array_size(int req_handle, int array_size)
```

- `req_handle` : (IN) Request handle of a prepared statement
- `array_size` : (IN) Binding array size

Return Value

- Error code (0 : success)

Error Code

- `CCI_ER_REQ_HANDLE`

cci_blob_free**Description**

The `cci_blob_free` function frees memory of **BLOB** structure.

Syntax

```
int cci_blob_free (T_CCI_BLOB blob)
```

Return Value

- Error code (0 : success)

Error Code

- `CCI_ER_INVALID_LOB_HANDLE`

cci_blob_new**Description**

The `cci_blob_new` function creates an empty file where **LOB** data is stored and returns Locator referring to the data to *blob* structure.

Syntax

```
int cci_blob_new(int conn_handle, T_CCI_BLOB* blob, T_CCI_ERROR* error_buf)
```

- *conn_handle* : (IN) Connection handle
- *blob* : (OUT) **LOB** Locator
- *error_buf* : (OUT) Error buffer

Return Value

- Error code (0 : success)

Error Codes

- `CCI_ER_CONNECT`
- `CCI_ER_COMMUNICATION`
- `CCI_ER_NO_MORE_MEMORY`
- `CCI_ER_DBMS`
- `CCI_ER_INVALID_LOB_HANDLE`

cci_blob_write**Description**

The `cci_blob_read` function reads as much as data from *start_pos* to *length* of the **LOB** data file specified in *blob*.; then it stores it in *buf* and returns it.

Syntax

```
int cci_blob_read(int conn_handle, T_CCI_BLOB blob, long start_pos, int length, const char
*buf, T_CCI_ERROR* error_buf)
```

- *conn_handle* : (IN) Connection handle
- *blob* : (IN) **LOB** Locator
- *start_pos* : (IN) Index location of **LOB** data file
- *length* : (IN) **LOB** data length from buffer
- *error_buf* : (OUT) Error buffer

Return Value

- Size of read value (>=0 : success)
- Error code (< 0 : error)

Error Codes

- CCI_ER_INVALID_LOB_READ_POS
- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_COMMUNICATION
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_DBMS
- CCI_ER_INVALID_LOB_HANDLE

cci_blob_size**Description**

The **cci_blob_size** function returns data file size that is specified in *blob*.

Syntax

```
long long cci_blob_size(T_CCI_BLOB* blob)
```

- *blob* : (IN) **LOB** Locator

Return Value

- Size of **BLOB** data file (>=0 : success)
- Error code (<0 : error)

Error Code

- CCI_ER_INVALID_LOB_HANDLE

cci_blob_write**Description**

The **cci_blob_write** function reads as much as data from *buf* to *length* and stores it from *start_pos* of the **LOB** data file specified in *blob*.

Syntax

```
int cci_blob_write(int conn_handle, T_CCI_BLOB blob, long start_pos, int length, const
char *buf, T_CCI_ERROR* error_buf)
```

- *conn_handle* : (IN) Connection handle
- *blob* : (IN) **LOB** Locator

- *start_pos* : (IN) Index location of **LOB** data file
- *length* : (IN) Data length from buffer
- *error_buf* : (OUT) Error buffer

Return Value

- Size of written value (> =0 : success)
- Error code (<0 : error)

Error Codes

- **CCI_ER_CON_HANDLE**
- **CCI_ER_CONNECT**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_NO_MORE_MEMORY**
- **CCI_ER_DBMS**
- **CCI_ER_INVALID_LOB_HANDLE**

cci_clob_free

Description

The **cci_clob_free** function frees memory of **CLOB** structure.

Syntax

```
int cci_clob_free (T_CCI_CLOB clob)
```

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_INVALID_LOB_HANDLE**

cci_clob_new

Description

The **cci_clob_new** function creates an empty file where **LOB** data is stored and returns Locator referring to the data to *clob* structure.

Syntax

```
int cci_clob_new(int conn_handle, T_CCI_CLOB* clob, T_CCI_ERROR* error_buf)
```

- *conn_handle* : (IN) Connection handle
- *clob* : (OUT) **LOB** Locator
- *error_buf* : (OUT) Error buffer

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_CONNECT**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_NO_MORE_MEMORY**

- CCI_ER_DBMS
- CCI_ER_INVALID_LOB_HANDLE

cci_clob_write

Description

The `cci_clob_write` function reads as much as data from `start_pos` to `length` in the **LOB** data file specified in `clob`; then it stores it in `buf` and returns it.

Syntax

```
int cci_clob_read(int conn_handle, T_CCI_CLOB clob, long start_pos, int length, const char
*buf, T_CCI_ERROR* error_buf)
```

- `conn_handle` : (IN) Connection handle
- `clob` : (IN) **LOB** Locator
- `start_pos` : (IN) Index location of **LOB** data file
- `length` : (IN) **LOB** data length from buffer
- `error_buf` : (OUT) Error buffer

Return Value

- Size of read value (≥ 0 : success)
- Error code (< 0 : Error)

Error Codes

- CCI_ER_INVALID_LOB_READ_POS
- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_COMMUNICATION
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_DBMS
- CCI_ER_INVALID_LOB_HANDLE

cci_clob_size

Description

The `cci_clob_size` function returns data file size that is specified in `clob`.

Syntax

```
long long cci_clob_size(T_CCI_CLOB* clob)
```

- `clob` : (IN) **LOB** Locator

Return Value

- Size of **CLOB** data file (≥ 0 : success)
- Error code (< 0 : error)

Error Code

- CCI_ER_INVALID_LOB_HANDLE

cci_clob_write

Description

The `cci_clob_write` function reads as much as data from *buf* to *length* and then stores the value from *start_pos* in **LOB** data file specified in *clob*.

Syntax

```
int cci_clob_write(int conn_handle, T_CCI_BLOB clob, long start_pos, int length, const
char *buf, T_CCI_ERROR* error_buf)
```

- *conn_handle* : (IN) Connection handle
- *clob* : (IN) **LOB** Locator
- *start_pos* : (IN) Index location of **LOB** data file
- *length* : (IN) Data length from buffer
- *error_buf* : (OUT) Error buffer

Return Value

- Size of written value (>=0 : success)
- Error code (<0 : Error)

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_COMMUNICATION
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_DBMS
- CCI_ER_INVALID_LOB_HANDLE

cci_close_req_handle

Description

The `cci_close_req_handle` function closes the request handle obtained by [cci_prepare\(\)](#).

Syntax

```
int cci_close_req_handle(int req_handle)
```

- *req_handle* : (IN) Request handle

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_COMMUNICATION

cci_col_get

Description

The `cci_col_get` function gets an attribute value of collection type. If the name of the class is C, and the domain of *set_attr* is set (multiset, sequence), the query looks like as follows:

```
SELECT a FROM C, TABLE(set_attr) AS t(a) WHERE C = oid;
```

That is, the number of members becomes the number of records.

Syntax

```
int cci_col_get (int conn_handle, char *oid_str, char *col_attr, int *col_size, int
*col_type, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *col_size* : (OUT) Collection size (-1 : null)
- *col_type* : (OUT) Collection type (set, multiset, sequence : u_type)
- *err_buf* : (OUT) Database error buffer

Return Value

- Request handle

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_col_seq_drop

Description

The **cci_col_seq_drop** function drops the index-th (base:1) member of the sequence attribute values. The following example shows how to drop the first member of the sequence attribute values.

```
cci_col_seq_drop(con_id, oid_str, seq_attr, 1, err_buf);
```

Syntax

```
int cci_col_seq_drop (int conn_handle, char *oid_str, char *col_attr, int index,
T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *index* : (IN) Index
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_col_seq_insert

Description

The **cci_col_seq_insert** function inserts a member at the index-th (base:1) position of the sequence attribute values. The following example shows how to insert "a" at the first position of the sequence attribute values.

```
cci_col_seq_insert(con_id, oid_str, seq_attr, 1, "a", err_buf);
```

Syntax

```
int cci_col_seq_insert (int conn_handle, char *oid_str, char *col_attr, int index, char *value, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *index* : (IN) Index
- *value* : (IN) Sequential element (string)
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_col_seq_put

Description

The **cci_col_seq_put** function replaces the index-th (base:1) member of the sequence attribute values with a new value. The following example shows how to replace the first member of the sequence attributes values with "a".

```
cci_col_seq_put(con_id, oid_str, seq_attr, 1, "a", err_buf);
```

Syntax

```
int cci_col_seq_put (int conn_handle, char *oid_str, char *col_attr, int index, char *value, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *index* : (IN) Index
- *value* : (IN) Sequential value
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT

- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_col_set_add

Description

The **cci_col_set_add** function adds a member to the set attribute values. The following example shows how to add "a" to the set attribute values.

```
cci_col_set_add(con_id, oid_str, set_attr, "a", err_buf);
```

Syntax

```
int cci_col_set_add ( int conn_handle, char *oid_str, char *col_attr, char *value,
T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *value* : (IN) Set element
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_col_set_drop

Description

The **cci_col_set_drop** function drops a member from the set attribute values. The following example shows how to drop "a" from the set attribute values.

```
cci_col_set_drop(con_id, oid_str, set_attr, "a", err_buf);
```

Syntax

```
int cci_col_set_drop (int conn_handle, char *oid_str, char *col_attr, char *value,
T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *col_attr* : (IN) Collection attribute name
- *value* : (IN) Set element (string)
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

cci_col_size

Description

The `cci_col_size` function gets the size of the set (seq) attribute.

Syntax

```
int cci_col_size (int conn handle, char *oid_str, char *col_attr, int *col_size,
T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `oid_str` : (IN) oid
- `col_attr` : (IN) Collection attribute name
- `col_size` : (OUT) Collection size (-1 : NULL)
- `err_buf` : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- `CCI_ER_CON_HANDLE`
- `CCI_ER_CONNECT`
- `CCI_ER_OBJECT`
- `CCI_ER_DBMS`

cci_connect

Description

A connection handle to the database server is assigned and it tries to connect to the server. If it has succeeded, the connection handle ID is returned; if fails, an error code is returned.

Syntax

```
cci_connect(char *ip, int port, char *db_name, char *db_user, char *db_password)
```

- `ip` : (IN) A character string representing the IP address of the server (host name)
- `port` : (IN) Broker port (the port configured in the `$CUBRID/conf/cubrid_broker.conf` file)
- `db_name` : (IN) Database name
- `db_user` : (IN) Database user name
- `db_passwd` : (IN) Database user password

Return Value

- Success : Connection handle ID (int)
- Failure : Error code

Error Codes

- `CCI_ER_NO_MORE_MEMORY`
- `CCI_ER_HOSTNAME`
- `CCI_ER_CON_HANDLE`
- `CCI_ER_DBMS`
- `CCI_ER_COMMUNICATION`
- `CCI_ER_CONNECT`

cci_connect_with_url

Description

The `cci_connect_with_url` function connects a database by using connection information passed with a url string argument. If CUBRID HA is enabled in CCI, you must specify the connection information of the standby server, which is used for failover when failure occurs, in the url string argument of this function. If it has succeeded, the ID of connection handle is returned; if it fails, an error code is returned.

Syntax

```
int cci_connect_with_url (char *url [, char *db_user, char *db_password ])

<url> ::=
cci:CUBRID:<host>:<db name>:<db user>:<db password>:[?<properties>]

<properties> ::= <property> [&<property>]
<property> ::= autocommit=<autocommit_mode>
<property> ::= althosts=<alternative_hosts> [ &vertime=<time>]
<property> ::= login_timeout=<milli_sec>
<property> ::= query_timeout=<milli_sec>
<property> ::= disconnect_on_query_timeout=true|false

<alternative_hosts> ::= <host>:<port> [,<host>:<port>]

<host> ::= HOSTNAME | IP_ADDR
<time> ::= SECOND
<milli_sec> ::= MILLI SECOND
```

- *url* : (IN) A character string that contains server connection information
- *host* : A host name or IP address of the master database
- *db_name* : A name of the database
- *db_user* : A name of the database user
- *db_password* : A database user password
- *autocommit=true/false* : Configures the value of auto-commit upon database connection
- *althosts=standby_broker1_host, standby_broker2_host, ...* : Specifies the broker information of the standby server, which is used for failover when it is impossible to connect to the active server. You can specify multiple brokers for failover, and the connection to the brokers is attempted in the order listed in **althosts**.
- *vertime* : An interval between the attempts to connect to the active broker in which failure occurred. After a failure occurs, the system connects to the broker specified by *althosts* (failover), terminates the transaction, and then attempts to connect to the active broker of the master database at every **vertime**. The default value is 600 seconds.
- *login_timeout* : Timeout value (unit: msec.) for database login. Upon timeout, a **CCI_ER_LOGIN_TIMEOUT** error is returned. The default value is 0, which means infinite postponement.
- *query_timeout* : Timeout value (unit: msec.) for query request. Upon timeout, a message to cancel requesting a query transferred to server is sent. The return value can depend on the **disconnect_on_query_timeout** configuration; even though the message to cancel a request is sent to server, that request may succeed.
- *disconnect_on_query_timeout* : Configures a value whether to immediately return an error of function being executed upon timeout. The default value is **false**. When timeout for query request occurs and it this value is **true**, a **CCI_ER_QUERY_TIMEOUT** error is returned after a cancellation message is sent and a socket is closed. In this case, a user must explicitly terminate database connection by using the **cci_disconnect** function. If the value is **false**, it waits until a response for query request from a server after a cancellation is sent.
- *db_user* : (IN) A name of the database user
- *db_passwd* : (IN) A database user password

Return Value

- Success : Connection handle ID (int)
- Failure : Error code

Error Codes

- **CCI_ER_NO_MORE_MEMORY**

- **CCI_ER_HOSTNAME**
- **CCI_ER_INVALID_URL**
- **CCI_ER_CON_HANDLE**
- **CCI_ER_CONNECT**
- **CCI_ER_DBMS**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_LOGIN_TIMEOUT**

Example

```
--connection URL string when a property(althosts) specified for HA
URL=cci:CUBRID:192.168.0.1:33000:demodb:::althosts=192.168.0.2:33000,192.168.0.3:33000

--connection URL string when properties(althosts,rctime) specified for HA
URL=cci:CUBRID:192.168.0.1:33000:demodb:::althosts=192.168.0.2:33000,192.168.0.3:33000&rctime=600
```

Remark

Because a colon (:) and a question mark (?) are used as a separator in URL string, it is not allowed to include them for password of URL string. To use them, you must specify a user name (*db_user*) and a password (*db_passwd*) as a separate parameter.

cci_cursor

Description

The **cci_cursor** function moves the cursor specified in the request handle to access the specific record in the query result executed by [cci_execute\(\)](#). The position of cursor is moved by the values specified in the *origin* and *offset* values. If the position to be moved is not valid, **CCI_ER_NO_MORE_DATA** is returned.

Syntax

```
int cci_cursor(int req_handle, int offset, T_CCI_CURSOR_POS origin, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle
- *offset* : (IN) Offset to be moved
- *origin* : (IN) Variable to represent a position. The type is **T_CCI_CURSOR_POS**. **T_CCI_CURSOR_POS** enum consists of **CCI_CURSOR_FIRST**, **CCI_CURSOR_CURRENT**, and **CCI_CURSOR_LAST**.
- *err_buf*: (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_REQ_HANDLE**
- **CCI_ER_NO_MORE_DATA**
- **CCI_ER_COMMUNICATION**

Example

```
//the cursor moves to the first record
cci_cursor(req, 1, CCI_CURSOR_FIRST, &err_buf);

//the cursor moves to the next record
cci_cursor(req, 1, CCI_CURSOR_CURRENT, &err_buf);

//the cursor moves to the last record
cci_cursor(req, 1, CCI_CURSOR_LAST, &err_buf);
```

```
//the cursor moves to the previous record
cci_cursor(req, -1, CCI_CURSOR_CURRENT, &err_buf);
```

cci_cursor_update

Description

The **cci_cursor_update** function updates *cursor_pos* from the value of the *index*-th column to *value*. If the database is updated to **NULL**, *value* becomes **NULL**. For update conditions, see [cci_prepare\(\)](#). The data type of *value* for *a_type* is shown in the table below.

a_type	value Type
CCI_A_TYPE_STR	char*
CCI_A_TYPE_INT	int*
CCI_A_TYPE_FLOAT	float*
CCI_A_TYPE_DOUBLE	double*
CCI_A_TYPE_BIT	T_CCI_BIT*
CCI_A_TYPE_SET	T_CCI_SET
CCI_A_TYPE_DATE	T_CCI_DATE*
CCI_A_TYPE_BIGINT	int64_t (For Windows : __int64)
CCI_A_TYPE_BLOB	T_CCI_BLOB
CCI_A_TYPE_CLOB	T_CCI_CLOB

Syntax

```
int cci_cursor_update(int req_handle, int cursor_pos, int index, T_CCI_A_TYPE a_type, void *value, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle
- *cursor_pos* : (IN) Cursor position
- *index* : (IN) Column index
- *a_type* : (IN) *value* Type
- *value* : (IN) A new value
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_TYPE_CONVERSION
- CCI_ER_ATYPE

cci_datasource_borrow

Description

The **cci_datasource_borrow** function obtains CCI connection to be used in **T_CCI_DATASOURCE** structure.

Syntax

```
T_CCI_CONN cci_datasource_borrow (T_CCI_DATASOURCE * datasource, T_CCI_ERROR * err)
```

- *datasource* : **T_CCI_DATASOURCE** structure pointer in which CCI connection exists
- *err* : Error code and message returned upon error occurrence

Return Value

- Success : CCI connection handler identifier
- Failure : -1

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

cci_datasource_create

Description

The **cci_datasource_create** function creates DATASOURCE of CCI.

Syntax

```
T_CCI_DATASOURCE * cci_datasource_create (T_CCI_PROPERTIES * properties, T_CCI_ERROR * err)
```

- *properties* : **T_CCI_PROPERTIES** structure pointer in which configuration of structure pointer is stored
- *err* : Error code and message returned upon error occurrence

Return Value

- Success : **CCI_DATASOURCE** structure pointer created
- Failure : **NULL**

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

cci_datasource_destroy

Description

The **cci_datasource_destroy** function destroys DATASOURCE of CCI.

Syntax

```
void cci_datasource_destroy (T_CCI_DATASOURCE * datasource)
```

- *datasource* : **T_CCI_DATASOURCE** structure pointer to be destroyed

Return Value

None

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_release](#)

cci_datasource_release**Description**

The `cci_datasource_release` function returns CCI connection released in `T_CCI_DATASOURCE` structure.

Syntax

```
int cci_datasource_release (T_CCI_DATASOURCE * data_source, T_CCI_CONN conn)
```

- *datasource* : `T_CCI_DATASOURCE` structure pointer which returns CCI connection
- *conn* : CCI connection handler identifier released

Return Value

- Success : 1
- Failure : 0

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)

cci_disconnect**Description**

The `cci_disconnect` function disconnects all request handles created for *conn_handle*. If a transaction is being performed, the handles are disconnected after [cci_end_tran\(\)](#) is executed.

Syntax

```
int cci_disconnect(int conn_handle, T_CCI_ERROR * err_buf)
```

- *conn_handle* : (IN) Connection handle
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- `CCI_ER_CON_HANDLE`
- `CCI_ER_DBMS`
- `CCI_ER_COMMUNICATION`

`cci_end_tran`

Description

The `cci_end_tran` function performs a commit or rollback on the current transaction. At this point, all open request handles are terminated and the connection to the database server is disabled. However, even after the connection to the server is disabled, the connection handle remains valid.

You can configure the default value of auto-commit mode by using `CCI_DEFAULT_AUTOCOMMIT` (broker parameter) upon startup of an application. If configuration on broker parameter is omitted, the default value is `ON`; use the [cci_set_autocommit\(\)](#) function to change auto-commit mode within an application. If auto-commit mode is `OFF`, you must explicitly commit or roll back transaction by using the [cci_end_tran\(\)](#) function.

Syntax

```
int cci_end_tran(int conn_handle, char type, T_CCI_ERROR *err_buf)
```

- `conn_handle`: (IN) Connection handle
- `type`: (IN) `CCI_TRAN_COMMIT` or `CCI_TRAN_ROLLBACK`
- `err_buf`: (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- `CCI_ER_CON_HANDLE`
- `CCI_ER_DBMS`
- `CCI_ER_COMMUNICATION`
- `CCI_ER_TRAN_TYPE`

Remark

Auto-commit mode is supported for SELECT statements. To apply this mode, you must add `SELECT_AUTO_COMMIT=ON` to the `cubrid_broker.conf` file. However, auto-commit is performed only at the point at which the result set for all n query statements is fetched from the server when there are n prepared statements. An example is as follows:

Example 1

```
$sql1 = "select * from db_user";
$sql2 = "select * from db class where owner name = ?";

$result = cubrid_execute($con, $sql1); // 1 select handle. fetch completed - autocommit
if ($result) {
    while ($row = cubrid_fetch ($result))
    {
        echo ($row[0]);

        $req = cubrid_prepare ($con, $sql2);
        cubrid_bind ($req, 1, $row[0]);
        $res = cubrid_execute ($req); // 1 select handle. fetch completed - autocommit
    }
}
```

Example 2

```

$sql1 = "select * from db user";
$sql2 = "select * from db class where owner name = ?";

$req = cubrid_prepare ($con, $sql2);
$result = cubrid_execute($con, $sql1); // 2 handle. fetch completed for only 1 handle -
no autocommit
if ($result) {
    while ($row = cubrid_fetch ($result))
    {
        echo ($row[0]);

        cubrid_bind ($req, 1, $row[0]);
        $res = cubrid_execute ($req); // fetch completed for all select handles -
autocommit
    }
}

```

Example 3

```

$sql1 = "select * from db user";
$sql2 = "insert into a values (?)";

$result = cubrid_execute($con, $sql1); // 1 select handle. fetch completed - autocommit
if ($result) {
    while ($row = cubrid_fetch ($result))
    {
        echo ($row[0]);

        $req = cubrid_prepare ($con, $sql2);
        cubrid_bind ($req, 1, $row[0]);
        $res = cubrid_execute ($req); // no autocommit for insert
    }
}

```

Example 4

```

$sql1 = "select * from db user";
$sql2 = "insert into a values (?)";

$req = cubrid_prepare ($con, $sql2);
$result = cubrid_execute($con, $sql1); // no autocommit for insert because no fetch
if ($result) {
    while ($row = cubrid_fetch ($result))
    {
        echo ($row[0]);

        cubrid_bind ($req, 1, $row[0]);
        $res = cubrid_execute ($req); // no autocommit for insert
    }
}

```

cci_execute**Description**

The **cci_execute** function executes the SQL statement (prepared statement) that has executed [cci_prepare\(\)](#). A request handle, *flag*, the maximum length of a column to be fetched, and the address of a **T_CCI_ERROR** construct variable in which error information being stored are specified as arguments.

The function of retrieving the query result from the server through a *flag* can be classified as synchronous or asynchronous. If the flag is set to **CCI_EXEC_QUERY_ALL**, a synchronous mode (*sync_mode*) is used to retrieve query results immediately after executing prepared queries if it is set to **CCI_EXEC_ASYNC**, an asynchronous mode (*async_mode*) is used to retrieve the result immediately each time a query result is created. The *flag* is set to **CCI_EXEC_QUERY_ALL** by default, and in such cases the following rules are applied.

- The return value is the result of the first query.

- If an error occurs in any query, the execution is processed as a failure.
- For a query composed of in a query composed of q1 q2 q3 if an error occurs in q2 after q1 succeeds the execution, the result of q1 remains valid. That is, the previous successful query executions are not rolled back when an error occurs.
- If a query is executed successfully, the result of the second query can be obtained using [cci_next_result\(\)](#).

max_col_size is a value that is used to determine the maximum length of a column to be sent to a client when the type of the column of the prepared query is **CHAR**, **VARCHAR**, **NCHAR**, **VARNCHAR**, **BIT** or **VARBIT**. If this value is 0, full length is fetched.

Syntax

```
int cci_execute(int req_handle, char flag, int max_col_size, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle of a prepared SQL statement
- *flag* : (IN) Exec flag (**CCI_EXEC_ASYNC** or **CCI_EXEC_QUERY_ALL**)
- *max_col_size* : (IN) The maximum length of a column to be fetched when it is a string data type in bytes. If this value is 0, full length is fetched.
- *err_buf* : (OUT) Database error buffer

Return Value

- Success
- **SELECT** : Returns the number of results in sync mode returns 0 in async mode.
- **INSERT, UPDATE** : Returns the number of tuples reflected.
- Others queries : 0
- Failure : Error code

Error Codes

- **CCI_ER_REQ_HANDLE**
- **CCI_ER_BIND**
- **CCI_ER_DBMS**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_QUERY_TIMEOUT**
- **CCI_ER_LOGIN_TIMEOUT**

cci_execute_array

Description

If more than one value are bound to the prepared statement, this gets the values of the variables to be bound and executes the query by binding each value to the variable.

To bind the data, call the [cci_bind_param_array_size\(\)](#) function to specify the size of the array, bind each value to the variable by using the [cci_bind_param_array\(\)](#) function, and execute the query by calling the [cci_execute_array\(\)](#) function.

You can get three execution results by calling the [cci_execute\(\)](#) function. However, the [cci_execute_array\(\)](#) function returns the number of queries executed by the *query_result* variable. You can use the following macro to get the information about the execution result. However, note that the validity check is not performed for each parameter entered in the macro. After using the *query_result* variable, you must delete the *query_result* by using the [cci_query_result_free\(\)](#) function.

Macro	Return Type	Meaning
CCI_QUERY_RESULT_RESULT	int	the number of results
CCI_QUERY_RESULT_ERR_MSG	char*	error message about query

 CCI_QUERY_RESULT_STMT_TYPE int(T_CCI_CUBRID_STMT enum) type of query statement

Syntax

```
int cci_execute_array(int req_handle, T_CCI_QUERY_RESULT **query_result, T_CCI_ERROR
*err_buf)
```

- *req_handle* : (IN) Request handle of a prepared SQL statement
- *query_result* : (OUT) Query results (the number of executed queries)
- *err_buf* : (OUT) Database error buffer

Return Value

- Success : The number of executed queries
- Failure : Negative number

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_BIND
- CCI_ER_DBMS
- CCI_ER_COMMUNICATION
- CCI_ER_QUERY_TIMEOUT
- CCI_ER_LOGIN_TIMEOUT

Example

```
char *query =
    "update participant set gold = ? where host year = ? and nation code = 'KOR'";
int gold[2];
char *host_year[2];
int null_ind[2];
T_CCI_QUERY_RESULT *result;
int n_executed;
...

req = cci_prepare (con, query, 0, &cci_error);
if (req < 0)
{
    printf ("prepare error: %d, %s\n", cci_error.err code, cci_error.err msg);
    goto handle_error;
}

gold[0] = 20;
host_year[0] = "2004";

gold[1] = 15;
host_year[1] = "2008";

null_ind[0] = null_ind[1] = 0;
error = cci_bind_param_array_size (req, 2);
if (error < 0)
{
    printf ("bind_param_array_size error: %d\n", error);
    goto handle_error;
}

error =
    cci_bind_param_array (req, 1, CCI_A_TYPE_INT, gold, null_ind, CCI_U_TYPE_INT);
if (error < 0)
{
    printf ("bind param array error: %d\n", error);
    goto handle_error;
}

error =
    cci_bind_param_array (req, 2, CCI_A_TYPE_STR, host_year, null_ind, CCI_U_TYPE_INT);
if (error < 0)
```

```

{
    printf ("bind param array error: %d\n", error);
    goto handle_error;
}

n_executed = cci_execute_array (req, &result, &cci_error);
if (n_executed < 0)
{
    printf ("execute error: %d, %s\n", cci_error.err_code,
           cci_error.err_msg);
    goto handle_error;
}
for (i = 1; i <= n_executed; i++)
{
    printf ("query %d\n", i);
    printf ("result count = %d\n", CCI_QUERY_RESULT_RESULT (result, i));
    printf ("error message = %s\n", CCI_QUERY_RESULT_ERR_MSG (result, i));
    printf ("statement type = %d\n",
           CCI_QUERY_RESULT_STMT_TYPE (result, i));
}
error = cci_query_result_free (result, n_executed);
if (error < 0)
{
    printf ("query result free: %d\n", error);
    goto handle_error;
}
error = cci_end_tran(con, CCI_TRAN_COMMIT, &cci_error);
if (error < 0)
{
    printf ("end tran: %d, %s\n", cci_error.err_code, cci_error.err_msg);
    goto handle_error;
}
}

```

cci_execute_batch

Description

In CCI, multiple jobs can be processed simultaneously when using DML queries such as **INSERT/UPDATE/DELETE**. [cci_execute_array\(\)](#) and [cci_execute_batch\(\)](#) functions can be used to execute such batch jobs. Note that prepared statements cannot be used in the [cci_execute_batch\(\)](#) function.

Executes `sql_stmt` as many times as `num_sql_stmt` specified as a parameter and returns the number of queries executed with the `query_result` variable. You can use the macro ([CCI_QUERY_RESULT_RESULT](#), [CCI_QUERY_RESULT_ERR_MSG](#), [CCI_QUERY_RESULT_STMT_TYPE](#)) available in the [cci_execute_array\(\)](#) function to get the information about the execution result.

However, note that the validity check is not performed for each parameter entered in the macro. After using the `query_result` variable, you must delete the query result by using the [cci_query_result_free\(\)](#) function.

Syntax

```
int cci_execute_batch(int conn_handle, int num_sql_stmt, char **sql_stmt,
T_CCI_QUERY_RESULT **query_result, T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `num_sql_stmt` : (IN) The number of `sql_stmts`
- `sql_stmt` : (IN) SQL statement array
- `query_result` : (OUT) The results of `sql_stmt`
- `err_buf` : (OUT) Database error buffer

Return Value

- Success : The number of executed queries
- Failure : Negative number

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_DBMS
- CCI_ER_COMMUNICATION
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_CONNECT
- CCI_ER_QUERY_TIMEOUT
- CCI_ER_LOGIN_TIMEOUT

Example

```

char **queries;
T CCI_QUERY_RESULT *result;
int n_queries, n_executed;
...

count = 3;
queries = (char **) malloc (count * sizeof (char *));
queries[0] =
    "insert into athlete(name, gender, nation_code, event) values('Ji-sung Park', 'M',
'KOR', 'Soccer')";
queries[1] =
    "insert into athlete(name, gender, nation code, event) values('Joo-young Park', 'M',
'KOR', 'Soccer')";
queries[2] =
    "select * from athlete order by code desc for orderby_num() < 3";
//calling cci_execute_batch()
n_executed = cci_execute_batch (con, count, queries, &result, &cci_error);
if (n_executed < 0)
{
    printf ("execute_batch: %d, %s\n", cci_error.err_code,
        cci_error.err_msg);
    goto handle_error;
}
printf ("%d statements were executed.\n", n_executed);

for (i = 1; i <= n_executed; i++)
{
    printf ("query %d\n", i);
    printf ("result count = %d\n", CCI_QUERY_RESULT_RESULT (result, i));
    printf ("error message = %s\n", CCI_QUERY_RESULT_ERR_MSG (result, i));
    printf ("statement type = %d\n",
        CCI_QUERY_RESULT_STMT_TYPE (result, i));
}

error = cci_query_result_free (result, n_executed);
if (error <
0)
{
    printf ("query result free: %d\n", error);
    goto handle_error;
}

```

cci_execute_result

Description

The **cci_execute_result** function gets the execution results (e.g. statement type, result count) performed by [cci_execute\(\)](#). The results of each query are retrieved by [CCI_QUERY_RESULT_STMT_TYPE](#) and [CCI_QUERY_RESULT_RESULT](#). The query results used must be deleted by [cci_query_result_free](#).

Syntax

```
int cci_execute_result(int req_handle, T_CCI_QUERY_RESULT **query_result, T_CCI_ERROR
*err_buf)
```

- *req_handle* : (IN) Request handle of a prepared SQL statement
- *query_result* : (OUT) Query results
- *err_buf*: (OUT) Database error buffer

Return Value

- Success : The number of queries
- Failure : Negative number

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_COMMUNICATION

Example

```
T_CCI_QUERY_RESULT *qr;
...
cci_execute( ... );
res = cci_execute_result(req_h, &qr, &err_buf);
if (res < 0) {
    /* error */
}
else {
    for (i=1 ; i <= res ; i++) {
        result_count = CCI_QUERY_RESULT_RESULT(qr, i);
        stmt_type = CCI_QUERY_RESULT_STMT_TYPE(qr, i);
    }
    cci_query_result_free(qr, res);
}
```

cci_fetch

Description

The **cci_fetch** function fetches the query result executed by [cci_execute\(\)](#) from the server-side CAS and stores it to the client buffer. The [cci_get_data\(\)](#) function can be used to identify the data of a specific column from the fetched query result.

Syntax

```
int cci_fetch(int req_handle, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle
- *err_buf*: (OUT) Database error buffer

Return Value

- Error code (0 : success)

cci_fetch_buffer_clear

Description

The **cci_fetch_buffer_clear** function clears the records temporarily stored in the client buffer.

Syntax

```
int cci_fetch_buffer_clear(int req_handle)
```

- *req_handle* : (IN) Request handle

Return Value

- Error code (0 : success)

Error Code

- CCI_ER_REQ_HANDLE

cci_fetch_sensitive

Description

The `cci_fetch_sensitive` function sends changed values for sensitive columns when the results are sent to the client from the server. If the results by *req_handle* are not sensitive, they are same as the ones by [cci_fetch\(\)](#). The return value of CCI_ER_DELETED_TUPLE means that the given row has been deleted.

Syntax

```
int cci_fetch_sensitive(int req_handle, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_NO_MORE_DATA
- CCI_ER_COMMUNICATION
- CCI_ER_DBMS
- CCI_ER_DELETED_TUPLE

cci_fetch_size

Description

The `cci_fetch_size` function determines the number of records sent by [cci_fetch\(\)](#) from the server to the client.

Syntax

```
int cci_fetch_size(int req_handle, int fetch_size)
```

- *req_handle* : (IN) Request handle
- *fetch_size* : (IN) Fetch size

Return Value

- Error code (0 : success)

Error Code

- CCI_ER_REQ_HANDLE

cci_get_autocommit

Description

The `cci_get_autocommit` function returns the auto-commit mode currently configured.

Syntax

```
int cci_get_autocommit (int conn_handle)
```

- `conn_handle` : Connection handle

Return Value

- 1 : Auto-commit ON
- 0 : Auto-commit OFF

Error Code

- None

cci_get_bind_num

Description

The `cci_get_bind_num` function gets the number of input bindings. If the SQL statement used during preparation is composed of multiple queries, it represents the number of input bindings used in all queries.

Syntax

```
int cci_get_bind_num(int req_handle)
```

- `req_handle` : (IN) Request handle for a prepared SQL statement

Return Value

- The number of input bindings

Error Codes

- `CCI_ER_REQ_HANDLE`

cci_get_class_num_objs

Description

The `cci_get_class_num_objs` function gets the number of objects of the `class_name` class and the number of pages being used. If the flag is configured to 1, an approximate value is fetched; if it is configured to 0, an exact value is fetched.

Syntax

```
int cci_get_class_num_objs(int conn_handle, char *class_name, int flag, int *num_objs, int *num_pages, T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `class_name` : (IN) Class name
- `flag` : (IN) 0 or 1
- `num_objs` : (OUT) The number of objects
- `num_pages` : (OUT) The number of pages
- `err_buf` : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_COMMUNICATION
- CCI_ER_CONNECT

CCI_GET_COLLECTION_DOMAIN**Description**

If *u_type* is set, multiset or sequence type, this macro gets the domain of the set, multiset or sequence. If *u_type* is not a set type, the return value is the same as *u_type*.

Syntax

```
#define CCI_GET_COLLECTION_DOMAIN(u_type)
```

Return Value

- Type (CCI_U_TYPE)

cci_get_cur_oid**Description**

The **cci_get_cur_oid** function gets the OID of the currently fetched records if **CCI_INCLUDE_OID** is configured in execution. The OID is represented as a string for a page, slot or volume.

Syntax

```
int cci_get_cur_oid(int req_handle, char *oid_str_buf)
```

- *conn_handle* : (IN) Request handle
- *oid_str_buf* : (OUT) OID string

Return Value

- Error code (0 : success)

Error Code

- CCI_ER_REQ_HANDLE

cci_get_data**Description**

The **cci_get_data** function gets the *col_noth* value from the currently fetched result. The *type* of the *value* variable is determined according to the given *type* parameter, and the value or the pointer is copied to the value variable accordingly.

For a value to be copied, the memory for the address to be transferred to the *value* variable must have been previously assigned. Note that if a pointer is copied, a pointer in the application client library is returned, so the value becomes invalid next time the **cci_get_data()** function is called.

In addition, the pointer returned by the pointer copy must not be freed. However, if the type is **CCI_A_TYPE_SET**, the memory must be freed by using the [cci_set_free\(\)](#) function after using the set because the set is returned after the

T_CCI_SET type memory is allocated. The following table shows the summary of *type* parameters and data types of their corresponding *values*.

type	value Type	Meaning
CCI_A_TYPE_STR	char**	pointer copy
CCI_A_TYPE_INT	int*	value copy
CCI_A_TYPE_FLOAT	float*	value copy
CCI_A_TYPE_DOUBLE	double*	value copy
CCI_A_TYPE_BIT	T_CCI_BIT*	value copy (pointer copy for each member)
CCI_A_TYPE_SET	T_CCI_SET*	memory alloc and value copy
CCI_A_TYPE_DATE	T_CCI_DATE*	value copy
CCI_A_TYPE_BIGINT	int64_t* (For Windows : __int64*)	value copy
CCI_A_TYPE_BLOB	T_CCI_BLOB	memory alloc and value copy
CCI_A_TYPE_CLOB	T_CCI_CLOB	memory alloc and value copy

Syntax

```
int cci_get_data(int req_handle, int col_no, int type, void *value, int *indicator)
```

- *req_handle* : (IN) Request handle
- *col_no* : (IN) One-based column index. It starts with 1.
- *type* : (IN) Data type (defined in the **T_CCI_A_TYPE**) of *value* variable
- *value* : (OUT) Variable address for data to be stored
- *indicator* : (OUT) NULL indicator (-1 : NULL)
- if *type* is **CCI_A_TYPE_STR** : -1 is returned in case of NULL; the length of character string stored in *value* is returned, otherwise.
- if *type* is **CCI_A_TYPE_STR** : -1 is returned in case of NULL, 0 is returned, otherwise.

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_REQ_HANDLE**
- **CCI_ER_TYPE_CONVERSION**
- **CCI_ER_COLUMN_INDEX**
- **CCI_ER_ATYPE**

Remark

- For **LOB** type, if the **cci_get_data()** function is called, meta data with the **LOB** type column (Locator) is displayed. To call data of the **LOB** type column, the [cci_blob_read\(\)](#) function should be called.

cci_get_db_parameter

Description

The **cci_get_db_parameter** function gets a parameter specified in the database. The data type of *value* for *param_name* is shown in the table below.

param_name	value Type	note
CCI_PARAM_ISOLATION_LEVEL	int*	get/set

CCI_PARAM_LOCK_TIMEOUT	int*	get/set
CCI_PARAM_MAX_STRING_LENGTH	int*	get only

Syntax

```
int cci_get_db_parameter(int conn_handle, T_CCI_DB_PARAM param_name, void *value,
T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *param_name* : (IN) System parameter name
- *value* : (OUT) Parameter value
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_PARAM_NAME
- CCI_ER_DBMS
- CCI_ER_COMMUNICATION
- CCI_ER_CONNECT

cci_get_db_version

Description

The `cci_get_db_version` function gets the Database Management System (DBMS) version.

Syntax

```
int cci_get_db_version(int conn_handle, char *out_buf, int out_buf_size)
```

- *conn_handle* : (IN) Connection handle
- *out_buf* : (OUT) Result buffer
- *out_buf_size* : (IN) *out_buf* size

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_COMMUNICATION
- CCI_ER_CONNECT

cci_get_query_timeout

Description

The `cci_get_query_timeout` function returns timeout configured for query execution.

Syntax

```
int cci_get_query_timeout (int req_handle)
```

- *conn_handle* : (IN) Request handle

Return Value

- Success : Timeout value configured in current request handle (unit: msec.)
- Failure : Error code

Error Code

- CCI_ER_REQ_HANDLE

cci_get_result_info

Description

If the prepared statement is **SELECT**, the **T_CCI_COL_INFO** struct that stores the column information about the execution result can be obtained by using this function. If it is not **SELECT**, **NULL** is returned and the *num* value becomes 0.

You can access the **T_CCI_COL_INFO** struct directly to get the column information from the struct, but you can also use a macro to get the information, which is defined as follows. The address of the **T_CCI_COL_INFO** struct and the column index are specified as parameters for each macro. The macro can be called only for the **SELECT** query. Note that the validity check is not performed for each parameter entered in each macro. If the return type of the macro is *char**, do not free the memory pointer.

Macro	Return Type	Meaning
CCI_GET_RESULT_INFO_TYPE	T_CCI_U_TYPE	column type
CCI_GET_RESULT_INFO_SCALE	short	column scale
CCI_GET_RESULT_INFO_PRECISION	int	column precision
CCI_GET_RESULT_INFO_NAME	char*	column name
CCI_GET_RESULT_INFO_ATTR_NAME	char*	column attribute name
CCI_GET_RESULT_INFO_CLASS_NAME	char*	column class name
CCI_GET_RESULT_INFO_IN_NON_NULL	char(0 or 1)	whether a column is NULL

Syntax

```
T_CCI_COL_INFO* cci_get_result_info(int req_handle, T_CCI_SQLX_CMD *cmd_type, int *num)
```

- *req_handle* : (IN) Request handle for a prepared SQL statement
- *cmd_type* : (OUT) Command type
- *num* : (OUT) The number of columns in the **SELECT** statement (if *cmd_type* is **SQLX_CMD_SELECT**)

Return Value

- Success : Result info pointer
- Failure : **NULL**

Example

```
col_info = cci_get_result_info (req, &cmd_type, &col_count);
if (col_info == NULL)
{
    printf ("get result info error: %d, %s\n", cci error.err code,
           cci error.err msg);
    goto handle error;
}
for (i = 1; i <= col_count; i++)
{
    printf ("%12s = %d\n", "type", CCI_GET_RESULT_INFO_TYPE (col_info, i));
    printf ("%12s = %d\n", "scale",
           CCI_GET_RESULT_INFO_SCALE (col_info, i));
}
```

```

printf ("%12s = %d\n", "precision",
        CCI_GET_RESULT_INFO_PRECISION (col_info, i));
printf ("%12s = %s\n", "name", CCI_GET_RESULT_INFO_NAME (col_info, i));
printf ("%12s = %s\n", "attr_name",
        CCI_GET_RESULT_INFO_ATTR_NAME (col_info, i));
printf ("%12s = %s\n", "class name",
        CCI_GET_RESULT_INFO_CLASS_NAME (col_info, i));
printf ("%12s = %s\n", "is non null",
        CCI_GET_RESULT_INFO_IS_NON_NULL (col_info,i) ? "true" : "false");

```

CCI_GET_RESULT_INFO_ATTR_NAME

Description

The `CCI_GET_RESULT_INFO_ATTR_NAME` macro gets the actual attribute name of the *index*-th column of a prepared `SELECT` statement. If there is no name for the attribute (constant, function, etc), "" (empty string) is returned. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid. You cannot delete the returned memory pointer with `free()`.

Syntax

```
#define CCI_GET_RESULT_INFO_ATTR_NAME(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) pointer to the column information fetched by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- Attribute name (char*)

CCI_GET_RESULT_INFO_CLASS_NAME

Description

The `CCI_GET_RESULT_INFO_CLASS_NAME` macro gets the *index*-th class name of a prepared `SELECT` statement. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid. You cannot delete the returned memory pointer with `free()`. The return value can be `NULL`.

Syntax

```
#define CCI_GET_RESULT_INFO_CLASS_NAME(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) Column info pointer by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- Class name (char*)

CCI_GET_RESULT_INFO_IS_NON_NULL

Description

The `CCI_GET_RESULT_INFO_IS_NON_NULL` macro gets a value indicating whether the *index*-th column of a prepared `SELECT` statement is nullable. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_GET_RESULT_INFO_IS_NON_NULL(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) Column info pointer by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- 0 : nullable
- 1 : non NULL

CCI_GET_RESULT_INFO_NAME**Description**

The `CCI_GET_RESULT_INFO_NAME` macro gets the *index*-th column name of a prepared `SELECT` statement. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid. You cannot delete the returned memory pointer with `free()`.

Syntax

```
#define CCI_GET_RESULT_INFO_NAME(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) Column info pointer to [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- Column name (char*)

CCI_GET_RESULT_INFO_PRECISION**Description**

The `CCI_GET_RESULT_INFO_PRECISION` macro gets the *index* -th precision of a prepared `SELECT` statement. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_GET_RESULT_INFO_PRECISION(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) Column info pointer by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- Precision (int)

CCI_GET_RESULT_INFO_SCALE**Description**

The `CCI_GET_RESULT_INFO_SCALE` macro gets the *index*-th column's scale of a prepared `SELECT` statement. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_GET_RESULT_INFO_SCALE(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) Column info pointer by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- scale (int)

CCI_GET_RESULT_INFO_TYPE

Description

The `CCI_GET_RESULT_INFO_TYPE` macro gets the *index*-th column type of a prepared `SELECT` statement. It does not check whether the specified argument, *res_info*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_GET_RESULT_INFO_TYPE(T_CCI_COL_INFO* res_info, int index)
```

- *res_info* : (IN) pointer to the column information fetched by [cci_get_result_info](#)
- *index* : (IN) Column index

Return Value

- Column type (`T_CCI_U_TYPE`)

CCI_IS_SET_TYPE, CCI_IS_MULTISSET_TYPE, CCI_IS_SEQUENCE_TYPE, CCI_IS_COLLECTION_TYPE

Description

The `CCI_IS_SET_TYPE`, `CCI_IS_MULTISSET_TYPE`, `CCI_IS_SEQUENCE_TYPE`, and `CCI_IS_COLLECTION_TYPE` macros check whether *u_type* is set, multiset or sequence type.

Syntax

```
#define CCI_IS_SET_TYPE(u_type)
#define CCI_IS_MULTISSET_TYPE(u_type)
#define CCI_IS_SEQUENCE_TYPE(u_type)
#define CCI_IS_COLLECTION_TYPE(u_type)
```

Return Value

- `CCI_IS_SET_TYPE`
 - 1 : set
 - 0 : not set
- `CCI_IS_MULTISSET_TYPE`
 - 1 : multiset
 - 0 : not multiset
- `CCI_IS_SEQUENCE_TYPE`
 - 1 : sequence
 - 0 : not sequence
- `CCI_IS_COLLECTION_TYPE`
 - 1 : collection (set, multiset, sequence)
 - 0 : not collection

cci_is_updatable

Description

The `cci_is_updatable` function checks whether the SQL statement, which executed [cci_prepare\(\)](#), is updatable. If it is updatable, 1 is returned.

Syntax

```
int cci_is_updatable(int req_handle)
```


- *req_handle* : (IN) Request handle for a prepared SQL statement

Return Value

- 1 : updatable
- 0 : not updatable

Error Code

- CCI_ER_REQ_HANDLE

cci_next_result

Description

The **cci_next_result** function gets results of next query if **CCI_EXEC_QUERY_ALL** flag is set upon [cci_execute\(\)](#). The information about the query fetched by next_result can be obtained with [cci_get_result_info](#). If next_result is executed successfully, the database is updated with the information of the current query.

The error code **CAS_ER_NO_MORE_RESULT_SET** means that no more result set exists.

Syntax

```
int cci_next_result(int req_handle, T_CCI_ERROR *err_buf)
```

- *req_handle* : (IN) Request handle of a prepared statement
- *err_buf* : (OUT) Database error buffer

Return Value

- Success
- **SELECT** (sync mode) : the number of results, (async mode) : 0
- **INSERT, UPDATE** : the number of records reflected
- Others : 0
- Failure : Error code

Error Codes

- CCI_ER_REQ_HANDLE
- CCI_ER_DBMS
- CCI_ER_COMMUNICATION

cci_oid

Description

CCI_OID_DROP : Deletes the given oid.

CCI_OID_IS_INSTANCE : Checks whether the given oid is an instance oid.

CCI_OID_LOCK_READ : Sets a read lock on the given oid.

CCI_OID_LOCK_WRITE : Sets a write lock on the given oid.

Syntax

```
int cci_oid(int conn_handle, T_CCI_OID_CMD cmd, char *oid_str, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *cmd* : (IN) CCI_OID_DROP, CCI_OID_IS_INSTANCE, CCI_OID_LOCK_READ, CCI_OID_LOCK_WRITE
- *oid_str* : (IN) oid
- *err_buf* : (OUT) Database error buffer

Return Value

- CCI_OID_IS_INSTANCE
- 0 : non-instance
- 1 : instance
- 0 : error
- CCI_OID_DROP, CCI_OID_LOCK_READ, CCI_OID_LOCK_WRITE
- Error code (0 : success)

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_OID_CMD
- CCI_ER_OBJECT
- CCI_ER_DBMS

cci_oid_get**Description**

The `cci_oid_get` function gets the attribute values of the given oid. `attr_name` is an array of the attributes, and it must end with NULL. If `attr_name` is NULL, the information of all attributes is fetched. The request handle has the same form as when the SQL statement "SELECT attr_name FROM oid_class WHERE oid_class = oid" is executed.

Syntax

```
int cci_oid_get(int conn_handle, char *oid_str, char **attr_name, T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `oid_str` : (IN) oid
- `attr_name` : (IN) A list of attributes
- `err_buf` : (OUT) Database error buffer

Return Value

- Success : Request handle
- Failure : Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_CONNECT

cci_oid_get_class_name**Description**

The `cci_oid_get_class_name` function gets the class name of the given oid.

Syntax

```
int cci_oid_get_class_name(int conn_handle, char *oid_str, char *out_buf, int out_buf_len, T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `oid_str` : (IN) oid
- `out_buf` : (OUT) Out buffer

- *out_buf_len* : (IN) *out_buf* length
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code

Error Codes

- **CCI_ER_CON_HANDLE**
- **CCI_ER_CONNECT**
- **CCI_ER_OBJECT**
- **CCI_ER_DBMS**

cci_oid_put

Description

The **cci_oid_put** function configures the *attr_name* attribute values of the given oid to *new_val_str*. The last value of *attr_name* must be **NULL**. Any value of any type must be represented as a string. The value represented as a string is applied to the database after being converted depending on the attribute type on the server. To insert a **NULL** value, configure the value of *new_val_str[i]* to **NULL**.

Syntax

```
int cci_oid_put(int conn_handle, char *oid_str, char **attr_name, char **new_val_str,
T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *attr_name* : (IN) A list of attribute names
- *new_val_str* : (IN) A list of new values
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_CON_HANDLE**
- **CCI_ER_CONNECT**

cci_oid_put2

Description

The **cci_oid_put2** function sets the *attr_name* attribute values of the given oid to *new_val*. The last value of *attr_name* must be **NULL**. To insert a **NULL** value, set the value of *new_val[i]* to **NULL**.

The type of *new_val[i]* for *a_type* is shown in the table below.

Type of *new_val[i]* for *a_type*

Type	value type
CCI_A_TYPE_STR	char*
CCI_A_TYPE_INT	int*
CCI_A_TYPE_FLOAT	float*
CCI_A_TYPE_DOUBLE	double*

CCI_A_TYPE_BIT	T_CCI_BIT*
CCI_A_TYPE_SET	T_CCI_SET
CCI_A_TYPE_DATE	T_CCI_DATE*
CCI_A_TYPE_BIGINT	int64_t (For Windows : __int64)

Syntax

```
int cci_oid_put2(int conn_handle, char *oidstr, char **attr_name, void **new_val, int
*a_type, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *oid_str* : (IN) oid
- *attr_name* : (IN) A list of attribute names
- *new_val* : (IN) A new value array
- *a_type* : (IN) *new_val* type array
- *err_buf* : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT

Example

```
char *attr_name[array_size]
void *attr_val[array_size]
int a_type[array_size]
int int_val

...
attr_name[0] = "attr_name0"
attr_val[0] = &int_val
a_type[0] = CCI_A_TYPE_INT
attr_name[1] = "attr_name1"
attr_val[1] = "attr_val1"
a_type[1] = CCI_A_TYPE_STR

...
attr_name[num_attr] = NULL

res = cci_put2(con_h, oid_str, attr_name, attr_val, a_type, &error)
...
```

cci_prepare

Description

The **cci_prepare** function prepares SQL execution by acquiring request handle for SQL statements. If a SQL statement consists of multiple queries, the preparation is performed only for the first query. With the parameter of this function, an address to **T_CCI_ERROR** where connection handle, SQL statement, *flag*, and error information are stored.

CCI_PREPARE_UPDATABLE or **CCI_PREPARE_INCLUDE_OID** can be configured in *flag*. If **CCI_PREPARE_UPDATABLE** is configured, updatable result set is created and **CCI_PREPARE_INCLUDE_OID** is automatically configured. However, not all updatable result sets are created even though **CCI_PREPARE_UPDATABLE** is configured. So you need to check if the results are updatable by using [cci_is_updatable](#) after preparation.

The conditions of updatable queries are as follows:

- A query must be **SELECT**.
- OID must be contained in the query result.
- The column to be updated must be the one that belongs to the table specified in the **FROM** clause.

Syntax

```
int cci_prepare(int conn_handle, char *sql_stmt, char flag, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *sql_stmt* : (IN) SQL statement
- *flag* : (IN) prepare flag (**CCI_PREPARE_INCLUDE_OID** or **CCI_PREPARE_UPDATABLE**)
- *err_buf* : (OUT) Database error buffer

Return Value

- Success : Request handle ID (int)
- Failure : Error code (negative)

Error Codes

- **CCI_ER_CON_HANDLE**
- **CCI_ER_DBMS**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_STR_PARAM**
- **CCI_ER_NO_MORE_MEMORY**
- **CCI_ER_CONNECT**
- **CCI_ER_QUERY_TIMEOUT**
- **CCI_ER_LOGIN_TIMEOUT**

cci_prepare_and_execute

Description

The **cci_prepare_and_execute** function executes the SQL statement immediately and returns a request handle for the SQL statement. A request handle, SQL statement, the maximum length of a column to be fetched, error code, and the address of a **T_CCI_ERROR** construct variable in which error information being stored are specified as arguments. *max_col_size* is a value to configure the maximum length of a column to be sent to a client when the column of a SQL statement is **CHAR**, **VARCHAR**, **NCHAR**, **VARCHAR**, **BIT**, or **VARBIT**. If this value is 0, full length is fetched.

Syntax

```
int cci_prepare_and_execute(int conn_handle, char *sql_stmt, int max_col_size, int *exec_retval, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Request handle
- *sql_stmt* : (IN) SQL statement
- *max_col_size* : (IN) The maximum length of a column to be fetched when it is a string data type in bytes. If this value is 0, full length is fetched.
- *exec_retval* : (OUT) Error code
- *err_buf* : (OUT) Database error buffer

Return Value

- Success : Request handle ID (int)
- Failure : Error code

Error Codes

- **CCI_ER_CON_HANDLE**

- CCI_ER_DBMS
- CCI_ER_COMMUNICATION
- CCI_ER_STR_PARAM
- CCI_ER_NO_MORE_MEMORY
- CCI_ER_CONNECT
- CCI_ER_QUERY_TIMEOUT

cci_property_create

Description

The `cci_property_create` function creates `T_CCI_PROPERTIES` structure to configure DATASOURCE of CCI.

Syntax

```
T_CCI_PROPERTIES * cci_property_create ()
```

Return Value

- Success : `T_CCI_PROPERTIES` structure pointer in which memory is allocated
- Failure : `NULL`

See Also

- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

cci_property_destroy

Description

The `cci_property_destroy` function destroys `T_CCI_PROPERTIES` structure.

Syntax

```
void cci_property_destroy (T_CCI_PROPERTIES * properties)
```

- *properties* : `T_CCI_PROPERTIES` structure pointer to be destroyed

Return Value

None

See Also

- [cci_property_create](#)
- [cci_property_get](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

cci_property_get

Description

The `cci_property_get` function gets the value of `T_CCI_PROPERTIES` structure.

Syntax

```
char * cci_property_get (T_CCI_PROPERTIES * properties, char *key)
```

- `properties` : `T_CCI_PROPERTIES` structure pointer which gets value corresponding to `key`
- `key` : String pointer of property to be obtained

Return Value

- Success : String pointer of value corresponding to `key`
- Failure : NULL

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_set](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

cci_property_set

Description

It configures a property value in `T_CCI_PROPERTIES` structure. The property name and its meaning that can be configured in the structure are as follows:

- **pool_size** : Maximum number of connection (default : 10)
- **max_wait** : Maximum waiting time to get connection (default : 1000 msec.)
- **pool_prepared_statement** : Whether to enable statement pooling (default: false)
- **login_timeout** : Login timeout time (default : 0 (unlimited))
- **query_timeout** : Query timeout time (default : 0(unlimited))
- **disconnect_on_query_timeout** : Whether to terminate connection when execution is discarded due to query execution timeout (default : no)
- **default_autocommit** : Auto-commit mode refreshed whenever `cci_datasource_borrow` is called (true or false)
- **default_isolation** : Transaction isolation level refreshed whenever `cci_datasource_borrow` is called
- **default_lock_timeout** : lock_timeout refreshed whenever `cci_datasource_borrow` is called

If you configure **default_autocommit**, **default_isolation**, or **default_lock_timeout** value, connection for autocommit, isolation, or lock_timeout based on current configured value is returned when `cci_datasource_borrow` is called. If you do not configure it, connection for autocommit, isolation, or lock_timeout is returned with keeping the value that a user changed before.

default_isolation has one of the following configuration values. For details on isolation level, see "CUBRID SQL Guide > Transaction and Lock > Transaction Isolation Level."

isolation_level	Configuration Value
SERIALIZABLE	"TRAN_SERIALIZABLE"
REPEATABLE READ CLASS with REPEATABLE READ INSTANCES	"TRAN_REP_CLASS_REP_INSTANCE" or "TRAN_REP_READ"

REPEATABLE READ CLASS with READ COMMITTED INSTANCES	"TRAN_REP_CLASS_COMMIT_INSTANCE" or "TRAN_READ_COMMITTED" or "TRAN_CURSOR_STABILITY"
REPEATABLE READ CLASS with READ UNCOMMITTED INSTANCES	"TRAN_REP_CLASS_UNCOMMIT_INSTANCE" or "TRAN_READ_UNCOMMITTED"
READ COMMITTED CLASS with READ COMMITTED INSTANCES	"TRAN_COMMIT_CLASS_COMMIT_INSTANCE"
READ COMMITTED CLASS with READ UNCOMMITTED INSTANCES	"TRAN_COMMIT_CLASS_UNCOMMIT_INSTANCE"

Syntax

```
int cci_property_set (T_CCI_PROPERTIES * properties, char * key, char * value)
```

- *properties* : T_CCI_PROPERTIES structure pointer in which *key* and *value* are stored
- *key* : String pointer of property name
- *value* : String pointer of property value

Return Value

- Success : 1
- Failure : 0

See Also

- [cci_property_create](#)
- [cci_property_destroy](#)
- [cci_property_get](#)
- [cci_datasource_borrow](#)
- [cci_datasource_create](#)
- [cci_datasource_destroy](#)
- [cci_datasource_release](#)

CCI_QUERY_RESULT_ERR_MSG

Description

The CCI_QUERY_RESULT_ERR_MSG macro gets error messages for the [cci_execute_batch](#) query. If there is no error message, "" (empty string) is returned. It does not check whether the specified argument, *query_result*, is NULL, and whether *index* is valid.

Syntax

```
#define CCI_QUERY_RESULT_ERR_MSG(T_CCI_QUERY_RESULT* query_result, int index)
```

- *query_result* : (IN) Query results of [cci_execute_batch](#)
- *index* : (IN) Column index (base : 1)

Return Value

- Error message

cci_query_result_free

Description

The `cci_query_result_free` function deletes query result.

Syntax

```
int cci_query_result_free(T_CCI_QUERY_RESULT* query_result, int num_query)
```

- *query_result* : (IN) Query results of [cci_execute_batch](#)
- *num_query* : (IN) The number of arrays in *query_result*

Return Value

- Error code (0 : success)

Example

```
T_CCI_QUERY_RESULT *qr;
char **sql stmt;

res = cci_execute_array(conn, &qr, &err_buf);

cci_query_result_free(qr, res);
```

CCI_QUERY_RESULT_RESULT

Description

The `CCI_QUERY_RESULT_RESULT` macro gets the result count of the [cci_execute_batch](#) query. It does not check whether the specified argument, *query_result*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_QUERY_RESULT_RESULT(T_CCI_QUERY_RESULT* query_result, int index)
```

- *query_result* : (IN) Query results of [cci_execute_batch](#)
- *index* : (IN) Column index (base : 1)

Return Value

- result count

CCI_QUERY_RESULT_STMT_TYPE

Description

The `CCI_QUERY_RESULT_STMT_TYPE` macro gets the statement type of the [cci_execute_batch](#) query. It does not check whether the specified argument, *query_result*, is `NULL` and whether *index* is valid.

Syntax

```
#define CCI_QUERY_RESULT_STMT_TYPE(T_CCI_QUERY_RESULT* query_result, int index)
```

- *query_result* : (IN) Query results of [cci_execute_batch](#)
- *index* : (IN) Column index (base : 1)

Return Value

- statement type (`T_CCI_SQLX_CMD`)

cci_savepoint

Description

The **cci_savepoint** function configures savepoint or performs transaction rollback to a specified savepoint. If *cmd* is set to **CCI_SP_SET**, it configures savepoint and if it is set to **CCI_SP_ROLLBACK**, it rolls back transaction to specified savepoint.

Syntax

```
int cci_savepoint(int conn_handle, T_CCI_SAVEPOINT_CMD cmd, char* savepoint_name,
T_CCI_ERROR *err_buf)
```

- *conn_handle*: (IN) Connection handle
- *cmd*: (IN) CCI_SP_SET or CCI_SP_ROLLBACK
- *savepoint_name*: (IN) Savepoint name
- *err_buf*: (OUT) Database error buffer

Return Value

- Error code

Example

```
con = cci connect( ...);
.../* query execute */

/* sets a savepoint named "savepoint1"
cci_savepoint(con, CCI_SP_SET, "savepoint1", err_buf);

... /* query execute */

/* rolls back the set savepoint to "savepoint1" */
cci_savepoint(con, CCI_SP_ROLLBACK, "savepoint1", err_buf);
```

cci_schema_info

Description

The **cci_schema_info** function gets schema information. If it is performed successfully, the results are managed by the request handle and can be fetched by `fetch` and `getdata`. If you want to retrieve a *class_name* of *attr_name* by pattern matching, configure the *flag*.

Two flags, **CCI_CLASS_NAME_PATTERN_MATCH** and **CCI_ATTR_NAME_PATTERN_MATCH**, are used for pattern matching. You can configure these two *flags* by using the OR operator (`|`). Performance may significantly decrease if pattern matching is used.

The following table shows records composition of each *type*.

Record Composition of Each Type

Type	Column Order	Column Name	Column Type
CCI_SCH_CLASS	1	NAME	char*
	2	TYPE	short 0 : system class 1 : vclass 2 : class 3 : proxy
CCI_SCH_VCLASS	1	NAME	char*
	2	TYPE	short 1 : vclass 3 : proxy

CCI_SCH_ATTRIBUTE	1	NAME	char*
	2	DOMAIN	int
	3	SCALE	int
	4	PRECISION	int
	5	INDEXED	int 1 : indexed
	6	NON_NULL	int 1 : non null
	7	SHARED	int 1 : shared
	8	UNIQUE	int 1 : unique
	9	DEFAULT	void*
	10	ATTR_ORDER	int base : 1
	11	CLASS_NAME	char*
	12	SOURCE_CLASS	char*
	13	IS_KEY	short 1 : key
CCI_SCH_CLASS_METHOD	1	NAME	char*
	2	RET_DOMAIN	int
	3	ARG_DOMAIN	char*
CCI_SCH_METHOD_FILE	1	METHOD_FILE	char*
CCI_SCH_super class	1	CLASS_NAME	char*
	2	TYPE	short
CCI_SCH_SUBCLASS	1	CLASS_NAME	char*
	2	TYPE	short
CCI_SCH_CONSTRAINT	1	TYPE 0 : unique 1 : index 2 : reverse unique 3 : reverse index	int
	2	NAME	char*
	3	ATTR_NAME	char*
	4	NUM_PAGES	int
	5	NUM_KEYS	int
	6	PRIMARY_KEY 1 : primary key	short
	7	KEY_ORDER	short base : 1
CCI_SCH_TRIGGER	1	NAME	char*
	2	STATUS	char*
	3	EVENT	char*
	4	TARGET_CLASS	char*

	5	TARGET_ATTR	char*
	6	ACTION_TIME	char*
	7	ACTION	char*
	8	PRIORITY	float
	9	CONDITION_TIME	char*
	10	CONDITION	char*
CCI_SCH_CLASS_PRIVILEGE	1	CLASS_NAME	char*
	2	PRIVELEGE	char*
	3	GRANTABLE	char*
CCI_SCH_ATTR_PRIVILEGE	1	ATTR_NAME	char*
	2	PRIVILEGE	char*
	3	GRANTABLE	char*
CCI_SCH_PRIMARY_KEY	1	CLASS_NAME	char*
	2	ATTR_NAME	char*
	3	KEY_SEQ	short base : 1
	4	KEY_NAME	char*
CCI_SCH_IMPORTED_KEY	1	PKTABLE_NAME	char**
Used to retrieve primary key columns that are referred by a foreign key column in a given table. The results are sorted by PKTABLE_NAME and KEY_SEQ. If this type is specified as a parameter, a foreign key table is specified for <i>class_name</i> , and NULL is specified for <i>attr_name</i> .	2	PKCOLUMN_NAME	char**
	3	FKTABLE_NAME	char**
	4	FKCOLUMN_NAME	char**
	5	KEY_SEQ	char**
	6	UPDATE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*
	7	DELETE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*
	8	FK_NAME	char**
	9	PK_NAME	char**
CCI_SCH_EXPORTED_KEYS	1	PKTABLE_NAME	char**
Used to retrieve primary key columns that are referred by all foreign key columns. The results are sorted by FKTABLE_NAME and KEY_SEQ. If this type is specified as a parameter, a primary key table is specified for <i>class_name</i> , and NULL is specified for <i>attr_name</i> .	2	PKCOLUMN_NAME	char**
	3	FKTABLE_NAME	char**
	4	FKCOLUMN_NAME	char**
	5	KEY_SEQ	char**
	6	UPDATE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*

	7	DELETE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*
	8	FK_NAME	char**
	9	PK_NAME	char**
CCI_SCH_CROSS_REFERENCE Used to retrieve foreign key information when primary keys and foreign keys in a given table are cross referenced. The results are sorted by FKTABLE_NAME and KEY_SEQ. If this type is specified as a parameter, a primary key is specified for <i>class_name</i> , and a foreign key table is specified for <i>attr_name</i> .	1	PKTABLE_NAME	char**
	2	PKCOLUMN_NAME	char**
	3	FKTABLE_NAME	char**
	4	FKCOLUMN_NAME	char**
	5	KEY_SEQ	char**
	6	UPDATE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*
	7	DELETE_ACTION -cascade=0 -restrict=1 -no action=2 -set null=3	Int*
	8	FK_NAME	char**
	9	PK_NAME	char**

Pattern match

CCI_SCH_TYPE	Class name	ATTR_name
CCI_SCH_CLASS (VCLASS)	O	none
CCI_SCH_ATTRIBUTE (CLASS ATTRIBUTE)	O	O
CCI_SCH_CLASS_PRIVILEGE	O	none
CCI_SCH_ATTR_PRIVILEGE	X	O
CCI_SCH_PRIMARY_KEY	O	none

If the pattern flag is not configured, exact string matching is used for the given class or attribute name. Therefore, there is no result if NULL is given. If the name of the class or attribute is NULL when the pattern flag is configured, the result is the same as when "%" is used.

Note TYPE column of CCI_SCH_CLASS and CCI_SCH_VCLASS : The proxy type is added. When used in OLEDB, ODBC or PHP, vclass is represented without distinguishing between proxy and vclass.

Syntax

```
int cci_schema_info(int conn_handle, T_CCI_SCHEMA_TYPE type, char *class_name, char *attr_name, char flag, T_CCI_ERROR *err_buf)
```

- *conn_handle* : (IN) Connection handle
- *type* : (IN) Schema type
- *class_name* : (IN) Class name or NULL
- *attr_name* : (IN) Attribute name or NULL
- *flag* : (IN) Pattern matching flag (CCI_CLASS_NAME_PATTERN_MACTH or CCI_CLASS_NAME_PATTERN_MATCH)

- *err_buf*: (OUT) Database error buffer

Return Value

- Success : Request handle
- Failure : Error code

Error Codes

- **CCI_ER_CON_HANDLE**
- **CCI_ER_DBMS**
- **CCI_ER_COMMUNICATION**
- **CCI_ER_SCHEMA_TYPE**
- **CCI_ER_NO_MORE_MEMORY**
- **CCI_ER_CONNECT**

cci_set_allocators

Description

The **cci_set_allocators** function registers the memory allocation/release functions used by users. By executing this function, you can use user-defined functions for every memory allocation/release jobs being processed in CCI APIs. If you do not use this function, system functions (malloc, free, realloc, and calloc) are used.

Syntax

```
int cci_set_allocators(CCI_MALLOC_FUNCTION malloc_func, CCI_FREE_FUNCTION free_func,
CCI_REALLOC_FUNCTION realloc_func, CCI_CALLOC_FUNCTION calloc_func)
```

- *malloc_func* : (IN) Pointer of externally defined function corresponding to malloc
- *free_func* : (IN) Pointer of externally defined function corresponding to free
- *realloc_func* : (IN) Pointer of externally defined function corresponding to realloc
- *calloc_func* : (IN) Pointer of externally defined function corresponding to calloc

Return Value

- Error code (0 : success)

Error Codes

- **CCI_ER_NOT_IMPLEMENTED**

Example

```
/*
   How to build: gcc -Wall -g -o test cci test cci.c -I${CUBRID}/include -
L${CUBRID}/lib -lcascci
*/

#include <stdio.h>
#include <stdlib.h>
#include "cas cci.h"

void *
my_malloc(size_t size)
{
    printf ("my malloc: size: %ld\n", size);
    return malloc (size);
}

void *
my_calloc(size_t nm, size_t size)
{
    printf ("my calloc: nm: %ld, size: %ld\n", nm, size);
}
```

```

    return calloc (nm, size);
}

void *
my_realloc(void *ptr, size_t size)
{
    printf ("my_realloc: ptr: %p, size: %ld\n", ptr, size);
    return realloc (ptr, size);
}

void
my_free(void *ptr)
{
    printf ("my_free: ptr: %p\n", ptr);
    return free (ptr);
}

int
test_simple (int con)
{
    int req = 0, col_count = 0, i, ind;
    int error;
    char *data;
    T_CCI_ERROR cci_error;
    T_CCI_COL_INFO *col_info;
    T_CCI_SQLX_CMD cmd_type;
    char *query = "select * from db class";

    //preparing the SQL statement
    req = cci_prepare (con, query, 0, &cci_error);
    if (req < 0)
    {
        printf ("prepare error: %d, %s\n", cci_error.err_code,
            cci_error.err_msg);
        goto handle_error;
    }

    //getting column information when the prepared statement is the SELECT query
    col_info = cci_get_result_info (req, &cmd_type, &col_count);
    if (col_info == NULL)
    {
        printf ("get_result_info error: %d, %s\n", cci_error.err_code,
            cci_error.err_msg);
        goto handle_error;
    }

    //Executing the prepared SQL statement
    error = cci_execute (req, 0, 0, &cci_error);
    if (error < 0)
    {
        printf ("execute error: %d, %s\n", cci_error.err_code,
            cci_error.err_msg);
        goto handle_error;
    }
    while (1)
    {
        //Moving the cursor to access a specific tuple of results
        error = cci_cursor (req, 1, CCI_CURSOR_CURRENT, &cci_error);
        if (error == CCI_ER_NO_MORE_DATA)
        {
            break;
        }
        if (error < 0)
        {
            printf ("cursor error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
            goto handle_error;
        }

        //Fetching the query result into a client buffer
        error = cci_fetch (req, &cci_error);
    }
}

```

```

    if (error < 0)
    {
        printf ("fetch error: %d, %s\n", cci_error.err_code,
                cci_error.err_msg);
        goto handle error;
    }
    for (i = 1; i <= col count; i++)
    {
//Getting data from the fetched result
        error = cci get data (req, i, CCI A TYPE STR, &data, &ind);
        if (error < 0)
        {
            printf ("get data error: %d, %d\n", error, i);
            goto handle_error;
        }
        printf ("%s\t", data);
    }
    printf ("\n");
}

//Closing the request handle
error = cci close req handle (req);
if (error < 0)
{
    printf ("close_req_handle error: %d, %s\n", cci_error.err_code,
            cci_error.err_msg);
    goto handle error;
}

//Disconnecting with the server
error = cci_disconnect (con, &cci_error);
if (error < 0)
{
    printf ("error: %d, %s\n", cci_error.err_code, cci_error.err_msg);
    goto handle error;
}

return 0;

handle error:
if (req > 0)
    cci_close_req_handle (req);
if (con > 0)
    cci_disconnect (con, &cci_error);

return 1;
}

int main()
{
    int con = 0;

    if (cci_set_allocators (my_malloc, my_free, my_realloc, my_calloc) != 0)
    {
        printf ("cannot register allocators\n");
        return 1;
    };

    //getting a connection handle for a connection with a server
    con = cci_connect ("localhost", 33000, "demodb", "dba", "");
    if (con < 0)
    {
        printf ("cannot connect to database\n");
        return 1;
    }

    test simple (con);
    return 0;
}

```


cci_set_autocommit

Description

The `cci_set_autocommit` function configures the auto-commit mode of current database connection. It is only used to turn ON/OFF of auto-commit mode. When this function is called, every transaction being processed is committed regardless of configured mode.

Note that `CCI_DEFAULT_AUTOCOMMIT`, broker parameter configured in the `cubrid_broker.conf` file, determines whether it is in auto-commit mode upon program startup.

Syntax

```
int cci_set_autocommit (int conn_handle, CCI_AUTOCOMMIT_MODE autocommit_mode)
```

- `conn_handle` : (IN) Connection handle
- `autocommit_mode` : (IN) Configures the auto-commit mode. It has one of the following value: `CCI_AUTOCOMMIT_FALSE` or `CCI_AUTOCOMMIT_TRUE`.

Return Value

- Error code (0 : success)

Error Code

- `CCI_ER_CON_HANDLE`

cci_set_db_parameter

Description

The `cci_set_db_parameter` function configures a system parameter. For the type of *value* for *param_name*, see [cci_get_db_parameter\(\)](#).

Syntax

```
int cci_set_db_parameter (int conn_handle, T_CCI_DB_PARAM param_name, void* value, T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `param_name` : (IN) System parameter name
- `value` : (IN) Parameter value
- `err_buf` : (OUT) Database error buffer

Return Value

- Error code (0 : success)

Error Codes

- `CCI_ER_CON_HANDLE`
- `CCI_ER_PARAM_NAME`
- `CCI_ER_DBMS`
- `CCI_ER_COMMUNICATION`
- `CCI_ER_CONNECT`

cci_set_element_type

Description

The `cci_set_element_type` function gets the element type for the set fetched by `CCI_A_TYPE_SET` with [cci_get_data\(\)](#).

Syntax

```
int cci_set_element_type(T_CCI_SET set)
```

- *set* : (IN) cci set pointer

Return Value

- Type

cci_set_free**Description**

The `cci_set_free` function releases the memory assigned to `T_CCI_SET` gotten by `CCI_A_TYPE_SET` with `cci_get_data()`.

Syntax

```
void cci_set_free(T_CCI_SET set)
```

- *set* : (IN) cci set pointer

Return Value

- None

cci_set_get**Description**

The `cci_set_get` function gets the index-th data for the set fetched by `CCI_A_TYPE_SET` with `cci_get_data()`. The data type of *value* for *a_type* is shown in the table below.

a_type	value Type
<code>CCI_A_TYPE_STR</code>	<code>char**</code>
<code>CCI_A_TYPE_INT</code>	<code>int*</code>
<code>CCI_A_TYPE_FLOAT</code>	<code>float*</code>
<code>CCI_A_TYPE_DOUBLE</code>	<code>double*</code>
<code>CCI_A_TYPE_BIT</code>	<code>T_CCI_BIT*</code>
<code>CCI_A_TYPE_DATE</code>	<code>T_CCI_DATE*</code>
<code>CCI_A_TYPE_BIGINT</code>	<code>int64_t*</code> (For Windows : <code>__int64*</code>)

Syntax

```
int cci_set_get(T_CCI_SET set, int index, T_CCI_A_TYPE a_type, void *value, int *indicator)
```

- *set* : (IN) cci set pointer
- *index* : (IN) Set index (base : 1)
- *a_type* : (IN) Type
- *value* : (OUT) Result buffer
- *indicator* : (OUT) Null indicator

Return Value

- Error code

cci_set_isolation_level

Description

The `cci_set_isolation_level` function sets the transaction isolation level of connections. All further transactions for the given connections work as `new_isolation_level`.

Note If the transaction isolation level is set by `cci_set_db_parameter()`, only the current transaction is affected. When the transaction is complete, the transaction isolation level returns to the one set by CAS. You must use `cci_set_isolation_level()` to set the isolation level for the entire connection.

Syntax

```
int cci_set_isolation_level(int conn_handle, T_CCI_TRAN_ISOLATION new_isolation_level,
T_CCI_ERROR *err_buf)
```

- `conn_handle` : (IN) Connection handle
- `new_isolation_level` : (IN) Transaction isolation level

```
int cci_set_isolation_level(int conn_handle, T_CCI_TRAN_ISOLATION new_isolation_level,
T_CCI_ERROR *err_buf)
```

Return Value

- Error code

Error Codes

- CCI_ER_CON_HANDLE
- CCI_ER_CONNECT
- CCI_ER_ISOLATION_LEVEL
- CCI_ER_DBMS

cci_set_make

Description

The `cci_set_make` function makes a set of a new `CCI_A_TYPE_SET` type. The created set is sent to the server as `CCI_A_TYPE_SET` by `cci_bind_param()`. The memory for the set created by `cci_set_make()` must be freed by `cci_set_free()`. The type of `value` for `u_type` is shown in the table below.

Syntax

```
int cci_set_make(T_CCI_SET *set, T_CCI_U_TYPE u_type, int size, void *value, int *indicator)
```

- `set` : (IN) cci set pointer
- `u_type` : (IN) Element type
- `size` : (IN) Set size
- `value` : (IN) Set element
- `indicator` : (IN) Null indicator array

Return Value

- Error code

cci_set_max_row

Description

The `cci_set_max_row` function configures the maximum number of records for the results of the `SELECT` statement executed by `cci_execute`. If the `max` value is 0, it is the same as not setting the value.

Syntax

```
int cci_set_max_row(int req_handle, int max)
```

- *req_handle* : (IN) Connection handle
- *max* : (IN) The maximum number of rows

Return Value

- Error code

Example

```
req = cci_prepare( ... );
cci_set_max_row(req, 1);
cci_execute( ... );
```

cci_set_query_timeout**Description**

The **cci_set_query_timeout** function configures timeout value for query execution.

The timeout value configured by **cci_set_query_timeout** affects **cci_prepare**, **cci_execute**, **cci_execute_array**, an **cci_execute_batch** functions. When timeout occurs in the function and if the **disconnect_on_query_timeout** value configured in **cci_connect_with_url** connection URL is yes, it returns the **CCI_ER_QUERY_TIMEOUT** error.

These functions can return the **CCI_ER_LOGIN_TIMEOUT** error in case that **login_timeout** is configured in the connection URL, which is an argument of **cci_connect_with_url** function; this means that login timeout happens between application client and server (CAS) during re-connection.

It is going through the process of re-connection between application client and server (CAS) when an application restarts or it is re-scheduled. Re-scheduling is a process that an application server chooses an application client, and starts and stops connection in the unit of transaction. If **KEEP_CONNECTION**, Broker parameter, is OFF, it always happens; if AUTO, it can happen depending on its situation. For details, see "Performance Tuning > Broker Configuration > Parameter by Broker."

Syntax

```
int cci_set_query_timeout (int req_handle, int milli_sec);
```

- *req_handle* : (IN) Request handle
- *milli_sec* : (IN) Timeout (unit: msec.)

Return Value

- Success : Request handle ID (int)
- Failure : Error code

Error Code

- **CCI_ER_REQ_HANDLE**

cci_set_size

Description

The `cci_set_size` function gets the number of elements for the set fetched by `CCI_A_TYPE_SET` with [cci_get_data\(\)](#).

Syntax

```
int cci_set_size(T_CCI_SET set)
```

- `set` : (IN) cci set pointer

Return Value

- Size