



ACTUATE.  
The BIRT Company™



BIRT iHub

# System Administration Guide

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 - 2014 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](http://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 Content Services, BIRT Data Analyzer, BIRT for Statements, BIRT iHub, BIRT Metrics Management, 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](http://www.zlib.net)): zlib. Adobe Systems Incorporated: Flash Player, Source Sans Pro font.  
Amazon Web Services, Incorporated: Amazon Web Services SDK. Apache Software Foundation ([www.apache.org](http://www.apache.org)): Ant, Axis, Axis2, Batik, Batik SVG library, Commons Command Line Interface (CLI), Commons Codec, Commons Lang, Commons Math, Crimson, Derby, Hive driver for Hadoop, Kafka, log4j, Pluto, POI ooxml and ooxml-schema, Portlet, Shindig, Struts, Thrift, Tomcat, Velocity, Xalan, Xerces, Xerces2 Java Parser, Xerces-C++ XML Parser, and XML Beans. Daniel Bruce ([www.entypo.com](http://www.entypo.com)): Entypo Pictogram Suite. Castor ([www.castor.org](http://www.castor.org)). ExoLab Project ([www.exolab.org](http://www.exolab.org)), and Intalio, Inc. ([www.intalio.org](http://www.intalio.org)): Castor. Alessandro Colantonio: CONCISE. Day Management AG: Content Repository for Java. Eclipse Foundation, Inc. ([www.eclipse.org](http://www.eclipse.org)): Babel, Data Tools Platform (DTP) ODA, Eclipse SDK, Graphics Editor Framework (GEF), Eclipse Modeling Framework (EMF), Jetty, and Eclipse Web Tools Platform (WTP). Dave Gandy: Font Awesome. Gargoyle Software Inc.: HtmlUnit. GNU Project: GNU Regular Expression. Groovy project ([groovy.codehaus.org](http://groovy.codehaus.org)): Groovy. Guava Libraries: Google Guava. HighSlide: HighCharts. [headjs.com](http://headjs.com): head.js. Hector Project: Cassandra Thrift, Hector. Jason Hsueth and Kenton Varda ([code.google.com](http://code.google.com)): Protocole Buffer. H2 Database: H2 database. Groovy project ([groovy.codehaus.org](http://groovy.codehaus.org)): Groovy. IDAutomation.com, Inc.: IDAutomation. IDR solutions Ltd.: JBIG2. InfoSoft Global (P) Ltd.: FusionCharts, FusionMaps, FusionWidgets, PowerCharts. Matt Inger ([sourceforge.net](http://sourceforge.net)): Ant-Contrib. Matt Ingenthron, Eric D. Lambert, and Dustin Sallings ([code.google.com](http://code.google.com)): Spymemcached. International Components for Unicode (ICU): ICU library. JCraft, Inc.: JSch. jQuery: jQuery. Yuri Kanivets ([code.google.com](http://code.google.com)): Android Wheel gadget. LEAD Technologies, Inc.: LEADTOOLS. The Legion of the Bouncy Castle: Bouncy Castle Crypto APIs. Bruno Lowagie and Paulo Soares: iText. MetaStuff: dom4j. Microsoft Corporation (Microsoft Developer Network): CompoundDocument Library. Mozilla: Mozilla XML Parser. MySQL Americas, Inc.: MySQL Connector. Netscape Communications Corporation, Inc.: Rhino. nullsoft project: Nullsoft Scriptable Install System. OOPS Consultancy: XMLTask. OpenSSL Project: OpenSSL. Oracle Corporation: Berkeley DB, Java Advanced Imaging, JAXB, JDK, Jstl, Oracle JDBC driver. PostgreSQL Global Development Group: pgAdmin, PostgreSQL, PostgreSQL JDBC driver. Progress Software Corporation: DataDirect Connect XE for JDBC Salesforce, DataDirect JDBC, DataDirect ODBC. Quality Open Software: Simple Logging Facade for Java (SLF4J), SLF4J API and NOP. Rogue Wave Software, Inc.: Rogue Wave Library SourcePro Core, tools.h++. Sam Stephenson ([prototype.conio.net](http://prototype.conio.net)): prototype.js. Sencha Inc.: Ext JS, Sencha Touch. Shibboleth Consortium: OpenSAML, Shibboleth Identity Provider. Matteo Spinelli: iscroll. StAX Project ([stax.codehaus.org](http://stax.codehaus.org)): Streaming API for XML (StAX). SWFObject Project ([code.google.com](http://code.google.com)): SWFObject. ThimbleWare, Inc.: JMemcached. Twitter: Twitter Bootstrap. VMWare: Hyperic SIGAR. Woodstox Project ([woodstox.codehaus.org](http://woodstox.codehaus.org)): Woodstox Fast XML processor (wstx-asl). World Wide Web Consortium (W3C)(MIT, ERCIM, Keio): Flute, JTIty, Simple API for CSS. XFree86 Project, Inc.: ([www.xfree86.org](http://www.xfree86.org)): xvfb. ZXing Project ([code.google.com](http://code.google.com)): ZXing.

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

Document No. 131215-2-530303 October 26, 2014

# Contents

<b>About <i>BIRT iHub System Administration Guide</i> .....</b>	<b>v</b>
---	----------

## Part 1

### **Architecture and configuration**

#### Chapter 1

<b>Understanding Actuate BIRT iHub architecture .....</b>	<b>3</b>
---	----------

Understanding BIRT iHub architecture .....	4
--	---

Using a third-party RDBMS with BIRT iHub system .....	4
---	---

About the cluster and volume schemas .....	4
--	---

Managing backup, recovery, and failover .....	5
---	---

Understanding the BIRT iHub process model .....	5
---	---

Understanding process flow in a BIRT iHub system .....	6
--	---

iHub security overview .....	10
------------------------------	----

Understanding SAML .....	11
--------------------------	----

About the BIRT iHub SAML implementation .....	11
---	----

Understanding SSL .....	12
-------------------------	----

Understanding SSO .....	13
-------------------------	----

Administering BIRT iHub System .....	13
--------------------------------------	----

Using JDBC to connect to the BIRT iHub database .....	15
---	----

About international character sets .....	15
--	----

Administrative reports .....	15
------------------------------	----

Supported operating systems .....	16
-----------------------------------	----

#### Chapter 2

<b>Configuring a BIRT iHub cluster .....</b>	<b>17</b>
--	-----------

Installing a BIRT iHub cluster node .....	18
---	----

Preparing the BIRT iHub cluster environment .....	18
---	----

Creating the shared configuration directory .....	20
---	----

Sharing the folders that all cluster nodes access .....	20
---	----

Configuring two nodes to communicate with each other .....	24
--	----

Specifying a logon account for the Actuate iHub 3 service on a cluster node .....	27
---	----

#### Chapter 3

<b>Configuring BIRT iHub to use an alternative database .....</b>	<b>31</b>
---	-----------

#### Chapter 4

<b>Configuring BIRT iHub security .....</b>	<b>33</b>
---	-----------

Securing data in a BIRT iHub volume .....	34
---	----

BIRT iHub secure communications .....	34
Understanding HTTPS .....	35
Understanding SSL .....	35
BIRT iHub SSL communication process .....	36
BIRT iHub SSL support .....	36
Understanding SAML .....	37
SAML communication process .....	37
BIRT iHub SAML support .....	37
Using SSL .....	38
Using SSL with IDAPI .....	38
Using SSL with JSAPI .....	39
Using SSL and external user management .....	40
Using Visualization Platform with SSL .....	40
Using SSL for communication with the volume metadata database .....	49
Managing SSL files .....	52
Using a commercial SSL certificate .....	54

## Part 2

# BIRT iHub System Console

## Chapter 5

### **Understanding System Console .....** **1**

About System Console .....	2
Viewing clusters, nodes, and system administrators .....	3
Logging in to System Console .....	3
Using the System Console configuration user interface .....	4
About Monitoring .....	8

## Chapter 6

### **Managing clusters .....** **11**

About clusters .....	12
Creating and configuring a cluster .....	13
About the cluster configuration categories .....	14
Adding cluster nodes to a cluster .....	15
Understanding Cluster Configuration .....	26
Performing management tasks for the entire cluster .....	26
Performing management tasks for an individual cluster node .....	27
Performing management tasks for a service running on a cluster node .....	28
Adding a volume .....	31
Adding or updating a storage location .....	35
Understanding the volume menu .....	36
Selecting the metadata database type .....	37
Configuring alerts .....	40

Viewing the list of alerts .....	40
Adding an alert .....	41
Enabling e-mail notification .....	43
Editing, deleting, disabling, and enabling an alert .....	45
Configuring Single Sign-On .....	47
Configuring User Management .....	49
Configuring LDAP Adapter .....	49
Configuring RSSE SOAP Service .....	60
Updating the license .....	61
About Configuration File .....	61
Editing an existing cluster .....	63
Managing a cluster node .....	64
Viewing the list of clusters .....	66
Filtering the list of clusters using a search string .....	66
Filtering non-starred clusters .....	68
Deleting clusters .....	69
Viewing cluster resources .....	71
Viewing diagnostic log files using Logs .....	72
Viewing system-level information using Trends .....	74
About the Actuate Viewer menu .....	76
Viewing Current Activity .....	78
Using the ControlConsole utility to free memory .....	81
About BIRT iHub service and resource group properties .....	84
Reporting service template properties .....	85
Viewing service template properties .....	86
Integration service Template properties .....	88
Understanding resource groups .....	88
Understanding resource group properties pertaining to a template .....	89
Understanding resource group properties pertaining to all templates .....	93
Configuring data source connections in BIRT iHub .....	94
Configuring an Apache web server for load balancing and proxying .....	95
<b>Chapter 7</b>	
<b>Managing system administrators .....</b>	<b>99</b>
About Settings .....	100
Viewing System Information .....	100
Working with system administrators .....	101
Creating a system administrator .....	101
Customizing the notification e-mail template file .....	103
Editing a system administrator .....	104
Deleting system administrators .....	104
Configuring Email Settings .....	106

## Part 3

# Managing and backing up

## Chapter 8

<b>Licensing BIRT iHub</b> .....	<b>111</b>
Understanding licensing types .....	112
Understanding licensing options .....	113
Installing BIRT iHub System license files .....	114
About the license file .....	115
Collecting machine information for a license .....	116
About modifying a license .....	119
About modifying the data collection option .....	119
Understanding CPU binding .....	119
Configuring CPU binding on Windows .....	120
Binding to specific CPU cores .....	120
Binding to multiple-core CPU cores .....	121
Binding an Actuate process to a processor .....	121
About processors and hyperthreading .....	122
Configuring CPU binding on Linux .....	123
Checking BIRT iHub bound processors .....	124
Determining the number of processors BIRT iHub System uses .....	124
Understanding CPU binding validation while BIRT iHub is running .....	125
Understanding CPU binding validation when a volume comes online .....	126
Understanding CPU binding validation when running BIRT iHub processes .....	126
Configuring e-mail for CPU license problems .....	126

## Chapter 9

<b>Backing up BIRT iHub System</b> .....	<b>129</b>
Performing a BIRT iHub System backup .....	130
Managing the backup and recovery of BIRT iHub metadata and data files .....	130
Using RDBMS and file system backup utilities .....	131
Backing up and restoring a BIRT iHub System that uses a PostgreSQL database .....	132
Backing up BIRT iHub System using pg_dump .....	133
Restoring BIRT iHub System using pg_restore .....	135
<b>Index</b> .....	<b>137</b>

# A b o u t B I R T i H u b S y s t e m A d m i n i s t r a t i o n G u i d e

---

*BIRT iHub System Administration Guide* includes the following chapters:

- *About BIRT iHub System Administration Guide..* Provides an overview of this guide, Actuate BIRT iServer documentation, and the typographical conventions in this book.
- *Part 1. Architecture and configuration.* Describes BIRT iHub architecture and customizing the set-up of BIRT iHub.
- *Chapter 1. Understanding Actuate BIRT iHub architecture.* Describes BIRT iHub architecture, the iHub System process model, and system administration, including new utilities and third-party relational database management systems (RDBMS) used to store iHub system and volume metadata.
- *Chapter 2. Configuring a BIRT iHub cluster.* Describes how to set up a BIRT iHub cluster node in a Windows environment.
- *Chapter 3. Configuring BIRT iHub to use an alternative database.* If you want to configure BIRT iHub to use a pre-existing PostgreSQL or Oracle database instead of the out-of-the-box PostgreSQL database, contact Actuate Support.
- *Chapter 4. Configuring BIRT iHub security.* Explains how to secure data in a BIRT iHub volume, BIRT iHub secure communications, and how to use SSL.
- *Part 2. BIRT iHub System Console.* Describes how to use BIRT iHub System Console.
- *Chapter 5. Understanding System Console.* Introduces BIRT iHub System Console.
- *Chapter 6. Managing clusters.* Describes creating, configuring, and managing a BIRT iHub cluster. Describes BIRT iHub cluster monitoring features.

- *Chapter 7. Managing system administrators.* Describes creating and working with System administrators.
- *Part 3. Managing and backing up.* Describes licensing, backup, and utilities.
- *Chapter 8. Licensing BIRT iHub.* Describes licensing options, license key installation, and CPU-binding policies for BIRT iHub.
- *Chapter 9. Backing up BIRT iHub System.* Describes how to back up and restore BIRT iHub volume metadata and data.



# Part One



**Architecture and configuration**



# 1

## Understanding Actuate BIRT iHub architecture

This chapter contains the following topics:

- Understanding BIRT iHub architecture
- Understanding the BIRT iHub process model
- iHub security overview
- Administering BIRT iHub System

---

## Understanding BIRT iHub architecture

Actuate BIRT iHub stores metadata containing information about the cluster and volume configurations in a relational database management system (RDBMS). In the default configuration, Actuate BIRT iHub uses the open-source, third-party database, PostgreSQL. iHub also supports using other alternative, third-party database systems, such as Oracle or a pre-existing PostgreSQL instance.

iHub stores metadata in the following schemas:

- **Cluster**  
Contains settings related to BIRT iHub configuration, such as nodes, templates, volumes, users, and roles.
- **Volume**  
Contains settings related to volume configuration, such as folders, files, and other objects.

### Using a third-party RDBMS with BIRT iHub system

Actuate installs BIRT iHub cluster and volume schemas in the default PostgreSQL RDBMS installation. Installing the schemas in a pre-existing PostgreSQL RDBMS or alternative RDBMS, such as Oracle, requires running a SQL script containing the appropriate Data Definition Language (DDL) statements and switching to the alternative RDBMS. Following sections of this book provide detailed, step-by-step instructions on how to perform these operations.

Only the metadata that specifies the cluster and volume configuration are in the database. Designs, documents, and other data objects are stored in the file system.

### About the cluster and volume schemas

Actuate supports read-only operations on the cluster and volume metadata stored in the tables of the OOTB or other third-party database. Actuate does not support the addition, deletion, or modification of these metadata tables.

Actuate does permit the creation of additional indexes on these tables. For example, a customer can create an index on the job completion notices table to expedite database processing.

Actuate does not recommend any customization of the cluster schema. Any customization that the customer does on the volume schema must be redone when migrating to or reinstalling BIRT iHub. BIRT iHub does not track any database objects that a customer creates. Actuate reserves the right to change the structure of the schemas in future releases.

## Managing backup, recovery, and failover

The BIRT iHub administrator uses third-party RDBMS tools to manage the backup, recovery, and failover capabilities of the cluster and volume metadata database. The administrator uses standard operating system or other third-party tools to manage the backup and recovery of the data files.

The third-party database schemas that contain cluster and volume metadata are critical components of a BIRT iHub system. To guard against data loss, the database administrator must back up the cluster and volume schemas and all related file data to ensure recoverability in the event of failure. Please consult Actuate Support at the time of installation if you have any questions about the backup, recovery, or failover procedures necessary to protect a BIRT iHub system against the possibility of catastrophic failure.

For more information on backing up an BIRT iHub installation, see Chapter 9, “Backing up BIRT iHub System,” later in this book.

In BIRT iHub, there is no concept of volume failover, since each node in a cluster can operate on all the volumes. Configuring cluster and volume database failover is the responsibility of the third-party RDBMS administrator. The database administrator must use the facilities available in the RDBMS to configure this failover capability. Consult the third-party RDBMS documentation for detailed information on how to use native system tools to configure backup, recovery, and failover operations for the externally managed cluster and volume database.

Documentation for a PostgreSQL RDBMS is available at:

<http://www.postgresql.org/docs>

Documentation for an Oracle RDBMS is available at:

<http://www.oracle.com/technetwork/database/enterprise-edition/documentation/index.html>

---

## Understanding the BIRT iHub process model

The Actuate BIRT iHub platform uses a multi-threaded, multi-process model, running single instances of the following components on each iHub node:

- iHub servlet container
  - Provides the run-time environment for client applications, such as BIRT iHub System and Visualization Platform. Client applications communicate with BIRT iHub System using SOAP-based messaging.
- Process Management Daemon (PMD)
  - Distributes service requests among available BIRT iHub services and nodes.

In addition, the BIRT iHub platform supports multiple instances of the following services on each node:

- **Factory**  
Executes requests to generate documents and perform server-side printing.
- **View**  
Supports viewing documents in DHTML and other output formats, such as CSV and PDF. Handles requests to download files from a volume.
- **Integration (Actuate Integration service or AIS)**  
Coordinates the running of information object files that extract data from multiple data sources.

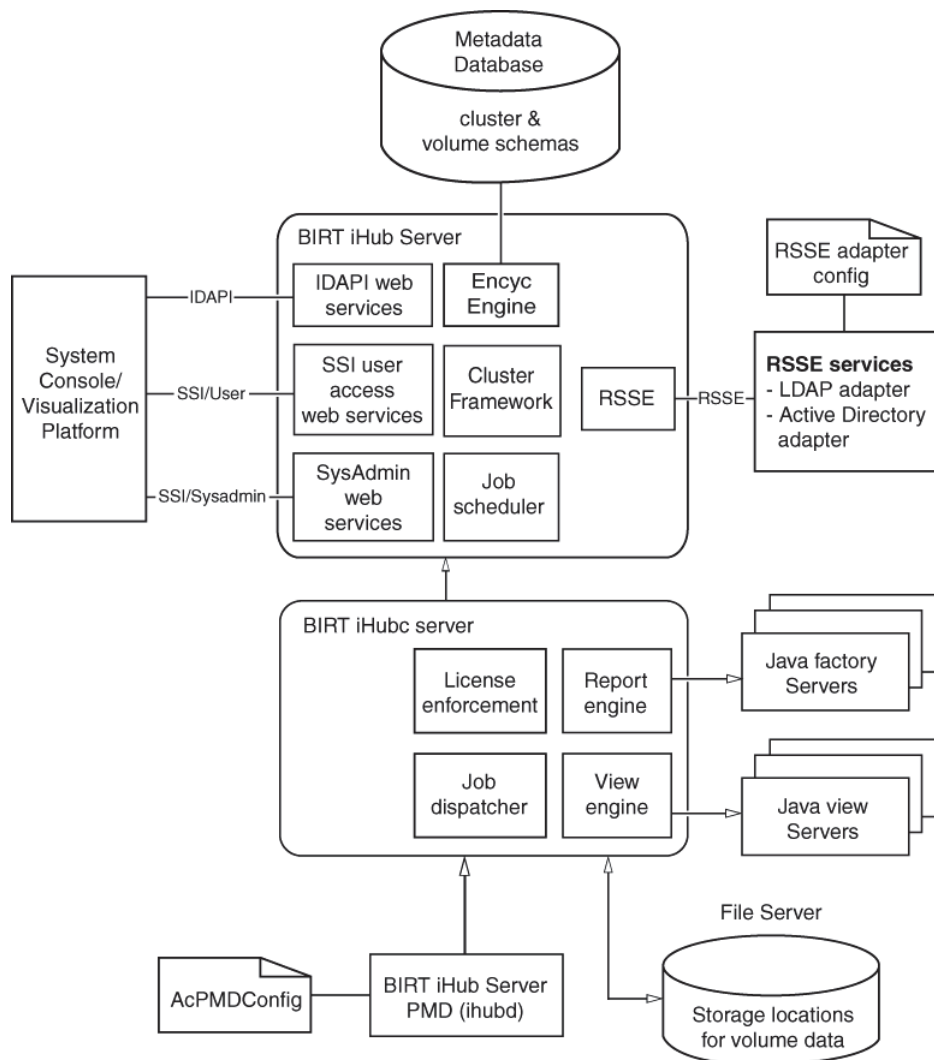
This loosely coupled BIRT iHub architecture model provides the following maintenance and performance benefits:

- Startup and shutdown of a BIRT iHub node is fast because it is independent of the RDBMS that manages a volume. The database server can remain online when shutting down a BIRT iHub node and is available when the node starts up.
- Controlling the sequence of a volume startup is not necessary. All volumes are either already online in the database server or come online as the database server starts.
- Downtime to apply a patch or diagnostic fix for a BIRT iHub node is reduced. The RDBMS does not have to be shutdown.

## **Understanding process flow in a BIRT iHub system**

Figure 1-1 illustrates the BIRT iHub system architecture for a multi-volume, out-of-the-box (OOTB) PostgreSQL database configuration. In this configuration, the iHub administrator starts and stops an iHub instance by running scripts from the command line or using the graphical user interface (GUI) available in System Console.

Client applications, such as System Console and Visualization Platform, run in a servlet container. Single-Sign-On (SSO) security using Security Assertion Markup Language (SAML) provides access to the BIRT iHub system.



**Figure 1-1** BIRT iHub Release 3 system architecture

BIRT iHub supports administering security internally through iHub system services or externally using Report Server Security Extension (RSSE) services such as LDAP or Active Directory. Client applications communicate with BIRT iHub through SOAP messaging using the Actuate Information Delivery API (IDAPI.)

The Process Management Daemon (PMD) or ihubd process handles the services and processes defined on the cluster node in acpmdconfig.xml. The iportal process running in Visualization Platform routes messages to different nodes in

the cluster in a round-robin manner when the Message Distribution service (MDS) is enabled, which is the default setting. To modify MDS processing, set the `MDS_ENABLED` property to false in the `web.xml` file in `iHub/web/iportal/WEB-INF`.

When a message reaches iHub, an administrative or provisioning message is handled locally by the node. Report generation and viewing requests are dispatched by a built-in load balancing program based on cluster service configuration, current load on java factory and viewing processes, and available work units specified for each node.

When a BIRT iHub node receives a request, iHub deserializes the SOAP message, performs the appropriate action, and sends a response in the form of a SOAP message back to the application. For example, BIRT iHub receives a request to run a design, such as a BIRT design, immediately or as a scheduled job. BIRT iHub communicates with the internal framework and the cluster and volume metadata database as necessary to locate the design and identify the resources required to run it.

The reporting engine selects a Java Factory service to run the BIRT design and checks job status. BIRT iHub uses an asynchronous Java Factory service to generate a temporary document or a synchronous Java Factory service to generate a scheduled document.

The View service renders the document in DHTML format, or converts the output to other supported formats, such as CSV or PDF, and handles requests to download files from the volume. The View service sends the document to the requesting application for viewing.

A design that uses an information object utilizes Actuate Integration service (AIS) to extract and cache data from an external data source and perform the following processing:

- Run a query to extract data from an external data source.
- Cache data in iHub System for high availability and to reduce load on the network, data source, and volume by avoiding repetitive data retrieval operations.

The PostgreSQL RDBMS runs as a service in Windows or a process in Linux. The RDBMS can be configured to start automatically or run manually, using a script similar to the BIRT iHub startup script.

iHub stores cluster and volume metadata in the third-party RDBMS, communicating with the RDBMS as necessary using JDBC. iHub uses the physical file system to read and store designs, documents, and other iHub objects as data in volume storage locations.

The out-of-the-box (OOTB) iHub PostgreSQL installation configures the volume database on the local disk to increase the reliability and performance of file input and output (I/O) operations. PostgreSQL discourages creating databases



accessed using a Network File Systems (NFS) for these reasons. For more information, see section 17.2.1 Network File Systems at the following URL:

<http://www.postgresql.org/docs/8.3/static/creating-cluster.html>

The iHub OOTB PostgreSQL RDBMS starts multiple instances to handle connections for running queries that access metadata. In database jargon, PostgreSQL uses a process-per-user, client/server model. For more information, refer to the PostgreSQL documentation at the following URL:

<http://www.postgresql.org/docs/8.4/static/connect-estab.html>

A cluster node is a machine running a BIRT iHub instance. The system administrator adds a node to a cluster to scale BIRT iHub System to the necessary processing requirements. Every cluster node must have network access to the following directory and resources to join the cluster:

- The shared configuration directory
- Cluster resources, such as printers, database systems, and disk storage systems

Each node gets its configuration from a template in `acserverconfig.xml`, which is located in a shared configuration home directory along with the license file, `acserverlicense.xml`.

The `acserverconfig.xml` file contains the server templates as well as other configuration parameters specifying the host names, volume names, port numbers, printers, and services used by nodes in the cluster. When the Process Management Daemon (PMD) starts up, it reads these configurations and exposes them to the process environment variable list. When a node joins a cluster, it configures itself using its template.

After installation and configuring the appropriate environment variables in `acpmdconfig.xml`, the system administrator launches the installed BIRT iHub image from System Console or the command line by passing the necessary arguments or creating a script to execute the command. Nodes with the same cluster ID, running on the same sub-net, automatically detect and join each other to form the cluster. This feature is known as elastic iHub clustering.

The cluster automatically detects the on-off status of any node. Single-point node failure does not affect the availability of other nodes.

The cluster communicates across the network using standard HTTP/IP addressing. The Process Management Daemons (PMDs) located on each node coordinate processing among the available BIRT iHub services based on message type to balance the workload across the nodes.

This loosely coupled model provides the following improvements to intra-cluster messaging:

- Each node in the cluster is relatively independent and identical in terms of components and functionality. Intra-cluster messages are limited to messages for cluster membership and load balancing.
- Operations like design execution and viewing typically require intermediate information from the volume metadata database. This information is now directly retrieved from or updated in the RDBMS, eliminating internal messages to services on other nodes.

This increased scalability of operations at the cluster level can create bottlenecks in the metadata database. Important factors to consider when configuring nodes and ancillary resources include estimating processing power and access to hardware and software resources, such as printers and database drivers.

BIRT iHub instances running on multiple machines maintain cluster and volume metadata in a database, which controls access to shared volume data. This data can be on machines that are not running BIRT iHub, but must be shared and accessible to each iHub instance.

This loosely coupled cluster model provides the following maintenance and performance benefits:

- Startup and shutdown of a BIRT iHub node is fast because it is independent of the RDBMS that manages the cluster and volume. An RDBMS can remain online when shutting down a BIRT iHub node. The RDBMS is available when the node starts up.
- Controlling the sequence of node and volume startup is not necessary. All nodes and volumes are either already online or come online as the RDBMS starts.
- Downtime to apply a patch fix or a diagnostic fix for a BIRT iHub node is reduced. The RDBMS, including the OOTB PostgreSQL database server, does not have to be shutdown. In a BIRT iHub cluster, the patch or diagnostic fix can be applied to one node at a time.

This operational model lends itself well to grid, cloud, and other data-center types of deployments.

---

## iHub security overview

The following sections describe the basic elements of iHub security:

- Security Assertion Markup Language (SAML)
- Secure Sockets Layer (SSL)
- Single sign-on (SSO)

## Understanding SAML

Security Assertion Markup Language (SAML) is an open standard that provides a secure means of authenticating and authorizing a user to resources on the internet. SAML eliminates the need of a user to provide an authentication credential such as a password, to every internet resource the user accesses. The following entities participate in SAML-based communication:

- User  
The party accessing a resource on the internet.
- Service provider  
The application, resource, or service the user wants to access.
- Identity provider  
The entity that authenticates the user to the service provider.

A SAML-enabled identity provider and a SAML-enabled service provider can communicate with each other. The user has an account with the identity provider. The identity provider maintains a list of users and can authenticate them.

The user attempts to access a service provider. The service provider contacts the identity provider. The identity provider authenticates the user, sending an assertion, or message containing information about the user, back to the service provider. The service provider determines that the assertion is valid, then allows the user access.

SAML-based user-authentication activity is transparent to the user. Any service provider that can communicate with an identity provider with which the user has an account can obtain user authentication and grant access to the user. The user authenticates once, with the identity provider. Then, the user can potentially access all the service providers that communicate with the identity provider.

## About the BIRT iHub SAML implementation

BIRT iHub provides its own out-of-the-box (OOTB) SAML identity provider (IdP) implementation. BIRT iHub does not support other third-party SAML identity provider software. BIRT iHub uses Shibboleth IdP 2.4.0 to implement this feature with some customization.

BIRT iHub implements single sign-on (SSO) in SAML 2.0, using the Shibboleth OpenSAML API for System Console, Visualization Platform, and BIRT Analytics. BIRT iHub uses the default authentication method, which specifies PasswordProtectedTransport and PreviousSession.

BIRT iHub enables SSL by default, specifying port 8001 in the `acpmdconfig.xml` file. BIRT iHub stores SSL certificates in the `AC_SERVER_HOME/shared/credential` folder.

BIRT iHub implements 2048 RSA private key encryption. Actuate supports the following secure protocol configurations:

- SSL V3 and TLSV 1.0
- TLSV 1.1 and TLSV 1.2

BIRT iHub disables SSL V2, client-initiated renegotiation for ihubd, and TLS compression.

## Understanding SSL

SSL, or Secure Sockets Layer, is a protocol for securely exchanging information on the internet. SSL provides the means to:

- Encrypt the data exchanged between two parties.
- Establish trust between the two parties by proving identity.

SSL uses the following components:

- **Public key**  
Available to the public. Used to:
  - Encrypt a message sent to the owner of the private key
  - Verify the digital signature on a digital certificate.
- **Private key**  
Available only to the owner of the private key. Used to:
  - Decrypt a message encrypted by the public key.
  - Sign a digital certificate using a digital signature.
- **Digital certificate**  
Contains information such as the certificate issuer, certificate expiration date, information about the certificate owner, and the certificate owner's public key. The certificate is signed by the certificate issuer, using a digital signature.
- **Certification authority (CA)**  
A trusted agency that validates information about an entity such as an online business, then creates a digital certificate for the entity.

As an example, a client user wants to purchase a product from an online business. The client user contacts the online business over the internet. The online business sends its digital certificate back to the client, which contains a digital signature. The private key is required to create the digital signature. The client uses the public key to verify the signature. If the public key can verify the signature, the response indicates that the corresponding private key was used to sign the certificate.

The client must verify that the signer of the certificate is in fact a trusted party, such as a Certification Authority. The client possesses a list of root certificates, each signed by a trusted party or Certification Authority. The client compares the signature on the digital certificate that the online business sent to the client against the signature on the root certificate of the Certification Authority whose name appears on the certificate the online business sent. If the signatures match, the identity of the online business is authenticated. The client and the online business now have a secure connection.

## Understanding SSO

Single sign-on (SSO) refers to any authentication mechanism that supports a user logging in once to gain access to multiple applications. Single sign-on authenticates the user for every application the user is authorized to access and eliminates the need for the user to log in again when switching to another application during a session.

---

## Administering BIRT iHub System

Administering a BIRT iHub System includes the following tasks:

- Setting up users, user groups, folders, files, and other administrative tasks  
An administrator creates, configures, and manages users, user groups, files, and folders, including assigning and updating privileges and managing group memberships. User and user group privileges selectively control access to the volume and its data objects.

- Scheduling jobs to run designs and generate documents  
Each node in an iHub cluster has a job scheduler and dispatcher. The job dispatcher sends jobs to the local resource group factories.

In this loosely coupled cluster model, the dispatcher sends a job from the pending queue to available factories, balancing the load across the cluster. Multiple job schedulers running on the nodes in a cluster allows the system to scale processing to handle a large number of scheduled jobs at the same time.

- Reviewing logs and auditing the information to diagnose system problems  
BIRT iHub Logging and Monitoring System (LMS) can capture usage and error information in log files to assist a system administrator in evaluating resource usage and troubleshooting problems. The usage and error logging applications are open framework applications, which are available as DLLs in Windows and shared libraries in Linux.

It is best to set up the iHub system so that System Console and the LMS server are on a separate computer. This allows system administrators to access

system logging information when iHub is busy, and also reduces the memory requirements on the iHub computer.

- **Configuring a cluster using automated installation programs**  
The system administrator can run the installation program to configure BIRT iHub. Each cluster node gets its configuration from a template in `acsserverconfig.xml`, located in a shared configuration directory. Nodes with the same cluster ID, running on the same sub-net, automatically detect and join each other to form the cluster.
- **Using BIRT iHub Integration Technology scripts and tools to develop client applications and extend functionality**  
The Actuate Information Delivery application programming interface (IDAPI) supports integrating and administering BIRT iHub using extensible markup language (XML) and the simple object access protocol (SOAP). Using the IDAPI, developers can create an application that performs such tasks as scheduling a custom event or running a Report Server Security Extension (RSSE) application to manage users and user groups in an external system, such as LDAP or Active Directory.

A BIRT iHub administrator uses the System Console, Visualization Platform, command-line utilities, and Integration Technology components to perform these tasks.

Please consult the following iHub documentation for more information on how to administer the system using these components:

- *Installing BIRT iHub and BIRT Analytics*  
Provides detailed instructions on how to use automated installation programs and command-line utilities to install BIRT iHub modules, such as Visualization Platform, System Console, Metrics Management, and BIRT Analytics.
- *System Administrator Guide*  
Describes how to use System Console to perform tasks such as managing a cluster, connecting to a database, adding a volume, setting up user accounts and groups, updating the license, and configuring other BIRT iHub properties, such as logging levels, notification, and printing. Also describes BIRT iHub system architecture, and backup and recovery operations.
- *Managing Volumes and Users*  
Describes how to use Visualization Platform to perform volume administration tasks such as executing designs, viewing documents, and scheduling jobs.
- *Using Visualization Platform*

Describes how to use Visualization Platform to perform volume administration tasks such as executing designs, viewing documents, and scheduling jobs.

- *Application Integrator Guide*  
Provides information about application programming using the SOAP-based Actuate Information Delivery API (IDAPI), including a Java developer guide and sections on logging and using the Java Report Server Security Extension (RSSE).

## Using JDBC to connect to the BIRT iHub database

BIRT iHub uses JDBC for connecting to the metadata database. The BIRT iHub run-time JRE environment supports Java 6.0 and 7.0. Any JDBC driver must be compatible with these JRE versions.

iHub requires a JDBC driver that complies with the JDBC 3.0 specification or later. The function `Driver.jdbcCompliant()` must return `TRUE`. `DatabaseMetadata.getJDBCMajorVersion()` must return 3 or greater than 3.

A system administrator, who decides to customize BIRT iHub to connect to a database other than the OOTB PostgreSQL database, must ensure that the JDBC driver returns adequate information about the types on the database. At a minimum, the database must return the common data types, such as integer, floating point, and character. If the database does not return these common data types, then the database administrator must customize the database mapping framework to specify the types.

The JDBC driver must also support the following features:

- Scrollable cursor
- Retention of a cursor after commit
- Update using a prepared cursor

When using connection pooling, the tracing functionality of the JDBC driver captures connection pool run-time statistics.

## About international character sets

BIRT iHub operates on the assumption that the metadata database is configured to run with UTF-8 encoding. Any other database encoding scheme requires configuring the connection parameters to specify the database encoding. The driver must handle the conversion to UCS2.

## Administrative reports

The default iHub volume contains sample BIRT reports that provide information using the metadata and data extracted from the OOTB database, including job

schedule, file, and user tracking and usage and error logging. Installing the sample volume is an option in a custom installation.

## **Supported operating systems**

Actuate BIRT iHub supports the following operating systems:

- Windows
- Linux



# Configuring a BIRT iHub cluster

This chapter discusses the following topics:

- Installing a BIRT iHub cluster node
- Preparing the BIRT iHub cluster environment

---

## Installing a BIRT iHub cluster node

A node is a machine running a BIRT iHub instance. The system administrator adds a node to a BIRT iHub cluster to improve availability and throughput and scale the installation to meet processing requirements.

Every cluster node must have network access to the following directory and resources to join the cluster:

- The shared configuration directory
- Cluster resources, such as printers, database systems, and disk storage systems

Each node gets its configuration from a template in `acserverconfig.xml`, which is located in the shared configuration directory along with the license file, `acserverlicense.xml`. The `acserverconfig.xml` file also contains other configuration parameters specifying the host names, volume names, port numbers, printers, and services used by nodes in the cluster.

When the Process Management Daemon (PMD) starts up, it reads these configurations and exposes the settings to the process environment variable list. When a node joins a cluster, the node configures itself using the designated template.

---

## Preparing the BIRT iHub cluster environment

This section makes the following assumptions:

- The system administrator installed System Console and a BIRT iHub instance on one computer. This machine contains the shared configuration directory, which all nodes in the cluster access. This section refers to the machine containing the shared configuration directory as node1.
- The system administrator installed a BIRT iHub instance on another computer, referred to in this section as node2.
- The system administrator created a cluster using System Console. You must create a cluster using System Console before adding any nodes to the cluster. See “Creating and configuring a cluster” in Chapter 6, “Managing clusters.”

The system administrator performs the following tasks on the cluster node machines to support clustering:

- Creates the shared configuration directory
- Shares the folders that all cluster nodes access
- Configures two nodes to communicate with each other

- Specifies a logon account for the Actuate iHub 3 service on a cluster node, if necessary

This section provides examples of these tasks in the Windows environment in an installation performed using the graphical installer.

The following list describes these tasks in more detail within the context of the sequence in which the administrator performs them:

- Before adding the first node, the node containing the shared configuration directory, to a cluster in System Console, the system administrator checks the Actuate iHub 3 Service Log On property for whether it specifies an account having administrator privileges, and if necessary, specifies a logon account for the Actuate iHub 3 service that has administrator privileges.
- Before adding a second node to a cluster in System Console, the system administrator performs the following tasks:
  - On node1:
    - Creates a folder for the shared configuration directory
    - Shares the configuration and storage folders that all cluster nodes must access
    - Obtains the machine host name and IP address
    - Adjusts any network security configuration, such as a firewall, that can prevent shared access between machines
    - Tests the network accessibility of the machine
    - Checks Actuate iHub 3 Service Log On property for whether it specifies an account having administrator privileges, and if necessary, specifies a logon account for the Actuate iHub 3 service that has administrator privileges
  - On node2:
    - Obtains the machine host name and IP address
    - Adjusts any network security configuration, such as a firewall, that can prevent shared access between machines
    - Tests the network accessibility of the machine
    - Checks Actuate iHub 3 Service Log On property for whether it specifies an account having administrator privileges, and if necessary, specifies a logon account for the Actuate iHub 3 service that has administrator privileges
- Before adding a third or subsequent node to a cluster in System Console, the system administrator performs the following tasks on the new node:
  - Obtains the machine host name and IP address

- Adjusts any network security configuration, such as a firewall, that can prevent shared access between machines
- Tests the network accessibility of the machine
- Checks Actuate iHub 3 Service Log On property for whether it specifies an account having administrator privileges, and if necessary, specifies a logon account for the Actuate iHub 3 service that has administrator privileges

For a demonstration showing how to add three nodes to a cluster, see “Adding cluster nodes to a cluster” in Chapter 6, “Managing clusters.”

AC\_SHARED\_HOME is a variable that represents the folder that contains the shared configuration directory, to which all nodes in a cluster share access. This section makes reference to the following AC\_SHARED\_HOME variable settings:

- In a default BIRT iHub installation on Windows, performed using the installer, in which the install folder is C:\Actuate\BIRTiHubVisualization, AC\_SHARED\_HOME represents the following path:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
```

- In a default command-line installation on Windows, in which the install folder is C:\Actuate, AC\_SHARED\_HOME represents the following path:

```
C:\Actuate\iHub3\modules\BIRTiHub\iHub\shared
```

- In a default BIRT iHub installation on Linux, in which the install folder is /opt/actuate/iHub3, AC\_SHARED\_HOME represents the following path:

```
/opt/actuate/iHub3/modules/BIRTiHub/iHub/shared
```

The following sections provide examples of the operations necessary to support clustering in the Windows environment in an installation performed using the installer.

## Creating the shared configuration directory

The system administrator creates the folder for the shared configuration directory on node1 before adding the second node to the cluster.

### How to create the shared configuration directory

On node1, in AC\_SHARED\_HOME, create a new folder for the cluster to use as the shared configuration directory. For example, create a folder named config\_cluster.

## Sharing the folders that all cluster nodes access

In a BIRT iHub installation, cluster nodes must have read-write access to the following folders in AC\_SHARED\_HOME on node1:

- `config_cluster`  
The shared configuration directory. System Console populates this folder when the system administrator adds the second node to the cluster.

- `storage`  
Contains the files for the sample volume, Default Volume.

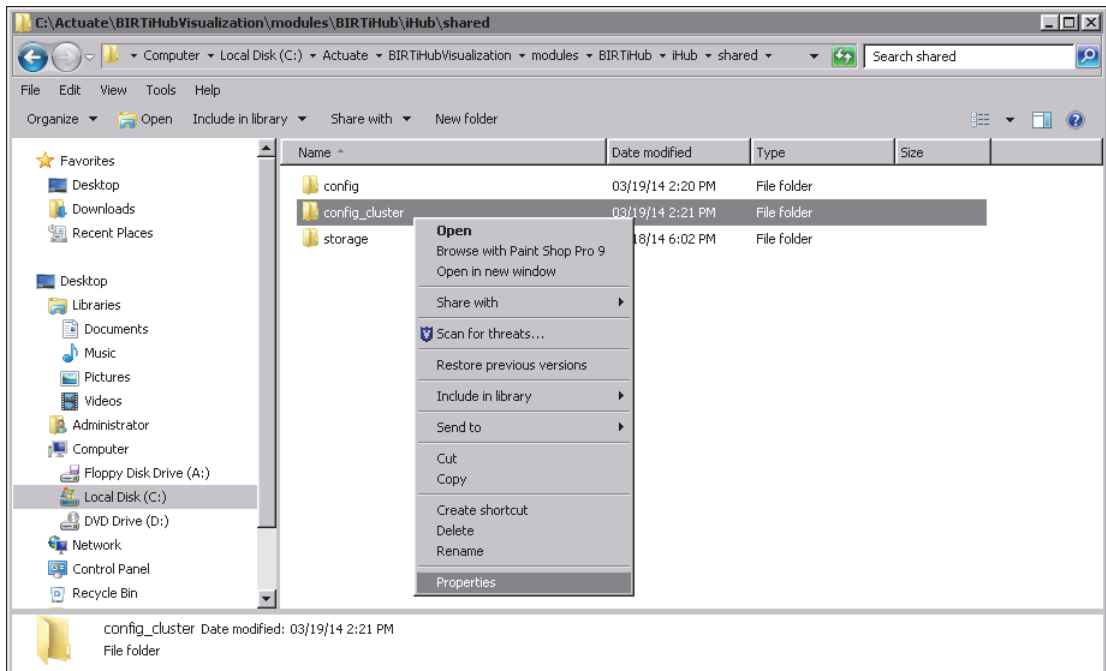
The system administrator shares these folders before adding the second node to the cluster.

The following instructions provide a basic example of the operations required to configure network sharing. It is the responsibility of the system administrator performing this task to make sure that all settings conform to the security policies in force for the environment.

### How to share the `\config_cluster` and `\storage` folders

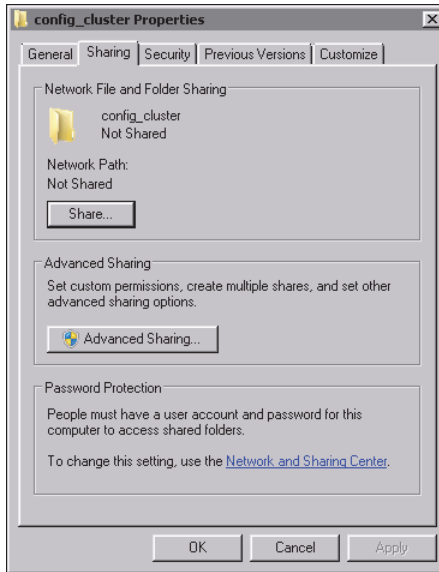
To give a cluster node read-write access to these resources on node1 perform the following tasks:

- 1 Using Windows Explorer on node1, right-click the `config_cluster` folder, and choose Properties, as shown in Figure 2-1.



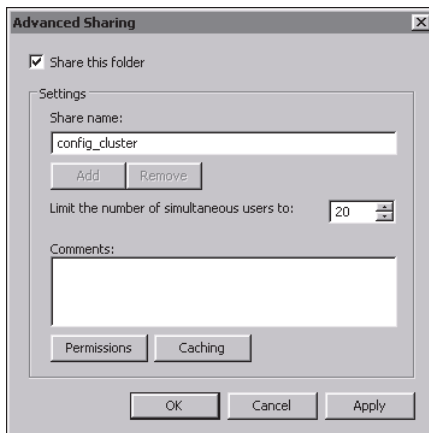
**Figure 2-1** Choosing Properties

- 2 On `config_cluster` Properties, choose Sharing, as shown in Figure 2-2. On Sharing, choose Advanced Sharing.



**Figure 2-2** Choosing Advanced Sharing

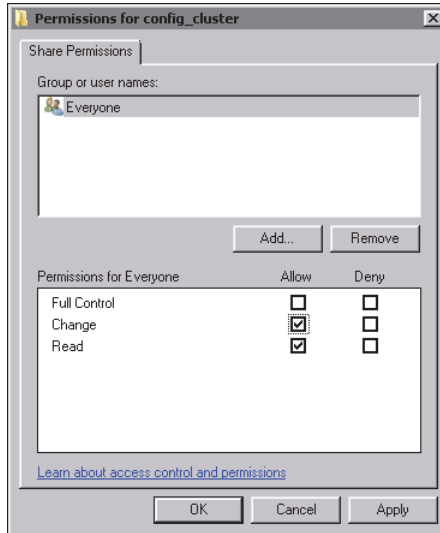
- 3 On Advanced Sharing, select Share this folder, as shown in Figure 2-3.



**Figure 2-3** Selecting Share this folder

On Advanced Sharing, choose Permissions.

- 4 On Permissions for `config_cluster`, in Share Permissions, select Allow for Change and Read, as shown in Figure 2-4.

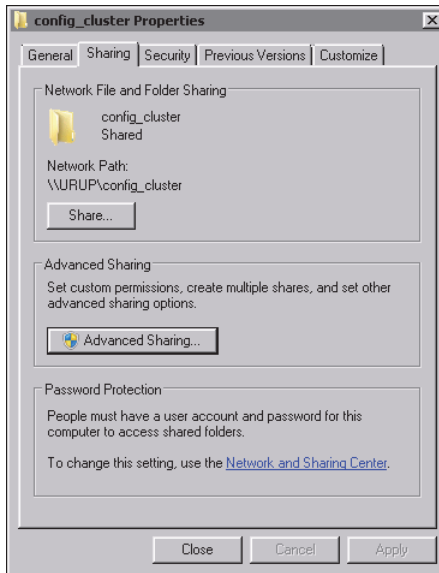


**Figure 2-4** Selecting Change and Read permission

Choose OK.

On Advanced Sharing, choose OK.

On config\_cluster Properties, take note of the Network Path, as shown in Figure 2-5. You specify this path when adding the node to the cluster in System Console. Choose Close.



**Figure 2-5** Taking note of the Network Path

- 5 Repeat steps 1 through 4 for the storage folder that contains the sample volume files. Make sure that all settings conform to the security policies in force for the environment.

In step 4, take note of the Network Path appearing on storage Properties—Sharing. You specify this path when enabling Default Volume in System Console after adding the second node to the cluster.

Close Windows Explorer.

## Configuring two nodes to communicate with each other

Before adding a node to a cluster, perform the following tasks to support communication between the node containing the shared configuration directory and the node you are going to add to the cluster.

- Turn off a Windows firewall
- Obtain the machine name and IP address of each machine
- Test the network connection between the two machines

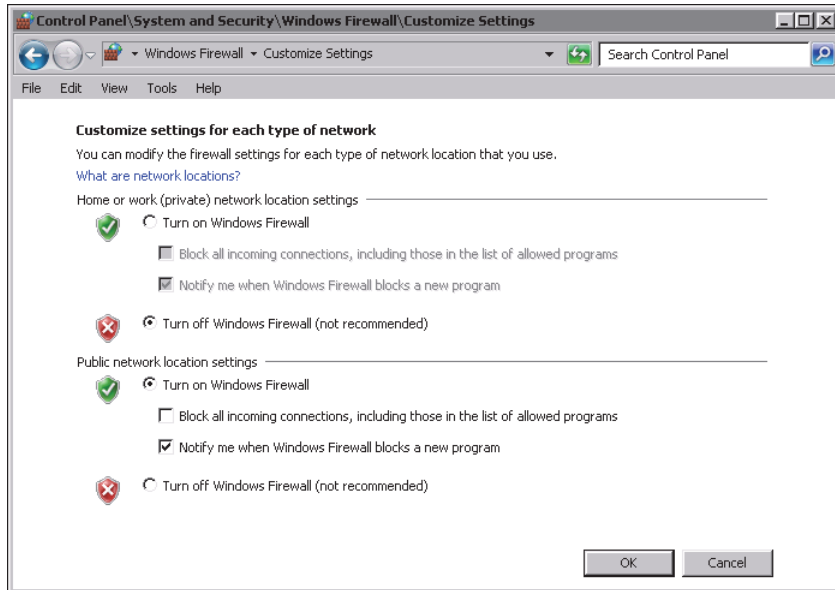
### How to turn off a Windows firewall

Perform the following steps on both node1 and node2:

- 1 Choose Start→Control Panel→System and Security→Windows Firewall.



- 2 On Windows Firewall, choose Turn Windows Firewall on or off. Make sure that the firewall settings conform to the security policies in force for the environment.
- 3 On Customize Settings, in Home or work (private) network location settings, choose Turn off Windows Firewall, as shown in Figure 2-6.



**Figure 2-6** Turning off the home or work network location firewall  
Choose OK.

### How to display a computer's IP address

To obtain the host names of node1 and the computer on which you will install the cluster node, perform the following tasks on node1 and node2:

- 1 Choose Start→Programs→Accessories→Command Prompt.
- 2 In Command Prompt, type the following command:

```
ipconfig /all
```

Press Enter. The host name appears, as shown in Figure 2-7. In this example, the host name for node1 is urup.

```

cs, Administrator: Command Prompt
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ipconfig /all

Windows IP Configuration

Host Name . . . . . : URUP
Primary Dns Suffix . . . . . :
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : actuate.com

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : actuate.com
    Description . . . . . : Intel(R) PRO/1000 MT Network Connection
    Physical Address. . . . . : 00-0C-29-7F-C9-87
    DHCP Enabled. . . . . : Yes
    Autoconfiguration Enabled . . . . : Yes
    IPv4 Address. . . . . : 192.168.41.140(Preferred)
    Subnet Mask . . . . . : 255.255.255.0

```

**Figure 2-7** Displaying the host name

- 3 Write the host names and IP addresses of the computers to be clustered, as shown in Table 2-1.

**Table 2-1** Host names and IP addresses of computers to be clustered

iHub	Host name	IP address
Node1	urup	192.168.41.140
Node2	kozu	192.168.41.138

### How to test the connection between computers

Perform the following steps on both computers:

- 1 In Command Prompt, type the ping command followed by the IP address or host name of the other computer. For example, type the following command to ping a computer named kozu:

```
ping kozu
```

Press Enter.

If your computer reaches the other computer, Command Prompt displays a series of replies, as shown in Figure 2-8.

```

cs, Administrator: Command Prompt

C:\Users\Administrator>ping kozu

Pinging kozu.actuate.com [192.168.41.138] with 32 bytes of data:
Reply from 192.168.41.138: bytes=32 time<1ms TTL=128
Reply from 192.168.41.138: bytes=32 time<1ms TTL=128
Reply from 192.168.41.138: bytes=32 time<1ms TTL=128
Reply from 192.168.41.138: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.41.138:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Administrator>_

```

**Figure 2-8** Receiving a reply to a ping command

- 2 Close Command Prompt.

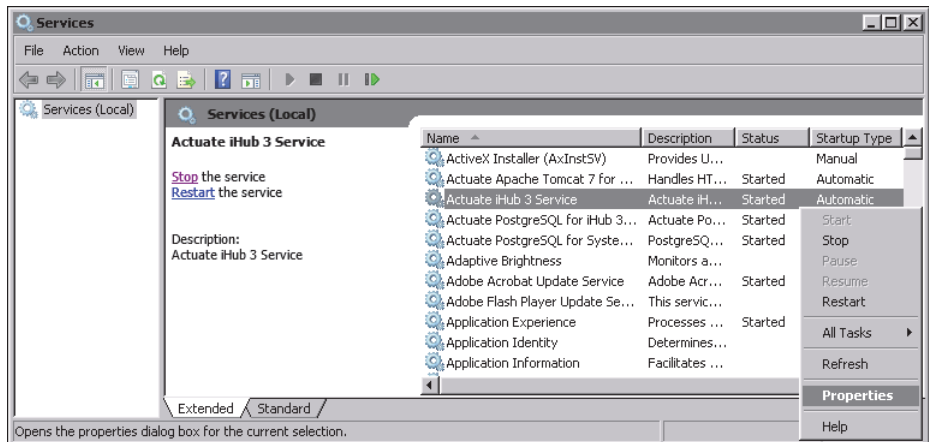
## Specifying a logon account for the Actuate iHub 3 service on a cluster node

Before adding the node to the cluster, the system administrator checks the Actuate iHub 3 Service Log On property for whether it specifies an account having administrator privileges. If the Log On property does not specify an account having administrator privileges, the system administrator performs the following tasks:

- Stops the Actuate iHub 3 service
- Specifies a logon account for the Actuate iHub 3 service that has administrator privileges
- Restarts the Actuate iHub 3 service

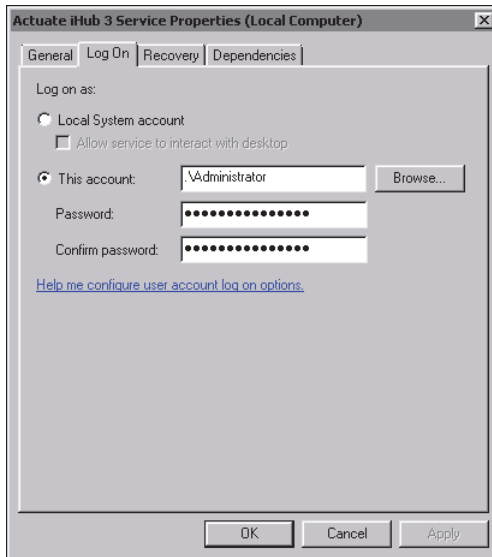
### How to check the Actuate iHub 3 Service Log On property

- 1 Choose Start—Control Panel—System and Security—Administrative Tools—Services. On Services, right-click Actuate iHub 3 Service, and choose Properties, as shown in Figure 2-9.



**Figure 2-9** Choosing Actuate iHub 3 Service properties

- 2 On Actuate iHub 3 Service Properties, choose Log on. If This account already specifies an account having administrator privileges, as shown in the example in Figure 2-10, you do not need to specify a logon account for the Actuate iHub 3 service. Choose Cancel on Actuate iHub 3 Service Properties, and close Services. Otherwise, perform the tasks described in “How to specify a logon account for the Actuate iHub 3 service,” which follows.

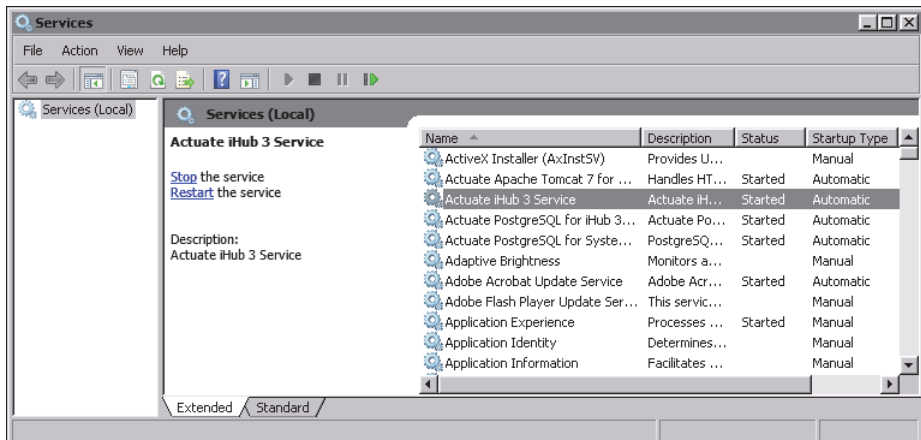


**Figure 2-10** Checking the Log On property

### How to specify a logon account for the Actuate iHub 3 service

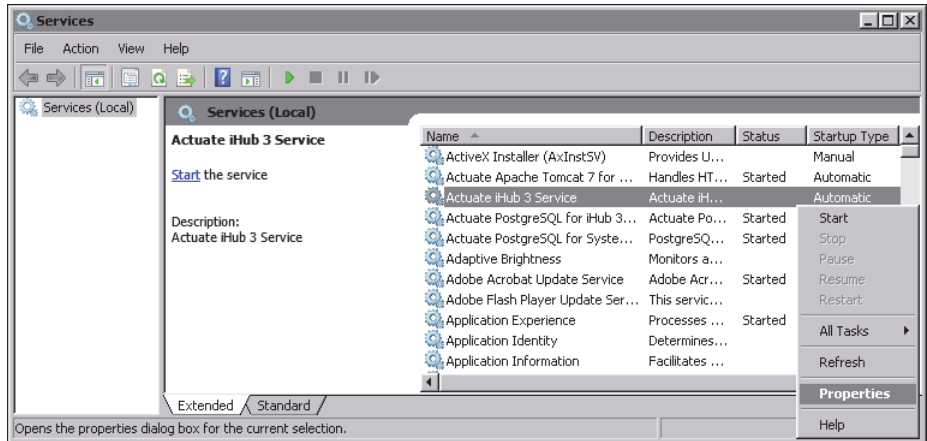
To specify a logon account for the Actuate iHub 3 service, perform the following tasks on the machine, in this example, node2:

- 1 Choose Start—Control Panel—System and Security—Administrative Tools—Services. On Services, select Actuate iHub 3 Service. Then, choose Stop the service, as shown in Figure 2-11.



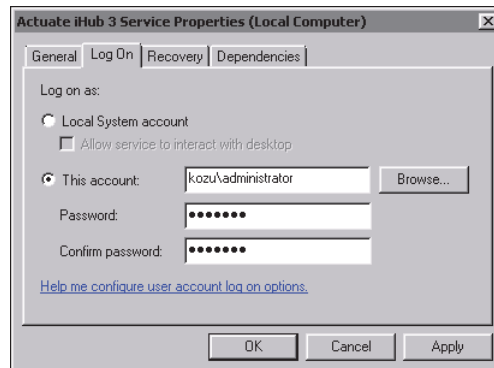
**Figure 2-11** Stopping the Actuate iHub 3 service

- 2 On Services, right-click Actuate iHub 3 Service, and choose Properties, as shown in Figure 2-12.



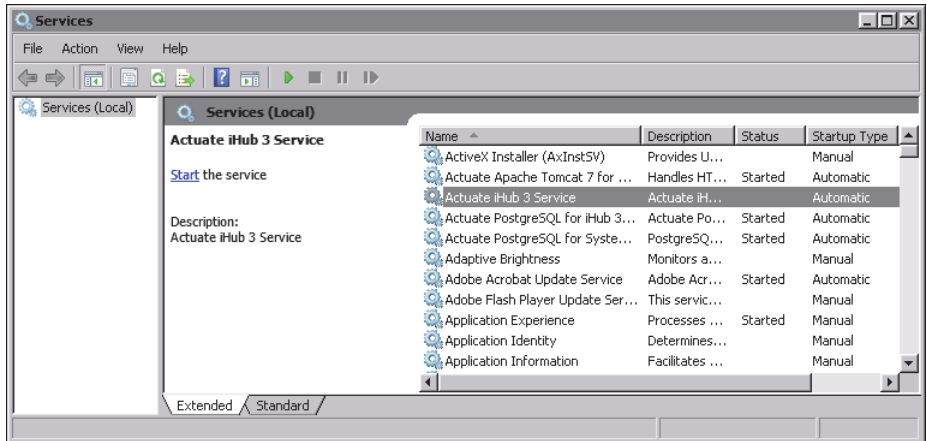
**Figure 2-12** Choosing Properties for the Actuate iHub 3 service

- 3 On Actuate iHub 3 Service Properties, perform the following tasks:
  - 1 Choose Log On.
  - 2 In Log On, select This account, and specify an account that has administrator privileges, such as <machine name>\administrator.
  - 3 In Password and Confirm password, type the password for the account.
  - 4 Choose Apply. Figure 2-13 shows Actuate iHub 3 Service Properties, using kozu as an example machine name in This account.



**Figure 2-13** Specifying an account with administrator privileges  
Choose OK.

- 4 On Services, select Actuate iHub 3 Service, and choose Start the service, as shown in Figure 2-14.



**Figure 2-14** Starting the Actuate iHub 3 service

# Configuring BIRT iHub to use an alternative database

If you want to configure BIRT iHub to use a pre-existing PostgreSQL or Oracle database instead of the out-of-the-box PostgreSQL database, contact Actuate Support.





# 4

## Configuring BIRT iHub security

This chapter discusses the following topics:

- Securing data in a BIRT iHub volume
- BIRT iHub secure communications
- Using SSL

---

## Securing data in a BIRT iHub volume

All files stored in a BIRT iHub volume are subject to a standard security procedure, which restricts file access to authorized users. The iHub security model is based on user groups and privileges. The iHub administrator creates groups for various job functions in an organization, such as finance, marketing, and sales. The privileges, or permissions, to perform certain operations, such as read, write, and execute, are assigned to individual users and to user groups. Administrators assign users to groups. Through these groups, users acquire the privileges to perform particular operations on folders and files.

With this level of security, each user has access to files and folders on a need-to-know basis. For security at a more detailed level, BIRT iHub provides the following types of security:

- Page-level security, which controls user access to particular sections or pages in a report. This security feature requires the Page Level Security option on iHub. To access pages of the published report, a user requires the Secure Read privilege. Read privilege on the report provides access to the entire document.
- Data security, which controls user access to a particular set of data provided by a BIRT data object. This security feature is part of the core iHub functionality. To access data in the published data object, a user requires the Secure Read privilege. Read privilege on the data object provides access to the entire data object.

The security procedure that manages users and their access to files and folders in an iHub volume is implemented using one of BIRT iHub's user management solutions. Page-level security and data security, however, require implementation in BIRT Designer Professional in addition to the licensed options for BIRT iHub.

---

## BIRT iHub secure communications

BIRT iHub secures data in transit to prevent unauthorized third parties from accessing information passing between iHub services and between the iHub server and the end user. Data in transit use the following security protocols:

- Hypertext Transfer Protocol Secure (HTTPS) is used for communication between the user's web browser and both the Administration Console and BIRT iHub Visualization Platform services.
- Secure Socket Layer (SSL) for communications between:
  - The BIRT iHub server processes and the JDBC data sources
  - The BIRT iHub server processes and the metadata database.

- Security Assertion markup Language (SAML) provides a secure means of authenticating and authorizing a user to access a volume.

## Understanding HTTPS

Hypertext Transfer Protocol Secure (HTTPS) is a protocol used to enable secure communication over a computer network. When HTTPS is at the beginning of a URL web address such as `https://www.actuate.com`, the web browser attempts to activate an encrypted connection using the Secure Sockets Layer (SSL).

A server requires two keys and one certificate to build the HTTPS connection. These establish an SSL handshake to encrypt the data between the client and the server. HTTPS also encrypts the URL's query parameters, headers, and cookies.

After the client completes the SSL handshake, the web browser securely uses the same features available to an HTTP connection, such as hyperlinks and JavaScript.

## Understanding SSL

Secure Sockets Layer (SSL) is a protocol for securely exchanging information, such as passwords, credit card numbers, medical, and other private data, over a computer network. SSL provides the means to:

- Encrypt data exchanged between two parties.
- Verify the identity of the server and optionally the client.
- Verify the integrity of the message, that no tampering of the data occurred.

SSL uses the following components:

- Private key  
Available only to the owner of the private key. Used to decrypt a message encrypted by the public key.
- Public key  
Available to the public. Used to encrypt a message sent to the owner of the private key.
- Digital certificate  
Contains information such as the certificate issuer, certificate expiration date, information about the certificate owner, and the certificate owner's public key. The certificate is signed by a certificate authority, using a digital signature.
- Certification authority (CA)  
A trusted agency that validates information about an entity such as an online business, and then signs a digital certificate for the entity to use.

## **BIRT iHub SSL communication process**

A client uses a web browser to access a server hosting Visualization Platform. The secured connection starts when a client visits the SSL secured server using a URL web address beginning with HTTPS. The client receives the server's digital certificate identifying the server and including the server's public key. The client web browser checks the following characteristics of the certificate:

- That the domain in the certificate matches the domain of the server
- That the certificate is trusted or signed by a trusted certificate authority
- That the certificate is not expired
- That the certificate is not revoked
- That the encryption cipher chosen by the server is supported

After accepting the server's certificate, the client uses the public key from the server's certificate to encrypt a message and then sends the message to the server. The message requests that the server generate a session key, also known as a shared secret key. At the same time, the client uses the data in the message to generate the same session key that the client expects the server to generate.

The server uses its private key to decrypt the message from the client. Then, the server uses the data in the message to generate a session key, identical to the session key the client generated. The client and the server use the generated session key to encrypt data, decrypt data, and verify data integrity using checksums.

The server sends a message encrypted using the generated session key back to the client. This message completes the SSL handshake and confirms that data travels securely between both sides of the connection.

## **BIRT iHub SSL support**

BIRT iHub enables SSL by default, specifying port 8001 in the `acpmdconfig.xml` file. BIRT iHub stores SSL certificates in the `AC_SERVER_HOME/shared/credential` folder.

BIRT iHub implements 2048 RSA private key encryption. Actuate supports the following secure protocol configurations:

- SSL V3 and TLSV 1.0
- TLSV 1.1 and TLSV 1.2

BIRT iHub disables SSL V2, client-initiated renegotiation for `ihubd`, and TLS compression.

## Understanding SAML

Security Assertion Markup Language (SAML) is an open standard that provides a secure means of authenticating a user and authorizing the user's access to resources on the internet. SAML eliminates requiring an authentication credential such as a password to every internet resource the user accesses. The following entities participate in SAML-based communication:

- User  
The party accessing a resource on the internet. The user has an account with the identity provider.
- Service provider  
The application, resource, or service the user wants to access
- Identity provider (IdP)  
The entity that authenticates the user to the service provider. The identity provider maintains a list of users.

### SAML communication process

A SAML-enabled service provider communicates with the SAML-enabled identity provider to authenticate the user in the following way:

The user attempts to access a service provider. The service provider contacts the identity provider. The identity provider authenticates the user, sending an assertion, or message containing information about the user, to the service provider. The service provider checks that the assertion is valid, and then allows the user access.

SAML-based user-authentication activity is transparent to the user. Any service provider that can communicate with an identity provider with which the user has an account can obtain user authentication and grant access to the user. The user authenticates once, with the identity provider. Then, the user can access all the service providers that communicate with the identity provider.

### BIRT iHub SAML support

System Console, Visualization Platform, and BIRT Analytics use SAML to authenticate and authorize users.

BIRT iHub provides its own SAML IdP implementation, which is a customized implementation of Shibboleth IdP 2.4.0 OpenSAML API supporting SAML 2.0. This implementation uses the default authentication method, which specifies PasswordProtectedTransport and PreviousSession. BIRT iHub does not support other third-party SAML identity provider software.

---

## Using SSL

Using an SSL certificate with BIRT iHub enables verification of the identity of the BIRT iHub server and encryption of data sent between the web browser and BIRT Visualization Platform. The certificate installed with BIRT iHub is self-signed and is for demonstration purposes only. A self-signed certificate is not signed by a Certification Authority (CA). A CA verifies that a certificate is valid and that no tampering has occurred. Using this demonstration SSL certificate shows a warning in the web browser. Use the self-signed certificate to test the creation of an SSL-based connection between the web browser and the BIRT iHub server.

Using self-signed certificates during server testing is common practice for web developers. To test a secure SSL connection, generate a root certificate that signs and validates the self-signed certificate. This root certificate can then be installed and trusted in web browsers that connect to BIRT iHub. Root certificates from many certificate authorities are preinstalled in operating systems. These root certificates offer temporary SSL certificates that can also be used for testing SSL data security.

An SSL certificate has the following general characteristics:

- Domain name, such as `actuate.com`. The name confirms the server is associated with the domain name of the web site.
- Expiration date. After this expiration date, the certificate will not be trusted.
- Certificate authority signature. The certificate authority distributes a public root certificate that, when trusted, can validate an SSL certificate. Most commercial certificate authorities distribute a public root certificate to computer operating systems. Check that this is the case with your certificate authority.
- The server's public key, used to send encrypted information to the server.

## Using SSL with IDAPI

The Actuate Information Delivery application programming interface (IDAPI) supports integrating and administering BIRT iHub using extensible markup language (XML) and the simple object access protocol (SOAP). Using the IDAPI, developers create applications that perform such tasks as uploading and downloading files, generating a document and scheduling document generation, sending an e-mail notification when a job completes, and using external libraries.

By default, BIRT iHub supports SSL secured SOAP services on port 8001. This port number is set in the `acserverconfig.xml` file in the `SOAPDispatchService` element. In a default Windows installation, the location of this file is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\
config
```

The default values for the port numbers defined in the SOAPDispatchService are:

```
<SOAPDispatchService
  EnableRequestService="true"
  ProvisioningSOAPPort="8010"
  SOAPDispatchSOAPPort="8000"
  ProvisioningSOAPSSLPort="8011"
  SOAPDispatchSOAPSSLPort="8001"/>
```

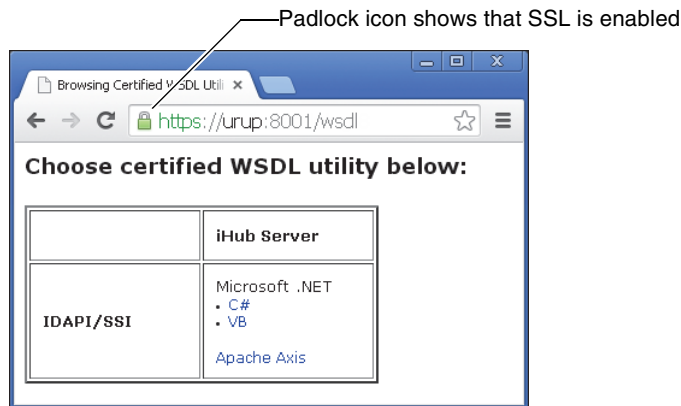
After enabling SSL for the Visualization Platform, test the SSL secured SOAP port using a URL of the following format:

```
https://<servername>:8001/wsdl
```

For example, for a server named urup, use the following URL:

```
https://urup:8001/wsdl
```

This request asks which Web Service Description Language (WSDL) utilities are available. The response is a list of available SOAP APIs and their implementations, as shown in Figure 4-15. The green padlock symbol in the browser address field confirms that SSL security is enabled.



**Figure 4-15** WSDL utility secured with SSL

For more information about using IDAPI, see *Integrating Applications into BIRT iHub*.

## Using SSL with JSAPI

The Actuate JavaScript API (JSAPI) is a set of JavaScript classes that support authenticating users, connecting to data sources, interacting with the user, generating reports, and viewing reports. These classes support using the HTTPS protocol and SSL security.

The JSAPI library is available from any iHub Visualization Platform client installation or Actuate BIRT Deployment Kit. The URL for the library is:

```
http://127.0.0.1:8700/iportal/jsapi
```

- 127.0.0.1:8700 is the host name and TCP port for an available Actuate web application host.
- /iportal is the context root for the web application.
- /jsapi is the default location of the JSAPI libraries.

A script tag in an HTML page loads the JSAPI library, as shown in the following code:

```
<script type="text/javascript" src="http://127.0.0.1:8700/iportal/  
  jsapi">  
</script>
```

After enabling SSL for the Visualization Platform, access the JSAPI library securely using the following URL:

```
https://127.0.0.1:8701/iportal/jsapi
```

The following code uses HTTPS in the script tag that loads the JSAPI library:

```
<script type="text/javascript" src="https://127.0.0.1:8701/  
  iportal/jsapi">  
</script>
```

## Using SSL and external user management

A BIRT iHub system using external tools to manage Visualization Platform users supports connecting to a Lightweight Directory Access Protocol (LDAP) or Active Directory server using SSL. By default, BIRT iHub only connects to an LDAP or Active Directory server that has a signed certificate. To connect to a server that does not have a signed certificate, use the Java keytool utility to add that certificate as a trusted certificate. For information on using the Java keytool utility, see the following URL:

```
http://docs.oracle.com/javase/6/docs/technotes/tools/windows/  
  keytool.html
```

## Using Visualization Platform with SSL

Use SSL to validate the identity of the BIRT iHub Visualization Platform server and to encrypt the data between the client web browser and the BIRT iHub server. To use SSL with the BIRT Visualization Platform, disable SAML and Message Distribution Service (MDS). In the web.xml file located at \iHub\web\iportal\WEB-INF\. In a default Windows installation, this location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\web\  
  iportal\WEB-INF
```



Change the following values and restart the Actuate iHub 3 Service:

- Set the SAMLEntityID parameter to an empty value. For example:

```
<context-param>
  <description>The SP ID for SAML SSO</description>
  <param-name>SAMLEntityID</param-name>
  <param-value></param-value>
</context-param>
```

- Set the Message Distribution Service MDS\_ENABLED parameter to false. For example:

```
<context-param>
  <!-- true or false: Enable or disable MDS -->
  <param-name>MDS_ENABLED</param-name>
  <param-value>>false</param-value>
</context-param>
```

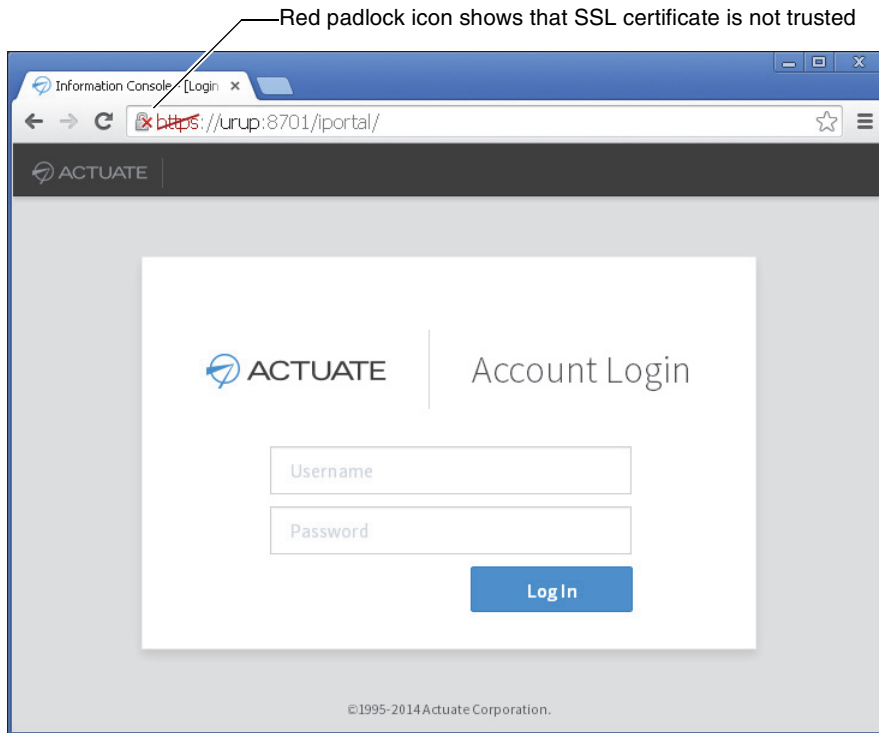
After restarting the Actuate iHub 3 Service, access the Visualization Platform using a URL of the following format:

```
https://<servername>:8701/iportal/
```

For example, for a server named urup, use the following URL:

```
https://urup:8701/iportal/
```

Figure 4-16 shows the secured SSL connection to the Visualization Platform using HTTPS. The default certificate included with the installation of Visualization Platform is not signed by a certification authority and the browser identifies it as not trusted. When you use your own signed and trusted SSL certificate, the web browser trusts your certificate.



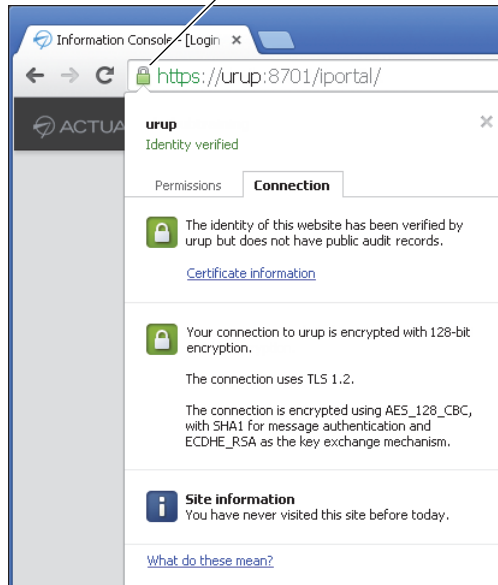
**Figure 4-16** Using HTTPS to access the Visualization Platform

For testing purposes, install and trust the self-signed demonstration certificate included in the default installation of Visualization Platform. This certificate is `birtihub.crt` located in the following folder:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\
config\credentials
```

Each operating system has a different method to install a trusted certificate. For example, in Windows, install this certificate into the Trusted Root Certification Authorities certificate store. Figure 4-17 shows the same web URL as Figure 4-16 after setting the demonstration certificate to be trusted.

Green padlock icon shows that SSL certificate is trusted



**Figure 4-17** Using HTTPS with a trusted certificate

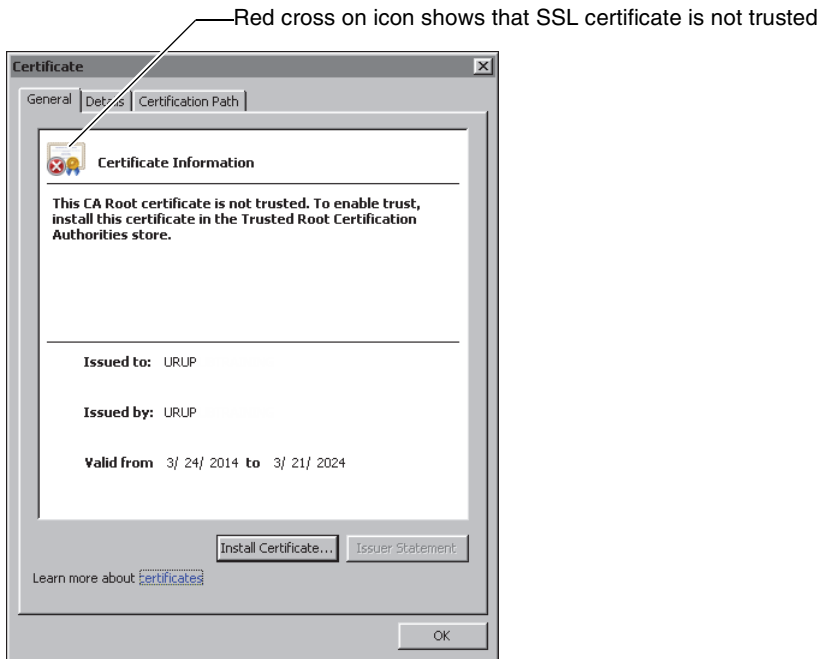
### How to install and trust the demonstration SSL certificate on Windows

This example shows how to install your own root certificate for testing purposes. This procedure applies to browsers other than Mozilla Firefox. Firefox uses a different mechanism to trust certificates. Refer to the Firefox documentation to set up a trusted certificate on Firefox.

- 1 Using Windows Explorer, navigate to the following folder:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\
config\credentials
```

- 2 Open the birtihub.crt file. Certificate—General appears, as shown in Figure 4-18.



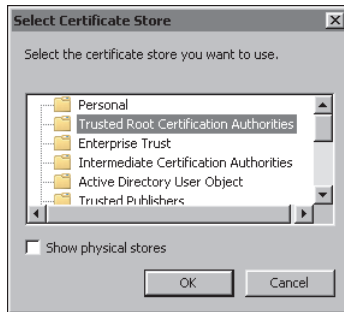
**Figure 4-18** Opening an untrusted root certificate

- 3 Choose Install Certificate. Certificate Import Wizard appears, as shown in Figure 4-19.



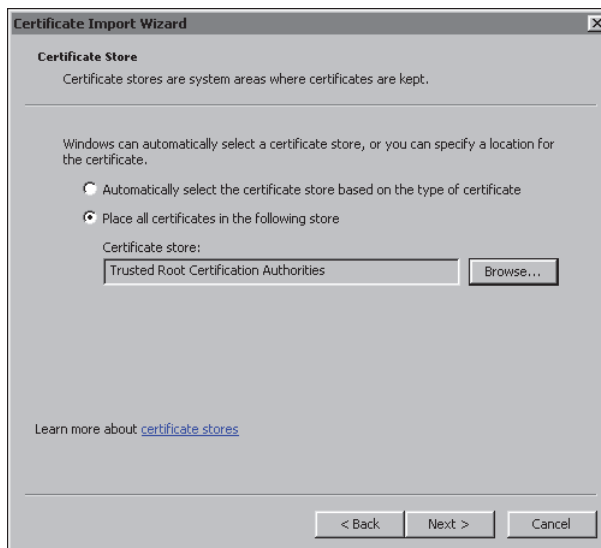
**Figure 4-19** Installing a root certificate

- 4 Choose Next. Certificate Store appears.
- 5 Enable Place all certificates in the following store. Choose Browse. Select Certificate Store appears.
- 6 Select Trusted Root Certification Authorities, as shown in Figure 4-20. Choose OK.



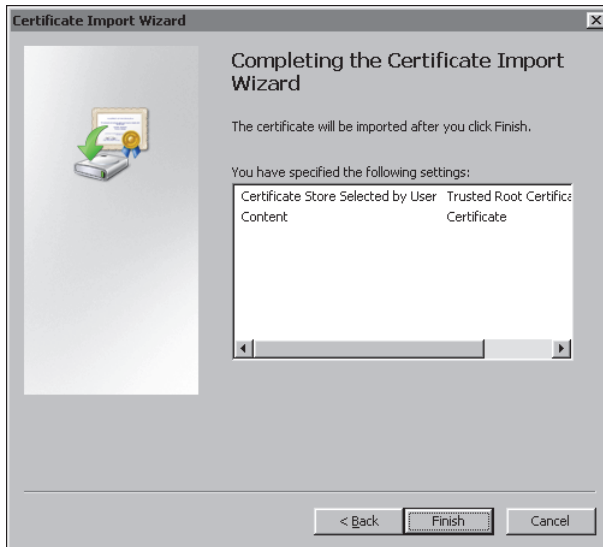
**Figure 4-20** Selecting a store to install the root certificate

- 7 In Certificate Store, choose Next, as shown in Figure 4-21.



**Figure 4-21** Selecting where to install the root certificate

In Completing the Certificate Import Wizard, choose Finish, as shown in Figure 4-22.



**Figure 4-22** Finishing the installation of the root certificate

If you receive a security warning asking if you want to install this certificate, choose Yes.

When you receive an alert that the import was successful, choose OK.

Choose OK to close Certificate—General.

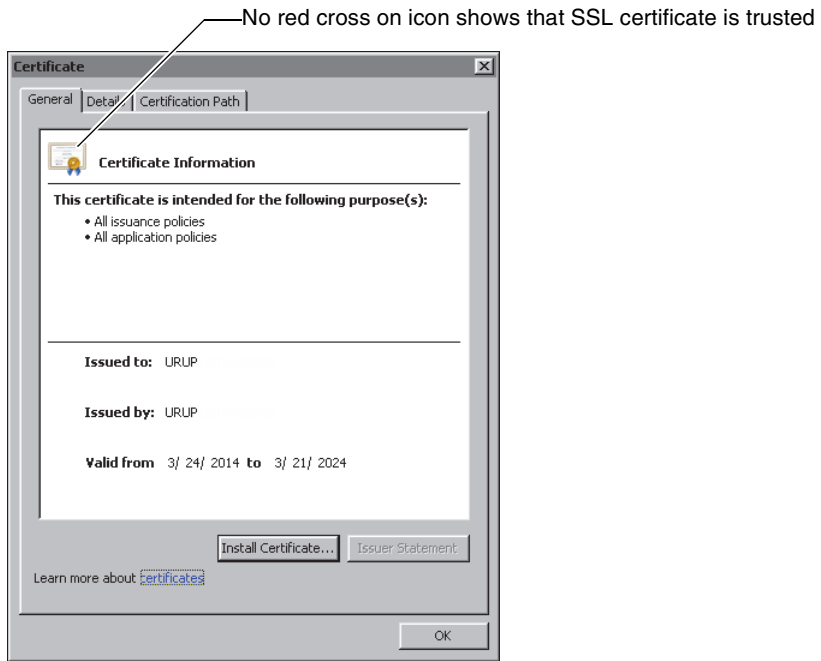
### **How to verify that the HTTPS connection is trusted**

This example shows how to verify that the HTTPS connection to Visualization Platform is trusted. This procedure applies to browsers other than Mozilla Firefox. Firefox uses a different mechanism to check trusted certificates. Refer to the Firefox documentation to check the HTTPS connection on Firefox.

- 1** Using Windows Explorer, navigate to the following folder:

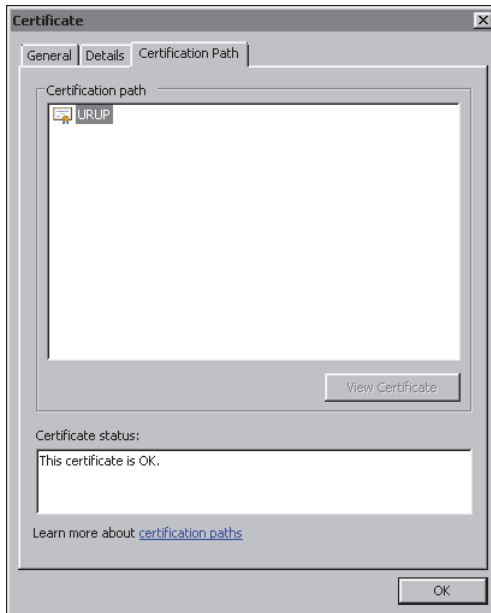
```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\
config\credentials
```

- 2** Open the birtihub.crt file. The certificate should look similar to Figure 4-23.



**Figure 4-23** Verifying the installation of the root certificate

- 3 Choose Certification Path. Verify that the certificate status is OK, as shown in Figure 4-24.



**Figure 4-24** Verifying the certificate status

Choose OK.

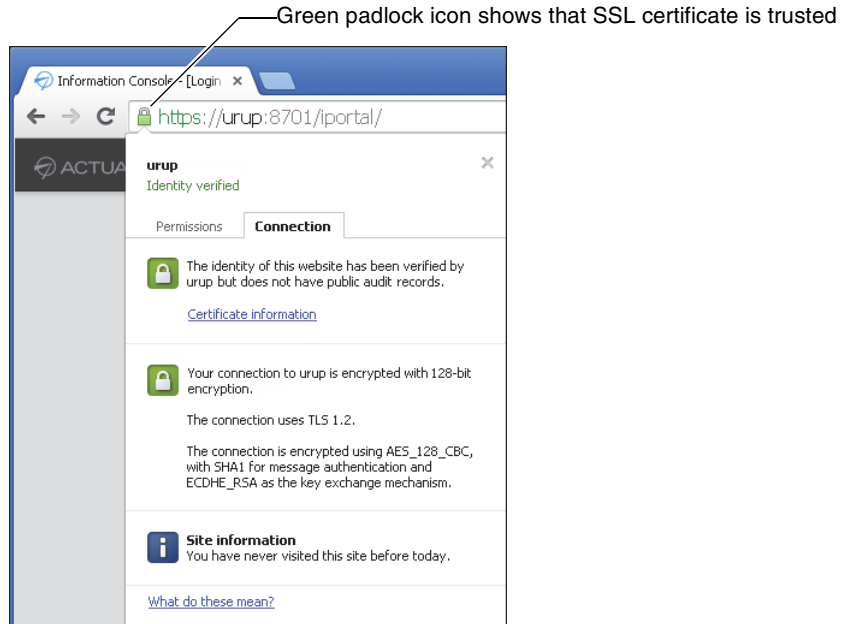
- 4 Open a web browser such as Google Chrome. Type the a URL of the following format, replacing servername with the name of your server. Do not use localhost as the name of the server.

`https://servername:8701/iportal/`

Visualization Platform appears, using HTTPS.

- 5 Choose view site information. Choose Connection as shown in Figure 4-25.





**Figure 4-25** Verifying a secured SSL connection to Visualization Platform

## Using SSL for communication with the volume metadata database

You can encrypt the connection from BIRT iHub to the volume metadata database. By default, BIRT iHub uses a PostgreSQL database to contain volume metadata. In System Console, Metadata Database displays the properties for this PostgreSQL database. The database type for this database is ActuatePostgreSQL. Figure 4-26 shows the following properties for this database, installed on a machine named urup. An asterisk (\*) next to the property name means the property is required.

- Database server  
The host name of the machine containing the database.
- Database port  
The default port number for the default PostgreSQL database is 8432.
- Database name  
The name of the database. The default database name is ihub.
- Encryption Method  
One of the following methods:

- requestSSL  
BIRT iHub encrypts the login request and data using SSL. If the database server does not support SSL, the driver establishes an unencrypted channel.
- SSL  
BIRT iHub performs SSL certificate verification.
- noEncryption  
The channel between BIRT iHub and the metadata database passes unencrypted data.
- Username  
The database user name. The default user name is ihub.
- Password  
The database user name password. The default password of the ihub user is postgres.
- Test Connection  
Choose to verify that the BIRT iHub system can successfully connect to the metadata database.

The screenshot shows a configuration dialog box with the following fields and controls:

- \*Database Type: ActuatePostgreSQL (dropdown menu)
- \*Database Server: URUP (text input)
- \*Database Port: 8432 (text input)
- \*Database Name: ihub (text input)
- Encryption Method: noEncryption (dropdown menu)
- \*Username: ihub (text input)
- \*Password: [masked with 8 dots] (password input)
- Test Connection (button)
- Cancel (button)
- Save (button)

**Figure 4-26** Viewing OOTB PostgreSQL metadata database properties

The PostgreSQL data folder is the location of the certificate and keys used by the PostgreSQL server. If you change the certificate or keys, restart the PostgreSQL server. The SSL files for a default PostgreSQL database are in the following folder:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\data\
postgresql\data
```

Test the SSL connection of the PostgreSQL using the PostgreSQL interactive terminal (psql) command. This command is located in \postgresql\bin folder of the iHub installation. The default location of this software is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\
  postgresql\bin
```

See the documentation at the following URL for more information about configuring and securing a PostgreSQL database:

<http://www.postgresql.org/docs/>

### **How to verify that a PostgreSQL server supports an SSL connection**

The following example shows how to use the Windows command prompt to check if an SSL connection to a PostgreSQL database is possible. This example connects you to the default PostgreSQL server installed with iHub. Use the same computer as the PostgreSQL server. This server has a user name and password with the value postgres and a database table named ihub. If either the user name or password of your PostgreSQL server has changed, use the current user name and password.

- 1** In a command window, navigate to \postgresql\bin folder of the iHub installation. The default location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\
  postgresql\bin
```

- 2** Type the following command. Then press Enter:

```
psql postgresql://postgres@localhost:8432/ihub?sslmode=require
```

- 3** When prompted for a password, type the password for the postgres user. In this example, the password is postgres. Then press Enter. You should receive the following response.

```
psql (9.2.4)
WARNING: Console code page (437) differs from Windows code page
(1252)
      8-bit characters might not work correctly. See psql
reference
      page "Notes for Windows users" for details.
SSL connection (cipher: DHE-RSA-AES256-SHA, bits: 256)
Type "help" for help.

ihub=#
```

- 4** Type \q and press Enter to quit the terminal. You can see in the connection information that an SSL connection is established.

## Managing SSL files

The SSL certificates and keys used to secure BIRT iHub and the Visualization Platform are located in the `\shared\config\credentials` folder in the BIRT iHub installation folder. The default location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
\config\credentials
```

This location contains the iHub's digital certificate in the Privacy Enhanced Mail (PEM) format and the Java KeyStore (JKS) file, which is a repository of security certificates. BIRT iHub is configured to use these certificates in the following files:

- The `acpmdconfig.xml` file located in the iHub `\etc\` folder. The default location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\etc
```

The following settings in the `acpmdconfig.xml` file point to the PEM files.

```
<EnableSSLEngine>true</EnableSSLEngine>
<SSLCertificateFile>
  $AC_CONFIG_HOME$/credentials/birtihub_nopassphrase.pem
</SSLCertificateFile>
<SSLCertificateKeyFile>
  $AC_CONFIG_HOME$/credentials/birtihub_nopassphrase.pem
</SSLCertificateKeyFile>
<SSLRootCertificateFile>
  $AC_CONFIG_HOME$/credentials/birtihub_nopassphrase.pem
</SSLRootCertificateFile>
<SSLCipherSuite>ALL:!ADH:!EDH</SSLCipherSuite>
```

- The `acserverconfig.xml` file located in the iHub `\shared\config` folder. The default location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\
config
```

The following settings in the `acserverconfig.xml` file point to the JKS file.

```
<System
  KeyAlias="birtihub"
  ...
  KeystoreFile="$AC_CONFIG_HOME$/credentials/birtihub.jks"

  KeystorePass="!1!MsGLAyDce0TZhxvh1xDrTkG0Ea6hTslzaidAvxx5pfK!"
  ...
```

Although you can change the SSL key alias and keystore file, you must use the existing the keystore password defined in KeystorePass. Your JKS keystore must use the following password:

```
birtihub
```

If you change these SSL files, you must restart the Actuate iHub 3 Windows service.

You can use the Java keytool utility to view and create SSL certificates and keys. This utility is located in the \bin folder of the BIRT iHub installation of the Java SE Development Kit (JDK). The default location of the JDK is:

```
C:\Actuate3\BIRTiHubVisualization\modules\JDK64\bin
```

### **How to use the Java keytool utility to view the contents of the JKS file**

BIRT iHub generates sample SSL certificates that securely connect a web browser to BIRT iHub. To use SSL security in a production environment, you must replace these SSL certificates with certificates signed by a Certification Authority.

- 1** In a command window, navigate to the \credentials folder. The default location is:

```
C:\Actuate3\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared  
  \config\credentials
```

- 2** Type the following command and press enter:

```
keytool -list -v -keystore birtihub.jks -storepass birtihub
```

Information similar to the following example appears:

```
Keystore type: JKS
```

```
Keystore provider: SUN
```

```
Your keystore contains 1 entry
```

```
Alias name: birtihub
```

```
Creation date: Mar 24, 2014
```

```
Entry type: PrivateKeyEntry
```

```
Certificate chain length: 1
```

```
Certificate[1]:
```

```
Owner: CN=CH-IHUBTRAINING, OU=admin@localhost, O=Actuate, C=US,  
  ST=CA
```

```
Issuer: CN=CH-IHUBTRAINING, OU=admin@localhost, O=Actuate,  
  C=US, ST=CA
```

```
Serial number: 1ca757fc
```

```
Valid from: Mon Mar 24 09:14:30 PDT 2014 until: Thu Mar 21  
  09:14:30 PDT 2024
```

```
Certificate fingerprints:
    MD5: 90:15:F7:79:FB:0F:23:7E:BF:4C:CE:C3:FA:8A:84:91
    SHA1:
    E8:A8:2C:14:74:97:61:F2:F3:74:82:34:B3:AC:F0:A4:D7:4C:BA:0F
    Signature algorithm name: SHA256withRSA
    Version: 3
```

Extensions:

```
#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 66 8E E8 FC DF D8 6E 48 22 CD 61 E1 3E DB 58 90
    f.....nH".a.>.X.
0010: CD CC 6D F9 .m.
]
]
```

```
*****
*****
```

## Using a commercial SSL certificate

If you want to use a commercial SSL certificate instead of the sample certificate that installs with BIRT iHub, perform the following tasks:

- Change the settings in the `acpmdconfig.xml` file to point to the correct PEM files and enable SSL.
- In `web.xml`, turn off SAML.
- Install a commercial SSL certificate.

### How to install a commercial SSL certificate

- 1 Convert the commercial certificate to a `.pem` file.

```
openssl x509 -in MyCertificatesite.crt -out
MyCertificatesite.pem
```

- 2 Combine the private key with the commercial certificate.

```
type MyCertificatereq.key MyCertificatesite.pem >
MyCertificatesitefull.pem
```

- 3 Convert the `.pem` file to PKCS12 format. The password is `birtihub`.

```
openssl pkcs12 -export -in MyCertificatecertfull.pem -out
MyCertificatecertfull.p12
```

- 4 Copy the `p12` file to `\iHub\shared\config\credentials` (where `birtihub.jks` is located).

- 5 Merge the p12 file into the keystore. The password is birtihub in both cases.

```
keytool -importkeystore -srckeystore MyCertificatecertfull.p12  
-srcstoretype PKCS12 -keystore birtihub.jks
```

- 6 Verify that the commercial certificate and the key are in birtihub.jks.

```
keytool -list -keystore birtihub.jks
```

You should see a "1" alias set to "PrivateKeyEntry", for example:

```
1, Aug 12, 2014, PrivateKeyEntry,  
Certificate fingerprint (MD5):  
6A:0A:47:07:21:39:EF:50:5F:09:11:82:4D:E5:35:D8  
birtihub, Aug 7, 2014, PrivateKeyEntry,  
Certificate fingerprint (MD5):  
90:AD:77:5A:9F:4C:C5:4A:D1:83:E5:7C:66:8B:D2:E7
```

- 7 Optionally, rename the "1" alias.

```
keytool -changealias -alias 1 -destalias MyCertificate  
-keystore birtihub.jks
```

- 8 In acserverconfig.xml, change KeyAlias.

```
KeyAlias="MyCertificate"
```

- 9 Restart the Actuate iHub 3 Windows service.

- 10 In a web browser, navigate to <https://MyHost:8701/iportal>, where MyHost is the name of the computer on which iHub is installed. You must also change the cluster URL in System Console.

- 11 Check the SSL certificate. It should have the certification number of the commercial SSL certificate, not the SSL certificate that installs with BIRT iHub.





# Part **Two**



**BIRT iHub System Console**



# 5

## Understanding System Console

This chapter discusses the following topics:

- About System Console
- Viewing clusters, nodes, and system administrators
- Logging in to System Console
- About Monitoring

---

## About System Console

System Console provides a single graphical user interface (GUI) that supports an administrator creating and managing the resources used by the Actuate applications for an entity, such as a corporation or corporate division.

A system administrator creates a cluster for the entity in System Console. A cluster supports creating a scalable BIRT iHub system consisting of two or more machines, or cluster nodes, each running BIRT iHub using the same configuration details.

When setting up a cluster, the administrator defines and configures resources such as:

- Cluster nodes
- BIRT iHub metadata database
- Volumes the cluster uses
- BIRT iHub license
- Storage space for volume data on disk

In addition to creating and managing clusters, a system administrator uses System Console to perform the following operations:

- Create alerts  
System Console monitors a range of activity, conditions, and resources in a BIRT iHub System. An attribute identifies a monitored item. The system administrator can set a threshold on an attribute. When the monitored item reaches the threshold, for example, when the days until the product license expires reaches a specified number, System Console generates an alert.
- Create other system administrators  
A system administrator can create other system administrators.
- Configure security  
The system administrator selects where user authentication and authorization information are stored and provides information about the identity provider when creating a new cluster.

System Console contains the following features:

- Monitoring  
Monitor cluster activity.
- Clusters  
Create, update, and delete clusters. View logs and detailed information on system resource usage.

- Settings  
View system information, create System Administrator users, update the information for an existing System Administrator, and specify an e-mail server.

---

## Viewing clusters, nodes, and system administrators

System Console displays a list for each of the following System Console data items:

- Alerts
- Clusters
- Volumes

For each list, the system administrator can view the entire list, or a subset. By default, System Console displays the entire list.

The system administrator can filter each list to display only the items for which the item property System Console searches contains a string the administrator specifies. The administrator can filter the Clusters list to alternatively display only starred items.

Table 1-1 shows the list type, the list item property on which System Console performs a string search, and whether System Console filters the list type to show only monitored items.

**Table 1-1** Filtering System Console lists

System Console list	List item property on which System Console searches for specified string:	Does System Console display only monitored items?
Alerts	Cluster name	N/A
Clusters	Cluster ID	Yes
Volumes	Volume Name	N/A

---

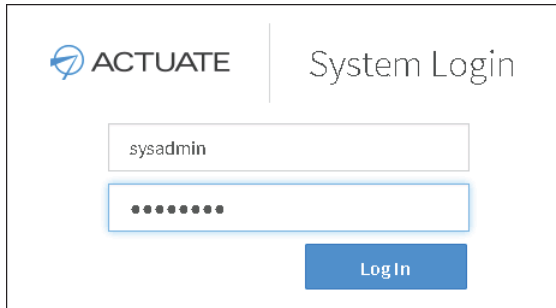
## Logging in to System Console

System Console is a standalone application, which the administrator logs into separately from a BIRT iHub application.

When logging in to System Console for the first time, you log in as the default system administrator, using a default password. The following section describes how to log in to System Console.

## How to log in to System Console

- 1 Open a new browser window. In the address bar, type the following URL:  
`http://localhost:8500/sysconsole`
- 2 On System Login, specify the following credentials, as shown in Figure 1-27:
  - Username: sysadmin
  - Password: system11



**Figure 1-27** Logging in to System Console as default system administrator

---

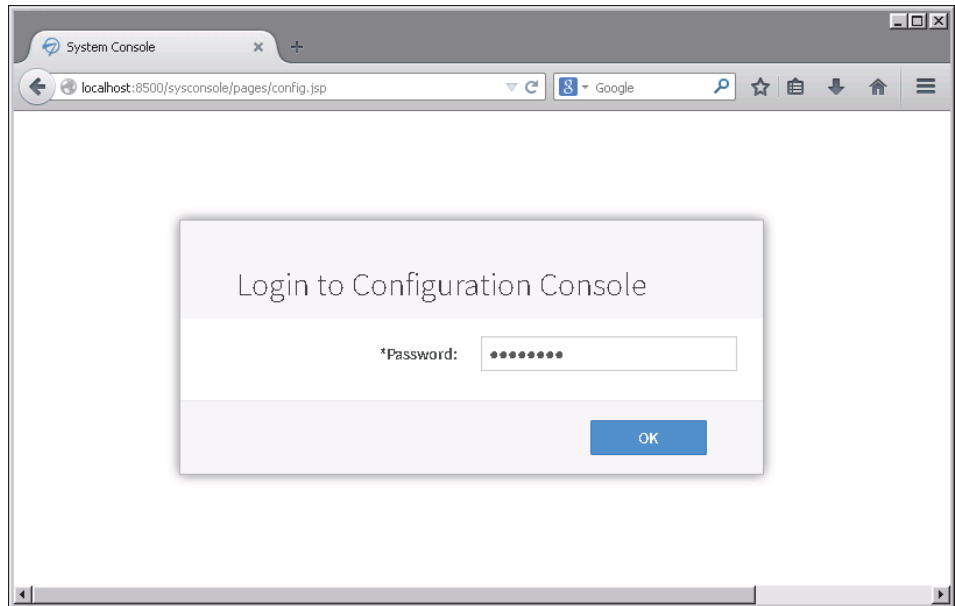
## Using the System Console configuration user interface

The System Console configuration user interface supports the following actions:

- Separating System Console from the System Console metadata database  
The name of the database System Console uses is umc. System Console uses umc to store metadata about objects System Console contains, for example, the name, ID, and description of each cluster, all volume names in System Console and the clusters to which the volumes belong, and the name and e-mail address of each system administrator. By default, the System Console install program installs and configures a PostgreSQL database for System Console on the same machine as System Console. The System Console configuration user interface supports System Console creating and using a PostgreSQL metadata database that resides on a different machine.
- Updating the System Console metadata database  
A scenario this functionality supports is if a database administrator changes the password for the umc database and System Console can no longer connect to umc, a system administrator can specify the new password on the System Console configuration user interface to re-establish connectivity.

## How to log in to the System Console configuration user interface

- 1 Open a new browser window. In the address bar, type the following URL and press Enter:  
`http://localhost:8500/sysconsole/pages/config.jsp`
- 2 On Login to Configuration Console, type the password for sysadmin, the default System Console administrator, as shown in Figure 1-28. The default password for sysadmin is system11.



**Figure 1-28** Accessing the System Console configuration user interface  
System Console appears, as shown in Figure 1-29.

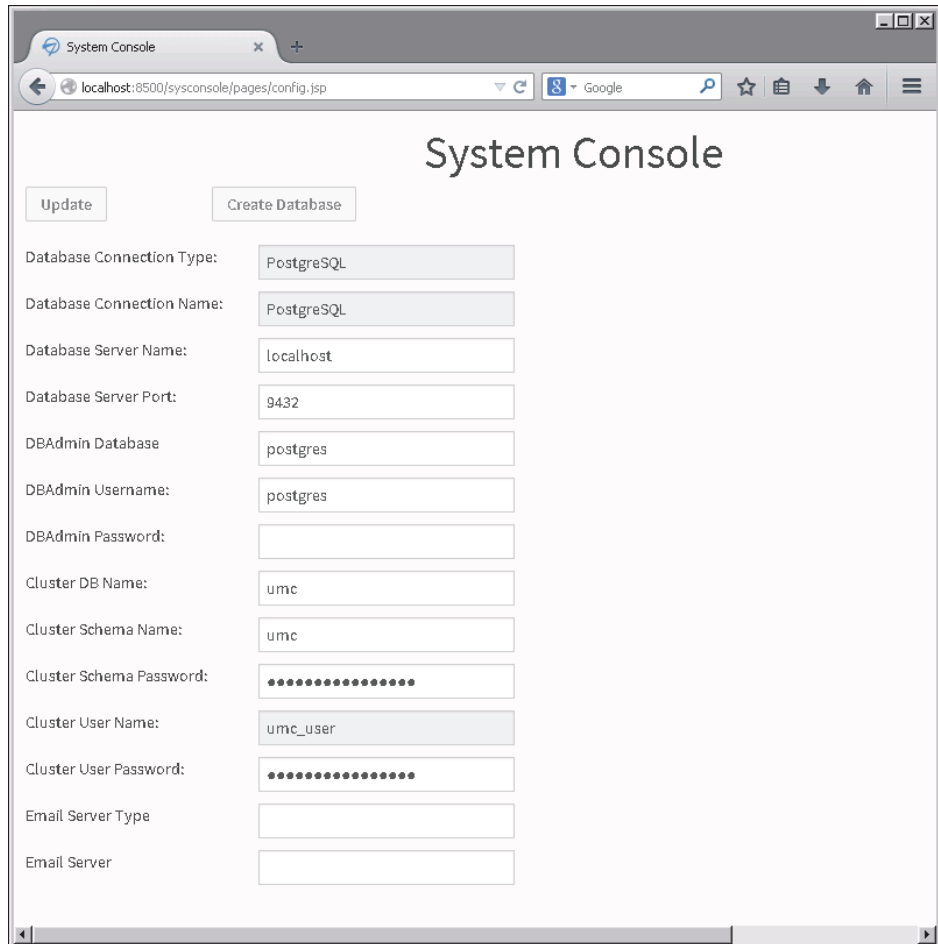
## How to create a PostgreSQL database for System Console

- 1 Log in to the System Console user interface, as described in “How to log in to the System Console user interface.”
- 2 On System Console, provide or accept the values for the following properties, as shown in Figure 1-29.
  - Database Connection Type  
Type of database to connect to. Must be PostgreSQL.
  - Database Connection Name  
Name of the database connection. Must be PostgreSQL.
  - Database Server Name

Name of the machine containing the database to which you want System Console to connect.

- Database Server Port  
Name of the port you want System Console to use to connect to the database.
- DBAdmin Database  
Name of the administrative connection database. This is the database that System Console connects to initially, before System Console creates the umc database and schema.
- DBAdmin Username  
Name of the superuser that logs in to the DBAdmin Database.
- DBAdmin Password  
Password for the superuser that logs in to the DBAdmin Database.
- Cluster DB name  
Name of the System Console metadata database. System Console creates this database when you choose Create Database. Accept the name umc.
- Cluster Schema name  
Name of the schema containing the System Console metadata. System Console creates this schema when you choose Create Database. Accept the name umc.
- Cluster Schema Password  
Type a password for the Cluster schema.
- Cluster User Name  
Name of the user with which System Console logs into the metadata database. Must be umc\_user. This user reads and writes to the System Console metadata database.
- Cluster User Password  
Type a password for the Cluster user.
- Email Server Type  
Type a name for the mail server. When System Console creates the umc database, System Console sends an e-mail notifying the system administrator if you specify values for Email Server Type and Email Server.
- Email Server  
Type the IP address or fully qualified domain name of the mail server. For example, type mailserver.companydomain.com.





**Figure 1-29** Specifying metadata database properties

**3** Choose Create Database.

#### **How to update the System Console metadata database**

- 1** Log in to the System Console user interface, as described in “How to log in to the System Console configuration user interface.”
- 2** On System Console you can update any of the following properties for an existing System Console metadata database:
  - Database Server Port
  - Cluster User Password
  - Email Server Type

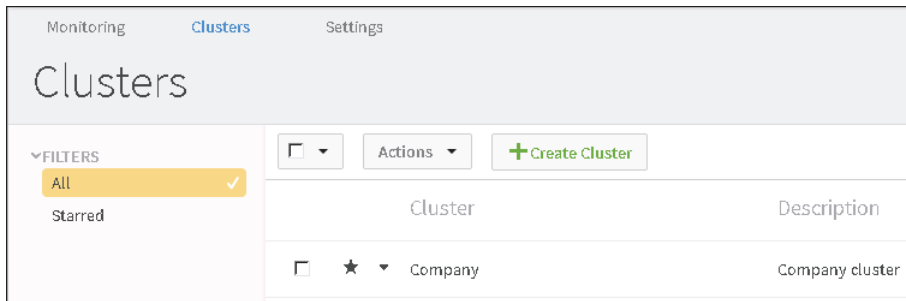
- Email Server
- 3 Make any desired changes and choose Update, as shown on Figure 1-29.

---

## About Monitoring

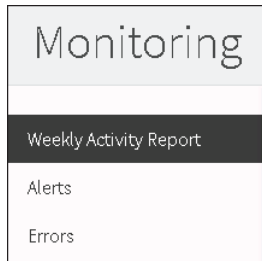
When the administrator logs in to System Console, System Console displays Monitoring. Monitoring consists of the following categories, which the administrator chooses from a side menu, as shown in Figure 1-31:

- Weekly Activity Report  
Displays the status of any starred cluster for each of the previous seven days. To star a cluster, choose Clusters. Then, on Clusters, select the star next to a cluster for which you want to view a Weekly Activity Report, as shown in Figure 1-30.



**Figure 1-30** Starring a cluster

- Alerts  
Displays an alert if the attribute that System Console is monitoring has reached the alert threshold. Alerts displays the following information for an alert:
  - Cluster  
Name of the cluster for which System Console generated the alert
  - Attribute Name  
Name of the attribute System Console is monitoring
  - Timestamp  
Date and time at which System Console generated the alert
  - Condition  
Specifies the threshold that when reached, triggers an alert



**Figure 1-31** Monitoring menu options

- **Current Value**  
Current value of the attribute System Console is monitoring
- **Object Type**  
Type of object System Console is monitoring, for example, SERVER
- **Process**  
Name of BIRT iHub process, if applicable
- **Server**  
Machine name of the cluster node
- **Volume**  
Current value of the attribute System Console is monitoring
- **Email Address**  
E-mail address to which iHub sends notification of an alert
- **Errors**  
Displays error messages a cluster generates. Errors displays the following information for an error:
  - **Cluster**  
Name of the cluster in which the error occurred
  - **Host Name**  
Machine name of the cluster node in which the error occurred
  - **Time Stamp**  
Date and time of the error
  - **PID**  
Process identification number
  - **Application**  
Name of the application

- Thread ID  
Identification number of the thread that was active when the error occurred
- Error Message  
Error message the error generated
- Error type  
Severity of the error

# Managing clusters

This chapter contains the following topics:

- About clusters
- Creating and configuring a cluster
- Editing an existing cluster
- Managing a cluster node
- Viewing the list of clusters
- Deleting clusters
- Viewing cluster resources
- Using the ControlConsole utility to free memory
- About BIRT iHub service and resource group properties
- Configuring an Apache web server for load balancing and proxying


---

## About clusters

The system administrator creates, configures, and manages clusters. Specifying a cluster allows a system administrator to rapidly configure and monitor Actuate applications deployed for use by an entity, such as a corporation or corporate division. A cluster consists of one or more computers, or cluster nodes, on which a BIRT iHub instance runs. A node configures itself using a template in `acsserverconfig.xml`, which is located on one node, in a shared configuration directory that all nodes in the cluster access.

In System Console, choose Clusters to perform the following tasks, as shown in Figure 2-1:

- Create a cluster
- Edit an existing cluster
- View and filter the list of clusters
- Delete clusters
- Select only starred clusters
- Star or unstar clusters
- Add a volume
- View log files
- View resource usage information



**Figure 2-1** Viewing Clusters

The following sections describe how to perform these tasks.

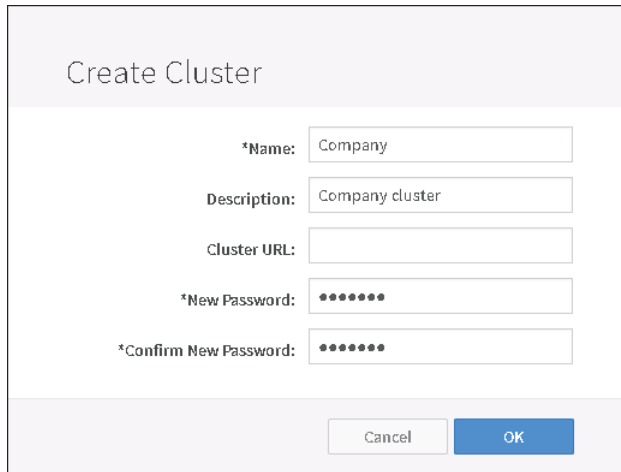
---

## Creating and configuring a cluster

The system administrator creates a cluster in System Console—Clusters. Then, the system administrator configures the cluster. The cluster must exist before the system administrator can perform any configuration tasks.

### How to create a cluster

- 1 Choose Create Cluster, as shown in Figure 2-1.



The screenshot shows a 'Create Cluster' dialog box. The title is 'Create Cluster'. Below the title are five input fields:

- \*Name:** A text box containing 'Company'.
- Description:** A text box containing 'Company cluster'.
- Cluster URL:** An empty text box.
- \*New Password:** A text box containing seven dots (•••••••).
- \*Confirm New Password:** A text box containing seven dots (•••••••).

At the bottom of the dialog are two buttons: 'Cancel' and 'OK'.

**Figure 2-2** Creating a basic cluster

- 2 In Create Cluster, set the following properties, as shown in Figure 2-2. A property name appearing with an asterisk (\*) next to the name is a required property.
  - **Name**  
Type a unique identifier for an cluster name, such as the company name.
  - **Description**  
Type a description for the cluster.
  - **Cluster URL**  
Optionally, type a URL for the cluster. The cluster URL specifies the location of a BIRT iHub proxy that performs load balancing by distributing requests among nodes. The proxy can be a server or a third-party load-balancing mechanism. Apache and Nginx are examples of third-party load-balancing solution providers. For information about using an Apache web server for load balancing and as a proxy for a BIRT iHub cluster, see “Configuring an Apache web server for load balancing and proxying,” later in this chapter.

- **New Password**  
Type a new password. Actuate recommends creating a password at least eight characters long, containing at least one lowercase letter, one uppercase letter, and one digit.
- **Confirm New Password**  
Type the new password again.

Choose OK.

System Console creates the cluster, and displays a message telling the system administrator that the cluster has been created and that the system administrator must add a cluster node for the cluster to be operational.

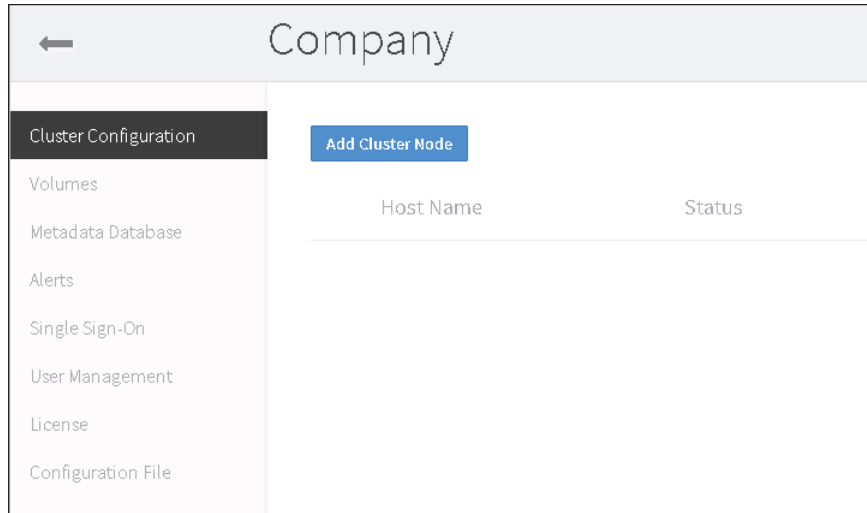
## **About the cluster configuration categories**

After System Console creates the cluster, System Console displays Cluster Configuration, which includes a side menu containing the following cluster configuration categories, as shown in Figure 2-3. The system administrator specifies property settings for each category to configure the cluster.

- **Cluster Configuration**  
Add a cluster node to BIRT iHub System. Completion of this task makes the cluster operational. The system administrator must complete this task to specify any other property settings for the cluster.
- **Volumes**  
Add a volume to the cluster.
- **Metadata Database**  
Specify the type of relational database management system (RDBMS) the cluster uses, such as PostgreSQL, or Oracle.
- **Alerts**  
Configure one or more alerts for the cluster. System Console monitors conditions and activity levels in the cluster. An alert is a notification triggered by a condition or an activity level crossing a particular threshold.
- **Single Sign-On**  
View or change SAML Identity Provider information for the cluster. View or change Service Provider information for the cluster. Add a Service Provider.
- **User Management**  
Specify settings for managing user authentication and authorization.
- **License**  
Update the license file for the cluster.



- Configuration File  
Update the shared configuration file for the cluster.



**Figure 2-3** Viewing menu of cluster configuration categories

## Adding cluster nodes to a cluster


This section demonstrates adding three nodes to a cluster named Company. The machine name of the first node the system administrator adds to the cluster is urup, the machine name of the second node is kozu, and the machine name of the third node the system administrator adds to the cluster is tupo. System Console and BIRT iHub are running on urup. urup also contains the shared configuration directory, which all nodes in the cluster access. The second and third nodes, kozu and tupo, each run a BIRT iHub instance. Neither kozu nor tupo run a System Console instance.

### How to add the first cluster node to a cluster

Before adding the first node, urup, to a cluster, the system administrator ensures that the logon account for the Actuate iHub service on the node has administrator privileges. See “Specifying a logon account for the Actuate iHub 3 service on a cluster node,” in Chapter 2, “Configuring a BIRT iHub cluster.”

After performing this task at the operating system level, the system administrator performs the following tasks in System Console:

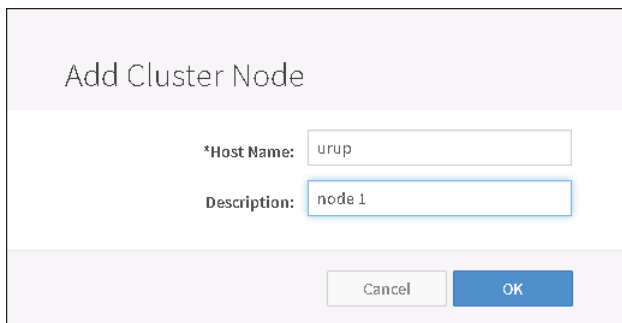
- 1 On Cluster Configuration, choose Add Cluster node, as shown in Figure 2-4.



**Figure 2-4** Choosing Add Cluster Node

2 On Add Cluster Node, set the following properties, as shown in Figure 2-5. A property name appearing with an asterisk (\*) next to the name is a required property.

- Host Name  
Type the cluster node computer name.
- Description  
Type a description for the node.

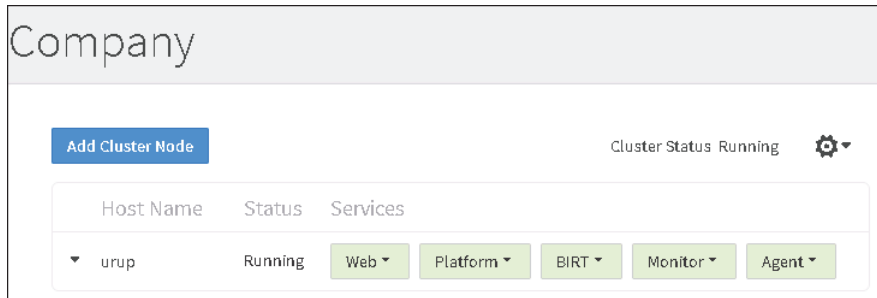


**Figure 2-5** Adding a cluster node

Choose OK.

System Console displays the following information about the cluster node, as shown in Figure 2-6:

- Host Name  
The machine name of the cluster node
- Status  
Status is either Running or Not Running
- Services  
The services running on the cluster node



**Figure 2-6** Viewing cluster node host name, status, and services

**How to add the second cluster node to a cluster and enable the default volume**

Before adding the second node, *kozu*, to the cluster, the system administrator performs the following tasks:

- On *urup*, the system administrator:
  - Creates a folder for the shared configuration directory and shares it.
  - Shares the folder containing the files for the out-of-the-box (OOTB) sample volume, Default Volume.

See “Creating the shared configuration directory,” and “Sharing the folders that all cluster nodes access,” in Chapter 2, “Configuring a BIRT iHub cluster.”

- On both *urup* and *kozu*, the system administrator:
  - Turns off the firewall.
  - Obtains the machine name and IP address, and pings each machine from the other machine to ensure the machines can communicate.

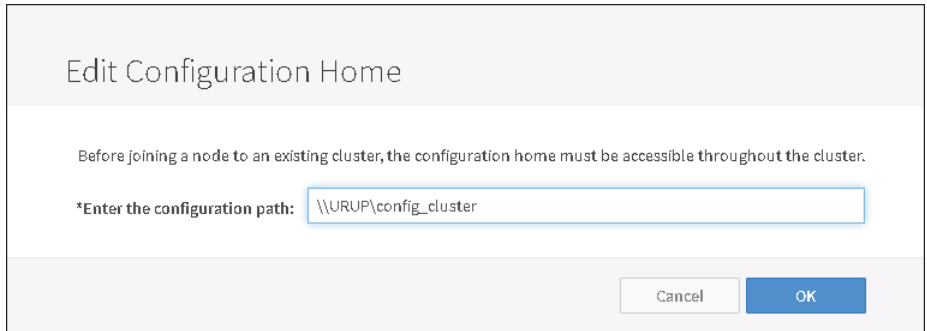
See “Configuring two nodes to communicate with each other,” in Chapter 2, “Configuring a BIRT iHub cluster.”

- On *kozu*, the system administrator ensures that the logon account for the Actuate iHub service on the node has administrator privileges. See “Specifying a logon account for the Actuate iHub 3 service on a cluster node,” in Chapter 2, “Configuring a BIRT iHub cluster.”

After performing these tasks at the operating system level, the system administrator performs the following tasks in System Console:

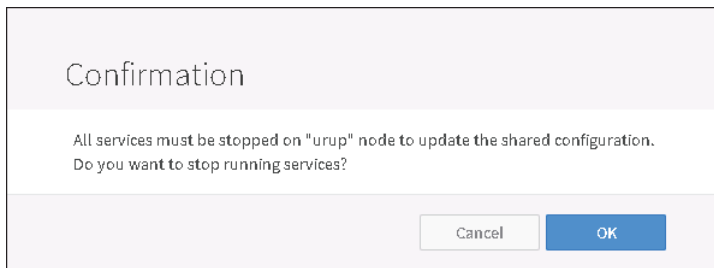
- 1 On Cluster Configuration, choose Add Cluster Node, as shown in Figure 2-6.
- 2 On Edit Configuration Home, in Enter the configuration path, type the path to the shared configuration directory, using UNC format, as shown in Figure 2-7. UNC format supports all nodes in the cluster finding the shared configuration directory. The path you type is the path that appears as the Network Path in Properties—Sharing for the shared configuration directory. In this example,

the shared configuration directory, `config_cluster`, is on a machine named URUP. Choose OK.



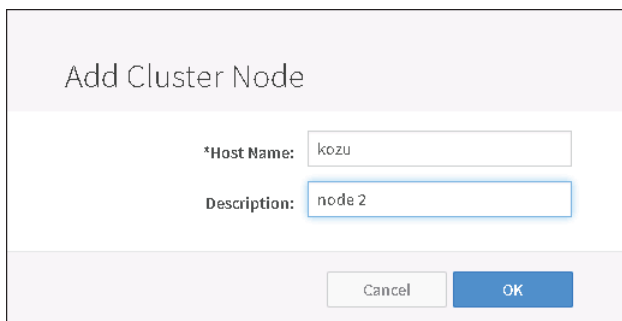
**Figure 2-7** Specifying the path of the shared configuration directory

- 3 On Confirmation, choose OK to stop the services on the previously added cluster node, `urup` in this example, as shown in Figure 2-8.



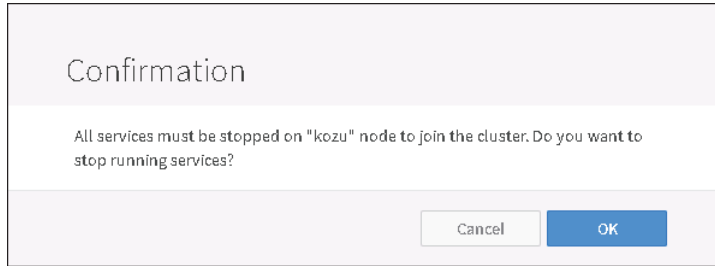
**Figure 2-8** Stopping the services on previously added cluster node

- 4 On Add Cluster Node, specify the machine name of the cluster node you are adding and optionally, a description, as shown in Figure 2-9.

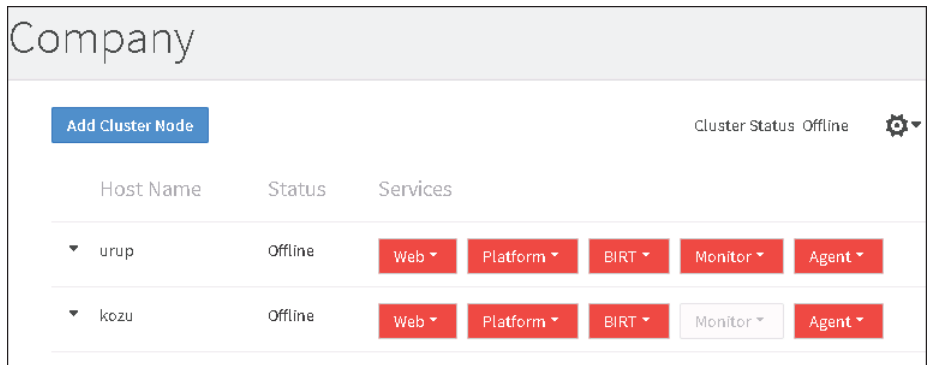


**Figure 2-9** Specifying name and description of node you are adding

- 5 On Confirmation, choose OK to stop the services on the node you are adding to the cluster, as shown in Figure 2-10.

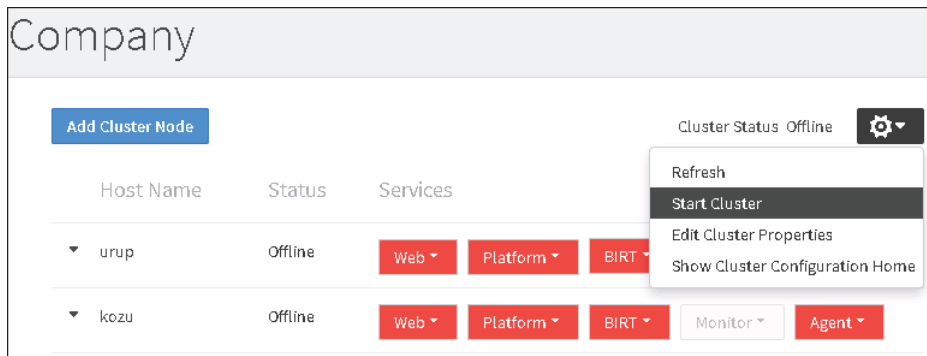


**Figure 2-10** Stopping the services on the node you are adding to the cluster. System Console adds the second node to the cluster, as shown in Figure 2-11. By default, the Monitor service runs only on the node having the shared configuration directory.

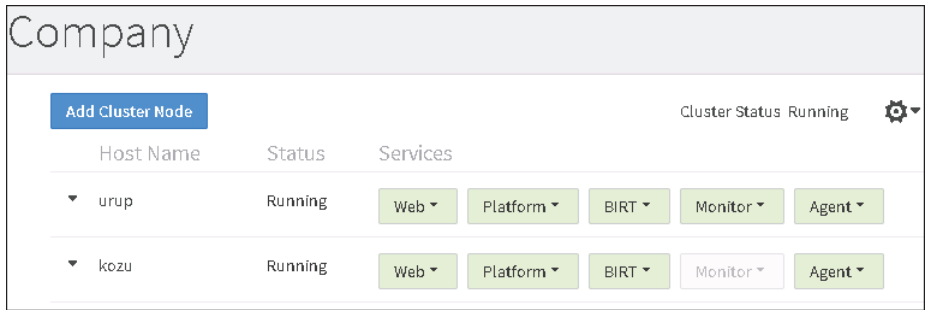


**Figure 2-11** Viewing the second node added to the cluster

- 6 Choose Start Cluster from the Manage Cluster menu, as shown in Figure 2-12. Then, choose Refresh from this menu to update the status of the services during the Start Cluster operation. Wait until all services that are red turn green before proceeding to the next step, as shown in Figure 2-13.




**Figure 2-12** Choosing to start the cluster



**Figure 2-13** Viewing the started services on both nodes


- 7 Choose Volumes from the side menu. Default Volume shows a status of 'Error'. Left-click the arrowhead icon in the first Storage Status box for Default Volume and choose Set Read Only, as shown in Figure 2-14.



**Figure 2-14** Choosing Set Read Only for Default Volume

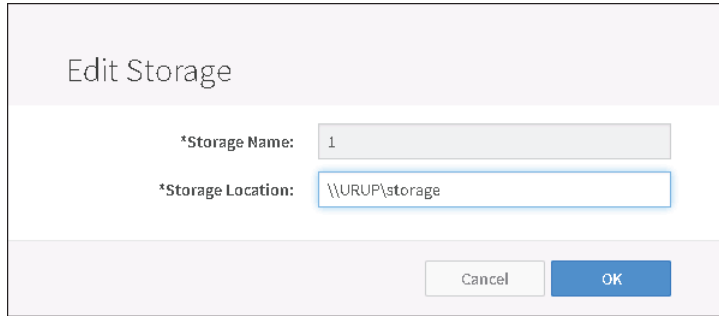
On Confirmation, choose OK to confirm that you want to change the Default Volume state to Read only.

- 8 On Volumes, left-click the arrowhead icon in the first Storage Status box for Default Volume and choose Edit, as shown in Figure 2-15.



**Figure 2-15** Choosing to edit Default Volume storage


- 9 On Edit Storage, in Storage Location, type the path to the Default Volume storage folder, storage, using UNC format, as shown in Figure 2-16. UNC format supports all nodes in the cluster finding this folder. The path you type is the path that appears as the Network Path in Properties—Sharing for the storage folder after sharing it. In this example, the Default Volume storage folder is on a machine named URUP. Choose OK.



The screenshot shows a dialog box titled "Edit Storage". It contains two input fields: "\*Storage Name:" with the value "1" and "\*Storage Location:" with the value "\\URUP\storage". At the bottom right, there are "Cancel" and "OK" buttons.

**Figure 2-16** Specifying the Default Volume storage folder

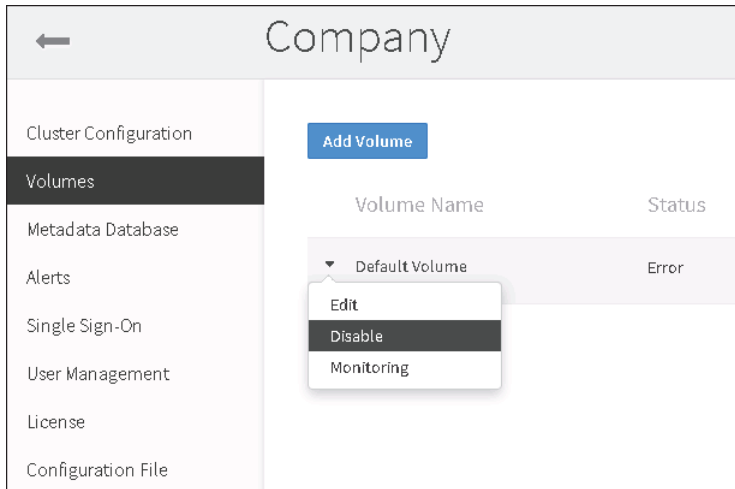
- 10 On Volumes, left-click the arrowhead icon in the first Storage Status box for Default Volume and choose Set Read/Write, as shown in Figure 2-17.



**Figure 2-17** Setting Default Volume to Read/Write status

On Confirmation, choose OK to confirm that you want to change the Default Volume state to Read/Write.

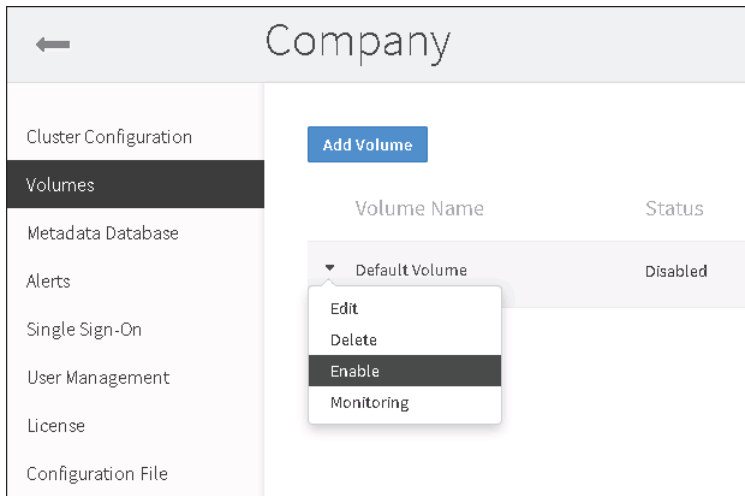
- 11 On Volumes, left-click the arrowhead icon next to Default Volume and choose Disable, as shown in Figure 2-18.



**Figure 2-18** Disabling Default Volume

On Confirmation, choose OK to confirm that you want to disable Default Volume.

- 12** On Volumes, left-click the arrowhead icon next to Default Volume and choose Enable, as shown in Figure 2-19.

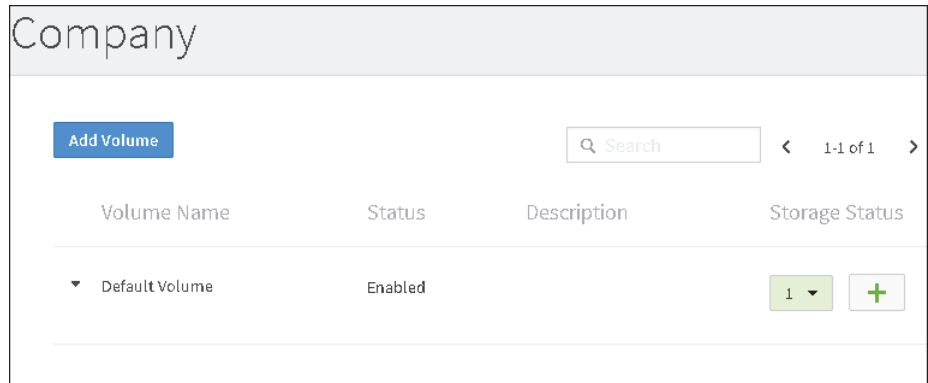


**Figure 2-19** Enabling Default Volume

On Confirmation, choose OK to confirm that you want to enable Default Volume.



Default Volume is enabled and ready for use, as shown in Figure 2-20.



**Figure 2-20** Viewing Enabled status of Default Volume

### How to add a third or subsequent node

Before adding the third node, tupo, or any subsequent node, to the cluster, the system administrator performs the following tasks:

- On tupo, the system administrator:
  - Turns off the firewall.
  - Obtains the machine name and IP address.
  - Ensures that the logon account for the Actuate iHub service on the node has administrator privileges.
- On both urup and tupo, the system administrator pings each machine from the other machine to ensure the machines can communicate.

For details on how to perform these tasks, see “Configuring two nodes to communicate with each other” and “Specifying a logon account for the Actuate iHub 3 service on a cluster node” in Chapter 2, “Configuring a BIRT iHub cluster.”

After performing these tasks at the operating system level, the system administrator performs the following tasks in System Console:

- 1 On Cluster Configuration, choose Add Cluster Node.
- 2 On Add Cluster Node, specify the machine name of the cluster node you are adding and optionally, a description, as shown in Figure 2-21.

Add Cluster Node

\*Host Name:

Description:

**Figure 2-21** Specifying name and description of node you are adding

- 3 On Confirmation, choose OK to confirm that you want to stop the services on tupo, as shown in Figure 2-22.

Confirmation

All services must be stopped on "tupo" node to join the cluster. Do you want to stop running services?

**Figure 2-22** Stopping services on the third node

System Console adds the third node to the cluster, as shown in Figure 2-23. By default, the Monitor service runs only on the node having the shared configuration directory.

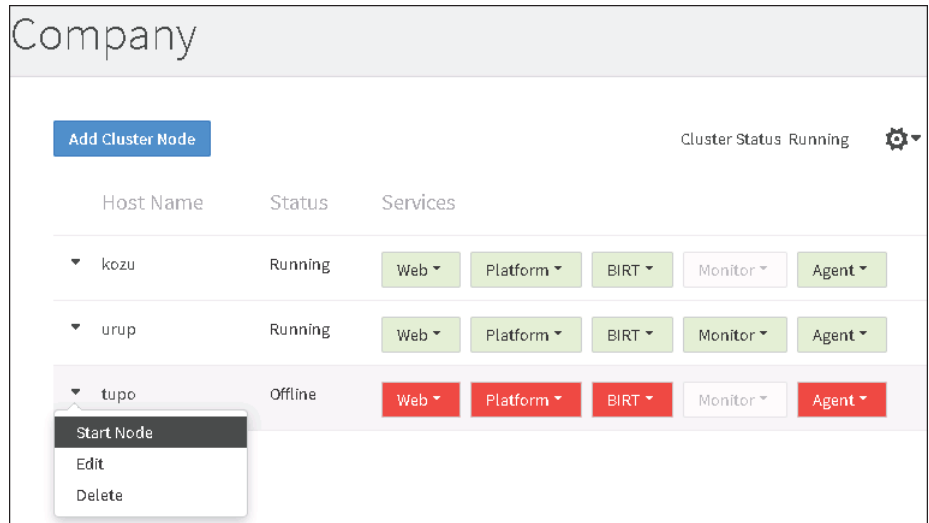
Company

Cluster Status Running

Host Name	Status	Services
▼ kozu	Running	Web ▼ Platform ▼ BIRT ▼ Monitor ▼ Agent ▼
▼ urup	Running	Web ▼ Platform ▼ BIRT ▼ Monitor ▼ Agent ▼
▼ tupo	Offline	Web ▼ Platform ▼ BIRT ▼ Monitor ▼ Agent ▼

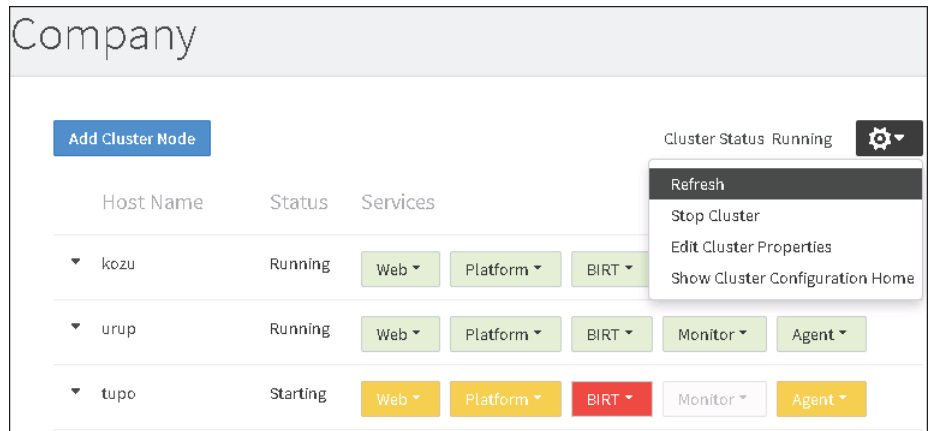
**Figure 2-23** Viewing the second node added to the cluster

- 4 Left-click the arrowhead icon next to tupo and choose Start Node, as shown in Figure 2-24. Then, choose Refresh from this Manage Cluster menu to update the status of the services during the Start Node operation. When all the services that are red turn green, the node is ready for use.

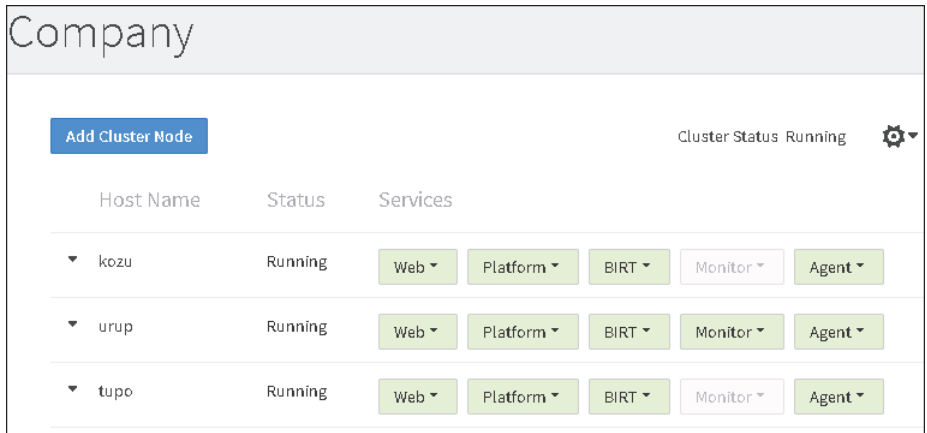


**Figure 2-24** Choosing to start the cluster

- 5 Choose Refresh from the Manage Clusters menu to update the status of the services during the Start Node operation, as shown in Figure 2-25. When the services display green, the node is ready for use, as shown in Figure 2-26. By default, the Monitor service runs only on the node containing the shared configuration directory, urup, in this example.



**Figure 2-25** Refreshing the status of services on the third node



**Figure 2-26** Viewing the running services in the cluster


## Understanding Cluster Configuration

In Cluster Configuration, the system administrator adds a cluster node to the cluster. Additionally, Cluster Configuration supports management tasks such as starting, stopping, and editing the properties of the following:

- The entire cluster
- An individual cluster node
- A service running on a cluster node

### Performing management tasks for the entire cluster

The system administrator chooses the cog-shaped icon to access the Manage Cluster menu, as shown in Figure 2-27.



**Figure 2-27** Accessing the Manage Cluster menu

The Manage Cluster menu consists of the following options:

- Refresh  
Refreshes the status of the services running on all cluster nodes.

- Stop or Start Cluster  
Stops or Starts all nodes in the cluster. If the cluster is running, or online, Stop Cluster displays in the Manage Cluster menu. If the Cluster is stopped, or offline, Start Cluster displays in the Manage Cluster menu.
- Edit Cluster Properties  
Displays Edit Cluster Properties. The system administrator can change any of the following cluster properties. Choose Stop Cluster to stop the cluster before changing Cluster URL.
  - Name
  - Description
  - Cluster URL
  - Password

After making any cluster property changes choose OK. If you changed the Cluster URL, choose Start Cluster to start the cluster after choosing OK.

- Show Cluster Configuration Home  
Displays the location of the shared configuration folder that the AC\_CONFIG\_HOME element specifies in the acpmdconfig.xml file on the cluster node, in UNC format. For example, the following syntax specifies the path to the shared configuration directory used in “How to add the second cluster node to a cluster and enable the default volume,” earlier in this chapter:

```
\\urup\config_cluster
```

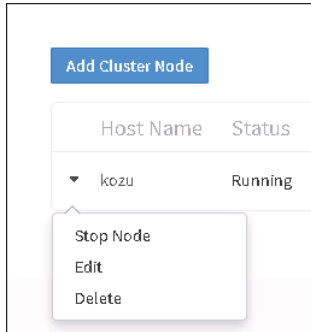
where urup is the name of the machine containing the shared configuration directory.

In a default BIRT iHub installation on Windows, performed using the installer, in which the install folder is C:\Actuate, the path AC\_CONFIG\_HOME specifies is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
  \config_cluster
```

## Performing management tasks for an individual cluster node

The system administrator chooses the arrowhead icon next to a cluster node name to access the cluster node menu, as shown in Figure 2-28.




**Figure 2-28** Accessing the cluster node menu

The following list describes the options on the cluster node menu:

- Stop or start node  
Stops or Starts the cluster node. If the cluster node is running, or online, Stop Node displays in the cluster node menu. If the cluster node is stopped, or offline, Start Node displays in the cluster node menu.
- Edit  
Displays Edit Cluster Node. The system administrator can change either of the following properties.
  - Host Name
  - Description
- Delete  
Deletes the node from the cluster.

## Performing management tasks for a service running on a cluster node

The system administrator chooses the arrowhead icon next to a service name to access the service menu. For example, Figure 2-29 shows the menu for the Web service.



**Figure 2-29** Accessing a service menu

The following list describes the options on any service menu except the BIRT menu. For more information about the BIRT service, see “About the BIRT service,” later in this chapter.

- **Stop or Start Service**  
Stops or Starts the service. If the service is running, the color of the icon for the service is green, and Stop Service displays in the service menu. If the service is stopped, the color of the icon for the service is red, and Start Service displays in the service menu.

- **Edit**  
Displays Edit <service name>. For example, when the system administrator chooses to edit the Web service, System Console displays Edit Web.

For each service, Edit <service name> displays the Startup Mode, Process Name, and Java Arguments properties, as shown in Table 2-1. The system administrator can change the Startup Mode and the Java Arguments properties. A property name appearing with an asterisk (\*) next to the name is a required property.

If you modify the Java heap size argument for a service, make sure not to specify a size that exceeds the amount of RAM on the node. On some Linux platforms, the LMServer process may encounter an error if the Java heap size you specify exceeds the amount of RAM available on the node.

**Table 2-1** Cluster node service properties

Service name	Startup mode	Process name	Java arguments
Web	Auto Start, Manual, Disable	ihubservletcontainer	-Xms256m -Xmx1024m -XX:PermSize=64M -XX:MaxNewSize=256m -XX:MaxPermSize=128m -Djava.awt.headless=true com.actuate.server.embeddedtomcat.EmbeddedTomcat
Platform	Auto Start, Manual, Disable	ihub	-Xms256m -Xmx1024m -XX:MaxPermSize=128m "-Djava.library.path=C:/Actuate/BIRTiHubVisualization/modules/BIRTiHub/iHub/bin" com.actuate.iserver.server.Server-ppid 1234
Platform	Auto Start, Manual, Disable	ihubc	-Spmd -Jux-team-win-vm -YCluster

*(continues)*

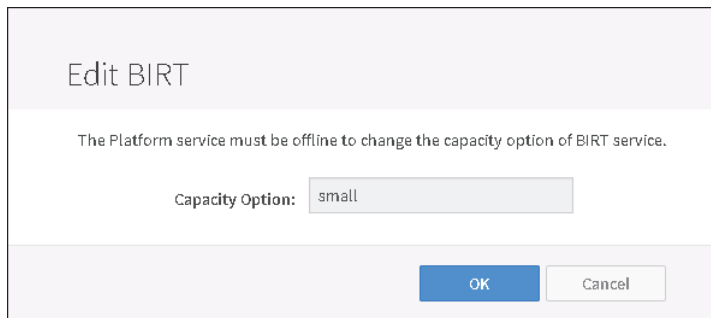
**Table 2-1** Cluster node service properties (continued)

Service name	Startup mode	Process name	Java arguments
Monitor	Auto Start, Manual, Disable	LMServer	-Xms386m -Xmx8g "- Dlog4j.configuration=file:C:/Actuate /BIRTiHubVisualization/modules /BIRTiHub/iHub/etc/lmservice- log4j.properties" "-Djava.library.path= C:/Actuate/BIRTiHubVisualization/modules /BIRTiHub/iHub/bin" com.actuate.lmservice.server.LMServer -p
Agent	Auto Start, Manual, Disable	LSTailer	-Xms64m -Xmx256m "- Dlog4j.configuration=file:C:/Actuate /BIRTiHubVisualization/modules /BIRTiHub/iHub/etc/lmservice- log4j.properties" "-Djava.library.path= C:/Actuate/BIRTiHubVisualization /modules/BIRTiHub/iHub/bin" com.actuate.lmservice.logging.logtailer .ProducerAgent

### About the BIRT service

Choosing the arrowhead icon next to BIRT displays a menu containing one option, Edit.

When the system administrator chooses Edit on the menu for BIRT, System Console displays Edit BIRT, as shown in Figure 2-30.



**Figure 2-30** Editing the BIRT service

Changing the Capacity Option changes the server configuration template that this cluster node uses to configure itself. AC\_CONFIG\_HOME \acserverconfig.xml contains the server configuration templates. The names of the default server configuration templates that acserverconfig.xml contains are



small, medium, large, and disable. Stop the Platform service before changing the name for Capacity Option in Edit BIRT.

For more information on server configuration templates, see “About BIRT iHub service and resource group properties,” later in this chapter.

## Adding a volume

A cluster can contain one or more volumes. For each volume, there is one database schema and one or more storage areas. The metadata database contains volume metadata, such as user and user group information. The storage area or areas contain volume consumable data, such as BIRT document content. When adding a volume, properties the system administrator specifies include schema name, storage area or areas, and the database user and password with which to connect to the metadata database.

A single BIRT iHub cluster can use only one security mechanism. For example, if the system administrator wants to use iHub User Management (default) as the user management setting for one volume and LDAP Adapter as the user management setting for a second volume, the system administrator must create a cluster for each volume. For more information on the user management setting, see “Configuring User Management,” later in this chapter.

Actuate recommends enabling e-mail notification before creating a new volume if you have not already enabled e-mail notification. You need e-mail notification enabled to successfully perform the following tasks:

- Create a volume, if also specifying an e-mail address for the volume administrator
- Edit an existing volume and selecting to reset the password

If you have not enabled e-mail notification, System Console displays an error message and does not allow you to complete these tasks. For information on enabling e-mail notification, see “Enabling e-mail notification,” later in this chapter.

This section demonstrates adding an example volume named `sales_volume` in the process of creating a two-node cluster.

### How to add a volume

- 1 Actuate recommends enabling e-mail notification. For information on enabling e-mail notification, see “Enabling e-mail notification,” later in this chapter.

- 2 Create a new folder at the location where you want to store the volume data. For example, create a new folder in AC\_SHARED\_HOME named sales\_storage. In a default BIRT iHub installation on Windows, performed using the installer, in which the install folder is

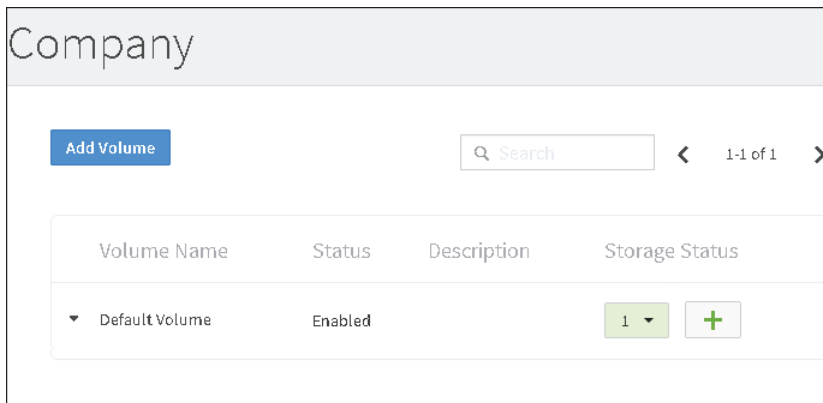
C:\Actuate, the path for AC\_SHARED\_HOME\sales\_storage is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
\sales_storage
```

Do not reuse a storage location for a new volume. For example, if AC\_SHARED\_HOME\sales\_storage was the storage folder for a previously existing volume, create a storage location for the new volume that has a path other than AC\_SHARED\_HOME

\sales\_storage. System Console allows using a subfolder of AC\_SHARED\_HOME\sales\_storage, such as AC\_SHARED\_HOME\sales\_storage\sales\_storage\_2.

- 3 On Volumes, choose Add Volume, as shown in Figure 2-31.



**Figure 2-31** Choosing Add Volume

- 4 Configure the following properties on Add Volume. Figure 2-32 shows the property values for an example volume, sales\_volume. An asterisk (\*) next to the property name means the property is required.
  - Volume Name  
Type a name for the volume.
  - Description  
Type a description for the volume.
  - Volume Administrator Email  
Type the e-mail address of the volume administrator. When you create a volume, System Console sends a notification e-mail containing the volume password to this address if you have enabled e-mail notification. For more

information, see “Enabling e-mail notification,” later in this chapter. If you leave Volume Administrator Email blank, BIRT iHub does not create a password for accessing the new volume in Visualization Platform. The BIRT iHub default user, Administrator, can log in to Visualization Platform to access the new volume without using a password. Then, in Visualization Platform, the administrator can choose My Profile and create a new password for accessing the volume.

- **Schema Name**  
Type a name for the volume schema that is 30 characters or less. BIRT iHub creates the volume and the volume schema at the same time.
- **Create New Schema**  
Select this property except under either of the following conditions:
  - You have already populated the schema using the Volume Data Store Administrator utility.
  - You are adding a volume for which the schema is already populated and the storage location already contains files.
- **Tablespace**  
Type the name of a tablespace for the volume schema. Alternatively, leave Tablespace blank to use the default tablespace.
- **DBA User**  
Type the name of the PostgreSQL superuser, postgres.
- **DBA password**  
Type the PostgreSQL superuser password. By default, the password is postgres.
- **Storage Location**  
Type the path of the volume storage folder you created in step 1.
- **Organization ID**  
Type an alphanumeric character string for the Organization ID. The LDAP adapter and RSSE implementation use the Organization ID to filter users and user groups. Alternatively, leave Organization ID blank. For more information on Organization ID, see “About managing volume access by users and user groups when using LDAP,” later in this chapter.
- **Encryption Key for Storage**  
Type the name of the Encryption key. Alternatively, leave Encryption Key blank.

On Add Volume, choose OK.

**Add Volume**

\*Volume Name:

Description:

Volume Administrator Email:

\*Schema Name:

Create New Schema:

Tablespace:

\*DBA User:

\*DBA Password:

\*Storage Location:

Organization ID:

Encryption Key for Storage:

**Figure 2-32** Adding a volume

If e-mail notification is enabled, BIRT iHub sends an e-mail notifying the volume administrator that BIRT iHub has created the volume. The e-mail contains the password with which to log into Visualization Platform to access sales\_volume, as shown in Figure 2-33.

```

Actuate iHub volume: "sales_volume" has been created

Login URL: http://URUP:8700/iportal/login.jsp?volume=sales\_volume

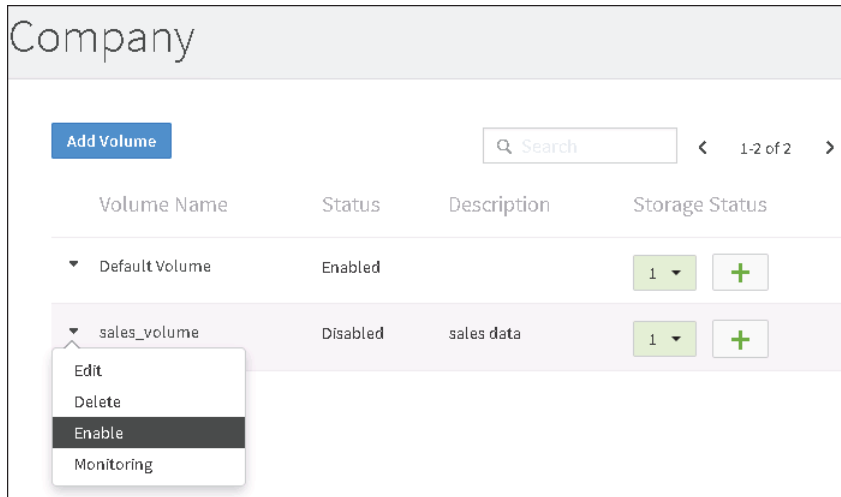
Username: Administrator

Password: kFbfMDM69u

```

**Figure 2-33** Viewing the notification e-mail that the volume is created

- 5 On Volumes, left-click the arrowhead icon next to the new volume name and choose Enable, as shown in Figure 2-34.



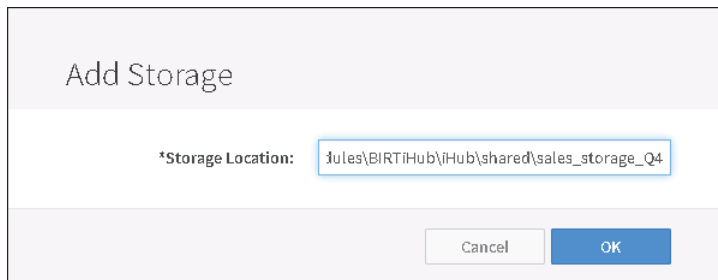
**Figure 2-34** Viewing the new volume in the list on Volumes

## Adding or updating a storage location

The system administrator can add a storage location for a volume. A single volume can use a maximum of 10 storage locations. The system administrator can also change the storage location for an existing volume.

### How to add a storage location for an existing volume

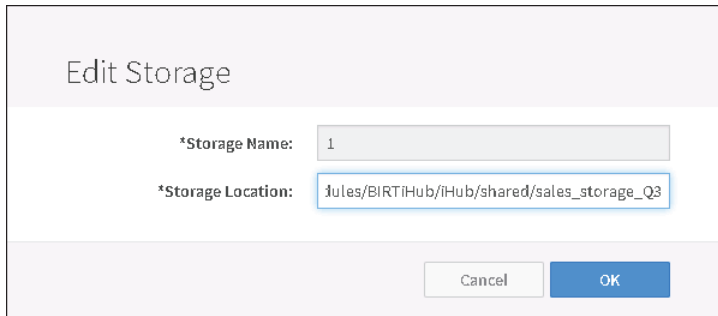
- 1 Create a new folder at the location where you want to add storage. Do not use a storage location that a volume has used previously. The path of the storage location must be new.
- 2 On Volumes, in the Storage Status column, left-click the plus sign (+) in the row containing the name of the volume for which you want to add storage.
- 3 In Add Storage, specify the new storage location in Storage Location, as shown in Figure 2-35. Choose OK.



**Figure 2-35** Adding a storage location for a volume

## How to change the storage location for an existing volume

- 1 Create a new folder at the location to which you want to change the storage location. Do not use a storage location that a volume has used previously. The path of the storage location must be new.
- 2 On Volumes, in the Storage Status column, left-click the arrowhead icon in the row containing the name of the volume for which you want to change the storage location, and choose Set Read Only.
- 3 On Edit Storage, in Storage Location, specify the path to the new storage location you created in step 1, as shown in Figure 2-36. Choose OK.



The screenshot shows a dialog box titled "Edit Storage". It has two input fields. The first is labeled "\*Storage Name:" and contains the text "1". The second is labeled "\*Storage Location:" and contains the text "fules/BIRTIHub/iHub/shared/sales\_storage\_Q3". At the bottom right of the dialog, there are two buttons: "Cancel" and "OK".

**Figure 2-36** Changing the storage location for a volume

- 4 On Volumes, in the Storage Status column, left-click the arrowhead icon in the row containing the name of the volume for which you changed the storage location, and choose Set Read/Write

## Understanding the volume menu

Left-click the arrowhead icon next to a volume name to display a menu containing the following options:

- Edit  
Supports changing the following volume properties:
  - Description
  - Organization ID
  - Encryption Key for Storage
- Delete  
Deletes the volume. Delete is a menu option only when the volume is offline.
- Enable or Disable  
Brings the volume online and takes it offline. If the status of the volume is Enabled, the menu option is Disable. If the status of the volume is Disabled, the menu option is Enable.

- **Monitoring**  
Displays a link named Server Resource. Choose Server Resource to open a new browser window, in which System Console uses Actuate Viewer to display a graph showing the last 48 hours of activity on this volume for each of the following statistics.
  - Response Time (milliseconds)
  - Number of Alerts

## Selecting the metadata database type

When the system administrator runs the BIRT iHub installation program, BIRT iHub installs the out-of-the-box (OOTB) PostgreSQL database to contain volume metadata. By default, Metadata Database displays the properties for the OOTB PostgreSQL database. The database type for this database is ActuatePostgreSQL. Figure 2-37 shows the following properties for the OOTB PostgreSQL database, installed on a machine named urup. An asterisk (\*) next to the property name means the property is required.

- **Database server**  
The host name of the machine containing the database.
- **Database port**  
The default port number for the OOTB PostgreSQL database is 8432.
- **Database name**  
The name of the database.
- **Encryption Method**
  - **requestSSL**  
BIRT iHub encrypts the login request and data using SSL. If the database server does not support SSL, the driver establishes an unencrypted channel.
  - **SSL**  
BIRT iHub performs SSL certificate verification.
  - **noEncryption**  
The channel between BIRT iHub and the metadata database passes unencrypted data.
  - **Username**  
The database user name.
  - **Password**  
The database user name password.

- Test Connection

Choose to verify that the BIRT iHub system can successfully connect to the metadata database.

The screenshot shows a configuration dialog box for the PostgreSQL metadata database. It contains the following fields and controls:

- \*Database Type: A dropdown menu set to "ActuatePostgreSQL" and a "Test Connection" button.
- \*Database Server: A text input field containing "URUP".
- \*Database Port: A text input field containing "8432".
- \*Database Name: A text input field containing "ihub".
- Encryption Method: A dropdown menu set to "noEncryption".
- \*Username: A text input field containing "ihub".
- \*Password: A text input field with masked characters (dots).
- Buttons: "Cancel" and "Save".

**Figure 2-37** Viewing OOTB PostgreSQL metadata database properties

Accept the default values for the ActuatePostgreSQL database type. Alternatively, choose a different database type to configure the properties for a pre-existing third-party database for storing volume metadata. In database type, select one of the following relational database management systems (RDBMS):

- PostgreSQL
- Oracle

Depending on the database type selected, provide the following database information:

- PostgreSQL
  - Database server  
Type the host name of the machine containing the database, such as localhost.
  - Database port  
Type a port number, or accept the default value, 8432 for out-of-the-box (OOTB) ActuatePostgreSQL, or 5432 for a pre-existing PostgreSQL database.
  - Database name  
Type a name for the database.
  - Encryption Method



- requestSSL  
BIRT iHub encrypts the login request and data using SSL. If the database server does not support SSL, the driver establishes an unencrypted channel.
- SSL  
BIRT iHub performs SSL certificate verification.
- noEncryption  
The channel between BIRT iHub and the metadata database passes unencrypted data.

- Schema name  
Type the name of the volume schema.

- Username  
Type the database user name.

- Password  
Type the database user name password.

After setting these options, choose Test Connection to verify that System Console can successfully connect to the database.

Choose Save and continue to Alerts.

- Oracle

- Database Server  
Type the host name of the machine containing the database, such as localhost.

- Database Port  
Type a port number, or accept the default value, 1521.

- Service Name  
Type a valid service name, such as orcl.actuate.com, that identifies the Oracle database server on which you want to install the volume metadata. Do not use just the system identifier (SID). Provide the complete reference to the server, including the domain. When using a service name, leave Tns Server Name and Tns Names File blank. When using a Transparent Network Substrate (TNS) service, leave service name blank.

- TNS Name  
Type the host name of the machine containing the TNSNAMES.ORA file if leaving Service name blank.

- TNS Names File  
Type the path to the TNSNAMES.ORA file if leaving Service name blank.

- Encryption Method
  - SSL  
BIRT iHub performs SSL certificate verification.
  - noEncryption  
The channel between BIRT iHub and the metadata database passes unencrypted data.
- Schema name  
Type the name of the volume schema.
- Username  
Type the database user name.
- Password  
Type the database user name password.

After setting these options, choose Test Connection to verify that System Console can successfully connect to the database.

Choose Save and continue to Alerts.

## Configuring alerts

System Console monitors a range of activity, conditions, and resources in a BIRT iHub System. An attribute identifies a monitored item. The system administrator can create an alert for any system attribute. Alerts supports the system administrator performing the following operations:

- Viewing the list of alerts
- Adding an alert
- Editing an alert
- Disabling and enabling an alert
- Deleting an alert

The following sections describe these operations.

### Viewing the list of alerts

View the list of alerts by choosing Alerts from the side menu, as shown in Figure 2-38. An alert contains the following information:

- Alert name  
Name of the alert.
- Attribute  
Name of the attribute identifying the item BIRT iHub monitors.

- **Condition**  
Condition that determines whether a monitored item reaches the alert threshold.
- **Threshold**  
Limit that when met, triggers an alert.
- **Enable**  
True if the alert is enabled, false if the alert is disabled.
- **Email**  
E-mail address to send notification of an alert.
- **Message**  
Message System Console sends when an alert occurs.

Set alerts for this cluster

You may set multiple alerts for the cluster. For each alert, select what feature you want the system to monitor (the counter), set a threshold and a value that will trigger the alert, and specify who should receive an email notice when the alert is triggered. Alerts are sent to the System Administrator's email address by default; you may delete this address and add other addresses (separated by commas) if different individuals should be informed when the cluster crosses a particular threshold.

[Add Alert](#)

Alert Name	Attribute	Condition	Threshold	Enable	Email	Message
▼ Volume Status Alert	Volume status	=	OFFLINE	true		Threshold limit reached.
▼ Server Status Alert	Server status	=	OFFLINE	true		Threshold limit reached.
▼ Server Used CPU Perc Alert	Percent of server CPUs used	>=	90	true		Threshold limit reached.
▼ Server Used RAM Perc Alert	Percent of server RAM used (MB)	>=	90	true		Threshold limit reached.

**Figure 2-38** Viewing the list of alerts

## Adding an alert

When adding an alert, the system administrator selects an attribute name from a list, and sets a value, or threshold, that when reached, causes System Console to trigger an alert. The alert displays on Monitoring, and System Console sends an e-mail to the e-mail address the system administrator specifies. The e-mail notifies the recipient that the attribute for the item System Console is monitoring has met the specified threshold.

The value for Threshold that the system administrator specifies for most Alert attributes is a number from 0 (zero) to 100. For these Alert attributes, the

administrator can specify one of the following values for determining whether the condition which triggers an alert has been met:

- equal to (=)
- greater than (>)
- greater than or equal to (>=)
- less than (<)
- less than or equal to (<=)

For the remainder of the Alert attributes, the administrator specifies a string value for Threshold and a condition value of equal to (=).

Table 2-2 displays the Condition and Threshold values that the administrator can specify for Alert attributes.

**Table 2-2** Alert attribute Condition and Threshold values

Alert attribute name	Threshold value data type	Permissible values for Condition	Permissible values for Threshold
Volume status	String	=	ONLINE, OFFLINE, ERROR
Server status	String	=	ONLINE, OFFLINE
Server needs restart	String	=	YES, NO
Integration service status on server	String	=	ONLINE, OFFLINE
Factory service status on server	String	=	ONLINE, OFFLINE
View service status on server	String	=	ONLINE, OFFLINE
Encyclopedia service status on server	String	=	ONLINE, OFFLINE
All other Alert attributes	Numeric	=, >, >=, <, <=	Any number from 0 through 100

This following section demonstrates adding an alert on the system attribute named Percent of server RAM used (MB).

#### How to add an alert

- 1 Choose Alerts from the Clusters side menu.

- 2 Choose Add Alert.
- 3 On Add Alert, perform the following tasks, as shown in Figure 2-39. An asterisk (\*) next to the property name means the property is required.
  - 1 In Attribute Name, select an attribute.
  - 2 In Condition, select a condition by which System Console determines whether the monitored item has reached the threshold.
  - 3 In Threshold, specify a value that triggers an alert when reached.
  - 4 In Email, specify an email address where System Console sends notification of an alert. You must enable e-mail notification. For more information, see “Enabling e-mail notification,” later in this chapter.
  - 5 In Message, type a message to display on Monitoring and to include in the notification e-mail when an alert is triggered, such as ‘Number of jobs running on the volume has reached the specified limit’.
  - 6 In Alert Name, type a name for the alert. Choose OK.

**Figure 2-39** Adding an alert

## Enabling e-mail notification

The system administrator modifies `acserverconfig.xml`, the configuration file that all cluster nodes share, adding properties supporting e-mail notification when the following events occur:

- A scheduled job in Visualization Platform completes.
- The system administrator creates another system administrator.
- The system administrator creates a volume.
- System Console triggers an alert.

## How to enable e-mail notification

- 1 On Clusters, choose to edit the cluster for which you want to enable e-mail notification.
- 2 Stop the cluster by performing the following steps:
  - 1 Choose Cluster Configuration from the side menu.
  - 2 On Cluster Configuration, left-click the cog icon and choose Stop Cluster from the Manage Cluster menu.
  - 3 Choose Refresh from the Manage Cluster menu. When all the services icons have turned red, continue to the next step.
- 3 Using Windows Explorer, navigate to AC\_CONFIG\_HOME. For example, if the system administrator created a folder for the shared configuration directory named config\_cluster, then in a default BIRT iHub installation on Windows, performed using the installer, in which the install folder is C:\Actuate, AC\_CONFIG\_HOME represents the following path:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
  \config_cluster
```

Create a backup copy of acserverconfig.xml. Then, open acserverconfig.xml in a text editor such as Notepad. In acserverconfig.xml, locate the following string:

```
<SMTPServers/>
```

Create a child element of <SMTPServers> named <SMTPServer>. Using the following example, provide values for the attributes of the <SMTPServer> element:

```
<SMTPServers>
  <SMTPServer
    Name="mailhost.actuate.com"
    SenderName="Notifications"
    SMTPHostName="mailhost.actuate.com"
    SenderAddress="support@actuate.com"/>
</SMTPServers>
```

The <SMTPServers> element appears in acserverconfig.xml as shown in Listing 2-1.

### Listing 2-1 acserverconfig.xml with configured <SMTPServer> element

---

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<Config>
  <System
    KeyAlias="birtihub"
  ...
  <UsageAndErrorLogging/>
```

```

    <SMTPServers>
      <SMTPServer
        Name="mailhost.actuate.com"
        SenderName="Notifications"
        SMTPHostName="mailhost.actuate.com"
        SenderAddress="support@actuate.com"/>
    </SMTPServers>
  </System>

```


Save acserverconfig.xml and exit the file.

- 4 Start the cluster by performing the following steps:
  - 1 Choose Cluster Configuration from the side menu.
  - 2 In System Console, on Clusters—Cluster Configuration, left-click the cog icon and choose Start Cluster from the Manage Cluster menu.
  - 3 Choose Refresh from the Manage Cluster menu. When all the services icons have turned green, the cluster is back online.
- 5 If your anti-virus software prevents processes from sending e-mail, the anti-virus software may block an alert notification e-mail. Configure your anti-virus software to allow processes such as java.exe, LMServer.exe, ihub.exe, and ihubc.exe to send e-mail.
- 6 Verify that you are receiving e-mail notification by performing any one of the following tasks. Completion of any of these tasks prompts System Console to send an e-mail notification.
  - Schedule a job in Visualization Platform to run immediately. For more information, see Chapter 3, “Scheduling and Managing Jobs,” in *Using Visualization Platform*.
  - Configure an alert. For example, add an alert having the following properties:
    - Attribute Name: Percent of server RAM used (MB)
    - Condition: Greater than or equal to
    - Threshold: 0 (zero)
 For more information on configuring an alert, see “Adding an alert,” earlier in this chapter.
  - Add a volume. For more information, see “Adding a volume,” earlier in this chapter.

## Editing, deleting, disabling, and enabling an alert

Choose the icon next to an alert on Clusters—Alerts to access the alert menu. This menu contains the following options, as shown in Figure 2-40:

- **Edit**  
Edit the alert.
- **Delete**  
Delete the alert.
- **Disable**  
If the alert is enabled, the menu contains Disable. If the alert is disabled, the menu contains Enable.



**Figure 2-40** Viewing the alert

When editing an existing alert, the system administrator can change any value except the attribute name and the alert name.

**How to edit an alert**

- 1 Point to the icon next to the name of an alert and choose Edit.
- 2 On Edit Alert, modify any properties as necessary, as shown in Figure 2-41. Choose OK.

**How to delete an alert**

Point to the icon next to the name of an alert and choose Delete.

**How to disable or enable an alert**

Disable an enabled alert by left-clicking the icon next to the name of an enabled alert and choosing Disable.

Enable a disabled alert by left-clicking the icon next to the name of a disabled alert and choosing Enable.



**Figure 2-41** Editing an alert

## Configuring Single Sign-On

Choose Single Sign-On to view the SAML identity and service provider information for the nodes in the cluster and optionally, to add a service provider, as shown in Figure 2-42. Service provider information for a cluster node becomes visible to the cluster when the node joins the cluster.

**Figure 2-42** Choosing iHub User Management

## Viewing the information in SAML Identity Provider (IdP) for this cluster

SAML Identity Provider (IdP) for this cluster specifies the following Security Assertion Markup Language (SAML) information:

- Entity ID  
The identity provider identifier. This is the value of the entityID attribute in the <EntityDescriptor> element in the identity provider metadata.
- Metadata URI  
The identifier for the identity provider metadata.
- Metadata path  
The path to the identity provider metadata on disk.

## Viewing and adding service provider information

Service Provider Information displays the information for each service provider on each node in the cluster. The system administrator can also add additional service providers using Add Service Provider.

By default, each node uses the following service providers:

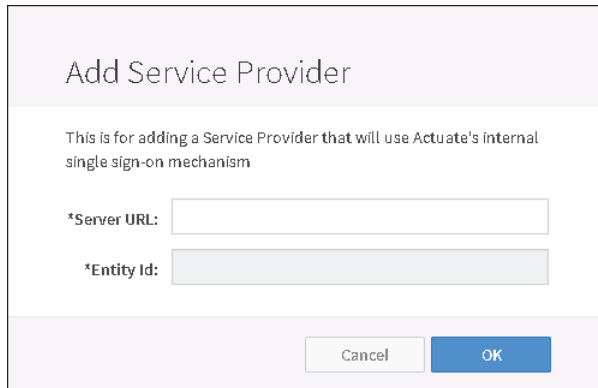
- icadv  
Provides access to Visualization Platform.
- iportal  
Provides access to Visualization Platform.

Choose the icon next to the service provider URL to view the following information for the service provider:

- Entity ID  
The service provider identifier. This is the value of the entityID attribute in the <md:EntityDescriptor> element in the service provider metadata.
- Server URL  
The URL of the login for a service provider. To enable https, set up a proxy that has https enabled.
- Metadata path  
The path of the metadata file for this service provider.
- Metadata URI  
The URI for the metadata for this service provider.
- ACS Post URL  
The URL for ACS Post.

Choose Add Service Provider to specify these properties, as shown in Figure 2-43.

- **Server URL**  
The URL for the service provider.
- **Entity ID**  
The service provider identifier.



The screenshot shows a dialog box titled "Add Service Provider". Below the title, it says "This is for adding a Service Provider that will use Actuate's internal single sign-on mechanism". There are two input fields: one labeled "\*Server URL:" and another labeled "\*Entity Id:". At the bottom right, there are two buttons: "Cancel" and "OK".

**Figure 2-43** Specifying Service Provider information

## Configuring User Management

The system administrator specifies settings for managing user authentication and authorization on User Management. Select among the following ways that BIRT iHub manages users for this cluster:

- iHub User Management (default)
- LDAP Adapter
- RSSE SOAP Service

iHub User Management is the default setting and requires no action.

### Configuring LDAP Adapter

Choose LDAP Adapter to configure settings for user management using a LDAP server. Settings for LDAP Adapter are grouped into the following sections:

- Search setting
- LDAP connection settings
- LDAP Performance Settings
- LDAP Mapping

The following sections describe these property groups.

## About Search setting

Search setting contains one property, Search Cache Only. Cache Only restricts any search for users and user groups that BIRT iHub performs to the open security cache, with the exception of user authentication. When performing user authentication, BIRT iHub always searches the external security source, such as the LDAP server. Searching the cache only improves performance because data retrieval from the cache is faster than from the external data source.

A user sync thread runs in the background, and refreshes the cache automatically, at an interval that the Performance Settings—Cache Timeout property specifies. To prevent BIRT iHub from refreshing the cache, set Performance Settings—Cache Timeout to -1 to prevent a user from ever expiring. If you want BIRT iHub to refresh the cache, Actuate recommends setting Performance Settings—Cache Timeout to 1440 minutes, which is 24 hours, or more, instead of the default 60 minutes.

To use the Search Cache Only feature, create a script that sends the SOAP request that executes the caching operation, an example of which is shown in Listing 2-2. For more information, see “Chapter 24, Actuate Information Delivery API operations,” in *Application Integrator Guide*.

### Listing 2-2 The SOAP request for the operation that loads the cache

---

```
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap
/envelope/" xmlns:act="http://schemas.actuate.com/actuate11">
  <soapenv:Header>
  <soapenv:Body>
    <act:Administrate>
      <act:AdminOperation>
        <act:UpdateOpenSecurityCache>
          <act:LoadAllUsers>true</act:LoadAllUsers>
          <act:LoadAllUserGroups>true</act:LoadAllUserGroups>
        </act:UpdateOpenSecurityCache>
      </act:AdminOperation>
    </act:Administrate>
  </soapenv:Body>
</soapenv:Envelope>
```

Perform this operation immediately after installing BIRT iHub, to load the open security cache. Subsequently, perform the operation to refresh the cache when information in the external data source has changed.

Actuate recommends selecting Search Cache Only only if you have a large number of users or user groups, when using the feature makes enough of a difference in performance to warrant the management task of refreshing the cache.

## Configuring LDAP connection settings

Configure LDAP Connection settings to connect to the LDAP or Active Directory server by providing values for each of the following settings in LDAP Connection settings, as shown in Figure 2-44. An asterisk next to the property name indicates that this property is required. Choose Test Connection to test the connection to the LDAP server after setting all the values in LDAP connection settings. A message displays, indicating whether the connection is successful.

- LDAP Server

Name of the machine hosting the LDAP or Active Directory server. BIRT iHub must be able to resolve this name. For example, when using a LDAP server, specify:

```
ldap.company.com
```

For an Active Directory server, an example value is:

```
ad.company.com
```

- LDAP Port

Port on which the LDAP or Active Directory server listens. Whether using a LDAP server or an Active Directory server, the default port is:

```
389
```

For an LDAP server with SSL (LDAPS), the default port is:

```
636
```

- User DN

Distinguished name of the user that can log in to the LDAP or Active Directory server. The distinguished name with which BIRT iHub binds to the LDAP server. For example, when using a LDAP server, specify:

```
ou=Engineering,dc=company,dc=com
```

For an Active Directory server, an example value is:

```
user@company.com
```

- Password

Password for the LDAP or Active Directory server.

- SSL

Enables connecting to a LDAP server or an Active Directory server with SSL. An out-of-the-box (OOTB) BIRT iHub installation only connects to an LDAP or Active Directory server that has a signed certificate. To connect to a server that does not have a signed certificate, use the Java keytool utility to add that certificate as a trusted certificate. For information on using the Java keytool utility, see: <http://docs.oracle.com/javase/6/docs/technotes/tools/windows/keytool.html>.

- **Active Directory**  
Supports a LDAP implementation using Active Directory. Select if implementing LDAP using Active Directory.
- **Recursive Groups**  
Supports nested group membership. Leave this property deselected if not using an Active Directory LDAP implementation.

The screenshot shows a web form titled "LDAP connection settings". It contains the following fields and controls:

- \*LDAP Server: [Text input field]
- \*LDAP Port: [Text input field with value "0"]
- \*User DN: [Text input field]
- \*Password: [Text input field]
- SSL:
- Active Directory:
- Recursive Groups:
- [Test Connection button]
- [LDAP Performance Settings button]

**Figure 2-44** Configuring LDAP connection settings

### Configuring LDAP Performance Settings

Choose LDAP Performance Settings to set the following properties, as shown in Figure 2-45. An asterisk next to the property name indicates that this property is required.

- **Timeout**  
The number of milliseconds before the time to perform an LDAP operation expires.
- **Maximum Pool Size**  
The maximum number of connections per connection identity that can be maintained concurrently.
- **Fetch Limit**  
The maximum number of entries to be returned from the directory.
- **Preferred Pool Size**  
The preferred number of connections per connection identity to maintain concurrently.

- **Cache Timeout**  
The number of minutes before BIRT iHub deletes cached data.

The image shows a dialog box titled "Performance Settings". It contains five input fields, each with a label and a value:

- \*Timeout: 300000
- \*Maximum Pool Size: 20
- \*Fetch Limit: 500
- \*Preferred Pool Size: 20
- \*Cache Timeout: 60

At the bottom of the dialog box, there are two buttons: "Cancel" and "OK".

**Figure 2-45** Setting Performance Settings properties

## Configuring LDAP Mapping

Configure LDAP Mapping to map BIRT iHub user data to the LDAP or Active Directory server by providing values for each of the following settings in LDAP Mapping, as shown in Figure 2-46. An asterisk next to the property name indicates that this property is required.

- **Prefix**  
For simple authentication, a string value that LDAP prepends to the name with which the user logs on to the server. For LDAP servers requiring distinguished name (DN) login, set this property to the appropriate value, followed by an equal sign (=). For example, specify:

`uid=`

When using an Active Directory server, leave Prefix blank.

- **Suffix**  
For simple authentication, a string value that LDAP appends to the name with which the user logs on to the server. For LDAP servers requiring distinguished name (DN) login, set this property to the appropriate chain of values, preceded by a comma (.). For example, specify:

`,ou=company users,dc=company,dc=com`

When using an Active Directory server, which requires logging in with an e-mail address, set Suffix to @ followed by the domain name of the Active Directory. For example, specify:

`@company.com`

### LDAP mapping

Enter the prefix and suffix for your LDAP principal DN below.

\*Prefix:

\*Suffix:

Next, enter the LDAP properties that should be matched to each of the following iHub values. Some properties can be mapped to more than one value; separate multiple values with commas

\*User Base DN:

\*User Login Name Attribute:

\*User Full Name Attribute:

User Description Attribute:

\*User Object:

User Search Filter:

\*Email Attribute:

\*Group Base DN:

Group Description Attribute:

\*Group Object:

Group Search Filter:

\*Member List Attribute:

\*Member ID Type:

Home Folder Attribute:

\*Default Home Folder:

User Volume Filter Attribute:

Group Volume Filter Attribute:

"Admin" Group:

**Figure 2-46** Configuring LDAP Mapping



- **User Base DN**

The root of the tree that BIRT iHub searches for user information. A user name must be unique for each distinguished name BIRT iHub searches. Separate multiple distinguished names with a semicolon. For example, when using a LDAP server, specify:

```
ou=Users, dc=east, dc=com; ou=Users, dc=west, dc=com
```

For an Active Directory server, an example value is:

```
OU=Users, DC=east, DC=com; OU=Users, DC=west, DC=com
```

- **User Login Name Attribute**

Attribute that specifies the user login name. Cannot contain a space. For example, when using a LDAP server, specify:

```
uid
```

When using an Active Directory server, specify:

```
sAMAccountName
```

Note that if the LDAP or Active Directory server contains a user login name longer than 255 characters, BIRT iHub reads only the first 255 characters. User login names longer than 255 characters are not supported.

- **User Full Name Attribute**

Attribute that specifies the user's full name. For example, when using a LDAP server, specify:

```
cn
```

```
or
```

```
displayName
```

When using an Active Directory server, specify:

```
cn
```

- **User Description Attribute**

Attribute specifying a description of the user. For example, whether using a LDAP server or an Active Directory server, specify:

```
description
```

- **User Object**

LDAP object class for users. For example, when using a LDAP server, specify:

```
person
```

When using an Active Directory server, specify:

```
user
```

- **User Search Filter**

Use this property to identify which users can access BIRT iHub. Use the format appropriate to the indicated provider. For example, create a group for BIRT iHub users on your LDAP server. Then, specify this group as a filter to ensure that BIRT iHub imports only users belonging to the group of BIRT iHub users. For example, when using a LDAP server, specify:

```
cn=birtUsers
```

When using an Active Directory server, an example value is:

```
memberOf:1.2.840.113556.1.4.1941:=CN=\\#QA,CN=Users,DC=actuate,DC=com
```

Be aware that for a distinguished name containing one or more special characters, LDAP stores the distinguished name with any special characters escaped with a backslash, so you must also escape any special character in the value you specify for User Search Filter with a backslash. For more information, see “About searching when Active Directory implements LDAP,” later in this chapter.

- **Email Attribute**

Attribute that stores a user’s e-mail address. For example, whether using a LDAP server or an Active Directory server, specify:

```
mail
```

- **Group Base DN**

The root of the tree that BIRT iHub searches for user group information. Separate multiple distinguished names with a semicolon. For example, when using a LDAP server, specify:

```
ou=Groups, dc=eastern, dc=com; ou=Groups, dc=western, dc=com
```

For an Active Directory server, an example value is:

```
CN=Groups, OU=east, DC=company, DC=com; DC=Groups, OU=west, DC=company, DC=com
```

- **Group Description Attribute**

Attribute specifying a description of the user group. For example, whether using a LDAP server or an Active Directory server, specify:

```
description
```

- **Group Object**

LDAP object class for user groups. For example, when using a Sun Directory LDAP server, specify:

```
groupofuniqueNames
```

For an Active Directory server, an example value is:

group

- **Group Search Filter**

Value with which to filter user groups. For example, when using a LDAP server, specify:

cn=Engineering\*

For either an LDAP Directory server or an Active Directory server, a more advanced example is:

( & ( businessCategory=Sales ) ( cn=a\* ) )

For an Active Directory server, an example value is:

member:1.2.840.113556.1.4.1941:=CN=Vince  
Price,CN=Users,DC=actuate,DC=com

- **Member List Attribute**

The LDAP Role Member attribute. BIRT iHub uses this attribute to find a user in a group. Groups use this attribute to name a user to a group. For example, when using a Sun Directory LDAP server, specify:

uniqueMember

When using an Active Directory server, specify:

member

- **Member ID Type**

The LDAP Role Member. Specifies the type of a member in a group. Whether using a LDAP server or an Active Directory server, specify the type as:

DN

or

LoginID

- **Home Folder Attribute**

Attribute key that maps to a user's home folder. For example, when using a LDAP server or an Active Directory server, specify:

companyHomeFolder

When using an Active Directory server, leave this property blank.

- **Default Home Folder**

Value that specifies the default parent folder of a user's home folder.

If no Home Folder Attribute exists, BIRT iHub uses this property to construct the user's home folder. For example, whether using a LDAP server or an

Active Directory server, specifying `/home` results in a home folder of `/home/bHill` for a user named `bHill`.

- **User Volume Filter Attribute**  
Specifies an attribute, for example, `employeeType`, that BIRT iHub uses to determine which users have access to a volume. Requires the Multi-Tenant license option. For more information, see “About managing volume access by users and user groups when using LDAP,” later in this chapter.
- **Group Volume Filter Attribute**  
Specifies an attribute, for example, `businessType`, that BIRT iHub uses to determine which user groups have access to a volume. Requires the Multi-Tenant license option. For more information, see “About managing volume access by users and user groups when using LDAP,” later in this chapter.
- **“Admin” Group**  
Specifies the name of a group of users to whom BIRT iHub gives Administrator-level privileges in Visualization Platform. When using a LDAP or Active Directory server for user management, BIRT iHub does not use the default Administrators user group in Visualization Platform—iHub Administration. For example, whether using a LDAP server or an Active Directory server, specify:

```
volumeAdministrators
```

### **About managing volume access by users and user groups when using LDAP**

The LDAP mapping attribute **User Volume Filter Attribute** identifies the users that can access a particular volume. The LDAP mapping attribute **Group Volume Filter Attribute** identifies user groups that can access a particular volume.

Whether using a LDAP server or an Active Directory server, the value the system administrator specifies for **User Volume Filter Attribute** is the name of an attribute having a value that is shared by a group of users to which the system administrator wants to give access to a particular volume. The system administrator specifies this attribute value for the Organization ID when creating a volume.

As an example, `employeeType` is an attribute for a user on a LDAP or Active Directory server. All users for which the value of `employeeType` is `Sales` can access a volume having an Organization ID of `Sales`.

Likewise, the value the system administrator specifies for **Group Volume Filter Attribute** is the name of an attribute having a value that is shared by a group of user groups to which the system administrator wants to give access to a particular volume. The system administrator specifies this attribute value for the Organization ID when creating a volume.

As an example, `businessType` is an attribute for a user group on a LDAP or Active Directory server. All user groups for which the value of `businessType` is Insurance can access a volume having an Organization ID of Insurance.

Multiple volumes can have the same Organization ID. When creating a volume, Organization ID can be only one value. If the system administrator specifies both User Volume Filter Attribute and Group Volume Filter Attribute, and on the LDAP or Active Directory the value for these two attributes is the same for a given user and user group, and the system administrator specifies this value as the Organization ID when creating a volume, both the user and the user group can access the volume.

As an example, the system administrator specifies `employeeType` for User Volume Filter Attribute and `businessType` for Group Volume Filter Attribute. If the value for each of these attributes is Sales, and the system administrator specifies Sales for the Organization ID when creating a volume, then the users for which the value of `EmployeeType` is Sales and the user groups for which the value of `businessType` is Sales can access the volume.

## About searching when Active Directory implements LDAP

Active Directory requires that the following characters be escaped with a backslash (`\`) if used in a Distinguished Name (DN):

- Comma (,)
- Backslash character (`\`)
- Pound sign (#)
- Plus sign (+)
- Less than symbol (<)
- Greater than symbol (>)
- Semicolon (;)
- Double quote (")
- Equal sign (=)
- Leading or trailing spaces

If any of these characters appear in a component of a DN, Active Directory stores the character escaped. For example, Active Directory stores the following DN:

```
memberOf=CN=\#QA, CN=Users, DC=actuate, DC=com
```

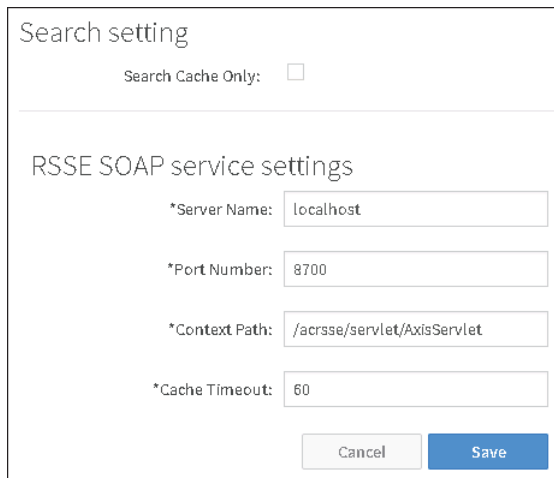
For Active Directory to recognize this DN in a search, you must escape the backslash escape character with another backslash. The following example query returns the users belonging to the #QA group:

```
memberOf=CN=\\#QA, CN=Users, DC=actuate, DC=com
```

## Configuring RSSE SOAP Service

Choose RSSE SOAP Service to configure and view properties for user management using a RSSE web service application for a volume. RSSE SOAP Service is an appropriate choice if you manage user information using an external data source that does not implement LDAP. Configure the following properties for RSSE SOAP Service, as shown in Figure 2-47:

- Search setting
  - Contains Search Cache Only. Restricts searching to only the BIRT iHub metadata database
- RSSE SOAP service settings
  - Contains the following properties:
    - Server Name
      - Machine name of the server that runs the RSSE web service.
    - Port Number
      - Port number for the RSSE web service.
    - Context Path
      - Specifies the location of the RSSE web service for BIRT iHub to use when sending messages to the web service. The path for the default volume is /acrse/servlet/AxisServlet.
    - Cache Timeout
      - Number of minutes before BIRT iHub deletes cached data.

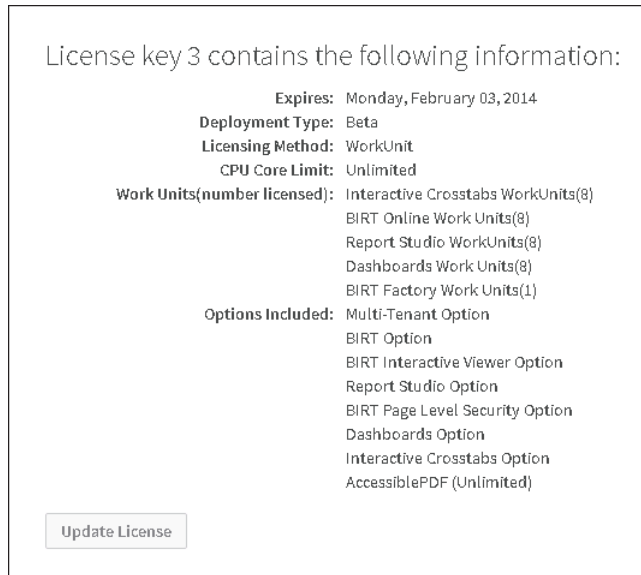


The screenshot shows a configuration dialog box with two main sections. The top section, titled "Search setting", contains a checkbox labeled "Search Cache Only" which is currently unchecked. The bottom section, titled "RSSE SOAP service settings", contains four text input fields: "\*Server Name" with the value "localhost", "\*Port Number" with the value "8700", "\*Context Path" with the value "/acrse/servlet/AxisServlet", and "\*Cache Timeout" with the value "60". At the bottom right of the dialog are two buttons: "Cancel" and "Save".

**Figure 2-47** Configuring security settings

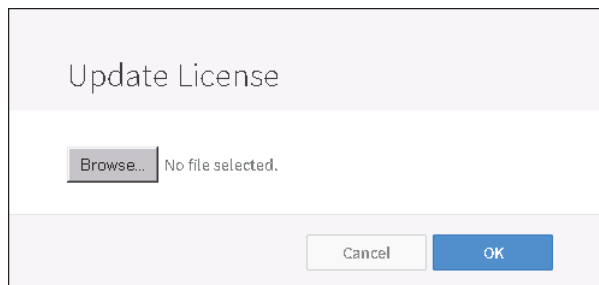
## Updating the license

Each BIRT iHub cluster uses a separate BIRT iHub license. Choose License to view the license options or update the license, as shown in Figure 2-48.



**Figure 2-48** Choosing License

Choose Update License to browse for and select the license file, as shown in Figure 2-49.



**Figure 2-49** Updating the license


## About Configuration File

The configuration file contains BIRT iHub System property settings and templates that all cluster nodes in a cluster use for configuration. The name of the configuration file is `acserverconfig.xml`. Choose Configuration File to perform the following operations, as shown in Figure 2-50:

- Update the configuration file  
Uploads a new acserverconfig.xml to the default location, AC\_CONFIG\_HOME. If the system administrator created a folder for the shared configuration directory named config\_cluster, in a default BIRT iHub installation on Windows, performed using the installer, AC\_CONFIG\_HOME points to the following location:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared  
  \config_cluster
```

- Download or edit the configuration file  
Supports saving AC\_CONFIG\_HOME\acserverconfig.xml to a new location or editing acserverconfig.xml.

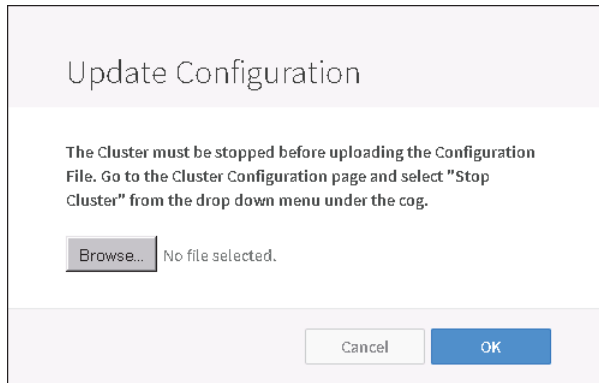


**Figure 2-50** Viewing the options on Clusters—Configuration File

#### How to update the configuration file

- 1 On Clusters, left-click the icon next to the cluster name and choose Edit to edit the cluster.
- 2 On Cluster Configuration, left-click the cog icon and choose Stop Cluster.
- 3 On Configuration File, choose Update Configuration File.
- 4 On Update Configuration File, choose Browse to navigate to the location of the acserverconfig.xml file to upload, as shown in Figure 2-51.



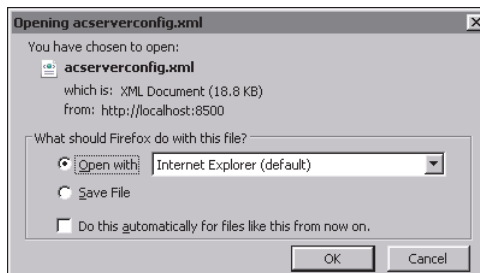


**Figure 2-51** Updating the configuration file

- 5 On File Upload, select the file and choose Open.
- 6 On Update Configuration, choose OK.
- 7 On Cluster Configuration, left-click the cog icon and choose Start Cluster.

#### How to edit or download the configuration file

- 1 On Configuration File, choose Download acserverconfig.xml File.
- 2 On Opening acserverconfig.xml, choose to open the file for editing, or save the file, as shown in Figure 2-52.



**Figure 2-52** Choosing to open or save acserverconfig.xml

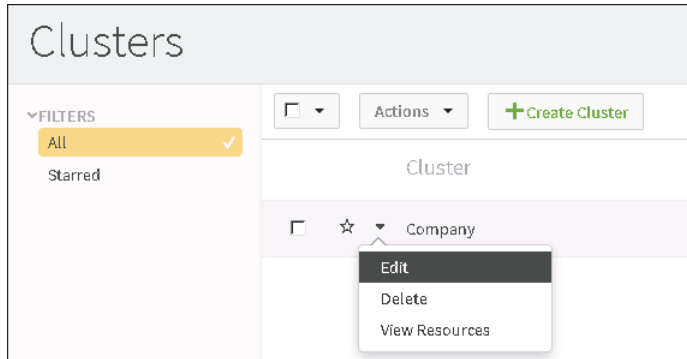
---

## Editing an existing cluster

The system administrator has access to the same properties when editing a cluster as when creating a cluster. Choosing to edit a cluster displays the side menu containing the same property categories the system administrator chooses from when creating a cluster.

### How to edit an existing cluster

On Clusters, left-click the arrowhead in an cluster and choose Edit, as shown in Figure 2-53.



**Figure 2-53** Choosing to edit a cluster

Make any necessary property changes in any cluster category.

---

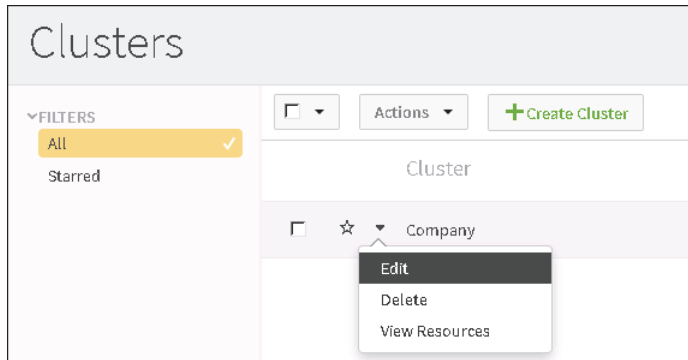
## Managing a cluster node

After adding a cluster node to a cluster, the system administrator can access the node and perform the following operations on the node:

- Starting and stopping
- Editing
- Deleting

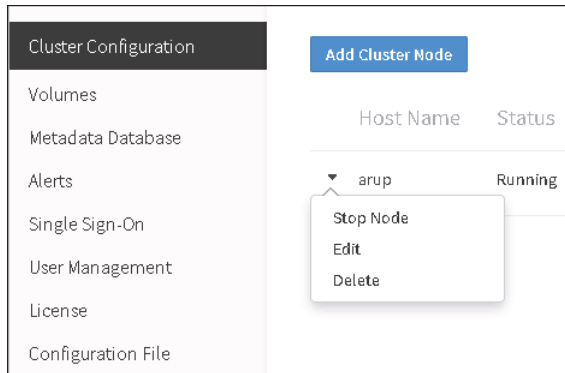
### How to access a cluster node

- 1 In System Console, navigate to Clusters. On Clusters, left-click the arrowhead next to the cluster name and choose Edit, as shown in Figure 2-54.



**Figure 2-54** Choosing to edit a cluster


- 2 On Cluster Configuration, left-click the arrowhead next to a node. System Console displays the operations available to perform on a cluster node, as shown in Figure 2-55.



**Figure 2-55** Accessing the cluster node menu

### How to stop and start a cluster node

On Cluster, left-click the arrowhead icon next to a cluster node. If the node is running, Stop Server displays in the cluster node menu, as shown in Figure 2-56. Choose Stop Node to stop the cluster node.

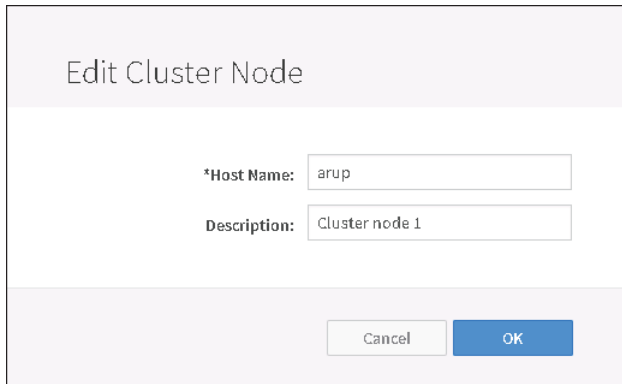


**Figure 2-56** Stopping a cluster node

If the node is not running, Start Node displays in the cluster node menu instead of Stop Node. Choose Start Node to start the cluster node.

#### How to edit a cluster node

- 1 On Cluster, left-click the arrowhead next to a cluster node. Choose Edit to edit the cluster node properties.
- 2 On Edit Cluster Node, make any necessary changes and choose OK, as shown in Figure 2-57.



**Figure 2-57** Editing cluster node properties

#### How to delete a cluster node

On Cluster, left-click the arrowhead next to a cluster node. Choose Delete to delete the cluster node. Confirm the deletion.

---

## Viewing the list of clusters

Clusters displays all existing clusters. The system administrator can view the entire list, or a subset of the list. Choose Clusters to view the entire list. The system administrator can filter the cluster list to display only the following cluster subsets:

- Clusters where the cluster ID contains a specified search string
- Clusters that are starred

The following sections describe how to filter the cluster list.

### Filtering the list of clusters using a search string

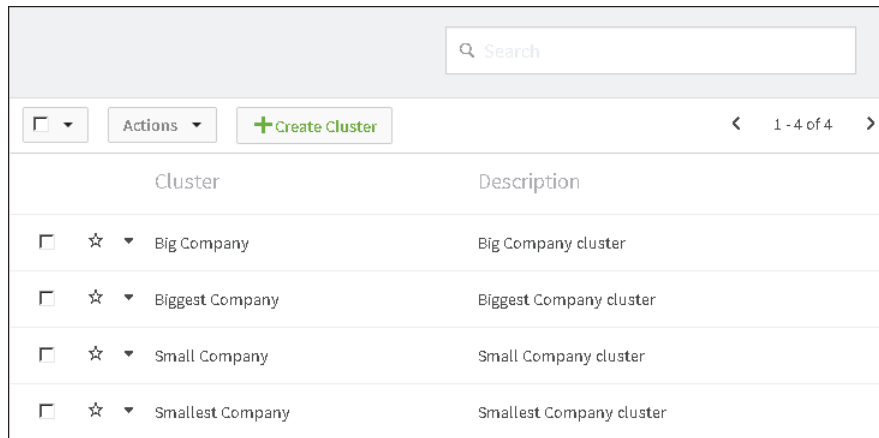
The system administrator can filter the cluster list to display only clusters containing a specified search string in the cluster ID. In Search, type the first letter

or letters of a cluster ID to display the clusters for which the name starts with the specified letter or letters. Search is not case-sensitive.

System Console supports using an asterisk (\*) in a search string as a wildcard character. The asterisk represents zero or more characters, excluding spaces and punctuation.

### How to filter the list of clusters using a search string

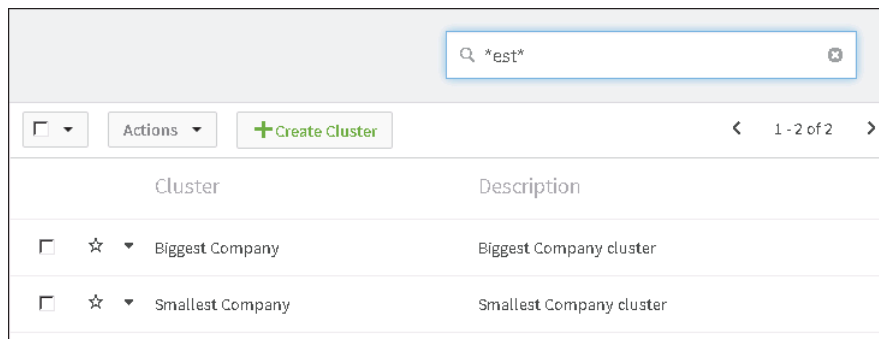
Using the cluster list shown in Figure 2-58 as an example, type `*est*` in Search to find all cluster IDs that contain this string. System Console automatically filters the list and displays the results, as shown in Figure 2-59.



The screenshot shows a web interface for managing clusters. At the top right is a search bar with a magnifying glass icon and the text "Search". Below the search bar is a toolbar with a checkbox, an "Actions" dropdown menu, a "+ Create Cluster" button, and a pagination indicator showing "1 - 4 of 4". The main content is a table with two columns: "Cluster" and "Description".

Cluster	Description
<input type="checkbox"/> ☆ ▾ Big Company	Big Company cluster
<input type="checkbox"/> ☆ ▾ Biggest Company	Biggest Company cluster
<input type="checkbox"/> ☆ ▾ Small Company	Small Company cluster
<input type="checkbox"/> ☆ ▾ Smallest Company	Smallest Company cluster

**Figure 2-58** Viewing the cluster list before filtering



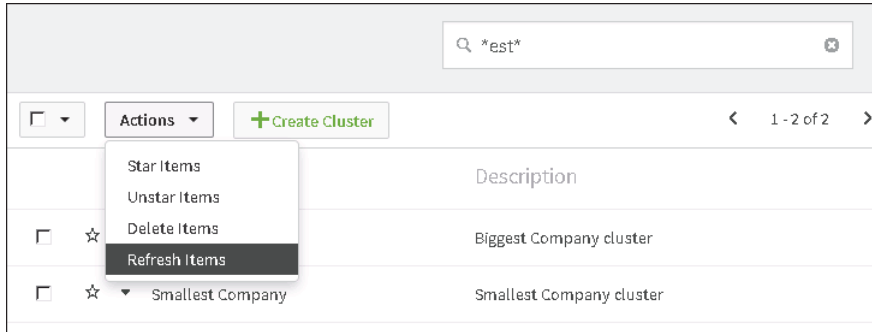
The screenshot shows the same web interface as Figure 2-58, but with the search bar containing the text `*est*`. The pagination indicator now shows "1 - 2 of 2". The table now only displays two rows, corresponding to the filtered results.

Cluster	Description
<input type="checkbox"/> ☆ ▾ Biggest Company	Biggest Company cluster
<input type="checkbox"/> ☆ ▾ Smallest Company	Smallest Company cluster

**Figure 2-59** Viewing the results of filtering the cluster list

### How to refresh the list of clusters

Refresh the list of clusters by left-clicking Actions and choosing Refresh items, as shown in Figure 2-60:



**Figure 2-60** Refreshing the clusters list



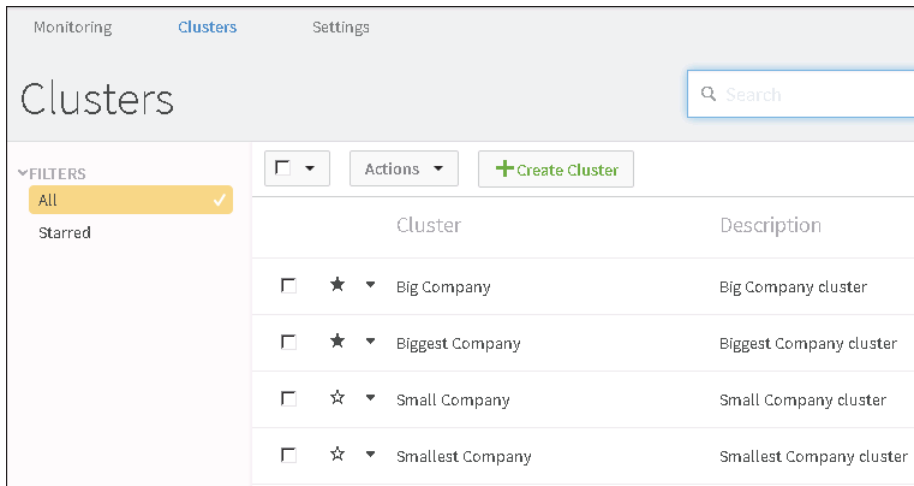
Alternatively, choose the x icon in the search text box.

## Filtering non-starred clusters

The system administrator can filter the cluster list to show only starred clusters. A starred cluster appears with a filled-in star next to the cluster ID. A starred cluster appears in the list of clusters on Monitoring—Weekly Activity Report.

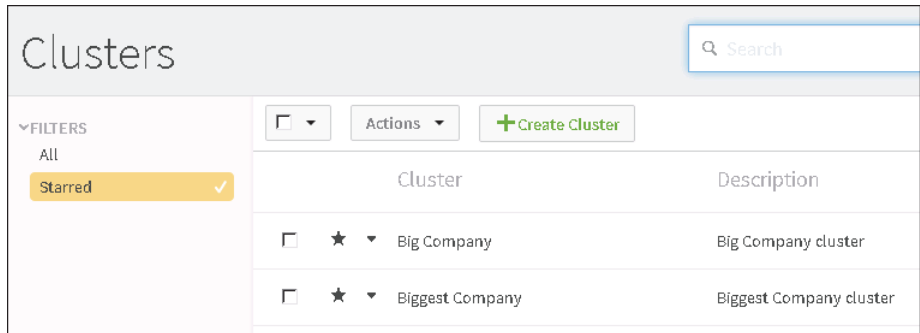
### How to filter non-starred clusters

- 1 Choose Clusters, as shown in Figure 2-61.



**Figure 2-61** Choosing Clusters

- 2 Select Starred. Only the starred clusters display, as shown in Figure 2-62.



**Figure 2-62** Viewing only the starred clusters

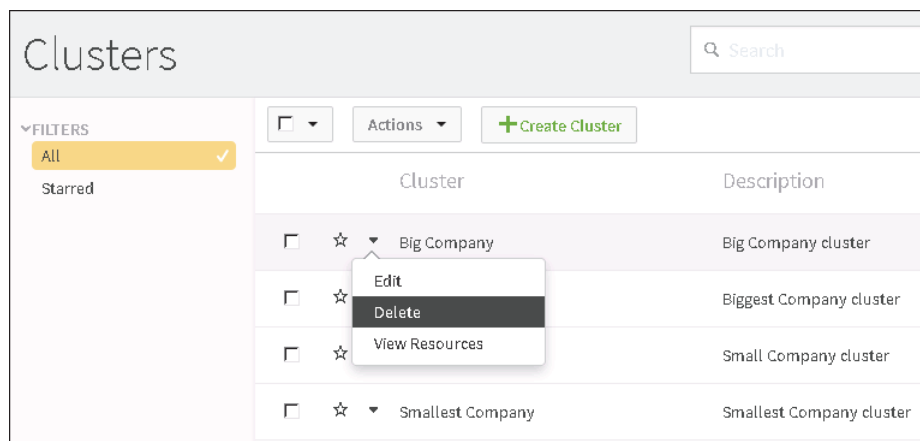
Choose All to re-display all clusters.

## Deleting clusters

The system administrator can delete one cluster only, or multiple clusters simultaneously.

### How to delete a single cluster

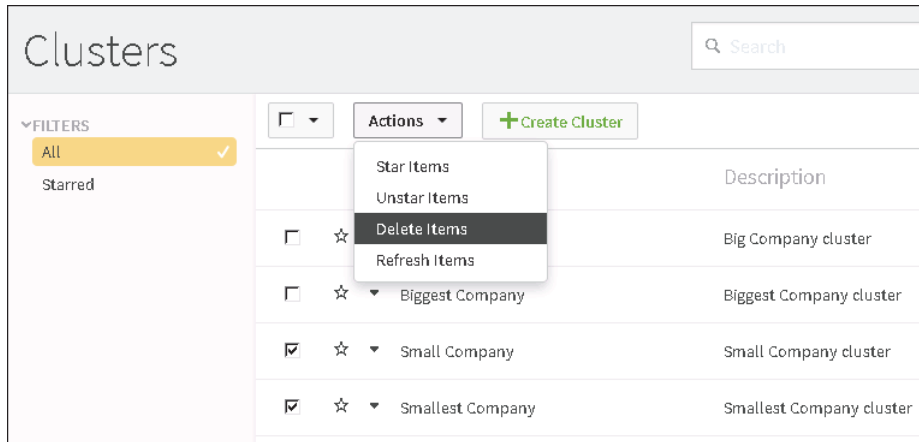
Left-click the arrowhead icon next to the cluster name in the list of clusters and choose Delete to delete a cluster, as shown in Figure 2-63.



**Figure 2-63** Deleting a single cluster

### How to delete one or more clusters

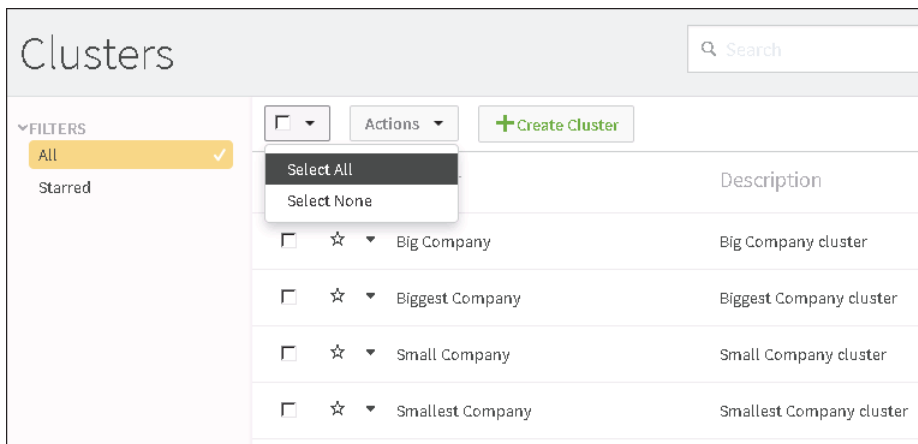
Select each cluster to delete, then choose Actions. In Actions, select Delete items, as shown in Figure 2-64.



**Figure 2-64** Deleting multiple clusters

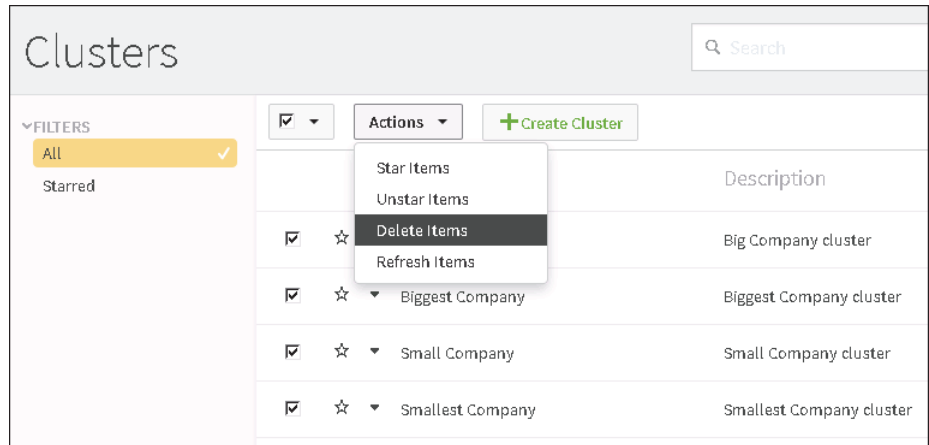
### How to delete all clusters

- 1 Left-click the arrowhead in the check box next to Actions and choose Select All, as shown in Figure 2-65.
- 2 Left-click the arrowhead in Actions and choose Delete Items, as shown in Figure 2-66. Left-clicking a check box toggles between selected and deselected.



**Figure 2-65** Selecting all clusters





**Figure 2-66** Deleting all clusters

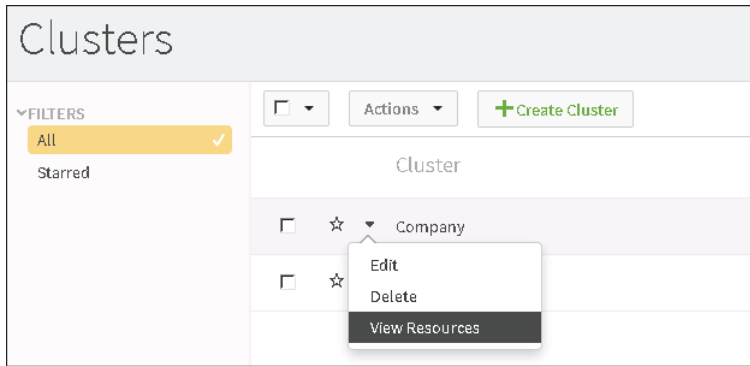
## Viewing cluster resources

System Console provides detailed information on activity and resource usage in a cluster. System Console groups this information into the following categories:

- Logs  
Diagnostic logs the cluster creates
- Trends  
System-level information
- Current Activity  
Number of active requests the cluster is processing

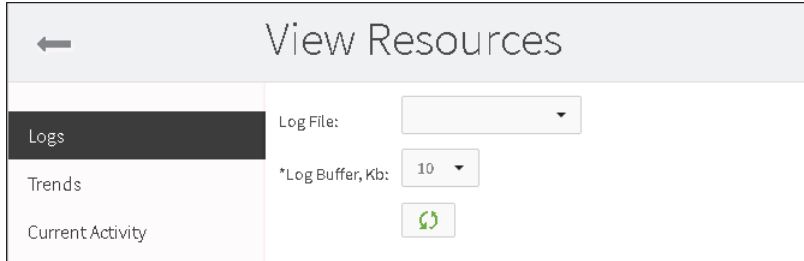
### How to access cluster activity and resource usage information

- 1 On Clusters, left-click the arrowhead icon next to a cluster name and choose View Resources, as shown in Figure 2-67.



**Figure 2-67** Choosing View Resources

- 2 On View Resources, choose a category from the side menu, as shown in Figure 2-68.



**Figure 2-68** Choosing a resource information category

The following sections describe the resource information categories.

## Viewing diagnostic log files using Logs

The system administrator can view the diagnostic logs a cluster creates. Each log is an aggregate of the most recent log entries from every node in the cluster. An aggregate log exists for each of the following log types:

- ihub  
System log entries
- ihubc  
Volume log entries
- ihubd  
Process Management Daemon (PMD) log entries
- jsrvrihub  
View service log entries

- **Imserver**

View monitoring service log entries

The system administrator chooses the type of log to view, and how much of the log to view.

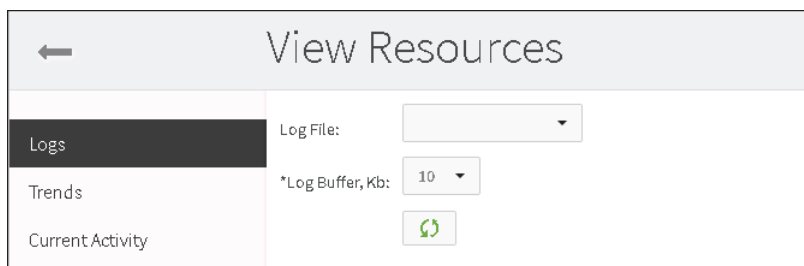
Each log entry contains the host name of the cluster node generating the log entry, the process ID, and the product name. Using the following jsrvrihub log entry as an example,

```
hn:URUP|pid:4920|prod:iHub|[Thread 1] 2013-11-25 08:49:51 UTC-0800  
com.sun.xml.internal.bind.v2.runtime.reflect.opt.OptimizedAccess  
sorFactory.get()
```

URUP is the host name, 4920 is the process id, and iHub is the product name.

### How to view the diagnostic log

- 1 On Clusters, left-click the arrowhead icon next to a cluster name and choose View Resources.
- 2 From the side menu, choose Logs, as shown in Figure 2-69.



**Figure 2-69** Choosing Logs

- 3 On View Resources—Logs, perform the following tasks:
  - 1 In Log File, select the type of log file to view. For example, choose jsrvrihub.log.
  - 2 In Log Buffer, Kb, select how many kilobytes of the log to view. For example, choose 10 KB.

The aggregated diagnostic log appears, as shown in Figure 2-70.



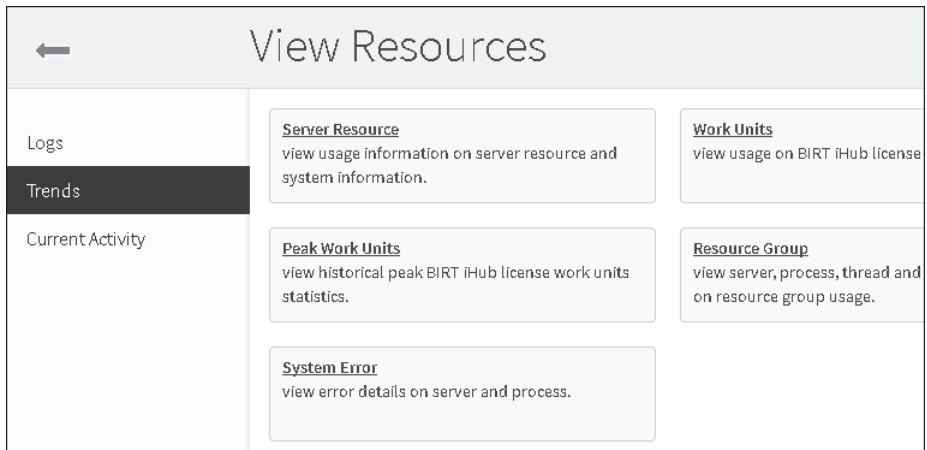
**Figure 2-70** Viewing jsrvrihub.log

## Viewing system-level information using Trends

System Console monitors each cluster and makes a broad range of system-level information on a cluster available for viewing, including statistics on cluster resource usage, performance, and load, presented in graphical and tabular formats.

### How to view system-level information

- 1 On Clusters, left-click the arrowhead icon next to a cluster name and choose View Resources to access the list of system-level information categories.
- 2 From the side menu, choose Trends, as shown in Figure 2-71.



**Figure 2-71** Choosing Trends

**3** Choose one of the following system-level information categories on View Resources—Trends. A new browser window opens, in which Actuate Viewer displays one or more graphs or tables showing the information for the chosen category.

■ **Server Resource**

For each of the following statistics, System Console displays a graph showing the last 48 hours of activity for every node in the cluster. Each node appears as a line in the graph:

- **Response Time (milliseconds)**  
Left-click anywhere on a line in this graph to generate a set of node-specific graphs showing the last 48 hours of activity for each of the following statistics:
  - **Response Times (milliseconds)**
  - **Running Requests**
  - **Total Completed OnDemand Requests Today**
- **CPU Usage (%)**
- **Memory Usage (%)**
- **Disk Usage (%)**

When viewing the CPU Usage, Memory Usage, or Disk Usage graphs, left-click anywhere on a line in any one of these graphs to generate a set of node-specific graphs showing the last 48 hours of activity for each of the following statistics:

- **Memory, CPU and Disk Usage**
- **Processes CPU Usage (%)**
- **Processes Memory Usage (MB)**

■ **Work Units**

Displays a graph showing usage for each of the following work units over the past 48 hours:

- **BIRT Online**
- **BIRT Factory**
- **Dashboards**
- **Interactive Crosstabs**
- **Report Studio**

- **Peak Work Units**  
Displays a table showing the dates that each work unit reached a particular threshold, and the number of times the work unit reached the threshold on each date.
- **Resource Group**  
View process, thread and queue statistics on resource group. Displays a table showing the following information:
  - **Time stamp**  
Date and time at which BIRT iHub used the resource group
  - **Server name**  
Name of the server using the resource group
  - **Resource group**  
Name of the resource group
  - **Number of processes**  
Number of Factory processes the resource group can allocate for executing jobs
  - **Busy threads**  
Number of threads running under a resource group
  - **Queue size**  
Size of the queue where jobs wait to use the resource group
- **System Error**  
Displays a table listing system errors and their details.


## About the Actuate Viewer menu

System Console displays each of the system-level information category views using Actuate Viewer. Actuate Viewer provides a menu that supports working with the contents of each system-level information category view.



Choose the icon in the banner of a system-level information view to access a menu containing the following options:

- **Disable Interactivity**  
Disables interactivity with a chart or table in a system-level information category view. By default, Actuate Viewer displays a chart or table in a system-level category view with interactivity enabled. Interactivity enables control over any column or chart in the view. Left-click a column name or chart in the view. Choose the ellipse to display a menu from which to choose options such as filtering, sorting, alignment, among others. Figure 2-72 shows the menu for a column.

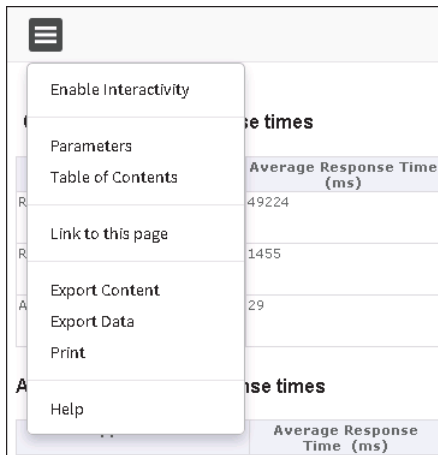


**Figure 2-72** Viewing the menu of column control options

- **Parameters**  
Supports refreshing the view. Choose Parameters. Then, choose Run Report.
- **Table of Contents**  
Supports immediately navigating to a location in the view.
- **Hide/Show Item**  
Supports hiding or showing one or more elements in a view, such as a chart, table, table column, or column label.
- **Link to this page**  
Provides a link to paste in an e-mail or message, or the HTML to embed in a web-page.
- **Export Content**  
Supports downloading the view in one of the following formats:
  - Excel (XLS)
  - Excel (XLSX)
  - PDF
  - PostScript (PS)
  - PowerPoint (PPT)
  - PowerPoint (PPTX)
  - Word (DOC)
  - Word (DOCX)
  - XHTML

- **Export Data**  
Supports downloading the view data in one of the following file formats:
  - Comma (CSV)
  - Pipe (PSV)
  - Tab (TSV)
  - Semicolon (SSV)
- **Print**  
Supports printing the view.
- **Help**  
Accesses help for Actuate Viewer.

Figure 2-73 shows an example of the menu, from the Request/Response Time by Category and Application view.



**Figure 2-73** Viewing the menu common to each system information view

## Viewing Current Activity

Current Activity shows the Node Active Request Summary, which lists the number of active requests each cluster node is processing. Choose a node name in the list to see a detailed status report for the node. The Node Status report displays statistics for the following categories, as shown in Figure 2-75:

- **Resource Status**  
Displays the following resource usage statistics:
  - **RAM**  
Percentage of total memory the node is using



- CPU  
Percentage of total CPU the node is using
- Disk  
Percentage of Disk space node is using
- Process Status  
Displays the following information about BIRT Service processes and BIRT iHub processes on the node:
  - PID  
Process ID number
  - Process Name  
Name of the process
  - Resource Group  
The resource group running the process
  - File Descriptors  
Number of file descriptors the process is using
  - RAM  
Amount of memory in megabytes that the process is using
  - CPU  
Percentage of CPU the process is using
  - Total threads  
Total number of threads the process is using
  - Busy Worker Threads  
Number of worker threads performing work
  - Acquired WUs  
Number of acquired work units
  - Queue Length  
Length of the queue
  - Last Updated Time  
The time at which Process Status information was captured
- Active Requests Status  
Displays the following information about requests that are active on the node:
  - PID  
Process ID number

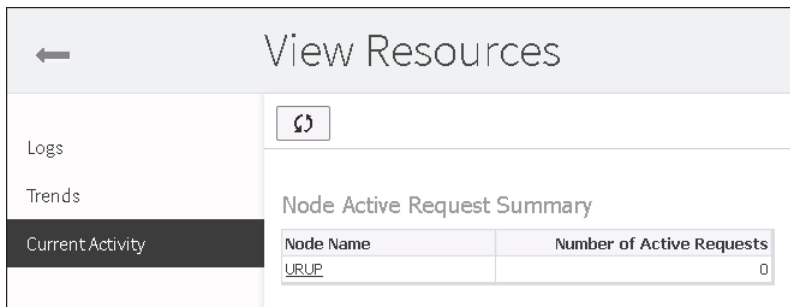
- Request ID  
Request ID number
- Volume  
Volume on which the request is active
- User  
User who initiated the request
- Operation  
Type of operation the request is performing
- File  
File the request is using
- Status  
Status of the request
- Submission time  
Time the user submitted the request
- Running time  
Status of the request
- Action  
Action the system is taking for the request

#### How to view the list of active requests

1 On Clusters, left-click the arrowhead icon next to a cluster name and choose View Resources.

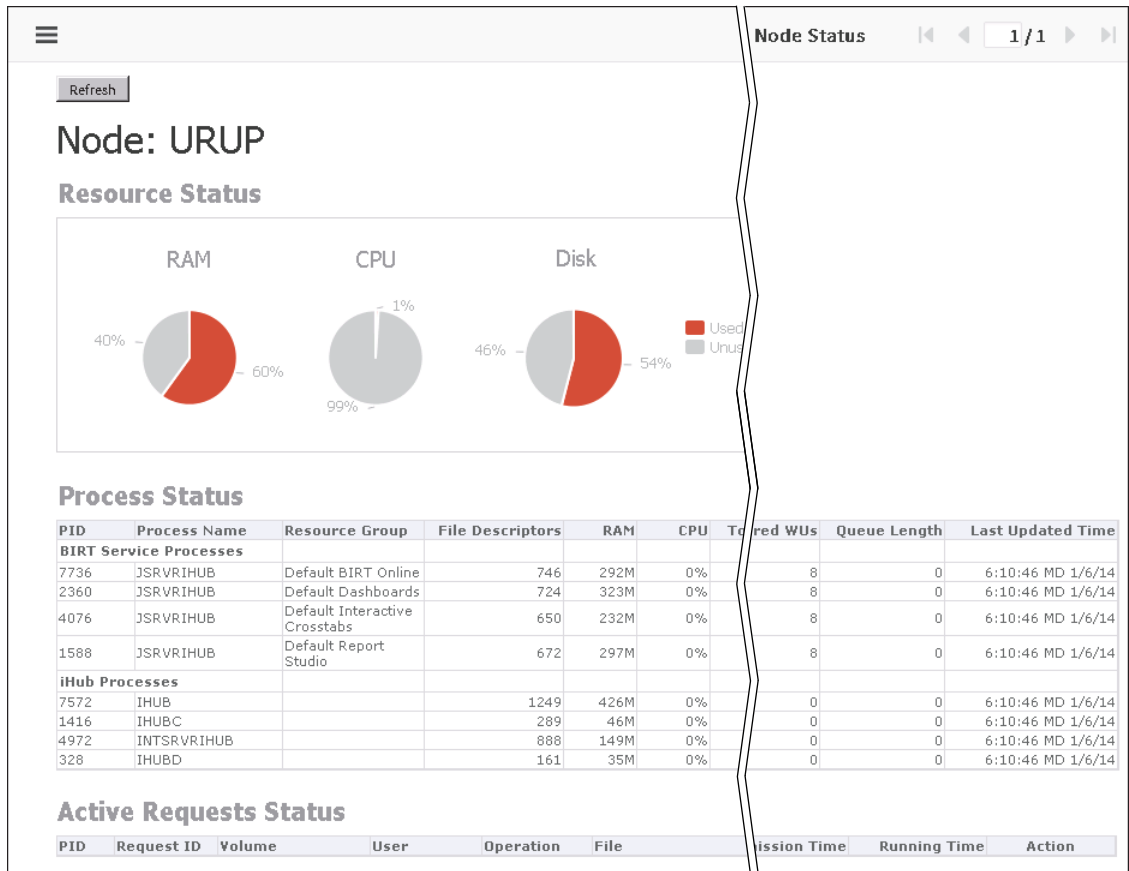


2 On View Resources, choose Current Activity to access the Node Active Request Summary. Then left-click the Refresh button to update the Node Active Request Summary, as shown in Figure 2-74.



**Figure 2-74** Viewing the list of active requests

- On Node Active Request Summary, choose the node name for which you want to view a node status report. For example, choosing the node name URUP in Figure 2-74 generates the report shown in Figure 2-75.



**Figure 2-75** Viewing the Node Status report

## Using the ControlConsole utility to free memory

LMServer maintains an in-memory database. By default, this database holds up to 15 days worth of monitoring data. This database can fill up, causing LMServer to run out of memory.

ControlConsole is a Java utility which performs various logging and monitoring operations. ControlConsole accepts a number of command options, including the following options for freeing memory:

- gc  
Runs Java garbage collection.
- dataunload <number of days of monitoring data to keep>  
Releases monitoring data from memory that is older than the number of days you specify.

### How to run the ControlConsole utility

#### 1 Create the following environment variables:

- AC\_CONFIG\_HOME  
Location of the shared configuration directory for a cluster. If the system administrator created a folder for the shared configuration directory named config\_cluster, and assuming the BIRT iHub installation folder is C:\Actuate\BIRTiHubVisualization, the default location for AC\_CONFIG\_HOME is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub
  \shared\config_cluster
```

- AC\_DATA\_HOME  
Location of logs to which BIRT iHub processes write and System Console monitors. Assuming the BIRT iHub installation folder is C:\Actuate\BIRTiHubVisualization, the default location for AC\_DATA\_HOME is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\data
```

- JAVA\_HOME  
Location of the Java JDK and JRE. Assuming the BIRT iHub installation folder is C:\Actuate\BIRTiHubVisualization and the BIRT iHub install process installed a JDK, the default location for JAVA\_HOME is:

```
C:\Actuate\BIRTiHubVisualization\modules\JDK64
```

#### 2 Append the following value to the PATH environment variable:

```
%JAVA_HOME%\bin
```

#### 3 Open a command prompt. Navigate to AC\_SERVER\_HOME\Jar. The lmsutility.jar file contains the ControlConsole utility. Run ControlConsole by executing the Java command, including the classpath of the ControlConsole utility, followed by:

```
-h <hostname> <command option>
```

For example, on a machine named urup, run ControlConsole with the help command option, which lists the ControlConsole command options:

```
java -classpath .;lmsutility.jar
  com.actuate.lmservice.utils.ControlConsole -h urup help
```

Listing 2-3 shows this command and its output.

**Listing 2-3** Executing the ControlConsole help command

---

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>java
  -classpath .;lmsutility.jar com.
actuate.lmservice.utils.ControlConsole -h urup help
Usage: lmsctrl -h hostname command
      where command is one of the follow:
      shutdown lmserver | lstailer
      ping lmserver
      gc
      start logging | monitoring
      stop logging | monitoring
      restart logging | monitoring
      memprofile
      datareload
      dataunload no_of_days_tokeep
      nightllysnapshot
      snapshotcsv dateinyyyy-MM-dd
      help
```

Example:

```
lmsctrl -h svr1 shutdown lmserver
lmsctrl -h svr1 restart monitoring
```

- 4** Run Java garbage collection by executing the gc command, as shown in Listing 2-4.

**Listing 2-4** Running Java garbage collection

---

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>java
  -classpath .;lmsutility.jar com.
actuate.lmservice.utils.ControlConsole -h urup gc
Send control command to host:urup
```

LMServer execute gc command issued

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>
```

- 5** Release monitoring data from memory that is older than the number of days you specify by running ControlConsole with the dataunload <number of days of data to keep> command option. Listing 2-5 shows running this command option to keep three days worth of monitoring data.

**Listing 2-5** Releasing monitoring data older than number of days you specify

---

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>java
  -classpath .;lmsutility.jar com.
```

```
actuate.lmservice.utils.ControlConsole -h urup dataunload 3  
Send control command to host:urup
```

```
dataunload command successful
```

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>
```

- 6 View the amount of memory used and the amount of memory free by running ControlConsole with the memprofile command option. Listing 2-6 shows running ControlConsole with this command option.

**Listing 2-6** Viewing the amount of memory used and amount of memory free

---

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>java  
-classpath .;lmsutility.jar com.  
actuate.lmservice.utils.ControlConsole -h urup memprofile  
Send control command to host:urup
```

```
MEMORY USED: 25 MB MEMORY FREE: 353 MB Completed operation in  
: 2 Sec
```

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\Jar>
```

---

## About BIRT iHub service and resource group properties

The shared configuration file, `acserverconfig.xml`, contains the server configuration templates that the nodes in a cluster use. Each node uses a single server configuration template. The names of the default templates are `small`, `medium`, `large`, and `disable`. The system administrator can assign the `disable` configuration template to a cluster node to prevent the node from performing any processing. By default, all the resource groups in the `disable` template are configured to run no Factory processes. For example, assign the `disable` template to a node that runs only the Monitor service.

Table 2-3, Table 2-4, and Table 2-5 describe the Reporting, Viewing, and Integration service properties that the `small`, `medium`, and `large` server configuration templates contain. Each property description in a table shows when a property change takes effect, and includes the default value for the property for each of the templates.

The shared server configuration file, `acserverconfig.xml`, contains the templates. In a typical BIRT iHub installation, the location of `acserverconfig.xml` is: `AC_CONFIG_HOME`.

## Reporting service template properties

**Table 2-3** Viewing Reporting service template properties

Property Name	Description	Takes effect	Small	Medium	Large
EnableGeneration Service	Enable factory service	Immediate	true	true	true
SyncJobQueueSize	Job queue size for synchronous reports	Immediate	100	100	100
SyncJobQueueWait	Job queue timeout for transient reports	Immediate	600 seconds	600 seconds	600 seconds
MaxSyncJobRuntime	Maximum execution time for on demand execution requests	Immediate	300 seconds	300 seconds	300 seconds
MaxSyncRequestTime	Maximum execution time for all request types except on-demand execution	Immediate	300 seconds	300 seconds	300 seconds
TransientReportTime Out	Disk cache timeout for transient reports	Server Restart	30 minutes	30 minutes	30 minutes
TransientReportCache Size	Disk cache size for transient reports	Immediate	100 MB	100 MB	100 MB
TransientStoreMax CacheEntries	Maximum memory cache entries for transient reports	Immediate	10000	10000	10000
MaxBIRTDataResult setBufferSize	Maximum result set buffer size for BIRT data object generation query	Server Restart	128 MB	256 MB	512 MB
DatamartArchive Limit	Number limit for datamart files	Server Restart	10	20	40
MaxDatamartArchive Size	Maximum memory size for each datamart file	Server Restart	30720 KB	30720 KB	30720 KB
PersistentArchive Limit	Number limit for persistent BIRT document files	Server Restart	10	20	40

*(continues)*

**Table 2-3** Viewing Reporting service template properties (continued)

Property Name	Description	Takes effect	Small	Medium	Large
MaxPersistentArchiveSize	Maximum memory size for each persistent BIRT document file	Server Restart	1024 KB	1024 KB	1024 KB
SynchReportingWeight	Weight of this server for load balancing on demand execution requests	Immediate	100	200	400

## Viewing service template properties

**Table 2-4** Viewing Viewing service template properties

Property Name	Description	Takes effect	Small	Medium	Large
EnableViewingService	Enable viewing service	Immediate	true	true	true
JavaServerClientMaxConnections	Maximum number of connections from client to encyc server	Server Restart	8	8	8
JavaServerClientMinConnections	Minimum number of connections from client to encyc server	Server Restart	3	3	3
JavaServerClientConnectionTimeout	The timeout value for the connections from client to encyc server	Server Restart	300 seconds	300 seconds	300 seconds
FileCacheTimeout	Cache timeout for search results, table of contents and image files	Server Restart	86400 seconds	86400 seconds	86400 seconds
MaxConcurrentRequests	Maximum queue size per process for requests	Immediate	128	128	128
BIRTReportDesignCacheTimeout	Cache timeout for BIRT designs	Server Restart	1800 seconds	1800 seconds	1800 seconds
OnDemandServerViewMessageTimeout	Timeout for on demand and viewing messages	Server Restart	300 seconds	300 seconds	300 seconds



**Table 2-4** Viewing Viewing service template properties (continued)

Property Name	Description	Takes effect	Small	Medium	Large
TransientArchiveFileCacheTimeout	Expiration timeout for transient BIRT documents and datamarts in the archive file cache	Server Restart	1200 seconds	1200 seconds	1200 seconds
PersistentArchiveFileCacheTimeout	Expiration timeout for persistent BIRT documents and datamarts in the archive file cache	Server Restart	7200 seconds	7200 seconds	7200 seconds
DatamartArchiveLimit	Number limit for datamart files	Server Restart	5	10	20
TransientArchiveLimit	Number limit for transient BIRT document files	Server Restart	10	20	40
MaxDatamartArchiveSize	Maximum memory size for each datamart file	Server Restart	30720KB	30720KB	30720KB
PersistentArchiveLimit	Number limit for persistent BIRT document files	Server Restart	10	20	40
MaxTransientArchiveSize	Maximum memory size for each transient BIRT document file	Server Restart	512KB	512KB	512KB
MaxPersistentArchiveSize	Maximum memory size for each persistent BIRT document file	Server Restart	1024KB	1024KB	1024KB
BIRTImageCacheTimeout	Cache timeout for images and charts from BIRT designs, documents and datamarts	Server Restart	86400 seconds	86400 seconds	86400 seconds
BIRTReportDesignCacheTotalNumberOfEntries	Maximum number of BIRT designs to cache	Server Restart	50	50	50

*(continues)*

**Table 2-4** Viewing Viewing service template properties (continued)

Property Name	Description	Takes effect	Small	Medium	Large
ViewingWeight	Weight of this server for load balancing viewing requests	Immediate	100	200	400

## Integration service Template properties

**Table 2-5** Viewing Integration service template properties

Property Name	Description	Takes effect	Small	Medium	Large
EnableIntegrationService	Enable Integration service	Immediate	true	true	true
PagePoolSize	Page pool size	Server Restart	2000 pages	2000 pages	2000 pages
BufferPoolSize	Buffer pool size	Server Restart	18000 pages	18000 pages	18000 pages
StartArguments	Start parameters for Integration service processes	Server Restart	-Xms64M -Xmx128M com.nimble.nie.server.Server	-Xms128M -Xmx256M com.nimble.nie.server.Server	-Xms256M -Xmx512M com.nimble.nie.server.Server

## Understanding resource groups

A resource group controls the Factory processes that BIRT iHub uses to run a synchronous or asynchronous job. A resource group specifies a set of Factory processes reserved to run only those jobs assigned to the group.

A design that runs unscheduled runs synchronously, as soon as possible in the foreground. iHub does not store the generated document in the volume. A scheduled job runs asynchronously in the background. iHub stores the generated document in the volume. Whether you generate a document by running a design scheduled or unscheduled, you can view, navigate, and search the generated document.

BIRT iHub uses the following default resource groups:

- Default BIRT Factory  
Used to run a BIRT design (.rptdesign) as a scheduled job. Also used to print a BIRT document (.rptdocument).

- **Default BIRT Online**  
Used for running a BIRT design unscheduled and viewing the generated BIRT document.
- **Default Dashboards**  
Used for running a BIRT dashboard (.dashboard) or gadget (.gadget) design unscheduled and viewing the generated document.
- **Default Interactive Crosstabs**  
Used for running a Data Object Store (.data) design unscheduled and viewing the generated document.
- **Default Report Studio**  
Used when creating, modifying, and viewing documents using Report Studio.

## **Understanding resource group properties pertaining to a template**

Each template contains an XML element named `<ServerResourceGroupSetting>` for each resource group. This element contains attributes that affect the configuration of the resource group only for that template.

A `<ServerResourceGroupSetting>` element contains the following properties for a resource group:

- **Name**  
Name of the resource group.
- **Activate**  
Possible values are true or false. A value of true activates the resource group, enabling the node using this configuration template to use the Factory processes this resource group provides. A value of false deactivates the resource group. By default, Activate is true for all resource groups.
- **MinFactory**  
Minimum number of factories a resource group can use.
- **StartArguments**  
The Java Runtime Environment (JRE) start arguments, or Java command-line options, for the resource group.

Any change to a resource group property requires a cluster restart to take effect.

The StartArguments property appearing in each `<ServerResourceGroupSetting>` element include the following JRE start arguments:

- **Heap limit option**  
Specifies the amount of heap the Java process can use. For example, `-Xmx512M` specifies that the Java process can use 512 MB of heap. Too large a heap can slow garbage collection because there is more heap to scan. This property affects Java view server memory usage.
- **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.

The following list describes the start arguments that each StartArguments property specifies, for the small, medium, and large templates. Additionally, if the <ServerResourceGroupSetting> element contains a MinFactory property, the list includes the value for this property.

- **Small**
  - **Default BIRT Online**  
This resource group contains the following start arguments:  

```
-Xmx512M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.
  - **Default Dashboards**  
This resource group contains the following start arguments:  

```
-Xmx512M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
```

```
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.

- **Default Interactive Crosstabs**

This resource group contains the following start arguments:

```
-Xmx512M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

- **Default BIRT Factory**

This resource group contains the following start arguments:

```
-Xmx512M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

- **Default Report Studio**

This resource group contains the following start arguments:

```
-Xmx512M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

- **Medium**

- **Default BIRT Online**

This resource group contains the following start arguments:

```
-Xmx1024M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.

- **Default Dashboards**

This resource group contains the following start arguments:

```
-Xmx1024M -XX:MaxPermSize=256m -XX:-UsePerfData
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true
-Djava.protocol.handler.pkgs=com.actuate.javaserver
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.

- **Default Interactive Crosstabs**

This resource group contains the following start arguments:

```
-Xmx1024M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

- **Default BIRT Factory**

This resource group contains the following start arguments:

```
-Xmx1024M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

- **Default Report Studio**

This resource group contains the following start arguments:

```
-Xmx1024M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

- **Large**

- **Default BIRT Online**

This resource group contains the following start arguments:

```
-Xmx2048M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.

- **Default Dashboards**

This resource group contains the following start arguments:

```
-Xmx2048M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

The minimum number of factories this resource group can use is 1.

- **Default Interactive Crosstabs**

This resource group contains the following start arguments:

```
-Xmx2048M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

- **Default BIRT Factory**

This resource group contains the following start arguments:

```
-Xmx2048M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

- **Default Report Studio**

This resource group contains the following start arguments:

```
-Xmx2048M -XX:MaxPermSize=256m -XX:-UsePerfData  
-Djava.awt.headless=true -Djava.net.preferIPv4Stack=true  
-Djava.protocol.handler.pkgs=com.actuate.javaserver  
.protocol com.actuate.javaserver.Server
```

## **Understanding resource group properties pertaining to all templates**

The shared configuration file, `acserverconfig.xml`, contains an XML element named `<ResourceGroup>` for each resource group. This element contains attributes that affect the configuration of the resource group for all templates.

A `<ResourceGroup>` element contains the following properties for a resource group:

- **Name**

Name of the resource group.

- **Type**

Type of job the resource group supports. Resource groups support the following job types:

- **View**

Supports running a job synchronously, or unscheduled, and viewing a document.

- **Async**

Supports running a job asynchronously, or scheduled, and printing a document.

- **Volume**

Supports specifying a particular volume. Default value is all volumes.

- **Disabled**

Supports disabling and enabling the resource group. Value is true or false. Default value is false.

- **Reserved**  
Supports reserving a synchronous resource group for targeted requests.
- **ReportType**  
Supports running a report type of JavaReport.
- **Description**  
The resource group description.
- **MaxPriority**  
Used by a resource group having a Type of Async. The maximum priority a job can have. A job with a higher priority than another job runs first. The range of possible values is 0 through 1000.
- **MinPriority**  
Used by a resource group having a Type of Async. The minimum priority a job can have. A job with a lower priority than another job runs after the higher priority job. The range of possible values is 0 through 1000.
- **WorkUnitType**  
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.

## Configuring data source connections in BIRT iHub

A connection configuration file is an XML file that specifies the data source connection properties to use when BIRT 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 in the template settings in `acserverconfig.xml`.

To specify the location of a data configuration file in `acserverconfig.xml`, add the `ConnConfigFile` property containing the configuration file path as shown in the following code:

```
<Config>
  <Templates>
    <Template ConnConfigFile="config_file_path">
      . . .
    </Template>
  </Templates>
</Config>
```

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 2-7 externalizes the file path to the connection profile, C:\PostgreSQL.profile.

**Listing 2-7** BIRT connection configuration file example

---

```
<oda-data-source
  extensionID="org.eclipse.birt.report.data.oda.jdbc" name="JDBC
  Data Source - PostgreSQL" id="783">
  <property name="odaDriverClass">com.actuate.jdbc.postgresql.
    PostgreSQLDriver
  </property>
  <property name="odaURL">jdbc:actuate:postgresql://DBSRV1-W2K
  </property>
</oda-data-source>
<ConnectOptions Type=".eclipse.birt.report.data.oda.jdbc_ JDBC
  Data Source - PostgreSQL ">
  <Property PropName="OdaConnProfileStorePath">C:\MyPath
  </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\_PostgreSQL.

---

## Configuring an Apache web server for load balancing and proxying

An system administrator can configure an Apache web server to perform both load balancing and general proxying against a BIRT iHub cluster. Following are the minimum steps necessary to perform this task.

- 1 Create a cluster using System Console—Clusters.
- 2 Set up an Apache HTTP server with mod\_headers, mod\_proxy, mod\_proxy\_balancer, and mod\_proxy\_http enabled.
- 3 Use the following example XML to create an Apache VirtualHost.

```
<VirtualHost _default_:*>
  # required
  ProxyRequests Off
  ProxyPreserveHost On

  # add cookie for sticky session
```

```

Header add Set-Cookie "IHUB_ROUTE=.%{BALANCER_WORKER_ROUTE}e;
    path="/" env=BALANCER_ROUTE_CHANGED
# load balancing on 2 internal nodes ("node-01" and "node-
02")
<Proxy balancer://ihub/>
    BalancerMember http://node-01:8700 route=1
    BalancerMember http://node-02:8700 route=2
    ProxySet stickysession=IHUB_ROUTE
</Proxy>


# do not forward balancer-manager requests
ProxyPass /balancer-manager !

# forward all requests to load balancer
ProxyPass / balancer://ihub/
</VirtualHost>

```

**4** Set the cluster URL to the Apache website address by performing the following tasks:

- 1 On System Console—Clusters, edit the cluster created in step 1 by left-clicking the icon next to the cluster name.
- 2 On Cluster Configuration, left-click the icon next to the Cluster Status indicator and choose Edit Cluster Properties, as shown in Figure 2-76.




**Figure 2-76** Choosing to edit cluster properties

- 3 On Edit Cluster Properties, type the Apache website address in Cluster URL, for example, `http://mywebsite.com`, as shown in Figure 2-77.


**Figure 2-77** Specifying the Apache website address  
Choose OK.

- 5 Restart the Apache Web server to apply changes, if necessary.
- 6 Restart the cluster by performing the following tasks:
  - 1 On Cluster Configuration, left-click the icon next to the Cluster Status indicator and choose Stop Cluster, as shown in Figure 2-78.



**Figure 2-78** Stopping the cluster

- 2 On Cluster Configuration, left-click the icon next to the Cluster Status indicator and choose Start Cluster, as shown in Figure 2-79.



**Figure 2-79** Starting the cluster

As an alternative to restarting the cluster from Clusters—Cluster Configuration, you can restart the Actuate PostgreSQL for iHub 3 Service on every cluster node in the cluster.

Users can now access BIRT iHub from a web browser using the following URL:

`http://mywebsite.com/iportal`

# 7

## Managing system administrators

This chapter contains the following topics:

- About Settings
- Viewing System Information
- Working with system administrators
- Configuring Email Settings

---

## About Settings

In System Console, choose Settings to perform the following tasks, as shown in Figure 3-1:

- View System Information.
- Create a System Administrator user.
- Edit a System Administrator user.
- Delete a System Administrator user.



**Figure 3-1** Choosing Settings


The following sections describe how to perform these tasks.

---

## Viewing System Information

On Settings, choose System Information to view the following information, as shown in Figure 3-2:

- JSAPI Version  
The JavaScript API version number
- Product version  
The System Console build number

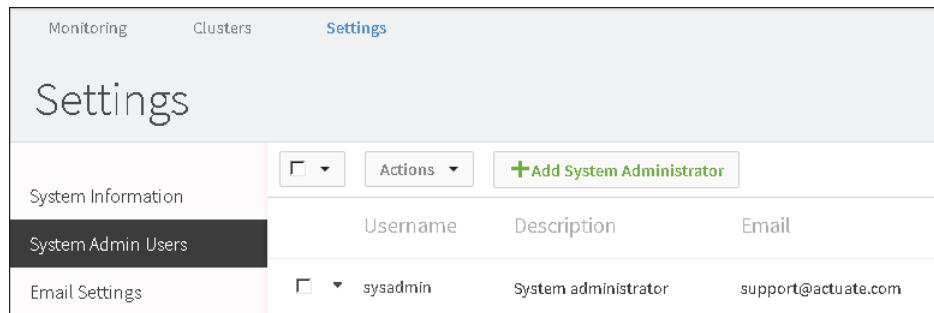


**Figure 3-2** Viewing System Information

---

## Working with system administrators

On Settings, choose System Admin Users to work with System Administrator users, as shown in Figure 3-3.



**Figure 3-3** Working with System Administrator users

Working with System Administrator users includes creating, deleting, and editing a System Administrator user. The following sections describe how to perform these tasks.

### Creating a system administrator

A system administrator can create other System Administrator users.

#### How to create a system administrator

- 1 On Settings—System Admin Users, choose Add System Administrator to create a new System Administrator user, as shown in Figure 3-3.
- 2 Specify the following properties on Add System Administrator, as shown in Figure 3-4. A property name appearing with an asterisk (\*) next to the name is a required property.
  - Username  
Type the name of the new system administrator.
  - Email  
Type the new administrator e-mail address.
  - Description  
Type a description for the new administrator.
  - Language  
Choose the language System Console uses.

- Is Locked Out  
Type true or false.
- Password  
In Password, type a password for the new user. The password must be at least eight characters long, and must contain at least one lower case letter, one upper case letter, and one digit. If you create a user using IDAPI, there is no restriction on what the password can be.
- Confirm Password  
Type the password again.

Choose OK. Confirm that you want to save changes.

**Figure 3-4** Adding a system administrator

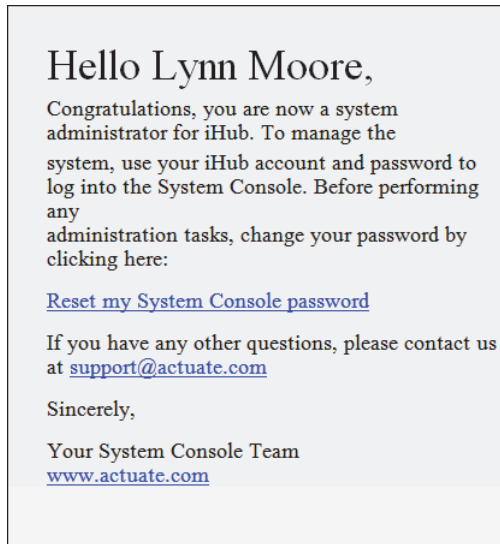
The new System Administrator user appears in the user list, as shown in Figure 3-5.

Username	Description	Email	Is Locked Out
<input type="checkbox"/> Lynn Moore	System Admin South	lmoore@company.com	false
<input type="checkbox"/> sysadmin	System administrator	support@actuate.com	false

**Figure 3-5** Viewing the new system administrator in the list



If you have e-mail notification enabled, BIRT iHub sends a notification e-mail to the new system administrator, as shown in Figure 3-6. For information on enabling e-mail notification, see “Enabling e-mail notification” in Chapter 6, “Managing clusters.”



**Figure 3-6** Viewing the notification e-mail

## Customizing the notification e-mail template file

The name of the e-mail template file that BIRT iHub uses to create the new system administrator notification e-mail is `EmailFormat.properties`. You can edit this file to customize it for your entity.

### How to customize the system administrator notification e-mail template file

- 1 In a default BIRT iHub installation on Windows, performed using the graphical installer, in which the install folder is `C:\Actuate`, navigate to the following file:

```
C:\Actuate\BIRTiHubVisualization\modules\SystemConsole\tomcat  
  \webapps\sysconsole\WEB-INF\lib\resource.jar
```

- 2 Using a file utility such as WinZip, open `com\actuate\umc\data\resource\EmailFormat.properties` and make any necessary modifications.


Save and exit the file.

## Editing a system administrator

When editing a system administrator, the properties are the same as when creating a system administrator, and the administrator can change any property.

### How to edit a system administrator

- 1 Left-click the arrowhead icon next to the System Administrator user in the list of users and choose Edit to edit a System Administrator, as shown in Figure 3-7.



**Figure 3-7** Choosing to edit a system administrator

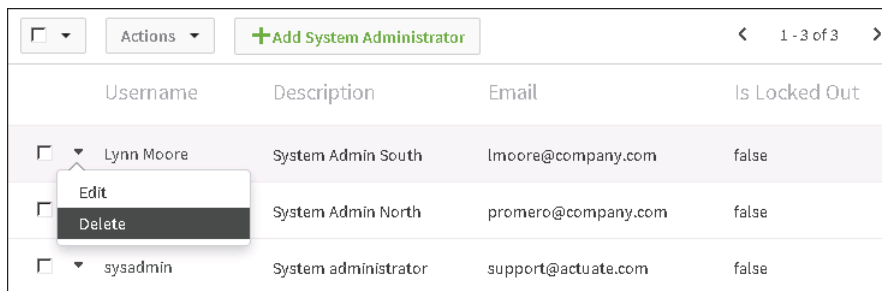
- 2 Make any necessary changes on Edit <System Administrator name> and choose Save.

## Deleting system administrators

The system administrator can delete one System Administrator user only, or multiple users simultaneously.

### How to delete a single administrator

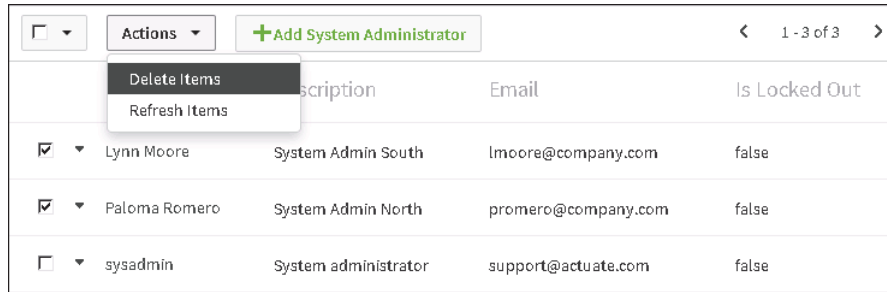
Left-click the arrowhead icon next to the System Administrator user in the list of administrators and choose Delete to delete an administrator, as shown in Figure 3-8.



**Figure 3-8** Deleting a system administrator

### How to delete selected administrators

Check the boxes next to the administrators you want to delete. Then, left-click Actions and choose Delete Items, as shown in Figure 3-9. Choose Refresh Items to deselect selected administrators.



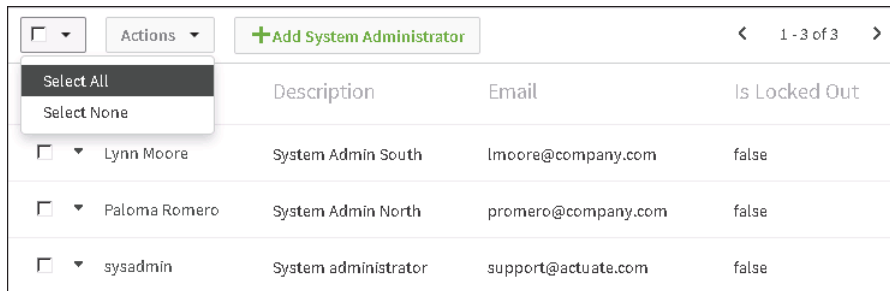
<input type="checkbox"/>	Actions	<a href="#">+Add System Administrator</a>	< 1 - 3 of 3 >		
		Description	Email	Is Locked Out	
<input checked="" type="checkbox"/>	▼	Lynn Moore	System Admin South	lmoore@company.com	false
<input checked="" type="checkbox"/>	▼	Paloma Romero	System Admin North	promero@company.com	false
<input type="checkbox"/>	▼	sysadmin	System administrator	support@actuate.com	false

**Figure 3-9** Deleting one or more users

### How to delete all administrators except one

Do not delete all system administrators. Delete all but one system administrator by performing the following tasks:

- 1 Left-click the arrowhead icon in the check box next to Actions and choose Select All, to select all administrators, as shown in Figure 3-10. Alternatively, choose Select None to deselect all administrators.



<input type="checkbox"/>	Actions	<a href="#">+Add System Administrator</a>	< 1 - 3 of 3 >		
		Description	Email	Is Locked Out	
<input type="checkbox"/>	▼	Lynn Moore	System Admin South	lmoore@company.com	false
<input type="checkbox"/>	▼	Paloma Romero	System Admin North	promero@company.com	false
<input type="checkbox"/>	▼	sysadmin	System administrator	support@actuate.com	false

**Figure 3-10** Selecting all system administrators

- 2 Deselect any system administrator, for example, sysadmin.
- 3 Left-click Actions and choose Delete Items, as shown in Figure 3-11. Alternatively, select Refresh Items to deselect administrators.

	Description	Email	Is Locked Out
<input checked="" type="checkbox"/>	Lynn Moore	lmoore@company.com	false
<input checked="" type="checkbox"/>	Paloma Romero	promero@company.com	false
<input type="checkbox"/>	sysadmin	support@actuate.com	false

**Figure 3-11** Deleting all administrators except sysadmin

## Configuring Email Settings

On Settings, choose Email Settings to configure the following properties, as shown in Figure 3-12.

- **Email Server Type**  
Type a name for the mail server. For example, type mail.smtp.host.
- **Email Server**  
Type the IP address or fully qualified domain name of the mail server. For example, type mailserver.companydomain.com.
- **Email Address**  
Type the e-mail address that appears in the From line of the e-mail notification. For example, type support@companydomain.com. BIRT iHub also sends a notification to this address when the mail server cannot deliver an e-mail notification to a user.

**Figure 3-12** Configuring an e-mail server

The system administrator can change monitoring to not send out alerts. If the system administrator chooses to have alerts, but does not set up an e-mail server configuration, BIRT iHub continuously generates errors saying that an e-mail

server configuration cannot be found. This message is classified as a severe error. To remedy this problem, edit the `lmserviceconfig.properties` file, and change the following setting from:

```
lmservice.monitoring.alert.notify.email.enable=true
```

to:

```
lmservice.monitoring.alert.notify.email.enable=false
```

After changing this setting to false, BIRT iHub still logs errors, but does not send e-mail alerts to the system administrator.

In a default BIRT iHub installation on Windows, performed using the graphical installer, in which the install folder is `C:\Actuate`, `lmserviceconfig.properties` is in the following location:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared  
  \config\lmsrvr
```



# Part **Three**

---

**Managing and backing up**





# Licensing BIRT iHub

This chapter discusses the following topics:

- Understanding licensing types
- Understanding licensing options
- Installing BIRT iHub System license files
- Understanding CPU binding

---

## Understanding licensing types

BIRT iHub System licensing supports running BIRT iHub with sets of features grouped as license options. You enable BIRT iHub System options using one or more of the following types of license models:

- **CPU Core**

Specifies the maximum number of CPU Cores that BIRT iHub System can use. Any number of users can access the licensed options on the system provided adequate licensing and capacity exists.
- **Work Unit (WU)**

Specifies BIRT iHub features and functionality using an aggregate model. This plan defines each BIRT iHub System resource as a work unit.

Similar to CPU Core licensing, but defined at a more granular level. With Work Unit Licensing, the customer can license just the precise amount of capacity needed for application requirements. Any number of users can access the licensed options provided sufficient capacity has been purchased.
- **Capacity Edition**

A BIRT on Demand licensing option that provides a pre-packaged amount of dedicated capacity for a customer application. With instance licensing, the customer does not need to count named users. Multiple instances can be combined to meet capacity needs.
- **Subscription**

A monthly or annual payment option that permits the use of the licensed software that includes maintenance. Offered with some of the other licensing models.

A subscription license is not a perpetual license. Once the subscription term expires, the software cannot be used.
- **Software as a Service (SaaS)**

Some products are offered as a Software as a Service (SaaS) option, providing customers with a solution without incurring the acquisition and management costs of hardware and traditional perpetual licenses.
- **BIRT Analytics**

A combination of a base license, a data-volume component measured in data rows, and a specified number of named-user licenses.
- **Packages**

Some options are offered as packages to customers for convenience and cost-saving benefits. These packages can be used in conjunction with individually selected options.

In a CPU Core and Work Unit licensing, Actuate currently uses the Standard Performance Evaluation Corporation (SPEC) standard benchmark for measuring machine capacity to establish the value of iHub’s performance on any given system.

## Understanding licensing options

Table 4-1 lists and describes BIRT iHub System license options. BIRT iHub System options are separately licensed products. Some license options require other options to be licensed before their functionality is available to users. Table 4-1 also describes these prerequisites.

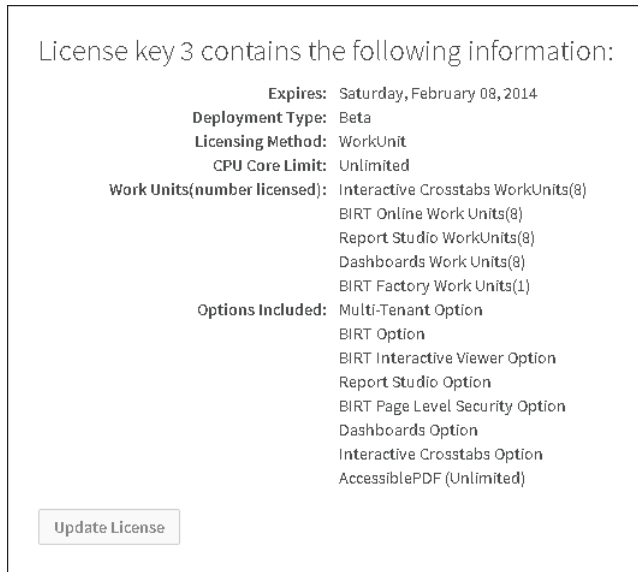
**Table 4-1** BIRT iHub System license options

Option	Description	Supported releases
Accessible PDF	Allows generating Accessible PDF output documents.	iHub
BIRT	Allows users to publish and run a BIRT design using BIRT iHub. This option is a requirement for Page Level Security and Interactive Viewer.	11, iHub
Dashboards	Enables the operational and analytical dashboard capabilities of BIRT iHub. This option is a requirement for the Metrics Management option.	iHub
Interactive Viewer	Allows a user who has the BIRT Option to use BIRT Interactive Viewer to view and interact with a BIRT document.	11, iHub
		<i>(continues)</i>
Metrics Management	Allows a user to create, execute, and view dashboard files, manage key performance indicators (KPI), publish briefing books, and export generated data. This option requires the Dashboards option.	11, iHub
Multi-Tenant	Allows a BIRT iHub System user to access more than one volume. This option is available with an Unlimited User CPU License.	11, iHub
Page Level Security	Controls access to structured content available on the web. This option works for designs created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or user group.	11, iHub

**Table 4-1** BIRT iHub System license options (continued)

Option	Description	Supported releases
Report Studio	Allows a user to create BIRT content using a web-based development environment. The BIRT option is required to view content.	iHub

To determine the license options installed on iHub, log in to System Console, and choose Show License. The license options appear, as shown in Figure 4-1.



**Figure 4-1** iHub License options

## Installing BIRT iHub System license files

Actuate provides a license file to use when installing BIRT iHub System. New customers receive an e-mail containing a temporary BIRT iHub license file to use for the initial installation after Actuate processes the order. The temporary BIRT iHub System license expires 45 days after installation.

Actuate license enforcement for BIRT iHub requires a single, shared license for all nodes in a cluster. The name for the BIRT iHub license file uses the following format:

`Actuate_iHub_key_xxxxxxx.xml`

XXXXXXX is a unique seven-digit number generated by Actuate Licensing when it creates the license file.

Actuate BIRT iHub System customers perform an initial installation using a temporary license. After installing BIRT iHub System using the temporary license, the login screen displays two messages.

The following message appears on the login screen for a license that has an expiration date:

Reminder

```
Your BIRT iHub license expires in [the number of days] days, on
[the specified date]. When the current license expires, the
iHub will shut down and require a new license to restart.
Please contact Actuate to purchase a new license.
```

The following message about how to obtain the second license file from Actuate Licensing appears until you install the new license issued by Actuate Licensing:

Reminder

```
One or more iHubs in your BIRT iHub System are in violation of
the node locked BIRT iHub license. After the grace period
expires, the iHubs that violate the node locked BIRT iHub
license cannot be restarted. Please contact Actuate Licensing
(licensing@actuate.com or http://www.actuate.com/licensing), or
your representative, and request a new license file for the
iHub nodes that are in violation. Please restart the iHubs on
the nodes after updating the license key file.
```

You have 45 days to apply for and install the license file after you install Actuate BIRT iHub System.

After installing BIRT iHub System, the installation informs a customer requiring a license to obtain the machine ID information on which Actuate BIRT iHub is running and transmit this information to Actuate Licensing. The machine ID is displayed in the reminder message. You can also use the utility, `acmachineid`, to obtain the machine ID. For information on how to use the `acmachineid` utility, see “How to use the `acmachineid` utility,” later in this chapter.

After receiving the machine ID information, Actuate Licensing issues a new Actuate BIRT iHub System license file.

## About the license file

This license file specifies the available BIRT iHub license options and node-key information for the cluster nodes. This license file must be in a shared location, specified by the `<AC_CONFIG_HOME>` attribute of the `<Server>` element in the `acpmdconfig.xml` file of each node, and accessible to all nodes in the cluster.

A node key associates a BIRT iHub node in a cluster with the machine ID. The node-key licensing mechanism restricts the iHub node installation to that machine.

On startup, each node in the cluster checks the shared license file, verifies the installed options, and determines whether its node key, which is generated at run time, matches the license information. If the node key matches, the node joins the cluster. Otherwise, it shuts down with an error if the node-lock-violation grace period has been exceeded.

A license file with an expiration date remains valid until the specified date. If the license file is about to expire, the system reminds you that the file expires on the specified date when you log in to System Console or Visualization Platform. Reminders also appear in the system log file. To arrange for a permanent license file, or if you have a problem with an expiring file, please contact Actuate Licensing at [licensing@actuate.com](mailto:licensing@actuate.com).

When upgrading a cluster node or installing BIRT iHub on a new machine, the customer must request a new license and supply the machine ID of the new machine.

## Collecting machine information for a license

After installing BIRT iHub System using a temporary license file, such as an evaluation license, you must collect information about the machines running BIRT iHub software and send it to Actuate Licensing. During the installation process, the install program prompts you to provide the location of the license file. After providing the location of the license file, the install program issues a prompt similar to the following message.:

```
The iHub system license file is locked to the machines that are
used in the iHub system. The following machine id must be used
to request a node key license file from Actuate:
```

```
IORRHEHs6S5UCsEtrdVu6jOixmzvFY3BbOqXLiwsWQGdceJmKYYaEu0j18lQxjM
sYCxka3hVkdZFGwkmQMxb+hgKaz4om2vLUcS0ocYTA7Ta6VTMavLFQo7bEjRyr
olwxAKu0Vr4NA6o8uWCzjGZXX8KrjViSUoRoj70hWOY=
```

```
Please contact Actuate Licensing (licensing@actuate.com or
http://www.actuate.com/licensing), or your representative, and
request a node locked iHub system license.
```

```
The machine id required for the node locked iHub system license
can also be generated by using the acmachineid utility that can
be found in the ACTUATE_HOME\iHub\bin folder.
```

The format of the alphanumeric string for the machine ID and location of the license file are different depending on the operating system. On a Windows system, the unique identifier for the network card is the source of the machine ID. You must have at least one network card enabled on the BIRT iHub machine to be able to obtain the machine ID.

After installing BIRT iHub, you must run the utility, `acmachineid`, from the command line to generate the machine ID information. Copy the machine ID in

the command prompt to a file or e-mail message and send it to Actuate Licensing. Actuate Licensing processes your request and sends the new license file for BIRT iHub System.

### **How to use the acmachineid utility**

Use the acmachineid utility to obtain the machine ID information by performing the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME\bin.
- 2 Type the following command and press Enter:

```
acmachineid
```

The utility provides output in the following format:

```
C:\Actuate\iHub3\modules\BIRTiHub\iHub\bin>acmachineid
STATUS:          OK
GEN_VERSION:     23 Augusta R1 Development
GEN_BUILD:       DEV131216
SERVERNAME:      W7OOTB
MACHINEID:
    u0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67Gb08VjcFs6pcuAgbt
    DaZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwtLykwJisqMS0EhYe5sC
    YoKjG+XL2UEEnL2GGhLtI9fJUMyzZORS1WPPRyrolwxAKu0Vr61qxoMC8Khvp
    y1HGtKvIhgEcasrUerKE674lnutEXTIVo+
```

Send the licenselocation.txt file automatically generated by the execution of the acmachineid utility to Actuate Licensing as an email attachment.

### **How to obtain a license file**

To register as a new user and obtain a new license file for a product, go to the Actuate Support web site at the following location:

```
http://support.actuate.com/SiteRegister
```

Enter the new user account and license key request information.

A maintenance customer should have login information for the Actuate e.Support web site. If you do not have access, please contact Actuate Licensing at [licensing@actuate.com](mailto:licensing@actuate.com).

If you are not a direct Actuate customer, contact the partner or distributor who supplies the product for the license file. If you have a problem obtaining a license file from this source, please contact Actuate Licensing at [licensing@actuate.com](mailto:licensing@actuate.com).

### **Updating the BIRT iHub System license file**

After performing an installation of BIRT iHub System and transmitting the required machine ID information to obtain a license, Actuate sends an e-mail containing an attached .txt (TXT) file. Replace the .txt extension with a .zip (ZIP) extension and open the file. This ZIP file contains the following files:

- `readme.txt`  
Instructions for installing BIRT iHub System using a license file and for obtaining a license file
- `Actuate_iHub_key_XXXXXXX.xml`  
BIRT iHub System license

An Actuate license file is an XML file. Actuate Licensing sends this XML file inside a ZIP file with its extension altered to TXT because transmitting a file with a .xml extension can cause problems in an e-mail system.

### How to install the license file

To install the license file, perform the following steps:

- 1 Extract the contents of the ZIP file to a location on your local file system.
- 2 Log in to System Console. For example, open a browser and type the following URL address, `http://localhost:8500/sysconsole`, and log in using the system administrator ID and password.
- 3 In Clusters—License, choose Update License.
- 4 On Update License, choose Choose File and browse to navigate to the location where you extracted the contents of the ZIP file. Select the Actuate BIRT iHub System license file and choose Save to apply the license.

If iHub requires a system restart to update the license file, the following message appears:

```
The license file cannot be applied without a server restart.
Please copy the license file to the iHub license file
location and restart the iHub system.
```

- 5 Restart any node where the node-key configuration changed.

If you change the machine on which you installed BIRT iHub, you must re-apply to Actuate Licensing for a new license file. If you replace the network card on some machines, such as a Windows system, you may have to obtain a new license file since the unique identifier for the network card may be the source of the machine ID. If you have a license file installed and a reminder message appears when logging into System Console, contact Actuate Licensing and provide the current BIRT iHub System license file with the output from the machine ID utility.

`Actuate_iHub_key_XXXXX.xml` will contain the node key information for the stand-alone machine or all machines in a cluster. There is no separate node license file for each machine.

Listing 4-1 shows a sample of the node key information the license contains obtained from `acmachineid` output submitted to Actuate Licensing.

#### **Listing 4-1** Viewing license node key information

---

```
<NodeKeys>
```



```

<NodeKey
  MachineId="E0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKws2t1W7qM67GbO8
  VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwtLykDa/
  wJisqMS0EhYe5sCYoKjG+XL2UEEnL2GGhLtI9fJUMYzZORKk23jrxaswUDig
  Ksvlc1A6q8UbmrrAYHD8GgtptuiAmxWt4xjEM6rq1msNEW/4Vjm40Kx1kSv"
  ServerName="W7CLSTRNODE1" />
<NodeKey
  MachineId="I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKws2t1W7qM67GbO8
  VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwtLykDa/
  wJisqMS0EhYe5sCYoKjG+XL2UEEnL2GGhLtI9fJUMYzZORKk23jrxaswUDsg
  Ksvlc1A6q8UbmrrAYHD8GgtptuiAmxWt4xjEM6rq1msNEW/4ViMC0KDBkSn"
  ServerName="W7CLSTRNODE2" />
</NodeKeys>

```

## About modifying a license

If you decide later to license additional BIRT iHub options, the existing license file becomes invalid. You must install a new license file. Contact Actuate Licensing for the new license file.

## About modifying the data collection option

If you are evaluating the BIRT iHub product and do not wish to participate in the collection and reporting of product usage information, you must disable the collection of this information.

### How to disable data collection and reporting

1 Shut down the iHub 3 Service.

- On Windows, use Services to shut down the Actuate iHub 3 Service.
- On Linux, run `BIRTiHubEvaluation/stopiHub.sh`.

2 Remove the activityfilter filter and filtermapping elements from the file:

```

BIRTiHubEvaluation\modules\BIRTiHub\iHub\web\portal\WEB-INF
\web.xml .

```

3 Start the iHub 3 Service.

- On Windows, use Services to start the Actuate iHub 3 Service.
- On Linux, run `BIRTiHubEvaluation/modules/BIRTiHub/startiHub.sh`.

---

## Understanding CPU binding

BIRT iHub System supports CPU binding on a machine with an appropriate CPU-based license. CPU binding restricts a process or processes to run on a subset of CPU cores. If you bind the BIRT iHub System to a subset of CPU cores,

only those CPU cores count toward the total number of licensed CPU cores. The CPU limit in the license file applies to all CPU cores for all machines in the cluster. Depending on the operating system and specific system command, you can restrict other processes from running on the processor to which you bind a process.

You can bind BIRT iHub processes to a specific set of processors on a machine that runs a Windows or Linux operating system. The default configuration does not bind BIRT iHub to a set of processors. In the default configuration, all processors on a BIRT iHub machine count toward the maximum number of licensed CPU cores.

To bind BIRT iHub to a set of processors, bind the iHub Daemon (ihubd) to the processors. ihubd starts all BIRT iHub processes. The processes inherit the binding from ihubd.

In a cluster, BIRT iHub counts only the processors on nodes that join the cluster and run the ihubc process. An ihubc process runs when a node is online. BIRT iHub counts the number of processors on a machine when the first ihubc process starts.

This section contains information on the following topics:

- Configuring CPU binding on Windows
- Configuring CPU binding on Linux
- Checking BIRT iHub bound processors
- Configuring e-mail for CPU license problems

## **Configuring CPU binding on Windows**

You can perform the following types of CPU binding on Windows:

- Binding to specific CPU cores
- Binding to multiple-core CPU cores
- Binding an Actuate process to a processor

The following sections describe these features.

### **Binding to specific CPU cores**

On a multiple-CPU machine running the Windows operating system, the operating system assigns an ID number to each processor. Windows Task Manager lists the IDs of the available processors. The numbering starts at 0.

If you are not licensed to use all the CPU cores on the iHub 3 host machine, you must bind iHub 3 to the appropriate number of CPUs. To bind iHub3 to a set of processors, you modify the acpmdconfig.xml file in AC\_SERVER\_HOME\etc.

For example, to bind processors to four logical cores, add the following line to `acpmdconfig.xml`:

```
<DaemonCPUaffinity>0,1,2,3</DaemonCPUaffinity>
```

After restarting iHub, `ihubc.log` shows this setting in the System Bound field, for example:

```
System Bound           : 4
```

### How to configure CPU binding on a Windows machine

The following example shows the settings for a two CPU (0,1) machine running Windows, which uses only one CPU (0) for BIRT iHub:

- 1 To set CPU affinity, edit the `<DaemonCPUaffinity>` element in `acpmdconfig.xml` located in `AC_SERVER_HOME\etc` as follows:

```
<DaemonCPUaffinity>0</DaemonCPUaffinity>
```

- 2 Restart iHub.

- 3 Use Process Explorer to verify all BIRT iHub processes, including `ihub`, `ihubc`, and `ihubd` are using only CPU 0. Alternatively, verify the CPU binding by checking the Processor Affinity of the BIRT iHub process using Task Manager.

### Binding to multiple-core CPU cores

You can also perform multiple-core CPU binding, similar to the way you bind to a single CPU, as described in the previous section. To BIRT iHub, each core appears as a logical CPU.

For example, on a dual-core, two-CPU system, setting the `DaemonCPUaffinity` value to `0,1` binds BIRT iHub to both cores on the first CPU. Setting the value to `0,2` binds BIRT iHub to one core on each CPU. Setting the value to `0` binds BIRT iHub to one core on the first CPU.

Actuate does not recommend restricting BIRT iHub processing on a multiple-core CPU machine to one core for licensing purposes. BIRT iHub System achieves significant performance gains on a multiple-core CPU machine.

For example, BIRT iHub scales nearly perfectly from 1 to 2 cores and gets 50% better throughput on a dual-core system than on a two-CPU system.

### Binding an Actuate process to a processor

If you bind a BIRT iHub `ihubd` process to a subset of CPU cores on a machine, you can also bind the Factory, View, Integration, and Caching processes to a specific CPU. Under some conditions, binding an Actuate process to a specific CPU can enhance performance. Binding an Actuate process to a CPU has no effect on the CPU calculations BIRT iHub performs to determine the maximum number of licensed CPU cores.

If you bind a process to a CPU core, you must bind the CPU core to both BIRT iHub ihubd and the process. BIRT iHub writes to the error log and stops the process if you bind a process to a CPU that you do not bind to BIRT iHub ihubd.

To bind a BIRT iHub ihubd process to CPU processors, use the ProcessorAffinity element in the acserverconfig.xml file for BIRT iHub. List the IDs for the CPU cores to which to bind a process as Item subelements in the following ProcessorAffinity elements:

- To bind Factory processes, specify the CPU IDs in the ProcessorAffinity element within the ReportingService element.
- To bind View processes, specify the CPU IDs in the ProcessorAffinity element within the ViewingService element.
- To bind Integration processes, specify the CPU IDs in the ProcessorAffinity element within the IntegrationService element.

You must also ensure that you bind the specified CPU cores to ihubd for the BIRT iHub machine. For example, on a four-CPU machine, the following ProcessorAffinity example binds View processes to CPU IDs 0 and 2:

```
<ViewingService
  EnableViewingService="true"
  <ProcessorAffinity>
    <Item>0</Item>
    <Item>2</Item>
  </ProcessorAffinity>
/>
```

## About processors and hyperthreading

Some Intel processors use hyperthreading, a technology that counts each physical processor as a specific number of logical processors. The operating system and any programs running on the machine see the number of logical processors, not the number of physical processors.

When a machine uses hyperthreading, Windows Task Manager lists the logical processors, not the physical ones. You specify the number of logical processors in the environment variable. When a machine uses hyperthreading, BIRT iHub calculates the number of bound processors by dividing the number of bound logical processors by the number of logical processors for each physical processor. If the result contains a decimal component, BIRT iHub uses the next highest integer. For example, it rounds 4.3 to 5. In the following example, a machine has four physical processors. With hyperthreading enabled, each physical processor corresponds to two logical processors. The machine has the following logical processors available:

- Physical processor 0 corresponds to logical processors 0 and 1.
- Physical processor 1 corresponds to logical processors 2 and 3.

- Physical processor 2 corresponds to logical processors 4 and 5.
- Physical processor 3 corresponds to logical processors 6 and 7.

If you bind BIRT iHub to the five logical processors 0, 2, 3, 6, and 7, it calculates the number of bound processors as:

$$5/2 = 2.5$$

BIRT iHub rounds this number up to determine that you have three bound processors.

## Configuring CPU binding on Linux

The following section describes how to perform various CPU-binding operations in the Linux environment.

The `ihubd` process is the root parent process for all other BIRT iHub processes, so CPU binding can be done only for `ihubd`. Binding must be done before starting the `ihubd` process. Binding a running `ihubd` process has no effect.

On a multiple-CPU machine running the Linux operating system, the operating system assigns an ID number to each processor. The numbering starts at 0.

If you are not licensed to use all the CPU cores on the iHub 3 host machine, you must bind iHub 3 to the appropriate number of CPUs. To bind iHub3 to a set of processors, you modify the `acpmdconfig.xml` file in `AC_SERVER_HOME/etc`. For example, to bind processors to four logical cores, add the following line to `acpmdconfig.xml`:

```
<DaemonCPUaffinity>0,1,2,3</DaemonCPUaffinity>
```

### How to configure CPU binding on Linux

The following example shows the settings for a four CPU (0,1,2,3) machine running Linux, which uses only two CPUs (0,1) for BIRT iHub:

- 1 In Linux, log in as root and use the `less` command to view CPU core information, as shown in the following example:

```
less /proc/cpuinfo
```

Use `<Ctrl>+Z` to suspend the command.

- 2 To verify CPU cores, use the `cat` command for more detailed information, as shown in the following example:

```
sudo cat /proc/cpuinfo
```

- 3 Use the `taskset` command, referencing the process ID (PID) to verify processor affinity, as shown in the following example:

```
taskset -p -c PID
```

The taskset command binds to the number of logical cores. If a machine does not have hyperthreading enabled, you will not see any difference between the physical and logical cores.

Using the -c option provides a processor affinity mask in list form rather than a bit mask for a PID specified using the -p option. For example, a typical affinity list generated by these arguments is 0,2,3-5.

- 4 Use the `ps | grep` commands to verify the current core settings for all running processes, as shown in the following example, where `ihub_id` is the user that runs the BIRT iHub processes:

```
ps -e -o pid,cpuid,comm,user | grep ihub_id
```

- 5 To bind CPU cores on Linux, perform the following tasks:

- 1 Stop BIRT iHub, including `ihubd`. Make sure no processes are running by typing the following command, where `ihub_id` is the user that runs the BIRT iHub processes:

```
ps -e -o pid,cpuid,comm,user | grep ihub_id
```

- 2 In a typical installation, using a text editor such as `vi`, open the `acpmdconfig.xml` file located in `AC_SERVER_HOME/shared/config` and edit the `<DaemonCPUaffinity>` element as follows:

```
<DaemonCPUaffinity>0,1</DaemonCPUaffinity>
```

- 3 Restart BIRT iHub. Make sure that all processes started and verify the core to which each process is bound by running the following command, where `ihub_id` is the user that runs the BIRT iHub processes:

```
ps -e -o pid,cpuid,comm,user | grep ihub_id
```

## Checking BIRT iHub bound processors

BIRT iHub performs the following bound processor checks:

- The number of processors a cluster uses
- The set of bound processors

## Determining the number of processors BIRT iHub System uses

When the `ihubd` process starts the first `ihubc` process on a machine, the `ihubd` determines the number of processors to which BIRT iHub is bound and stores the list of bound processors.

If you change the processor binding, BIRT iHub does not recognize the changes until you shut down all `ihubc` processes on the machine and restart one of the `ihubc` processes.

For example, a cluster that has a maximum licensed CPU limit of nine processors consists of two nodes, machine A and machine B.

The machines have the following configuration:

- Machine A has four processors with no processor binding. All the processors can run Actuate processes. BIRT iHub manages a volume.
- Machine B has eight processors with BIRT iHub bound to five processors. There is no ihubc process running on the machine, only the ihubd process.

The cluster counts four processors, the processors on machine A. If you start an ihubc process on machine B, BIRT iHub on machine A counts the five bound processors on the machine and increases the cluster processor count to nine, four on machine A and five on machine B.

If you bind the ihubd process on machine B to six processors, the change has no effect until you shut down all the running ihubc processes on machine B and restart an ihubc process on machine B.

After you stop the ihubc processes and restart an ihubc process on machine B, BIRT iHub System detects that the number of processors in the cluster is ten, which is greater than the maximum number of nine licensed processors. When the number of CPU cores exceeds the number of CPU cores your license permits, BIRT iHub does not start and returns an error message to System Console.

## **Understanding CPU binding validation while BIRT iHub is running**

When BIRT iHub is running, each ihubc process periodically compares the list of processors to which it is bound with the list to which it was bound when it started. If the lists differ:

- BIRT iHub writes a message with the processor information to the log file. The message contains the maximum number of processors the BIRT iHub license file permits and the following information:
  - Current and original number of bound processors
  - Current and original list of bound processors
- If configured, BIRT iHub sends an e-mail message to the administrator. The message states that the BIRT iHub System will shut down in one hour if the list of bound processors is not corrected. The e-mail message contains the information that BIRT iHub sends to the log file.

You must rebind the ihubc process to the same processors to which it was originally bound. During the next hour, any attempt to use the ihubc services fails and a message is written to the appropriate log file. If the list of processors is not restored after an hour, each BIRT iHub in the cluster shuts down and writes an error to its log file.

After updating a CPU-limit license, the system administrator must perform a complete restart of the system to refresh the list of processors that BIRT iHub uses to periodically compare the list of currently bound processors to the list to which

it was bound when it started. Before restarting, the system administrator must also edit the `acpmdconfig.xml` file to adjust the CPU affinity to the new settings specified in the `<DaemonCPUaffinity>` element.

## Understanding CPU binding validation when a volume comes online

BIRT iHub uses a separate `ihubc` process to manage each volume on a machine. When you take a volume online, `ihubd` starts an `ihubc` process.

When `ihubd` starts an `ihubc` process, the `ihubd` compares the list of processors to which the `ihubc` process is bound to the original list of processors to which the `ihubd` is bound. If the lists differ:

- The `ihubc` process writes an error to its log file and shuts down.
- BIRT iHub does not take the volume online.  
A message in the configuration states that the binding of the new process differs from the original binding of the parent process.

## Understanding CPU binding validation when running BIRT iHub processes

Each Factory and View process periodically compares its list of bound processors with the list of processors to which it was bound at startup. If the lists differ, the process writes an error to its log file and shuts down.

## Configuring e-mail for CPU license problems

BIRT iHub System can send e-mail messages to an administrator if a change in processor binding violates the maximum number of licensed CPU cores for BIRT iHub System. To send e-mail about a CPU license problem, set up BIRT iHub System by completing the following tasks in this order:

- 1 Configure every BIRT iHub node to send e-mail.
- 2 Specify the administrator e-mail address for BIRT iHub System.

Specify an administrator e-mail address as the value for the Account to receive administrative e-mail parameter. Set the value by logging into System Console, and choosing Settings—System Admin Users. Choose the `sysadmin` user name. In Edit `sysadmin`, type the email address for the `sysadmin` user. Choose OK.

For example, the following e-mail address sends e-mail to a user named `admin` at a company for which the domain is `mycompany`:

```
admin@mycompany.com
```

- 3 Restart BIRT iHub System. Restarting applies the changes after you set or change the e-mail address.



For more information on configuring BIRT iHub System to send e-mail messages, see Chapter 6, “Managing clusters,” later in this book.



# Backing up BIRT iHub System

This chapter discusses the following topics:

- Performing a BIRT iHub System backup
- Backing up and restoring a BIRT iHub System that uses a PostgreSQL database

---

## Performing a BIRT iHub System backup

When performing a backup, it is important to note that there are two types of data:

- **Metadata**  
Information about BIRT iHub cluster and volume settings and data objects stored in third-party relational database management system (RDBMS) schemas.
- **Data**  
BIRT iHub cluster and volume data objects, such as designs, documents, and information objects, stored as files in storage locations on disk, and the acserverconfig.xml file containing BIRT iHub System configuration settings.

The administrator must back up all BIRT iHub System metadata and data to ensure the recoverability of the system in the event of failure. The third-party database that contains BIRT iHub metadata is a critical component of BIRT iHub System. The system administrator must take all necessary precautions to ensure that this database is properly backed up and available to safeguard cluster and volume metadata. Please consult Actuate Support at the time of installation if you have any questions about the backup, recovery, or failover procedures necessary to protect against the possibility of catastrophic failure.

### Managing the backup and recovery of BIRT iHub metadata and data files

A complete backup of BIRT iHub System must include the following items:

- A database backup of the cluster and volume schemas containing the metadata
- A copy of the folders from all volume storage file locations containing file data
- A copy of the acserverconfig.xml file containing BIRT iHub configuration information

In the Windows BIRT iHub environment, the default AC\_SERVER\_HOME path is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub
```

The default location for volume storage folders is AC\_SERVER\_HOME\shared. The absolute path of this location is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
```

The default `acserversconfig.xml` file path is `AC_SERVER_HOME\shared\config`. The absolute path of this location is:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared
  \config
```

Back up the metadata in the RDBMS at the same time that you back up the data files in the volume storage locations. A carefully coordinated backup ensures that a one-to-one correspondence exists between each entry in the volume metadata database and the data files.

The metadata backup on the RDBMS must be done before the backup of the data files in the volume storage locations. Files that are partially created when the metadata backup begins are either not yet registered in the database or are marked incomplete in the database. The metadata database does not retain a record of incomplete files.

When contacting Actuate Support to troubleshoot problems, it is best to provide a snapshot of the BIRT iHub System configuration, including the following items and information:

- A database backup of the cluster and volume metadata schemas
- The names of the volume schemas and users that iHub uses to connect to the RDBMS
- A copy of the `acserversconfig.xml` file containing BIRT iHub configuration information
- A copy of BIRT iHub logs

## Using RDBMS and file system backup utilities

The administrator must perform the metadata backup using the tools provided or supported by the RDBMS. Copying the physical files of a database at the operating system level while an RDBMS is running does not create a valid backup.

Most RDBMS backup tools can be scripted and run while BIRT iHub is using the database. PostgreSQL and Oracle also provide graphical administration tools in addition to command-line tools. This chapter provides instructions on how to perform a backup in the PostgreSQL RDBMS environment as a reference example. For more information on using other RDBMS systems and tools to back up and restore BIRT iHub schemas, see the vendor documentation.

### How to perform an BIRT iHub System backup

To back up BIRT iHub System, perform the following tasks:

- 1 Make sure that the autoarchive file purging process is not running.
- 2 Make an online backup of the cluster and volume schemas using the tools provided by the RDBMS.

- 3 Back up the volume data and system configuration files using the tools available in the operating system environment.

A metadata backup is consistent with a data backup only if the file purging process that runs during an autoarchive operation does not occur between the time you back up the metadata and the time you back up the data. In System Console, the administrator can specify when the file purging process runs. Configure the following time-related file purging properties to times that do not conflict with the time when the backup operation runs:

- Purge deleted files time  
Specifies the time when the file purging process runs to permanently delete expired files
- Expiration time of deleted files  
Specifies the length of time that must elapse before the file purging process permanently deletes an expired file

---

## Backing up and restoring a BIRT iHub System that uses a PostgreSQL database

PostgreSQL provides the pgAdmin graphical administration tool or the `pg_dump` and `pg_restore` command-line utilities to back up and restore a database. These PostgreSQL utilities run on the client not the server.

To back up a volume in the out-of-the-box (OOTB) PostgreSQL RDBMS environment, the administrator performs the following operations:

- Backs up cluster and volume schemas containing the BIRT iHub System metadata using the pgAdmin graphical administration tool or the `pg_dump` PostgreSQL command-line utility
- Backs up volume data and BIRT iHub System configuration files using operating system copy commands

Note that a backup of a PostgreSQL database is not portable across all operating systems.

To restore BIRT iHub System in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Restores BIRT iHub System cluster and volume schemas containing the BIRT iHub System metadata using the pgAdmin graphical administration tool or the `pg_restore` PostgreSQL command-line utility
- Restores volume data and BIRT iHub System configuration files using operating system copy commands

The following sections describe how to back up and restore a BIRT iHub System that uses the OOTB PostgreSQL database to store the metadata. These demonstrations serve as a detailed reference example. Other supported database systems, such as Oracle provide similar utilities and require comparable operational procedures.

## Backing up BIRT iHub System using pg\_dump

To back up BIRT iHub System using the PostgreSQL pg\_dump utility, perform the following tasks:

- Create a folder to contain the metadata and volume data backup files.
- Back up BIRT iHub System metadata using the pg\_dump utility.
- Back up the acserverconfig.xml file and volume data folders to the backup folder.

Create a folder to contain the metadata, configuration file, and volume data backup files outside the BIRT iHub data installation environment. To provide protection against single-point media failure, it is best to store the backup files at a storage location that is physically separate from the BIRT iHub System and volume data locations.

The following example shows a typical pg\_dump command used to export the contents of the iHub cluster and volume schemas to a backup file:

```
pg_dump.exe --host dbhost --port 8432 --username "postgres"  
--format custom --blobs --verbose --file "C:\Actuate  
\BIRTiHubVisualization\backup\ihub.backup" "dbname"
```

This pg\_dump command example uses the following arguments:

- host  
Specifies the host name of the machine where the PostgreSQL server is running, such as dbhost.
- port  
Specifies the port where the server listens for connection requests.
- username  
Specifies the user name for the connection to the PostgreSQL server, such as postgres.
- format  
Specifies the output format. The value custom creates a compressed archive that can be used as input to pg\_restore.
- blobs  
Specifies including blob data types.

- **verbose**  
Specifies the level of command-line status messages.
- **file**  
Specifies the output file, such as `ihub.backup`.
- **dbname**  
Replace this string in the example with the database name, such as `ihub`.
- **n or name**  
Species the schema name. Use multiple `-n` arguments to specify a list. Use wildcard notation to specify a character pattern, such as `ac_*`, to specify all volumes names that start with the prefix `ac_`. If `-n` is not specified, `pg_dump` exports all non-system schemas.

Alternatively, run the command at the volume schema level to back up individual volume schema to a separate archive. To run a backup using a script, set up auto-login using a `.pgpass` file. The file should contain connection information in the following format:

```
hostname:port:database:username:password
```

More information about setting up a scripted backup using a `.pgpass` file is available at:

<http://www.postgresql.org/docs/9.2/static/libpq-pgpass.html>

Back up BIRT iHub System metadata using `pg_dump` by performing the following tasks.

### How to run `pg_dump` from a command prompt

- 1 Open a command prompt.
- 2 Navigate to the following location:  

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\
  postgresql\bin
```
- 3 Execute the following command. Substitute your machine name for `urup` in this example:

```
pg_dump.exe --host urup --port 8432 --username "postgres"
  --format custom --blobs --verbose --file "C:\Actuate
  \BIRTiHubVisualization\backup\ihub.backup" "ihub"
```

This operation backs up the entire `ihub` database. If the `-n` argument specifying a specific schema or list of schemas is not specified, `pg_dump` exports all database schemas. Alternatively, you can back up only one volume schema by using the `-n` argument to specify a particular schema.

- 4 If prompted to enter a password, type the postgres superuser password.



`pg_dump` executes, writing status messages to the command prompt.

After backing up the BIRT iHub System metadata, back up the `acserverconfig.xml` file and volume data directories to the backup directory by performing the following tasks.

### How to back up the volume data folders

- 1 Open Windows Explorer and navigate to the config folder that contains `acserverconfig.xml` file. In a default BIRT iHub installation, this file is located in `AC_SERVER_HOME\shared\config`. For example:

```
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub\shared\config
```

- 2 Select `acserverconfig.xml`, right click, and choose Copy. Copy the file to the backup location. For example:

```
C:\Actuate\BIRTiHubVisualization\backup
```

- 3 Navigate to the folder or folders that contain volume data files, such as `AC_SERVER_HOME\shared\storage`. Right-click the folder, and choose Copy. Copy this folder to the backup location. For example:

```
C:\Actuate\BIRTiHubVisualization\backup
```

## Restoring BIRT iHub System using `pg_restore`

To restore a backed-up BIRT iHub System, perform the following tasks:

- Take BIRT iHub System offline.
- Delete the `acserverconfig.xml` file in `AC_SERVER_HOME\shared\config` and the volume data folder in `AC_SERVER_HOME\shared`.
- Copy the backed-up `acserverconfig.xml` file to `AC_SERVER_HOME\shared\config` from the backup folder and the volume data folder from the backup folder to `AC_SERVER_HOME\shared`.
- Restore the BIRT iHub system and volume metadata using the PostgreSQL `pg_restore` utility.
- Take BIRT iHub System online.

Alternatively, the administrator can restore an individual volume by selectively backing up and restoring only the related volume schema and data.

To begin a restore operation, take BIRT iHub System or the individual volume offline.

### How to restore the backed-up data folders

- 1 In Windows Explorer, navigate to `AC_SERVER_HOME\shared\config`.
- 2 Select `acserverconfig.xml`, right-click, and choose Delete. Confirm the deletion.

- 3 In AC\_SERVER\_HOME\shared, delete the volume folder, for example, AC\_SERVER\_HOME\shared\storage. Confirm the deletion.
- 4 In Windows Explorer, navigate to the following location:  
C:\Actuate\BIRTiHubVisualization\backup  
Select acserverconfig.xml, right-click, choose Copy, and copy this file to AC\_SERVER\_HOME\shared\config.
- 5 In C:\Actuate\BIRTiHubVisualization\backup, right-click the volume folder, for example, storage. Choose Copy, and copy this folder to AC\_SERVER\_HOME\shared.

Restore BIRT iHub System schemas using the command-line version of pg\_restore. The pg\_restore utility runs using arguments similar to the pg\_dump utility.

The following example shows a typical pg\_restore command used to import the contents of a backup file to the BIRT iHub System database:

```
pg_restore -h dbhost -p 8432 -U postgres -d db_name_
ihub_ihub.backup
```

Run pg\_restore from the command line by performing the following tasks.

#### **How to run pg\_restore from a command prompt**

- 1 Open a command prompt.
- 2 Navigate to the following location:  
C:\Actuate\BIRTiHubVisualization\modules\BIRTiHub\iHub  
  \postgresql\bin
- 3 Enter the following command. Substitute your machine name for urup in this example:

```
pg_restore.exe --host urup --port 8432 --username postgres
--dbname ihub --clean --verbose "C:\Actuate
\BIRTiHubVisualization\backup\ihub.backup"
```

Press Enter.

- 4 If prompted, type the postgres superuser password. Press Enter.  
pg\_restore executes, writing status messages to the command prompt.

Take BIRT iHub System or, alternatively, the individual volume online.

More information about backing up and restoring a volume schema using the PostgreSQL pg\_dump and pg\_restore utilities is available at:

<http://www.postgresql.org/docs/9.2/static/backup.html>

# Index

## A

- AC\_CONFIG\_HOME element 115
- AC\_SERVER\_HOME variable 130
- AC\_SHARED\_HOME variable 20
- access permissions. *See* privileges
- AccessiblePDF licensing option 113
- accessing
  - configuration parameters 9
  - Encyclopedia volumes 13, 113
  - iHub features 112
  - metadata 9
  - metadata databases 15
  - RDBMS documentation 5
  - shared resources 21
  - web-based content 113
- acmachineid utility 115, 117
- acserverconfig.xml 9, 30
- acserverlicense.xml 30
- Active Directory property 52
- Active Directory servers 14, 52, 59
- active requests 78, 80
- Actuate Viewer 76
- adding
  - cluster nodes 27, 14, 15, 17, 23
  - Encyclopedia volumes 14, 31
  - JDBC drivers 15
  - license files 115
  - licensing options 119
  - passwords. *See* passwords
  - system administrators 101
- “Admin” Group property 58
- administration console applications
  - See also* specific application
- administration tools (iHub) 131
- administrative reports 15
- administrators
  - adding cluster nodes and 9, 18, 19, 27
  - backing up Encyclopedia and 130, 131, 132
  - changing system 104
  - creating system 101
  - deleting 104, 105
  - deleting clusters and 69
  - failover procedures and 5
  - installing alternate databases and 15
  - launching iHub images and 9
  - managing cluster nodes and 27, 64, 65
  - managing iHub clusters and 12, 14, 26, 63
  - managing iHub services and 28
  - managing iHub System and 13, 2, 100
  - managing metadata databases and 5
  - obtaining licenses and 112
  - preventing data loss and 5
  - receiving CPU violation notices 126
  - scheduling file purging processes and 132
  - setting properties for 101
- Alert Name property 43
- alerts
  - changing 45, 46
  - configuring 14, 40
  - creating 2, 41, 42
  - displaying 8, 40
  - enabling or disabling 46
- Alerts list 3
- alternative databases. *See* third-party databases
- Apache web servers 95
- application programming interfaces 14
- applications
  - developing client 14
  - managing resources for 2
  - managing users and 60
  - monitoring 12
  - restricting CPU processes for 119
  - running client 6
  - running iHub processes and 5
- Attribute Name property 43
- authentication 6, 2, 14, 47, 49
- authorization 2, 14, 49
- autoarchive file purging processes 131
- auto-login scripts (pgpass) 134

## B

- backing up
  - configuration files 135
  - data files 131

- backing up (*continued*)
    - data folders 135
    - database schemas 134
    - Encyclopedia volumes 131
    - iHub System 130, 133
    - metadata 131
    - PostgreSQL databases 132
    - system databases 134
  - backup conflicts 132
  - backup procedures 5
  - backup utilities 131
  - BIRT Designer Professional 113
  - BIRT iHub. *See* iHub System
  - BIRT iHub Encyclopedia. *See* Encyclopedia volumes
  - BIRT Interactive Viewer licensing option 113
  - BIRT iServer. *See* iServer releases
  - BIRT licensing option 113
  - BIRT onDemand licensing option 112
  - BIRT Page Level Security licensing option 113
  - BIRT reports 15
    - See also* reports
  - BIRT service 30
  - BIRTImageCacheTimeout property 87
  - BIRTReportDesignCacheTimeout property 86
  - BIRTReportDesignCacheTotalNumberOfEntries property 87
  - briefing books 113
  - browsers. *See* web browsers
  - BufferPoolSize property 88
- ## C
- Cache Timeout property 53, 60
  - changing
    - alerts 45, 46
    - cluster node properties 66
    - cluster properties 27, 63, 64
    - configurations 62, 63
    - CPU bindings 124, 126
    - database encoding 15
    - database schemas 4
    - iHub service properties 29
    - license file names 117
    - licensing options 119
    - metadata databases 4
    - network cards 118
    - server configuration templates 30
    - system administrator 104
  - character encoding 15
  - client applications 5, 6, 14
    - See also* specific application
  - client/server models 9
  - cluster activity information 71, 80
  - Cluster Configuration page (System Console) 14, 26
  - cluster IDs 9, 66
  - cluster nodes
    - See also* clusters
    - accessing 64
    - adding 27, 14, 15, 17, 23
    - associating with machine IDs 115
    - changing properties for 66
    - configuring 14, 20
    - deleting 28, 66
    - failing 9
    - getting host names for 25
    - installing 18, 19
    - managing 27, 64, 65
    - networked environments and 9
    - running iHub instances and 9, 18
    - setting properties for 16, 28
    - starting or stopping 10, 27, 28, 65
    - viewing active requests for 78, 80
    - viewing activity information for 75
  - cluster schemas 4
    - See also* system schemas
  - clusters
    - accessing configurations for 9
    - accessing resources for 21
    - adding nodes to 28, 14, 15, 17, 23
    - administrative tasks for 5, 12, 26
    - changing node-key configurations for 118
    - changing properties for 27, 63, 64
    - configuring 12, 13, 14
    - creating 2, 13
    - deleting 69, 70
    - determining number of processors for 120, 124
    - exceeding CPU licenses for 125
    - getting IP addresses for 25
    - licensing options for 114, 115, 120
    - monitoring 3, 8

- removing nodes from 28
- retrieving volume metadata and 10
- running iHub processes and 5
- running iHub services and 6
- running jobs and 13
- sending e-mail over 126
- setting properties for 13, 14
- setting up failover procedures for 5
- setting up iHub environment for 18–26
- shutting down 6
- starting 6
- storing metadata for 4
- testing connections for 26
- turning off firewalls for 24
- updating configurations for 62
- viewing activity and resource usage for 71
- viewing diagnostic logs for 72, 73
- viewing error messages for 9
- viewing list of 66, 67
- viewing node key information for 118
- viewing service provider information for 47, 48
- viewing system information for 74, 76
- Clusters page (System Console) 12, 66
- collecting machine information 115, 116, 117
- command line utilities 14, 81, 117, 132
- Condition property 43
- Configuration Console 114
- configuration file
  - updating 15
- configuration files 61
  - See also* configurations
- configuration home directory 27
- configuration keys 95
- configuration parameters 9, 84
  - See also* parameters
- configuration template properties 84
- configuration templates 9, 30
- configurations
  - accessing information about 4
  - accessing resources and 21
  - adding JDBC drivers and 15
  - backing up 130, 135
  - changing database encoding and 15
  - CPU binding and 120, 122, 123
  - editing 62, 63
  - failover procedures and 5
  - file I/O operations and 8
  - licensing iHub and 115
  - restoring 132
  - running iHub clusters and 9, 14, 18, 17
  - troubleshooting 131
  - updating 62
- configuring
  - alerts 14, 40
  - Apache web servers 95
  - cluster nodes 14, 20
  - clusters 12, 13, 14
  - connections 94
  - e-mail settings 106
  - iHub System 10, 14
  - LDAP adapters 49
  - LDAP mappings 53
  - network sharing 19, 21
- ConnConfigFile property 94
- connection configuration files 94
- connection information 94
- connection pooling 15, 52
- connection profiles 94
- connection properties 94
- connections
  - accessing metadata databases and 9, 15
  - backing up multiple schemas and 134
  - changing database encoding and 15
  - configuring 94
  - externalizing 94
  - installing alternate databases and 15
  - running cluster nodes and 9
  - setting LDAP server 51
  - testing 26, 50, 38
- console applications
  - See also* specific application
- Context Path property 60
- ControlConsole utility 81, 82
- coprocessors. *See* CPUs
- copying
  - database files 131
- CPU binding 119–126
- CPU binding validation 121, 123, 125, 126
- CPU-based licenses
  - determining number of 120
  - exceeding number of 125, 126
  - obtaining 112, 113
  - viewing information about 125

## CPUs

- binding iHub processes to 120, 123
- binding to multiple core 121
- determining number of 121, 124
- hyperthreading and 122
- licensing options for. *See* CPU-based licenses
- restricting iHub processes for 119
- running encycsrvr processes and 124
- stopping iHub processes for 124
- testing connections to 26
- viewing information about 123
- viewing maximum number of 125
- viewing processor IDs for 120

## creating

- iHub clusters 2, 13
- passwords. *See* passwords
- shared files and folders 20
- system administrators 101

## custom events 14

## customer licenses 117

*See also* licensing options

## customizing

- cluster schemas 4
- database encoding 15
- metadata databases 4

# D

## dashboards 113

## Dashboards licensing option 113

## data

- backing up Encyclopedia and 130
- exporting 113
- preventing loss of 5
- recovering 5
- restoring 132
- retrieving 8, 10
- sharing across multiple machines 10
- storing volume 32

## data extraction processes 8

## data files 131

## data objects 8, 130

## data repositories. *See* Encyclopedia volumes

## data sources 8, 94

## Data Store Upgrader. *See* Encyclopedia Data Store Upgrader

## data stores

*See also* database schemas

## data types 15

## database drivers. *See* drivers

## database encoding 15

## Database name property 49, 37

## database objects 4

## Database port property 49, 37

## database schemas

### backing up 134

### customizing 4

### installing iHub System and 4

### preventing data loss for 5

### storing metadata and 4

## Database server property 49, 37

## database servers. *See* servers

## databases

*See also* specific database type

### accessing documentation for 5

### accessing shared data and 10

### backing up 134

### changing metadata 4

### configuring failover procedures for 5

### connecting to 15

### copying files in 131

### installing schemas for 4

### integrating with iHub 8

### managing 5

### monitoring data and 81

### retrieving data from 10

### searching 60

### setting properties for 49, 37, 38

### shutting down iHub clusters and 10

### specifying type 14, 37

### storing metadata and 4

### testing connections to 50, 38

## DatamartArchiveLimit property 85, 87

## default database encoding 15

## default directories. *See* directories

## Default Home Folder property 57

## default ports. *See* ports

## deleting

### alerts 46

### cluster nodes 28, 66

### clusters 69, 70

### Encyclopedia volumes 36

### product files. *See* uninstalling

- system administrators 104, 105
- deploying
  - iHub System 14, 12
- design files 8
  - See also* designs
- designs 8, 94, 113
- developing
  - client applications 14
- diagnosing system problems 13
- diagnostic fixes 6
- diagnostic log files 72, 73
- directories
  - cluster configurations and 9, 17
  - restoring volume data 135
  - shared resources and 21
  - storing Encyclopedia volumes and 35
  - storing volume data and 32
  - viewing configuration home 27
- directory paths 17, 95, 130
- directory servers. *See* LDAP servers
- disk storage location backups 131
- display formats. *See* output formats
- displaying
  - alerts 8, 40
  - clusters 66, 67
  - CPU information 123, 125
  - diagnostic logs 72, 73
  - host names 25
  - IP addresses 25
  - licensing information 114, 115
  - licensing options 61
  - node key information 118
  - processor IDs 120
  - reports 113
  - service provider information 47, 48
  - system information 74, 76, 100
- distinguished names 59
- distributed iHub System. *See* clusters
- document files 8
  - See also* documents
- documentation
  - administering iHub System and v, 14
  - Network File Systems 9
  - PostgreSQL client/server models 9
  - third-party RDBMS tools and 5
- documents 8, 113
  - See also* reports

- downloading
  - configuration files 62, 63
  - System Console category views 77
- drivers
  - accessing metadata databases and 15
  - adding JDBC 15
  - changing database encoding and 15
  - configuring iHub System and 10
  - encryption and 50, 37
  - running third-party databases and 8
- dual-core CPUs 121

## E

- elastic iHub clustering 9
- e-mail 41, 126
  - See also* notifications
- e-mail addresses 106
- Email Attribute property 56
- Email property 43
- e-mail settings 106
- EnableGenerationService property 85
- EnableIntegrationService property 88
- EnableViewingService property 86
- encoding 15
- encryption 50, 37
- encryption keys 33
- Encryption Method property 49, 37
- Encyclopedia Data Store Administrator. *See* Volume Data Store Administrator
- Encyclopedia Data Store Upgrader 86, 88
- Encyclopedia databases. *See* volume databases
- Encyclopedia metadata. *See* volume metadata
- Encyclopedia partitions. *See* volume partitions
- Encyclopedia processes. *See* encycsrvr processes
- Encyclopedia schemas. *See* volume schemas
- Encyclopedia volumes
  - accessing multiple 113
  - adding to clusters 14, 31, 32
  - backing up 131
  - controlling access to 13
  - CPU binding and 126
  - customizing metadata databases for 4
  - customizing schemas for 4

- Encyclopedia volumes (*continued*)
  - deleting 36
  - disabling and enabling 36
  - installing sample 16
  - managing 5
  - monitoring 3, 8
  - preventing data loss for 5
  - running iHub processes and 10
  - running third-party databases and 15
  - saving data objects for 8
  - saving metadata for 4
  - setting file-sharing permissions for 22
  - setting properties for 31, 32, 36
  - setting up failover procedures for 5
  - starting 6, 36
  - storing 35, 36
  - taking snapshots of 131
  - viewing activity information for 37
  - viewing information for 37
- encycsrvr processes 124, 125, 126
- Entity ID property 49
- error information 13, 9, 76
- error logging application 13
- error logging reports 15
- error messages 9
- escaped characters 59
- evaluation licenses 116
- events 14
- expired licenses 115, 116
- exporting
  - data 113
- external connection profiles 94
- external data sources 8, 60
- external management systems 14

## F

- Factory service 6, 8, 85
- Factory service processes 122, 126
- failover procedures 5
- features 112
- Fetch Limit property 52
- file I/O operations 8
- file names 114
- file paths 17, 95, 130
- file system backups 131
- file systems 8

- FileCacheTimeout property 86
- files
  - backing up data 131
  - copying 131
  - installing license 114, 118
  - restoring 135
  - sending licensing information and 118
  - setting up network sharing for 21
  - storing BIRT-specific 8
  - updating license 14, 117, 118
  - viewing information about 15
- filtering cluster list 66, 67, 68
- firewalls 19, 24
- folders
  - See also* directories
  - backing up data 135
  - restoring 135
  - setting up network sharing for 21
  - storing Encyclopedia volumes and 35, 36
  - storing volume data and 32
  - viewing shared configuration 27
- freeing memory 81

## G

- generating
  - backup files 133
  - machine IDs 116, 117
  - reports 8
- Generation service. *See* Factory service
- getJDBCMajorVersion function 15
- graphs. *See* charts
- Group Base DN property 56
- Group Description Attribute property 56
- Group Object property 56
- Group Search Filter property 57
- Group Volume Filter Attribute property 58

## H

- help topics
  - See also* online documentation
- Home Folder Attribute property 57
- host names 25, 16
- HTML documentation
  - See also* online documentation
- HTTP connections 9



hyperlinks. *See* URLs  
hyperthreading 122, 124

## I

I/O operations 8  
IBM DB2 databases. *See* DB2 databases  
IDAPI applications 7, 14  
iHub clusters. *See* clusters  
iHub Encyclopedia. *See* Encyclopedia volumes  
iHub images 9, 14  
iHub Integration Technology 14  
iHub licensing options 113, 119  
iHub Logging and Monitoring System 13  
iHub processes  
    adding cluster nodes and 9  
    binding to CPUs 120, 121  
    binding to Linux servers 123  
    binding to Windows systems 120, 121, 122  
    monitoring 8  
    restricting number of running 119  
    running 5, 6  
    starting 120  
    stopping 124  
    verifying CPU bindings for 121, 123, 124, 125, 126  
iHub run-time environment 5, 15  
iHub servers  
    binding to CPUs 119–126  
    exceeding CPU licenses for 125  
    getting machine ID for 115, 116, 117  
    monitoring 8  
    running as clusters. *See* clusters  
    sending requests to 7  
iHub service configuration parameters 84  
iHub services 6, 7, 9, 28  
    *See also* specific service  
iHub servlet container 5  
iHub System  
    accessing features 112  
    administering 13, 2, 100  
    backing up 130, 133  
    changing CPU bindings and 124, 126  
    changing machines for 118  
    checking bound processors for 123, 124–126  
    communication protocol for 7  
    configuring 10, 14  
    customizing volume schemas and 4  
    deploying 14, 12  
    extending functionality of 14  
    installing license files for 114, 118  
    integrating with RDBMS databases 8, 15  
    maintaining 6, 112  
    monitoring 2, 71  
    preventing data loss for 5  
    restoring 132, 135, 136  
    running on multiple machines 10  
    running processes for. *See* iHub processes  
    security mechanisms for 7  
    shutting down 6  
    specifying number of CPUs for 112  
    starting 6  
    storing metadata for 4  
    testing connections for 26, 50, 38  
    viewing diagnostic logs for 72  
    viewing licensing information for 114, 115  
images. *See* iHub images  
incomplete files 131  
indexed searches 4  
Information Console  
    providing run-time environment for 5  
    running 6  
Information Delivery API 7, 14  
installation  
    cluster nodes 18, 19  
    database schemas 4  
    Encyclopedia sample volume 16  
    JDBC drivers 15  
    license files 114, 118  
instance licenses 112  
Integration service 6, 8, 88  
Integration service processes 122  
Integration Technology. *See* iHub Integration Technology  
IntegrationService element 122  
Intel processors 122  
Interactive Viewer licensing option 113  
intra-cluster messaging 9  
IP addresses 25  
iServer System. *See* iHub System

## J

- Java developer guide 15
- Java Factory service. *See* Factory service
- Java Report Server Security Extension. *See* Report Server Security Extension
- Java Runtime Environment 15
- Java Virtual Machines. *See* JVM libraries
- JavaServerClientConnectionTimeout property 86
- JavaServerClientMaxConnections property 86
- JavaServerClientMinConnections property 86
- JDBC drivers 8, 15
- jdbcCompliant function 15
- job dispatcher 13
- job information 15
- job schedulers 13
- jobs 8, 13, 85
- JRE environment 15

## L

- language-specific licenses. *See* locales
- LDAP Adapter options (System Console) 49
- LDAP connection settings 51
- LDAP Mapping configurations 53
- LDAP Performance Settings 52
- LDAP Port property 51
- LDAP Server property 51
- LDAP servers 14, 49, 51
- license file names 114, 117
- license files
  - adding licenses and 119
  - installing 114, 118
  - obtaining 115, 116, 117
  - reapplying for 118
  - receiving e-mail about 117
  - running iHub clusters and 115
  - selecting 61
  - specifying location of 116
  - updating 14, 117, 118
- license keys 118
- License Management Server 81
- licensed CPUs 125
  - See also* CPU-based licenses
- licenses 112, 114
- licensing information 118, 125, 126

- licensing options 61, 112, 113, 119
- licensing support (Actuate) 117, 118
- Lightweight Directory Access Protocol. *See* LDAP
- Linux servers
  - binding iHub processes to 123
  - running logging applications on 13
  - verifying core settings for 124
  - verifying CPU bindings for 123
- LM Server 81
- LMS. *See* Logging and Monitoring System
- load balancing (clusters) 9, 13, 95
- log files 72, 125
- Logging and Monitoring System 13
- logging applications 13
- logging in to
  - System Console 3, 4
- logical processors 121, 122, 124
- losing data 5

## M

- machine capacity 113
- machine IDs 115, 116
- machine information 116, 117
- mail servers 106
- maintenance 6, 112
- maintenance customers 117
- Management Console. *See* System Console
- manuals. *See* documentation
- MaxBIRTDDataResultsetBufferSize property 85
- MaxConcurrentRequests property 86
- MaxDatamartArchiveSize property 85, 87
- Maximum Pool Size property 52
- MaxPersistentArchiveSize property 86, 87
- MaxSyncJobRuntime property 85
- MaxSyncRequestTime property 85
- MaxTransientArchiveSize property 87
- Member ID Type property 57
- Member List Attribute property 57
- memory 81, 113
- Message property 43
- metadata
  - See also* system metadata; volume metadata
  - accessing 9

- autoarchiving and 132
- backing up 131
- defined 130
- restoring 132
- storing 4
- metadata databases 4, 5, 10, 14
  - See also* RDBMS databases
- metadata tables 4
- Metrics Management licensing option 113
- Microsoft SQL Server databases. *See* SQL Server databases
- Microsoft Windows systems. *See* Windows systems
- migration
  - See also* upgrades
- monitoring data 81
- Monitoring page (System Console) 8, 43
- multiple-core CPU binding 121
- Multi-Tenant licensing option 113

## N

- native system tools 5
- network administrators. *See* administrators
- network cards 116, 118
- Network File Systems 8
- network sharing configurations 19, 21
- networked environments
  - obtaining licenses for 113, 116, 118
  - running cluster nodes and 9
  - sending data over 8
- NFS (Network File Systems) 8
- node keys 115
- node-key configurations 118
- node-key licensing 115, 116, 118
- nodes. *See* cluster nodes
- notifications 9, 106
  - See also* e-mail

## O

- ODA data sources 95
- onDemand licensing option 112
- OnDemandServerViewMessageTimeout property 86
- online documentation
  - administering iHub System and v, 14
- online help. *See* online documentation

- operating systems 16
- options (licensing) 112, 113, 119
- Oracle databases
  - accessing documentation for 5
  - setting properties for 39
- output 8
- output formats 8, 77

## P

- packages (licensing option) 112
- page-level security 113
- PagePoolSize property 88
- parameters
  - configuring clusters and 9
  - setting iHub service 84
- partitions (volume) 8
- Password property 50, 37, 51
- passwords
  - creating 14, 102
- patches 6
- paths 17, 95, 130
- PDF documentation
  - See also* online documentation
- performance
  - CPU binding and 121
  - file I/O operations and 8
  - iHub architecture models and 6
  - iHub clusters and 10
- performance statistics 113
- permanent licenses 116
- permissions. *See* privileges
- PersistentArchiveFileCacheTimeout property 87
- PersistentArchiveLimit property 85, 87
- pg\_dump command line options 133
- pg\_dump utility 132, 133, 134
- pg\_restore command line options 136
- pg\_restore utility 132, 135, 136
- pgAdmin database administration tool 132, 136
- pgpass files 134
- ping command 26
- PMD. *See* Process Management Daemon
- pool. *See* connection pooling
- Port Number property 60
- ports 49, 37, 60

- PostgreSQL administration utilities 132, 136
- PostgreSQL command line utilities 132
- PostgreSQL databases
  - accessing documentation for 5
  - accessing metadata in 9
  - backing up 132
  - file I/O operations and 8
  - installing iHub and 4
  - running volume backup and restore operations for 132
  - setting properties for 49, 37, 38
  - starting 8
- Preferred Pool Size property 52
- Prefix property 53
- printers 10
- privileges
  - accessing Encyclopedia and 13
  - accessing shared resources and 22
- Process Management Daemon
  - CPU binding and 120, 121, 123
  - distributing SOAP requests and 7
  - running iHub clusters and 9
  - running iHub processes and 5
  - starting encycsrvr processes and 124, 126
  - viewing diagnostic logs for 72
- processes. *See* iHub processes
- processor affinity 121, 122, 123
- processor IDs 120, 122, 123
- ProcessorAffinity element 122
- processors. *See* CPUs
- program files. *See* product files
- properties
  - adding Encyclopedia volumes and 31, 32, 36
  - adding service providers and 49
  - changing iHub cluster 27, 63, 64
  - changing iHub service 29
  - configuring RSSE SOAP Service 60
  - connecting to data sources and 94
  - creating system administrators and 101
  - file sharing 21
  - sending e-mail and 106
  - setting iHub cluster 13, 14
  - setting iHub cluster node 16, 28, 66
  - setting LDAP server 51, 52, 53
  - setting metadata database 49, 37, 38
- protecting data. *See* security

- proxy servers 95
- publishing reports 113
- purging processes 131

## Q

- queries 8, 9

## R

- RDBMS command line utilities 132
- RDBMS databases 4, 131
  - See also* databases; third-party databases
- RDBMS documentation 5
- RDBMS tools 5, 131
- rebinding encycsrvr processes 125
- recovery operations 5
- Recursive Groups property 52
- relational database management systems 4
  - See also* RDBMS databases
- relational databases. *See* databases
- removing. *See* deleting
- renaming license files 117
- report design files 8
  - See also* report designs
- report designs 8, 94, 113
- report document files 8
  - See also* report documents
- report documents 8, 113
  - See also* reports
- report files
  - See also* specific report file type
  - backing up 131
  - restoring 135
  - setting up network sharing for 21
  - storing 8
  - viewing information about 15
- Report Server Security Extension 14
- Report Server Security Extension services 7
- Report Studio licensing option 114
- report viewers 113
- reporting servers. *See* iHub servers
- reporting services. *See* iHub services
- reporting system. *See* iHub System
- ReportingService element 122
- reports
  - displaying 113
  - generating 8

- installing sample 15
- publishing 113
- repositories. *See* Encyclopedia volumes
- requests 5, 7, 78
  - See also* SOAP-based messages
- resource groups 76
- resource usage information 71
- resources
  - accessing 21
  - creating cluster nodes and 9
  - defining as work unit 112
  - evaluating usage 13
  - managing 2
  - monitoring 2, 71
  - retrieving volume metadata and 10
- restoring iHub System 132, 135, 136
- restricting iHub processes 119
- RSSE applications 7, 14, 60
- RSSE SOAP service settings 60
- running
  - acmachineid utility 117
  - console client applications 6
  - ControlConsole utility 82
  - encycsrvr processes 124, 125, 126
  - iHub processes 5, 6, 119
  - iHub services 6
  - jobs 8, 13, 85
  - pg\_dump utility 134
  - pg\_restore utility 136
  - PostgreSQL databases 8
  - queries 8, 9
  - report designs 113
  - third-party databases 8

## S

- SaaS licensing option 112
- SAML Identity Provider information 14, 47, 48
- sample reports 15
- scalability 2
- scheduling jobs 13
- schemas
  - backing up 134
  - customizing 4
  - installing iHub System and 4
  - preventing data loss for 5

- storing metadata and 4
- scripts
  - extending iHub functionality and 14
  - iHub setup. *See* iHub distribution setup
    - script
    - running RDBMS backup tools 131
- Search setting property 60
- searching
  - Active Directory servers and 59
  - clusters 66, 67, 68
  - metadata database 60
- security 6, 20, 2
- security settings 14
- sending licensing information 118, 125, 126
- server configuration template properties 84
- server configuration templates 9, 30
- Server Name property 60
- Server URL property 49
- servers
  - See also* iHub servers
  - configuring Apache web 95
  - configuring mail 106
  - exceeding CPU licenses for 125
  - installing stand-alone. *See* stand-alone servers
  - licensing information and 118
  - running as clusters. *See* clusters
  - shutting down iHub clusters and 6
  - service provider information 47, 48
  - service providers 48
  - services (reporting). *See* iHub services
  - servlet container (iHub) 5
  - shared licenses 114, 116
  - shared resources 21
  - shutting down. *See* stopping
  - simple object access protocol. *See* SOAP-based messages
  - single sign-on authentication 14, 47
  - single-point node failures 9
  - snapshots 131
  - SOAP APIs 39
  - SOAP-based messages 5, 7, 15
    - See also* requests
  - Software as a Service licensing option 112
  - SSL encryption 50, 37
  - SSL property 51
  - stand-alone servers 118

- starred clusters 68
- StartArguments property 88
- starting
  - Encyclopedia volumes 6, 36
  - iHub cluster nodes 10, 27, 28, 65
  - iHub clusters 6
  - iHub processes 120
  - iHub services 29
  - iHub System 6
  - pg\_dump utility 134
  - pg\_restore utility 136
  - PostgreSQL databases 8
- stopping
  - cluster nodes 10, 27, 28, 65
  - Encyclopedia volumes 36
  - iHub processes 124
  - iHub services 29
  - iHub System 6
- storage locations (Encyclopedia) 35, 36
- subscription licenses 112
- Suffix property 53
- Sun ONE LDAP servers. *See* LDAP servers
- SynchReportingWeight property 86
- SyncJobQueueSize property 85
- SyncJobQueueWait property 85
- system. *See* iHub System
- system administrators 2, 101
  - See also* administrators
- System Console
  - adding Encyclopedia and 32, 35, 36
  - administering iHub and 14, 2, 61, 71, 100
  - documentation for 14
  - filtering cluster list in 66, 67, 68
  - licensing and 118, 125
  - logging in to 3, 4
  - managing iHub clusters and 12, 14, 16
  - managing users and 49
  - monitoring tasks and 8, 40
  - providing run-time environment for 5
  - running 6
  - sending e-mail and 126
  - starting iHub and 6
  - viewing data items in 3
  - viewing system information and 74, 76, 100
- system data store. *See* system schemas

- System Data Store Administrator. *See* Cluster Data Store Administrator
- system databases
  - backing up 134
  - changing 4
  - connecting to 15
  - installing 4
  - storing metadata and 4
- system failover procedures 5
- system metadata
  - See also* metadata
  - maintaining 10
  - storing 4
- system metadata schemas. *See* system schemas
- system schemas
  - See also* database schemas
  - customizing 4
  - storing configuration metadata and 4
- system tools 5

## T

- tables (metadata databases) 4
- templates. *See* server templates
- temporary documents 8, 85
- temporary licenses 114, 116
- Test Connection property 50, 38
- testing
  - connections 26, 50, 38
  - JDBC drivers 15
- text files (licenses) 117
- third-party command line utilities 132
- third-party database schemas. *See* database schemas
- third-party databases 4, 15
  - See also* databases
- third-party RDBMS documentation 5
- third-party RDBMS tools 5, 131
- Threshold property 43
- Timeout property 52
- transient reports. *See* temporary documents
- TransientArchiveFileCacheTimeout property 87
- TransientArchiveLimit property 87
- TransientReportCacheSize property 85
- TransientReportTimeOut property 85

TransientStoreMaxCacheEntries property 85  
troubleshooting 13, 131  
types. *See* data types

## U

Uniform Resource Locators. *See* URLs

Unlimited User CPU License 113

updating

- configuration file 15
- configurations 62
- licenses 14, 61, 117, 118

upgrades

- Encyclopedia volumes and 4
- licensing options and 116

URLs

- Network File Systems documentation 9
- PostgreSQL administration utilities 136
- PostgreSQL client/server models 9
- RDBMS documentation 5
- service providers 49
- System Console 4

usage information 13

usage logging application 13

usage reports 15

user authentication 6, 2, 14, 47, 49

user authorization 2, 14, 49

User Base DN property 55

User Description Attribute property 55

User DN property 51

User Full Name Attribute property 55

user groups 13

User Login Name Attribute property 55

User Object property 55

User Search Filter property 56

User Volume Filter Attribute property 58

Username property 50, 37

users

- accessing Encyclopedia and 13
- adding as system administrator 101
- managing 14, 15, 49

## V

validating CPU binding 121, 123, 125, 126

View monitoring service 73

View service 6, 8, 72, 86

View service processes 122, 126

viewers 113

viewing

- alerts 8, 40
- clusters 66, 67
- CPU information 123, 125
- diagnostic logs 72, 73
- host names 25
- IP addresses 25
- licensing information 114, 115
- licensing options 61
- node key information 118
- processor IDs 120
- reports 113
- service provider information 47, 48
- system information 74, 76, 100

Viewing service. *See* View service

ViewingService element 122

ViewingWeight property 88

Visualization Platform 48

- administering iHub System and 14
- documentation for 14, 15

volume administrators. *See* administrators

volume data 10, 31, 130

*See also* data

volume data directories 135

volume databases

*See also* metadata databases

- changing 4
- connecting to 15
- file I/O operations and 8
- installing 4
- retrieving data from 10
- viewing incomplete files in 131

volume failover procedures 5

volume metadata

*See also* metadata

- backing up 131
- maintaining 10
- restoring 132
- storing 4

volume metadata schemas. *See* volume schemas

volume partitions 8

volume schemas

*See also* database schemas

- backing up 134
- customizing 4

- volume schemas (*continued*)
  - storing configuration metadata and 4
- Volumes
  - viewing diagnostic logs for 72
- volumes. *See* Encyclopedia volumes

## W

- web administrators. *See* administrators
- web browsers
  - accessing structured content and 113
  - accessing System Console and 4
- web pages 113
- web servers 95
- web service applications 60
- web services 29
- wildcard characters 67
- Windows systems
  - backing up volume data folders on 135
  - binding iHub processes to 120, 121, 122

- network cards and 116, 118
  - obtaining IP addresses for 25
  - restoring iHub data folders on 135
  - running logging applications on 13
  - setting up firewalls for 19
  - setting up network sharing for 20, 21
  - turning off firewalls for 24
- work unit licenses 112, 113
- work units 75, 112
- WSDL utilities 39

## X

- XML files 94, 118
  - See also* configuration files

## Z

- ZIP files 117