



bedework

Bedework Reference Manual

Bedework version 3.4.1.1

Last modified: June 4, 2008

<http://www.bedework.org>

Bedework Reference Manual

The Bedework Manual contains an overview of the system, instructions for customizing and installing a production version of Bedework, and detailed descriptions of the various components making up the system.

Table of Contents

Chapter 1 Overview

Calendar collections and folders.....	5
Subscriptions.....	6
Views.....	6
1.1 Public Events Calendaring.....	6
Public Event Administration: Users & Groups.....	7
Groups:.....	7
Group structure and access control:.....	8
Administrative users:.....	9
Authentication.....	9
1.2 Calendar Suites.....	10
1.3 Personal & Group Calendaring.....	11

Chapter 2 Getting Started

2.1 The Quickstart.....	12
Packaged with the quickstart.....	12
System requirements.....	12
Instructions.....	12
Logging.....	14
Notes.....	14
2.2 Taking it further: initializing the database.....	14

Chapter 3 Deploying Bedework

3.1 Prerequisites.....	16
Summary.....	16
Requirements	16
Supported databases.....	16
Unsupported databases.....	17
Install the quickstart and test your environment.....	17
3.2 Prepare a localized version of Bedework.....	17
Prepare your build and runtime properties.....	17

The properties file.....	19
Install.....	19
Global.....	19
Application.....	20
The options file.....	20
The name property.....	21
The tzid property.....	21
The systemid property.....	21
Other properties.....	21
Copy and update the skins (templates).....	21
Stylesheets and Path discovery.....	21
3.3 Set up a production database.....	23
Language settings.....	23
Configure your applications context.xml.....	24
Configure Tomcat 5.5.x.....	25
Allow no roles.....	25
Add Jdbc drivers.....	26
Allow directory browsing.....	26
Uri encoding.....	26
JVM parameters.....	27
Configure hibernate.....	27
Setting your SQL dialect.....	27
Build the schema and initialize.....	28
3.4 Authentication.....	29
3.5 Build and deploy.....	30
3.6 Add a super user.....	30
3.7 Add administrative groups and users.....	30
3.8 Add Calendar Suites.....	31
3.9 Create initial public calendars, subscriptions, and views.....	32
3.10 Access rights and groups.....	32
Using ldap for groups and account validation.....	33
3.11 Upgrading.....	34
Chapter 4 Performance and tuning	
4.1 Hibernate and caching.....	35
Chapter 5 Bedework and Portals	
5.1 Introduction.....	36
5.2 Limitations of JSR168.....	36
Binary – or ics – downloads.....	37

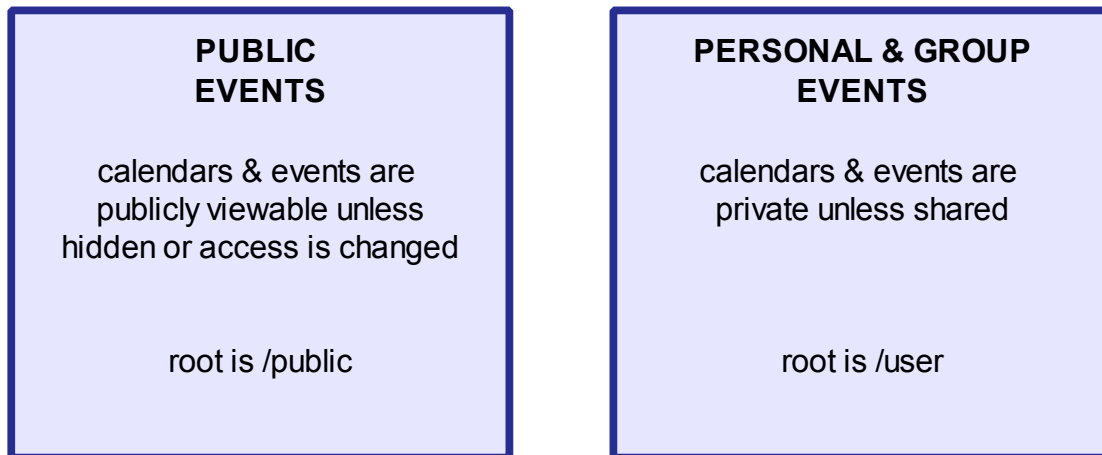
Popups.....	37
5.3 Hope on the horizon.....	37
Chapter 6 Access Control	
6.1 An overview.....	38
6.2 Definitions.....	38
Principals.....	38
ACE.....	38
ACL.....	38
6.3 Scheduling and freebusy access.....	39
Simplifying:.....	39
Chapter 7 Personal Calendars	
7.1 Overview.....	41
Default calendars.....	41
Subscriptions and views.....	41
7.2 Scheduling Meetings.....	42
Calendar users and addresses.....	42
Special Calendars: Inbox and Outbox	42
Special Calendars: Trash.....	43
Special Calendars: Deleted.....	43
Automatic processing.....	43
Scheduling resources.....	43
Chapter 8 Timezones	
8.1 System timezones.....	44
Building timezone information.....	44
Chapter 9 Bedework and CalDAV	
9.1 Summary.....	46
9.2 The Bedework implementation.....	46
9.3 CalDAV clients.....	47
9.4 Unsupported features.....	47
Recurrence features.....	47
Recurrence id ranges.....	47
Chapter 10 Other features	
10.1 Freebusy URL.....	48
10.2 Real-time scheduling	49

Chapter 1 Overview

Calendar collections and folders.

All calendar entities, events, tasks (todos) etc are stored within calendar structures consisting of 'calendar collections' and 'folders'. The distinction between folders and calendar collections is a requirement of CalDAV, one of the calendaring systems supported by bedework. A folder may not directly contain calendar objects such as events. It may contain calendar collections or other folders. A calendar collection may only contain calendar objects and may not contain folders.

The Bedework system is divided into two main spaces: the public events space, and the personal and group calendaring space. Public events are stored below a public calendar root folder and personal calendars are below the user calendar root folder.



There are a number of types of calendar collection:

1. **Collection:** the usual kind of collection of events, todos etc.
2. **Trash:** when events are 'deleted' they are moved here if possible. This gives users the opportunity to undo the delete by moving it back.
3. **Deleted:** this is used when the user does not have delete access to an entity. For example, a user might be subscribed to films and want to delete some of them from the calendar view. An annotation is created and placed in this special calendar and acts as a mask to suppress the entity. This may be changed in the future for some filter based approach.
4. **Busy:** this type of calendar will be used to hold busy time information which will be included in a free/busy display. For example, if the user always has lunch at 12midday

to 1pm, then a recurring object could be placed here so that the calendar view is not cluttered by such entries.

5. **Inbox:** this is where incoming meeting requests or published event information will appear.
6. **Outbox:** scheduling requests targeted at users on other systems will be placed here for subsequent processing.

Subscriptions

Subscriptions are made to calendars. At the moment bedework only supports subscriptions to calendars within the same system. In time we hope to allow subscriptions to external calendar systems. A subscription must have a unique name for the owner and may reference either a folder or a calendar collection. A calendar must have a subscription before it can be made visible in the personal user or public events client.

By default a user has a single subscription to their 'home' folder, for example, a user 'janet' would have a subscription to "/user/janet". There are no default subscriptions in the public space. They must all be explicitly created by administrators. Associated with each subscription is some styling information which allows for coloring or other stylistic preferences.

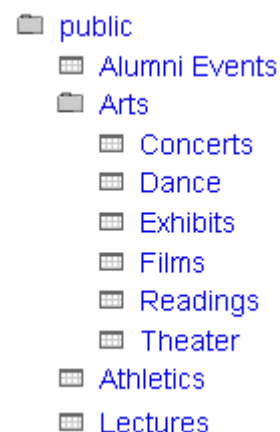
Views

Views are named collections of subscriptions. They are of most use in the public events system though they are used in personal calendars. Public events administrators can create a number of views. Typically they are used to group together subscriptions are to hold all subscriptions (as in an "All" view).

1.1 Public Events Calendaring

The point of a public events system is, in general, to disseminate event information to the widest and most appropriate audience possible. It is important that the public events are easily navigable and straightforward to find, subscribe to, and create. To achieve this, the calendar tree should be topical, consisting of high-level aggregate categories such as "Arts", "Lectures", and "Athletics".

Illustration 1 shows an example calendar structure for public events. Noticeably absent from the tree are folders named by



organization or department. As we have learned in web information architecture, users should not need to know what group within an organization's hierarchy sponsors an item to find it. However, an organization should still be able to pull a “departmental” view of the calendar information and present only that to their intended audience: this is the purpose of Calendar Suites. These are used to display a customized and filtered view of public events, but are not intended for group calendaring – which takes place within the Personal & Group calendaring space.

We believe the most flexible architecture for delivering public events is to create a large central pool of events organized by topic and to draw from this pool both for individual user and departmental access. Users who are interested in all lectures across your enterprise will be able to subscribe to a single source for these events, while departments may display only the lectures sponsored by them from the central pool.

As stated above Bedework's calendar trees are divided into folders and calendar collections. Only calendar collections may be added to folders. Events may only be added to calendar collections. In Illustration 1, it would not be possible to add an event to the root “Arts” folder; rather, an administrator must decide to which art calendar an event belongs. To keep this decision making process from becoming onerous, calendar collections within the top-level categories should be course-grained, keeping the selection of appropriate calendars for events as simple as possible.

Public Event Administration: Users & Groups

Groups:

In Bedework's current system, public events are administered centrally through the administrative web client. Every administrator is a member of a group, and events are owned not by administrative users directly, but by a special “event owner” associated with the group. Each group also has a group owner that is a specific administrative user; this role does not currently provide any functionality, but may in time (such as managing group membership).




Illustration 2: Group Structure

The event owner is a user, not found in the organizations user space, that owns all the events for a group. Having such an owner allows all group members to see and edit events within the group. Also, the event owner is the “user” who's preferences define the behavior of a calendar suite. The system ensures that these event owners are distinct from real users by prefixing them with a string, by default “agrp_”.

Within the administrative client events are filtered by the current event owner so that administrators can only see and edit events for their current group.

The super user(s) can switch to any group making all events available.

In addition, if the system is appropriately configured, locations, contacts and categories are also filtered making them editable only by the group. They are however, always readable. The other extreme of having no editable locations, contacts and categories has been requested and something near that can be achieved by creating them in a special group. The intent is to create locations so that they conform to a particular code of practice. Not allowing administrators to create locations at all may become a problem when events take place off campus.

Group structure and access control:

Access control is inherited from the top down the group tree. Therefore, it is best to create a single, top-level group for access to /public and then add all other administrative groups to it. By default, Bedework comes with a top-level group named “campusAdminGroups”. All other groups should be made members of this group to inherit write-content access on /public. Unlike calendars, groups are organizational, e.g. “Arts”, “SOE” (school of engineering), or “Athletics”.

While it is possible to close branches of the /public sub-tree from access to all administrators by creating a different group hierarchy, we discourage doing this. If public calendars are kept topical, an administrator from any group can add events to any calendar in the tree. However, an administrator can only see events created by his or her group.

It is important to remember that “departmental” calendaring (group calendaring) should be kept away from the /public tree – which is only intended for public events. Events that must not be seen by any but a subset of your community are not public (such events are to be distinguished from those intended for a subset of your community that are ok for others to see). These should be kept in the personal space, or if your need dictates it, in a top-level /dept branch of the tree which you will need to create. We believe in actual practice, most group calendaring will best be managed within the personal and group space.

Administrative users:

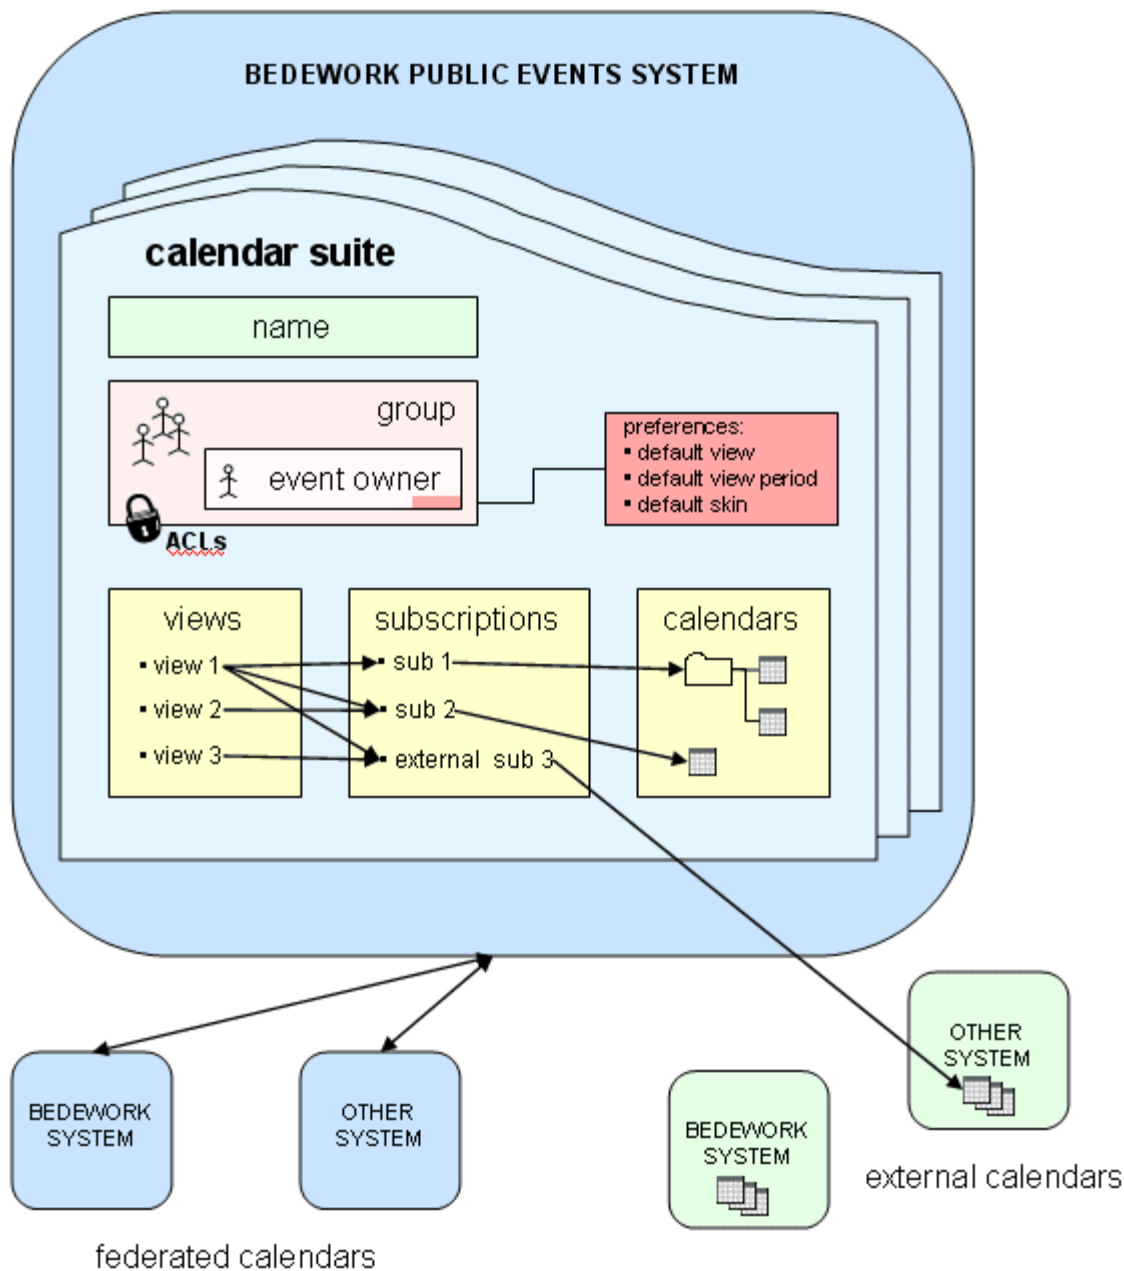
To add an administrative user to the system, simply add the user to a group. When a user is removed from all groups, the user will be removed from the system. Because the public event space is distinct from the personal calendaring space, administrative users are managed in Bedework's database by default (though they need not be).

Once in the system, an administrative user may be given roles from the "Manage public event administrators" menu item under "User management". The current meaningful roles are 'super-user' and 'publicEvent' administrator. The publicEvent role allows write-content access to the /public tree (assuming the user is in a child group of the top-level campusAdminGroups and the ACLs are not changed); the super-user role allows full access to the Bedework admin client.

Authentication

Authentication is not managed by the Bedework system internally, but by the servlet (or J2ee) container (see section 2.4, Authentication). Administrative roles and authentication are separate; a user who can authenticate, but who is not an administrative user, will see a "No Access" message upon logging into the admin client. A user who has been added as an administrator but who cannot authenticate cannot log in at all. Both conditions – authentication and authorization – must be met for administrators to work in the admin client.

1.2 Calendar Suites



The intent of calendar suites is to provide a customized view and appearance for significant groups within an organization, such as a school or large department. A calendar suite is associated with an administrative group which has an event owner (an internal bedework account) whose preferences are used to set up the skin, default view, view period, views and subscriptions.

1.3 Personal & Group Calendaring

Personal and group calendaring is the second calendaring space provided by Bedework. It is the realm of desktop calendar clients, personal schedules, shared calendars, meetings, invitations, to dos, and so on. Much of the calendaring standards focus on this aspect of calendaring, and there are many commercial examples in this space (e.g. Microsoft Exchange with Outlook).

Many of the mechanisms used by the personal calendars, such as preferences, views and subscriptions, are also used by the public user clients to configure their appearance. The major differences between public and personal events lie on the way events are shared and access is managed. A further major difference not fully handled in this release is the need for multi-language support in public events, a need nowhere near as pressing in personal and group calendaring.

Personal calendars also have the problem of interacting with external clients and handling events generated elsewhere, for example through meeting requests.

The personal calendar space is rooted of the main user calendar, named by default “/user”. (these names can be configured at build and deployment time to be more appropriate for non-English speaking institutions.)

Below this is a folder for each user of the system for which there has been some activity. Unless an organization chooses to build some form of feed, users who have never had meeting requests and have never logged on will have no presence in bedework. Their user calendars will be created automatically the first time an event is sent to the account or they log on. Personal calendaring is dealt with more fully in Personal Calendars – Page 41

Chapter 2 Getting Started

2.1 The Quickstart

To try out Bedework, begin by [downloading](#) and running the quickstart package. You can get the most recent release from the Bedework website: <http://www.bedework.org> .

The quickstart release distribution is provided for those who want to get familiar with the calendar quickly and easily, without having to compile and deploy code, and without having to set up a database. It is preloaded with data so you can see how the system looks in production.

Packaged with the quickstart

1. Bedework (version 3.4)
2. Tomcat 5.5 (apache-tomcat-5.5.17)
3. Hypersonic SQL 1.7 (hsqldb-1.7.3.3)
4. Apache Ant 1.6 (apache-ant-1.7)

System requirements

1. [JDK 1.5](#)
2. JAVA_HOME environment variable must be set
3. Nothing else should be running on port 8080 and 8887.

Instructions

For the commands below, the '<ANT>' command depends on your operating system:

Windows: ant.bat
UNIX: ./ant

Note, it is important to use the form “./ant” to avoid running a version on your path.

1. To start the calendar, open a console window and cd to the quickstart directory. Start hsql by typing:

```
<ANT> hsqldb  
(e.g. ./ant hsqldb)
```

Next start Tomcat by typing:

```
<ANT> tomcatstart  
(e.g. ./ant tomcatstart)
```

You may need to enter each command in a separate console window.

2. Once Hypersonic and Tomcat are running, you can access the calendar with your browser by entering the following URL:

<http://localhost:8080/bedework>

Note: this link will only work on the system on which the quickstart is running.

Which will give you a choice of three applications: the public events calendar (exposed in two calendar suites), the personal calendar, and a public events administration client.

3. The valid users and passwords are set in <tomcat-dir>/conf/tomcat-users.xml. In the quickstart distribution all passwords are set to *bedework*.

Administrative client login

There are users with administrative privileges:

- calowner: the lowest level - can enter public events
- caladmin: a superuser.

Personal client login:

There are a number of users with no special privileges, these can be used as test users with the personal calendar client. They are:

- caluser1, caluser2, caluser3, and
- caltest01, caltest02, caltest03, caltest04, caltest05, caltest06, caltest07, caltest08.

4. Both Tomcat and hsqldb should be stopped when you are done. If the commands are still running (hsqldb should be, and possibly Tomcat), type

```
CTRL-C
```

in the window in which the command was started.

Optionally, Tomcat can be stopped by typing

```
<ANT> tomcatstop
```

Logging

Log messages largely appear in the tomcat log in <tomcat-dir>/logs. This release uses log4j for most of the application logging. The distributed log4j.xml appears in calendar/resources in the quickstart and in common/classes in the tomcat directory. It is configured to maintain a rolling log file, server.log and also append output to the console.

In addition, a socket appender is defined which allows chainsaw to be used to watch the log output.

Notes

Calendar binaries are included so there is no need to compile any calendar source code. Also, the database is pre-loaded, so there is no need to run any database scripts.

Thanks to the uPortal team (<http://www.ja-sig.org>) for the inspiration and the quickstart release on which ours is based.

2.2 Taking it further: initializing the database

If you've successfully run the quickstart release with the instructions above, you may wish to initialize the database to clean out the example data and explore the Bedework system with a clean slate.

These instructions explain how to do this with the Hypersonic database packaged with the quickstart. *If you are ready to use a different database and set up your production environment, please see Chapter 3.*

1. Unzip the dump/restore utility in a useful location, for example in the root of the quickstart directory:

```
unzip bedework/dist/dumpres.zip
```

This will create a directory named "dumpres".

2. Stop HSQL (if running) and rename hsqldb-1.7.3.3/demo to move the demo database out of the way. (You may also simply delete it.)
3. Restart HSQL, and Hypersonic will create a new, empty demo database.
4. From within the dumpres directory, generate the tables (on linux, make sure the bwrn.sh script is executable first: `chmod +x bwrn.sh`):

```
cd dumpres
```

```
./bwrun.sh schema-export  
(windows: bwrun.bat schema-export)
```

5. Finally, initialize the database with `dumpres/data/initbedework.xml`:

```
./bwrun.sh initdb  
(windows: bwrun.bat initdb)
```

`initbedework.xml` creates a suggested calendar structure, a MainCampus calendar suite, a few administrative groups, and everything else you should require to get up and running with a clean instance of Bedework. No events are created.

`initbedework.xml` creates exactly two administrative users:

1. a super-user "caladmin" with password "bedework", and
2. a normal admin "calowner" with password "bedework"

Use these userids for access to the administrative web client after initializing with `initbedework.xml`.

For more information about the dump/restore utility and its use, please refer to chapter 3.3.

Chapter 3 Deploying Bedework

3.1 Prerequisites

Summary

This section describes setting up your build environment and provides guidance on creating your initial calendars, subscriptions, views, and administrative users. For instructions on customizing the look and layout of your production system, please see the Bedework Design Guide on the Bedework website: www.bedework.org.

Requirements

- JDK 1.5 – specifically, Java 1.5.0_11 or later
- JAVA_HOME environment variable must be set
- Hardware for testing: most current desktop or laptops will be adequate.
- Hardware for production: A server class machine generally with at least 1Gig allocated to the jvm (see settings below).
- An adequate database system. In particular, the database probably has to support unicode.
- To implement the Bedework calendaring system, it is useful to understand the following:
 - [Java servlets: http://java.sun.com/products/servlet/docs.html](http://java.sun.com/products/servlet/docs.html)
 - [Servlet containers](#) (e.g. [Tomcat](#), [JBoss](#))
 - Authentication is local to your site - some Java programming may be necessary to accomplish this.

Supported databases.

Bedework uses hibernate (<http://www.hibernate.org>) as a persistence engine. Bedework therefore should run on any database supported by hibernate. In reality, there are some problems, though not insurmountable, with the support of some systems which may require hand-editing of the schema. In future releases we hope to minimize those issues. The current list of hibernate supported databases can be found on their site. We hope, eventually to support a wide range of database systems. The list below reflects the systems on which we have successfully deployed and run.

Databases may be unsupported at the moment (or permanently) but it may still be possible to massage the schema enough to make bedework run. We try to indicate why they are not

supported and some will eventually move into the supported category.

Currently, there are versions of bedework deployed on at least:

- Hsql – the database we provide with the quickstart. It is unclear how safe it is to use as a production database. It appears to be used by various systems to provide a persistence mechanism for messages etc (jboss?) so it might be usable on a small scale.
- MySQL version 5
- Oracle Version 10 and perhaps Version 9 with Version 10 jdbc drivers.

Unsupported databases

Because of our use of hibernate, we don't support databases they don't support. Other problems may be unresolved at the moment but in time we will discover a workable solution.

- MySQL version 4: Has problems in a number of areas. These are unlikely to ever be resolved. Version 4 is now old.
- MS Sql Server: Partially supported but requires at least hand-editing of the schema. Has not been tried with 3.3.1 onwards. At least one problem remains, Sql Server does not follow the ANSI standard for unique indexes in that null=null for Sql Server but null is never equal to null in ANSI standard databases. This breaks some of the unique indexes in bedework.
- Postgres: We hope to do more work in the near future supporting postgres. At this stage it is possibly mostly a matter of unfamiliarity.

Install the quickstart and test your environment

Before attempting any customization, please test your environment by running the quickstart release. It is always wise to test your changes incrementally; test each small change to make certain you understand its effects. Doing these two things will help you understand the system and will provide useful information to the Bedework support community if you run into trouble and wish to ask for help.

3.2 Prepare a localized version of Bedework

Prepare your build and runtime properties

Bedework uses ant for the build and deploy process and a number of property files are used to control that process. Ant properties have the characteristics that once set they cannot be

modified so to override default settings you need to set them earlier in the process.

The build first looks for a property file called **bedework.build.properties** in your home directory. Your home directory in unix is usually

```
/home/userid
```

And in windows is typically

```
C:\Documents and Settings\userid
```

So, for example, the local properties file in windows would live here:

```
C:\Documents and Settings\userid\bedework.build.properties
```

You can set properties in this file which will override the default settings. In particular you can tell the build system the location of the configuration you want to build with the setting (e.g):

```
# Location of our bedework property files
org.bedework.config.properties=${user.home}/bwbuild/myconfig.properties
org.bedework.config.options=${user.home}/bwbuild/myconfig.options.xml
```

This would cause the build to include the files

```
<home>/bwbuild/myconfig.properties and
<home>/bwbuild/myconfig.options.xml
```

The default settings mean the system builds and deploys using the files

```
bedework/config/configs/democal.properties and
bedework/config/configs/democal.options.xml.
```

Do not change the original files. Please keep them for reference. Make copies named appropriately and set the properties above to use them.

The properties file is mostly associated with the deployment process while the options.xml file is for runtime properties. Gradually this demarcation is being cleaned up so that the properties file will eventually not be included on the class path. Some properties are needed at deployment and at run time. These are copied during deployment from the properties file into the options file.

The properties file.

The properties file is divided into sections with different property prefixes.

Install

The section prefixed “org.bedework.install” defines which applications are to be installed. This consists of a list of application names.

For each name there should be a corresponding section prefixed with “org.bedework.app.<name>” and also a corresponding section in the options file.

The default configuration comes with a number of application configurations. When creating a configuration it is appropriate to simplify by removing unneeded applications. Remove the name from the install property and delete the appropriate section.

Applications in the default configuration are:

- **bwconfig** – (not in the list but a configuration section is present). This is a try at a web based configuration application.
- **CalAdmin** – the administration application
- **Events** – public events (unauthenticated access to calendars)
- **SoEDept** – example calendar suite
- **UserCal** – personal events
- **Pubcaldav** – public events caldav server. (unauthenticated access to calendars)
- **Usercaldav** - personal caldav server -
- **caldavTest** – a test application for caldav
- **test** – a test suite
- **dumpres** – the dump restore utility
- **restoreFrom2p3px** – the dump restore utility configured to restore from uwcalendar 2.3.x

Global

The section prefixed “org.bedework.global” defines properties global to the whole deployment process.

Application

The section with properties prefixed “org.bedework.app.<name>” are the application deployment properties, one section per named application.

Two properties define the project and type of application. The value of the property “org.bedework.app.<name>.project” defines which project the application is a part of. Currently these can be

- “caldav” - a caldav server
- “caldavTest” - a caldav test package
- “webapps” - a web client
- “dumprestore” - a dump/restore application
- “freebusy” - the freebusy aggregator

The value of the property “org.bedework.app.<name>.type” corresponds to the name of a subdirectory in bedework/deployment, e.g. webpublic, webadmin, etc. So to define the administrative client named CalAdmin of type webadmin we have the fragments:

```
org.bedework.install.app.names=...,CalAdmin,...  
...  
org.bedework.app.CalAdmin.project=webapps  
org.bedework.app.CalAdmin.type=webadmin
```

Multiple versions of each application type may be deployed, each configured differently. This is of importance for calendar suites (departmental calendars).

The options file.

This xml file contains run time properties and is divided into sections much like the properties file. Some values may be copied out of the properties file if they affect both the deployment and run time. Most of the options are used to set field values in named classes so that the application will load the settings once only with a single call.

It is important to set the system properties for a new system. These are found in the “syspars” section of the properties file. A number can be left with the default values and some are not yet implemented. The properties it is particularly important to set (and their default settings are:

```
<name>bedework</name>  
<tzid>America/New_York</tzid>
```

```
<systemid>demobedework@cal.mysite.edu</systemid>
```

The name property

This property appears in the system table as the primary key.

The tzid property

The tzid is the default timezone to be used for times and dates.

The systemid property

The systemid is used when generating uids for calendar entities. This name should be related to your site for ease of identification and if you run multiple systems should be different for each. In addition it takes part in the creation and interpretation of calendar user addresses which appear in attendees. The part following “@” will probably be the domain to which imip messages are addressed (in some as yet undefined manner).

Calendar user addresses take the form of a “mailto:” uri so that user “testuser01” on a system configured as above would have a calendar user address of

```
mailto:testuser01@cal.mysite.edu
```

Other properties

There are also a number of names used when creating default calendars. These should be set to some appropriate localized value.

The size settings are mostly unused at the moment.

The property

```
<userauthClass>org.bedework.calcore.hibernate.UserAuthUWDbImpl  
</userauthClass>
```

defines which class handles administrative groups for the administrative client. The group class setting are explained in the “Access rights and groups” section below.

Copy and update the skins (templates)

Now is the time to localize the skins. This can be as simple as replacing the title graphics and text, or as involved as modifying the global layout and behavior of the front-end. Please see the Bedework Design Guide for instructions on updating layout and styles.

Stylesheets and Path discovery

Your XSL stylesheets and associated template images and resources (e.g. css files) are located

on a web server at the url specified by the property "app.<name>.root" (and "app.<name>.cal.suite" for the public client) in the config file (e.g. config/configs/myconfig.properties). Each of the three web clients (public/guest, personal, and admin) has an associated approot where these files are to be found by the system. For example, given the values in the properties file:

```
org.bedework.app.CalAdmin.root=http://somewebsserver/bedework-3-1-admin
org.bedework.app.Events.root=http://somewebsserver/bedework-3-1-guest
org.bedework.app.Events.cal.suite=MainCampus
org.bedework.app.SoeDept.root=http://somewebsserver/bedework-3-1-guest
org.bedework.app.SoeDept.cal.suite=SoeDept
org.bedework.app.UserCal.root=http://somewebsserver/bedework-3-1-personal
```

the administrative client and user client stylesheets will be found at the given url while the Events public (guest) client will have its stylesheets located at <http://somewebsserver/bedework-3-1-guest.MainCampus>

Note: the above property values currently have to be copied into the <approot> element in the options.xml file.

Placing the stylesheets and resources on a separate web server (we suggest/encourage you to use your primary web server) will make them significantly more convenient to access and manipulate. Note that if you choose to serve a client over https, you will need to specify the same protocol for the approot to avoid mixed content messages. Note that the stylesheets are *only* loaded at first reference (or if an administrator explicitly flushes the stylesheets).

Under this is a 3 tier structure based on:

- locale
- user agent
- stylesheet name.

The top two are normally named "default". The filters will work down the structure trying a specific name first then trying default. For example, in the locale "fr_CA" a path

```
{ $appRoot } / fr_CA / default / ...
```

could lead to a set of French stylesheets while

```
{ $appRoot } / default / default / ...
```

could lead to English.

The discovered 'real' path is cached with the 'idealized' path as the key so that subsequent lookups for the same path will proceed without the discovery phase.

A deploy time property can specify that the target server has directory browsing disallowed, in which case the filters search for a marker file called `xsldir.properties`

3.3 Set up a production database

The quickstart is set up to use `hsqldb` and this provides a good way to try out the system. To move to another database system you will need to configure the build to affect the following:

- Configure `tomcat` or each context to use another database
- Configure `hibernate` to use the appropriate dialect
- Rebuild
- Configure the `dump/restore`

Note: if you'd like to set up an initialized database in `hsqldb` for testing, you can skip to "Build Schema and Initialize" on page 10 and just run `./bwrn schema-export` followed by `./bwrn initdb` against an empty `hsqldb` database. To create this, stop `hypersonic` (if running) and rename (or throw away) the directory named `<hsqldb-dir>/demo`. Restart `hsqldb` and the directory will be recreated in an empty state. See [Build Schema and Initialise](#) for more information about the `schema-export` and `initdb` commands.

Language settings.

Your database needs to support at least unicode. For most European and American systems this may not be obvious at first but eventually problems will occur with certain special characters. For most non-European languages unicode or some other multibyte support is an absolute necessity.

Each database has its own settings for language support. Many have built in traps for the unwary (most of us). For example, in `Postgres` `UNICODE` apparently means `UTF8`. This is incorrect for Chinese for example.

In addition, there are no checks that the application, e.g the calendar, has the same settings as the database. Databases will interpret the byte stream according to their configuration even if that does not match the configuration of the calendar server. Care in matching up all the components is obviously needed. Those components are at least:

- The client (browser, caldav client etc)
- The servlet container (`jboss`, `tomcat`)
- The servlet (bedework)

- Jdbc settings
- The database

Configure your applications context.xml

Tomcat 5.5.x applications are configured by setting up context.xml files for each application.

Each bedework web application has one or more predefined context.xml files located in the applications war/META-INF directory. Currently all are almost identical and look like this:

```
<Context path="@CONTEXT-ROOT@" reloadable="false">
  <Resource name="jdbc/calDB" auth="Container"
    type="javax.sql.DataSource"
    driverClassName="org.hsqldb.jdbcDriver"
    url="jdbc:hsqldb:hsqldb://localhost:8887"
    username="sa"
    password=""
    maxActive="8"
    maxIdle="4"
    maxWait="-1"
    defaultAutoCommit="false" />

  <!-- Disables restart persistence of sessions -->
  <Manager pathname="" />
</Context>
```

The “@CONTEXT-ROOT@” is an example of a token which will be replaced during deployment by a value from the properties file. You probably only need one of these files, if you have the above token in your file it will be replaced with the correct context root.

To use your own copy of the context.xml file, create such a copy in an accessible location and then update your properties file to refer to it. The relevant property looks like (for example)

```
...
org.bedework.app.CalAdmin.tomcat.context.xml=war/META-INF/context.xml
...
org.bedework.app.Events.tomcat.context.xml=war/META-INF/publiccontext.xml
...
```

and you would change them to something like:

```
...
```



```
org.bedework.app.CalAdmin.tomcat.context.xml=<home>/bwbuild/prod/context.xml
...
org.bedework.app.Events.tomcat.context.xml=<home>/bwbuild/prod/context.xml
...
```

Configure Tomcat 5.5.x

With the above configuration you do not need to configure a database in tomcat's server.xml. However some changes are needed. These are:

- Allow no roles
- Add jdbc drivers
- Allow directory browsing
- Jvm parameters

Allow no roles

Tomcat's handling of security constraints was at some point changed to be absolutely compliant with the servlet specification. Unfortunately, in this regard at least, the specification is somewhat impractical. Within the web.xml security-constraint element is a role-name element which is set to "*" as in this example:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>CalendarAdmin</web-resource-name>
    <description>Public events Administration</description>
    <url-pattern>/*</url-pattern>
  </web-resource-collection>
  <auth-constraint>
    <role-name>*</role-name>
  </auth-constraint> ...
```

The entry

```
<role-name>*</role-name>
```

used to mean ANY role. It now means by default any role defined in the web.xml. There is no way defined in the servlet specification to have security constraints without roles.

However, it appears that the behavior is configurable in tomcat. In the server.xml there is the

definition of the Realm used for authentication. Add a setting for the allRolesMode attribute, in the quickstart it will look like

```
<Realm className="org.apache.catalina.realm.UserDatabaseRealm"
        resourceName="UserDatabase"
        allRolesMode="authOnly" />
```

and you should be able to login without any role assigned.

Add Jdbc drivers

You will need to ensure that the driver jar is added to the common/lib directory.

Allow directory browsing

This change is only required if you are using tomcat to host your stylesheets. Our recommendation is to host them on your public web server. If you do wish to host them on tomcat you must decide whether you want directory browsing enabled or disabled. Enabled allows bedework to discover the stylesheets by checking for the existence of directories. Otherwise marker files must be created. The tomcat default is to disallow directory browsing. If you wish to allow browsing open the file conf/web.xml and find the setting for the listing property:

```
<init-param>
    <param-name>listings</param-name>
    <param-value>>false</param-value>
</init-param>
```

Change the value "false" to "true" to allow browsing of directories.

Uri encoding

By default tomcat uses ISO-8859-1. If you want to use Tomcat with UTF-8, you must add the following parameter in your server.xml file: **URIEncoding="UTF-8"**.

```
<Connector port="8080" maxHttpHeaderSize="8192"
           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
           enableLookups="false" redirectPort="8443"
           acceptCount="100"
           connectionTimeout="20000" disableUploadTimeout="true"
           URIEncoding="UTF-8" />
```

JVM parameters

The quickstart jvm parameters may be adequate for small scale testing or development. For production use you will certainly need to modify the JVM parameters. Below is a sample set of options (from a jboss configuration).

```
JAVA_OPTS="$JAVA_OPTS -server -Xms1792m -Xmx1792m"
JAVA_OPTS="$JAVA_OPTS -XX:NewSize=512m -XX:MaxNewSize=512m"
JAVA_OPTS="$JAVA_OPTS -XX:SurvivorRatio=2"
JAVA_OPTS="$JAVA_OPTS -XX:PermSize=128m -XX:MaxPermSize=128m"
JAVA_OPTS="$JAVA_OPTS -XX:+UseParallelGC -XX:+UseAdaptiveSizePolicy"
JAVA_OPTS="$JAVA_OPTS -XX:+AggressiveHeap"
```

Configure hibernate

Setting your SQL dialect

Configuring hibernate to use the appropriate dialect is done in your configuration file.

For example, if your **bedework.build.properties** specifies a properties file at **<home>/bwbuild/myconfig.properties** (see [Setting up your local configuration](#)) then in that file you need to set a few properties

```
org.bedework.global.hibernate.dialect=yourDialect
```

and at least

```
org.bedework.app.dumpres.hibernate.dialect=yourDialect
```

The value **yourDialect** is a defined hibernate SQL dialect such as `org.hibernate.dialect.HSQLDialect` or `org.hibernate.dialect.MySQL5Dialect`. The dialect is a class defined on the class path. Hibernate defines a number of 'standard' dialects.

For example, the global property if using MySQL5 would be

```
org.bedework.global.hibernate.dialect=org.hibernate.dialect.MySQL5Dialect
```

A list of dialects understood by Hibernate can be found at:

http://www.hibernate.org/hib_docs/v3/reference/en/html/session-configuration.html#configuration-optional-dialects

For the dialects included with bedework, look at the `org.hibernate.dialect` classes in the Hibernate jar file (for example, to use MySQL 5 which is not currently listed in the on-line Hibernate documentation, use `org.hibernate.dialect.MySQL5Dialect`).

Before building the system, if you are building for a database other than the quickstart hsqldb, you will need to make the appropriate jdbc drivers available. Place the driver jar file in the directory bedework/lib/jdbc and it will appear on the class path for the schema and the dump and restore applications.

Build the schema and initialize

With all that in place it is now possible to create a schema and initialize the database. The deploy process created a zip file in the bedework/dist directory which can be unwrapped to run the schema build. The default name for that application is dumpres.zip. This file contains a Unix shell script named bwrun.sh and a Windows bat file named bwrun.bat.

Unzip that file to a convenient location and in linux ensure the file called bwrun.sh is executable (**chmod +x bwrun.sh**).

To create a schema file for your system enter:

```
./bwrun schema
```

This produces a file named "schema.sql".

The hibernate schema tool can also create all the tables and constraints directly in the target database. If you wish to use this option then enter:

```
./bwrun schema-export
```

The tables will be created in your database (a file named "schema.sql" is also created). This process is non-destructive, so if you wish to start with a clean database, you will need to manually drop your tables. A word of warning here; while the schema-export option is useful there is little or no error checking taking place. Check back through the output for errors.

To initialize an empty database use initdb, e.g.

```
./bwrun initdb -ndebug -indexroot path-for-lucene
```

This uses an initial data file wrapped up with the dump restore application (data/initbedework.xml). (See also the "restore" target below to restore a different file.)

The optional parameter -ndebug will turn off debugging output. This makes things faster for large restores.

The indexroot parameter is used to tell the restore where your lucene indexes should be built. In the quickstart this is a relative path because the quickstart might be restored anywhere, but in general this should be an absolute path.

For testing, a directory in your home directory might be appropriate, for production use perhaps a /data directory, e.g. /data/calendar-3.3.1/lucene/indexes

This directory must be destroyed and recreated before restoring the data.

Note: initbedework.xml creates a super-user "caladmin" with password "bedework". However, there is a potential problem here if you are trying to bring yourself up in your local user address space; you need an administrator that can log into the admin client. To simplify this, we will probably add a property to set the administrative user you wish to use in future releases.

To dump the database data use

```
./bwrun dump <filename>
```

The application can be used to produce a regular nightly xml dump of the data. In future releases the dump/restore will be reformatted to facilitate restoration of a single users data.

To restore the data use

```
./bwrun restore <filename> -ndebug -indexroot path-for-lucene
```

where <filename> is the name of a dump file produced by the dump process. The tables must be empty to restore the data.

The indexroot parameter is needed if you want to change the one set in the data. Remember to empty (or destroy and recreate) the directory.

3.4 Authentication

The calendar uses container-based authentication as defined by the [Java servlet specification](#). There is no authentication code within the calendar system.

Authentication can be managed by the servlet container in a number of ways which are currently beyond the scope of this document. Tomcat's implementation can be configured, at its simplest, in its <tomcat>/conf/tomcat-users.xml file. The tomcat website provides details on configuring tomcat to use other forms of authentication, including ldap and databases. (see also [Tomcat SSL HowTo](#)).

An alternative which has been implemented is to use filter based [Yale/JA-SIG CAS](#).

3.5 Build and deploy

Build the calendar with the command:

```
./ant clean.deploy.debug
```

(“.debug” will provide debugging output in the server log. You can remove it and rebuild when you are convinced you are production ready.)

```
./ant clean.deploy
```

Do not specify “ant”. The form above executes the ant contained in the quickstart which has some extra features installed. If you use the form “ant” you will execute the ant installed in your system which may be an earlier version and/or not have the required libraries.

This will create a number of WAR files in **<bedwork>/dist/** including for example:

```
cal.war, caladmin.war, and ucal.war
```

If you are using the Tomcat directory within the quickstart distribution, your application war files are now deployed. Otherwise, collect the war files and drop them in your container.

3.6 Add a super user

At some point you will switch from the quickstart authentication to your sites authentication. Before this happens you need to ensure you have an administrative super user that exists within your own authentication domain.

Two ways to achieve this are to:

1. Create a user in your domain, e.g. caladmin, that is already set up as a superuser in the bedework quickstart.
2. Create a bedework superuser with an account that already exists in your authentication domain.

Option 2 is probably the most appropriate and as the deployer it is probably appropriate to use your own account, at least initially.

3.7 Add administrative groups and users

In the quickstart administrative groups are stored in the calendar database. Administrative groups are intended to be separate from user groups to allow different access rights to be

defined for administrative users. (See "Access rights and groups" below)

1. Within the /caladmin application, select "Admin Groups: Add"
2. Give the group a name, a description, provide a userId for the owner, and set the "Events Owner" to `agrp_groupName`, or something meaningful that will easily identify which group the event belongs to. (Note: you should leave the prefix on group owners; the prefix is defined in the properties file as the property `org.bedework.app.CalAdmin.admingroups-prefix`)
3. Once the group is added, add members by providing a userId that should map to the userId used to authenticate to the administrative client.

3.8 Add Calendar Suites

Note: You cannot run a public client unless an associated calendar suite has been defined in the admin client. These steps document setting up an example calendar suite; note that the name for the suite must match the "cal.suite" property in your configuration properties file.

For example, the quickstart defines a calendar "Events" and the property "org.bedework.app.Events.cal.suite" has the value "MainCampus".

- Within the /caladmin application, create a group which will effectively own the new suite. Only administrative users who are members of this group will be able to administer this suite. For example we can create a group called "myDept" with owner set to your user account and the events owner set to "agrp_mydept"
- Add users and/or groups to that administrative group.
- Go back to the main menu, select "Manage Calendar Suites" and add a new suite. Give it a name, set the group, e.g. myDept from above and set the root calendar to be "/public".
- Give administrators write access. At least add yourself.
- Log out and log in again and select the group you just created if given a choice.
- Currently, the system creates a default subscription you need to remove. Go to manage views, select "All" and remove the single subscription. Next, via the main menu, go to manage subscriptions, select that subscription and delete it.
- Add subscriptions to those public calendars you want to see through one or more views. Set "display" to yes as you add them.

- Via the main menu, select Manage Views and add any other views you want, adding subscriptions to those views.
- To set the default view (and other preferences) for the new calendar suite, go to “Manage Preferences” on the main menu. Until you are familiar managing views, you may wish to set the default view to “All”..

3.9 Create initial public calendars, subscriptions, and views

- Within the /caladmin application, select: “Manage calendars” from the Main Menu. You will see a default set of calendars that you can manipulate.
- When you are satisfied with your collection of calendars, select “Manage subscriptions” from the Main Menu. You will see a default set of subscriptions that map to top level calendars. A one-to-one relationship does not have to exist between top-level calendars and subscriptions, but it is useful for initially setting up your views. Set any subscription preferences here, including css style.
- Finally, create views by selecting “Manage views” from the Main Menu. Create views by adding subscriptions to the view. Again, a one-to-one relationship does not have to exist between top-level calendars and views, but in many cases, this is what you'll want to expose in your public calendar system. The views named here will present themselves in the pull down list of views in the public client.
- By default, there is a view named “All” that contains all existing subscriptions. The “All” view is set as the default view for the special user “public-user”. If you want to change the default view, enter the special user “public-user” in the “Edit user preferences (enter userid):” field in the User Management section of the Main Menu; you will be presented with a form in which you can change the default view.

3.10 Access rights and groups

Access to resources in Bedework is controlled by an access control system based on WebDAV and CalDAV. A user’s access is based on their identity and group membership. The system allows a different group structure for administrative and user access. This ensures that a given user, when logged in to the user client, does not have any special administrative rights which may lead to unfortunate consequences, such as deleting a public calendar by mistake, or adding private events to public calendars.

Two system parameters determine this behavior, the initial values are set in the xml options file, democal.options.xml in the quickstart. The relevant elements in the syspars section are:


```
<admingroupsClass>org.bedework.calcore.hibernate.AdminGroupsDbImpl
</admingroupsClass>
<!--
<usergroupsClass>org.bedework.calcore.ldap.UserGroupsLdapImpl
</usergroupsClass>
-->
<usergroupsClass>org.bedework.calcore.hibernate.GroupsDbImpl
</usergroupsClass>
```

Note the second of the three is commented out. These settings define which class will be used to manage groups for the administrative and user clients. The first class is an implementation which uses a the Bedework database to store the administrative groups.

The last setting is a dummy class which does nothing but which could use the database to allow testing or perhaps even for sites which don't want to use a site-wide directory.

The commented out setting is a class which uses ldap to determine group membership. In the options file is a section labeled "user-ldap-group" which configures this class. This class should also be usable with Active Directory.

Using ldap for groups and account validation

Authentication is performed outside of bedework, either by the container (e.g. tomcat) or by a filter mechanism or some other approach.

However, bedework still requires access to directory services of some kind to determine group membership and to check whether accounts are valid. As indicated above, a class is available to perform group lookup and account validation using an ldap directory.

The current implementation of this class assumes the directory is readable. In production this is unlikely to be adequately secure so we will be working on changes to improve the security. One possibility is to define a custom resource in the container to provide access to the naming context. That will allow some of the authentication parameters to be moved out into the container configuration.

In addition the class named above may not carry out all the checks appropriate for your organization. In particular you may want to implement a class that overrides the `validUser` method. The default carries out no checks with your system directory. A deployed version should probably check the directory to ensure the user does exist.

3.11 Upgrading

Bedework has been running at your site for some time and new versions have appeared and now it's time to upgrade. How do you migrate all that data from one version to the next?

The dump/restore utility is the tool to use. The intent is that any given version of the restore process will be able to read dump files produced by earlier versions.

Currently we expect to be able to read data dumped from UWCalendar 2.3.2 (with slight massaging) and any version of bedework. Later on we intend dropping UWCalendar support but it will still be possible to migrate through a two step process.

Chapter 4 Performance and tuning

4.1 Hibernate and caching

The performance of your system is likely to be very dependent upon the caching strategies used. For small to medium deployments the default settings may be perfectly adequate.

For large systems you may need to spend some time reading the documentation available at the hibernate site and the sites of the various cache implementors.

Since release 3.4.1 Bedework allows different cache regions for each application type. This allows the deployer to adopt aggressive caching strategies for the public events clients for example, while having less caching or much shorter flushing intervals for personal clients where immediate visibility of changes is necessary.

There are two parameters in the options xml file which affect caching. The default settings are

```
<cachingOn>true</cachingOn>
<cachePrefix>bwpubevents</cachePrefix>
```

These settings turn caching on and set the cache prefix to “ bwpubevents”. This is used in the ehcache settings file to distinguish cache settings.

The default cache for Bedework is currently Ehcache. The latest versions of this cache do allow clustering. Other caching schemes are available such as the Jboss cache. The hibernate site offers details on the alternatives.

Chapter 5 Bedework and Portals

5.1 Introduction

The intent is to have an application that can live inside and outside of a portal and be equally functional in both contexts. Some organizations that run bedework have no intention, at least currently, of running a portal. Others do not wish to make any of their portal applications visible outside of that framework.

To achieve this we decided to use the apache portals-struts bridge. This is a widely used package which, in theory, allows a struts based application to run inside a portal. In addition, that same deployed package is available outside of the portal if you so wish.

In reality, the portal context places a number of restrictions on the application. Some of these are easily met, (the action/render urls). Others restrictions cause problems, (see below).

While it is possible that organizations will want some different styling for the portal version, as much as possible we try to have a single stylesheet for both. Bedework is a complex application so inheriting styles from the portal may not work too well. At the very least however, much of the header and footers should be dropped for the portal.

5.2 Limitations of JSR168

To allow bedework to run in as many portals as possible we need a common framework. That framework in the portal world is JSR168 – the portlet specification.

Unfortunately, this specification left out some key pieces. The worst omission is the ability in some controlled way to set the content type. This prevents us from using a number of standard web practices in a portal agnostic fashion. We are unable to use Ajax, create pop-up windows or even to initiate a binary (or ics) download.

Most (possibly all) portals provide ways around these restrictions. However, because they are not part of any specification they are different for most portals. Thus, after the 3.3.1 release we reimplemented a number of features in the stylesheets which previously used popup windows. This does not prevent someone with sufficient knowledge of their portal framework from reimplementing such features using portal specific mechanisms. As our intent is to have an easily deployable portlet it does prevent us from distributing bedework with those features already in place.

Binary – or ics – downloads.

This still leaves us with the problem of binary downloads. In bedework this shows up in the downloading of ics files. At the moment it appears that the only portal agnostic solution to this problem is single-signon. Bedework as distributed will build links to the servlet version bypassing the portal altogether.

Without single signon, this leads to the authentication prompt and, at the moment, a failure to download. Immediately following this with another attempt leads to a successful download.

With single signon, this process should work.

Popups

The possible solutions to popups are related to that for binary downloads. The portlet specification states (PLT.15.4) that “data stored in the `HttpSession` by servlets or JSPs is accessible to portlets”. It appears it should be possible to share the current portlet session with the servlet and there is some suggestion that the current tomcat distribution allows this.

5.3 Hope on the horizon

The portlet specification is being worked upon and version 2 is on its way. This appears to deal with a number of the outstanding issues as well as handling a number of other areas passed over in the first version.

However, we will have to wait for portals to implement the new specification which may take some time.

Chapter 6 Access Control

6.1 An overview

Bedework has an access control system based in large part on the WebDAV ACL specification [RFC3744] with CalDAV extensions. Access rights are inherited down the tree of folders, calendars and calendar entities and can be given to or denied from users and groups and other principals.

Setting access rights is complex and difficult, even for apparently simple cases. For that reason later versions of bedework will incorporate 'wizard' style tools to help users set access based on their intentions rather than requiring them to determine the appropriate set of ACEs.

6.2 Definitions

Principals

A principal represents an entity which might need to be denied or given access or with which we might want to interact. For example, user principals usually represent real people and group principals represent groups. Other principals might be host principals, ticket principals or resource principals.

A resource principal might be, for example, a room allowing us to invite a room to a meeting, thus booking that room.

ACE

An ACE is an Access Control Entry. This defines grant or deny access for a single principal. It consists of a principal type, the principal itself – perhaps in a modified form and the denied or granted access rights. Special principal types include *unauthenticated*, *authenticated* and *others*.

ACL

An Access Control List. This is a list of ACEs which together define the access for a given entity. Internally only the ACEs set explicitly for an entity are stored with that entity. The full ACL consists of all inherited ACEs together with any explicit ACEs.

6.3 Scheduling and freebusy access

Access control for scheduling is based on CalDAV calendar access and CalDAV scheduling.

CalDAV introduced the **read-freebusy** privilege which allows principals to read freebusy information for the resource with which it is associated.

CalDAV scheduling introduces further privileges which are associated with a principal, not with any specific resource, and are actually attached to the users inbox. The privileges are **scheduling** which includes

- **schedule request**: allow (named principals) to request meetings
- **schedule reply**: allow (named principals) to reply to meeting requests
- **schedule freebusy** - allow (named principals) to send freebusy requests

If user *douglm* sets the schedule privilege on the inbox for say user *testuser01* then user *testuser01* can send and reply to meeting requests and see *douglm*'s freebusy through the web ui and the caldav server.

NOTE – this is important, in addition to those privileges, *douglm* must set **read-freebusy** on those calendars that he wants *testuser01* to access for freebusy info.

The reasoning is that we may want to block or allow scheduling operations as a whole, but if we allow them we may want to provide a different view of our freebusy to different users.

So if I as *douglm* have calendars "*work*" and "*athome*" I might not allow freebusy access to my "*athome*" calendar to colleagues because I don't want them to see how empty my home-life is. Alternatively I might disallow access to "*work*" to my wife for the opposite reason.

Simplifying:

It is possible to simply access for most users by setting access at the **/user** level, as this will be inherited by all subfolders and calendars.

What you set depends upon what you want as an organization.

The simplest approach is to add the following access to **/user**

```
authenticated: schedule, read-freebusy
```

which should allow any authenticated user to schedule and reply to meeting requests and see anybody elses freebusy.

The administrative interface currently doesn't display **/user** in the manage calendars page but you can use a url like:

```
http://localhost:8080/caladmin/calendar/fetchForUpdate.do?calPath=/user
```

(with your host of course)

A little more open is to allow that access for others INSTEAD of authenticated - that would allow scheduling via CalDAV form anywhere.

Chapter 7 Personal Calendars

7.1 Overview.

Personal calendars allow users to carry out all the normal calendaring functions as well as providing a customized view of public events through subscription to public calendars. Currently access to their calendar is through the web clients provided with bedework. In the near future we hope CalDAV will be available giving access to Apple, Sunbird and Outlook users.

Default calendars

A new user will have a set of default calendars created. One of these is the default calendar for events, the name can be configured at deployment but the default is "calendar". In addition a set of special calendars are created, "Trash", "Inbox" and "Outbox". The inbox and outbox are used for scheduling meetings and supporting bedework's implementation of itip. The Trash calendar is used in the usual way to store deleted events before complete removal from the system.

Subscriptions and views.

The default state for a personal calendar user is to have one view with a name determined by the "defaultUserViewName" syspars setting (default "All"). This view contains one subscription to the user root collection at "/user/<account>" with the user account as the name. Only the default calendar with the name given by the "userDefaultCalendar" syspar is created.

Other special calendars, such as Trash, Deleted, Inbox etc are only created as needed.

The initial default setting is normally created at the first service api open for that user. Because all (non-special) calendars in a subscription are visible, if a user creates a new calendar it will automatically be visible in the default view. Thus the initial default state is relatively simple for users to manage and will probably be sufficient for most users.

When a user explicitly subscribes to a calendar, such as a public calendar or one shared by another user, that subscription will be automatically added to the default view. Once again, this default behavior is suitable for most users.

7.2 Scheduling Meetings

The scheduling features of bedework are almost complete; there is certainly sufficient support to carry out simple scheduling. The flow of meeting requests and responses is defined by the relevant rfc's (2446, 2447) and the CalDAV scheduling extensions. The process can be complicated and subject to delays depending upon the location of the attendees.

Calendar users and addresses

The potential attendees for a meeting may be internal to the system, that is they are bedework users, or external. This is determined by their **calendar user address (CUA)** which usually looks like an email address. The bedework system is identified by one or more email domains, for example **cal.mysite.edu** and an address in those domains is considered internal otherwise it is external.

For example, with the above domain, jim@cal.mysite.edu is internal, jim@thatsite.edu is external.

In addition, bedework supports **principals** which look something like `"/principals/resources/vcc311"`. These are generally used to handle resources and locations but user principals are also mapped on to the email form of the calendar user address.

Bedework also allows the configuration option of preserving the domain part of a **CUA**. If we are not preserving the domain then a bedework CUA of jim@cal.mysite.edu would map on to a bedework user **jim**. For single domain systems this is more convenient.

Special Calendars: Inbox and Outbox

Each user has an inbox and an outbox. These are calendars with some special characteristics. Incoming scheduling requests always go to the inbox. They may arrive there via CalDAV, through uploading meeting requests or via an email interface. Scheduling responses to **external** users will go to the outbox. They may be immediately processed and at some point removed from the outbox.

To respond to a meeting request, users select each unprocessed request in the inbox, indicate their response and then submit that response.

To initiate a meeting request, use the add meeting link, add the attendees then continue on to set the details. Meeting requests must have one or more attendees and one originator (usually the current user) who will be added as an attendee.

The inbox and outbox will be created automatically when required. The actual names are configured during the build process so may be localized.

Special Calendars: Trash

This is an attempt to provide the same feature as is available in mail clients. Deleted events are moved into the trash which can be emptied when desired. There are some problems with the use of the trash which have not yet been resolved so events will not always go into the trash.

Recurring instances must be in the same calendar as the master, they cannot be trashed. A delete will always create an exdate (or remove the rdate or override).

Deleting events we have full access to but which we don't own will also cause an immediate delete.

Special Calendars: Deleted

This is here to allow users to subscribe to a calendar to which they have only read access but still be able to 'delete' events they do not want to see. For example, a user may subscribe to a 'films' calendar and delete those they are not interested in. On deletion of such an event we add an entry to the deleted calendar which acts as a mask on retrieval.

Note that there is a fix in the stylesheets which hide the delete action if the subscription is marked unremovable. The assumption is that such subscriptions require that the events also be unremovable. These subscriptions can be used for class lists etc.

Automatic processing

There are user preferences which indicate meeting requests can be automatically processed. If time is available for the incoming request it will be accepted, otherwise it will be declined. It is also possible to indicate that acceptances will be automatically processed; a meeting will have the attendee status updated automatically when the incoming response is an acceptance.

Scheduling resources

Bedework supports simple scheduling of resources. This is enabled with a degree of automatic processing of meeting requests to special resources. For example, if a room has the principal */principals/resources/vcc311* then that principal can be added as an attendee to a meeting which is intended to be in that room. The aggregated freebusy for the attendees will be displayed which includes the free time in that room.

The meeting request will be processed and added to the room's calendar, effectively booking the room for that period.

Chapter 8 Timezones

Timezones are an important and at times awkward feature of calendaring. An approach a calendar system could adopt would be to ignore timezones completely and display events using local time while storing times as UTC. This loses information implied by the presence of a timezone so bedework attempts to maintain all timezones set and imported by the user.

There is no official registry of timezones. The closest to such a registry is the Olson database which chooses to name timezones according to their continent and nearest large city, for example America/New_York.

Note: At the moment we have had problems generating timezones from the Olson database. Instead we are using the timezones packaged with ical4j. The instructions below require updating.

8.1 System timezones

Bedework provides a set of timezones, the system timezones, which are available to all users of the system. The distributed system comes with timezones derived from the Olson database but any set of timezones could be used. The administrative application provides a means for replacing the system timezones by uploading an xml formatted data file.

Building timezone information.

First we need to convert the Olson database into a set of ics files. The data itself is available from <ftp://elsie.nci.nih.gov/pub>

The vzic program available via <http://www.dachaplin.dsl.pipex.com/vzic/> is used to convert the zone information, Download and unpack the latest source and set the appropriate variables.

We set them as follows:

```
OLSON_DIR=wherever the data was unpacked.  
PRODUCT_ID=--//BEDEWORK//NONSGML Bedework Calendar system//US  
TZID_PREFIX=
```

The TZID_Prefix can be set to a value which indicates which system the timezone originated from, e.g. “/bedework.org/”. Having built vzic (make) and run it (./vzic) a directory named zoneinfo should be built. This is copied into the bwtolls/resources directory.

A copy of this generated data is available at <http://bedework.org/downloads/data/>

The next step is to convert the data into an xml form for upload. Change directory into the bwtools project and run the following (all one line) which will create a file of xml timezone information:

```
java -cp bin/bw-tztools-3.2.jar:lib/log4j-1.2.8.jar  
org.bedework.tools.timezones.Timezones -dir  
projects/bwtools/resources/zoneinfo -f tz.xml
```

A copy of this is generated file is also available at the link above. Finally we need to replace the system timezones in the production system.

Log in to the administrative client as a super user, go to "[Upload and replace system timezones](#)", browse to the generated timezones file, then upload.

Chapter 9 Bedework and CalDAV

9.1 Summary

CalDAV (rfc4791) provides a protocol for interaction between calendar clients and servers, much like iMap provides such a protocol for email. CalDAV is built on top of WebDAV which is an HTTP based protocol. As a result, CalDAV inherits all of the advantages and disadvantages of those protocols.

What CalDAV adds to WebDAV is largely reporting but also some restrictions on the placement and handling of calendaring entities.

WebDAV is a protocol which is oriented towards document sharing. This works well enough as long as we remember the unit of information in CalDAV is not a calendar but a calendaring entity such as an event or task.

So, for example, we cannot use the PUT method to store an entire calendar consisting of many events or tasks. Nor, using GET, can we retrieve an entire calendar. Typically, to date, WebDAV sharing of calendar information has involved viewing an entire calendar as a single document which is retrieved, updated and then stored. This is not the case with CalDAV.

9.2 The Bedework implementation

Bedework, at it's core, is not file system based, but uses a database for storage, retrieval and indexing. Events are not stored as a byte for byte image of rfc2445 calendar components but are stored in a relational database as rows and columns in tables.

CalDAV is not the only method used to access the data and there exists a certain tension between the needs of the different access methods.

Bedework is intended to be ultimately a complete implementation of CalDAV. At the moment we support most of the CalDAV operations (with varying degrees of success) but it is still a work in progress. As more clients become available and more experience is gained in practical use of the protocol a practical working subset supported by all clients and servers will probably emerge.

The quickstart configuration has two CalDAV servers, a public unauthenticated server and the authenticated version used for personal calendars. As CalDAV is a WebDAV based protocol it is possible to retrieve appropriately permitted personal information via the unauthenticated server. This allows users to share their freebusy information with the world

if they so wish.

9.3 CalDAV clients

Currently the following clients support CalDAV to a greater or lesser degree.

- Evolution (for linux)
- Mozilla lightning and sunbird (all platforms)
- OSAF Chandler (all platforms)

In addition Apple has announced support to be released later this year.

9.4 Unsupported features

Some of these unsupported features reflect lack of support for some rfc features – others difficulty in providing support for CalDAV specific features. The list is also incomplete but over time will probably get shorter but more accurate.

Recurrence features

Recurrence id ranges

Recurrence id ranges take the values THISANDFUTURE or THISANDPRIOR. As yet this feature is not supported and CalDAV queries or updates using this feature will have uncertain results.

Chapter 10 Other features

10.1 Freebusy URL

Bedework supports the use of a freebusy url to provide clients access to freebusy information. Users who wish to make their freebusy information available need to set schedule-freebusy and read-freebusy access for the intended audience, for example to unauthenticated users, authenticated users or specific users or groups.

The freebusy url is currently being worked on by a CalConnect technical committee which will publish recommendations on request parameters and the form of the url so the implementation in bedework may change.

Currently it takes the following form for the caldav servers:

```
<host-and-port>/<context>/fbsvc<parameters>
```

where:

- context is one of the following in the quickstart but may differ in production:
 - ucal for the authenticated user calendar
 - cal for the public unauthenticated calendar
 - ucaldav for the authenticated CalDAV server
 - caldav for the unauthenticated CalDAV server
- parameters is a list of request parameters as defined below.

Request parameters		
Name	Value	Meaning
user	A user account e.g. testuser01	A user account known to the system
cua	A calendar user address, e.g. mailto:testuser01@mysite.edu	This is the form that is usually found for an attendee.
start	Start date yyyy-mm-dd	Start of the period of interest
end	End date yyyy-mm-dd	End of the period of interest

The start and end may be omitted in which case a default period of information around the

current time will be returned. The period is limited internally to a reasonable value.

Freebusy information is returned as a VFREEBUSY object.

While the information returned may be the same from each service there are some differences in usage. Currently, the web based services only support forms based authentication, the CalDAV services support basic and digest only.

An example url might be:

```
http://myhost.edu/ucaldav/fbsvc?start=2007-09-06&end=2007-09-16
```

10.2 Real-time scheduling

This is currently under development and unsupported by most systems. Bedework allows the definition of known hosts (or more properly domains) which support realtime scheduling. When a calendar user address (CUA) is encountered which is not within the current system the default behavior is to mail the itip request to the user.

However, if the domain is in the list of known systems, the sever will attempt realtime scheduling with that server. This allows us to display at least the freebusy information for a user and perhaps send a scheduling itip message and receive the response.