

ActuateOne™

One Design
One Server
One User Experience

Configuring BIRT iHub

Information in this document is subject to change without notice. Examples provided are fictitious. No part of this document may be reproduced or transmitted in any form, or by any means, electronic or mechanical, for any purpose, in whole or in part, without the express written permission of Actuate Corporation.

© 1995 - 2013 by Actuate Corporation. All rights reserved. Printed in the United States of America.

Contains information proprietary to:

Actuate Corporation, 951 Mariners Island Boulevard, San Mateo, CA 94404

www.actuate.com

The software described in this manual is provided by Actuate Corporation under an Actuate License agreement. The software may be used only in accordance with the terms of the agreement. Actuate software products are protected by U.S. and International patents and patents pending. For a current list of patents, please see <http://www.actuate.com/patents>.

Actuate Corporation trademarks and registered trademarks include:

Actuate, ActuateOne, the Actuate logo, Archived Data Analytics, BIRT, BIRT 360, BIRT Analytics, The BIRT Company, BIRT Data Analyzer, BIRT iHub, BIRT Performance Analytics, Collaborative Reporting Architecture, e.Analysis, e.Report, e.Reporting, e.Spreadsheet, Encyclopedia, Interactive Viewing, OnPerformance, The people behind BIRT, Performancesoft, Performancesoft Track, Performancesoft Views, Report Encyclopedia, Reportlet, X2BIRT, and XML reports.

Actuate products may contain third-party products or technologies. Third-party trademarks or registered trademarks of their respective owners, companies, or organizations include:

Mark Adler and Jean-loup Gailly (www.zlib.net): zLib. Adobe Systems Incorporated: Flash Player. Amazon Web Services, Incorporated: Amazon Web Services SDK, licensed under the Apache Public License (APL). Apache Software Foundation (www.apache.org): Ant, Axis, Axis2, Batik, Batik SVG library, Commons Command Line Interface (CLI), Commons Codec, Crimson, Derby, Hive driver for Hadoop, Pluto, Portals, Shindig, Struts, Tomcat, Xalan, Xerces, Xerces2 Java Parser, and Xerces-C++ XML Parser. Castor (www.castor.org), ExoLab Project (www.exolab.org), and Intalio, Inc. (www.intalio.org): Castor. Day Management AG: Content Repository for Java. Eclipse Foundation, Inc. (www.eclipse.org): Babel, Data Tools Platform (DTP) ODA, Eclipse SDK, Graphics Editor Framework (GEF), Eclipse Modeling Framework (EMF), and Eclipse Web Tools Platform (WTP), licensed under the Eclipse Public License (EPL). Gargoyle Software Inc.: HtmlUnit, licensed under Apache License Version 2.0. GNU Project: GNU Regular Expression, licensed under the GNU Lesser General Public License (LGPLv3). HighSlide: HighCharts. Jason Hsueh and Kenton Varda (code.google.com): Protocol Buffer. IDAutomation.com, Inc.: IDAutomation. IDRsolutions Ltd.: JBIG2, licensed under the BSD license. InfoSoft Global (P) Ltd.: FusionCharts, FusionMaps, FusionWidgets, PowerCharts. Matt Inger (sourceforge.net): Ant-Contrib, licensed under Apache License Version 2.0. Matt Ingenthron, Eric D. Lambert, and Dustin Sallings (code.google.com): Spymemcached, licensed under the MIT OSI License. International Components for Unicode (ICU): ICU library. jQuery: jQuery, licensed under the MIT License. Yuri Kanivets (code.google.com): Android Wheel gadget, licensed under the Apache Public License (APL). LEAD Technologies, Inc.: LEADTOOLS. The Legion of the Bouncy Castle: Bouncy Castle Crypto APIs. Bruno Lowagie and Paulo Soares: iText, licensed under the Mozilla Public License (MPL). Microsoft Corporation (Microsoft Developer Network): CompoundDocument Library. Mozilla: Mozilla XML Parser, licensed under the Mozilla Public License (MPL). MySQL Americas, Inc.: MySQL Connector. Netscape Communications Corporation, Inc.: Rhino, licensed under the Netscape Public License (NPL). OOPS Consultancy: XMLTask, licensed under the Apache License, Version 2.0. Oracle Corporation: Berkeley DB, Java Advanced Imaging, JAXB, JDK, Jstl. PostgreSQL Global Development Group: pgAdmin, PostgreSQL, PostgreSQL JDBC driver. Progress Software Corporation: DataDirect Connect XE for JDBC Salesforce, DataDirect JDBC, DataDirect ODBC. Rogue Wave Software, Inc.: Rogue Wave Library SourcePro Core, tools.h++. Sam Stephenson (prototype.conio.net): prototype.js, licensed under the MIT license. Sencha Inc.: Ext JS, Sencha Touch. ThimbleWare, Inc.: JMemcached, licensed under the Apache Public License (APL). World Wide Web Consortium (W3C) (MIT, ERCIM, Keio): Flute, JTIty, Simple API for CSS. XFree86 Project, Inc.: (www.xfree86.org): xvfb. ZXing authors (code.google.com): ZXing, licensed under the Apache Public License (APL).

All other brand or product names are trademarks or registered trademarks of their respective owners, companies, or organizations.

Document No. 130131-2-530303 January 28, 2013

Contents

About Configuring BIRT iHub	ix
--	-----------

Chapter 1

Performing basic configuration tasks	1
---	----------

Understanding Configuration Console	2
---	---

Launching Configuration Console	2
---------------------------------------	---

Troubleshooting problems launching the console	3
--	---

Using Simple view	4
-------------------------	---

Displaying and updating the license	4
---	---

Stopping and starting iHub System	6
---	---

Configuring diagnostic logging	7
--------------------------------------	---

Understanding diagnostic logging	7
--	---

Configuring a diagnostic logging categories	11
---	----

Changing the Configuration Console password	13
---	----

Configuring network settings	14
------------------------------------	----

Configuring e-mail notification settings	14
--	----

Configuring the URL for e-mail notification	16
---	----

Using Advanced view	17
---------------------------	----

Setting Configuration Console options	19
---	----

Changing server, template, volume, or partition options	20
---	----

About the servers list	22
------------------------------	----

About the server templates list	22
---------------------------------------	----

About the volumes list	23
------------------------------	----

About the partitions list	23
---------------------------------	----

About the printers list	23
-------------------------------	----

Viewing release information	23
-----------------------------------	----

Chapter 2

Configuring an Encyclopedia volume	25
---	-----------

Understanding an Encyclopedia volume	26
--	----

Configuring a partition	27
-------------------------------	----

Adding a partition	27
--------------------------	----

Viewing partitions	29
--------------------------	----

Working with databases, schemas, and volumes	29
--	----

Configuring a metadata database	30
---------------------------------------	----

Configuring a schema	33
----------------------------	----

Configuring an Encyclopedia volume	35
--	----

Adding an Encyclopedia volume	35
-------------------------------------	----

Editing Encyclopedia volume properties	38
--	----

Configuring an Encyclopedia volume partition	40
Supporting application-level partitioning	44
Specifying a hyperlink in a PDF document	44
Configuring events	46
Configuring advanced volume properties	48
Printing a summary of advanced volume properties	48
Retrying failed asynchronous jobs	49
Changing the status of an Encyclopedia volume	49
Taking an Encyclopedia volume offline	50
Disabling an Encyclopedia volume	50
Enabling an Encyclopedia volume	51
Taking an Encyclopedia volume online	51
Removing an Encyclopedia volume	52
Renaming the default Encyclopedia volume	53
Removing a partition	54

Chapter 3

Using diagnostic, usage, and error logging 55

Configuring diagnostic logging	56
Changing the default log file location in acserverconfig.xml	58
Configuring additional iHub diagnostic logging properties	59
About iHub diagnostic logging	59
About logging for Encyclopedia database queries	60
Configuring usage and error logging	62
Configuring usage logging	63
About types of recorded events	65
About the usage logging extension	65
Understanding a usage log entry	66
Configuring error logging	69
Understanding an error log entry	70
Configuring usage and error logging file settings	72

Chapter 4

Configuring e-mail notification 75

About e-mail notification	76
Adding an SMTP server to the iHub environment	76
Using SMTP server load balancing	79
Using multiple SMTP servers in a cluster	79
Specifying the Message Distribution service for e-mail notification	80
Setting up sendmail e-mail notification	81
Configuring the notification list size and To: line	81
Handling e-mail notification errors	83
Handling SMTP e-mail errors	83
Handling sendmail e-mail errors	83

Logging e-mail notification errors	84
Customizing the e-mail message	84
Sending e-mail notification in a cluster	84
Sending e-mail notices from nodes using a template	85
Using the e-mail message template	86
Working with e-mail template elements	87
Using variable attributes	88
Using HTML in the e-mail template	89

Chapter 5

Working with services 91

About BIRT iHub and PostgreSQL services	92
Using iHub services	92
Configuring service properties	94

Chapter 6

Configuring the View service 97

Configuring View service settings	98
About setting View service properties in iHub	99
Enabling the View service	99
About diagnostic logging	99
Configuring the message distribution weight for a node	99
Configuring the View service for BIRT documents	101
Configuring Java process communication	101
Configuring message time-out	102
Configuring the thread pool	103
Configuring BIRT document and design caching	104

Chapter 7

Configuring the Factory service 113

About the Factory service	114
About setting Factory service properties in iHub	114
Enabling the Factory service	115
About diagnostic logging	115
Configuring the Factory service for general use	115
Configuring the transient document cache	116
Setting the transient document cache location	116
Setting the transient document cache size	116
Setting the transient document time-out	116
Configuring the number of entries in the transient document cache	117
Configuring the Factory service for synchronous jobs	118
Configuring the maximum run time	118
Configuring the queue size	119
Configuring the queue time-out	119

Configuring the message distribution weight of a node	120
Configuring the Factory service for BIRT documents	121
Configuring the Factory service base port and range of ports	122
Recycling Java Factories	123
Configuring BIRT caching	124
Configuring the data result set buffer	124
Configuring the BIRT design cache	126
Configuring the BIRT document cache	127
Configuring the maximum rows in a BIRT chart	128

Chapter 8

Using resource groups	131
Introducing resource groups	132
Configuring a resource group	133
Allocating Factories for a resource group	134
Limiting Java Factories for a resource group	134
Setting start arguments for the JRE	136
Adding a resource group	137
Using a resource group	140
Selecting a resource group for a job	140
Prioritizing a job	141
About Factory and View service requirements	141
Managing loads with resource groups	142
Understanding the Java View service	142
Stopping a resource group	142
Deleting a resource group	143
Using resource groups programmatically	144

Chapter 9

Clustering	145
About a cluster configuration	146
About distributing services to a cluster	147
About the configuration home directory	149
About the primary configuration files	150
About acpmdconfig.xml	150
About acserverconfig.xml	152
Creating a cluster	153
Creating an initial cluster	154
Configuring heartbeat messaging	156
Configuring Message Distribution service properties	157
Adding and modifying server templates	160
Modifying a server template	160
Creating a new server template	160

Adding a node to a cluster	161
About node configuration	161
About cluster configuration	162
Testing a cluster	162
Starting and stopping a node	162
About starting a node	163
About stopping a cluster or a node	165
Removing a node from a cluster	165
Managing a cluster	166
Accessing partitions	166
Specifying the Encyclopedia location	167
About cluster option requirements	167
Handling file system failure	168
Configuring the cluster administrator e-mail account	168
Managing console configurations and load balancing	169
Using the consoles directly	169
Using the consoles through a firewall	169
Using multiple console installations	170
About load balancing	170
Deploying load balancing	171
 Chapter 10	
Configuring Integration and Caching services	173
About the Integration service	174
Managing Integration service resources	174
Enabling the Integration service	174
About diagnostic logging	175
Setting port numbers for process communication	175
Specifying load settings	176
Managing Actuate SQL query execution	176
Specifying query settings	178
Specifying query optimization settings	179
Setting resource management properties	180
Setting query statistics logging properties	181
Using information objects	182
Setting up caching	182
About setup in Actuate Information Object Designer	182
About setup in iHub	183
Updating cache files	183
About Actuate Caching service	183
Configuring the Caching service	185
Enabling the Caching service	186
About diagnostic logging	187

Configuring process management properties	187
Configuring process parameters	187
Configuring communication	188
Configuring requests	188
Configuring bulk load	189
About bulk loading files to the cache	190
Configuring Actuate Caching service and NetOSI file type	191
 Chapter 11	
Configuring iHub security	193
Understanding the Report Server Security Extension	194
Working with RSSE	194
About iHub and RSSE application interaction	195
Installing iHub using open security	196
Installing iHub with the LDAP option	196
Installing the RSSE web service application	201
Using the RSSE application with a service provider	201
Installing an RSSE application on Tomcat	201
Configuring the Encyclopedia volume to run RSSE	202
Configuring open security	202
Specifying RSSE service start-up parameters	204
Understanding LDAP configuration	205
Mapping Encyclopedia volume management information to LDAP objects	206
Mapping Encyclopedia volume user properties	206
Mapping roles	207
Mapping groups	207
Mapping channels	207
Mapping pass-through security information	208
Setting ldapconfig_<volume>.xml parameters	209
Understanding an LDAP directory structure	216
Mapping LDAP objects to users	217
Converting an Encyclopedia volume to use an RSSE application	218
Converting internal IDs to external names	219
Converting information from external to internal	219
Caching external security information	220
Working with RSSE page-level security	220
About the RSSE page-level security example	220
Creating an access control list	221
 Chapter 12	
Archiving files	223
Understanding online archiving	224
About the online archive driver	224

About the online archive driver files	224
Retaining file attributes during archiving	224
Preserving file dependencies during archiving	225
Using the online archive driver	225
Specifying the online archive configuration file	226
Modifying the online archive configuration file	226
Creating and specifying the startup script	229
Managing file purging	231
Locating the archived files	232

Chapter 13

Printing documents 233

Understanding printing	234
Printing on Windows	234
Printing on Linux	235
Adding a printer on Linux	236
About Xprinter environment variables	236
Obtaining the PPD file	236
Sending a document to the printer	237
Managing the printing environment	238
Changing a path to a printer	238
Removing a printer	239
Configuring advanced printing parameters	239
Disabling automatic printer configuration	239
Configuring fonts	240
Understanding PostScript font conversion issues	241
Mapping fonts for charts in documents	241
Locating fonts when generating BIRT documents	242

Chapter 14

Connecting to data sources 245

About data source connections	246
About drivers	246
About information object connections	246
Using a connection configuration file	246
Changing a configuration file	247
Specifying the location of the connection configuration file	248
Configuring a cluster to use a connection configuration file	248
Defining environment variables	249
Connecting to data sources	250
Connecting to a DB2 data source	250
Defining DB2 environment variables	251
Using DB2 libraries on HP-UX	251

Checking a connection to a DB2 instance	252
About using XML Extender	252
Connecting to an Informix data source	252
Connecting to an Oracle data source	253
Defining Oracle environment variables	253
About the Oracle listener processes	254
Setting the maximum column length	254
Connecting to custom data sources	254
Installing a custom open data access (ODA) driver	254
Using custom Java-based data source connections	256
Using custom relational data sources with the Integration service	256
Specifying connection types	257
Using a connection pool	258
ODBC/JDBC connection pooling	259
Configuring BIRT JDBC connection pooling	259
Setting the Encyclopedia engine connection pool reap interval	262
 Chapter 15	
Setting miscellaneous properties	263
Changing locale, encoding, and time zone	264
Modifying general volume properties	265
Changing ports used by iHub	266
Viewing and modifying general server properties	268
About application container process	270
About SOAP settings	271
Changing iHub system start-up parameters	272
Configuring general system properties	274
About General	274
Setting start-up arguments for the Encyclopedia server JVM	276
Starting and stopping iHub	277
Stopping iHub	277
Starting iHub	278
Setting conversion queue and e-mail queue resource management properties	278
Configuring dates	280
Using Printable Summary to view system properties	281
Using Printable Summary to view template properties	282
Setting Filetype driver properties	283
 Index	285

About Configuring BIRT iHub

Configuring BIRT iHub discusses how to set up BIRT iHub and an Encyclopedia volume. The chapters in this guide are:

- *About Configuring BIRT iHub.* This chapter provides an overview of this guide.
- *Chapter 1. Performing basic configuration tasks.* This chapter discusses how to perform basic tasks such as logging in.
- *Chapter 2. Configuring an Encyclopedia volume.* This chapter covers how to add, remove, and back up an Encyclopedia volume.
- *Chapter 3. Using diagnostic, usage, and error logging.* This chapter describes how to configure and consolidate logs.
- *Chapter 4. Configuring e-mail notification.* This chapter describes how to configure e-mail notification about the completion of iHub jobs.
- *Chapter 5. Working with services.* This chapter discusses iHub architecture and services.
- *Chapter 6. Configuring the View service.* This chapter discusses how to test and improve performance iHub viewing.
- *Chapter 7. Configuring the Factory service.* This chapter describes how to tune the Factory service and control Factory processes using resource groups.
- *Chapter 8. Using resource groups.* This chapter describes how to control Factory services using resource groups.
- *Chapter 9. Clustering.* This chapter describes how to create and manage a cluster, add and delete a node, and modify server templates.
- *Chapter 10. Configuring Integration and Caching services.* This chapter describes how to optimize performance of information object-based reporting using Actuate Caching service.

- *Chapter 11. Configuring iHub security.* This chapter covers the Report Server Security Extension.
- *Chapter 12. Archiving files.* This chapter describes online archiving and job completion notice purging.
- *Chapter 13. Printing documents.* This chapter describes setting up printing from iHub and customizing fonts.
- *Chapter 14. Connecting to data sources.* This chapter describes how to connect iHub to data sources, such as Oracle and DB2.
- *Chapter 15. Setting miscellaneous properties.* This chapter describes how to set iHub properties such as locales, ports, and process communication settings.

1

Performing basic configuration tasks

This chapter contains the following topics:

- Understanding Configuration Console
- Using Simple view
- Using Advanced view

Understanding Configuration Console

Configuration Console provides the administrator with a convenient graphical user interface (GUI) for configuring the iHub system. Configuration Console provides two configuration perspectives:

- **Simple view**
Displays immediately after logging in to Configuration Console. In Simple view, the administrator can configure the basic functionality required to get iHub running quickly, such as starting or stopping iHub, updating a license, and specifying diagnostic logging, network, and e-mail notification settings.
- **Advanced view**
Accessed from Simple view. In Advanced view, the administrator has access to all iHub settings available in Configuration Console, such as options for system, server, template, volume, partition, resource group, and printer settings. The administrator can perform any task supported in Simple view in Advanced view. Some settings, such as advanced template settings, may not be available through the console. These settings require manually updating XML configuration files to fine-tune iHub.

Launching Configuration Console

The administrator accesses Configuration Console using a browser.

How to launch Configuration Console

- 1 In Windows, launch Configuration Console from the Start menu by choosing the following options:
Start→Programs→Actuate→BIRT iHub Configuration Console
Or, in a browser, launch Configuration Console on any platform by typing the following URL:
<http://localhost:8900/acadmin/config>
- 2 Log in to Configuration Console using the password specified during installation.
- 3 Accept the default language and time zone, or choose another language and time zone. Figure 1-1 shows the login page for Configuration Console.
- 4 Choose Log In.

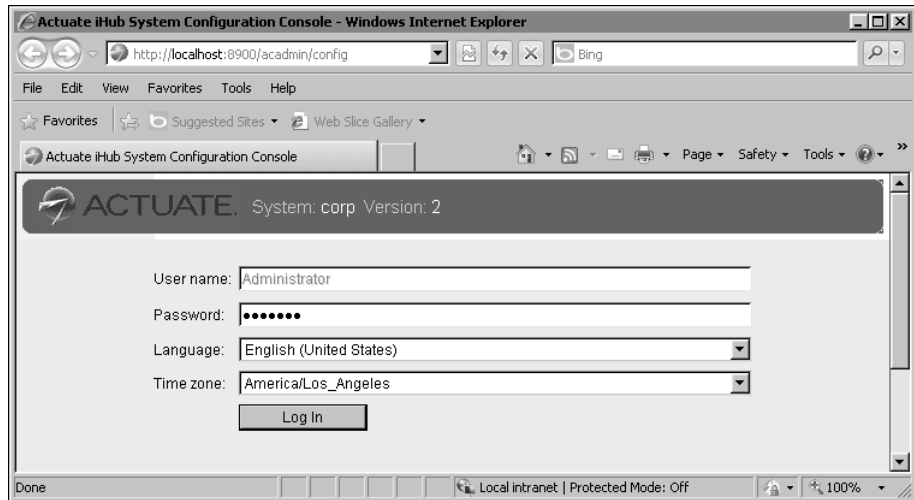


Figure 1-1 Logging in to Configuration Console

Troubleshooting problems launching the console

To launch Configuration Console, the Actuate BIRT iHub service must be running. By default, the service starts each time the machine reboots.

If the service does not start automatically, use Services in Control Panel to start the service on Windows. Alternatively, you can stop and restart the service on any platform using commands on the command line.

How to start and stop the iHub service on Windows

- 1 Open Command Prompt from the Start menu as follows:
Start→Accessories→Command Prompt
- 2 In Command Prompt, type the following command to stop the service:
`net stop "Actuate BIRT iHub 2"`
- 3 To start the service, type:
`net start "Actuate BIRT iHub 2"`

How to start and stop the iHub service on Linux

To start the iHub service each time the machine reboots, log in as root, and type:

```
./AcServer/bin/update_rclocal.sh
```

If you want to start the iHub service manually, perform the following tasks:

- 1 Navigate to the AC_SERVER_HOME/bin directory. For example, type:
`cd /home/actuate/AcServer/bin`

- 2 To run the script to start the iHub service, type:

```
./start_srvr.sh
```

The term AC_SERVER_HOME refers to the iHub installation directory. By default, iHub installs in the following directories:

Windows: C:\Program Files (x86)\Actuate\iHub2

Linux: <installation directory>/AcServer

Using Simple view

In Simple view, the administrator specifies options for the following property categories:

- License information
Display or update a license.
- iHub System
Start or stop iHub System.
- Diagnostic logging
Configure diagnostic logging levels, directory location, file size, and number of log files.
- Account settings
Change and confirm a new administrator password.
- Network settings
Specify hostname or IP address, and Management Console, Information Console, and Web Service API (IDAPI) port numbers.
- E-mail notification settings
Specify SMTP server name, hostname or IP address, listen port, SMTP greeting, and sender name and e-mail address.

Displaying and updating the license

In Simple view, the administrator checks license options and updates the iHub license by choosing Show license and Update License, as shown in Figure 1-2. In Advanced view, the administrator checks license options by choosing System→Properties→License, then chooses Update License to update the iHub license.




Figure 1-2 Viewing and updating the license file

To display the options covered by the license, choose Show License. Figure 1-3 shows the option list for an evaluation license.

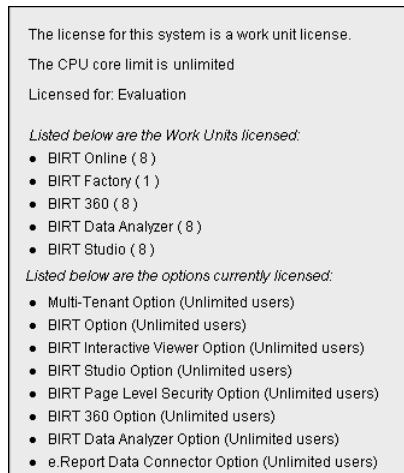


Figure 1-3 Viewing iHub license options

iHub supports the following license options:

- Multi-Tenant option
- BIRT option
- BIRT Interactive Viewer option
- BIRT Studio option
- BIRT Page Level Security option
- BIRT 360 option
- BIRT Data Analyzer option
- e.Report Data Connector option

For more information on license options, see *Installing BIRT iHub for Windows* or *Installing BIRT iHub for Linux*.

How to update an iHub System license

- 1 In License information, choose Update license.
- 2 License file appears, as shown in Figure 1-4.



Figure 1-4 Selecting a new license file

- 3 Choose Browse and select a new license file.
Choose OK. A message says that the license updated successfully. Close License file.
- 4 Restart iHub.
 - 1 In a stand-alone system, choose Stop system.
 - 2 Choose Start system.

The update takes effect when iHub restarts.

Stopping and starting iHub System

The administrator can stop and start iHub in both simple or advanced views. In Simple view, the list of user interface symbols, shown in Figure 1-5, indicate which fields require restarting a stand-alone iHub or, in a cluster, the entire iHub system. The asterisk symbol, *, indicates a required field that cannot be left blank. The exclamation point, !, indicates a field that takes a default value when left blank.




Figure 1-5 Interpreting iHub restart indicators

How to perform a system-wide restart

Simple view indicates the status of the iHub system in Version information, as shown in Figure 1-6.

- 1 To stop a running iHub system, choose Stop system, as shown in Figure 1-6. In Version information, the status changes from online to offline, as shown in Figure 1-7.




Figure 1-6 Performing a system-wide shutdown

- 2 To restart iHub, choose Start system, as shown in Figure 1-7.




Figure 1-7 Performing a system-wide startup

Configuring diagnostic logging

iHub creates a log file when an internal process starts or other event occurs, such as an error or warning. The administrator can configure the level of diagnostic information that iHub writes to a log file when an error, warning, or other event occurs. The administrator can use this information to monitor system performance and troubleshoot problems.

Understanding diagnostic logging

The administrator configures diagnostic logging by specifying settings for the following properties:

- **Level**
Specifies how much information iHub writes to a log, as shown in Figure 1-8. Select one of the following options to control the amount of information written to the diagnostic logs:

- Severe
Writes error messages only, as shown in Listing 1-1.
- Warning
Writes warning and error messages only.
- Informational
Writes warning and error messages including warning and error code descriptions.
- Fine
Writes warning and error messages with descriptions and limited diagnostic information.
- Finest
Writes the most detailed warning, error, descriptive, diagnostic, and tracing information possible, as shown in Listing 1-2.

Diagnostic logging

Level: Warning

Directory: \$AC_DATA_HOME\$/server/ *

Size: 10000 KB

Number of log files: 3

Update logging

Figure 1-8 Viewing the default level of detail for diagnostic logging

Listing 1-1 contains an excerpt from a log that describes a problem caused by an application blocking port 25, the SMTP port. Setting the diagnostic logging level to Severe results in only minimal information about the problem.

Listing 1-1 Excerpt from a log containing only severe error messages

```
****0000008000*acmail.cpp*03438*09000*2012FEB26*10:18:56
AcMailDefaultTransport: cannot open socket to
exchangesvr.abcbank.com:25
Error:10053
****0000008000*acmail.cpp*03438*09000*2012FEB26*10:18:59
AcMailDefaultTransport: cannot open socket to
exchangesvr.abcbank.com:25
Error:10053
```

Listing 1-2 contains an excerpt from a log that describes the same problem. Setting the diagnostic logging level to Finest results in the maximum possible information.

Listing 1-2 Excerpt from a log containing the finest level of information

```
****0000008000*acmail.cpp*03943*06000*2012FEB26*10:37:49
Message assigned to server iHub mail server
****0000008000*acmail.cpp*03438*09000*2012FEB26*10:37:49
AcMailDefaultTransport: cannot open socket to
exchangesvr.abcbank.com:25
Error:10053
****0000008000*mailexcept.cpp*00322*08000*2012FEB26*10:37:49
AcSMTPMailImpl::TryServer
SMTP server name: iHub mail server
Error code: 13018
Error description: SMTP: Could not connect to the SMTP Server
on the specified port.
Parameter: exchangesvr.abcbank.com:25
****0000008000*mailloadbal.cpp*00073*06000*2012FEB26*10:37:49
Moved to next server.
Current Server Index: 0
Current Backup Server Index: 0
****0000008000*mailloadbal.cpp*00386*06000*2012FEB26*10:37:49
Current Index: 0
Current Backup Index: 0
Using backup servers: false
Server "iHub mail server": quota=0, state=Retryable
****0000008000*mailloadbal.cpp*00095*06000*2012FEB26*10:37:49
Using backup servers.
```

Setting diagnostic logging to higher information levels can impact system performance and is recommended only for troubleshooting problems.

■ Directory

Specifies the log directory location. By default, iHub writes logs to the following directory, as shown in Figure 1-8:

```
AC_DATA_HOME/server/log
```

The administrator can configure iHub to write the logs to another directory in AC_DATA_HOME or other location on the network, provided permissions do not restrict access.

A log file name contains embedded information about the process event. For example, the view server created the following log file, indicating the process name, server, date, and time:

```
viewsrv11.exe.3736.urup.2012 MAR08_09_35_02_Pacific_Standard_
Time.1.log
```

■ Size

Specifies the maximum log file size in kilobytes (KB). When the log file size reaches the limit, iHub starts writing to a new log file. To retain enough

historical information to investigate a problem, particularly when using a verbose log level, such as Finest, increase the limit.

- **Number of log files**

Specifies the maximum number of log files that iHub creates since the system last started. The administrator can use this property to manage log space usage in an environment with limited disk space.

When a log file reaches the limit specified by Size, iHub creates another log file, until reaching the maximum number of log files specified in Number of log files. iHub then deletes the earliest log file and creates a new one. To conserve disk space usage in the log file directory, specify a lower limit.

Stopping BIRT iHub service breaks the log generation cycle. When BIRT iHub service restarts, a new cycle begins. iHub creates the specified number of log files, but does not overwrite any pre-existing files.

How to configure diagnostic logging

- 1 In Diagnostic logging, perform the following tasks:

- 1 In Level, accept the default level, Warning, or choose another level for all processes that write to the log files. For example, choose Finest, as shown in Figure 1-9.

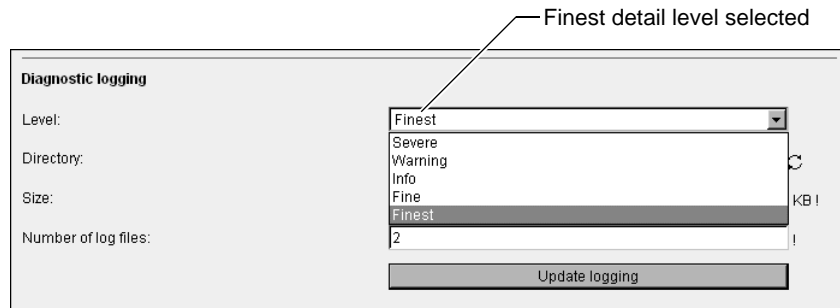


Figure 1-9 Selecting the finest detail level of detail for a log

- 2 In Directory, accept the default directory, AC_DATA_HOME/server/log, or specify a different directory for iHub to write log files. Ensure that the user account running iHub has permission to write to this location.
- 3 In Size, accept the default size limit for log files, 10000 KB, or specify a different limit.
- 4 In Number of log files, accept the default, 3, or specify the maximum number of log files allowed in the log directory.
- 5 Choose Update logging, as shown in Figure 1-10.

Update logging

Diagnostic logging

Level:

Directory: *

Size: KB

Number of log files:

Figure 1-10 Default diagnostic logging properties

2 If you change Directory, restart iHub.

Configuring a diagnostic logging category

A configuration change in the default Simple view applies to all categories of logging. In the default Simple view, Category does not appear in Diagnostic logging.

Advanced view—Server Configuration Templates—Settings supports changing diagnostic logging for specific types of logs. If the administrator configures one logging type differently from the others in Advanced view, Category appears in Simple view, as shown in Figure 1-11.

Category list

Diagnostic logging

Category:

- General
- Factory
- Integration
- Caching
- Viewing

Level:

Directory:

Size: KB

Number of log files:

Figure 1-11 Viewing a list of categories in Simple view

The Category list includes the following log types:

- General
- Factory
- Integration

- Caching
- Viewing

Table 1-1 describes these categories, provides examples, and lists the log file names.

Table 1-1 Diagnostic logging categories and log files

Category	Logged events	Log file
General	Logs Encyclopedia volume events: <ul style="list-style-type: none"> ■ Validating login requests ■ Creating Encyclopedia volume folders ■ Adding files ■ Managing e-mail notification Records other events that occur outside Factory, Integration, Caching, and Viewing service processes	encycsrvr11
Factory	Logs Factory service events: <ul style="list-style-type: none"> ■ Running designs ■ Generating queries ■ Printing documents 	facsrvr11
Integration	Logs Integration service events regarding information objects that use data from multiple data sources	intsrvr11
Caching	Logs Caching service events regarding information object cache	cachesrvr11
Viewing	Logs Caching service events: <ul style="list-style-type: none"> ■ Viewing documents in DHTML format ■ Converting documents to formats, such as Excel and PDF ■ Handling requests to download files from an Encyclopedia volume 	viewsrv11

How to configure diagnostic logging by category

- 1 In Diagnostic logging, in Category, choose a log category to configure. For example, choose General.
- 2 In Level, Directory, Size, and Number of log files, accept the default values or specify new values.
- 3 Choose Update logging.
- 4 If you change Directory, restart iHub.

Configuration Console writes setting changes to the acserverconfig.xml file in AC_DATA_HOME/config/iHub2. Table 1-2 lists the property names that appear in Configuration Console with the corresponding parameter names in acserverconfig.xml, including default settings, ranges, and when the property change takes effect.

Table 1-2 Diagnostic logging parameters

Property name	Parameter name	Default	Range	Takes effect
Log level	LogLevel	8000	0 - 9000	Immediate
Log size	LogSize	10000 KB		Immediate
Number of log files	NumLogs	3		Immediate

For more information about diagnostic logging, see Chapter 3, “Using diagnostic, usage, and error logging.”

Changing the Configuration Console password

During installation, the installer specifies the Configuration Console password. In Simple view, the administrator can change this password in Account settings. In Advanced view, the administrator chooses System→Properties to the system password.

How to change the password for configuring iHub

- 1 In Account settings, type the old and new passwords.
- 2 In Confirm system password, type the new password again.
- 3 Choose Change password, as shown in Figure 1-12.

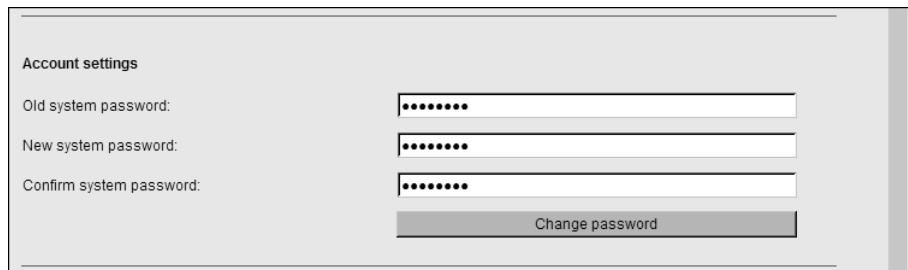
The image shows a screenshot of a software window titled "Account settings". Inside the window, there are three text input fields. The first field is labeled "Old system password:" and contains seven dots. The second field is labeled "New system password:" and also contains seven dots. The third field is labeled "Confirm system password:" and contains seven dots. Below these three fields is a button labeled "Change password". The window has a light gray background and a standard title bar.

Figure 1-12 Changing the administrator password for Configuration Console

Configuring network settings

iHub communicates with Management Console, Information Console, and web service APIs through network ports. In Network settings, the administrator can change the following items:

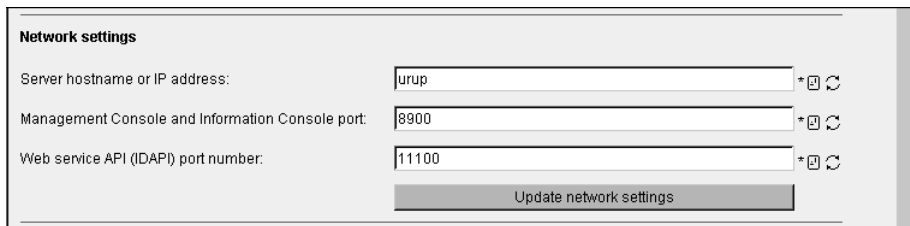
- iHub hostname or IP address
- Management Console and Information Console port
- Web service API (IDAPI) port

Any change to a network setting requires restarting iHub.

If you change the port numbers assigned by the installation program, test that these ports work with these console and web service applications. By default, the Microsoft Vista firewall blocks port 8900. Ensure this port is unblocked.

How to change network settings

- 1 In Network settings, in Server hostname or IP address, type the new iHub host name or IP address, as shown in Figure 1-13.



The screenshot shows a window titled "Network settings". It has three rows of input fields. The first row is "Server hostname or IP address:" with the text "urup" entered. The second row is "Management Console and Information Console port:" with the text "8900" entered. The third row is "Web service API (IDAPI) port number:" with the text "11100" entered. Each input field has a small icon to its right. At the bottom of the window is a button labeled "Update network settings".

Figure 1-13 Network settings

- 2 In Management Console and Information Console port, type the port number used by the application container to listen for requests from the consoles.
- 3 In Web service API (IDAPI) port, type the port number used by the Actuate Information Delivery API.
- 4 Choose Update network settings.
- 5 Restart iHub.

Configuring e-mail notification settings

iHub can notify users by e-mail that a job completed or failed. The e-mail includes a link to the document in HTML format. iHub can attach the document in other formats to the e-mail. A template, which the administrator can modify, specifies the message content, as shown in Figure 1-14.

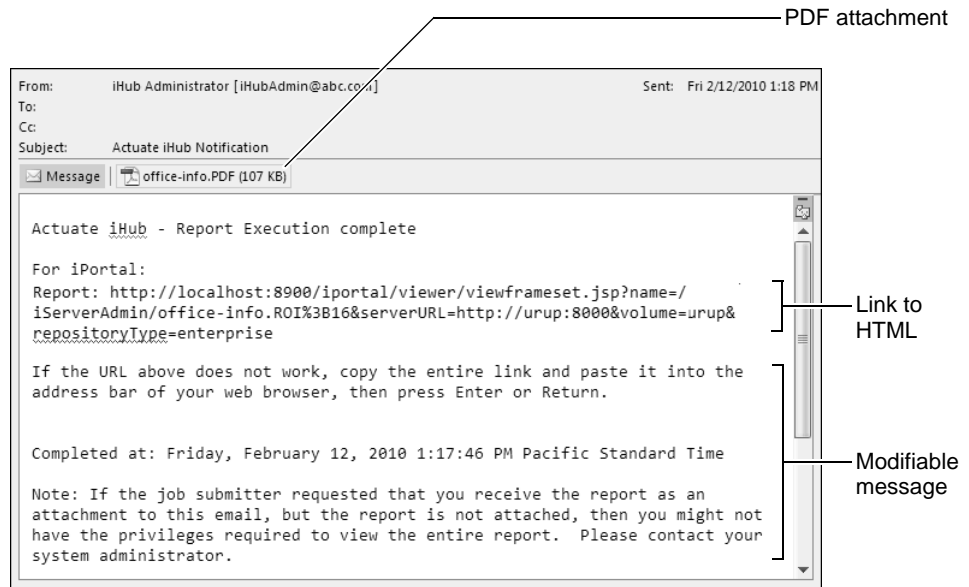


Figure 1-14 Viewing an example of an SMTP e-mail notice

The e-mail message template, `acnotification.xml`, is located in `AC_SERVER_HOME\etc`.

How to configure e-mail notification settings

To use a Simple Mail Transfer Protocol (SMTP) e-mail server for iHub e-mail notification, configure the properties in Server e-mail notification settings.

- 1 In SMTP Server name, type an arbitrary name for the new mail server that appears in the list of SMTP servers in Configuration Console. For example, type:
`iHub mail server`
- 2 In Hostname or IP Address, type the IP address or the fully qualified domain name of the mail server. For example, type:
`exchangesvr.abcbank.com`
- 3 In Listen port, accept the default port number, that iHub uses for e-mail notification, or specify the number of a free port. In Figure 1-15, the default listen port number is 25.
- 4 In SMTP greeting, accept the default, which is no entry. iHub sends HELO and appends the blank suffix to the greeting during protocol exchanges with the mail server. Alternatively, specify a suffix to the greeting. iHub sends HELO <suffix>.

- 5 In Sender e-mail address, specify the e-mail address that appears in the From line of the e-mail notification. iHub also sends an alert to this address when the mail server cannot deliver an e-mail notification to a user. For example, type:

iHubAdmin@abc.com

- 6 In Sender name, specify the name that appears in the From line of the e-mail notification. For example, type:

iHub Administrator

Figure 1-15 shows Server e-mail notification settings.

Server e-mail notification settings	
SMTP Server name	iHub mail server * 🌐
Hostname or IP Address	exchangesvr.abcbank.com *
Listen port	25 !
SMTP greeting	!
Sender e-mail address	iHubAdmin@abc.com *
Sender name	iHub Administrator !
<button>Update e-mail settings</button>	

Figure 1-15 Server e-mail notification settings

- 7 Choose Update e-mail settings.
- 8 Restart iHub.

For more information about e-mail notification, including additional advanced settings, see Chapter 4, “Configuring e-mail notification.”

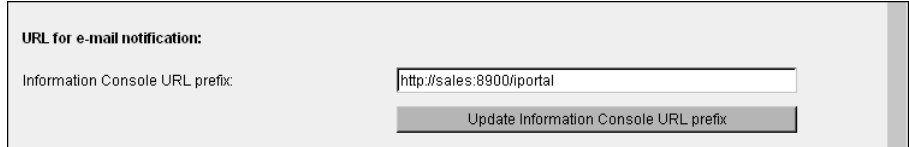
Configuring the URL for e-mail notification

The administrator can configure iHub to send an e-mail to an Information Console user about a completed job. When a user receives a default e-mail notice about a completed job, the e-mail message contains a link to the generated document.

How to configure the URL for e-mail notification

- 1 In URL for e-mail notification, in Information Console URL prefix, specify the machine name, port, and context root of Information Console, as shown in Figure 1-16. This property is also configurable in Email notification on Advanced view in Volume—Properties.

Information Console URL prefix can also be used to specify a hyperlink, such as a drill-through hyperlink, containing the appropriate Information Console context string in a PDF document. For more information on customizing e-mail notifications and Information Console URL prefix, see Chapter 4, “Configuring e-mail notification,” later in this book.



URL for e-mail notification:

Information Console URL prefix:

Figure 1-16 Adding a URL prefix to e-mail notices

In this example, the machine name is sales, the port is 8900, and the context root of Information Console is ortal.

2 Choose Update Information Console URL prefix.

A document opens in Information Console when a user clicks the link. When the user clicks a link in a completion notice, browser security settings can sometimes prevent the display of the page. To avoid this problem, set browser security to medium low. For example, in Internet Explorer, choose Tools→Internet Options→Security→Trusted Sites→Medium Low.

Using Advanced view

From Simple view, choose Advanced view, as shown in Figure 1-17.

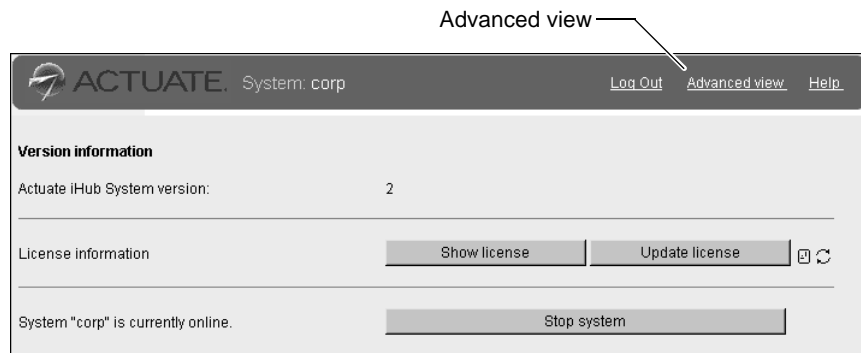


Figure 1-17 Choosing Advanced view from Simple view

To go back to Simple view, choose Simple view, as shown in Figure 1-18.

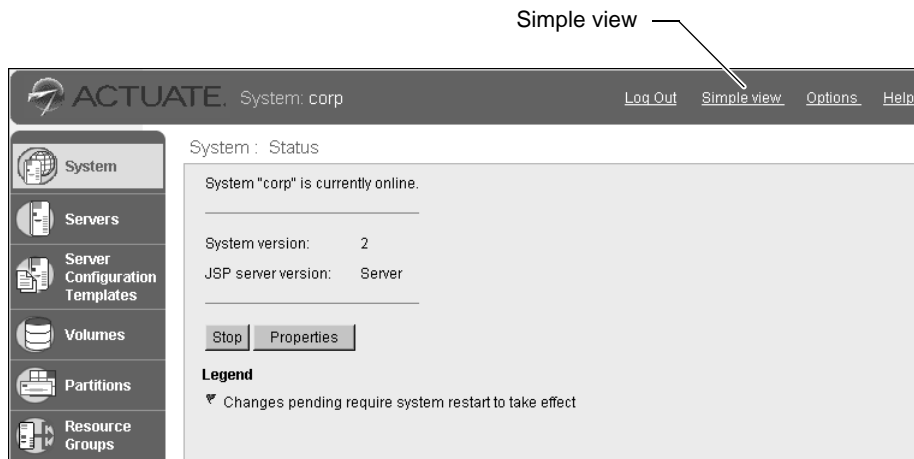


Figure 1-18 Configuring iHub in Advanced view

In addition to Simple view, Advanced view also contains the following items:

- **Log Out**
Logs out of Configuration Console.
- **Options**
Provides options for configuring the Configuration Console display.
- **Help**
Provides access to iHub system documentation in a browsable format. For more information about installing, accessing, and navigating documentation in Help, see *Installing BIRT iHub for Windows* or *Installing BIRT iHub for Linux*.

Figure 1-18 shows Advanced view with System selected. The left menu provides access to user interfaces that allow the administrator to view and change the settings for the following iHub system property categories:

- **System**
System name, usage and error logging, notification, regional settings, license, and other advanced properties.
- **Server**
Server, application container, Process Management Daemon (PMD), SOAP, and enabled service settings.
- **Server Configuration Templates**
Factory, Message Distribution, Viewing, Integration, Caching services. Other advanced options, such as NetOSI File Types, filetype driver information, diagnostic logging, process management, and other settings.
- **Volumes**
Metadata database, schema, and volume property settings.

- **Partitions**
Partition name, path, status, and other settings.
- **Resource Groups**
Name, description, status, report type, volume, work unit type, start arguments, priority settings and template assignments for resource groups associated with information object and BIRT jobs.
- **Printers**
Name, path, spool command, and other template settings.

For more information about configuring iHub system properties, see the related chapters later in this book. The remaining sections in this chapter describe the options for configuring Configuration Console display settings.

Setting Configuration Console options

The administrator can change the way Configuration Console displays the following settings:

- Regional settings, such as locale and time zone
- The number and order of columns that appear in server, template, volume, partition, or printer lists
- About, providing version information on installed iHub system software

To view or change display options, choose Options in Advanced view, as shown in Figure 1-19.

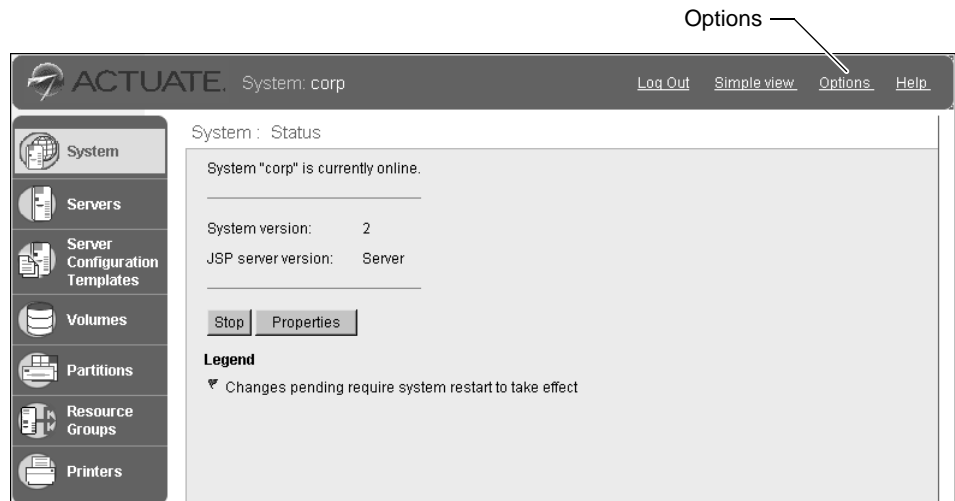


Figure 1-19 Setting options

Choose an option to modify. Figure 1-20 shows the General option selected.

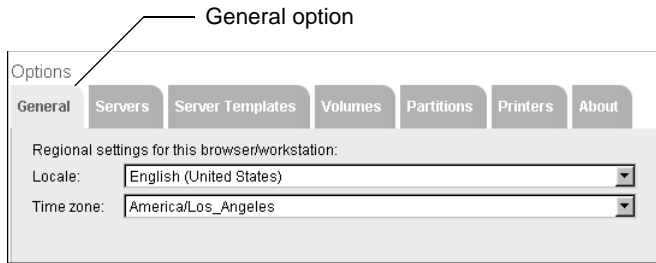


Figure 1-20 Selecting the General option

In Options—General, change the regional settings to correspond to the current locale and time zone, if different from the settings specified during installation.

Changing server, template, volume, or partition options

In Options—Servers, Server Templates, Volumes, Partitions, or Printers, change the number and order of the columns that appear in server, template, volume, partition, or printer lists by performing the following tasks:

- Use the left and right arrows to move column names between Available columns and Selected columns.
- Use the up and down arrows to change the order in which the information appears.

Figure 1-21 shows the list of available and selected columns in Options—Volumes.

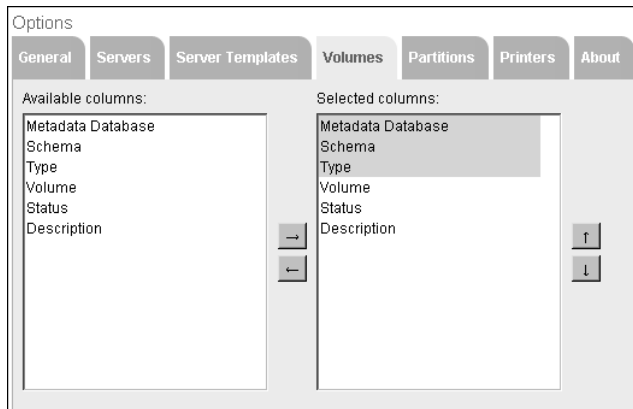


Figure 1-21 Modifying columns in Options—Volumes

In Selected columns, the order in which the columns appear vertically determines the horizontal order in the displayed list. You cannot hide or change the order of highlighted columns in Selected columns. For example, you cannot hide or change the order for the Metadata Database, Schema, or Type columns in Options—Volumes.

How to change the list of servers, volumes, or partitions

- 1 In Advanced view of Configuration Console, choose Options.
Options—General appears.
- 2 Choose Servers, Server Templates, Volumes, Partitions, or Printers.
The available and selected columns for the list appears.
- 3 Change the column display options:
 - To add columns to the list, select the columns in Available columns. Choose the right arrow to move the selected columns to Selected columns.
 - To remove columns from the list, select the columns in Selected columns, and choose the left arrow.
 - To change the order in which the columns appear in the list, select a column in Selected columns, and choose the up or down arrow to reposition the column in the list, as shown in Figure 1-21.

Choose OK.

In the left menu of Advanced view, when you choose Servers, Server Configuration Templates, Volumes, Partitions. Resource Groups, or Printers, a related list appears. Figure 1-22 shows an example of the Volumes list.

In Figure 1-22, the list displays the Metadata database, Schema, Volume, Status, and Description columns for each volume, such as corp. The changes made in Options—Volume alter the number and order of columns that appear in this list and similar lists accessed from the left menu of Advanced view.

Volumes					
Act upon selected items					
Metadata Database	Schema	Type	Volume	Status	Description
Default ActuatePostoreSQL_MetadataDatabase	ac_corp	Volume	<input type="checkbox"/> corp	ONLINE	
	ac_corp_system	System			
Legend					
Changes pending require volume restart to take effect					

Figure 1-22 Viewing the volumes list

About the servers list

Table 1-3 describes the columns available in Options—Servers that display in the servers list.

Table 1-3 Columns available for display in the node list

Column name	Description
Name	Name of the server, typically the machine name.
Template	Name of the template that defines the configuration properties for the server.
Status	Current status of the machine. Status can be master, online, offline, or stopping. Status also lists the following services configured for iHub on the specified machine: <ul style="list-style-type: none">■ M for Message Distribution service■ F for Factory service■ V for View service■ C for Caching service■ I for Integration service
Description	Description field from the iHub definition.
Operating system and version	Server machine's operating system and version.
Actuate version	iHub release number.
Current requests	Current number of active requests.

A flag appears when you must restart iHub to apply updates to the configuration.

About the server templates list

Table 1-4 describes the columns available in Options—Server Templates that display in the server configuration templates list.

Table 1-4 Columns available for display in the server templates list

Column name	Description
Template	Name of the template that defines the configuration properties for the server
Server	Name of the server, typically the machine name

About the volumes list

Table 1-5 describes the columns available in Options—Volumes that display in the Encyclopedia volume list.

Table 1-5 Columns in the Encyclopedia volumes list

Column name	Description
Metadata Database	Names of the metadata databases in this iHub system.
Schema	Names of the schemas. One database can contain one or more schemas.
Type	Type of schema, either Volume or System.
Volume	Names of the volumes. One schema can contain one or more volumes.
Status	Current status of the volume. The status of the volume is online or offline.
Description	Description field from the volume definition.

About the partitions list

Table 1-6 describes the columns available in Options—Partitions that display in the partitions list.

Table 1-6 Columns in the partitions list

Column name	Description
Name	Name of partition.
Status	Current status of the partition. The status of the partition is active or inactive.
Volume	Names of the volumes associated with the partition if assigned.

About the printers list

Name is the only available column in Options—Printers. Name is the name of the printer.

Viewing release information

In Options—About, you can view version information about the iHub system software installation, as shown in Figure 1-23.

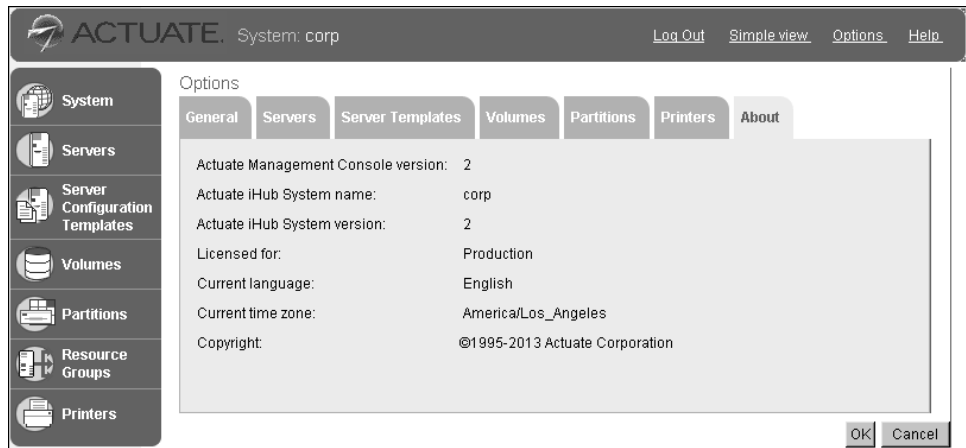


Figure 1-23 Viewing iHub System general information

2

Configuring an Encyclopedia volume

This chapter contains the following topics:

- Understanding an Encyclopedia volume
- Configuring a partition
- Working with databases, schemas, and volumes

Understanding an Encyclopedia volume

An Encyclopedia volume consists of files such as BIRT designs, documents, and information objects. iHub creates a default Encyclopedia volume during installation. The installation process names the default Encyclopedia volume the same as the machine name, as shown in Figure 2-1.

The screenshot shows the 'System : Properties' dialog box with the 'General' tab selected. The left sidebar contains icons for System, Servers, Server Configuration Templates, Volumes, Partitions, Resource Groups, and Printers. The main area contains the following fields:

- System name:** corp
- System password:** [masked]
- System password confirm:** [masked]
- System Heartbeat:**
 - Heartbeat send period: 30 sec
 - Heartbeat failure period: 90 sec
- System default volume:**
 - Volume: corp
- System schema:**
 - System schema: ac_corp_system

At the bottom, there are three buttons: OK, Cancel, and Apply. Below the fields, there are three lines of text:

- * These fields are required and cannot be left blank
- [icon] These fields require system restart to take effect
- (!) These fields will take default value if left blank

Figure 2-1 Viewing general system properties

The default Encyclopedia volume directory is AC_DATA_HOME/encyc. By default, AC_DATA_HOME on a Windows system is Actuate\iHub\data. On a Linux system, AC_DATA_HOME is AcServer/data.

In an out-of-the-box (OOTB) installation, iHub stores Encyclopedia volume metadata, such as information about users, roles, groups, and job schedules, in the PostgreSQL database installed with iHub. The PostgreSQL database resides in AC_DATA_HOME/encyc. iHub stores configuration metadata in the database separately from data such as designs, documents, information objects, and other iHub data objects, which are stored in the file system. Files containing data have the .dat file-name extension in the file system, regardless of their Actuate file type. By default, the .dat files reside in AC_DATA_HOME/encyc/file.

Configuring a partition

An iHub partition is a physical disk location used to store Encyclopedia volume data files. Every Encyclopedia volume must have an iHub partition.

An Encyclopedia volume supports using multiple iHub partitions. You can expedite input and output processing by using multiple partitions that operate across separate physical disks. iHub transparently manages the allocation of files among partitions. Encyclopedia volumes cannot share a partition.

The default Encyclopedia volume resides in the primary partition, DefaultPartition. The default path of the primary partition is AC_DATA_HOME /encyc. Figure 2-2 shows the path of the primary partition on Windows.



Figure 2-2 Default Partition path

The administrator can change the location of the primary partition, but cannot remove it. If no other partitions exist, the primary partition stores all designs, documents, and other data files.

Adding a partition

Add a partition before attempting to add a new Encyclopedia volume to iHub.

How to add a partition to iHub

- 1 Make a directory for the partition on the physical drive of the machine or storage device that iHub can access. For example, make a directory called encyc2 in AC_DATA_HOME.
- 2 From Advanced view of Configuration Console, choose Partitions.
- 3 In Partitions, choose Add partition.
- 4 In Partition name, specify a name. For example, name the partition, Partition2.

In Partition Path, specify the fully qualified path to the partition directory, as shown in Figure 2-3. Choose OK.

Partitions > Add Partition

Partition name: *

Template Name Partition Path

urup *

* These fields are required and cannot be left blank

OK Cancel

Figure 2-3 Adding a partition

- 5 In Partitions, choose the new partition, Partition2, from the list of partitions, as shown in Figure 2-4.

Partitions

Add Partition

Name	Status	Volume
DefaultPartition	Active	corp
Partition2		

Figure 2-4 Choosing the new partition

- 6 In Server Settings, choose Test, as shown in Figure 2-5.

Partitions > Partition2 : Server Settings

Template Name Partition Path

urup Change... Test *

* These fields are required and cannot be left blank

OK Cancel Apply

Figure 2-5 Testing the new partition

If the test succeeds, the message in Figure 2-6 appears. If the test fails, check that the directory named in the partition path exists.

Message from webpage

The test is successful.

OK

Figure 2-6 Choosing OK after successful partition path test

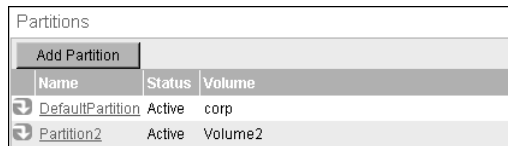
Choose OK.

Viewing partitions

After the administrator assigns a partition to an Encyclopedia volume, Partitions displays one of the following status conditions and the name of the Encyclopedia volume.

- **Active**
An Encyclopedia volume is using the partition.
- **Phaseout**
An Encyclopedia volume is moving data out of the partition.
- **Unused**
The partition is assigned to an Encyclopedia volume, but it is not using the partition.

From Advanced view of Configuration Console, choose Partitions to view the partitions that are available, as shown in Figure 2-7.



Partitions		
Add Partition		
Name	Status	Volume
DefaultPartition	Active	corp
Partition2	Active	Volume2

Figure 2-7 Displaying the list of partitions

Working with databases, schemas, and volumes

iHub stores information related to users, roles, groups, files and folders, or metadata, in a third-party relational database management system (RDBMS). By default, iHub uses the PostgreSQL RDBMS for this purpose. iHub also supports the use of an Oracle, DB2, and SQLServer database.

A database can contain one or more schemas. An Encyclopedia volume schema can relate to one or more volumes. When installing BIRT iHub with either the OOTB PostgreSQL or a third-party database, the install program creates two schemas, one for Encyclopedia metadata, and one for the system metadata. In Advanced view of Configuration Console, on Volumes, the metadata database appears in the first column, the schemas the database contains appear in the second column, and the volume the Encyclopedia volume schema relates to appears in the third column, as shown in Figure 2-8.

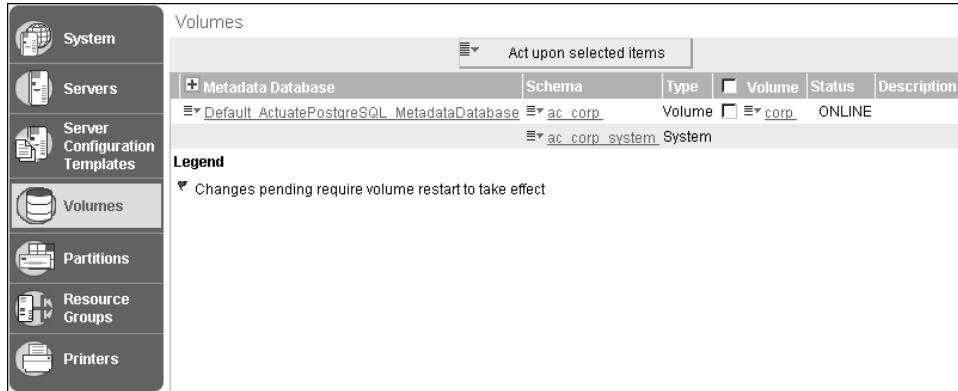


Figure 2-8 Viewing Volumes on Advanced view of Configuration Console

Configuring a metadata database

On Volumes, an administrator can create, configure, or delete a metadata database used to contain Encyclopedia volume schema.

How to add a metadata database

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Volumes.
- 2 Point to the icon next to the Metadata Database heading and choose Add new metadata database, as shown in Figure 2-9.

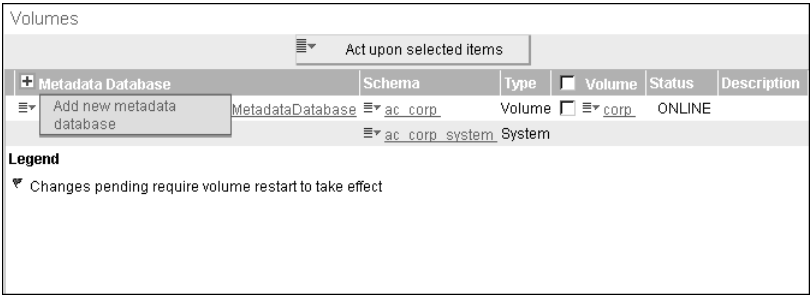


Figure 2-9 Choosing to add a new metadata database

- 3 On New Metadata Database, as shown in Figure 2-10, perform the following tasks:
 - 1 In Metadata database name, type a name for the metadata database.
 - 2 In Database type, select the type of database to create.
 - 3 In Database server, type the host name of the machine containing the database, such as localhost.

- 4 In Database name, type a name for the database.
- 5 In Connection login, type the database user name.
- 6 In Connection password, type the database user name password.
- 7 In Database port, specify a port number, or accept the default value.

Choose OK.

Figure 2-10 Adding a new metadata database

How to edit metadata database properties

- 1 On Volumes, point to the icon next to a database and choose Properties.
- 2 On Metadata Database, the settings are the same as when you add a new metadata database, except that when editing database properties, Metadata Database includes the option to confirm the connection password if you change it. Choose Advanced.
- 3 On Advanced, choose MetadataDatabaseProperties.
- 4 On MetadataDatabaseProperties, as shown in Figure 2-11, you can specify the following additional database properties:
 - Type of database
 - Database tablespace name
 - Database tablespace path
 - Super user name
 - Super user password

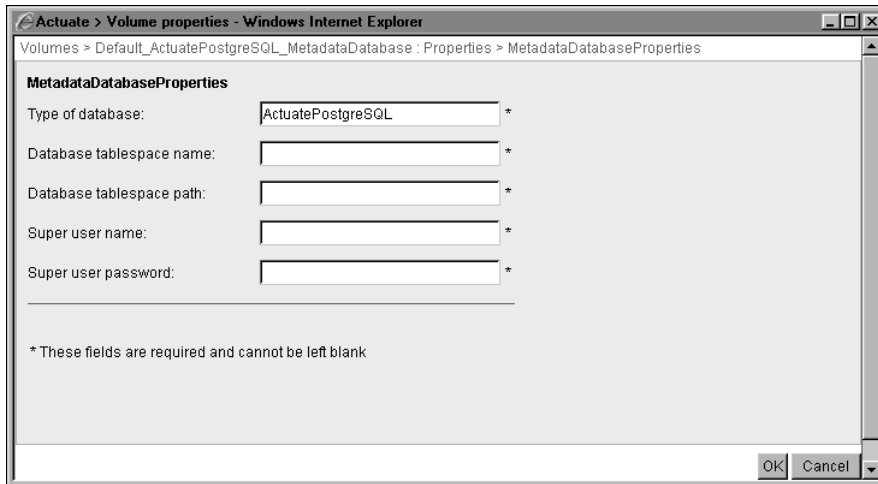


Figure 2-11 Setting additional metadata database properties

- 5 On Advanced, choose Database Connection Pool Manager Settings.
- 6 On Database Connection Pool Manager Settings, as shown in Figure 2-12, you can specify the following database connection properties:
 - Initial size of the connection pool
 - Minimum number of open connections in the connection pool
 - Maximum number of open connections in the connection pool
 - Maximum number of open connections in the internal connection pool
 - Maximum number of open connections in the system connection pool
 - Maximum number of open connections in the cluster connection pool
 - Maximum idle time for a connection before removing it from the connection pool
 - Max connection wait time in seconds
 - Max query execution time in seconds (use zero value to indicate no limit)
 - PreparedStatement cache size per connection (use zero value to indicate no caching)

To acquire a new JDBC connection for each database query, type a value of 0 in Minimum number of open connections in the connection pool and in Maximum number of open connections in the connection pool.

Database Connection Pool Manager Settings

Initial size of the connection pool: ! ⌂ ↻

Minimum number of open connections in the connection pool: ! ⌂ ↻

Maximum number of open connections in the connection pool: ! ⌂ ↻

Maximum number of open connections in the internal connection pool: ! ⌂ ↻

Maximum number of open connections in the system connection pool: ! ⌂ ↻

Maximum number of open connections in the cluster connection pool: ! ⌂ ↻

Maximum idle time for a connection before removing it from the connection pool: Seconds ! ⌂ ↻

Max connection wait time in seconds: ! ⌂ ↻

Max query execution time in seconds (use zero value to indicate no limit): !

PreparedStatement cache size per connection (use zero value to indicate no caching): ! ⌂ ↻

☐ ⌂ ↻ These fields require server restart to take effect
(!) These fields will take default value if left blank

OK Cancel

Figure 2-12 Configuring database connection pool manager settings

Configuring a schema

Additionally, the administrator can add, configure, or remove a schema used to contain Encyclopedia volume metadata.

How to add a schema

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Volumes.
- 2 On Volumes, point to the icon next to a metadata database and choose either Add system schema, or Add volume schema. Figure 2-13 shows Add volume schema. The procedure for adding either schema type is the same. The default metadata database is Default ActuatePostgreSQL MetadataDatabase.




Figure 2-13 Choosing to add a schema

- 3 On New Volume Schema, as shown in Figure 2-14, perform the following tasks:
 - 1 In Schema name, type a name for the new schema. The name must be less than 30 characters.
 - 2 In Schema owner name, type the schema owner name.
 - 3 Type and confirm a password for the schema owner.
 - 4 In Database superuser, type the database superuser name. For the PostgreSQL RDBMS that installs with iHub by default, the PostgreSQL superuser name is postgres.
 - 5 In Database superuser password, type the password that the installer specified for the database superuser during the iHub installation.

Choose OK.

Volumes > New Volume Schema

Schema

Metadata Database: Default_ActuatePostgreSQL_MetadataDatabase

Schema type: Volume

Schema name: *

Database schema name: *

Database schema password: *

Database schema password confirm:

Please enter the database superuser credentials to make changes to the database schema for the encyclopedial volume.

Database superuser: *

Database superuser password: *

* These fields are required and cannot be left blank

OK Cancel Apply

Figure 2-14 Adding a new schema

How to edit schema properties

- 1 On Volumes, point to the icon next to a schema and choose Properties, as shown in Figure 2-15.




Figure 2-15 Choosing to edit schema properties

- 2 In Schema, you can change the schema owner name and password, as shown in Figure 2-16.

The screenshot shows the 'Schema' configuration window for 'ac_corp'. It contains the following fields and values:

- Metadata Database: Default_ActuatePostgreSQL_MetadataDatabase
- Schema type: Volume
- Schema name: ac_corp
- Database schema name: ac_corp *
- Database schema password: ***** *
- Database schema password confirm: ***** *

A note at the bottom states: '* These fields are required and cannot be left blank'. At the bottom right, there are buttons for 'Test', 'OK', 'Cancel', and 'Apply'.

Figure 2-16 Editing schema properties

- 3 To test the connection, choose Test.
Choose OK.

Configuring an Encyclopedia volume

The administrator can also add, configure, or remove a volume contained in a schema within a metadata database.

Adding an Encyclopedia volume

Unless a partition, unassigned to an Encyclopedia volume, already exists, you must first create a new partition to assign to the Encyclopedia volume. Next, add the Encyclopedia volume to iHub by performing the following tasks.

How to add an Encyclopedia volume

- 1 Log in to Configuration Console and choose Advanced view. In Advanced view, choose Volumes.

- 2 On Volumes, point to the icon next to a schema and choose Add Volume, as shown in Figure 2-17.




Figure 2-17 Choosing to add a volume

- 3 On New Volume—General, perform the following steps:
 - 1 Type a name for the new volume. For example, type Volume2, as shown in Figure 2-18.




Figure 2-18 Specifying general properties

- 2 In Primary partition, accept the default or select another unassigned partition. For example, accept Partition2.
- 4 Choose Partitions, and start the partition for the new Encyclopedia volume by performing the following steps:
 - 1 In Available partitions, select a partition, then move it to Selected by choosing the right arrow.
 - 2 In Selected partitions, select the partition. Choose Start, as shown in Figure 2-19.

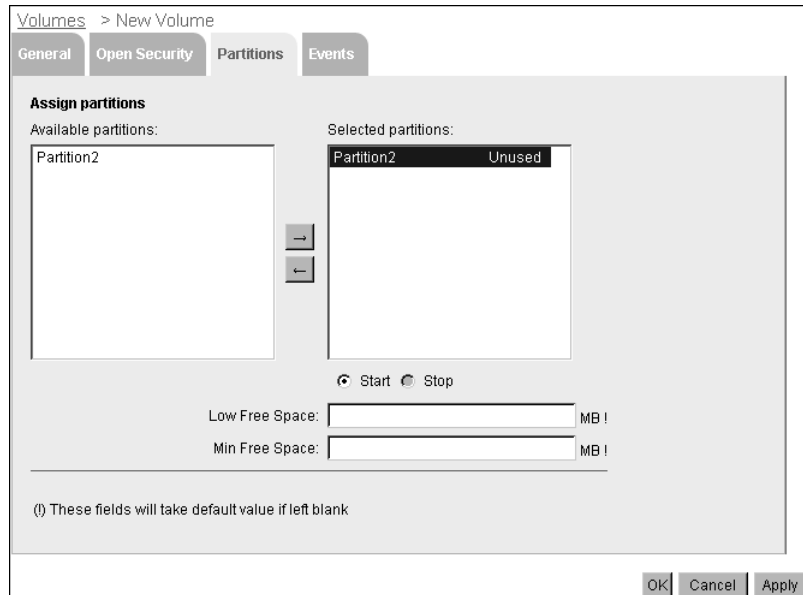


Figure 2-19 Starting the partition

Choose OK.

- 5 In Volumes, point to the arrow next to the new volume name, and choose Take online, as shown in Figure 2-20.

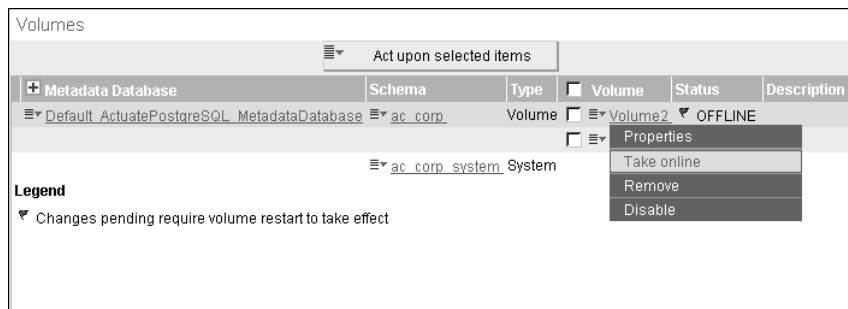


Figure 2-20 Taking a new volume online

- 6 In Volumes, check that the status of the new volume changes to ONLINE, as shown in Figure 2-21.




Figure 2-21 Confirming that the volume is online

If the volume does not go online, check for insufficient free disk space for the partition and consider configuring the free space threshold.

Editing Encyclopedia volume properties

On Volumes, an administrator can edit Encyclopedia volume properties in Volumes—Properties—General.

How to edit Encyclopedia volume properties

- 1 On Volumes, point to the icon next to a volume name and choose Properties, as shown in Figure 2-22.




Figure 2-22 Choosing Encyclopedia volume properties

- 2 On Properties, the available properties are the same as when you add a new volume, with the exception that you can set advanced properties, as shown in Figure 2-23.

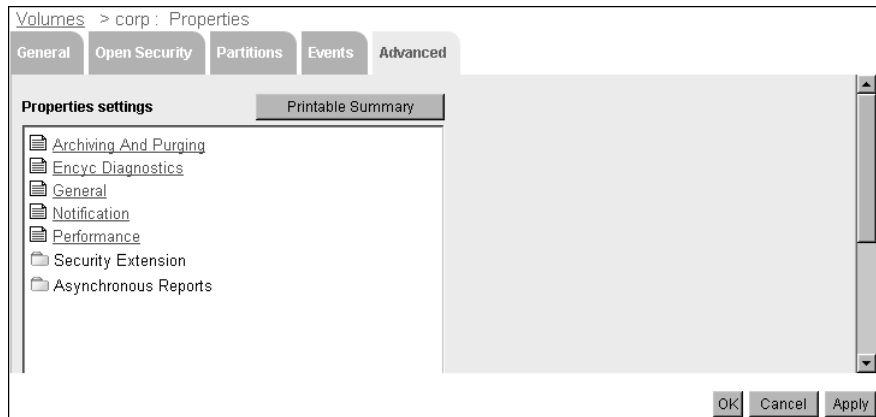


Figure 2-23 Viewing advanced volume properties

Table 2-1 lists some of the properties that appear on Volumes—Properties—General.

Table 2-1 New Encyclopedia volume general properties

Property	Description
Volume name	Name of the Encyclopedia volume.
Description	Optional description of the Encyclopedia volume.
Schedule for purging notices	Time or times at which iHub deletes job completion notices. The value format is a semicolon-separated list of times, in ascending order. Use a 24-hour format. For example: 03:15;16:15
Schedule for purging deleted files	Time or times at which iHub removes deleted files from the system. The value format is a semicolon-separated list of times, in ascending order. Use a 24-hour format. For example: 03:15;16:15
Primary partition	Name of the primary partition for the Encyclopedia volume.
Min Free Space for the primary partition	Minimum amount of free disk space that the primary partition requires. If the free space falls below the specified minimum, iHub does not create the file.
Use archive service	Specify the command-line string iHub runs to start the SOAP-based Inline Archive Driver application.

(continues)

Table 2-1 New Encyclopedia volume general properties (continued)

Property	Description
Information Console URL prefix	<p>URL prefix, for example <code>http://sales:8900/iportal</code>, to add to the hypertext link in an e-mail notification sent to Information Console users. Use also to specify a hyperlink, such as a drill-through hyperlink, containing the Information Console context string in a generated PDF document.</p> <p>For more information on customizing e-mail notifications, see Chapter 4, “Configuring e-mail notification,” later in this book. For more information on configuring Information Console URL prefix to specify the context root for a hyperlink in a PDF document, see “Specifying a hyperlink in a PDF document,” later in this chapter.</p>

The following examples describe how to configure various general volume properties.

Scheduling purging of job completion notices

The administrator can disable purging of job completion notices, but this action can cause too many notices to build up. The best practice is to configure iHub to purge job completion notices during times of light use and at least one hour before or after autoarchiving.

Use 24-hour clock time points, separated by a semicolon. For example, the default value: `2:15` runs a purge task at 2:15 A.M. every day. A blank disables the job and notice purging.

Scheduling purging of deleted files

The physical file for an Encyclopedia volume file appearing in Management Console resides in the Encyclopedia volume file folder as a `.dat` file. For a file a user deletes from a volume in Management Console, iHub purges the corresponding `.dat` file from the volume file folder at the time or times the Schedule for purging deleted files property specifies. Specify multiple purge times using 24-hour clock time points, separated by a semicolon.

Configuring an Encyclopedia volume partition

View and modify the partitions assigned to an Encyclopedia volume in Volumes—Properties—Partitions.

Assigning a partition to an Encyclopedia volume partition

To associate a partition with a volume, assign the partition to the volume, as shown in Figure 2-24.



Figure 2-24 Assigning a partition to an Encyclopedia volume

How to assign a partition to an Encyclopedia volume

- 1 In Volumes—Properties—Partitions, use the left and right arrows to move iHub partition names between Available partitions and Selected partitions.

Available partitions lists the partitions that are available for use by an Encyclopedia volume. Selected partitions lists the partitions iHub uses to store Encyclopedia volume files.
- 2 Use the up and down arrows to change the order of the partitions in Selected partitions.

Primary indicates which of the selected partitions is the primary partition. Active indicates which of the selected partitions are actively in use.
- 3 Select a partition in Selected partitions and choose Start or Stop to change the status of a partition.

Choosing Start activates an inactive partition. You must activate a partition before using it for an Encyclopedia volume. Choosing Stop changes the status of a partition to Phasing out or Not in use. When a partition is stopped, you cannot use it.
- 4 Select a partition in Selected partitions and specify free disk space properties.

Configuring free space on an Encyclopedia volume partition

To prevent possible Encyclopedia volume corruption and operational problems with the machine due to lack of free disk space, iHub shuts down the volume when the disk space available to the partition for the volume falls below a minimum, 128MB by default. The administrator can change the minimum by configuring Min Free Space. Increase Min Free space if iHub, or any other application making use of the same physical disk, consumes disk space at a rapid rate, or if fragmentation occurs.

In Volume—Properties—General, the administrator can configure the minimum free disk space for a partition by typing a new value in Min Free Space for the primary partition. For example, increase the value to 256, as shown in Figure 2-25.

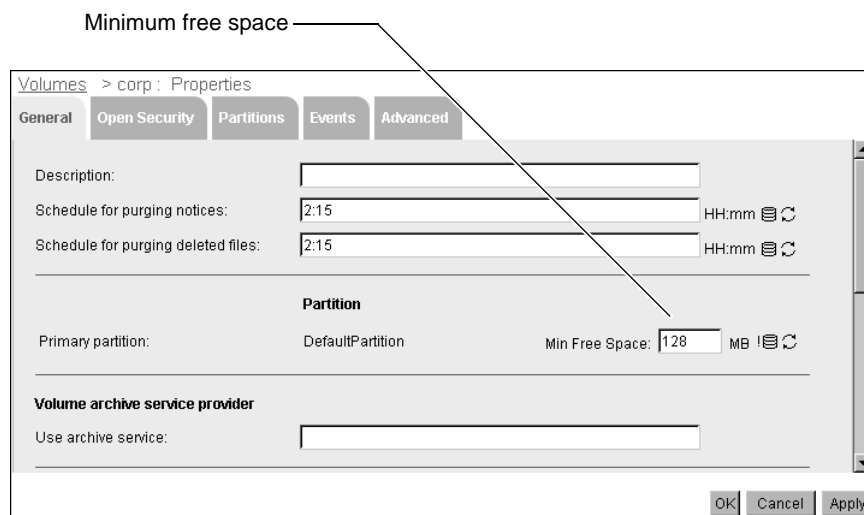


Figure 2-25 Increasing the minimum free space for a partition

Table 2-2 lists these additional properties that appear on Volumes—Properties—Partitions.

Table 2-2 Volume partition properties

Property	Value
Low Free Space	Amount of free space, in megabytes (MB), below which iHub displays a warning message. If a user tries to create a file in a partition with less than the specified low amount of free disk space, file creation succeeds, and a warning message is displayed. The default value is 512 MB.

Table 2-2 Volume partition properties

Property	Value
Min Free Space	Minimum amount, in megabytes (MB), of free space that the Encyclopedia volume partition must maintain. If a user tries to create a file that would put the partition under its minimum free space limit, iHub does not create the file. The default value is 128 MB.

Moving an Encyclopedia volume to another partition

Volumes—Properties—General lists the primary partition for the volume. Volumes—Properties—Partitions, shown in Figure 2-26, lists all the partitions available to the Encyclopedia volume for storing volume files.

The volume uses the primary partition for volume administrative information and to store volume files. To move an Encyclopedia volume to another partition, the administrator must copy the Encyclopedia volume files to the new partition, and then update the old partition path to the new location.




Figure 2-26 Moving an Encyclopedia volume

How to move an Encyclopedia volume

- 1 Set up one or more partitions.
- 2 Put the volume offline.
- 3 For each partition, specify the partition location on Partitions—Template Settings.

- 4 Copy the Encyclopedia volume files from the location specified by the old partition path to the location specified by the new partition path. In a cluster, all machines must be able to access the path specified by the partition.
- 5 Take the volume online.

Supporting application-level partitioning

In a cluster, you can use Encyclopedia volumes to support application-level partitioning by configuring different applications to use different Encyclopedia volumes. In some cases, such as when an application service provider hosts services, a separate physical structure is part of the logical design. For example, business needs can require the separation of billing information from sales information. To use multiple Encyclopedia volumes without changing the logical design of the application, an application designer can use links in the application folder structure to redirect users to the appropriate Encyclopedia volume.

To make two Encyclopedia volumes work as a single unit, you need a single source of user information, such as user names, passwords, and role membership. The application designer can externalize Encyclopedia volume user information using the Actuate Open Security feature and centralize the user information in an external security source. For example, iHub can store the Encyclopedia volume user information in an LDAP server. Multiple Encyclopedia volumes can use the information from the LDAP server.

Specifying a hyperlink in a PDF document

The context string in which a job runs is not available in the background job that generates a PDF document. To specify a hyperlink, such as a drill-through hyperlink, containing the appropriate Information Console context string in a PDF document, the administrator must specify the Information Console URL prefix by performing one of the following tasks:

- In Configuration Console, update the setting for Information Console URL prefix in Volumes→Properties.
- In `acserverconfig.xml`, manually edit the setting for `IPortalURLPrefix`.

How to update the setting for Information Console URL prefix

- 1 Open Configuration Console. Choose Advanced view.
- 2 In Advanced view, navigate to Volumes→Properties.
- 3 In Information Console URL prefix, type the context string, as shown in Figure 2-27. For example, type:

```
http://urup:8900/iportal
```


4 Choose OK.

Configuration Console updates the setting for the system variable, IPortalURLPrefix, in acserverconfig.xml in AC_DATA_HOME\config\iHub2.

Volumes > corp : Properties

General Open Security Partitions Events Advanced

Description:

Schedule for purging notices: HH:mm

Schedule for purging deleted files: HH:mm

Partition

Primary partition: DefaultPartition Min Free Space: MB

Volume archive service provider

Use archive service:

Metadata database and schema

Metadata database name: Default_ActuatePostgreSQL_MetadataDatabase

Database schema name: urup

Email notification

E-mail notification template partition:

Use Information Console for e-mail notifications ☒

Information Console URL prefix:

Figure 2-27 Specifying Information Console URL prefix

How to manually edit the setting for IPortalURLPrefix

- 1 In Windows Explorer, navigate to AC_DATA_HOME\config\iHub2.
- 2 Open acserverconfig.xml in a text editor.
- 3 In acserverconfig.xml, edit the <Volume> setting, IPortalURLprefix, to contain the appropriate Information Console context string, as shown in Listing 2-1. For example, type:

```
IPortalURLPrefix="http://urup:8900/portal"
```
- 4 Save the file.

Listing 2-1 shows an example of an edited setting for IPortalURLprefix in acserverconfig.xml.

Listing 2-1 Creating a .bat file to point iHub to portserv.exe

```
<Volumes>
  <Volume
    Name="urup"
    EmailURLType="IPortal"
    RSSESOAPPort="8900"
    PrimaryServer="$$ServerName$$"
    IPortalURLPrefix="http://urup:8900/portal"
```

It is not necessary to restart iHub after changing the setting for IPortalURLPrefix. In AC_SERVER_HOME/etc, acmetadescription.xml specifies that any change to IPortalURLPrefix takes effect immediately, as shown in the following code excerpt:

```
<Variable
  Name="IPortalURLPrefix"
  Type="String"
  Access="Public"
  Required="False"
  DisplayName="IPortal URL Prefix"
  TakesEffect="Immediate"/>
```

Configuring events

In Volumes—Properties—Events, you set values used with event schedules, enable custom event schedules, and specify the web service that a custom event uses. Figure 2-28 shows Volumes—Properties—Events.

The screenshot shows the 'Events' tab of the 'Volumes > corp : Properties' dialog box. The 'Polling' section has three input fields: 'Polling interval' set to 5, 'Polling duration' set to 300, and 'Lag time' set to 60, each with a 'min !' unit indicator. Below this, the 'Enable custom events' checkbox is checked. The 'Custom event web service configuration' section has three input fields: 'IP address' set to 'localhost', 'Soap port' set to '8900', and 'Context string' set to '/acevent/servlet/AxisServlet'. Each of these fields has a refresh icon (circular arrow) to its right. At the bottom of the configuration area, there is a note: 'These fields require volume restart to take effect' with a refresh icon, and a sub-note: '(!) These fields will take default value if left blank'. At the very bottom of the dialog box are 'OK', 'Cancel', and 'Apply' buttons.

Figure 2-28 Specifying event-based job configuration values

When you deploy a BIRT design to an Encyclopedia volume, you must provide iHub with access to Java classes that the design uses. You package these classes as JAR files that a BIRT iHub Java factory process accesses. There are two ways to deploy Java classes:

- Deploy the JAR files to the Encyclopedia volume.
Supports creating specific implementations for each volume in iHub. This method of deployment requires packaging the Java classes as a JAR file and attaching the JAR file as a resource to the design file. You treat a JAR file as a resource in the same way as a library or image. Using this method, you publish the JAR file to iHub every time you make a change in the Java classes.
- Deploy the JAR files to the following iHub subdirectory:
AC_SERVER_HOME\resources
This method uses the same implementation for all volumes in BIRT iHub. You do not have to add the JAR file to the design Resource property. Deploying JAR files to an iHub /resources folder has the following disadvantages:
 - You must restart iHub after deploying the JAR file.
 - The JAR file is shared across all volumes. This can cause conflicts if you need to have different implementations for different volumes.

Table 2-3 describes the properties that appear on Volumes—Properties—Events.

Table 2-3 Event-based job configuration properties

Property	Value
Polling interval	The frequency in minutes that iHub checks for a system event.
Polling duration	The duration in minutes that iHub checks for an event. If the event does not occur within the allotted time, iHub marks it as expired. A user can customize this value when creating an event-driven schedule. This value applies to all types of system events.
Lag time	The minutes that iHub scans for completed jobs to determine if an event occurred. For example, if you submit an event-based schedule with the default event lag time, iHub checks the status of jobs for the previous 60 minutes. If the event occurred within the previous 60 minutes, it sets the event status to satisfied.

(continues)

Table 2-3 Event-based job configuration properties (continued)

Property	Value
Enable custom events	A flag that enables iHub custom event processing for a scheduled job. If the value is true, the service is enabled. If you change the value to false, all the existing instances of scheduled jobs using the custom event fail. This configuration value also affects the EnableEventService property value in the Actuate IDAPI GetVolumeProperties response.
IP address	The server name or IP address where the custom event service resides. The default value is localhost.
Soap port	The number of a valid, used port for the custom event service. iHub uses an application container to host web services applications.
Context string	The context string of the request URL for sending a message to the custom event service. Default value is /acevent/servlet/AxisServlet.

Configuring advanced volume properties

Volumes—Properties—Advanced contains a variety of miscellaneous volume property settings that an administrator can use to make performance-tuning adjustments for a mix of features. Figure 2-29 shows the advanced properties list.

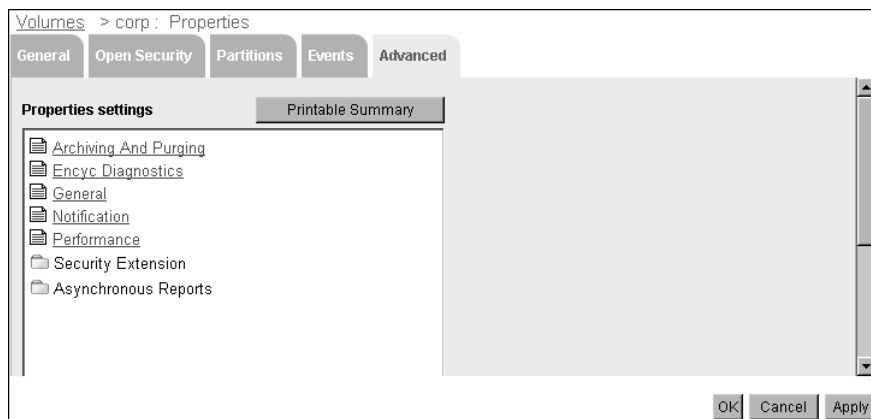


Figure 2-29 Specifying advanced Encyclopedia volume property settings

Printing a summary of advanced volume properties

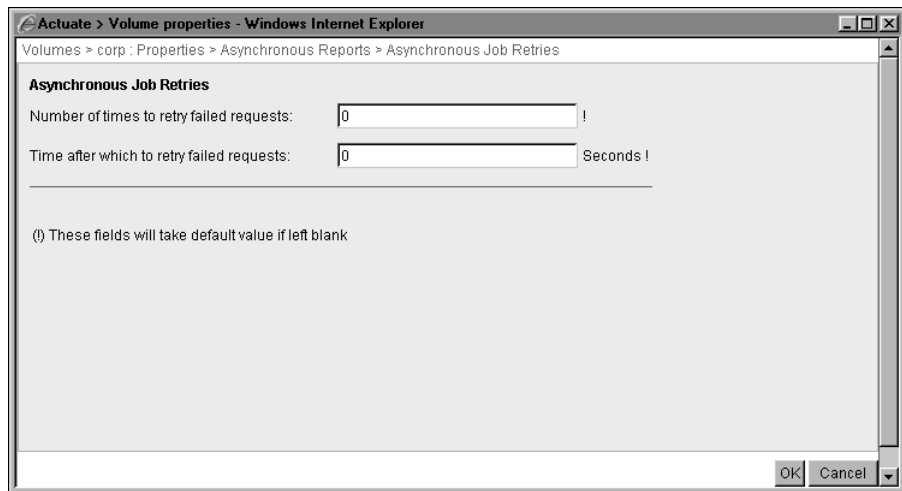
In Volumes—Properties—Advanced, you can view or print a summary of Encyclopedia volume properties and their values. First, select the category

properties, then choose Printable Summary to print the information. You can select one of following categories of properties:

- Archiving And Purging
- Encyc Diagnostics
- General
- Notification
- Performance
- Security Extension
- Asynchronous Reports

Retrying failed asynchronous jobs

In Volumes—Properties—Advanced—Asynchronous Reports—Asynchronous Job Retries, specify how to retry running a scheduled job, as shown in Figure 2-30.



The screenshot shows a web browser window titled 'Actuate > Volume properties - Windows Internet Explorer'. The address bar shows the path: 'Volumes > corp : Properties > Asynchronous Reports > Asynchronous Job Retries'. The main content area is titled 'Asynchronous Job Retries' and contains two input fields. The first field is labeled 'Number of times to retry failed requests:' and has a value of '0'. The second field is labeled 'Time after which to retry failed requests:' and has a value of '0', followed by the text 'Seconds !'. Below these fields is a note: '(!) These fields will take default value if left blank'. At the bottom right of the window are 'OK' and 'Cancel' buttons.

Figure 2-30 Specifying values for retrying failed asynchronous jobs

If the job retry options are set to retry a job if it fails, the job remains active if the node the job is running on fails. For example, if the node crashes, iHub tries to run the job again when the node restarts.

Changing the status of an Encyclopedia volume

The status of an Encyclopedia volume can be disabled, enabled, offline, or online. You change the status of a volume on Volumes, in Advanced view of Configuration Console.

Taking an Encyclopedia volume offline

To disable or remove a volume, you must first take it offline.

How to take an Encyclopedia volume offline

- 1 From Advanced view of Configuration Console, choose Volumes.
- 2 In Volumes, point to the arrow next to a volume name and choose Take offline, as shown in Figure 2-31.

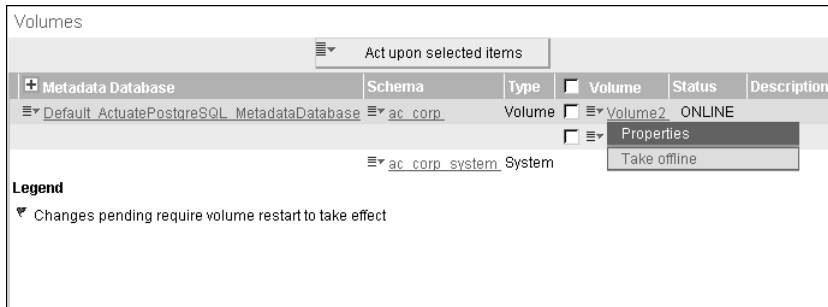


Figure 2-31 Taking a volume online or taking a volume offline

- 3 On Volume Offline Grace Period, choose OK, as shown in Figure 2-32. You can adjust the time a volume takes to go offline in Grace period.




Figure 2-32 Choosing OK to take the volume offline

Disabling an Encyclopedia volume

The difference between an offline volume and a disabled volume is that iHub takes an offline volume online when you start the iHub system or the iHub service. A disabled volume does not come online when you start the iHub system or service. The volume remains disabled.

How to disable a volume

- 1 In Volumes, point to the icon next to a volume name and choose Take offline.
- 2 Point to the icon next to the offline volume name and choose Disable, as shown in Figure 2-33. Confirm that you want to disable the volume.




Figure 2-33 Disabling a volume

The volume status changes to DISABLED.

Enabling an Encyclopedia volume

Before you can take a disabled volume online, you must enable it. To enable a volume, point to the icon next to a disabled volume name and choose Enable, as shown in Figure 2-34.

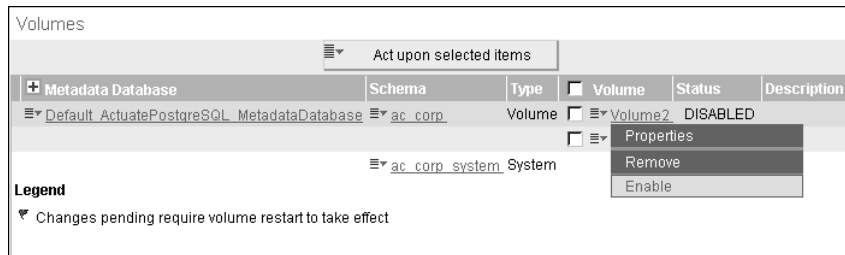


Figure 2-34 Enabling a volume

The volume status changes to OFFLINE.

Taking an Encyclopedia volume online

When the status of a volume is OFFLINE, you can take the volume online. To be available to users, an Encyclopedia volume must be online. To take a volume online, on Volumes, point to the icon next to the offline volume name and choose Take online, as shown in Figure 2-35.




Figure 2-35 Taking a volume online

Removing an Encyclopedia volume

When you remove an Encyclopedia volume from iHub, the following changes occur:

- Encyclopedia volume users no longer see the volume in the list of volumes available to iHub or a cluster.
- Encyclopedia volume users cannot log in to the Encyclopedia volume.
- iHub rejects requests to access the content of the Encyclopedia volume.
- The volume no longer appears in the configuration file.
- iHub changes an Encyclopedia volume's resource group volume assignment to All volumes and disables the resource group.

Deleting an Encyclopedia volume does not delete the partitions assigned to the volume.

How to remove an Encyclopedia volume from iHub

To completely remove an Encyclopedia volume from iHub, perform the following tasks:

- 1 Before removing an Encyclopedia volume, note which partitions the volume uses.
 - 1 From Advanced view of Configuration Console, choose Volumes.
 - 2 On Volumes, point to the icon next to a volume name and choose Properties.
 - 3 In Volumes—Properties, choose Partitions.
 - 4 In Partitions, note the list of partitions that appears in Selected partitions. From the side menu, choose Volumes.
- 2 Point to the icon next to a volume name and choose Take offline.

In Volume Offline Grace Period, choose OK.

The volume's status changes to OFFLINE.
- 3 On Volumes, point to the icon next to a volume name and choose Remove, as shown in Figure 2-36.

A confirmation dialog box appears.
- 4 To delete the Encyclopedia volume from the machine, choose OK.

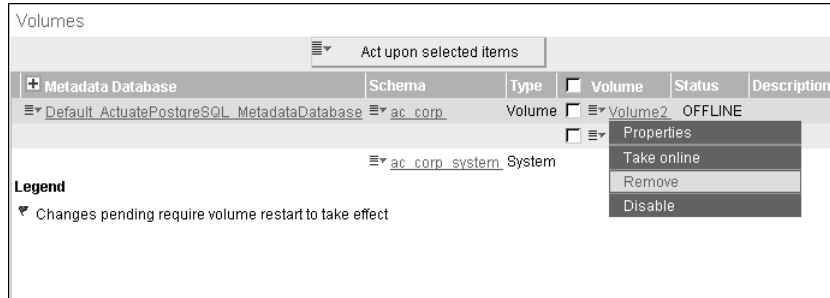


Figure 2-36 Removing an Encyclopedia volume

Renaming the default Encyclopedia volume

The login page of Management Console lists the default Encyclopedia volume. In System—Properties—General, change the System default volume, as shown in Figure 2-37.

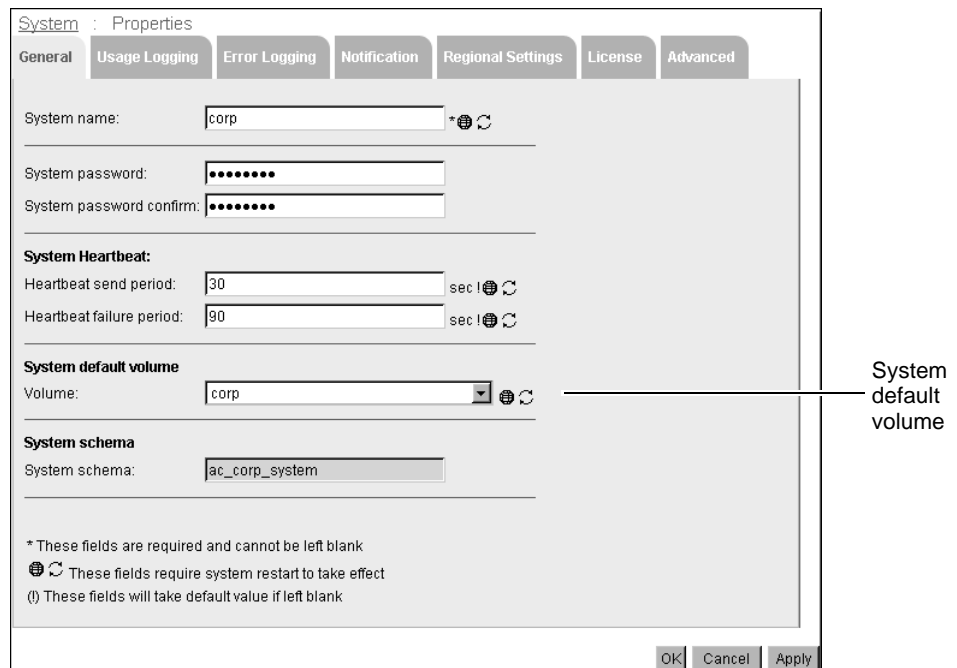


Figure 2-37 Specifying the default Encyclopedia volume

Removing a partition

Before you remove a partition, you must check the list of partitions in Configuration Console to see which partitions are available and which are already assigned to Encyclopedia volumes. Assign a different partition to the Encyclopedia volume to which the partition is assigned, or back up the Encyclopedia volume.

If the partition is a secondary partition, you delete the partition by changing its configuration. iHub starts a background process of moving files from one partition to another. The operation fails if the other partitions drop below their minimum free space limits. When the process finishes, the state of the secondary partition changes to inactive, and you can remove it.

How to remove a partition

- 1 From Advanced view of Configuration Console, choose Volumes.
- 2 On Volumes, point to the arrow next to the Encyclopedia volume that uses the partition and choose Properties.
- 3 In Volumes—Properties, choose Partitions.
- 4 On Volumes—Properties—Partitions, in Selected partitions, select the partition to remove. Select Stop. Choose OK.

iHub moves data from the partition to other available volume partitions. If there is not enough room to store the data from the partition being deleted, a failure message appears.

- 5 After iHub finishes phasing out the partition, repeat steps 2 and 3.
- 6 On Volumes—Properties—Partitions, in Selected partitions, select the partition to remove. Choose the left arrow to move the partition from Selected partitions to Available partitions.
- 7 Choose Partitions from the side menu.
- 8 In Partitions, confirm that the partition is not assigned to an Encyclopedia volume and that it is unused. Then, point to the arrow next to the partition name to display the drop-down list for the partition.
- 9 Choose Delete. To confirm the deletion, choose OK, as shown in Figure 2-38.




Figure 2-38 Deleting a partition

Using diagnostic, usage, and error logging

This chapter contains the following topics:

- Configuring diagnostic logging
- Configuring usage and error logging

Configuring diagnostic logging

iHub performs various types of diagnostic logging with configurable levels of detail. An administrator can experiment with the available levels to collect information on the iHub system to detect problems and improve performance. Setting log files to higher levels of information collection increases disk space usage and can decrease iHub performance.

The administrator configures diagnostic logging for iHub process by choosing Change in Server Configuration Templates—Settings, as shown in Figure 3-1.



Figure 3-1 Changing diagnostic logging settings

Server Configuration Templates—Settings allows the administrator to specify the following diagnostic logging settings, as shown in Figure 3-2:

- Enable logging
Enable or disable diagnostic logging.
- Log level
Level of detail in the log.

- Log directory
Location of the log.
- Log size
Maximum size of the log before iHub starts overwriting the log.
- Number of log files
Maximum number of logs that iHub creates.

Actuate > Server properties - Windows Internet Explorer

Server Configuration Templates > orup : Template properties

General log ☒ Enable
 Level: Warning
 Directory: \$AC_DATA_HOME\$/server/log
 Size: 10000 KB
 Number of log files: 3

Factory log ☒ Enable
 Level: Warning
 Directory: \$AC_DATA_HOME\$/server/log
 Size: 10000 KB
 Number of log files: 3

Integration log ☒ Enable
 Level: Warning
 Directory: \$AC_DATA_HOME\$/server/log
 Size: 10000 KB
 Number of log files: 3

Caching log ☒ Enable
 Level: Warning
 Directory: \$AC_DATA_HOME\$/server/log
 Size: 10000 KB
 Number of log files: 3

Viewing log ☒ Enable
 Level: Warning
 Directory: \$AC_DATA_HOME\$/server/log
 Size: 10000 KB
 Number of log files: 3

☐ These fields require server restart to take effect
 (!) These fields will take default value if left blank

OK Cancel

Figure 3-2 Configuring diagnostic logging

By default, diagnostic logging is enabled for all iHub processes. The administrator can disable diagnostic logging for a process by deselecting Enable.

Changing the default log file location in acserverconfig.xml

iHub writes log entries by default to AC_DATA_HOME\server\log. The log location is no longer specified in acserverconfig.xml.

To specify a different location or manually update a previously edited log setting, the administrator must modify acserverconfig.xml to contain a LogDirectory attribute that points to a non-default location by performing the following tasks:

- 1 Open AC_DATA_HOME\config\iHub2\acserverconfig.xml in a text editor.
- 2 Insert or edit a LogDirectory attribute that specifies the path to a different log folder.

For example, in Listing 3-1, an inserted LogDirectory attribute specifies the non-default log location, C:\Actuate\iHub\log, for the following properties:

- Template
- Reporting service
- Viewing service
- Integration service
- Caching service

Listing 3-1 acserverconfig.xml edited log file location settings

```
<Template
  Name="urup"
  PMDPort="8100"
  ActuateBuild="220A130126"
  LogDirectory="C:\Actuate\iHub\log"
  ActuateVersion="2"
  ServerSOAPPort="11100"
  StartArguments="-Xmx2048M -Xms1024M -XX:MaxPermSize=128m
  com.actuate.iserver.server.Server"
  AppContainerPort="8900"
  <ReportingService
    LogDirectory="C:\Actuate\iHub\log"
    EnableGenerationService="true"/>
  <SOAPDispatchService
    EnableRequestService="true"
    SOAPDispatchSOAPPort="8000"/>
  <ViewingService
    EnableRender="false"
    LogDirectory="C:\Actuate\iHub\log"
```

```

        EnableViewingService="true"
        MSAnalysisTunnelPort="9021"/>
    <IntegrationService
        LogDirectory="C:\Actuate\iHub\log"
        EnableIntegrationService="true"/>
    <CachingService
        LogDirectory="C:\Actuate\iHub\log"
        EnableCachingService="true"/>
    ...
</Template>

```

Configuring additional iHub diagnostic logging properties

In Server Configuration Templates—Settings, the administrator can also configure the following additional iHub property settings, as shown in Figure 3-3:

- iHub diagnostic logging
- Logging for Encyclopedia database queries

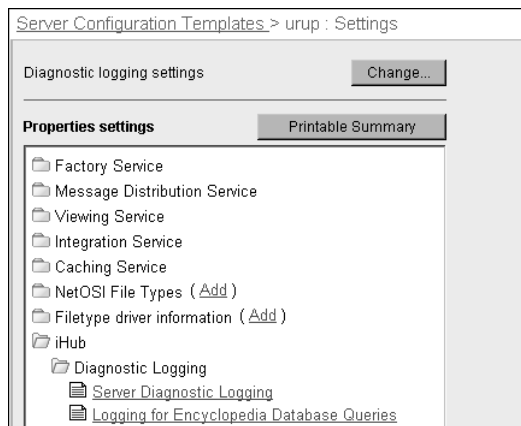


Figure 3-3 Choosing iHub diagnostic logging

About iHub diagnostic logging

The following parameters define the levels of iHub diagnostic logging detail you can specify.

How to configure additional diagnostic logging settings

- 1 In Server Configuration Templates, choose the name of a server template, as shown in Figure 3-4.




Figure 3-4 Choosing a server template name

- 2 On Server Configuration Templates—Settings, expand iHub, expand Diagnostic Logging, then choose Server Diagnostic Logging, as shown in Figure 3-3.
- 3 On Server Diagnostic Logging, accept the default value for each setting, as shown in Figure 3-5, or alternatively, specify a new value.

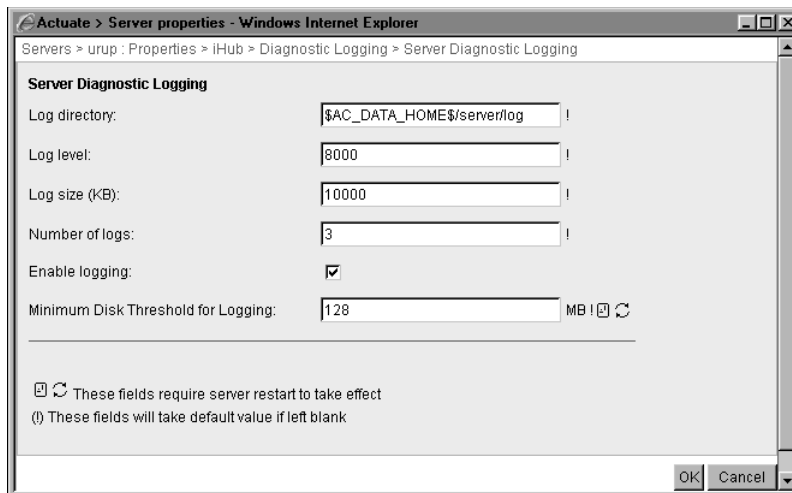


Figure 3-5 Specifying server diagnostic logging property settings

Choose OK.

- 4 If you change Minimum Disk Threshold for Logging, restart iHub.

About logging for Encyclopedia database queries

In Logging for Encyclopedia Database Queries, the administrator can specify the level of diagnostic log detail that iHub provides when an SQL statement fails to execute successfully. The following parameters define the levels of logging detail you can specify:

- **LogSQLQuery**
iHub logs the query sent to the database.

- **LogQueryPrepareTime**
iHub logs the time spent doing a prepare against the database.
- **LogQueryExecutionTime**
iHub logs the time spent executing the query against the database.

How to configure logging for Encyclopedia database queries

- 1 In Server Configuration Templates, choose the name of a server template, as shown in Figure 3-6.




Figure 3-6 Choosing a server template name

- 2 On Server Configuration Templates—Settings, expand iHub, Diagnostic Logging, then choose Logging for Encyclopedia Database Queries, as shown in Figure 3-7.

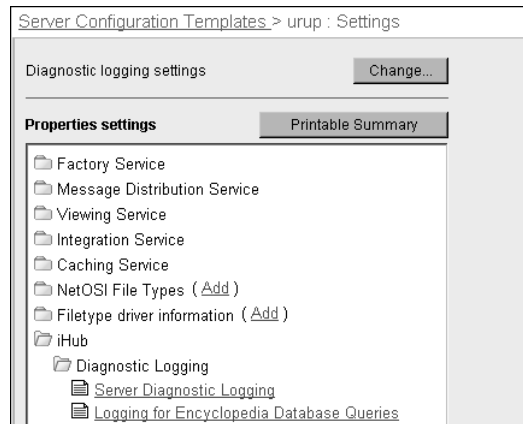


Figure 3-7 Choosing Logging for Encyclopedia Database Queries

- 3 In Logging for Encyclopedia Database Queries, select an option from Select a specific option and move it to Diagnostic Logging for Encyclopedia Database Queries, using the right arrow, as shown in Figure 3-8.

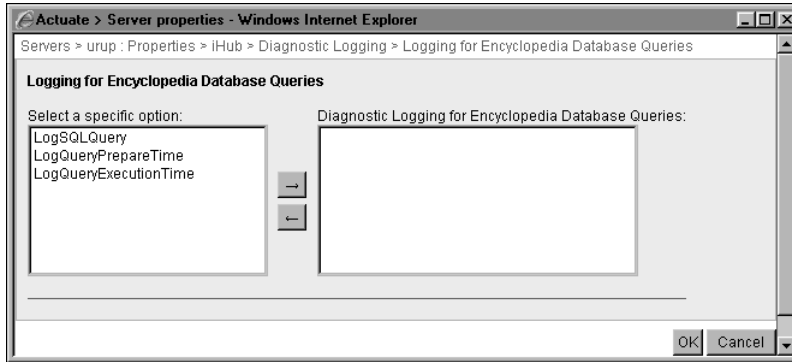


Figure 3-8 Choosing a logging level

Choose OK.

For more information about diagnostic logging properties, see “Configuring diagnostic logging” in Chapter 1, “Performing basic configuration tasks.”

The environment variable `AC_SERVER_FILE_CREATION_MASK` on Linux, and registry key on Windows, specify the permissions for log files and directories in `/UsageErrorLogs` and `/iHub/log`. The default value in Linux is owner read-write (066 UMASK), and in Windows is 0 (read-only).

Configuring usage and error logging

iHub can capture usage monitoring and error information in log files. To understand resource usage and troubleshoot problems, use usage and error logging.

In Advanced view, choose System. Choose Properties, then choose Usage Logging. Figure 3-9 shows Usage Logging.

The following Actuate usage and error logging applications that ship with iHub record usage activity and error information in a log file:

- **Usage Logging Extension**
Writes the usage information to a comma-delimited (.csv) file, `usage_log.csv`
- **Error Logging Extension**
Writes the usage information to a comma-delimited (.csv) file, `error_log.csv`

The usage logging and error logging applications are open framework applications. These applications are DLLs in Windows and shared libraries in Linux. By default, the usage and error logging extensions reside in the following location:

`AC_SERVER_HOME/bin`

iHub creates a primary log directory that contains the usage and error log records for the default volume in the following location:

`AC_SERVER_HOME/UsageErrorLogs/primary`

iHub creates secondary log directories for additional volumes as required in the following location:

`AC_SERVER_HOME/UsageErrorLogs/secondary_$(VOLUMENAME)`

The directories for usage and error log files are not configurable.

In an iHub cluster, the asynchronous design execution log entry is in the log of the Encyclopedia process node. For a synchronous or transient document generation, the log entry is in the log file of the node that runs the design. Before you configure error logging for a cluster, ensure that each node in a cluster has the same logging applications installed.

A developer can customize the way the DLL or shared library handles the usage and error log information. iHub Integration Technology provides a reference implementation for this purpose.

iHub supports consolidating log information into a database. iHub Integration Technology provides a reference implementation for the log consolidator application. For more information, see *Using BIRT iHub Integration Technology*.

Configuring usage logging

Configure usage log settings for the following user operations:

- Document viewing
A viewing log records document-viewing events that use the View process.
- Document printing
A printing log records document-printing events.
- Document generation
A Factory log records document-generation events.
- Document deletion
A deletion log records internal deletion operations, such as deletions that occur from archiving volume files.
- Admin
An admin log records Encyclopedia volume management user operations, such as:
 - Managing users, roles, groups, and channels
 - Uploading and deleting files
 - Changing file and folder properties

- Data integration
A data integration log records AIS events, such as information object jobs.

How to configure usage logging

To configure usage logging, perform the following tasks:

- 1 Log in to Configuration Console and choose Advanced view.
- 2 In System—Status, choose Properties.
System—Properties—General appears.
- 3 Choose Usage Logging.

System—Properties—Usage Logging appears, as shown in Figure 3-9.

System : Properties

General Usage Logging Error Logging Notification Regional Settings License Advanced

Viewing log: ☐ Enable
Logging level: Standard

Printing log: ☐ Enable
Logging level: Standard

Factory log: ☐ Enable
Logging level: Standard

Deletion log: ☐ Enable
Logging level: Standard

Admin log: ☐ Enable
Logging level: Standard

Data Integration log: ☐ Enable
Logging level: Standard

Usage logging extension name: UsrActivityLoggingExt

These fields require system restart to take effect
(!) These fields will take default value if left blank

OK Cancel Apply

Figure 3-9 Configuring usage logging

- 4 On Usage Logging, select the usage logging information you want to capture from the following list of logging options:
 - Viewing
 - Printing

- Factory
- Deletion
- Admin
- Data Integration

1 Select Enable to activate the logging option.

2 Select Standard or Detail for the logging level.

For viewing, deletion, and printing logging, standard and detail information are the same in the logging application that ships with iHub.

For Factory logging, detailed information includes design parameters. Logging detailed Factory information, instead of standard Factory information, decreases performance.

5 In Usage logging extension name, enter the name of the usage logging extension.

UsrActivityLoggingExt is the name of the default usage logging extension. Do not append DLL.

Choose OK.

About types of recorded events

For each type of event, you can set the logging level to Standard or Detail. If you are using the default usage logging extension, UsrActivityLoggingExt, the logging level does not affect how the file records the following types of events:

- Document viewing
- Document printing
- Document deletion

If you set the logging level for document generation or factory events to Detail, the usage log includes design parameters. Setting the logging level to Detail for document generation events decreases performance.

Before you set up system usage for a cluster, you must ensure that each node in a cluster has the same logging applications installed.

About the usage logging extension

You can use the default usage logging extension, UsrActivityLoggingExt, or you can specify the name of your custom DLL or shared library that generates the log files. The default location of the usage logging extension is AC_SERVER_HOME /bin. Type the name of the DLL or shared library without the file-name extension. For example, on a Windows system, if the DLL is CustomUsage.dll, type:

CustomUsage

Understanding a usage log entry

Each usage log entry is a comma-separated list containing up to 40 fields of information about an event. The following example describes a delete user event:

```
3272649170,5,1,3272649170,3272649170,-,-,0,Administrator,3,
urup,urup,urup,User,testUser,-,-,-,-,-,-,-,-,-,
2,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
```

A dash in the usage log entry means the field is not used. The usage log organizes the entry fields into the following information groups:

- Fields 1 through 10 contain general information:
 - Fields 1, 4, and 5 contain the log file time stamp, start time, and finish time. The time is in seconds since 00:00:00, Jan. 1, 1901, GMT.
 - Field 2 contains the event type. The numeric values in Table 3-1 indicate the event types.

Table 3-1 Event types and the corresponding event values

Event type	Event value
DocumentGeneration	1
DocumentPrinting	2
DocumentViewing	3
DocumentDeletion	4
Admin	5
Query	6
Search	7

- Field 3 contains the event result. The event result value is either 1 or 0, indicating success or failure.
- Fields 6 through 8 contain document output information, indicating the file name, version, and file size. The document output group information appears only with document events.
- Fields 9 and 10 contain execution information, indicating the user name and the iHub subsystem where the operation executed. The numeric values in Table 3-2 indicate the iHub subsystems.

Table 3-2 iHub subsystems and the corresponding ID numbers

Subsystem	ID number
ReportEngine	1
ViewEngine	2

Table 3-2 iHub subsystems and the corresponding ID numbers

Subsystem	ID number
EncycEngine	3
IntegrationEngine	4
Cache	5

- Fields 11 through 25 contain operational information in string format, including the Encyclopedia volume, iHub, and cluster names. Fields 26 through 40 contain operational information in numeric format.

The values in these fields depend on the value for the event type in field 2. Table 3-3 summarizes some of the information available for each event type at Standard level.

Table 3-3 Examples of information that is available about the different types of events

Event type	Event value	Operation data available
Document generation	1	<p>String fields 11 through 21 display the following information:</p> <ul style="list-style-type: none">- ,executable name, executable version, volume name, server name, cluster name, resource group name, node running request, page count, job name, request ID <p>Numeric fields 26 through 29 display the following information:</p> <p>number of pages,submit time, job type, job priority</p>
Document printing	2	<p>String fields 11 through 18 display the following information:</p> <p>page numbers printed, volume name, printer name, server name, clustername, node sent to, file type, server request id</p> <p>Numeric fields 26 through 29 display the following information:</p> <p>number of pages printed, submit time, job type, job priority</p>

(continues)

Table 3-3 Examples of information that is available about the different types of events (continued)

Event type	Event value	Operation data available
Document viewing	3	String fields 11 through 18 display the following information: output format, page numbers, volume name, server name, cluster name Numeric field 26 displays the number of pages viewed.
Administrative	5	String fields 11 through 13 display the following information: volume name, server name, cluster name Numeric field 26 displays an operation ID for an administration event. The following list provides the event name for each operation ID: <ul style="list-style-type: none"> ■ 1 Create ■ 2 Delete ■ 3 Modify ■ 4 Login
Actuate Integration service	6	String fields 11 through 14 display the following information: volume name, server name, cluster name, server request id Numeric fields 26 and 27 display the following information: request wait time, request generation time
Search	7	String fields 11 through 15 display the following information: document format, page numbers, volume name, server name, cluster name Numeric field 26 displays the number of pages viewed.

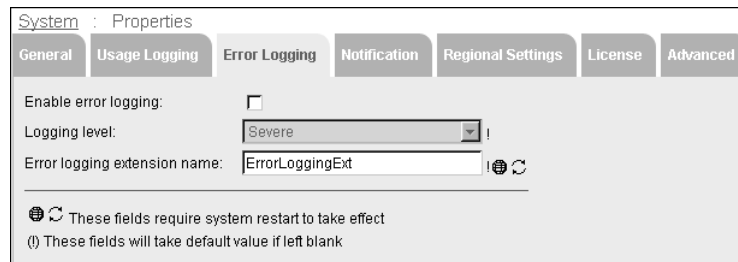
Configuring error logging

The error log, `error_log.csv`, is a comma-separated values (CSV) file. If you use the default error logging extension, `ErrorLoggingExt`, you can set the logging level to the following settings:

- **Information**
The error log records messages that trace iHub behavior.
- **Warning**
The error log records warnings. The errors do not necessarily affect the operation of iHub.
- **Severe**
The error log records errors that can result in iHub failure if you do not correct them.
- **Fatal**
The error log records critical errors from which iHub cannot recover and that can result in failure.

How to configure error logging

- 1 Log in to Configuration Console and choose Advanced view.
- 2 In System—Status, choose Properties.
System—Properties—General appears.
- 3 Choose Error Logging.
System—Properties—Error Logging appears, as shown in Figure 3-10.



The screenshot shows the 'System : Properties' window with the 'Error Logging' tab selected. The 'Enable error logging' checkbox is unchecked. The 'Logging level' dropdown is set to 'Severe'. The 'Error logging extension name' text box contains 'ErrorLoggingExt'. Below these fields, a note states: 'These fields require system restart to take effect'. A smaller note below that says: '(!) These fields will take default value if left blank'.

Figure 3-10 Configuring error logging

- 4 Select Enable error logging.
- 5 Select the error logging level you want to capture from the following list of options:
 - **Information**
 - **Warning**

- Severe
 - Fatal
- 6** In Error logging extension name, enter the name of the error logging extension.
- ErrorLoggingExt is the name of the default error logging extension.
- Choose OK.

Understanding an error log entry

Each error log entry is a comma-separated list containing up to 12 fields about an error-related event. The following example describes an error in a submit job event:

```
3272648796,2,3230,SubmitJob,Administrator,"Invalid start time or
end time.",urup,urup,urup,-,-,-
```

A dash in the usage log entry means the field is not used. The error log organizes the entry fields into the following information groups:

- Fields 1 through 9 contain general information:
 - Field 1 contains the log file time stamp. The time is in seconds since 00:00:00, Jan. 1, 1901, GMT.
 - Field 2 contains the error severity level, an integer between 1 and 4. The numeric values in Table 3-4 indicate the level.

Table 3-4 Error severity levels and the corresponding values

Error severity level	Value
Information	1
Warning	2
Severe	3
Fatal	4

- Field 3 contains the error ID code.
- Field 4 contains the service name, indicating the subsystem where the error occurred such as the Factory, Encyclopedia, View, or Request service.
- Field 5 indicates the Encyclopedia volume user.
- Field 6 contains the error message.
- Field 7 contains the Encyclopedia volume name.
- Field 8 contains the iHub cluster name.

- Field 9 contains the iHub node name.
- Depending on the error, fields 10 through 12 can contain information such as a file name and ID number.

Table 3-5 summarizes some of the information available in fields 10 through 12 for an error log entry at standard level.

Table 3-5 Information that is available for error log entries at the standard level

Type of error	Operation data available
Encyclopedia volume user activity	Fields 10 through 12 can contain error parameters such as the following items: <ul style="list-style-type: none"> ■ Object name ■ ID number
Volume online or offline	Fields 10 and 11 contain the following data: <ul style="list-style-type: none"> ■ Volume name ■ Operation type either online or offline
iHub node start or stop	Field 10 contains the iHub name.
Service enable or disable	Fields 10 and 11 contain the following data: <ul style="list-style-type: none"> ■ Server name ■ List of services
Archive service error	Fields 10 through 12 contain error parameters.
Encyclopedia volume job purging Field 4 is Job Purge	Fields 10 through 12 contain error parameters.
Encyclopedia volume health monitoring Field 4 is Encyclopedia Health Monitor	Fields 10 through 12 contain error parameters.

Table 3-6 lists the general categories of iHub error messages.

Table 3-6 Categories of iHub error messages

Error ID range	Error description
0001 - 1000	System errors such as Out of memory or Low thread count
1001 - 3000	iHub errors such as Corrupt encyclopedia or Transient storage full Within this error category, the following sub-categories exist: <ul style="list-style-type: none"> ■ 1001 - 2000 Actuate internal datastore ■ 2001 - 3000 Actuate internal

(continues)

Table 3-6 Categories of iHub error messages (continued)

Error ID range	Error description
3001 - 6000	User errors such as Permission denied Within this error category, the following sub-categories exist: <ul style="list-style-type: none">■ 3001 - 4000 Encyclopedia engine■ 4001 - 5000 Report engine■ 5001 - 6000 View engine
6001 - 12000	<ul style="list-style-type: none">■ 6001 - 7000 SOAP engine■ 7001 - 8000 Process management daemon■ 8001 - 9000 Cluster engine■ 10001 - 11000 Server configuration■ 11001 - 12000 XML parsing
12001 - 13000	Viewing server errors
13000 - 14000	AcMail exceptions
100001 - 100600	Actuate Information service
100601 - 100699	Actuate Caching service
100700 - 150000	Shared by Actuate Information service and Actuate Caching service

Configuring usage and error logging file settings

The administrator configures usage and error logging files in System—Properties—Advanced—Usage and Error Logging—Usage And Error Log File Settings, shown in Figure 3-11.

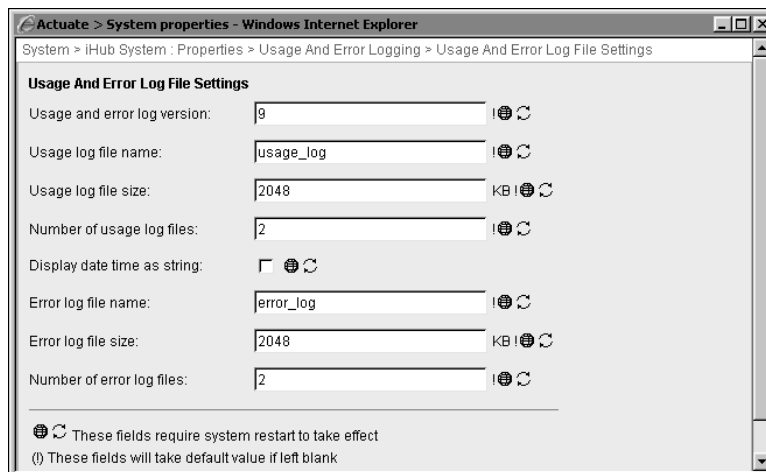


Figure 3-11 Specifying settings for usage and error log files

Table 3-7 describes the properties that appear on System—Properties—Advanced—Usage and Error Logging—Usage And Error Log File Settings.

Table 3-7 Usage and error log file parameters

Property name	Parameter name	Description
Usage and error log version	UsageAndErrorLogVersion	Use this setting for backward compatibility.
Usage log file name	UsageLogFile Name	Base name for the usage log file, which iHub sends to the usage logging application. The sample usage logging application places the log file in the log subdirectory that contains the usage log records for the volume. The directory for a usage log file is not configurable.
Usage log file size	UsageLogFileSize	If iHub uses multiple log files, this value is the maximum size of each log file.
Number of usage log files	NumberOfUsageLogFiles	Maximum number of usage log files. The usage logging application uses this value to create log file names, such as usage_log.csv and usage_log_1.csv.
Display date time as string	DateTimeAsString	Format of the date and time field for usage and error log entries. The format is either a string in the format mm/dd/yyyy or an unsigned long that specifies the number of seconds since January 1, 1970. The default value is false. The default value uses the unsigned long format.
Error log file name	ErrorLogFileName	Base name for the error log file. The sample error logging application places the log file in the log subdirectory that contains the error log records for the volume. The directory for a error log file is not configurable.
Error log file size	ErrorLogFileSize	Maximum size of an error log file. If iHub uses multiple log files, this is the maximum size of each log file. The error logging application uses this value. A custom error logging application can use a different value.
Number of error log files	NumberOfErrorLogFiles	Maximum number of error log files.

4

Configuring e-mail notification

This chapter contains the following topics:

- About e-mail notification
- Adding an SMTP server to the iHub environment
- Specifying the Message Distribution service for e-mail notification
- Setting up sendmail e-mail notification
- Configuring the notification list size and To: line
- Handling e-mail notification errors
- Customizing the e-mail message

About e-mail notification

The administrator can configure iHub to send e-mail notification to users and groups about completed jobs. After configuring iHub to send e-mail notifications, users trigger notification by setting an option in the schedule of a job.

By default, iHub uses Simple Mail Transfer Protocol (SMTP) to send the e-mail. iHub must connect to an SMTP server that the administrator configures using Configuration Console. To perform e-mail notification, iHub must have the View service enabled. On Linux, iHub also supports sendmail.

To set up SMTP e-mail notification in Advanced view, the administrator enables SMTP and specifies a maximum message size.

Adding an SMTP server to the iHub environment

When using multiple SMTP mail servers, iHub can balance e-mail loads to improve performance, as described in the following procedure.

How to add an SMTP server to the iHub environment

- 1 In Advanced view of Configuration Console, choose System.
- 2 In System—Status, choose Properties, as shown in Figure 4-1.

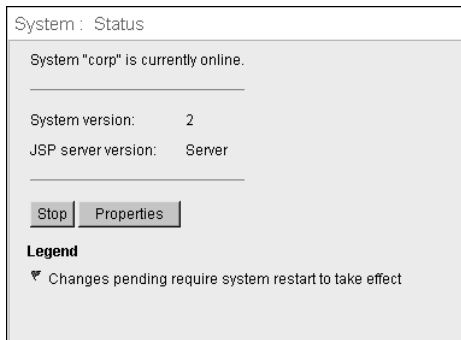


Figure 4-1 Choosing Properties on System—Status

- 3 In Properties, choose Notification.
- 4 In Notification, perform the following tasks:
 - 1 Enable SMTP.
 - 2 Accept the maximum message size, 5120, or specify a different limit for the size of the message and attachment in kilobytes.
 - 3 Choose Add SMTP Server, as shown in Figure 4-2.

System : Properties

General Usage Logging Error Logging **Notification** Regional Settings License Advanced

Maximum mail message size: 5120 KB

☒ Enable SMTP

SMTP Servers Add SMTP Server

Name

These fields require system restart to take effect
 (!) These fields will take default value if left blank

OK Cancel Apply

Figure 4-2 Adding an SMTP server to the iHub environment
 New SMTP Server appears, as shown in Figure 4-3.

SMTP Servers > New SMTP Server

SMTP Server name: *

Hostname or IP Address: *

Listen port: 25 !

Sender e-mail address: *

Sender name: !

SMTP greeting: !

Mailing weight: 100 !

* These fields are required and cannot be left blank
 (!) These fields will take default value if left blank

OK Cancel

Figure 4-3 Setting properties of the SMTP e-mail server

- 5 In New SMTP Server, specify the properties SMTP Server Name through SMTP greeting, as shown in Figure 4-4. For more information about configuring SMTP properties, see “Configuring e-mail notification settings” in Chapter 1, “Performing basic configuration tasks.”

SMTP Servers > New SMTP Server

SMTP Server name: iHub mail server *

Hostname or IP Address: exchangesvr.abcbank.com *

Listen port: 25 !

Sender e-mail address: iHubAdmin@abc.com *

Sender name: iHub Administrator !

SMTP greeting: !

Mailing weight: 100 !

* These fields are required and cannot be left blank
 (!) These fields will take default value if left blank

OK Cancel

Figure 4-4 Specifying properties of a new SMTP server

- 6 In Mailing weight, specify the relative weight to use to determine message routing, then choose OK. If the value is zero, iHub uses the SMTP server only after receiving errors from all other SMTP servers in the SMTP server list. Otherwise, it uses load balancing.

The display name of the new SMTP server that you specified in the first step appears in Notification.

Choose OK.

- 7 Restart iHub System.

Table 4-1 lists the property names that appear in Configuration Console with the corresponding parameter names in AC_SERVER_HOME/etc/acmetadescription.xml, indicating default settings, ranges, and when a property change takes effect.

Table 4-1 SMTP e-mail notification parameters

Property name	Parameter name	Default	Range	Takes effect
Hostname or IP address	SMTPHostName			Immediate
Listen port	SMTPPort	25	1 - 65535	Immediate
Sender display name	SenderName			Immediate
Sender e-mail address	SenderAddress			Immediate
SMTP greeting	Greeting			Immediate
SMTP mailing weight	MailingWeight	100	0 - 1000000	Immediate

Using SMTP server load balancing

iHub uses load balancing to send e-mail notifications through SMTP mail servers based on the availability of processing resources. For example, the administrator configures load balancing by setting the mailing weight parameter values of the first mail server to 40, the second to 30, and the third to 30. The settings establish a relative load-balancing ratio of 4:3:3 among the three mail servers. For every 100 messages iHub distributes, it sends 40 to the first mail server, 30 to the next, and 30 to the last.

When a cluster node receives an error sending an e-mail notice to one mail server, and succeeds in sending the notice to another mail server, the notice counts as part of the load-balancing quota for the mail server that failed. The notice also counts towards the quota of the mail server that succeeds, unless its quota for that round is already exhausted.

Using multiple SMTP servers in a cluster

In a cluster, iHub distributes e-mail notification requests only among the nodes that have the View service enabled. Nodes must have access to the mail servers to send e-mail notices. The example in Figure 4-5 configures two SMTP servers.

System : Properties

General Usage Logging Error Logging Notification Regional Settings License Advanced

Maximum mail message size: 5120 KB ⓘ ⌂

☒ Enable SMTP ⓘ ⌂

SMTP Servers Add SMTP Server

Name		
Actuate Mktg	Edit	Delete
Testing	Edit	Delete

ⓘ ⓘ These fields require system restart to take effect
(!) These fields will take default value if left blank

Figure 4-5 Viewing e-mail notification properties

To add SMTP servers to the iHub environment, the administrator modifies the SMTP server setup.

How to modify the SMTP server setup

- 1 In Configuration Console, in System—Status, choose Properties.
System—Properties appears.
- 2 Choose Notification.
System—Properties—Notification appears.
- 3 To modify the SMTP server setup, perform the following tasks.
 - 1 Specify a maximum message size for messages that iHub sends.
 - 2 Select Enable SMTP to use SMTP e-mail servers for e-mail notification, if necessary. Deselect this option to disable the use of SMTP e-mail servers for e-mail notification.
 - 3 Choose Add SMTP Server to add an SMTP e-mail server to the list that iHub uses for e-mail notification.
 - 4 In the list of SMTP e-mail servers, choose Delete to delete an SMTP e-mail server that iHub uses for notification.
 - 5 In the list of SMTP e-mail servers, choose Edit to edit the parameters for an SMTP e-mail server that iHub uses for notification.
- 4 Choose OK.

Specifying the Message Distribution service for e-mail notification

The administrator can use the simple view or the advanced view to configure iHub to send the e-mail notice to an Information Console user about a completed job. The e-mail message can contain a URL that includes a hyperlink to a design or document and the location of the Message Distribution service (MDS) for connecting to an iHub where the document resides. In the following example, the value of serverURL specifies the MDS:

```
http://sales:8900/iportal/newrequest  
/index.aspx?__requestType=scheduled & __executableName=  
/forecast.rptdesign%3B1&serverURL=http://  
end2243:8000&volume=volume1
```

By default, iHub uses any node in the cluster that is online and has the Message Distribution service enabled.

In the example, the serverURL parameter and value are:

```
serverURL=http://end2243:8000
```

Include the http:// prefix when you specify the serverURL.

Setting up sendmail e-mail notification

A sendmail e-mail notification originates from the iHub user account. To send an iHub e-mail notice from a Linux system, ensure that the account that runs the iHub can access the sendmail program from each node. Enable the View service on each node. The sendmail program is in `/usr/lib/sendmail`.

To test the e-mail configuration, send an e-mail message to a user from the iHub e-mail account.

How to test e-mail notification

- 1 Log in to the account that runs iHub.
- 2 Send the e-mail message using the following command:

```
/usr/lib/sendmail mail-address < message.txt
```

where
 - `message.txt` contains the test message.
 - `mail-address` is the user's account name as registered with iHub.
- 3 If the message arrives, the account setup is correct. If the message does not arrive, perform one or more of the following tasks:
 - Check the e-mail address for typographical errors.
 - Log in to an account other than the account that handles iHub e-mail and try sending e-mail to the user.
 - Compare the user's account name that is registered with your e-mail program to the account name registered in the Encyclopedia volume. These two account names must match exactly.
- 4 If the e-mail does not arrive after you complete the tasks in step 3, and you continue to have problems sending e-mail from the iHub account, contact Actuate Support for assistance.

Configuring the notification list size and To: line

In System—Properties—Advanced—Notification, you can specify how iHub notifies users and administrators of events by setting the following properties, as shown in Figure 4-6.

- Maximum number of recipients per e-mail message
Parameter name: `MaxMailRecipients`
Specifies the maximum number of e-mail addresses to which iHub can address a single e-mail message. If the number of e-mail recipients exceeds the

value of this parameter, iHub divides the list into smaller lists and sends the same e-mail message to each of the smaller lists.

The maximum number of e-mail recipients for an iHub e-mail message cannot exceed the maximum number of e-mail recipients limit for the e-mail server. If the e-mail server receives a request that contains more e-mail recipients than the server permits, it does not send the e-mail message. The e-mail server sends a failure notice to iHub. iHub does not attempt to resend the message.

The maximum number of e-mail recipients is a system attribute. The default value is 10,000 e-mail addresses. The minimum value is 100. The maximum value is 100,000.

- Use dummy line in place of empty To: line
Parameter name: UseDummyToLine

Indicates whether to use the value of Dummy To: line in an e-mail notice that iHub sends if both the To: and Cc: values are empty. Applicable to SMTP and sendmail configurations.

- Dummy To: line
Parameter name: DummyToLine

Text to use in the To: line of an e-mail notice from iHub if Use dummy line in place of empty To: line is selected and both the To: and Cc: values in the e-mail notice from iHub are empty.

Actuate > System properties - Windows Internet Explorer

System > iHub System : Properties > Notification

E-mail Transmission Limits

Maximum number of recipients per e-mail message: 10000

Empty To: Line In E-mail Message

Use dummy line in place of empty To: line: ☒

Dummy To: line: (names withheld)

These fields require system restart to take effect
(!) These fields will take default value if left blank

OK Cancel

Figure 4-6 Specifying advanced notification property values

Handling e-mail notification errors

If an error occurs when iHub sends an e-mail message, iHub writes a message to the diagnostic log. Under most conditions, iHub also writes a message to the error log. Except under certain conditions, iHub tries to resend the e-mail.

Handling SMTP e-mail errors

When configured to use SMTP, iHub categorizes errors that occur when sending e-mail messages as either retryable or non-retryable. iHub supports resending e-mail messages when a retryable e-mail error occurs.

If a cluster node receives an error when sending an e-mail message to an SMTP server, it sends the e-mail message using the next configured SMTP server. If the node that sends the e-mail message receives non-retryable errors when sending the e-mail message to all SMTP servers, it writes a message to the error log.

If the node that sends the e-mail message receives an error from all the SMTP servers, and at least one of the errors is a retryable error, the node waits for the retry interval of one second and tries to send the message to the SMTP server that returns a retryable error. If the SMTP server returns an error again, the node writes a message to the error log.

If iHub receives one of the following types of errors, it attempts to resend the e-mail message:

- Cannot connect to SMTP server
- Network connection fails
- Wait for response times out while sending message

For each error iHub receives, it writes a message to the diagnostic log.

If iHub receives one of the following types of errors, it does not attempt to resend the e-mail message:

- Unrecognized recipient
- Message too large
- Message contains too many recipients

Handling sendmail e-mail errors

When iHub uses sendmail on a Linux system, iHub makes one attempt to send an e-mail message. If iHub receives an error, it writes a message to the diagnostic log.

In a cluster, iHub distributes e-mail notice requests among the nodes that have the View service enabled. Each node uses its iHub e-mail account. If the cluster node receives an error, the node writes a message to the diagnostic log.

Logging e-mail notification errors

By default, iHub logs e-mail notification error messages.

- If an error occurs while an Encyclopedia volume is sending an e-mail notice request to a cluster node, the volume writes an error message to its diagnostic log.
- If a cluster node receiving a request to send an e-mail notice encounters an error, the node writes an error message to its diagnostic log. For example, when the combined size of the message text and the attachment exceeds the maximum message size, iHub does not send the e-mail. The diagnostic log in `AC_DATA_HOME/server/log` describes failed e-mail notification attempts.

iHub localizes all errors that it writes to the log to the default locale for the machine.

Customizing the e-mail message

The administrator can customize the e-mail message that iHub sends to notify users about successful and failed jobs. For each Encyclopedia volume, the default location for the e-mail message template is `AC_SERVER_HOME/etc`. In a cluster, unless otherwise configured, each View process uses the notification template in the local `AC_SERVER_HOME/etc` directory when processing e-mail notification.

Sending e-mail notification in a cluster

iHub distributes e-mail notification requests among nodes in a cluster that have the View service enabled. In System—Properties—Advanced—Cluster Operation—Administrative, specify administrative e-mail account information for a cluster node administrator, as shown in Figure 4-7.

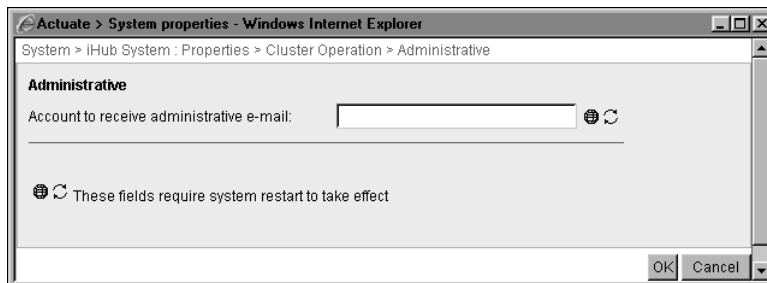


Figure 4-7 Specifying e-mail account information for cluster node administrator

You can configure iHub to send e-mail notices from multiple nodes using a single template.

Sending e-mail notices from nodes using a template

iHub uses an e-mail notification template to create the job completion notification e-mail. In a cluster, unless otherwise configured, each View process uses the notification template in the local AC_SERVER_HOME/etc directory.

If iHub sends e-mail notices from multiple nodes, the nodes use a single template in a partition that all nodes can access. The administrator creates the partition, copies acnotification.xml to the directory the partition specifies, then specifies the partition on Volumes—Properties—General in E-mail notification template partition, as shown in Figure 4-8.

The screenshot shows the 'Volumes > corp: Properties' dialog box with the 'General' tab selected. The 'Email notification' section is highlighted, showing the 'E-mail notification template partition' dropdown menu. A callout line points to this dropdown with the text 'Specify partition'. Other fields include 'Description', 'Schedule for purging notices', 'Schedule for purging deleted files', 'Primary partition', 'Min Free Space', 'Volume archive service provider', 'Metadata database and schema', and 'Information Console URL prefix'. The 'OK', 'Cancel', and 'Apply' buttons are at the bottom right.

Figure 4-8 Specifying a partition for e-mail notification template file

If iHub cannot find acnotification.xml, or the template file contains invalid formatting, iHub cannot send e-mail notices.

Using the e-mail message template

The e-mail message template, `acnotification.xml`, is an XML file that uses UTF-8 encoding. iHub constructs the e-mail about job completion notices based on the template. To customize the e-mail that iHub sends, modify the following default template:

```
<?xml version="1.0" encoding="UTF-8"?>
<notificationTemplate version="1.0">
  <successMessage>
    <subject>Actuate iHub Notification</subject>
    <body>
      <
-- Body Text Begin -->
      Actuate iHub - Report <insert variable="jobType"/>
      complete
      For Information Console:
      Report: <insert variable="reportLink"/>
      If the URL above does not work, copy the entire link and paste it
      into the address bar of your web browser, then press Enter or
      Return.
      Completed at: <insert variable="jobCompletion"></insert>
      Note: If the job submitter requested that you receive the report
      as an attachment to this email, but the report is not attached,
      then you probably do not have the privileges required to view
      the entire report. Please contact your system administrator.
      <
-- Body Text End -->
    </body>
  </successMessage>
  <failureMessage>
    <subject>Actuate iHub Notification</subject>
    <body>
      <
-- Body Text Begin -->
      Actuate iHub - Report <insert
        variable="jobType"/>
      failed.
      For Information Console:
      Report: <insert variable="reportLink"/>
      If the URL above does not work, copy the entire link and paste it
      into the address bar of your web browser, then press Enter or
      Return.
      Completed at: <insert variable="jobCompletion"></insert>
      Warning/Error:
      <insert variable="jobDetailedStatus"/>
      <
```

```
-- Body Text End -->
    </body>
  </failureMessage>
</notificationTemplate>
```

Working with e-mail template elements

The following list describes the e-mail template elements and text encoding information:

- **body**
Element that specifies the content of the e-mail body. The format of the body of the e-mail body can be in any format supported by the e-mail client. The body content can include insert elements.

Use CDATA sections to embed e-mail body content that is exempt from XML parsing rules. For example, if you specify the e-mail body content in HTML or any other markup language, the insert elements must be outside of CDATA sections.
- **failureMessage**
Parent element of the subject and body for an e-mail notice for a failed job.
- **insert element**
Optional element that inserts job and document information in subject and body message content. The value of the variable attribute specifies the information iHub inserts.

The element must appear outside any CDATA section used within the content of body or subject elements.
- **notificationTemplate**
Required root element of the e-mail notification template.
- **subject**
Element that specifies the content of the subject line. The subject content is plain text. The tag value or content can include insert elements to display values related to the job or document.

Use CDATA sections to embed e-mail body content that is exempt from XML parsing rules. The insert elements must be outside of CDATA sections.
- **successMessage**
Parent element of the subject and body elements for an e-mail notice for a successful job.

The following list describes the e-mail template attributes:

- **email-content-type**
Optional body element attribute that specifies the content type of the body content.

Use this attribute when constructing the e-mail message, because the e-mail protocol requires that you specify the content type.

The value for this attribute is either text/plain and text/html.
- **variable**
Required insert element attribute. Specifies the information to insert in the e-mail subject or body.
- **version**
Required notification template attribute. Specifies the version number of the notification template file.

Using variable attributes

Table 4-2 describes the valid values of variables for insertion in e-mail notices.

Table 4-2 Variable values for e-mail notices

Variable	Value
jobName	Job name.
jobSubmitter	Job submitter's user name.
jobStatus	Status of job: Completed or Failed.
jobDetailedStatus	Detailed status of the job from the job status page.
jobType	Type of job: Execution or Printing.
jobHeadline	The job's headline.
jobCompletion	Date and time of job completion.
reportDocumentName	Document name. Available for a successful job.
reportDocumentVersionName	Document version name. Available for a successful job.
reportDocumentVersionNumber	Document version number. Available for a successful job.
reportLink	Hyperlink to the document in the Encyclopedia volume for a successful job. For a failed job, the link accesses the job status page.

The following example uses the insert element's reportLink variable to display the URL to the document in the e-mail notice that iHub sends:

```
Report: <insert variable="reportLink" />
```

Using HTML in the e-mail template

To use HTML in the successMessage or failureMessage elements, set the message's body email-content-type attribute to text/html.

```
<body email-content-type="text/html">
```

Enclose the HTML in CDATA sections to exclude the HTML from XML parsing rules. The insert elements must be outside CDATA sections. The following example shows a successMessage element with HTML formatting:

```
<?xml version="1.0" encoding="UTF-8" ?>
  <notificationTemplate version="1.0">
    <successMessage>
      <subject>
        Report Delivery Notification:
        <insert variable="jobHeadline"/>
      </subject>
      <body email-content-type="text/html">
        <![CDATA[
          <html>
            <body>
              <h2>
                ]]>
                <insert variable="jobHeadline"/>
                <![CDATA[
                  </h2>
                ]]>
                Version <insert variable=
                  "reportDocumentVersionNumber"/>
                  of report <insert variable="reportDocumentName"/>
                  is now available online.
                <![CDATA[
                  <a href="
                ]]>
                  <insert variable="reportLink"/>
                <![CDATA[
                  ">Go to Report</a>
                  <p>
                    <table border="2pt;">
                      <tr><td>Report Submitter: </td><td>
                ]]>
                <insert variable="jobSubmitter"/>
```

```

<[!CDATA[
    </td></tr>
    <tr><td>Report Generation Date: </td>
    <td>
]]>
<insert variable="jobCompletion"/>
<[!CDATA[
    </td>
    </tr>
    </table>
    <br><br>
</body>
</html>
]]>
</body>
</successMessage>

```

Working with services

This chapter contains the following topics:

- About BIRT iHub and PostgreSQL services
- Using iHub services
- Configuring service properties

About BIRT iHub and PostgreSQL services

The BIRT iHub and the PostgreSQL services must run before other iHub services and processes can start. During iHub installation, accepting the default settings starts the BIRT iHub service automatically when the computer reboots. Accepting the default settings installs the PostgreSQL database. The iHub service appears in the Windows Control Panel, as shown in Figure 5-1.

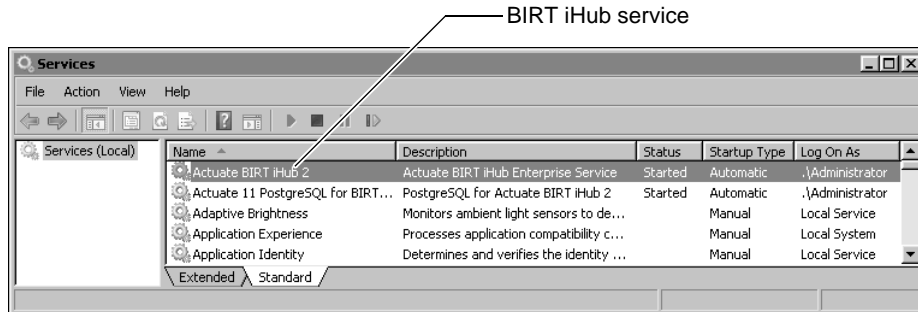


Figure 5-1 Viewing the iHub service

The PostgreSQL for BIRT iHub service also appears in the Windows Control Panel, as shown in Figure 5-2.

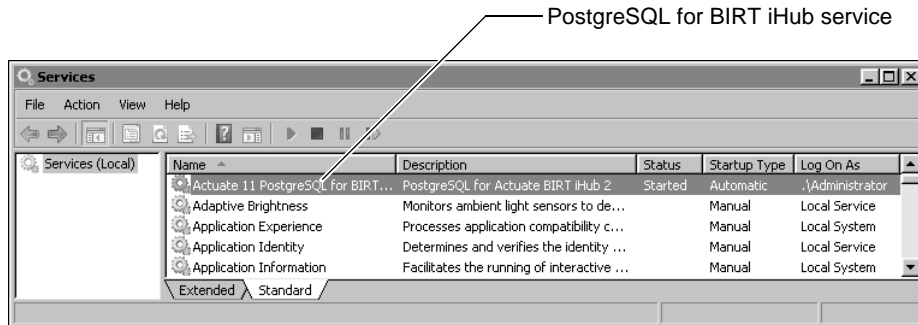


Figure 5-2 Viewing the PostgreSQL for BIRT iHub service

Using iHub services

Understanding iHub services and processes is important for taking advantage of the many configuration possibilities. The administrator can configure iHub as a stand-alone server or as a node in a cluster of iHubs. Figure 5-3 shows a stand-alone configuration that runs iHub services on a single machine. Services publish access methods to the iHub processes through the Simple Object Access Protocol (SOAP) interface.

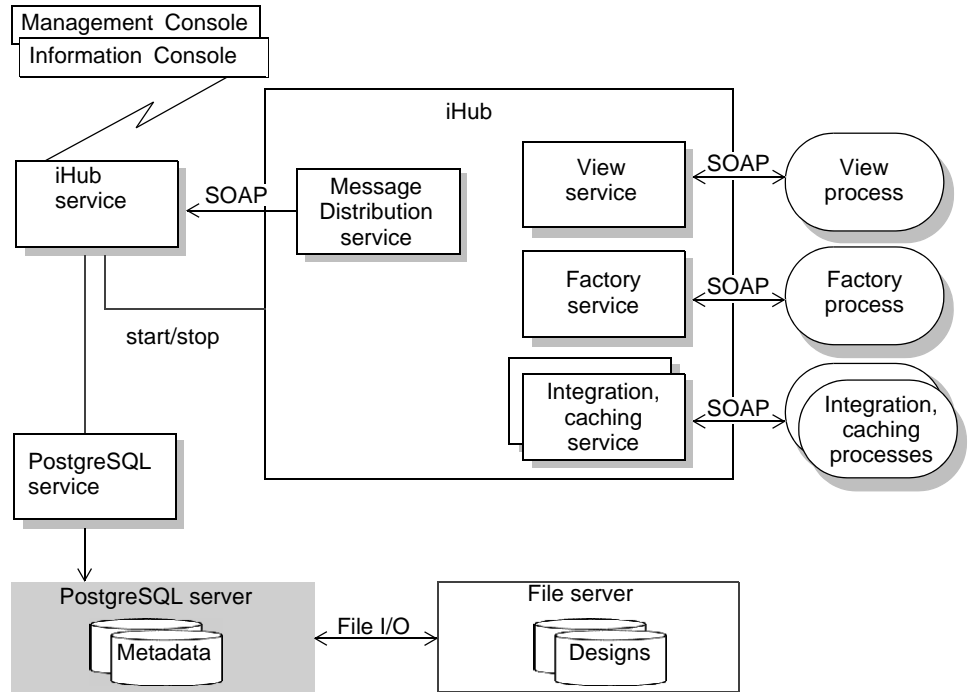


Figure 5-3 iHub architecture

Applications access processes using the Actuate Information Delivery API (IDAPI). IDAPI applications can communicate with cluster nodes and with iHub in a stand-alone configuration. A custom application that uses IDAPI to administer an iHub node sends messages to the cluster through the Message Distribution service.

The following list describes services running inside iHub:

- **Message Distribution**
Participates in load balancing by dispatching generation and print requests to other cluster nodes. In a stand-alone iHub or a cluster, the Message Distribution service dispatches SOAP requests that come to iHub from clients, users, and applications.
- **View**
Facilitates viewing documents in DHTML format, converting output to formats such as Excel and PDF, and handling requests to download files from an Encyclopedia volume. The Java view process, which handles BIRT designs, runs inside the View process.

- **Factory**
Executes requests to generate queries and documents and to print output on an iHub printer. The Java factory process, which handles BIRT designs, runs inside the Factory process.
- **Integration and Caching**
Coordinates the running of an information object (IOB) file that uses data from multiple data sources. Controls the Actuate Caching process that manages the information object cache and enables caching data from an information object IOB.

To control how iHub uses Factory processes, the administrator creates and manages resource groups using Configuration Console. Resource groups allocate Factory processes to handle prioritized requests for documents.

Configuring service properties

The administrator configures services in the advanced view of Configuration Console.

How to access service properties

Access service properties settings by performing the following tasks:

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Server Configuration Templates.
- 2 On Server Configuration Templates, choose a template name, as shown in Figure 5-4.

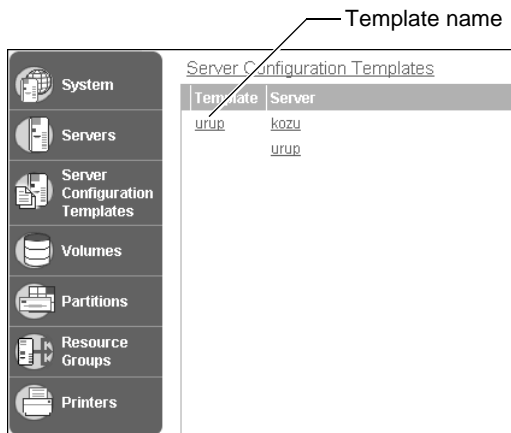


Figure 5-4 Choosing a template name

- 3 On Server Configuration Templates—Settings, service property folders appear, as shown in Figure 5-5.

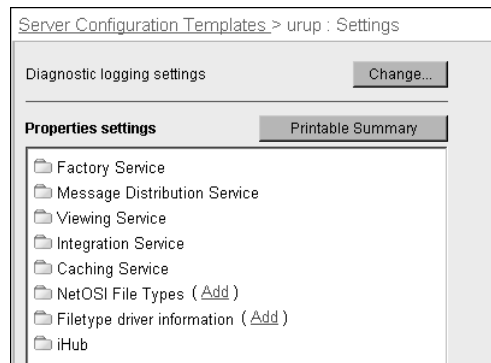


Figure 5-5 Accessing service property folders

Configuring the View service

This chapter contains the following topics:

- Configuring View service settings
- Configuring the View service for BIRT documents

Configuring View service settings

The View service provides the basic framework to render and display a document. Viewing starts when a user runs a design to create a document or selects an existing document in either Information Console or Management Console. The View service controls viewing a page or the table of contents, and searching a document.

A user navigates to Server Configuration Templates—Settings by choosing Server Configuration Templates from the side menu in the advanced view of Configuration Console. Then, on Server Configuration Templates, the user chooses a template name, as shown in Figure 6-1.

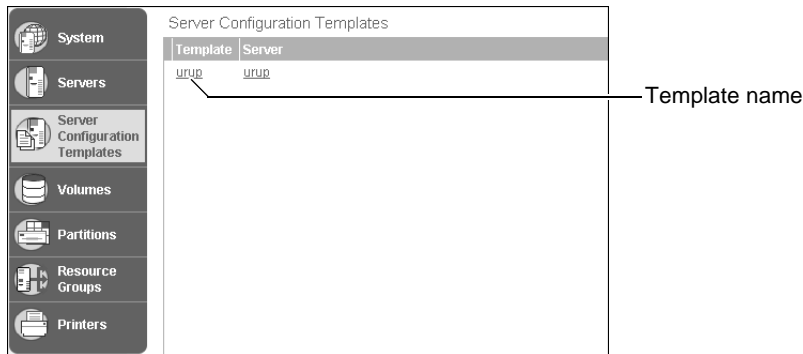


Figure 6-1 Choosing the template name

In Server Configuration Templates—Settings, View service properties that the administrator configures include properties for BIRT designs and documents, as shown in Figure 6-2.

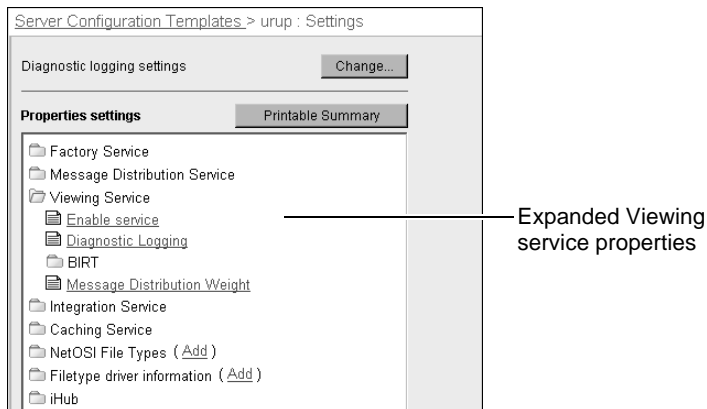


Figure 6-2 Viewing service properties in Settings

About setting View service properties in iHub

The properties that appeared in Servers—Properties—Viewing service in previous iServer releases appear in Server Configuration Templates—Settings in iHub. Set View service properties in Server Configuration Templates—Settings as follows:

- To set viewing weight, choose Viewing Service—Message Distribution Weight.

Enabling the View service

The administrator can enable or disable the View service from Server Configuration Templates—Settings, or by setting the EnableViewingService parameter in AC_DATA_HOME/config/iHub2/ acserverconfig.xml to true or false.

How to enable the View service

- 1 Expand Viewing Service and choose Enable service, as shown in Figure 6-2.
- 2 For Enable viewing service, accept the default value, which is selected, as shown in Figure 6-3.



Figure 6-3 Enabling or disabling the View service

About diagnostic logging

The administrator can configure diagnostic logging by expanding Viewing Service, and choosing Diagnostic Logging. For more information, see Chapter 3, “Using diagnostic, usage, and error logging.”

Configuring the message distribution weight for a node

To improve performance in an iHub cluster, the administrator can configure the viewing weight for each node to control how iHub distributes the load across the cluster. By default, the cluster master attempts to balance the load by routing requests to the View service of the node that is the least busy.

Setting the viewing weight of a node to a higher number than other nodes routes more requests to that node. Setting the viewing weight to a lower number tends to conserve node resources.

How to configure the message distribution weight for a node

- 1 In Server Configuration Templates—Settings, expand Viewing Service and choose Message Distribution Weight.
- 2 In Weight of this server for load balancing viewing requests, accept the default, 100, as shown in Figure 6-4. Alternatively, specify a different value.

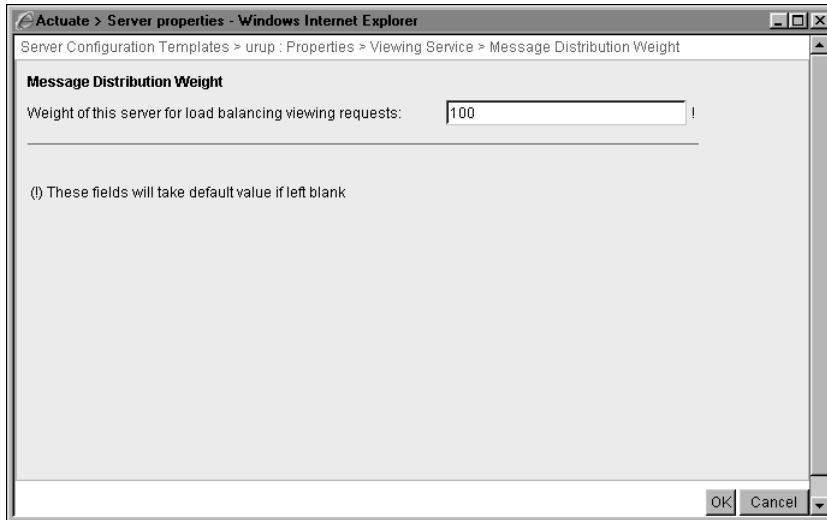


Figure 6-4 Specifying load balancing viewing requests weight

- 3 Choose OK.

The following snippet from the default `acmetadescription.xml` includes the `acmetadescription.xml` parameter name, the corresponding display name of the Configuration Console property, when a property change takes effect, and the default value:

```
Name="Weight of this server for load balancing viewing requests"
DisplayName="Viewing weight"
TakesEffect="Immediate"
DefaultValue="100"
```

Configuring the View service for BIRT documents

The iHub architecture incorporates the Java View service running within the View service. The Java View service handles user requests to view BIRT documents. Figure 6-5 shows the list of Java View service configuration categories. This section describes how to configure properties in these categories, except BIRT JDBC Connection Pool, which is discussed later in this book.

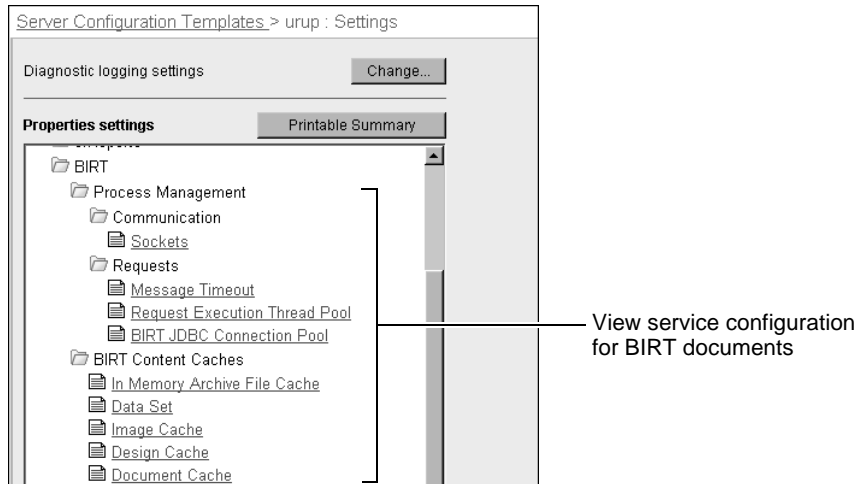


Figure 6-5 Viewing BIRT properties for the View service

Configuring Java process communication

The View service engages in Java process communication when a user views a BIRT document. In an environment that restricts port usage, the administrator can specify and change the base port for the View service and maximum range of other ports used for SOAP communication.

How to configure Java process communication

- 1 Expand Viewing Service, BIRT, Process Management, and Communication, then choose Sockets, as shown in Figure 6-5.
- 2 In Base port number for processes, accept the default base port, 21000, as shown in Figure 6-6. Alternatively, change the base port to a value in the range 1025 through 65535.

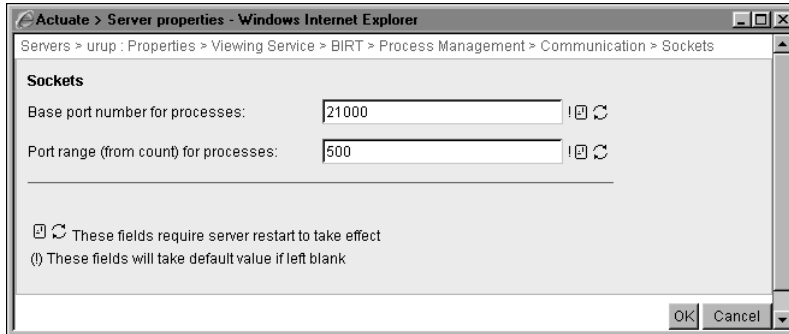


Figure 6-6 Specifying Java View service settings

- 3 In Port range (from count) for processes, accept the default value, 500, or change the value to a number in the range 0 through 64510.

The following snippets from the default `acmetadescription.xml` include the `acmetadescription.xml` parameter names, the corresponding display names of the Configuration Console properties, when property changes take effect, default values, and ranges:

```
Name="SocketBaseForJavaProcesses"
DisplayName="Base port number for processes"
TakesEffect="ServerRestart"
DefaultValue="21000"
Range="1025 65535"
```

```
Name="SocketCountForJavaProcesses"
DisplayName="Port range (from count) for processes"
TakesEffect="ServerRestart"
DefaultValue="500"
Range="0 64510"
```

Configuring message time-out

The administrator can set a message time-out period. If the View service does not respond in time to an on-demand request for a BIRT document, iHub rejects the request. Base the message time-out period on the expected response time of iHub for on-demand generation and loading of the temporary document.

Timeout for generation of gadgets in BIRT 360 specifies the number of seconds that iHub has to generate a gadget. If this time limit is reached, iHub cancels the gadget-generation task.

How to configure message time-out

- 1 In Timeout for generation of gadgets in BIRT 360, accept the default value, 300, or change the value to a different number. iHub uses the default, 300, when you set the message time-out to blank, as shown in Figure 6-7. Choose OK.

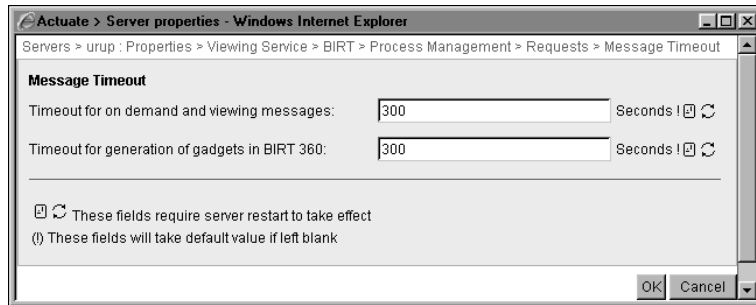


Figure 6-7 Setting time-out periods

2 Restart iHub.

The following snippets from the default `acmetadescription.xml` include the `acmetadescription.xml` parameter names, the corresponding display names of the Configuration Console properties, when property changes take effect, default values, and ranges:

```
Name="OnDemandServerViewMessageTimeout"
DisplayName="Timeout for on demand and viewing messages"
TakesEffect="ServerRestart"
DefaultValue="300"
UnitOrFormat="Seconds"
```

```
Name="GadgetGenerationTimeOut"
DisplayName="Timeout for generation of gadgets in BIRT 360"
UnitOrFormat="Seconds"
TakesEffect="ServerRestart"
DefaultValue="300"
```

Configuring the thread pool

The administrator can limit the number of threads and size of the queue each on-demand server uses for performing the following tasks:

- Generating temporary BIRT documents
- Processing requests to view temporary and persistent BIRT documents

To limit the number of threads, the administrator sets maximum number of threads in each on-demand server. The thread batch size is 5. iHub opens five new threads as needed until the total number reaches the limit. For example, setting the limit to 16 causes iHub to open the new threads in batches, 5, 5, 5, and 1. Increasing the value of this property increases CPU and memory consumption.

To set the size of the queue, the administrator sets task queue size in each on-demand server.

How to configure the thread pool

- 1 Expand Viewing Service, BIRT, Process Management, and Requests, as shown in Figure 6-5, then choose Request Execution Thread Pool.

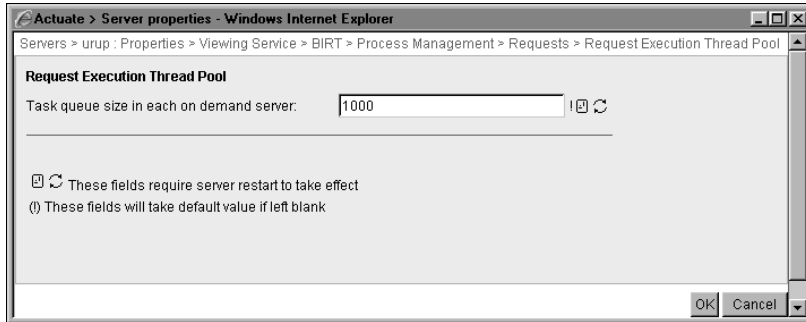


Figure 6-8 Configuring the thread pool

- 2 In Task queue size in each on demand server, accept the default, 1000, or type another value.

Choose OK.

- 3 Restart iHub.

The following snippet from the default acmetadescription.xml includes the acmetadescription.xml parameter name, the corresponding display name of the Configuration Console property, when property changes take effect, and default value:

```
Name="OnDemandServerQueueSize"
DisplayName="Task queue size in each on demand server"
TakesEffect="ServerRestart"
DefaultValue="1000"
```

Configuring BIRT document and design caching

iHub caches documents and designs to respond quickly to viewing requests. The View service maintains the following types of caches for BIRT documents and designs:

- In-memory archive
- Result set buffer
- Image cache
- Design cache
- Document cache
- Page count cache

The Encyclopedia volume stores persistent documents and designs until removed by the user. Temporary documents are not stored in the Encyclopedia volume. The administrator can configure caching to improve viewing performance for a particular environment.

Configuring the BIRT document in-memory archive

The administrator can configure caching to improve performance of viewing small, temporary, and persistent BIRT documents. iHub caches BIRT documents of a size equal to or smaller than the value appearing in Maximum memory size of each BIRT report archive. When a user requests a document of a size larger than this value, iHub temporarily stores the document on the hard drive.

The total memory size used to cache BIRT document and datamart files value determines how much memory iHub can use to cache documents. Increasing this value caches more BIRT documents in memory, consuming more memory. Decreasing this value temporarily stores more BIRT documents on the hard drive.

To optimize performance, the administrator tries to prevent storing most documents on the hard drive by making configuration decisions based on the size of the average document and the size of the cache.

The administrator can also set the time-out period for temporary and persistent document caches. When the time elapses, iHub clears the cache.

How to configure the BIRT document in-memory archive

- 1 Expand Viewing Service, BIRT, and BIRT Content Caches, as shown in Figure 6-5, and choose In Memory Archive File Cache.
- 2 In Total memory size used to cache BIRT document and datamart files, accept the default value, 153600 KB, or type another value, as shown in Figure 6-9. To disable memory-based caching, type a negative value or 0.

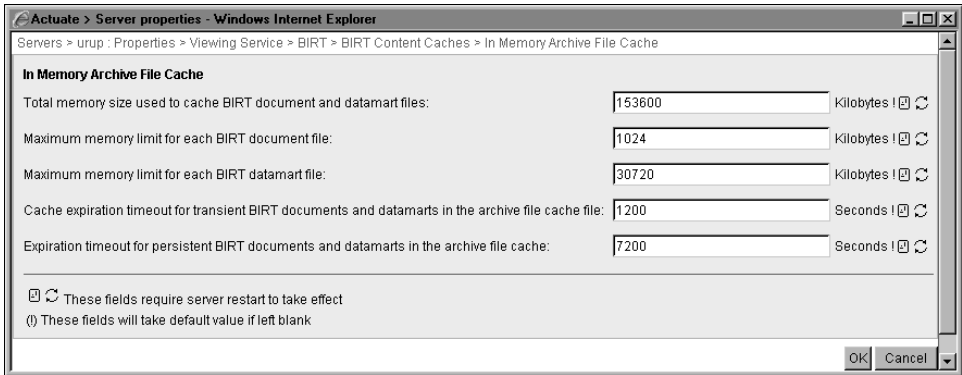


Figure 6-9 Configuring the BIRT document in-memory archive

- 3 In Maximum memory limit for each BIRT document file, accept the default value, 1024 KB. Alternatively, type a value that represents the size of average

document you want iHub to cache. To disable memory-based caching, type a negative value or 0.

- 4 In Maximum memory limit for each BIRT datamart file, accept the default value, 30720 KB. Alternatively, type a value that represents the size of average datamart you want iHub to cache. To disable memory-based caching, type a negative value or 0.
- 5 In Cache expiration timeout for transient BIRT documents and datamarts in the archive file cache file, accept the default, 1200 seconds, or type another value.
- 6 In Expiration timeout for persistent BIRT documents and datamarts in the archive file cache, accept the default, 7200, or type another value.

Choose OK.

7 Restart iHub.

The following snippets from the default `acmetadescription.xml` include the `acmetadescription.xml` parameter names, the corresponding display names of the Configuration Console properties, when property changes take effect, and default values:

```
Name="TotalArchiveMemory"
DisplayName="Total memory size used to cache BIRT document and
  datamart files"
TakesEffect="ServerRestart"
UnitOrFormat="Kilobytes"
DefaultValue="153600"
```

```
Name="MaxMemoryPerArchive"
DisplayName="Maximum memory limit for each BIRT document file"
UnitOrFormat="Kilobytes"
TakesEffect="ServerRestart"
DefaultValue="1024"
```

```
Name="MaxMemoryPerDatamartArchive"
DisplayName="Maximum memory limit for each BIRT datamart file"
UnitOrFormat="Kilobytes"
TakesEffect="ServerRestart"
DefaultValue="30720"
```

```
Name="TransientArchiveFileCacheTimeout"
DisplayName="Cache expiration timeout for transient BIRT documents
  and datamarts in the archive file cache file"
TakesEffect="ServerRestart"
UnitOrFormat="Seconds"
DefaultValue="1200"
```

```

Name="PersistentArchiveFileCacheTimeout"
DisplayName="Expiration timeout for persistent BIRT documents and
    datamarts in the archive file cache"
TakesEffect="ServerRestart"
UnitOrFormat="Seconds"
DefaultValue="7200"

```

Configuring the BIRT result set buffer

The Java View service handles on-demand document generation requests and multiple concurrent requests, one at a time. Configuring the size of the buffer that stores the result sets for a data set in a BIRT document can improve the response to on-demand requests to generate BIRT documents. An administrator typically changes the size of this buffer from the default to another value under the following conditions:

- The data sets of most BIRT designs are larger than the default buffer size.
- The JVM start arguments include specification of a sufficient heap size to handle the new value.

Increasing the size of the result set buffer for a BIRT data object query increases the memory used to store the final result set, and decreases the disk space used. If the result set is larger than this value, iHub writes the data to disk.

How to configure the BIRT result set buffer

- 1 Expand Viewing Service, BIRT, and BIRT Content Caches, as shown in Figure 6-5, and choose Data Set.
- 2 In Maximum buffer size for BIRT Data Object query result set in BIRT 360, accept the default, 8 MB, as shown in Figure 6-10. Alternatively, type a different value.

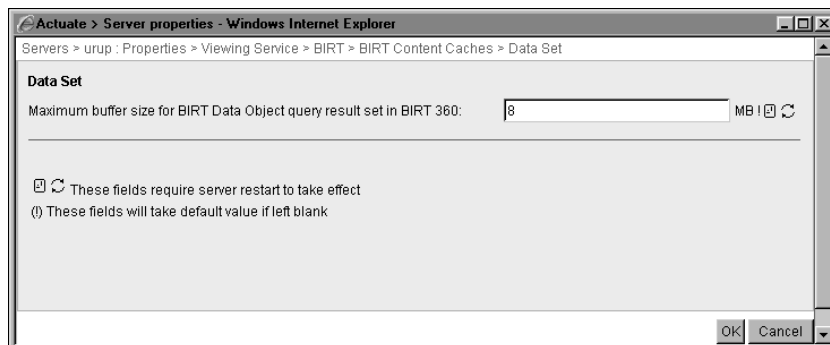


Figure 6-10 Configuring the buffer size for query result set
Choose OK.

- 3 Restart iHub.

The following snippet from the default `acmetadescription.xml` includes the `acmetadescription.xml` parameter name, the corresponding display name of the Configuration Console property, when a property change takes effect, the default value, and range:

```
Name="MaxBIRTDataSetBufferSize"
DisplayName="Maximum buffer size for BIRT Data Object query result
  set in BIRT 360"
UnitOrFormat="MB"
TakesEffect="ServerRestart"
DefaultValue="8"
Range="1 256"
```

Configuring the BIRT image cache

The administrator can configure the cache of images in persistent BIRT documents by setting cache time-out for images and charts from BIRT designs, documents and datamarts. Increasing this value increases the size of the memory buffers for caching images. Graphic-intensive documents load faster, but iHub uses more memory.

How to configure the BIRT image cache

- 1 Expand Viewing Service, BIRT, and BIRT Content Caches, as shown in Figure 6-5, and choose Image Cache.
- 2 In Cache timeout for images and charts from BIRT designs, documents and datamarts, accept the default, 86400 seconds, which is one day, as shown in Figure 6-11. Alternatively, type another value greater than 0. A value of 0 or less causes iHub to use a hard-coded value of 5.

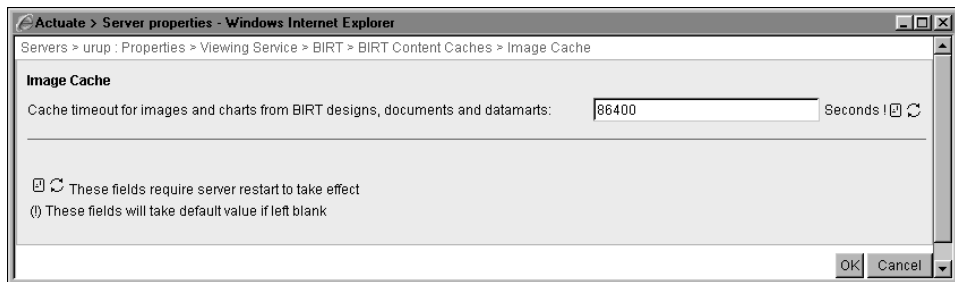


Figure 6-11 Configuring the BIRT report image cache

Choose OK.

- 3 Restart iHub.

The following snippet from the default `acmetadescription.xml` includes the `acmetadescription.xml` parameter name, the corresponding display name of the

Configuration Console property, when a property change takes effect, and the default value:

```
Name="BIRTImageCacheTimeout"  
DisplayName="Cache timeout for images and charts from BIRT  
designs, documents and datamarts"  
TakesEffect="ServerRestart"  
UnitOrFormat="Seconds"  
DefaultValue="86400"
```

Configuring the BIRT design cache

By default, iHub caches a BIRT design, including access privileges. Caching benefits users who access the design concurrently. Users who request access to the same design share the cached design if they have the required privileges. Performance can improve because iHub does not have to repeatedly load the design. Generally, the fewer number of designs iHub needs to load, the better the response time.

By configuring cache time-out for BIRT designs, the administrator can control how long the design remains in cache. A cached design persists in memory until time-out occurs. Use this time-out setting to control memory usage.

By configuring the maximum number of BIRT designs, the administrator controls the number of designs in the cache. When the cache reaches this limit, design caching stops. Use this setting to improve load performance.

How to configure the BIRT design cache

- 1 Expand Viewing Service, BIRT, and BIRT Content Caches, as shown in Figure 6-5, and choose Design Cache.
- 2 In Cache timeout for BIRT designs, accept the default, 1800 seconds or 30 minutes, as shown in Figure 6-12. Alternatively, type another value.

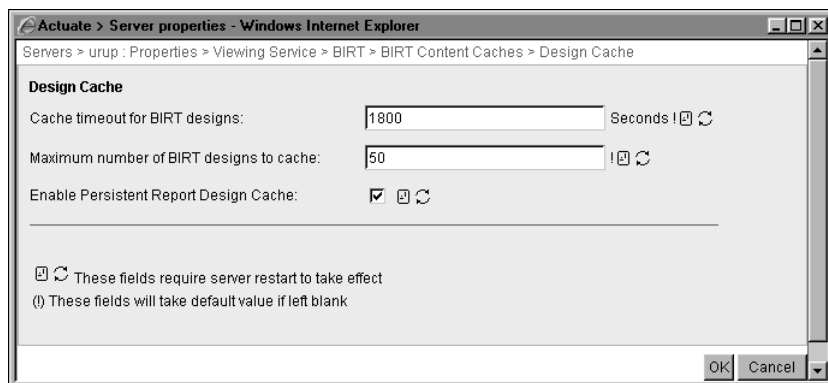


Figure 6-12 Configuring the BIRT design cache

- 3 In Maximum number of BIRT designs to cache, accept the default, 50, or type another value that limits the number of designs in the cache.
- 4 For Enable Persistent Report Design Cache, accept the default value of selected. Alternatively, disable the cache by deselecting this option.

Choose OK.

5 Restart iHub.

The following snippets from the default `acmetadescription.xml` include the `acmetadescription.xml` parameter names, the corresponding display names of the Configuration Console properties, when property changes take effect, and default values:

```
Name="BIRTReportDesignCacheTimeout"  
DisplayName="Cache timeout for BIRT designs"  
TakesEffect="ServerRestart"  
UnitOrFormat="Seconds"  
DefaultValue="1800"
```

```
Name="BIRTReportDesignCacheTotalNumberOfEntries"  
DisplayName="Maximum number of BIRT designs to cache"  
TakesEffect="ServerRestart"  
DefaultValue="50"
```

```
Name="EnablePersistentDesignCache"  
DisplayName="Enable Persistent Report Design Cache"  
TakesEffect="ServerRestart"  
DefaultValue="true"
```

Configuring the BIRT document cache

By default, iHub caches a BIRT document, including access privileges. Caching benefits users who access the document concurrently. Users who request access to the same document share the cached document if they have the required privileges. Performance can improve because iHub does not have to repeatedly load the document. Generally, the fewer number of documents iHub needs to load, the better the response time. iHub caches BIRT documents in the BIRT document in-memory archive cache. To access the cache, iHub creates a handle. If you enable the BIRT document cache, iHub caches this handle in the BIRT document cache. If you do not enable the BIRT document cache, iHub creates a new handle every time a user chooses to view a document. Enabling the BIRT document cache results in a faster response time but uses more memory, because iHub maintains the BIRT document cache in memory.

How to configure the BIRT document cache

- 1 Expand Viewing Service, BIRT, and BIRT Content Caches, as shown in Figure 6-5, and choose Document Cache.

- 2 For Enable caching of BIRT document and datamart handles, accept the default value, selected, as shown in Figure 6-13. Alternatively, deselect the option.

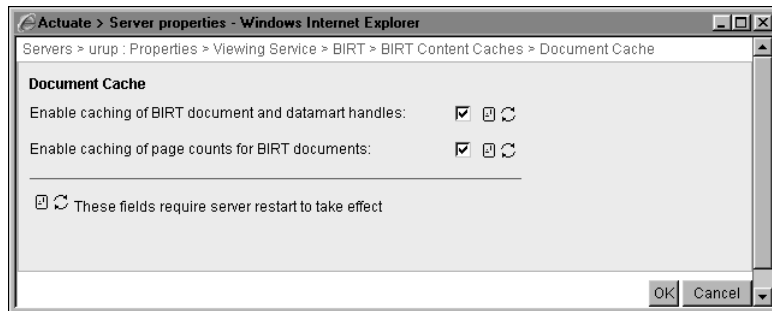


Figure 6-13 Enabling or disabling the BIRT document cache

- 3 For Enable caching of page counts for BIRT documents, accept the default value, selected, as shown in Figure 6-13. Alternatively, deselect the option.

Choose OK.

- 4 Restart iHub.

The following snippets from the default `acmetadescription.xml` include the `acmetadescription.xml` parameter names, the corresponding display names of the Configuration Console properties, when a property change takes effect, and the default values:

```
Name="BIRTReportDocumentCacheEnabled"
DisplayName="Enable caching of BIRT document and datamart handles"
TakesEffect="ServerRestart"
DefaultValue="True"
```

```
Name="BIRTReportPageCountCacheEnabled"
DisplayName="Enable caching of page counts for BIRT documents"
TakesEffect="ServerRestart"
DefaultValue="True"
```


7

Configuring the Factory service

This chapter contains the following topics:

- About the Factory service
- Configuring the Factory service for general use
- Configuring the Factory service for BIRT documents

About the Factory service

The Factory service provides the basic engine for running a design, and printing a document on the server-side. An asynchronous Factory generates scheduled documents or queries. A synchronous Factory generates temporary documents. To generate temporary documents, iHub must enable both the Factory and View services.

In Server Configuration Templates—Settings, Factory service properties that the administrator configures include properties for BIRT designs and documents, as shown in Figure 7-1.

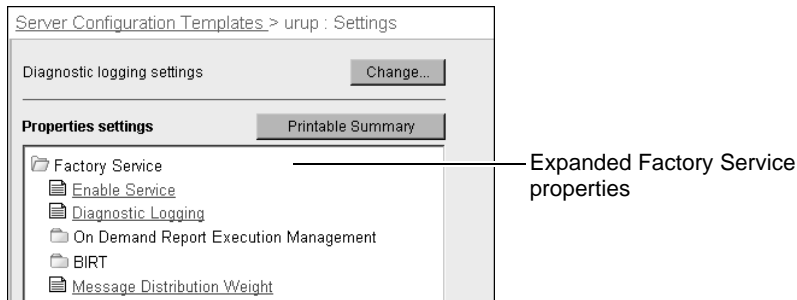


Figure 7-1 Configuring the Factory service in Settings

About setting Factory service properties in iHub

The properties that appeared in Servers—Properties—Factory Service in previous iServer releases appear in Server Configuration Templates—Settings starting with iServer Release 11 and continuing in iHub. Set Factory service properties in Server Configuration Templates—Settings as follows:

- To set transient report storage properties, choose Factory Service—On Demand Report Execution Management—Transient Report Management.
- To set synchronous report execution properties, choose Factory Service—On Demand Report Execution Management—On Demand Execution Queue.
- To set cluster load-balancing properties, choose Factory Service—Message Distribution Weight.
- To set the maximum number of factories for a resource group, choose Resource Groups, then choose <resource group>—Properties—Template Assignments.

Enabling the Factory service

The administrator can enable or disable the Factory service from Server Configuration Templates—Settings, or by setting the `EnableGenerationService` parameter in `AC_DATA_HOME/config/iHub2/acserverconfig.xml` to true or false.

How to enable the Factory service

- 1 Expand Factory Service and choose Enable Service, as shown in Figure 7-1.
- 2 For Enable factory service, accept the default value, which is selected, as shown in Figure 7-2.




Figure 7-2 Enabling or disabling the Factory service

About diagnostic logging

The administrator can configure diagnostic logging by expanding Factory Service, and choosing Diagnostic Logging. For more information, see Chapter 3, “Using diagnostic, usage, and error logging.”

Configuring the Factory service for general use

The following sections describe Factory service configuration tasks for general use:

- Configuring the transient document cache
- Configuring the number of entries in the transient document cache
- Configuring the Factory service for synchronous jobs

Figure 7-3 shows the selections the administrator makes in Server Configuration Templates—Settings to configure properties for general use.

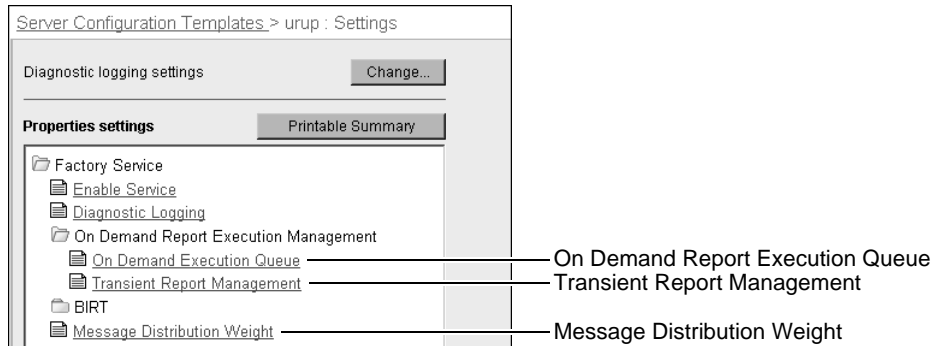


Figure 7-3 Configuring synchronous and transient report properties

Configuring the transient document cache

When a user requests the Factory service to generate a temporary document, iHub stores the document in an external, disk-based cache. By default, after 30 minutes, iHub deletes the file. The administrator can set the size of the cache for temporary documents, location of the cache, and cache time-out.

Setting the transient document cache location

By default, the Factory uses the following directory location for the cache:

`AC_DATA_HOME/server/tmp`

The administrator can specify a different path, locally or on the network. iHub must have access to the location. Actuate recommends keeping `/tmp` on a local disk drive. Accessing the `/tmp` directory over the network can impede iHub performance.

Setting the transient document cache size

The transient document cache is disk-based. By setting the cache size and time-out, the administrator controls how much disk space the Factory uses to cache temporary documents. When the cache reaches the limit, iHub sends a message to users indicating a storage problem. Increase the size of the cache if users receive this message. Decrease the size of the cache to conserve disk space. Setting the cache size to 0 permits only one document at a time in the cache.

Setting the transient document time-out

The administrator can control how long documents remain in the cache and on the disk by setting the transient document time-out. The document remains in the cache until the time expires, then iHub clears the cache.

When configuring the time-out, consider the time required for document generation and viewing by users. For example, to generate and view one-page

documents, try a time-out value of 10 minutes. To estimate the document generation time, add the values of `MaxSyncJobRuntime` and `SyncJobQueueWait`.

Configuring the number of entries in the transient document cache

The administrator can limit how many temporary documents the Factory caches when users generate documents. Maximum memory cache entries for transient reports is the upper limit on how many temporary documents the cache stores. When the cache reaches the limit, iHub sends a message to the user indicating a storage problem.

If users generate too many temporary documents too quickly, a message about failure to register the document appears. To correct the problem, increase the value of maximum memory cache entries for transient reports. Specify a value larger than the typical maximum number of documents generated within the transient document time-out period. Increasing the time-out and maximum memory cache entries for transient reports can increase iHub memory usage. When changing value of maximum memory cache entries for transient reports, consider adjusting the disk cache size for transient reports property.

How to configure the transient document cache

- 1 In Server Configuration Templates—Settings, expand Factory Service. Expand On Demand Report Execution Management, and choose Transient Report Management, as shown in Figure 7-3.
- 2 In Location of disk cache for transient reports, accept the default, or type another path, as shown in Figure 7-4.

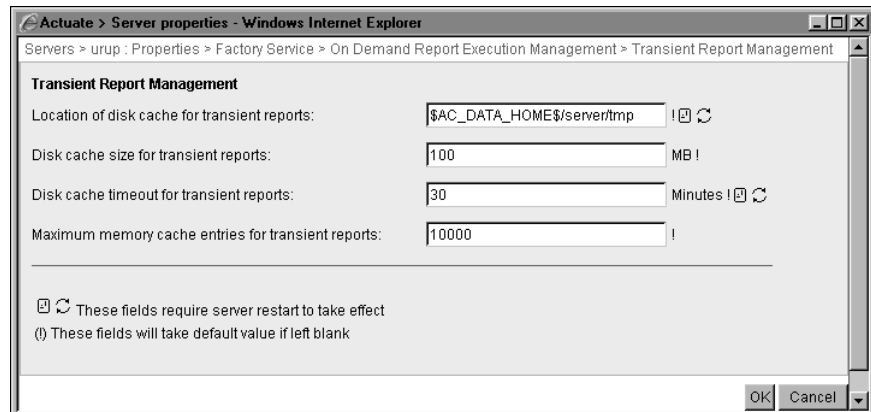


Figure 7-4 Changing the transient document cache size

- 3 In Disk cache size for transient reports, accept the default, 100 MB. Alternatively, type another value in megabytes.

- 4 In Disk cache timeout for transient reports, accept the default, 30 minutes. Alternatively, increase the value to give users more time to complete their work. Decrease the value to free disk space sooner.
- 5 In Maximum memory cache entries for transient reports, accept the default, 10000, as shown in Figure 7-4. Alternatively, type a value greater than 0 to change the number of entries permitted in the transient document cache. Setting the value to 0 prevents the generation of transient documents.

Choose OK.

- 6 Restart iHub if you change the cache location or time-out values.

Table 7-1 lists the property names that appear in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, which indicate default settings, ranges, and when a property change takes effect.

Table 7-1 Transient document storage parameters

Property name	Parameter name	Default	Range	Takes effect
Disk cache size for transient reports	TransientReport CacheSize	100 MB	0 - 99999	Immediate
Location of disk cache for transient reports	TransientReport CacheLocation			Server Restart
Disk cache timeout for transient reports	TransientReport TimeOut	30 Minutes	1 - 1440	Server Restart
Maximum memory cache entries for transient reports	TransientStoreMax CacheEntries	10000		Immediate

Configuring the Factory service for synchronous jobs

A user best practice is to use a schedule to generate large documents, but not necessarily short ones. The administrator can configure synchronous job settings to prevent problems that can occur when users generate large documents unscheduled. The Factory generates unscheduled documents synchronously. A synchronous job is the Factory process that generates a temporary document.

Configuring the maximum run time

Max synchronous job runtime limits the time a Factory can take to generate a document. Such a limit can prevent the generation of a huge unscheduled document from dominating Factory resources and degrading the response time of iHub. If an organization has mostly short documents, decreasing the value of the maximum execution time for on-demand execution requests can improve performance. If an organization has a high ratio of system resources to users, increasing the value of the maximum execution time for on-demand execution requests permits users to generate large, unscheduled documents.

Setting the value of Maximum execution time for on demand execution requests too high can delay generation of small, unscheduled documents, which can cause requests to back up in the queue and time out, and users to experience delays. Setting the value to 0 prevents the generation of any unscheduled documents.

Configuring the queue size

Each synchronous resource group has a job queue. When Factories are busy, a request for a Factory service to generate a document waits in the queue. The administrator sets the number of requests that the queue can hold using Job queue size for synchronous reports.

If a user receives a message that the synchronous job queue is full, consider increasing the queue size or resources, such as CPU power. Setting the size of the queue too large can cause the accumulation of too many requests in the queue. A request can time out in the queue.

Configuring the queue time-out

Job queue time-out for transient reports is the period of time, in seconds, that a request for an unscheduled document remains in the queue. Requests wait in the queue for 600 seconds, or 10 minutes, by default, then the request times out and fails. The user who requested the document receives a message that the job expired in the queue. The administrator can change the queue time-out by setting Job queue time-out for transient reports.

How to configure the Factory service for synchronous jobs

- 1 In Server Configuration Templates—Settings, expand Factory Service. Expand On Demand Report Execution Management, and choose On Demand Execution Queue, as shown in Figure 7-3.
- 2 In Maximum execution time for on demand execution requests, accept the default, 300 seconds, or 5 minutes, as shown in Figure 7-5. Alternatively, type the maximum number of seconds a Factory can spend per document.

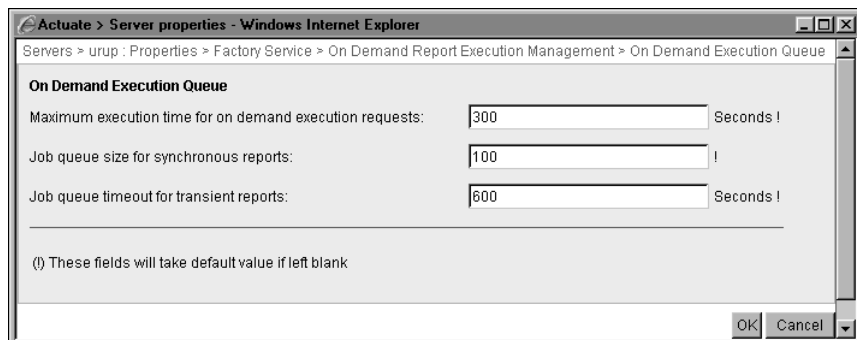


Figure 7-5 Configuring Max synchronous job runtime

- 3 In Job queue size for synchronous reports, accept the default, 100.
Alternatively, type a value greater than 0 to increase or decrease the queue size. A value of 0 causes requests for documents to fail when a Factory is not available to generate the document.
- 4 In Job queue timeout for transient reports, accept the default, 600.
Alternatively, type a value greater than 0 to increase or decrease the wait time. A value of 0 rejects requests to run a document unless a Factory is available immediately.

Table 7-2 lists the property names that appear in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating default settings, ranges, and when a property change takes effect.

Table 7-2 Synchronous document execution parameters

Property name	Parameter name	Default	Range	Takes effect
Maximum execution time for on demand execution requests	MaxSyncJobRuntime	300 Seconds	0 - 86400	Immediate
Job queue size for synchronous reports	SyncJobQueueSize	100	0 - 99999	Immediate
Job queue timeout for transient reports	SyncJobQueueWait	600 Seconds	0 - 999	Immediate

Configuring the message distribution weight of a node

By setting the Message Distribution Weight property, the administrator can assign priority to nodes for synchronous Factory service. By default, the number of Factories, compared to other nodes in the cluster, determines the weight of a node. The number of Factories in a node usually reflects the number of processors. The administrator can change the default weight of a node to take other factors, such as CPU speed, into consideration. The administrator can also change the weights of nodes if the synchronous Factories of a node are too idle or too busy.

The message distribution weight of a node is relative to the message distribution weights of other nodes. A node having a lower weight than other nodes receives fewer requests for unscheduled documents than the others. The message distribution weight of a node applies to all synchronous resource groups defined for the node.

How to set the Message Distribution Weight property

- 1 In Server Configuration Templates—Settings, expand Factory Service, then choose Message Distribution Weight, as shown in Figure 7-3.
- 2 In Weight of this server for load balancing on demand execution requests, as shown in Figure 7-6, accept the default, 100. Alternatively, type a value

relative to the weights of other nodes in the cluster. Type a multiple of 100 to increase the weight of a node. Type a value below 100 to decrease the weight of a node.

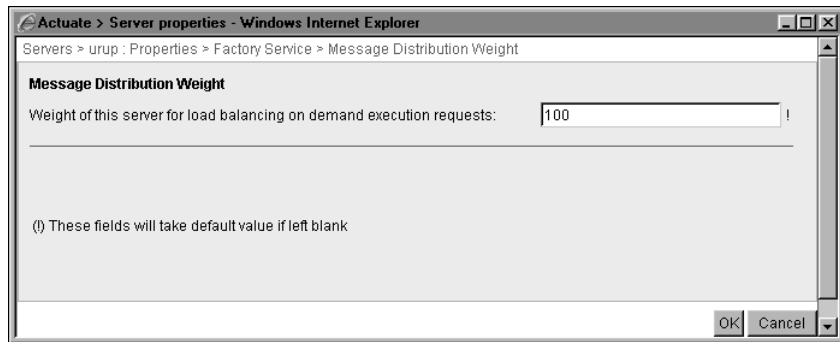


Figure 7-6 Setting the Message Distribution Weight property

Choose OK.

Table 7-3 lists the property name that appears in Configuration Console with the corresponding parameter name in `acmetadescription.xml`, indicating the default setting and when a property change takes effect.

Table 7-3 Message distribution weight parameter

Property name	Parameter name	Default	Takes effect
Weight of this server for load balancing on demand execution requests	SynchReportingWeight	100	Immediate

Configuring the Factory service for BIRT documents

When generating BIRT documents, the administrator can configure the Factory service for managing resources and optimizing performance. The following sections discuss this topic:

- Configuring the Factory service base port and range of ports
- Recycling Java Factories
- Configuring BIRT caching
- Configuring the maximum rows in a BIRT chart

Figure 7-7 shows the selections the administrator makes in Server Configuration Templates—Settings to configure properties for BIRT documents.

Configuring the Factory service base port and range of ports

The Factory service engages in Java process communication when a user generates a BIRT document. In an environment that restricts port usage, the administrator can specify and change the base port for the Factory service and the maximum range of other ports used for SOAP communication.

How to configure the Factory service base port and range of ports

- 1 In Server Configuration Templates—Settings, expand Factory Service, then expand BIRT. Expand Process Management. Expand Communication, and choose Sockets, as shown in Figure 7-7.

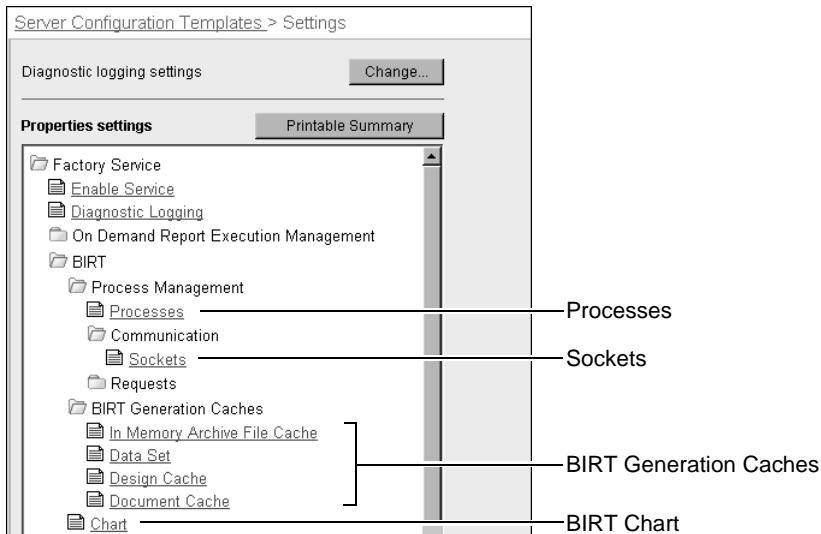


Figure 7-7 Configuring properties for BIRT documents

- 2 In Base port number for BIRT factory processes, accept the default, 21500, as shown in Figure 7-8. Alternatively, type another port number.

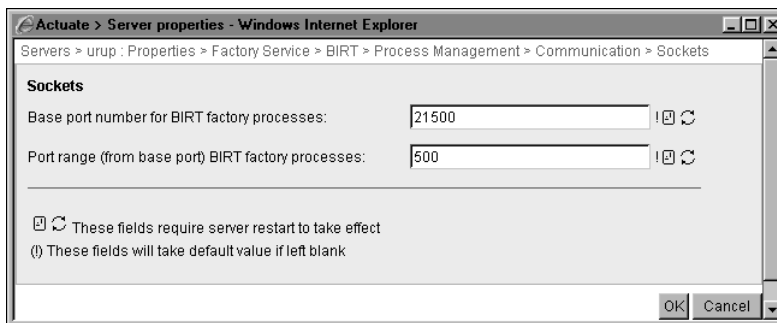


Figure 7-8 Specifying a base port for the Factory

- 3 In Port range (from base port) BIRT factory processes, accept the default 500, or type another maximum port number.

Choose OK.

- 4 Restart iHub.

Table 7-4 lists the property names that appear in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating default settings, ranges, and when a property change takes effect.

Table 7-4 Java process communication parameters

Property name	Parameter name	Default	Range	Takes effect
Base port number for BIRT factory processes	SocketBase	21500	1025 - 65535	Server
	ForJavaProcesses			Restart
Port range (from base port) BIRT factory processes	SocketCount	500	0 - 64510	Server
	ForJavaProcesses			Restart

Recycling Java Factories

By setting Number of Requests Before Recycling Processes to a value greater than 0, the administrator can limit the number of requests for BIRT documents that a Factory can handle. After the Factory reaches the limit for handling requests, the Factory shuts down, freeing resources. A new Factory emerges. Setting the value of Number of Requests Before Recycling Processes low restarts Factories more frequently than setting the value high. Actuate generally recommends the default value, 0, which disables recycling. In the event of a resource shortage that can occur over time, for example, due to connectivity problems, enable recycling.

How to enable or disable Java Factory recycling

- 1 In Server Configuration Templates—Settings, expand Factory Service, then expand BIRT. Expand Process Management, then choose Processes, as shown in Figure 7-7.
- 2 In Number of requests before recycling Java Factory processes, accept the default, 0, as shown in Figure 7-9. This action disables recycling. Alternatively, to enable recycling, type a non-zero value as high as the maximum number of requests a Java Factory handles concurrently.




Figure 7-9 Enabling Java Factory recycling

Choose OK.

3 Restart iHub.

Table 7-5 lists the property name that appears in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating the default setting and when a property change takes effect.

Table 7-5 Java process management parameter

Property name	Parameter name	Default	Takes effect
Number of requests before recycling Java Factory processes	JavaProcessRecycleCount	0	Server Restart

Configuring BIRT caching

The administrator can configure the Factory service to change the BIRT data result set buffer, and the design and document caches by choosing **Server Configuration Templates—Settings—Factory Service—Java Factory Service—Caching**, as shown in Figure 7-7. These topics are discussed in the following sections:

- Configuring the data result set buffer
- Configuring the BIRT design cache
- Configuring the BIRT document cache

Configuring the data result set buffer

A memory-based buffer stores the data result set for a BIRT data object generation query. The Factory sorts, groups, and aggregates data in the buffer. By default this buffer is 128 MB. If the buffer is too small, iHub writes the data result sets to disk. Change this property to tune BIRT document generation performance on iHub under the following conditions:

- The data set of most documents is larger than the default size.
- You have configured sufficient JVM heap size to handle the size of the buffer.

Consider the number of concurrent requests for BIRT documents that the Java Factory can handle when configuring the buffer size. The Factory handles scheduled, asynchronous BIRT document generation requests one at a time.

How to configure the data result set buffer

- 1 In Server Configuration Templates—Settings, expand Factory Service, BIRT, and BIRT Generation Caches, as shown in Figure 7-7, then choose Data Set.
- 2 In Maximum result set buffer size for BIRT data object generation query, accept the default buffer size, 128 MB, as shown in Figure 7-10. Alternatively, type a larger buffer size to accommodate larger data result sets.




Figure 7-10 Configuring the BIRT report data set buffer

Choose OK.

- 3 In Server Configuration Templates—Settings, expand Factory Service, BIRT, and BIRT Generation Caches, as shown in Figure 7-7, then choose In Memory Archive File Cache.
- 4 In Maximum memory limit for each BIRT document file, accept the default, 8192 KB, or specify more space for larger documents, as shown in Figure 7-11.

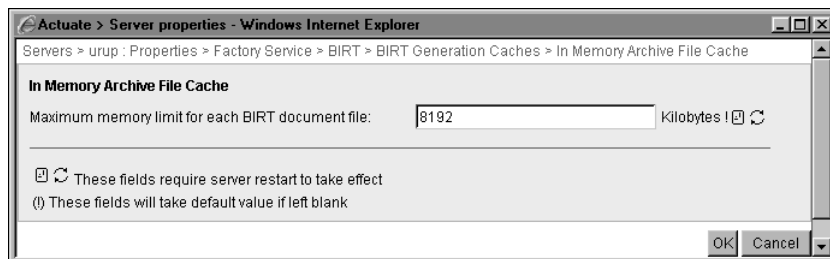


Figure 7-11 Setting maximum memory limit for a BIRT document

Choose OK.

- 5 Restart iHub.

Table 7-6 lists the property name that appears in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating the default setting, range, and when a property change takes effect.

Table 7-6 BIRT data set buffer parameter

Property name	Parameter name	Default	Range	Takes effect
Maximum result set buffer size for BIRT data object generation query	MaxBIRTData ResultsetBufferSize	10 MB	1 - 256	Server Restart
Maximum memory limit for each BIRT document file	MaxMemory PerArchive	8192 KB		Server Restart

Configuring the BIRT design cache

The administrator can configure how long the Factory keeps BIRT designs in the cache by setting a time-out value. When the time expires, the Factory clears the design from the cache. While a design remains in the cache, users who request the design share the cached file. iHub checks permissions of users to access the design. A cached design has a lifetime equal to the time-out value. Increasing the time-out value keeps designs in cache longer. The higher the time-out value, the more likely users are to access a stale design. The lower the time-out value, the sooner the Factory clears the cache.

The administrator can also configure the capacity of the cache by setting the maximum number of entries allowed in the cache. When the cache reaches capacity, the Factory stops caching designs.

How to configure the BIRT design cache

- 1 In Server Configuration Templates—Settings, expand Factory Service, BIRT, and BIRT Generation Caches, as shown in Figure 7-7, then choose Design Cache.
- 2 In Cache timeout for BIRT designs, accept the default, 1800 seconds, or 30 minutes, as shown in Figure 7-12. Alternatively, type a new value in seconds.

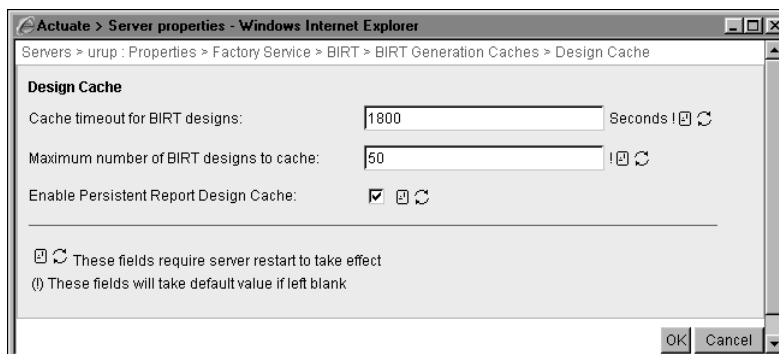


Figure 7-12 Configuring the time-out for the BIRT design cache

- 3 In Maximum number of BIRT designs to cache, accept the default, 50, or type a new value.
- 4 For Enable Persistent Report Design Cache, accept the default value of selected. Alternatively, disable the cache by deselecting this option.

Choose OK.

- 5 Restart iHub.

Table 7-7 lists the property names that appear in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating default settings and when a property change takes effect.

Table 7-7 BIRT report design cache parameters

Property name	Parameter name	Default	Takes effect
Cache timeout for BIRT designs	BIRTReportDesign CacheTimeout	1800 Seconds	Server Restart
Maximum number of BIRT designs to cache	BIRTReportDesign CacheTotalNumber OfEntries	50	Server Restart
Enable Persistent Report Design Cache	EnablePersistentDesign Cache	True	Server Restart

Configuring the BIRT document cache

By default, iHub caches a BIRT document, including access privileges. Caching benefits users who access the document concurrently. Users who request access to the same document share the cached document if they have the required privileges. Performance can improve because iHub does not have to repeatedly load the document. Generally, the fewer number of documents iHub needs to load, the better the response time. iHub caches BIRT documents in the BIRT document in-memory archive cache. To access the cache, iHub creates a handle. If you enable the BIRT document cache, iHub caches this handle in the BIRT document cache. If you do not enable the BIRT document cache, iHub creates a new handle every time a user chooses to view a document. Enabling the BIRT document cache results in a faster response time but uses more memory, because iHub maintains the BIRT document cache in memory.

How to disable caching BIRT documents

- 1 In Server Configuration Templates—Settings, expand Factory Service, BIRT, and BIRT Generation Caches, as shown in Figure 7-7, then choose Document Cache.
- 2 In Document Cache, accept the default. By default, Enable caching of BIRT document and datamart handles is selected, as shown in Figure 7-13. Alternatively, deselect the option to disable the cache.

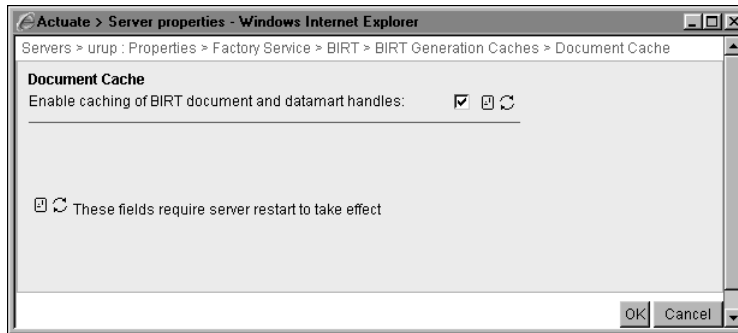


Figure 7-13 Enabling the BIRT document cache

Choose OK.

3 Restart iHub.

Table 7-8 lists the property name that appears in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating the default setting and when a property change takes effect.

Table 7-8 BIRT report document cache parameter

Property name	Parameter name	Default	Takes effect
Enable caching of BIRT document and datamart handles	BIRTReportDocumentCacheEnabled	True	Server Restart

Configuring the maximum rows in a BIRT chart

By default, BIRT charts display all data when rendered. The administrator can limit the number of rows displayed to prevent the rendering of huge charts from causing performance problems. Actuate recommends using the default setting, 0, for Maximum number of rows for a BIRT chart that displays all data when rendered.

How to limit the rows in a BIRT chart

- 1 In Server Configuration Templates—Settings, expand Factory Service and BIRT, then choose Chart, as shown in Figure 7-7.
- 2 In Maximum number of rows for generating a chart, accept the default, 0, as shown in Figure 7-14. Alternatively, type a positive integer that represents the maximum number of rows.

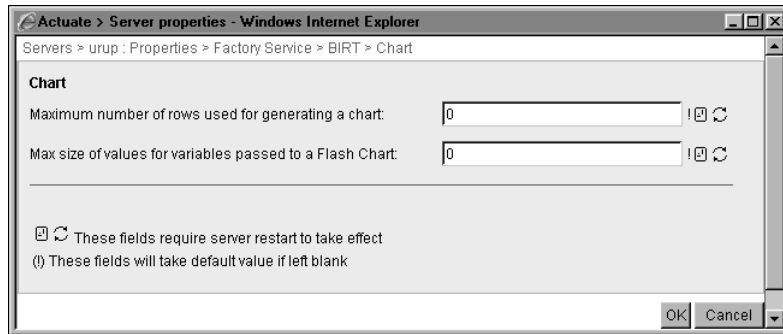


Figure 7-14 Configuring the maximum number of rows in a BIRT chart

- 3 In Max size of values for variables passed to a Flash Chart, accept the default, 0. Alternatively, type a non-zero value to limit the size of variables used in a Flash chart that appears in a BIRT document.

Choose OK.

- 4 Restart iHub.

Table 7-9 lists the property name that appears in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating the default setting and when a property change takes effect.

Table 7-9 Chart parameter

Property name	Parameter name	Default	Takes effect
Maximum number of rows for generating a chart	BIRTChartMaxRows	0	Server Restart
Max size of values for variables passed to a Flash Chart	BIRTChartMaxVariableSize	0	Server Restart

Using resource groups

This chapter contains the following topics:

- Introducing resource groups
- Configuring a resource group
- Adding a resource group
- Using a resource group
- Stopping a resource group
- Deleting a resource group
- Using resource groups programmatically

Introducing resource groups

A resource group controls the Factory processes that iHub uses to run a design and distribute a document. A resource group allocates a set of Factory processes in a stand-alone iHub or in a cluster for executing jobs assigned to the resource group. You choose synchronous or asynchronous resource groups, depending on whether you schedule the design to run.

A design that runs unscheduled runs synchronously, as soon as possible in the foreground. iHub does not store the generated document in the Encyclopedia volume. You can view, navigate, and search the document. You must enable both the Factory and View services to run designs unscheduled.

A scheduled job runs asynchronously in the background. iHub stores job schedules and information about completed jobs in the Encyclopedia volume.

Table 8-1 maps the types of resource groups to user operations in Management Console.

Table 8-1 Types of resource groups mapped to user operations

Default resource group	Resource group type	User operations
Default BIRT 360	View	Run a BIRT dashboard (.dashboard) or gadget (.gadget) design unscheduled and view the generated document.
Default BIRT Data Analyzer	View	Run a Data Object Store (.data) design unscheduled and view the generated document.
Default BIRT Factory	Async	Schedule a Java design (.rptdesign) to run right now, later, on a recurring basis, or when triggered by an event. Print a Java document.
Default BIRT Online	View	Run a BIRT design unscheduled and view the generated document.
Default BIRT Studio	View	Used when creating, modifying, and viewing documents using BIRT Studio.

You typically use resource groups for the following purposes:

- To control the load balancing of servers in a cluster
- To prioritize asynchronous jobs
- To specify which node in a cluster runs designs of a particular type

- To improve performance of a BIRT application, for example, by passing an argument to the application to decrease the heap size
- To quickly run a synchronous design using the Default Sync or Default BIRT Online resource group

When users run a design unscheduled in an Encyclopedia volume that uses multiple resource groups, iHub selects a resource group based on load balancing and on the file type of the design. If iHub cannot find a resource group with an available Factory, the job fails. Using the Default Sync or Default BIRT Online resource group can ensure the availability of a Factory for running a design unscheduled.

Configuring a resource group

You can access resource group properties from the advanced view by choosing Resource Groups, as shown in Figure 8-1.

By configuring resource group properties, you can control the following operations:

- Prevent iHub from running jobs sent to a resource group.
- Determine which Encyclopedia volumes use the Factories of a resource group.
- Set the maximum number of Factory processes that an iHub can run.
- Specify the file types that Factories allocated to a resource group can run.
- Prioritize jobs for an asynchronous resource group.
- Specify start arguments for Java Runtime Environment (JRE) for running Java designs.

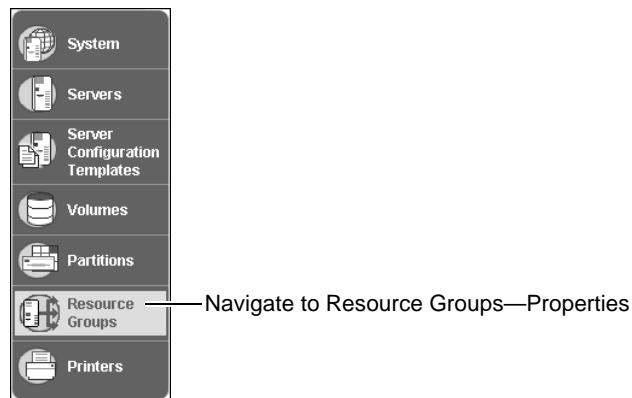


Figure 8-1 Navigating to resource group properties


Allocating Factories for a resource group

Valid values for Max factory property range from 1 to twice the number of processors on your machine. When setting this property, consider the Factories allocated for other resource groups. Also, consider using `FactoryIdleTimeout` to make unused Factories available. Increasing the number of Factories impacts performance. Setting the value to 0 disables the resource group for the server. By default, iHub allocates one Factory process to each default resource group during installation.

Limiting Java Factories for a resource group

A Factory process assigned to a resource group stays idle until iHub routes a job to the resource group. If there are no idle Factory processes for the resource group, a new Factory process starts when a job arrives. Using the Min factories property, you can optimize performance by preventing iHub from starting excessive Java Factories to run BIRT designs.

How to change properties of a resource group

- 1 From the advanced view of Configuration Console, choose Resource Groups. Figure 8-2 shows the default resource groups that iHub creates at installation.
- 2  Point to the arrow next to the resource group name, for example Default BIRT Factory, and choose Properties, as shown in Figure 8-3.

Properties—General appears, as shown in Figure 8-4.

Default resource groups






Resource Groups						
Add Resource Group						
Name	Status	Description	Type	Volume	Priority	
 Default BIRT 360	Enabled	Default resource group for BIRT 360 Dashboard.	View	<All>		
 Default BIRT Data Analyzer	Enabled	Default resource group for BIRT Data Analyzer	View	<All>		
 Default BIRT Factory	Enabled	Default resource group for BIRT Factory jobs	Async	<All>	0 - 1000	
 Default BIRT Online	Enabled	Default resource group for BIRT reports	View	<All>		
 Default BIRT Studio	Enabled	Default resource group for BIRT Studio.	View	<All>		

Figure 8-2 System Resource Groups lists resource groups






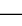
Resource Groups						
Add Resource Group						
Name	Status	Description	Type	Volume	Priority	
 Default BIRT 360	Enabled	Default resource group for BIRT 360 Dashboard.	View	<All>		
 Default BIRT Data Analyzer	Enabled	Default resource group for BIRT Data Analyzer	View	<All>		
 Default BIRT Factory	Enabled	Default resource group for BIRT Factory jobs	Async	<All>	0 - 1000	
 Properties	Enabled	Default resource group for BIRT reports	View	<All>		
 Disable	Enabled	Default resource group for BIRT Studio.	View	<All>		
 Delete						

Figure 8-3 Changing resource group properties

Resource Groups > Default BIRT Factory : Properties

General Template Assignments

Name: Default BIRT Factory

Disabled: ☐

Description: Default resource group for BIRT Factory jobs

Type: Async

Report Type: JavaReport

Volume: <All>

Work unit type: BIRT Factory

Start Arguments: -Xmx512M -XXMaxPermSize=256m -XX-UsePerfData -Dj:

Priority: Min 0 Max 1000

(!) These fields will take default value if left blank

OK Cancel Apply

Appears only when report type is Async

Figure 8-4 General properties of Default Java Async resource group

On General, you can perform the following tasks:

- Select Disabled to prevent iHub from running jobs sent to this resource group.
- All Encyclopedia volumes use the Factories of a default resource group. For a resource group that the administrator creates, select <All>, or select a particular volume to use this resource group.
- For a resource group whose report type is JavaReport, such as Default BIRT Online and Default BIRT Factory, Start Arguments appears. In Start Arguments, accept the default start arguments or change them to suit your Java Runtime Environment.
- For an Async resource group, change the minimum and maximum job priority ranges. Valid values are 0–1,000, where 1,000 is the highest priority. The minimum must be less than the maximum.

On Properties, choose Template Assignments.

- 3 Template Assignments, shown in Figure 8-5, lists the properties and names of the nodes belonging to the resource group that starts the Factory service.

Resource Groups > Default BIRT Factory : Properties

General Template Assignments

Template Name	Activate	Max factories	Min factories	Start Arguments
urup	<input checked="" type="checkbox"/>	1 *	0 *	

* These fields are required and cannot be left blank

Figure 8-5 Resource group server assignments

4 On Template Assignments, you can perform the following tasks:

- Select Activate to make that server a member of the resource group so it can use the resource group's Factory processes.
- Select the maximum number of Factory processes to assign to the resource group. For a resource group having a report type of JavaReport, type the minimum number of Factory processes also.
- For a resource group having a report type of JavaReport, enter start arguments for the Java Runtime Environment.

Restart the cluster node or master, or the stand-alone server if you change the file types or start arguments. You can enable or disable resource groups, change activation status, or change the number of Factory processes in a resource group without restarting iHub.

Setting start arguments for the JRE

The Default BIRT Factory resource group has the following start arguments by default:

- **Heap limit option**
Specifies the amount of heap the Java process can use. Too large a heap can slow garbage collection because there is more heap to scan. This property affects Java view server memory usage. Actuate sets this option to `-Xmx512M` to accommodate generating typical BIRT documents. For example, `-Xmx256m` specifies that the Java process can use 512 MB of heap.
- **MaxPermSize**
PermSize is additional heap space, separate from the space the Heap limit option specifies. The heap space that PermSize specifies holds reflective data for the JVM, such as class and method objects. By specifying MaxPermSize without also specifying PermSize, heap size does not increase unless an application needs more heap.
- **Headless graphics option**
Includes the Java graphics environment in lieu of a native graphics environment when set to true. For example, `-Djava.awt.headless=true` specifies including the Java graphics environment.
- **Protocol library specification**
For example, `Djava.protocol.handler.pkgs=com.actuate.javaserver.protocol` specifies the package name in which the Actuate protocol handler class can be found.
- **Java server entry point specification**
For example, `com.actuate.javaserver.Server` specifies the Java server main class.

You can change the start arguments for a Java resource group, as shown in Figure 8-6.

The screenshot shows the 'General' tab of the 'Template Assignments' configuration window for the 'Default BIRT Factory' resource group. The fields are as follows:

- Name: Default BIRT Factory
- Disabled: ☐
- Description: Default resource group for BIRT Factory jobs
- Type: Async
- Report Type: JavaReport
- Volume: <All>
- Work unit type: BIRT Factory (selected from a dropdown)
- Start Arguments: -Xmx512M -XX:MaxPermSize=256m -Djava.awt.headless=
- Priority: Min 0, Max 1000

A callout line points to the 'Start Arguments' field with the text: 'Appears only when report type is JavaReport'.

(!) These fields will take default value if left blank

Figure 8-6 General properties of the Default BIRT Factory resource group

Adding a resource group

You create new resource groups using Configuration Console. From the Advanced view, select Resource Groups, then choose Add Resource Group. First you set general properties in Resource Groups—New Resource Group—General. You can specify the following properties on Resource Groups—New Resource Group—General:

- The name of the resource group.
- Whether to enable or disable the resource group.
- A description of the resource group.
- The type of job that a resource group supports. The types of jobs from which to choose are asynchronous, synchronous, and view.
- The Encyclopedia volumes that can use the resource group's Factory processes.
- Start arguments for the JRE.
- Work unit type. The property specifies the type of processing this resource group can perform. For example, generating a BIRT document asynchronously requires the BIRT Factory work unit type. Generating a BIRT document immediately requires the BIRT Online work unit type.

- The job priority range for an asynchronous resource group.
The synchronous process executes designs as soon as possible and cannot assign priority to a job. If you set the type to Sync or View, iHub does not display Priority.

Figure 8-7 and Figure 8-8 contrast the differences between the general properties of an asynchronous resource group that processes Java reports and the general properties of a synchronous or view type resource group that processes Java reports.

Resource Groups > New Resource Group

General Template Assignments

Name: Custom resource group for Java report types *

Disabled: ☒

Description:

Type: Async

Report Type: JavaReport

Volume: <All>

Work unit type: BIRT Factory

Start Arguments: -Xmx512M -XX:MaxPermSize=256m -Djava.awt.headless=

Priority: Min 0 Max 1000

* These fields are required and cannot be left blank
(!) These fields will take default value if left blank

Sync, Async, or View

Java report

Appears only when type is Async

Figure 8-7 New Async resource group for Java reports

Resource Groups > New Resource Group

General Template Assignments

Name: Custom BIRT report on quarterly sales *

Disabled: ☒

Description:

Type: View

Report Type: JavaReport

Volume: <All>

Work unit type: BIRT Online

Start Arguments: -Xmx512M -XX:MaxPermSize=256m -Djava.awt.headless=

* These fields are required and cannot be left blank
(!) These fields will take default value if left blank

Appears only when report type is JavaReport

Figure 8-8 View resource group for Java reports

In Template Assignments, you set properties for a new resource group in the same way as when you change properties for an existing resource group.

How to add a resource group

- 1 From the Advanced view of Configuration Console, choose System Resource Groups.
- 2 On Resource groups, choose Add resource group.

Resource Groups—New Resource Group—General appears, as shown in Figure 8-9.

Resource Groups > New Resource Group

General Template Assignments

Name: *

Disabled: ☒

Description:

Type:

Report Type:

Volume:

Work unit type:

Start Arguments:

* These fields are required and cannot be left blank
(!) These fields will take default value if left blank

Figure 8-9 Adding a synchronous resource group

- 3 Supply the following parameter values:
- Type a name for the resource group. Choose the name carefully because you cannot change it later.
 - Deselect the Disabled option, so that cluster nodes assigned to the resource group can run jobs that users send to the group. The default status of a new resource group is Disabled.
 - Type a description of the resource group.
 - Select the type of resource group Async, Sync, or View. You cannot change the type later.
 - To specify that all Encyclopedia volumes can use the Factories of the resource group, select All. Alternatively, select the name of the Encyclopedia volume that can use the Factory processes.
 - Select a work unit type.
 - Modify Start Arguments if necessary.

- Type a minimum and maximum priority for jobs sent to an asynchronous resource group.

Choose OK.

- 4 On Resource Groups, choose the resource group you just created.

On Properties, choose Template Assignments.

- 5 In Template Assignments, change the following settings, as shown in Figure 8-10:

- Select Activate to make that server a member of the resource group so it can use the resource group's Factory processes.
- Select the maximum number of Factory processes to assign to the resource group.
- For resource groups that support Java reports, type the minimum number of Factory processes.
- Specify Start Arguments to optimize performance.

Template Name	Activate	Max factories	Min factories	Start Arguments
urup	<input checked="" type="checkbox"/>	0 *	0 *	

* These fields are required and cannot be left blank

Figure 8-10 Making server assignments to the new resource group

Choose OK.

- 6 Restart the cluster node or stand-alone server.

Using a resource group

A user can specify a resource group in a job schedule to assign priority to a job. To specify a resource group to handle a job, specify the group as described in this section.

Selecting a resource group for a job

Log in to Management Console and choose Files and Folders. Next, point to the arrow next to the name of a design and choose Schedule. On Schedule—Schedule, select a resource group from Resource Group, as shown in Figure 8-11.

corp > Home > administrator > MyCustomers (RPTDESIGN) (Version 1) : Schedule

Schedule Parameters Output Privileges Datamart Security Channels Notification Print

Job name:

TimeZone:

Run job:

- ☒ Right now
- ☐ Once: date # time (M/d/yyyy h:mm a)
- ☐ Recurring: time (h:mm a)
- ☐ Advanced: - ☐ Wait for event: Event name:

Priority: ☐ Low (200) ☐ Medium (500) ☐ High (800) ☒ Other (1 - 1000): (1)

Resource Group:

Executable version: MyCustomers.rptdesign

- ☒ Always use version of MyCustomers.rptdesign

Retry failed jobs:

- ☒ Use volume default
- ☐ Retry times; wait hours minutes between attempts
- ☐ Do not retry

(1) This job will use the lower priority of this setting and the one assigned to you in your user profile.

Prioritizes the job

Available resource groups

Figure 8-11 Specifying a resource group for a job

Prioritizing a job

You specify the priority that the Factory process gives a job when you create a resource group. iHub tries to match the priority setting in a resource group with the priority setting in the job schedule. The following conditions can occur:

- If no match is found, the job stays in a pending state until you change the priority specifications in the resource group.
- If the only available resource group is disabled, iHub sets the job to pending until you enable an asynchronous resource group for the Encyclopedia volume.
- If multiple resource groups have the same job priority settings and an available Factory process, iHub chooses a resource group that best balances the load.

About Factory and View service requirements

After you create a custom resource group for a cluster, you need to understand the Factory and View service requirements for nodes assigned to the resource group. It is also important to consider the effect that using resource groups has on performance.

You must enable only the View service on iHub to fulfill requests to view persistent documents. Persistent documents are documents that are saved on a volume. You must enable both the View and Factory services on iHub to fulfill all other requests to generate and view documents.

Managing loads with resource groups

You need to understand and avoid the following problems that can affect performance:

- **Unequal loads**
In a cluster, iHub uses load-balancing mechanisms to distribute jobs among the cluster nodes. The load-balancing mechanisms attempt to maximize performance. Creating custom resource groups can restrict the capability to maximize performance and adversely affect the document-generation performance of a cluster.
- **Configuration problems**
If a cluster node is a member of a resource group that is not configured to run all the executable design file types, designs can fail. For example, if a design requires access to a database, and the database driver is not installed on a node, the design fails when that node attempts to run the design.

Understanding the Java View service

The Java View service is capable of handling multiple on-demand design execution and document viewing requests concurrently. By default, the service can handle 10 requests concurrently and can queue up to 1,000 requests.

Initially, configure a Java View resource group for a two-CPU system, to use just a few factories, and then tune the system under load.

Stopping a resource group

Stop a resource group from running jobs by performing one of the following actions:

- **Disabling the resource group**
- **Setting the resource group's maximum number of Factory processes to 0 for all nodes that belong to the resource group**
- **Removing active membership of nodes**

Disabling an asynchronous resource group is the same as setting the maximum number of Factory processes to 0. If all resource groups are disabled, jobs sent to

the disabled asynchronous resource groups go into a pending state until a resource group becomes available.

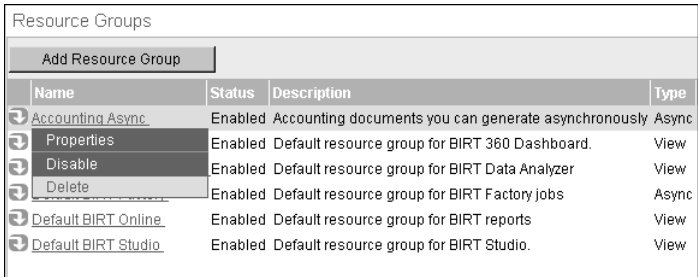
If you disable a synchronous resource group, it processes jobs that are currently being executed and those that are waiting. If you set the maximum number of Factory processes to 0 on any of the nodes, jobs wait to be executed until they time out.

Removing all active memberships from a resource group is the same as setting the maximum number of Factory processes to 0 on all nodes in a cluster.

If you remove a resource group from an Encyclopedia volume, you must assign a resource group with available Factory processes to the Encyclopedia volume. Otherwise, you cannot run a job. For example, if you change a resource group Encyclopedia volume assignment from volume1 to volume2, you must ensure that another resource group can handle the jobs volume users create in volume1. If you remove an Encyclopedia volume assigned to a resource group, iHub changes the Encyclopedia volume’s resource group assignment to All volumes and disables the resource group.

Deleting a resource group

Delete a resource group by pointing to the arrow next to the resource group and choosing Delete, as shown in Figure 8-12.









Resource Groups			
Add Resource Group			
Name	Status	Description	Type
 Accounting Async	Enabled	Accounting documents you can generate asynchronously	Async
 Properties	Enabled	Default resource group for BIRT 360 Dashboard.	View
 Disable	Enabled	Default resource group for BIRT Data Analyzer	View
 Delete	Enabled	Default resource group for BIRT Factory jobs	Async
 Default BIRT Online	Enabled	Default resource group for BIRT reports	View
 Default BIRT Studio	Enabled	Default resource group for BIRT Studio.	View

Figure 8-12 Deleting a resource group

Deleting a resource group produces the following results, depending on the state of related jobs:

- If a job is already running on a Factory assigned to a resource group that you delete, the job completes.
- If a scheduled job is assigned to a deleted resource group, the job either fails when iHub runs the job, or it remains in a pending state, depending on the job status at the time of deletion.

You can delete a pending job on Jobs—Pending.

- If an unscheduled job is assigned to the deleted resource group, the job fails when iHub runs the job.

If a job is running on a Factory assigned to a resource group that you delete, the job completes.

Using resource groups programmatically

Developers can write applications using resource groups with Actuate Information Delivery API (IDAPI). Using IDAPI, you can:

- Send a job directly to a resource group, bypassing the priority settings for asynchronous jobs and the Encyclopedia volume settings for synchronous jobs.
- Change the resource group configuration. For example, you can enable or disable resource group membership or change the number of Factory processes in a resource group.

Using IDAPI to direct jobs to a specific resource group bypasses the cluster's load-balancing mechanisms. As a result, applications that use IDAPI to run designs can adversely affect the performance of the cluster.

9

Clustering

This chapter contains the following topics:

- About a cluster configuration
- About distributing services to a cluster
- About the configuration home directory
- About the primary configuration files
- Creating a cluster
- Adding and modifying server templates
- Adding a node to a cluster
- Starting and stopping a node
- Removing a node from a cluster
- Managing a cluster
- Handling file system failure
- Configuring the cluster administrator e-mail account
- Managing console configurations and load balancing

About a cluster configuration

In the simplest configuration, an iHub cluster consists of two iHub nodes. A node is a single machine in a cluster. You can add nodes to a cluster to scale iHub System to your requirements. You install iHub on each node in a cluster. The node gets its configuration from a template in `acserverconfig.xml`, which is located in a shared configuration home directory. After the node is configured, it joins the cluster.

Figure 9-1 shows the relationships between a cluster and its Encyclopedia volumes, application servers, and database servers.

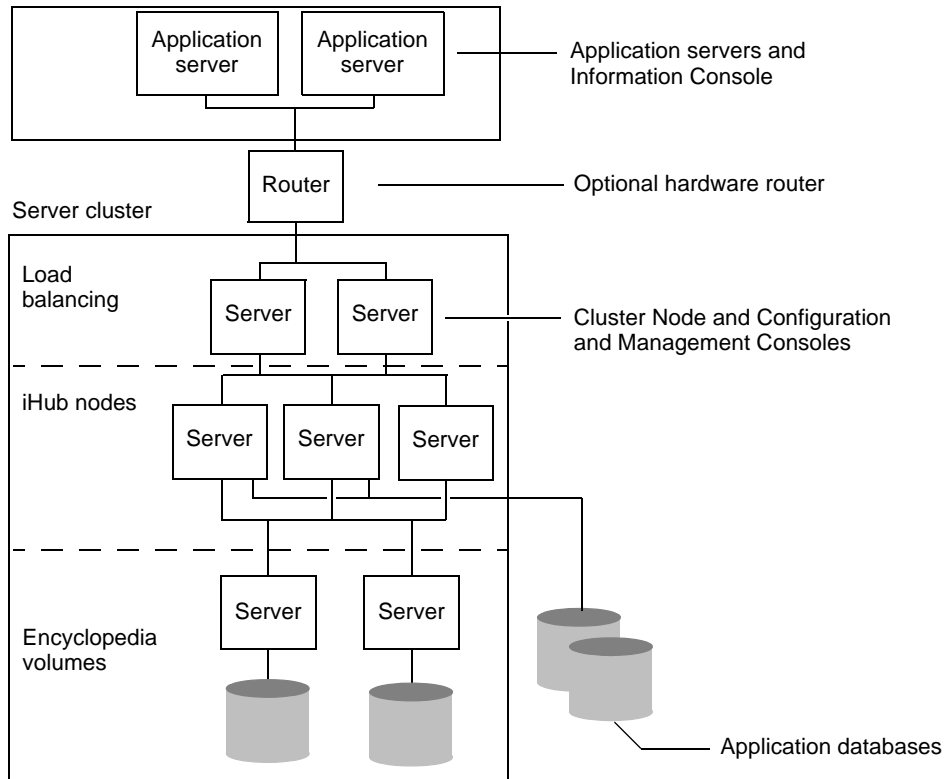


Figure 9-1 An iHub cluster

In this example, the cluster uses a network router to create a single virtual IP address to distribute the load-balancing requests that come into two nodes. The Actuate Information, Configuration, and Management Consoles support distributing requests to multiple machines, which handle load balancing in the cluster. Figure 9-1 shows the following cluster sections:

- In the load-balancing section, the requests are routed to a node that performs load balancing.
- In the nodes section, iHub generates documents and delivers them to clients for viewing.
- In the Encyclopedia volumes section, iHub running on multiple machines maintains Encyclopedia volume management information and controls access to the volumes. The volumes can be on machines that are not running iHub but are accessible to a machine. A node shares all volumes.

In Figure 9-1, separate machines handle separate functions. You can combine these functions on one machine.

About distributing services to a cluster

You can control the configuration of a cluster and each node in a cluster, including:

- iHub services and settings
- Services and settings for a cluster node
- Cluster membership

You can enable one or more services in each server template.

In a cluster, you can use templates to configure the nodes instantiated in an iHub System, as shown in Figure 9-2. Each node is a computer running iHub, configured through the template definitions stored in `acserverconfig.xml`, which is accessible by all nodes through a shared directory. The following letters represent the available services in Figure 9-2:

- M is the Message Distribution service.
- V is the View service.
- F is the Factory service.
- I is the Integration service.
- C is the Caching service.

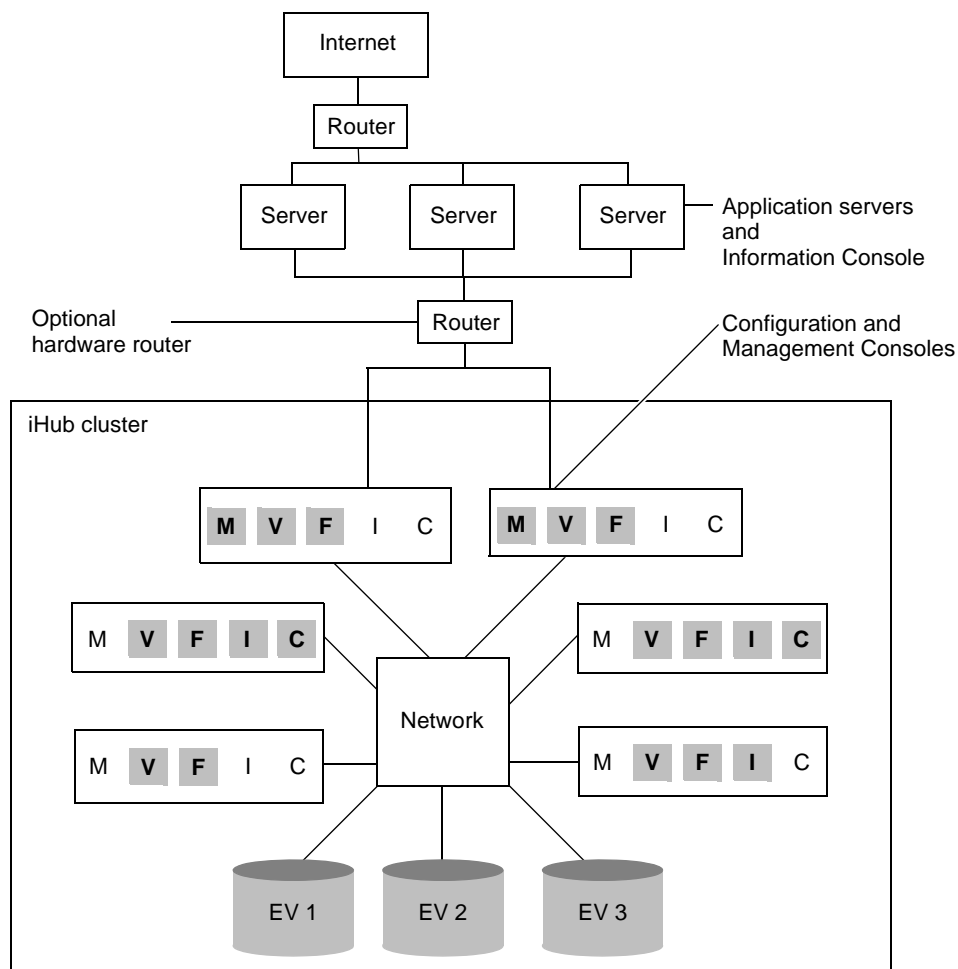


Figure 9-2 An iHub cluster configuration

A letter in a gray box represents an enabled service. The other letters represent disabled services. When you enable or disable a service in a template, the cluster nodes that use different templates can have different functionality.

Two nodes have the Message Distribution service enabled in the cluster in the preceding example. In the example, the two machines share one virtual IP address to communicate with the router.

Information Console or iHub perform load balancing independent of load-balancing capabilities in the router. All requests to the iHub cluster go to one of the two redirector nodes, which are the nodes with the Message Distribution service enabled. The redirector nodes dispatch the requests to other nodes in the cluster based on message type and cluster load.

The configuration file, `acserverconfig.xml`, located in a shared location provides centralized management of the cluster configuration through the use of server templates and other parameters. A server template is a set of parameters that specifies the configuration properties of a node or set of nodes.

The `acpmdconfig.xml` file for each node, located in the `AC_SERVER_HOME/etc` directory on the node machine, has its `<AC_CONFIG_HOME>` element set to point to the location for the shared `acserverconfig.xml` file. The `<AC_TEMPLATE_NAME>` element specifies which template the node uses.

One or more nodes in the cluster manage request message routing. The Message Distribution service uses HTTP to communicate between nodes in a cluster. An iHub cluster supports multicast network communication.

About the configuration home directory

The configuration home directory is the directory that holds the configuration file, `acserverconfig.xml`, and the licensing file, `acserverlicense.xml`.

The `acserverconfig.xml` file contains the information for connecting to Encyclopedia volumes and printers, message distribution state, configuration settings for Factory, View, Integration, and Caching services, the location of partitions, and licensing and open security information.

In a cluster, `acserverconfig.xml` plays a central role in the operation of cluster nodes. The file includes server templates for various server roles with each template containing configuration information for connecting to Encyclopedia volumes, printers, and services for all the nodes in the cluster. When a node joins a cluster, it first configures itself using its template in the `acserverconfig.xml` file located in the shared directory.

The `acserverconfig.xml` file is located by default in `AC_DATA_HOME/config/iHub2`. This directory is referred to as the configuration home directory. Make the configuration home directory sharable to allow cluster nodes access to this file. To specify the location of this directory for each node, modify the value of the `<AC_CONFIG_HOME>` element located in `acpmdconfig.xml`, which by default is located in `AC_SERVER_HOME/etc` of the node.

How to change the location of the configuration home directory

To change the location of the configuration home directory containing `acserverconfig.xml` and `acserverlicense.xml`, perform the following tasks:

- 1 Shut down all cluster nodes.
- 2 Stop the Actuate BIRT iHub service on each node.
- 3 Back up the configuration home directory.
- 4 Move the configuration folder to the new destination.

- 5 Share the configuration folder.
- 6 For every node dependent on this configuration home directory, update the <AC_CONFIG_HOME> element located in the node's acpmdconfig.xml file.
- 7 Start Actuate BIRT iHub service for each node.

About the primary configuration files

There are two main configuration files:

- acpmdconfig.xml
Located by default in AC_SERVER_HOME/etc
- acserverconfig.xml
Located by default in AC_DATA_HOME/config/iHub2

In a cloud configuration, server templates give cloud the flexibility to change configurations at the launch time. Acpmdconfig.xml contains the configurations that are node specific. When the Process Management Daemon (PMD) starts up, it reads these configurations first and exposes them to the process environment variable list.

The acserverconfig.xml file contains other cluster and node configuration parameters, which specify the host names, port numbers, volume names, and server templates to be used by the nodes. The following sections elaborate on these configuration files.

About acpmdconfig.xml

Acpmdconfig.xml sets the environment variables at the operating system level. The administrator can specify the node configuration settings in acpmdconfig.xml. Alternatively, the administrator can set the necessary environment variables in setsrvrenv and run this script before restarting an instance.

Using acpmdconfig.xml, set <AC_CONFIG_HOME> to the directory that contains acserverconfig.xml and acserverlicense.xml. Modify the <AC_TEMPLATE_NAME> element to specify which template this iHub uses.

When creating an image for a set of machines, create a template name, such as SharedTemplate, in acserverconfig.xml. In acpmdconfig.xml, set <AC_TEMPLATE_NAME> to the template name. The administrator can create different templates that use different images for machines with varying computing capacities and resources. Listing 9-1 shows acpmdconfig.xml for a machine named urup, a node containing the configuration home directory.

Listing 9-1 The acpmdconfig.xml file

```
<PMDConfig>
  <!--Actuate system Type -->
  <System>Cluster</System>
  <Mode>Default</Mode>
  <PMDConfigFileVersion>2</PMDConfigFileVersion>
  <!--Daemon SOAP endpoint information -->
  <DaemonSOAPPort>8100</DaemonSOAPPort>
  <!-- Disk Thresholds are in MB -->
  <MinDiskThreshold>100</MinDiskThreshold>
  <LowDiskThreshold>300</LowDiskThreshold>
  <!--Server information -->
  <Server>
    <Startup>Auto</Startup>
    <AC_TEMPLATE_NAME>urup</AC_TEMPLATE_NAME>
    <AC_DATA_HOME>C:\Actuate\iHub\data</AC_DATA_HOME>
    <AC_CONFIG_HOME>C:\Actuate\iHub\data\config\iHub2
      </AC_CONFIG_HOME>
    <AC_JRE_HOME>C:\Program Files (x86)\Common Files\Actuate
      \22.0\JDK160\jre</AC_JRE_HOME>
    <AC_JRE64_HOME>C:\Program Files (x86)\Common Files\Actuate
      \22.0\JDK160_64\jre</AC_JRE64_HOME>
    <AC_JAVA_HOME>C:\Program Files (x86)\Common Files\Actuate
      \22.0\JDK160</AC_JAVA_HOME>
    <AC_ODBC_HOME>C:\Program Files (x86)\Common Files\Actuate
      \22.0\odbc</AC_ODBC_HOME>
    <AC_SERVER_IP_ADDRESS>urup</AC_SERVER_IP_ADDRESS>
    <AC_SOAP_DISPATCH_ADDRESS>urup</AC_SOAP_DISPATCH_ADDRESS>
    <AC_DOC_BASE>http://www.actuate.com/documentation/iHub2
      </AC_DOC_BASE>
    <AC_ICU_DATA>C:\Program Files (x86)\Actuate\iHub2\bin
      </AC_ICU_DATA>
  </Server>
  <!-- Servlet Container information -->
  <ServletContainer>
    <Startup>Auto</Startup>
    <JavaOpts
      Args="-Xms128m -Xmx512m -XX:MaxPermSize=128m"/>
    </ServletContainer>
</PMDConfig>
```

Use this configuration file to set the environment variables at the operating system level. Table 9-1 shows the configuration setting in acpmdconfig.xml.

Table 9-1 acpmdconfig.xml node-specific configuration settings

Configuration name	Description
AC_TEMPLATE_NAME	Template name
AC_DATA_HOME	Location of data directory
AC_CONFIG_HOME	Location of acserverconfig.xml and the license file
AC_JRE_HOME	Location of Java run-time environment (JRE)
AC_JRE64_HOME	Location of 64-bit Java run-time environment (JRE)
AC_ODBC_HOME	Location of ODBC resources
AC_SERVER_IP_ADDRESS	iHub IP address
AC_SOAP_DISPATCH_ADDRESS	iHub dispatcher IP address
AC_ICU_DATA	Location of ICU library

About acserverconfig.xml

This configuration file can specify one or more templates to provide flexibility when instantiating iHub nodes in an environment where machines have varying resources. Acserverconfig.xml provides access to the following elements:

- System
- FileSystems
- MetadataDatabases
- Schemas
- Volumes
- Templates
- Resource groups
- Printers
- ServerList

When starting a cluster, the acserverconfig.xml file must be in a shared directory.

Listing 9-2 The acserverconfig.xml file

```
<Config>
  <System
    ClusterID="_4_ffffefdfc_ce4fdb2c_e4c"
    SystemName="corp"
    DefaultLocale="en_US"
```

```

        DefaultEncoding="windows-1252"
        ConfigFileVersion="13"
        EncyclopediaOwnerID="_6_fffe9dfc_ce4fdb2c_e4c"
        SystemDefaultVolume="corp"
        ClusterDatabaseSchema="ac_corp_system"
        DefaultCLocaleOnWindows="true"
        EncyclopediaVolumeServer="urup">
    <UsageAndErrorLogging/>
    <SMTPServers/>
</System>
...
<Templates>
    <Template
        Name="urup"
        PMDPort="8100"
        ActuateBuild="220A130126"
        ActuateVersion="2"
        ServerSOAPPort="11100"
        AppContainerPort="8900"
    ...
    </Template>
    <Template
        ...
    </Template>
</Templates>
</Config>

```

Creating a cluster

To create a cluster, the administrator first installs a stand-alone iHub or uses an existing installation. Next, the administrator shares the configuration home directory, so other servers joining the cluster can access it. Then, the administrator adds new nodes to form the cluster.

There are two methods of adding a new node to the cluster:

- Through the use of the configuration file, `acpmdconfig.xml`
- Through the installation wizard, when performing a custom cluster-node installation

Every cluster node must have network access to the following directory and resources to join the cluster:

- The shared configuration home directory
- Cluster resources, such as printers, database systems, disk storage systems, and Encyclopedia volume directories

The administrator can configure nodes, using server templates, to run different services and to process different types of requests. Important factors to consider when configuring nodes include processing power and access to hardware and software resources, such as printers and database drivers.

From the configuration console the administrator can add resources, such as partitions, Encyclopedia volumes, and resource groups to the cluster.

Creating an initial cluster

- 1 Install two stand-alone iHubs. In this example, the server names are urup and kozu.
- 2 On urup, share the following folders so that they are available to other servers:
 - \config\iHub2
 - \encyc
- 3 If urup and kozu are both Windows machines, turn off the Windows firewall on both urup and kozu.
- 4 Confirm that urup and kozu can access each other on the network. In a command prompt on each computer, type the ping command followed by the IP address or host name of the other computer. Verify that the receiving computer replies to the sending computer.
- 5 Shut down urup and kozu by performing the following tasks on each machine:
 - 1 Log in to Configuration Console.
 - 2 On the simple view, choose Stop system, as shown in Figure 9-3.

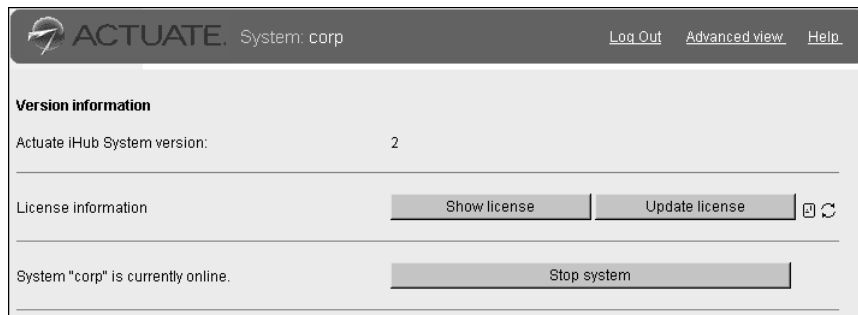


Figure 9-3 Stopping the system

- 6 In a Windows environment, using Administrative Tools—Services, stop the Actuate BIRT iHub services for urup and kozu on each server machine.

7 On urup, make a backup of acserverconfig.xml, then open it and perform the following tasks:

1 Locate the <ServerFileSystemSettings> element under the <Template> element.

2 Under <ServerFileSystemSettings>, locate:

```
<ServerFileSystemSetting  
    Name="DefaultPartition"  
    Path="$AC_DATA_HOME$/encyc"/>
```

Using Universal Naming Convention (UNC) format, change the value of the Path attribute to specify the location of DefaultPartition.

DefaultPartition contains the Encyclopedia data files. By default, the DefaultPartition path is:

AC_DATA_HOME\encyc

To specify this path using UNC format, type:

\\urup\encyc

3 Locate the <ConnectionProperties> element under the <MetadataDatabase> element.

4 Under <ConnectionProperties> locate:

```
<ConnectionProperty  
    Name="server"  
    Value="localhost"/>
```

5 Change Value from localhost to the name of the machine on which the volume resides, in this example “urup”. For urup, the volume is located on localhost, but in a cluster setting the administrator must use the machine name, since it is not on a localhost from the reference point of the other nodes.

6 Save acserverconfig.xml file.

8 On kozu, open acpmdconfig.xml located in AC_SERVER_HOME/etc and perform the following tasks:

1 Change <AC_CONFIG_HOME> to point to the path specified by <AC_CONFIG_HOME> in acpmdconfig.xml on urup. For example, if <AC_CONFIG_HOME> on urup is set to AC_DATA_HOME\config\iHub2, type \\urup\iHub2 as the value for <AC_CONFIG_HOME> on kozu.

2 Change <AC_TEMPLATE_NAME> to the name specified by <AC_TEMPLATE_NAME> on urup.

9 Save the modified acpmdconfig.xml file.

10 Start the Actuate BIRT iHub service for urup.

11 After urup starts, start the Actuate BIRT iHub service for kozu.

12 Log in to Configuration Console for either node.

13 Choose Advanced view.

14 In Advanced view, select Servers from the side menu.

Servers displays the nodes, urup and kozu. Both servers are using the urup template, as shown in Figure 9-4.

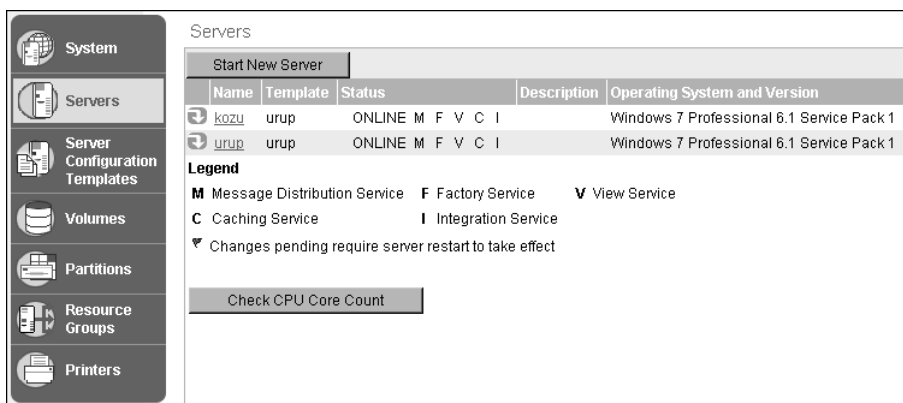


Figure 9-4 A two-node cluster

When adding a node to an already existing cluster, it is not necessary to shut down the cluster before adding the new node.

Configuring heartbeat messaging

Nodes in a cluster use heartbeat messaging to monitor the status of the other nodes in the cluster. On System—Properties, the administrator configures System Heartbeat, as shown in Figure 9-5.

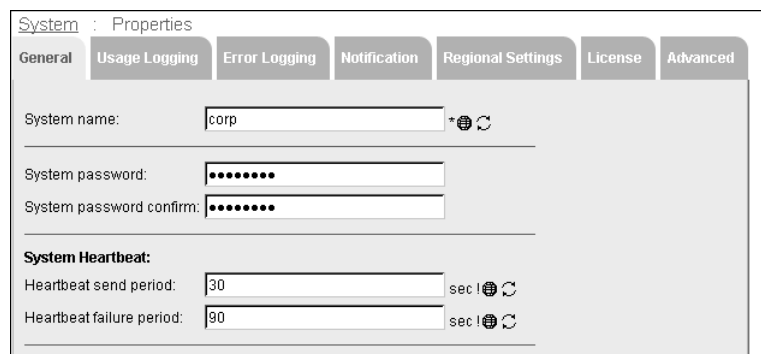


Figure 9-5 Configuring System Heartbeat

System Heartbeat consists of properties for measuring messaging frequency. System Heartbeat properties are:

- Heartbeat send period
The interval to send a heartbeat message, typically 30–40 seconds. Change this property to increase or decrease the number of heartbeat messages.
- Heartbeat failure period
The period in which the cluster nodes determine that another node is within the heartbeat failure period, typically 90–100 seconds. If the monitoring nodes do not receive a response within this period, the monitored node is assumed to be down.

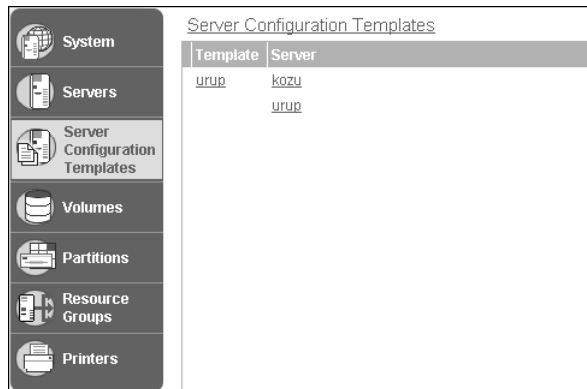
Configuring Message Distribution service properties

In BIRT iHub, the administrator configures the Message Distribution service for a server template rather than for an individual node.

How to configure the Message Distribution service for a server template

To configure the Message Distribution service (MDS) for a particular template, perform the following tasks:

- 1 Log on to Configuration Console on urup, and choose Advanced view.
- 2 From the side menu, choose Server Configuration Templates. Select the template for which you want to configure the Message Distribution service, as shown in Figure 9-6.



Template	Server
urup	kozu
urup	urup

Figure 9-6 Server templates

- 3 In Advanced, expand the Message Distribution Service folder and the Process Management folder, as shown in Figure 9-7.

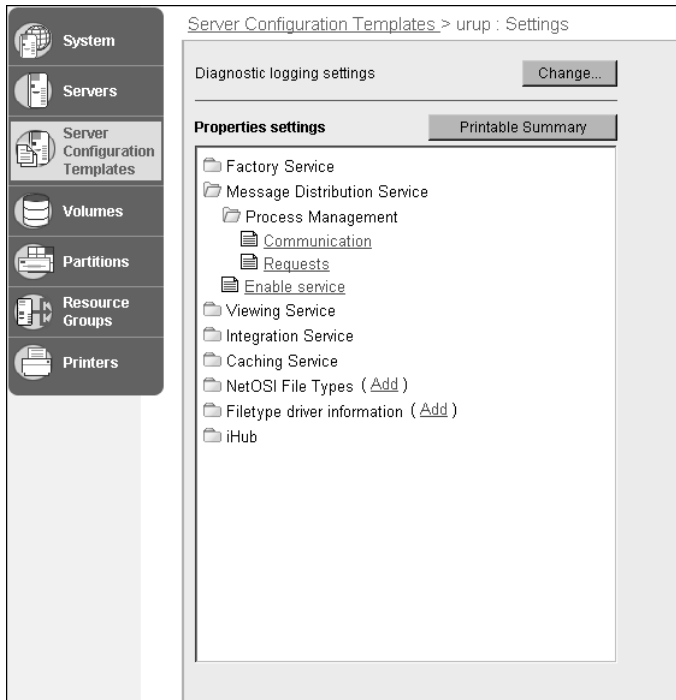


Figure 9-7 Message Distribution Service

- 4 In Message Distribution Service—Process Management—Communication, as shown in Figure 9-8, accept or set the value for the Message Distribution service port.

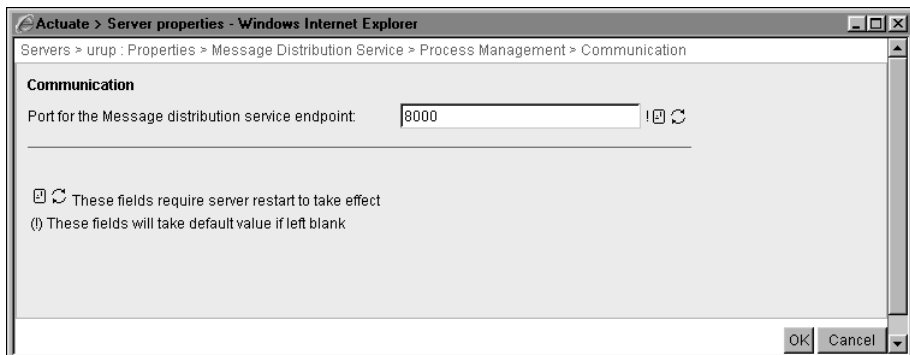


Figure 9-8 Setting communication properties

- 5 In Message Distribution Service—Enable service, enable or disable the Message Distribution service by selecting or deselecting Enable request service, as shown in Figure 9-9.




Figure 9-9 Enabling and disabling the Message Distribution service

- 6 In Message Distribution Service—Process Management—Requests, accept or set the value for the total number of concurrent requests processed by iHub through the Message Distribution service, as shown in Figure 9-10. When the maximum number of requests is reached, iHub refuses to process any new incoming requests.

If you do not use the recommended default value of 1000, perform throughput tests after changing the value. If the value is 0, iHub does not accept any incoming requests. If the value is too low, iHub does not utilize all system resources when handling the maximum number of client requests. Increasing the value can reduce total throughput.

At peak usage times, iHub might need the maximum available system resources, such as system memory and temporary disk space, to handle requests.

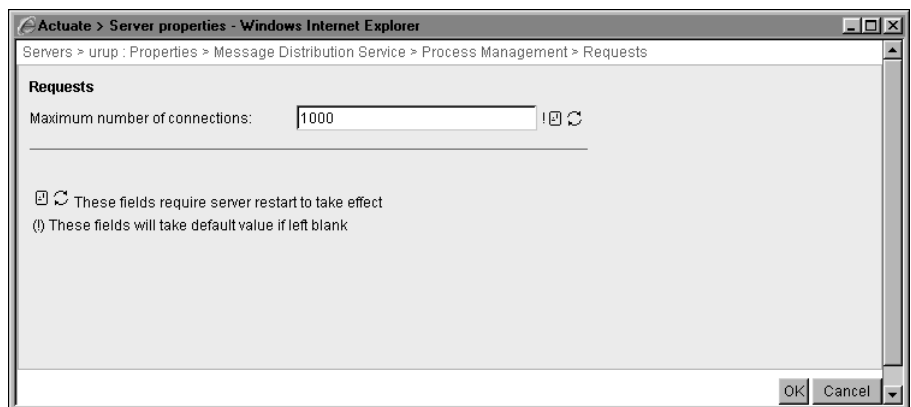


Figure 9-10 Specifying the maximum number of connections

You can also modify these settings through `acserverconfig.xml`, for each template.

Adding and modifying server templates

Every cluster node takes its configuration from the template to which it is assigned. To change a node configuration, the administrator either modifies the assigned template or creates a new template and assigns it to the node. The administrator creates a server template definition in the `acserverconfig.xml` file located in the shared configuration home directory.

The administrator can create a server template for any possible server configuration. The following sections cover the process of modifying an existing template, and creating a new template.

Modifying a server template

Back up the original `acserverconfig.xml` located in the shared configuration home directory, then open the file and perform the following tasks:

- 1 Shut down the nodes that use the template you want to modify.
- 2 Open `acserverconfig.xml`.
- 3 Under the `<Templates>` tag, locate the template you want to modify.
- 4 Locate and modify the elements associated with the parameters you want to configure.
- 5 After making the modifications to the template, save `acserverconfig.xml`.
- 6 Start the nodes, which now use the modified templates.

Creating a new server template

The easiest way to add a new template is to copy an existing template, and then modify it for new functionality.

- 1 Back up the original `acserverconfig.xml` file.
- 2 Open `acserverconfig.xml`.
- 3 Locate an existing template that closely approximates the new template you want to create.
- 4 Copy the template, then paste it under the last `</Template>` tag above the `</Templates>` tag.
- 5 Modify the template name attribute, to give the new template its own unique name.
- 6 Modify the parameters of the elements to produce the functionality you desire.
- 7 Save `acserverconfig.xml`.

The template is now ready for use by cluster nodes. To use the new template, modify the <AC_TEMPLATE_NAME> for the node in acpmdconfig.xml file to the new template name.

Adding a node to a cluster

To add a new node to the cluster perform the following tasks:

- 1 Shut down the node.
- 2 Open the acpmdconfig.xml file.
- 3 Modify <AC_CONFIG_HOME> to point to the shared configuration home directory that contains acserverconfig.xml for the cluster.
- 4 Modify <AC_TEMPLATE_NAME> to use a server template from the available server templates listed in acserverconfig.xml.
- 5 Save the acpmdconfig.xml file.
- 6 Start the new node. The node will automatically contact the acserverconfig.xml for the configuration profile and join the cluster, as shown in Figure 9-11.

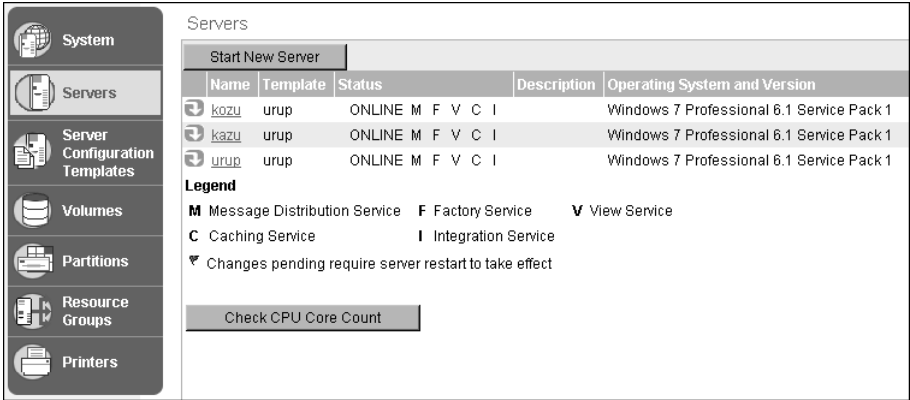


Figure 9-11 A three-node cluster

If the particular server configuration that you want to use for the new node is not available in the templates listed in acserverconfig.xml, create a new template. For more information on the process of creating a new server template, see “Creating a new server template,” earlier in this chapter.

About node configuration

Every node is configured by its server template. To reconfigure a node, the administrator can either change the template that the node is using, or modify the

existing template. Use caution when modifying a template, because you are modifying the configuration of all the nodes that subscribe to the template.

You configure templates to perform tasks, such as generating and printing documents. Ensure that the cluster meets the following requirements:

- Configure access to printers and databases from templates that run the Factory service.
- Configure resource groups.

About cluster configuration

You can make the following configuration changes to a cluster:

- Add more nodes to run Factory and View services to handle an increased workload.
- Add new nodes to the cluster to increase robustness, and take over the functionality in the case of a node failure.

Testing a cluster

You can test a cluster using a URL, such as the following one:

```
http://host1:8900/<context root>/login.jsp?serverURL=  
http://host2:8000&daemonURL=http://host3:8100
```

where

- context root is acadmin, the name for Management Console.
- 8900 is the default port number for the application container running Management Console.
- 8000 is the default port number for iHub.
- 8100 is the default port number for the iHub service on Windows and the iHub service Process Management Daemon on Linux.

This URL tests a three-node cluster that runs Management Console on host1, a node iHub on host2, and the iHub service on host3.

Starting and stopping a node

When iHub starts, it gathers the following licensing information:

- Validity of the license file
- Encyclopedia volume limitations

- CPU limitations, if applicable
- License expiration date, if applicable

You start or stop the nodes in a cluster using Servers in Configuration Console.

You can start and stop a node independent of the cluster.

How to stop a single cluster node

- 1 From the advanced view of Configuration Console, choose Servers.
- 2 In the list of servers, point to the arrow next to the iHub name that you want to stop.
- 3 Choose Stop, as shown in Figure 9-12. iHub changes the status of the server.

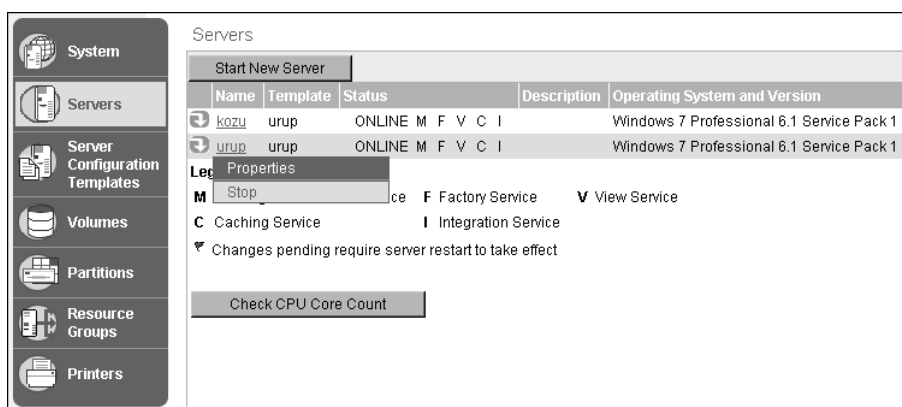


Figure 9-12 Starting or stopping a single node

About starting a node

When you start an iHub node, the following events occur:

- The node configures itself by contacting acserverconfig.xml located in a shared configuration home directory.
- The node joins the cluster.
- The services and resource groups of the node become available to the cluster.
- The status of the node changes to online.

How to start a node

To start an offline node which is already part of the cluster, perform the following tasks:

- 1 From the advanced view of Configuration Console, choose Servers. Point to the arrow next to the offline node, and choose Start, as shown in Figure 9-13.

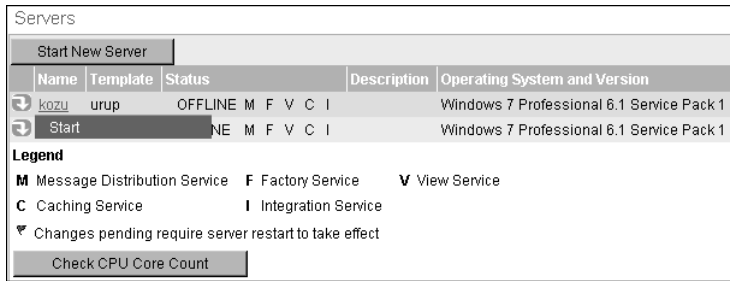


Figure 9-13 Choosing Start

- 2 In Servers—Start server, accept the default values for Host Name, iHub Process Manager Port Number, and Server Template, as shown in Figure 9-14. Alternatively, you can change these values.

The screenshot shows the 'Start server: kozu' form. It has four fields: 'Server name:' with value 'kozu', 'Host Name or IP Address:' with value 'kozu', 'iHub Process Manager Port Number:' with value '8100', and 'Server template name:' with value 'urup'. Each field has a required field icon (asterisk) and a refresh icon. A note at the bottom states: '* These fields are required and cannot be left blank' and '* These fields require server restart to take effect'.

Servers > Start server: kozu

Server name: kozu

Host Name or IP Address: kozu * [refresh]

iHub Process Manager Port Number: 8100 * [refresh]

Server template name: urup * [refresh]

* These fields are required and cannot be left blank
* These fields require server restart to take effect

Figure 9-14 Starting a node

- 3 Choose Start. After a few moments, the node comes online using the chosen server template.

How to start a new node

To start an offline node that is not currently part of the cluster, but is configured to be part of the cluster, perform the following tasks:

- 1 In the advanced view of Configuration Console, choose Servers, then choose Start New Server.
- 2 In Servers—Start New Server, enter the values for the server name, host name, iHub process manager port number, and choose a server template, as shown in Figure 9-15.

The screenshot shows the 'Start New Server' form. It has four fields: 'Server name:' with value 'kozu', 'Host Name or IP Address:' with value 'kozu', 'iHub Process Manager Port Number:' with value '8100', and 'Server template name:' with value 'urup'. Each field has a required field icon (asterisk) and a refresh icon. A note at the bottom states: '* These fields are required and cannot be left blank' and '* These fields require server restart to take effect'.

Servers > Start New Server

Server name: kozu

Host Name or IP Address: kozu * [refresh]

iHub Process Manager Port Number: 8100 * [refresh]

Server template name: urup * [refresh]

* These fields are required and cannot be left blank
* These fields require server restart to take effect

Figure 9-15 Starting a new node

- 3 Choose Start. After a few moments, the node comes online using the chosen server template, and joins the cluster.

About stopping a cluster or a node

Shutting down a cluster differs from stopping a node in a cluster. To shut down a node, from the node's Configuration Console, on System—Status, choose Stop. When you stop a node, the following events occur:

- The node leaves the cluster.
- The services and resource groups on the node are no longer available to the cluster.
- The status of the node changes to offline.

To shut down a cluster, you shut down all the nodes forming the cluster.

Removing a node from a cluster

After you remove the node from the cluster, iHub can operate in a stand-alone configuration only if you originally installed it in a stand-alone configuration. When you remove a node from a cluster, the following events occur:

- iHub removes information about the node from the cluster configuration information.
- When you log in to Configuration Console for the cluster, you cannot access the node.

To remove a node from a cluster, complete the following tasks in the following order:

- Log in to Configuration Console for the cluster.
- Reconfigure the cluster to handle functionality assigned to the template the node used. For example, ensure that remaining nodes in the cluster can assume responsibility for the following functionality:
 - Running Message Distribution, View, Factory, Integration, and Caching services
 - Fulfilling resource group requirements
 - Other functionality that the node supports
- Stop the system on the cluster node. When you take the node offline, the cluster can no longer access its services.

- Remove the node from the cluster by performing the following tasks:
 - 1 Stop the system on the node containing the shared configuration home directory.
 - 2 In Windows Explorer, navigate to AC_DATA_HOME\config\iHub2. Open acserverconfig.xml in a text editor, such as Notepad.
 - 3 Find the <ServerList> element, located towards the end of the file.
 - 4 Select and delete the <ServerInfo> element for the cluster node that you are removing.
 - 5 Save the modified acserverconfig.xml file.
 - 6 Select Start to start the server system, as shown in Figure 9-16.



Figure 9-16 Starting the server

- 7 Choose Servers from the side menu. The name of the removed cluster node no longer appears in the server list.

Removing a node from the cluster does not change the configuration information for the cluster. For example, cluster nodes can still access the following resources, even after you remove the node from the cluster:

- Printers the cluster uses
- Partitions that refer to directories on the hard drive that the cluster can access

Managing a cluster

It is important to consider the following when managing a cluster:

- Accessing partitions
- Specifying the Encyclopedia volume location
- Cluster option requirements

Accessing partitions

All cluster nodes must have access to resources that the cluster uses to read and write to every partition in the cluster. In Add Partition, you specify a path to the

partition, as shown in Figure 9-17. In BIRT iHub, the partition path is set for each template rather than for each individual node.

Template Name	Partition Path
urup	\urup\encyc2

Figure 9-17 Adding a partition

Specifying the Encyclopedia location

The default installation of a stand-alone iHub configures an Encyclopedia volume, but the installation of a cluster node does not include an Encyclopedia volume. From Configuration Console, you can add a new volume. In iHub, every server serves all volumes.

You specify the location of an Encyclopedia volume for a cluster in a template. When the node accesses the template to which it was assigned, it also receives the location of the partitions.

About cluster option requirements

A cluster uses a single license file. All nodes in a cluster have the same iHub System options enabled. Any node that you add to the cluster uses only the options that the cluster uses. For example, if you add an iHub that was in a stand-alone configuration with page-level security to the cluster, and the cluster does not have page-level security enabled, the new node does not support page-level security.

To support multiple Encyclopedia volumes in a cluster, you need the Multi-Tenant option or an equivalent. For example, you need the Multi-Tenant option to use both Encyclopedia volumes in a cluster that you create from multiple stand-alone iHubs, each with their own Encyclopedia volume.

When iHub is offline, you can replace the license file in the configuration home directory.

Handling file system failure

iHub handles file system failure on stateless and stateful file systems. This overview uses Network File System (NFS) as an example of a stateless network file system and Common Internet File System (CIFS) as an example of a stateful network file system.

iHub handles some file system failures by retrying file I/O. Retrying file I/O works when a file system failure is transparent to iHub. For example, on an NFS-based network storage system, a file system failure can be transparent to iHub. Retrying file I/O is insufficient in a configuration where file system failure is not transparent to iHub, such as on a Windows-based CIFS file system.

On a stateless file system such as an NFS-based file system, iHub can handle a network storage system failure. The machine detects that the connection to a file system is lost and attempts to reconnect. When the file system recovers, the machine re-establishes a connection to the file system. If the connection to the file system does not time out during failure, iHub does not detect the failure.

On a stateful file system such as a Microsoft Windows-based CIFS network file system, a machine using the file system tracks file system connection states, including open files and locks. If the file system connection breaks, the machine loses connection state information. The CIFS client machine must manually reestablish file system connections. iHub can re-establish file system connections on a stateful network file system.

iHub identifies a file system failure as a failure of the following file I/O functions:

- Reading the configuration lock file
- Reading the Encyclopedia volume lock file
- Reading or writing to an Encyclopedia volume

Failure to read the configuration lock file affects the cluster nodes. The other two I/O failures affect the Encyclopedia volume.

Configuring the cluster administrator e-mail account

In System—Properties—Advanced—Cluster Operation—Administrative, shown in Figure 9-18, you can specify administrative e-mail account information.

The AdminEmail property specifies the account that receives administrative e-mail, such as the notice iHub sends when licensing problems occur.

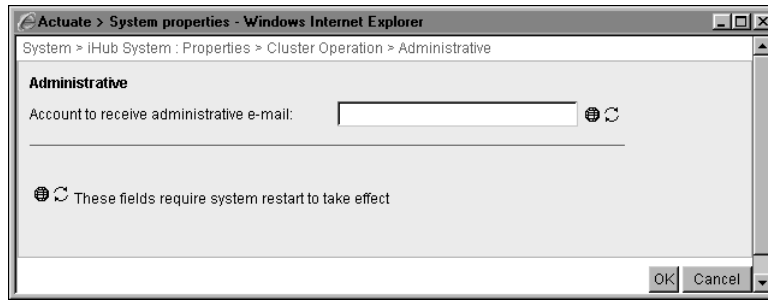


Figure 9-18 Specifying administrative e-mail account information

Managing console configurations and load balancing

You can configure iHub and Configuration Console and Management Console in several ways. The following sections describe the most common configurations.

Using the consoles directly

This configuration places both iHub and its Configuration and Management Consoles behind a corporate firewall, as shown in Figure 9-19.




Figure 9-19 Using the consoles directly

Perform a stand-alone installation to set up this configuration.

Using the consoles through a firewall

Figure 9-20 shows iHub and Configuration and Management Consoles deployed behind a second firewall.




Figure 9-20 Using the consoles through a firewall

Deploy the load-balancing proxy separately from the application container, which exists in the demilitarized zone (DMZ). You can use a proxy from a third-party vendor or the Actuate proxy.

To install this configuration, you must complete the following tasks in this order:

- Install iHub, integrated with the consoles. Ensure that each console uses the same context root name. The default name is acadmin.
- Deploy the Actuate or third-party proxy on a machine that exists in the DMZ.
- Configure the proxy with the list of available nodes.
Management and Configuration Consoles should be installed on the nodes.

Using multiple console installations

You can have multiple console installations in a cluster. The proxies maintain the session stickiness and distribute requests to the available installations of iHub consoles. The load balancer that iHub uses in this scenario does not have to support session stickiness.

Figure 9-21 illustrates a configuration with multiple installations of iHub consoles.

To set up this configuration, complete the following tasks in this order:

- Install the cluster, with the integrated consoles.
- Deploy the Actuate or third-party proxy on different machines than those where you installed iHub consoles.
- Configure the proxy with the list of available nodes with iHub consoles.

If you remove a node from a cluster, install the consoles for that iHub if they are not already installed.

About load balancing

iHub consoles ensure high availability and distribute tasks for efficient processing using load balancing. The installation of iHub consoles include a lightweight web application that distributes requests to the available iHub nodes in a cluster.

When deployed, this proxy enables load balancing among Management Console instances. The proxy is available as `mgmtconsoleproxy.war` (web archive) in the downloaded iHub product package. To use the proxy, install it under its own context root on your application server.

The easiest way to customize the proxy for your installation is to modify the WAR file to use a modified `web.xml` file for each installation.

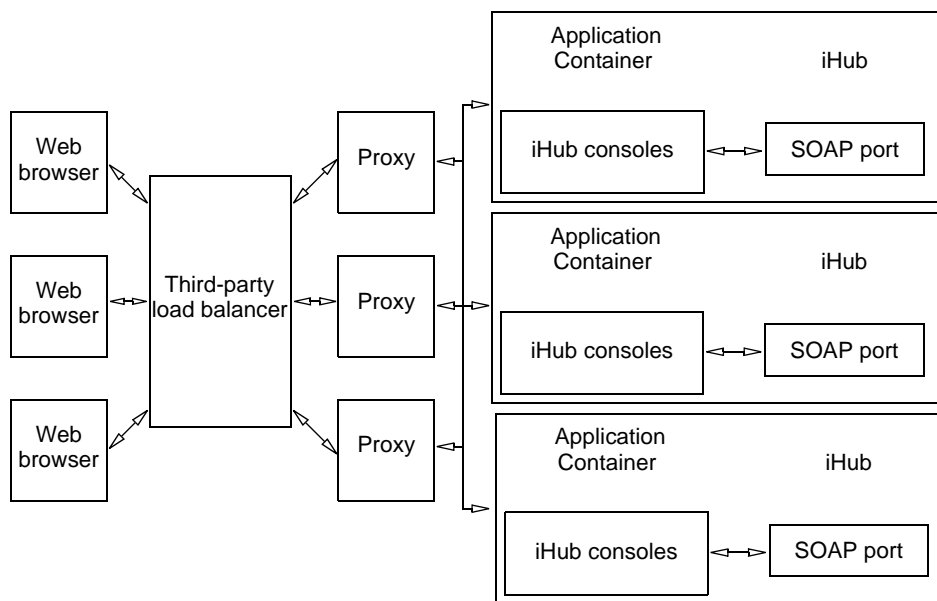


Figure 9-21 Using multiple iHub consoles

Deploying load balancing

Before you use the WAR file to deploy the Management Console load-balancing application, AcProxy, customize the WAR file for your installation as follows:

- Extract the configuration file, web.xml, from the WAR file.
- Customize web.xml.
- Create a new WAR file using the customized web.xml file.

How to extract the configuration file from the WAR file

- 1 Create a temporary directory, such as C:\temp\AC on a Microsoft Windows server or /temp/ac on a Linux server.
- 2 Decompress the mgmtconsoleproxy.war file onto the temporary directory.

For example, on Windows, open a Command window and type the following commands, replacing the X: drive letter with a drive letter appropriate to your system:

```

mkdir C:\temp\AC
copy X:\mgmtconsoleproxy.war .
jar -xf mgmtconsoleproxy.war

```

The AcProxy files appear in the temporary directory.

Using a Linux server, type the following commands:

```
mkdir /temp/ac  
cp /dev/dsk/iHubProduct/mgmtconsoleproxy.war .  
jar -xf mgmtconsoleproxy.war
```

The AcProxy files appear in the temporary directory.

- 3 If you used the temporary path in step 3, the file location is:

Windows—C:\temp\AC\WEB-INF\web.xml

Linux—/temp/ac/WEB-INF/web.xml

How to customize web.xml for your installation

- 1 Using a text editor that accepts UTF-8 encoding, edit web.xml to configure AcProxy for your application server.
- 2 Specify the list of available nodes with iHub consoles in the SERVER_LIST context parameter. The list contains all iHub console URLs that AcProxy uses to balance requests. The default list is empty. Add your URLs in a comma-separated list. The following code is a sample SERVER_LIST entry:

```
<context-param>  
  <param-name>SERVER_LIST</param-name>  
  <param-value>http://hostname1:8900/,http://hostname2:8900/  
  </param-value>  
</context-param>
```

- 3 Save and close web.xml.

How to customize the WAR file for your installation

To create a new WAR file, type the following command:

```
jar -cf ../newmgmtconsproxy.war *
```

This command creates Newmgmtconsproxy.war in the /temp directory. This file is a new WAR file for AcProxy, and it contains the modified configuration values.

Use Newmgmtconsproxy.war to deploy to your application servers instead of Mgmtconsoleproxy.war.

Configuring Integration and Caching services

This chapter contains the following topics:

- About the Integration service
- Managing Integration service resources
- Using information objects
- About Actuate Caching service
- Configuring the Caching service

About the Integration service

Whenever possible, the Integration service uses a single database connection for a query. For example, if a design uses information object data sources that depend on multiple data connection definition files (.dcd), the Integration service may use only one database connection if the DCDs have the same values for the User name, Password, and Server properties and the query can be completely pushed to the database. For concurrent queries, however, the Integration service creates multiple database connections.

Managing Integration service resources

The administrator configures Integration service properties on Server Configuration Templates—Advanced to control how iHub and an Integration handle data from a data source. These settings also determine how the process manages RAM memory and disk-based memory when processing information object data. Figure 10-1 shows the properties in Server Configuration Templates—Settings—Integration Service.

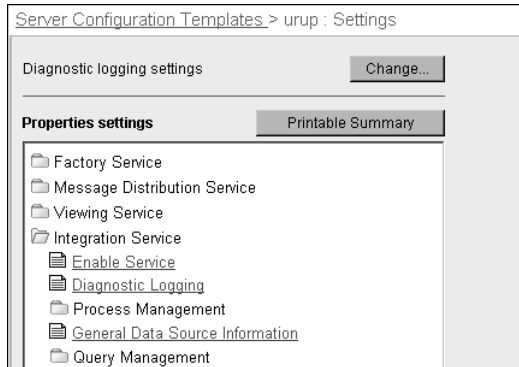


Figure 10-1 Viewing Integration service properties

Enabling the Integration service

The administrator can enable or disable the Integration service from Server Configuration Templates—Settings, or by setting the `EnableIntegrationService` parameter in `AC_DATA_HOME/config/iHub2/acserverconfig.xml` to true or false.

How to enable the Integration service

- 1 Expand Integration Service, as shown in Figure 10-1, and choose Enable Service.

- 2 For Enable Integration service, accept the default value, which is selected, as shown in Figure 10-2.

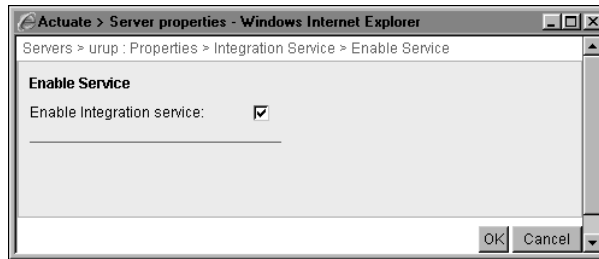


Figure 10-2 Enabling or disabling the Integration service

About diagnostic logging

The administrator can configure diagnostic logging by expanding Integration Service, and choosing Diagnostic Logging. For more information, see Chapter 3, “Using diagnostic, usage, and error logging.”

Setting port numbers for process communication

In Server Configuration Templates—Settings—Integration Service—Process Management—Communication, as shown in Figure 10-3, the administrator accepts or sets the value for each of the following properties:

- Port for Integration server message endpoint
Port for Integration service process communication with Encyclopedia volume processes
- Port for Integration server query endpoint
Port for Integration service process communication and iHub when running information objects to pass queries and data

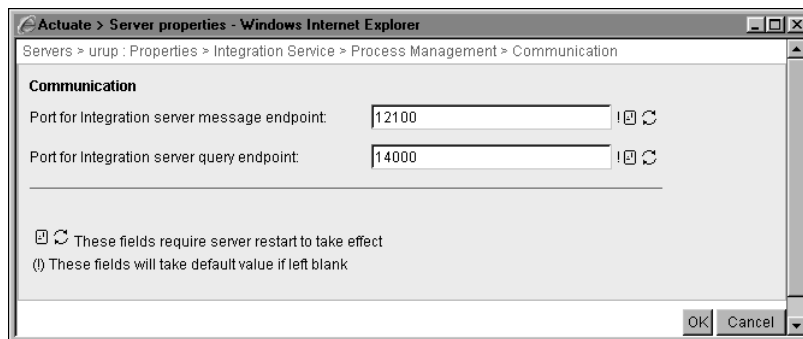


Figure 10-3 Setting port numbers for Integration service process communication

Specifying load settings

In Server Configuration Templates—Settings—Integration Service—Process Management—Requests, as shown in Figure 10-4, the administrator accepts or sets the value for each of the following properties:

- Max SOAP requests
Maximum number of connections for communicating with Encyclopedia volume processes
- Max SOAP request threads
Maximum number of threads for communicating with volume processes

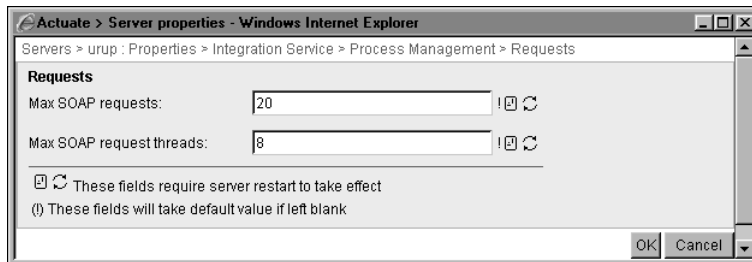


Figure 10-4 Specifying load settings

Managing Actuate SQL query execution

In Server Configuration Templates—Settings—Integration Service—General Data Source Information, shown in Figure 10-5, the administrator sets properties to specify the data source, including database collation.

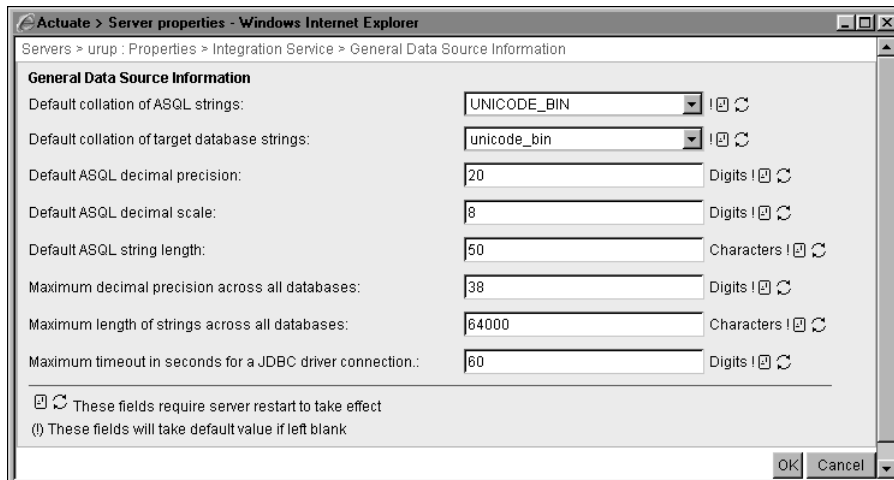


Figure 10-5 Specifying Integration service general data source information

Collation is an algorithm for ordering strings. When an Actuate SQL query executes, the type of collation determines the result of sort and comparison operations. Databases support one or more collations, usually determined by the database locale. The Integration service, however, supports only the Unicode and ASCII code-point collations, which order strings based on the Unicode or ASCII numbers corresponding to each character.

Based on properties the administrator sets on Servers—Properties—Advanced—Integration Service—General Data Source Information, the database collation and the Integration service collation determine which operations are sent to the database and which the Integration service performs. These properties are:

- **Default collation of ASQL strings**
Specifies database collation of an Actuate SQL query. You can specify one of the following values:
 - **UNICODE_BIN**, the default, sets unicode code point order (binary order). All characters are different from one another and are sorted by their unicode values.
 - **ASCII_CI** sets code point order. Uppercase characters have the same value as lowercase characters.
- **Default collation of target database strings**
Specifies the type of collation used by the database. Refer to your database documentation to determine the appropriate category for your database collation. You can specify one of the following values for the Default collation of target database strings property:
 - **unicode_bin**, the default, specifies that the collation of target strings is the same as the Integration service collation, **UNICODE_BIN**.
 - **ascii_ci** specifies that the collation of target strings is the same as the Integration service collation, **ASCII_CI**.
 - **null** specifies that the collation of target strings does not correspond to either the **UNICODE_BIN** or **ASCII_CI** used by the Integration service collation. Each character has a unique value.
 - **null_sensitive** also specifies that the collation of target strings does not correspond to either **UNICODE_BIN** or **ASCII_CI** used by the Integration service collation, but more than one character can have the same value, for example 'E' = 'e'.
- **Default ASQL decimal precision**
Specifies the maximum number of digits after the decimal point in an Actuate SQL query.

- **Default ASQL decimal scale**
Specifies the maximum scale for NUMERIC and DECIMAL types in an Actuate SQL query. For example, 15 represents decimals that can have up to 15 digits in all, including decimals after the decimal point.
- **Default ASQL string length**
Specifies the maximum size in an Actuate SQL query for CHAR, VARCHAR, and LONGVARCHAR data types.
- **Maximum decimal precision across all databases**
Specifies the AIS maximum decimal precision.

If the value is greater than the cache database maximum precision, the cache database might return truncation errors when creating a cache due to a mismatch between the precision of the AIS data and the maximum precision of a cache database. To avoid truncation errors, use the Actuate SQL CAST() function to change the precision of the data.
- **Maximum length of strings across all databases**
Specifies the maximum string length of AIS strings.

If the value is greater than the cache database maximum string length, the cache database might return truncation errors due to the mismatch between the length of the AIS data and the maximum string length of a cache database. To avoid truncation errors, use the Actuate SQL CAST() function to change the length of the data.
- **Maximum timeout in seconds for a JDBC driver connection**
Specifies how long iHub waits to connect to a database over JDBC. The JDBC driver never times out when connecting to the database if you set this value to 0.

Specifying query settings

In Server Configuration Templates—Settings—Integration Service—Query Management—Query Execution, shown in Figure 10-6, the administrator accepts or sets the value for each of the following properties:

- **Maximum run time for the query**
Specify a non-zero value to set a limit on the amount of time a query can run.
- **Maximum query size in rows**
Specify a non-zero value to set a limit on the number of rows a query returns.
- **Max memory per query**
Specify a non-zero value to set a limit on the amount of memory a query can consume.

- **Max Fetch Scroll Memory Size per Query**
Defines the maximum size of the memory buffer that binds to the ODBC per query. Change this value if there is sufficient memory on the machine. Use trial and error to determine the ideal setting.
- **Max Rows Fetched Per Fetch Scroll**
Determines the total number of ODBC rows retrieved for each call of SQLFetchScroll and sets a limit on the memory usage per execution. Use trial and error to determine the ideal number of rows to retrieve per call.
- **Max ODA Fetch Size Per Query**
Sets the fetch buffer size when fetching data from an ODA driver.

These settings are per connection, therefore all queries on the particular connection use the same setting. Using a value other than the default on upper limit could result in increased memory usage by iHub.

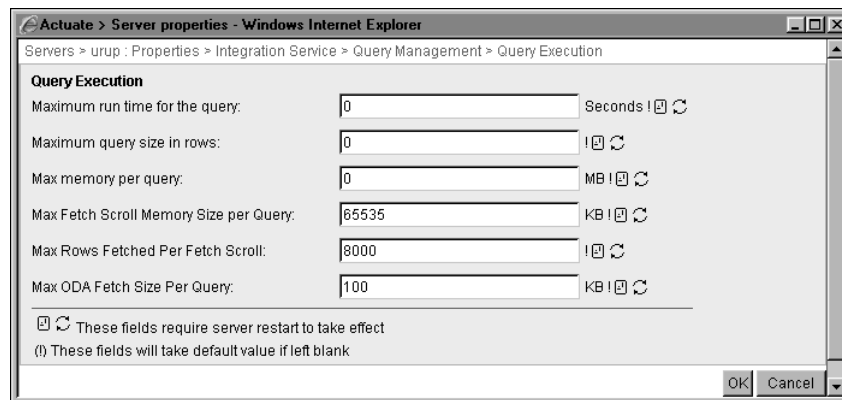


Figure 10-6 Specifying Integration service query settings

Specifying query optimization settings

In Server Configuration Templates—Settings—Integration Service—Query Management—Query Optimization, shown in Figure 10-7, the administrator sets properties to improve the overall system performance. iHub reduces the execution time and system resource usage for queries across multiple data sources.

The properties are:

- **Enable cost-based optimization**
Determines whether to apply the cost-based query optimization.
Change the default value when the query plans generated with cost-based optimization are inefficient or when you cannot provide the cost information for the information objects.

- Minimum rows to trigger creation of an index during materialize operation
Creates an index on materialization that improves performance of searching for matched tuples in the materialization.

When the materialization has more rows than this parameter value, iHub creates the index.

Increase the default value when iHub executes too many concurrent queries, exhausting the memory for indexing and impacting overall system performance.

Reduce the number of queries that invoke the creation of an index, thereby slowing the execution of some queries and reducing system resource usage.

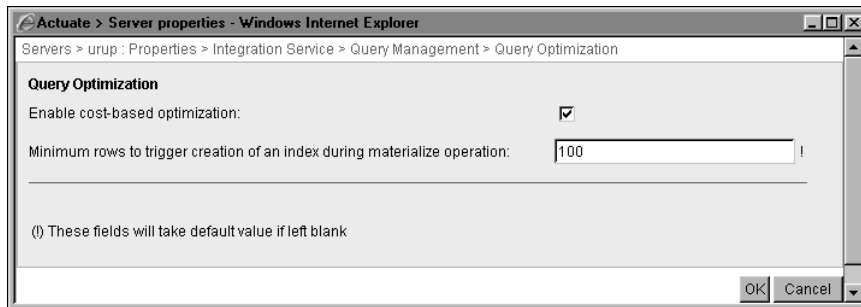


Figure 10-7 Specifying Integration service query optimization settings

Setting resource management properties

In Server Configuration Templates—Settings—Integration Service—Query Management—Resource Management, shown in Figure 10-8, the administrator accepts or sets values for each of the following Integration service resource properties:

- Temporary storage path for spill to disk
Path to the directory that holds temporary files used by an Integration service process. The default directory is the temp subdirectory of the iHub home directory.
- Minimum Disk Threshold for spill partition
Amount of memory that must be consumed before iHub writes overflow to disk.
- Page pool size
The maximum page size required to process the information object data.
- Buffer pool size
The maximum buffer pool size required to process information.

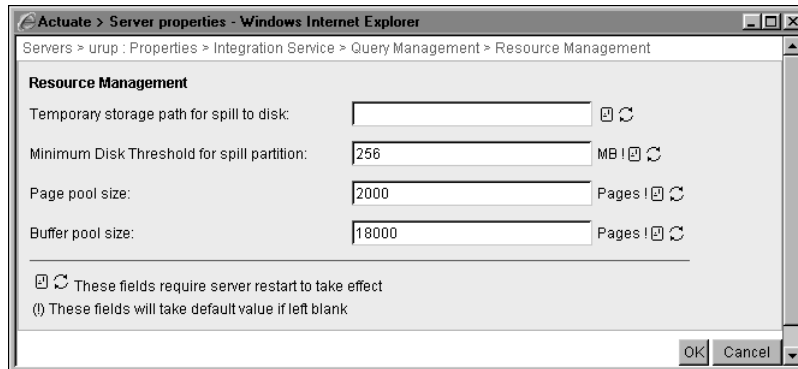


Figure 10-8 Specifying the Integration service resource property values

The Integration service process uses disk-based files to store temporary data when processing an information object that requires a large amount of memory.

When allocating disk space for the directory, consider the maximum amount of memory required to process the information object data and the maximum number of concurrent information objects that iHub can run.

Setting query statistics logging properties

In Server Configuration Templates—Settings—Integration Service—Query Management—Query Statistics, shown in Figure 10-9, the administrator accepts or sets the value for each of the following query statistics logging properties:

- **Enable query statistics logging**
Enables and disables logging of the statistics.
- **Query statistics log level**
 - Standard logs query-level statistics.
 - Info logs query- and operator-level statistics.
 - Detail logs query-, operator-, and database-level statistics.

The cumulative execution time is logged at the Info level.




Figure 10-9 Specifying query statistics logging settings

Using information objects

Information architects use information objects to control information retrieved from data sources, consolidate information from multiple data sources, and cache data from remote data sources for offline use.

To ensure data is available for an information object job, such as a BIRT design that uses information objects, you can cache data. When caching data, iHub uses one or more external databases and manages each information object cache independently. Future queries using that information object use the data stored in the cache database.

iHub stores cached information objects in one or more of the following databases:

- DB2
- Oracle
- MS SQL

Caching information object data optimizes the timeliness of data, the load constraints of operational data stores, and response time. For example:

- Response time is faster when you execute large volumes of requests against data in the cache database instead of querying production data sources.
- Performance can improve when you:
 - Schedule population of the cache during non-peak traffic hours.
 - Populate the cache incrementally instead of retrieving all the rows of the output of an information object.
 - Retrieve increments to the output since the last time the cache was updated.

Setting up caching

In iServer Release 9, the Encyclopedia volume can cache data using a single cache connection to a database. In iServer Release 10 and later, you can set up multiple cache connections to multiple databases.

About setup in Actuate Information Object Designer

In iServer Release 9, you run SQL scripts to create an information cache definition (ICD) file. In iServer Release 10 and later, an information architect creates an information cache definition file using Actuate Information Object Designer. For each ICD, the information architect specifies the name of its cache data table, the names of the columns of that table, and any table indexes.

About setup in iHub

From Configuration Console, you configure properties on Server Configuration Templates—Settings—NetOSI File Types—ICD to communicate between the information object and Actuate Caching service (ACS) using SOAP. From Management Console, you activate the cache when scheduling a job. iHub executes the ICD schedule and populates the cache table. When you run a query against the information object, iHub retrieves data from the cache table.

You need to provide the information architect with the location and connection properties to use for caching, so the information architect can publish either information objects or designs that contain them to iHub. The information objects become available in the Encyclopedia volume, and designs can use them as data sources.

Updating cache files

iHub supports caches and cache definitions created in Actuate 11 and later releases. Ensure that the ICD files in the Encyclopedia volume are updated regularly. To update an information object cache, run or schedule the ICD as a job using Management Console.

When iHub updates the cache data, it updates all of the data from the data source, not just the new or changed data. An information object can access updated information object data when the job completes. Until the job completes, the information object uses the existing cache data. When the information architect deletes cache definitions from an Information Object Designer project, you must delete the cache definitions from the Encyclopedia volume.

To move an information object project to a new location within the same Encyclopedia volume, use Management Console to move a folder to the new location.

About Actuate Caching service

Actuate Caching service manages one or more configurable ACS databases, and performs functions such as adding and dropping databases, tables, and indexes, and inserting data into databases. Actuate Integration service (AIS) and ACS are persistent, multithreaded processes that accept multiple simultaneous requests. A Factory process communicates directly with AIS through the ODA interface, not through the Message Distribution service (MDS).

In a stand-alone configuration, iHub runs single AIS and ACS processes. In a cluster, each node can run AIS and ACS processes. A cluster distributes information object jobs among the cluster nodes with AIS enabled. You start and stop the AIS and ACS processes by enabling or disabling the AIS and ACS

services. When using ACS in a cluster, Actuate recommends enabling both ACS and the Factory service on the cluster node.

iHub uses the AIS to run an information object. Figure 10-10 shows the communication among components when using a cache database.




Figure 10-10 iHub component communication

Both the Factory process and an Actuate desktop application access an information object and the Actuate Integration service using the ODA interface. The AIS ODA driver supports the following design-time and run-time interfaces of ODA:

- The ODA driver requests a connection to an iHub, then creates a connection to that host if it has AIS enabled.
- For a Factory process, the ODA driver uses the AIS server specified by the configuration. The default configuration specifies that all Factory processes use the AIS server.

The following types of files play a role in caching information objects:

- DCD (data connection definition)
Contains properties to identify and connect to a particular external data source, such as a database or an external application.
Properties consist of data source type, connection properties, and pass-through security type.
- ICD (information object cache definition)
Contains configuration information for caching data that a data source map SMA or information object IOB uses. Each cache is stored in an external DBMS configured for the Encyclopedia volume.
Properties consist of SMA information and cache state information.

- IOB (information object)

Presents a view of a data source, a logical set of data from other maps or information objects. Supports row-level security.

Properties consist of schema information such as parameter and column names, other properties such as a query, editing state, and caching configuration.

- SMA (data source map)

Represents a single set of data from a data source, such as a database table.

Properties consist of schema information such as parameter and column names, information to access a data source, and caching configuration.

You cannot use more than one version of a DCD, ICD, IOB, or SMA file in an Encyclopedia volume. If an information object uses another information object or map as input, that relationship information is internal to the information object and is specified only by name. The information object does not use Encyclopedia version information, such as file ID or version number.

If you enable the Factory service and disable the ACS on a cluster node, it must use another node's Caching service. To specify a node having the Caching service enabled, use either one of the following methods to modify the configuration of the node that has the disabled Caching service:

- Change the hostname parameter in Servers—Properties—Advanced—NetOSI File Types—ICD to the name of the machine with the ACS enabled. The default hostname value is localhost.
- Remove the ICD file type from the server node using Servers—Properties—Advanced—NetOSI File Types—ICD (Delete).

Configuring the Caching service

The administrator can improve the database performance by using the Actuate Caching service (ACS) to cache information object data in the iHub environment. Caching information object data eliminates repetitive queries, reducing the load on the network and data source.

In Servers—Properties—Caching Service, the administrator can change the following properties:

- Caching service port
iHub uses this port for communication between the ACS process and Encyclopedia volume.

- **Caching service request settings**
The maximum number of SOAP requests and request threads for communicating with an Encyclopedia volume process.
- **Bulk load settings**
Path to the directory of the intermediate files ACS uses to perform a bulk load to a data source, such as an Oracle, DB2, or SQL Server database. iHub saves temporary files only when an ACS bulk load job fails. Keep temp files instructs iHub to always or never save temporary files or only save these files when a job fails.

Figure 10-11 shows the selections the administrator makes in Server Configuration Templates—Settings to configure properties for the Caching service.

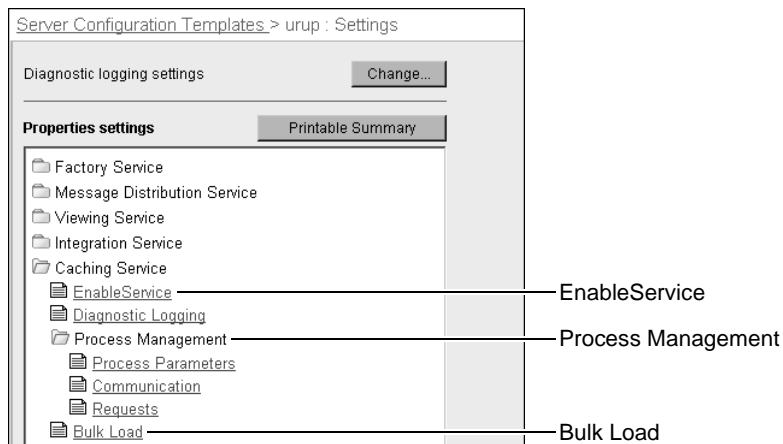


Figure 10-11 Configuring properties for the Caching service

The properties that appeared in Servers—Properties—Caching Service in previous iServer releases appear in Server Configuration Templates—Settings in iHub. Set Caching service properties in Server Configuration Templates—Settings as follows:

- To set Caching service port, choose Caching Service—Process Management—Communication.
- To set Caching service request settings, choose Caching Service—Process Management—Requests.
- To set bulk load settings, choose Caching Service—Bulk Load.

Enabling the Caching service

The administrator can enable or disable the Caching service from Server Configuration Templates—Settings, or by setting the EnableCachingService

parameter in AC_DATA_HOME/config/iHub2/acserverconfig.xml to true or false.

How to enable the Caching service

- 1 Expand Caching Service and choose EnableService, as shown in Figure 10-11.
- 2 For Enable Caching service, accept the default value, which is selected, as shown in Figure 10-12.




Figure 10-12 Enabling or disabling the Caching service

About diagnostic logging

The administrator can configure diagnostic logging by expanding Caching Service, and choosing Diagnostic Logging. For more information, see Chapter 3, “Using diagnostic, usage, and error logging.”

Configuring process management properties

The administrator configures process management properties for the Caching service by setting property values in the following process management categories:

- Process parameters
- Communication
- Requests

Configuring process parameters

In Process Parameters, the administrator sets the start-up parameters for Caching service processes property.

How to configure process parameters

- 1 Expand Caching Service, Process Management, as shown in Figure 10-11, and choose Process Parameters.
- 2 In Start parameters for Caching service processes, accept the default value, or alternatively, modify the value, as shown in Figure 10-13.




Figure 10-13 Setting the start parameters for the Caching service
Choose OK.

3 Restart iHub.

Configuring communication

In Communication, the administrator sets the Port for Information Object Caching server messages property.

How to configure communication

- 1 Expand Caching Service, and Process Management, as shown in Figure 10-11, and choose Communication.
- 2 In Port for Information Object Caching server messages, accept the default value, 11550, or alternatively, specify a different value, as shown in Figure 10-14.

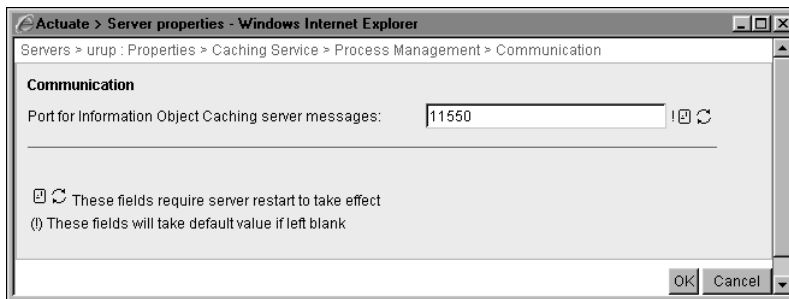


Figure 10-14 Configuring the Communication property
Choose OK.

3 Restart iHub.

Configuring requests

In Requests, the administrator sets the following properties:

- Max SOAP requests
- Max SOAP request threads

How to configure requests

- 1 Expand Caching Service, and Process Management, as shown in Figure 10-11, and choose Requests.
- 2 In Max SOAP requests, accept the default value, 20, or alternatively, specify a different value, as shown in Figure 10-15.

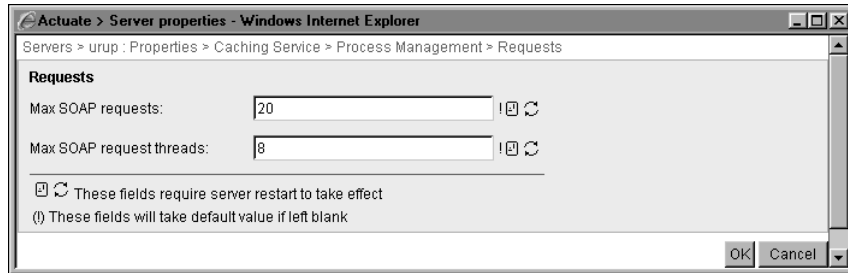


Figure 10-15 Configuring the Requests property

- 3 In Max SOAP request threads, accept the default value, 8, or alternatively, specify a different value.

Choose OK.

- 4 Restart iHub.

Configuring bulk load

In Bulk Load, the administrator sets the following properties:

- Keep temporary bulk load files
- Client bulk load path

How to configure bulk load

- 1 Expand Caching Service, and choose Bulk Load, as shown in Figure 10-11.
- 2 In Keep temporary bulk load files, accept the default value, or alternatively, specify a different value, as shown in Figure 10-16.

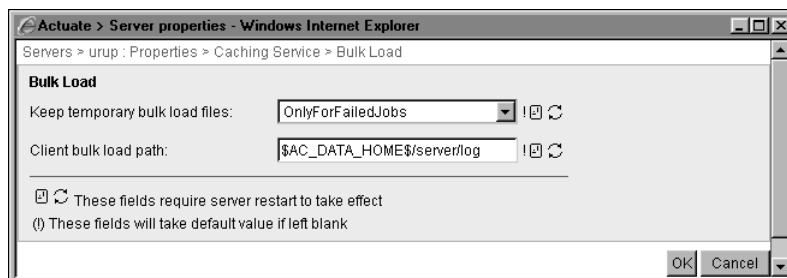


Figure 10-16 Configuring the Bulk Load property

- 3 In Client bulk load path, accept the default path, or alternatively, specify a different path.

Choose OK.

- 4 Restart iHub.

Table 10-1 lists the Caching service properties, and relates the property names that appear in Configuration Console with the corresponding parameter names in `acmetadescription.xml`, indicating default settings, ranges, and when a property change takes effect.

Table 10-1 Actuate Caching service parameters

Property name	Parameter name	Default	Range	Takes effect
Client bulk load path	BulkLoadPath			Server Restart
Enable caching service	EnableCaching Service	False		Fixed
Keep temporary bulk load files	KeepTempFiles	OnlyForFailed-Jobs	OnlyForFailed-Jobs Always Never	Server Restart
Max SOAP requests	MaxConnections	20	2 - 1024	Server Restart
Max SOAP request threads	MaxThreads	8	2 - 200	Server Restart
Start parameters for Caching service processes	StartArguments	-Xmx128M com.nimble.ccs .Server		Server Restart
Port for Information Object Caching server messages	SOAPPort	11550	1024 - 65535	Server Restart

About bulk loading files to the cache

Use database tools, such as `sqlldr` and `db2cmd`, for Oracle and DB2 respectively, for bulk loading. Install these tools on the node where the caching server is enabled. iHub also supports bulk loading through the JDBC driver using `jdbc insert`. To prevent problems with client tools, such as the SQL Server Bulk Copy Program (BCP), specify a fully qualified name for a database object, using `database.schema.object` notation, as shown in the following example of BCP command-line syntax:

```

bcp { [[database_name.] [schema].] {table_name | view_name} |
    "query"}
    {in | out | queryout | format} data_file
    [-S server_name[\instance_name]] [-U login_id] [-P password]

```

Configuring Actuate Caching service and NetOSI file type

The Encyclopedia volume information object cache definition, ICD, stores information object data in a cache. iHub uses properties that you configure on Server Configuration Templates—Settings—NetOSI File Types—ICD to communicate between the information object and Actuate Caching service (ACS) using SOAP.

Choose Server Configuration Templates—Settings—NetOSI File Types (Add) to add a new NetOSI interface.

Figure 10-17 shows Server Configuration Templates—Settings—NetOSI File Types (Add).

Actuate > Server properties - Windows Internet Explorer

Name:

Factory-side Parameters for Third-Party Service

Version of third-party service: ! Ⓜ Ⓜ

Command line arguments:

Use attachment: ☐

Third-Party Service Parameters

SOAP port number: !

Application context: Ⓜ Ⓜ

Target hostname: Ⓜ Ⓜ

SOAP message timeout: Seconds !

* These fields are required and cannot be left blank
 Ⓜ Ⓜ These fields require server restart to take effect
 (!) These fields will take default value if left blank

OK Cancel

Figure 10-17 Adding a NetOSI interface

Table 10-2 describes the properties for adding a NetOSI interface.

Table 10-2 Properties for adding a NetOSI interface

Property	Value
Name	Name of the NetOSI file type. iHub displays the name in the list of file types.
Version of third-party service	Actuate recommends contacting Actuate Customer Support before you change this value.
Command line arguments	Command line arguments for the open server driver. iHub uses these arguments when the open server SOAP interface and a Factory process need to start the open server service.
Use attachment	Indicator of the type of file transfers between Factory process and open server service. iHub uses this setting in conjunction with the open server SOAP interface. Select this property to send open server files as attachments instead of specifying a link to the path. If the Factory process and open server service are on different machines, linking to the files is impossible. You must use attachment mode.
SOAP port number	Port number to connect with the open server service that is running on an application server. The default value, -1, uses the default port.
Application context	String used as the application context when sending a SOAP message to the open server service deployed on the application server.
Target hostname	IP address or name of the host machine that hosts the open server service.
SOAP message timeout	Seconds to wait for SOAP messages between the Factory process and open server service. The default value, -1, disables time-out. If the value is larger than the document request time-out, the Factory process aborts the request. If the time-out value is too small, the Factory process breaks the connection before the open server service is able to respond.

You can delete a NetOSI file type in Server Configuration Templates—Settings—NetOSI File Types—ICD (Delete).

Configuring iHub security

This chapter contains the following topics:

- Understanding the Report Server Security Extension
- Installing iHub using open security
- Understanding LDAP configuration
- Working with RSSE page-level security

Understanding the Report Server Security Extension

The Report Server Security Extension (RSSE) supports open security and page-level security. Open security is the framework that a developer uses to create an interface to an external security source, such as Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory. Using the interface, the Encyclopedia volume controls access using information from the external security source.

Using page-level security, a developer can create an RSSE application that associates security IDs in the access control list (ACL) of a design to Encyclopedia volume users or roles. To use page-level security when working with BIRT designs and documents, obtain a license for the BIRT Page Level Security option. iHub Integration Technology contains reference implementations of RSSE applications, which include source code and JAR files.

iHub supports the following types of SOAP-based RSSE applications that you can install with iHub:

- **External authentication**
Authenticates users in the Encyclopedia volume based on an external, third-party security system, such as LDAP.
- **External registration**
Control access to Encyclopedia volume items based on an external, third-party security system, such as LDAP. With this strategy, you externalize users, roles, groups, and user properties.
- **Page-level security**
Controls user access to sensitive information in a document by implementing page-level security.

To set up an iHub that uses RSSE to connect to an LDAP Directory Server, the administrator must first install an iHub specifying the custom security source option. Then, using Configuration Console, the administrator configures an Encyclopedia volume to use a web service that supports RSSE processing.

Working with RSSE

After installing iHub using the custom security source option, iHub is ready to use RSSE implementation. To prepare an LDAP security source to interoperate with iHub, configure the security source by populating it with the appropriate security information.

In Configuration Console, configure the Encyclopedia volume for RSSE by specifying the web service parameters in Volume—Properties—Open Security. The SOAP-based RSSE application runs as a web service from the iHub

application container. Developers can create a custom RSSE application that uses other data sources. For more information on creating custom RSSE applications, see the Server Integration Technology RSSE reference implementations.

About iHub and RSSE application interaction

Figure 11-1 illustrates the communication between iHub, the RSSE application as a web service, and the external security source.

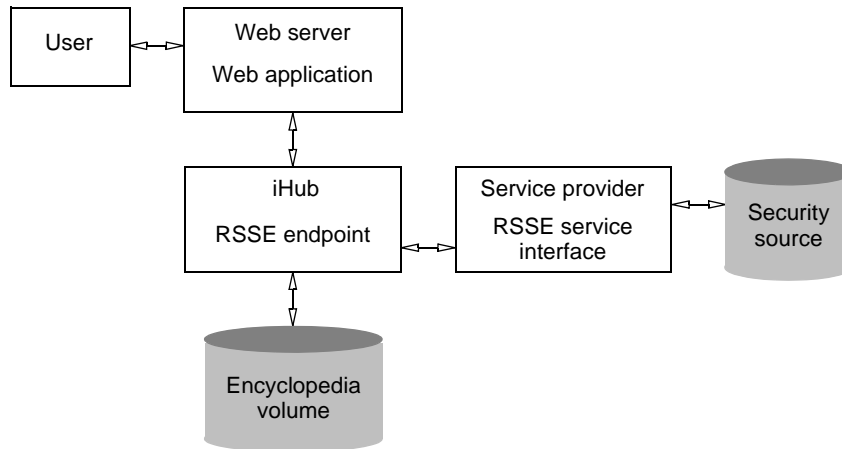


Figure 11-1 iHub communication flow diagram

The service runs from the application container that ships with iHub. The installation program installs the application container with the following components:

- Management and Configuration Consoles
- RSSE web services application

The application container, iHub Consoles, and the RSSE web services application are installed with iHub under the AC_SERVER_HOME directory in the following subdirectories:

- The application container is in the servletcontainer directory.
- iHub Consoles are in the servletcontainer/mgmtconsole directory.
- The RSSE web services application is in the servletcontainer/webapps/acsse directory.

iHub implements the RSSE application as a Java web service using Apache Axis.

Installing iHub using open security

After installing and configuring an LDAP directory server, install a new iHub with an Encyclopedia volume that uses LDAP and configure the SOAP-based Java RSSE web service.

Installing iHub with the LDAP option

The following example installs an iHub that sets up an Encyclopedia volume that uses an LDAP directory server and a SOAP-based Java RSSE web service in a Windows environment.

How to install iHub using the RSSE option

If you downloaded iHub, run ActuateBIRTiHub.exe. If you have an ftp distribution, run setup.exe. The welcome message appears, as shown in Figure 11-2. Choose Next.



Figure 11-2 Viewing the welcome message

- 1 Follow the setup instructions in the Installing BIRT iHub for Windows or Installing BIRT iHub for Linux manual, except choose a custom setup type, select all features, and select the stand-alone installation option. When you reach Select Security Source, select Use an LDAP Directory Server, as shown in Figure 11-3. Choose Next.

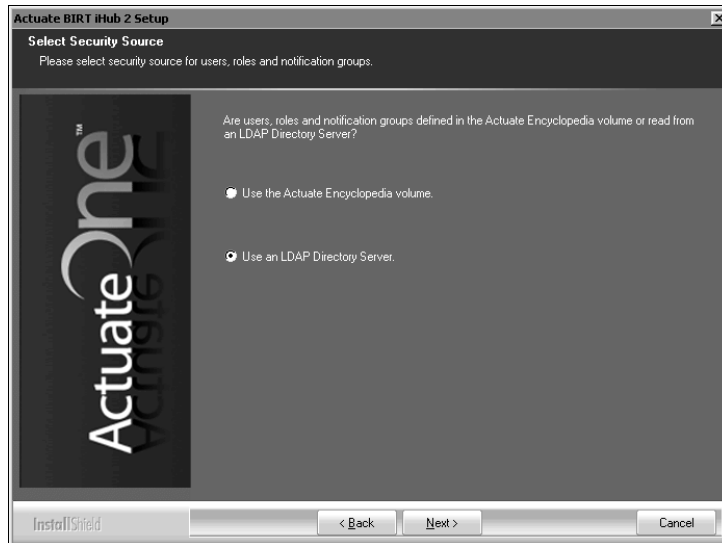


Figure 11-3 Selecting a security source

- 2 In LDAP Configuration, specify the server name where the LDAP directory server is running and the port number where it listens, as shown in Figure 11-4. Choose Next.

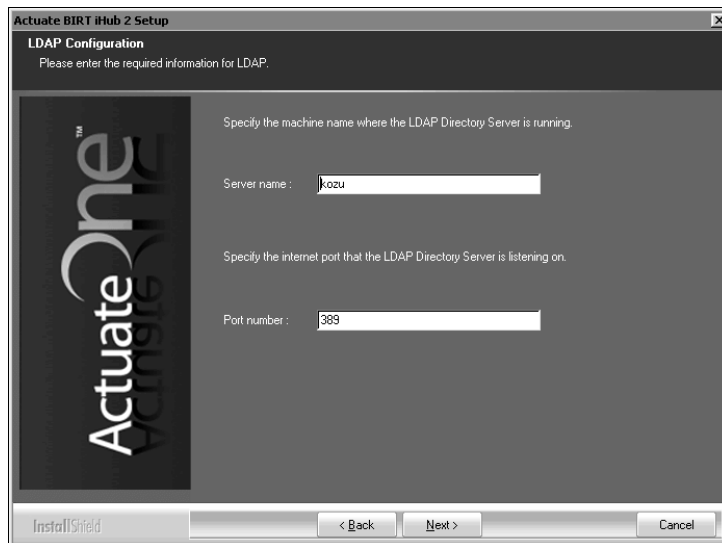


Figure 11-4 Specifying the LDAP directory server name and port

- 3 Specify the query account and password, as shown in Figure 11-5. The query account is for anonymous operations to the LDAP directory server, such as validation. Choose Next.

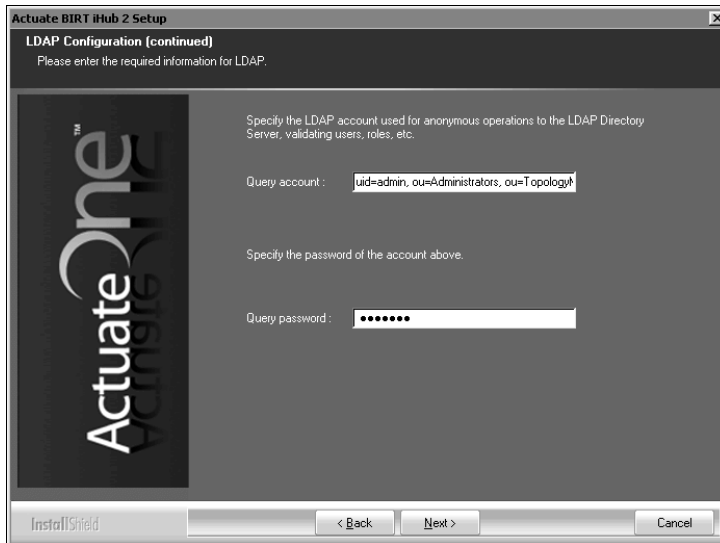


Figure 11-5 Specifying the query account and password

- 4 Specify the name of the organization, as shown in Figure 11-6. Choose Next.



Figure 11-6 Specifying the organization name

- 5 Specify the base domain used for queries of users and the object used as a filter for queries of users, as shown in Figure 11-7. Choose Next.

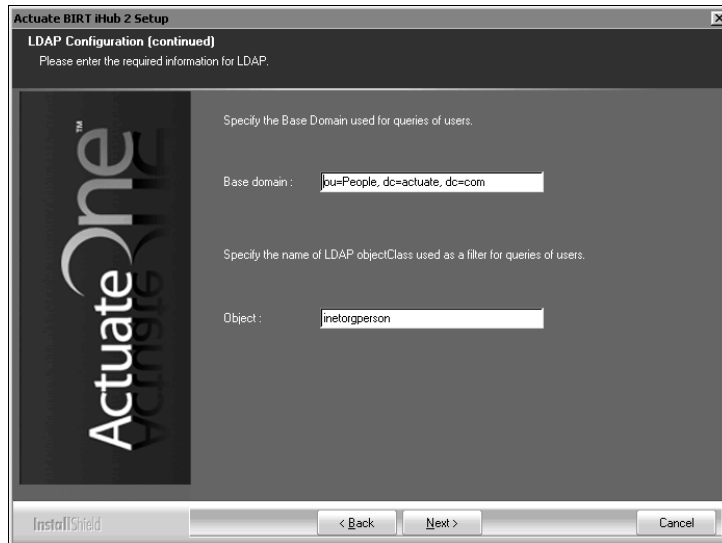


Figure 11-7 Specifying the base domain and object for queries of users

- 6 Specify the base domain used for queries of roles and specify the object used as a filter for queries of roles, as shown in Figure 11-8. Choose Next.

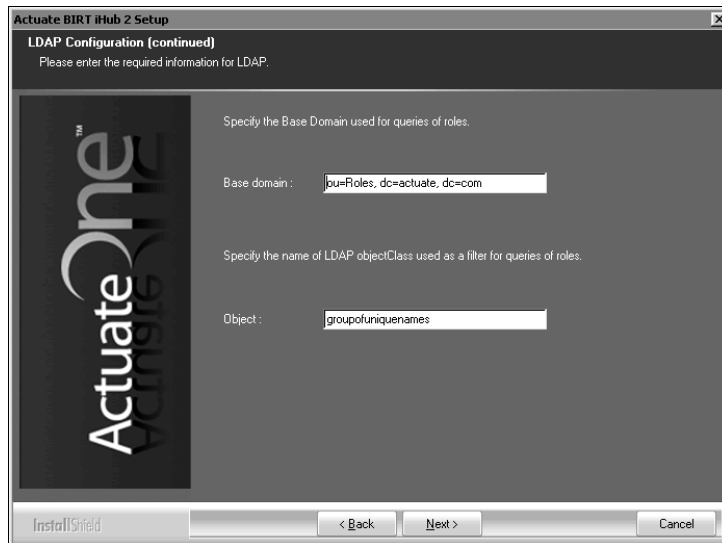


Figure 11-8 Specifying the base domain and object for queries of roles

- 7 Specify the base domain used for queries of groups and specify the object used as a filter for queries of groups, as shown in Figure 11-9. Choose Next.

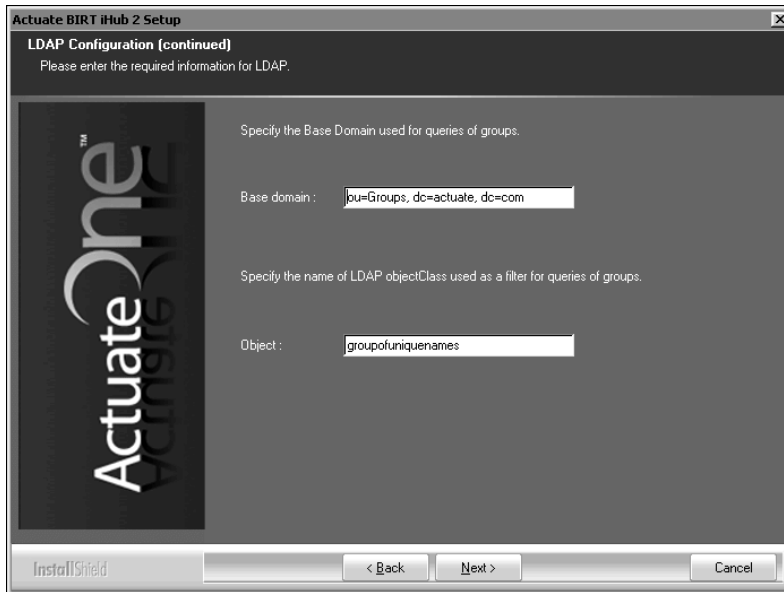


Figure 11-9 Specifying the base domain and object for queries of groups

- 8 Specify the name of a user with an administrator role and the name of a user with an operator role, as shown in Figure 11-10. Choose Next.

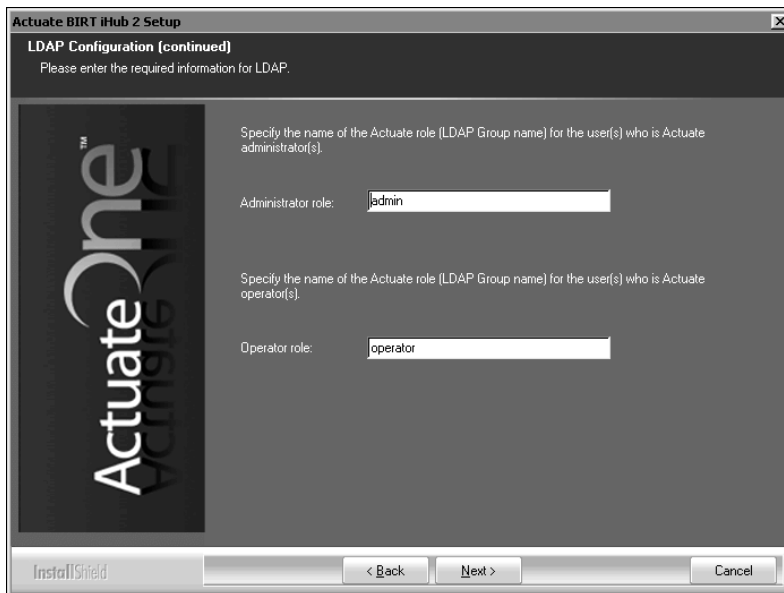


Figure 11-10 Specifying the names of users with administrator and operator roles

After the installation, inspect the `ldapconfig_<VolumeName>.xml` configuration file, which the installation program creates in `AC_SERVER_HOME/etc`.

Installing the RSSE web service application

For an Encyclopedia volume, the RSSE web service configuration information is in the `ldapconfig_<volume>.xml` file, where `<volume>` is the name of the Encyclopedia volume using the RSSE web service application. For example, the configuration file for the `sales1` volume is:

```
ldapconfig_urup.xml
```

The `ldapconfig_<volume>.xml` configuration file is in the following directory:

```
AC_SERVER_HOME/etc
```

Each volume that uses the RSSE web service application must have a configuration file.

The iHub application container that runs the RSSE web service application is in the following directory:

```
AC_SERVER_HOME/servletcontainer/webapps/acrsse
```

To configure the RSSE web service application, change the values in the `ldapconfig_<volume>.xml` file, and restart the iHub application container that runs the RSSE web service application. For more information about setting `ldapconfig_<volume>.xml` parameters, see “Setting `ldapconfig_<volume>.xml` parameters,” later in this chapter.

Using the RSSE application with a service provider

The RSSE web service application runs using a service provider that supports Apache Axis, such as Apache Tomcat. Use the RSSE web service application with an Apache Tomcat servlet container by performing the following steps:

- Install the RSSE web service application on the Tomcat servlet container.
- Configure the Encyclopedia volume to use the RSSE web service application.

The following sections briefly describe how to use the RSSE web service application with a Tomcat servlet container.

Installing an RSSE application on Tomcat

To install the web service RSSE application on a Apache Tomcat servlet container, copy the files from the iHub application container directory to the Tomcat directory.

On a Windows system, the default installation places the RSSE web services application in the following directory:

```
AC_SERVER_HOME\servletcontainer\webapps\acrsse
```

Copy the acrsse directory to the Tomcat webapps directory. For release 7.0 of Tomcat, the directory is:

```
\Program Files\Apache Tomcat 7.0\webapps
```

Stop and restart Tomcat to update the configuration. Check the HTTP server log files to ensure proper start-up of the application.

Configuring the Encyclopedia volume to run RSSE

Configure each Encyclopedia volume to run an RSSE web service application separately. A SOAP-based Java RSSE application runs as a web service in the iHub servlet container. To run SOAP-based RSSE applications on multiple Encyclopedia volumes, configure a separate location for each RSSE application.

Install an RSSE application on an Encyclopedia volume to run in its own location on iHub by performing the following tasks:

- Make a copy of the AC_SERVER_HOME/servletcontainer/webapps/acrsse directory.

For example, copy the directory to the following location:

```
AC_SERVER_HOME/servletcontainer/webapps/acrsse_AUTH
```

- Copy the file, rsseAcl.jar, to the lib directory of the servlet container in the following location:

```
AC_SERVER_HOME/servletcontainer/webapps/acrsse_AUTH/WEB-INF/lib
```

- Extract the file, class.properties, from the archive file, rsseACL.jar, to the following location:

```
AC_SERVER_HOME/servletcontainer/webapps/acrsse_AUTH/WEB-INF  
/classes/com/iHub/rsse/wsdl
```

If necessary, create the subdirectories, /classes/com/iHub/rsse/wsdl, manually or using an archive extraction tool create the subdirectories when extracting the class.properties file.

- Use a source code editor to open the class.properties file and change the code reference specification to:

```
class=com.iHub.rsse.authSample.SampleRSSE
```

Configuring open security

Use Configuration Console to enable a RSSE web service application for an Encyclopedia volume. The following parameters appear on Volumes—Properties—Open Security, as shown in Figure 11-11:

- Do not enable
Disables open security and the RSSE application.

- **Enable as a web service**
Enables open security and the RSSE application.
- **Cache timeout**
Sets the maximum time, in minutes, before iHub deletes cached open security data. The minimum cache time-out period is 1 minute. Set to -1 to specify that the cache never expires.
- **IP address**
Sets IP address or machine name of the server that runs the RSSE web service.
- **Soap port**
Establishes the port number for the RSSE web service.
- **Context string**
Specifies the location of the RSSE web service for iHub to use when sending messages to the web service. The path for the default Encyclopedia volume is /acrsse/servlet/AxisServlet.

Volumes > corp : Properties

General Open Security Partitions Events Advanced

Enable/Disable ⓘ ⌛

☐ Do not enable

☒ Enable as web service

Cache

Cache timeout: 60 min ⓘ ⌛

Web service

IP address: localhost ⓘ ⌛

Soap port: 8900 ⓘ ⌛

Context String: /acrsse/servlet/AxisServlet ⓘ ⌛

* These fields are required and cannot be left blank

ⓘ ⌛ These fields require volume restart to take effect

(!) These fields will take default value if left blank

Figure 11-11 Specifying open security property values

How to enable open security for an Encyclopedia volume

- 1 Create a configuration file for the RSSE web service application.
The configuration file maps the Encyclopedia volume management information to LDAP security directory information.
- 2 From the advanced view of Configuration Console, choose Volumes.
- 3 In Volumes, point to the arrow next to the Encyclopedia volume name and choose Properties.
- 4 In Volumes—Properties—General, choose Open Security, as shown in Figure 11-11.

- 5 In Volumes—Properties—Open Security, choose Enable as a web service.
- 6 Specify web service parameter values as necessary. Choose OK.
- 7 On Volumes, point to the arrow next to the Encyclopedia volume name and choose Put offline.
- 8 Log out of Configuration Console. Stop and start the iHub application container using the StartMC and StopMC scripts in AC_SERVER_HOME/bin. Log back in to Configuration Console. Take the Encyclopedia volume online.

Test whether the RSSE web service application is working by logging in to the Encyclopedia volume or, if using a page-level security, by deploying a design file to the Encyclopedia volume.

If you are using page-level security and change the assignments in the users.acs file, be sure to wait for the volume cache time-out period or recycle the volume before checking to see if the changes are effective. Using Configuration Console, you can re-configure the volume cache time-out period by going to Servers—Properties—View Service. On View Service, in View process cache for executables, the default value for Max cache time-out is 3600 seconds.

In View service, in Extended viewing cache, you can also set Cache level to Cache only page-level security requests to retain previously viewed documents that have page-level security enabled. In Extended viewing cache, the default value for Cache time-out is 1200 seconds.

Specifying RSSE service start-up parameters

Volumes—Properties—Advanced—Security Extension—Service, shown in Figure 11-12, allows the administrator to supply an operating system command that iHub uses to start the RSSE service. The administrator can also specify license options that iHub assigns to users by default when upgrading an RSSE-enabled volume from an iHub version that has no user-based licensing.

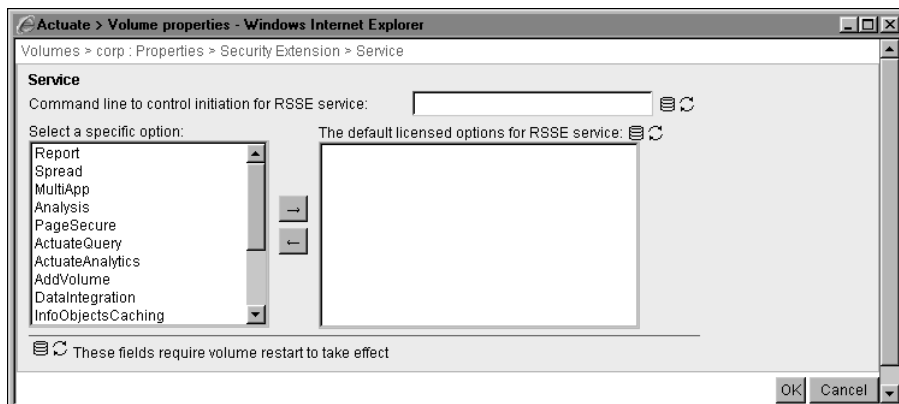


Figure 11-12 Specifying RSSE start-up command and default license options

Table 11-1 describes the parameters that appear on Volumes—Properties—Advanced—Security Extension—Service.

Table 11-1 RSSE service start-up parameters

Parameter	Description
Command line to control initiation for RSSE service	Command that iHub issues to start the web service for RSSE, if the RSSE service is not running. If iHub cannot connect to the web service for the RSSE service, iHub starts an operating system shell and runs the value of this parameter as a command-line request.
The default licensed options for RSSE service	Supports specifying license options that iHub assigns to all users by default when upgrading an RSSE-enabled volume from an iHub version that has no user-based licensing. The administrator is free to change license option assignments after the upgrade. Assigning default license options in The default licensed options for RSSE service does not affect the list of available options from which to select when assigning options to a user in Management Console.

Understanding LDAP configuration

When iHub uses the RSSE application, Encyclopedia volume users that are defined in the LDAP server must have a password. The user information must use:

- A single LDAP organization
- The appropriate LDAP object classes

The `ldapconfig_<volume>.xml` file defines the mapping between Encyclopedia volume user information and LDAP information. Encyclopedia volume user management information typically maps to LDAP information in the following way:

- The Encyclopedia volume users, groups, and roles map to LDAP object classes. For example, you can specify that Actuate users map to the LDAP `inetorgperson` object class and that roles and groups map to the LDAP `groupofuniquenames` object class.
- Individual users, roles, and groups are LDAP objects. For example, after specifying how Actuate users map to the LDAP `inetorgperson` object class of an organizational unit, you create LDAP users based on the `inetorgperson` object class. In LDAP, you can create users named MBarajas and JThompson based on the `inetorgperson` object class. MBarajas and JThompson are Encyclopedia volume users.

- Actuate user properties, such as e-mail address and home folder, are specified in LDAP attribute-value pairs of the LDAP object class. For example, after specifying that the Actuate users are based on the inetorgperson object class, you specify that Actuate e-mail maps to the mail attribute of the LDAP inetorgperson object class. You create object class attributes to store Actuate user properties.

In LDAP, you can also create a set of Encyclopedia volume users, notification groups, and roles by creating the LDAP objects within different LDAP organizational units and specifying LDAP distinguished names to point to the different LDAP organizational units.

Mapping Encyclopedia volume management information to LDAP objects

The following topics discuss how to map particular types of Encyclopedia volume management information:

- Mapping Encyclopedia volume user properties
- Mapping roles
- Mapping groups
- Mapping channels
- Mapping pass-through security information

Mapping Encyclopedia volume user properties

Encyclopedia volume users and their properties map to attributes of an LDAP object. User properties include login name, password, e-mail address, and default privilege template. For example, if Actuate user information maps to an LDAP inetorgperson object class, the user properties map to the LDAP inetorgperson object class attributes.

All LDAP user object attributes are string attributes. To specify multiple values, the LDAP user object attributes for a privilege template list and an Encyclopedia volume channel membership list must be multivalued. For example, to specify that a user belongs to multiple volume channels in LDAP, the user object Encyclopedia volume channel attribute must be multivalued with an Actuate channel name for each value.

The Encyclopedia volume login name is the user ID attribute of the LDAP user object, called the uid attribute. The volume password is the password attribute of the LDAP user object. When using an RSSE application, a user must log in to an Encyclopedia volume using a password.

To indicate that an Actuate user is a member of an Actuate role or an Actuate notification group, add the LDAP Actuate user as a member of the appropriate LDAP Actuate group or role.

Actuate uses the default value for an Actuate user property in `ldapconfig_<volume>.xml` when:

- The LDAP server does not contain a definition for the LDAP user object attribute.
- The LDAP attribute for a user object does not contain any values.

Mapping roles

Encyclopedia volume roles map to an LDAP object class, such as `groupofuniquenames`. The name Actuate displays is the LDAP object's common name attribute, called the LDAP `cn` attribute.

When using the RSSE application with an LDAP server, you cannot nest roles. Roles are an LDAP object. To indicate that an Actuate user is a member of one or more Actuate roles, add the LDAP Actuate user object as a member of the LDAP Actuate roles.

iHub uses the members specified in the LDAP Actuate role objects when it performs authorization functions for Actuate roles. iHub also uses the LDAP role objects when it lists roles used to specify privileges.

To specify privileges to access a file or folder in the Encyclopedia volume for an LDAP role object, first create the Actuate role object in the LDAP directory. Then, in the Encyclopedia volume, specify privileges for the Actuate role in the Encyclopedia volume file or folder.

Information Console supports using different levels of user functionality based on a user's membership in particular Encyclopedia volume roles. To use the Information Console roles and functionality levels, create corresponding LDAP roles and add users to the roles.

Mapping groups

Encyclopedia volume groups map to an LDAP object class, such as `groupofuniquenames`. The name Actuate displays in the Encyclopedia volume is the LDAP object `cn` attribute.

If the e-mail notification group maps to the LDAP `groupofuniquenames` object class, the LDAP group objects do not require Actuate-specific attributes. To specify that an Actuate user is a member of an Encyclopedia volume group, first create an LDAP group object. Then, add the LDAP user object as a member of the LDAP group object.

Mapping channels

The Encyclopedia volume stores channel names. An LDAP user object attribute specifies the list of channels to which an Actuate user subscribes. The RSSE application does not verify that the Encyclopedia volume's channel names match

the LDAP user object channel attribute values. You must ensure that the channel names match the values in the LDAP user object channel attribute.

Mapping pass-through security information

When a user runs an information object that uses pass-through security, iHub requires a database user name and password. When using the RSSE application, you specify two LDAP user object attributes that iHub uses with pass-through security. One attribute contains the value for the database user name and the other attribute contains the database password. The following example shows pass-through security parameters in the LDAP configuration file:

```
<ConnectionPropertyList>
  <ConnectionProperty>
    <Name>username</Name>
    <Value>acdbname</Value>
  </ConnectionProperty>
  <ConnectionProperty>
    <Name>password</Name>
    <Value>acdbpassword</Value>
  </ConnectionProperty>
</ConnectionPropertyList>
```

The ConnectionPropertyList contains ConnectionProperty elements. The ConnectionProperty name-value pairs specify the LDAP user object attributes that contain the database user name and password that iHub uses when a design accesses data through an information object.

The value for username is the LDAP user attribute that contains the database user name. In the following example, iHub uses the value in the LDAP user's dbname attribute as the database user name:

```
<ConnectionProperty>
  <Name>username</Name>
  <Value>dbname</Value>
</ConnectionProperty>
```

The value for password is the LDAP user attribute that contains the database password.

In the following example, iHub uses the value in the LDAP user's dbpassword attribute as the database password:

```
<ConnectionProperty>
  <Name>password</Name>
  <Value>dbpassword</Value>
</ConnectionProperty>
```

Setting ldapconfig_<volume>.xml parameters

The RSSE application uses a mapping file, ldapconfig_<volume>.xml, to map Encyclopedia volume management information to LDAP objects and object attributes.

In the Actuate ldapconfig_<volume>.xml file, a parameter is an XML element. Specify the value for a parameter as shown in the following example:

```
<parameter-name>value 1, value 2</parameter-name>
```

where

- The parameter name is one of the valid parameter names specified in ldapconfig_<volume>.xml.
- A comma separates multiple parameter values.

ldapconfig_<volume>.xml can contain comments. Enclose comments in <- - and - -> tags, as shown in the following example:

```
<
--This is the port number on which the LDAP server is listening.--
>
```

Table 11-2 contains example values for parameters that appear in ldapconfig_<volume>.xml.

Table 11-2 ldapconfig_<volume>.xml parameters

Parameter	Description and example
AdminRole	Actuate role attribute value that indicates that an LDAP user object can perform Encyclopedia volume management. <AdminRole> actuateAdmin </AdminRole>
AllRole	LDAP role object name that maps to the All role in the Encyclopedia volume. Use the All role to grant privileges to all Encyclopedia volume users. <AllRole> actuateAll </AllRole>
GroupBase DN	Base LDAP distinguished name used to locate the LDAP Actuate notification group object in queries of notification group names. <GroupBaseDN> ou=Groups, dc=actuate, dc=com </GroupBaseDN>

(continues)

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
Group Object	LDAP object class that the RSSE application uses to find Actuate notification group names. <pre><GroupObject> groupofuniquenames </GroupObject></pre>
GroupTo Notify	Name of the LDAP notification group that receives notification about all iHub requests in the manner of the administrator user when the Encyclopedia volume uses default, internal security. The GroupBaseDN parameter defines the base DN of this group name. <pre><GroupToNotify> specialGroup </GroupToNotify></pre> <p>When combined with the GroupBaseDN value, this parameter specifies the LDAP Actuate notification group object. iHub uses that object for LDAP notification. For example: "cn=AdminGroup, ou=Actuate Groups, o=actuate.com"</p>
Operator Role	LDAP role object name that maps to the Encyclopedia volume Operator role. A user must have this role name to perform Encyclopedia volume Operator functions, such as configuring autoarchive operations. <pre><OperatorRole> actuateOperator </OperatorRole></pre>
Port	Internet port on which the LDAP server listens. The default value is 389. <pre><Port> 389 </Port></pre>
Query Account	LDAP account that the RSSE application uses for query operations to the LDAP server. <pre><QueryAccount> uid=actuate, ou=Administrators, ou=TopologyManagement, o=NetscapeRoot </QueryAccount></pre> <p>The RSSE application uses this account to validate users, roles, ACLs, and other Encyclopedia volume user information. For example: "uid=admin, ou=Administrators, ou=TopologyManagement, o=NetscapeRoot"</p>

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
Query Password	Password for the LDAP account specified by the QueryAccount parameter. <pre><QueryPassword> Actu8 </QueryPassword></pre>
RoleBaseDN	Base LDAP distinguished name that the RSSE application uses to locate the LDAP role object in queries of roles. <pre><RoleBaseDN> ou=AcRoles, dc=actuate, dc=com </RoleBaseDN></pre>
RoleObject	LDAP object class that the RSSE application uses to find Actuate role names. <pre><RoleObject> groupofuniquenames </RoleObject></pre>
Server	Name of the LDAP server that the RSSE application and iHub use. Use the fully qualified name, including the domain name. You can use the server's IP address. The default value is the name of the machine. <pre><Server> helium.actuate.com </Server></pre>
UserBaseDN	LDAP distinguished name that the RSSE application uses to locate the LDAP user object. When you add a user's name as a prefix to a base-distinguished name, the resulting name uniquely identifies the user in the external data source. Most base-distinguished names consist of the organizational unit or a series of organizational units and an organization. <pre><UserBaseDN> ou=People, dc=actuate, dc=com </UserBaseDN></pre>
Channel Subscription ListAttr	LDAP attribute that specifies the channels to which an Actuate user subscribes. In the LDAP directory server, the attribute has multiple values with a single channel name for each value. <pre><ChannelSubscriptionListAttr> actuateChannelList </ChannelSubscriptionListAttr></pre>

(continues)

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
Channel Subscription ListDefault	<p>Value to use for ChannelSubscriptionListAttr when LDAP does not contain a value for that attribute.</p> <pre><ChannelSubscriptionListDefault> portfolio update, sales forecasts </ChannelSubscriptionListDefault></pre> <p>The value is a comma-separated list of channel names. For example: "portfolio update, sales forecasts"</p>
Privilege Template Attr	<p>LDAP attribute that specifies which privilege template to use for files and folders that an Encyclopedia volume user creates.</p> <pre><PrivilegeTemplateAttr> actuateDefaultPriv </PrivilegeTemplateAttr></pre>
Privilege Template Default	<p>Value to use for PrivilegeTemplateAttr when LDAP does not contain a value for that attribute.</p> <p>The value is a comma-separated list of user or role privileges. This example gives read and visible privileges to a role called viewing only and gives read, write, execute, and delete privileges to a user named jbob.</p> <pre><PrivilegeTemplateDefault> viewing only~rv, jbob=rwed </PrivilegeTemplateDefault></pre> <p>A user permission is a user login name followed by "=" and a zero (0) or more permission characters. A role permission is a role name followed by tilde (~) followed by a zero or more permission characters. The following table is a list of the privilege characters and their meanings:</p> <ul style="list-style-type: none"> r = read w = write e = execute d = delete v = visible s = secure read g = grant <p>To specify a privilege template that lists multiple users or roles in the LDAP directory server, the attribute must be multi-valued with a single user or role for each value.</p>

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
Attach Report InEmailAttr	<p>LDAP attribute that specifies an Actuate user's preferred form of e-mail notification.</p> <pre><AttachReportInEmailAttr> actuateEmailForm </AttachReportInEmailAttr></pre> <p>The e-mail can contain either a copy of the document or a link to the document.</p>
Attach Report InEmail Default	<p>Value to use for AttachReportInEmailAttr when LDAP does not contain a value for that attribute. The value is either included or linked. If the value is included, the user receives the document as an attachment to the notice, if possible. If the value is linked, the user receives a link to the document. The default value in ldapconfig_<volume>.xml is linked.</p> <pre><AttachReportInEmailDefault> linked </AttachReportInEmailDefault></pre>
Email Address Attr	<p>Name of the LDAP user attribute that specifies an Encyclopedia volume user's e-mail address that iHub uses to send e-mail. For some object classes, such as inetorgperson, an e-mail attribute exists in the standard LDAP schema.</p> <pre><EmailAddressAttr> mail </EmailAddressAttr></pre>
SendEmail Attr	<p>LDAP user attribute that specifies when to send an e-mail notification message to notify an Actuate user of the completion of a job.</p> <pre><SendEmailAttr> actuateEmailWhen </SendEmailAttr></pre>
SendEmail Default	<p>Value to use for SendEmailAttr when the LDAP directory server does not contain a value for that attribute.</p> <pre><SendEmailDefault> never </SendEmailDefault></pre> <p>Use one of the following values: never, always, failures, or successes.</p> <p>never—Do not modify.</p> <p>always—Notify of failures and successes.</p> <p>failures—Notify of failures only.</p> <p>successes—Notify of successes only.</p> <p>The default value in ldapconfig_<volume>.xml is never.</p>

(continues)

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
Failure Notice Expiration Attr	LDAP attribute that specifies how long iHub keeps a user's notices about failed jobs in the completed notice folder of the Encyclopedia volume. The value is a number of minutes. A value of 0 (zero) means that iHub does not keep notices about failed jobs. A value of -1 means that iHub keeps the notices indefinitely.
Failure Notice Expiration Default	Value to use for FailureNoticeExpirationAttr when LDAP does not contain a value for that attribute. The value is a number of minutes. The default value in ldapconfig_<volume>.xml is 0.
SendNotice Attr	LDAP user attribute that specifies when to notify a user about the completion of a job by placing a notice in the completed notice folder of the Encyclopedia volume.
SendNotice Default	Value to use for SendNoticeAttr when LDAP does not contain a value for that attribute. Use one of the following values: never, always, successes, or failures. never—Do not modify. always—Notify of failures and successes. failures—Notify of failures only. successes—Notify of successes only. The default value in ldapconfig_<volume>.xml is always.
Home Folder Attr	LDAP attribute that specifies a user's home folder in the Encyclopedia volume. There is no default value. <pre><HomeFolderAttr> actuateHomeFolder </HomeFolderAttr></pre>
MaxJob PriorityAttr	LDAP attribute that specifies a user's maximum request priority. The value is the maximum request priority that the user can set for a document print or generation request in the Encyclopedia volume. In LDAP, the value must be an integer between 0 and 1000. <pre><MaxJobPriorityAttr> actuateMaxPriority </MaxJobPriorityAttr></pre>

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
MaxJob Priority Default	<p>Value to use for MaxJobPriorityAttr when LDAP does not contain a value for that attribute.</p> <p>The value must be an integer between 0 and 1000.</p> <p>The default value in ldapconfig_<volume>.xml is 500.</p> <pre><MaxJobPriorityDefault> 500 </MaxJobPriorityDefault></pre>
UserObject	<p>Name of the LDAP object class that the RSSE application uses to find Actuate user names.</p> <p>An example of an LDAP object class is inetorgperson.</p> <pre><UserObject> inetorgperson </UserObject></pre>
Success Notice Expiration Attr	<p>LDAP attribute that specifies how long to keep a user's success completion notices in the completed notice folder of the Encyclopedia volume.</p> <p>The value is a number of minutes. A value of 0 (zero) discards notices about successful jobs. A value of -1 keeps success notices indefinitely.</p> <pre><SuccessNoticeExpirationAttr> actuateSuccessNoticeExpiration </SuccessNoticeExpirationAttr></pre>
Success Notice Expiration Default	<p>Value to use for SuccessNoticeExpirationAttr when LDAP does not contain a value for that attribute.</p> <p>The value is a number of minutes. The default value in ldapconfig_<volume>.xml is 0.</p> <pre><SuccessNoticeExpirationDefault> 0 </SuccessNoticeExpirationDefault></pre>
View Preference Attr	<p>LDAP attribute that specifies the user's default viewing preference.</p> <p>Use one of the following values: default or dhtml.</p> <pre><ViewPreferenceAttr> actuateViewingPref </ViewPreferenceAttr></pre>

(continues)

Table 11-2 ldapconfig_<volume>.xml parameters (continued)

Parameter	Description and example
View Preference Default	<p>Value to use for ViewPreferenceAttr when LDAP does not contain a value for that attribute.</p> <p>Specify the default viewing mode using one of the following values: default or dhtml.</p> <p>The default value in ldapconfig_<volume>.xml is default.</p> <pre><ViewPreferenceDefault> default </ViewPreferenceDefault></pre>
Connection PropertyList	<p>Values to use for information object pass-through security. When using pass-through security, iHub requires a database user name and password. The ConnectionPropertyList element contains two ConnectionProperty elements. Each ConnectionProperty element contains a Name and Value element.</p> <p>The values for the ConnectionProperty Name elements are username and password.</p> <p>The value for username is the LDAP user attribute that contains the database user name.</p> <p>ConnectionProperty Value for password is the LDAP user attribute that contains the database password.</p> <pre><ConnectionPropertyList> <ConnectionProperty> <Name>username</Name> <Value>dbname</Value> </ConnectionProperty> <ConnectionProperty> <Name>password</Name> <Value>dbpassword</Value> </ConnectionProperty> </ConnectionPropertyList></pre>

Understanding an LDAP directory structure

Figure 11-13 shows a simple LDAP directory structure that stores information about the following iHub objects:

- Encyclopedia volume users
- Security roles
- Notification groups

In the LDAP directory, Actuate is the LDAP domain that the RSSE application uses. The domain Actuate contains LDAP objects used by the RSSE application.

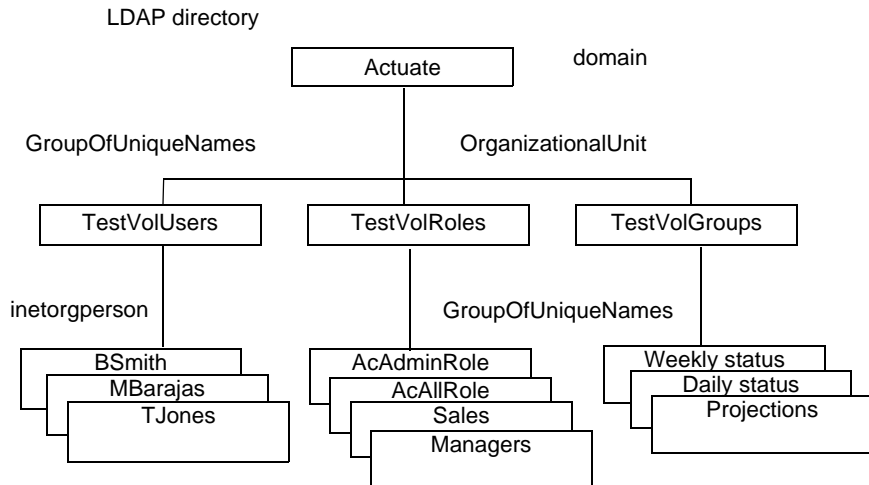


Figure 11-13 LDAP directory structure

TestVolUsers is an instance of groupofuniquenames. TestVolUsers contains a list of the Encyclopedia volume users. The users are instances of the inetorgperson LDAP object.

TestVolRoles and TestVolGroups are instances of the OrganizationalUnit LDAP object. Within TestVolRoles are instances of the GroupOfUniqueNames LDAP object with names AcAdminRole, AcAllRole, Sales, and Managers. Each GroupOfUniqueNames object contains references to Actuate users that are instances of the inetorgperson LDAP object.

Within TestVolGroups are instances of the GroupOfUniqueNames LDAP object with names Weekly status, Daily status, and Managers. Each GroupOfUniqueNames object contains references to Actuate users that are instances of the inetorgperson LDAP object.

Mapping LDAP objects to users

Use the following parameters in the LDAP RSSE application to map the LDAP objects to Encyclopedia volume users:

- UserObject parameter maps the LDAP object to the Encyclopedia volume user.

```
<UserObject>inetorgperson</UserObject>
```

- UserBaseDN parameter maps the instance of the LDAP object to the Encyclopedia volume user.

```
<UserBaseDN>cn=TestVolUsers, dc=actuate, dc=com</UserBaseDN>
```

In the configuration file, use the following properties to map the LDAP objects to Encyclopedia volume security roles:

- **RoleObject** parameter maps the LDAP object to the Encyclopedia volume security role.

```
<RoleObject>groupofuniquenames</RoleObject>
```

- **RoleBaseDN** parameter maps the instance of the LDAP object to the Encyclopedia volume security role.

```
<RoleBaseDN>ou=TestVolRoles, dc=actuate, dc=com</RoleBaseDN>
```

In the configuration file, use the following properties to map the LDAP objects to Encyclopedia volume notification groups:

- **GroupObject** parameter maps the LDAP object to the Encyclopedia volume notification group.

```
<GroupObject>groupofuniquenames</GroupObject>
```

- **GroupBaseDN** parameter maps the instance of the LDAP object to the Encyclopedia volume notification group.

```
<GroupBaseDN>ou=TestVolGroups, dc=actuate, dc=com  
</GroupBaseDN>
```

In the configuration file, use the following properties to map the LDAP objects to the special Encyclopedia volume roles:

- **AdminRole** parameter maps the instance of the LDAP object to the Encyclopedia volume administrator security role.

```
<AdminRole>AcAdminRole</AdminRole>
```

- **AllRole** parameter maps the instance of the LDAP object to the Encyclopedia volume administrator security role.

```
<AllRole>AcAllRole</AllRole>
```

Converting an Encyclopedia volume to use an RSSE application

When you configure iHub to use the default Encyclopedia volume security, the Encyclopedia volume stores all security information. iHub uses an identifier ID and does not use the name of the user, role, or group when it assigns privileges and sets other administrative options. To use the RSSE application and an external security source, replace the internal IDs with the user, role, or group name. When the administrator enables open security as a web service and restarts the volume, iHub replaces references to Actuate user, role, and group IDs with their corresponding names in the following Encyclopedia volume management information:

- Privilege rules, or access control lists (ACLs) for files and folders in the volume
- Privilege rules in scheduled jobs
- Privilege rules for a volume's channels
- Job completion notification settings
- Scheduled jobs
- Completed jobs

The RSSE application matches the users, roles, and groups in the Encyclopedia volume to users, roles, and groups in the external security source. When the administrator disables the open security as a web service and restarts the volume, iHub modifies these references. For each user name reference in the Encyclopedia volume, iHub looks up the corresponding Encyclopedia volume ID and changes the reference to an ID. If there is no corresponding ID, iHub removes the reference.

Converting internal IDs to external names

To convert existing Encyclopedia volume management information to a form that the RSSE application can use, complete the following tasks in this order:

- 1 Open Properties for the selected Volume.
- 2 In Open Security, select "Enable as web service". Choose OK.
- 3 Take the volume online.

When the volume is taken online, iHub converts the Encyclopedia volume references from internal users, roles, and groups to the corresponding external names.

Converting information from external to internal

To convert an Encyclopedia volume from using external information to using internal information, complete the following tasks in this order:

- 1 Open Properties for the selected Volume.
- 2 In Open Security, select "Do not enable". Choose OK.
- 3 Take the selected volume offline.
- 4 Take the selected volume online.

When the volume is taken online, iHub converts the Encyclopedia volume references from external users, roles, and groups to the corresponding internal IDs.

Caching external security information

When using an external security application, iHub caches external security source information, including:

- External user properties
- Roles of a user under external user registration
- Groups and group memberships

Control the maximum time that the cache holds the information by setting open security parameters. Use the RSSE Cache Timeout parameter to control how long the cache stores security information. The default value is 60 minutes. The RSSE Cache Timeout setting can be set to -1 to specify that the RSSE cache never expires.

Working with RSSE page-level security

The following topics describe how to use the example RSSE page-level security application:

- About the RSSE page-level security example
- Installing the RSSE web service application
- Creating an access control list

About the RSSE page-level security example

Using the RSSE framework, a developer can create an RSSE service that manages page-level security by retrieving a user's access control list (ACL) externally. By default, when a secure design asks for the ACL of a user, the Encyclopedia volume returns a list that includes the user ID and the roles to which the user belongs. Frequently, the information in iHub security does not match the information in a database used by a secure design. An RSSE page security application can translate an iHub ACL to a design-specific ACL.

iHub Integration Technology contains an example of how external page-level security works using RSSE. This example is located under the Java Report Server Security Extension directory in the subdirectory, `Page_Security_Example`.

To use page-level security, license the BIRT Page Level Security option. The example RSSE page-level security application uses a text file, `users.acls`, that maps Encyclopedia volume user names to other external security IDs. Use the sample application as a basis for a custom RSSE application.

Page-level security protects a report from unauthorized access by comparing the user ID and user membership in Encyclopedia volume roles to an access control

list (ACL) for the report document. If the user's name appears in the ACL, the user can view the protected pages.

The RSSE application uses a file, `users.acls`, to translate an Encyclopedia volume user ID or role to one or more security IDs, which iHub uses to check against the Actuate report page ACL when a report uses page-level security.

For example, if a report shows information about the sales reps in the following city offices:

- NYC
- Boston
- Philadelphia

and the file, `user.acls`, contains the following access control list specifications:

```
user1=NYC
user2=Boston
user3=Philadelphia
```

then user1 has access to the pages with information about NYC office, user2 to the Boston office, and user3 to the Philadelphia office.

In this example, when user1 tries to read a report with page-level security enabled, the RSSE application returns a list of security IDs that contain user1 and user1's roles. iHub checks the user ID, roles, and RSSE list against the Actuate report page ACL. iHub lets user1 view any report page where a security ID that the RSSE application returns matches a report page security ID.

Creating an access control list

The example file, `user.acls`, stores a user's access control list (ACL) using the following format:

```
Username=ac11, ac12, ..
```

The username field matches the name of user in the Encyclopedia volume. An equal (=) sign separates the user name from the ACL list. The ACL list can contain 0, 1, or more ACL specifications as shown in the following code example:

```
user1=ac11, ac12, ac13, ac14
user2=ac15, ac16,\
    ac17, ac18
user3=ac19
user4=ac110
```

If there is more than one ACL in the list, separate each ACL using a comma. The scanner reading the file, `users.acls`, eliminates any white space or backslash. All the username specifications in the example are legal. A list can contain users that do not appear in the Encyclopedia. The information for these users is ignored.

The default maximum length of an ACL is 64 kilobytes (KB). If you use an ACL longer than 64 KB, specify a longer maximum length using the RSSE Max ACL Buffer Size parameter on Volumes—Properties—Open Security in Configuration Console.

12

Archiving files

This chapter contains the following topics:

- Understanding online archiving
- Using the online archive driver

Understanding online archiving

iHub ships with a configurable, Java-based Encyclopedia volume archive driver that you can use to archive files from an Encyclopedia volume. Developers can create custom online archive drivers using the Actuate Information Delivery API (IDAPI). The source code and build files for the online archive driver reside in the online archive driver directory after installing IDAPI.

Using the online archive driver or your customized version requires purchasing the Online Archive option.

About the online archive driver

The iHub online archive driver copies expired Encyclopedia volume files to a second Encyclopedia volume that serves as a file archive.

The online archive driver supports the following features:

- Preserving file attributes, such as description, version name, and security information
- Preserving file dependencies
- Deleting empty folders from the source Encyclopedia volume

Configure iHub to use the online archive driver from Configuration Console. Configure and perform archiving from Management Console.

About the online archive driver files

iHub installs the following online archive driver files:

- A script runs the driver and sets environment variables. On Linux, the script is `aconlinearchive.sh`. On Windows, the script is `aconlinearchive.bat`. The script resides in `AC_SERVER_HOME/bin`.
- The Java JAR files `aconlinearchive.jar` and `aconlinearchiveDEP.jar` in `AC_SERVER_HOME/drivers`.
- An XML configuration file, `onlinearchive.cfg`, is a reference implementation. The default location for the configuration file is the `AC_SERVER_HOME/etc` directory.

Retaining file attributes during archiving

When configuring the online archive driver, specify whether the driver retains the following general file and security information:

- Time stamp

- Owner
- Permissions
- Dependencies

If the online archive driver does not retain the permissions information for an archived file, the owner of the file has full access. If the driver retains security information, including owner and privilege information about an archived file, the driver performs the following functions:

- The access control list (ACL) of the file refers to a user or role, and the driver attempts to match the reference to a name in the archive. When using an existing Encyclopedia volume user or role, the driver does not update the properties of the user or role with information from the original volume.
- If the archive uses RSSE external registration for user and role information, the online archive driver uses the RSSE information to update the user and role privileges to access the file.
 - If the original owner of the file is not defined in the archive, the user named in the configuration file of the driver becomes the owner of the file when it is added to the archive.
 - If a user or role is in the original ACL, but not in the archive, they do not get privileges to access the file when it is added to the archive.
- If the user or role is not present in the archive, the driver creates them, assigning the properties of the original user or role to them, except the driver does not enable login for a user or role it creates.

Preserving file dependencies during archiving

When the online archive driver archives a file that has a dependency on another file, the driver can archive both files and preserve the dependency information. When the driver retains file dependency, and an archived file depends on another file, the driver copies both files to the archive and creates the same relative file structure. The driver does not delete the file on which the archived file depends from the original Encyclopedia volume, unless the driver archives both files.

If more than one archived file depends on the same file, the driver retains the original file dependency and typically retains only one copy of the file on which the archived files depend. The driver archives two copies of the file on which the other files depend, one for each file that is archived in each session.

Using the online archive driver

To use the online archive driver, first create an Encyclopedia volume for storing archived files. Next, specify the configuration file that contains the configuration

parameters for the driver. Use the default configuration file, shown in Listing 12-1, or modify the file. Finally, specify the startup script file that sets the environment variables and launches the online archive driver.

Specifying the online archive configuration file

The online archive driver uses an XML configuration file. The default name of the configuration file is `onlinearchive.cfg`. The default location for the configuration file is the `AC_SERVER_HOME\etc` directory.

In Linux, you can specify the configuration file using the `-c` command-line option. For example, specify the following command for the Encyclopedia volume parameter, Use archive service, on System Volumes—Properties—General:

```
aconlinearchive.sh -c /local/iHub/archiveconfig.xml
```

In a production system, you can configure an Encyclopedia on a different volume or on a separate node in a cluster to function as the repository for the archive. You copy the reference implementation of `onlinearchive.cfg` in `AC_SERVER_HOME/etc`, and rename it using the following format:

```
onlinearchive_<volume>.cfg
```

where `<volume>` is the name of the Encyclopedia volume that runs the archive driver.

Modifying the online archive configuration file

Listing 12-1 shows the default online archive configuration file. You can modify the parameters described in Table 12-1.

Listing 12-1 Default online archive configuration file

```
<?xml version="1.0" encoding="UTF-8"?>
<archiveconfig>
  <
    --TargetServer, TargetSOAPPort: [Required]          -->
  <
    -- Name or IP of server and port for connecting to  -->
  <
    -- the SOAP dispatcher service of the target volume -->
    <TargetServer>localhost</TargetServer>
    <TargetSOAPPort>8000</TargetSOAPPort>

  <
    --ArchiveVolume: [Required]                          -->
  <
    -- Name of target volume to copy archived files to  -->
    <ArchiveVolume>DefaultVolume</ArchiveVolume>
```

```

    <
--AdminUser, AdminPassword: [Required]                                -->
    <
--    Name and password of a user in the target volume                -->
    <
--        that belongs to the Administrator role                      -->
    <AdminUser>administrator</AdminUser>
    <AdminPassword></AdminPassword>

    <
--RetainTimestamp: [Optional, default: false]                        -->
    <
--        Whether timestamp of archived file is preserved            -->
    <RetainTimestamp>>false</RetainTimestamp>

    <
--RetainOwner: [Optional, default: false]                            -->
    <
--        Whether Owner of archived file is preserved                -->
    <RetainOwner>>false</RetainOwner>

    <
--RetainPermission: [Optional, default: false]                      -->
    <
--        Whether Permission (ACL) of archived file is              -->
    <
--        preserved                                                    -->
    <RetainPermission>>false</RetainPermission>

    <
--CopyDependOnFile: [Optional, default: true]                        -->
    <
--        Whether files depended on by archived file are            -->
    <
--        copied                                                        -->
    <CopyDependOnFile>true</CopyDependOnFile>

    <
--CreateUserRole: [Optional, default: true]                          -->
    <
--        Whether to create missing user or roles in target          -->
    <
--        volume in order to retain Owner or Permissions             -->
    <CreateUserRole>true</CreateUserRole>

```

```

    <
--ArchiveRoot: [Optional, default: /]                                -->
    <
--    Root encyc folder for all archived files                        -->
    <ArchiveRoot></ArchiveRoot>

    <
--CreateArchiveSubFolder: [Optional, default: true]                -->
    <
--    Whether to create a timestamp dependent subfolder -->
    <
--    under ArchiveRoot for each archive session                    -->
    <CreateArchiveSubFolder>true</CreateArchiveSubFolder>

    <
-- LogLevel: [Optional, default: Summary]                          -->
    <
--    Level of detail in log file. Valid values are: -->
    <
--    Summary, Detail and Trace -->
    <LogLevel>Summary</LogLevel>

</archiveconfig>

```

Table 12-1 describes the parameters in the online archive configuration file for the Encyclopedia volume that stores the archive.

Table 12-1 Online archive configuration file parameters

Parameter	Description
TargetServer	Machine name or IP address to use to connect to the Encyclopedia volume that holds the archived files. Required.
TargetSOAPPort	Port number that the iHub SOAP Dispatch service uses to connect to the Encyclopedia volume. Required.
ArchiveVolume	Encyclopedia volume name. The default value is DefaultVolume. Required.
AdminUser	Encyclopedia volume user name. The user must belong to the Administrator role.
AdminPassword	Password of the user specified by the AdminUser parameter. Required.
RetainTimestamp	Flag for archiving the file's time stamp. The default value is false.

Table 12-1 Online archive configuration file parameters

Parameter	Description
RetainOwner	Flag for archiving the name of the file's owner. The default value is false.
RetainPermission	Flag for archiving the file's permissions. The default value is false.
CopyDependOnFile	Flag for archiving the file's dependency list. The default value is true.
CreateUserRole	Flag for creating missing user or roles in the archive Encyclopedia volume to retain the file's owner or permission information. The default value is true.
ArchiveRoot	Encyclopedia volume archive session root folder. The default value is /, the Encyclopedia volume root folder.
CreateArchiveSubFolder	Flag for creating a time stamp dependent subfolder under ArchiveRoot for each archive session. The default value is true.
LogLevel	The level of detail for information in the log. Valid values are Summary, Detail, and Trace. The default value is Summary.

The online archive driver applies changes to the configuration file when iHub runs the driver. Changes do not affect any archiving already in process.

Creating and specifying the startup script

Create a script to start your online archive driver, or use the implementation script and driver as shown in Listing 12-2.

Listing 12-2 Sample online archive driver startup script

```
@echo off

REM Actuate Online Archive Driver Startup Script
set ARCHIVE_DRIVER_JRE=%AC_JRE_HOME%
.
.
.
set DRIVER_JAR_PATH=%AC_SERVER_HOME%\drivers
set JAVA_EXE=%ARCHIVE_DRIVER_JRE%\bin\java.exe
if exist "%JAVA_EXE%" goto JavaOK
```

```

echo *** Java home not set correctly; trying Java.Exe in PATH ***
set JAVA_EXE=java.exe
:JavaOK
set DRIVER_JAR=%DRIVER_JAR_PATH%\aconlinearchive.jar
set DRIVER_LIB_JAR=%DRIVER_JAR_PATH%\aconlinearchiveDEP.jar
set DRIVER_CP=%DRIVER_JAR%;%DRIVER_LIB_JAR%
start "Online Archive Service" "%JAVA_EXE%" -cp "%DRIVER_CP%"
com.actuate.onlinearchivedriver.Main %1 %2 %3 %4 %5 %6 %7 %8

```

To specify the script file, in the advanced view of Configuration Console, choose Volumes. Choose Properties from the drop-down menu of a volume. For Volume archive service provider, in Use archive service on Properties—General, type the path and name of the script, as shown in Figure 12-1.

Volumes > corp : Properties

General Open Security Partitions Events Advanced

Description:

Schedule for purging notices: HH:mm

Schedule for purging deleted files: HH:mm

Partition

Primary partition: DefaultPartition Min Free Space: MB

Volume archive service provider

Use archive service:

Metadata database and schema

Metadata database name: Default_ActuatePostgreSQL_MetadataDatabase

Database schema name: ac_corp

Email notification

E-mail notification template partition:

Use Information Console for e-mail notifications ☐

Information Console URL prefix:

These fields require volume restart to take effect
 (!) These fields will take default value if left blank

OK Cancel Apply

Figure 12-1 Volumes—Properties—General

The default Volume archive service uses the configuration batch file aconlinearchive.bat, shown in Figure 12-1, that is located in the AC_SERVER_HOME/bin directory.

Managing file purging

In Volumes—Properties—Advanced—Archiving and Purging, shown in Figure 12-2, you can specify expiration times for Encyclopedia volume items.

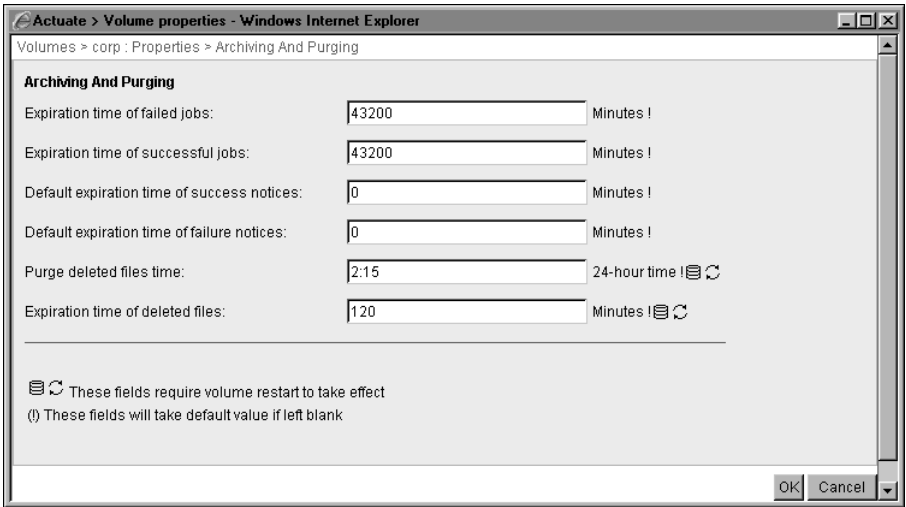


Figure 12-2 Specifying Encyclopedia volume item expiration times

Table 12-2 describes the parameters that appear on System Volumes—Properties—Advanced—Archiving and Purging. Purging notices reduces the size of the Encyclopedia volume and contributes to better performance.

Table 12-2 Archiving and purging parameters

Parameter	Description
Expiration time of failed jobs	Length of time before iHub can delete a job completion notice for a failed job from Jobs-Completed. iHub purges the notice during the scheduled completion notice purge time if no completion notice for the failed job exists on a volume channel, and the age of the notice exceeds the value of this parameter. The lower the value of this parameter, the better iHub performs. If an administrator specifies 0 for this parameter, iHub uses the default value for the volume.

(continues)

Table 12-2 Archiving and purging parameters (continued)

Parameter	Description
Expiration time of successful jobs	Length of time before iHub can delete a job completion notice for a successful job from Jobs-Completed. iHub purges the notice during the scheduled completion notice purge time if no completion notice for the successful job exists on a volume channel, and the age of the notice exceeds the value of this parameter. The lower the value of this parameter, the better iHub performs. If an administrator specifies 0 for this parameter, iHub uses the default value for the volume.
Default expiration time of success notices	The default value for the volume specifying the length of time before iHub can delete a job completion notice for a successful job from a user's personal channel. This is the same property as Purge success notices after n days n hours, in Management Console, Volume—Properties—Archiving and Purging. Setting this property in one console sets the property in the other console. iHub does not delete the notice if the value for this parameter is 0, the default value.
Default expiration time of failure notices	The default value for the volume specifying the length of time before iHub can delete a job completion notice for a failed job from a user's personal channel. This is the same property as Purge failure notices after n days n hours, in Management Console, Volume—Properties—Archiving and Purging. Setting this property in one console sets the property in the other console. iHub does not delete the notice if the value for this parameter is 0, the default value.

Locating the archived files

The online archive driver creates a folder in the archive Encyclopedia volume and places files from an online archive session in this folder. In the driver configuration file, you specify the name of a root folder that contains the directories for all the online archive sessions. The directory structure within the folder matches that of the original Encyclopedia volume. The subfolder name containing the files from the online archive session consists of the start date and time of the online archive session in the following format:

```
YYYY_mm_dd.hh_mm_ss
```

For example, if you specify the online archive content folder as /archive2009, for an autoarchive session starting on 8:25 pm June 2, 2009, the driver copies the file /documents/sales/commission.rptdocument to the archive Encyclopedia volume in the following location:

```
/archive2009/2009_06_02.20_25_14/documents/sales  
/commission.rptdocument
```

Printing documents

This chapter contains the following topics:

- Understanding printing
- Sending a document to the printer
- Managing the printing environment
- Configuring fonts

Understanding printing

This chapter describes the requirements and operations necessary to set up document printing in the Windows and Linux environments.

Printing on Windows

The installation process on Windows sets up printers already configured on iHub as system printers. To connect additional system printers after installation, first connect them to client computers using Printers and Faxes in the Control Panel. Next, in Configuration Console—Advanced view, you add printers to the server using Printers—Add Printer, shown in Figure 13-1.

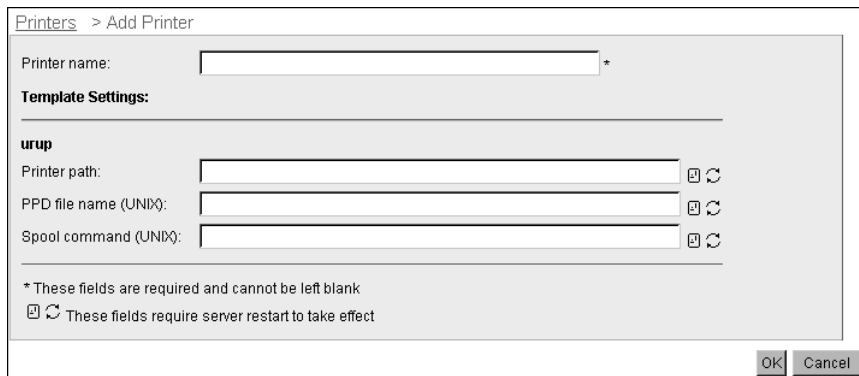


Figure 13-1 Adding a printer

To add an iHub printer, connect the printer to client computers using Printers and Faxes in the Control Panel. Next, add the printer as an iHub printer. Lastly, restart iHub.

How to add a printer as an iHub printer

- 1 From the Advanced view of Configuration Console, choose Printers. Printers lists the printers available to iHub, as shown in Figure 13-2. In Printers, choose Add Printer.

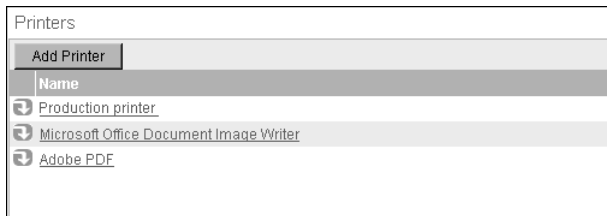


Figure 13-2 List of printers

In Printers—Add Printer, you can change the following server settings for each machine:

- **PPD File Name**
PPD file name for Linux operating systems. Do not specify a path.
- **Printer name**
Name of the printer.
- **Printer path**
Path to the printer from the machine.
- **Spool command**
Spool command for Linux. Actuate recommends using the copy before printing option. For example, on an AIX system, use the following command:

```
lp -c -d
```

On Linux, use the following command:

```
lpr -P
```

Choose OK.

2 Restart iHub.

In an iHub cluster, each machine maintains printer property information. When adding a printer to a cluster, you must specify the printer parameter values for each node in the cluster.

Printing on Linux

Actuate sends output to any Windows printer and to PostScript Level 1 or 2 printers. If the PostScript font is not installed on the system, the PostScript file, generated at print time, contains only the references to the missing font name.

You must install required PostScript fonts on either the printer or iHub; otherwise, the printer substitutes a font or printing fails. If you install the PostScript font on iHub, font information embeds in the PostScript file, and the document prints with the correct font.

iHub uses Adobe Type-1 font technology and a font configuration file for printing. The configuration file is in the following directory:

```
AC_SERVER_HOME/operation/print/fonts/fonts.supported
```

Throughout this documentation, AC_SERVER_HOME refers to the folder that the installer chose as the location for binary files during the iHub installation. By default, this location is AC_SERVER_HOME/AcServer in a Linux environment. To print documents, check that XVFBDISPLAY variables, start_svr.sh and display_value, are set in pmd11.sh.

Adding a printer on Linux

For iHub to recognize a printer in Solaris, Linux, HP, and IBM environments, the printer name must appear in the following locations:

- Linux—a subdirectory of the /etc/lp/printers directory.
- HP—the /etc/lp/member directory.
- IBM—the /etc/qconfig file.

How to add a printer on Linux

You add a printer to the server in Configuration Console—Advanced view, using Printers—Add Printer, shown in Figure 13-3. Specify the following information:

- The printer path.
- Name of the PostScript printer definition (PPD) file. Do not include the path.
- The name of the spool command.

Restart iHub after adding printers.

Printers > Add Printer

Printer name: *

Template Settings:

renfield

Printer path:

PPD file name (UNIX:):

Spool command (UNIX:):

urup

Printer path:

PPD file name (UNIX:):

Spool command (UNIX:):

* These fields are required and cannot be left blank

* These fields require server restart to take effect

OK Cancel

Figure 13-3 Adding a printer to iHub in Linux

About Xprinter environment variables

The XPHOME and XPPATH environment variables for Xprinter on Linux specify the location of the iHub Xprinter home and the path to Xprinter files, respectively.

Obtaining the PPD file

If the manufacturer of your printer does not provide a PPD, download the Windows PostScript drivers and extract the PPD file. Put the PPD file in the

AC_SERVER_HOME/operation/print/ppds. List the file in operation/print/ppds/driver_mapping.

For example, to use a the PPD file hp9000.ps with the iHub printer HP LaserJet 9000, the PPD file must be in AC_SERVER_HOME/operation/print/ppds/ and in the same directory, the driver_mapping file must contain an entry similar to the following one:

```
hp9000.ps: "HP LaserJet 9000"
```

After updating the Linux system, restart the cluster node for the changes to take effect. Specify the name of the PPD in Configuration Console—Advanced view by choosing Printers—Add Printer.

Sending a document to the printer

A user can send a document to the printer from Management Console. You must schedule a job to perform server-side printing. On Files and Folders, select Schedule for the document you want to print. Choose Print, and specify the printer and other options, as shown in Figure 13-4.

The screenshot shows a window titled "corp > Home > administrator > Sales by Territory (RPTDESIGN) (Version 1) :". Inside, there's a "Schedule" tab selected among others like Parameters, Output, Privileges, Datamart Security, Channels, Notification, and Print. A checkbox "Print the output document on the server:" is checked. Below it, the "Printer:" dropdown is set to "Microsoft Office Document Image Writer". Other fields show "Manufacturer: Microsoft", "Model: Microsoft Office Document Image Writer Driver", and an empty "Description:" and "Location:" fields. There's a "Print to file:" checkbox and a text field. Under "Override default settings for", there are checkboxes for "Scale:", "Resolution" (set to "100 X 100"), "Mode:" (with radio buttons for "B&W" and "Color"), "Number of copies" (set to "1") with a "Collate:" checkbox, "2-Sided printing:" (with radio buttons for "1-Sided Print", "Flip on short edge", and "Flip on long edge"), "Page size:", and "Paper tray:" (set to "Default tray"). A "Print format: PostScript" section is expanded, showing "Page range" (empty), "Page style" (set to "Auto"), "BIDI processing" (checked), "Text wrapping" (checked), "Font substitution" (checked), and "Chart DPI" (set to "192"). At the bottom right are "OK" and "Cancel" buttons.

Figure 13-4 Scheduling a job to print a document

Managing the printing environment

The following sections describe how to manage the iHub document printing environment.

Changing a path to a printer

In Printers, the administrator can specify the path to a printer.

How to set the path to a printer

In the advanced view of Configuration Console, set the path to a printer by performing the following steps:

- 1 Select Printers from the advanced view.
- 2 Point to the arrow next to the printer name in the list, as shown in Figure 13-5.

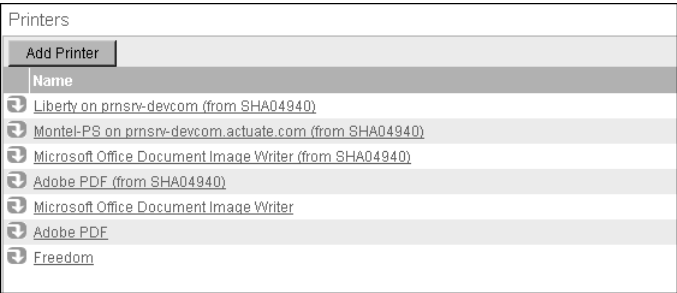


Figure 13-5 Choosing a printer name

- 3 Choose Template settings to change the path to a printer on Windows or to a PPD file name of a printer on Linux.
- 4 On Properties, as shown in Figure 13-6, in Printer path, specify the path to the printer. For a Linux system, specify the PPD file name and the Spool command.

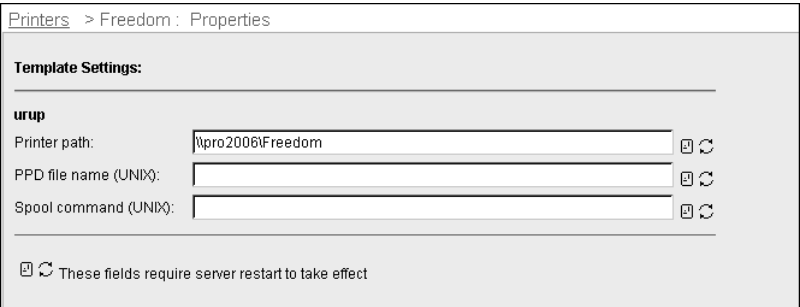


Figure 13-6 Specifying the path to a printer

Removing a printer

In Printers, the administrator can remove a printer.

How to remove a printer

To remove a printer from iHub, complete the following tasks:

- 1 If the printer is the default printer for an Encyclopedia volume, assign a new default printer.
- 2 On Printers, point to the arrow next to the printer name and choose Delete.

Configuring advanced printing parameters

You can configure iHub to limit the list of printers that users see. Configure a subset of printers in your enterprise as iHub printers. In a reasonably sized list, users can find a printer quickly to print a document. Displaying a long list of printers in Management Console can slow the application.

In System—Properties—Advanced—Printing, set Maximum number of printers to automatically configure, as shown in Figure 13-7. You can specify up to 50 printers. At start-up, iHub configures the specified number of printers.

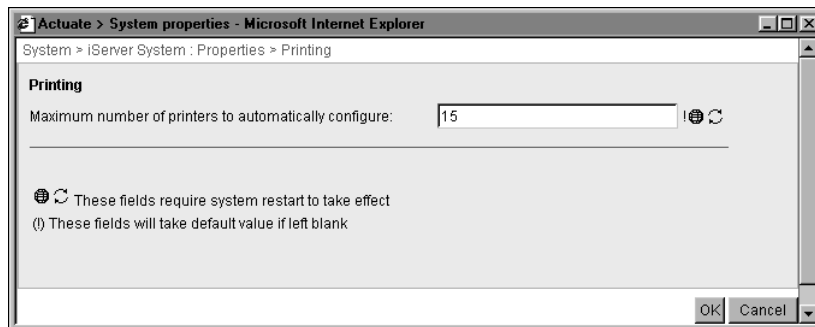


Figure 13-7 Specifying advanced printing property values

Disabling automatic printer configuration

During restart, iHub typically deletes printer information and adds the names of printers known to the operating system running iHub. Setting Maximum number of printers to automatically configure to 0 disables automatic printer configuration. iHub does not update the list of printers during restart. For example, set the value to 0 if iHub uses printers that are not detected by the operating system, such as a printer on a cluster node.

Configuring fonts

iHub supports TrueType and PostScript fonts. iHub ships with the following Latin-1 fonts:

Arial Black	Arial Narrow Italic	Palatino Italic
Arial Black Italic	Arial Narrow Bold Italic	Times
Arial	Courier New	Times Bold
Arial Bold	Courier New Bold	Times Italic
Arial Italic	Courier New Italic	Times Bold Italic
Arial Bold Italic	Courier New Bold Italic	Verdana
Arial Narrow	Frutiger-Roman	Verdana Bold
Arial Narrow Bold	Johnny-Bold	Verdana Italic
		Verdana Bold Italic

On Linux systems, iHub installs Latin-1 fonts, AFM files, and at least one font set of each non-CJK language. Non-CJK languages are those languages other than Chinese, Japanese, and Korean. Each font set consists of the normal, italic, bold, and bold-italic variations of the font. For non-CJK languages, Actuate installs a localized version of a font type similar to the Arial font, including the corresponding AFM and PFA files.

Actuate works with, but does not supply the Chinese, Japanese, and Korean (CJK) fonts listed in Table 13-1.

Table 13-1 Supported CJK language fonts

Language	Font 1	Font 2
Simplified Chinese	SimSun	SimHei
Traditional Chinese	MingLiu	MS Hei
Japanese	Mincho	Gothic
Korean	Batang	Dotum

If you add a CJK font to the font properties file, the font name must be in the native language. The font's AFM file is required. iHub does not embed CJK fonts; therefore, you must print the PostScript file using a printer on which the font is installed.

iHub can embed non-CJK fonts. If a font is embedded in the output, the font's PFA file is required. For example, for a non-CJK font, the Linux font name must be the same as the value for the FontName parameter in the PFA file.

Understanding PostScript font conversion issues

When converting a font using Fontographer 4.1 from Macromedia, choose File→GenerateFontFile. In the dialog box, select Advanced and in the Encoding section, select Adobe Standard. Do not change any other options.

If you used Fontographer 4.1 from Macromedia, check the font’s KernPair value. There is a problem in Fontographer 4.1 in which it uses incorrect KernPair values when converting some TTF fonts. After converting a font with Fontographer and installing the font using fontutils, go to the font metrics directory and use the UNIX grep utility to search for the string KPX in the font’s AFM file.

```
grep -c KPX mynewfont.afm
```

Open the .afm file in a text editor, search for the StartKernPairs variable, and check the value. If the StartKernPairs number is not the same as the number displayed using the grep utility, replace the number with the one displayed from the grep command.

Mapping fonts for charts in documents

A developer can specify fonts of text components in a chart, such as the title, data points, and axis label. To render charts properly, the fonts specified in the chart must be available to the iHub Java virtual machine (JVM).

The procedure for making fonts available differs by platform.

- In Windows, install fonts in the standard fonts folder for the operating system. For example, on Windows, from Control Panel, open Fonts. Choose Fonts and then choose File→Install new font.
- In Linux, copy fonts to the JAVA_HOME/jre/lib/fonts directory.

iHub uses a default font for chart elements that do not specify a font. The default font depends on the locale setting of the iHub that generates the chart. Table 13-2 lists default fonts by language. To achieve consistent appearance in output in Linux and Windows, the default fonts must be available to the JVM.

Table 13-2 Default fonts by language

Language	Default fonts
Languages other than Chinese, Japanese, or Korean (non-CJK languages)	arial.ttf arialbd.ttf arialbi.ttf ariali.ttf times.ttf timesbd.ttf timesbi.ttf timesi.ttf

(continues)

Table 13-2 Default fonts by language (continued)

Language	Default fonts
Simplified Chinese	simsum.ttc
Traditional Chinese	mingliu.ttc
Japanese	msgothic.ttc msmincho.ttc
Korean	gulim.ttc batang.ttc

Locating fonts when generating BIRT documents

To render or generate BIRT design output, iHub looks for font configuration information in the `AC_SERVER_HOME/Jar/BIRT/platform/plugins/org.eclipse.birt.report.engine.fonts` directory.

BIRT designs use five different types of font configuration files. The font configuration file naming convention includes information about the rendering format, the system platform, and the system locale, as shown in the following template:

```
fontsConfig_<Format>_<Platform>_<Locale>.xml
```

BIRT supports the following levels of font configuration files:

- For all rendering formats:

These files have no format specifier in their names. The configuration files are divided into three sub-levels:

 - The default configuration file:

```
fontsConfig.xml
```
 - Configuration files for a specific platform, for example:

```
fontsConfig_Windows_XP.xml
```
 - Configuration files for a specific platform and locale, for example:

```
fontsConfig_Windows_XP_zh.xml
```

```
fontsConfig_Windows_XP_zh_CN.xml
```
- For certain formats only

These files include the format specifier in their names. These configuration files are divided into three sub-levels:

 - The default configuration file for a format, for example:

```
fontsConfig_pdf.xml
```

- Configuration files for a format for a specific platform:
`fontsConfig_pdf_Windows_XP.xml`
- Configuration files for a format for a specific platform and locale:
`fontsConfig_pdf_Windows_XP_en_AU.xml`

iHub first looks for the font location in the font configuration files. If the font is not found, iHub searches for the font in the system-defined font folder. An exception to this search order occurs if iHub encounters a `<fonts-path>` section in one of the font configuration files. iHub searches each path coded in the `<fonts-path>` section instead of searching the system-defined font folder.

14

Connecting to data sources

This chapter contains the following topics:

- About data source connections
- Using a connection configuration file
- Defining environment variables
- Connecting to data sources
- Using a connection pool

About data source connections

Actuate Customer Support publishes the Supported Products and Obsolescence Policy document that describes the data sources, drivers, operating systems, and other software requirements for connecting iHub to data sources. Actuate Supported Products and Obsolescence Policy, available on the Actuate Support site at the following URL, also contains information about the required patches:

<http://support.actuate.com/documentation/spm>

iHub connects to data sources when generating documents and using the Actuate Caching service (ACS). The design developer specifies data source connection information in the design, or in an external connection configuration file. In most cases, iHub and the database run on different computers for load balancing purposes, but this division is not mandatory. Running iHub on the database host can improve performance.

About drivers

The iHub installation process installs and configures DataDirect Connect for ODBC 5.3 drivers and JDBC 3.7 SP1 drivers. You can also use third-party drivers to connect to data sources, but you must license, install, and configure them.

To make a JDBC connection with an Oracle database, iHub uses the DataDirect JDBC driver that installs with iHub only.

About information object connections

You can use information objects in designs. Information objects make ODBC and JDBC connections, so you need to be able to connect underlying data sources to databases in addition to the obvious data sources required by your design.

Using a connection configuration file

A connection configuration file is an XML file, such as the one shown in Listing 14-1, in UTF-8 or ASCII encoding. The file specifies the data source connection properties to use when iHub runs a design. Having the data source connection information for a design in an external file makes it convenient to modify. You change the connection information without altering the design. You specify the location of the file using Configuration Console.

You can use an external connection configuration file to define a data source for the Actuate Caching service and for a data connection definition (.dcd) file, which contains information object connection properties for a data source. You can also

use an external connection configuration file for connecting data sources to designs.

You can create an external connection profile to a data source used by a design. Changes to the profile are automatically picked up by the design. The settings in a connection configuration file override any connection configuration properties in the connection profile. The sample connection configuration file in Listing 14-1 externalizes the file path to the connection profile, C:\SqlServer.profile.

Listing 14-1 BIRT connection configuration file example

```
<oda-data-source
  extensionID="org.eclipse.birt.report.data.oda.jdbc" name="JDBC
  Data Source - SQL Server" id="783">

  <property name="odaDriverClass">com.actuate.jdbc.sqlserver.
    SQLServerDriver
  </property>

  <property name="odaURL">jdbc:actuate:sqlserver://DBSRV1-W2K
  </property>
</oda-data-source>

<ConnectOptions Type=".eclipse.birt.report.data.oda.jdbc_ JDBC
  Data Source - SQL Server ">

  <Property PropName="OdaConnProfileStorePath">C:\Myopath
  </Property>
</ConnectOptions>
```

In a BIRT design, the configuration key used to specify a data source is the unique ID of the ODA data source extension and data source name defined in the BIRT design or library. You must concatenate the string as follows:

extensionID + "_" + data source name

For example, the key is org.eclipse.birt.report.data.oda.jdbc_SQL Server.

Changing a configuration file

The Factory process reads the configuration file for the configuration key values when the process starts. After changing a configuration file, you must restart Factory processes for changes to take effect. Only Factory processes that start after changes in the configuration file use the new information. To ensure that design files use updated configuration file information, confirm that no document generation jobs are running and stop Factory processes that are running before you change the configuration file. After changing the file, iHub starts a Factory process for the next document generation job.

Specifying the location of the connection configuration file

There is no default location for the connection configuration file. To use a connection configuration file, create the file and then specify its name and location using the ConnConfigFile parameter in Configuration Console.

From Server Configuration Templates—Settings, expand iHub, then choose Database Connection Configuration File. Specify the location of the file using the Configuration file for database connections and search path parameter shown in Figure 14-1.

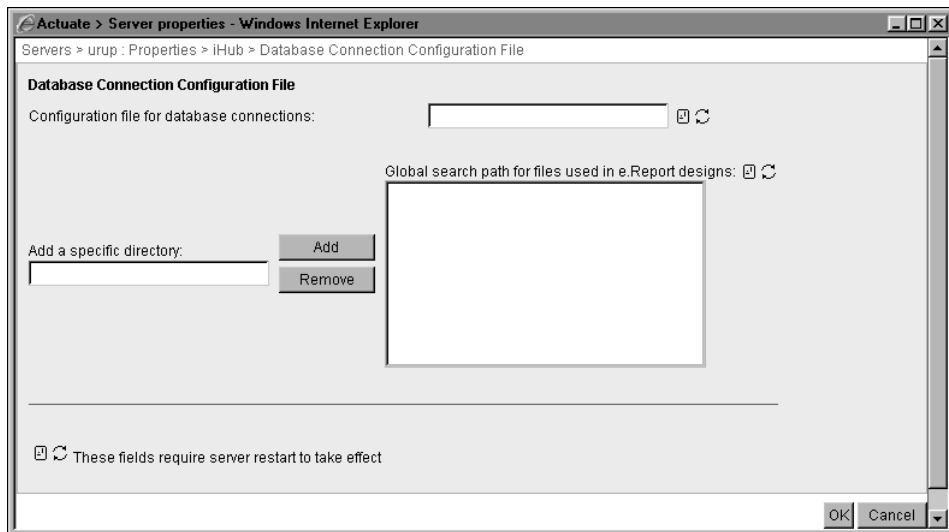


Figure 14-1 Specifying the location of a connection file

On Linux, the value of the parameter can be a path and file name only. On Windows, it can be either a path and file name or a URL. For example:

```
\\server1\configs\serverconfig.xml
```

or

```
http://myserver/configs/testconfig.xml
```

If you do not specify a value for the configuration file parameter, iHub uses the data source connection properties in the design.

Configuring a cluster to use a connection configuration file

In a cluster of iHubs, each node must have access to the connection configuration file. The path can be a local absolute path on each machine and must be specified

for each iHub. If you use a single copy of the file for a cluster, put it in a shared location, then specify the path to that shared location for all iHubs in the cluster.

Defining environment variables

The installation program defines environment variables for data sources that you specify during installation, but if you want to connect to data sources later, you need to set these variables.

On Linux, insert the library path environment definition in the file `pmd11.sh` in `AC_SERVER_HOME/bin`. The library path variable names are:

- `LD_LIBRARY_PATH` for Linux
- `SHLIB_PATH` for HP-UX

On Windows, you define system, not user, environment variables. iHub processes are system processes, and cannot access the user environment. Windows uses `PATH` to search for libraries and executable files.

In some cases, environment variables or registry entries affect the capability to display documents in different languages.

Table 14-1 lists data source-related Actuate registry keys on Windows, and environment variables on Linux and Windows.

Table 14-1 Environment variables

Key or variable name	Description
AC_DBMS_INFORMIX_MAXVARLEN	Maximum column length Actuate uses with Informix data sources. Default value is 4000.
AC_DBMS_ODBC_MAXVARLEN	Maximum column length Actuate uses with ODBC data sources. Default value is 8000.
AC_DBMS_ORACLE_MAXVARLEN	Maximum column length Actuate uses with Oracle data sources. Default value is 4000.
AC_DBMS_PROGRESS_MAXVARLEN	Maximum column length Actuate uses with Progress data sources. Default value is 4000.
DB2COMM	DB2 information.
DB2DIR	The path to the DB2 client installation.

(continues)

Table 14-1 Environment variables (continued)

Key or variable name	Description
DB2INSTANCE	Specifies the DB2 instance name.
DISPLAY	Linux. Specifies the X Windows server used by the machine.
DLC	Progress 9.1 installation directory.
INFORMIXDIR	The directory where the Informix product is installed.
INFORMIXSERVER	Specifies the name of the Informix data source.
INSTHOME	Inherited from shell and used by third-party processes.
ORACLE_HOME	The path to the Oracle installation.
SYBASE	The path to the Sybase installation.
SYBASE_OCS	The path to the Sybase OpenClient installation.

Connecting to data sources

The following sections describe how to connect to various external data sources from iHub System.

Connecting to a DB2 data source

iHub can connect from designs to IBM DB2 data sources on UNIX and Windows using run-time client version 9.1, 9.5, and 9.7. To connect to DB2 data sources, you must perform the following tasks:

- Define appropriate environment variables.
- Specify the database environment.
- Supply an account name and password.
- Specify the protocol your site uses.
- Configure the DB2 client's locale setting to match the locale of the machine.

The Actuate interface to DB2 clients supports DB2 stored procedures. The following limitations of DB2 stored procedures exist:

- DB2 does not support large object (LOB) data types between versions. See the DB2 documentation for information about DB2 support for LOB data types.

- You must bind the DB2 CLI packages from the client system to the DB2 server using `db2cli.lst` in some situations. For example, after you apply a Fix Pak to a client or server, or on a system where an Actuate design uses DB2 9.1 client software with a DB2 9.1 server.

Defining DB2 environment variables

For UNIX and Microsoft Windows operating systems, you must define the following environment variables:

- **DB2INSTANCE and DB2DIR**
Define these variables to connect to a DB2 instance. `DB2INSTANCE` specifies the instance name, and `DB2DIR` is the path to the DB2 client installation.
- **DB2CODEPAGE**
`DB2CODEPAGE` is a DB2-specific environment variable. For Windows and UNIX, Actuate uses `DB2CODEPAGE` to determine the DB2 database's client locale. On Windows systems, `DB2CODEPAGE` is a registry setting. On UNIX systems, `DB2CODEPAGE` is an environment variable. At execution time when Actuate connects to a DB2 data source, the active code page is in effect for the duration of the connection. All data is interpreted based on this code page. If this variable is not set, Actuate's DBMS module determines the client locale setting from the operating system locale.

Use `db2set` to set `DB2CODEPAGE`. For example, the following command sets the DB2 database client to retrieve data in UTF-8 format:

```
db2set DB2CODEPAGE=1208
```

When using the `db2set` command, add the location of `db2set` to the environment variable `PATH`. For example, if `db2set` is in `$DB2DIR/adm` and you use `db2set` in the `pmd11.sh` shell script, add the location of `db2set` to the environment variable `PATH` in `pmd11.sh`.

About the HP-UX11i Version 1 library path

The definition of the HP-UX11i in the environment variable `SHLIB_PATH` must put the path to ODBC and DB2 libraries before the `AC_SERVER_HOME` path. For example, if `ODBC/lib` and `DB2DIR/lib` are paths to the ODBC and DB2 libraries, use the following `SHLIB_PATH` definition:

```
SHLIB_PATH=$ODBC/lib:$DB2DIR/lib:$AC_SERVER_HOME/lib:usr/local  
/bin:
```

Using DB2 libraries on HP-UX

If you did not specify DB2 data source information during installation of iHub, and you now want to use DB2 data source connections from the machine, you must create a symbolic link to the DB2 library in the DB2 client installation directory.

Linking to a DB2 library on HP-UX

Create a symbolic link DB2CLI.sl in AC_SERVER_HOME/lib to the DB2 library libdb2.sl.

Checking a connection to a DB2 instance

To check that a connection exists between the machine and the DB2 instance, use the command-line utility DB2. This utility comes with the DB2 software and is available for Linux, and Windows.

To start the DB2 utility, open a command-line window and enter DB2 at the command prompt. At the DB2 prompt, enter the command to connect to a DB2 database.

```
connect to <database> user <user name>
```

The <database> is the name of the DB2 database, and <user name> is the DB2 database user. You are prompted for a password. Enter the password for the user. DB2 displays the connection information in the command window when a connection is made. Enter quit to terminate the session.

About using XML Extender

The XML Extender component of DB2 provides data types to store XML documents and DTDs in DB2 databases as either an XML column or XML collection. Actuate retrieves the XML column data as string data. DB2 stores an XML collection as a set of tables. Actuate retrieves data from the set of DB2 tables.

Connecting to an Informix data source

When accessing data from Informix using the included DataDirect ODBC Informix data driver, set the UseDelimitedIdentifier property to 1. To specify the property value for connection-string based designs, add the property UseDelimitedIdentifier=1 to the end of the connection string. To specify the property value for DSN-based connections, perform the appropriate task for your operating system.

On a Windows system, go to the Windows ODBC registry entry:

```
HKLM\Software\ODBC\ODBC.ini\<DSN Name>
```

and set the value of UseDelimitedIdentifier to 1.

On a Linux system, edit the odbc.ini file, and add UseDelimitedIdentifier=1 to the DSN entry.

Connecting to an Oracle data source

To connect with Oracle data sources, you must perform the following tasks:

- Install the proper Oracle client software on the server running the design.
- Supply a connection string.
- Define the appropriate environment variables.
- Ensure that a listener process is running on the database host.

When configuring the Oracle client, you must configure the client's locale setting to match the locale of the machine.

Defining Oracle environment variables

You must set the following environment variables when connecting to an Oracle data source:

- ORACLE_HOME
- NLS_LANG

About ORACLE_HOME

For Oracle databases on Linux platforms, the account running iHub processes must have a definition for the standard Oracle environment variable ORACLE_HOME. You can provide this definition in a login script such as .cshrc or .profile, or you can include it in the scripts that start the iHub processes.

For Oracle databases on Windows, also ensure that the definition of the environment variable PATH indicates the location of the dynamic link library that selects the proper database.

About NLS_LANG

The Oracle Linux environment variable and Windows registry setting NLS_LANG specifies the Oracle locale, that consists of the language, territory, and character set. The default value for NLS_LANG is American_America.US7ASCII. The administrator must ensure the NLS_LANG setting is correct for the information in the Oracle database.

On Linux systems, add the environment variable NLS_LANG to the pmd11.sh script, the iHub request server startup script. On Windows servers, the Oracle installer configures NLS_LANG.

The following example sets NLS_LANG for simplified Chinese on Linux:

```
export NLS_LANG  
NLS_LANG="Simplified Chinese_China.ZHS16GBK"
```

Double quotes are required when setting a value that contains spaces.

On Microsoft Windows server operating systems, set the NLS_LANG registry value in the key HKEY_LOCAL_MACHINE \Software\Oracle\Home.

For information about the NLS_LANG values, see the Oracle documentation.

About the Oracle listener processes

iHub interacts with an Oracle data source through the Oracle listener process. If iHub is having trouble communicating with the Oracle data source, the link to the listener process possibly failed. You can frequently solve such problems by stopping and starting the listener process.

Setting the maximum column length

The default maximum column length Actuate uses with Oracle data sources is 4000 characters. Problems can occur when generating an Actuate document if the design uses an Oracle data source column with a large column length. You can use a registry setting on Windows or an environment variable on Linux to set a smaller maximum column length used by Actuate.

To change the maximum character length:

- On Windows, change the value of the string value name MaxVarLen. The value name is in the registry key.
`HKEY_CURRENT_USER\Software\Actuate\DBMS\Oracle`
- On Linux, set the environment variable AC_DBMS_ORACLE_MAXVARLEN.

Connecting to custom data sources

iHub runs designs using third-party software to connect to data sources. The following sections describe how to install custom data source connection software:

- Installing a custom open data access (ODA) driver
- Installing a custom Eclipse DTP ODA driver plug-in
- Using custom Java-based data source connections
- Using custom relational data sources with the Integration service

Installing a custom open data access (ODA) driver

iHub supports open data access (ODA) drivers that a design uses to retrieve data from a data source. To install an ODA driver, place the ODA driver files in the AC_SERVER_HOME/oda directory. AC_SERVER_HOME refers to the folder that the iHub installer chose as the location for binary files during the iHub installation. By default, this location is /Program Files (x86)/Actuate/iHub2 in a

Windows environment, and <installation directory>/AcServer in a Linux environment.

Each ODA driver must be in a separate directory. For example, if you have two ODA drivers, the installation directories would be similar to the following directories:

- AC_SERVER_HOME/oda/oda-driver1
- AC_SERVER_HOME/oda/oda-driver2

iHub does not support installing ODA drivers in a directory below the oda directory. For example, if you have two ODA drivers, Driver1 and Driver2, iHub does not support the following ODA directory structure for the two drivers:

- AC_SERVER_HOME/oda/CustomDrivers/Driver1
- AC_SERVER_HOME/oda/CustomDrivers/Driver2

The directory name for an ODA driver in the AC_SERVER_HOME/oda directory must match the driver name specified in an Actuate Basic information object (.dox).

When installing an ODA driver, you must also install and configure any software the ODA driver requires to access a data source. For example, you must install and configure any database connection software the ODA driver uses to connect to a database. You must also ensure the ODA driver can access the required software.

Install the configuration file and files required for the ODA driver on the machine that runs the design. Each ODA driver requires a run-time configuration file, odaconfig.xml. The configuration file must be in the ODA driver directory. You can locate the ODA library files outside of the Actuate installation directory. Specify the location in odaconfig.xml.

The iHub installation process puts the Actuate Data Integration service driver in the AC_SERVER_HOME/oda directory. The installation also provides an example of an ODA flat file driver that operates with iHub in ACTUATE_HOME/oda/examples/FlatFileExample.

Installing a custom Eclipse DTP ODA driver plug-in

A design can use a custom Eclipse Data Tools Platform (DTP) open data access (ODA) driver to retrieve data from a data source. The DTP ODA API supports building a custom Eclipse plug-in that accesses data from standard and custom data sources.

Install a plug-in in the AC_SERVER_HOME/MyClasses/eclipse/plugins directory. You may need to manually create the eclipse/plugins subdirectory if it does not already exist in MyClasses. In the Linux installation of iHub, there is a common folder, acshared/MyClasses, for installing a customized DTP ODA driver.

Do not place a custom ODA plug-in in the AC_SERVER_HOME/oda/eclipse/plugins directory. This directory is reserved only for plug-ins provided by the Actuate product installation. Installing a plug-in in the AC_SERVER_HOME/MyClasses/eclipse/plugins directory ensures that the Actuate install and uninstall processes do not remove the custom plug-in.

About using an ODA driver

When running a design using an ODA driver, the Factory process loads the driver during the document generation. If the Factory process cannot load the specified driver, iHub logs an error message and document generation terminates.

An ODA driver cannot share a library file with another ODA driver or data source connection software. Each ODA driver must have a separate copy of the library file for dedicated use. For example, on a Windows system, if a database library uses a DLL to connect to a data source, and an ODA driver uses the same DLL, you create a copy of the DLL file and use the copy with the ODA driver.

If you change an ODA driver's configuration such as a setting in the configuration file or an ODA driver library, the Factory process uses the updated configuration information and the updated library during the next document generation.

iHub can cache an ODA driver. If a Factory process uses a cached ODA driver, the Factory process checks the last modified time of the configuration file and the cached driver's run-time libraries before document generation. If you modify the configuration file or the driver since last loading the driver, the Factory process releases the cached driver and reloads it.

Using custom Java-based data source connections

To connect BIRT designs to JDBC data sources, place the .jar files for the custom database driver in:

```
AC_SERVER_HOME\Jar\birt\platform\plugins\  
org.eclipse.birt.report.data.oda.jdbc<version>\drivers
```

The iHub Integration service also looks in this location for drivers to access JDBC data sources.

Using custom relational data sources with the Integration service

iHub connects to Actuate information object data sources using the BIRT iHub Integration service (AIS). AIS works with JDBC drivers that are compliant with DataDirect Connect for JDBC version 3.7 Service Pack 1 and later.

AIS uses configuration files that define the data source connections.

- AC_SERVER_HOME/etc/data_integration/datasources.xml contains a list of custom relational data sources.

- AC_SERVER_HOME/etc/data_integration/<database>/mappings.xml contains the data source mappings for a custom relational data source.

Specifying connection types

iHub installation provides an example of a datasources.xml configuration file that specifies a MySQL Enterprise connection type as an example. The datasources.xml file specifies the JDBC connection type as shown in Listing 14-2.

Listing 14-2 Specifying JDBC connection types in datasources.xml

```
<DataSourceConfig>

    <DataSourceHosts>
    </DataSourceHosts>

    <ConnectionTypes>

        <
-- Example: MySQL Enterprise connection type
   (requires MySQL driver) -->
        <ConnectionType Name="MySQL_41">
            <JDBCDriver DriverName="com.mysql.jdbc.Driver">
                <ConnectionString>
                    jdbc:mysql://%server%:%port%/%database%?
                </ConnectionString>
                <ConnectionProperties>
                    <Property Name="user">%username%</Property>
                    <Property Name="password">%password%</Property>
                </ConnectionProperties>
                <LibraryPath>
                    <
-- Fill in JAR location below -->
                    <Location></Location>
                </LibraryPath>
            </JDBCDriver>
            <ConnectionParams>
                <ConnectionParam Name="server"
                    Display="Server"
                    Type="string"
                    ValueIsCaseSensitive="false"/>
                <ConnectionParam Name="database"
                    Display="Database"
                    Type="string"
                    ValueIsCaseSensitive="false"/>
                <ConnectionParam Name="username"
                    Display="User name"
                    Type="string"/>
            </ConnectionParams>
        </ConnectionType>
    </ConnectionTypes>
</DataSourceConfig>
```

```

        <ConnectionParam Name="password"
            Display="Password"
            Type="masked"/>
        <ConnectionParam Name="port"
            Display="Port"
            Type="integer"
            Optional="true"
            DefaultValue="3306"/>
    </ConnectionParams>
</ConnectionType>

</ConnectionTypes>

<DatabaseTypes>
    <
-- Example: MySQL data base type -->
    <
--
        <DatabaseType Name="MySQL_41" DisplayName="MySQL 4.1"
            ConnectionType="MySQL_41"
            DataSourceMapper="BaseMySQLMapper"
            DataSourceMapping="MySQL"/>
    -->
    </DatabaseTypes>
</DataSourceConfig>

```

Specify the location of the JDBC driver in <Library Path>, as shown in Listing 14-2. Place the driver classes or Java archive (.jar) files accessing JDBC data sources in the /drivers directory, as described earlier in this document.

Using a connection pool

Using the ODA Java interface, a developer can create an ODA driver that pools connections. For example, when a Factory process uses an ODA driver, the driver can create a connection to the data source. When the Factory process requests another connection to the data source, the ODA driver can return a new connection or reuse the previous connection instance. Connection pooling optimizes application performance and improves scalability.

You can use an application server, such as a J2EE application server, to implement a connection pool to a data source that supplies data to iHub. In a typical J2EE application server environment, the connection pool uses a data access object (DAO) to provide a common interface between iHub and the data storage system. A DAO separates the application logic from the data access logic to provide a re-usable, persistent connection.

Figure 14-2 shows a J2EE application server configuration that supports multiple Information Console client sessions. Each viewed document contains data extracted by an iHub from a data source using a shared connection.




Figure 14-2 Implementing a connection pool using a J2EE application server

You need the Deployment Kit option to enable connection pooling for Actuate Java Components.

ODBC/JDBC connection pooling

Connection pooling allows applications to share a connection for greater efficiency. BIRT iHub connects to an external data source using either an out-of-the-box (OOTB) JDBC driver or an Actuate information object.

Configuring BIRT JDBC connection pooling

In Server Configuration Templates>Settings>Factory Service and View Service, the administrator configures the BIRT JDBC connection pool by configuring the following properties:

- **BIRT JDBC Connection Pool Size**
Caches concurrent connections, so the Factory can reuse them. The default is 10. Increase the pool size to establish more connections. To disable the connection pool, set BIRT JDBC Connection Pool Size to 0. The larger the pool size, the better performance in a highly concurrent system and the greater the memory consumption.
- **BIRT JDBC Connection Pool Timeout**
Specifies the time-out interval for cached connections. The default is 3600 seconds. Increasing the available connections increases performance, but also increases memory consumption.

- **BIRT Connection Validation Interval**

Controls the frequency of BIRT connection validation. Connection validation prevents a query from attempting to use a broken connection. The default value of -1 disables connection validation.

The interval is expressed in seconds. If the value is positive, iHub first determines whether the last validation occurred in the specified interval.

If not, iHub performs validation by calling `java.sql.Connection.isValid()`. All DataDirect JDBC drivers support `java.sql.Connection.isValid()`. Because some JDBC driver do not support `java.sql.Connection.isValid()`, the administrator must restart iHub after rebooting the metadata database to refresh the JDBC connection pool even if the interval has not changed.

Figure 14-3 shows Server Configuration Templates>Settings>Factory Service>BIRT JDBC Connection Pool.

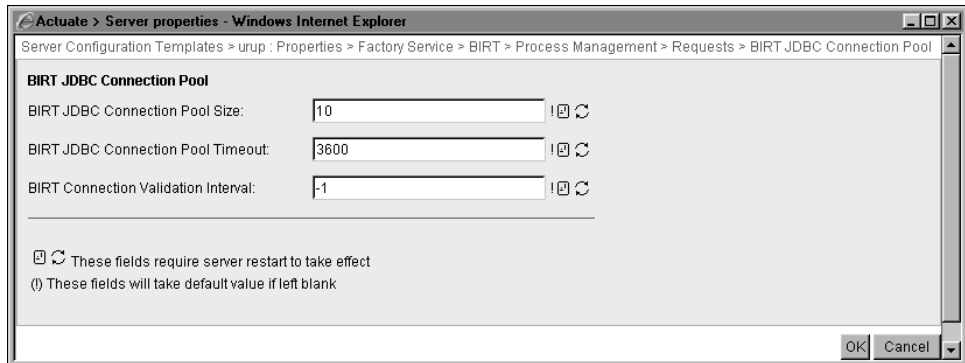


Figure 14-3 Specifying BIRT JDBC connection pool properties

How to configure BIRT JDBC connection pool properties

To configure BIRT JDBC connection pool properties, perform the following tasks:

- 1 Log in to Configuration Console, choose Advanced view, then choose Server Configuration Templates.
- 2 On Server Configuration Templates, choose a template.
- 3 On Settings, expand Factory Service or Viewing Service, then, expand BIRT, Process Management, and Requests, as shown in Figure 14-4 or Figure 14-5. Choose BIRT JDBC Connection Pool.
- 4 In BIRT JDBC Connection Pool Size, accept the default value, 10 connections. Alternatively, decrease or increase the pool size by typing a number not less than 0 or greater than 65535.
- 5 In BIRT JDBC Connection Pool Timeout, accept the default value, 3600 seconds. Alternatively, increase or decrease the time-out period by typing a different number.

- 6 In BIRT Connection Validation Interval, accept the default value, -1, to leave connection validation disabled. Alternatively, enable connection validation by specifying a period of time in seconds.

Figure 14-3 shows BIRT JDBC Connection Pool set to default values.

- 7 Restart iHub after setting any of these values.

Figure 14-4 shows BIRT JDBC Connection Pool in Server Configuration Templates>Settings>Factory Service.

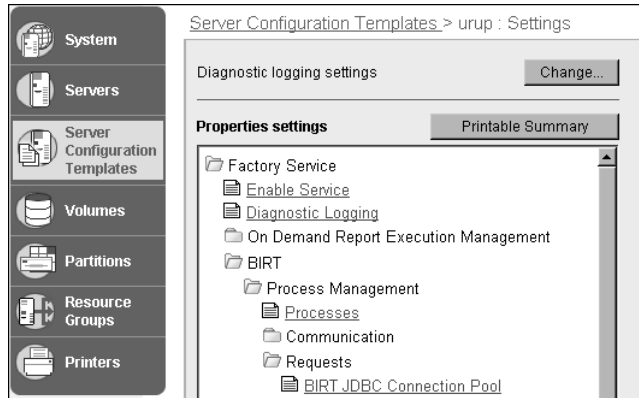


Figure 14-4 Choosing BIRT JDBC Connection Pool for Factory service

Figure 14-5 shows BIRT JDBC Connection Pool in Server Configuration Templates>Settings>View Service.

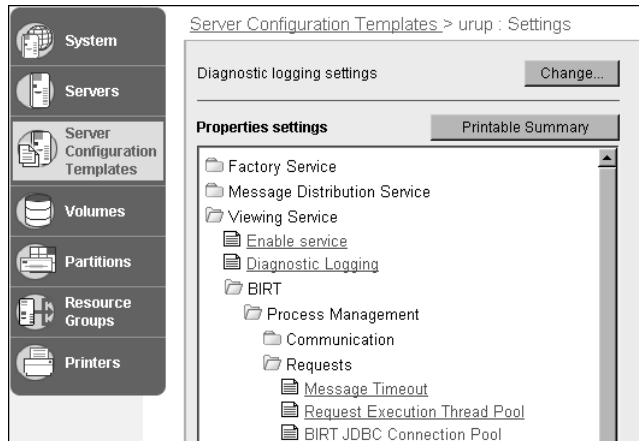


Figure 14-5 Choosing BIRT JDBC Connection Pool for View service

Setting the Encyclopedia engine connection pool reap interval

The connection pool reap interval specifies the time, in seconds, between runs of the maintenance thread. When the maintenance thread runs, it discards any connections that remain unused for longer than the specified time.

The administrator sets the default connection pool reap interval by performing the following tasks:

- 1 Log in to Configuration Console and choose Advanced view.
- 2 Choose Server Configuration Templates. On Settings, expand the iHub folder. Choose Encyclopedia Engine.
- 3 In Default connection pool reap interval in seconds, accept the default value, 600, as shown in Figure 14-6. Alternatively, type a new value, from 0 to 86400. To disable the pool maintenance thread, set the interval to 0.

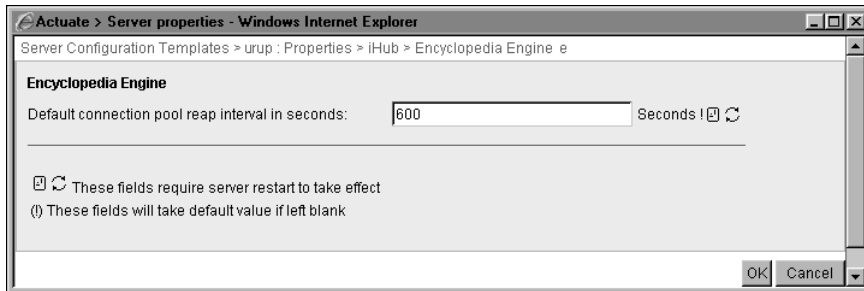


Figure 14-6 Specifying default connection reap interval

Setting miscellaneous properties

This chapter contains the following topics:

- Changing locale, encoding, and time zone
- Modifying general volume properties
- Changing ports used by iHub
- Viewing and modifying general server properties
- Changing iHub system start-up parameters
- Configuring general system properties
- About General
- Setting start-up arguments for the Encyclopedia server JVM
- Starting and stopping iHub
- Setting conversion queue and e-mail queue resource management properties
- Configuring dates
- Using Printable Summary to view system properties
- Configuring dates
- Using Printable Summary to view template properties
- Setting Filetype driver properties

Changing locale, encoding, and time zone

The language and regional settings of your computer operating system must be compatible with the default locale settings that you select. Incompatibilities can cause problems with the character set used for entries from your keyboard and the language used in the iHub graphical user interface.

You can specify a time zone:

- When you log in. This value overrides the settings in Options—General.
- On System—Properties—Regional Settings.
- On Options—General. This value sets the locale and time zone of the browser and workstation. The machine that runs the browser stores this setting locally.

Figure 15-1 shows the default regional settings for iHub or a cluster in System—Properties—Regional Settings.



Figure 15-1 Configuring default regional settings

Figure 15-2 shows Options—General in the advanced view of Configuration Console, where you specify the locale and time zone to use for iHub.




Figure 15-2 Specifying locale and time zone in Options—General

The following parameters appear in Options—General:

- **Locale**
The user selects a locale during login. This locale appears in the Configuration and Management Consoles. If a user does not specify a value for this parameter during login, iHub reads the locale from the user's web browser cookie.
- **Time zone**
The user selects a time zone during login. This time zone appears in the Configuration and Management Consoles. If a user does not specify a value for this parameter during login, iHub reads the time zone from the user's web browser cookie.

The TimeZones.xml file in the following directory stores information about time zones:

```
\Program Files (x86)\Actuate\iHub2\servletcontainer\mgmtconsole
  \WEB-INF
```


TimeZones.xml stores time zones in the format that the tz database, or zoneinfo database prescribes.

The LANG environment variable on Linux specifies the machine's language code.

Modifying general volume properties

You can view or modify the general property values for an Encyclopedia volume in Volumes—Properties—General using the following procedure.

How to modify general property values for an Encyclopedia volume

- 1 Log in to Configuration Console and choose Advanced view.
- 2 From the side menu, choose Volumes.
- 3  Choose the arrow next to the volume name, then choose Properties.

- 4 In Volumes—Properties—General, shown in Figure 15-3, you can modify the schedule for purging notices, purging deleted files, the volume archive service provider, and the partition for the e-mail notification template.

Choose OK.

The properties you can set in Volumes—Properties—General are the same whether you are modifying an existing volume or creating a new volume, with the exception that you can specify the volume name when creating a volume.

Volumes > corp : Properties

General Open Security Partitions Events Advanced

Description:

Schedule for purging notices: HH:mm

Schedule for purging deleted files: HH:mm

Partition

Primary partition: DefaultPartition Min Free Space: MB

Volume archive service provider

Use archive service:

Metadata database and schema

Metadata database name: Default_ActuatePostgreSQL_MetadataDatabase

Database schema name: ac_corp

Email notification

E-mail notification template partition:

Use Information Console for e-mail notifications

Information Console URL prefix:

These fields require volume restart to take effect

(!) These fields will take default value if left blank

OK Cancel Apply

Figure 15-3 Modifying general property values for an Encyclopedia volume

Changing ports used by iHub

Table 15-1 lists a number of key ports. You can view the port numbers appearing on Servers—Properties—General. To change these port numbers, you must edit

acserverconfig.xml. You can view or change the port numbers that you access from Server Configuration Templates—Settings.

Table 15-1 Setting iHub ports

Port	Configuration Console location
Daemon listen port	Servers—Properties—General
The application container process listen port for the Management and Information Consoles	Servers—Properties—General
Server port	Servers—Properties—General
Server port base	Servers—Properties—General
Message distribution, which the Message Distribution service uses	Server Configuration Templates—Settings—Message Distribution Service—Process Management—Communication
Base port for Java Factory server	Server Configuration Templates—Settings—Factory Service—BIRT—Process Management—Communication—Sockets
Base port for Java View server	Server Configuration Templates—Settings—Viewing Service—BIRT—Process Management—Sockets

Table 15-2 contains a list of iHub ports set during installation, default port numbers, and ranges of values.

Table 15-2 iHub ports

Name	Display name	Default	Range	Changeable
AppContainerPort	Application container process listen port	8900	1 - 65535	Yes
CustomEventServicePort	Custom Event Service Port	8900	1 - 65535	Yes
NWPPort	Network process port for integration server queries	14000	1024 - 65535	Yes
PMDPort	Daemon listen port	8100	None	No

(continues)

Table 15-2 iHub ports (continued)

Name	Display name	Default	Range	Changeable
ReportEngineHeartBeatPort	Port number for receiving Factory server heartbeat	11101	None	No
RSSESOAPPort	RSSE service port	8900	1 - 65535	Yes
ServerSOAPPort	Port number for iHub internal SOAP endpoint	11100	None	No
ServerSOAPPortBase	Base port number for iHub internal SOAP endpoint	13500	None	No
SMTPPort	Listen port	25	1 - 65535	Yes
SOAPDispatchSOAPPort	Message Distribution service port	8000	1 - 65535	Yes
SOAPPort	Port for caching server operations	11550	1024 - 65535	Yes
SOAPPort	Port for integration server operations	12100	1024 - 65535	Yes
SocketBaseForJavaProcesses	Base port for Java Factory server	21500	1025 - 65535	Yes
SocketBaseForProcesses	Base port number for processes	18500	1025 - 65535	Yes
SocketBaseForJavaProcesses	Base port for Java View server	21000	1025 - 65535	Yes

Viewing and modifying general server properties

Choose a machine name from the list of servers on Servers in the advanced view of Configuration Console. In Servers—Properties—General, shown in Figure 15-4, you view general property values for a machine.

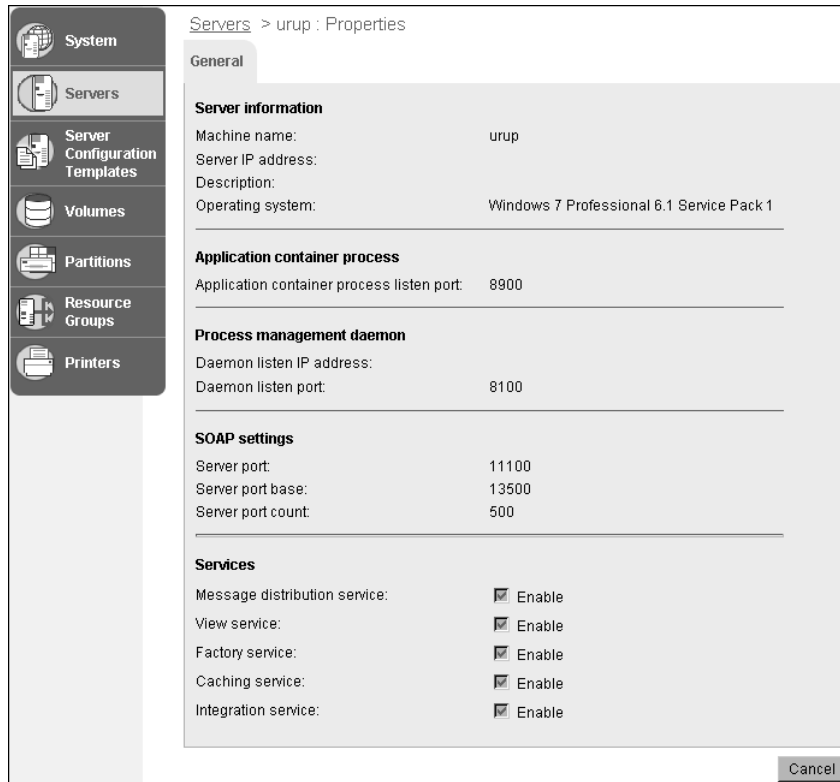


Figure 15-4 Viewing general properties for a server

The administrator can modify the following general server properties from Server Configuration Templates—Settings:

- Application container process listen port
- Server port
- Server port base
- Server port count

Services are enabled by default in the Template element in `acserverconfig.xml`.

The administrator modifies general server properties by expanding Server Configuration Templates—Settings—iHub—Process Management—Communication, and choosing the Application Container and Internal SOAP Endpoint options, as shown in Figure 15-5.

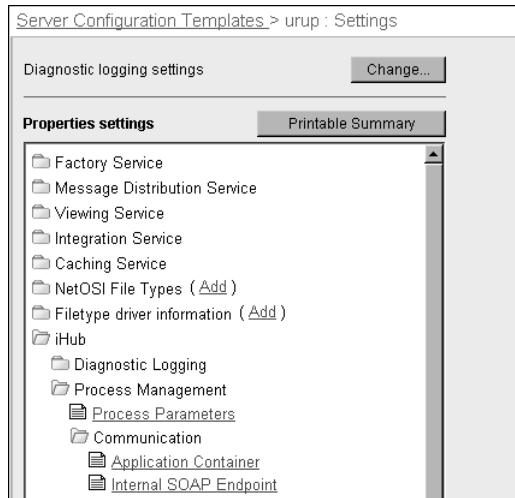


Figure 15-5 Choosing to view or modify general server properties

About application container process

This property specifies the application container port number. iHub uses an application container to host web services applications. Start and stop the application container process using the StartMC and StopMC scripts in AC_SERVER_HOME/bin.

How to set the application container process listen port

- 1 On Server Configuration Templates—Settings, expand iHub, Process Management, and Communication, then choose Application Container, as shown in Figure 15-5.
- 2 In Application container process listen port, accept the default, 8900, as shown in Figure 15-6. Alternatively, specify a different value.

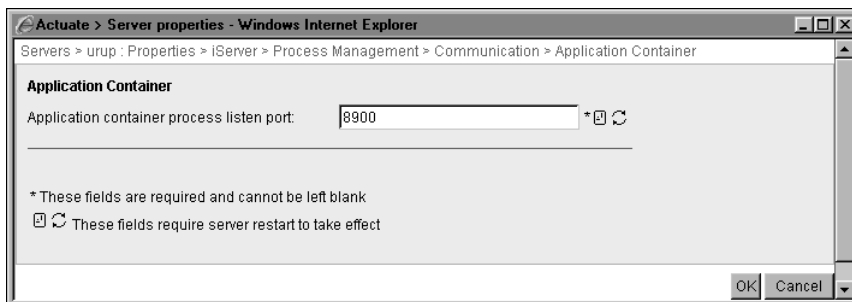


Figure 15-6 Viewing or changing the application container process listen port

About SOAP settings

On Server Configuration Templates—Settings—iHub—Process Management—Communication—Internal SOAP Endpoint, the administrator can specify the following properties:

- Port number for iHub internal SOAP endpoint
Web service API (IDAPI) port and internal server port for communication with other nodes in a cluster. iHub uses this port to receive information, such as heartbeat messages.
- Base port number for iHub internal SOAP endpoint
Beginning of a range of port numbers that iHub attempts to use for internal SOAP messages. Used with server port count.
- Port range (from base) for iHub internal SOAP endpoint
The range of port numbers of ports used for exchanging internal SOAP messages starting from the server port base.
- Port number for iHub Encyclopedia engine internal SOAP endpoint

How to set the SOAP settings properties

- 1 On Server Configuration Templates—Settings—iHub—Process Management—Communication, choose Internal SOAP Endpoint, as shown in Figure 15-5.
- 2 On Internal SOAP Endpoint, as shown in Figure 15-7, perform the following tasks:
 - 1 In Port number for iHub internal SOAP endpoint, accept the default, 11100. Alternatively, specify a different value.
 - 2 In Base port number for iHub internal SOAP endpoint, accept the default, 13500. Alternatively, specify a different value.
 - 3 In Port range (from base) for iHub internal SOAP endpoint, accept the default, 500. Alternatively, specify a different value.
 - 4 In Port number for iHub encyclopedia engine internal SOAP endpoint, accept the default, 14100. Alternatively, specify a different value.
 - 5 In Port number for iHub encyclopedia engine internal SOAP endpoint for servicing iHub components requests, accept the default, 14200. Alternatively, specify a different value.
 - 6 In Java encyc server purge thread limit, accept the default, 2. Alternatively, specify a different value.
- 3 If you change any property values, restart iHub.

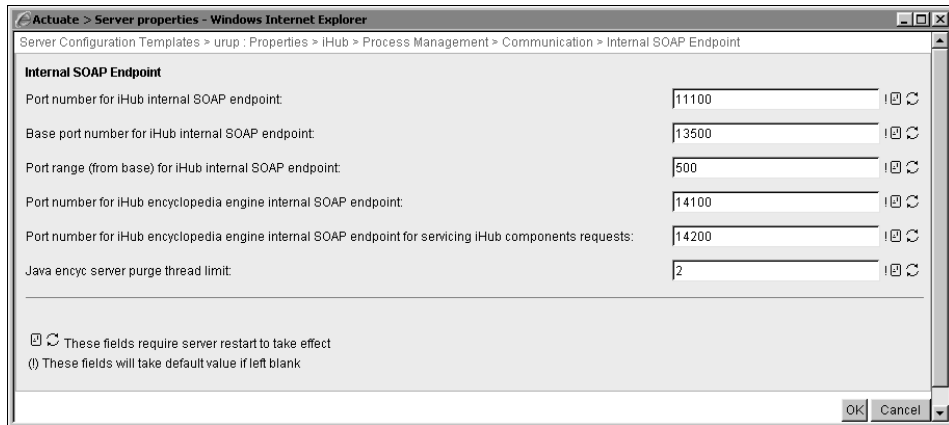


Figure 15-7 Specifying SOAP settings

Changing iHub system start-up parameters

Before starting the iHub system, you can specify an alternate configuration file as well as an alternate template for this iHub node to use. When you start the iHub system, iHub will use these resources.

How to modify start-up parameters for iHub

- 1 Put iHub offline by choosing Stop on System—Status, as shown in Figure 15-8.

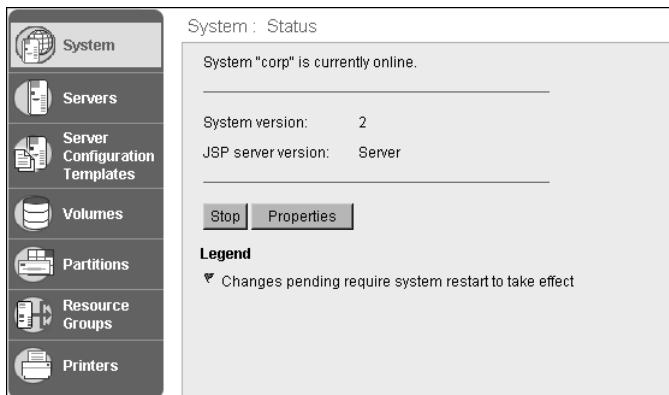


Figure 15-8 Putting iHub offline

- 2 On System—Status, choose Modify Start Parameters, as shown in Figure 15-9.




Figure 15-9 Choosing Modify Start Parameters

- 3 On System—Modify Start Parameters, set start-up parameters by performing the following tasks:
 - 1 In Configuration home, type the path of the folder containing an alternate acserverconfig.xml. This folder must also contain acserverlicense.xml. By default, the location of acserverlicense.xml is AC_DATA_HOME/config/iHub2.
 - 2 In Template name, type the name of the template that you want iHub to use, as shown in Figure 15-10.

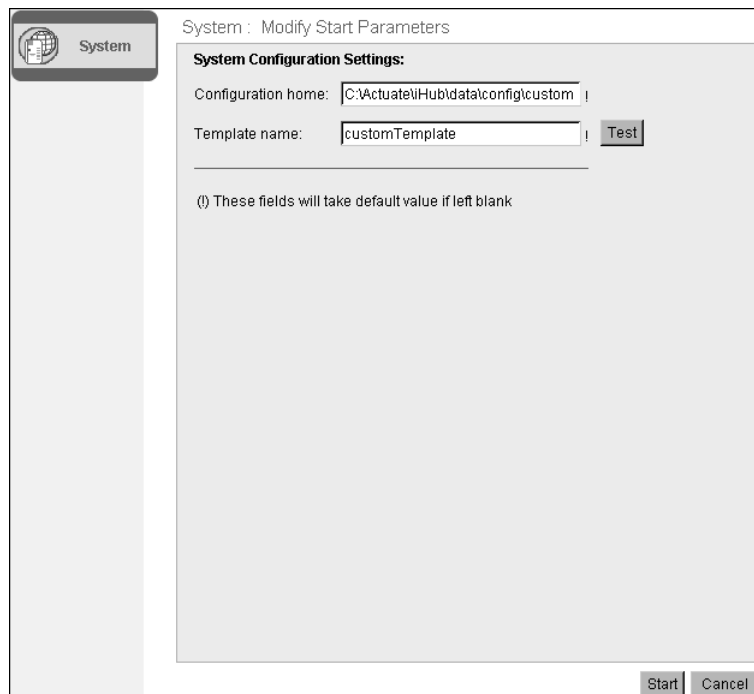


Figure 15-10 System—Modify Start Parameters

- 3 Choose Test to verify that iHub recognizes the parameters. A message box appears, confirming that the test is successful. Choose OK, as shown in Figure 15-11.



Figure 15-11 Confirming that the parameters are valid

- 4 On System—Modify Start Parameters, choose Start.

Configuring general system properties

In System—Properties—General, the administrator can view or change the following general system properties, as shown in Figure 15-12:

- System name of the cluster node or stand-alone iHub
- The password for logging in to Configuration Console

Figure 15-12 Changing general system properties

If you change the system name, restart iHub.

About General

In System—Properties—Advanced, the administrator can choose General to view or set values for the following properties:

- Cluster schema
Name of the iHub system schema
- Maximum number of user properties cache entries
Tunable property for holding a user's expanded role information in a cache
- Enable Multi-system mode
Supports joining multiple clusters together

- Startup arguments for the cluster JVM
Arguments passed to the JVM that communicates with the system schema

How to set general properties

- 1 On System—Properties—Advanced, choose General, as shown in Figure 15-13.

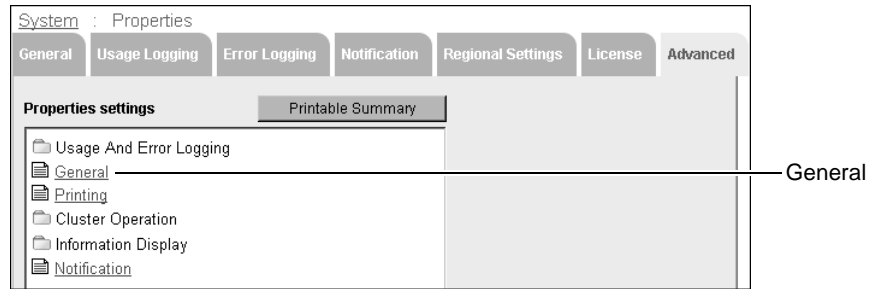


Figure 15-13 Viewing System—Properties—Advanced

- 2 On General, perform the following tasks:
 - 1 Accept the default values for Cluster Schema, as shown on Figure 15-14. Alternatively, specify a different schema.
 - 2 In Maximum number of user properties cache entries, accept the default value, 500, or alternatively, specify a different value.
 - 3 For Enable Multi-system mode, accept the default value of unchecked, or alternatively, select the property.
 - 4 In Startup arguments for the cluster JVM, accept the default arguments. Alternatively, modify these arguments.

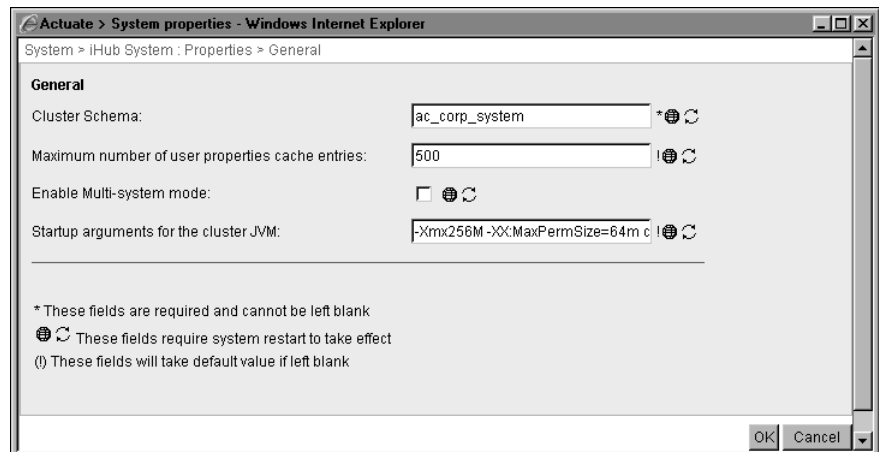


Figure 15-14 Configuring General

Setting start-up arguments for the Encyclopedia server JVM

In Server Configuration Templates—Settings—iHub, the administrator can choose to view or edit the start-up arguments for the Encyclopedia server JVM.

How to set server configuration properties

- 1 Expand iHub and Process Management, then choose Process Parameters, as shown in Figure 15-15.

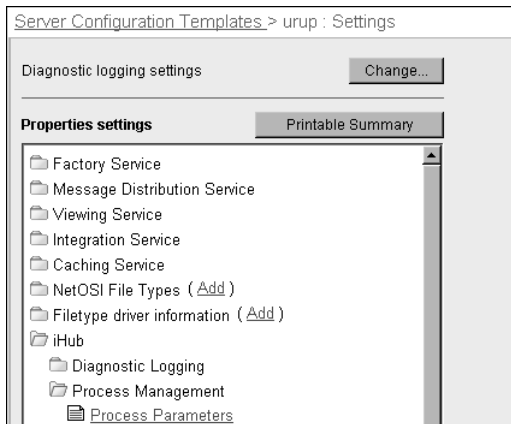


Figure 15-15 Choosing Process Management properties

- 2 In Start parameters for iHub encyclopedia engine, as shown in Figure 15-16, accept the default parameters. Alternatively, modify these parameters.

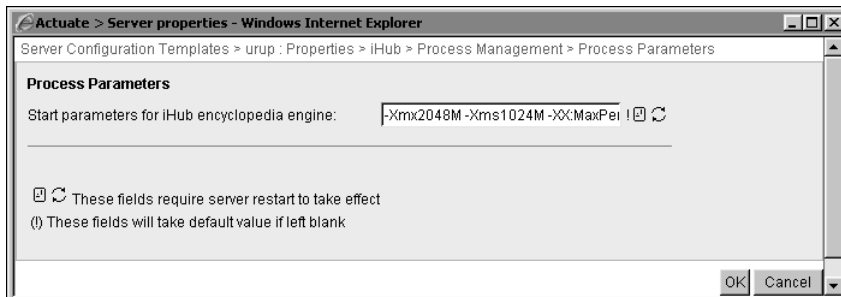


Figure 15-16 Viewing or setting start parameters for the Encyclopedia engine

If you change the property, restart iHub.

Starting and stopping iHub

System—Status displays iHub status and provides controls for starting and stopping iHub, cluster-level management, and cluster creation, as shown in Figure 15-17.




Figure 15-17 Viewing System—Status

System—Status displays the following information about iHub:

- iHub name and its current status
- Release version of the iHub
- Name and version of the application server

For example, in Figure 15-17, System—Status indicates that a machine is online. Actuate is running.

Earlier in this book you learned about the possible iHub states that appear in the simple view of Configuration Console. The same states can appear in System—Status. Two of these states are online and offline.

Stopping iHub

When iHub is online, you can access the following functionality in System—Status:

- Choose Stop to stop iHub. Choosing Stop for a node stops this node only.
- Choose Properties to view System—Properties and modify the following property values, which are described later:
 - General properties
 - System usage and error logging
 - E-mail notification
 - Default regional settings
 - License information

Starting iHub

When iHub is offline, you can start iHub from System—Status by choosing Start system, as shown in Figure 15-18.




Figure 15-18 Starting the system

Setting conversion queue and e-mail queue resource management properties

Conversion queue and e-mail queue resource management controls thread and queue length allocations for the following processes:

- Conversion queue
 - Number of conversion threads

The number of threads available for document conversion tasks, such as conversion to PDF. Specify a number between 1 and 50, inclusive. Actuate recommends setting this number to a value that is equivalent to the number of Async Factories plus 1 to 2 if all factories are expected to be filled during peak hours.
 - Max conversion queue length

The maximum number of document conversion tasks in the queue. Specify a number between 1000 and 1000000, inclusive.
- E-mail queue
 - Number of e-mail notification threads

The number of threads available for e-mail notification tasks. Specify a number between 1 and 50, inclusive.
 - Max e-mail notification queue length

The maximum number of e-mail notification tasks allowed to wait for execution. Specify a number between 1000 and 1000000, inclusive.

Restart iHub after setting any of these values.

How to configure conversion queue and e-mail queue resource management properties

To configure conversion queue and e-mail queue resource management properties, perform the following tasks:

- 1 Log in to Configuration Console, choose Advanced view, then choose Server Configuration Templates.
- 2 On Server Configuration Templates, choose a template.
- 3 On Settings, expand iHub folder. Choose Conversion Queue and E-mail Queue Resource Management, shown in Figure 15-19.

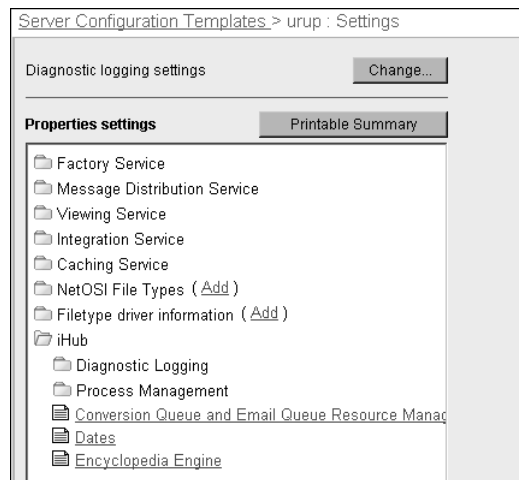


Figure 15-19 Choosing Conversion Queue and E-mail Queue Resource Management properties

- 4 In Number of conversion threads, accept the default value, 1. Alternatively, increase the number of threads available for document conversion by typing a number not greater than 50.
- 5 In Max conversion queue length, accept the default value, 10000. Alternatively, decrease or increase the queue length by typing a number not less than 1000 or greater than 1000000.
- 6 In Number of e-mail notification threads, accept the default value, 1. Alternatively, increase the number of threads available for e-mail notification by typing a number not greater than 50.
- 7 In Max e-mail notification queue length, accept the default value, 10000. Alternatively, decrease or increase the queue length by typing a number not less than 1000 or greater than 1000000.

Figure 15-20 shows Conversion Queue and E-mail Queue Resource Management set to default values.

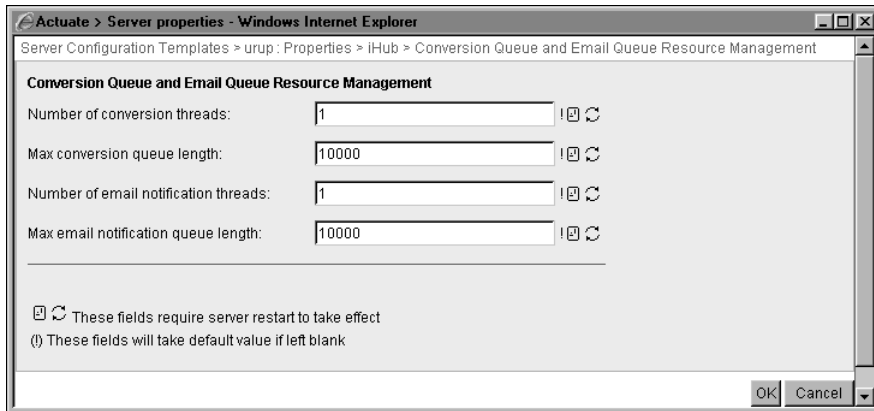


Figure 15-20 Specifying conversion queue and e-mail queue resource management properties

- 8 Restart iHub after setting any of these values.

Configuring dates

In Server Configuration Templates—Settings—iHub, the administrator can choose Dates, to specify how iHub formats dates.

How to set date formatting for iHub

- 1 Choose iHub—Dates, as shown in Figure 15-21.

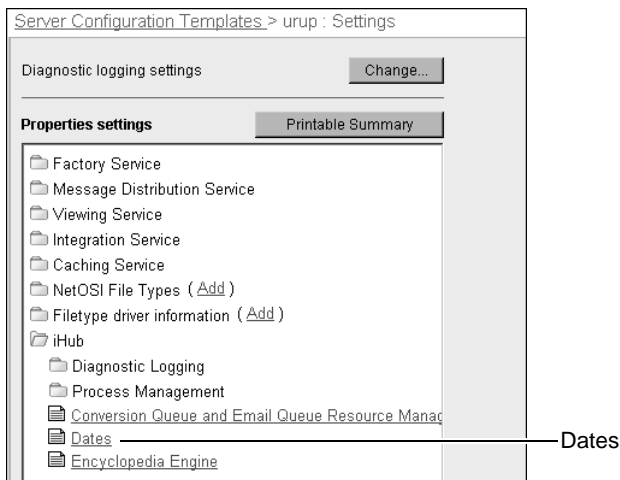


Figure 15-21 Viewing iHub properties

- Specify values for the properties appearing on Dates, as shown in Figure 15-22. Alternatively, accept the default values. Table 15-3 describes these properties.

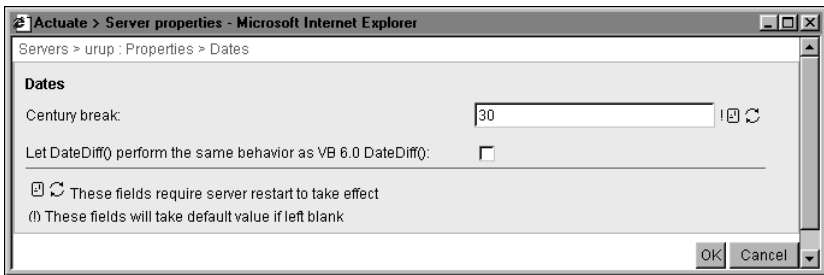


Figure 15-22 Specifying iHub date property values

Table 15-3 Date properties

Property	Description
Century break	<p>Indicator of how iHub converts 2-digit year values into 4-digit years. Using the default value 30, iHub converts a date using a 2-digit year in the following manner:</p> <ul style="list-style-type: none"> A value less than 30 becomes 20xx. A value of 30 and greater becomes 19xx. <p>For example, the date 3-5-19 becomes 3-5-2019, and 3-5-57 converts to 3-5-1957.</p>

Using Printable Summary to view system properties

System>Properties>Advanced>Printable Summary provides an HTML-formatted document containing a text summary of system property settings for an administrator to view or print. This summary contains the settings for the following system property groups:

- System, including General, Locale, Printing, Cluster Operations, Information Display, and Notification
- Usage and Error Logging

How to use Printable Summary to view system properties

The administrator can view this summary of system property settings by performing the following tasks:

- Log in to Configuration Console, choose Advanced view, then choose System.
- On System, choose Properties.

- 3 On Properties, choose Advanced.
- 4 Choose Printable Summary

An HTML-formatted document appears, displaying a detailed list of system properties, as shown in Figure 15-23.

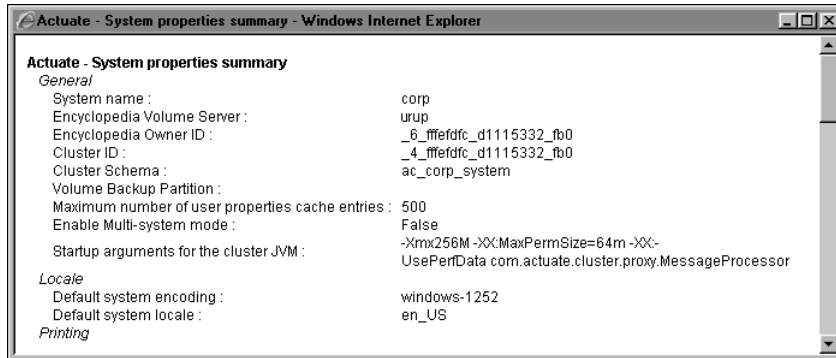


Figure 15-23 Viewing the system properties summary

Using Printable Summary to view template properties

Server Configuration Templates>Settings>Printable Summary provides an HTML-formatted document containing a text summary of template property settings for an administrator to view or print. This summary contains the settings for the following template property groups:

- Template, including Diagnostic Logging, Process Management, Queue Resource Management, and other miscellaneous properties
- View, Factory, Message Distribution, Caching, and Integration services

How to use Printable Summary to view template properties

The administrator can view this summary of template property settings by performing the following tasks:

- 1 Log in to Configuration Console, choose Advanced view, then choose Server Configuration Templates.
- 2 On Server Configuration Templates, choose a template.
- 3 On Settings, choose Printable Summary.

An HTML-formatted document appears, displaying a detailed list of the template and View, Factory, Message Distribution, Caching, and Integration service properties, as shown in Figure 15-24.

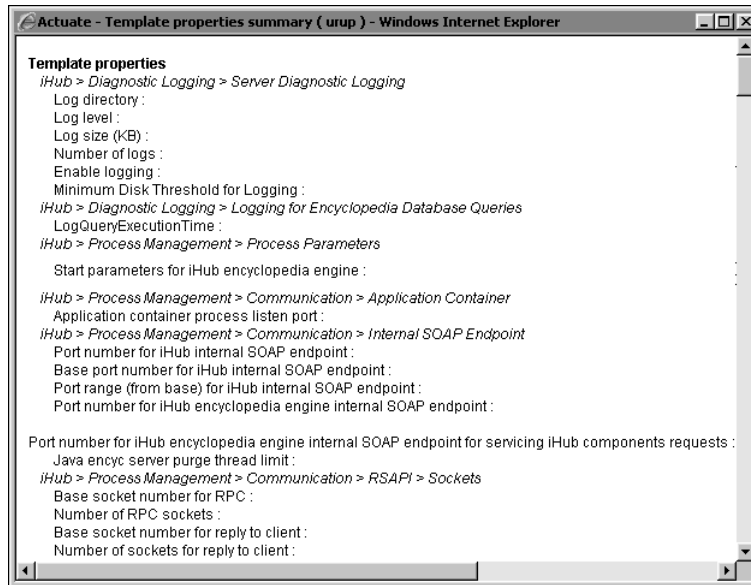


Figure 15-24 Viewing the template properties summary

Setting Filetype driver properties

The administrator specifies the following properties for a file-type driver:

- **Name**
Open Systems Interconnection (OSI) file-type driver name
- **OSI driver path**
Specifies the network link used to access the OSI driver
- **OSI parameter file location**
Specifies the location of the OSI driver parameter file used to configure network settings

How to set Filetype driver properties

To set Filetype driver properties, perform the following tasks:

- 1 Log in to Configuration Console, choose Advanced view, then choose Server Configuration Templates.
- 2 On Server Configuration Templates, choose a template.
- 3 On Settings, next to Filetype driver information, choose Add, as shown in Figure 15-25.

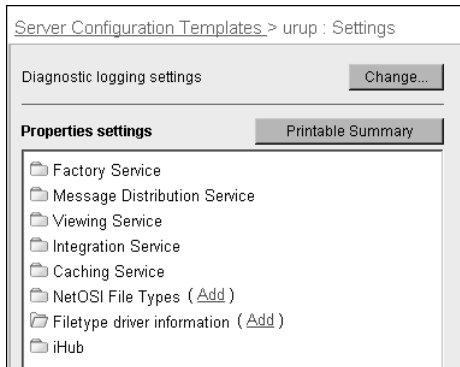


Figure 15-25 Choosing to add Filetype driver information

Figure 15-26 shows Add.

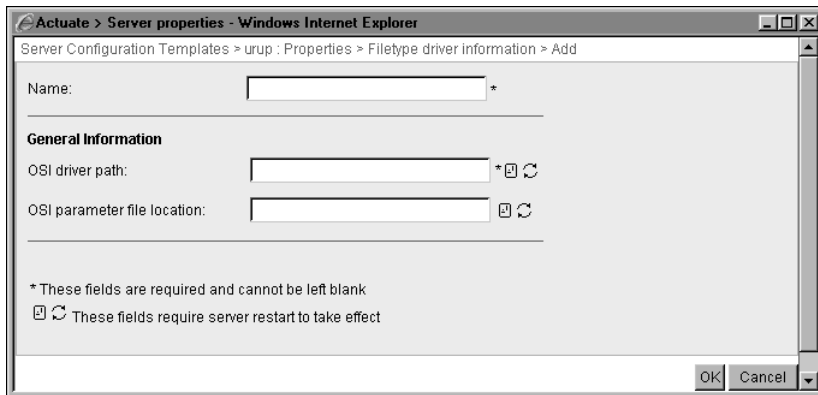


Figure 15-26 Specifying Filetype driver information

- 4 In Name, type the name of the OSI file-type driver.
 - 5 In OSI driver path, type the full path to the OSI driver.
- In OSI parameter file location, type the full path to the OSI driver parameter file.

Index

A

- About page (Options) 23
- AC_CONFIG_HOME variable 149, 150, 152
- AC_DATA_HOME directory 26
- AC_DATA_HOME variable 152
- AC_DBMS_INFORMIX_MAXVARLEN variable 249
- AC_DBMS_ODBC_MAXVARLEN variable 249
- AC_DBMS_ORACLE_MAXVARLEN variable 249, 254
- AC_DBMS_PROGRESS_MAXVARLEN variable 249
- AC_ICU_DATA variable 152
- AC_JRE_HOME variable 152
- AC_JRE64_HOME variable 152
- AC_ODBC_HOME variable 152
- AC_SERVER_FILE_CREATION_MASK variable 62
- AC_SERVER_HOME variable 254
- AC_SERVER_IP_ADDRESS variable 152
- AC_SOAP_DISPATCH_ADDRESS variable 152
- AC_TEMPLATE_NAME variable 150, 152, 161
- access control list buffer 222
- access control lists 220, 221, 225
- access permissions. *See* privileges
- accessing
 - archived files 232
 - configuration files 149, 150
 - custom data sources 254, 256
 - Encyclopedia volumes 26, 167, 194
 - external data sources 184, 194, 195, 246
 - iHub processes 93
 - information objects 184
 - in-memory archive cache 110, 127
 - Java classes 47
 - license files 167
 - log files 62, 63
 - online archive driver files 226
 - partitions 166
 - remote data sources 182
 - server templates 149
- accounts
 - accessing Oracle databases and 253
 - changing passwords for 13, 274
 - configuring administrative 84, 168
 - setting up e-mail 81
- ACL Buffer Size setting 222
- ACLs. *See* access control lists
- acmetadescription.xml 78
- acnotification.xml 85
- acpmdconfig.xml 149, 150
- AcProxy application 170, 171, 172
- ACS. *See* Caching service
- ACS database. *See* cache database
- acserverconfig.xml 149, 152, 273
- Active Directory servers 194
- active requests 22
- active volume partition 41
- Actuate Java Components 259
- Actuate version property 22
- Add Partition page 27, 166
- Add Printer page 235, 236, 237
- Add SMTP Server property 80
- adding
 - cluster nodes 153, 156, 161, 167
 - clusters 153–156
 - e-mail accounts 81
 - Encyclopedia volumes 26, 27, 35–38
 - external connection profiles 247
 - hyperlinks 88
 - metadata databases 30
 - metadata schemas 33
 - NetOSi interfaces 191
 - partitions 27–28
 - resource groups 94, 134, 137–140, 144
 - server templates 160
 - SMTP servers 76–78, 79, 80
 - system printers 234, 236
 - virtual IP addresses 146
- addresses (e-mail) 16, 81, 213
- admin logs 63, 68
- AdminEmail parameter 168

- administrative events 68
- Administrative page (Cluster Operation) 84, 168
- administrator e-mail account 84, 168
- Administrator role 209, 218, 228
- AdminPassword parameter 228
- AdminRole parameter 209, 218
- AdminUser parameter 228
- Advanced page (Server Templates) 174
- Advanced page (System Properties) 274, 275
- Advanced page (Volume Properties) 48
- Advanced view (Configuration Console) 2
- AFM files 240
- AIS. *See* Integration service
- alerts 16
- All security role 209, 218
- AllRole parameter 209, 218
- Apache Axis servers 195, 201
- Apache Tomcat servlet container 201
- AppContainerPort parameter 267
- application container 170, 195, 270
- Application Container page 270
- application container process listen port 267, 270
- Application context property 192
- application servers 14, 146, 172, 258, 277
 - See also* servers; web servers
- application-level partitioning 44
- applications
 - accessing external security sources and 194, 195
 - accessing iHub processes and 93
 - accessing information objects and 184
 - allocating disk space and 42
 - changing default ports and 14
 - configuring Encyclopedia volumes for 44
 - configuring web service 201, 202
 - hosting web services 270
 - installing RSSE 195, 201, 202
 - logging system information and 62, 63, 65
 - pooling connections and 258
 - running reports and 144
- archive applications 39
- archive cache 105, 110, 125, 127
- archive driver 224, 225, 229
- archive service 71, 226, 230
- archive service provider 266

- ArchiveRoot parameter 229
- ArchiveVolume parameter 228
- Archiving and Purging page 231
- archiving report files 224, 225, 232
- ASCII_CI collation 177
- asynchronous execution log entries 63
- asynchronous Factory processes 114
- Asynchronous Job Retries page 49
- asynchronous jobs 132
 - See also* scheduled jobs
- asynchronous resource groups 138, 142
- attachments 192, 213
 - See also* e-mail
- AttachReportInEmailAttr parameter 213
- AttachReportInEmailDefault parameter 213
- authentication 194
- authorization 207
- autoarchiving. *See* archiving
- Axis servers 195, 201

B

- background jobs. *See* scheduled jobs
- backward compatibility 73, 99, 114, 186
- bar charts. *See* histograms
- base port for Java factory server 267, 268
- base port for Java view server 267, 268
- base port for SOAP messaging 268, 271
- Base port number property 101, 122
- bcp utility (SQL) 190
 - See also* bulk load settings
- binary files 254
- BIRT 360 documents 102, 107
- BIRT 360 resource groups 132
- BIRT Content Caches properties
 - data set settings 107
 - design cache settings 109
 - document cache settings 110
 - image cache settings 108
 - in-memory archive settings 105
- BIRT Data Analyzer resource groups 132
- BIRT designs. *See* report designs
- BIRT documents. *See* report documents
- BIRT Factory resource groups 132
- BIRT Generation Caches properties
 - data sets 125
 - design cache 126

- document cache 127
 - in memory archive 125
- BIRT iHub. *See* iHub
- BIRT image cache 104, 108
- BIRT in-memory archive cache 105, 110, 127
- BIRT Online resource groups (default) 132
- BIRT Process Management properties 101, 121, 123
- BIRT reports 101, 121, 132, 242
 - See also* reports
- BIRT result set buffer 107, 124
- BIRT Studio resource groups 132
- BIRTChartMaxRows parameter 129
- BIRTChartMaxVariableSize parameter 129
- BIRTImageCacheTimeout parameter 109
- BIRTReportDesignCacheTimeout
 - parameter 110, 127
- BIRTReportDesignCacheTotalNumberOf
 - Entries parameter 110, 127
- BIRTReportDocumentCacheEnabled
 - parameter 111, 128
- BIRTReportPageCountCacheEnabled
 - parameter 111
- body element 87
- browsers. *See* web browsers
- buffer size 179, 180, 222
- buffers 107, 124
- Bulk Load property 189
- bulk load settings 186, 189
- bulk loading utilities 190
- BulkLoadPath parameter 190

C

- cache
 - accessing in-memory archive 110, 127
 - accessing remote data sources and 182
 - bulk loading files to 190
 - creating 178
 - disabling 127
 - managing 182
 - running View service and 104
 - setting transient document size 116
 - timing out. *See* cache time-out settings
- cache database 178, 182, 183
 - See also* Caching service
- cache definition files 182, 183, 184, 185, 191
- cache definitions, deleting 183
- cache entries (user properties) 275
- cache time-out settings
 - BIRT designs 109
 - BIRT document in-memory archive 106
 - BIRT images 108
 - open security data 203
 - transient documents 116
- caching
 - BIRT designs 104, 109, 110, 126
 - BIRT documents 104, 105, 107, 110, 127
 - BIRT report images 108
 - external security information 220
 - information objects 182, 183, 185, 246
 - ODA drivers 256
 - open security data 203
 - page-level security requests 204
 - session data 275
 - temporary files 105, 106, 116, 117
- Caching processes 187
- Caching property 124
- Caching server operations port 268
- Caching service
 - caching information objects and 182
 - configuring 183, 185
 - defining data sources for 246
 - described 94
 - enabling 186, 187
 - saving temporary files for 186
 - setting ports for 185, 188
 - starting or stopping 183
 - viewing diagnostic information for 12, 187
 - viewing status of 22
- Caching Service page 185
- Cannot connect to SMTP server errors 83
- CDATA sections (e-mail templates) 87, 89
- Century break property 281
- changing
 - archive driver configurations 226, 229
 - cluster configurations 162
 - cluster node configurations 160, 165
 - configuration files 22, 247
 - configuration home directory 149
 - decimal precision 178
 - Encyclopedia properties 38
 - Encyclopedia volume names 39
 - external connection profiles 247

- changing (*continued*)
 - Factory processes 136
 - file types 136
 - iHub names 14
 - IP addresses 14
 - licensing options 205
 - locales 264
 - log file configurations 11
 - metadata database properties 31
 - metadata schema properties 34
 - ODA driver configurations 256
 - partitions 40, 266
 - passwords 13, 31, 274
 - port numbers 14
 - property list items 20
 - resource group names 139
 - resource group properties 134
 - server templates 160, 162
 - SOAP settings properties 271
 - startup parameters 272–274
 - system names 274
 - time zones 264
 - volume names 53
- channel attribute 206, 207
- channel membership lists 206
- channels 206, 207, 211
- ChannelSubscriptionListAttr parameter 211
- ChannelSubscriptionListDefault parameter 212
- character encoding 86, 87
- character encoding. *See* encoding
- character sets 264
- character strings. *See* strings
- Chart page 128
- charts
 - mapping fonts for 241
 - reducing size of 128
 - rendering images for 108
- CIFS file systems 168
- CJK fonts 240
- classes 47
- Client bulk load path property 190
- client locale settings 251
- cloud configurations 150
- cluster nodes
 - See also* clusters
 - accessing Encyclopedia and 167
 - accessing resources and 166
 - adding 153, 156, 161, 167
 - assigning resource groups to 135, 139, 140, 141, 142
 - changing configurations for 160, 165
 - communicating with 271
 - configuring 147, 150, 154, 160, 161
 - deleting 165
 - displaying properties for 21, 22, 23
 - distributing requests to 142, 149, 170
 - instantiating 152
 - monitoring 156
 - removing 170
 - renaming 274
 - setting priority for 120
 - setting viewing weights for 100
 - starting 163, 164
 - stopping 163, 165, 277
 - taking offline 165
- cluster schemas 275
- clusters
 - See also* cluster nodes
 - accessing configurations for 149, 150
 - adding volume partitions and 44, 166
 - caching information objects and 183, 185
 - configuring 146, 162
 - creating 153–156
 - defining connections for 248
 - displaying status of 22
 - generating reports and 142
 - iHub licensing requirements for 167
 - installing console applications in 170
 - joining 153
 - load-balancing operations for 99, 114, 132
 - logging information for 63, 65
 - managing 149, 166
 - moving Encyclopedia and 44
 - running iHub processes and 132, 147
 - sending notifications over 79, 83, 84, 168
 - setting iHub ports for 14, 271
 - setting up printers for 235
 - shutting down 165
 - specifying regional settings for 264
 - starting 152
 - testing 162
- cn attribute 207

- code
 - e-mail body content 87, 89
 - online archive driver 224
 - RSSE applications and 194
- code pages 251
- collation 177
- command-line options
 - DB2 utility 252
 - Inline Archive Driver 39
 - online archive driver 226
 - open server drivers 192
- Common Internet File System. *See* CIFS file systems
- common name (cn) attribute 207
- Communication page 158, 175, 188
- completed jobs 16, 47, 132
 - See also* jobs
- completion notices 39, 40, 213, 214, 232
 - See also* notifications
- concurrent requests 107, 142, 159
- Configuration Console
 - archiving and 230, 231
 - changing display settings for 19–23
 - changing passwords for 13, 274
 - changing property lists for 21
 - cluster nodes and 156, 163, 164, 165
 - clusters and 165
 - configuring 169
 - data connections and 248, 260, 262
 - diagnostic information and 10, 11, 12, 56, 59, 61
 - distributing requests for 146
 - e-mail notifications and 15, 16, 80, 81, 84
 - Encyclopedia metadata and 29, 30, 31
 - Encyclopedia metadata schemas and 33, 34
 - Encyclopedia volumes and 35, 52, 265
 - error information and 69
 - file-type drivers and 283
 - iHub licenses and 6
 - iHub partitions and 27, 41, 54, 166
 - iHub services and 94, 98, 114, 174, 185
 - iHub system properties and 274, 281
 - iHub system status and 277
 - information object cache and 185
 - installing 170
 - logging in to 2
 - mail servers and 76, 78, 80
 - network settings and 14
 - Open Security applications and 202
 - overview 2
 - placing behind firewalls 169
 - release information and 277
 - resource group properties and 133
 - resource management properties and 279
 - RSSE security applications and 194
 - server properties and 268, 269
 - server templates and 157
 - starting 2, 3
 - startup parameters and 272
 - system printers and 234, 236, 238
 - template properties and 282
 - toggling between views in 17
 - usage information and 64
- configuration files
 - accessing 149, 150
 - adding resource groups and 144
 - allocating disk space and 42
 - caching information objects and 183, 184, 185, 190
 - changing online archive driver 226, 229
 - changing server templates and 160
 - creating e-mail messages and 14, 80, 81, 86, 87, 88
 - customizing 172
 - editing 22, 247
 - enabling iHub services and 99, 115
 - externalizing user information and 194, 205, 209
 - extracting from WAR files 171
 - handling file system failures and 168
 - installing drivers and 255, 256
 - installing fonts and 242
 - installing RSSE applications and 201, 202, 203
 - removing cluster nodes and 166
 - setting connections in 246–249, 257
 - setting environment variables and 150
 - setting location of 149
 - storing report documents and 105
 - updating 22
- configuration home directory 149
- configuration keys 247
- configuration lock files 168

- configuration parameters. *See* parameters
- configuration settings 149, 150, 151
 - See also* server templates
- configuring
 - administrative accounts 84
 - cluster nodes 150, 154, 160, 161
 - clusters 146, 162
 - Configuration Console 169
 - connections 248, 257
 - diagnostic logging 7, 10, 11, 12, 56–62
 - e-mail notifications 14, 16, 80, 81, 86
 - Encyclopedia volumes 29, 35, 48, 265
 - error logs 69–71, 72
 - event-based jobs 46–47
 - external security sources 194
 - iHub 92, 169
 - iHub services 94, 98, 114, 174, 185
 - J2EE application servers 259
 - JDBC connection pooling 32, 260
 - LDAP servers 205, 209
 - Management Console 169
 - metadata database schemas 33–35
 - metadata databases 30–32
 - online archive driver 224, 226, 229
 - resource groups 133–137
 - RSSE applications 201, 202
 - SMTP servers 15–16, 80
 - system printers 234, 235, 238, 239
 - usage logs 63–67, 72
- ConnConfigFile parameter 248
- connection configuration files 246–249
- connection definition files 184, 246
- connection information 149
- Connection login property 31
- Connection password property 31
- connection pool reap interval 262
- connection pools 32, 258–262
- connection profiles 247
- connection properties 246, 247
- ConnectionProperty parameter 208, 216
- ConnectionPropertyList parameter 208, 216
- connections
 - accessing custom data sources and 254, 256
 - accessing external data sources and 184, 246
 - accessing information objects and 183, 246
 - configuring 248, 257
 - losing 168
 - maintenance threads and 262
 - opening DB2 databases and 250–252
 - opening Informix databases and 252
 - opening JDBC data sources and 32, 257
 - opening ODBC data sources and 250
 - opening Oracle databases and 253–254
 - optimizing 258
 - overview 246, 250
 - running queries and 174, 179
 - sending notifications and 76
 - setting environment variables for 249
 - specifying type 257
 - timing out 178
- consolidating log information 63
- content-type attribute 88, 89
- Context string property 48
- Conversion Queue process 278, 279
- conversion threads 278, 279
- converting
 - dates 281
 - Encyclopedia volumes 218–220
 - PostScript fonts 241
 - user information 205, 218
- cookies 265
- coprocessors. *See* CPUs
- CopyDependOnFile parameter 229
- copying
 - online archive driver configurations 226
 - report files 44, 235
 - RSSE application files 201, 202
 - server templates 160
- cost-based optimization 179
- CPUs 120, 142
- CreateArchiveSubFolder parameter 229
- CreateUserRole parameter 229
- creating
 - access control lists 221
 - cluster nodes 153, 156, 161, 167
 - clusters 153–156
 - e-mail accounts 81
 - e-mail notification templates 85–90
 - e-mail notifications 76, 84
 - external connection profiles 247
 - hyperlinks 88
 - metadata databases 30

- metadata schemas 33
- partitions 27–28
- resource groups 94, 134, 137–140, 144
- server templates 160
- virtual IP addresses 146
- WAR files 172
- critical errors 69, 70
- Current requests property 22
- custom data sources 254, 256
- custom event service port 267
- custom events 48
- CustomEventServicePort parameter 267
- customizing
 - archive drivers 224
 - configuration files 172
 - e-mail messages 84, 86
 - load balancing proxy 171, 172
 - resource groups 142
 - RSSE applications 195, 220

D

- daemon listen port 267
- dashboards 132
- .dat files 26, 40
- data
 - restricting retrieval of 178, 182
 - retrieving 254, 255
 - sorting 177
- data access objects (DAOs) 258
- Data Analyzer resource group 132
- data buffer size 179, 180, 222
- data buffers 107, 124
- data cache. *See* cache
- data connection definition files 184, 246
- data files 26, 27, 40
- data integration logs 64
- data object stores 132
- data repository 226
- Data Set page 107, 125
- data sets 107, 124
- data source map files 185
- data sources
 - accessing external 184, 246
 - caching information objects and 246
 - connecting to 246, 249, 250, 254, 258
 - limiting data returned by 178, 182
 - retrieving data from 254, 255
 - running information objects and 256
 - specifying 176
- Data Tools Platform (Eclipse) 255
- data types 178, 250, 252
- database collation 177
- Database Connection Configuration File
 - page 248
- Database Connection Pool Manager Settings
 - page 32
- database drivers. *See* drivers
- Database name property 31
- Database port property 31
- database schema names 23, 34
- database schema owners 34, 35
- database schemas 29, 33, 185
- Database server property 30
- database servers 146
 - See also* servers
- Database superuser password property 34
- Database superuser property 34
- Database type property 30
- databases
 - caching data and 182
 - caching information objects and 184
 - configuring connections for 32, 248
 - configuring metadata 23, 26, 30–32
 - consolidating log information and 63
 - installing ODA drivers for 255
 - installing PostgreSQL 92
 - load-balancing and 246
 - setting environments for 249
 - setting up accounts for 253
 - sorting data and 177
 - storing Encyclopedia metadata and 29
- DataDirect JDBC drivers 256
- DataDirect ODBC drivers 246
- datamarts 105, 108, 111
- dates 73, 280
- Dates properties 280
- DateTimeAsString parameter 73
- DB2 databases
 - connecting to 250–252
 - getting connection information for 252
 - performing bulk loads for 186, 189, 190
 - setting up environments 249, 251
 - storing XML collections in 252

- DB2 utility 252
- DB2CODEPAGE variable 251
- DB2COMM variable 249
- DB2DIR variable 249, 251
- DB2INSTANCE variable 250, 251
- db2set command 251
- DBMS module 251
- dbname attribute 208
- dbpassword attribute 208
- .dcd files 184
- decimal precision 177, 178
- default collation 177
- default fonts 240, 241
- default installation settings (iHub) 92
- default locale 264
- default message size 76
- default notification template 86
- default port configurations 267
- default regional settings 264
- default resource groups 134, 143
- default time zone 264
- default volume directory 26
- default volume name 26, 53
- default volume partition 27
- DefaultPartition path 155
- deleted files, purging 39, 40, 266
- deleting
 - cache definition files 185
 - cache definitions 183
 - cluster nodes 165, 170
 - completion notices 39, 40, 232
 - Encyclopedia volumes 52, 143
 - NetOSi file types 192
 - partitions 27, 54
 - pending jobs 143
 - report files 40
 - resource groups 143
 - SMTP mail servers 80
- deletion events 65
- deletion logs 63, 65
- deletion notices 39, 40, 232
- dependencies. *See* file dependencies
- deploying
 - Java classes 47
 - load-balancing proxy 170, 171, 172
 - report designs 47
- Description property 22, 39
- design cache 104, 109, 110, 126
- Design Cache page 109, 126
- design file types 142
- designs 47, 114, 132, 242, 246
- Detail logging level 65
- dhtml mode 215, 216
- diagnostic information 7–13, 56–62
- diagnostic log files 9, 10, 62, 83, 84
- diagnostic log properties 10, 56, 59
- diagnostic logging categories 11, 12
- directories
 - accessing archived files in 232
 - accessing configuration files and 149
 - accessing Encyclopedia data and 26
 - accessing online archive driver files and 226
 - caching temporary files and 116
 - deploying JAR files to 47
 - installing ODA drivers and 255
 - installing RSSE applications and 195
 - logging diagnostic information and 9, 10
 - logging system information and 63
 - saving message templates and 84
 - setting privileges for 62
 - starting clusters and 152
 - storing temporary files and 180
- directory paths. *See* paths
- Directory property 9
- directory servers. *See* LDAP servers
- disabled volumes 50
- disk space
 - configuring for volume partitions 42
 - logging diagnostic information and 10
 - monitoring 42
 - running queries and 107, 181
 - setting 39, 42
 - specifying cache size and 116
- disk threshold property 180
- dispatcher service 228
- Display date time as string property 73
- display names 78
- DISPLAY variable 250
- display_value variable 235
- displaying
 - active requests 22
 - charts 241
 - dates 280

- Encyclopedia metadata 29
 - licensing options 5
 - primary volume partition 43
 - release information 22, 277
 - report documents 103, 105
 - reports 17, 98, 220
 - server templates 22
 - status information 22, 23, 277
 - system information 22
 - system property settings 281
 - template property settings 282
 - distinguished names (LDAP) 209, 211
 - distributed iHub System. *See* clusters
 - Distribution service. *See* Message Distribution service
 - DLC variable 250
 - DLLs 63, 65, 253, 256
 - document cache
 - BIRT reports 104, 110, 111, 127
 - transient files 116
 - Document Cache page 110, 127
 - document conversion tasks 278, 279
 - document deletion events 65
 - document generation events 63, 65, 67
 - document printing events 63, 65, 67
 - Document Type Definitions (DTDs) 252
 - document viewing events 63, 65, 68
 - documentation ix
 - documents
 - See also* PDF documents; reports
 - caching 104, 105, 107, 110, 127
 - configuring Java processes for 101, 122
 - configuring thread pool for 103, 104
 - configuring usage logging for 63
 - displaying 98, 103, 105
 - generating 107, 118, 247
 - locating fonts for 242
 - mapping fonts for 241
 - selecting resource groups for 132
 - sharing 110, 127
 - downloading
 - files 12, 93
 - PostScript drivers 236
 - driver properties 283
 - drivers
 - See also* specific type
 - accessing information objects and 256
 - adding resource groups and 142
 - archiving and 39, 224, 229
 - configuring JDBC data sources for 256
 - configuring ODA interfaces for 184, 254–256, 258
 - performing bulk loads and 190
 - printing reports and 236
 - running third-party 246
 - setting connections for 246
 - setting fetch buffer size for 179
 - setting location of 258
 - timing out 178
 - DTDs (Document Type Definitions) 252
 - DTP ODA API 255
 - Dummy To line property 82
 - DummyToLine parameter 82
 - dynamic hypertext markup language. *See* DHTML
 - dynamic link libraries 63, 65, 253, 256
- ## E
- e.reports 278
 - Eclipse Data Tools Platform 255
 - editing. *See* changing
 - e-mail
 - See also* notifications
 - addressing 16, 81
 - attaching files to 192, 213
 - customizing messages for 84, 86
 - externalizing user information and 213, 214
 - formatting content 87
 - sending 81
 - testing configurations for 81
 - e-mail accounts 81, 84, 168
 - email-content-type attribute 88, 89
 - e-mail error messages 83–84
 - e-mail message templates 14, 84, 85–90
 - E-mail notification template partition 85
 - E-mail Queue process 278, 279
 - EmailAddressAttr parameter 213
 - embedding fonts 240
 - Enable custom events property 48
 - Enable error logging property 69
 - Enable logging property 56
 - Enable SMTP property 76, 80

- Enable viewing service property 99
- EnableCachingService parameter 186, 190
- EnableEventService parameter 48
- EnableGenerationService parameter 115
- EnableIntegrationService parameter 174
- EnableViewingService parameter 99
- encoding 86, 87
- Encyclopedia engine 262, 271, 276
- Encyclopedia volume definition field 23
- Encyclopedia volume health errors 71
- Encyclopedia volume job purging errors 71
- Encyclopedia volume user activity errors 71
- Encyclopedia volumes
 - accessing 26, 167, 194
 - adding 26, 27, 35–38
 - allocating Factory processes for 135, 139
 - archiving items in 224, 225, 232
 - assigning partitions to 36, 39, 41
 - assigning to applications 44
 - authenticating users for 194
 - changing properties for 38, 265, 266
 - changing status of 49, 50
 - cluster configurations and 146, 167
 - configuring 29, 35, 265
 - deploying designs to 47
 - disabling and enabling 50, 51
 - disk locations for 27
 - displaying available partitions for 29, 43
 - displaying diagnostic information for 12
 - displaying metadata for 29
 - displaying properties for 21
 - displaying status of 23
 - expiring items in 231
 - externalizing user information for 44, 194
 - freeing disk space for 42
 - integrating multiple 44
 - logging in to 206
 - logging user operations for 63
 - mapping security information for 203, 206–208, 220
 - moving 43
 - naming 23, 26, 36, 39
 - removing files from 40
 - removing from system 50, 52, 143
 - removing resource groups and 143
 - renaming 53
 - RSSE applications and 194, 202, 203, 218–220
 - sending SOAP requests to 176
 - setting properties for 23, 48
 - setting up metadata database for 30–32
 - starting 37, 51
 - storing files in 105, 132, 185
 - taking offline 50
- engines 66
 - See also* Encyclopedia engine
- environment variables 150, 249
- error codes 71
- error information 62, 70, 84
- error log consolidator 63
- error log consolidator. *See* log consolidator
- error log entries 70
- Error log file name property 73
- Error log file size property 73
- error log files
 - accessing 63
 - configuring 69–71, 72
 - naming 73
 - sending notifications and 83
 - setting logging levels for 69
 - setting maximum size 73
 - specifying number of 73
- error logging applications 62
- error logging extension 62, 69, 70
- Error Logging page 69
- error logging properties 69, 72
- error messages 71
- error severity levels 70
- error_log.csv 62, 69
- ErrorLogFileName parameter 73
- ErrorLogFileSize parameter 73
- ErrorLoggingExt extension 69, 70
- errors
 - caching data and 178
 - logging levels and 8, 69
 - sending e-mail and 79, 83–84
- event lag time 47
- event service 48
- event service port 267
- event types 66, 67
- event-based jobs 46–47
- events 47, 65, 81
- Events page (Volumes Properties) 46, 47

- Exchange servers 83
- executable files 142, 204, 249
- extended viewing cache 204
- external authentication application 194
- external connection profiles 247
- external data sources 184, 246, 259
 - See also* external security sources
- external databases 182
- external registration application 194
- external security applications 194, 195, 220
- external security sources
 - accessing 194, 195
 - archiving and 225
 - assigning privileges for 207
 - associating with specific channels 207
 - caching information for 220
 - centralizing user information in 44
 - configuring 194
 - converting user information for 218, 219
 - creating notification groups for 207
 - integrating with RSSE applications 218
 - mapping security IDs to 220
 - mapping user information for 205
 - mapping volume information to 203, 206–208
 - running information objects and 208
- external security systems 194
- externalized Encyclopedia volumes 219
- externalizing user information 44, 205, 218, 225

F

- Factory processes
 - allocating 132, 134, 136
 - caching information objects and 183, 184
 - changing configuration files and 247
 - controlling 94, 132, 134
 - disabling resource groups and 134, 142
 - logging information for 63, 65
 - prioritizing jobs and 141
 - recycling 123
 - running ODA drivers and 256, 258
 - running open server service and 192
 - setting base port for 122
- factory server base port 267, 268
- factory server heartbeat port 268

- Factory service
 - assigning to cluster nodes 120
 - assigning to Encyclopedia 135, 139
 - assigning to resource groups 134, 136
 - configuring 114, 115, 121
 - described 94
 - displaying reports and 124, 142
 - enabling or disabling 115
 - generating charts and 128
 - generating reports and 132
 - generating temporary documents and 114, 116
 - running jobs and 118, 133, 134
 - sending requests to 123
 - setting properties for 114
 - viewing diagnostic information for 12, 115
 - viewing status of 22
- Factory Service page 114
- FactoryIdleTimeout parameter 134
- failed file systems 168
- failed jobs 49, 87, 133, 143
- failed notices 16, 82, 83, 89, 232
- failureMessage element 87, 89
- FailureNoticeExpirationAttr parameter 214
- FailureNoticeExpirationDefault parameter 214
- fatal errors 69, 70
- fetch operations 179
- file attributes 224
- file dependencies 225, 229
- file names 9
- file paths. *See* paths
- file system failures 168
- file types 184
- files
 - See also* specific types
 - accessing license 167
 - accessing volume metadata and 26
 - allocating disk space and 42
 - archiving 224, 225, 232
 - assigning privileges to 207
 - copying 44, 235
 - deleting 40
 - downloading 12, 93
 - locating archived 232
 - preserving dependencies for 225

- files (*continued*)
 - preserving owner information for 229
 - sending as attachments 192, 213
 - sharing 109, 110, 126, 127
 - storing 105, 132, 185
 - updating license 4, 6
- Filetype driver properties 283
- finding user names 215
- firewalls 14, 169
- flat file drivers 255
- folder names 232
- folders
 - See also* directories; paths
 - assigning privileges to 207
 - deploying JAR files to 47
 - mapping to LDAP servers 214
 - running online archive driver and 229, 232
- font file search paths 243
- font files 240, 242
- font information 240
- font styles 240
- FontName parameter 240
- Fontographer software 241
- fonts
 - configuring 240, 242
 - converting 241
 - embedding 240
 - generating PostScript output and 235
 - generating reports and 242
 - installing 235, 240
 - mapping for charts 241
 - missing 235
- format-specific font files 242
- formats
 - date values and 280
 - e-mail messages 87
 - log file entries 73
 - output 242
- fragmentation (partitions) 42
- free disk space 42
- functions 178

G

- GadgetGenerationTimeout parameter 103
- gadgets 102, 132
- garbage collection 136

- General Data Source Information page 176, 177
- General page (Options) 264
- General page (Resource Groups) 135
- General page (Servers Properties) 268
- General page (System Properties) 274
- General page (Volume Properties) 38, 39, 43, 265
- generating
 - gadgets 102
 - large documents 118
 - log files 10
 - query output 124
 - report designs 242
 - reports 107, 118, 142
 - system property summaries 281
 - template property summaries 282
 - temporary report documents 103, 114
- generation logs 63, 65, 67
- generation requests 107, 119, 142, 247
 - See also* jobs
- graphics. *See* images
- graphics environments 136
- graphs. *See* charts
- Greeting parameter 78
- GroupBaseDN parameter 209, 218
- GroupObject parameter 210, 218
- GroupToNotify parameter 210

H

- Headless graphics property 136
- heap 107, 136
- heartbeat failure period 157
- heartbeat messaging 156
- heartbeat port 268
- heartbeat send period 157
- home folder 206, 214
- HomeFolderAttr parameter 214
- hosting web services 270
- Hostname or IP Address property 15, 78
- hostname parameter 185
- hosts 246
- HP-UX servers 236, 249, 251
- HTML code 87, 89
- HTML documents 281, 282
- HTTP servers 149

hyperlinks 16, 88

See also URLs

hypertext markup language. *See* HTML code

hypertext transfer protocol. *See* HTTP servers

I

I/O. *See* input; output

I/O failures 168

IBM DB2 databases. *See* DB2 databases

IBM servers 236

.icd files. *See* cache definition files

ICD page 183

IDAPI applications 93

See also Information Delivery API

identifier IDs 218

iHub

adding Encyclopedia volumes to 35

adding volume partitions to 27

allocating disk space and 42

authorizing external users for 207

caching information objects and 183

changing passwords for 13

changing startup parameters for 272–274

changing system names and 274

configuring 92, 169

consolidating log information and 63

deploying designs to 47

disabling Encyclopedia and 50

displaying current state 277

displaying current status 22, 277

displaying properties for 268

displaying release information for 277

handling file system failures for 168

installing PostScript fonts on 235

installing to LDAP servers 196–201

localizing font information and 240

logging diagnostic information and 10, 60

logging error information for 69, 70, 73

optimizing performance for 246

placing behind firewalls 169

purging job notifications and 40

removing Encyclopedia volumes from 52

running queries and 180

running RSSE applications and 194, 195

sending notifications over 76, 79, 80, 83

setting font search paths for 243

setting ports for 14, 162, 266

setting properties for 22, 23, 268, 269

specifying regional settings for 264

starting or stopping 165, 277, 278

iHub 11 service 3, 92

iHub definition fields 22

iHub Diagnostic Logging property 59

iHub error messages 71

iHub node start or stop error type entries 71

iHub process port 268

iHub processes 14, 92

See also specific process

iHub services

See also specific service

configuring 94, 98, 114, 174, 185

disabling 148

distributing among cluster nodes 147, 165

enabling 147, 148

getting status of 22

running 92, 93, 154

setting ports for 162

starting 3, 92, 272

iHub System

changing passwords for 13

displaying information about 23, 277

fonts shipped with 240

generating property summaries for 281

incorrect character sets and 264

monitoring performance of 7

removing printers from 239

restarting 7

scaling 146

setting locales and time zones for 264, 265

setting properties for 274, 275

setting up printers for 234, 239

starting 6

stopping 6, 7

updating licenses for 4, 6

validating licenses for 162

iHub System errors 71

iHub System events 47

image cache 104, 108

images 108, 150

immediate jobs. *See* on-demand reports;

unscheduled jobs

in-memory archive 105, 125

- in-memory archive cache 110, 127
- In Memory Archive File Cache page 105, 125
- inactive volume partitions 54
- indexed searches 180
- Information Console
 - changing default ports and 14
 - distributing requests and 146
 - mapping external users and 207
 - setting port numbers for 14
 - setting up notifications for 80
 - viewing reports and 17
- Information Console URL prefix property 16, 40
- Information Delivery API
 - accessing processes and 93
 - adding resource groups and 144
 - running event-based jobs and 48
 - setting port numbers for 14, 271
- information object cache database. *See* cache database
- information object cache definition files. *See* cache definition files
- Information Object Designer 182
- information object projects 183
- information objects
 - accessing 184
 - allocating disk space for 181
 - caching 182, 183, 185, 246
 - configuring LDAP servers for 208, 216
 - retrieving data and 182
 - running jobs for 183
 - setting connections for 246, 256
 - setting pool size for 180
 - viewing diagnostic information for 12
- informational messages (logs) 8, 69, 70
- Informix databases 249, 250, 252
- INFORMIXDIR variable 250
- INFORMIXSERVER variable 250
- Inline Archive Driver application 39
- input 27, 168, 185
- input sources. *See* data sources
- insert element 87
- installation
 - console applications 170
 - custom fonts 240
 - ODA drivers 254–256
 - online archive driver 224
 - PostgreSQL database 92
 - PostScript fonts 235
 - RSSE applications 195, 201, 202
- INSTHOME variable 250
- Integration processes
 - handling requests for 176, 183
 - setting ports for 175, 267, 268
 - setting temporary file paths for 180
- Integration service
 - accessing 184
 - configuring 174
 - connecting to data sources and 174, 256
 - described 94
 - enabling or disabling 174
 - logging information for 64, 68
 - managing resources for 180
 - running information objects and 184, 256
 - running queries and 176
 - setting decimal precision for 178
 - setting string length for 178
 - starting or stopping 183
 - viewing diagnostic information for 12, 175
 - viewing status of 22
- Integration service driver 255
- Integration service events 64, 68
- Internal SOAP Endpoint page 271
- internationalization. *See* locales
- Internet Service Providers 201
- .job files 185
- IP address property 48
- IP addresses
 - cluster configurations 146
 - custom event service 48
 - iHub 14
 - open server service 192
 - RSSE web service 203
 - SMTP mail servers 78
- iServer cache definitions 182
- ISPs 201
- isValid function 260

J

- J2EE application servers 258
- JAR files 47, 194, 224
- Java classes 47
- Java Components (Actuate) 259

- Java documents 132
 - See also* Java reports
- Java factories 123, 134
- Java factory server 267, 268
- Java graphics environments 136
- Java processes 101, 122, 136
- Java reports 135, 136
- Java resource groups 137
- Java RSSE application 202
 - See also* RSSE applications
- Java Runtime Environment 135, 136
- Java server entry point property 136
- Java view process 93
- Java view server 136, 267, 268
- Java View service 101, 107, 142
- Java Virtual Machines. *See* JVMs
- Java web service 195, 196
- Java-based archive driver 224
- JavaProcessRecycleCount parameter 124
- JDBC connection pooling 32, 259, 260
- JDBC data sources 32, 257
- JDBC drivers 178, 190, 246, 256, 258
- jobCompletion variable 88
- jobDetailedStatus variable 88
- jobHeadline variable 88
- jobName variable 88
- jobs
 - assigning resource groups to 140–144
 - caching information objects and 183
 - checking status of 47
 - configuring event-based 46–47
 - configuring Factory processes for 117, 118, 134
 - deleting pending 143
 - deleting resource groups and 143
 - disabling 135, 142
 - failing 133, 143
 - generating reports and 114, 132
 - mapping to LDAP servers 215
 - printing reports and 237
 - prioritizing 135, 140, 141
 - retrying 49
 - running for designs 133
 - sending notices for. *See* notifications
 - setting polling interval for 47
 - setting queue size for 119
- jobStatus variable 88

- jobSubmitter variable 88
- jobType variable 88
- JRE software 135, 136
- JVM heap size 107, 124
- JVM start-up parameters 276
- JVMs 241, 275

K

- KeepTempFiles parameter 190
- KernPair variable 241

L

- Labels page
 - See also* e.Analysis reports
- Lag time property 47
- language codes 265
- language settings 253, 264
- language-specific reports. *See* locales
- large object data types 250
- large reports 118
- LD_LIBRARY_PATH variable 249
- LDAP directory structures 216
- LDAP security sources 194, 203
- LDAP servers
 - assigning privileges to 207
 - associating with specific channels 207
 - configuring 205, 209
 - creating notification groups for 207
 - installing iHub on 196–201
 - mapping volume information to 203, 206–208
 - running queries on 210
 - setting port numbers for 210
 - setting viewing preferences for 215, 216
 - specifying 211
 - storing user information on 44, 194, 205, 207
- libraries
 - configuring ODA drivers and 256
 - creating symbolic links to 251
 - logging error information and 62, 63
 - logging usage information and 62, 63, 65
 - running DB2 databases and 251
 - running Oracle databases and 253
- Library Path element 258
- license files 6, 167

- licensing information 149, 162
- licensing options
 - assigning 204, 205
 - running online archive driver and 224
 - updating iHub 4, 6
- licensing requirements 167
- Lightweight Directory Access Protocol. *See* LDAP
- linking to reports 16, 88
- Linux systems
 - adding printers for 236, 238
 - configuring connections for 248, 252
 - deploying AcProxy application for 172
 - installing iHub on 196
 - installing ODA drivers on 255
 - logging diagnostic information and 62
 - logging system information and 62
 - printing and 235, 236, 237
 - rendering charts for 241
 - running archive driver on 224
 - running iHub services on 162
 - sending notifications over 76, 81, 83
 - setting environment variables for 249, 253
 - starting iHub services on 3
 - testing connections for 252
 - testing e-mail configurations for 81
- listen port 15, 78, 268
- Listen port property 78
- lists 20
- load balancing
 - cluster configurations and 99, 114, 132, 147, 148, 170
 - data source connections and 246
 - on demand requests and 120
 - running jobs and 142, 144
 - sending notifications and 79
- load balancing application 170, 171, 172
- load balancing proxy 170
- load configuration properties 99
- LOB data types 250
- Locale parameter 265
- locales
 - changing 264
 - connecting to databases and 251, 253
 - displaying charts and 241
 - generating BIRT output and 242
 - installing fonts for 240, 241
- localization. *See* locales
- lock files 168
- log consolidator 63
- Log directory property 57
- log file names 9
- log files
 - accessing 63
 - changing configurations for 11
 - generating 10
 - naming 73
 - sending notifications and 83, 84
 - setting directory for 9, 58
 - setting number of 10, 73
 - setting privileges for 62
 - setting size 9, 57, 73
 - setting up diagnostic 10, 12, 56
 - setting up error 69–71
 - setting up system usage 63–67
- Log level property 56
- Log size property 57
- LogDirectory attribute 58
- logging
 - diagnostic information 7–13, 56–62
 - error information 69, 84
 - query statistics 181
 - usage information 63, 64
- logging applications 62, 63, 65
- logging extensions 62, 65, 69, 70
- Logging for Encyclopedia Database Queries
 - page 60, 61
- logging in to Configuration Console 2
- logging in to Encyclopedia volumes 206
- logging levels
 - diagnostic logs 7
 - error logs 69
 - online archive driver 229
 - query statistics 181
 - usage logs 65
- logging libraries 65
- login information 206, 212
- login page (iHub consoles) 264
- login scripts (Oracle tables) 253
- LogLevel parameter 229
- LogQueryExecutionTime parameter 61
- LogQueryPrepareTime parameter 61
- LogSQLQuery parameter 60

long report documents 116
 See also large reports
Low Free Space property 42

M

machine names 14, 22
mail. *See* e-mail
mail servers. *See* SMTP servers
MailingWeight parameter 78
Management Console
 archiving and 232
 changing default volume for 53
 configuring 169
 distributing requests for 146, 170
 information object cache and 183
 installing 170
 placing behind firewalls 169
 printing from 237, 239
 resource groups and 132, 140
 setting ports for 14, 162
map files 185
MAPI error handling 83
mapping LDAP security information 206–208
markup languages 87
Max factory property 134
MaxBIRTDataResult parameter 126
MaxBIRTDataResultsetBufferSize parameter 108
MaxConnections parameter 190
Maximum message size property 80
Maximum number of recipients property 81
Maximum number of threads property 103
MaxJobPriorityAttr parameter 214
MaxJobPriorityDefault parameter 215
MaxMailRecipients parameter 81
MaxMemoryPerArchive parameter 106, 126
MaxMemoryPerDatamartArchive parameter 106
MaxPermSize parameter 136
MaxSyncJobRuntime parameter 117, 120
MaxThreads parameter 190
MaxVarLen variable 254
memory
 caching image files and 108
 caching report files and 105, 106, 110, 127

 running information objects and 174, 181
 running queries and 178, 180
 setting heap size and 136
 setting thread batch size and 103
Message contains too many recipients errors 83
Message Distribution service
 configuring 157
 described 93
 enabling 158
 routing requests and 149
 sending notifications and 80
 setting port for 158, 267, 268
 viewing status of 22
message distribution service port 268
Message Distribution Weight property 99, 100, 114, 120
message routing 78, 149
 See also SOAP
message size 76, 80
 See also e-mail; notifications
message templates. *See* e-mail message templates
Message Timeout properties 102
Message too large errors 83
metadata 26, 29, 33
metadata database 23, 30–32
 See also Oracle databases; PostgreSQL databases
Metadata database name property 30
Metadata Database page 31
metadata database schema names 23, 34
metadata database schema properties 34
metadata database schemas 33–35
MetadataDatabaseProperties page 31
Microsoft Exchange e-mail accounts 83
Microsoft SQL databases. *See* SQL Server databases
Microsoft Windows. *See* Windows systems
Min factories property 134
Min Free Space for primary partition property 39
Min Free Space property 43
Minimum Disk Threshold for Logging property 60
missing fonts 235

Modify Start Parameters page 273, 274
moving Encyclopedia volumes 43
multicast messaging 149
multilingual reporting. *See* locales
multiprocessor systems 142
Multi-system mode 275
multi-valued attributes 206

N

Name property 22

names

- changing iHub 14
- changing resource group 139
- changing system 274
- changing volume 53
- finding user 215
- getting notification group 210
- getting role 211
- mapping user 220

naming

- database schemas 23
- Encyclopedia volumes 23, 26, 36, 39
- log files 73
- metadata database schemas 34
- metadata databases 23, 30
- resource groups 139
- server templates 150
- servers 22
- volume partitions 27

nesting security roles 207

NetOSi file type parameters 191

NetOSi file types 192

NetOSi File Types page 183, 185, 191

NetOSi interfaces 191, 192

Network connection fails errors 83

network file systems 168

network routers 146

network storage system failover 168

networks 14, 149, 153

New Metadata Database page 30

New Resource Group page 137, 139

New Schema page 34

New SMTP Server page 77

New Volume page 36

NFS-based file systems 168

NLS_LANG variable 253

nodes. *See* cluster nodes

non-retryable e-mail errors 83

notices. *See* notifications

notification group names 210

notification groups

- See also* notifications

- externalizing 206

- mapping LDAP objects to 218

- mapping to external security sources 218

- mapping to LDAP objects 206, 207, 210

Notification page (System Properties) 76, 80, 81

notification templates 14, 84, 85–90

notification threads 278, 279

notifications

- adding hyperlinks to 88

- allocating queue length for 278, 279

- configuring administrative accounts for 84, 168

- configuring SMTP servers for 14

- configuring UNIX systems for 81

- configuring URLs for 16

- configuring Windows systems for 80–81

- creating 76, 84

- entering HTML code for 89

- failing 16, 82, 83

- logging errors for 83, 84

- purging 39, 40, 231, 232, 266

- resending 82, 83

- sending over LDAP servers 213, 214

- sending over multiple servers 85

- sending over SMTP servers 15, 76

- setting expiration dates for 214, 215, 231, 232

- setting number of recipients for 81

- setting up 80, 81

notificationTemplate element 87

null-sensitive collation 177

Number of error log files property 73

Number of log files property 10, 57

Number of usage log files property 73

NumberOfErrorLogFiles parameter 73

NumberOfUsageLogFiles parameter 73

numbers 177, 178

NWPPort parameter 267

O

- object aging. *See* archiving
- ODA connection pools 258–259
- ODA data sources 247, 254
- ODA driver plug-in 255
- ODA drivers 179, 184, 254, 256, 258
- ODA interfaces 183, 184, 258
- odaconfig.xml 255
- ODBC connection pooling 259
- ODBC data sources 179, 249
- ODBC drivers 246
- On Demand Execution Queue page 119
- On Demand Report Execution Management
 - properties 114, 117, 119
- on-demand generation requests 107, 119, 142
- on-demand server processes 103
- on-demand viewing requests 102, 142
 - See also* unscheduled jobs
- OnDemandServerQueueSize parameter 104
- OnDemandServerViewMessageTimeout
 - parameter 103
- online archive driver 224, 225, 229
- online archive driver configurations 224, 226
- online archive driver files 226
- online documentation ix
- open data access data sources. *See* ODA data sources
- open data access drivers. *See* ODA drivers
- open database connectivity. *See* ODBC
- Open Security applications
 - See also* external security sources; RSSE applications
 - caching information for 220
 - centralizing user information in 44
 - configuring 202
 - converting Encyclopedia for 218–220
 - creating notification groups for 207
 - logging in to Encyclopedia and 206
 - mapping channel information and 207
 - mapping volume information for 206–208
 - nesting security roles in 207
 - testing 204
- Open Security cache 203
- Open Security functionality 194
- Open Security page 194, 202, 204
- open server drivers 192
- open server files 192
- open server service 192
- opening
 - report designs 109, 126
 - report documents 17, 110, 127
- Operating system and version property 22
- operating system information 22
- operating systems 22, 264
 - See also* UNIX systems; Windows systems
- Operator security role 210
- OperatorRole parameter 210
- optimizing
 - connections 258
 - job distribution 142
 - load operations 109, 110, 126, 127, 142
 - on-demand requests 107
 - performance. *See* performance
 - queries 179
- options (licensing). *See* licensing options
- Options link (Advanced view) 19
- Options page 265
- Oracle databases
 - connecting to 253–254
 - performing bulk loads for 186, 189, 190
 - setting column lengths for 254
 - setting environments for 249, 250, 253
 - storing Encyclopedia metadata and 29
- Oracle listener processes 254
- ORACLE_HOME variable 250, 253
- OSI driver path property 283
- OSI file types. *See* NetOSI file types
- OSI parameter file location property 283
- output
 - embedding fonts in 240
 - generating query 124
 - handling file system failures and 168
 - locating fonts for 242
 - printing 235, 240
 - processing 27
 - rendering charts and 241
- output formats 242
- overflow memory 180
- overriding system options 264
- owner information 229

P

- page count 104, 111
- page-level security
 - enabling 194
 - setting cache for 204
 - testing 204
 - viewing reports and 194, 220
- page-level security application 220–222
- parameters
 - adding mail servers and 78, 80
 - caching BIRT designs and 110, 127
 - caching BIRT images and 108
 - caching information objects and 190
 - caching reports and 106, 111, 128
 - caching temporary documents and 118
 - changing iHub startup 272–274
 - configuring Factory processes and 134
 - configuring Java processes and 102, 123
 - configuring on-demand server processes and 104, 108
 - configuring online archive driver and 228
 - controlling external cache and 220
 - enabling iHub services and 99, 115
 - logging diagnostic information and 60
 - logging error information and 73
 - logging usage information and 65, 73
 - mapping LDAP objects and 209, 217
 - recycling Java factories and 124
 - rendering charts and 129
 - running jobs and 120, 121
 - sending notifications and 81
 - setting iHub ports and 267
 - setting message time-out values and 103
 - setting result set buffer and 108, 125
 - starting Encyclopedia server JVMs 276
 - starting JRE and 136
 - starting RSSE service and 205
- Partition name property 27
- Partition Path property 28
- partitions
 - accessing 166
 - adding 27–28
 - assigning to Encyclopedia 36, 39, 41
 - changing 40, 266
 - configuring 27
 - deleting 27, 54
 - moving Encyclopedia among 43
 - naming 27
 - removing Encyclopedia volumes and 52
 - sending notifications and 85
 - setting disk space for 39, 42
 - setting paths for 28, 166
 - specifying primary 39
 - starting 36, 41
 - stopping 41
 - testing 28
 - updating paths for 43
 - viewing primary 43
 - viewing properties for 21
 - viewing status of 29
- Partitions page (Advanced view) 27, 29
- Partitions page (Options) 21, 23
- Partitions page (Volumes Properties) 40, 43
- pass-through security 208, 216
- passwords
 - accessing Encyclopedia and 31
 - accessing external security sources and 205, 211
 - changing 13, 31, 274
 - mapping to LDAP objects 206
 - running Configuration Console and 2
 - running information objects and 208
- PATH variable 249, 251, 253
- paths
 - connection configurations 248
 - database files 186
 - fonts and 243
 - system printers 235, 236, 238
 - temporary files 180
 - transient document cache 116
 - volume partitions 28, 43, 166
- pending jobs 141, 142, 143
 - See also* jobs
- performance
 - See also* performance tuning
 - adding resource groups and 142, 144
 - allocating Factory processes and 134
 - caching information objects and 185
 - defining connections and 246
 - generating reports and 118
 - logging usage information and 65
 - purging notifications and 231
 - running queries and 179

- setting thread batch size and 103
 - sharing report files and 109, 110, 126, 127
 - viewing reports and 99, 105
- performance tuning 48, 124
- permissions. *See* privileges
- PermSize property 136
- persistent documents 142
- PersistentArchiveFileCacheTimeout
 - parameter 107
- personal channels. *See* channels
- personal folder. *See* home folder
- PFA files 240
- plain text formats 87
- PMD. *See* Process Management Daemon
- PMDPort parameter 267
- Polling duration property 47
- Polling interval property 47
- pool. *See* connection pools
- port numbers 267
- Port parameter 210
- Port range for processes property 102, 123
- port usage restrictions 101, 122
- ports
 - application container listen 267, 270
 - Caching service 185, 188, 268
 - custom event service 48, 267
 - daemon listen 267
 - Factory processes 122
 - factory server heartbeat 268
 - iHub 162
 - iHub services 162
 - iHub System 266
 - integration server queries 267
 - Integration service 175, 268
 - internal SOAP messages 268
 - Java processes 101, 268
 - LDAP servers 210
 - Management Console 162
 - Message Distribution service 158, 268
 - metadata database 31
 - open server service 192
 - RSSE web service 203, 268
 - SMTP server 15, 78, 268
 - SOAP dispatcher service 228, 268
 - SOAP endpoint 268
- PostgreSQL databases 26, 29, 92
- PostgreSQL service 92
- PostgreSQL superuser names 34
- PostScript drivers 236
- PostScript files 235
- PostScript fonts 240, 241
- PostScript printer definitions 236
- PPD files 236
- preserving file dependencies 225, 229
- primary log directory 63
- Primary partition property 39
- primary volume partitions 27, 39, 43
 - See also* partitions
- print jobs 237
 - See also* printing
- print requests. *See* print jobs
- Printable Summary properties 281, 282
- printer information 239
- printer lists 239
- printer names 235
- printers
 - adding 234, 236
 - configuring 234, 235, 238, 239
 - disabling automatic configuration for 239
 - removing 239
 - sending output to 235
- Printers page (Configuration Console) 234
- printing
 - charts 241
 - Encyclopedia property lists 48
 - Java documents 132
 - reports 114, 234, 235, 237, 238
- printing events 63, 65, 67
- printing logs 63, 65
- Printing page 239
- prioritizing jobs 135, 140, 141
- privilege templates 206, 212
- privileges
 - accessing log files and 62
 - archiving and 225, 229
 - assigning to all users 209
 - assigning to external users 207
 - caching 109
 - mapping to LDAP objects 212
- PrivilegeTemplateAttr parameter 212
- PrivilegeTemplateDefault parameter 212
- process listen port 267, 270
- Process Management Daemon 150, 162

- Process Management properties (Caching service) 187
- Process Parameters page 187, 276
- processes. *See* iHub processes
- Processes page 123
- processors. *See* CPUs
- Progress databases 249, 250
- properties
 - archiving operations 224, 226
 - BIRT content caches 105, 107, 108, 109, 110, 127
 - BIRT Generation cache 125, 127
 - Caching service 183, 185, 186, 190
 - cluster nodes 156
 - data source connections 246, 259
 - database collation 177
 - diagnostic logs 10, 56, 59
 - disk space allocation 42
 - e-mail notifications 81
 - Encyclopedia server JVMs 276
 - Encyclopedia volumes 38, 48, 266
 - error logs 69, 72
 - external data sources 184
 - external user information 206
 - Factory service 114
 - file-type drivers 283
 - iHub server 268, 269
 - iHub services 94
 - iHub System 274, 275
 - Integration service 174
 - LDAP object mappings 206
 - Message Distribution service 157
 - metadata database 31
 - metadata schemas 34
 - Open Security applications 202
 - performance-tuning 48
 - previous releases and 99, 114
 - print operations 239
 - printing system 281
 - printing template 282
 - process management 101, 121, 123
 - purging operations 231
 - query management 178, 179
 - query statistics logs 181
 - resource groups 133, 134
 - resource management 278, 279
 - running jobs 47, 120, 121

- search operations 180
- SMTP mail servers 76, 77, 80
- usage logs 64, 72
- Viewing service 98, 99
- Properties page (Resource Groups) 135
- Properties page (System) 156, 274
- Properties page (Volumes) 38, 266
- property lists 20, 48
- property sheets. *See* Properties page; property lists
- Protocol library specification property 136
- proxy servers 170
- purging
 - deleted files 39, 40, 266
 - notifications 39, 40, 231, 232, 266

Q

- queries
 - accessing information objects and 183
 - accessing JDBC data sources and 32
 - generating output for 124
 - generating statistics for 181
 - logging diagnostic information for 60, 61
 - optimizing 179
 - retrieving data with 178, 210
 - retrying 49
 - running 174, 176, 178
 - setting collation type for 177
 - setting decimal precision and scale for 178
 - setting properties for 178
 - setting result set buffer for 107
 - sorting data for 177
- Query Execution page 178
- Query Optimization page 179
- Query Statistics page 181
- QueryAccount parameter 210
- QueryPassword parameter 211
- queue length allocations 278, 279
- queue size 103, 104, 119
- queue time-out settings 119

R

- RDBMS (relational database management systems) 29, 34
- recycling Java Factory 123
- redirector nodes 148

- Regional Settings page 264
- registering
 - external users 194
 - temporary reports 117
- registry keys 249
- relational database management systems. *See* RDBMS
- relational databases. *See* databases
- release information 22, 277
- removing. *See* deleting
- renaming
 - cluster nodes 274
 - Encyclopedia volumes 53
 - resource groups 139
- report design cache 104, 109, 110, 126
- report design file types 142
- report designs 47, 114, 132, 242, 246
- report document cache
 - BIRT reports 104, 110, 111, 127
 - transient files 116
- report documents
 - caching 104, 105, 107, 110, 127
 - configuring Java processes for 101, 122
 - configuring thread pool for 103, 104
 - configuring usage logging for 63
 - displaying 98, 103, 105
 - generating 107, 118, 247
 - mapping fonts for 241
 - selecting resource groups for 132
 - setting fonts for 242
 - sharing 110, 127
- report engines 66
- report executables. *See* report object executable files
- report files
 - See also* specific type
 - allocating disk space and 42
 - archiving 224, 225, 232
 - assigning privileges to 207
 - copying 44, 235
 - deleting 40
 - downloading 12, 93
 - locating archived 232
 - preserving dependencies for 225
 - preserving owner information for 229
 - sending as attachments 192, 213
 - sharing 109, 110, 126, 127
 - storing 105, 132, 185
- report generation logs 63, 65, 67
- report generation requests 107, 119, 247
 - See also* jobs; on-demand generation requests
- report object design files. *See* report designs
- report parameters. *See* parameters
- report server. *See* iHub
- report server API. *See* RSAPI applications
- Report Server Security Extension 194, 195
 - See also* RSSE applications
- reportDocumentName variable 88
- reportDocumentVersionName variable 88
- reportDocumentVersionNumber variable 88
- ReportEngineHeartBeatPort parameter 268
- reporting servers. *See* iHub; servers
- reporting system. *See* iHub System
- reportLink variable 88
- reports
 - controlling access to 194, 220
 - displaying 17, 98, 105, 220
 - generating 107, 118, 142
 - limiting requests for 123
 - opening 17
 - printing 114, 234, 235, 237, 238
 - retrieving access control lists for 220
 - saving 105
- repository 226
- repository. *See* Encyclopedia volumes
- Request Execution Thread Pool page 104
- Request Management property 159
- requests
 - See also* jobs
 - allocating iHub processes for 94, 176, 189
 - displaying reports and 100, 101
 - distributing 146, 148, 149, 170
 - generating documents and 119, 247
 - generating on-demand reports and 107, 119
 - handling concurrent 142
 - limiting number of 123
 - mapping to LDAP servers 214
 - running information object services and 183
 - setting listening port for 14
 - setting number of concurrent 159
 - setting number of connections for 176

- requests (*continued*)
 - setting number of threads for 176, 186
 - viewing number of active 22
- Requests page 176, 188
- resending e-mail notifications 82, 83
- resource group types 132, 139
- resource groups
 - adding 94, 134, 137–140, 144
 - allocating Factory processes for 134, 136
 - assigning to clusters 135, 139, 140, 141, 142
 - assigning to Encyclopedia 135, 143
 - configuring 133–137
 - customizing 142
 - deleting 143
 - disabling 134, 139, 142
 - mapping to user operations 132
 - naming 139
 - optimizing load balancing and 142
 - overview 132, 140
 - removing cluster nodes and 165
 - removing members from 143
 - renaming 139
 - running jobs and 119, 133, 134, 140–144
 - selecting 140
 - setting number of factories for 114
 - stopping 142
- Resource Groups page 133, 134
- Resource Management page 180
- Resource Management properties 278, 279
- Resource property 47
- resources
 - accessing 154
 - creating cluster nodes and 152, 162
 - deploying designs and 47
 - handling requests and 159
 - logging usage information for 62
 - running queries and 179, 180
- restarts 7
- result set buffer 107
- result sets 107, 124
- RetainOwner parameter 229
- RetainPermission parameter 229
- RetainTimestamp parameter 228
- retrying jobs 49
- rights and privileges. *See* privileges
- .roi files. *See* report object instance files
- RoleBaseDN parameter 211, 218
- RoleObject parameter 211, 218
- roles
 - archiving and 225, 229
 - assigning privileges with 207
 - associating security IDs with 194
 - externalizing 206
 - getting 211
 - mapping LDAP objects to 218
 - mapping to external security sources 218
 - mapping to LDAP objects 206, 207, 209
 - nesting 207
 - setting privileges for 212
 - validating 220
- routers 146
- RSSE applications
 - See also* Open Security applications
 - archiving and 225
 - configuring Encyclopedia for 194, 202, 203, 218–220
 - converting user information for 218
 - creating 194
 - customizing 195, 220
 - installing 195, 201, 202
 - mapping volume information for 206–208
 - retrieving access control lists for 220
 - running as web service 194, 195, 201–205
 - running on LDAP servers 205, 209
 - running sample 194, 220
- RSSE Cache Timeout parameter 220
- RSSE Max ACL Buffer Size parameter 222
- RSSE service 203, 220
- RSSE service port 268
- RSSE service startup parameters 204
- RSSESOAPPort parameter 268
- running
 - iHub services 154
 - information objects 184
 - jobs 114, 132, 133
 - online archive driver 225
 - queries 174, 176, 178
 - report designs 114, 134, 142, 246
 - RSSE sample applications 194, 220

S

sample RSSE applications 194, 220

- saving
 - message templates 84
 - reports 105
 - temporary files 186
- Schedule for deleted files property 39
- Schedule for purging deleted files property 40
- Schedule for purging notices property 39
- scheduled jobs 48, 49, 132, 143
 - See also* jobs
- schema database names 23, 34
- schema database owner names 34, 35
- Schema name property 34
- Schema owner name property 34
- schemas 29, 33, 185, 275
- search events 68
- search paths (fonts) 243
- searching 98, 180
 - See also* indexed searches
- secondary volume partitions 27, 54
 - See also* partitions
- security
 - accessing external sources for 194, 195
 - caching external source information for 220
 - caching information objects and 185
 - enabling page-level 194
 - mapping for information objects 208
 - opening reports and 17
 - running sample applications for 194, 220
 - storing information for 218
- security applications 194
 - See also* Open Security applications; page-level security application
- security IDs 194, 220
 - See also* access control lists
- security interface 194
- security roles
 - archiving and 225, 229
 - assigning privileges with 207
 - associating security IDs with 194
 - externalizing 206
 - getting 211
 - mapping LDAP objects to 218
 - mapping to external security sources 218
 - mapping to LDAP objects 206, 207, 209
 - nesting 207
 - setting privileges for 212
 - validating 220
- security sources. *See* external security sources
- SELECT statements. *See* queries
- Selected columns items (Options) 21
- SendEmailAttr parameter 213
- SendEmailDefault parameter 213
- Sender display name property 78
- Sender e-mail address property 16, 78
- Sender name property 16
- SenderAddress parameter 78
- SenderName parameter 78
- sending mail. *See* e-mail; notifications
- sendmail utility 81, 83
- SendNoticeAttr parameter 214
- SendNoticeDefault parameter 214
- server clusters. *See* clusters
- Server Configuration Templates page 59, 94, 157
- Server Diagnostic Logging property 60
- server names 14, 22, 274, 277
- server nodes. *See* cluster nodes
- Server parameter 211
- server ports 267
- server template names 150
- server templates
 - accessing 149
 - adding 160
 - changing 160, 162
 - defined 149
 - enabling iHub services in 147
 - removing cluster nodes and 165
 - setting parameters for 152
 - viewing 22
- SERVER_LIST parameter 172
- ServerFileSystemSetting parameter 155
- servers
 - See also* iHub
 - allocating Factory processes for 132, 134
 - assigning to resource groups 136, 139, 140
 - configuring as cluster nodes 150, 154, 160, 161
 - configuring as clusters 146, 162
 - configuring stand-alone 92, 132, 165, 274
 - deploying load-balancing proxy to 172
 - displaying information about 22, 277
 - displaying properties for 21, 22, 23

- servers (*continued*)
 - displaying status of 22
 - logging diagnostic information for 59
 - setting connection pools for 258
 - setting ports for 14
 - setting properties for 268, 269
 - setting queue size for 103, 104
 - setting thread pool for 103, 104
 - shutting down 165
- Servers column list (Options) 22
- Servers page (Configuration Console) 163
- Servers page (Options) 21
- ServerSOAPPort parameter 268
- ServerSOAPPortBase parameter 268
- ServersTemplates column list (Options) 22
- serverURL parameter 80
- service enable or disable error type entries 71
- Service page 204, 205
- service providers 201
- services. *See* iHub services; web services
- session cache 275
- Settings page (Server Configurations)
 - Caching Service properties 186
 - diagnostic logging 56, 59, 60
 - Factory Service properties 114
 - iHub properties 269, 276
 - Integration service properties 174
 - service property folders on 95
 - Viewing service properties 98, 99
- settings. *See* properties
- severe errors 69, 70
- shared connections 258
- shared files 109, 110, 126, 127
- shared libraries 63, 65
- SHLIB_PATH variable 249, 251
- shutting down. *See* stopping
- SIDs. *See* security IDs
- Simple Object Access Protocol. *See* SOAP
- Simple view (Configuration Console) 2, 11
- single-machine configurations. *See* stand-alone configurations
- Size property 9
- .sma files 185
- SMTP greeting property 15, 78
- SMTP listen port 15, 268
- SMTP mailing weight property 78
- SMTP servers
 - adding 76–78, 79, 80
 - changing setup for 80
 - configuring 15–16, 80
 - removing 80
 - sending notifications over 15, 76, 79, 83
 - setting up clusters for 80
- SMTPHostName parameter 78
- SMTPPort parameter 78, 268
- SOAP base port 268, 271
- SOAP dispatch port 268
- SOAP dispatcher service 228
- SOAP interface 92
- SOAP internal endpoint port 268, 271
- SOAP message timeout property 192
- SOAP messages 101, 122, 183, 192, 271
- SOAP port for caching 190, 268
- SOAP port for integration service 268
- SOAP port number property 192
- SOAP port parameters 48
- Soap port property 48
- SOAP requests. *See* requests
- SOAP-based archive application 39
- SOAP-based RSSE applications 194, 202
- SOAPDispatchSOAPPort parameter 268
- SOAPPort parameter 190, 268
- SocketBaseForJavaProcesses parameter 102, 123, 268
- SocketBaseForProcesses parameter 268
- SocketCountForJavaProcesses
 - parameter 102, 123
- Sockets page 101, 122
- sort operations 177
- source code
 - e-mail body content 87, 89
 - online archive driver 224
 - RSSE applications and 194
- spool command (UNIX) 235, 236
- SQL Server Bulk Copy Program 190
- SQL Server databases 186, 189, 190
- SQL statements. *See* queries
- stand-alone configurations 92, 132, 165, 274
- start arguments (JRE) 136
- Start Arguments property 135
- StartArguments parameter 190
- starting
 - application container process 270

- archive driver 229
 - Caching service processes 187
 - cluster nodes 163, 164
 - clusters 152
 - Configuration Console 2, 3
 - DB2 utility 252
 - Encyclopedia volumes 37, 51
 - iHub 277, 278
 - iHub services 3, 92, 272
 - iHub System 6
 - Inline Archive Driver application 39
 - Java resource groups 137
 - Java Runtime Environment 135, 136
 - Java Virtual Machines 276
 - RSSE service 204
 - volume partitions 36, 41
 - web services 205
 - StartKernPairs variable 241
 - startup parameters
 - Caching service 187
 - Encyclopedia engine 276
 - iHub 272–274
 - RSSE service 204
 - startup script (archive driver) 229, 230
 - stateful file systems 168
 - stateless file systems 168
 - status information 23, 277
 - Status page (System) 165, 277
 - Status property 22
 - stopping
 - application container process 270
 - cluster nodes 163, 165, 277
 - clusters 165
 - iHub 165, 277
 - iHub System 6, 7
 - resource groups 142
 - volume partitions 41
 - strings
 - defining configuration keys and 247
 - logging system information and 73
 - mapping volume information and 206
 - ordering 177
 - retrieving XML data and 252
 - running event-based jobs and 48
 - setting database column length and 254
 - setting default collation for 177
 - setting length 178
 - subject element 87
 - successMessage element 87, 89
 - SuccessNoticeExpirationAttr parameter 215
 - SuccessNoticeExpirationDefault parameter 215
 - Sun ONE LDAP servers. *See* LDAP servers
 - superusers 34
 - Sybase databases 250
 - SYBASE variable 250
 - SYBASE_OCS variable 250
 - symbolic links 251
 - SynchReportingWeight parameter 121
 - synchronous execution log entries 63
 - synchronous jobs 114, 118, 132
 - See also* unscheduled jobs
 - synchronous processes 114, 138
 - synchronous resource groups 119, 138, 143
 - SyncJobQueueSize parameter 120
 - SyncJobQueueWait parameter 117, 120
 - system database. *See* Oracle databases; PostgreSQL databases
 - System default volume property 53
 - System Heartbeat property 156
 - system names 274
 - system printers 234, 235, 236, 239
 - system resources. *See* resources
 - system variables. *See* environment variables
- ## T
- Target hostname property 192
 - TargetServer parameter 228
 - TargetSOAPPort parameter 228
 - Task queue size property 103, 104
 - Template Assignments page 135, 139, 140
 - template files 84, 85, 88
 - template properties 282
 - Template property 22
 - templates
 - adding HTML code to 89
 - configuring cluster nodes and 147, 149
 - creating e-mail messages and 85
 - e-mail. *See* e-mail message templates
 - server configurations. *See* server templates
 - specifying 273
 - temporary files
 - bulk load operations and 189

- temporary files (*continued*)
 - caching 105, 106, 116, 117
 - failing to register 117
 - generating 103, 114
 - saving 186
 - setting directory paths for 180
 - setting time-out values for 119
 - storing 105, 114, 132
 - viewing logging information for 63
- testing
 - clusters 162
 - DB2 connections 252
 - e-mail configurations 81
 - Open Security applications 204
 - startup parameters 273
 - volume partitions 28
- text 241
- text strings. *See* strings
- third-party drivers 246
- third-party processes 250
- third-party security systems 194
- third-party services 192
- thread batch size 103
- thread pool 103, 104
- threads 278, 279
- throughput tests 159
- time formats 73
- Time zone property 265
- time zones 264
- time-out intervals
 - design cache 109, 126
 - document cache 105, 106
 - external security cache 220
 - gadget generation 102
 - image cache 108
 - JDBC drivers 178
 - Open Security cache 203
 - temporary reports 119
 - transient document cache 116
 - viewing operations 102, 103
- Tomcat servers 201
 - See also* Apache Axis servers
- Tomcat servlet container 201
- TotalArchiveMemory parameter 106
- transient document cache 116, 117
- transient files. *See* temporary files
- Transient Report Management page 117

- TransientArchiveFileCacheTimeout
 - parameter 106
- TransientReportCacheLocation
 - parameter 118
- TransientReportCacheSize parameter 118
- TransientReportTimeOut parameter 118
- TransientStoreMaxCacheEntries
 - parameter 118
- translations. *See* locales
- troubleshooting iHub system 7, 9
- TrueType fonts 240, 241
- truncation errors 178

U

- uid attribute 206
- UNICODE_BIN collation 177
- universal hyperlinks. *See* hyperlinks
- UNIX language code variable 265
- UNIX systems
 - adding printers for 236, 238
 - configuring database connections for 248, 250, 252
 - deploying AcProxy application for 172
 - installing fonts for 240
 - installing iHub on 196
 - logging diagnostic information and 62
 - logging system information and 62
 - printing and 235, 236, 237
 - rendering charts for 241
 - running archive driver on 224, 226
 - running iHub services on 162
 - sending notifications over 76, 81, 83
 - setting environment variables for 249, 251, 253
 - starting Configuration Console on 2, 3
 - starting iHub services on 3
 - testing connections for 252
 - testing e-mail configurations for 81
- Unrecognized recipient errors 83
- unscheduled jobs 118, 132, 144
- unscheduled reports 132
- unsent e-mail 85
- Update Information Console URL
 - property 17
- updating
 - cache definition files 183

- configuration files 22
- license files 4, 6
- partition paths 43
- upgrades 204
- URLs
 - connection configurations 248
 - customizing AcProxy application and 172
 - linking to reports and 17, 80
 - running event-based jobs and 48
 - sending notifications and 16, 40, 80, 89
 - starting Configuration Console and 2
 - testing clusters and 162
- Usage And Error Log File Settings page 72, 73
- Usage and error log version property 73
- usage log consolidator 63
- usage log consolidator. *See* log consolidator
- Usage log file name property 73
- Usage log file size property 73
- usage log files
 - accessing 63
 - configuring 63–67, 72
 - naming 73
 - setting logging levels for 65
 - setting maximum size 73
 - specifying number of 73
 - viewing entries in 66
- usage logging applications 62
- usage logging extension 62, 65
- Usage Logging page 62, 64
- usage logging properties 64, 72
- usage monitoring information 62
- usage_log.csv 62
- UsageAndErrorLogVersion parameter 73
- UsageLogFileName parameter 73
- UsageLogFileSize parameter 73
- Use archive service property 39
- Use dummy line property 82
- UseDelimitedIdentifier property 252
- UseDummyToLine parameter 82
- user ID attribute 206
- user IDs 220
- user names 208, 215, 221
- user operations 63, 132
- user permission 212
- user sessions 275
- UserActivityLoggingExt application 65
- UserBaseDN parameter 211, 217
- UserObject parameter 215, 217
- users
 - authenticating 194
 - externalizing information for 44, 205, 207
 - mapping LDAP objects to 217
 - mapping security IDs for 220
 - mapping to external security sources 218
 - mapping to LDAP objects 206
 - registering external 194
 - retrieving access control lists for 220
- UsrActivityLoggingExt application 65

V

- variable attribute 88
- variables. *See* environment variables
- version attribute 88
- version information 22, 88, 277
- View process cache for executables
 - property 204
- View processes 85
- view server base port 267, 268
- View service
 - configuring 98, 99, 101
 - described 93
 - displaying reports and 99, 101, 142
 - enabling or disabling 99
 - generating temporary reports and 103, 114, 132
 - logging diagnostic information for 12, 99
 - logging usage information for 63, 65
 - sending notifications and 76, 79, 85
 - setting cache type for 104
 - setting message time-outs for 102, 103
 - viewing status of 22
- View service processes 63, 84
- viewing
 - active requests 22
 - charts 241
 - dates 280
 - Encyclopedia metadata 29
 - licensing options 5
 - primary volume partition 43
 - release information 22, 277
 - report documents 103, 105
 - reports 17, 98, 220

- viewing (*continued*)
 - server templates 22
 - status information 22, 23, 277
 - system information 22
 - system property settings 281
 - template property settings 282
- viewing events 63, 65, 68
- viewing log files 63, 65
- viewing preferences 215, 216
- Viewing service. *See* View service
- Viewing service page 99
- viewing weight property 99
- ViewPreferenceAttr parameter 215
- ViewPreferenceDefault parameter 216
- virtual IP addresses 146, 148
- volume cache time-out property 204
- volume data files 27
- volume metadata 26, 29, 33
- volume metadata database 30
- volume metadata schema properties 34
- volume metadata schemas 33–35
- Volume name property 39
- volume names 23, 26, 36, 39, 53
 - See also* Encyclopedia volumes
- volume online or offline error type entries 71
- volume partitions. *See* partitions
- Volumes page (Advanced view) 35, 38
- Volumes page (Options) 21, 23

W

- Wait for response times out errors 83
- wait intervals. *See* time-out intervals
- WAR files 171, 172
- warnings 8, 69, 70
- web browsers
 - displaying reports in 17
 - setting security levels for 17
 - setting title bar for 111
 - starting Configuration Console from 2
- web servers 14
 - See also* servers
- Web service API port 14, 271
- web service applications
 - configuring 201, 202

- hosting 270
- installing 195, 201, 202
- integrating with Encyclopedia 202
- running 201
- web services
 - running open security applications and 196
 - running RSSE applications and 194, 195
 - setting IP addresses for 203
 - setting location of 203
 - setting port number for 203
 - starting 205
- Windows systems
 - adding printers for 238
 - configuring database connections for 248, 250, 252
 - configuring network ports and 14
 - deploying AcProxy application for 171
 - installing iHub on 196
 - installing ODA drivers on 255
 - installing RSSE applications on 201
 - logging diagnostic information and 62
 - logging system information and 62
 - printing on 234, 236
 - rendering charts for 241
 - running archive driver on 224
 - running iHub services on 162
 - sending notifications over 83
 - setting environment variables for 249, 251, 253
 - starting Configuration Console on 2, 3
 - starting iHub services on 3
 - testing connections for 252

X

- X Windows servers 250
- XML documents 252
- XML Extender 252
- Xprinter environment variables 236
- XVFBDISPLAY variable 235

Y

- year conversions 281