



Nesstar Server Administrator Guide

Version 3.50

This guide is meant to provide you with a working knowledge of how to publish and protect your statistical metadata and data using the Nesstar Server.

Contents:

- [Introduction](#)
 - [Hardware/Software Requirements](#)
 - [Nesstar Server Components](#)
 - [The J2EE Server](#)
 - [The Web Server](#)
 - [The Embedded Database](#)
 - [The Statistical Engine](#)
 - [Embedded Clients](#)
 - [Configuration Tool](#)
 - [User Management Tool](#)
 - [Object Browser](#)
 - [Nesstar WebView](#)
 - [Other Clients](#)
- [Server Administration](#)
 - [Starting and stopping](#)
 - [Log files](#)
 - [Windows 2000/NT error handling](#)
 - [Running your server offline](#)
- [Database Administration](#)
 - [MySQL](#)
 - [SQL Server](#)
 - [Oracle](#)
- [Publishing](#)
 - [Publishing Statistical Studies](#)
 - [Publishing Studies Using the Nesstar Publisher](#)
 - [Publishing Studies Using the Automatic Study Deployer](#)
 - [Publishing Studies Using the Object Browser](#)
 - [Organising Statistical Studies in a Tree View](#)
 - [Publishing Other Web Resources](#)
- [Security](#)
 - [Access Control](#)
 - [The Default Access Control Policy](#)
 - [Managing User Accounts](#)
 - [Customising the Access Control Policy](#)
 - [Enabling SSL \(Secure Socket Layer\)](#)

Introduction

Hardware/Software Requirements

- Windows XP or Windows Server 2003 or Windows 2000 or NT4 (if you have problems with the installation under NT4 check that your video drivers are up to date) all with the most recent Service Pack updates.
- 512M of RAM (1G is preferable). You may optimise the amount of memory your server is allocated by setting the "Java Heap Size" under Advanced Settings in the configuration tool.
- 1Ghz CPU or better. A dual-processor machine can be useful, especially during upgrading when the new and the old server might be running side-by-side.

Nesstar Server Components

The major components of the Nesstar Server are:

The J2EE Server

The Nesstar Server is based on [JBoss](#). JBoss is an Open Source application server that implements the [Java 2 Enterprise Edition](#) specification from SUN Microsystems.

The Nesstar Server code is mostly written in Java. The server comes with an embedded copy of the [Sun's Java Virtual Machine](#).

The Nesstar Server is basically a container of objects that corresponds to statistical metadata and data concepts (studies, variables, etc.). This objects are defined in an object-oriented model, expressed as an UML Class Diagram. A copy of the model (in Poseidon/Argo UML or XMI format) is available on request.

The Web Server

The Nesstar server contains a standard web server ([Tomcat](#)). The server has a default [Home Page](#), which is the point of entry to new users. The **port** and **host** are specified during configuration (immediately after installation) and are stored in **./config/nesstar.properties** (If the standard port 80 is used, then the port can be omitted from the URL). This web server supports secure pages (see [secure sockets layer](#) for details).

The homepage provides links to the embedded clients (Nesstar WebView and the Object Browser) and to useful documents, such as this guide. It can also be customized to

provide links to your own web content. See the section on [Publishing Other Web Resources](#) for more details. Further information can be found in the [Configuration guide](#).

The Embedded Database

The Nesstar Server comes pre-packaged and pre-configured with a version of the [MySQL](#) database server. The database is used to store the persistent state of the Nesstar objects and to provide efficient search and retrieval functionality. This can be disabled in favour of another database in the Configuration Tool. Further information can be found in the [Database Administration](#) section and also in the [Configuration guide](#).

The Statistical Engine

The Nesstar Server contains a version of the [NSDstat](#) statistical data engine and converter. The Statistical Engine provides all the statistical processing functionality of the server, e.g.: online variable tabulation, descriptive statistics, dynamic data subsetting and merging across multiple files.

Embedded Clients

Configuration Tool

The [Configuration Tool](#) allows the Administrator to edit most of the server configuration properties and to perform basic administrative operations as starting/stopping the server.

User Management Tool

The [User Management Tool](#) allows you to create and manage user accounts.

Object Browser

The Object Browser is a cut-down interface which allows you to inspect and manipulate the objects contained in any Nesstar Server through a normal Web Browser.

It can be accessed [with](#) or [without](#) frames and is useful to server administrators to perform certain tasks that are not available through the Configuration Tool or other clients (e.g. deleting a Dataset remotely). All tasks performed in the object browser are subject to the same access control policies in place on the server.

The Object Browser can inspect and manipulate any Nesstar object on any Nesstar Server, subject to the access control policy.

The parameters are:

Object Browser Parameters	
Parameter	Function
url	the URL of the object to browse, if a local reference is given (ex: <i>/obj/server</i>) it is interpreted as relative to the URL of the server where the Object Browser is hosted.
action and FRAMES	With FRAMES=false&action=LIST the Object Browser will produce a plain HTML interface. With FRAMES=true it will produce a frame interface with the left frame showing a tree of the browsed objects with their properties and methods.

Nesstar WebView

The [Nesstar WebView](#) provides a convenient end-user, web-based, interface to the server contents.

The basic Nesstar WebView look and feel can be customised using the [Configuration Tool](#).

More extensive customisation can be performed by directly modifying the Java Server Pages (JSP) sources held in the directory
./jboss/server/default/deploy/webview.war/.

Other Clients

The [Web Console](#) provides convenient access to all web-based interfaces currently installed in your server, typically:

admin	security User Management Tool module
console	Web Console module
jmx-console	Java Bean Console module
root	Server Root pages module
web-console	JBoss™ Application Server
webview	Nesstar WebView interface module

Server Administration

Note: Unless stated otherwise all file paths are relative to the server installation (root) directory, e.g. **Nesstar-Server-3.xx**.

The main configuration settings are defined in the file **./config/nesstar.properties**. This serves as a single point of entry to the most essential properties of your server. Further configuration files can be found in the **./config/instdata** directory. To configure your server, run the batch file **configure** in the server **./bin** directory (or follow the shortcut on the start menu in Windows). This will launch the Nesstar Server Configuration Tool. After having changed your properties, clicking on **[Apply settings]** will install the new properties and re-start the server. You can also perform other tasks, such as defining your security settings and restarting the server, from the configuration tool.

See the [Configuration Tool Guide](#) for more information.

Please note that:

- re-installation does not affect any data that was previously published on your server;
- uninstallation (**./bin/uninstall.exe**) will not remove data that was previously published on your server (Windows only);
- Once installed, you cannot change the name of your server. Please ensure that you assign the server with the desired name before you complete the installation.

Starting and stopping

If you chose to install short cuts as default during installation (Windows only) you can use these to start and stop the Nesstar server. They are usually located in the server installation directory and the start menu. You can also use the Configuration Tool to start and stop your Server and launch the server welcome page.

If you installed your server (and/or MySQL) as an NT Service (Windows only), these can be started and stopped via

- **"Start/Settings/Control Panel/Services"** for Windows NT;
- **"Start/Settings/Control Panel/Administrative Tools/Computer Management/Services and Applications"** for Windows 2000/XP;
- **"Start/Administrative Tools/Services"** for Windows XP.

Please note that:

- you must be a system administrator to install NT services;
- you may need to specify the account logon details in the Nesstar Server service Properties when

connecting to an external database. To do this, right-click the service, select Properties, Log On. Specify **"This account"** instead of the default **"Local System account"**.

Alternatively, you can use the batch file **run** or **launch** in the server **./bin** directory to start your server in a command window. To stop your server either press **control-c** in the active command window (or close the command window) or use the batch file **stop** in the server **./bin** directory.

If you switch from command line to NT service - or vice-versa - for either the Server or MySQL, it is advised to run the **./bin/cleanup** batch file first to make sure that existing services or demons are properly shut down.

It is also possible to re-start the server remotely using the Object Browser Server Methods.

Log files

The server maintains a number of log files in the **./jboss/server/default/log** directory. Note that these can be accessed remotely through the Object Browser using the `Server.GetFile` method.

File name	Contents of file	Can be configured in
service.log	Lists the date and time of each start and stop of the server	N/A
boot.log	Logs only startup information from the jboss application server	N/A
server.log	Contains information on all server activity including any error messages	./jboss/server/default/conf/log4j.xml
error.log	Similar to server.log but only logs	./jboss/server/default/conf/log4j.xml

	nesstar error messages	
keepalive.log	Similar to server.log but only logs nesstar keep alive messages	./jboss/server/default/conf/log4j.xml
converter.log	Contains information on parsing your extension code	N/A
parser.log	Contains information on parsing your security policy	N/A
StatEngine.log	Logs the results of statistical data analysis only	N/A
NsdstatConverter.log	Logs the results of statistical data downloads only	N/A
CubeEngine.log	Logs the results of statistical cube access only	N/A
localhost_access YYYY-MM-DD.log	Logs web engine access	./jboss/server/default/deploy/jbossweb-tomcat.sar/server.xml

Windows 2000/NT error handling

To perform its functions the server uses some win32 specific software. Sometimes one of these modules might throw an error (this will not affect the server itself which will keep on running) and by default an error window (a "Dr.Watson" window) will appear on the screen. By following the procedure below you can change this behaviour so that whenever a module crashes you will hear a 'beep' but no window will be displayed.

- Open an MSDOS Console
- Run: `drwt$sn32`
- Uncheck 'Visual notification'
- Exit
- Run: `drwt$sn32 -i` to make the new settings effective.

Running your server offline

If you wish to disable external access to your server, you can switch it to work in offline mode, which means other machines on the network/internet cannot access it. To do this, open the configuration tool, switch to Advanced mode. Change the 'Work Offline' combo box to 'true'.

Database Administration

For database administration see also the [configuration guide](#).

MySQL

Advanced Settings

The following settings can be altered by going to "Advanced Settings" in the config tool and scrolling down to the MySQL section:

- **MySQL Stopwords File**
The default location of the MySQL stopword file is `./mysql/stopwords`. You can specify an alternative path or edit the existing one. This stores all words that MySQL will treat as insignificant in your search for full-text searches.
- **MySQL Key Buffer Size**
The default key buffer size is 64 MBytes. This is the size of the buffer used for index blocks in MySQL. You may increase this if you experience slow searching speeds, or decrease if you need to free up the memory on your computer.
- **MySQL Table Cache Limit**
The default table cache limit is 256. This is the maximum number of open file descriptors that MySQL will keep in memory at any one time. Alter this if you need to increase the size to accommodate other database applications hosted in the same MySQL server.
- **MySQL Max Packet Size**
The default max packet size is 100 MBytes. This is the maximum size of one packet of information in MySQL for both incoming and outgoing packets. Increase this if you have a significantly large amount of text data in one field.

Administration/Browsing

The recommended MySQL clients are MySQL Administrator and MySQL Query Browser graphical clients (download from [MySQL Home](#)). The alternative is to use the command line interface.

Using MySQL Administrator/Query Browser

When setting up the graphical MySQL clients, you will need to point it to the hostname of the machine on which the server was installed. The database name is nesstar, the username is root, and the password is the administrator password entered during installation.

Using Command line interface

To use the command line interface, locate the mysql.bat file in the mysql directory in your server root. Executing this will log you into mysql. Here are some tips for using mysql:

- Before doing operations on the Nesstar database, you must switch to it using 'use nesstar;' (or your JDBC Database name).
- List the tables in the nesstar database using 'show tables;'.
The command line will accept SQL statements terminated with a semi-colon (i.e. select * from studyejb where id = 'uk.ac.data-archive.ddi.4213';).

SqlServer

Here are some general notes on administration of SQL Server, learned through experience:

- when executing sql scripts directly on the database, log in to the query analyser using the same credentials as the server will be using - otherwise, unless the create statements directly specify the username for whom to create the object, it will be created for user 'dbo', which may mean it is not available to the nesstar user
- remember to allow SQL Server authentication under the Security tab in the Properties of the nesstar database.
- make sure the nesstar user is setup to use the correct nesstar database - not some alternate database
- when creating databases - if you plan to use any resources in other databases (i.e. views which union a local table and a remote table) make sure when creating the database that the collation types match
- when creating a database, its always a good idea to schedule log shrinkages.

Oracle 9i/10g

Oracle Administration

For Oracle, note that in the Jdbc Url, **nesstar** refers to the Oracle 8.x **Service Name** (or **SID** in previous versions of Oracle) , e.g.

- jdbc:oracle:thin:@//**localhost**:1521/**nesstar**

Launch the **Enterprise Manager Console** and connect to the new database you created as the **Database Name**. Then inside the database tree create a new user as the **Jdbc User Name** with standard database administration privileges. Note that Oracle automatically creates a database **Schema** name which defaults to the **Jdbc User Name** during the database connection.

Technical issues with Oracle

- for multi-byte language support please make sure that your databases default Character Set is set up to use Unicode (AL32UTF8).
- the data type **VARCHAR** is restricted to 4000 bytes only.
- indexes cannot be created on varchars exceeding 3209 characters.

Publishing

You can use the Nesstar Server:

- To publish statistical studies
- To organise statistical studies in a tree view
- To publish additional web resources such as HTML pages, PDF documents, etc.

Publishing Statistical Studies

Studies published in the same server must have different identification numbers (IDs). The study ID must be specified in the **ID** attribute of the **<codeBook>** element in the DDI description file, e.g.:

```
<codeBook ID="uk.ac.data-archive.ddi.2568">
```

It is a good idea to make sure that this ID is not only unique in your server but it is also globally unique. Just like ISBN numbers make it possible to search for the same book

in different libraries (as well as to distinguish between different books that might have been categorised under the same local ID in different libraries) the globally unique study IDs will be used in the future to correctly identify studies published on any Nesstar server around the world.

There are a number of ways of producing a globally unique identifier for studies. The convention we advise you to adopt is:

`publisher_internet_domain_in_reverse_order ".ddi." local_identifier`

So for example the study published by the UK Data Archive, whose Internet domain is *data-archive.ac.uk*, and internally catalogued as study number 2568 would have the ID: *uk.ac.data-archive.ddi.2568*.

The server comes packaged with a free study named **uk.ac.data-archive.ddi.2568** in **./data/sample** courtesy of the [ESDS](#). You may use this study to test your server. You are allowed to publish it as long as you do not alter it in any way.

Publishing Studies Using the Nesstar Publisher

The [Nesstar Publisher](#) is a specialised client that allows data publishers to create, edit and check metadata, import data from different popular statistical data formats and easily upload data and metadata to Nesstar Servers. See the [Nesstar Publisher Help](#) for more information.

Note: Publishing on Local Networks (Windows only)

If the server you are publishing to is accessed via a mapped network drive and its Data Directory is not inside the **Server Home** directory then the **Network Path To Data Directory** must be specified using its complete network path during configuration. E.g. to publish to server **\\RemotePC01\Nesstar-Server** with its data stored in **\\RemotePC01\Microdata**, you must specify **\\RemotePC01\Microdata** as the Network Path To Data Directory during configuration.

If the Data Directory is remote to the server itself, e.g. **\\RemotePC02\Microdata**, you will also need to change the nt service logon so that the [service](#) can access the remote drive.

Publishing Studies Using the Automatic Study Deployer

If you already have study metadata files in DDI format and data files in NSDStat format you can publish them by simply copying them into the directory **data**. The deployer will automatically publish the study creating a Study object and

all the other objects that are described in the DDI file (Variables, etc.). Once published the data files will be automatically renamed according to their fileName ID's and the metadata according to its codeBook ID.

To remove a study simply delete its DDI file from the **data** directory. The study object (together with all the other objects described in the DDI file) will be automatically removed from the server. All the study data files will be automatically removed from the data directory as well.

Publishing Studies Using the Object Browser

To publish a study you can use the [StudyHome's addStudy\(File ddi\)](#) method.

If you want the created Study to be inserted in a catalogue, point the Object Browser to the catalogue (for example: the [root catalogue](#)) and use the **AddDataset(File file)** method.

To publish the data associated with the study, you must point the Object Browser to the newly created Study (for example the [sample study](#)), click on **files** then the individual **Datafile** you want to add and then use **AddDataFile(File NSDstat)** in order to add one data file at a time.

To remove a study use the Study's **Delete** method.

Organising Statistical Studies in a Tree View

In order to categorise and to make searchable the contents of your server, you will need to create one or more searchable folders (currently represented by the Catalogues objects) and insert your studies (or any other server object, including other folders) into them.

The folder/catalogue is the basic unit of organisation provided by the Nesstar server. A folder can for example be used to hold studies of a similar nature, or on similar topics. Their role is similar to that of the file directories on your PC. Nesstar clients allow users to browse the contents of a folder (open the folder to display its contents) as well as to search its contents.

When the server is first installed a single root folder is created, whose label (its short human readable name) is the same as the server name.

You can use the [Nesstar Publisher](#) to create additional folders, organise them in a hierarchy and link your studies in the most appropriate folder.

Publishing Other Web Resources

As the Nesstar Server embeds a standard web server ([Tomcat](#)) you can use it to publish any kind of static or dynamic WWW resources, including HTML, images and other static files, and Java Server Pages.

The root directory of the web server is **`./jboss/server/default/deploy/root.war`**. Any files placed in this directory or one of its subdirectories can be accessed directly over the web using a normal web browser. For example this guide corresponds to the file [`./jboss/server/default/deploy/root.war/nesstar/doc/admin_guide.jsp`](#)

Security

Access Control

In order to protect the data (and the servers themselves) from unauthorised access, the server includes an Access Control Unit (ACU). The ACU enforces the access control policy defined in the server Access Control Policy file. The policy specifies the access control conditions that apply to each resource. Complex conditions can be defined relative to the kind of user, the purpose in using the data (commercial, research, etc.), the kind of operation required, and the legal conditions accepted by the users. On the base of this information the ACU can protect the resources and also "drive" the user in the process of acquiring the necessary rights (for example the user might be required to login, to agree to some conditions, etc.). The ACU can be customized by defining additional site-specific conditions and/or connections to the site user database.

The Default Access Control Policy

The Access Control Policy file installed by default defines a few basic types of users with the following rights:

- anonymous: can browse and search the studies metadata
- authorised and guest: can also perform statistical operations
- fully authorised: can also download and subset data
- publisher: can also publish data and metadata
- administrator: can perform any operation

Managing User Accounts

You can create, delete or edit users using the [User Management Tool](#).

Customising the Access Control Policy

The Nesstar ACU can be customised to implement different access control policies. For more information refer to the [Nesstar ACU Customisation Guide](#).

Enabling SSL (Secure Socket Layer)

The default installation provides only a non-secure connection between the client and server. If you wish to provide a secure encrypted communication route (**https**), you must enable the secure sockets protocol. To do this simply change the value of the property **Secure Sockets** in **./config/nesstar.properties** to **true**. This will grant both secure and non-secure access. You may also choose to completely disable non-secure access by setting the property **Open Sockets** to **false**. Then, next time you run **install** the script will automatically

1. modify your **./jboss/server/default/deploy/jbossweb-tomcat.sar/META-INF/jboss-service.xml** configuration file
2. create your server certificate using the java keytool
3. export the server certificate using the java keytool to **./config/nesstar.cer**
4. import the server certificate into your java keystore using the keytool.

Whenever you wish to create a new certificate, you must first remove any existing certificate from the **config** directory. The default secure connection will be through **https://<host>:<ssl_port>**, where **ssl_port** is defined in **nesstar.properties**.

You must then provide Nesstar Publisher/Explorer users with this certificate, which they must import into their own java key store. To assist the user, a simple batch file is provided for this purpose in the batch file **./bin/importcert**.

The previous steps added the apache server certificate to the list of trusted certificates in the application's keystore. These are the minimum requirements, however, clients must always accept this certificate each time they establish a secure connection unless your certificate has been signed by a Certificate Authority. The following is a list of Trusted Certificate Authorities.

- [Thawte Inc.](#)
- [VeriSign Inc.](#)
- [RSA Data Security Inc.](#)

See also the [Tomcat SSL Configuration](#) guide for further information.

If you have any suggestions or questions please contact: Nesstar Support

Copyright © 2001-2006 NSD