



@enterprise 7.0

Installation and Configuration

September 2009

Document Version 7.1

Groiss Informatics GmbH

Contents

1	System Requirements	4
1.1	Server	4
1.2	Java	4
1.3	Database	4
1.4	Client	5
2	Installation	6
2.1	Database Preparation	6
2.1.1	Oracle	6
2.1.2	MS SQL-Server	7
2.1.3	DB2	7
2.1.4	Derby	7
2.2	Extract and Install	7
2.3	Installing a Service	9
2.4	Using an Application Server or Servlet Container	9
3	Configuration	11
3.1	License	11
3.2	HTTP-Server	11
3.2.1	Defining Allowed and Denied Hosts or Networks	12
3.2.2	Access Control	13
3.3	Database	15
3.4	Directories	17
3.5	Logging	17
3.6	Classes	18
3.7	Localization	18
3.8	Communication	20
3.9	Cluster	21
3.10	DMS	22
3.11	Search	23
3.12	Tuning	24
3.13	SSL	25
3.14	Password Policy	26
3.14.1	General Policy Settings	26

3.14.2	Default Policy Checker Settings	26
3.14.3	Your Own Checker Class	28
3.15	Calendar	28
3.16	Time Management	28
3.17	ACLCache	29
3.17.1	Configuration of ACLCache	29
3.18	Change Administrator Password	30
3.19	Initialize Database Schema	30
3.20	Parameters without GUI	30
4	Clustered @enterprise System	33
4.1	Overview and Principles of the Clustered Architecture	33
4.2	Cluster and Nodes	34
4.3	Configuring a clustered @enterprise System	34
4.3.1	Platform Configuration	34
4.3.2	Installation of a nonclustered System	35
4.3.3	Transport Mechanisms for Cache Coherence Service	35
4.3.4	Adapting the @enterprise Configuration	37
4.4	Operation of a clustered system	39
4.4.1	Monitoring	39
4.4.2	Load Balancing	39
4.4.3	Event Handling	40
5	Setting up an Archive Schema	41
A	Database Performance Hints under Oracle	42
A.1	Preliminaries	42
A.2	Key Operating Parameters of the Database	42
A.3	Optimizer	45
A.4	Storage	46
A.4.1	Disks	46
A.4.2	Parameters for Tablespaces	46
A.5	One owns Tables and Queries	47

1 System Requirements

1.1 Server

@**enterprise** 7.0 is available for several platforms. The only requirement is the availability of a Java JDK Version 1.5 or higher. The following operating systems are supported:

- Windows NT, 2000, XP, or 2003
- Solaris
- AIX
- Linux

The server should have at least 512MB RAM for @**enterprise** and 100MB free disk space.

1.2 Java

On the server a Java Development Kit (JDK) must be installed. It can be downloaded from Sun (<http://java.sun.com>) or from another vendor. On Sun's web-site a list of Java ports to other platforms is available.

Please note, that for using @**enterprise** the Java Interpreter *and* the Java-Compiler is necessary, both are part of the JDK. The needed version of the JDK is 1.5.0 or higher.

On clients we require the Java plugin to be installed in order to run the @**enterprise** process editor.

1.3 Database

We support the following databases: ORACLE, MS SQL-Server, IBM's DB2, Firebird, and Derby. The database can be installed on another server, @**enterprise** connects via the network to the database.

The following database versions are required:

- ORACLE 9i or higher
- SQL-Server 2000

1.4. CLIENT

- DB2 6.1 or higher on Windows or AIX
- Firebird Version 1.5 or higher
- Derby 10.1.2.1 or higher
- MySQL 5.0 (experimental)

1.4 Client

When using the Web-Client you need only a Web-browser, the following versions are required:

- Netscape 7.2 or higher
- MS Internet Explorer 6 or higher
- Firefox 1.0 or higher

2 Installation

2.1 Database Preparation

@enterprise needs a database with one user. In the following we briefly describe the necessary steps for creating a database user for the three supported database. Please consult the database manuals or the local experts for further information about database setup and creation of a user.

2.1.1 Oracle

You need a database user with the following rights:

```
CREATE SESSION
ALTER SESSION
CREATE TABLE
CREATE VIEW
```

The user must have access to a *tablespace* and the right to add data there.

Example:

```
SQL> create user ep identified by eppwd
2 default tablespace system;
User created.
SQL> grant create session, alter session to ep;
Grant succeeded.
SQL> grant create table,create view to ep;
Grant succeeded.
SQL> grant unlimited tablespace to ep;
Grant succeeded.
SQL>
```

If you use full-text search the *IndexRefreshTimer* needs an additional right:

```
grant execute on ctxsys.ctx_ddl to ep;
```

Hint: If you got the message *Could not get Session ID. Probably no right on V\$SESSION*, you have to do following steps in Oracle:

1. *Login as sys:* sqlplus sys as sysdba

2. *Assign grant*: `grant select on v_$session to <username>;`

<username> is the name of the user, which is entered as @enterprise database user.

If the use of WfXML2 functionality is intended with Oracle as the underlying DBMS, you must select the Oracle LOBs database type in the configuration (and not the Oracle LONGs one). Since Oracle supports just one LONG column per table, the tables for WfXML2 functionality will not be generated when LONGs are used instead of LOBs.

Hints for the performance of Oracle-based @enterprise installations can be found in appendix [A](#).

2.1.2 MS SQL-Server

@enterprise requires a case insensitive installation of MS SQL-Server.

When creating a SQL-Server database, use the option 'ANSI NULL is default'. You can specify it in the database property panel or by execution of a stored procedure after installation.

```
sp_dboption <dbname>,'ANSI null default', true
```

<dbname> must be replaced with the name of your database. The procedure results in behavior consistent with the ANSI standard regarding the handling of NULL values.

The database user for @enterprise must have the right to create tables, for example via the role `db_owner`.

If you use full-text search, please ensure that MSSEARCH service is running and automatic population (for creating indices of full-text catalog) is activated.

2.1.3 DB2

When using DB2 you have to create an operating system user. Afterwards a database user is created with the rights *connect to database* and *create tables*. Then you create a database schema, for which the user is authorized.

2.1.4 Derby

Derby doesn't need any preparation.

2.2 *Extract and Install*

This section describes how to install the @enterprise stand-alone server. @enterprise is distributed on CD or can be downloaded from our web site. It is packed in one single file named `setup70.jar`. The installation can be started with a double-click on the file. The installation of a Java JDK 1.5 (or higher) is required.

If *.jar files are not associated with Java on your machine, or if you don't have a GUI available, please start the setup on a command line:

2.2. EXTRACT AND INSTALL

```
java -jar setup70.jar
```

The setup process consists of the following steps:

1. Verify if this is the version of **@enterprise** that you want to install and start the setup by clicking on *OK*.
2. Specify the directory of the Java compiler and interpreter.
3. Installation directory: The directory where the system will be installed.
4. Choose the port on which the **@enterprise** server will run.
5. If your server operating system is MS Windows you can install a service.
6. Now setup shows you information about how you can start the server and continue the setup process.
7. Setup will try to start the server and open a browser for you. If this fails and if you did not install a service, you have to start the server manually by executing the batch file (avw.sh or avw.bat).

If your browser didn't already do it, please navigate to <http://localhost:port/>, where *port* is the port number that you have chosen during the previous setup steps. The rest of the installation is done with the browser.

8. The first screen is the Welcome-screen, click on *Start Setup* to start the configuration.
9. On the next screen you specify a logical name for the server (server ID), a server number (an integer value for distributed installations), the license key and the server's default language.
10. Now you can load a database JDBC driver. Use the *JDBC Driver Help Page* for information about different databases and their JDBC drivers.
11. On the next screen you have to specify some database parameters. We suggest to use the help function (the question mark next to *Database Type*) to fill the Database Type, JDBC Driver Class, and JDBC URL fields with valid values.
 - Database Type: The database; you can select ORACLE, DB2, MS SQL-Server, Firebird, or Derby.
 - JDBC Driver Class: Java class that contains the driver. Take a look at the table on page 16 for a list of driver classes.
 - JDBC URL: URL for the database. The syntax of this string depends on the JDBC driver used. See the examples on page 16 or consult the documentation of the driver.
 - Database Userid: The ID of the user with whom you want to connect to the database.
 - Database Password: Password for the database user with the ID that you entered above.

2.3. INSTALLING A SERVICE

- Number of Connections: Default number of database connections.
 - Session Environment: You can specify SQL-commands, which are executed for each connection after connecting, for example: `set TEXTSIZE 1000000`
12. Now the database and driver will be tested. Optionally you can test if your database can store unicode characters.
 13. The next step is the creation of the database tables. The time may vary depending on your server's speed and the database that you use.
 14. After initializing the database, some internal services have to be started.
 15. On the next screen the password of the system administrator can be specified.
 16. Now a user and an organizational unit can be created. The following roles will be given to this user: *all*, *home* in the inserted organizational unit, and *sys*.
 17. If you want, you can load an example process now.
 18. Congratulation! You finished the setup of **@enterprise**. Click on *Login* to go to the login page, where you can immediately start to use **@enterprise**.

By completing the previous steps you finished the setup of **@enterprise**. If you want to change the configuration or configure advanced settings, take a look at chapter 3.

2.3 Installing a Service

In Windows you can configure a stand-alone installation of **@enterprise** to run as service. This can be done while installing (see the previous section) or later with calling the program `installep.bat` in the directory `service`.

To uninstall the service, call

```
javaservice -uninstall servicename
```

from the `service` directory.

If you want to change start parameters (like `-Xms` and others) or the classpath, we suggest to modify the `installep.bat` file and reinstall the service (uninstall and install).

2.4 Using an Application Server or Servlet Container

If you want to run **@enterprise** in an application server (e.g., IBM's *WebSphere*) or a servlet container (e.g., Apache's *Tomcat*) you need the **@enterprise** web application archive file named `ep70.war`. Deploy this file in your server. Afterwards, open your browser and navigate to `http://host:port/context-root/`, where `host` and `port` must be the right values for accessing your server and `context-root` is the context root that you chose when deploying the file. See section 2.2 step 8 for details about the rest of the installation.

The `ep70.war` archive is especially prepared to be used in a servlet container like *Tomcat*. It contains a `jar` file named `j2eesmallnoservlet.jar`, which is a smaller version of

2.4. USING AN APPLICATION SERVER OR SERVLET CONTAINER

`j2eesmall.jar` of the stand-alone **@enterprise** server. This file is required if you run **@enterprise** in *Tomcat*. If you use an application server this file will usually not be required. If you encounter problems deploying or starting **@enterprise** in an application server, remove the `j2eesmallnoservlet.jar` file from `ep70.war` and try again. You can open the archive with a zip tool (e.g., *WinZip*) and remove the `jar` file from the directory `WEB-INF/lib`.

3 Configuration

This chapter describes advanced configuration parameters of **@enterprise**. You can change the data that you entered at setup as well as additional configuration here. Open the configuration area in the system administration by clicking on *Configuration* in the menu on the left side.

In order to save your changes, you must use the *Update* button, which is available on every configuration page. Nearly all settings require to restart the server for changes to become active. The only exceptions are *HTTP-Server* → *Access Control*, *Database* → *Number of Connections* and *Maximum Number of Connections*, and *Logging* → *Trace Level*.

In the following we describe the different parameter groups. Each of them is represented by an entry in the configuration menu. If you use a german server installation and encounter problems understanding the english terms used in this manual, we suggest to create and use an administrator with english language (the *sys* right is required in order to enter the administration).

3.1 License

The first screen contains license information:

- **License Key:** Your license key. If you want to change your license key after you finished the setup, you can enter the new key here.

3.2 HTTP-Server

This screen contains the setup of the HTTP server:

- **Server IP Port:** HTTP port on which the server runs.
- **Minimum Number of Threads:** Number of threads, which are started on startup.
- **Maximum Number of Threads:** Maximum number of threads, which will be used for HTTP requests.
- **Allowed Hosts or Networks:** A list of hosts and networks can be specified. These hosts can access the HTTP server. The syntax of this field is described below in section [3.2.1](#).

3.2. HTTP-SERVER

HTTP-Server

Server IP Port: 8000

Minimum Number of Threads: 2

Maximum Number of Threads: 25

Allowed Hosts or Networks:

Denied Hosts or Networks:

URL-Prefix and Handler Classes: /secure/aww.method.com.groiss.servlet.Dispatcher
/aww.method.com.groiss.servlet.Dispatcher
/servlet.method.com.groiss.servlet.Dispatcher
/aww.class.com.dec.gi.distrib.RMIFileHandler
/webdav.com.groiss.webdav.WebDAVHandler

Access Control:

Update

Figure 3.1: @enterprise Configuration

- **Denied Hosts or Networks:** Analogous to above.
- **URL-Prefix and Handler Classes:** Here the servlets are listed, which are used for executing Java-methods on the server. You can add your own servlet, but do not modify or delete the entries already there, if you don't really know what you are doing.
- **Access Control:** We provide a mechanism which allows to grant or deny access to method-URLs based on a combination of IP-addresses and rights. The syntax of access rules and their semantics is described below in section 3.2.2.

3.2.1 Defining Allowed and Denied Hosts or Networks

To restrict access to the HTTP server to selected hosts or address ranges you can declare an *allow* and a *deny* list. The evaluation is as follows: If the allow-list is empty, access is allowed from every host except them in the deny-list. If the allow-list is not empty, access is allowed from the hosts and networks in the allow list minus the hosts (and networks) in the deny list.

Both lists contain pairs of IP-Address and netmask separated with spaces.

See the following example:

```
10.205.112.0/255.255.255.0
```

This entry in the allow-list means, access from the network 10.205.112.* is allowed. A bit set to 1 in the netmask means that this bit must be equal in the given IP-address and the address of the accessing host.

The following list used for the allow-list causes that access from hosts 10.205.112.4 and .8 is allowed.

```
10.205.112.4/255.255.255.255 10.205.112.8/255.255.255.255
```

3.2.2 Access Control

The access control mechanism affects the Dispatcher servlet which serves URLs targeting java methods. Rules can be specified which restrict access to certain URLs based on a combination of IP-address and **@enterprise** rights.

To activate the access control, the corresponding service must be added to the services in *Classes* → *Services*:

```
com.groiss.avw.contrib.URLChecker uc
```

Configuration

The access control property consists of a comma-separated list of rules. Each rule combines an IP-specifier, an URL-prefix and a set of rights separated by spaces. Each of the components can be a wildcard in the form of an asterisk.

Accordingly, the syntax of the ruleset is:

```
{ ( ip-specifier | "*" ) SPACE (url-prefix | "") SPACE ( "*" | "DENY" | (
    right { SPACE right }* ) COMMA }*
```

The IP-specifier consists of a host-mask and a net-mask separated by a "/". It can be used to specify a single host or a subnet in the following way:

10.205.112.22/255.255.255.255	designates the single host 10.205.112.22
10.205.112.0/255.255.255.0	designates all hosts in the subnet 10.205.112.*
10.0.0.0/255.0.0.0	designates all hosts in subnet 10.*.*.*
*	this wildcard designates all hosts

Technically, the IP-address of a requestor matches an IP-specifier of the form host-mask/net-mask, when (ip-address XOR host-mask) AND net-mask equals 0.

An URL-prefix consists of the first characters of a fully qualified method name (package, class, method). The URL-prefixes are case sensitive.

com.groiss	designates all calls to methods in classes in packages located in com.groiss or below
com.groiss.org.PasswdAuth	designates all calls to methods in the class com.groiss.org.PasswdAuth
*	this wildcard designates all methods regardless of origin

The set of rights is a space separated list of IDs of **@enterprise** rights. The right IDs are case sensitive.

3.2. HTTP-SERVER

set_agent	designates all users who have the right set_agent
admin stat	designates all users who have the right admin and / or the right stat
*	wildcard designating that rights are not needed
DENY	special dummy right id, can be used to deny access

Examples for Rules

The following examples show how those three designations can be combined to form a rule:

127.0.0.0/255.0.0.0 * *

Access from local host subnet is not restricted.

10.205.112.26/255.255.255.0 * DENY

Access from 10.205.112.26 is not allowed.

10.205.112.0/255.255.255.0 com.groiss.org.PasswdAuth *

Login of hosts from subnet 10.205.112.0 is allowed.

10.205.112.0/255.255.255.0 * internal

All operations of hosts from this subnet are allowed if users have the right internal.

* com.groiss DENY

Access to com.groiss.** classes and methods is denied to every host.

* com.my.appl admin,customer

Access to com.my.appl.** classes and methods is allowed if users have the right admin or customer.

* * DENY

Deny everything from everywhere.

Semantics

The validation of a list of rules in the *Access Control* property is as follows:

If the property is empty, nothing is filtered.

Otherwise all rules are checked in the order they are defined until a rule matches according to IP-specifier and URL-prefix. For a matching rule, the validation depends on the set of rights of the rule. We distinguish two cases:

- **Existing Session** (user already logged in):
The intersection of the rights of the user and the rights given in the rule is computed. If the intersection is empty, access is denied (an exception is thrown), else the rule succeeds and access is granted.
- **No Session** (user not yet logged in):
If the set of rights of the rule consists of a single DENY element, then access is denied (an exception is thrown), else the rule succeeds and access is granted.

If no rule at all matched, access is granted. This can be avoided if the last rule is "* * DENY".

Other Operational Considerations

Access Control gets reconfigured if the configuration is changed. This is also logged at log level 1 to allow one to find incorrect rules. Normal operations of *Access Control* are logged at log level 3.

Access Control is not automatically aware of additional rights given to a user or role or to the revocation of rights from them. In order to know about the constellation, the affected users must log out and log in again or the configuration must be saved (thereby reconfiguring *Access Control*). Caching of user rights in the *Access Control* mechanism is logged at log level 2.

3.3 Database

We suggest to use the help function (the question mark next to *Database*) to fill the Database, JDBC Driver Class, and JDBC URL fields with valid values for a selectable database.

- **Database:** The database; you can select ORACLE, DB2, MS SQL-Server, Firebird, or Derby.
- **JDBC Driver Class:** Java-Class, that contains the driver. See the table on page 16 for a list of driver classes.
- **JDBC URL:** URL for the database. The syntax of this string depends on the JDBC driver used. See the examples on page 16 or consult the documentation of the driver.
- **Database Userid:** The ID of the user with whom you want to connect to the database.
- **Database Password:** Password for the database user with the ID that you entered above.
- **Number of Connections:** Default number of database connections.
- **Maximum Number of Connections:** The maximum number of database connections that can be created.
- **Session Environment:** You can specify SQL-commands, which are executed for each connection after connecting, for example: `set TEXTSIZE 1000000`
- **Reconnect Try Interval (sec.):** Interval in seconds for reconnect tries to the database.
- **Reconnect Tries:** Number of reconnect tries.
- **Query Timeout (sec.):** Number of seconds after which a query times out.

Table 3.1 shows the recommended drivers for the databases, their class names and JDBC URLs (you can directly view and use this table in @enterprise by clicking on the help link next to *Database*).

3.3. DATABASE

DBMS	Driver Vendor	Driver Kind	Class and URL
DB2 UDB	IBM	Universal	COM.ibm.db2.jdbc.app.DB2Driver jdbc:db2:'dbname'
DB2 Z/OS	IBM	OS390	COM.ibm.db2os390.sqlj.jdbc.DB2SQLJDriver jdbc:db2os390:'location-name'
Derby	Apache	Embedded	org.apache.derby.jdbc.EmbeddedDriver jdbc:derby:ep;create=true
Firebird SQL 1.5	Firebird	JCA	org.firebirdsql.jdbc.FBDriver jdbc:firebirdsql:'host'/3050:'dbalias'
MS-SQLServer	Inetsoftware	Una2000	com.inet.tds.TdsDriver jdbc:inetdae:'host':1433?sql7=true
MS-SQLServer	jTDS Project	jTDS	net.sourceforge.jtds.jdbc.Driver jdbc:jtds:sqlserver://'host':1433
Oracle LOBs	Oracle	Thin (>=10g)	oracle.jdbc.OracleDriver jdbc:oracle:thin:@'host':1521:'SID'
Oracle LOBs	Oracle	OCI	oracle.jdbc.OracleDriver jdbc:oracle:oci:@'TNSNAME'
Oracle LONGs	Oracle	Thin	oracle.jdbc.OracleDriver jdbc:oracle:thin:@'host':1521:'SID'
Oracle LONGs	Oracle	OCI	oracle.jdbc.OracleDriver jdbc:oracle:oci:@'TNSNAME'
MySQL	MySQL	Connector/J (3.1)	com.mysql.jdbc.Driver jdbc:mysql://'host':port/'database'
PostgreSQL	PostgreSQL	Native PostgreSQL Driver	org.postgresql.Driver jdbc:postgresql://'host':port/'database'

Table 3.1: JDBC-Drivers

3.4 Directories

Here you can define some directories that **@enterprise** will use. The *Directory of Form Classes* and *Directory for Temporary Files* must exist.

- **Home-Directory:** This is the root directory for all relative paths, if you leave it empty the current directory of the start script is used.
- **HTTP-Server Document-Root:** Directory where the documents for the HTTP server reside.
- **Directory of Form Classes:** Directory, where the system writes the form classes.
- **Directory for Temporary Files:** Directory for temporary files.

3.5 Logging

- **Logfile:** Name of file, where **@enterprise** writes log information.
- **Error File:** This file is a centralized collection of errors. They will also appear in the general logfile. You can leave this field empty if you don't want a separate file for errors to be created. Anyway, we recommend to define an error file.
- **Logger Class:** If you write your own logging mechanism you can specify the class name here. The class must implement the interface `com.groiss.log.ILogger`.
- **Trace Level:** 0, 1, 2 or 3:
 - 0 Errors are logged.
 - 1 HTTP requests are logged (timestamp, user, IP-address, and URL).
 - 2 SQL-statements and process-oriented logging.
 - 3 The full HTTP-headers, parameters of prepared statements and other information for debugging purposes.

Don't use the options 2 or 3 in production for extended periods of time, because it generates a lot of data.

- **Log on console:** The log information is written to the standard output stream.

To include database session ids in the log, it is necessary, that the database user `SYS` executes the following grant:

```
grant select on v_$session to ep;
```

3.6 Classes

- **Authorization Class:** `@enterprise` allows the usage of different authorization mechanisms. The Java class used is specified here. The default class (part of the distribution) is `com.groiss.org.PasswdAuth`.
- **Settings Class:** A class defining some global settings can be defined here. For details see the `@enterprise` Programming Guide.
- **Notification Provider Class:** The class for the notification mechanism.
- **Archiving Class:** The class used for archiving process instances, must implement the interface `com.dec.avw.core.AVWArchiver2`.
- **Error-Formatter Class:** You can write an error formatter class that will be used to display errors. The class must implement the `com.groiss.gui.ErrorFormatter` interface.
- **Services:** The list of services that the system starts. You can add your own services but should not modify or delete the entries already there, if you don't really know what you are doing.

3.7 Localization

- **List of Locales:** Here you can define a comma-separated list of locales that will be used by the server. If you don't define anything here, the server will use the following default locales: `en_GB`, `en_US`, `de_DE`, `de_AT`, and `de_CH`.
- **Language:** Defines the language for the user interface. Language is defined in ISO language code, for example `de` for German.
- **Country:** ISO country code, for example `AT` for Austria.
- **Variant:** A default variant to use. You can define free variants in the list of locales (e.g., regions, companies, etc.).
- **Character Set:** Here you can define the character set that will be used for communication with the web clients. HTTP responses of `@enterprise` will contain this character set in the HTTP header. If you don't define anything here, the default character set will be used (usually this is `cp1252`). Valid other settings are for example `UTF-8` or `ISO8859-1`. If you change the character set to other values, please back up your `/conf/avw.conf` file before (so that you can use the old file again if the new setting doesn't work).
- **Date Format:** Format mask for date input and output. See the table below for a description of the possible values.
- **Date-Time Format:** Format mask for date and time.
- **Default-Unit for displaying TimeIntervals:** Default-Unit in seconds, minutes, hours, days and weeks.

- **Applet Look-and-Feel:** Specify the look-and-feel of the process editor, values are: metal, windows, or motif.
- **Max. Table Length:** Specify a natural number. For tables of size greater than this number the user is asked before the table is shown.
- **Items per Page:** This defines the maximum number of entries in worklists when paging is enabled. Paging can be enabled by adding the string `<Attrib key="paging" value="true"/>` to *worklist* node in the file *standard.xml*.
- **Keep object changes (days):** @enterprise stores every master data change. Here you can specify the number of days these changes are kept. No entry means, changes are kept forever.
- **Remove user sessions after (days):** Number of days after which inactive user sessions will be deleted.
- **Use browser language:** If this option is set, the system uses the language settings of the browser instead of the settings in the user table of @enterprise.
- **Open form on process start:** In the process start mask there is a checkbox where the user can decide to see the process form immediately after process start. Here you can define the default value of this checkbox.
- **Readonly fields as text:** Displays readonly form fields as simple text instead of readonly input fields.
- **Inherit Ids to subprocesses:** Don't create Ids for subprocesses - use the parent processes' Ids instead.
- **All tasks can be selected as adhoc:** This makes all tasks selectable as adhoc tasks.
- **Select List Look-and-Feel:** You can choose the Look-and-Feel for select list:
 - *Table:* The select list will be displayed as table.
 - *Select List:* The select list will be displayed as select list.
- **Select List Search Option:** The search option for searching in a select list:
 - prefix
 - Substring
- **Enable application-spanning process definition:** If this option is set, it is possible to define processes with application-spanning elements (i.e. Forms, Tasks, Subprocesses and Roles as Agents).

Table 3.2 shows possible values for the date and time format masks.

The count of pattern letters determine the format.

(Text): 4 or more pattern letters—use full form, < 4—use short or abbreviated form if one exists.

(Number): the minimum number of digits. Shorter numbers are zero-padded to this amount. Year is handled specially; that is, if the count of 'y' is 2, the year will be truncated to 2 digits.

3.8. COMMUNICATION

Symbol	Meaning	Presentation	Example
G	era designator	(Text)	AD
y	year	(Number)	1996
M	month in year	(Text & Number)	July & 07
d	day in month	(Number)	10
h	hour in am/pm (1 12)	(Number)	12
H	hour in day (0 23)	(Number)	0
m	minute in hour	(Number)	30
s	second in minute	(Number)	55
S	millisecond	(Number)	978
E	day in week	(Text)	Tuesday
D	day in year	(Number)	189
F	day of week in month	(Number)	2 (2nd Wed in July)
w	week in year	(Number)	27
W	week in month	(Number)	2
a	am/pm marker	(Text)	PM
k	hour in day (1 24)	(Number)	24
K	hour in am/pm (0 11)	(Number)	0
z	time zone	(Text)	Pacific Standard Time
'	escape for text	(Delimiter)	'
”	single quote	(Literal)	'

Table 3.2: Values for Date and Time Format Masks

(Text & Number): 3 or more—use text, less than 3—use number.

Any characters in the pattern that are not in the ranges of [`'a'..'z'`] and [`'A'..'Z'`] will be treated as quoted text. For instance, characters like `':', '.', ' ', '#'` and `'@'` will appear in the resulting time text even if they are not embraced within single quotes.

3.8 Communication

- **SMTP Host:** Server for outgoing E-Mails (host name or IP address).
- **Mail Sender:** The mail address that will appear in the *from* field of mails that the system sends.
- **Administrator E-Mail Address:** E-Mail address of the system administrator.
- **RMI Port:** Port number of RMI (Remote Method Invocation) listener. Needed for Java-Clients.
- **Enable RMI class loading:** Enables class loading via RMI, this is needed when working with forms and the Java-Client.
- **Enable full RMI access:** When starting an RMI session the system must authorize a user. All RMI calls will be performed as this user then. If you enable full RMI access, you can call all available API methods independent of the user's rights.
- **Allow plain communication over RMI:** Allows plain (unencrypted) communication for RMI connections.

- **Export Port for plain RMI communication:** If specified, this is the port used for RMI traffic.
- **Use SSL for login sequence at RMI communication:** Use SSL to encrypt the login sequence for RMI communication.
- **Crypt RMI communication with SSL:** Encrypt the whole communication over RMI.
- **Export Port for RMI over SSL:** If specified, the encrypted RMI communication uses this port.
- **Enable Admin-shell:** If not enabled, access from the admin-shell is denied.
- **Require SSL for Admin-shell:** If enabled, unencrypted communication is denied.
- **Allowed Hosts or Networks for Admin-shell:** Specifies a network restriction pattern. See the parameter *Allowed Hosts or Networks* in section 3.2.
- **Enable Wf-XML:** Defines if this server is Wf-XML enabled. Possible values are *off*, *active*, or *passive*. For further details on how to set up and use Wf-XML, please take a look at the section *Communication with other Systems → Wf-XML* of the **@enterprise** Application Development Guide.
- **WfXML Org. Unit:** Default Wf-XML Organizational Unit.
- **WfXML User:** Default Wf-XML User.
- **WfXML Server:** Defines the default Wf-XML server.
- **WfXML access log for:** Defines the objects, which will be logged. You can select between
 - ServiceRegistry
 - Factory
 - Instance
 - Activity
 - Observer.
- **Size of log:** Max. size of the logfile.

To enable RMI communication you must at least enable either *Allow plain communication over RMI*, *Use SSL for login sequence at RMI communication*, or *Crypt RMI communication with SSL*.

If a timer doesn't catch an exception, **@enterprise** sends a mail to the system administrator and deactivates the timer.

3.9 Cluster

See section 4.3.4 in the chapter about clusters for details about configuring clusters.

3.10 DMS

- **Show extensions:** Show the document name extension, e.g., .doc or .txt.
- **Versioning:** *Not automatically* disables automatic version creation. *On agent change* creates a version if a different user edits the document (so, if the same user edits a document multiple times, no documents are created). *On every change* creates a version every time the document is edited.
- **Inherit permission list:** When this option is checked, the permission lists of a folder is inherited to the contents of the folder.
- **Basic-Auth in WebDAV:** Check this checkbox if you want to allow Basic-Auth authentication in WebDAV. If this is disabled, not logged in users will not be able to access WebDAV.
- **Open docs in new window:** If checked, documents will be opened in new windows.
- **Maximum Document Size (in bytes):** You can define a maximum size for DMS documents here. **@enterprise** will not allow users to create documents that are bigger than this value. If you don't define a maximum size, there will be no size restriction for DMS documents. Anyway, also databases can limit the maximum size.
- **DMS Storage Class:** You can specify your own DMS storage class here. The class must implement the interface `com.groiss.dms.IStore`.
- **DMS Archiving Class:** Class for archiving documents, must implement the interface `com.groiss.dms.DMSArchiver`.
- **Standard Table Model / Table Handler:** A class can be specified, which is used for displaying the document tables. Take a look at the section *Using the DMS API → Adapting Folder and Table Model* of the **@enterprise** Application Development Guide for further details about table model classes.
- **Signature Class:** The signature implementation. If the default implementation of **@enterprise** should be used, `com.groiss.security.impl.KeySignature` has to be entered here. The default implementation signs documents by using a key-pair (private-, public key) which is protected by a password. You can also specify your own signature class which must implement the interface `com.groiss.security.Signature`.
- **Signature Types:** It is possible to distinguish between different types of signatures (e.g., read, approved, etc.). Here you can enter a comma separated list of strings where each string represents a signature type.
- **Full-text Search:** With the help of this parameter the state of the full-text search can be determined: There are three possible states:
 - **Switched Off:** No full-text search is used at all.
 - **String Search in Form Fields:** The database doesn't support full-text search. Therefore the required string can be searched in a table containing all string values of form fields.

– **Activated:** The full-text search of the current database is used.

- **Do not display hidden documents:** If this option is checked, users cannot see any hidden documents (beginning with a point in the filename) in the DMS .
- **Character Set for Text Files:** Here you can enter the character set for text files, if the content of these files is not displayed correctly, e.g. the content of the file has ANSI charset, but the server charset is UTF-8 - for this purpose set the character set for text files to the value *CP1252* (if client is running under Windows only).

3.11 Search

- **Maximum Table Size on Server (rows):** Maximum table size the server will handle. If the table size exceeds this value, the operation is cancelled and an error message is produced.
- **Cache Interval (minutes):** Specifies, how long a query result resides in cache.
- **Maximum Number of Cached Queries:** Number of queries in cache.
- **Maximum Number of Simultaneous Queries:** Number of threads, that concurrently compute query results.
- **Maximum Number of Startable Queries:** Length of queue of queries waiting for execution (waiting for a free thread).
- **Process Relations:** It is possible to define a relationship between process instances. The relation is defined as *ProcessRelation(ProcessInstance p1, ProcessInstance p2, String reltype)*. The relation can be maintained via API or with the task-function *addRelation*. The available relation types can be defined in the field *Process Relations*. For each relation type a pair of id and name is defined, name and id separated by whitespace. A comma separates the pairs. The id is stored in the database relation, the name is used in the user interface.
- **Reporting schema:** The filename of the reporting shema can be entered here. The default schema is in the file *reporting.xml*.
- **Exact Id Shortsearch only:** If this checkbox is activated, you have to enter the right Id to get a correct result.
- **Shortsearch includes subject:** If this checkbox is activated, the subject will be included in shortsearch.
- **Shortsearch includes Fieldvals:** If this checkbox is activated, fieldvals will be included in shortsearch.
- **Order Process-Ids by OID:** In worklist and Reporting processes will be sorted by OID, if this checkbox is activated.

For more information on process relations read the corresponding chapter of the **@enterprise** Application Development Guide.

- **Show all rows, even when no View Right:** If this checkbox is activated and the user who uses search-engine has no view right on DMS-object, he will get all rows as result.
- **Open Forms in Edit Mode:** If this checkbox is activated, DMS forms in reporting result will be opened in edit mode, i.e. forms can be changed.
- **Search case-insensitive by default:** If this checkbox is activated, the checkbox *Ignore Case* on process search mask is activated by default.
- **Default Subjectsearchttype:** Here you can define the standard type for subject search in *Process Search* - see user manual for further information.
- **Default process-id searchttype:** The same as *Default Subjectsearchttype*, but for id.

3.12 Tuning

With the following parameters the system's performance can be influenced.

- **Ignore reference roles:** If you don't use reference roles (see the Administration Manual for details) you can set this option.
- **Ignore hierarchic roles:** If you don't need hierarchic roles you can set this option.
- **Ignore personal substitutions:** If this is set, personal substitutions are ignored.
- **Ignore role substitutions:** If this is set, role substitutions are ignored.
- **Worklist-Cache at server startup:** Specify, whether the worklist cache should be used. *Activated* means that the cache is used; *Started (but not active)* means that data structures are maintained, but the cache is not used for worklist construction; *Switched off* means that the cache is not used and data structures are not maintained.
- **Do not cache Seen Objects:** If this checkbox is activated, seen objects will not be cached anymore.
- **Reload classes:** Reloads classes without server restart if possible. This should be used only in development environments.
- **Statement statistics:** Creates statistics of database statements. If enabled, you will see how often statements have been executed and how much time they consumed (total and average). You can find these statistical information in *Admin-Tasks* → *Servermonitor* → *DB-Statements-Details*. Don't activate statement statistics for long time periods in production environments because they may need a lot of resources and therefore slow down your server.
- **File Cache Size (in Bytes):** Here you can define the size of the web-server file cache. The default value is 1000000 (1MB). If no value is entered, **@enterprise** uses the default value.

- **Allow automatic take:** Allows users to take tasks automatically if they perform a function directly on an entry in the role-worklist or suspension worklist. This will only work if you add additional functions to the GUI of these worklists (e.g., the finish function).
- **Archive Schema:** See section 5 for details about archive schemas and how to set them up.
- **Organizational hierarchy mandatory:** If this is enabled, process instances can only be assigned to roles and users who belong to organizational units that appear in the organizational tree of the process' application.

3.13 SSL

- **Server SSL Port:** Port of the HTTPS server.
- **KeyStore File:** The Java KeyStore is a binary file, which holds the keys and certificates of the system and the certificates of trusted organizations, so called trust anchors. The KeyStore is the central “database” for certificate management. Ensure that there exists a backup of the KeyStore of **@enterprise**.
- **KeyStore Password:** To access a KeyStore a password (with a minimum length of 6 characters) is needed.
- **Password for Server Certificate:** The Java API to access the KeyStore is not able to handle different keys with different key passwords. So a system key password has to be configured to access the keys. This password has a minimum length of 6 characters.
- **Client certificates for HTTPS:** This parameter determines how a secure SSL connection can be established by a client. There are three possibilities:
 - **are not requested:** If this option is selected, SSL connections are established in any case.
 - **are required:** If this option is selected, SSL connections are established only if the client has a valid certificate for authorization.
 - **are requested:** If this option is selected, the establishment of SSL connections depends on the content of the response: if the response contains a valid client certificate the SSL connection is established automatically; if the response contains no valid client certificate a login mask will be displayed to the user and after a successful login the SSL connection will be established.
- **Client certificates for RMI over SSL:** Like the option above, but for RMI clients.

3.14 Password Policy

The parameters in this section are separable in 3 main groups, which are explained in the following paragraphs.

Note: No parameter of these groups is needed to be set, quite the contrary is recommended. If a too strict password policy is established - especially with the parameters of group 2 -, a brute-force attack may be effective in a small amount of time, because of the insufficient number of possible passwords.

So, if you don't want to set a parameter let the input field blank.

3.14.1 General Policy Settings

The following parameters do not focus on the password itself but on the password change- and login-management. These parameters are:

- **Period of Validity (in days):** Defines the password's period of validity in days.
- **Inform user before password expires (in days):** Defines the days before the validation time is expired where the user will get a warning, that his password will expire.
- **Maximal Number of Unsuccessful Logins till Account is Deactivated:** A unsuccessful login is defined as a login attempt of an existing user id with a non valid password. If the specified number of unsuccessful logins are performed between two valid sessions of the specific user, the account is deactivated and the user will get a specific error message on the next login.
- **One-way Hash Algorithm to Use:** The password is stored in encrypted form by using a one-way-hash function. In former releases this algorithm was the Unix Crypt algorithm. Now one of three different algorithms can be chosen.
 - **Unix Crypt:** Is limited to 8 bytes input (that means 8 characters), so it is not recommended to use Unix Crypt furthermore. Nevertheless, to ensure compatibility it is supported further on.
 - **SHA (Secure Hash Algorithm):** Takes a plain string of any length and produces a 160-bit hash output. SHA is said to be secure and is the default value if nothing is configured.
 - **MD5 (Message Digest 5):** Takes a plain string of any length and produces a 128-bit hash output. MD5 is said to be secure and calculates the hash value faster than SHA.

3.14.2 Default Policy Checker Settings

The release is delivered with a default password checker which ensures proper passwords and which is highly configurable. If you need extended configuration options, it is possible to implement a special password checker.

The following parameters of the default checker can be changed to specify the minimum requirements for a password. The default values are 0!

- **Minimal Length of Password:** Specifies the minimal length of a password. As an example, if the parameter is set to 8, the password "soccer" is not accepted, but "icehockey" is. (Recommended:4)
- **Maximal Length of Password:** Specifies the maximal length of a password. As an example, if the parameter is set to 8, the password "hello_its_me" is not accepted, but "hello" is. (Recommended:8)
- **Minimal Number of Letters in Password:** Specifies the minimal number of letters in the password. As an example, if the parameter is set to 1, the password "1234" is not accepted, but "a1234" is. (Recommended:1)
- **Minimal Number of Capital Letters:** Specifies the minimal number of capital letters in the password. As an example, if the parameter is set to 1, the password "hello" is not accepted, but "Hello" is. (Recommended:1)
- **Minimal Number of Lowercase Letters:** Specifies the minimal number of lowercase letters in the password. As an example, if the parameter is set to 1, the password "HELLO" is not accepted, but "hELLO" is. (Recommended:1)
- **Minimal Number of Digits:** Specifies the minimal number of digits in the password. As an example, if the parameter is set to 1, the password "Hello" is not accepted, but "Hello1" is. (Recommended:1)
- **Minimal Number of Special Characters:** Specifies the minimal number of special characters in the password. Special characters are defined as any character which does not belong to any of the following character classes: uppercase characters, lowercase characters, digits, space characters. As an example, if the parameter is set to 1, the password "hello" is not accepted, but "hello*" is. (Recommended:0)
- **Minimal Number of Different Characters:** Specifies the minimal number of different characters in the password. As an example, if the parameter is set to 3, the password "aaaa2222" is not accepted, but "aabb2222" is. (Recommended:3)
- **Maximal Sequence of Same Character:** Specifies the maximal sequence of the same character in the password. As an example, if the parameter is set to 3, the password "aaaa" is not accepted, but "aaabaaa" is. (Recommended:2)
- **Maximal Length of Substring (generated from user data):** Specifies the maximal length of any substring in the password, which exists in the user's first name, surname or id. As an example, if the parameter is set to 3 and the user's id is "testuser", the password "stus" is not accepted, but "tes ser" is. This check is case insensitive. (Recommended:2)
- **Number of Old Passwords to Check:** Specifies the number of old password to check password reuse. As an example, if the parameter is set to 3 and the user changed his password in the order "hello", "itsMe", "myPassword" and "letMeIn", the password "itsMe" is not accepted, but "hello" is. (Recommended:5)

Note: The history check can only cover old passwords, which have been logged in the database. These old passwords are deleted by the LogTask Timer, so if the timer

has deleted all old passwords according to his configuration, the history check can't be performed correctly. The result will indicate a correct password, although the password may have been reused and even be equal to the previous.

Note: If parameters are set in a way that an inconsistent policy is specified, the users may not be able to change their passwords. So please care about the following rules for the parameters:

maximal length \geq minimal length

minimum capitals + minimum lowercase characters + minimum digits + minimal special characters \leq maximal length

minimum letters + minimum digits + minimal special characters \leq maximal length

minimum different characters \leq maximal length

3.14.3 Your Own Checker Class

- **Checker Class:** If the default password checker does not satisfy your requirements, you can enter your own password checker class here. The class must implement the `com.groiss.passwd.Checker` interface.

3.15 Calendar

- **Calendar Class:** Here you can define a class for displaying the holidays in the calendar. It must implement the `com.groiss.cal.Holidays` interface.
- **Show Default Resource:** If this checkbox is checked the user can use a simple resource form for assigning resources to calendar appointments.
- **Resource Classes:** It is possible to use arbitrary forms for calendar resources. The names (incl. package name) of the classes for this forms can be entered into this field. After this the self defined resources can be used in the calendar.
- **Non Working Days:** In this list it is possible to select one or more non working days, which will be needed for example in escalations.

3.16 Time Management

The parameters in this section are needed for setting time management specific properties.

- **Default time unit:** Time unit shown in histograms, duration statistics, etc. You can select between:
 - Seconds

- Minutes
- Hours
- Days
- **Max. length of time histogram:** The maximum length of the time histogram can be set here (in objects).
- **Prune probability:** Here you can set the range of the red area, which has the default range from 0% to 95%. The default value is 0.95.
- **Process deadline probability:** You can select an appropriate deadline based on reliability requirements to new process. The default value is 95% (=0.95).

3.17 ACLCache

In @enterprise it is possible to speed up the rights check by activating the ACLCache (see section 3.17.1). The cache improves the speed of the ACL.hasRight() method calls. The results of calls to method ACL.hasRight() are cached, and the cache is consulted before accessing the database. The cache is organized as an expirable and size bounded LRU cache.

The items have a maximum lifespan associated with them. If an item has been found in the cache, but has expired its lifespan, it is removed from the cache and is reported as being not in the cache. This behavior ensures, that cached right checks do not become unduly outdated. The value lifespan is configurable whereas the default value is 5 minutes.

The cache has also a maximum number of cached elements associated with it. If this number would be exceeded by the insertion of a new cached item, the least recently used item is removed from the cache, thereby ensuring a size bound while providing good hit rate.

Actually, there are two caches, one which stores acl-entries for specific objects and one which stores acl-entries for classes. The parameters for size and lifespan can be configured separately for those two caches.

3.17.1 Configuration of ACLCache

To activate the ACLCache, the following service entry has to be included in the *services* property of the conf-file (*avw.conf*):

```
\ncom.groiss.server.ACLCacheServer aclcache
```

In the administration console (*Administration* → *Configuration*), a section for the ACLCache can be found (provided the service entry has been made in the configuration).

There are five properties to configure:

- **Activated:** Check, if the ACLCache should be activated
- **Max. Number of object specific rights:** Size of the object specific rights cache (in objects).

- **Lifetime of object specific rights (sec.):** Lifetime of rights in the object specific rights cache.
- **Max. Number of class rights:** Size of the class rights cache (in objects).
- **Lifetime of class rights (sec.):** Lifetime of rights in the class rights cache.

3.18 Change Administrator Password

With this link you can change the password of the *sysadm* user. The default password is *digital* (after a default installation of **@enterprise**).

3.19 Initialize Database Schema

This function executes the database initialization again (like when you perform an **@enterprise** installation). This could be useful for setting up **@enterprise** if the database creation failed during setup or if you didn't create the database schema during setup.

If your system is correctly installed, don't execute this function! It might affect and possibly destroy existing data.

3.20 Parameters without GUI

In this section we describe parameters that cannot be configured with the GUI. If you want to add or change them, please open the **@enterprise** configuration file (you can find it in the server root under *conf/avw.conf*), modify the parameters, and restart the server.

- **database.direct.access:** Set this value to 1 to activate the query tool, which you can access in *Admin-Tasks* → *Server* → *Query Tool*. In order to deactivate the query tool, set the value to 0.
- **avw.history.editable:** Set this value to 1 to activate the supplement task, which allows to edit forms in the process-history. In order to deactivate the supplement task, set the value to 0.
- **http.ip-address:** The default-behavior of multiple network-interfaces: the HTTP-server runs on all interfaces. With this parameter you can restrict the interfaces by entering an ip-adress, where the server should run.
- **pred_applet.ext_jars:** Here you can enter a comma-separated list of jar-files, which will be loaded additionally by the process editor. For example you can get an I18N-support for other languages as supported by **@enterprise**.
- **user.select.attrs** and **user.select.attrNames:** The columns of the table in tab *User* of the function *Reassign* (see User Manual - Chapter *Functions of the Worklist Component*) can be modified with these two parameters. The default behaviour is that surname, firstname and id of a user are displayed in the table.
The possible values for parameter *user.select.attrs* are:

- column-names
- getDefaultDept()

It is necessary to use the parameters *user.select.attrs* AND *user.select.attrNames* to ensure correct behaviour.

Example:

```
user.select.attrs=surname,firstName,id,getDefaultDept()
user.select.attrNames=@surname@,@firstName@,@id@,@dept@
```

The table displays the surname, firstname, id and the OU of each user now.

- **avw.goback.abort:** With this parameter it is possible to allow or deny using the function *GoBack* in worklist, if the current step is within an AND- or OR-parallelism. In the following the values of this parameter are explained:
 - 0: The function *GoBack* is not allowed in a parallelism. Within the parallelism the behaviour is like the first step in a process.
 - 1: The function *GoBack* is allowed, if the rights *Abort Step* and/or *Edit Process Instances* are assigned to a user. This value should be the default setting.

For further information about rights, please take a look in the *System Administration Guide - The @enterprise right system*.

Example:

We assume that *GoBack* is allowed and user A and user B have got the rights *Abort Step* and *Edit Process Instances*.

A process contains an AND-parallelism whereas the first branch has a step *andpar1* and the second branch has a step *andpar2*. User A gets task *andpar1* and user B gets *andpar2*. If user A activates the function *GoBack* and send the task to a previous step, task *andpar2* will be removed from the worklist of user B.

- **avw.java.compiler:** With this parameter you can specify the path for the java-compiler (*javac*).
- **avw.decimal.separator:** Allows to set the decimal separator for parsing and displaying decimal numbers. The default is "."
- **httpd.jetty.maxformcontentsize.kb:** Maximal size of form content jetty will accept in KB; default = -1 means jetty native value will be used; 0 means no limit
- **httpd.jetty.headerbuffersize.kb:** Size of jetty buffer for http request headers in KB; default = 8; jetty 6.1.7 default is 4; jetty 6.0.1 default was 8
- **httpd.jetty.requestbuffersize.kb:** Size of jetty buffer for http requests in KB; default = 0; jetty 6.1.7 default is 8; jetty 6.0.1 default was 32

3.20. PARAMETERS WITHOUT GUI

- **httpd.jetty.responsebuffersize.kb:** Size of jetty buffer for http responses in KB; default = 0; jetty 6.1.7 default is 24; jetty 6.0.1 default was 64

4 Clustered @enterprise System

4.1 Overview and Principles of the Clustered Architecture

The clustered architecture supersedes the previous distributed architecture. The aim of the new architecture is to allow for

- increased scalability,
- increased availability,
- easier configuration,
- more flexible operation.

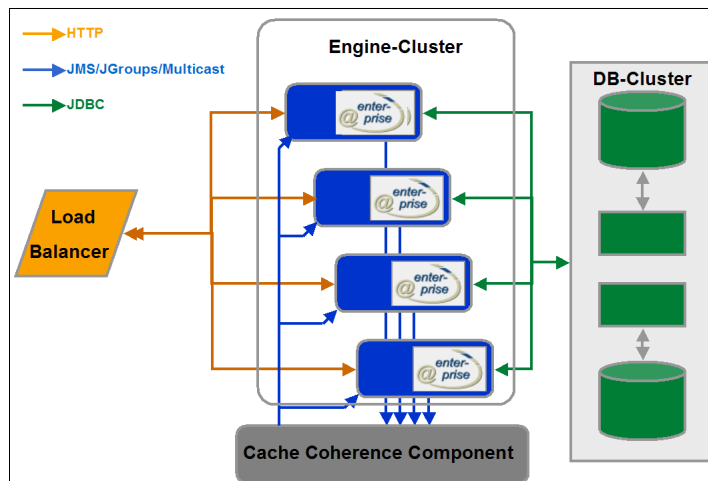


Figure 4.1: Cluster Architecture

Figure 4.1 shows the principal layout of such a cluster. The logical architecture consists of a set of @enterprise engines (termed "nodes") which access a common database and are operated in a peer to peer mode to a large extent. A load balancing mechanism is employed to ensure even load distribution within the cluster. Consistency between the caches in the nodes is ensured by a cache coherence service.

While there are no single points of failure within the cluster nodes, we require the database to be available and scalable to an extent that imposes no bottlenecks for the rest of the system.

4.2 Cluster and Nodes

As already mentioned, a node is a single Java Virtual Machine instance. In a typical production environment, there will be one node running on a single physical machine. In a development or test environment, more than one node could be running on one machine (without enhanced scalability and availability).

The cluster is represented by a single entry in the Server section of the administration. Each node is identified by a Node-Id which must of course be unique within the cluster. Nodes can enter and leave the cluster at runtime. New nodes can be added to the cluster on the fly.

4.3 Configuring a clustered @enterprise System

The clustering of an @enterprise system will typically comprise of the following actions

- configuration of the underlying platforms in terms of hardware, operating system, network and database connectivity and JVM,
- installation of a single (nonclustered) @enterprise system,
- selection of the appropriate transport mechanism for the cache coherence service, its configuration and startup if necessary,
- distribution of the @enterprise installation directory to the nodes,
- adapting the @enterprise configuration,
- starting the nodes.

Details for each of the steps can be found in the following sections.

4.3.1 Platform Configuration

The nodes of an @enterprise cluster can run on a heterogeneous platform as far as the hardware and operating system is concerned. While it is also possible to use different versions of the JVM/JDK it is strongly recommended to use the same principal version for each node (that is e.g. either 1.4 or 1.5, not a mixture of both). If your installation must use different versions, intense testing is strongly advisable.

The requirements for the minimal technical layout of the nodes do not differ from the layout of a single machine. A possible exception are the network interface requirements. It may be advisable to use different physical network interfaces and interconnections for client connections, database connections and possibly for the cache coherence service.

4.3.2 Installation of a nonclustered System

No special issues are arising here because of the cluster. Just install a plain @enterprise system and make sure that it is working.

4.3.3 Transport Mechanisms for Cache Coherence Service

The cache coherence mechanisms task is to propagate cache relevant events within the cluster in order to keep the caches current. For the time being, the following event types are propagated:

- **Workitems:** Changes in the worklist (new items, finished items, ...)
- **Substitution:** Changes in substitutions of users (new substitute, period of substitution starts or ends)
- **Seen Objects:** Items that are new to a user.

We provide the following choice of transport mechanisms to account for different needs of an installation:

- Unreliable Multicast via UDP
- Reliable Multicast via JGroups
- Java Message Service (JMS)

Unreliable Multicast via UDP

While this mechanism is easy to configure and poses virtually no overhead, it is recommended primarily just for development or test installations, due to possible loss of packets. A installation which uses dedicated physical network interfaces and interconnections for cache coherence service might also use unreliable multicast with good results, but one should be aware of the susceptibility to errors. This transport mechanism uses features present in the Java platform, no deployment or startup is needed.

The following configuration parameters are relevant and must be identical on all cluster nodes:

- **Multicast-IP-Address:** Must be a valid multicast address. No two clusters should use the same multicast address. Be aware of other applications using multicast in your configuration. For specification and assignments of multicast addresses, refer to <http://www.isi.edu/in-notes/iana/assignments/multicast-addresses>. Monitoring of multicast packets is quite easy with tcpdump ("tcpdump ip multicast").
- **Multicast IP Port:** Port to send and receive multicast packets. Must be available on the machine.
- **Multicast TTL:** Determines the scope of multicast packets on the network. For clustered systems with small "network diameter" this should be 1.

- **BufferSize (Bytes):** Size of reception buffer in bytes. Recommended value is at least 30000 Bytes.

When specifying these values, be aware of possible address space collisions with a multi-cast based client notification service or cache coherence services of other clusters.

Reliable Multicast via JGroups

JGroups is an open source communications library for reliable group communication. It is written in Java (www.jgroups.org). It is deployed in the @enterprise engine itself and needs no external processes running. It is started automatically. The library itself consists of a single Java archive named `jgroups-all.jar` and uses the apache commons logging facility (`commons-logging.jar`), which must be explicitly added to the classpath (mere placement in the lib directory is not sufficient).

The following configuration parameters are relevant and must be identical on all cluster nodes:

- **Groupname:** JGroups has the notion of communication groups. A member must state the groups he belongs to. Can be an arbitrary string, we recommend to use `epgroup` or to use the name of the server entry in the cluster.
- **Properties:** This parameter specifies the location of a configuration file in XML-syntax.

The recommended configuration is located in the `classes/jgroups/ccs.xml` file. Since the whole JGroups protocol stack is configured through it, it looks rather complicated. But in normal situations, just a handful of key parameters need to be changed. Such parameters are clearly marked in the `ccs.xml` file. The parts of the configuration to be changed are the multicast IP address `mcast_addr`, the multicast port number `mcast_port`, and the time to live `ip_ttl`. For the multicast address and multicast port we refer to the previous section about unreliable multicast, for the time to live we recommend either 1 as the packets should only reach the other @enterprise node which are placed in the network vicinity. If network components are between the nodes of the cluster, it might be necessary to increase this value to 32. In case of doubt, consult your local network administrator. Please avoid any interference within @enterprise (e.g. client notification service with multicast) when selecting multicast parameters.

The other properties in the file should not be changed without intimate knowledge about JGroups.

Java Message Service (JMS)

The usage of JMS for the transport of cache coherence messages can be characterized as follows. The publish subscribe paradigm is used. Per node there is one subscriber and one publisher. All nodes subscribe to the same topic. No message selectors are used. We use nonpersistent, auto-acknowledged, nontransacted messages and nondurable subscribers. JMS does not run within an @enterprise JVM, it must be configured and started separately. The following configuration parameters are relevant and must be identical on all cluster nodes:

- **JMS Provider URL:** The URL name of the JMS provider. For SWIFTMQ this is something like `smqp://<jms-server-name-or-ip-address>:4001/timeout=10000`.
- **JMS ContextFactory:** Name of the Java class for construction of the JNDI-Context. For SWIFTMQ this is `com.swiftmq.jndi.InitialContextFactoryImpl`
- **JMS TopicConnectionFactory:** Java class name for the topic factory of the JMS provider. For SWIFTMQ this is `TopicConnectionFactory`.
- **JMS Topic:** The name of the topic used for communication. Such topics must typically be created within an JMS provider by the administrator.
- **JMS Time to Live (ms):** The JMS provider is free to throw away messages which are older than this timespan. Should be in the range of 30 to 120 seconds.
- **JMS Username:** Name of the user which is utilized for communication with the JMS provider. If this parameter is left empty, an anonymous connection is established. User administration is specific for each JMS provider.
- **JMS Password:** Password for the user mentioned before.

More than one cluster can use a JMS provider, if the names of the topics are kept unique for each cluster. Do not use the same topic name for client notification via JMS and for client notification if you are using the same physical provider for both purposes.

4.3.4 Adapting the @enterprise Configuration

Configuration: Under Admin Tasks / Configuration / Classes / Services, an entry for the cache coherence service must be added as the last service:
`com.groiss.dbcache.coherence.CoherenceService cs`

The following configuration entries are needed in a clustered node under Admin Tasks / Configuration / Cluster

- **Clustering enabled:** Must be checked.
- **Server Name:** Name of the server. Must be the same on each node of one cluster.
- **Node-Id:** Id of the cluster node. Must be unique within the cluster.
- **Performance Factor:** Relative performance factor of the node. Depends largely on CPU power of the node. A node with a factor of 2 is expected to support twice the users of a node with factor 1. The load balancer makes use of the factor to distribute user sessions according to the relative power of the nodes.
- **Clustercheck Tolerance (sec.):** The clustercheck timer sets nodes to inactive where the last heartbeat has not been received for a while (tolerance time). Default value for this property is 30 seconds, changes take effect immediately.

- **Heartbeat Tolerance (sec.):** The load-balancing mechanism excludes nodes from which no heartbeat has been received for more than a specified amount of time. This does not imply that the nodes are set to inactive (this is the job of the clustercheck timer). Here you can set a tolerance time for heartbeats. It is recommended to set the value to two times of the maximum heartbeat timer interval of all nodes. Default value of this property is 10 seconds, changes take effect immediately.
- **Coherence Strategy:** Currently there is just one strategy supported: Notification. Do not confuse this with the client notification mechanism. While the things share the same name, they have nothing in common. In the future, other strategies might be provided as well.
- **Transportlayer for Coherence:** Choose the appropriate transport mechanism like described above.

Ports: If you do run several nodes on one machine (e.g. for testing purposes), ensure that distinct network port numbers for the HTTP server, the HTTPS server and the RMI-mechanism are used.

Directories: If your nodes run on the same machine or access the same remote file systems, be sure to configure each of the nodes with distinct destinations for the log file and the error log file as well as a distinct temporary directory.

Timers: Timers require special consideration in a cluster. There might be timers which should run on each node, and there might be timers that should only be running on one dedicated node of the cluster. The former timers must just be marked by checking the box `runEveryWhere` on the timer edit form.

The latter ones must be marked by NOT checking the box and require special action. In a clustered system one of the nodes assumes responsibility for running the timers. Transparent failover is provided.

To enable this functionality, make sure that two timers are started on each node:

- **HeartBeat:** Should be running on each of the nodes. Periodically writes a timestamp to the database. Used to monitor cluster nodes. During normal operation, there is exactly one update of a single row followed by a commit per heartbeat (and node). The heartbeat mechanism uses a dedicated database-connection when more than five database connections have been configured for the node eliminating hold-ups from finding a connection and overhead from frequently releasing and reacquiring the connection. Recommended periods are in the range of 3 to 10 seconds. Because of these short heartbeat intervals it is recommended to use a dedicated timer thread by assigning a unique thread-id (e.g. "heartbeat") to the timer. This avoids the possible delay of the heartbeat by other (longer-running) timers, thereby getting the heartbeat info to the database as fast as possible.
- **ClusterCheck:** Should be running on each of the nodes. Periodically checks health state of the cluster. Recommended periods are in the range of 120 to 600 seconds. There are two aspects to check for. First, if a node fails to update its timestamp within the tolerance time defined in the *Clustercheck Tolerance* parameter, its state is set to

not running. Second, if none of the nodes runs the timers which are started just once for the whole cluster, one node must assume this role.

4.4 Operation of a clustered system

4.4.1 Monitoring

A cluster health monitor which displays the state for each of the nodes can be accessed via Admin Tasks / Distribution / Check Server States.

The fields displayed are:

- **Hostname:** Name of the cluster.
- **Node-Id:** Id of the node.
- **Start Time:** Time of startup of this node.
- **Last HeartBeat:** Timestamp of last heartbeat made by this node.
- **Running:** Marks if the node is running.
- **Runs ClusterTimers:** Marks if the node is the one which runs the cluster timers.
- **Connected Users:** Current number of users.
- **Performance Factor:** The performance factor of the node.
- **Load Coefficient:** The current load coefficient (number of users divided by performancefactor).

4.4.2 Load Balancing

Principle A client which wants to obtain a load balanced session should first connect to a special URL on an arbitrary running cluster node. There, the client will be redirected to the least loaded node (HTTP-Client) or can obtain hostname and RMI-portnumber of this node (RMI-Client).

HTTP-Clients The URL for getting a load balanced session for an HTTP Client is:

```
http://<host>:<port>/<context-root>/  
servlet.method/com.groiss.avw.html.HTMLNodes.redirect
```

The client will be redirected immediately to the server with the lightest load.

RMI-Clients Use the same mechanism as mentioned above. A client should open an URL-Connection to:

```
http://<host>:<port>/<context-root>/  
servlet.method/com.groiss.avw.html.HTMLNodes.redirectJavaClient
```

A single line is returned containing the hostname and the RMI-port of the server, separated by a colon. This data can be used in the client to obtain a session to that node.

4.4.3 Event Handling

Event handlers are executed on the node where the event has been raised.

5 Setting up an Archive Schema

Large amounts of data can decrease the performance of @enterprise. With the *archive schema* we provide a mechanism to move finished processes to another database schema. This can speed up database operations. In the current version of @enterprise, the archive schema is supported only for Oracle.

It works as follows: A separate database schema is used to store historic data from the three tables `avw_stepinstance`, `avw_forinstance`, and `avw_formversion`. The timer *Archive-Timer* moves all processes that have been finished n days ago to the archive schema. The number of days is taken from the parameter field in the timer entry mask. You can find these processes using the process or extended search when you check the "Add Archive" checkbox. Reactivating a process will move it back to the standard database schema.

For installing an archive schema perform the following steps:

1. Create a new schema, in Oracle a database user.
2. Insert the schema name in the @enterprise configuration parameter "Archive Schema" (group Tuning).
3. Restart the server.
4. Create the tables and views using the following URL:
`http://<host>:<port>/<context-root>/servlet.method/com.dec.avw.timertask.ArchiveTimer.createArchiveSchema`

It is necessary that your standard database user has the rights to create tables and views for the archive schema. For example you can temporary give the dba right to this user.

5. Activate the archive timer and supply values for the timer interval. The timer argument is a single integer n : Processes finished since n days will be moved into the archive schema.

A Database Performance Hints under Oracle

A.1 Preliminaries

The statements in this chapter refer to an @enterprise installation with an Version 8 Oracle DBMS. It is assumed that no atypical characteristics concerning either data distributions or data volumes or transaction volumes like extremely long worklists or BLOBs dominate the system. Further we assume that no other significant workload besides the @enterprise-service is processed on the system (dedicated hardware).

For successful performance improvements, the most crucial issue is to correctly identify and pinpoint system bottlenecks. Applying tuning actions without having a specific hint about the kind or reason for unacceptable performance is not target-oriented. It is essential to isolate and contain the problem area (database, @enterprise server, CPU, memory, network, own application classes, specific user operations). One should apply all means and tools which are offered by the underlying platform to check performance parameters or monitor them on a regular basis. Because of the wide variety of the platforms concerning this specific area, we refer the reader to the appropriate systems documentation.

We assume that the reader has some basic familiarity about the architecture of Oracle and is somewhat acquainted with its significant mechanisms.

A.2 Key Operating Parameters of the Database

The following parameters are vitally important for an efficient operation of the database. They all can be found in the **ini.ora** file.

DB_BLOCK_SIZE States the size of the data blocks in the DB. In most environments the default value is 2048 bytes. For @enterprise the value should be increased to 4096 or 8192. The change should reduce IO-overhead and has no other significant implications. Unfortunately, the value can't be changed in an existing data base, one would be forced to apply a complete export/import cycle to apply a modification.

DB_FILE_MULTIBLOCK_READ_COUNT Determines how many blocks are read during a full table scan. The value should be dimensioned in such a way, that the product of **DB_BLOCK_SIZE** and **DB_FILE_MULTIBLOCK_READ_COUNT** equals the size of the operating system buffer (often 64K). The value can be changed during operations but is applied only at the next startup of the database instance.

DB_BLOCK_BUFFERS States the size of the database block buffer caches in units of blocks. It is an extremely crucial parameter. The default values of Oracle are way too small. For an application system with the characteristics of @enterprise (mostly interactive users in OLTP, insignificant batch processing) one should configure the cache size to achieve a hit rate above 95% to 98% in regular operations. Regular monitoring is essential. One could apply the following queries (as user SYSTEM) to determine current hit rates:

```
select
  SUM(DECODE(Name, 'consistent gets', Value, 0)) Consistent,
  SUM(DECODE(Name, 'db block gets', Value, 0)) Dbblockgets,
  SUM(DECODE(Name, 'physical reads', Value, 0)) Physrds,
  ROUND(((SUM(DECODE(Name, 'consistent gets', Value, 0))+
    SUM(DECODE(Name, 'db block gets', Value, 0)) -
    SUM(DECODE(Name, 'physical reads', Value, 0)) )/
    (SUM(DECODE(Name, 'consistent gets', Value, 0))+
    SUM(DECODE(Name, 'db block gets', Value, 0))))
    *100,2) Hitratio
from V$SYSSTAT;
```

```
column HitRatio format 999.99
select Username,
  Consistent_Gets,
  Block_Gets,
  Physical_Reads,
  100*(Consistent_Gets+Block_Gets-Physical_Reads)/
  (Consistent_Gets+Block_Gets) HitRatio
from V$SESSION, V$SESS_IO
where V$SESSION.SID = V$SESS_IO.SID
and (Consistent_Gets+Block_Gets)>0
and Username is not null;
```

If an unsatisfactory hit rate is measured, **DB_BLOCK_BUFFERS** should be increased in steps of 15% to 25%, until hit rate levels out. Meaningful measurements are only possible in real production mode and not immediately after the startup phase of the instance when the cache is still cold.

It is common knowledge, that the buffer cache should not be increased beyond certain thresholds. Each word of main memory that is allocated exclusively for the buffer cache can be in high demand by other system components. In no way the machine should be

A.2. KEY OPERATING PARAMETERS OF THE DATABASE

pressed to swapping or paging activities. After every expansion of buffer cache size, measurements with a warm cache are called for in combination with keeping an eye on paging or thrashing. Memory expansions should be considered at such points.

SHARED_POOL_SIZE Determines the size of the shared pool in the System Global Area (SGA). Oracle defaults are often found to be too small.

A rule of thumb says that 15% to 20% of the shared pool should stay free.

The current size can be calculated as follows:

```
select value from v$parameter where name='shared_pool_size';
```

The free space is returned by this query:

```
select name, bytes from v$sgastat where name='free memory';
```

Key elements in the shared pool are the library cache and the data dictionary. Miss rates for both components can be determined with the help of the following queries. In the library cache miss rates of under 1% and of under 5% in the data dictionary are commonly seen as appropriate.

```
column "Executions" format 9,999,999,990
column "Cache Misses Executing" format 9,999,999,990
column "Data Dictionary Gets" format 9,999,999,999
column "Get Misses" format 9,999,999,999
column "% Ratio" format 999.99
```

```
select sum(pins) "Executions",
       sum(reloads) "Cache Misses Executing",
       (sum(reloads)/sum(pins)*100) "% Ratio"
from v$librarycache;
```

```
select sum(gets) "Data Dictionary Gets",
       sum(getmisses) "Get Misses",
       100*(sum(getmisses)/sum(gets)) "% Ratio"
from v$rowcache;
```

If higher miss rates are measured, we advise a similar procedure like in the case of the `DB_BLOCK_BUFFERS` parameter.

SORT_AREA_SIZE Size of the area in the main memory which is reserved for each user for in-memory sorting operations. If disk-based sorts make up for more than 5% to 10% of the in memory sorts, then `SORT_AREA_SIZE` should be increased. The current configuration can be determined with:

A.3. OPTIMIZER

```
select substr(name,1,25) Name,  
       substr(value,1,15) Value  
from V$PARAMETER  
where Name = 'sort_area_size';
```

Statistics about the number of sorts, separately for main memory and disk based sorts are implemented by:

```
select substr(name,1,25) Name,  
       substr(value,1,15) Value  
from V$SYSSTAT where name like 'sort%';
```

LOG_BUFFER Size of the redo log buffer in the SGA.
The current size can be obtained by:

```
select substr(name,1,25) Name,  
       substr(value,1,15) Value  
from V$SGA  
where Name = 'Redo Buffers';
```

If redo log space requests are issued in the database, there might be a bottleneck here. The following query investigates this:

```
select substr(name,1,25) Name,  
       substr(value,1,15) Value  
from v$sysstat  
where name = 'redo log space requests';
```

The value should approximate zero. If this is not the case, one should increase the LOG_BUFFER parameter in steps of 50% to 100%. It might be advisable to increase the shared pool size by the same (absolute) amount.

A.3 Optimizer

Cost based optimization is the way to go with Oracle. In general, better query plans can be generated than pure rule based optimization could achieve.

To activate the cost based optimizer, the parameter OPTIMIZER_MODE in init.ora must be set to CHOOSE. It is also necessary to statistically analyze the data distribution and index selectivity.

Oracle offers commands of the form analyze table <mytable> compute statistics. One can supplement statistics for an entire schema using execute dbms_utility.analyze_schema('USER','COMPUTE');. The 'USER' element should be replaced by the name of the @enterprise data base user.

It is highly advisable to run this command from time to time. In any case, it should be run periodically during the first period of production use and additionally when significant

configuration changes (new applications, other data volumes) take place. The analysis is quite resource intensive and should not be applied during peak operational hours. Sufficient temporary tablespace must be provided, also. A practical trade-off between statistical accuracy and resource consumption can be achieved through use of 'ESTIMATE' instead of 'COMPUTE'. In this case the system takes samples of the data and does not go through the entire volume. A good strategy might be to establish a batch-job which issues this schema analysis commands on a regular (weekly) basis.

A.4 Storage

A.4.1 Disks

The main performance issues in the disk subsystem are the separation of random access and sequential access and further to isolate individual sequential accesses.

More precisely, separate the redo-logs, the after image files and the rollback segments, and put them on individual disks without any further activity.

Further split up SYSTEM and TEMPORARY tablespaces from the rest of the system.

Tables with particular high activity on them are AVW_STEPINSTANCE, AVW_FOLLOWS and AVW_FORMVERSION. A good measure would be to place them together with their indices on separate tablespaces, to be able to place them on specific disks and to distribute the load on multiple devices. Another possible strategy would be the division of index space and table data space in different tablespaces.

It is not possible to give general advice without deeper knowledge of the operational characteristics. Nevertheless, for an installation with significant size, we strongly recommend to devote some thoughts to this issues and to divert from the default configuration.

An overview about IO distribution over the individual datafiles can be gained by:

```
select DF.Name File_Name,
       FS.Phyblkrd Blocks_Read,
       FS.Phyblkwrt Blocks_Written,
       FS.Phyblkrd+FS.Phyblkwrt Total_IOs
from V$FILESTAT FS, V$DATAFILE DF
where DF.File#=FS.File#
order by FS.Phyblkrd+FS.Phyblkwrt desc;
```

A.4.2 Parameters for Tablespaces

Appropriate default storage parameters for the tablespaces would be:

```
alter tablespace AVW default storage
(initial 256k next 256k maxextents 200 pctincrease 0);
```

Instead of AVW, state the tablespaces which are used to store the @enterprise tables and indexes, in particular the default tablespace of the @enterprise database user. For some tables which can be assumed to have a greater size than that (50MB) like AVW_STEPINSTANCE,

AVW_FOLLOWS and AVW_FORMVERSION, the storage parameters can be changed in full operation mode; e.g.:

```
alter table <mytable> storage(next 1M maxextents 1200);
```

With this statement, table <mytable> can use 1000 additional extents, each being 1 MB in size when one assumes that 200 extents were already used. It is generally advisable to use zero as value for `pctincrease`, to avoid exponentially increasing storage demand for extents.

A.5 One owns Tables and Queries

For own tables which are used to store application relevant data, exactly the same considerations like for system tables according to table placement and to storage parameters should be made. In particular, popular access paths should be supported by appropriate (multi-column) indexes.

Queries of application tables should generate a result set as small as possible. It is recommended to use a two phase approach for queries with potentially large result sets. First, the number of tuples (`count(*)`) should be determined. If this number exceeds a certain threshold, it is time to give the user a chance to decide upon further execution of the query. The user could apply additional constraints to the search condition which would further confine the result set, or she could explicitly get the whole large result set (and thereby accepting higher response time and workload on the server).

For medium sized tables, which are often scanned in their entirety, table level caching could be advantageous:

```
alter table mytable cache;
```

Clearly, sufficient space in form of `DB_BLOCK_BUFFERS` must be provided.

Criteria in queries should be used in such a way that indexes get used. Strive for point queries or at least for multipoint queries with high selectivity. (its better to use `a='b'` than a like `'b%'` which is in turn better than a like `'%b%'`).

Of uttermost importance is the usage of the @enterprise transaction cache mechanism, which works for all subclasses of `SQLObject`. Access to such objects should be done through `receiver.get(oid)` and not via `receiver.get('oid=xxx')`;

Performance friendly formulation of application queries (especially such statements which are executed quite often) call for generation, interpretation and perhaps modification of the execution plans. Measures could be the definition of additional indices or clustering on a physical level or semantic preserving reformulation of the query or explicit incorporation of query optimization hints.

Concerning these issues we refer to the 'Oracle8 Tuning' and 'Oracle8 Concepts' and 'Oracle8 Application Developers Guide' manuals. Consider the possibilities of `TKPROF` and `EXPLAIN PLAN`. The logfile of @enterprise may have valuable first hints like duration of SQL statements.

It is much better to run complex queries in their entirety on the DB-server than to overflow the server with lots and lots of simple individual queries and to stick their results together in the @enterprise server. This is due to relatively high startup and communication overhead and context switches between the two servers.