



@enterprise 8.0

Installation and Configuration

December 2017

Groiss Informatics GmbH

Groiss Informatics GmbH

Strutzmannstraße 10/4
9020 Klagenfurt
Austria

Tel: +43 463 504694 - 0
Fax: +43 463 504594 - 10
Email: support@groiss.com

Document Version 8.0.22989

Copyright © 2001 - 2017 Groiss Informatics GmbH.
All rights reserved.

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Groiss Informatics GmbH does not warrant that this document is error-free.

No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of Groiss Informatics GmbH.

@enterprise is a trademark of Groiss Informatics GmbH, other names may be trademarks of their respective companies.

Contents

1	System Requirements	6
1.1	Platform	6
1.2	Java	6
1.3	Database Management Systems	6
1.4	Client	7
2	Installation	8
2.1	Database Preparation	8
2.1.1	Oracle	8
2.1.2	MS SQL-Server	10
2.1.3	DB2	11
2.1.4	PostgreSQL	11
2.1.5	Derby	12
2.2	Extract and Install	12
2.2.1	Bootstrap in stand-alone server (Jetty)	13
2.3	Installing a Service	14
2.3.1	Components of the Framework	14
2.3.2	Migrating to the new framework	15
2.4	Using an Application Server or Servlet Container	17
3	Configuration	18
3.1	License	18
3.2	HTTP server	18
3.2.1	Defining allowed and denied hosts or networks	22
3.2.2	Access Control	22
3.3	Database	25
3.4	Directories	26
3.5	Logging	28
3.6	Classes	28
3.7	Localization	29
3.8	Communication	31
3.9	Cluster	33
3.10	Workflow	33
3.11	DMS	34

3.12	Search	35
3.13	Tuning	37
3.13.1	ACLCache	38
3.14	Security	40
3.15	Password policy	40
3.15.1	General Policy Settings	40
3.15.2	Default Policy Checker Settings	41
3.15.3	Your Own Checker Class	43
3.16	Calendar	43
3.17	Process cockpit	44
3.18	Time management	44
3.19	Change administrator password	45
3.20	Initialize database scheme	45
3.21	Parameters without GUI	45
4	Patching and Upgrading your Installation	50
4.1	Patching the Installation	50
4.1.1	Manual Patch Method	50
4.1.2	Automatic Patch Method	51
4.2	Upgrading/Patching an @enterprise Application	52
4.3	Performing an Upgrade of @enterprise	52
4.4	Migration of deprecated DBMS data types	53
4.4.1	Migration of Oracle data type LONG	53
4.4.2	Migration of deprecated MS SQL-Server data types	54
5	Clustered @enterprise System	56
5.1	Overview and Principles of the Clustered Architecture	56
5.2	Cluster and Nodes	57
5.3	Configuring a clustered @enterprise System	57
5.3.1	Platform Configuration	57
5.3.2	Installation of a nonclustered System	58
5.3.3	Transport Mechanisms for Cache Coherence Service	58
5.3.4	Adapting the @enterprise Configuration	60
5.4	Operation of a clustered system	62
5.4.1	Monitoring	62
5.4.2	Load Balancing	63
5.4.3	Event Handling	63
6	Setting up an Archive Schema	64
7	@enterprise and Datasources	65
7.1	Configuration of @enterprise	65
7.2	Configuration of a Datasource in Tomcat6	65
7.3	Configuration of a Datasource in Jetty 6.1	66
7.4	Considerations for pooled Datasources	68

A Database Performance Hints under Oracle	69
A.1 Preliminaries	69
A.2 Key Operating Parameters of the Database	69
A.3 Optimizer	72
A.4 Storage	73
A.4.1 Disks	73
A.4.2 Parameters for Tablespaces	73
A.5 One owns Tables and Queries	74

1 System Requirements

1.1 Platform

@**enterprise** 8.0 is available for several platforms. For the operation of a server, a Java Runtime Environment (JRE) of Version 1.5 or higher is required (version 1.6 is recommended). The following operating systems are supported:

- Windows Variants (NT, 2000, XP, 2003, 2008, Vista, Win7)
- Solaris
- AIX
- Linux

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

1.2 Java

To develop @**enterprise** applications, a Java Development Kit (JDK) version 1.6 or higher must be installed. It is available for download from the Oracle web site (<http://java.oracle.com>) or from another vendor. At the Oracle web site, a list of Java ports to other platforms is available.

On clients where the process editor applet is to be used, we require the Java Plugin (version 1.6 or higher) to be installed in the browser.

1.3 Database Management Systems

We support the following DBMSs: Oracle, MS SQL Server, IBM's DB2, PostgreSQL, Derby. MySQL and Firebird is supported experimentally.

The following database versions are required:

- Oracle 9i or higher
- MS SQL Server 2000 or higher

1.4. CLIENT

- DB2 9.7 or higher on Windows or AIX
- PostgreSQL V8.4 or higher
- Derby 10.5.3.0 or higher
- MySQL 5.0 (experimental)
- Firebird Version 2.5 or higher (experimental)

The database can be installed on the same machine as @**enterprise** or on another networked server.

1.4 Client

In order to use the the Web-Client, a Web-Browser is all that is needed. Supported Products and Versions are:

- MS Internet Explorer 7 or higher (limited support for version 6)
- Firefox 3.0 or higher
- Safari 5.0 or higher
- Chrome 15 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 supported databases.

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 also have access to a *tablespace* and the permission to add data there.

Example (*EP_USER* is the name of the @**enterprise** database user):

```
create user EP_USER identified by <password> default tablespace users;
grant create session, alter session to EP_USER;
grant create table, create view to EP_USER;
grant unlimited tablespace to EP_USER;
```

Since Oracle 11g, a default profile mechanism with resource limitations and password expiration settings might lead to immediate lockout when getting the password wrong or to lockout after a password expiration intervall. It is recommended to check the applicable profile parameters and to change them appropriately for the @**enterprise** database user. An unlimited profile can be created with:

```
create profile EP_UNLIMITED_PROFILE limit
  composite_limit unlimited
  connect_time unlimited
  cpu_per_call unlimited
  cpu_per_session unlimited
  failed_login_attempts unlimited
  idle_time unlimited
```


2.1. DATABASE PREPARATION

```
logical_reads_per_call unlimited
logical_reads_per_session unlimited
password_grace_time unlimited
password_life_time unlimited
password_lock_time 1
password_reuse_max unlimited
password_reuse_time unlimited
password_verify_function null
private_sga unlimited
sessions_per_user unlimited;
```

The specific requirements for your site may vary, in case of doubt check with your local DBA. The profile can be assigned to the user with:

```
alter user EP_USER profile EP_UNLIMITED_PROFILE;
```

Other useful commands for account administration and trouble shooting are:

- *Check the users profile:*

```
select profile from dba_users
where username = 'EP_USER';
```

- *Check the properties of the profile:*

```
select resource_name, limit from dba_profiles
where profile=
(select profile from dba_users
where username = 'EP_USER');
```

- *Check the users account state:*

```
select username,profile,account_status,expiry_date from dba_users
where username='EP_USER';
```

- *Unlock a users account:*

```
alter user EP_USER account unlock;
```

- *Unexpire an account or change a users password:*

```
alter user EP_USER identified by <password>;
```

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 EP_USER;

2.1. DATABASE PREPARATION

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

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

If the use of full-text search or WfXML2 functionality is intended with Oracle as the underlying DBMS, you must select the Oracle LOBs database type in the configuration (and not the legacy mode with Oracle LONGs). 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.

Oracle offers the possibility to set the semantic of varchar/varchar2 datatypes (BYTE or CHAR). The decision of setting the correct type could be necessary, if UTF-8 texts should be stored. An example could be that a field has a length-restriction of 100 characters and the text to be stored contains 100 characters with 2 umlauts. Because of UTF-8 encoding the text will grow up to 102 Byte and could not be stored anymore.

For this purpose you can change the semantic on two ways:

- global by using following statement:

```
alter system set NLS_LENGTH_SEMANTICS='CHAR' scope=both;
```

- per session (db-session-environments in **@enterprise** configuration) by using following statement:

```
alter session set NLS_LENGTH_SEMANTICS='CHAR';
```

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`.

It is advisable to use Statement-Level Snapshot Isolation in order to avoid shared locks by readers. Enable it with:

```
ALTER DATABASE <dbname> SET READ_COMMITTED_SNAPSHOT ON;
```

Note that no other users are permitted to be in the database when you issue this command and that the feature is available only in SQLSserver 2005 or higher.

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*. Set the character set of the database to UTF-8 and the standard size of the buffer pool and table space to 16 KB. Then you create a database schema, for which the user is authorized.

2.1.4 PostgreSQL

During installation of PostgreSQL choose at least the following components:

- Database Server
 - Data Directory
 - National Language Support
- User interfaces
 - psql
 - pgAdmin III (optional admin GUI)
- Database Drivers : JDBC Driver

In the "initialize database cluster" dialog, it is advisable to use UTF-8 as encoding and to check "accept connection on all interfaces" if remote connections to the database are needed.

In the data subdirectory of the installation directory, edit `pg_hba.conf` to allow access from remote machines if desired. Example:

```
host all all 10.10.10.0/24 md5
```

In the data subdirectory of the installation directory edit `postgresql.conf`. Make sure that parameter `default_with_oids` is turned off:

```
default_with_oids = off
```

It may be advisable to restrict the search path to the current user schema with:

```
search_path = "$user"
```

After a configuration reload or restart of the server, the PgAdmin III gui or the psql command line can be used to

- Create a User Account (Login Role in Postgres Terminology)

```
CREATE ROLE ep LOGIN PASSWORD 'eppasswd'  
NOINHERIT VALID UNTIL 'infinity';
```
- Create a Schema (the id of the schema and of the role must be identical):

```
CREATE SCHEMA ep AUTHORIZATION ep;
```

To activate support for the soundex search, use the following command (the `fuzzystrmatch.sql` file is located in the `share/contrib` directory).

```
psql -U postgres {dbname} < fuzzystrmatch.sql
```

2.1.5 Derby

Derby doesn't need any preparation. Derby is perfectly suited for development purposes and test deployments. For heavyweight multiuser installations the use of Derby may not be advisable.

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 `setup80.jar`. The installation can be started with a double-click on the file. The installation of a Java JRE 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:

```
java -jar setup80.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 (`ep.sh` or `ep.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 27 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 27 or consult the documentation of the driver. @enterprise allows to configure datasources too. For further information take a look in chapter 7.
 - 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.
 - 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. If a schema of a (previous) @enterprise version exists, setup cannot be continued at this point!
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.

2.2.1 Bootstrap in stand-alone server (Jetty)

Since @enterprise 8.0 the bootstrap mechanism is used, which builds the classpath automatically. This mechanism allows to keep the batch- and/or shell-file simple and clear. Following configurations are possible in classpath using *com.groiss.component.Bootstrap* in *ep.bat* or *ep.sh*:

The java property *-Dep.bootstrap.path* can be changed optionally, so additional paths can be added to classpath with following behavior:

- *classes*: all files within this folder are added to classpath
- *lib*: all files with extension **.jar* are added to classpath
- **.jar*: the corresponding file is added to classpath
- all other paths are scanned for a *classes* and *lib* directory and the corresponding entries will be added to classpath. If these directories are not available, the entered directory will be added to the classpath.

Hint: The first entered path (leftmost) of property *-Dep.bootstrap.path* is loaded first, the rightmost path is loaded at last. The jar-files of the *lib* directory are loaded in alphabetical order.

Example:

```
%JAVACMD% -Xms16m -Xmx128m -Djava.awt.headless=true
-Dep.bootstrap.path=C:/eproot;C:/extension/classes;../libs/lib;C:/myjar.jar;.
com.groiss.component.Bootstrap conf\avw.conf
```

- *C:/eproot* is scanned for a *classes* and *lib* directory
- *C:/extension/classes* is added to the classpath
- *../libs/lib* results in adding all included jar files to classpath (scanned relative to root-path)
- *C:/myjar.jar* is added to the classpath
- *.* means, that the root-path is scanned for a *classes* and *lib* directory. If these directories are not available, all elements of the root-path will be added to the classpath.

Hint: If property *-Dep.bootstrap.path* is set, only these paths/files will be considered, i.e. the default behavior of *@enterprise* classpath will be disabled.

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 *install.bat* in the directory *service*.

@enterprise uses the framework *Java Service Wrapper* from Tanuki Software.

2.3.1 Components of the Framework

The service framework consists of the following files in the service subdirectory:

Common wrapper files

- `install.bat`: Used to install the service. Needs to be executed once or after changes in the service configuration.
- `uninstall.bat`: Used to uninstall the service. Needs to be executed before certain changes in the service configuration get effective.
- `wrapper.dll`, `wrapper.exe` and `wrapper.jar`: The core components of the wrapper.
- `license-wrapper.txt`: The license that governs the use of the wrapper.
- `run.bat`: Can be used to execute the wrapped program in the foreground (not as service) for debugging purpose.

Specific service template for @enterprise

`wrapper.conf`

Utility to send a CTRL-BREAK signal to a process

`SendSignal.exe`

Use

`SendSignal <pid>`

to get a threaddump. The `<pid>` is the PID of the Java/JVM process, not of the wrapper. The thread-dump is captured to `services/wrapper.log` file.

2.3.2 Migrating to the new framework

Since @enterprise 8.0 the new framework from Tanuki Software is used, the old one was *JavaService* from objectweb.org. Using the new framework has the following benefits:

- No restart problems under Windows 2003.
- Startup parameters can be changed easily in a plain textfile via `services/wrapper.conf`, there is no need for registry manipulation.
- Threaddumps can be captured to the logfile of the new service framework
 - `sendsignal.exe` in `./services`
 - dumps into `services/wrapper.log`

Obey following steps to migrate to the new framework:

1. Save the older service definition: Using `regedit`, go to `HKLM/System/CurrentControlSet/Services/<ServiceName>` and export this subtree to a file (say `epserviceold.reg`).
2. Edit the `wrapper.conf` file: Use the corresponding entries from the saved registry entries as a guideline here:

2.3. INSTALLING A SERVICE

- *Specify the Java directory:* Replace `%javadir_placeholder%` in the following line with the directory containing the desired java version:

```
wrapper.java.command=%javadir_placeholder%/bin/java
```

- *Adapt the classpath:* Replace `%classpath_placeholder%` in the following line with the desired classpath:

```
wrapper.java.classpath.2=%classpath_placeholder%
```

- *Memory size:* The value of the `-Xms` and `-Xmx` parameter of the JVM for heap sizing can be specified in the following way:

```
# Initial Java Heap Size (in MB)
```

```
wrapper.java.initmemory=32
```

```
# Maximum Java Heap Size (in MB)
```

```
wrapper.java.maxmemory=128
```

- *Additional parameters to the JVM:* Can be specified via following parameters:

```
wrapper.java.additional.<nr>
```

- *Configure servicename and description:* In the following 3 entries replace `%servicename_placeholder%` with the service name. A different service name is recommended, so you can keep the old service definition until you are sure the new one is working.

```
wrapper.console.title=%servicename_placeholder%
```

```
wrapper.ntservice.name=%servicename_placeholder%
```

```
wrapper.ntservice.displayname=%servicename_placeholder%
```

Give a meaningful description of the service via following parameter:

```
wrapper.ntservice.description
```

- *Dependencies in Startup (optional):* Via `wrapper.ntservice.dependency.<nr>` parameters, one can introduce dependencies upon other services (e.g. databases) for startup.
3. Stop the old service: Use the service manager to stop the service.
 4. Install the new service: Execute `<epdir>/service/install.bat`
 5. Start the new service: Start the new service using the service manager. Check the logfile, make sure the new service is running properly. If there are any problems, check the logfile `services/wrapper.log` and the parameters.
 6. Delete the old service: `<epdir>/service/javaservice -uninstall <oldservicename>`
 7. Remove obsolete files: `javaservice.exe` and `install.ep.bat` can be removed from the `<epdir>/service/directory`.

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 `ep80.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 `ep80.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 `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 `ep80.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`.

Hint: The database name of the Derby JDBC-URL (the 'ep' part in `jdbc:derby:ep;create=true`) in embedded mode is relative to either the current directory or relative to the directory specified in the `derby.system.home` system-property (if this is present). In a scenario of deploying multiple **@enterprise** systems as different web-applications in one servlet-container, with each of the systems using a dedicated embedded derby instance, use unique path names to the database files per web-application (e.g. `jdbc:derby:databases/app1/derbydb;create=true` and `jdbc:derby:databases/app2/derbydb;create=true`).

If Tomcat is used as Servlet Container and UTF-8 encoding should be used, you have to set following attributes in Tomcat's `server.xml`:

- `URIEncoding="UTF-8"`
- `useBodyEncodingForURI="true"`

Typically:

```
<Connector port="8080"
maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
enableLookups="false" redirectPort="8443" acceptCount="100"
debug="0" connectionTimeout="20000"
URIEncoding="UTF-8" useBodyEncodingForURI="true"
disableUploadTimeout="true" />
```

Hint: Since **@enterprise** 8.0 the Servlet API 2.5 is used! This can lead to compatibility problems with some application server, e.g. Tomcat in versions less than 6.0.

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 *Save* button in toolbar, which is available on every configuration page. After activating this button, the changes are stored in the file *avw.conf*, which can be found in the folder *classes* of **@enterprise** installation directory. When changing settings via GUI, no server restart is necessary, excepting the notification icon appears (yellow triangle)!

Hint: The parameter definition and their groups are defined in *properties.xml*. This file should not be changed!

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 - avw.license:** 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 - httpd.port:** HTTP port on which the server runs.
- **Minimum number of threads - httpd.minthread:** Number of threads, which are started on startup.

3.2. HTTP SERVER

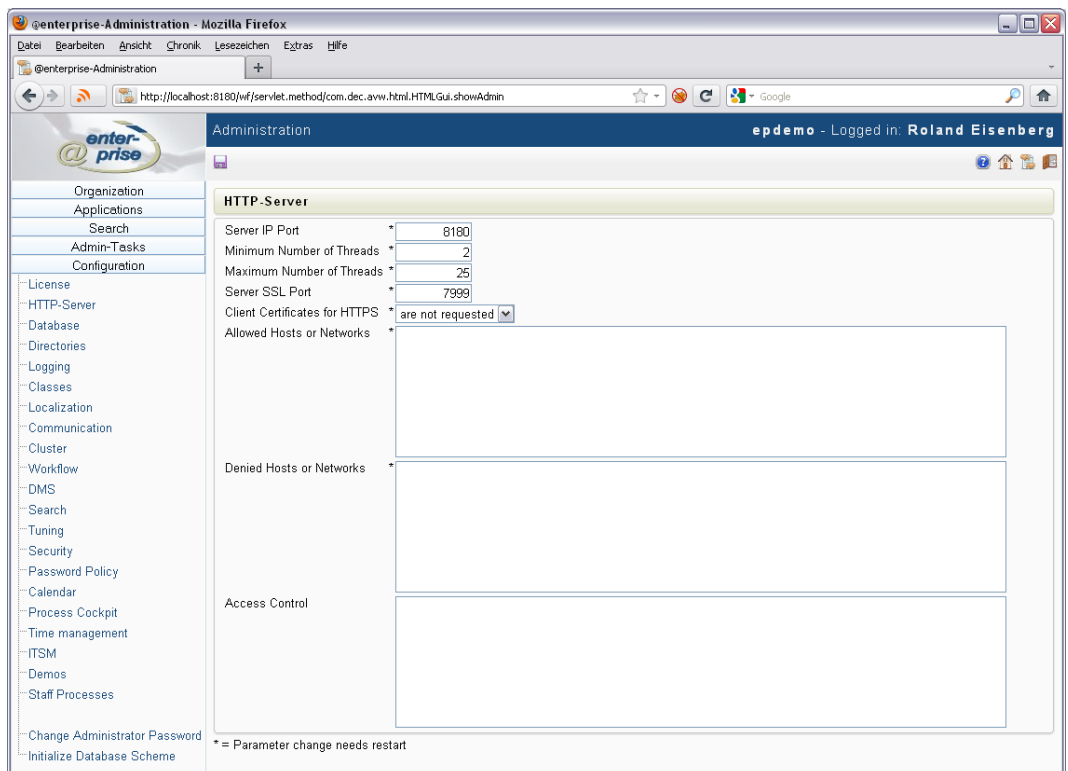


Figure 3.1: @enterprise Configuration

- **Maximum number of threads - `httpd.maxthreads`:** Maximum number of threads, which will be used for HTTP requests.

Hint: If Apple Safari is used in combination with SSL, it is recommended to set an adequate high number for *Minimum Number of Threads* and *Maximum Number of Threads*.

- **Server SSL port - `ssl.port`:** Port of the HTTPS server.
- **Client certificates for HTTPS - `ssl.requireclientcertificate`:** 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.
- **Allowed hosts or networks - `httpd.hosts.allow`:** 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.
- **Denied hosts or networks - `httpd.hosts.deny`:** Analogous to above.
- **Access control - `urls.allowed`:** 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.
- **Excluded cipher suites in HTTPS - `httpd.jetty.sslconnector.excludeciphersuites`:** Vulnerable SSL cipher suites can be excluded from use in HTTPS with following line:

```
httpd.jetty.sslconnector.excludeciphersuites=  
TLS_RSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_RSA_WITH_AES_256_CBC_SHA,  
\r\nTLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,  
\r\nTLS_ECDH_RSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_ECDH_RSA_WITH_AES_256_CBC_SHA,  
\r\nTLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,  
\r\nTLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,
```

3.2. HTTP SERVER

```
\r\nTLS_DHE_RSA_WITH_AES_128_CBC_SHA,  
\r\nTLS_DHE_RSA_WITH_AES_256_CBC_SHA,  
\r\nTLS_DHE_DSS_WITH_AES_128_CBC_SHA,  
\r\nTLS_DHE_DSS_WITH_AES_256_CBC_SHA,  
\r\nSSL_RSA_WITH_3DES_EDE_CBC_SHA,  
\r\nTLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA,  
\r\nTLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA,  
\r\nTLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA,  
\r\nTLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA,  
\r\nSSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA,  
\r\nSSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA,  
\r\nSSL_RSA_WITH_DES_CBC_SHA,  
\r\nSSL_DHE_RSA_WITH_DES_CBC_SHA,  
\r\nSSL_DHE_DSS_WITH_DES_CBC_SHA,  
\r\nSSL_RSA_EXPORT_WITH_RC4_40_MD5,  
\r\nSSL_RSA_EXPORT_WITH_DES40_CBC_SHA,  
\r\nSSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA,  
\r\nSSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA,  
\r\nSSL_RSA_WITH_NULL_MD5,  
\r\nSSL_RSA_WITH_NULL_SHA,  
\r\nTLS_ECDH_ECDSA_WITH_NULL_SHA,  
\r\nTLS_ECDH_RSA_WITH_NULL_SHA,  
\r\nTLS_ECDHE_ECDSA_WITH_NULL_SHA,  
\r\nTLS_ECDHE_RSA_WITH_NULL_SHA,  
\r\nSSL_DH_anon_WITH_RC4_128_MD5,  
\r\nTLS_DH_anon_WITH_AES_128_CBC_SHA,  
\r\nTLS_DH_anon_WITH_AES_256_CBC_SHA,  
\r\nSSL_DH_anon_WITH_3DES_EDE_CBC_SHA,  
\r\nSSL_DH_anon_WITH_DES_CBC_SHA,  
\r\nTLS_ECDH_anon_WITH_RC4_128_SHA,  
\r\nTLS_ECDH_anon_WITH_AES_128_CBC_SHA,  
\r\nTLS_ECDH_anon_WITH_AES_256_CBC_SHA,  
\r\nTLS_ECDH_anon_WITH_3DES_EDE_CBC_SHA,  
\r\nSSL_DH_anon_EXPORT_WITH_RC4_40_MD5,  
\r\nSSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA,  
\r\nTLS_ECDH_anon_WITH_NULL_SHA,  
\r\nTLS_KRB5_WITH_3DES_EDE_CBC_SHA,  
\r\nTLS_KRB5_WITH_3DES_EDE_CBC_MD5,  
\r\nTLS_KRB5_WITH_DES_CBC_SHA,  
\r\nTLS_KRB5_WITH_DES_CBC_MD5,  
\r\nTLS_KRB5_EXPORT_WITH_RC4_40_SHA,  
\r\nTLS_KRB5_EXPORT_WITH_RC4_40_MD5,  
\r\nTLS_KRB5_EXPORT_WITH_DES_CBC_40_SHA,  
\r\nTLS_KRB5_EXPORT_WITH_DES_CBC_40_MD5
```

The cipher suite used by a client can be seen via the URL

```
...servlet.method/com.dec.avw.html.HTMLAdmin.clientInfo.
```

Dealing with sporadic SSL-Handshake problems is greatly eased by setting the `javax.net.debug` system property in the java command line, eg.:

```
-Djavax.net.debug=ssl:defaultctx:sslctx:handshake:verbose
```

This generates considerable amounts of log data, usage is only advisable when client connection issues via HTTS arise. An on-line assessment of your SSL parameters can be obtained at <https://www.ssllabs.com/ssldb/index.html>

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 the ones 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-Addresses and net-mask separated by spaces, commas or new-lines. Both IPV4 and IPV6 addresses are permissible. A net-mask should be given in the CIDR style in form of an integer specifying the number of bits of the network-part. For IPV4 addresses, the traditional dotted notation is also permitted. See the following example:

```
10.205.112.0/255.255.255.0
10.205.224.0/24
2001:0db8:0010::/48
```

This entries in the allow-list means, access from the networks `10.205.112.*`, `10.205.224.*` and `2001:0db8:0010:*` is allowed. When entering IPV6 addresses directly in the config-file, bear in mind that each colon (:) must be escaped by preceding it with a backslash.

The following list used for the allow-list causes that access from hosts `10.205.112.4`, `10.205.224.8` and `2001:DB8:0010:0:8:800:200C:417A` is allowed.

```
10.205.112.4/32
10.205.224.8/255.255.255.255
2001:DB8:0010:0:8:800:200C:417A/128
```

3.2.2 Access Control

The access control mechanism affects only the Dispatcher 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 an ip-address and a net-mask separated by a "/". Both IPV4 and IPV6 addresses are permissible. A net-mask should be given in the CIDR style in form of an integer specifying the number of bits of the network-part. For IPV4 addresses, the traditional dotted notation is also permitted. It can be used to specify a single host or a subnet in the following ways

10.205.112.22/255.255.255.255	designates the single host 10.205.112.22
10.205.224.22/32	designates the single host 10.205.224.22
10.205.112.0/255.255.255.0	designates all hosts in the subnet 10.205.112.*
10.205.224.0/24	designates all hosts in the subnet 10.205.224.*
10.0.0.0/255.0.0.0	designates all hosts in subnet 10.*.*.*
11.0.0.0/8	designates all hosts in subnet 11.*.*.*
2001:0db8:0010::/48	designates all hosts in subnet 2001:0db8:0010:*
::ffff:0a0a:0a0a/128	designates a single IPV4 hosts 10.10.10.10
*	this wildcard designates all hosts

Technically, the IP-address of a requester matches an IP-specifier when the network prefix denoted by the netmask matches.

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 <code>com.groiss</code> or below
com.groiss.org.PasswdAuth	designates all calls to methods in the class <code>com.groiss.org.PasswdAuth</code>
*	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.

set_agent	designates all users who have the right <code>set_agent</code>
admin stat	designates all users who have the right <code>admin</code> and / or the right <code>stat</code>
*	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:

3.2. HTTP SERVER

127.0.0.0/8 * *

Access from local host subnet is not restricted.

10.205.112.26/32 * DENY

Access from 10.205.112.26 is not allowed.

10.205.112.0/24 com.groiss.org.PasswdAuth *

Login of hosts from subnet 10.205.112.0 is allowed.

10.205.112.0/24 * 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 - database:** The database; you can select ORACLE, DB2, MS SQL-Server, Firebird, or Derby.
- **JDBC Driver Class - database.driver.class:** Java-Class, that contains the driver. See the table on page 27 for a list of driver classes.
- **JDBC URL - database.url:** URL for the database. The syntax of this string depends on the JDBC driver used. See the examples on page 27 or consult the documentation of the driver.
- **Database Userid - database.user:** The ID of the user with whom you want to connect to the database.
- **Database password - database.password:** Password for the database user with the ID that you entered above.
- **Number of connections - database.connections:** Default number of database connections.
- **Maximum number of connections - database.connections.max:** The maximum number of database connections that can be created.
- **Session environment - database.session.env:** You can specify SQL-commands, which are executed for each connection after connecting, for example: `set TEXTSIZE 1000000`
- **Connection properties - database.connection.properties:** You can specify e.g. SSL properties to establish a secure connection to database. The value of this property is a list of property declarations separated by `\r\n`. Note that the = sign must be escaped by `\` when editing directly in *avw.conf*.

e.g. `database.connection.properties=my.prop\=a.value\r\nyour.prop\=another.value`

- **Reconnect try interval (sec.) - database.waitFor.seconds:** Interval in seconds for reconnect tries to the database.
- **Reconnect tries - database.waitFor.count:** Number of reconnect tries.

- **Query timeout (sec.) - `database.query.timeout`:** Number of seconds after which a query times out.
- **DB connection reservation warning interval (secs) - `database.connection.busy.warning.secs`:** Long-lasting reservations of DB connections can be logged and also monitored via the Server Monitor (Aged DB connections). Information includes thread-name, timestamp and stacktrace at moment of reservation. Monitoring information in the logfile will occur in 2 minute intervals. Each long-lasting reservation is logged not more than once. Following values can be defined:
 - -1 : do not monitor (default, behavior like before)
 - ≥ 0 : do monitor; log /report all connections being reserved longer than the specified time interval (seconds)
- **Element count in IN-Expressions - `ep.inlists.splitsize`:** `@enterprise` uses SQL "IN-lists" - SQL expressions of the form WHERE att IN (val1, val2, ..., valn) as one form of optimizing database access.

The API provides the `com.groiss.store.BulkQuery` class as a convenient means to utilize this fast kind of access.

Since database systems usually put restrictions on the textual length of SQL-statements and also on the number of elements in such IN-lists, `@enterprise` splits queries with long IN-expressions into several queries. This configuration parameter can be used to control the maximum number of elements of an IN-expression.

The default value is 250 elements. Increasing the parameter leads to fewer partial queries and fewer roundtrips to the database, but also to longer statements and IN-lists with the possibility to hit the limits imposed by your DBMS.

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*).

3.4 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 - `avw.base.dir`:** This is the root directory for all relative paths, if you leave it empty the current directory of the start script is used.
- **Directory of form classes - `avw.formclassdir`:** Directory, where the system writes the form classes.
- **Directory for temporary files - `Httpd.tempDir`:** Directory for temporary files.

3.4. DIRECTORIES

DBMS	Driver Vendor	Driver Kind	Class and URL
DB2 UDB	IBM	Data Server Driver	com.ibm.db2.jcc.DB2Driver jdbc:db2://host:50000/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 (V2005+)	Inetsoftware	Una2000	com.inet.tds.TdsDriver jdbc:inetdae:'host':1433?sql7=true
MS-SQLServer (V2005+)	jTDS Project	jTDS	net.sourceforge.jtds.jdbc.Driver jdbc:jtds:sqlserver://host:1433/'dbname'
MS-SQLServer (V2005+)	Microsoft	V3.0+	com.microsoft.sqlserver.jdbc.SQLServerDriver jdbc:sqlserver://host:1433;database='dbname'
MySQL (V5.0, experimental)	MySQL	Connector/J (3.1)	com.mysql.jdbc.Driver jdbc:mysql://host:'port'/dbname'
Oracle LOBs	Oracle	Thin (V10g+)	oracle.jdbc.OracleDriver jdbc:oracle:thin:@host:1521:'SID'
Oracle LOBs	Oracle	OCI	oracle.jdbc.OracleDriver jdbc:oracle:oci:@'TNSNAME'
PostgreSQL (V8.1+)	PostgreSQL	Native	org.postgresql.Driver jdbc:postgresql://host:'port'/database'
MS-SQLServer (-V2000)	Inetsoftware	Una2000	com.inet.tds.TdsDriver jdbc:inetdae:'host':1433?sql7=true
MS-SQLServer (-V2000)	jTDS Project	jTDS	net.sourceforge.jtds.jdbc.Driver jdbc:jtds:sqlserver://host:1433/'dbname'
MS-SQLServer (-V2000)	Microsoft	V3.0+	com.microsoft.sqlserver.jdbc.SQLServerDriver jdbc:sqlserver://host:1433;database='dbname'
Oracle LONGs (deprecated)	Oracle	Thin	oracle.jdbc.OracleDriver jdbc:oracle:thin:@host:1521:'SID'
Oracle LONGs (deprecated)	Oracle	OCI	oracle.jdbc.OracleDriver jdbc:oracle:oci:@'TNSNAME'

Table 3.1: JDBC-Drivers

3.5 Logging

- **Log file - `logger.logfile`:** Name (path) of file, where **@enterprise** writes log information. If file not exists, a new one will be created. By activating the button *Show* the logfile is opened in a new window.
- **Restart log - `logger.restart.logfile`:** Here you can define how often the log file should be initialized - daily (at midnight) or at startup only.
- **Number of logs - `logger.keep.logfile`:** The number of stored log files.
- **Error file - `logger.errorfile`:** This file is a centralized collection of errors. Errors 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.
- **Restart error log - `logger.restart.errorfile`:** see *Restart Log*
- **Number of error logs - `logger.keep.errorfile`:** see *Number of Logs*
- **Logger Class - `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 - `logger.trace`:** 0, 1, 2 or 3:
 - 0 Errors are logged.
 - 1 HTTP requests are logged (time stamp, 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 - `logger.logOnConsole`:** 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 - `HttpdAuth.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 - `settings.class`:** A class defining some global settings can be defined here. For details see the **@enterprise** Programming Guide.

- **Notification Provider Class - `avw.notification_provider_class`:** The class for the notification mechanism, must implement the interface `com.dec.avw.notification.NotificationKit`. This mechanism allows to notify RMI-based Java clients in an asynchronous manner about changes in worklists.
- **Archiving Class - `avw.archiveclass`:** The class used for archiving process instances, must implement the interface `com.dec.avw.core.AVWArchiver2`.
- **Error-Formatter Class - `avw.error.formatter`:** 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 - `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 - `Locale.list`:** 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 - `Locale.language`:** Defines the language for the user interface. Language is defined in ISO language code, for example `de` for German.
- **Country - `Locale.country`:** ISO country code, for example `AT` for Austria.
- **Variant - `Locale.variant`:** A default variant to use. You can define free variants in the list of locales (e.g., regions, companies, etc.).
- **Character set - `avw.charset`:** The character set for text files uploaded to @enterprise. The default is character set of the server. You will have to change it, if the server charset differs from the charset used on the clients. For example, if the server uses UTF-8 and clients use windows-1252 the `charset=UTF8` header results in incorrect displaying the special characters in the document. 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 does not work).

The used character set of the database is independent of the entered character set on this mask. What characters can be stored in the database is dependent of the used character set in database and the JDBC driver which is responsible for interpreting in Java strings.
- **Decimal separator - `avw.decimal.separator`:** Set the separator for floating-point numbers (default is `.`)
- **Date format - `DateFormat`:** Format mask for date input and output. See the table below for a description of the possible values.
- **Date-Time Format - `DateTimeFormat`:** Format mask for date and time.

- **Default unit for displaying time intervals - `calutil.defaultunit`:** Default-Unit in seconds, minutes, hours, days and weeks.
- **Applet look and feel - `applets.lookandfeel`:** Specify the look-and-feel of the process editor, values are: metal or windows.
- **Max. table length - `table.maxsize`:** Specify a natural number. For tables of size greater than this number the user is asked before the table is shown.
- **Items per page - `table.pagesize`:** This defines the maximum number of entries in tables (e.g. 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*.
- **Max. paging table length - `table.paging.maxsize`:** For paged tables of size greater than this number the user is asked before the table is shown or the search function in toolbar must be used.
- **Keep object changes (days) - `keep.log.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) - `avw.keep.user.sessions`:** Number of days after which inactive user sessions will be deleted.
- **Use browser language - `locale.from.browser`:** If this option is set, the system uses the language settings of the browser instead of the settings in the user table of `@enterprise`.
- **Select List look and feel - `selectlist.lookandfeel`:** 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 - `selectlist.search`:** The search option for searching in a select list:
 - Starts with: at the begin of a string
 - Substring: within a string
- **Enable Wiki link syntax - `ep.wikilinks.enable`:** If this checkbox is activated, links can be entered in the description of an ActivityInstance in wiki syntax: `[[link | text]]`. Please note that links must be relative.
- **Use shiny GUI - `ep.style.shiny`:** If this checkbox is activated, a modern GUI with new icons and some state-of-the-art features like gradients, box-shadows, etc. is used. Known problems (in browsers) are noted in `@enterprise` FAQ which can be found on <http://www.groiss.com>.
- **Form Toolbar - `ep.form.toolbarshape`:** This parameter allows to specify the shape of the toolbar for forms: Possible values are TEXT, ICON, and BOTH.

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

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.

(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 - mail.smtphost:** Server for outgoing E-Mails (host name or IP address).
- **Mail sender - mail.sender:** The mail address that will appear in the *from* field of mails that the system sends.
- **Administrator email address - avw.adminemail:** One ore more email addresses of the system administrator separated by comma.
- **RMI port - avw.rmi.port:** Port number of RMI (Remote Method Invocation) listener. Needed for Java-Clients.
- **Enable RMI class loading - avw.dyn_class_load.enabled:** Enables class loading via RMI, this is needed when working with forms and the Java-Client.

- **Enable full RMI access - `avw.rmi.enablefull`:** 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 - `avw.plain_rmi`:** Allows plain (unencrypted) communication for RMI connections.
- **Export port for plain RMI communication - `avw.rmi.plain_exportport`:** If specified, this is the port used for RMI traffic.
- **Use SSL for login sequence at RMI communication - `avw.secure_login_rmi`:** Use SSL to encrypt the login sequence for RMI communication.
- **Crypt RMI communication with SSL - `avw.permanent_secure_rmi`:** Encrypt the whole communication over RMI.
- **Export port for RMI over SSL - `avw.rmi.ssl_exportport`:** If specified, the encrypted RMI communication uses this port.
- **Client certificates for RMI over SSL - `ssl_requireclientcertificate_over_rmi`:** This parameter determines how a secure SSL connection can be established by a RMI 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.
- **Require SSL for Admin Shell - `ep.adminshell.sslonly`:** If enabled, unencrypted communication is denied.
- **Allowed hosts or networks for Admin Shell - `ep.adminshell.allowedips`:** Specifies a network restriction pattern. See the parameter *Allowed Hosts or Networks* in section [3.2](#).
- **Enable Wf-XML - `avw.wfxml.enabled`:** 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 OU - `wfxml2.orgunit`:** Default Wf-XML Organizational Unit.
- **WfXML User - `wfxml2.user`:** Default Wf-XML User.
- **WfXML Server - `wfxml2.server`:** Defines the default Wf-XML server.

- **WFXML access log for - wfxml2.log.objects:** Defines the objects, which will be logged. You can select between
 - ServiceRegistry
 - Factory
 - Instance
 - Activity
 - Observer.
- **Size of log - wfxml2.log.size:** Max. size of the logfile.
- **SMTP default properties - ep.mail.smtp.defaultprops:** Define default properties for SMTP mail communication (see <http://javamail.java.net/nonav/docs/api/>). In particular the following properties are useful in dealing with network problems: *mail.smtp.connectiontimeout* and *mail.smtp.timeout*.
- **IMAP default properties - ep.mail.imap.defaultprops:** Define default properties for IMAP mail communication (see <http://javamail.java.net/nonav/docs/api/>). In particular the following properties are useful in dealing with network problems: *mail.imap.connectiontimeout* and *mail.imap.timeout*.
- **POP3 default properties - ep.mail.pop3.defaultprops:** Define default properties for POP3 mail communication (see <http://javamail.java.net/nonav/docs/api/>). In particular the following properties are useful in dealing with network problems: *mail.pop3.connectiontimeout* and *mail.pop3.timeout*.

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 5.3.4 in the chapter about clusters for details about configuring clusters.

3.10 Workflow

- **Open form on process start - avw.start.with.form:** 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.
- **Inherit Ids to subprocesses - avw.inherit.ids:** Don't create Ids for subprocesses - use the parent processes' Ids instead.

- **Enable application-spanning process definition - `avw.procdef.appl_spanning`:** If this option is set, it is possible to define processes with application-spanning elements (i.e. Forms, Tasks, Subprocesses and Roles as Agents).
- **Allow automatic take - `avw.autotake`:** 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). If the process-form of such a task is edited, the current editor is written in table `avw_currenteditor` and is visible in the process-instance history.
- **Organizational hierarchy mandatory - `avw.orgtree.necessary`:** 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.

3.11 DMS

- **Show extensions - `avw.dms.showextensions`:** Show the document name extension, e.g., `.doc` or `.txt`.
- **Versioning - `avw.dms.versioning_strategy`:** *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 - `avw.dms.bequest_acl`:** When this option is checked, the permission lists of a folder is inherited to the contents of the folder.
- **Basic-Auth in WebDAV - `avw.dms.allow_basicauth`:** 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 - `avw.dms.newwindow`:** If checked, documents will be opened in new windows.
- **Maximum document size (in bytes) - `avw.dms.max_doc_size`:** 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 - `IStore.class`:** You can specify your own DMS storage class here. The class must implement the interface `com.groiss.dms.IStore`.
- **DMS Archiving Class - `DMSArchiver.class`:** Class for archiving documents, must implement the interface `com.groiss.dms.DMSArchiver`.
- **Standard table model / Table handler - `avw.dms.standard_tablemodel`:** 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 - `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 - `avw.security.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 - `avw.dms.textsearch.state`:** 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 - `avw.dms.hide_hidden_docs`:** 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 - `avw.dms.textfile_charset`:** 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).

Hint: Internet Explorer 7 opens MS Office files in read-only mode by default. By editing the registry entry `OpenDocumentsReadWriteWhileBrowsing` this behaviour can be adapted (open file in r/w mode). For further information please take a look on <http://support.microsoft.com/kb/870853>

3.12 Search

- **Maximum table size on server (rows) - `query.maxtable`:** 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) - `monitoring.cacheinterval`:** Specifies, how long a query result resides in cache.
- **Maximum number of cached queries - `monitoring.cachesize`:** Number of queries in cache.
- **Maximum number of simultaneous queries - `monitoring.maxparallel`:** Number of threads, that concurrently compute query results.

- **Maximum number of startable queries - monitoring.maxqueue:** Length of queue of queries waiting for execution (waiting for a free thread).
- **Default process Id search type - avw.reporting.defaultIdSearch:** Here you can define the standard type for id search in *Process Search* - see user manual for further information.
- **Default subject search type - avw.reporting.defaultSubjectSearch:** The same as *Default process Id search type*, but for subject.
- **Process relations - avw.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 (see syntax beneath). A comma separates the pairs. The id is stored in the database relation, the name is used in the user interface.

Definition syntax:

```
releid SPACE relname { "," releid SPACE relname } *
```

- **Search case-insensitive by default - avw.reporting.defaultIgnoreCaseSearch:** If this checkbox is activated, the checkbox *Ignore Case* on process search mask is activated by default.
- **Exact Id short search only - avw.reporting.exactIdShortSearch:** If this checkbox is activated, you have to enter the right Id to get a correct result.
- **Short search includes subject - avw.reporting.shortSearchSubject:** If this checkbox is activated, the subject will be included in shortsearch.
- **Short search includes field values - avw.reporting.shortSearchFieldvals:** If this checkbox is activated, fieldvals will be included in short search. It is necessary that full-text search is activated and the checkbox *Useable in DMS* must be activated for each formtype which should be found.
- **Order process Ids by OID - monitoring.orderProcessId:** 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 - avw.reporting.showNoAcces:** 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.
- **Use underscore (_) as SQL-wildcard - avw.reporting.underscoreIsWildcard:** If this checkbox is activated, underscores are allowed as SQL wildcard. If activated, it is not possible to search for '_' unless you escape it yourself.

- **Use smart search algorithm for multi-field searches - list.smartsearch:** If activated, it will be searched globally in all specified fields of a table by using OR-joins. This parameter takes effect on
 - short search in select-list and select-table
 - DOJO object select
 - short search in object-table (e.g. @enterprise administration master data tables)
 - short search in form-table
 - short search in select-list of function *Change Agent*

Example: The search fields are firstname and surname of the user table. In short search field of the user table the string *Roland Eisenberg* has been entered. The sql-condition would be following:

```
(lower(firstname) like(Roland%) or lower(surname) like(Roland%))  
and (lower(firstname) like(Eisenberg%) or lower(surname) like(Eisenberg%))
```

It is also possible to activate/deactivate this behaviour for each table by setting following attribute in configuration file (e.g. myappl.xml):

```
<Attrib key="smartSearch" value="true|false"/>
```

- **Show time in date conditions - avw.reporting.showTimeInDateConditions:** If activated, date and time for datefields on process-/document-searchmask and Reporting condition mask can be entered and on all these masks the appropriate checkbox is activated by default. If this checkbox is deactivated, only the date can be entered as value (without time).

3.13 Tuning

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

- **Ignore reference roles - avw.tune.ignorerefroles:** If you don't use reference roles (see the Administration Manual for details) you can set this option.
- **Ignore hierarchic roles - avw.tune.ignorehierroles:** If you don't need hierarchic roles you can set this option.
- **Ignore personal substitutions - avw.tune.ignorereprs:** If this is set, personal substitutions are ignored.
- **Ignore role substitutions - avw.tune.ignorerolereprs:** If this is set, role substitutions are ignored.

- **Worklist Cache at server startup - `avw.wlcache.state`:** 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 - `avw.wlcache.exemptseenobjects`:** If this checkbox is activated, seen objects will not be cached anymore.
- **Defer loading of finished parents - `avw.wlcache.parents.defer.missing`:** If checked, loading of missing parents (for finished parfors, scopes, processes) can be deferred.
- **Reload classes - `avw.class.reload`:** Reloads classes without server restart if possible. This should be used only in development environments.
- **Statement statistics - `avw.stmt.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) - `file.cache.size`:** 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.
- **Archive schema - `avw.archive.schema`:** See section 6 for details about archive schemas and how to set them up.
- **Clear ThreadContext after run of TimerEntries and Batchjobs - `ep.timer.batch.clearthread`:** Timer's and Batchjob's ThreadContext are cleaned, if this checkbox is activated.

3.13.1 ACLCache

In **@enterprise** it is possible to speed up the rights check by activating the ACLCache. 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.

There are five properties to configure the ACLCache:

- **Permission Cache activated - `aclcache.active`:** Check, if the ACLCache should be activated
- **Max. number of object specific rights - `aclcache.objectrights.maxelems`:** Size of the object specific rights cache (in objects).
- **Lifetime of object specific rights (sec.) - `aclcache.objectrights.lifespan.secs`:** Lifetime of rights in the object specific rights cache.
- **Max. number of class rights - `aclcache.classrights.maxelems`:** Size of the class rights cache (in objects).
- **Lifetime of class rights (sec.) - `aclcache.classrights.lifespan.secs`:** Lifetime of rights in the class rights cache.
- **Use partition optimized query for permission checks - `acl.separate.targetquery`:** ACL evaluations can be tuned by using separate queries for `objectscope = 3` versus `objectscope <> 3`. For this purpose activate this parameter.
- **ACL list: Permission Cache integration - `acl.list.cache.usage`:** This parameter allows to define how ACL list interacts with ACLCache. Following options are available:
 - None: List does not interact with cache
 - Check only: Cache is consulted, no results are inserted into cache
 - Insert positive results only: Cache is consulted, only positive results are inserted
 - Full: Cache is consulted, all results (positive and negative) are inserted into cache)
- **ACL list: Max. number of OIDs in IN-Clause - `acl.list.target.splitsize`:** The split size for the target set of an ACL.list-query. If the size of the target set is not greater than the split size, `@enterprise` can filter by using a SQL IN-Clause with the target oid's, otherwise a more general filter will be used which may result in a larger result set for that query. In both cases a single SQL statement will be executed. Please note that there are database specific restrictions concerning the number of literals within an IN-Clause and also the textual length of an SQL statement.
- **ACL list: Always restrict by OID for the following classes - `acl.list.target.splitclasses`:** A comma-separated list of fully qualified class names. For those classes, the target set should be splitted so that more than one SQL statement will be executed which always filter by target oid's using an IN-Clause. This is useful if a lot of object specific permissions exists for such a class so that the more general filter would cause a huge result set.

3.14 Security

- **KeyStore file - ssl.keystore:** 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 - ssl.keystore_pwd:** To access a KeyStore a password (with a minimum length of 6 characters) is needed.
- **Password for server certificate - prk.passwd:** 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.
- **Allow condition definitions on client (SQL Injection possible) - ep.allow.clientcondition:** If this checkbox is activated, SQL injection for SelectList, SelectTable and DOJO select is allowed. This parameter should be used for compatibility only. In future **@enterprise** version SQL Injection will not be possible anymore! More details about the correct usage can be found in the section *AJAX components integrated in @enterprise* in *Application Development Guide* which contains an example for DOJO selects (it is quiet the same procedure for SelectList and SelectTable).
- **Check view right at showing forms - ep.formwrapper.checkright:** If enabled, the right check for viewing forms is activated. What are the conditions to view a form (for more details see section *Rights* in *System Administration Guide*):
 - If the form is assigned to a task, the right *view_procinst*, *set_agent* or *proc_inst* to the parent of the task is needed.
 - If you are the agent of the current task or (if role is agent) the role of the task is part of the appropriate organizational unit and the task is not finished yet.
 - If the form is stored in a DMS folder, the *view* right is needed.

3.15 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.

3.15.1 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) - `passwdpolicy.days_password_valid`:** Defines the password's period of validity in days.
- **Inform user before password expires (in days) - `passwdpolicy.days_warning_before`:** 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 until account is deactivated - `passwdpolicy.max_count_invalid_logins`:** 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 - `passwdpolicy.algorithm`:** 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 the following different algorithms can be chosen.
 - **SHA-256 (Secure Hash Algorithm):** Takes a plain string of any length and produces a 256-bit hash output. Please note that JAVA 1.6.18 or higher is required for SHA-256!
 - **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. It is not recommended to use this algorithm anymore due to vulnerability!
 - **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.

3.15.2 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 - `passwdpolicy.min_length`:** 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 - `passwdpolicy.max_length`:** 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 - `passwdpolicy.min_letters`:** 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 - `passwdpolicy.min_capitals`:** 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 - `passwdpolicy.min_lowercase`:** 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 - `passwdpolicy.min_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 - `passwdpolicy.min_others`:** 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 - `passwdpolicy.min_different_chars`:** 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 - `passwdpolicy.max_char_adjacent`:** 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 - `passwdpolicy.max_sub_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 - `passwdpolicy.history_steps`:** 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.
- **Minimal number of whitespace characters - `passwdpolicy.min_whitespace`:** Specifies the number of min. allowed whitespace characters.

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

3.15.3 Your Own Checker Class

- **Checker Class - `passwdpolicy.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.

3.16 Calendar

- **Holiday Class - `avw.calendar.class`:** Here you can define a class for displaying the holidays in the calendar. It must implement the `com.groiss.cal.Holidays` interface.
- **Allow email address as attendee - `cal.date.attendees.mailuser`:** If this checkbox is checked, it is possible to add email-participants to an appointment.
- **iMIP - `calendar.imip`:** If this checkbox is activated, iMIP will be used. In @enterprise calendar notifications contain *iCalendar*-files. iMIP offers the possibility to process status information of an appointment.
- **iMIP email address - `calendar.imip.email`:** Email-address which is used for communicate with the participants; participants will reply to this email-address!
- **Show default resource - `cal.show.defaultres`:** If this checkbox is checked the user can use a simple resource form for assigning resources to calendar appointments.
- **Resource classes - `cal.resources`:** 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 Day - `cal.nonworkingdays`:** In this list it is possible to select one or more non working days, which will be needed for example in escalations.

3.17 Process cockpit

This section contains the parameters for the process cockpit (see details in the *User Manual*).

- **Root folder - ep.cockpit.rootfolder:** The path to the root folder of the process cockpit.
- **Reports for all processes - ep.cockpit.allreports:** A comma separated list of Report-Id's. These reports are displayed in the process cockpit for every process.
- **Show overdue processes of last n days - ep.cockpit.deadline.days:** Specifies the number of days, which are used for the calculation of process deadline violations per process definition.
- **Show last n instances - ep.cockpit.recent:** Specifies the number of instances, which are displayed in tab *Runtime* of table *Recently Started* in process cockpit.
- **Common processes - ep.cockpit.commonproc:** A comma separated list of formtypes (Formtype-Id + Version, e.g. jobform_1) which contain the formfield *area*. The forms are used to assign process instances of common processes (for example project) to a cockpit entry.

3.18 Time management

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

- **Default time unit - timemgmt.timeunit:** Time unit shown in histograms, duration statistics, etc. You can select between:
 - Seconds
 - Minutes
 - Hours
 - Days
- **Max. length of time histogram - timemgmt.histogram.length:** The maximum length of the time histogram can be set here (in objects).
- **Prune probability - timemgmt.pruning.prob:** 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 - timemgmt.deadline.prob:** You can select an appropriate deadline based on reliability requirements to new process. The default value is 95% (=0.95).
- **Create start histograms - timemgmt.create.starthistos:** If this parameter is activated, beside the other histograms additionally a start histogram will be created. For more details about the histogram please take a look in *System Administration Guide* - section *Time Management* (in chapter *The Process Editor*).

3.19 Change administrator password

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

3.20 Initialize database scheme

This function opens the setup of @enterprise. 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.

3.21 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.
- **ep.scripts.enable:** If this checkbox is activated, groovy scripts are allowed (e.g. in methods of tasks, etc.).
- **ep.adminshell.enable:** If set to false, access from the admin-shell is denied.
- **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-address, 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 behavior is that surname, first-name 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 behavior.

Example:

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

The table displays the surname, first name, 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*).
- **java.class.path:** Here you can set the classpath, which is used by @enterprise.
- **webdav.drive:** The webdav drive can be specified with this parameter, which represents the root (the same letter like set in WebDrive properties). If the value *off* is entered, WebDrive will not be used anymore. You have to reconnect to the @enterprise Server after changing this parameter.
- **webdav.show.subject:** Set this value to *1* for activating the subject in file-name of process. This parameter is effectless, when using spaces in the process-id (PID). Useful information about instance ids can be found in the System Administration Manual - chapter *Processes*.
- **webdav.worklist.name:** Here you can set a name (ASCII-string) of a folder for displaying the worklist entries. Default string is *wl*.

- **webdav.delete.allow:** This parameter must contains a java pattern. If no parameter is set, every document/folder can be deleted. If the pattern string has a wrong syntax, nothing can be deleted. If the parameter has a correct syntax, only the documents/folders depending on the pattern can be deleted.

Example:

```
webdav.delete.allow=.*\\.(tmp|temp)\\z
```

Only files with suffix *tmp* or *temp* can be deleted via webdav in this example.

- **dms.hide.common** and **dms.hide.userfolder:** With these parameters you can hide the *common-* and/or *userfolder*. The parameter **dms.hide.userfolder=1** fades out the userfolder and the folder of the substituted person. Default is *0* - folders are displayed.
- **ep.html.noescape:** If this parameter is set to *1*, the old behavior of HTML encoding will be used like in **@enterprise** versions less than or equal 7.0 (compatibility mode).
- **ep.choice.showsingle:** Set this parameter to *0* for displaying no choice-mask anymore when one branch of a choice-object is active only. If set to *1*, choice-mask always is displayed. Default is *0*.
- **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
- **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
- **httpd.jetty.plainconnector.usenio:** Set this parameter to *1* to use NIO connector for jetty http (not https)
- **httpd.jetty.sslconnector.usenio:** Set this parameter to *1* to use NIO connector for jetty https (not http)
- **cal.applications:** A list of classes can be defined here to activate/deactivate additional calendar-components, e.g. if *com.groiss.calendar.CalendarAppl* is removed, no appointments can be added anymore (default: *com.groiss.calendar.CalendarAppl*, *com.groiss.calendar.wf.DueTasks*, *com.groiss.calendar.wf.FinishedTasks*).
- **ep.forms.generate.with.javac:** Set this parameter to *0* to generate form classes without javac (*asm* byte code library is used). You can switch to old behavior like in **@enterprise** versions before 8.0 by setting this parameter to *1*.

<i>avw.redirectpath</i>	<i>redirection</i>
""	no redirection
/wf	redirect to /wf/a/x.txt?p=1
old/docs	redirect to /wf/old/docs/a/x.txt?p=1
/old/docs	redirect to /wf/old/docs/a/x.txt?p=1
/wf/old/docs	redirect to /wf/old/docs/a/x.txt?p=1

Table 3.3: **Examples for redirection paths**

- **avw.redirectpath:** Allow to redirect requests without context path (in embedded Jetty). When set to a value, all requests with a path not starting with the context path will be redirected to redirectpath. If *avw.redirectpath* does not start with the contextpath, the contextpath is prepended. Examples can be found in table 3.3.
- **avw.readonly.text:** If set to true, read-only form fields are displayed as simple text instead of read-only input fields (only for HTML forms).
- **Httpd.documentRoot:** Directory where the documents for the HTTP server reside.
- **ep.ignore.mainform.tid:** If set to true, the tid of the mainform will not be increased, if in sub(sub)forms an update of the mainform is done.
- **ep.hidden.oid:** A hidden field named oid is added to the HTML page of forms (forms of type HTML or XHTML), if this parameter is set to true.
- **ep.strings.allownonxml:** There are ASCII and UNICODE characters which are not allowed in a XML (defined in XML standard 1.0). If set to true, Non-XML characters are allowed in Strings in **@enterprise**. The default setting is *false*.
- **avw.reactivate.nopreproc:** When reactivating a process, the preprocessing method of the reactivated activities is executed, if parameter is set to false. If this behaviour is not wished, set the parameter to true.
- **ep.restricted.mode:** If set to true, @enterprise is running in restricted mode which means:
 - Setting of password for *sysadm* is not allowed.
 - Some configuration parameter are readonly (e.g. License Key).
 - Applications cannot be installed with function *Install Application*.
 - The XML import generates classes.
 - For preprocessing, postcondition, expression and compensation *SystemAction* and classes of the application classpath can be called - not allowed are e.g. *StoreEJB.execute()*, Groovy, etc.
 - It is not allowed to define an application classpath outside of the @enterprise installation.
- **ep.allow.formupload:** If this parameter is set to true, the form template on filesystem can be adapted on page 2 of form wizard (see System Administration Guide).

3.21. PARAMETERS WITHOUT GUI

- **ep.store.formhandler.behaviour:** If value is set to *view*, `onInsert()`, `onUpdate()` and `onDelete()` of view form event handler will be called. If a handler of the view form is found, this one is used and the form handler of the base form is ignored. If no handler is entered for view form, the handler of base form is used.
- **ep.history.forms.readonly:** If true, forms are shown in readonly mode when process detail is opened from search (default is false).

4 Patching and Upgrading your Installation

This chapter describes the patching and upgrading mechanism of **@enterprise**. For this purpose an explanation of the used terms is necessary first:

- **Version:** A version is the number of the **@enterprise** version, e.g. 8.0.
- **Build:** Represents a combination of a version and the revision number. The revision number can be found e.g. in **@enterprise** changelog. Example for build number: 8.0.6778
- **Upgrade:** This term is used, if the version of **@enterprise** has been increased, e.g. upgrade from **@enterprise** 7.0 to 8.0
- **Patch:** This term is used, if the build number has been increased, e.g. updating from **@enterprise** 8.0.6770 to 8.0.6778.

4.1 Patching the Installation

To assure the quality and reliability of your installation, bug fixes and enhancements for **@enterprise** are provided in the form of patches. The main starting point for obtaining such patch files is our download area reachable via **<http://www.groiss.com/download.html>**. There are two alternative ways to incorporate the patches in your installation. We will first describe the manual procedure and then go into the details of the automatic patch facility.

4.1.1 Manual Patch Method

An **@enterprise** patch consists of one or more files. Those patched files should replace the corresponding files in the appropriate directories of the installation. Further, the patch files may include a procedure to accomplish maintenance steps on the database. The following steps should be performed:

1. Stop the server.
2. To play it really safe, backup the database and the installation files.

3. Copy the patch files into the corresponding target directories.
4. Start the server with the *-upgrade* option. Any required database steps will be performed.
5. Start the server again without the *-upgrade* option.

4.1.2 Automatic Patch Method

To apply patches to an **@enterprise** installation or **@enterprise**application, a automatic mechanism can be used.

The mechanism is based on patch archives which can be obtained by visiting our download area at <http://www.groiss.com/download.html>. There you can find a cumulated patch, allowing you to apply the latest bugfixes.

A patch archive is a bundle which incorporates the individual new versions of the files to be patched, along with two additional files (*version,changes*) which describe the patch and the needed actions. The technical format of a patch archive is a ZIP-archive.

Because not every patch is applicable for each target system, the build number of the current installation will be compared with the patch build number. If the system is not at a minimum required build, or the system is newer than the patch, the patch will not be applied. If the currently installed build cannot be determined, the patch will not be applied.

The automatic patch deployment mechanism is conservatively designed and therefore prevents the loss of files which are in your current installation. Before a patch archive is applied, a backup of all affected files will be performed. Each time you apply a patch, a separate backup-folder will be created. The content of the backup folder will remain on your hard-disk, even if the patch has been successfully applied.

All actions will be appended to a file (*patch.log*) where the full patch-history can be seen. In case of an error, a rollback will be performed, leaving all files unchanged. Both, the backup folder and the log-file, will be created in the *patches* folder.

The following steps initiate an automated patch procedure:

1. Download the patch archive.
2. To play it really safe, backup the database.
3. Copy the patch archive file to the *./patches* folder of the installation.
4. Optionally check the effects of the patch archive prior to applying the patch. The system administration provides a function in the System-Control section. This is especially handy to identify files that have been overwritten in the local installation. If there are any clashes, you will have to make sure that your local changes are not lost by manually reapplying them after the patch action or by updating the local files according to the changes in the original base files.
5. The patch procedure can then be initiated by:
 - If you are using **@enterprise** in standalone mode (Jetty), you have to start the server with the *-upgrade* option.

- If your installation runs in an application-server (e.g. Apache Tomcat) use the provided patch-script (patch.bat/path.sh) to apply the patch. **Please note:** your application-server, or at least the @enterprise application must not be running while applying the patch.
6. After the patch has been successfully applied, the patch archive will be moved to the corresponding backup folder.

4.2 Upgrading/Patching an @enterprise Application

Applications in @enterprise can be kept up to date using the same mechanisms as for the base system itself.

The upgrade/patch of an application contains of a set of files which must be replaced in an installation. @enterprise also offers the possibility to execute further actions via an upgrade method. Typical actions include:

- XML import: Master data of the application can be adapted.
- Execution of database scripts.
- Other java methods.

The *upgrade* method is part of the application class which has to implement the interface *com.groiss.wf.ApplicationAdapter*. The application will be upgraded/patched automatically, if you start your server with disabled logins or the *-upgrade* option. Further information can be found in the API of @enterprise (*ApplicationAdapter.getVersion()* and *ApplicationAdapter.upgrade()*).

The steps for performing a manual application upgrade/patch are the same as manually installing a patch for the base system.

Patches for applications can also make use of the automatic patch mechanism. Place your patch-file in the *patches* folder in your application directory. The filename must match *patch*.zip*. When starting the server with disabled login or the *-upgrade* option, these patches will be applied too. Please note that it's not recommended to patch more than one application at a time. For further information on how to build your own patch-archives see the programmers manual.

4.3 Performing an Upgrade of @enterprise

This section describes the steps needed to upgrade from a prior @enterprise version to the current one (e.g. 7.0 to 8.0):

1. Backup your old installation and database.
2. Extract the content of *setup80.jar* into a new directory. We recommend to use the initial setup wizard for this purpose.

3. Copy your existing configuration file (avw.conf), forms directory, required jar-files (e.g. JDBC driver) of the *lib*-directory (e.g. ojdbc4.jar) to the corresponding directories of the new version.
4. To perform an interactive upgrade, start the server and login as **sysadm**. You should now be redirected to the upgrade page where you can initiate the necessary upgrade procedure for the database.
5. To perform an unattended upgrade, start the server with the *-upgrade* option. Any required database upgrades will be performed. The server will be stopped automatically. **Please note:** If you are using PostgreSQL as your DBMS, you must use this upgrade method.
6. (Re-)Start the server.

4.4 Migration of deprecated DBMS data types

4.4.1 Migration of Oracle data type LONG

This section describes the migration steps for existing Oracle 9i and newer installations, because since Oracle 9i LONG is deprecated. Existing installations can continue to use the old data types (not recommended), but new installations should use the new data types. Following steps must be performed for migration:

Hint: Please note that the migration is at your own risk and can take a long time!

1. Backup your old database!
2. Write down the indices of the affected tables. Following query helps to get a list of the affected indices wrapped in executable *alter index* statements:

```
select 'alter index '||index_name||' rebuild;' as command
from user_indexes where table_name in (
  select table_name
  from user_tab_columns
  where (table_name like 'AVW%' or table_name like 'FORM%')
  and data_type in ('LONG', 'LONG RAW')
)
order by index_name;
```

Do not execute the *alter index* statements in this step!

3. For each table which contains one of the previous mentioned data types an *alter table* statement must be performed like in following way:
 - alter table <tn> modify (<longcolname> clob default empty_clob());
 - alter table <tn> modify (<longrawcolname> blob default empty_blob());

4.4. MIGRATION OF DEPRECATED DBMS DATA TYPES

It is a little bit cumbersome to get the affected tables. For this purpose, the following query helps to get a list of affected tables wrapped in executable *alter table* statements as mentioned above:

```
select
  'alter table '||table_name||
  ' modify ('||column_name||' '||
  (case when data_type='LONG' then 'CLOB' when data_type = 'LONG RAW'
    then 'BLOB' else 'ERROR' end)|| ' default empty_'||
  (case when data_type='LONG' then 'CLOB' when data_type = 'LONG RAW'
    then 'BLOB' else 'ERROR' end)||'());'
  as command
  from user_tab_columns
  where (table_name like 'AVW%' or table_name like 'FORM%')
  and data_type in ('LONG', 'LONG RAW')
  order by table_name, column_name;
```

4. Perform the *alter index* statements which have been created in step 2) to rebuild the indices.

4.4.2 Migration of deprecated MS SQL-Server data types

Since MS SQL-Server 2005 Microsoft has been set some data types to deprecated and replaced them by new ones. The following list contains the deprecated and the appropriate new data types:

- text has been replaced by varchar(max)
- ntext has been replaced by nvarchar(max)
- image has been replaced by varbinary(max)

Existing installations can continue to use the old data types, but new installations should use the new data types. This section describes the recommended migration steps for existing installations with MS SQL-Server 2005 and newer:

Hint: Please note that the migration is at your own risk and can take a long time!

1. Backup your old database!
2. For each table which contains one of the previous mentioned data types an *alter table* statement must be performed like in following way:
 - alter table <tn> alter column <textcolname> varchar(max);
 - alter table <tn> alter column <ntextcolname> nvarchar(max);
 - alter table <tn> alter column <imagecolname> varbinary(max);

It is a little bit cumbersome to get the affected tables. For this purpose, the following query helps to get a list of affected tables wrapped in executable *alter table* statements as mentioned above:

4.4. MIGRATION OF DEPRECATED DBMS DATA TYPES

```
select
'alter table '+table_name+
' alter column '+column_name+' '+
(case when data_type='text' then 'varchar(max)'
      when data_type='ntext' then 'nvarchar(max)'
      when data_type = 'image' then 'varbinary(max)' else 'ERROR' end)+
','
as command
from information_schema.columns
where (table_name like 'AVW%' or table_name like 'FORM%')
and data_type in ('text','ntext', 'image')
order by table_name, column_name;
```

5 Clustered @enterprise System

5.1 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.

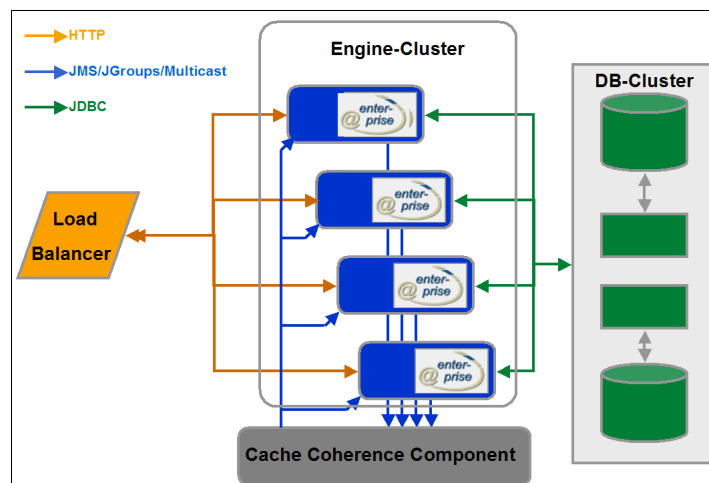


Figure 5.1: Cluster Architecture

Figure 5.1 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.

5.2 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.

5.3 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.

5.3.1 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 1.5 or 1.6, 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.

5.3.2 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.

5.3.3 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 needed under *Configuration* → *Coherence* and must be identical on all cluster nodes. These parameters are available only, if *Standard Multicast* is used as *Transportlayer for Coherence*:

- **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 multicast 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 needed under *Configuration* → *Coherence* and must be identical on all cluster nodes. These parameters are available only, if *JGroups* is used as *Transportlayer for Coherence*:

- **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 `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 non-persistent, auto-acknowledged, non-transacted messages and nondurable subscribers. JMS does not run within an @enterprise JVM, it must be configured and started separately.

Apache ActiveMQ (<http://activemq.apache.org>) meets all requirements and is known to be reliable, but virtually any JMS implementation should be suitable. Hints for configuring a particular JMS may be obtained via the @enterprise support.

The following configuration parameters are needed under *Configuration* → *Coherence* and must be identical on all cluster nodes. These parameters are available only, if *JMS* is used as *Transportlayer for Coherence*:

- **JMS Provider URL:** The URL name of the JMS provider. For ActiveMQ this is something like `tcp://localhost:61616?wireFormat.maxInactivityDuration=0`.
- **JMS ContextFactory:** Name of the Java class for construction of the JNDI-Context. For ActiveMQ this is `org.apache.activemq.jndi.ActiveMQInitialContextFactory`.
- **JMS TopicConnectionFactory:** Java class name for the topic factory of the JMS provider. For ActiveMQ this is `ConnectionFactory`.
- **JMS Topic:** The name of the topic used for communication. Such topics must typically be created within an JMS provider by the administrator. For ActiveMQ this can also be a dynamic topic like `dynamicTopics/avw`.
- **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.

5.3.4 Adapting the @enterprise Configuration

Configuration: Under *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 *Configuration* / *Cluster* or *avw.conf*:

- **Clustering enabled - avw.cluster.activated:** Must be checked.
- **Server name - avw.servername:** Name of the server. Must be the same on each node of one cluster.
- **Node Id - avw.node.id:** Id of the cluster node. Must be unique within the cluster.

- **Performance factor - `avw.node.perffactor`:** 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.
- **Member of load balancing - `avw.node.loadbalancing.member`:** If set to *YES* (by default), the loadbalancing function for this node is active, i.e. the node is a potential target for clients which request loadbalanced sessions. On nodes which serve special purposes and should not receive logins from "ordinary" clients, this parameter should be unchecked.
- **Clustercheck tolerance (sec.) - `avw.node.clustercheck.tolerance.secs`:** 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.) - `avw.node.heartbeat.tolerance.secs`:** 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 - `avw.dbcache.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.
- **Transport layer for Coherence - `avw.dbcache.coherence.transport`:** 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 *Run on each Node* 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.

5.4 Operation of a clustered system

5.4.1 Monitoring

A cluster health monitor which displays the state for each of the nodes can be accessed via Admin-Tasks / Server / Running Nodes Monitor.

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.
- **ClusterTimers:** Marks, if the node is the one which runs the cluster timers.
- **Load:** Current number of connected users.
- **Performance Factor:** The performance factor of the node.
- **Load Coefficient:** The current load coefficient (number of users divided by performance factor).
- **Load Balanced:** Marks, if the node is member of loadbalancing (see section [5.4.2](#)).
- **Logins enabled:** Marks, if new logins are allowed on the current node. Logins can be enabled/disabled with the toolbar-function *Disable/Enable Login*.

- **Current Session:** Shows, if current sessions should be kept, renewed or aborted. At startup this is always set to "keep".
- **Successor Nodes:** The id of the successor node is displayed where the clients of the "switched off" node should be logged in, if login is restricted (see column *Logins enabled*) and current session should not be kept. At server startup this column is always empty.

Hint: *Current Session* and *Successor Nodes* are used by @enterprise Java Clients only!

5.4.2 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(S)-Client) or can obtain appropriate initial URLs of this node (RMI-Client).

HTTP(S)-Clients The URL for getting a load balanced session for an HTTP(S) Client is:
`http[s]://<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`

Three URLs are returned, each in a separate line. The URLs can be used by the client to obtain an appropriate session to the node. The first URL is the one for HTTP clients, the second one is the URL for RMI clients, the third one is the URL for the HTTPS clients. This data can be used in the client to obtain a session to that node.

5.4.3 Event Handling

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

6 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_forminstance`, and `avw_formversion`. The timer *ArchiveTimer* 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. 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.

4. 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.

7 @enterprise and Datasources

This chapter describes the configuration of datasources in @enterprise and gives example configurations for *Tomcat6* and *Jetty 6.1*.

Before version 8.0, @enterprise could use the traditional method to acquire connections to the database via the DriverManager. From version 8.0 and onward, datasources are an alternative way to obtain database connections.

7.1 Configuration of @enterprise

To use a datasource in @enterprise, the JNDI-path of the datasource must be specified instead of the JDBC-URL. Instead of e.g.

```
jdbc:derby://localhost:1527/ep;create=true
```

use something like

```
jdbc/DerbyDB
```

Do not include the initial path (`java:/comp/env/`) in the datasource path, it will be prepended automatically. For using the datasource it is not needed to provide a JDBC-driver or to fill in the following configuration items:

- Database Userid
- Database Password

The other database related configuration items are still needed and used.

7.2 Configuration of a Datasource in Tomcat6

This section describes how Tomcat6 can be configured:

1. Put the JAR-file of the JDBC-driver into the *lib* directory of Tomcat.
2. Deploy the @enterprise WAR-file (e.g. using ep80 as the contextpath)
3. Go to `../conf/<service>/<host>` directory and put a `<contextpath>.xml` file there.

- `<service>`: The name of the Tomcat service (as in the service element in the `../conf/server.xml` file); usually Catalina
- `<host>`: The name of the Tomcat host (as in the host element in the `../conf/server.xml` file); usually localhost
- `<contextpath>`: The contextpath where @enterprise is deployed

So, you would end up with a file named `../conf/Catalina/localhost/ep80.xml`. In this file specify the datasource as a resource within the context element:

```
<Context>
  <Resource
    name="jdbc/DerbyDB"
    auth="Container"
    type="javax.sql.DataSource"
    factory="org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory"
    maxActive="12"
    username="derby"
    password="derby"
    driverClassName="org.apache.derby.jdbc.ClientDriver"
    url="jdbc:derby://localhost:1527/ep;create=true"
  />
</Context>
```

The value of the *name* attribute must match the path of the datasource in the @enterprise configuration file.

Details for the other parameters can be found in the Tomcat documentation.

4. The following step may not be needed. Include the reference to the resource in the `web.xml` descriptor of @enterprise application:

```
<resource-ref>
  <description>DB Connection</description>
  <res-ref-name>jdbc/DerbyDB</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

The content of the *resource-ref-name* element must match the path of the datasource in the @enterprise configuration file.

5. Restart Tomcat, start the @enterprise application and begin to setup @enterprise.

7.3 Configuration of a Datasource in Jetty 6.1

This section describes the configuration of @enterprise running as web-application in a jetty-installation (@enterprise does not start jetty as embedded web-server!):

1. Create a *myjetty.xml* file that activates the needed jetty-plus features. Using the *jetty.xml* and *jetty-plus.xml* files as model. Add the following lines to the server configuration element:

```
<Array id="plusConfig" type="java.lang.String">
  <Item>org.mortbay.jetty.webapp.WebInfConfiguration</Item>
  <Item>org.mortbay.jetty.plus.webapp.EnvConfiguration</Item>
  <Item>org.mortbay.jetty.plus.webapp.Configuration</Item>
  <Item>org.mortbay.jetty.webapp.JettyWebXmlConfiguration</Item>
  <Item>org.mortbay.jetty.webapp.TagLibConfiguration</Item>
</Array>
```

Add the configuration classes also to the *addLifeCycle* call element:

```
<Call name="addLifeCycle">
  <Arg>
    <New class="org.mortbay.jetty.deployer.WebAppDeployer">
      <Set name="contexts"><Ref id="Contexts"/></Set>
      <Set name="webAppDir">
        <SystemProperty name="jetty.home" default="."/>/webapps
      </Set>
      <Set name="parentLoaderPriority">>false</Set>
      <Set name="extract">>true</Set>
      <Set name="allowDuplicates">>false</Set>
      <Set name="defaultsDescriptor">
        <SystemProperty name="jetty.home" default="."/>/etc/webdefault.xml
      </Set>
      <Set name="configurationClasses"><Ref id="plusConfig"/></Set>
    </New>
  </Arg>
</Call>
```

2. Uncompress the @enterprise WAR-file (e.g. using ep80 as contextpath).
3. Put the JAR-file of the JDBC-driver into the *lib* directory of the web-application.
4. Go to the *../webapps/ep80/WEB-INF* directory and put a *jetty-env.xml* file there. In this file specify the datasource as a resource within a context element:

```
<Configure class="org.mortbay.jetty.webapp.WebAppContext">
  <New id="DerbyDB" class="org.mortbay.jetty.plus.naming.Resource">
    <Arg>jdbc/DerbyDB</Arg>
    <Arg>
      <New class="org.apache.derby.jdbc.ClientDataSource">
        <Set name="databaseName">ep</Set>
        <Set name="portNumber">1527</Set>
```

```
<Set name="serverName">localhost</Set>
<Set name="user">derby</Set>
<Set name="password">derby</Set>
</New>
</Arg>
</New>
</Configure>
```

The value of the first *Arg* element must match the path of the datasource in the @enterprise configuration file.

5. Restart Jetty with following parameter and begin to setup @enterprise:

```
java -jar start.jar etc/myjetty.xml
```

7.4 Considerations for pooled Datasources

@enterprise still uses its own connection pool, even when the datasource is a pooled one. We have better control over the connection this way, and can provide all features of the @enterprise pool itself (session environment, automatic reconnect, ...).

This strategy imposes two requirements for a pooled datasource:

- It should never expect to get the connection back or destroy connections in use. In a DBCP connection pool, this can be implemented via

```
removeAbandoned="false"
```

- The pool size should be large enough to provide the max. number of connections specified in the @enterprise configuration (see chapter 3) increased by at least 2 (for internal connections used by the engine itself).

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:

```
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 data files 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.