

ExtenDB Administration Guide

ExtenDB Parallel Server

Standard Edition & Workgroup Edition

Version 1.1

October 2006



ExtenDB Administration Guide

Table of Contents

Table of Contents	2
1 Introduction	4
1.1 Overview	4
1.2 References	4
2 Installation	5
2.1 Operating System	5
2.1.1 Additional Linux Kernel Settings	5
2.1.1.1 Number of Open Files	6
2.1.1.2 Read-Ahead	6
2.1.1.3 Access Time	6
2.2 Java	6
2.3 ExtenDB Parallel Server	7
2.3.1 Linux	7
2.3.1.1 From the RPM File	7
2.3.1.2 From extendb-1.1.tar.gz	8
2.3.1.3 Workgroup Edition Agents	9
2.3.2 MS-Windows	10
2.3.2.1 Workgroup Edition Agents	10
2.4 Underlying Database	11
2.4.1 Logging Considerations	11
2.4.2 Initializing Underlying System	11
2.4.3 Network	12
2.4.4 Configuring the Underlying Database	12
2.4.5 Starting the Underlying Database Server Process	13
2.4.6 Database User	13
2.4.6.1 Coordinator	14
2.4.7 Verify	14
3 ExtenDB Configuration	16
3.1 The xdb.config File	16
3.1.1 Sample xdb.config File	17
3.1.2 Workgroup Edition	19
3.1.2.1 Agents	20
3.2 The Metadata Database	22
3.2.1 Manual Mode	22
3.3 Starting the Coordinator and Agents	23
3.3.1 Coordinator	23
3.3.2 Agents	24
3.4 Creating User Databases	24
3.4.1 Example	25
3.5 Testing the Database	25
3.6 Starting and Stopping Databases	26
3.7 Dropping Databases	26
3.8 Planning	27
3.9 Multi-Language and Unicode Support	28
3.10 The xdb.config Reference	29
3.10.1 Licensing Configuration	29
3.10.2 Server Settings	30
3.10.3 Metadata Database Settings	31

ExtenDB Administration Guide

3.10.4	JDBC and Pool Settings	32
3.10.5	Multi-Language Support	34
3.10.6	Configuration for Underlying Database	35
3.10.6.1	Temp Table Handling	35
3.10.6.2	SQL Command Templates	36
3.10.6.3	Other settings	38
3.10.6.4	Gateway Settings for Administering Underlying Databases	39
3.10.6.5	XDBLoader settings	41
3.10.7	Data Type Mapping	42
3.10.8	Function Mapping	42
3.10.9	Logging	43
4	Users and Privileges	45
4.1	Introduction	45
4.2	Users	45
4.3	Privileges	45
5	Redundancy, Backup and Recovery	46
5.1	Redundancy	46
5.2	Load Balancing	46
5.3	Backup & Recovery	48
6	Command Reference	50
6.1	cmdline	50
6.2	createdb	52
6.3	createmddb	54
6.4	dropdb	55
6.5	execdb	56
6.6	XDBAgent	57
6.7	XDBdbstart	57
6.8	XDBdbstop	58
6.9	XDBServer	59
6.10	XDBshutdown	59
6.11	AddNode	60
6.12	DropNode	60
6.13	XDBLoader and xdbimpex	61
7	Isolation Levels and Locking	62
8	Troubleshooting	63
8.1	Installation and Configuration	63
8.2	Executing	63
9	Appendices	65
9.1	Appendix A – Metadata Database Schema	65

ExtenDB Administration Guide

1 Introduction

1.1 Overview

This document describes how to install and configure the ExtenDB Parallel Server. In addition, it includes a command reference for administering the system.

PostgreSQL, SAP-DB and MaxDB from SAP, MySQL and other trademarks and copyrights are the property of their respective owners.

1.2 References

ExtenDB Server Planning Guide
ExtenDB Import & Export Utilities
ExtenDB SQL Reference

ExtenDB Administration Guide

2 Installation

2.1 Operating System

ExtenDB can run under Linux or Windows, or any other platform that supports Java.

ExtenDB makes use of threading to a great extent, so if using Linux, please use a 2.6+ kernel or later to ensure that the Native POSIX Threading Library (NPTL) is used.

2.1.1 Additional Linux Kernel Settings

There are some additional kernel configuration values that should be modified. The underlying database requires ample shared memory, so the `kernel.shmmax`, `kernel.shmall` and `kernel.sem` values should be increased. This can be set in the `/etc/sysctl.conf` file.

The value of `kernel.shmmax` refers to the maximum size of shared memory segments in bytes, while `kernel.shmall` is the total amount of shared memory available. These values should be set fairly high; you can start with 50% of available memory and monitor the system to adjust. If a lot of system paging occurs, lower this value. Conversely, increase it if there is still a lot of available memory afterwards.

The value `kernel.shmni` is for the system wide maximum number of shared memory segments. 4096 is a reasonable value.

The kernel setting `kernel.sem` maps to four parameters: SEMMSL SEMMNS SEMOPM SEMMNI

- SEMMSL: maximum num of semaphores per id
- SEMMNS: maximum number of semaphores in system (SEMMNI*SEMMSL)
- SEMOPM: maximum num of ops per semop call
- SEMMNI: maximum number of semaphore identifiers

Sample values for 2 GB of shared memory to be set in `/etc/sysctl.conf` (adjust `shmmax` and `shmall`, depending on desired shared memory):

```
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.shmall = 2147483648
kernel.sem = 1000 128000 100 128
```

After modifying the file, execute `sysctl -p` to have these values take effect.

ExtenDB Administration Guide

2.1.1.1 Number of Open Files

Depending on your configuration, you may run into errors involving a limit to the number of open files that your operating system allows the user to have. This is particularly likely when you have several nodes in the cluster.

To change that, increase the limit of the number of files that can be open by modifying the `/etc/security/limits.conf` file, increasing the value for `nofile`.

If you do not do this, you may encounter error messages like "unable to send to nodes".

2.1.1.2 Read-Ahead

In data warehousing, tables are often scanned entirely, so having the operating system read ahead can significantly boost query times. You can increase the read-ahead size for your RAID devices with the `blockdev` command:

```
blockdev -setra 16384 <device>
```

2.1.1.3 Access Time

Very often whenever files are accessed, the operating system will update the last time of read or write access for bookkeeping. To turn off this extra unnecessary overhead in Linux, modify the `/etc/fstab` file after you create your dedicated data partitions, and add the "noatime" attribute.

2.2 *Java*

ExtenDB requires a Java Runtime Edition, version 1.5 or later. You can download the JRE from <http://java.sun.com>.

ExtenDB Administration Guide

2.3 ExtenDB Parallel Server

The instructions for installing ExtenDB Parallel Server vary, depending on the target operating system.

We will discuss configuring the actual ExtenDB system itself in the next chapter.

2.3.1 Linux

2.3.1.1 From the RPM File

1. Login as root.
2. We will install the RPM file.

```
rpm -ivh extendb-1.1-0.noarch.rpm
```

including the full path to the rpm as necessary. This will create the extendb group and user, and install in the directory `/usr/local/extendb-1.1`.

At this point, it will create the following subdirectories: bin, lib, config, and log. The bin directory contains some scripts that are wrappers to make it easier to execute ExtenDB programs, which are all java-based. The lib directory contains external jar libraries that are required by ExtenDB. Configuration information that must be customized is found in the `xdb.config` file in the config directory. Finally, a log directory is created for containing the server log files.

3. Configure environment. Modify `/usr/local/extendb-1.1/extendb_env.sh`, if necessary. It includes the following lines:

```
export XDBPATH=/usr/local/extendb-1.1

export CLASSPATH=$XDBPATH/lib/postgresql-8.1-407.jdbc3.jar

#export PATH=$PATH:$XDBPATH/bin
```

The first line just defines the XDBPATH environment variable, which is referenced in the scripts and must be set properly.

The second line defines CLASSPATH, which is used by the script when executing java programs to find additional needed external libraries. This should be set to the JDBC jar file you are using to interact with the underlying database. In the example above, it is set to a PostgreSQL JDBC driver jar. By default, a PostgreSQL driver is included. If using another database, you must obtain the appropriate driver and copy it onto the coordinator and reference it here.

ExtenDB Administration Guide

The third line can be uncommented if you wish to include the `$XDBPATH/bin` directory in the user's environment as well.

4. Users executing `exten` programs should source the `/usr/local/exten-1.1/exten_env.sh` file to have their environment set correctly:

```
source /usr/local/exten-1.1/exten_env.sh
```

or

```
./usr/local/exten-1.1/exten_env.sh
```

The `exten` user and other users may want to reference this file in a appropriate profile file, like `~/.bash_profile`. These users should also be made part of the `exten` group, if they want to execute anything other than `cmdline`. For most commands, it is required that you execute them as the user `exten`.

2.3.1.2 From `exten-1.1.tar.gz`

Instead of an `rpm` file, a tarball (`.tar.gz` file) may be used instead to accommodate manual installations.

5. If not already, change to user root:

```
su -
```

6. Create a user group `exten`:

```
groupadd exten
```

7. Create a user `exten`:

```
useradd -g exten exten
```

8. All the files will be installed under `/usr/local/exten-1.1`. Extract the downloaded gzipped tar file to `/usr/local/`

```
tar xvzf exten1_1.tar.gz -C /usr/local
```

At this point, it will create the following subdirectories: `bin`, `lib`, `config`, and `log`. The `bin` directory contains some scripts that are wrappers to make it easier to execute ExtenDB programs, which are all java-based. The `lib` directory contains external jar libraries that are required by ExtenDB. Configuration information that must be customized is found in the `xdb.config` file in the `config` directory. Finally, a `log` directory is created for containing the server log files.

9. Set ownership of the files correctly:

ExtenDB Administration Guide

```
chown extendb -R /usr/local/extendb-1.1
chgrp extendb -R /usr/local/extendb-1.1
```

10. We only allow the extendb user to execute programs, except for cmdline. We also want to set other permissions:

```
chmod 700 /usr/local/extendb-1.1/bin/*.sh
chmod 775 /usr/local/extendb-1.1/log
chmod 755 /usr/local/extendb-1.1/bin/cmdline.sh
chmod 600 /usr/local/extendb-1.1/config/*
```

11. Configure environment. Modify `/usr/local/extendb-1.1/extendb_env.sh`. It includes the following lines:

```
export XDBPATH=/usr/local/extendb-1.1

export CLASSPATH=$XDBPATH/lib/postgresql-8.1-407.jdbc3.jar

#export PATH=$PATH:$XDBPATH/bin
```

The first line just defines the XDBPATH environment variable, which is referenced in the scripts and must be set properly.

The second line defines CLASSPATH, which is used by the script when executing java programs to find additional needed external libraries. This should be set to the JDBC jar file you are using to interact with the underlying database. In the example above, it is set to a PostgreSQL JDBC driver jar. By default, a PostgreSQL driver is included. If using another database, you must obtain the appropriate driver and copy it onto the coordinator and reference it here.

The third line can be uncommented if you wish to include the `$XDBPATH/bin` directory in the user's environment as well.

12. Users executing extendb programs should source the `/usr/local/extendb-1.1/extendb_env.sh` file to have their environment set correctly:

```
source /usr/local/extendb-1.1/extendb_env.sh
```

or

```
./usr/local/extendb-1.1/extendb_env.sh
```

The extendb user and other users may want to reference this file in a appropriate profile file, like `~/.bash_profile`. These users should also be made part of the extendb group, if they want to execute anything other than cmdline. For most commands, it is required that you execute them as the user extendb.

2.3.1.3 Workgroup Edition Agents

ExtenDB Administration Guide

Workgroup Edition Only

For the Workgroup Edition only, repeat the procedure on each node that will participate in the database cluster. There is no need to do this with the Standard Edition.

2.3.2 MS-Windows

There is no automatic installation program for windows.

1. Create an exten`db` user, and login as exten`db`.
2. Choose an installation directory. Unzip `extendb1_1.zip`, which will create a subdirectory named `extendb1_1`.

At this point, it will create the following subdirectories: `bin`, `lib`, `config`, and `log`. The `bin` directory contains some scripts that are wrappers to make it easier to execute ExtenDB programs, which are all java-based. The `lib` directory contains external jar libraries that are required by ExtenDB. Configuration information that must be customized is found in the `xdb.config` file in the `config` directory. Finally, a `log` directory is created for containing the server log file.

3. Add the JDBC Driver required by your underlying database to the CLASSPATH environment variable. By default, the PostgreSQL 8.1 driver is included, so CLASSPATH can be set to `c:\extendb1_1\lib\postgresql-8.1-407.jdbc3.jar`. If using another database server, you can copy the driver to the coordinator and then reference the full path at the end of your CLASSPATH variable, including the ";" separator required in windows, if necessary. The scripts installed in the `bin` directory will already reference the jars found in the `lib` directory. They also will look for classes referred to by the CLASSPATH environment variable.

To access the environment variables in Windows: open a Windows Explorer, right click on My Computer, select Properties, select the Advanced tab, and click the Environment Variables button.

4. Set and export XDBPATH in your environment to be the base directory, at `C:\extendb1_1`.
5. Add `C:\extendb1_1\bin` to your PATH environment variable.

2.3.2.1 Workgroup Edition Agents

Workgroup Edition Only

For the Workgroup Edition only, repeat the procedure on each node that will participate in the database cluster.

ExtenDB Administration Guide

2.4 Underlying Database

ExtenDB uses PostgreSQL 8.1 as the underlying database on each of the nodes. The configuration options of ExtenDB are quite flexible and can support other databases, but PostgreSQL is the recommended database engine and the one that ExtenDB will support.

PostgreSQL should be fully and properly installed on all of the nodes that are going to make up the ExtenDB cluster. It is also recommended to install PostgreSQL on all of the nodes exactly the same way. This will make administration easier and allow you to use the `execdb` utility to simplify any administrative tasks.

It is important that your environment is set up properly so that ExtenDB can work with the underlying database and use its utilities. In the case of PostgreSQL, it is a good idea to add the PostgreSQL bin directory to your PATH environment variable, to allow access to all of its programs. You should add this to the appropriate profile file for `extenbd`, like `~extenbd/.bash_profile` Example:

```
export PATH=$PATH:/usr/local/pgsql/bin
```

2.4.1 Logging Considerations

Like other databases, PostgreSQL uses logging for point in time recovery, which they call Write Ahead Logging (WAL).

The location of these files in a subdirectory called `pg_xlog` in the data directory (the data directory is the one specified with `initdb`). Although it is not required, you may want to consider keeping these files on a separate disk, for performance reasons to avoid thrashing. This can be done by either setting up a symbolic link, or by creating a new disk partition and mounting it at `pg_xlog`, below the data directory.

2.4.2 Initializing Underlying System

PostgreSQL requires execution of the `initdb` command for initialization. Login as the user `postgres`, and execute a command like the one below. Note that if using MS-Windows, you cannot execute this as a user who has administrative privileges.

The `-D` option must be included to indicate the location of the data. In the example below, the data directory is initialized at `/XDBDATA`

```
/usr/local/pgsql/bin/initdb -D /XDBDATA
```

The `initdb` command must be executed on each node that will make up the ExtenDB cluster.

ExtenDB Administration Guide

2.4.3 Network

We must configure each of the underlying database instances to communicate with one another for the ExtenDB cluster to function properly.

With PostgreSQL, that means starting the postmaster process with the “-i” option. Without that option, only local connections will be accepted.

The next step is to configure the `pg_hba.conf` file, which is used to determine which users and clients can connect to the database. The file is found in the data directory specified by the `initdb` command earlier.

More detailed information about configuring the file can be found in the PostgreSQL documentation, but we include an example entry below:

```
# TYPE DATABASE USER CIDR-ADDRESS METHOD
host all all 192.168.75.0/24 md5
```

The above entry allows any connection from the `192.168.75.*` subnet, provided that the user name and password are valid, using the `md5` authentication method. It is a good idea to put the nodes on its own subnet, with access to the underlying databases only occurring through ExtenDB.

2.4.4 Configuring the Underlying Database

PostgreSQL offers many configuration and tuning options to help database administrators improve the performance of their system for the particular environment that it is running in. In this case, we want to tune the database for a decision support environment.

Available configuration options can be found in a file named `postgres.conf` in the data directory. The most important options that you’ll want to be concerned about appear below.

```
shared_buffers = 6000
work_mem = 16384
maintenance_work_mem = 65536
effective_cache_size = 90000

wal_buffers = 64
checkpoint_segments = 128
checkpoint_timeout = 900

constraint_exclusion = on
max_connections = 200
```

The first group of parameters should have each value multiplied by the number of gigabytes of memory. We advise you to purchase as much memory as possible. If you have 8 GB of memory per node, multiply the first group of parameters by 8.

ExtenDB Administration Guide

More details about these options appear in the PostgreSQL documentation. In the first group, the first three parameters have a direct effect on memory allocated for certain tasks. The parameter `shared_buffers` is not the total memory nor cache used for PostgreSQL; PostgreSQL relies on the OS cache. The parameter `work_mem` is used for operations like sorting and aggregation, while the parameter `maintenance_work_mem` is used for creating indexes, foreign keys, and vacuum and analyze. You may want to consider setting this value higher initially while loading up the database and building indexes, and then lowering it later. The PostgreSQL query planner is only influenced by `effective_cache_size`, it does not actually influence allocated memory.

The second group of values is used for the Write Ahead Log, and can impact performance in particular when loading the database.

The parameter `constraint_exclusion` is off by default and must be set to on in order to take advantage of PostgreSQL 8.1's feature that allows you to partition into subtables to reduce the amount of data that must be scanned for statements with qualifying WHERE clauses.

The parameter `max_connections` determines how many simultaneous connections can connect to the PostgreSQL database. It is a good idea to raise this, since ExtenDB creates pools of connections to the database, and you want to make sure you have enough connections. The exact amount for this to use may be influenced by how you configure the ExtenDB `xdb.config` file (described later), but it is a good idea to make sure that this is set sufficiently high.

2.4.5 Starting the Underlying Database Server Process

PostgreSQL offers the `pg_ctl` wrapper to start the postmaster and run it as a background process. For example, if `initdb` used `/XDBDATA` as the location of PostgreSQL data, you can start the postmaster process using the command below:

```
/usr/local/pgsql/bin/pg_ctl start -D /XDBDATA -l logfile
```

The `-l` option allows you to specify a log file for the PostgreSQL log.

2.4.6 Database User

A database user must be created on each instance that will be used by ExtenDB for connecting to the databases. This single user will always be used for connecting to the underlying database. The username and password will be required later when configuring ExtenDB's `xdb.config` file, so please make note of it. You should use the same username and password on all instances.

In the example below, we create a database user named `extendb`.

ExtenDB Administration Guide

```
/usr/local/pgsql/bin/createuser -d -E -a extenldb -U postgres -P
```

After executing this, it will prompt you for a password, and ask you to retype it. Please note this password for later. It may also give you a third password prompt. This is because of the `-U` option, where we are executing the command as database user `postgres`, the `postgres` "superuser". Just enter the `postgres` password you created when PostgreSQL was installed.

Repeat the execution of `createuser` on each node.

2.4.6.1 Coordinator

When using PostgreSQL with authentication, a password will be required for authentication. ExtenDB makes use of PostgreSQL command line utilities, so we create a `.pgpass` file in the `extenldb` user's home directory. This is used by PostgreSQL to provide passwords to connect to other servers. This only needs to be done on the coordinator.

Login as user `extenldb`.

The file's access must be restricted to the user, in this case `extenldb`. After creating the file, you need to restrict access via `chmod 600 ~extenldb/.pgpass`.

Note that on Microsoft Windows, the file is named `%APPDATA%\postgresql\pgpass.conf`, where `%APPDATA%` is the Application Data subdirectory of the user's profile. If the "Application Data" subdirectory is not viewable on your system, you may need to change the options in the Windows Explorer to allow viewing of hidden files and folders. From there, create a `postgresql` subdirectory, and then create a `pgpass.conf` file in the `postgresql` directory.

It is important that you setup `pgpass.conf` (or `.pgpass`), otherwise, executing ExtenDB scripts like `createmddb.sh` will appear to hang, because the particular PostgreSQL program will be trying to prompt for a password.

The lines in the file are to appear in the format:

```
hostname:port:database:username:password
```

Wildcards may be used. An example appears below:

```
*:*:*:extenldb:password
```

where `password` is the password we used with `createuser`. The ExtenDB process that will later run under the `extenldb` user can now connect to all nodes and use PostgreSQL utilities without requiring a password from the user. Note that user "extenldb" here is a database user, not an operating system user.

2.4.7 Verify

ExtenDB Administration Guide

Before proceeding to configuring ExtenDB, it is a good idea to verify that the underlying databases and network have been installed and configured properly. **Doing this now will help make troubleshooting your ExtenDB installation easier by eliminating the likelihood of configuration issues with the underlying database.**

Verification can be done by creating a test database on each, by running `createdb` from the coordinator.

```
/usr/local/pgsql/bin/createdb -h node1 -U extendb test
```

Here, `node1` is the host name or IP address of one of the nodes that will be in the ExtenDB cluster. Run this command from the coordinator node for each node.

If there is a problem, verify that the postmaster was started with `-i` on the node, that the node has a valid `pg_hba.conf` file, and that the `extendb` user on the coordinator has a valid `.pgpass` (or `pgpass.conf`) file. If you modify the `pg_hba.conf` file, you must restart postmaster.

3 ExtenDB Configuration

This chapter discusses how to configure ExtenDB, including creating and using databases. The first part assumes a first time configuration after installation, and the chapter concludes with more information about multi-language support and an `xdb.config` file reference.

If you have not done so already, login as the `extendb` user to modify.

3.1 The `xdb.config` File

A configuration file must be created that will determine how the ExtenDB Server is run. When the ExtenDB server is run, it reads from a file named `xdb.config` in the `config` directory for system defaults. The file needs to be configured properly to initialize the metadata database and user-created databases.

There are a myriad of options, but reasonable defaults have been set, with PostgreSQL targeted as the underlying database.

Permissions for `xdb.config` should be set to be readable only by the `extendb` user, since the file contains username and password information for connecting to the underlying databases.

The options for `xdb.metadata.*` determine where the metadata database resides. The metadata database contains information about all of the user created databases and schema info like tables, columns, data types, indexes, and constraints. Make sure that the `xdb.metadata` values are set properly, before trying to execute `createmddb`, which will create the metadata database.

This `xdb.config` file is also where you define the nodes that are used in the ExtenDB cluster, specifying the host or IP address of each node. The username and password should be set to the same values as those you used when you created the database user earlier (`createuser` with PostgreSQL).

The included default `xdb.config` file contains the most important options you may want to change. A detailed reference of all of the available configuration options appears at the end of this chapter in section 3.10, along with descriptions of mapping data types and defining and overriding SQL functions.

ExtenDB Administration Guide

3.1.1 Sample xdb.config File

To give a better idea of what a real `xdb.config` file looks like, a sample one appears below for a 4 node system. In addition, a sample `xdb.config` is found in the `config` directory of your installation.

Note that some lines contain template variables that are enclosed with curly braces, like `{dbhost}` and `{database}`. These are dynamically substituted by the ExtenDB server per database as needed- there is no need for you to replace these here.

Be sure and modify the username and password information properly for the underlying databases, as well as the hostname or IP address of each node in the database for `xdb.node.1.dbhost` through `xdb.node.4.dbhost`. Also, if logging is desired, modify the File parameter of each logger to be an absolute path to the desired target file location.

The example below is for a four-node system. Note that you can use the same host or IP address for all if you would like to create a "cluster" on a single system to just familiarize yourself with ExtenDB. It will assign a unique database name to each "node," creating 4 underlying databases. Also, you may choose to create fewer than 4 nodes if you wish. Just change the `xdb.nodecount` and comment out or remove the appropriate `xdb.node.n.dbhost` entries.

In this file 3 different logs are referenced, as can be seen by the `log4j` properties. There is a main server log, a query log (to log `SELECT` statements), and a long query log (for logging long `SELECT` statements).

```
#####
# xdb.config
#
# ExtenDB configuration file
#####

###
### Server settings
###

xdb.port=6453
xdb.maxconnections=10

###
### Node & JDBC Pool configuration
###

### Set defaults for all nodes and MetaData database.
### These can be overridden.
### Note that {dbhost} and {database} are template variables
### that will be substituted dynamically per database
```

ExtenDB Administration Guide

```
xdb.default.jdbcdriver=org.postgresql.Driver
xdb.default.jdbcstring=jdbc:postgresql://{dbhost}:{dbport}/{database}
xdb.default.dbport=5432
xdb.default.dbusername=extendb
xdb.default.dbpassword=password

### Connection thread defaults for each node
### Note that these are pooled, so the number of clients connected
### to the ExtenDB server can be greater than pool size.

xdb.default.threads.pool.initsize=5
xdb.default.threads.pool.maxsize=10

### Connectivity for MetaData database

xdb.metadata.database=XDBSYS
xdb.metadata.dbhost=localhost

### The number of nodes in cluster

xdb.nodecount=4

### Individual node info
### Set these to hostname or IP addresses of nodes

xdb.node.1.dbhost=node1
xdb.node.2.dbhost=node2
xdb.node.3.dbhost=node3
xdb.node.4.dbhost=node4

### Designate coordinator node
### In practice, the coordinator node should be the node where
### the ExtenDB database is running.

xdb.coordinator.node=1

###
### Logging Settings
###

### The log4j library is used.
### More info at http://logging.apache.org/log4j/docs/

# rootLogger. Log warnings and errors.
log4j.rootLogger=WARN, console

# Define other characteristics for console log
log4j.appender.console=org.apache.log4j.RollingFileAppender
log4j.appender.console.maxFileSize=500KB
log4j.appender.console.maxBackupIndex=10
log4j.appender.console.layout=org.apache.log4j.PatternLayout
log4j.appender.console.layout.ConversionPattern=%d{ISO8601} - %-5p %m%n
log4j.appender.console.File=/usr/local/extendb-1.1/log/console.log

# Log Server messages to the console logger
log4j.logger.Server=ALL, console

# Query logger.
```

ExtenDB Administration Guide

```
# This logs all queries sent to the database.
log4j.logger.query=INFO, QUERY
log4j.appender.QUERY=org.apache.log4j.RollingFileAppender
log4j.appender.QUERY.File=/usr/local/extendb-1.1/log/query.log
log4j.appender.QUERY.maxFileSize=500KB
log4j.appender.QUERY.maxBackupIndex=10
log4j.appender.QUERY.layout=org.apache.log4j.PatternLayout
log4j.appender.QUERY.layout.ConversionPattern=%d{ISO8601} - %m%n

# Uncomment this if you would like other SQL commands other
# than SELECT to be logged in the query logger as well
# (e.g. INSERT, UPDATE, DELETE).

#log4j.logger.command=INFO, QUERY

# A separate "long query" log may be defined to separately log queries
# that appear to be taking a long time.
# Specify the threshold in seconds at which queries will show up in the
# long query log.
xdb.longQuerySeconds=300

log4j.logger.longquery=INFO, LONGQUERY
log4j.appender.LONGQUERY=org.apache.log4j.RollingFileAppender
log4j.appender.LONGQUERY.File=/usr/local/extendb-1.1/log/longqry.log
log4j.appender.LONGQUERY.maxFileSize=500KB
log4j.appender.LONGQUERY.maxBackupIndex=10
log4j.appender.LONGQUERY.layout=org.apache.log4j.PatternLayout
log4j.appender.LONGQUERY.layout.ConversionPattern=%d{ISO8601} - %m%n
```

A request is determined to be "long" based on another `xdb.config` value, `xdb.longQuerySeconds`, which should be set to the number of seconds at which point it will be logged in the LONGQUERY log.

3.1.2 Workgroup Edition

Workgroup Edition Only

For the Workgroup Edition, some additional lines are needed in the `xdb.config` file on the coordinator. Please see the sample below:

Only for agent version

```
# Port for node's SocketCommunicator
xdb.node.1.port=6455
xdb.node.1.host=node1
xdb.node.2.port=6456
xdb.node.2.host=node2
xdb.node.3.port=6457
xdb.node.3.host=node3
xdb.node.4.port=6458
xdb.node.4.host=node4
```

```
### In practice, the coordinator node should be the node where
### the ExtenDB database is running.
```

ExtenDB Administration Guide

```
xdb.coordinator.host=node1
xdb.coordinator.port=6454

# Specify protocol types.
# Can use local connection between coordinator and node 1,
# since they are the same system
xdb.connector.0.1=0
xdb.connector.1.0=0
```

The first group assigns a port and host. Note that `xdb.node.n.host` differs from `xdb.node.n.dbhost` in that `dbhost` is where the underlying database is, and in theory could be different from the host where the agent is executing. As a practical matter, these values will be the same.

The second group configures the coordinator.

The third group specifies the connector type. The format is `xdb.connector.m.n.`, where `m` is the source node number and `n` is the destination node number. The default is 2, which is a channel connector. Setting it to 0 indicates that a local connector should be used, which is more efficient when one node is also a coordinator node. Note that node number 0 always refers to the coordinator.

3.1.2.1 Agents

Workgroup Edition Only

In the installation, in the `config` subdirectory exists two files, `xdb.config` and `xdb.config.agent`. On the non-coordinator nodes, rename `xdb.config.agent` to `xdb.config`, and edit that with appropriate settings. A sample `xdb.config` file for the agents appears below:

```
#####
#
# xdb.config - Agent
#####
#

###
### The coordinator host and port
###

xdb.coordinator.host=node1
xdb.coordinator.port=6454

###
### Logging settings
###

log4j.rootLogger=WARN, console

log4j.logger.Server=ALL, console

# A1 is set to be a ConsoleAppender.
log4j.appender.console=org.apache.log4j.RollingFileAppender
```

ExtenDB Administration Guide

```
# A1 uses PatternLayout.  
log4j.appender.console.layout=org.apache.log4j.PatternLayout  
log4j.appender.console.layout.ConversionPattern=%r [%t] %-5p %c %x - %m%n  
log4j.appender.console.File=/usr/local/extendb-1.1/log/agent.log
```

Note that this `xdb.config` file is much shorter than that of the coordinator's. Other than its own port number, the coordinator host and logging settings, all other configuration information is sent over from the coordinator, allowing for centralized maintenance of the configuration settings.

ExtenDB Administration Guide

3.2 The Metadata Database

It is a good idea to add the PostgreSQL bin directory to your PATH environment variable, to allow access to all of its programs. You may want to also add this the appropriate profile file for exten`db`, like `~extendb/.bash_profile` Example:

```
export PATH=$PATH:/usr/local/pgsql/bin
```

Also, please verify that XDBPATH environment variable has been set correctly to the base ExtenDB directory (`/usr/local/extendb-1.1`).

Before creating any user-defined databases, we must first create the metadata database to contain information about the user-created databases. Once it is created, user-created databases and all of their corresponding information about tables and columns, etc., will be stored there.

To create, use the `createmddb.sh` (`createmddb.bat` in Windows) command. It tries to create the required actual database on the underlying node and create all needed tables. The exact schema can be seen in the appendix.

This relies on configuration values stored in the `xdb.config` file, so be sure that it is set properly. Also, refer to the command reference in the next chapter that discusses `createmddb`. That means that you set all of the `xdb.metadata.*` properties as how you want them, and that `createmddb` will use these as your desired settings. The username and password used should be valid and have permission to create new databases and tables. You can use the username and password created earlier.

If everything was setup properly in `xdb.config`, you can just execute `createmddb.sh` without any arguments:

```
createmddb.sh
```

3.2.1 Manual Mode

You can also choose to create the metadata database manually with the `-m` option, instead of having `createmddb` do it for you. Instructions for doing this with PostgreSQL appear below.

If a database user was not created in section 2.4.5, you should do so now. The example below creates a user named `extendb` using PostgreSQL's `createuser` command.

```
createuser -d -E -a extendb
```

Next, create the database named XDBSYS, as designated in the `xdb.config` file as the metadata database by using the PostgreSQL command `createdb`.

ExtenDB Administration Guide

```
createdb XDBSYS -U extenldb
```

Now, initialize the ExtenDB metadata database. Note that `-m` is used here, indicating that the physical database already exists; we just need to initialize it and create the metadata tables in it. Make sure a valid username and password are set in the `xdb.metadata` options of the `xdb.config` file.

```
createmddb.sh -m
```

Now the metadata database is ready, which can be verified using the PostgreSQL command `psql`:

```
psql -U extenldb -d XDBSYS
```

You should plan on backing up the metadata database regularly. Whereas with other databases, an incremental backup may make the most sense, the metadata database will be relatively small, so a complete backup should be done.

3.3 Starting the Coordinator and Agents

3.3.1 Coordinator

We are now ready to start the XDBServer process so that we can create user databases. XDBServer can be started with no arguments if no databases have been created yet:

```
XDBServer.sh
```

This will start the server in the background. If there is a problem with XDBServer, please check the log files in the log directory and verify the configuration in `xdb.config`. Note that when executing the XDBServer process, you may need to modify the parameters that Java uses, increasing the maximum amount of memory specified in the `XDBServer.sh` launch script, depending on your requirements and system configuration.

In Windows, you will execute the `XDBServer.bat` file. It will remain executing in the window.

When executing XDBServer in the future, you can include a list of previously created databases to bring online with the `-d` option. That way, you will not need to separately execute `XDBdbstart`.

```
XDBServer.sh -d XTEST
```

ExtenDB Administration Guide

3.3.2 Agents

Workgroup Edition Only

If using the Workgroup Edition, you will want to start the agents on all of the nodes. Perform this as user `extendb`.

The ExtenDB agent is started by calling the `XDBAgent.sh` wrapper script in the `extendb` bin directory. It expects the `-n` argument, followed by the designated node number this will act as in the cluster. Example:

```
XDBAgent.sh -n 4
```

We recommend you start the server on the coordinator first (`XDBServer.sh`), before starting the agents. If the `XDBServer` process is stopped and restarted, it should reestablish connections with the running agents. Similarly, if an agent is stopped and restarted, `XDBServer` will detect that and reestablish a connection.

`XDBServer` will log the event that an agent has successfully connected to it, so if there is a problem in creating and using databases, please read the coordinator logs to pinpoint the source of the problem.

3.4 Creating User Databases

Now that the metadata database has been created and `XDBServer` is executing, you can create your own databases.

First, configure the individual nodes in the `xdb.config` file that you will be using if you have not already. You should have also installed the underlying database product like PostgreSQL on all of those nodes, with the database server running. Note that it is a good idea to have all of the nodes installed and configured exactly the same way to make administration easier. If you make any changes to `xdb.config`, you will need to restart `XDBServer`.

It should also be pointed out that you could also have a system where one of the nodes in the ExtenDB system both participates as a member of user database nodes as well as contains the metadata database.

You create a database by using the `createdb.sh` utility. When a database is created, it adds the appropriate information to the metadata database. The `createdb.sh` command will also try and create the database on the underlying nodes if you wish, which is recommended. Otherwise, use the `-m` (manual) option, which will only update the metadata database information without trying to create any databases on the nodes.

An ExtenDB DBA username needs to be specified for the database with the `-u` option. The command `createdb.sh` will then create this user automatically. A password can be specified with the `-p` option. If `-p` is not present, it will prompt the

ExtenDB Administration Guide

user for a password. To clarify, this is a user at the ExtenDB level, and not for communicating with the underlying databases. **Please note the username and password so that you will be able to access the database, such as when later using the cmdline utility.**

See the `createdb` command in the section of the Command Reference section of this document for more information.

3.4.1 Example

We show two examples.

The first example will create a database named XTEST using the template `xdb.gateway.createdb` (defaults are already set for PostgreSQL) setting in the `xdb.config` file, on 4 nodes. The physical databases will be named XTESTN1, XTESTN2... XTESTN4 on the nodes, corresponding to their node numbers.

```
createdb.sh -d XTEST -u admin -p <password> -n 1,2,3,4
```

Note that if you are prompted by a password even with `-p`, it is the underlying tool, like `psql` that is prompting you for a password. This means you are executing `createdb` under a user where a trusted PostgreSQL environment has not been configured. Be sure that it is configured for user `extendb`, and execute the command as user `extendb`.

Manual mode

If desired, the database can also be set up by hand, where the databases are first created on the individual nodes using PostgreSQL's `createdb` command. Care needs to be taken to name them with the desired database name followed by "N" and the node number (`<dbname>N1` on node 1, `<dbname>N2` on node 2, etc.). Once that is done, `createdb.sh` can be called using the `"-m"` option to wire it up and update the metadata information.

```
createdb.sh -d XTEST -u admin -p <password> -n 1,2,3,4 -m
```

3.5 Testing the Database

Once you have successfully created a database, you are ready to test using the command line utility.

Note that XDBServer will automatically bring the created database online and accept connections to it when you execute `createdb.sh`.

With the server running ok, execute `cmdline.sh`, specifying a database and valid username and password, such as:

```
cmdline.sh -d XTEST -u admin -p <password>
```

ExtenDB Administration Guide

If you were able to connect ok you should receive a command prompt like the following:

ExtenDB->

Try creating a table. The following command creates a table mytable1 and specifies that rows should be partitioned according to the column col1:

```
CREATE TABLE mytable1 (col1 INT) PARTITIONING KEY col1;
```

Try and insert some data:

```
INSERT INTO mytable1 VALUES (1);  
INSERT INTO mytable1 VALUES (2);
```

Select:

```
SELECT * FROM mytable1;
```

If everything looks ok, you can drop the table:

```
DROP TABLE mytable1;
```

3.6 Starting and Stopping Databases

The commands XDBdbstart and XDBdbstop communicate with the XDBServer process and can be used to bring ExtenDB databases online or offline.

Example:

```
XDBdbstart.sh -d XTEST -u admin -p <password>
```

```
XDBdbstop.sh -d XTEST -u admin -p <password>
```

3.7 Dropping Databases

You can drop databases using the dropdb command.

If the database is online, dropdb will fail, so you should first bring it offline with XDBdbstop:

```
XDBdbstop.sh -d XTEST -u admin -p <password>
```

An example for dropping the XTEST database appears below:

```
dropdb.sh -d XTEST -u admin -p <password>
```

ExtenDB Administration Guide

It will attempt to drop the databases on the underlying nodes, as well as clean out any metadata information in the XDBSYS database. Please see `dropdb` in the Command Reference section of this document for more information.

If dropping fails for some reason, you may want to try again with “-f” (force) option to continue and try and remove all metadata information from the metadata database, even if it failed to drop a database on a node.

3.8 Planning

You are now ready to create your own databases. **Please refer to important information in the Planning Guide for important information regarding determining your database schema and partitioning strategies before creating tables.** A poorly thought out schema will result in less than optimal performance.

ExtenDB Administration Guide

3.9 Multi-Language and Unicode Support

The ExtenDB Parallel Server supports international character sets, provided the chosen underlying database supports it as well and is configured properly. For the current version, however, the ExtenDB Metadata database does not support international identifiers, so all object names such as tables and columns must be standard identifiers using single-byte characters.

The following steps must be done in order to configure and support this properly

1. The underlying database needs to be configured properly. **Note: some databases require this to be configured only during installation, so take care to install the node databases properly.** As an example, in SAP-DB, the parameter `_UNICODE` should be set equal to YES. In PostgreSQL, unicode is enabled by default.
2. The JDBC Driver for the underlying database that ExtenDB uses may require additional parameters. In SAP-DB for example, if you wish to enable unicode support, the JDBC URL needs to include "unicode=true". To set the JDBC Driver string template, update the template urls in the `xdb.config`. Note, in SAP-DB that you cannot connect to non-Unicode databases if a connection string has `unicode=true`, so you cannot use both unicode and non-unicode databases simultaneously with the same ExtenDB server instance.
3. The ExtenDB server must be configured. If you intend to use international characters from some specific character set, you can specify its name in `xdb.config` configuration file, e.g.:

```
xdb.charset=windows-1252
```

The default for `xdb.charset` is ISO-8859-1.

If Unicode is desired, including support for various clients using different character sets, then add the following to the `xdb.config` file:

```
xdb.unicode=yes
```

4. If you have configured the ExtenDB server to use a specific character set or Unicode, you must specify the same for the ExtenDB Driver. Use the ExtenDB JDBC connection string to do this. For example:

```
jdbc:xdb:<dbname>:localhost:6453?user=<user>&password=<password>&xdb.unicode=yes
```

```
jdbc:xdb:<dbname>:localhost:6453?user=<user>&password=<password>&xdb.charset=windows-1252
```

ExtenDB Administration Guide

3.10 The *xdb.config* Reference

If you need to customize your particular installation, you can change the `xdb.config` file.

A table appears on the following pages describing all of the possible configuration options in `xdb.config`.

3.10.1 Licensing Configuration

Configuration Value	Default	Description
<code>xdb.license.companyname</code>	ExtenDB	The company name for the license. For the free 4-node version, ExtenDB is used.
<code>xdb.license.key</code>	*****	The license key indicating the length of time and number of nodes allowed under the current license. No license key is needed for the free 4-node version.

ExtenDB Administration Guide

3.10.2 Server Settings

Configuration Value	Default	Description
log4j.configuration		Optional configuration file for logging preferences in log4j format. Alternatively, configuration properties may also be specified directly in the xdb.config file. More information about log4j appears later in this chapter.
xdb.coordinator.node		The node to use for combining results from underlying nodes. In practice, it is the node that corresponds to where the ExtenDB server is running, whether or not it is a dedicated coordinator or not.
xdb.longQuerySeconds	300	The threshold in number of seconds at which a query is determined to be a long running query, and logged in the long query log, if enabled.
xdb.maxconnections	50	Maximum number of client connections to ExtenDB to allow at a time.
xdb.nodecount		The number of underlying nodes in the ExtenDB cluster
xdb.port	6453	The port number that ExtenDB will use to allow client processes to connect to. If you have more than one XDBServer running on the same coordinator node (one for development, one for testing, for example), make sure they use different ports.

ExtenDB Administration Guide

3.10.3 Metadata Database Settings

Configuration Value	Default	Description
xdb.metadata.jdbcdriver		The class name of the JDBC Driver to use with underlying metadata database
xdb.metadata.jdbcstring		A template JDBC url to use to connect to the underlying database. Parameter values may be dbhost and database. Example: jdbc:sapdb://{dbhost}/{database}?timeout=0
xdb.metadata.dbhost		The host name or IP address or the server that contains the metadata database. In practice, it will often be the same one as where the ExtenDB server runs.
xdb.metadata.dbport		The port to connect to for the underlying database
xdb.metadata.database		The name of the metadata database, e.g. XDBSYS
xdb.metadata.jdbcuser		The user to use when connecting to the metadata database
xdb.metadata.jdbcpassword		The password to use when connecting to the metadata database

ExtenDB Administration Guide

3.10.4 JDBC and Pool Settings

Configuration Value	Default	Description
xdb.default.jdbcdriver	org.postgresql.Driver	The default driver name to use for all connections
xdb.default.jdbcstring	jdbc:postgresql://{dbhost}:{dbport}/{database}	The default jdbc url to use for all connections
xdb.default.dbport	5432	The default port to use when connecting to the underlying database.
xdb.default.dbusername		The default username to use for all connections
xdb.default.dbpassword		The default password to use for all connections.
xdb.default.threads.pool.initsize	5	Default thread pool size for all nodes.
xdb.default.threads.pool.idle	600000	Default idle time in milliseconds.
xdb.default.threads.pool.maxsize	10	Default max thread pool size for all nodes.
xdb.default.threads.pool.timeout	60000	Default pool timeout for all nodes in milliseconds.
xdb.jdbc.coordinator.pool.initsize	xdb.default.threads.pool.initsize	The initial size of the coordinator connection pool.
xdb.jdbc.coordinator.pool.maxsize	xdb.default.threads.pool.maxsize * 0.8	The maximum size of the coordinator connection pool
xdb.jdbc.pool.maxsize	xdb.default.threads.pool.maxsize	Maximum number of connections per JDBC pool for underlying node. Note that this is an important value for managing simultaneous connections. You may still allow a large number of client connections via xdb.maxconnections, but you might want to limit how many simultaneous queries can execute on the underlying databases at the same time by limiting the pool here. In addition, depending on your underlying database, you might have licensing restrictions that dictate a smaller pool size. The ExtenDB Scheduler will handle sharing and managing of these pools.
xdb.jdbc.pool.initsize	xdb.default.threads.pool.initsize	Initial JDBC pool size
xdb.jdbc.pool.idle	xdb.default.threads.pool.idle	Default idle timeout value for JDBC Pool, in milliseconds. After this time, connections are released and pool is shrunk.
xdb.jdbc.pool.timeout	xdb.default.threads.pool.timeout	Maximum time to wait for available jdbc connection from pool
xdb.jdbc.pool.largequery.count	2	The number of connections to allow for "large queries". This

ExtenDB Administration Guide

		allows us to reserve some connections for low-cost commands, in effect reserving connections for fast operations without having to have them wait if all connections are being used executing large queries. This also provides a mechanism for allowing the DBA to be able to connect and administer the server if it is very busy.
<code>xdb.jdbc.pool.largequery.threshold</code>	25000	The cost at which a query will be designated as a "long" query. See <code>xdb.jdbc.largequery.count</code> for more details.
<code>xdb.node.n.dbhost</code>		The host address of the underlying database that the node is using, where <code>n</code> is the node id. In practice, the host will be same as the node itself, but that is not required.
<code>xdb.node.n.dbport</code>		The default port to use when connecting to the underlying database.
<code>xdb.node.n.jdbcdriver</code>		The JDBC driver to use for node <code>n</code> , where <code>n</code> is the node id.
<code>xdb.node.n.jdbcstring</code>		The JDBC URL template string used to connect to node <code>n</code> , where <code>n</code> is the node id.
<code>xdb.node.n.threads.pool.maxsize</code>	10	Maximum thread execution pool size for node <code>n</code> . In practice, this should be set to the same value as <code>xdb.jdbc.pool.initsize</code> .
<code>xdb.node.n.threads.pool.initsize</code>	5	Initial thread pool size for node id <code>n</code> .
<code>xdb.node.n.threads.pool.idle</code>	600000	In milliseconds, how long to allow a thread to be active with no activity before destroying it.
<code>xdb.node.n.threads.pool.timeout</code>	60000	In milliseconds, how long to wait on an available thread.
<code>xdb.nodeFetchSize</code>	1000	The fetch size to use on the underlying connection.

ExtenDB Administration Guide

3.10.5 Multi-Language Support

Configuration Value	Default	Description
xdb.charset	ISO-8859-1	The character set to use.
xdb.Unicode	false	xdb.unicode false Whether or not to use Unicode.

ExtenDB Administration Guide

3.10.6 Configuration for Underlying Database

3.10.6.1 Temp Table Handling

Configuration Value	Default	Description
xdb.sqlcommand.createTempTable.start	CREATE TEMP TABLE	Start of command for CREATE TABLE statement for creating temp table. Allows for alternate syntaxes like "CREATE TEMP TABLE". See also xdb.tempTablePrefix. This is for temp tables that the end user specifies.
xdb.sqlcommand.createTempTable.suffix	WITHOUT OIDS	Suffix to add at the end of CREATE statements for temp tables. This can be used to allow disabling of logging information on the underlying database and greatly improve performance, since temporary tables are used internally by ExtenDB. In SAP-DB, this value should be set to IGNORE ROLLBACK. This is for temp tables that the user specifies.
xdb.sqlcommand.createGlobalTempTable.start	CREATE TABLE	This is used when creating internal temp tables by the database for query processing. For PostgreSQL/Bizgres, real tables are used, in order to access them across sessions.
xdb.sqlcommand.createGlobalTempTable.suffix	WITHOUT OIDS	The suffix to use when creating an internal temp table used for query processing. Default is an empty string.
xdb.tempTablePrefix	TMPT	Temporary table prefix to use in underlying database. Various databases have different conventions, like "TEMP." or "#". Warning- be careful about assigning. On startup, ExtenDB will try and delete any tables that start with this name, in case permanent tables were used and tables were not cleaned up due to a server error.
xdb.allowtemptableindex	true	Whether or not the underlying database allows the support of indexes on temporary tables.
xdb.tempTableSelect	select tablename from pg_tables	If "fake" temp tables are used on the underlying database

ExtenDB Administration Guide

	<pre>where tablename LIKE '{xdb.tempTablePr efix}%'</pre>	(instead of actual temp tables), this value can be used to determine how to obtain a list of temp tables on the underlying nodes, to purge any tables at XDBServer start-up, in case there are any remaining from previous execution that were not cleaned due to an error.
--	---	---

3.10.6.2 SQL Command Templates

Some SQL commands like ALTER TABLE have syntaxes that are more likely to vary between different databases. Below are command templates that can be overridden. The defaults that appear below are for PostgreSQL/Bizgres.

Configuration Value	Default	Description
xdb.sqlcommand .altertable.addcolumn	add {colname}	For adding columns within an ALTER TABLE command
xdb.sqlcommand .altertable.addprimary	alter table {table} add constraint {constr_name} primary key({col_list})	For adding a primary key to a table
xdb.sqlcommand .altertable.addforeignkey	alter table {table} add constraint {constr_name} foreign key {col_list} references {reftable}({col_m ap_list})	For adding a foreign key to a table
xdb.sqlcommand .altertable.dropcolumn	alter table {table} drop {colname}	For dropping a column
xdb.sqlcommand .altertable.dropconstraint	drop constraint {constr_name}	Template for dropping a constraint within ALTER TABLE command.
xdb.sqlcommand .altertable.dropprimary	alter table {table} drop constraint {constr_name}	For dropping a primary key from a table
xdb.sqlcommand .altertable.modifycolumn	alter table {TableName} alter {ColumnName} type {ColumnType}	For modifying a column's type
xdb.sqlcommand .altertable.modifycolumn.dropdefault	alter {column} drop default	Used to indicate that a columns default should be removed
xdb.sqlcommand .altertable.modifycolumn.dropnotnull	alter {column} drop not null	Used to indicate a column should no longer be NOT NULL.
xdb.sqlcommand .altertable.modifycolumn.setdefault	alter {column} set default {expr}	Used to modify the default value of a column

ExtenDB Administration Guide

xdb.sqlcommand .altertable.modifycolumn.setnotnull	alter {column} set not null	Used to indicate a column should be set to not null
xdb.sqlcommand.analyze.template.table	ANALYZE {table}	The UPDATE STATISTICS or ANALYZE command template to run on the underlying database to update internal statistics on a table used by the optimizer.
xdb.sqlcommand.analyze.template.column	ANALYZE {table} ({column_list})	Other ANALYZE command template when columns are also specified.
xdb.sqlcommand.dropindex	drop index {index_list}	Command to use when dropping indexes
xdb.sqlcommand.renametable.template	ALTER TABLE {oldname} RENAME TO {newname}	Format of command to rename table. Example: For PostgreSQL (default): ALTER TABLE {oldname} RENAME TO {newname} For SAP-DB/MaxDB: RENAME TABLE {oldname} TO {newname}
xdb.sqlcommand .updatestatistics.template.table	VACUUM ANALYZE {table}	The UPDATE STATISTICS or ANALYZE command template to run on the underlying database to update internal statistics on a table used by the optimizer. Example template: UPDATE STATISTICS {table}
xdb.sqlcommand .updatestatistics.template.column	VACUUM ANALYZE {table} ({column_list})	The UPDATE STATISTICS or ANALYZE command template to run on the underlying database to update internal statistics on a table's column used by the optimizer. It will process those columns that were explicitly specified in the ExtenDB command. Example template: UPDATE STATISTICS COLUMN {column_list} FOR {table}
xdb.sqlcommand .updatestatistics.query	SELECT stadistinct FROM pg_statistic s, pg_class c, pg_attribute a WHERE s.starelid = c.oid AND s.staattnum = a.attnum AND c.relname = '{table}' AND a.attname = '{column}'	When calculating statistics, the server will try and run the corresponding command on the underlying database, but when finished, it may be able to determine the selectivity from the underlying database without having to resort to calculating it itself. If this parameter is set, it defines a command to obtain the statistics from the underlying database. Examples:

ExtenDB Administration Guide

		<p>PostgreSQL:</p> <pre>SELECT statdistinct FROM pg_statistic s, pg_class c, pg_attribute a WHERE s.starelid = c.oid AND s.staattnum = a.attnum AND c.relname = '{table}' AND a.attname = '{column}'</pre> <p>SAP-DB:</p> <pre>SELECT distinctvalues FROM {dbuser}.OPTIMIZERSTATISTICS WHERE tablename = '{table}' AND columnname = '{column}'</pre> <p>Ingres:</p> <pre>SELECT num_unique FROM iistats WHERE table_name = {table} AND column_name = {column}</pre>
xdb.sqlcommand.update.correlatedstyle	2	This is need for the UPDATE command to work properly for correlated updates. Set it to 1 for SAP-DB/MaxDB, 2 for PostgreSQL/Bizgres.
xdb.sqlcommand.vacuum.template.table	VACUUM {vacuum_type} {table}	For PostgreSQL and Bizgres, the command to execute for vacuuming.
xdb.sqlcommand.vacuum.analyze.template.table	VACUUM {vacuum_type} ANALYZE {table}	For PostgreSQL and Bizgres, the command to execute for vacuuming with analyze.
xdb.sqlcommand.vacuum.analyze.template.column	VACUUM {vacuum_type} ANALYZE {table} ({column_list})	For PostgreSQL and Bizgres, the command to execute for vacuuming with analyze with columns specified.

3.10.6.3 Other settings

xdb.combined.resultset.buffer	1000	Default read-ahead buffer per ResultSet when combining
xdb.connectiontest.statement	select 1	Statement to run against backend to verify that connection is still good.
xdb.connectiontest.createtable		Statement to run to (if not null) to create a table to run a query against to test the connection via xdb.connectiontest.statement.
xdb.index.useAscDesc	false	Whether or not it is ok to use ASC or DESC in indexes.
xdb.locks.readcommitted.mode	S	If using an isolation mode of read committed (the default), this can be fine tuned further. S (Strict) indicates that only

ExtenDB Administration Guide

		one UPDATE or DELETE statement may be executing at a time per table, which also helps prevent deadlocks. Setting this to L (Loose) allows for concurrent UPDATE and DELETE statements.
xdb.savepointType	S	T = subtransaction, S = Savepoints. Although savepoints from the user's point of view is currently not supported, this is used in working with the underlying database.
xdb.sort.case.sensitive	false	Whether or not the underlying database sorts in a case sensitive manner.
xdb.sort.nulls.style	2	How nulls are handled in sorting on the underlying database. 0-Nulls always at start 1-Nulls always at end 2-Null greater than not null 3-Null less than not null
xdb.sort.trim	true	Whether or not leading spaces are ignored in sorting
xdb.sql.usecrossjoin	true	Whether or not to use CROSS JOIN syntax for Cartesian products. If overridden to false, syntax used will be "table1, table2" instead.
xdb.subsecondPrecision	0	The number of digits for subsecond precision that the underlying database supports.
xdb.xrowid.type	DECIMAL(31,0)	The xrowid settings allow for customization for databases that support varying levels of precision. xrowid is the ExtenDB internal unique tuple identifier.
xdb.xrowid.SQLtype	3	java.sql.Types.DECIMAL
xdb.xrowid.length	0	
xdb.xrowid.precision	31	
xdb.xrowid.scale	0	

3.10.6.4 Gateway Settings for Administering Underlying Databases

Configuration Value	Default	Description
xdb.gateway.createdb	createdb -h {dbhost} -U {dbusername} -O {dbusername} {database}	Template command for creating a new database on underlying database
xdb.gateway.dropdb	dropdb -h {dbhost} -U {dbusername} {database}	Template command for dropping database on nodes
xdb.gateway.startdb		Template command for starting underlying database on nodes. Note that in PostgreSQL and Bizgres, there is no need to

ExtenDB Administration Guide

		bring individual databases online; databases are accessible provided the postmaster process is running
xdb.gateway.stopdb		Template command for stopping underlying database on nodes. Note that in PostgreSQL and Bizgres, there is no need to bring individual databases offline; databases are accessible provided the postmaster process is running

ExtenDB Administration Guide

3.10.6.5 XDBLoader settings

Configuration Value	Default	Description
xdb.loader.header.columnseparator		Optional separator for header for output file if exporting.
xdb.loader.header.template		Optional template for output file.
xdb.loader.footer.columnseparator		Optional separator for file footer for output file if exporting.
xdb.loader.footer.template		Optional file footer template if exporting.
xdb.loader.nodewriter.columninfo	{{columns}}	Column info template to use if column names explicitly specified on load
xdb.loader.nodewriter.columninfo.none		Template to use if no column names present
xdb.loader.nodewriter.delimiterinfo	DELIMITER AS '{delimiter}'	Template to be used within xdb.loader.nodewriter.template, for specifying passing along delimiter information.
xdb.loader.nodewriter.delimiterinfo.none	DELIMITER AS '\\\ '	delimiterinfo template to use when the user did not specify any delimiter. Default null.
xdb.loader.nodewriter.template	psql -h {dbhost} -p {dbport} -d {database} -U {dbusername} -a -e -E -c \"COPY {table} {columninfo} FROM STDIN WITH NULL AS '' {delimiterinfo}\"	This is used for bulk loading data into the underlying database, and describes the template of the command to use. Note that another template, delimiterinfo can be included here. See xdb.loader.nodewriter.delimiterinfo for more information
xdb.loader.nodewriter.rowdelimiter	\n	Row separator
xdb.loader.row.nullvalue		Null value indicator in loading data
xdb.loader.row.quote	(none)	Indicates that strings are to be quoted with the specified character when loading data.
xdb.loader.row.quote.escape	(none)	Quote escape character
xdb.loader.row.template	{value_list}	Row template for output file
xdb.loader.row.columnseparator	,	The default column separator character to use when loading in data.
xdb.use_load_for_step	y	Indicates if a native database bulk loader utility should be used for handling intermediate results. This is highly recommended, so this value should be set to "y".

ExtenDB Administration Guide

3.10.7 Data Type Mapping

ExtenDB also includes the ability to map SQL data types, to allow for flexibility with various underlying databases, since the different databases sometimes name things differently than standard ANSI. Below appears the data types supported and their default mappings, which can be overridden in the `xdb.config` file.

Numeric types:

```
xdb.sqltype.integer.map=INT
xdb.sqltype.smallint.map=SMALLINT
xdb.sqltype.boolean.map=BOOLEAN
```

Floating point types (parameter "length" available):

```
xdb.sqltype.float.map=FLOAT ({length})
xdb.sqltype.real.map=REAL ({length})
```

xdb.sqltype.double.map=DOUBLE PRECISION

Fixed point types (parameters "precision" and "scale" available):

```
xdb.sqltype.fixed.map=FIXED ({precision}, {scale})
xdb.sqltype.numeric.map=NUMERIC ({precision}, {scale})
xdb.sqltype.decimal.map=DEC ({precision}, {scale})
```

Character types (parameter "length" available):

```
xdb.sqltype.char.map=CHAR ({length})
xdb.sqltype.varchar.map=VARCHAR ({length})
xdb.sqltype.nchar.map=CHAR ({length}) UNICODE
xdb.sqltype.nvarchar.map=VARCHAR ({length}) UNICODE
```

Date and Time types:

```
xdb.sqltype.time.map=TIME
xdb.sqltype.date.map=DATE
xdb.sqltype.timestamp.map=TIMESTAMP
```

3.10.8 Function Mapping

ExtenDB's recognized SQL is ANSI-92 in nature, along with the most common functions found in most databases. However, it is possible to also use additional functions that are supported by your underlying database when issuing SQL commands. This also includes any stored procedures or user-defined functions used, with the caveat that these should usually not access any tables directly because each will be executed in isolation on the particular node.

ExtenDB Administration Guide

By default, any functions not recognized will be executed on the underlying database directly. In some queries, it is necessary for ExtenDB to know the return type. In those cases, it is best to define these in the `xdb.config` file.

In addition, it is possible to override the definition for an ExtenDB-recognized function and map it to the equivalent function on the underlying database.

To either define or override functions, use `xdb.sqlfunction`, followed by the function name, followed by following settings.

template	Used only if recognized function is being overridden or unknown function is being defined, maps the function to the underlying database
returntype	The return sql data type of the function
paramcount	The number of parameters the function takes
argn	Where n is 1, 2... the argument type

The SQL data types recognized are:

CHAR, VARCHAR

DATE, TIME, TIMESTAMP
BYTE, SMALLINT, INTEGER, BIGINT
ANYINT, FLOAT, REAL, DOUBLE, NUMERIC, DECIMAL

In addition, ANYCHAR, ANYDATETIME, ANYINT and ANYNUMBER are short-hand notations when more than one type is permissible:

ANYCHAR = CHAR|VARCHAR
ANYDATETIME = DATE|TIME|TIMESTAMP
ANYINT = BYTE|SMALLINT|INTEGER|BIGINT
ANYNUMBER = ANYINT|FLOAT|REAL|DOUBLE|NUMERIC|DECIMAL

For example, to define a function for `SUBDATE(date, number_of_days)` function:

```
xdb.sqlfunction.subdate.template=DATE({arg1})-INTERVAL '{arg2}
days'
xdb.sqlfunction.subdate.returntype=DATE
xdb.sqlfunction.subdate.paramcount=2
xdb.sqlfunction.subdate.arg1=DATE
xdb.sqlfunction.subdate.arg2=ANYNUMBER
```

3.10.9 Logging

ExtenDB Administration Guide

ExtenDB uses a popular library called log4j to implement its logging functionality. More detail can be found online here:

<http://logging.apache.org/log4j/docs/index.html>.

There are a few defined “loggers” that are used: console, Server, QUERY, and LONGQUERY. The console logger is used for errors and warnings. Server is used for significant server events. QUERY allows you to log all SQL requests to the database, which can be useful in troubleshooting. LONGQUERY allows you to log those requests which seem to be taking a long time to execute, which is useful for a DBA to get quickly to the source of which queries seem to be taking the most time to execute.

A request is determined to be “long” based on another `xdb.config` value, `xdb.longQuerySeconds`, which should be set to the number of seconds at which point it will be logged in the LONGQUERY log.

4 Users and Privileges

4.1 Introduction

The ExtenDB Parallel Server supports creation of users and privileges.

It is important to distinguish between users at the ExtenDB level, and those of the underlying databases. ExtenDB does not in turn try and create those same users on the underlying databases. It always accesses the underlying database with the single user defined in the `xdb.config` file. ExtenDB manages its own users and privileges for allowing access to the tables.

4.2 Users

There are 3 classes of users: DBA, RESOURCE, and STANDARD. DBA users have Database Administration privileges. RESOURCE users can create tables. STANDARD users cannot create tables, but can access the database.

Users can be created with the CREATE USER command, and can be modified and dropped with the ALTER USER and DROP USER commands, respectively.

4.3 Privileges

A user must be granted access to a table before being able to access it. By default, a user who creates a table has all privileges on that table.

Privileges can be set on tables by using the GRANT and REVOKE commands.

More details on using these commands can be found in the ExtenDB SQL Reference manual.

The following types of privileges are available:

- SELECT
- INSERT
- UPDATE
- DELETE
- REFERENCES
- INDEX
- ALTER

Note that in the current version, ExtenDB does not yet support ROLES.

5 Redundancy, Backup and Recovery

5.1 Redundancy

The current version of ExtenDB has no built-in redundancy, but this is a feature that will be added in the near future. Keep in mind that the component most likely to fail is going to be a hard disk, and by using a RAID configuration like RAID 0+1 or RAID-5, you are well protected against such a failure.

In any event, if node redundancy is required right away, it still is possible to create your own standby nodes, depending on the underlying database being used. Both the user-created databases and the metadata database should in that case be replicated.

For replication, you will need to rely on the tools available on the underlying database. For PostgreSQL, a non-commercial implementation is called Slony, about which more information can be obtained from <http://www.postgresql.org>.

If replicating, a stand-by node can be switched to be part of the ExtenDB cluster simply by changing the node information in `xdb.config`.

To make efficient use of the nodes in the cluster, you should consider creating the replicated copies of one node on another node. For example, node 1's databases are replicated to node 2, node 2's to node 3, and so on.

5.2 Load Balancing

ExtenDB provides some amount of "load balancing" by virtue of the fact that it parallelizes queries and leverages multiple nodes. This allows queries over large amounts of data to execute much faster than they would if they were just on a single system.

Also, above, we recommend creating stand-by databases on other nodes in the cluster that are also being used, for efficiency and cost savings, especially if OLTP activity is low.

Still, if standby nodes were created manually in your system and you wish to make use of them for querying for better throughput, it is possible to do so by hand, with some effort, however. (Built-in load balancing is planned for future support.)

For the current version, you can execute multiple coordinators while keeping the following in mind:

1. Your schema should be static. If doing schema changes, you should disable access temporarily to the second cluster until synchronized.

ExtenDB Administration Guide

2. An IP-based load balancer that supports sticky connections can be used to distribute the load amongst the coordinators.

ExtenDB Administration Guide

5.3 Backup & Recovery

How backups are performed will depend greatly on the underlying database you are using. It is best to rely on the tools of the underlying database to backup nodes. That allows you to do restores on individual nodes and achieve parallelism while performing backups, as opposed to just doing a complete dump of all the data on all the nodes to a single destination.

Many databases have the concept of full backups (backs up everything), incremental backups and log file backups. This allows for different backup schedules. For example, you may wish to do a complete backup of all of the nodes once a week, and incremental backups every evening, or after a nightly load.

If you are in an environment where ExtenDB houses a data warehouse or data mart that where no update or delete activity occurs, with just periodic loads, you can also have a backup schedule with periodic full backups of the database combined with backups of the regular import files.

Performing the backups can be done directly on the nodes using the database tools available for the underlying database. Alternatively, the `execdb` command can be used, which allows for the execution of the (nearly) exact same command on all of the underlying nodes. It makes use of the configuration value set for `xdb.config` file for the particular database product being used.

An example for backing up PostgreSQL locally on each host appears below, assuming a secure environment has been set up to use ssh (secure shell):

```
execdb.sh -c "ssh -h {dbhost} 'pg_dump -h {dbhost} -U {dbusername}
{database} -f /data/back/{database}.dump'" -d mydatabase -u extenldb -p password
```

An example for using SAP-DB appears below:

```
execdb -u DBA -p DBAPW -d PROD1 -c "dbmcli -n {dbhost} -d {database}
-u {dbusername},{dbpassword} -i backup.xdb"
```

where `backup.xdb` contains the SAP-DB commands:

```
util_connect
medium_put dataproduct1 /var/opt/sapdb/back/dataproduct1.backup FILE DATA
backup_start dataproduct11 DATA
util_release
```

To recover a database, there are a couple of scenarios to consider. Typically, the problem will just be on a single node due to a hardware or software failure. If that is the case, use the tools of the underlying database to restore a complete backup if necessary, and any incremental backups and logs, as the case may be. See the documentation for your particular database product for details on how to do this.

ExtenDB Administration Guide

Another scenario is that a recovery is required because of human error. In this case, all of the nodes may very well be affected and will need to be restored.

ExtenDB Administration Guide

6 Command Reference

In this section, commands used to administer the ExtenDB Parallel Server are described.

All of these are from classes in the ExtenDB java jar files, but can be accessed more conveniently via the script wrappers in the bin directory. If using Linux or other Unix variant, append a ".sh" at the end of the commands listed here. If using Windows, append ".bat" to execute.

Note that the scripts invoke java and specify the amount of memory to use for the JVM. In the event that you encounter the OutOfMemoryException, just increase the values specified for -Xmx.

Commands include connectivity options designated with the -o option. The two modifiers are mode and charset. Mode set to P indicates that the connection should be persisted, instead of using pooling. Charset is used for localization. For example:

```
-o mode=P charset=utf8
```

6.1 cmdline

```
cmdline.sh <connect> [-a] [-e] [-t] [-f inputfile]
           [-o connect_options]
```

<connect> is either a jdbcurl,

-j jdbc:xdb:<database>:<user>/<password>@<host>:<port>
or

```
-j dbc:xdb://<host>:<port>/<database>?user=<user>&password=<password>
or
[-h <host>] [-s <port>] -d <database> -u <user> [-p <password>]
```

-a : add delimiter. If output mode is NORMAL, it will append an extra delimiter at the end of the last column when doing SELECT queries.

-e : echo mode. Echoes any statements as it executes them

-t : has effect of SET OUPUT NORMAL
(turns off default table mode)

-f : input file to be executed, instead of interactive mode.

The `cmdline` utility is used to obtain a SQL command prompt and execute SQL commands like CREATE TABLE, SELECT and INSERT interactively. A complete list of SQL commands can be found in the SQL Reference manual.

The `cmdline` utility is mentioned here in order to point out that there are some additional administrative commands, which appear in the table below that can be used by the DBA.

ExtenDB Administration Guide

Command	Description
SHOW DATABASES	Lists all of the user-created ExtenDB databases
SHOW TABLES	Lists all of the tables that exist in the current database
SHOW VIEWS	Lists all of the views that exist in the current database
DESCRIBE <table>	Lists the columns and their definitions of the specified table
DESCRIBE <view>	Displays the view definition for the specified view.
SHOW INDEXES ON <table>	Lists all indexes for <table>
SHOW CONSTRAINTS ON <table>	Lists the following types of constraints for <table>: primary keys, foreign keys, foreign key references
SHOW USERS	Lists all defined users and their class
SHOW STATEMENTS	Lists all of the currently executing SQL statements
KILL <request_id>	Kills execution of the request id specified. Request ids can be obtained by executing the SHOW STATEMENTS command.

ExtenDB Administration Guide

6.2 *createdb*

```
createdb.sh -d dbname
           [-h host] [-s port]
           -u dbusername [-p dbpassword]
           -n nodelist [-i initscript]]
           [-m]
```

The `createdb` command is used to create ExtenDB databases.

There are two modes of operation, standard and manual. In standard mode it will try and create the physical underlying databases on all of the specified nodes using the command template defined by `xdb.gateway.createdb` in `xdb.config`, which should be customized depending on if you are using a database other than PostgreSQL.

In manual mode, specified by the `-m` option, it will not try and create the underlying databases, but you are required to create them all by hand properly first.

ExtenDB uses the naming convention of `<dbname>N<nodeid>` when naming the actual physical databases on the underlying nodes. So, if you run `createdb` in manual mode, you should first create all databases and their names properly before running `createdb` with `-m` to wire it up. This naming scheme means that you could create a logical multi-node system where all nodes are really on the same physical system- this is not recommended of course, but may be helpful in testing.

Note that some underlying databases have a limit to the number of characters that can be used when creating the database, so you may need to shorten the name you choose if it is rejected

The values of `dbusername` and `dbpassword` are used to create an ExtenDB DBA username and password for the database, and should be noted. If `-p` is not specified, the user will be prompted for a password. This user should subsequently be used for connecting to the database and for administration.

Note that if you are prompted by a password even with `-p`, it is the underlying tool, like `psql` that is prompting you for a password. This means you are executing `createdb` under a user where a trusted PostgreSQL environment has not been configured. Be sure that it is configured for user `extendb`, and execute the command as user `extendb`.

The `nodelist` is a comma-separated list of node ids that must be valid nodes as defined in the `xdb.config` file.

Optionally, the user may wish to include an initialization script to run on all of the underlying databases at the end of creation using the same mechanism of the `execdb` command.

ExtenDB Administration Guide

Note: in the current version, if XDBServer is running, it must be restarted after `createdb` is executed to be able to use the new database and allow users to connect to it.

If something goes wrong on one of the nodes during creation (a slightly different configuration on a node, underlying database server not running, etc), it might be easiest to fix the problem as follows: drop the database with the `dropdb.sh` command, and then try again to create.

ExtenDB Administration Guide

6.3 *createmddb*

```
createmddb.sh  
    [-i initscript]  
    [-m]
```

The `createmddb` command creates and initializes the metadata database.

It relies on the `xdb.metadata` values in the `xdb.config` file being used, so it is important that this file is configured properly before executing. It will try and create the database `xdb.metadata.database` on the system `xdb.metadata.dbhost` using the command template for `xdb.gateway.createdb` (underlying database dependent).

If successful, and the `-i` option is used, an initialization script will be used. That is required by some databases like SAP-DB, to configure the number of pages in the database and load system tables. This is not required for PostgreSQL.

After creating the database and running the optional initialization script, `createmddb` will create all ExtenDB metadata tables in the metadata database, connecting to it as determined by the `xdb.metadata.*` configuration values in the `xdb.config` file.

Using the `"-m"` option, manual mode, will just try and create the required tables without physically creating the database. This is useful if you want to create the metadata database yourself and then just need to initialize it by creating the required tables.

ExtenDB Administration Guide

6.4 *dropdb*

```
dropdb.sh -d dbname  
          [-h host] [-s port]  
          -u dbusername -p dbpassword [-f]
```

The `dropdb` command is used to drop databases.

The `dbusername` must be a DBA user who has privileges to drop the database.

The underlying databases are dropped as defined by the `xdb.gateway.dropdb` template in the `xdb.config` file.

If there is a problem dropping the database, retry with the `-f` option (force). It will continue to try and remove the metadata from the metadata database even after a failure to remove any underlying databases, and will continue to try and drop from all of the underlying nodes, even if it encounters an error on one.

ExtenDB Administration Guide

6.5 *execdb*

```
execdb.sh -d dbname
          [-h host] [-s port]
          -c command
          -u dbusername [-p dbpassword]
```

The *execdb* command is used to execute the same administrative command or command script on the underlying nodes of a database.

This command is very useful in systems with many nodes and can be used for things like: increasing the database cache size on all of nodes; running backups; changing the logging mode on all nodes. The exact commands to run for each of these will depend on your underlying database.

The command to execute is specified with the *-c* argument. Depending on the underlying database, you may have to specify a valid username and password.

When the command executes on the nodes, the command template specified may include variables that will be substituted:

TemplateVariable	Description
database	The name of the database on the underlying nodes. (Database name and N<nodeid> appended.)
dbhost	The the underlying node host/IP.
dbport	The port of the database server on the underlying node.
dbusername	A valid username on the underlying databases.
dbpassword	A valid password on the underlying databases.

An example for a simple backup of PostgreSQL using *execdb* appears below. In this particular example, since it is being executed from the coordinator, the *-f* option of *pg_dump* will store the backup files of all the nodes on the coordinator as well.

```
execdb.sh -c "pg_dump -h {dbhost} -U {dbusername} {database} -f
/data/back/{database}.dump" -d mydatabase -u extendb -p password
```

In Windows, a similar command might appear as:

```
execdb.bat -c "cmd /C pg_dump -h {dbhost} -U {dbusername} {database} -f
/data/back/{database}.dump" -d mydatabase -u extendb -p password
```

It is also helpful if you have configured a trusted environment using *ssh*, for example, where you can execute remote commands. In that case, we can modify the previous PostgreSQL example to have it run on the remote hosts directly, and save the files on the remote hosts.

```
execdb.sh -c "ssh -h {dbhost} 'pg_dump -h {dbhost} -U {dbusername}
{database} -f /data/back/{database}.dump'" -d mydatabase -u extendb -p password
```

ExtenDB Administration Guide

6.6 XDBAgent

Workgroup Edition Only

```
XDBAgent.sh -n nodelist
```

XDBAgent starts the ExtenDB agent on a node participating as in cluster in the Workgroup Edition.

Using XDBAgent on the nodes facilitates better scalability when more nodes are present in the cluster. Instead of the coordinator doing all the work in connecting directly with the underlying databases, each node can be responsible for one.

Each agent is started with `-n`, followed by its designated node number.

Like XDBServer, XDBAgent uses a `xdb.config` file for its configuration, but it is much smaller compared to XDBServer's. Once the agent connects to the coordinator, other configuration settings that are needed by the agent will be sent over by the coordinator.

It is recommended to start XDBServer on the coordinator before trying to start XDBAgent, but the agent can later be stopped and restarted without having to restart XDBServer.

6.7 XDBdbstart

```
XDBdbstart.sh -d dbname  
                [-h host] [-s socketport]  
                -u dbauser [-p dbapassword]  
                [-w waittimeout] [-x]
```

The XDBdbstart command is used to connect to an existing XDBServer that is already running and bring the database `dbname` online. Internally, it will tell XDBServer to initialize all necessary pools and begin accepting connections for that database.

Which XDBServer to connect to is determined by the host and port specified. If no host is specified, localhost will be used by default. If no port is specified, 6453 will be used by default.

A username and password is required to connect with an existing XDBServer process.

ExtenDB Administration Guide

Optionally, the `-x` can be included which will try and bring the underlying databases on the nodes online as well, determined by `xdb.gateway.startdb` in the `xdb.config` file. If a node's database server is not running, an error will result. Note that in PostgreSQL and Bizgres, there is no need to bring individual databases online; databases are accessible provided the postmaster process is running

An optional waittime may be included to determine how long to wait before failing if a node is inaccessible.

6.8 XDBdbstop

```
XDBdbstop.sh -d dbname
               [-h host] [-s socketport]
               -u dbusername [-p dbpassword] [-x]
```

The `XDBdbstop` command is used to connect to an existing `XDBServer` that is already running and bring the database `dbname` offline. Internally, it will tell `XDBServer` to free all related resources and stop accepting connections to that database.

The `XDBServer` to connect to is determined by the host and port arguments. If no host is specified, localhost will be used by default. If no port is specified, 6453 will be used by default.

The user and password must be valid for that particular database.

Optionally, the `-x` can be included which will try and bring the underlying databases on the nodes offline as well, determined by `xdb.gateway.stopdb` in the `xdb.config` file.

ExtenDB Administration Guide

6.9 XDBServer

XDBServer.sh [-d database_list] [-m] [-x]

XDBServer is executed to start the ExtenDB Parallel Server.

The main configuration for the server appears in its corresponding `xdb.config` file, which is found in `$XDBPATH/config`. Please see "The `xdb.config` File" section under Configuration in this document for more details.

When starting the XDBServer, a space-separated list of databases to bring online may be included with the `-d` option. A database must be brought online before clients can connect to it. If there already is an XDBServer instance running, ExtenDB databases can also be brought online with the `XDBdbstart` command.

The `-m` and `-x` options should not be needed if using PostgreSQL; there is no need to bring individual underlying databases online, databases are accessible provided the postmaster process is running

The `-m` option indicates that XDBServer should also first try and bring the underlying metadata database online first, before trying to connect to it.

The `-x` option indicates that all of those ExtenDB user databases specified in the database list should be brought online on the underlying nodes.

Note that when executing the XDBServer process, you may need to modify the parameters that Java uses, increasing the maximum amount of memory specified in the `XDBServer.sh` launch script.

6.10 XDBshutdown

```
XDBshutdown.sh [-h host ] [-s socketport]
                -u dbusername -p dbpassword
                [-d dblist]
[-x] [-m]
```

XDBshutdown is executed to shutdown an ExtenDB Parallel Server (XDBServer) process. It is not to be confused with `XDBdbstop`, which merely brings a database offline, while allowing the XDBServer process to continue executing.

The XDBServer to connect to is determined by the host and port specified. If no host is specified, localhost will be used by default. If no port is specified, 6453 will be used by default.

The user and password must be valid for that particular database.

ExtenDB Administration Guide

Optionally, the `-x` can be included which will try and bring the underlying databases on the nodes offline as well, determined by `xdb.gateway.stopdb` in the `xdb.config` file.

Similarly, if `-m` is present, it will try and shutdown the metadata database.

6.11 AddNode

```
AddNode.sh -d database
            [-h host ] [-s port]
            -u dbusername [-p dbpassword]
            -n nodelist
            [-m] [-i initscript]
```

`AddNode` allows for new nodes to be added to an existing database.

The nodes must be defined already in the `xdb.config` file, which may require restarting the ExtenDB Server process in order to read in the new node information.

If using the `-m` option for manual mode, a physical database must already exist on the underlying nodes, and named appropriately. For example, if adding nodes 5, 6, 7, and 8 to database PROD, the databases PRODN5, PRODN6, PRODN7, and PRODN8 must exist on nodes 5, 6, 7, and 8 respectively. If `-m` is not present, `AddNode` will try and create the underlying databases using the `xdb.gateway.createdb` definition in the `xdb.config` file.

When `AddNode` is run, it will copy all tables designated as replicated tables to all of the new nodes being added. Other tables, like partitioned tables, will not be redistributed across the new nodes. In order to take full advantage of the new nodes, they must be handled manually. A new table with the same table schema can be created for each of these. Then, the new table should be populated, either by being reloaded, or through an `INSERT..SELECT` command, taking data from the old table. Next, all indexes and primary and foreign keys should be created. Finally, the original tables should be dropped, and then the new tables should be renamed to refer to the old table names.

6.12 DropNode

```
DropNode.sh -d database
            [-h host ] [-s port]
            -u dbusername [-p dbpassword]
            -n nodelist
            [-f]
```

ExtenDB Administration Guide

`DropNode` allows for removing nodes from the database.

It will not allow you to drop nodes when there is an existing single node table or partitioned table on the node. These tables must first be moved or dropped. This is a precaution to prevent inadvertent loss of data. The node should contain no tables, or just replicated tables.

The `-f` option tries to force dropping the database, even if there was an error in updating the metadata of the database.

6.13 XDBLoader and xdbimpex

The `xdbimpex` utility allows for the importing and exporting of data, while `XDBLoader` is targeted exclusively for loading data.

There is a separate document, the *ExtenDB Import and Export Utilities* manual, which provides more detail about using these commands.

7 Isolation Levels and Locking

The four standard isolation levels are

SERIALIZABLE
REPEATABLE READ
READ COMMITTED
READ UNCOMMITTED

By default, ExtenDB uses Read Committed mode (a transaction only sees those rows from the beginning of the transaction until it completes). The ANSI SQL standard allows for a more restrictive isolation level than the one specified, and ExtenDB treats Read Uncommitted as Read Committed and Repeatable Read as Serializable.

Furthermore, even in Read Committed mode, by default ExtenDB will use an exclusive table lock for Update and Delete statements. This can be overridden with the `xdb.config` setting `xdb.locks.readcommitted.mode`. It is set to "S" (strict) by default, but can be overridden to "L" (loose), allowing for shared write locks on tables

If your particular environment does not have a lot of update activity, the default should be fine. Using mode "L" is useful for ETL processes where multiple threads are used to update the same table, which will result in much better performance. The downside of using mode "L" is the added risk that a deadlock may occur across nodes if multiple client sessions are updating the same rows in a transaction.

8 Troubleshooting

This section covers issues that you may encounter while using your ExtenDB cluster, and offers possible solutions.

8.1 Installation and Configuration

The script `createmddb.sh` appears to hang

This is due to a missing or misconfigured `.pgpass` or `pgpass.conf` file. Correct the problem, and try again.

“Template in use” error when running `createmddb.sh` or `createdb.sh`

This is an error message from the underlying PostgreSQL database server, and is caused when trying to create a new database when the template database is believed to be in use. Restart PostgreSQL, and try again.

8.2 Executing

Connections, Pooling, and Timeouts

ExtenDB utilizes various thread and connection pools, and depending on their settings and your workload, you may encounter a timeout issue.

For the client connecting to the ExtenDB server, keep in mind that there is a fixed limit to the maximum number of client connections. This is configured in the `xdb.config` file via the `xdb.maxconnections` setting, where you can override the default setting of 50.

ExtenDB in turn uses pooled connections for communicating with the underlying databases on the nodes. The number of connections used for each node is determined via `xdb.jdbc.pool.initsize` and `xdb.jdbc.pool.maxsize`. You may also have to change the settings in the underlying database that you are using to accept more connections, if you use large values here.

If the number of client connections is larger than these pools, the requests will remain on the request queue for a longer period of time. (Even if the number of requests is smaller than the pools, some “expensive” requests may be not be executed right away by the scheduler to try and both maximize throughput and be responsive for less expensive requests.)

ExtenDB Administration Guide

In addition to the pool sizes, the pools have timeouts. If an executing request cannot obtain the needed connections after the time specified in milliseconds by `xdb.jdbc.pool.timeout`, the request will timeout.

Closely related to the JDBC pools are the thread pools, with settings `xdb.default.threads.pool.maxsize`, `xdb.default.threads.pool.initsize`, `xdb.default.threads.pool.timeout`. A request will only be executed if there are enough threads available in the pool. Normally the thread pool and jdbc pools should have the same size values.

You may also receive timeouts under very heavy query loads with many concurrent sessions. You can try increasing the values of `xdb.messagemonitor.timeout.millis` and `xdb.messagemonitor.timeout.short.millis`.

“Cannot send data to nodes” error message

If you receive the “cannot send data to nodes” error message, it is likely that you have run into a memory resource issue. Try modifying the `XDBServer.sh` script, increasing the values for `MAXMEMORY`, and perhaps `STACKSIZE`.

If, however, you see this for all but the simplest queries, there probably is a permissions issue between the nodes. Make sure permissions are setup properly, including the `.pgpass` file and the usernames and passwords used.

OutOfMemory Exception

You have run into a memory resource issue. Try modifying the `XDBServer.sh` script, increasing the values for `MAXMEMORY`, and perhaps `STACKSIZE`.

Concurrent Performance Slow

The intended usage for ExtenDB is in a data-warehousing environment where not very much transaction activity is expected. Nonetheless, ExtenDB still can process hundreds of low-cost statements per second over multiple client sessions.

For an individual session, ExtenDB does add an extra hop and therefore latency. So, a single session will be much slower compared to a native PostgreSQL database for example. Keep in mind that individual session performance and total throughput are different things; over many sessions working concurrently, much greater total throughput can be achieved.

Please also read the chapter on isolation levels and locking. In particular, you can modify the setting `xdb.locks.readcommitted.mode` in the `xdb.config` file, setting it to “L”.

9 Appendices

9.1 Appendix A – Metadata Database Schema

```
create table xsystablespace (
  tablespaceid int not null,
  tablespacename varchar(255) not null,
  ownerid int not null,
  primary key(tablespaceid)
)
;
create unique index idx_xsystablespace_1
  on xsystablespace (tablespacename)
;
create table xsystablespacelocs (
  tablespaceid int not null,
  tablespacelocid int not null,
  filepath varchar(1024) not null,
  nodeid int not null,
  primary key(tablespaceid, tablespacelocid)
)
;
create unique index idx_xsystablespacelocs_1
  on xsystablespacelocs (tablespaceid, nodeid)
;
alter table xsystablespacelocs
  add foreign key (tablespaceid) references xsystablespace
  (tablespaceid)
;
create table xsysdatabases
(
  dbid int not null,
  dbname varchar(128) not null,
  dbusername varchar(128) not null,
  dbpassword varchar(128) not null,
  primary key (dbid)
)
;
create unique index idx_xsysdatabases_1
  on xsysdatabases (dbname)
;
create table xsysdbnodes
(
  dbnodeid int not null,
  dbid int not null,
  nodeid int not null,
  primary key (dbid, nodeid)
)
```

ExtenDB Administration Guide

```
;  
create unique index idxnodes1 on xsysdbnodes (dbnodeid)  
;  
alter table xsysdbnodes  
  add foreign key (dbid) references xsysdatabases (dbid)  
;  
create table xsystables  
(  
  tableid int not null,  
  dbid integer not null,  
  tablename char(255) not null,  
  numrows int not null,  
  partscheme smallint not null,  
  partcol char(255),  
  parthash int,  
  owner int,  
  parented int,  
  tablespaceid int,  
  clusteridx varchar(80),  
  primary key (tableid)  
)  
;  
alter table xsystables  
  add foreign key (dbid) references xsysdatabases (dbid)  
;  
alter table xsystables  
  add foreign key (parentid) references xsystables (tableid)  
;  
alter table xsystables  
  add foreign key (tablespaceid) references xsystablespaces  
(tablespaceid)  
;  
create table xsystabparts  
(  
  partid int not null,  
  tableid integer not null,  
  dbid integer not null,  
  nodeid int not null,  
  primary key (partid)  
)  
;  
alter table xsystabparts  
  add foreign key (tableid) references xsystables (tableid)  
;  
alter table xsystabparts  
  add foreign key (dbid, nodeid) references xsysdbnodes (dbid, nodeid)  
;  
create table xsystabparthash  
(  
  parthashid int not null,  
  tableid integer not null,  
  dbid integer not null,  
  hashvalue integer not null,  
  nodeid int not null,  
  primary key (parthashid)
```

ExtenDB Administration Guide

```
)
;
alter table xsystabparthash
  add foreign key (tableid) references xsystables (tableid)
;
alter table xsystabparthash
  add foreign key (dbid, nodeid) references xsysdbnodes (dbid, nodeid)
;

create table xsyscolumns
(
  colid int not null,
  tableid int not null,
  colseq smallint not null,
  colname varchar(255) not null,
  coltype smallint not null,
  collength int,
  colscale smallint,
  colprecision smallint,
  isnullable smallint not null,
  isserial smallint,
  defaultexpr varchar(255),
  checkexpr varchar(255),
  selectivity float,
  nativecoldef varchar(255),
  primary key (colid)
)
;
alter table xsyscolumns
  add foreign key (tableid) references xsystables (tableid)
;
create unique index idx_xsyscolumns_1
  on xsyscolumns (tableid, colseq)
;
create table xsysindexes
(
  idxid int not null,
  idxname varchar(80) not null,
  tableid int not null,
  keycnt smallint not null,
  idxtype char(1),
  tablespaceid int,
  issyscreated smallint not null,
  primary key (idxid)
)
;
alter table xsysindexes
  add foreign key (tableid) references xsystables (tableid)
;
alter table xsysindexes
  add foreign key (tablespaceid) references xsystablespaces
  (tablespaceid)
;
create table xsysindexkeys
(
```

ExtenDB Administration Guide

```
idxkeyid int not null,  
idxid int not null,  
idxkeyseq int not null,  
idxascdesc smallint not null,  
colid int not null,  
primary key (idxkeyid)  
)  
;  
alter table xsysindexkeys  
add foreign key (idxid) references xsysindexes (idxid)  
;  
alter table xsysindexkeys  
add foreign key (colid) references xsyscolumns (colid)  
;  
create unique index idx_xsysindexkeys_1  
on xsysindexkeys (idxid, idxkeyseq)  
;  
;  
create table xsysconstraints  
(  
constid int not null,  
constname varchar(128),  
tableid int not null,  
consttype char(1) not null,  
idxid int,  
issoft smallint not null,  
primary key (constid)  
)  
;  
alter table xsysconstraints  
add foreign key (tableid) references xsys tables (tableid)  
;  
alter table xsysconstraints  
add foreign key (idxid) references xsysindexes (idxid)  
;  
create table xsysreferences  
(  
refid int not null,  
constid int not null,  
reftableid int not null,  
refidxid int not null,  
primary key (refid)  
)  
;  
alter table xsysreferences  
add foreign key (constid) references xsysconstraints (constid)  
;  
alter table xsysreferences  
add foreign key (reftableid) references xsys tables (tableid)  
;  
alter table xsysreferences  
add foreign key (refidxid) references xsysindexes (idxid)  
;  
;  
create table xsysforeignkeys
```

ExtenDB Administration Guide

```
(
  fkeyid int not null,
  refid int not null,
  fkeyseq int not null,
  colid int not null,
  refcolid int not null,
  primary key (fkeyid)
)
;
alter table xsysforeignkeys
  add foreign key (refid) references xsysreferences (refid)
;
alter table xsysforeignkeys
  add foreign key (colid) references xsyscolumns (colid)
;
alter table xsysforeignkeys
  add foreign key (refcolid) references xsyscolumns (colid)
;
create unique index idx_xsysforeignkeys_1
  on xsysforeignkeys (refid, fkeyseq)
;
create table xsysusers (
  userid int not null,
  dbid int not null,
  username char(30) not null,
  userpwd char(32) not null,
  usertype char(8) not null,
  primary key (userid)
)
;
alter table xsysusers
  add foreign key (dbid) references xsysdatabases (dbid)
;
create unique index idx_xsysusers_1
  on xsysusers (dbid, username)
;
create table xsystabprivs (
  privid int not null,
  userid int,
  tableid int not null,
  selectpriv char(1) not null,
  insertpriv char(1) not null,
  updatepriv char(1) not null,
  deletepriv char(1) not null,
  referencespriv char(1) not null,
  indexpriv char(1) not null,
  alterpriv char(1) not null,
  primary key (privid)
)
;
alter table xsystabprivs
  add foreign key (userid) references xsysusers (userid)
;
alter table xsystabprivs
  add foreign key (tableid) references xsystables (tableid)
```

ExtenDB Administration Guide

```
;  
create unique index idx_xsystabprivs_1  
  on xsystabprivs (userid, tableid)  
;  
alter table xsystables  
  add foreign key (owner) references xsysusers (userid)  
;  
create table xsysviews (  
  viewid int not null,  
  dbid int not null,  
  viewname varchar(255),  
  viewtext varchar(7500))  
;  
create unique index idx_xsysviews_1  
  on xsysviews (viewid)  
;  
alter table xsysviews  
  add foreign key (dbid) references xsysdatabases (dbid)  
;  
create table xsysviewscolumns (  
  viewcolid int not null,  
  viewid int not null,  
  viewcolseqno int not null,  
  viewcolumn varchar(255),  
  coltype smallint not null,  
  collength int,  
  colscale smallint,  
  colprecision smallint,  
  primary key (viewcolid))  
;  
create unique index idx_sysviewscols_1  
  on xsysviewscolumns (viewid, viewcolseqno)  
;  
alter table xsysviewscolumns  
  add foreign key (viewid) references xsysviews (viewid)  
;  
create table xsysviewdeps (  
  viewid int not null,  
  columnid int not null,  
  tableid int not null)  
;  
alter table xsysviewdeps  
  add foreign key (viewid) references xsysviews (viewid)  
;  
create table xsyschecks (  
  checkid int not null,  
  constid int not null,  
  seqno int not null,  
  checkstmt varchar(8000),  
  primary key (checkid))  
;  
create unique index idx_xsyschecks_1  
  on xsyschecks (constid, seqno)  
;  
alter table xsyschecks
```

ExtenDB Administration Guide

```
add foreign key (constid) references xsysconstraints (constid)
;
```