

# ActuateOne™

One Design  
One Server  
One User Experience

**Installing BIRT iServer  
for Linux and UNIX**

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 - 2012 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)

[www.birt-exchange.com](http://www.birt-exchange.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 Data Analyzer, BIRT Performance Analytics, Collaborative Reporting Architecture, e.Analysis, e.Report, e.Reporting, e.Spreadsheet, Encyclopedia, Interactive Viewing, OnPerformance, Performancesoft, Performancesoft Track, Performancesoft Views, Report Encyclopedia, Reportlet, The people behind BIRT, 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. Apache Software Foundation ([www.apache.org](http://www.apache.org)): Axis, Axis2, Batik, Batik SVG library, Commons Command Line Interface (CLI), Commons Codec, Derby, Hive driver for Hadoop, Shindig, Struts, Tomcat, Xalan, Xerces, Xerces2 Java Parser, and Xerces-C++ XML Parser. 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. Codejock Software: Xtreme Toolkit Pro. 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), and Eclipse Web Tools Platform (WTP), licensed under the Eclipse Public License (EPL). Bits Per Second, Ltd. and Graphics Server Technologies, L.P.: Graphics Server. Gargoyles Software Inc.: HtmlUnit, licensed under Apache License Version 2.0. GNU Project: GNU Regular Expression, licensed under the GNU Lesser General Public License (LGPLv3). HighSlide: HighCharts. IDAutomation.com, Inc.: IDAutomation. Jason Hsueth and Kenton Varda ([code.google.com](http://code.google.com)): Protocol Buffer. IDRolutions Ltd.: JBIG2, licensed under the BSD license. ImageMagick Studio LLC.: ImageMagick. InfoSoft Global (P) Ltd.: FusionCharts, FusionMaps, FusionWidgets, PowerCharts. Matt Inger ([sourceforge.net](http://sourceforge.net)): Ant-Contrib, licensed under Apache License Version 2.0. Matt Ingenthron, Eric D. Lambert, and Dustin Sallings ([code.google.com](http://code.google.com)): Spymemcached, licensed under the MIT OSI License. International Components for Unicode (ICU): ICU library. jQuery: jQuery, licensed under the MIT License. Yuri Kanivets ([code.google.com](http://code.google.com)): Android Wheel gadget, licensed under the Apache Public License (APL). KL Group, Inc.: XRT Graph, licensed under XRT for Motif Binary License Agreement. LEAD Technologies, Inc.: LEADTOOLS. Bruno Lowagie and Paulo Soares: iText, licensed under the Mozilla Public License (MPL). Microsoft Corporation (Microsoft Developer Network): CompoundDocument Library. Mozilla: Mozilla XML Parser, licensed under the Mozilla Public License (MPL). MySQL Americas, Inc.: MySQL Connector. Netscape Communications Corporation, Inc.: Rhino, licensed under the Netscape Public License (NPL). OOPS Consultancy: XMLTask, licensed under the Apache License, Version 2.0. Oracle Corporation: Berkeley DB, Java Advanced Imaging, JAXB, JDK, Jstl. PostgreSQL Global Development Group: pgAdmin, PostgreSQL, PostgreSQL JDBC driver. Progress Software Corporation: DataDirect Connect XE for JDBC Salesforce, DataDirect JDBC, DataDirect ODBC. Rogue Wave Software, Inc.: Rogue Wave Library SourcePro Core, tools.h++. Sam Stephenson ([prototype.conio.net](http://prototype.conio.net)): prototype.js, licensed under the MIT license. Sencha Inc.: Ext JS. ThimbleWare, Inc.: JMemcached, licensed under the Apache Public License (APL). World Wide Web Consortium (W3C)(MIT, ERCIM, Keio): Flute, JTIty, Simple API for CSS. XFree86 Project, Inc.: ([www.xfree86.org](http://www.xfree86.org)): xvfb. ZXing authors ([code.google.com](http://code.google.com)): ZXing, licensed under the Apache Public License (APL).

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

Document No. 120201-2-430343 October 2, 2012

# Contents

<b>Introduction</b> .....	<b>vii</b>
Understanding ActuateOne .....	vii
About the BIRT iServer documentation .....	vii
Obtaining documentation .....	xi
Using PDF documentation .....	xi
Obtaining late-breaking information and documentation updates .....	xi
About obtaining technical support .....	xi
About supported and obsolete products .....	xii
Typographical conventions .....	xii
Syntax conventions .....	xii
About <i>Installing BIRT iServer for Linux and UNIX</i> .....	xiii

## Part 1

### Architecture

#### Chapter 1

<b>Understanding Actuate BIRT iServer architecture</b> .....	<b>3</b>
Understanding BIRT iServer architecture .....	4
Using a third-party RDBMS with an Encyclopedia volume .....	4
Customizing Encyclopedia volume databases .....	5
Installing and configuring iServer System .....	6
Managing the backup, recovery, and failover capabilities of the Encyclopedia volume database and data files .....	6
Managing an iServer cluster .....	8
Understanding the iServer System process model .....	8
Understanding process flow in a stand-alone iServer .....	9
Understanding process flow in an iServer cluster .....	11
Administering iServer System .....	15
About Migration and Administration Tools .....	16
Using JDBC to connect to an Encyclopedia volume database .....	18
API Compatibility .....	18
About international character sets .....	19
Administrative reports .....	19
Supported operating systems .....	19

## Part 2

# Installing

## Chapter 2

### **Installing BIRT iServer . . . . . 23**

Preparing to install BIRT iServer . . . . .	24
Creating a dedicated user account for installing and running iServer . . . . .	24
Backing up iServer system and Encyclopedia volume metadata . . . . .	24
About X frame buffer . . . . .	25
Installing X frame buffer . . . . .	26
About libstdc++ . . . . .	26
About run levels . . . . .	26
About Openmotif . . . . .	26
Performing a new installation . . . . .	26
Installing BIRT iServer Release 11 . . . . .	27
Accessing Information, Management, and Configuration Consoles . . . . .	40
Stopping and starting iServer and PostgreSQL processes . . . . .	42
Understanding the iServer installation environment . . . . .	43
About migrating an earlier iServer release to Actuate 11 . . . . .	43
About running different releases on the same machine . . . . .	44
About performance and disk space issues . . . . .	44
About upgrading an iServer with resource groups . . . . .	44
About the Java Software Development Kit . . . . .	44
Accessing JAR files for report generation . . . . .	45
Gathering LDAP information . . . . .	45
Following best practices . . . . .	46
Using a test environment . . . . .	46
Setting up a production staging area . . . . .	46
Setting up a production environment . . . . .	47

## Chapter 3

### **Installing BIRT iServer using an alternative database . . . . . 49**

Preparing to install BIRT iServer using an alternative database . . . . .	50
Creating a dedicated user account for installing and running BIRT iServer . . . . .	50
Creating the system and Encyclopedia volume schemas and iserver user in an alternative database . . . . .	50
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing PostgreSQL database . . . . .	51
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing Oracle database . . . . .	53
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing SQL Server database . . . . .	55

Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing DB2 database .....	57
Adding support for the digit wildcard character in iServer when the metadata database is DB2 .....	59
Filtering for a user name returns an empty result if name contains certain special characters .....	61
Backing up iServer system and Encyclopedia volume metadata .....	62
Installing an Encyclopedia volume that uses an alternative database .....	63

## Chapter 4

<b>Upgrading BIRT iServer .....</b>	<b>79</b>
Upgrading BIRT iServer .....	80
Preparing to upgrade BIRT iServer .....	80
Creating a dedicated user account for installing and running BIRT iServer .....	80
Backing up iServer system and Encyclopedia volume metadata .....	81
Performing an automatic in-place upgrade for an earlier major release .....	82
Running the in-place upgrade on an earlier major release .....	83
Performing an automatic in-place upgrade for an earlier minor release .....	93
Running the in-place upgrade on an earlier minor release .....	93
Performing a manual upgrade of an Encyclopedia volume schema for an earlier minor release .....	100
Performing a manual side-by-side upgrade .....	104
Installing a new Release 11 iServer .....	105
Performing a manual side-by-side migration .....	118
Working with iServer utilities .....	129
Working with Squirrel Data Exporter .....	129
Specifying Squirrel Data Exporter properties .....	129
Using the generated bulk-load script files .....	131
Working with Encyclopedia Data Store Administrator .....	133
Specifying Encyclopedia Data Store Administrator properties .....	133
Performing operations using Encyclopedia Data Store Administrator utility .....	140
Working with System Data Store Administrator .....	144
Specifying System Data Store Administrator properties .....	145
Performing operations using System Data Store Administrator utility .....	147
Working Encyclopedia Data Store Upgrader .....	147
Specifying Encyclopedia Data Store Upgrader properties .....	147

## Chapter 5

<b>Installing a BIRT iServer cluster .....</b>	<b>151</b>
Installing a BIRT iServer cluster node .....	152
Preparing to install an iServer cluster .....	153
Creating an administrative user account .....	153

Installing X-frame buffer .....	154
Performing a cluster installation using the wizard .....	155
Adding a node to a cluster .....	165
Finding the BIRT iServer home directory .....	170
About the Java Development Kit .....	171

## Chapter 6

<b>Installing BIRT iServer in a cloud .....</b>	<b>173</b>
Understanding a BIRT iServer cloud deployment .....	174
Deploying BIRT iServer in a cloud environment .....	174
Deploying an iServer distribution package .....	175
Running the setup script .....	176
Setting up iServer to use the out-of-the-box (OOTB) PostgreSQL database .....	178
Accessing Information, Management, and Configuration Consoles .....	179
Stopping and starting iServer and PostgreSQL processes .....	182
Setting up iServer to use an alternative database .....	182
Setting up iServer only .....	188
Setting up iServer only using the cloud deployment package .....	188
Configuring an external database .....	190
Using Configuration Console to configure the database connection .....	191
Using Configuration Console to configure the data partition and Encyclopedia volume .....	196
Setting up iServer to join an existing cluster .....	201
Understanding the cloud deployment environment .....	204
Specifying AC_SERVER_HOME .....	204
Specifying AC_DATA_HOME .....	205

## Chapter 7

<b>Installing Information Console .....</b>	<b>207</b>
Before you begin .....	208
About installing from an FTP download .....	208
About performing a full installation .....	208
Installing Information Console on Linux and UNIX .....	208
Using the script to install .....	209
Using the WAR file to install .....	214
General deployment tasks .....	214
Preparing the server .....	214
Preparing the WAR file .....	215
About clusters of servers .....	216
Avoiding cache conflicts after installing .....	216
Testing the installation .....	217

## Chapter 8

### **Installing iServer Integration Technology and Documentation . . . . . 219**

Installing BIRT iServer Integration Technology . . . . . 220

Installing the localization and documentation files . . . . . 222

## Part 3

### **Licensing**

## Chapter 9

### **Licensing BIRT iServer . . . . . 229**

Understanding licensing types . . . . . 230

Understanding licensing options . . . . . 231

Installing Actuate BIRT iServer System license files . . . . . 233

    About the license file . . . . . 235

    Collecting machine information for a license . . . . . 235

    About modifying a license . . . . . 238

Understanding CPU binding . . . . . 239

    Binding BIRT iServer to processors on a Sun Solaris machine . . . . . 240

        Binding to specific CPUs . . . . . 240

        Binding to multiple CPUs . . . . . 241

        Binding to multiple-core CPUs . . . . . 241

    Binding BIRT iServer to processors on an HP-UX 11i machine . . . . . 242

    Checking BIRT iServer bound processors . . . . . 242

        Determining the number of processors an iServer System uses . . . . . 242

        Understanding CPU binding validation while iServer is running . . . . . 243

        Understanding CPU binding validation when an Encyclopedia volume  
        comes online . . . . . 244

        Understanding CPU binding validation when running iServer processes . . . . . 244

    Configuring e-mail for CPU license problems . . . . . 244

## Part 4

### **Backing Up**

## Chapter 10

### **Backing up an Encyclopedia volume . . . . . 247**

Performing an Encyclopedia volume backup . . . . . 248

    Managing the backup and recovery of Encyclopedia volume metadata and data files . . . 248

    Using RDBMS and file system backup utilities . . . . . 249

    Avoiding conflict with the autoarchive file purging process . . . . . 250

Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database . . . . . 251

    Backing up an Encyclopedia volume using pg\_dump . . . . . 252

Restoring an Encyclopedia volume using pg_restore .....	256
Backing up and restoring an Encyclopedia volume that uses an Oracle database .....	259
Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp) .....	260
Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp) .....	262
Backing up and restoring an Encyclopedia volume that uses a DB2 database .....	264
Backing up an Encyclopedia volume .....	265
Restoring an Encyclopedia volume .....	267
<b>Index .....</b>	<b>271</b>



---

## Understanding ActuateOne

ActuateOne™ includes Release 11 of Actuate® Corporation's value-added products for the Eclipse BIRT open source project. ActuateOne institutes a paradigm shift in Business Intelligence technology from individualized tools to a suite of integrated capabilities within a single environment. ActuateOne is one designer, one server, one integrated environment providing a single unified user experience. A common architecture is precisely what today's information-rich global environment requires for development and deployment. This unified Actuate technology continues to enable information management and delivery while supporting advanced security, massive scalability, flexibility through programming, and reuse. ActuateOne realizes our corporate vision of a single user experience by providing extended new analytics capabilities reaching a broader spectrum of users. The new dashboard functionality supports building gadgets to enhance the visual presentation of information. Export to Excel® and other formats integrates Actuate output with other tools on the end-user desktop. Actuate's cloud-ready server supports elastic clustering for dynamic provision of uninterrupted efficient service.

Information, live demos, and endorsements about this release are available from [birt-exchange.com](http://birt-exchange.com) and [actuate.com](http://actuate.com). The Actuate site also makes "The Forrester Wave™: Open Source Business Intelligence (BI), Q3 2010" report freely available. The report recognizes Actuate and its value-added offerings for BIRT as a leader in Open Source Business Intelligence.



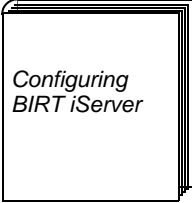
---

## About the BIRT iServer documentation

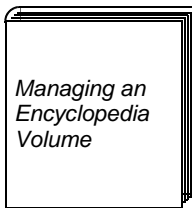
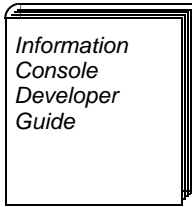
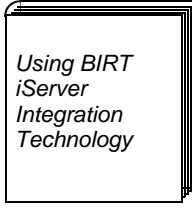
The printed and online documentation includes the materials described in Table I-1. You can obtain HTML and PDF files from the BIRT Exchange or Actuate Customer Support site.

Documentation updates are created in response to customer requirements and are available at both sites.

**Table I-1** BIRT iServer documentation

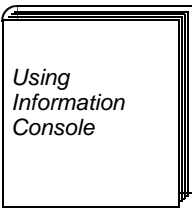
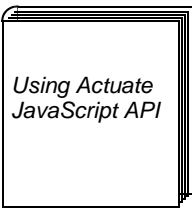
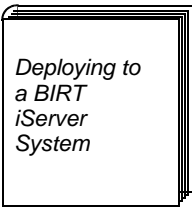


For information about this topic	See the following resource
Installing BIRT iServer for Linux and UNIX	 <p><i>Installing BIRT iServer for Linux and UNIX</i></p>
Installing BIRT iServer for Windows	 <p><i>Installing BIRT iServer for Windows</i></p>
Late-breaking information and documentation updates	Release notes and updated localization, HTML help, and PDF files posted on <a href="http://birt-exchange.com">birt-exchange.com</a> and Actuate <a href="#">Support</a>
Configuring BIRT iServer Use Configuration Console to: <ul style="list-style-type: none"><li>■ Add additional Encyclopedia volumes</li><li>■ Configure clusters of iServers</li><li>■ Tune iServer services and processes</li><li>■ Configure e-mail notification</li><li>■ Review and update license options</li><li>■ Open ports for iServer use</li><li>■ Manage iServer printers and resources</li><li>■ Configure diagnostic logging</li></ul>	 <p><i>Configuring BIRT iServer</i></p>

**Table I-1** BIRT iServer documentation (continued)

<b>For information about this topic</b>	<b>See the following resource</b>
<p>Managing an Encyclopedia Volume</p> <p>Use Management Console to:</p> <ul style="list-style-type: none"><li>■ Set up user accounts</li><li>■ Set up channels and notification groups</li><li>■ Assign security roles</li><li>■ Manage files and folders</li><li>■ Schedule, run, and manage reports</li><li>■ Back up the Encyclopedia volume</li><li>■ Use Actuate Open Security</li></ul>	 <p><i>Managing an Encyclopedia Volume</i></p>
<p>Information Console Developer Guide</p> <ul style="list-style-type: none"><li>■ Overview of Information Console concepts and web applications</li><li>■ Using, customizing, and configuring the Deployment Kit</li><li>■ Using code components for JSPs, URL parameters, JavaScript files, Java servlets, Java Beans, and security facilities</li></ul>	 <p><i>Information Console Developer Guide</i></p>
<p>Using BIRT iServer Integration Technology</p> <ul style="list-style-type: none"><li>■ Overview of Actuate web services and SOAP messaging</li><li>■ Managing an Encyclopedia volume</li><li>■ Developing API applications using Java or .NET</li><li>■ Using Java Report Server Security Extension (RSSE) APIs</li><li>■ Using logging, performance monitoring, and archiving features</li><li>■ Customizing the Actuate software installation process</li><li>■ Actuate Information Delivery API operations and data types reference</li></ul>	 <p><i>Using BIRT iServer Integration Technology</i></p>

(continues)

**Table I-1** BIRT iServer documentation (continued)

For information about this topic	See the following resource
<p>Using Information Console</p> <ul style="list-style-type: none"><li>■ Overview of Information Console concepts and online reporting</li><li>■ Accessing and managing files and folders; running designs</li></ul>	 <p><i>Using Information Console</i></p>
<p>Using Actuate JavaScript API</p> <ul style="list-style-type: none"><li>■ Overview of programming with Actuate JavaScript</li><li>■ Creating custom pages using Actuate JavaScript</li><li>■ Reference for BIRT JavaScript classes and methods</li></ul>	 <p><i>Using Actuate JavaScript API</i></p>
<p>Deploying to a BIRT iServer System</p> <p>Describes how to deploy designs and information objects to iServer</p>	 <p><i>Deploying to a BIRT iServer System</i></p>
<p>Actuate Glossary</p> <p>Definitions of product terminology</p>	 <p><i>Actuate Glossary</i></p>
<p>Adobe Acrobat Catalog</p> <p>A utility that can search all the documents in the Actuate manuals directory</p>	 <p><i>Adobe Acrobat Catalog</i></p>

## Obtaining documentation

Actuate provides technical documentation in PDF, HTML, and print formats. You can download PDF or view HTML versions of the documentation from [birt-exchange.com](http://birt-exchange.com). If you purchase the product, you can also download documentation using ftp as instructed in the e-mail from Actuate Distribution. When you install the files using the Online Documentation and Localization Resource Files program, if you accept the default location, the program loads the PDF in the `/home/Actuate/AcServer/Manuals` directory.

## Using PDF documentation

In each PDF version of a book, the table of contents and the index page numbers contain links to the corresponding topics in the text. In the table of contents, you access the link by positioning the pointer over the topic. In the index, you access the link by positioning the pointer over the page number.

The `/home/Actuate/AcServer/manuals` directory contains a file, `master-index.pdx`, which is an Adobe Acrobat Catalog utility that can search all the documents in the Actuate Manuals directory. This tool provides a convenient way to find information on a particular topic in Actuate documentation.

## Obtaining late-breaking information and documentation updates

The release notes contain late-breaking news about Actuate products and features. The release notes are available on the Actuate Support site at the following URL:

<http://support.actuate.com/documentation/releasenotes>

Updates to documentation in PDF form are available at the following URL:

<http://support.actuate.com/documentation>

If you are a new user, you must first register on the site and log in to view the release notes. [Birt-exchange.com](http://birt-exchange.com) and [actuate.com](http://actuate.com) also provide product update information.

---

## About obtaining technical support

You can contact Customer Support by e-mail or telephone. For contact information, go to the following URL:

<http://www.actuate.com/services/support/contact-support.asp>

---

## About supported and obsolete products

The Actuate Support Lifecycle Policy and Supported Products Matrix are available on the Actuate Support web site at the following URL:

<http://support.actuate.com/documentation/spm>

---

## Typographical conventions

Table I-2 describes the typographical conventions in this document.

**Table I-2** Typographical conventions

Item	Convention	Example
Code examples	Monospace	<code>Dim Text1 As String</code>
File names	Initial capital letter, except where file names are case-sensitive	<code>Detail.roi</code>
Key combination	A + sign between keys means to press both keys at the same time	<code>Ctrl+Shift</code>
Menu items	Capitalized, no bold	<code>File</code>
Submenu items	Separated from the main menu item with a small arrow	<code>File→New</code>
User input or user response	Monospace	<code>M*16*</code>
User input in XML and Java code	Monospace italics	<code>chkjava.exe</code> <code>cab_name.cab</code>

---

## Syntax conventions

Table I-3 describes the symbols used to present syntax.

**Table I-3** Syntax conventions

Symbol	Description	Example
[ ]	Optional item	[Alias<alias name>]
	Array subscript	matrix[ ]

**Table I-3** Syntax conventions

Symbol	Description	Example
{ }	Groups two or more mutually exclusive options or arguments when used with a pipe	{While   Until}
	Defines array contents	{0, 1, 2, 3}
	Delimiter of code block	public ACJDesigner( ) { }
	Separates mutually exclusive options or arguments in a group	Exit {Do   For   Function   Sub}
	Java OR operator	int length   4
< >	Argument you must supply	<expression to format>
	Delimiter in XML	<xsd:sequence>

---

## About *Installing BIRT iServer for Linux and UNIX*

*Installing BIRT iServer for Linux and UNIX* includes the following chapters:

- *Introduction*. Provides an overview of this guide, BIRT iServer documentation, and the typographical conventions in this book.
- *Part 1. Architecture*. Describes BIRT iServer architecture.
- *Chapter 1. Understanding Actuate BIRT iServer architecture*. Describes BIRT iServer architecture, the iServer System process model, and system administration, including new utilities and third-party relational database management systems (RDBMS) used to store iServer system and Encyclopedia volume metadata.
- *Part 2. Installing*. Describes how to install BIRT iServer.
- *Chapter 2. Installing BIRT iServer*. Describes how to install BIRT iServer using the out-of-the-box (OOTB) PostgreSQL relational database management system (RDBMS) in a Linux or UNIX environment.
- *Chapter 3. Installing BIRT iServer using an alternative database*. Describes how to install BIRT iServer using an alternative RDBMS, such as a pre-existing PostgreSQL or Oracle RDBMS, in a Linux and UNIX environment.
- *Chapter 4. Upgrading BIRT iServer*. Describes how to upgrade BIRT iServer in a Linux and UNIX environment.

- *Chapter 5. Installing a BIRT iServer cluster.* Describes how to install a BIRT iServer cluster node in Linux and UNIX.
- *Chapter 6. Installing BIRT iServer in a cloud.* Describes how to install BIRT iServer in a cloud environment using a ready-to-launch iServer image.
- *Chapter 7. Installing Information Console.* Describes how to install Actuate Information Console in Linux and UNIX.
- *Chapter 8. Installing iServer Integration Technology and Documentation.* Describes how to install BIRT iServer Integration Technology and Documentation in Linux and UNIX.
- *Part 3. Licensing.* Describes the licensing for BIRT iServer.
- *Chapter 9. Licensing BIRT iServer.* Describes licensing options, license key installation, and CPU-binding policies for BIRT iServer.
- *Part 4. Backing Up.* Describes how to back up a BIRT iServer.
- *Chapter 10. Backing up an Encyclopedia volume.* Describes how to back up and restore BIRT iServer Encyclopedia volume metadata and data.



# Part One



**Architecture**



# Understanding Actuate BIRT iServer architecture

This chapter contains the following topics:

- Understanding BIRT iServer architecture
- Understanding the iServer System process model
- Administering iServer System

---

## Understanding BIRT iServer architecture

Before Release 11, Actuate BIRT iServer used a proprietary relational database management system (RDBMS), known internally as the Squirrel database, to store the metadata related to iServer System and Encyclopedia volume configuration. In Release 11, Actuate replaced this out-of-the-box (OOTB) database with a customized version of the open-source, third-party database, PostgreSQL.

Actuate also adapted iServer to support alternative, customizable, third-party database installations. In Release 11 Service Pack 4, Actuate currently supports DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance.

In these third-party RDBMS, iServer stores metadata in the following schemas:

- System  
Contains settings related to iServer configuration, such as servers, templates, volumes, and partitions
- Encyclopedia volume  
Contains settings related to volume configuration, such as users, roles, groups, channels, folders, files, and other objects.

In Release 11, Actuate provides the following installation options:

- Install a new iServer with a PostgreSQL or other supported, alternative, third-party database
- Upgrade the Actuate proprietary database installation from a previous major release, such as Release 10 Service Pack 1, to a Release 11 version
- Upgrade an earlier Release 11 version to a newer Release 11 service pack, such as an upgrade from Release 11 Service Pack 3 to Service Pack 4
- Maintain a mix of Actuate PostgreSQL and third-party database implementations side-by-side in an iServer System installation

### Using a third-party RDBMS with an Encyclopedia volume

Actuate automatically installs the iServer system and Encyclopedia volume schemas in the OOTB PostgreSQL RDBMS installation. Installation of these schemas in a pre-existing PostgreSQL RDBMS or alternative RDBMS, such as DB2, Oracle, or SQL Server, requires manually running a SQL script containing the appropriate Data Definition Language (DDL) statements. The Installing section of this book contains chapters that provide detailed, step-by-step descriptions on how to perform these operations.

Actuate provides the iServer administrator with the ability to install the metadata for Encyclopedia volumes in databases in the same schema, separate schemas, or separate databases. By default, Actuate uses separate schemas for each Encyclopedia volume database, but provides the administrator with the option to specify whether to have volume databases share a schema. Actuate recommends using a separate schema for each Encyclopedia volume database for ease of administration.

In a PostgreSQL installation, the database administrator can manage an Encyclopedia volume either as an individual PostgreSQL database or as a schema in a PostgreSQL database. The same instance of a PostgreSQL server can manage multiple Encyclopedia volumes that use a mix of these configuration options. In PostgreSQL technical jargon, multiple databases managed by the same instance of the PostgreSQL server are in a PostgreSQL database cluster.

In Oracle, there is a one-to-one relationship between a database user and a schema. A schema is not a separate entity. An Actuate DB2, PostgreSQL, or SQL Server installation also requires this one-to-one relationship between a database user and a schema for consistency.

In a typical pre-existing RDBMS installation, the database administrator first creates a schema owner and a database user by running a SQL script. During iServer installation, the iServer system administrator provides the schema owner and database user credentials. The iServer installation program connects to the RDBMS, creates the necessary Encyclopedia volume database structures, then loads the metadata. The iServer application interacts with the third-party RDBMS using these database user credentials.

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

## **Customizing Encyclopedia volume databases**

Actuate supports read-only operations on the system and Encyclopedia volume metadata 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 system metadata database. Any customization that the customer does on the Encyclopedia volume database must be redone when migrating, reinstalling, or upgrading iServer. Actuate iServer does not track the objects that a customer creates. Actuate reserves the right to change the structure of the schema in future releases.

## Installing and configuring iServer System

The installation, configuration, and administration of an iServer System can include the following tasks:

- Install a new iServer using one of the following options:
  - Automated installation  
Run the installation program to configure iServer and the OOTB PostgreSQL database or an alternative, supported RDBMS.
  - Cloud deployment  
Deploy a prepared image of an installed iServer run-time environment. The administrator can create a customized image by generating an archive of an installed iServer run-time environment. Alternatively, an out-of-the-box (OOTB) image is available as a separate iServer distribution package for Windows. The administrator deploys the image by unbundling the archive or installing a virtual image on the target machine.
- Upgrade an earlier iServer installation to Release 11 using the installation program to overwrite automatically the earlier installation.  
Upgrade an earlier iServer System in place, such as Release 10 Service Pack 1, automatically migrating one or more Encyclopedia volumes.
- Upgrade an earlier iServer installation to Release 11 using the installation program to install iServer, then manually migrate Encyclopedia volume metadata and data from an earlier to the new installation.  
Upgrades an earlier iServer system in place, such as Release 10 Service Pack 1, without migrating any Encyclopedia volumes. During installation, the administrator chooses to migrate the volumes manually.  
After installation, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the existing iServer installation to either an in-place or a new, side-by-side instance of iServer Release 11 when upgrading from an earlier major release. When upgrading a Release 11 installation to a new Release 11 service pack, the administrator uses the Encyclopedia Data Store Upgrader utility. These utilities are Java programs run from the command line.

For more information on changes to iServer utilities in Release 11, see “About Migration and Administration Tools,” later in this chapter.

## Managing the backup, recovery, and failover capabilities of the Encyclopedia volume database and data files

The iServer administrator uses third-party RDBMS tools to manage the backup, recovery, and failover capabilities of the Encyclopedia volume database. The

iServer administrator uses standard operating system or other third-party tools to manage the backup and recovery of the data files.

Since iServer Release 11 no longer uses the internal proprietary Squirrel database to store Encyclopedia volume metadata, the automatic backup, recovery, and failover features available in earlier iServer releases are now obsolete. For information on the recommended procedures to back up and recover iServer system and Encyclopedia volume schemas in the Release 11 environment, see Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

In Actuate Release 11 and later, there is no concept of volume failover, since each node in a cluster can operate on all the volumes. Configuring system and Encyclopedia 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 failover capability.

In addition, consult the third-party RDBMS documentation for detailed information on how to use native system tools to configure backup, recovery, and failover operations for an externally managed Encyclopedia volume database.

Documentation for a PostgreSQL RDBMS is available at:

<http://www.postgresql.org/docs/8.4/static/release-8-4.html>

Documentation for an Oracle RDBMS is available at:

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

Documentation for Microsoft SQL Server RDBMS is available at:

<http://msdn.microsoft.com/en-us/sqlserver/bb671149>

Documentation for IBMDB2 RDBMS is available at:

<https://www-304.ibm.com/support/docview.wss?uid=swg27009474>

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the Encyclopedia volume schemas and all related file data to ensure the recoverability in the event of failure. For more information on backing up an iServer installation, see Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this operation to protect critical system metadata. The administrator can restore a corrupted or missing system schema using the System Data Store Administrator utility. For more information on this utility, see “Specifying System Data Store Administrator properties,” in Chapter 4, “Upgrading BIRT iServer,” later in this book.

An Actuate system administrator must take all necessary precautions to ensure that a database is properly backed up and available to safeguard Encyclopedia 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 an iServer cluster

In Actuate Release 11 and later, the concept of a master node no longer exists. Any node in a cluster has the ability to modify the shared server configuration file. The node performing these operations typically depends on which node the system administrator uses when connecting to iServer system through an administration console.

In Release 11 and earlier, iServer used multicasting to broadcast event information and synchronize operations in a cluster. Some cloud computing environments do not support multicasting. Starting in Release 11 Service Pack 1, iServer uses the third-party RDBMS as a shared repository for storing cluster information. This enhancement replaces multicasting as a way of managing cluster information.

---

## Understanding the iServer System process model

In Release 11, the Actuate BIRT iServer System platform uses a multi-threaded, multi-process model, running single instances of the following components on each iServer node:

- **Encyclopedia volume**  
Stores metadata in an OOTB (PostgreSQL) or alternative RDBMS and coordinates processing for designs, documents, information objects, and other iServer data objects stored in the file system.
- **Process Management Daemon (PMD)**  
Distributes service requests among available iServer services and nodes.
- **iServer servlet container**  
Provides the run-time environment for client applications, such as Actuate Information, Management, and Configuration Consoles. Client applications communicate with iServer System using SOAP-based messaging.

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

- **Factory**  
Executes requests to generate queries and documents and perform server-side printing.



- **View**  
Supports viewing documents in DHTML and other output formats, such as Excel and PDF. Handles requests to download files from an Encyclopedia volume.
- **Integration**  
Coordinates the running of information object (IOB) files that extract data from multiple data sources.
- **Caching**  
Controls the Actuate Caching process that manages an information object cache and enables caching of data retrieved from data sources.

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

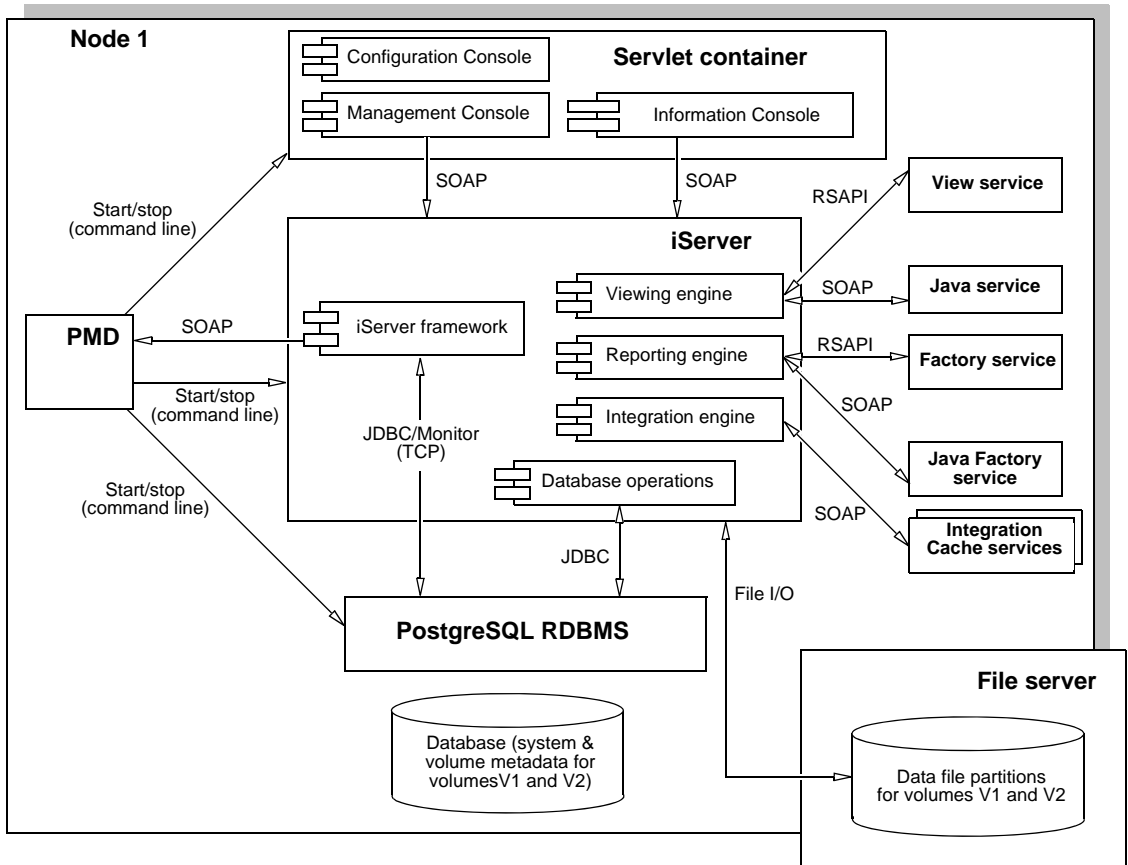
- Startup and shutdown of an iServer is fast because it is independent of the RDBMS that manages the Encyclopedia volume. The database server can remain online when shutting down an iServer and is available when the iServer starts up.
- Controlling the sequence of an Encyclopedia volume startup is not 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 an iServer is reduced. The RDBMS does not have to be shutdown.

## Understanding process flow in a stand-alone iServer

Figure 1-1 illustrates the iServer RDBMS process architecture for a stand-alone, two-volume, out-of-the-box (OOTB) PostgreSQL database configuration. In this configuration, the iServer administrator starts and stops an iServer instance by running scripts from the command line or using the graphical user interface (GUI) available in Configuration Console.

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

Client applications, such as Actuate Information, Management, and Configuration Consoles, run in a servlet container. These applications communicate with iServer using the Actuate Information Delivery API or IDAPI.



**Figure 1-1** iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

An IDAPI application uses a SOAP processor that serializes, or transforms, a remote procedure call by the application into an XML-based SOAP request to iServer to perform a web service. The application sends the message across the network using the Hypertext Transfer Protocol (HTTP) transport layer.

The Process Management Daemon (PMD) is a message distribution service that routes the request to iServer. iServer receives the request and deserializes the SOAP message. iServer performs the appropriate action and sends a response in the form of a SOAP message back to the application.

For example, iServer receives a request to run a design, such as a BIRT design, immediately or as a scheduled job. iServer communicates with the internal iServer framework and Encyclopedia volume metadata databases as necessary to locate the design and identify the resources required to run the design in the system.

The reporting engine selects a Java Factory service to run the BIRT design and checks job status. iServer 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 Excel and PDF, and handles requests to download files from the Encyclopedia volume. The View service sends the document to the requesting application for viewing.

A design that uses an information object utilizes the Integration and Caching services to perform the following processing:

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

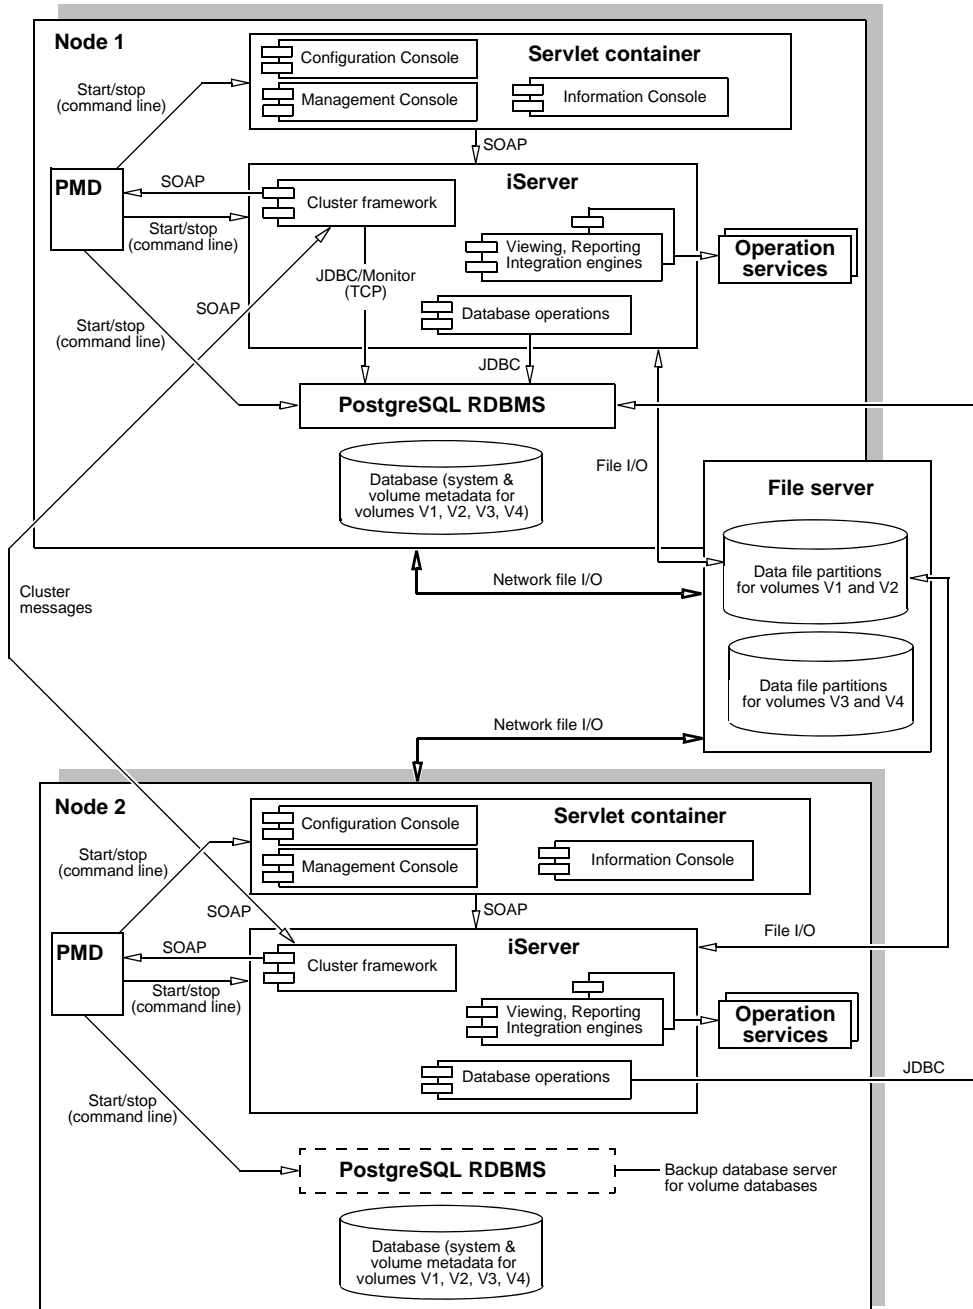
iServer stores system and Encyclopedia volume metadata in the third-party RDBMS, communicating with the RDBMS as necessary using JDBC. iServer uses the physical file system to read and store designs, documents, information objects, and other iServer objects as data in Encyclopedia volume partitions.

The out-of-the-box (OOTB) iServer PostgreSQL installation configures the Encyclopedia 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>

## **Understanding process flow in an iServer cluster**

Figure 1-2 illustrates the iServer RDBMS process architecture for a clustered, two-node, four-volume, OOTB database configuration. A node is a machine running an iServer instance.



**Figure 1-2** iServer RDBMS process architecture for a clustered, two-node, four-volume, OOTB database

The iServer OOTB PostgreSQL database server starts multiple instances to handle connections for running queries and accessing data. 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>

An iServer administrator adds a node to a cluster to scale iServer System to the necessary processing requirements. There are two methods of adding a new node to the cluster:

- Perform an automated, custom installation, using the wizard-driven installation program.
- Perform a manual installation or cloud deployment, using a prepared image of an installed iServer run-time environment.

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

- The shared configuration home directory
- Cluster resources, such as printers, database systems, 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 administrator launches the installed iServer image from the command line by passing the necessary arguments or creates 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 iServer clustering.

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

In the two-node cluster example, shown in Figure 1-2, client applications, such as Actuate Information, Management, and Configuration Consoles, run in a servlet container. These applications support distributing requests to multiple machines. The cluster communicates across the network using standard HTTP/IP addressing.

One or more nodes in the cluster manage the request message routing. The Process Management Daemons (PMDs) located on each node coordinate processing among available iServer services based on message type to balance load across the nodes.

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

- Each iServer 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 Encyclopedia volume metadata database. This information is now directly retrieved from or updated in the RDBMS, eliminating internal messages to Encyclopedia services on other nodes.

This increased scalability of operations at the iServer level can create bottlenecks at the RDBMS level. 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.

iServer instances running on multiple machines maintain iServer system and Encyclopedia volume metadata in databases and control access to shared volume data. The volume data can be on machines that are not running iServer, but must be shared and accessible to each iServer instance.

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

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

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

For more information about the cluster installation option, see Chapter 9, “Clustering,” in *Configuring BIRT iServer*.

---

# Administering iServer System

Administering an iServer System includes the following tasks:

- **Setting up users, roles, groups, channels, folders, files, and other administrative tasks**

An administrator creates, configures, and manages users, roles, groups, files, folders, and channels, including assigning and updating privileges, managing security role and group memberships, and providing access to channels. User, role, group, and channel privileges selectively control access to the Encyclopedia volume and its data objects.
- **Scheduling jobs to run designs and generate documents**

Each stand-alone iServer and node in an iServer cluster has a job scheduler and dispatcher. A job dispatcher send 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 allow iServer System to scale processing to handle thousands of scheduled jobs at the same time.
- **Reviewing logs and auditing the information to diagnose system problems**

iServer can capture usage and error information in log files to assist an administrator in evaluating resource usage and troubleshoot problems. The usage and error logging applications are open framework applications, which are available as DLLs in Windows and shared libraries in Linux or UNIX.
- **Configuring a cluster using automated installation programs and cloud computing base images**

The administrator can run the installation program to configure iServer or deploy a prepared image of an installed iServer run-time environment. Each cluster node gets its configuration from a template in `acserverconfig.xml`, located in a shared configuration home directory. Nodes with the same cluster ID, running on the same sub-net, automatically detect and join each other to form the cluster.
- **Using Actuate Server Integration Technologies scripts and tools to develop client applications and extend iServer functionality**

The Actuate Information Delivery application programming interface (IDAPI) supports integrating and administering iServer using extensible markup language (XML) and the simple object access protocol (SOAP). Using the IDAPI, developers can create applications that perform such tasks as scheduling a custom event, running an Report Server Security Extension (RSSE) application to manage users and roles in an external system such as an LDAP server, and installing and customizing usage and error logging and performance monitoring extensions.

A BIRT iServer administrator uses the Actuate Information, Management, and Configuration Consoles, command-line utilities, and Server Integration Technology components to perform these tasks.

Please consult the following Actuate iServer Release 11 documentation for more information on how to administer an iServer System using these components:

- *Installing BIRT iServer for Windows or Installing BIRT iServer for Linux and UNIX*  
Describes iServer System architecture. Provides detailed instructions on how to use automated installation programs and command-line utilities to install stand-alone iServer and clustered nodes that store Encyclopedia volume metadata in an external, third-party RDBMS, such as DB2, Oracle, PostgreSQL, or SQL Server. Also describes Actuate licensing policies and procedures and backup and recovery operations.
- *Managing an Encyclopedia Volume*  
Describes how to use Management Console and command-line options to perform tasks such as managing Encyclopedia volume user accounts, assigning privileges, scheduling jobs, and distributing documents.
- *Configuring BIRT iServer*  
Describes how to use Configuration Console to perform tasks such as managing an iServer cluster, adding Encyclopedia volumes to iServer, connecting to databases, updating the license, and configuring iServer properties, such as logging levels, e-mail notification, and printing from iServer.
- *Using BIRT iServer Integration Technology*  
Provides information about application programming using the SOAP-based Actuate Information Delivery API (IDAPI), including a Java developer guide and sections on logging, auto archiving, and using the Java Report Server Security Extension (RSSE).

## About Migration and Administration Tools

In Actuate BIRT iServer Release 11, the following utilities are obsolete and no longer exist:

- AcExport  
Formerly used to write a copy of the Encyclopedia volume metadata to a file, so the administrator can import the metadata into another release of iServer.
- AcImport  
Formerly used to populate an Encyclopedia volume with metadata previously written to an exported file.  
In an upgrade from an earlier major release to Release 11, use the Squirrel Data Exporter and the System and Encyclopedia Data Store Administrator utilities



that Actuate provides to migrate Encyclopedia volume metadata from the Squirrel database to an alternative RDBMS. Use the Encyclopedia Data Store Upgrader utility to upgrade an earlier Release 11 installation to a newer Release 11 service pack.

Back up the database using the utilities that the RDBMS provides. For example, PostgreSQL provides the `pg_dump` and `pg_restore` utilities to create and restore a database backup.

Use operating system or other third-party tools to back up and load designs, documents, information objects, and other iServer data objects stored in the file system. For more information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

- **AcToc**  
Formerly used to list the contents of an export directory. In Release 11, no comparable functionality exists.
- **AcVerify**  
Formerly used to validate an offline Encyclopedia volume and repair problems.  
  
In Release 11, use the tools available in the third-party RDBMS containing the Encyclopedia volume metadata to verify data integrity and make repairs.
- **AcExtern**  
Formerly used to convert Encyclopedia volume user security from internal to external registration to allow administration from another system, such as an LDAP server. In Release 11, use the iServer Integration Technology custom application as a reference to configure the Report Server Security Extension (RSSE) when implementing external registration.
- **AcIntern**  
Formerly used to convert Encyclopedia volume user security from external to internal registration from an LDAP or other system to iServer. In Release 11, no comparable functionality exists. For more information on how to install and configure RSSE in the Release 11 environment, refer to Chapter 11, “Configuring iServer security,” in *Configuring BIRT iServer* and Chapter 10, “Using Java Report Server Security Extension” in *Using BIRT iServer Integration Technology*.
- **AcMode**  
Formerly used to put an Encyclopedia volume in and out of online backup mode. In Release 11, a dynamic backup no longer requires putting the system into online backup mode. The administrator performs an Encyclopedia volume metadata backup using the tools provided by the third-party RDBMS,

which provides comparable features. The administrator uses standard operating system or other third-party tools to back up the data files.

- **AcEncycUpgrade**

Formerly used to convert an older Encyclopedia volume to the latest version. In a manual Release 11 upgrade process, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the existing iServer installation to either an in-place or a new, side-by-side instance of iServer Release 11. These utilities are Java programs run from the command line.

For more information on how to upgrade iServer and Encyclopedia volume to Release 11, refer to Chapter 4, "Upgrading BIRT iServer."

## **Using JDBC to connect to an Encyclopedia volume database**

iServer uses JDBC for connecting to an Encyclopedia volume database. The iServer run-time JRE environment uses Java 1.6. Any JDBC driver must be compatible with JRE version 1.6 or earlier.

iServer 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.

An administrator, who decides to customize iServer 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 is used to capture the connection pool run-time statistics.

## **API Compatibility**

Actuate Release 11 provides full backward compatibility with existing applications. Upgrading to an Actuate Release 11 iServer that utilizes an RDBMS has no impact on any applications that utilize Actuate APIs, such as IDAPI and RSSE.

## **About international character sets**

iServer operates on the assumption that the volume 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 iServer Encyclopedia 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 iServer Release 11 Service Pack 4 supports the following operating systems:

- Windows
- Solaris
- Linux



# Part **Two**

---

**Installing**



# Installing BIRT iServer

This chapter discusses the following topics:

- Preparing to install BIRT iServer
- Performing a new installation
- Understanding the iServer installation environment

---

## Preparing to install BIRT iServer

When installing BIRT iServer Release 11, the administrator must choose to use the out-of-the-box (OOTB) PostgreSQL relational database management system (RDBMS) or another data store, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance to store Encyclopedia volume metadata. This chapter describes how to install a new instance of BIRT iServer Release 11, using the out-of-the-box (OOTB) PostgreSQL RDBMS.

For more information about installing BIRT iServer using an alternative RDBMS, see Chapter 3, “Installing BIRT iServer using an alternative database,” later in this book. For more information about upgrading an existing BIRT iServer installation, see Chapter 4, “Upgrading BIRT iServer,” later in this book. For information about the new Release 11 BIRT iServer System architecture, see Chapter 1, “Understanding Actuate BIRT iServer architecture,” earlier in this book.

### Creating a dedicated user account for installing and running iServer

Before installing iServer, create a dedicated Linux user account for installing, running, and administering iServer. Having a dedicated user account isolates iServer-specific issues and events on a machine, making it easier to administer the environment. Use the same level of security that your site exercises for other system administrator and root accounts.

Installation of iServer under the root account is not recommended and not supported when using the PostgreSQL RDBMS. The PostgreSQL RDBMS must run using an unprivileged user ID to prevent compromising system security. If installed under the root account, the default installation is unable to set up the required iServer metadata schemas and Encyclopedia sample volume.

### Backing up iServer system and Encyclopedia volume metadata

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the schemas using the tools and resources of the third-party database system.

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult Actuate Support at the time of installation if you have any questions about these backup procedures to protect against the possibility of catastrophic failure. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 10, “Backing up an Encyclopedia volume,” later in this book.



When installing iServer, be sure to run the same versions of all products. Upgrade all products at the same time to maintain consistency in the versions you run.

If you are a purchasing customer, you can download iServer from an Actuate FTP software distribution site. If you are evaluating iServer, you can download iServer from BIRT Exchange at the following location:

<http://www.birt-exchange.com>

## About X frame buffer

Xvfb is an X Windows server that has neither a graphics card nor a physical graphics display. BIRT iServer uses the X server for font-rendering information and to generate graphics in reports. Typically, an X server requires a graphics card and physical graphics display on the iServer machine, but you can use Xvfb in place of these components.

The Xvfb software installed with BIRT iServer includes Type 1 fonts. Actuate maps these fonts to Microsoft Windows fonts for consistent graphics rendering on the various platforms.

The Xvfb software uses:

- X libraries installed on the BIRT iServer machine  
If you choose to install Xvfb, the installation script searches for the required libraries and displays a message if the install script cannot find the required libraries on the machine.
- Variables set to the path of the Xvfb libraries
  - XVFBDISPLAY variable in start\_svr.sh.
  - display\_value in pmd11.sh.

To view and print the reports from BIRT iServer, you need to set these variables only if you install Xvfb software.

The DISPLAY environment variable specifies the X Windows server used by the iServer machine. For example, if the iServer machine is running X Windows, it sets DISPLAY to the local machine:

```
# setenv DISPLAY :0.0
```

If you use a separate machine as the X Windows server, specify the machine name in the environment variable DISPLAY. The following example sets DISPLAY to use an X Windows server on a machine named urup:

```
# setenv DISPLAY urup:0.0
```

The original source code for Xvfb is included as a component of X11R6, but not in earlier X Window system releases.

## Installing X frame buffer

Actuate distributes Xvfb for the Sun and AIX operating systems, and installation and configuration of Xvfb is a BIRT iServer installation option in these environments. To use Xvfb in HP-UX, you must install Xvfb before you install BIRT iServer.

## About libstdc++

The libstdc++ library is a prerequisite for installing Actuate BIRT iServer on Linux and UNIX systems. This library is present by default on most systems. If it is not present, the administrator must install it before installing iServer.

When installing BIRT iServer on a Linux machine, the following message may appear if you have the 64-bit version of the c/c++ run-time installed:

```
Error: An error occurred in the license reading program. Please
      make sure you have all the recommended patches and the right
      c/c++ runtime environment installed on this machine.
```

If this message appears, install the following 32-bit c++ run-time libraries:

```
compat-libstdc++-33.i686
libstdc++.i686
```

## About run levels

The iServer installation process requires running Linux or UNIX at run level 5. This level supports networking and multi-user mode with a graphical window manager. Run level 5 is typically the default on most Linux or UNIX operating system distributions.

## About Openmotif

On Linux platforms, the Openmotif bundle is required and must be installed before installing BIRT iServer. BIRT iServer needs libXm.so.3 or libXm.so.4, which are part of Openmotif 2.2 and 2.3, respectively. If the BIRT iServer installation is unable to locate the required libXm.so.x library, create a symbolic link, as shown in the following example:

```
ln -s libXm.so.3 libXm.so.4
```

---

## Performing a new installation

Installing a new Release 11 BIRT iServer creates a default Encyclopedia volume without migrating data from a pre-existing volume. The default installation program performs the following operations:

- Installs and initializes iServer and the PostgreSQL relational database management system (RDBMS)
- Creates a database in the PostgreSQL RDBMS containing Encyclopedia volume data
- Creates the iserver user in the PostgreSQL RDBMS to access the system
- Creates the system and volume schema, initializing these schema with basic configuration information
- Creates the iServer configuration file, specifying system, volume, and connection information for the default installation

The default installation program also initializes the iserver, system, and volume user passwords to the PostgreSQL superuser password.

## Installing BIRT iServer Release 11

The following section describes how to install a new, stand-alone instance of BIRT iServer Release 11 in the UNIX and Linux operating system.

### How to perform a new installation of BIRT iServer Release 11 in UNIX and Linux

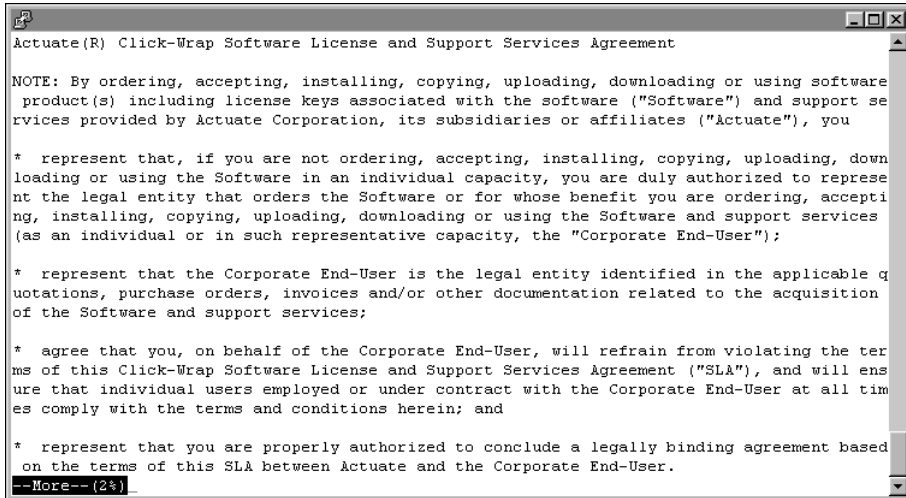
To reduce network traffic, you can install BIRT iServer on the same host machine as the iServer system database. You can also install BIRT iServer and the metadata database on a different machine to distribute processing across multiple machines.

Actuate recommends running the installation procedure from an account created exclusively for iServer administration. To install iServer, perform the following steps:


- 1 Download the required files from the FTP software distribution site. Extract the files.
- 2 To install the server files, execute the isinstall script:
 

```
sh ./isinstall.sh
```

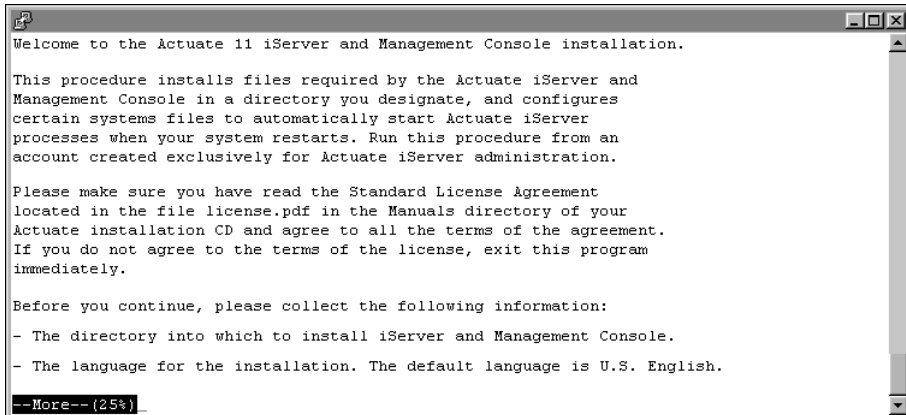
The script displays a series of prompts. Respond to the prompts as described in the following procedures.
- 3 The license agreement appears, as shown in Figure 2-1.
- 4 Read the license agreement, then press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 2-2.
- 5 The introduction to the installation appears, as shown in Figure 2-3.



**Figure 2-1** Reviewing the license agreement

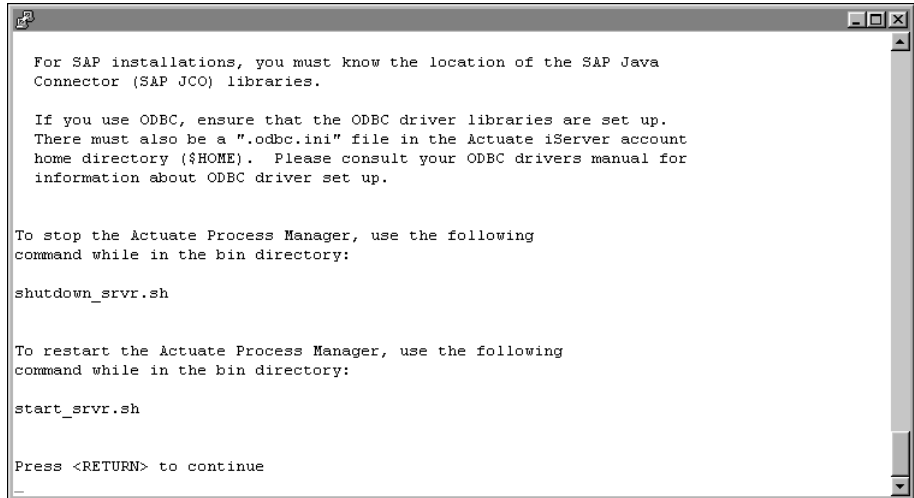


**Figure 2-2** Accepting the licensing terms



**Figure 2-3** Reviewing the introductory information

- 6 Press Return or Enter after finishing the review of the introductory information, as shown in Figure 2-4.



```
For SAP installations, you must know the location of the SAP Java
Connector (SAP JCO) libraries.

If you use ODBC, ensure that the ODBC driver libraries are set up.
There must also be a ".odbc.ini" file in the Actuate iServer account
home directory ($HOME). Please consult your ODBC drivers manual for
information about ODBC driver set up.

To stop the Actuate Process Manager, use the following
command while in the bin directory:

shutdown_srvr.sh

To restart the Actuate Process Manager, use the following
command while in the bin directory:

start_srvr.sh

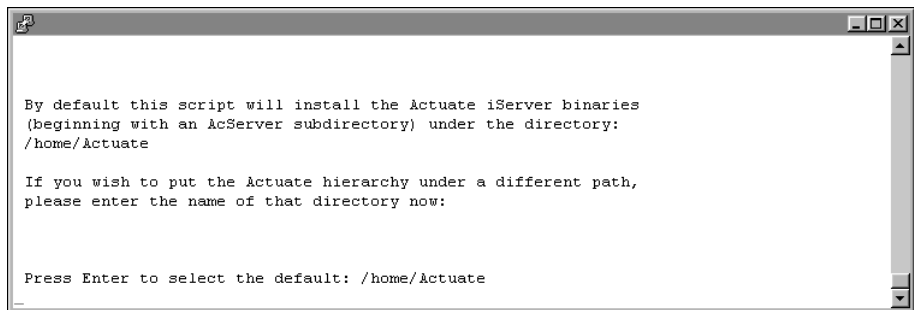
Press <RETURN> to continue
```

**Figure 2-4** Finishing the review of introductory information

- 7 Press Enter to accept the default installation directory for Actuate iServer binaries, as shown in Figure 2-5. Alternatively, type a different directory and press Enter.

The installation program creates the Actuate directory in the chosen location. iServer uses this location to resolve the path to all binaries that it launches.

The default path is /home/Actuate. This documentation uses the environment variable AC\_SERVER\_HOME to refer to \$HOME/AcServer in case the installer chooses a path that is different from the default path.



```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

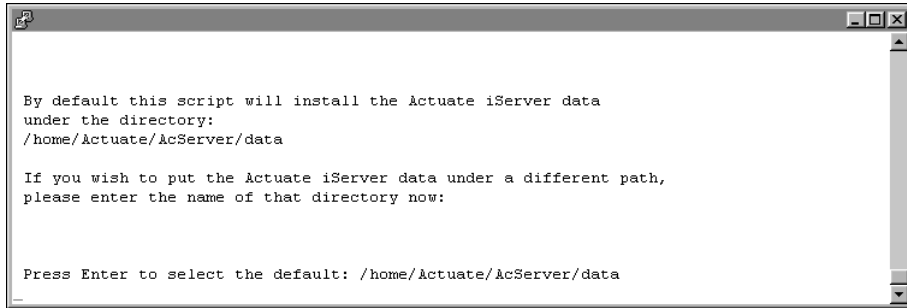
If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
```

**Figure 2-5** Specifying the installation directory

- 8 Press Enter to accept the default installation directory, AC\_SERVER\_HOME /data, for iServer data, as shown in Figure 2-6. Alternatively, choose a different directory for iServer data.

iServer uses this data location to store the iServer Encyclopedia volume data, including PostgreSQL metadata, logs, and other files.



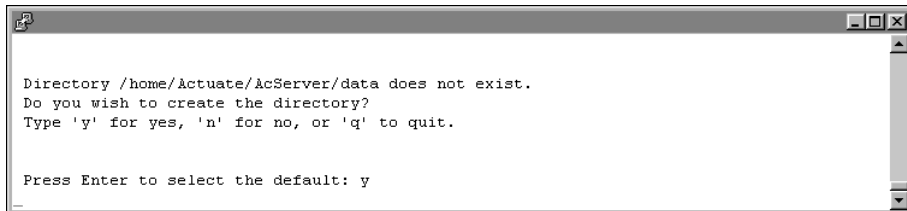
```
By default this script will install the Actuate iServer data
under the directory:
/home/Actuate/AcServer/data

If you wish to put the Actuate iServer data under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate/AcServer/data
```

**Figure 2-6** Specifying the data installation directory

- 9 Press Enter to accept the default option of creating the directory for data, as shown in Figure 2-7. Alternatively, type n for no, or q to quit, and press Enter.

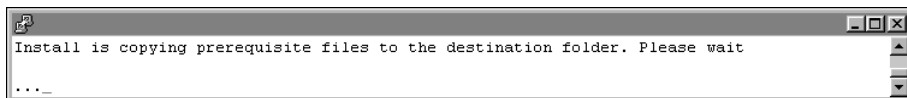


```
Directory /home/Actuate/AcServer/data does not exist.
Do you wish to create the directory?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 2-7** Creating the AC\_DATA\_HOME directory

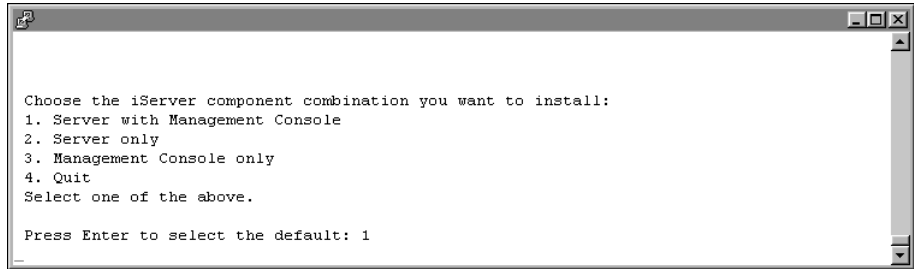
- 10 The installer copies prerequisite files to the destination folder, as shown in Figure 2-8. After the prerequisite files are copied, the installation continues.



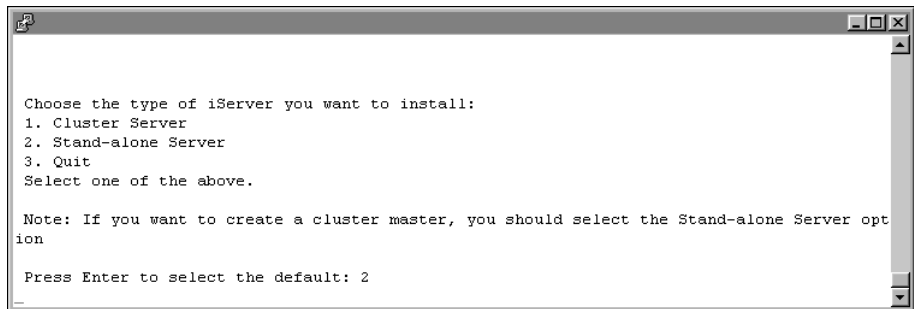
```
Install is copying prerequisite files to the destination folder. Please wait
..._
```

**Figure 2-8** Copying prerequisite files

- 11 Press Enter to accept the default iServer component combination, which includes Management Console, as shown in Figure 2-9. Alternatively, choose a different component combination and press Enter.
- 12 Press Enter to accept the default stand-alone Server installation, as shown in Figure 2-10. Alternatively, choose a different type of iServer to install. For information on how to install an iServer cluster, see Chapter 5, “Installing a BIRT iServer cluster.”




**Figure 2-9** Choosing the components to install




**Figure 2-10** Specifying the type of iServer to install

- 13 Type a name to use for the BIRT iServer System name, as shown in Figure 2-11. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.




**Figure 2-11** Specifying the BIRT iServer System name

- 14 Press Enter to choose the default embedded PostgreSQL database to store the Encyclopedia volume metadata, as shown in Figure 2-12.




**Figure 2-12** Choosing the embedded PostgreSQL

- 15 Press Enter to choose the default PostgreSQL superuser name, postgres, as shown in Figure 2-13. Alternatively, type a different PostgreSQL superuser name. This superuser administers the PostgreSQL relational database management system (RDBMS).




**Figure 2-13** Choosing the PostgreSQL superuser name

- 16 Type a PostgreSQL superuser password that conforms to the password security policy requirements within your organization, then press Enter, as shown in Figure 2-14.




**Figure 2-14** Typing the PostgreSQL superuser password

- 17 Re-enter the password for PostgreSQL superuser, then press Enter, as shown in Figure 2-15.



**Figure 2-15** Re-entering the PostgreSQL superuser password

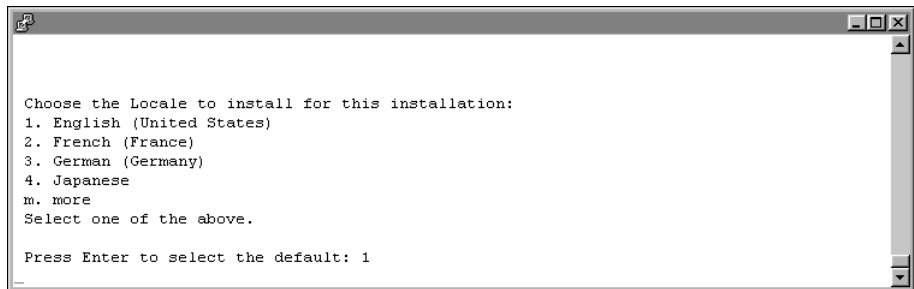
- 18 Press Enter to accept the default port on which PostgreSQL listens for requests, as shown in Figure 2-16. Alternatively, enter a different port number.



**Figure 2-16** Choosing the PostgreSQL port number

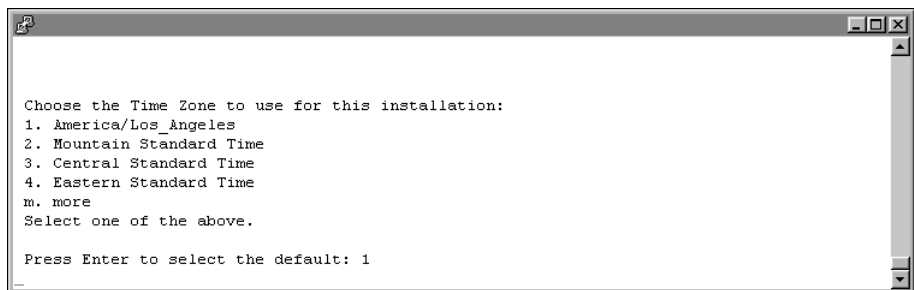


- 19** Press Enter to select the default locale, which is English, as shown in Figure 2-17. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.




**Figure 2-17** Specifying a locale

- 20** Press Enter to select the default time zone, which is America/Los\_Angeles, as shown in Figure 2-18. Alternatively, select another time zone from the numbered list.




**Figure 2-18** Specifying a time zone

- 21** To evaluate the product using the included evaluation software license press Enter, as shown in Figure 2-19. Alternatively, type 2, then type the path to a purchased license file.




**Figure 2-19** Specifying license type

- 22** Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 2-20. Alternatively, type a different IP address.




**Figure 2-20** Specifying the hostname that Management Console uses to contact the PMD

- 23 Press Enter to accept the default port number, 8100, where Process Management Daemon (PMD) listens for requests, as shown in Figure 2-21. Alternatively, type a different port number and press Enter.




**Figure 2-21** Specifying the port number on which the PMD listens

- 24 Press Enter to accept the default hostname, the name of the machine on which iServer runs, as shown in Figure 2-22. Alternatively, type a different hostname or IP address, then press Enter.




**Figure 2-22** Specifying the machine on which the iServer runs

- 25 Press Enter to accept the default port number where iServer listens to requests, as shown in Figure 2-23. Alternatively, type a different port number and press Enter.




**Figure 2-23** Specifying the port number on which the iServer listens

- 26 Specify the iServer administrator password, as shown in Figure 2-24. You use this password to log into the iServer Configuration Console.




**Figure 2-24** Specifying the iServer administrator password

- 27** Re-enter the password of the iServer administrator, as shown in Figure 2-25. You use this password to log in to Configuration Console.




**Figure 2-25** Re-entering the iServer administrator password

- 28** Press Enter to accept the default option to use a volume name for the Encyclopedia, as shown in Figure 2-26. Alternatively, type n for no to not use a volume name for the Encyclopedia, or q to quit the installation.




**Figure 2-26** Specifying whether to use an Encyclopedia volume name

- 29** Press Enter to accept the default Encyclopedia volume name, the BIRT iServer System name, as shown in Figure 2-27. Alternatively, type a different Encyclopedia volume name.




**Figure 2-27** Specifying the Encyclopedia volume name

- 30** Press Enter to accept the default option to start iServer automatically, as shown in Figure 2-28. Alternatively, type n for no.




**Figure 2-28** Specifying whether to start iServer automatically

- 31 Press Enter to accept the default option to not integrate LDAP with iServer, as shown in Figure 2-29. Alternatively, you can edit the setting.



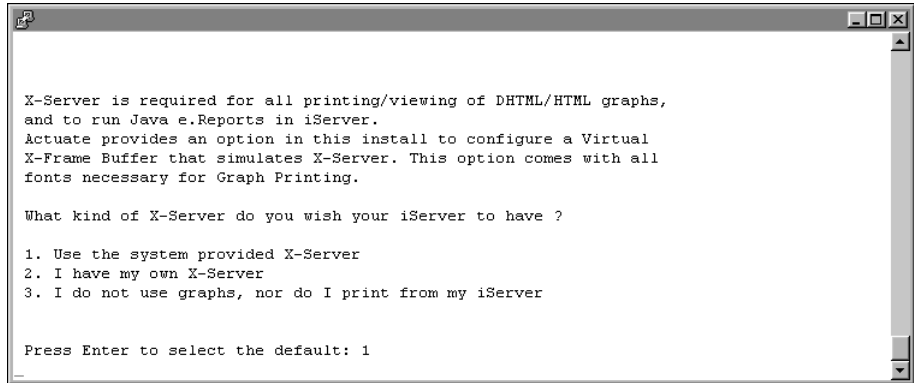
**Figure 2-29** Specifying whether to integrate LDAP with iServer

- 32 Press Enter to accept the default option to not use any database drivers/clients, as shown in Figure 2-30. Alternatively, type y for yes, and specify the database drivers/clients you want to use.




**Figure 2-30** Specifying whether to use database drivers/clients


- 33 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 2-31.
- 34 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 2-32. Alternatively, type a different IP address.
- 35 Press Enter to accept the default port number, 8100, on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 2-33. Alternatively, type a different port number.



**Figure 2-31** Specifying what kind of X-Server to use, if any




**Figure 2-32** Specifying the hostname that Management Console uses to contact the PMD




**Figure 2-33** Specifying the port number for the PMD to listen for requests from Management Console

- 36** Press Enter to accept the hostname or enter the IP address of the machine that Management Console uses to contact iServer, as shown in Figure 2-34. Alternatively, type a different IP address.




**Figure 2-34** Specifying the hostname that Management Console uses to contact iServer

- 37** Press Enter to accept the default port number, 8000, on which iServer listens for requests from Management Console, as shown in Figure 2-35. Alternatively, type a different port number.




**Figure 2-35** Specifying the port number on which iServer listens for requests from Management Console

- 38** Press Enter to accept the name of the default Encyclopedia volume to use with Management Console, as shown in Figure 2-36. Alternatively, type a different name for the Encyclopedia volume.



**Figure 2-36** Specifying the name of the default Encyclopedia volume


- 39** Press Enter to accept the default name, acadmin, for the HTTP server context root for Management Console configuration, as shown in Figure 2-37. Alternatively, type a different name.



**Figure 2-37** Specifying the name of the HTTP server context root

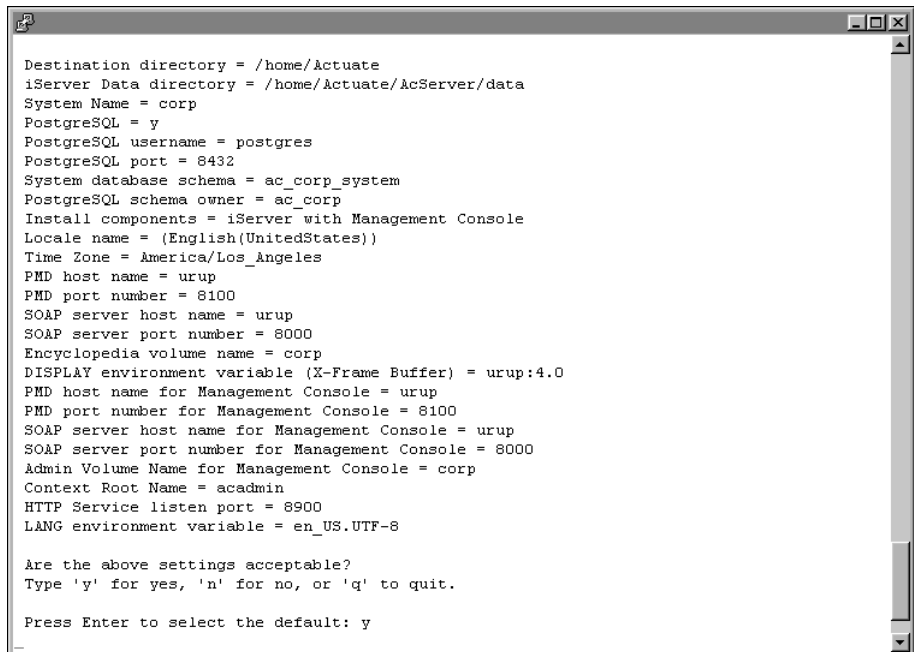
- 40** Press Enter to accept the default HTTP port number, 8900, on which the application container listens to requests, as shown in Figure 2-38. Alternatively, choose a different port.

You connect to the port from your browser when accessing various features of iServer.




**Figure 2-38** Specifying the application container listening port number

- 41 Review the settings, as shown in Figure 2-39, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.




**Figure 2-39** Reviewing the installation settings

- 42 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 2-40.



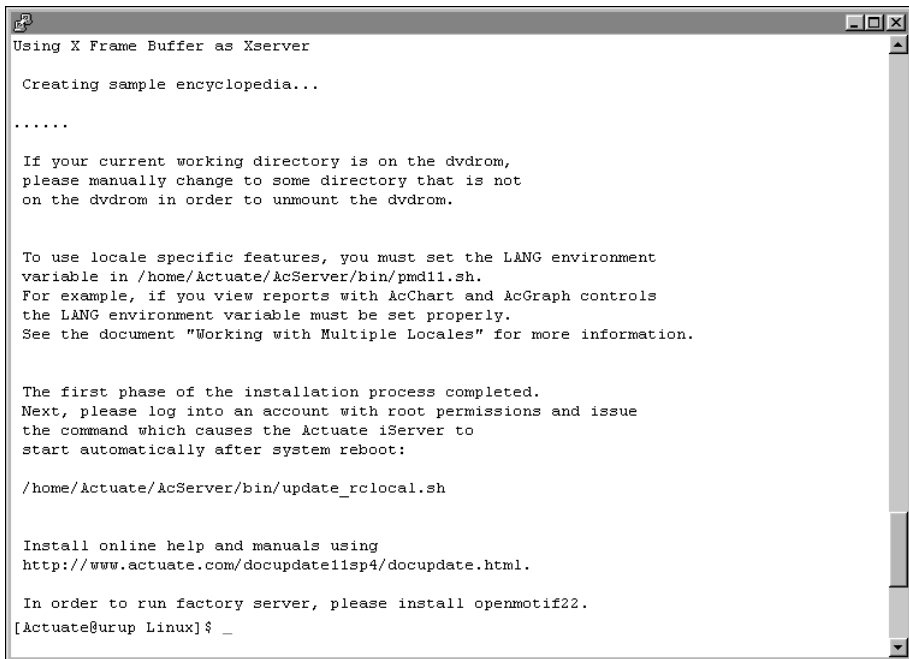
**Figure 2-40** Viewing iServer installation progress

- 43 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 2-41.



**Figure 2-41** Specifying whether to start iServer

- 44** When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 2-42.



**Figure 2-42** Viewing information about localization, logging in, and installing online help

## Accessing Information, Management, and Configuration Consoles

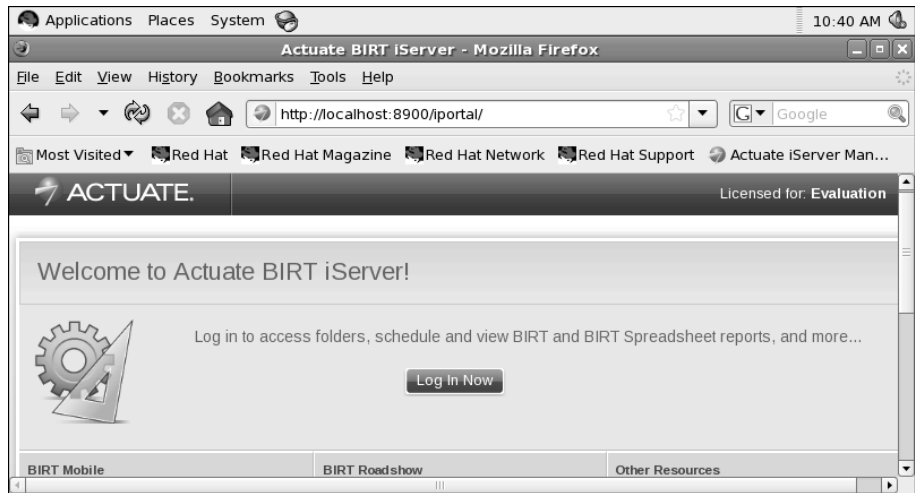
After the installation program finishes running, open a browser to log in to the following BIRT iServer 11 consoles to perform user and administrator tasks:

- **Information Console**  
Perform tasks such as accessing folders and viewing designs and documents.



To access Information Console, open a browser manually and enter the following URL, as shown in Figure 2-43:

`http://localhost:8900/iportal/`



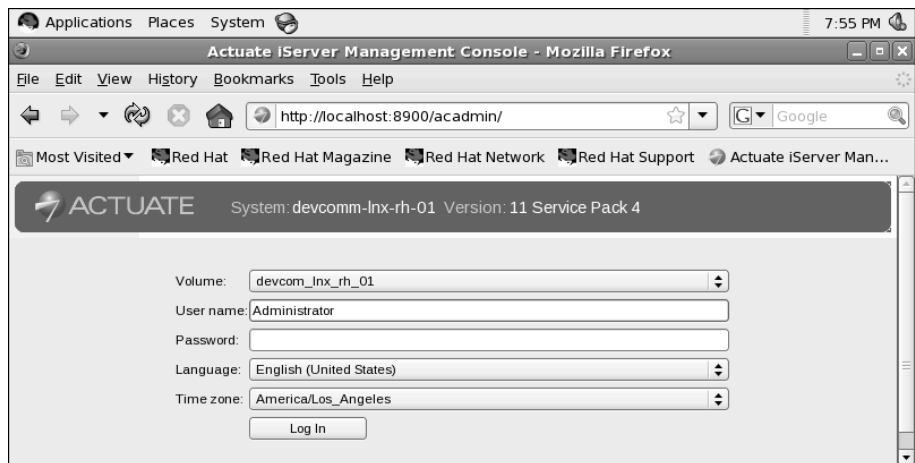
**Figure 2-43** Viewing Welcome to Actuate Information Console

- Management Console

Set up user accounts and schedule or run a design.

To access Management Console, open a browser manually and enter the following URL, as shown in Figure 2-44:

`http://localhost:8900/acadmin/`



**Figure 2-44** Logging in to Management Console

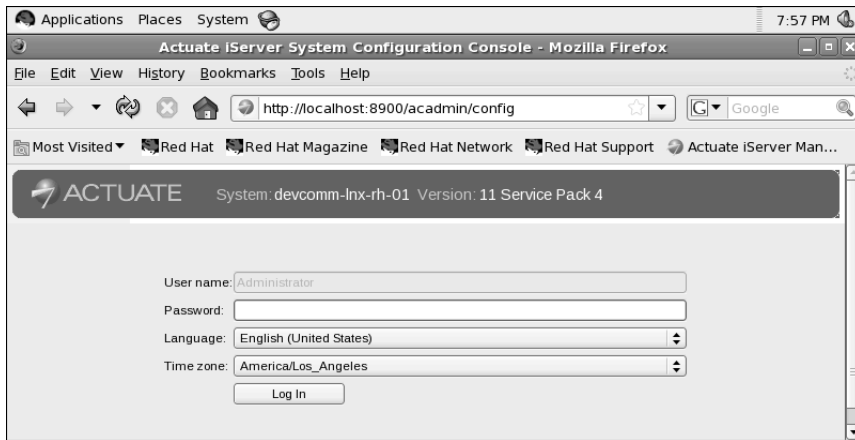
- Configuration Console

Perform administrative operations, such as the following tasks:

- Add an Encyclopedia volume
- Connect to a database
- Make modifications to iServer parameters and server templates
- Update the license

To access Configuration Console for administering iServer, open a browser manually and enter the following URL, as shown in Figure 2-45:

`http://localhost:8900/acadmin/config/`



**Figure 2-45** Logging in to Configuration Console

When starting PostgreSQL for Actuate iServer and Actuate BIRT iServer, the best practice is to start PostgreSQL then iServer. When stopping these programs, stop iServer then PostgreSQL, if necessary.

It is not necessary to shut down the database when starting and stopping iServer. iServer starts faster with the database already running and available in the background, which is particularly important in cluster and other high-performance operations.

The following sections describe how to perform these operations in the recommended order.

## Stopping and starting iServer and PostgreSQL processes

After installing iServer, an administrator can stop and start iServer and PostgreSQL processes from a command prompt using scripts provided in the installation.

### **How to stop and start BIRT iServer**

To stop iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./shutdown_srvr.sh
```

To start iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./startsrvr.sh
```

### **How to stop and start PostgreSQL for Actuate iServer**

To stop PostgreSQL, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin
- 2 Type the following command and press Enter:

```
sh ./stoppostgresql.sh
```

To restart PostgreSQL for Actuate iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./startpostgresql.sh
```

---

## **Understanding the iServer installation environment**

The following sections provide supplementary information about the iServer installation environment that is useful for an administrator to know.

### **About migrating an earlier iServer release to Actuate 11**

To upgrade an iServer earlier than Release 8 Service Pack 1, you must first upgrade to Release 8 Service Pack 1, 9, or 10 then upgrade to Release 11. To upgrade from an earlier release to Actuate 11 by replacing the older version, install Release 11 in the same directory as the earlier release, or perform a side-by-side installation, as described in Chapter 4, “Upgrading BIRT iServer.”

The upgrade program preserves any previous iServer configuration information and reuses the earlier settings. The installation program resolves any differences in default values between releases, ignoring old configuration defaults in favor of

new default values. For example, the installation does not prompt the user for port information and machine name. The installation detects the current port numbers and machine name and keeps those settings.

## **About running different releases on the same machine**

An Actuate 11 iServer can coexist on the same machine with an earlier major release, such as Release 10, and any associated minor release, such as Release 10 Service Pack 1.

Actuate does not support running multiple releases from the same version on a machine. For example, you cannot run Release 11 and Release 11 Service Pack 4 on the same machine.

To run different iServer major releases on the same machine, install the releases in separate directories. Change the default port settings for one release to enable running both versions at the same time.

## **About performance and disk space issues**

During an upgrade installation, the following operations can consume more disk space and take longer than a fresh installation:

- File comparison
- The copy operation to back up original files

During an upgrade, disk space requirements typically double. The installation routine copies files to the local machine for comparison between the original files and the new files. If you perform multiple upgrade installations, the installation routine consumes even more disk space for the backup files you need to restore previous installations.

## **About upgrading an iServer with resource groups**

When you upgrade BIRT iServer System, iServer creates a resource group on an iServer node that has the Factory service and the View service enabled. If no node has the Factory service and the View service enabled, iServer creates resource groups with zero Factory processes.

## **About the Java Software Development Kit**

The iServer installation routine installs the JDK files in:

`AC_SERVER_HOME/jdk160`

To use a different JDK with iServer, change the files in the installation directory or change the values of the following environment variables:

- AC\_JAVA\_HOME
- AC\_JVM\_HOME
- AC\_JRE\_HOME
- AC\_JRE64\_HOME

Using an earlier release of JDK can cause some Actuate features to fail or to work improperly. For example, using an earlier release of JDK can cause Actuate products to display Actuate report charts incorrectly.

The following types of Actuate report object executable files use AC\_JRE\_HOME and AC\_JVM\_HOME:

- Files containing charts use AC\_JVM\_HOME to locate the java.exe to generate the chart
- Files using the Actuate Java Object Interface use AC\_JVM\_HOME to locate the JVM DLL or library

## Accessing JAR files for report generation

To generate some reports, iServer requires access to jar files in the Jar directory of the iServer installation files. In UNIX and Linux, specify the CLASSPATH in the Process Management Daemon (PMD) startup script, pmd11.sh.

## Gathering LDAP information

An optional Open Security application ships with Actuate iServer Integration Technology. This application uses a Lightweight Directory Access Protocol (LDAP) security database to control access to the Encyclopedia volume. To use the Open Security application, you need to perform a custom installation and specify the following additional information:

- Name of the LDAP server and the port on which the LDAP server listens
- LDAP account and password used to query the LDAP server
- LDAP server organization that contains the Actuate users, roles, and groups
- LDAP base domain names and object classes that contain Actuate user, role, and group information
- LDAP group name assigned as the Actuate Encyclopedia volume administrator role

Actuate Open Security uses an LDAP configuration file to map the Encyclopedia volume user information to LDAP object attributes. For more information on

Actuate Open Security, see the reference implementations available in Actuate Server Integration Technology.

## Following best practices

Before deploying a new release in a production environment, Actuate recommends testing the installation.

### Using a test environment

Set up a test environment then migrate to Actuate 11 when the testing is complete. Earlier Actuate releases and Actuate 11 can coexist on the same machine. You must install products at different release levels in different folders. You cannot mix Actuate products from different release levels. For example, you cannot use Actuate 10 design tools with BIRT iServer Release 11.

How and when you upgrade to Actuate 11 depends on your site configuration and requirements. Complete the following general tasks in this order to determine how to upgrade your site to Actuate 11:

- Create a test environment for Actuate 11. The test environment can be on the same machine that hosts the earlier Actuate installation or on a separate machine.
- Install the software in the test environment and upgrade earlier versions of report designs and files. Also update any custom applications that you built using Actuate iServer Integration Technology. Verify that your applications work properly in the test environment.
- Ask application developers and a few users to perform some typical tasks in the test environment.
- Create a production staging area.
- Install the remaining Actuate 11 desktop products, if required, in production environments on the user workstations. Verify that the desktop products function properly.
- Schedule a low-impact time to switch to Actuate 11 to complete the transition.

### Setting up a production staging area

A production staging area is one that you can use for testing and also configure as the live production system. The production staging area can be a separate configuration on the live production machine or a separate machine. You can install all Actuate 11 products or the Actuate 11 server products and a subset of the desktop products.

If you plan to test Actuate 11 desktop products, identify which users to include in the final testing. Developers and users can then confirm that applications perform as expected in the Actuate 11 production staging environment.

Complete the following general tasks to test Actuate 11:

- Install BIRT iServer Release 11 software in a production staging area.
- Install Actuate 11 desktop software on the test user machines. Using separate folders, you can install Actuate 11 desktop software in conjunction with the earlier desktop software. Users can continue to use the existing Actuate software in production while testing the Actuate 11 desktop software.
- Verify that the Actuate 11 production staging environment works correctly.
- Install the remaining Actuate 11 desktop products, if you installed a subset earlier.
- Verify that all the Actuate 11 desktop products work correctly.
- Begin setting up a production environment, described in the following section.

## Setting up a production environment

When testing is complete, confirm that your applications work as expected in the Actuate 11 environment. Set up the production environment and schedule a date and time to switch from earlier versions to Actuate 11.

When you switch to Actuate 11, use the following procedure list as a general guideline:

- Shut down all Actuate servers.
- Back up earlier Actuate Encyclopedia volumes.
- Upgrade existing Encyclopedia volumes and archive volumes. Install upgraded designs and report files.

To upgrade from Actuate 10, for example, perform one of the following operations:

- If you are replacing your Actuate 10 system, upgrade Encyclopedia volumes by installing Actuate 11 in the same directory, overwriting Actuate 10.  
Uninstall BIRT iServer Release 11 from your production staging area and reinstall it in place of your existing Actuate 10 installation. The installer upgrades the Encyclopedia volumes when it starts.
- If you are running both Actuate 10 and Actuate 11, you can migrate volumes to the Actuate 11 location.  
Use the Actuate 11 Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate an Actuate 10 Encyclopedia volume to the Actuate 11 Encyclopedia location. Confirm that all partitions for Actuate 11 are in the active state before you use these utilities.

- Install upgraded design and report files.  
Encyclopedia volume data is located separately from iServer binaries in Release 11.
- Start BIRT iServer Release 11.
- Inform users that they can start using Actuate 11 design tool products.



# 3

## **Installing BIRT iServer using an alternative database**

This chapter discusses the following topics:

- Preparing to install BIRT iServer using an alternative database
- Installing an Encyclopedia volume that uses an alternative database

---

## Preparing to install BIRT iServer using an alternative database

When installing BIRT iServer Release 11, the administrator must choose to use the out-of-the-box (OOTB) PostgreSQL database or another data store, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance to store Encyclopedia volume metadata. This chapter describes how to install a new instance of BIRT iServer Release 11, using an alternative data store.

For all database systems other than OOTB PostgreSQL, the database administrator must create the system and Encyclopedia volume schemas and an iServer application user before installing BIRT iServer. During the iServer installation, the administrator provides the iServer system name, plus the system and Encyclopedia volume schema owner, and iServer application user credentials. The iServer installation program creates the necessary database structures, then loads the metadata.

### Creating a dedicated user account for installing and running BIRT iServer

Actuate recommends creating a dedicated user account for installing and running iServer. Having a dedicated user account isolates iServer-specific issues and events on a machine, making it easier to administer the environment.

If you exercise the same control over the user account for BIRT iServer as your site exercises for other system administrator and root accounts, you can maintain the same level of security for BIRT iServer. Actuate does not recommend installing iServer under the root account since the PostgreSQL server must be started and maintained under an unprivileged user ID to prevent compromising system security. If installed under the root account, the default installation is unable to set up the PostgreSQL schemas and Actuate Encyclopedia sample volume.

### Creating the system and Encyclopedia volume schemas and iserver user in an alternative database

Before installing BIRT iServer to use a pre-existing RDBMS, the database administrator must first run SQL scripts that contain the appropriate Data Definition Language (DDL) statements to create a database and the following schema owner and application user accounts with appropriate privileges:

- iServer system schema owner
- Encyclopedia volume schema owner
- iServer application user

Restrict schema and the iServer application user names to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z0-9]\*. Do not use a hyphen.

In an environment containing multiple Encyclopedia volume schemas, Actuate recommends using one iServer application user with privileges on all the schemas. This configuration allows iServer to maximize connection pooling and minimize the number of connections to the RDBMS.

## **Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing PostgreSQL database**

The following SQL scripts provide an example of DDL statements that create the database, schema owners, and an iServer application user role, then grant privileges in a pre-existing PostgreSQL server installation. These steps are not necessary when adding an Encyclopedia volume to an existing schema.

The PostgreSQL database administrator may need to modify these SQL command examples for the specific PostgreSQL installation. In the commands, substitute system and schema names appropriate to your environment.

### **Creating a database**

Connect to the PostgreSQL system database as a user with full administrator privileges, typically named postgres, and execute the following SQL commands to create a database named iserver:

```
CREATE DATABASE iserver
  WITH OWNER = "postgres"
  TEMPLATE = template0 ENCODING = 'UTF-8';
REVOKE ALL ON DATABASE iserver FROM PUBLIC;
```

In the iserver database, create the plpgsql procedural language by executing the following SQL command:

```
CREATE LANGUAGE plpgsql;
```

Plpgsql is a superset of PostgreSQL SQL that supports advanced programming features, such as variables, conditional expressions, iterative constructs, and events. If the language is already installed, an error message appears. If so, ignore the message.

### **Creating the system schema owner**

In an iServer installation, the system schema owner must have the same name as the system schema. The system schema owner has all privileges on the schema used for the system data store and can grant privileges to other users. The system schema owner must be able to create database objects, such as tables and indexes.

The following commands create a user role named ac\_corp\_system with appropriate privileges to connect to the previously created iserver database.

Connect to the PostgreSQL system database as a user with full administrator privileges and execute the following SQL commands:

```
CREATE ROLE ac_corp_system LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO ac_corp_system;
```

## Creating the Encyclopedia volume schema owner

In an iServer installation, the Encyclopedia volume schema owner must have the same name as the Encyclopedia volume schema. The Encyclopedia volume schema owner has all privileges on the schema used for the Encyclopedia volume data store and can grant privileges to other users. The Encyclopedia volume schema owner must be able to create database objects, such as tables and indexes.

The following commands create a user role named `ac_corp` with appropriate privileges to connect to the previously created `iserver` database. Connect to the PostgreSQL system database as a user with full administrator privileges and execute the following SQL commands:

```
CREATE ROLE ac_corp LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO ac_corp;
```

## Creating the iServer application user

iServer connects to the database as an application user. The application user requires only the privileges necessary to perform basic SQL Data Manipulation Language (DML) operations, such as `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. This user does not require privileges to create or modify the structure of the database.

The following SQL script provides an example of DDL statements that create the `iserver` user role in a pre-existing PostgreSQL database. Connect to the PostgreSQL system database as a user with full administrator privileges and execute the following SQL commands:

```
CREATE ROLE iserver LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO iserver;
```

## Creating the system schema

The system schema must have the same name as the system schema owner. The following commands create a system schema named `ac_corp_system`, owned by the user, `ac_corp_system`, then grant privileges to use that schema to the application user role named `iserver`. Connect to the `iserver` application database, not the PostgreSQL system database, as a user with full administrator privileges and execute the following commands:

```
CREATE SCHEMA ac_corp_system AUTHORIZATION ac_corp_system;
GRANT USAGE ON SCHEMA ac_corp_system TO iserver;
```

## Creating the Encyclopedia schema

In an iServer installation, the Encyclopedia schema must have the same name as the Encyclopedia schema owner. The following commands create an Encyclopedia volume schema named `ac_corp`, owned by the user, `ac_corp`, then grant privileges to use the schema to the application user role named `iserver`. Connect to the `iserver` application database, not the PostgreSQL system database, as a user with full administrator privileges and execute the following commands:

```
CREATE SCHEMA ac_corp AUTHORIZATION ac_corp;  
GRANT USAGE ON SCHEMA ac_corp TO iserver;
```

## Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing Oracle database

In Oracle, there is a one-to-one relationship between a user and a schema. A schema is not a separate entity. The iServer system schema owner has the same name as the system schema. The Encyclopedia volume schema owner also has the same name as the Encyclopedia volume schema.

The following SQL scripts provide an example of DDL statements that create the database, schema owners, and iServer application user, then grant privileges in a pre-existing Oracle database. These steps are not necessary when adding an Encyclopedia volume to an existing schema.

The Oracle database administrator may need to modify these SQL command examples for a specific Oracle installation. In the commands, substitute system and volume schema names appropriate to your environment.

## Creating the system schema owner

The iServer system schema owner has all privileges on the schema used for the system data store and can grant privileges to other users. The system schema owner must be able to create database objects, such as tables and indexes.

The following SQL script provides an example of DDL statements that create the iServer system schema owner and grant privileges in a pre-existing Oracle database:

```
DROP USER ac_corp_system CASCADE;  
CREATE USER ac_corp_system  
  IDENTIFIED BY password  
  DEFAULT TABLESPACE USERS  
  TEMPORARY TABLESPACE TEMP;  
GRANT CREATE TABLE TO ac_corp_system;  
GRANT CREATE VIEW TO ac_corp_system;  
GRANT CREATE SEQUENCE TO ac_corp_system;  
GRANT CREATE TYPE TO ac_corp_system;  
GRANT CREATE PROCEDURE TO ac_corp_system;  
GRANT CREATE OPERATOR TO ac_corp_system;  
GRANT CREATE TRIGGER TO ac_corp_system;
```

```
GRANT CREATE SESSION TO ac_corp_system;
ALTER USER ac_corp_system QUOTA UNLIMITED ON USERS;
COMMIT;
```

## Creating the Encyclopedia volume schema owner

The Encyclopedia volume schema owner has all privileges on the schema used for the volume data store and can grant privileges to other users. The Encyclopedia volume schema owner must be able to create database objects, such as tables and indexes.

The following SQL script provides an example of DDL statements that create the Encyclopedia volume schema owner and grant privileges in a pre-existing Oracle database:

```
DROP USER ac_corp CASCADE;
CREATE USER ac_corp
  IDENTIFIED BY password
  DEFAULT TABLESPACE USERS
  TEMPORARY TABLESPACE TEMP;
GRANT CREATE TABLE TO ac_corp;
GRANT CREATE VIEW TO ac_corp;
GRANT CREATE SEQUENCE TO ac_corp;
GRANT CREATE TYPE TO ac_corp;
GRANT CREATE PROCEDURE TO ac_corp;
GRANT CREATE OPERATOR TO ac_corp;
GRANT CREATE TRIGGER TO ac_corp;
GRANT CREATE SESSION TO ac_corp;
ALTER USER ac_corp QUOTA UNLIMITED ON USERS;
COMMIT;
```

## Creating the iServer application user

iServer connects to the database as an application user. The application user requires only the privileges necessary to perform basic SQL Data Manipulation Language (DML) operations, such as SELECT, INSERT, UPDATE, and DELETE. This user does not require privileges to create or modify the structure of the database.

The iServer installation process automatically grants the schema privileges required by the application user. The RDBMS database administrator does not have to configure these privileges manually.

The following SQL script provides an example of DDL statements that create the iserver user in a pre-existing Oracle database:

```
DROP USER iserver CASCADE;
CREATE USER iserver
  IDENTIFIED BY password
  DEFAULT TABLESPACE USERS
  TEMPORARY TABLESPACE TEMP;
```

```
GRANT CREATE SESSION TO iserver;  
ALTER USER iserver QUOTA UNLIMITED ON USERS;  
COMMIT;
```

## **Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing SQL Server database**

The following Transact-SQL scripts provide an example of DDL statements that create the database, schema owners, and iServer application user, then grant privileges in a pre-existing SQL Server database. These steps are not necessary when adding an Encyclopedia volume to an existing schema.

The SQL Server database administrator may need to modify these SQL command examples for the specific SQL Server installation. In the commands, substitute system and schema names appropriate to your environment.

### **Creating a database**

Connect to the SQL Server master database as a user with full system administrator, sysadmin, privileges and execute the following Transact-SQL commands to create a database named iserver:

```
USE master;  
GO  
CREATE DATABASE iserver  
    COLLATE SQL_Latin1_General_CP1_CI_AS;  
GO
```

Any database created for iServer processing must use a case-insensitive collation, such as SQL\_Latin1\_General\_CP1\_CI\_AS. The names of case-insensitive collations typically include the letters, CI.

### **Creating the system schema owner**

In an iServer installation, the system schema owner must have the same name as the system schema. The system schema owner has all privileges on the schema used for the system data store and can grant privileges to other users. The system schema owner must be able to create database objects, such as tables and indexes.

The following commands create a user named ac\_corp\_system to function as the system schema owner with appropriate privileges to connect to the previously created iserver database. Connect to the iserver database as a user with full administrator privileges and execute the following SQL commands:

```
USE iserver;  
CREATE LOGIN ac_corp_system  
    WITH PASSWORD = 'password';  
CREATE USER ac_corp_system  
    FOR LOGIN ac_corp_system  
    WITH DEFAULT_SCHEMA = ac_corp_system;
```

```

GO
GRANT CONNECT TO ac_corp_system;
GRANT CREATE TABLE TO ac_corp_system;
GRANT CREATE VIEW TO ac_corp_system;
GRANT CREATE FUNCTION TO ac_corp_system;
GRANT CREATE PROCEDURE TO ac_corp_system;
GO

```

In the SQL Server environment, the default schema does not have to exist when creating the user. The system administrator can create the schema later.

### **Creating the Encyclopedia volume schema owner**

In an iServer installation, the Encyclopedia volume schema owner must have the same name as the Encyclopedia volume schema. The Encyclopedia volume schema owner has all privileges on the schema used for the Encyclopedia data store and can grant privileges to other users. The Encyclopedia schema owner must be able to create database objects, such as tables and indexes.

The following commands create an Encyclopedia volume schema owner named `ac_corp` with appropriate privileges to connect to the previously created iServer database. Connect to the iServer database as a user with full administrator privileges and execute the following SQL commands:

```

USE iServer;
CREATE LOGIN ac_corp
    WITH PASSWORD = 'password';
CREATE USER ac_corp
    FOR LOGIN ac_corp
    WITH DEFAULT_SCHEMA = ac_corp;
GO
GRANT CONNECT TO ac_corp;
GRANT CREATE TABLE TO ac_corp;
GRANT CREATE VIEW TO ac_corp;
GRANT CREATE FUNCTION TO ac_corp;
GRANT CREATE PROCEDURE TO ac_corp;
GO

```

In the SQL Server environment, the default schema does not have to exist when creating the user. The system administrator can create the schema later.

### **Creating the iServer application user**

iServer connects to the database as an application user. The application user requires only the privileges necessary to perform basic SQL Data Manipulation Language (DML) operations, such as `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. This user does not require privileges to create or modify the structure of the database.



The following SQL script provides an example of DDL statements that create the `iserver` user in a pre-existing SQL Server database. Connect to the `iserver` database as a user with full administrator privileges and execute the following SQL commands:

```
USE iserver;
CREATE LOGIN iserver WITH PASSWORD = 'password';
CREATE USER iserver FOR LOGIN iserver;
GO
GRANT CONNECT TO iserver;
GO
```

### **Creating the system schema**

The system schema must have the same name as the system schema owner. The following commands create a system schema named `ac_corp_system` and grant ownership to the user named `ac_corp_system`. Connect to the `iserver` application database, not the SQL Server master database, as a user with full administrator privileges and execute the following commands:

```
USE iserver;
GO
CREATE SCHEMA ac_corp_system AUTHORIZATION ac_corp_system;
GO
```

### **Creating the Encyclopedia volume schema**

In an `iServer` installation, the Encyclopedia volume schema must have the same name as the Encyclopedia volume schema owner. The following commands create an Encyclopedia volume schema named `ac_corp` and grant ownership to the user named `ac_corp`. Connect to the `iserver` application database, not the SQL Server master database, as a user with full administrator privileges and execute the following commands:

```
USE iserver;
GO
CREATE SCHEMA ac_corp AUTHORIZATION ac_corp;
GO
```

### **Creating the system and Encyclopedia volume schemas and `iserver` user in a pre-existing DB2 database**

The following SQL scripts provide an example of DDL statements that create the database, schema owners, and `iServer` application user, then grant privileges in a pre-existing DB2 database. These steps are not necessary when adding an Encyclopedia volume to an existing schema.

The DB2 database administrator (DBA) may need to modify these SQL command examples for the specific DB2 installation. In the commands, substitute system and schema names appropriate to your environment.

## Creating user accounts

DB2 uses operating system accounts instead of internally defined database users. A database user must exist as an operating system user account, using the native security mechanisms that the operating system provides, before a user can be referenced in a DB2 system. Once a user exists in the operating system, the DB2 system administrator can assign privileges to that user using DDL statements.

## Creating a database

Actuate requires a DB2 database to support VARGRAPHIC columns. DB2 does not support UCS-2 or UTF-16 as the primary encoding for a database. DB2 also sizes the VARCHAR data type in bytes, not characters. To work around these issues, iServer uses VARGRAPHIC instead of VARCHAR. VARGRAPHIC stores UTF-16 data and sizes this data in characters.

Actuate also requires a DB2 database to use a case-insensitive collation, such as UCA500R1\_LEN\_S2, which is not the default. DB2 supports this functionality only in DB2 9.5 Fix Pack 1 and later versions.

To create the iserver database, connect to the DB2 system as a user with full administrator privileges and execute the following SQL command:

```
CREATE DATABASE iserver
  AUTOMATIC STORAGE YES
  USING CODESET UTF-8
  TERRITORY US
  COLLATE USING UCA500R1_LEN_S2
  PAGESIZE 8192
```

## Creating the system schema owner

The iServer system schema owner has all privileges on the schema used for the system data store and can grant privileges to other users. The system schema owner must be able to create database objects, such as tables and indexes.

To create the iServer system schema owner and grant privileges in a pre-existing DB2 database, execute the following command:

```
GRANT CONNECT, LOAD ON DATABASE TO USER ac_corp_system;
```

## Creating the Encyclopedia volume schema owner

The Encyclopedia volume schema owner has all privileges on the schema used for the volume data store and can grant privileges to other users. The Encyclopedia schema owner must be able to create database objects, such as tables and indexes.

To create the Encyclopedia volume schema owner and grant privileges in a pre-existing DB2 database, execute the following command:

```
GRANT CONNECT, LOAD ON DATABASE TO USER ac_corp;
```

## Creating the iServer application user

iServer connects to the database as an application user. The application user requires only the privileges necessary to perform basic SQL Data Manipulation Language (DML) operations, such as SELECT, INSERT, UPDATE, and DELETE. This user does not require privileges to create or modify the structure of the database.

The iServer installation process automatically grants the schema privileges required by the application user. The RDBMS database administrator does not have to configure these privileges manually.

To create the iserver user in a pre-existing DB2 database, execute the following command:

```
GRANT CONNECT ON DATABASE TO USER iserver;
```

## Creating the system schema

The system schema must have the same name as the system schema owner. To create a system schema named ac\_corp\_system and grant ownership to the user named ac\_corp\_system, execute the following command:

```
CREATE SCHEMA ac_corp_system AUTHORIZATION ac_corp_system;
```

## Creating the Encyclopedia volume schema

In an iServer installation, the Encyclopedia volume schema must have the same name as the Encyclopedia volume schema owner. To create an Encyclopedia volume schema named ac\_corp and grant ownership to the user named ac\_corp, execute the following command:

```
CREATE SCHEMA ac_corp AUTHORIZATION ac_corp;
```

## Adding support for the digit wildcard character in iServer when the metadata database is DB2

In a stand-alone iServer installation that uses DB2 to store Encyclopedia volume metadata, the pound sign ('#') in iServer is treated as a single alphanumeric wildcard character instead of a digit wildcard character. This limitation exists because DB2 does not support regular expressions in SQL syntax.

The administrator can add regular expression testing capability to DB2 by creating a User Defined Function, or UDF. The following article documents this approach:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-1011db2luwpatternmatch/index.html>

The solution consists of the following parts:

- Create a Java implementation that performs the regular-expression testing functionality. In the article, IBM provides the sample java code for such an

implementation. See the Implementation section in the article. IBM also provides the pre-built java package, `db2_regex.jar`, for download. See the Download section in the article.

- Create a UDF to access the external Java method. Use the following sample implementation as an example:

```
CREATE OR REPLACE FUNCTION REGEXP_LIKE(SOURCE VARCHAR(3000) ,
    REGEX VARCHAR(512) ,
    MODE VARCHAR(3))
RETURNS INTEGER
FENCED
NOT DETERMINISTIC
NO SQL
LANGUAGE JAVA
PARAMETER STYLE JAVA
EXTERNAL NAME
    'db2_regex:com.ibm.avalanche.udf.regex.Regexp.regexpLike'
NO EXTERNAL ACTION
@
```

You can create this method in a schema that the Encyclopedia volume schema owner accesses.

After creating the UDF, grant execute privileges to the volume schema user as well as to the iServer application user by executing the following DDL statements:

```
GRANT EXECUTE ON FUNCTION REGEXP_LIKE( VARCHAR(3000) ,
    VARCHAR(512) ,
    VARCHAR(3) )
TO USER /* volume schema user */ WITH GRANT OPTION

GRANT EXECUTE ON FUNCTION REGEXP_LIKE( VARCHAR(3000) ,
    VARCHAR(512) ,
    VARCHAR(3) )
TO USER /* iServer application user */
```

To support the encyclopedia engine smart search capability to use this regular expression UDF, update the `DB2.xml` file in the `AC_SERVER_HOME\etc\DataStores\DatabaseConfiguration` folder. Change the `MatchOpMapper` section from the following:

```
<MatchOpMapper SingleMatch="_"
    GreedyMatch="%"
    DigitMatch="_"
    EscapeTemplate="@ $"
    AdditionalSpecialChars="@ ">
</FunctionMappings>
```

```

        <FunctionMapping FunctionName="MATCH">
            $P0 LIKE $P1 ESCAPE '@'
        </FunctionMapping>
    </FunctionMappings>
</MatchOpMapper>

```

to the following:

```

<MatchOpMapper SingleMatch="."
    GreedyMatch=".*"
    DigitMatch="[0-9]"
    AdditionalSpecialChars="\^.$|() []*+?{},">
    <FunctionMappings>
        <FunctionMapping FunctionName="MATCH">
            REGEXP_LIKE
            ( $P0 , '^' || $P1 || '$$', 'c' ) > 0
        </FunctionMapping>
    </FunctionMappings>
</MatchOpMapper>

```

Note that you may need to prefix the UDF REGEXP\_LIKE with the schema name, depending on where the function is located.

## Filtering for a user name returns an empty result if name contains certain special characters

In Management Console—Users, if iServer uses a PostgreSQL database to store metadata, filtering on a user name returns an empty result when certain special characters are used in the filter string. For example, if you filter on a user name containing one or more of the letters, é, à, è, ü, ä, ö, ê, or ô, iServer does not find the name, using following default locale collation and type settings:

```

CREATE DATABASE iserver
    WITH OWNER = postgres
    ENCODING = 'UTF8'
    LC_COLLATE = 'C'
    LC_CTYPE = 'C'
    CONNECTION LIMIT = -1;

```

To resolve this problem, perform the following tasks:

- 1 Export the iserver database to a dump file.
- 2 Create a new database named iserver. If you are running iServer on a Windows machine, execute the following DDL statements to recreate the database:

```

CREATE DATABASE iserver
    WITH OWNER = "postgres"
    TEMPLATE = template0 ENCODING = 'UTF-8'
    LC_COLLATE = 'English, United States'

```

```
LC_CTYPE = 'English, United States'  
CONNECTION LIMIT = -1;
```

If you are running iServer on a Linux or UNIX machine, execute the following DDL statements to recreate the database:

```
CREATE DATABASE iserver  
WITH OWNER = "postgres"  
TEMPLATE = template0 ENCODING = 'UTF-8'  
TABLESPACE = pg_default  
LC_COLLATE = 'en_US.UTF8'  
LC_CTYPE = 'en_US.UTF8'  
CONNECTION LIMIT = -1;
```

3 Import the data back into the newly created iserver database.

## Backing up iServer system and Encyclopedia volume metadata

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the schemas using the tools and resources of the third-party database system.

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult Actuate Support at the time of installation if you have any questions about these backup procedures to protect against the possibility of catastrophic failure. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

When installing BIRT iServer, be sure to run the same versions of all products. Upgrade all products at the same time to maintain consistency in the versions you run.

If you are a purchasing customer, you can download iServer from an Actuate FTP software distribution site. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

<http://www.birt-exchange.com>

For information about the new Release 11 BIRT iServer System architecture, see Chapter 1, “Understanding Actuate BIRT iServer architecture,” earlier in this book

The following sections describe how to install BIRT iServer Release 11 using Oracle as an alternative data store.

---

## Installing an Encyclopedia volume that uses an alternative database

The following procedures use a pre-existing Oracle database and schema as the example. During the iServer installation, the administrator provides the following installation details and any related credentials:

- External Oracle database host name or IP address, such as urup
- iServer Encyclopedia volume name, which by default is the machine name, in this case, urup
- Oracle database iServer system schema name, such as ac\_corp\_system
- Oracle database Encyclopedia volume schema name, such as ac\_corp
- Oracle database iServer application user name, such as iserver

The iServer installation program creates the necessary volume database structures, then loads the metadata.

### How to install an Encyclopedia volume that uses an alternative database

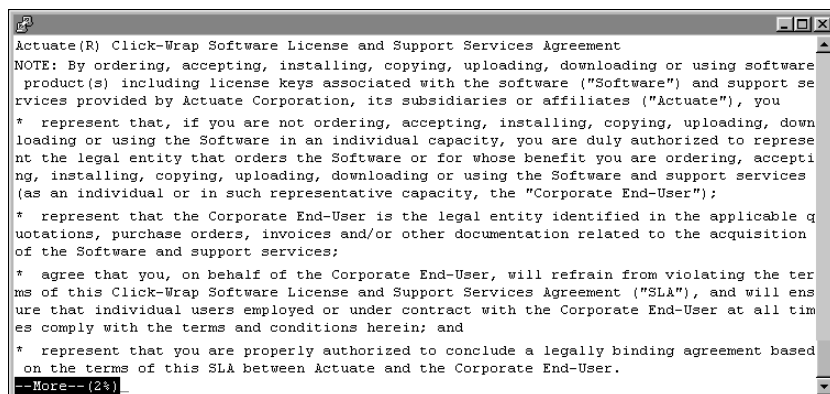
To install iServer, perform the following steps:

- 1 Download the required files from the FTP software distribution site. Extract the files.
- 2 To install the server files, execute the following isinstall script:

```
sh ./isinstall.sh
```


The script displays a number of prompts. Respond to the prompts as described in the following procedure.

- 3 The license agreement appears, as shown in Figure 3-1.



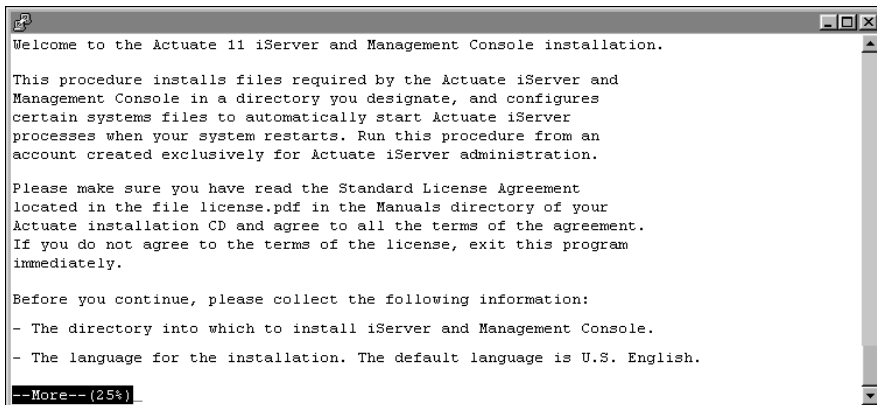
**Figure 3-1** Reviewing the license agreement

- 4 Read the license agreement and press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 3-2.



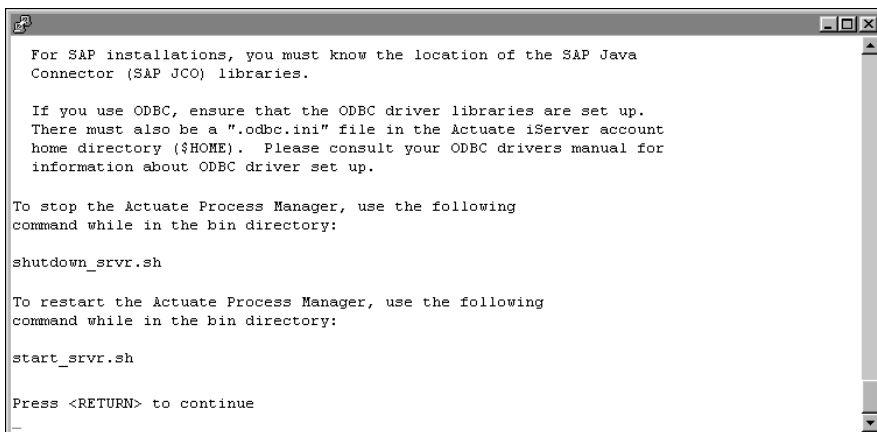
**Figure 3-2** Accepting the licensing terms

- 5 The introduction to the installation appears, as shown in Figure 3-3.



**Figure 3-3** Reviewing the introductory information

- 6 Press Enter after reviewing the introductory information, as shown in Figure 3-4.



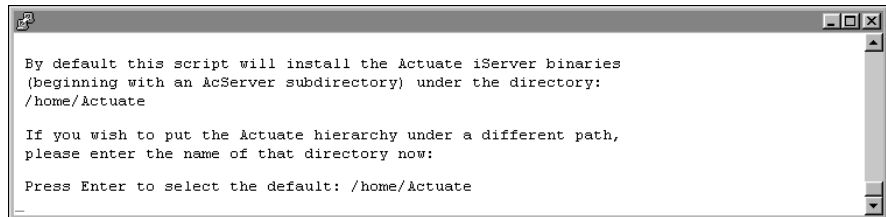
**Figure 3-4** Finishing the review of introductory information



- 7 Press Enter to accept the default location for the installation, as shown in Figure 3-5. Alternatively, type a different directory and press Enter.

The installation program creates the AcServer directory in the chosen location and installs the files.

iServer uses this location to resolve paths to all the binaries that it launches. The default path for this location is `$HOME/AcServer`, which is referred to in the iServer documentation by the environment variable, `AC_SERVER_HOME`.

A terminal window with a title bar containing a question mark icon and standard window controls (minimize, maximize, close). The text inside the terminal reads: "By default this script will install the Actuate iServer binaries (beginning with an AcServer subdirectory) under the directory: /home/Actuate. If you wish to put the Actuate hierarchy under a different path, please enter the name of that directory now: Press Enter to select the default: /home/Actuate".

```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

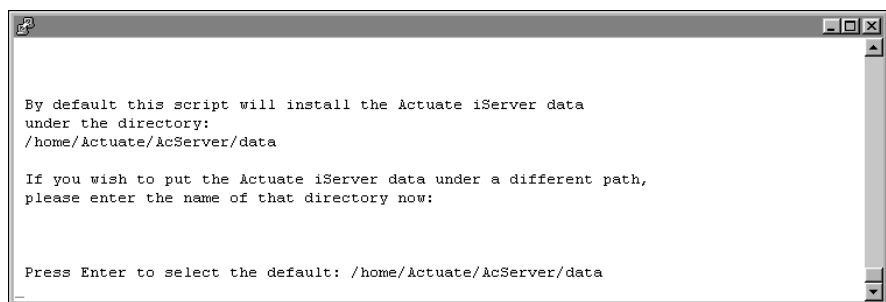
If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
```

**Figure 3-5** Specifying the installation directory

- 8 Press Enter to accept the default installation directory, `AC_SERVER_HOME /data`, for iServer data, as shown in Figure 3-6. Alternatively, choose a different directory for iServer data.

iServer uses this data location to store the iServer Encyclopedia volume data, logs, and other files. The default path is `AC_SERVER_HOME/data`, which is referred to in the iServer documentation by the environment variable `AC_DATA_HOME`.

A terminal window with a title bar containing a question mark icon and standard window controls (minimize, maximize, close). The text inside the terminal reads: "By default this script will install the Actuate iServer data under the directory: /home/Actuate/AcServer/data. If you wish to put the Actuate iServer data under a different path, please enter the name of that directory now: Press Enter to select the default: /home/Actuate/AcServer/data".

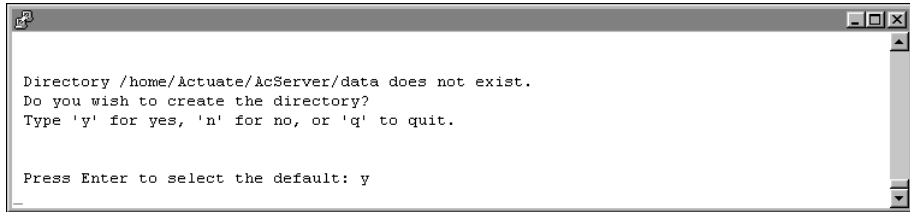
```
By default this script will install the Actuate iServer data
under the directory:
/home/Actuate/AcServer/data

If you wish to put the Actuate iServer data under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate/AcServer/data
```

**Figure 3-6** Specifying the data installation directory

- 9 Press Enter to accept the default option of creating the directory for data, as shown in Figure 3-7. Alternatively, type n for no, or q to quit, and press Enter.

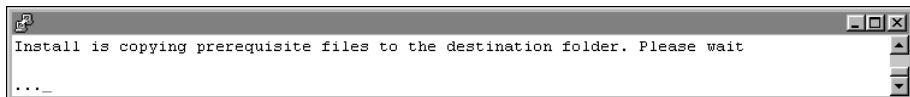


```
Directory /home/Actuate/AcServer/data does not exist.
Do you wish to create the directory?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 3-7** Creating the AC\_DATA\_HOME directory

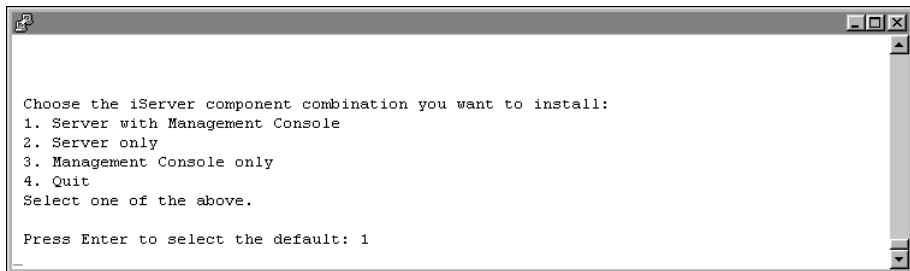
- 10 The installer copies prerequisite files to the destination folder, as shown in Figure 3-8. After copying the prerequisite files, the installation continues.



```
Install is copying prerequisite files to the destination folder. Please wait
..._
```

**Figure 3-8** Copying prerequisite files

- 11 Press Enter to accept the default iServer component combination, iServer with Management Console, as shown in Figure 3-9. Alternatively, choose a different iServer component combination and press Enter.

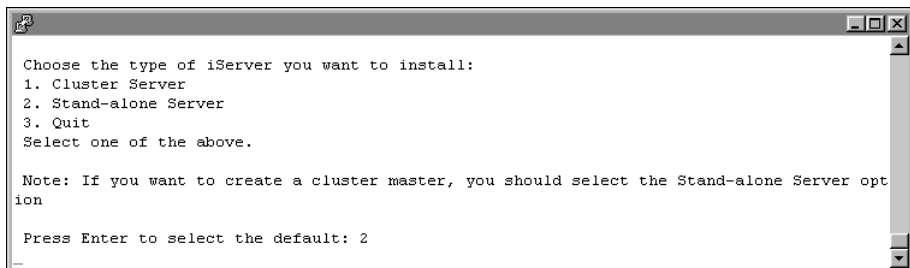


```
Choose the iServer component combination you want to install:
1. Server with Management Console
2. Server only
3. Management Console only
4. Quit
Select one of the above.

Press Enter to select the default: 1
```

**Figure 3-9** Choosing the iServer component combination

- 12 Press Enter to accept the default stand-alone iServer installation, as shown in Figure 3-10. Alternatively, choose a different type of iServer to install.



```
Choose the type of iServer you want to install:
1. Cluster Server
2. Stand-alone Server
3. Quit
Select one of the above.

Note: If you want to create a cluster master, you should select the Stand-alone Server option

Press Enter to select the default: 2
```

**Figure 3-10** Specifying the type of iServer to install


- 13 Type a name to use for the BIRT iServer System name, as shown in Figure 3-11. iServer assigns this name to the default Encyclopedia volume. Additionally,

iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.




**Figure 3-11** Specifying the BIRT iServer System name

- 14 Type n for no, and press Enter, as shown in Figure 3-12. You do not want to install the embedded PostgreSQL database if you are using an alternative database such as Oracle to store Encyclopedia volume metadata.




**Figure 3-12** Choosing to not use Embedded PostgreSQL

- 15 Type 2 to choose Oracle as the external database to work with iServer Encyclopedia and press Enter, as shown in Figure 3-13.




**Figure 3-13** Choosing an external database for iServer Encyclopedia

- 16 Specify the external Oracle database TNS server, if any, and press Enter, as shown in Figure 3-14. If there is no external Oracle database TNS Server, leave the field blank and press Enter.



**Figure 3-14** Specifying the external Oracle database TNS Server

- 17 Specify the name of the external Oracle database host name or IP address and press Enter, as shown in Figure 3-15.




**Figure 3-15** Specifying the external Oracle database server name

- 18 Specify the external Oracle database port number and press Enter, as shown in Figure 3-16. Typically, Oracle uses port 1521 as the database port.




**Figure 3-16** Specifying the external Oracle database port

- 19 Specify the external Oracle database service name, such as orcl.actuate.com, that identifies the Oracle database server on which you want to install the Encyclopedia volume metadata and press Enter, as shown in Figure 3-17.




**Figure 3-17** Specifying the external Oracle database service name

- 20 Specify the external Oracle database user name, such as iserver, and press Enter, as shown in Figure 3-18.




**Figure 3-18** Specifying the external Oracle database user name

- 21 Specify the external Oracle database user password and press Enter, as shown in Figure 3-19.




**Figure 3-19** Specifying the external Oracle database user password

- 22 Re-enter the external Oracle database user password and press Enter, as shown in Figure 3-20.




**Figure 3-20** Confirming the external Oracle database user password

- 23** Specify the system database schema for iServer, such as `ac_corp_system`, and press Enter, as shown in Figure 3-21.




**Figure 3-21** Specifying the system database schema

- 24** Specify the System database schema password, and press Enter, as shown in Figure 3-22.




**Figure 3-22** Specifying the system database schema password

- 25** Re-enter the system database schema password and press Enter, as shown in Figure 3-23.




**Figure 3-23** Confirming the system database schema password

- 26** Specify the external Oracle database schema for the Encyclopedia volume, such as `ac_corp`, and press Enter, as shown in Figure 3-24.




**Figure 3-24** Specifying the schema for the Encyclopedia volume

- 27** Specify the external Oracle database schema password, and press Enter, as shown in Figure 3-25.



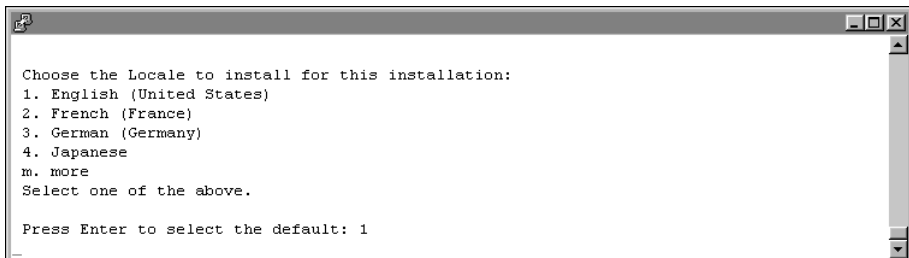
**Figure 3-25** Specifying the external Oracle database schema password

- 28** Re-enter the external Oracle database schema password and press Enter, as shown in Figure 3-26.



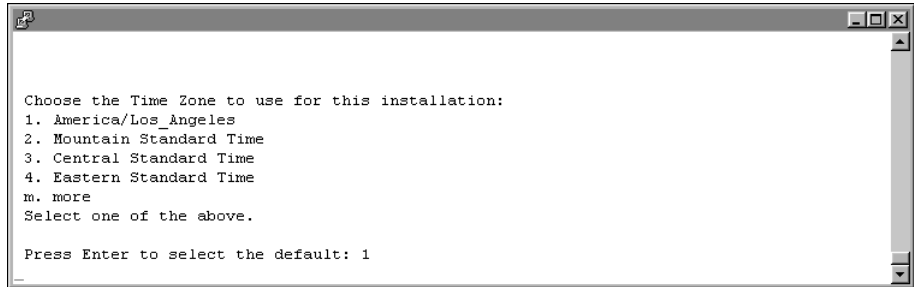
**Figure 3-26** Confirm the external Oracle database schema password

- 29** Press Enter to select the default locale, which is English (United States), as shown in Figure 3-27. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press enter.




**Figure 3-27** Specifying a locale

- 30** Press Enter to select the default time zone, America/Los\_Angeles, as shown in Figure 3-28. Alternatively, select another time zone from the numbered list.




**Figure 3-28** Specifying a time zone

- 31** To evaluate the product using the included evaluation software license, press Enter, as shown in Figure 3-29. Alternatively, type 2, then type the path to the license file you purchased.




**Figure 3-29** Specifying license type

- 32** Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 3-30. Alternatively, type a different IP address.




**Figure 3-30** Specifying the hostname that Management Console uses to contact the PMD

- 33** Press Enter to accept the default port number, 8100, where Process Management Daemon (PMD) listens for requests, as shown in Figure 3-31. Alternatively, type a different port number.




**Figure 3-31** Specifying the port number on which the PMD listens

- 34** Press Enter to accept the default host name, the name of your machine, as shown in Figure 3-32. Alternatively, type a different IP address.




**Figure 3-32** Specifying the machine on which the iServer runs

- 35** Press Enter to accept the default port number, 8000, where iServer listens for requests, as shown in Figure 3-33. Alternatively, type a different port number.




**Figure 3-33** Specifying the port number on which the iServer listens

- 36** Specify the iServer administrator password, as shown in Figure 3-34. You use this password to log in to the iServer Configuration Console.



**Figure 3-34** Specifying the iServer administrator password


- 37** Re-enter the password of the iServer administrator, as shown in Figure 3-35.



**Figure 3-35** Re-entering the iServer administrator password




- 38** Press Enter to accept the default option to use a volume name for the Encyclopedia, as shown in Figure 3-36. Alternatively, type n for no to not use a volume name for the Encyclopedia, or q to quit the installation.




**Figure 3-36** Specifying whether to use a volume name

- 39** Press Enter to accept the default Encyclopedia volume name, the name of your machine, as shown in Figure 3-37. Alternatively, type a different Encyclopedia volume name.




**Figure 3-37** Specifying the Encyclopedia volume name

- 40** Press Enter to accept the default option to start iServer automatically, as shown in Figure 3-38. Alternatively, type n for no.




**Figure 3-38** Specifying whether to start iServer automatically

- 41** Press Enter to accept the default option of not integrating LDAP with iServer, as shown in Figure 3-39. Alternatively, type n for no, or q to quit the installation.



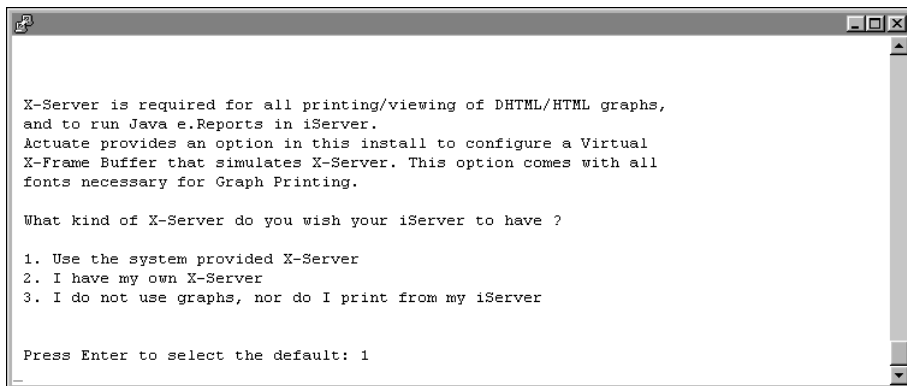
**Figure 3-39** Specifying whether to integrate LDAP with iServer

- 42** Press Enter to accept the default option to not use any database drivers/clients, as shown in Figure 3-40. Alternatively, type y for yes, and specify the database drivers/clients you want to use.




**Figure 3-40** Specifying whether to use database drivers/clients

- 43** Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 3-41.




**Figure 3-41** Specifying what kind of X-Server to use, if any

- 44** Press Enter to accept the default hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 3-42. Alternatively, type a different IP address.




**Figure 3-42** Specifying the machine host name that Management Console uses to contact the PMD

- 45** Press Enter to accept the default port number, 8100, on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 3-43. Alternatively, type a different port number.




**Figure 3-43** Specifying the port number for the PMD to listen for requests from Management Console

- 46 Press Enter to accept the default hostname, the name of your machine, that Management Console uses to contact iServer, as shown in Figure 3-44. Alternatively, type a different IP address.




**Figure 3-44** Specifying the name Management Console uses to contact iServer

- 47 Press Enter to accept the default port number, 8000, on which iServer will listen for requests from Management Console, as shown in Figure 3-45. Alternatively, type a different port number.




**Figure 3-45** Specifying the port number for listening for requests from Management Console

- 48 Press Enter to accept the default name of the Encyclopedia volume to use with Management Console, as shown in Figure 3-46. Alternatively, type a different name for the Encyclopedia volume.



**Figure 3-46** Specifying the name of the default Encyclopedia volume


- 49 Press Enter to accept the default name, `acadmin`, for the HTTP server's context root, as shown in Figure 3-47. Alternatively, type a different name.



**Figure 3-47** Specifying the name of the HTTP server context root

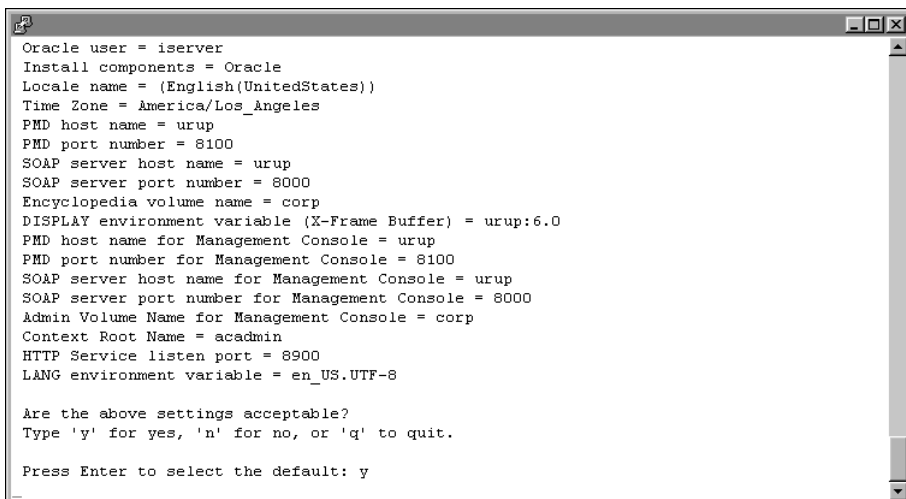
- 50 Press Enter to accept the default port number, `8900`, on which the application container listens for requests, as shown in Figure 3-48. Alternatively, choose a different port.

You connect to the port from your browser when accessing various iServer features.




**Figure 3-48** Specifying the application container listening port number

- 51 Review the settings, as shown in Figure 3-49, then specify whether to accept the settings. Press Enter to accept the default, `y` for yes. Alternatively type `n` for no, or `q` to quit.




**Figure 3-49** Reviewing your settings

- 52 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 3-50.



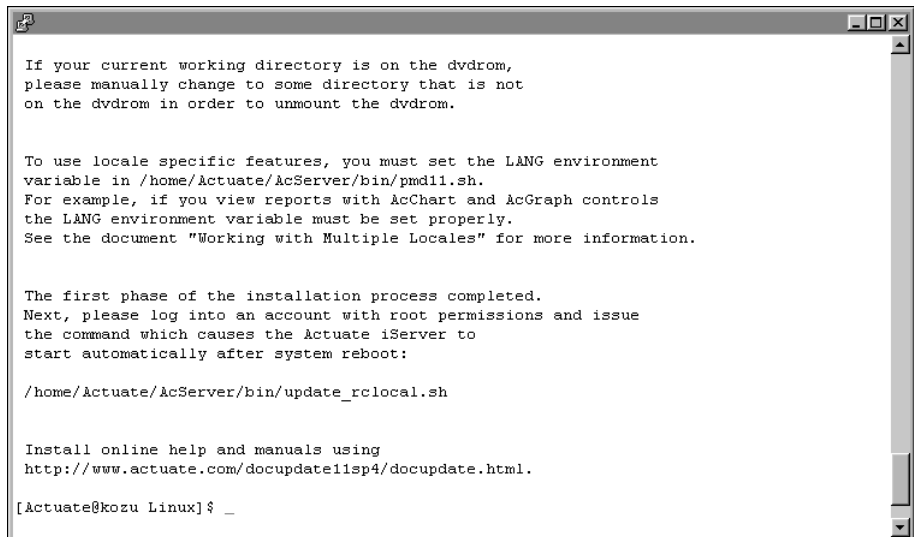
**Figure 3-50** Copying iServer files to your destination folder

- 53 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start the Process Management Daemon (PMD), as shown in Figure 3-51.



**Figure 3-51** Specifying whether to start iServer

- 54 The installation program provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 3-52.



**Figure 3-52** Specifying information about localization, logging in, and installing online help

After the installation program finishes running, open a browser to log in to the following BIRT iServer 11 consoles to perform user and administrator tasks:

- Information Console  
Perform tasks such as accessing folders and viewing designs and documents.
- Management Console  
Set up user accounts and schedule or run a design.
- Configuration Console  
Perform administrative operations, such as adding an Encyclopedia volume and making modifications to iServer parameters and server templates.

For more information on accessing BIRT iServer 11 consoles, see Chapter 2, “Installing BIRT iServer,” earlier in this book.

# Upgrading BIRT iServer

This chapter discusses the following topics:

- Upgrading BIRT iServer
- Preparing to upgrade BIRT iServer
- Performing an automatic in-place upgrade for an earlier major release
- Performing an automatic in-place upgrade for an earlier minor release
- Performing a manual side-by-side upgrade
- Working with iServer utilities

---

## Upgrading BIRT iServer

When upgrading to BIRT iServer Release 11, the administrator must choose to use the out-of-the-box (OOTB) PostgreSQL database or another data store, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance to store Encyclopedia volume metadata. The administrator upgrades iServer and the Encyclopedia volume database using one of the following options:

- Automatic in-place upgrade  
Upgrades an earlier iServer system in place for a major release, such as Release 10 Service Pack 1, or a minor release, such as Release 11 Service Pack 3, automatically migrating one or more existing Encyclopedia volumes during the installation process.
- Manual side-by-side upgrade  
Upgrades a new Release 11 iServer system in the environment side-by-side with an earlier iServer System, such as Release 10 Service Pack 1. During installation, the administrator chooses whether to install the default volume with sample designs and documents. After installation, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the volumes manually from the older system to the Release 11 iServer.

---

## Preparing to upgrade BIRT iServer

Before running the BIRT iServer upgrade program, the system administrator must prepare the environment by performing the following tasks:

- Create a dedicated user account with appropriate privileges for installing and running iServer.
- Back up iServer system and Encyclopedia volume metadata.

### **Creating a dedicated user account for installing and running BIRT iServer**

Actuate recommends creating a dedicated user account for installing and running iServer. Like other Linux and UNIX processes, the processes that perform BIRT iServer tasks run under a specific user account. Having a dedicated user account isolates iServer-specific issues and events on a machine, making it easier to administer the environment.

Before installing iServer, create the user account, making sure to configure it with privileges to access the relevant files and directories. If you exercise the same control over the user account for BIRT iServer as your site exercises for other



system administrator and root accounts, you can maintain the same level of security for BIRT iServer.

Installation of iServer under the root account is not recommended since the PostgreSQL server must be started and maintained under an unprivileged user ID to prevent compromising system security. If installed under the root account, the default installation is unable to set up the PostgreSQL schema and Actuate Encyclopedia sample volume.

## **Backing up iServer system and Encyclopedia volume metadata**

When upgrading iServer to Release 11, be sure to create a backup of the earlier Encyclopedia volume before performing the upgrade operation. In Release 11, there is no automatic rollback capability during an upgrade. In the event of a failure, the administrator must uninstall the new version of iServer, reinstall the previous version, and use the backup to restore the Encyclopedia volume to its previous state if a rollback becomes necessary.

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the Encyclopedia volume schemas using the tools and resources of the third-party database system. It is necessary to back up all Encyclopedia volume metadata and file data to ensure the recoverability of the volume in the event of failure.

In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this operation to protect critical system metadata. The administrator can restore a corrupted or missing system schema using the System Data Store Administrator utility. For more information on this utility, see “Specifying System Data Store Administrator properties,” later in this chapter.

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult Actuate Support at the time of installation if you have any questions about these backup procedures to protect against the possibility of catastrophic failure. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

When installing BIRT iServer, be sure to run the same versions of all products. Upgrade all products at the same time to maintain consistency in the versions you run.

To upgrade an iServer earlier than Release 8 Service Pack 1, you must first migrate to Release 8 Service Pack 1, 9, or 10, then upgrade to Release 11.

If you are a purchasing customer, you can download iServer from an Actuate FTP software distribution site. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

<http://www.birt-exchange.com>

For information about the new Release 11 BIRT iServer System architecture, see Chapter 1, “Understanding Actuate BIRT iServer architecture.”

The following sections describe how to install BIRT iServer Release 11 as an upgrade using the available installation options.

---

## Performing an automatic in-place upgrade for an earlier major release

In an automatic in-place upgrade from an earlier major release, such as Release 10 Service Pack 1, the installation program performs the following operations:

- Installs and initializes iServer and the OOTB PostgreSQL relational database management system (RDBMS)
- Creates the iserver user in the OOTB PostgreSQL RDBMS to access the system, initializing the iserver user password to the PostgreSQL superuser password
- Updates the iServer configuration file, specifying the volume, database, and connection information for the default volume

The automatic option also performs the following operations during installation:

- Creates the system and volume schema, initializing these schema with basic configuration information
- Creates a separate schema for each Encyclopedia volume
- Creates the iServer configuration file, specifying system, volume, and connection information for the default installation
- Migrates all Encyclopedia volumes automatically

After upgrading from an earlier major release, the administrator can use the iServer migration utilities to move the metadata content from the embedded PostgreSQL database to another supported RDBMS, such as SQL Server, Oracle, or DB2. For more information on using the iServer migration utilities, see “Working with iServer utilities,” later in this chapter.

Alternatively, iServer supports a manual, side-by-side upgrade to a new Release 11 iServer in a path separate from the earlier release. For more information, see “Performing a manual side-by-side migration,” later in this chapter.

The following section describes how to perform an automatic in-place upgrade of an earlier major BIRT iServer release, such as Release 10 Service Pack 1, to Release 11 Service Pack 4.

## Running the in-place upgrade on an earlier major release

The following procedure describes step-by-step how to perform an automatic upgrade in place of an earlier major version of iServer to BIRT Release 11.

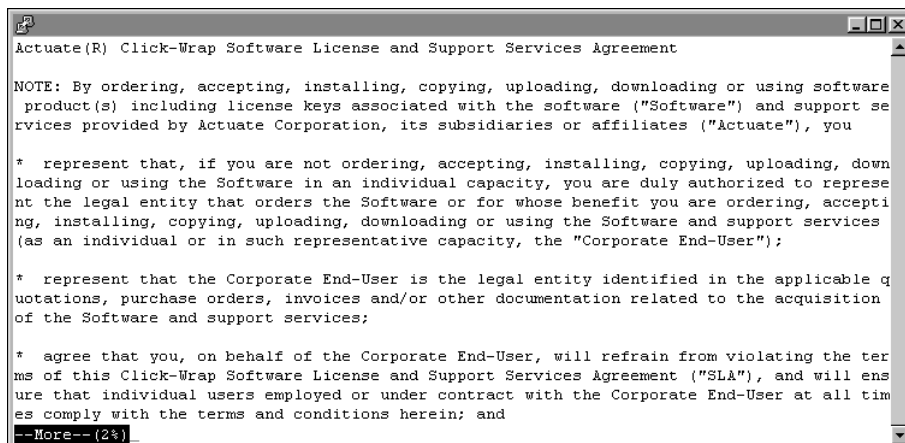
### How to run the in-place upgrade on an earlier major release

- 1 The installation program can encounter a problem over-writing a file linked with a running process. Stop all iServer processes before proceeding with the upgrade.
- 2 Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
  - encyc directories from all nodes
  - acserverconfig.xml in the /etc directory
  - acpmdconfig.xml in the /etc directory
  - RSSE code and associated files if you use the Open Security option
- 3 Download the required files from the FTP software distribution site. Extract the files.
- 4 To install the server files, execute the isinstall script:

```
sh ./isinstall.sh
```

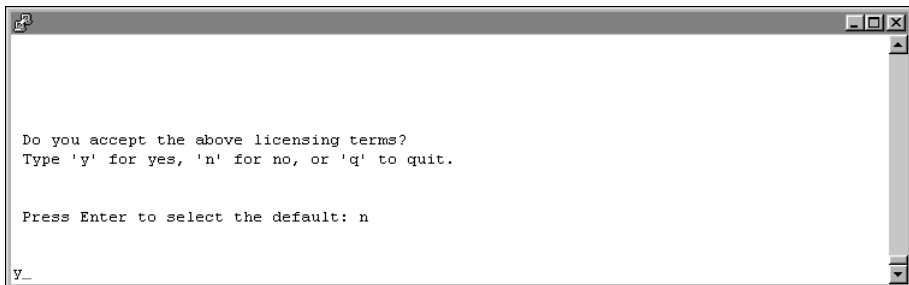
The script displays a number of prompts. Respond to the prompts as described in the following procedure.

- 5 The license agreement appears, as shown in Figure 4-1.



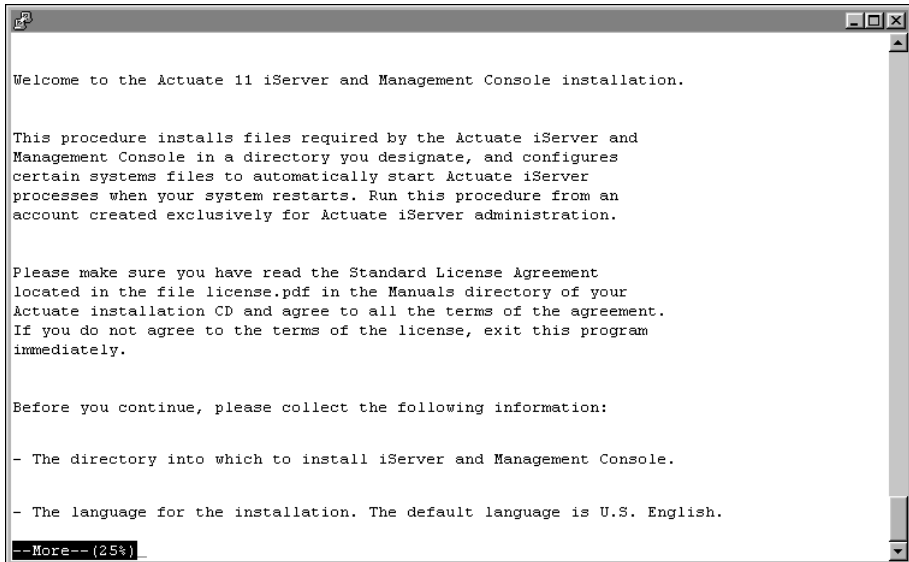
**Figure 4-1** Reviewing the license agreement

- 6 Read the license agreement, then press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 4-2.



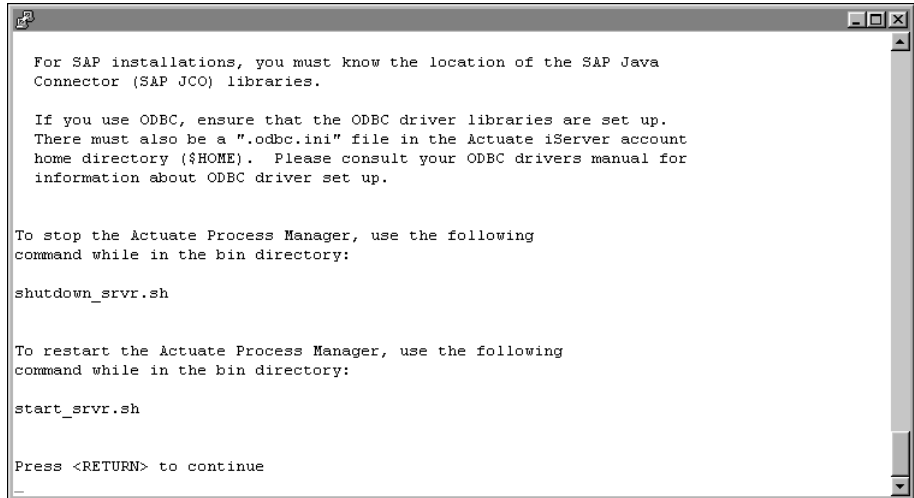
**Figure 4-2** Accepting the licensing terms

- 7 The introduction to the installation appears, as shown in Figure 4-3.



**Figure 4-3** Reviewing the introductory information

- 8 Press Enter after reviewing the introductory information, as shown in Figure 4-4.



```
For SAP installations, you must know the location of the SAP Java
Connector (SAP JCO) libraries.

If you use ODBC, ensure that the ODBC driver libraries are set up.
There must also be a ".odbc.ini" file in the Actuate iServer account
home directory ($HOME). Please consult your ODBC drivers manual for
information about ODBC driver set up.

To stop the Actuate Process Manager, use the following
command while in the bin directory:

shutdown_srvr.sh

To restart the Actuate Process Manager, use the following
command while in the bin directory:

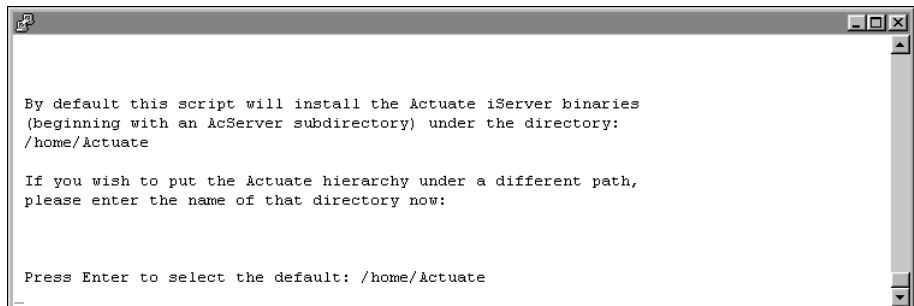
start_srvr.sh

Press <RETURN> to continue
```

**Figure 4-4** Finishing the review of introductory information

- 9 Press Enter to select the default directory as the location of the Actuate 10 installation, as shown in Figure 4-5. Alternatively, type the appropriate path to that location and press Enter.

iServer uses this location to resolve paths to all the binaries that it launches. The default path for this location is `$HOME/AcServer`, which is referred to in the iServer documentation by the environment variable `AC_SERVER_HOME`.



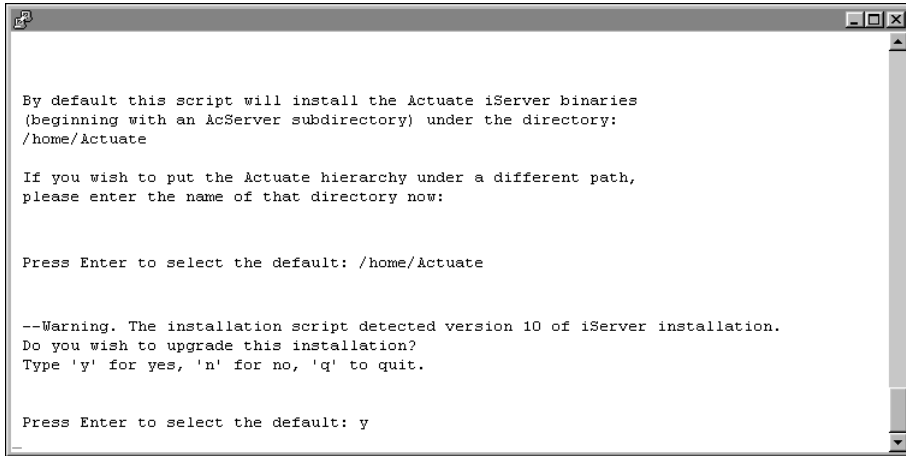
```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
```

**Figure 4-5** Specifying the installation directory

- 10 The installer detects the previous iServer version, and asks whether to perform an upgrade to the new iServer version, as shown in Figure 4-6. Press Enter to accept the default option of upgrading the older version. Alternatively, type n for no, or q to quit.
- 11 If the installer detects that the Actuate servers are currently running, it will try to stop the servers then continue with the installation, as shown in Figure 4-7. This process might take a few minutes.



```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

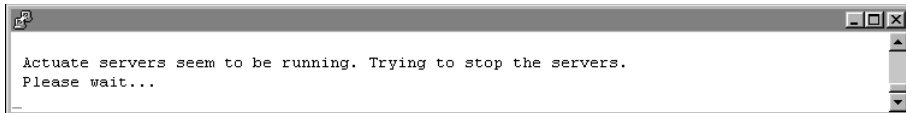
If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate

--Warning. The installation script detected version 10 of iServer installation.
Do you wish to upgrade this installation?
Type 'y' for yes, 'n' for no, 'q' to quit.

Press Enter to select the default: y
```

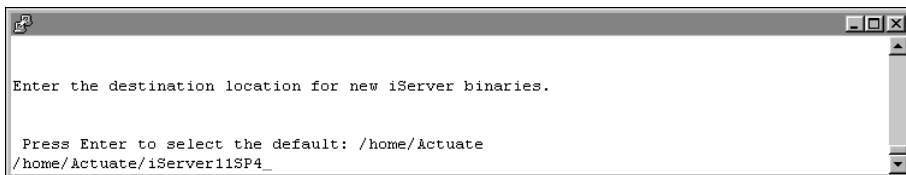
**Figure 4-6** Choosing to upgrade iServer



```
Actuate servers seem to be running. Trying to stop the servers.
Please wait...
```

**Figure 4-7** Shutting down the servers

- 12 Type the path to a new directory in which to install the iServer binaries, as shown in Figure 4-8.

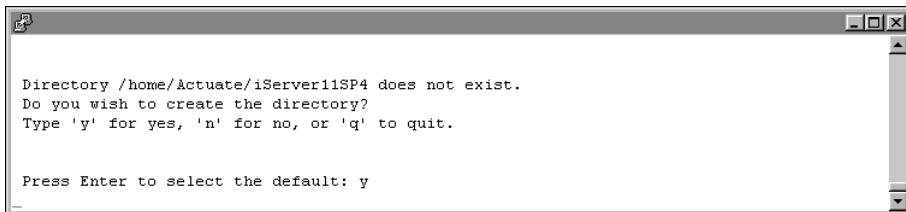


```
Enter the destination location for new iServer binaries.

Press Enter to select the default: /home/Actuate
/home/Actuate/iServer11SP4_
```

**Figure 4-8** Specifying a new location for the iServer binaries

- 13 Press Enter to accept the default option of creating the folder you specified in the previous step, as shown in Figure 4-9. Alternatively, type n for no, or q to quit, and press Enter.



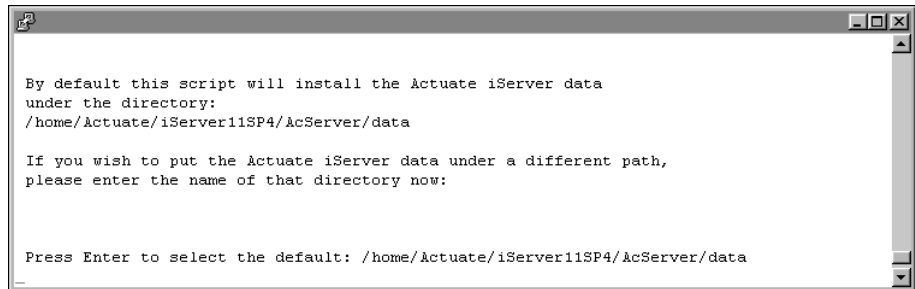
```
Directory /home/Actuate/iServer11SP4 does not exist.
Do you wish to create the directory?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 4-9** Creating the new installation directory

- 14 Press Enter to accept the default installation directory, AC\_SERVER\_HOME /data, to install iServer data, as shown in Figure 4-10. Alternatively, choose a different directory for iServer data.

iServer uses this data location to store the iServer Encyclopedia volume data, including PostgreSQL metadata, logs, and other files. The default path is AC\_SERVER\_HOME/data, which is referred to in the iServer documentation by the environment variable AC\_DATA\_HOME.



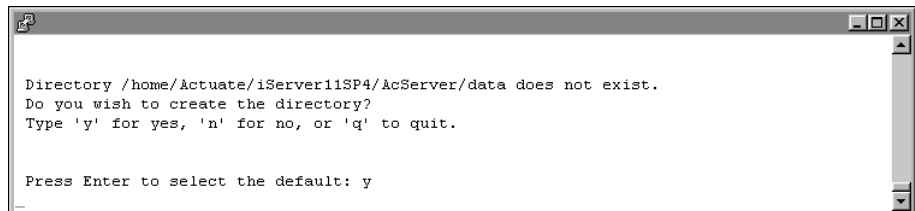
```
By default this script will install the Actuate iServer data
under the directory:
/home/Actuate/iServer11SP4/AcServer/data

If you wish to put the Actuate iServer data under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate/iServer11SP4/AcServer/data
```

**Figure 4-10** Specifying the data installation directory

- 15 The installer asks whether you want to create the directory for AC\_DATA\_HOME. Press Enter to accept the default option to create the folder, as shown in Figure 4-11. Alternatively, type n for no, or q to quit, and press Enter.

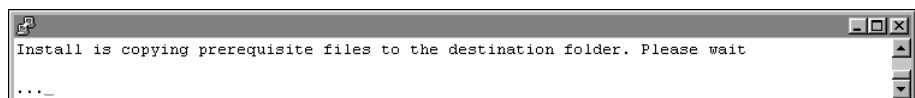


```
Directory /home/Actuate/iServer11SP4/AcServer/data does not exist.
Do you wish to create the directory?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 4-11** Creating the AC\_DATA\_HOME directory

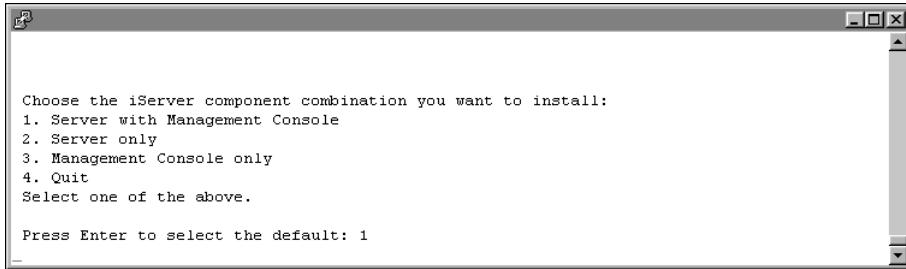
- 16 The installer copies the prerequisite files to the destination folder, as shown in Figure 4-12.



```
Install is copying prerequisite files to the destination folder. Please wait
..._
```

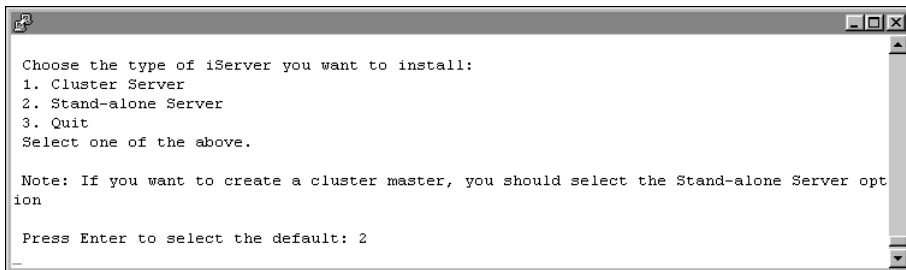
**Figure 4-12** Copying prerequisite files

- 17 Press Enter to accept the default iServer component combination, as shown in Figure 4-13. Alternatively, choose a different iServer component combination and press Enter.




**Figure 4-13** Choosing the iServer components to install

- 18** Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-14.




**Figure 4-14** Choosing the iServer installation type

- 19** Type a name to use for the BIRT iServer System name, as shown in Figure 4-15. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.



**Figure 4-15** Specifying the BIRT iServer System name


- 20** Press Enter to choose the default PostgreSQL superuser name, postgres, as shown in Figure 4-16. Alternatively, type a different PostgreSQL superuser name.



**Figure 4-16** Choosing the PostgreSQL superuser name




- 21 Type a PostgreSQL superuser password that conforms to the password security policy requirements within your organization, then press Enter, as shown in Figure 4-17.




**Figure 4-17** Typing the PostgreSQL superuser password

- 22 Re-enter the password for PostgreSQL superuser, as shown in Figure 4-18, and press Enter.




**Figure 4-18** Re-entering the PostgreSQL superuser password

- 23 Press Enter to accept the default port on which the PostgreSQL database server listens for requests, as shown in Figure 4-19. Alternatively, type a different port number.




**Figure 4-19** Entering the port number which PostgreSQL uses

- 24 Specify the iServer administrator password, as shown in Figure 4-20. You use this password to log in to the iServer Configuration Console.




**Figure 4-20** Specifying the iServer administrator password

- 25 Re-enter the password of the iServer administrator, as shown in Figure 4-21. You use this password to log in to Configuration Console.




**Figure 4-21** Re-entering the iServer administrator password

- 26** Press Enter to accept the default option to upgrade the iServer license file, as shown in Figure 4-22. Alternatively, press n for no, or q to quit.




**Figure 4-22** Accepting to upgrade the iServer license

- 27** To evaluate the product using the included evaluation software license press Enter, as shown in Figure 4-23. Alternatively, type 2, then type the path to a the license file you purchased.



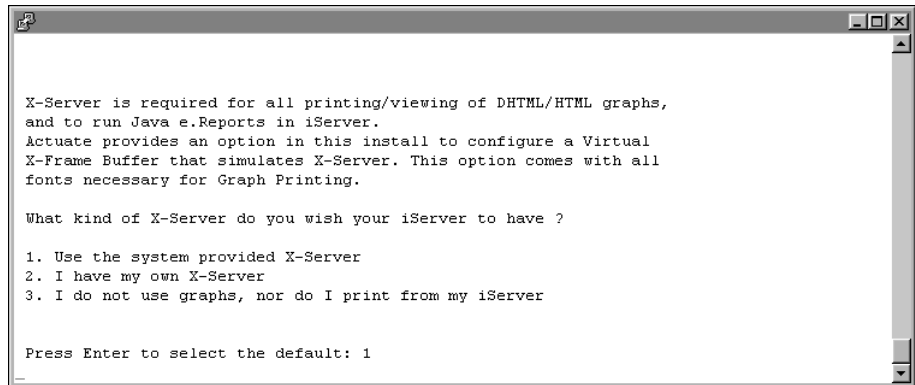
**Figure 4-23** Specifying license type

- 28** Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 4-24. Alternatively, type y for yes, specify the database drivers/clients you want to use, and press Enter.



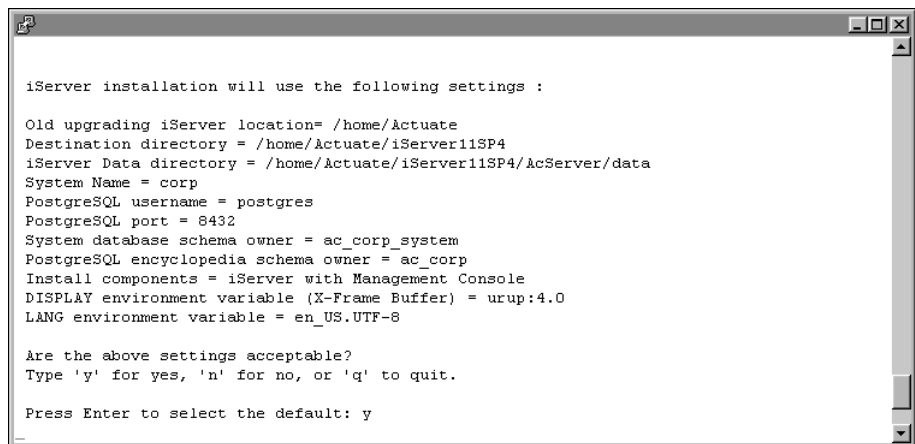
**Figure 4-24** Specifying whether to use database drivers/clients

- 29 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-25.




**Figure 4-25** Specifying what kind of X-Server to use

- 30 Review the settings, as shown in Figure 4-26, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.




**Figure 4-26** Reviewing the installation settings

- 31 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-27.



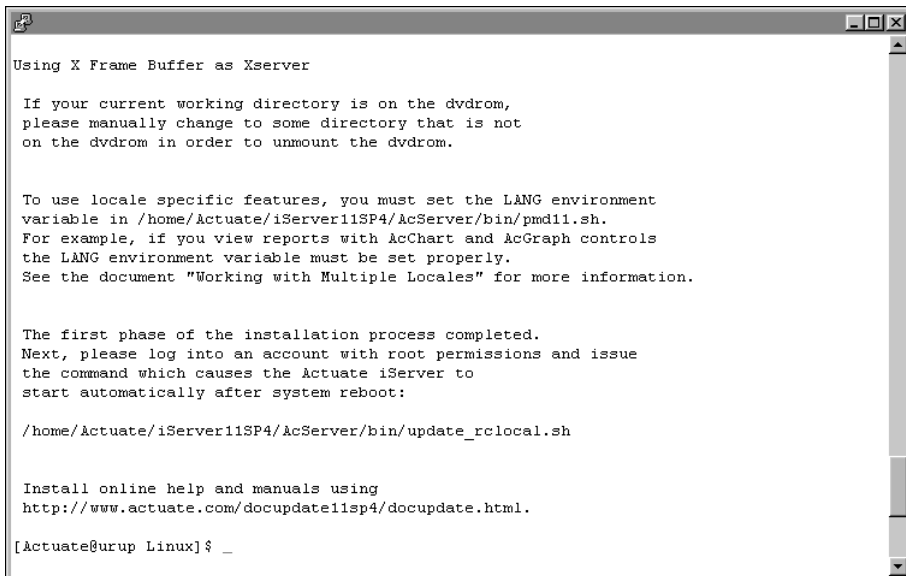
**Figure 4-27** Viewing iServer installation progress

- 32** At the end of the installation, the program asks if you want to start iServer. Accept the default, *y* for yes, to start iServer, as shown in Figure 4-28.



**Figure 4-28** Specifying whether to start iServer

- 33** When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-29.



**Figure 4-29** Viewing information about localization, logging in, and installing online help

- 34** Log in to Management Console. In Files and Folders, the data from the previous release appears.

---

## Performing an automatic in-place upgrade for an earlier minor release

When upgrading to iServer Release 11 Service Pack 4 in place from an earlier minor release, such as Service Pack 2 or 3, the administrator upgrades iServer using the automatic upgrade process. The automatic upgrade program performs the following tasks:

- Installs iServer in a new directory
- Updates the volume schema, enabling the new iServer version to work with your existing Encyclopedia volume or volumes
- Brings the volume or volumes online.

The upgrade from an earlier minor release automatically updates the metadata in the installed RDBMS. In Release 11 Service Pack 3 and 4, it is not necessary to perform a manual in-place upgrade of an Encyclopedia volume schema in any supported RDBMS after performing the automatic, wizard-based, iServer system upgrade.

## Running the in-place upgrade on an earlier minor release

The following procedure describes step-by-step how to perform an automatic upgrade in place of an earlier minor version of iServer Release 11, such as Service Pack 2 or 3.

### How to run the in-place upgrade on an earlier minor release

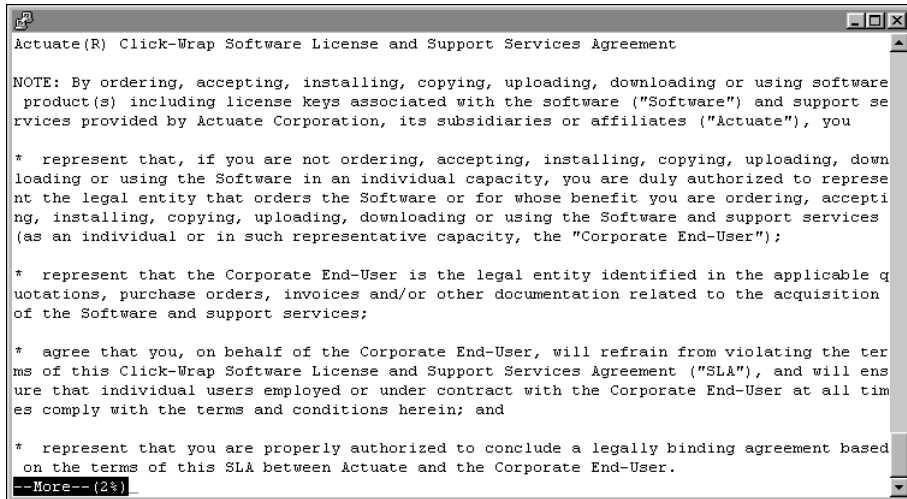
- 1 The installation program can encounter a problem over-writing a file linked with a running process. Be sure to stop all iServer processes before proceeding with the upgrade.
- 2 Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
  - encyc directories from all nodes
  - acserverconfig.xml in the /etc directory
  - acpmdconfig.xml in the /etc directory
  - RSSE code and associated files if you use the Open Security option
- 3 Download the required files from the FTP software distribution site. Extract the files.

- 4 To install the server files, execute the isinstall script:

```
sh ./isinstall.sh
```

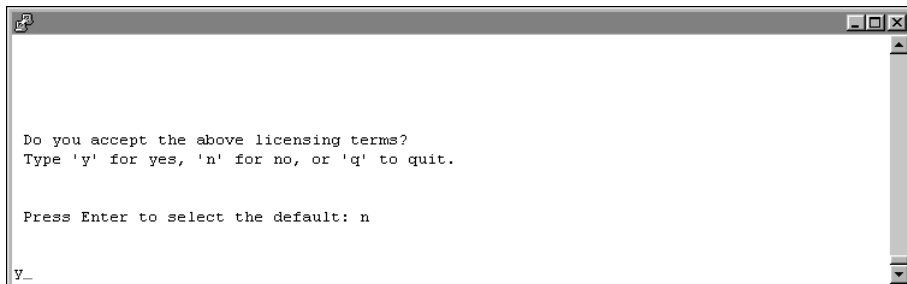
The script displays a number of prompts. Respond to the prompts as described in the following procedure.

- 5 The license agreement appears, as shown in Figure 4-30.



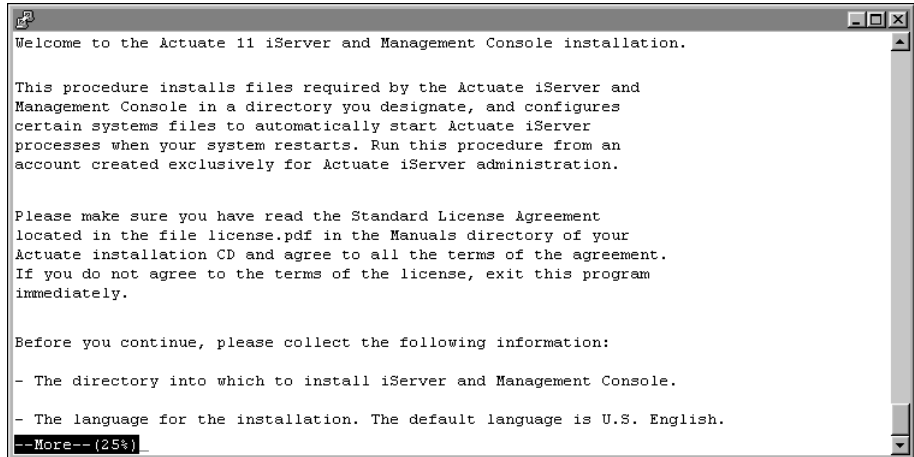
**Figure 4-30** Reviewing the license agreement

- 6 Read the license agreement, then press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 4-31.

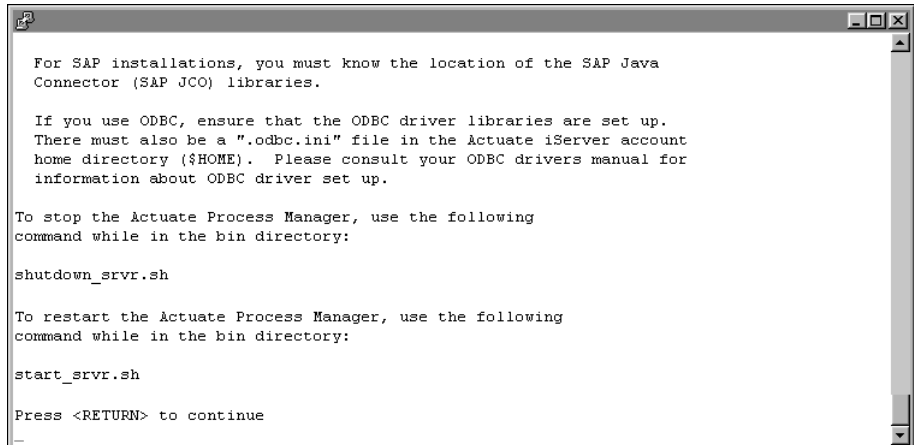


**Figure 4-31** Accepting the licensing terms

- 7 The introduction to the installation appears, as shown in Figure 4-32.
- 8 Press Enter after reviewing the introductory information, as shown in Figure 4-33.



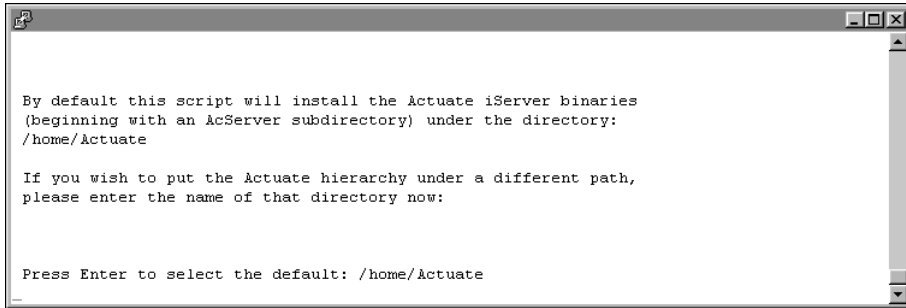
**Figure 4-32** Reviewing the introductory information



**Figure 4-33** Finishing the review of introductory information

- 9 Press Enter to select the default directory as the location of the earlier Release 11 installation. Alternatively, type the appropriate path to that location and press Enter, as shown in Figure 4-34.

iServer uses this location to resolve paths to all the binaries that it launches. The default path for this location is `$HOME/ACServer`, which is referred to in the iServer documentation by the environment variable `AC_SERVER_HOME`.

A terminal window with a standard Linux window title bar (minimize, maximize, close buttons). The text inside the terminal is as follows:

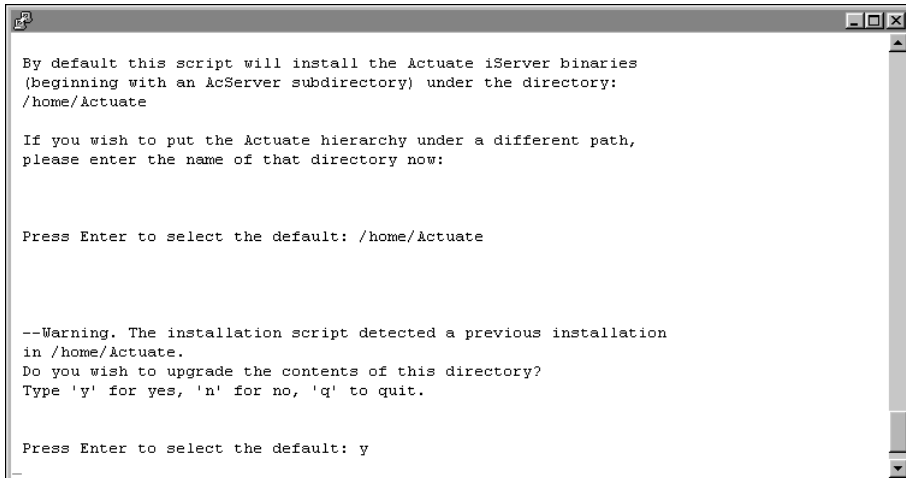
```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
```

**Figure 4-34** Specifying the installation directory

- 10 The installer detects the previous iServer version, and asks whether to perform an upgrade to the new iServer version, as shown in Figure 4-35. Press Enter to accept the default option of upgrading to the new version. Alternatively, type n for no, or q to quit.

A terminal window with a standard Linux window title bar. The text inside the terminal is as follows:

```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

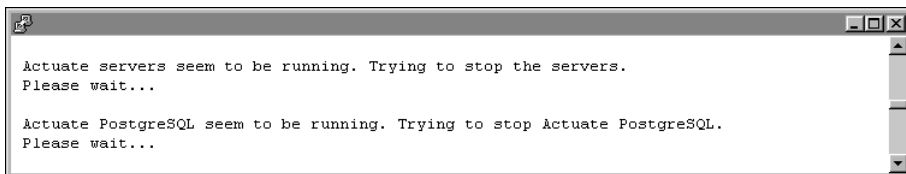
Press Enter to select the default: /home/Actuate

--Warning. The installation script detected a previous installation
in /home/Actuate.
Do you wish to upgrade the contents of this directory?
Type 'y' for yes, 'n' for no, 'q' to quit.

Press Enter to select the default: y
```

**Figure 4-35** Choosing to upgrade iServer

- 11 If the installer detects that the Actuate servers are currently running, it will try to stop the servers then continue with the installation, as shown in Figure 4-36. This process might take a few minutes.

A terminal window with a standard Linux window title bar. The text inside the terminal is as follows:


```
Actuate servers seem to be running. Trying to stop the servers.
Please wait...

Actuate PostgreSQL seem to be running. Trying to stop Actuate PostgreSQL.
Please wait...
```

**Figure 4-36** Shutting down the servers




- 12 Type the path to a new directory in which to install the iServer binaries, as shown in Figure 4-37.




**Figure 4-37** Specifying a new location for the iServer binaries

- 13 Press Enter to accept the default option of creating the folder you specified in the previous step, as shown in Figure 4-38. Alternatively, type n for no, or q to quit, and press Enter.



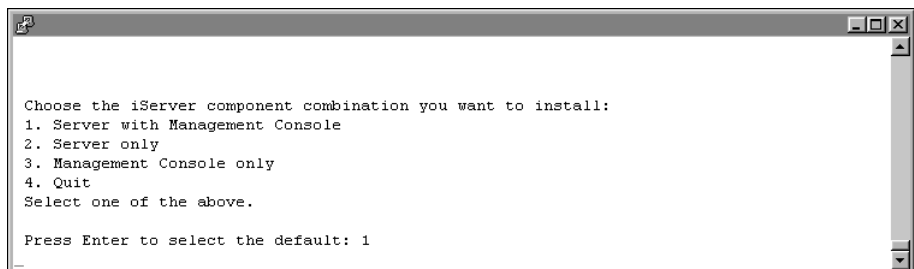
**Figure 4-38** Creating the new installation directory

The installer copies prerequisite files to the new installation directory, as shown in Figure 4-39.



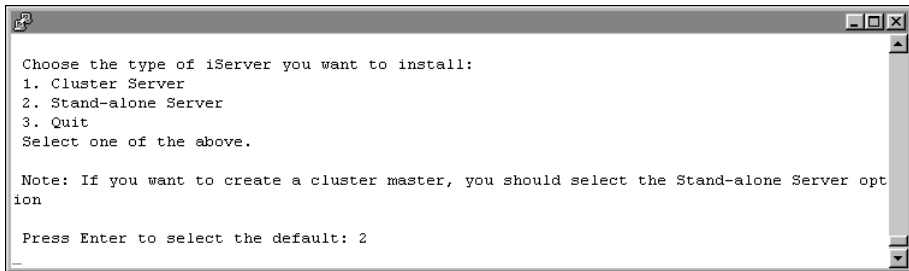
**Figure 4-39** Copying prerequisite files to the new installation directory

- 14 Press Enter to accept the default iServer component combination, as shown in Figure 4-40. Alternatively, choose a different iServer component combination and press Enter.




**Figure 4-40** Choosing the iServer components to install

- 15 Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-41.




**Figure 4-41** Choosing the iServer installation type

- 16 Specify the iServer administrator password, as shown in Figure 4-42. You use this password to log in to the iServer Configuration Console.




**Figure 4-42** Specifying the iServer administrator password

- 17 Re-enter the password of the iServer administrator, as shown in Figure 4-43. You use this password to log in to Configuration Console.



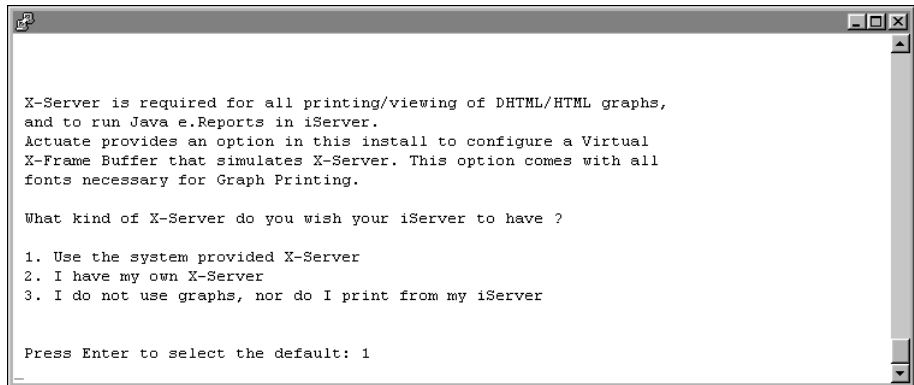
**Figure 4-43** Re-entering the iServer administrator password

- 18 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 4-44. Alternatively, type y for yes, specify the database drivers/clients you want to use, and press Enter.



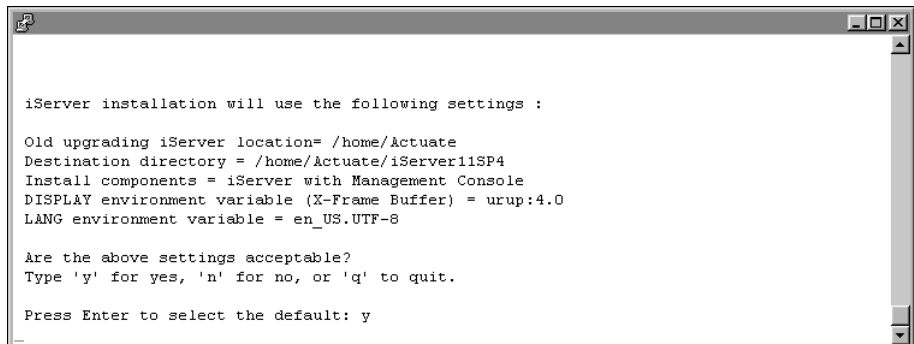
**Figure 4-44** Specifying whether to use database drivers/clients

- 19 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-45.




**Figure 4-45** Specifying what kind of X-Server to use, if any

- 20 Review the settings, as shown in Figure 4-46, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.




**Figure 4-46** Reviewing the installation settings

- 21 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-47.



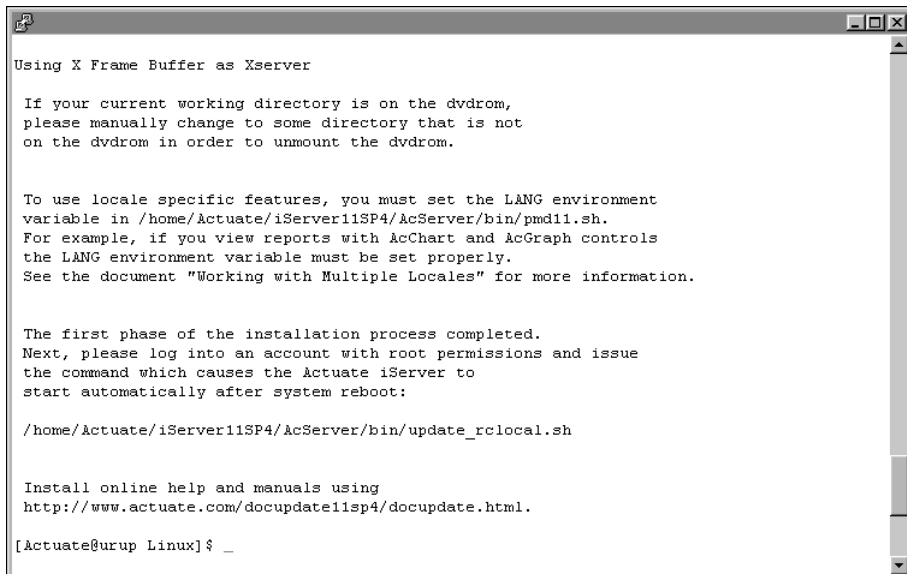
**Figure 4-47** Viewing iServer installation progress

- 22 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 4-48.



**Figure 4-48** Specifying whether to start iServer

- 23** When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-49.



**Figure 4-49** Viewing information about localization, logging in, and installing online help

- 24** Log in to Management Console. In Files and Folders, the data from the previous release appears.

## Performing a manual upgrade of an Encyclopedia volume schema for an earlier minor release

When upgrading BIRT iServer system from an earlier minor release, such as Release 11 Service Pack 3, you can perform a manual upgrade of an Encyclopedia volume, if necessary. Use the Encyclopedia Data Store Upgrader utility to complete an installation, for example, if the database is not online when the upgrade process runs.

The following section describes how to perform a manual upgrade to an in-place installation as an example.


**How to perform a manual, in-place upgrade of an Encyclopedia volume for an earlier minor release**

- 1 In a web browser, type:

`http://localhost:8900/acadmin/config`


- 2 Log in to the BIRT iServer Release 11 Configuration Console as Administrator. Choose Advanced view. Then, from the side menu, choose Volumes.

On Volumes, point to the icon next to the volume and choose Take offline, if necessary, as shown in Figure 4-50.




**Figure 4-50** Choosing to take volume offline

Choose OK to accept the default grace period to allow current transactions on the volume to complete, as shown in Figure 4-51.



**Figure 4-51** Choosing to take volume offline

- 3 On Volumes, point to the icon next to the volume schema and choose Properties, as shown in Figure 4-52.



**Figure 4-52** Choosing to view volume schema properties

- 4 On Properties, make note of the value that appears for Schema name. You must specify the schema name, not the database schema name, in a later step.

Schema name is `ac_corp`, as shown in Figure 4-53.

Volumes > Schema : ac\_corp

Schema

Metadata Database: Default\_Oracle\_MetadataDatabase

Schema type: Volume

Schema name: ac\_corp

Database schema name: ac\_corp \*

Database schema password: \*\*\*\*\* \*

\* These fields are required and cannot be left blank

\* These fields are required and cannot be left blank

Test OK Cancel Apply

**Figure 4-53** Making note of schema name

- 5 Run the Encyclopedia Data Store Upgrader utility by performing the following tasks:
  - 1 Add the following string to the PATH variable on your machine:  
`<AC_SERVER_HOME>/bin`  
For example, using the default value for AC\_SERVER\_HOME, add:  
`/home/Actuate/AcServer/bin`
  - 2 Navigate to AC\_SERVER\_HOME/bin.
  - 3 Run the `upgrade_encyclopedia_data_store.sh` file using the following command line syntax:

```
sh ./upgrade_encyclopedia_data_store.sh <schema name |  
property file name>
```

where `<schema name>` is the Encyclopedia volume schema name or `<property file name>` is the name of a pre-existing file containing the environment variable settings the utility requires. Type the schema password when prompted for it.

Alternatively, you can execute the utility by running the `upgrdedds.sh` script, using the following command-line syntax:

```
sh ./upgrdedds.sh <schema name>
```

The batch file performs the following operations, as shown in Listing 4-1:

- ❑ Checks to see if the administrator submitted a property file or a schema name on the command line when running the script.
- ❑ If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.

- Calls the `set_tools_environment.sh` script, which sets the environment variables.
- If the property file exists, the script executes the `EncyclopediaDataStoreUpgrader` utility using the name of the property file as an argument.
- If the property file does not exist, the script creates a property file, `upgrade_encyclopedia_data_store.properties`, which contains the environment variable settings.

**Listing 4-1**      `upgrade_encyclopedia_data_store.sh`

---

```
#!/bin/sh

if [ "x$1" = "x" ]; then
    echo "Usage: upgrade_encyclopedia_data_store.sh <schema
    name | property file name>"
    exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

# Check if argument 1 is a file
if [ -f "$1" ]; then
    # Use property file specified on command line
    PROPERTY_FILE="$1"
else
    # Create property file
    PROPERTY_FILE=upgrade_encyclopedia_data_store.properties
    echo "AC_SERVER_HOME = $AC_SERVER_HOME" > $PROPERTY_FILE
    echo "AC_DATA_HOME = $AC_DATA_HOME" >> $PROPERTY_FILE
    echo "USE_SERVER_CONFIG_FILE = true" >> $PROPERTY_FILE
    echo "CONFIG_SCHEMA_NAME = $1" >> $PROPERTY_FILE
fi

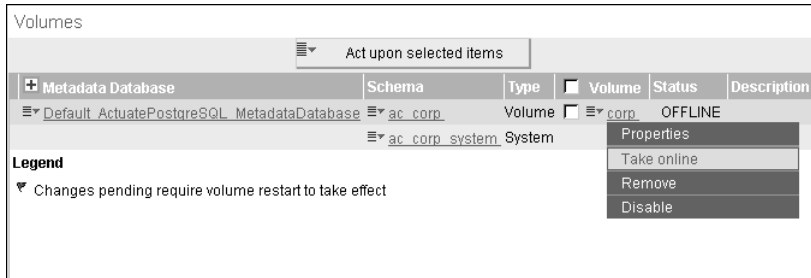
# Upgrade data store
java com.actuate.iserver.encyclopedia.datastore.admin
    .EncyclopediaDataStoreUpgrader "$PROPERTY_FILE"
```

You must run Encyclopedia Data Store Upgrader separately for every schema that you want to upgrade when upgrading from an earlier Actuate 11 version.

See Table 4-5 for descriptions of the required Encyclopedia Data Store Upgrader properties later in this chapter. See Table 4-6 for descriptions of the optional Encyclopedia Data Store Upgrader properties.


- In Advanced view of Configuration Console, choose Volumes from the side menu.

Point to the icon next to the volume name and choose Take online, as shown in Figure 4-54.



**Figure 4-54** Taking the volume online

The volume comes online, as shown in Figure 4-55.



**Figure 4-55** Viewing the status of the volume after it comes online

For more information about setting Encyclopedia Data Store Upgrader utility properties, see “Working with iServer utilities,” later in this chapter.

## Performing a manual side-by-side upgrade

In a manual side-by-side upgrade, the administrator installs a new Release 11 iServer in a path separate from the earlier release. The installation program performs the following operations using a PostgreSQL relational database management system (RDBMS) as the example:

- Installs and initializes iServer and the PostgreSQL relational database management system (RDBMS)
- Creates the iserver user in the PostgreSQL RDBMS to access the system, initializing the iserver user password to the PostgreSQL superuser password
- Creates the system and volume schema, initializing these schema with basic configuration information



- Creates the iServer configuration file, specifying system, volume, and connection information for the default installation

This manual option requires the system administrator to export the volume metadata and data from the earlier iServer volume to a temporary workspace, then import these resources into the new Release 11 installation. Actuate provides the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to perform these tasks.

The following sections describe the manual side-by-side upgrade process.

## Installing a new Release 11 iServer

The following procedure describes how to install a new Release 11 iServer in a side-by-side configuration with an earlier major release, such as Release 10 Service Pack 1, on the same machine. Installing on the same machine requires setting port numbers to values other than the defaults if you plan to continue to run the previous major release on the same machine.

The installation detects whether the earlier release is running in the environment and chooses other available port numbers automatically for the new installation. If you do not plan to run both releases simultaneously on the same machine and want to run using the default port numbers, shut down the Actuate iServer service for the earlier release before running the upgrade program.

### How to install a new Release 11 iServer


- 1 Download the required files from the FTP software distribution site. Extract the files.

- 2 To install the server files, execute the `isinstall` script:


```
sh ./isinstall.sh
```

The script displays a number of prompts. Respond to the prompts as described in the following procedure.

- 3 The license agreement appears, as shown in Figure 4-56.
- 4 Read the license agreement, then press Enter to continue the installation. At the prompt, type `y` for yes if you accept the licensing terms, as shown in Figure 4-57.

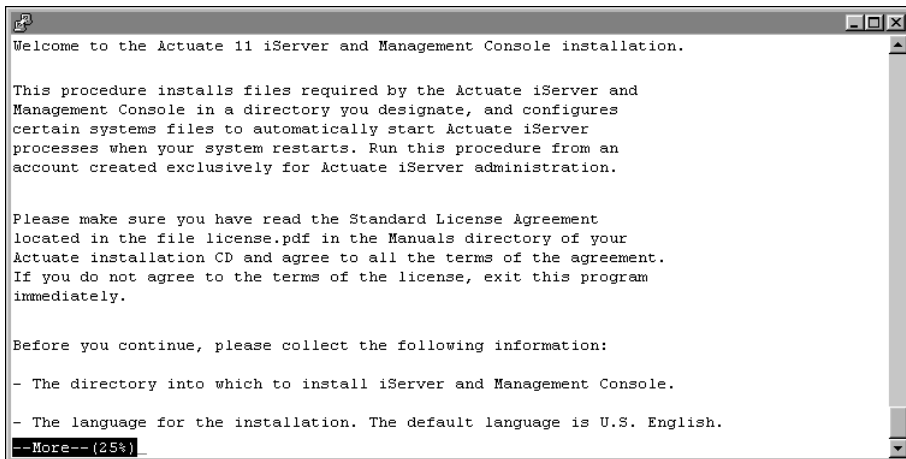


**Figure 4-56** Reviewing the license agreement



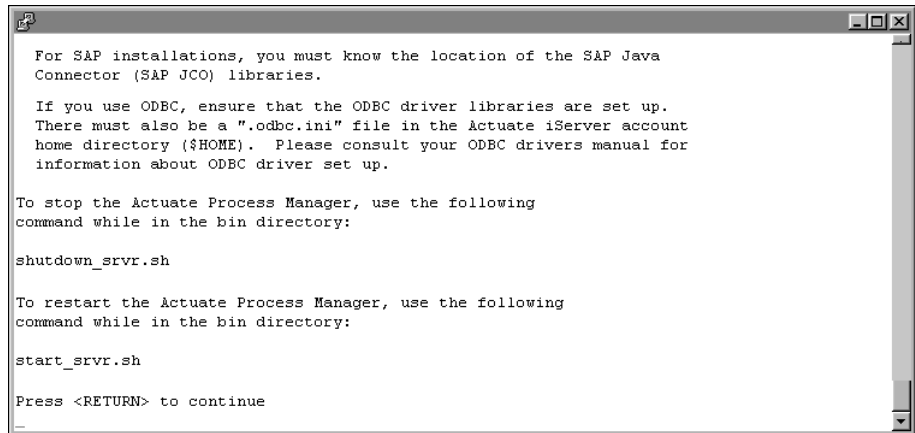
**Figure 4-57** Accepting the licensing terms

## 5 The introduction to the installation appears, as shown in Figure 4-58.



**Figure 4-58** Reviewing the introductory information

- 6 Press Enter after reviewing the introductory information, as shown in Figure 4-59.



```
For SAP installations, you must know the location of the SAP Java
Connector (SAP JCO) libraries.

If you use ODBC, ensure that the ODBC driver libraries are set up.
There must also be a ".odbc.ini" file in the Actuate iServer account
home directory ($HOME). Please consult your ODBC drivers manual for
information about ODBC driver set up.

To stop the Actuate Process Manager, use the following
command while in the bin directory:

shutdown_srvr.sh

To restart the Actuate Process Manager, use the following
command while in the bin directory:

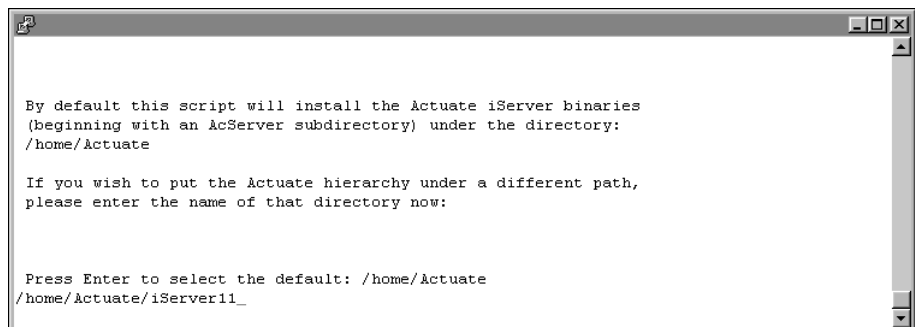
start_srvr.sh

Press <RETURN> to continue
_
```

**Figure 4-59** Finishing the review of introductory information

- 7 Type the name of a new directory in /home/Actuate, where the install program installs BIRT iServer Release 11. Specify both the path and the directory, and press enter, as shown in Figure 4-60.

iServer uses this location to resolve paths to all the binaries that it launches. The default path for this location is \$HOME/AcServer, which is referred to in the iServer documentation by the environment variable AC\_SERVER\_HOME. For example, using the new directory name shown in Figure 4-60, AC\_SERVER\_HOME refers to /home/Actuate/iServer11/AcServer.



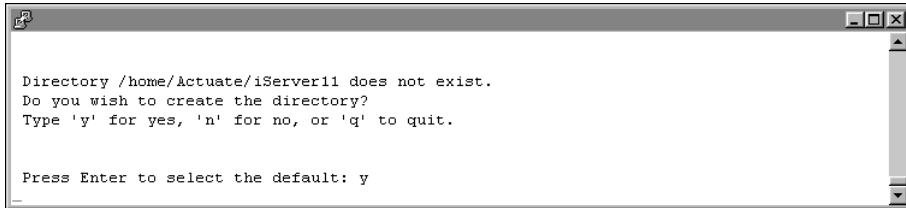
```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
/home/Actuate/iServer11_
```

**Figure 4-60** Specifying the installation directory

- 8 Press Enter to accept the default option of creating the folder you specified in the previous step, as shown in Figure 4-61. Alternatively, type n for no, or q to quit, and press Enter.

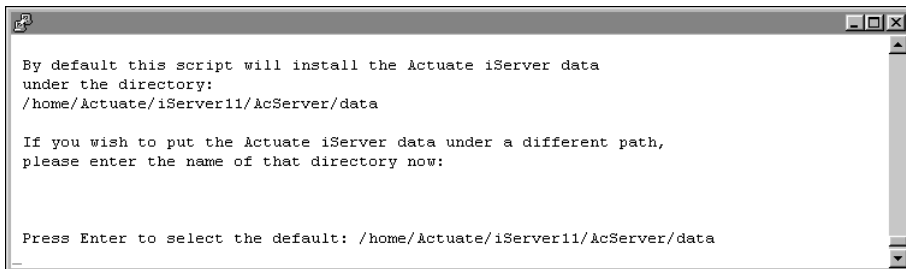


```
Directory /home/Actuate/iServer11 does not exist.  
Do you wish to create the directory?  
Type 'y' for yes, 'n' for no, or 'q' to quit.  
  
Press Enter to select the default: y
```

**Figure 4-61** Creating the new installation directory

- 9 Press Enter to accept the default installation directory, `AC_SERVER_HOME /data`, to install iServer data, as shown in Figure 4-62. Alternatively, choose a different directory for iServer data.

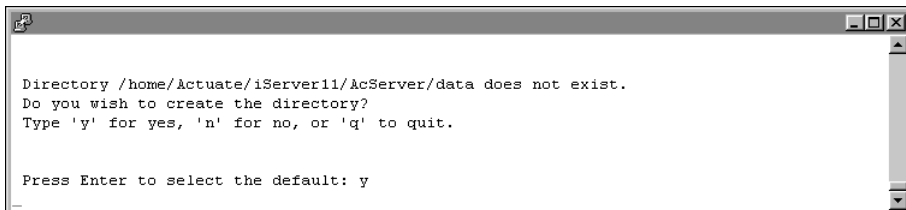
iServer uses this data location to store the iServer Encyclopedia volume data, including PostgreSQL metadata, logs, and other files. The default path is `AC_SERVER_HOME/data`, which is referred to in the iServer documentation by the environment variable `AC_DATA_HOME`.



```
By default this script will install the Actuate iServer data  
under the directory:  
/home/Actuate/iServer11/AcServer/data  
  
If you wish to put the Actuate iServer data under a different path,  
please enter the name of that directory now:  
  
Press Enter to select the default: /home/Actuate/iServer11/AcServer/data
```

**Figure 4-62** Specifying the data installation directory


- 10 The installer asks whether you want to create the directory for `AC_DATA_HOME`. Press Enter to accept the default option to create the folder, as shown in Figure 4-63. Alternatively, type `n` for no, or `q` to quit, and press Enter.



```
Directory /home/Actuate/iServer11/AcServer/data does not exist.  
Do you wish to create the directory?  
Type 'y' for yes, 'n' for no, or 'q' to quit.  
  
Press Enter to select the default: y
```

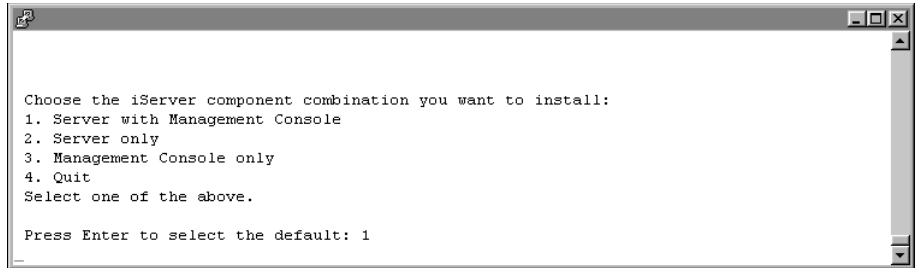
**Figure 4-63** Creating the `AC_DATA_HOME` directory

- 11 The installer copies the prerequisite files to the destination folder, as shown in Figure 4-64.



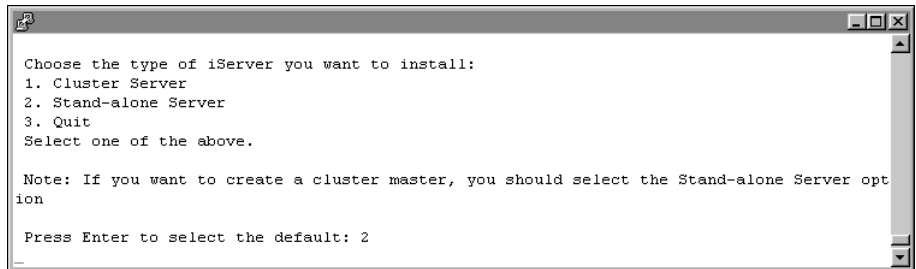
**Figure 4-64** Copying prerequisite files

- 12 Press Enter to accept the default iServer component combination, as shown in Figure 4-65. Alternatively, choose a different iServer component combination and press Enter.




**Figure 4-65** Choosing the iServer components to install

- 13 Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-66.




**Figure 4-66** Choosing the iServer installation type

- 14 Type a name to use for the BIRT iServer System name, as shown in Figure 4-67. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.




**Figure 4-67** Specifying the BIRT iServer System name

- 15 Press Enter to choose the default embedded PostgreSQL database to store the Encyclopedia volume metadata, as shown in Figure 4-68.




**Figure 4-68** Choosing the embedded PostgreSQL

- 16 Press Enter to choose the default PostgreSQL superuser name, postgres, as shown in Figure 4-69. Alternatively, type a different PostgreSQL superuser name.




**Figure 4-69** Choosing the PostgreSQL superuser name

- 17 Type a PostgreSQL superuser password that conforms to the password security policy requirements within your organization, then press Enter, as shown in Figure 4-70.




**Figure 4-70** Typing the PostgreSQL superuser password

- 18 Re-enter the password for PostgreSQL superuser, as shown in Figure 4-71, and press Enter.



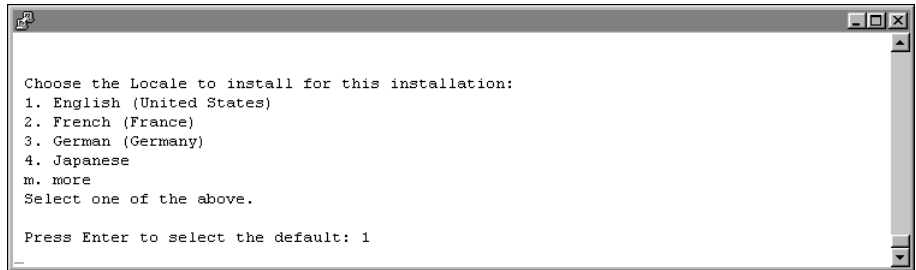
**Figure 4-71** Re-entering the PostgreSQL superuser password

- 19 Press Enter to accept the default port on which the PostgreSQL database server listens for requests, as shown in Figure 4-72. Alternatively, type a different port number.



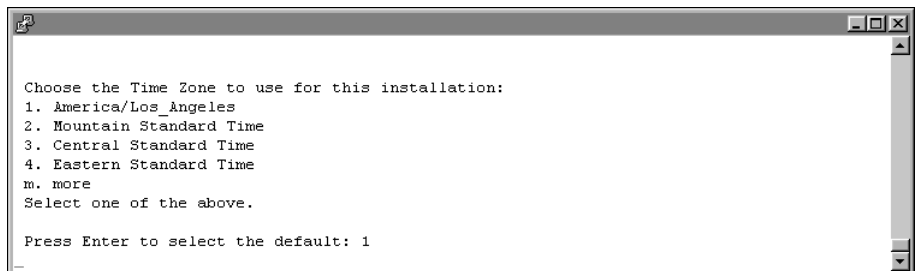
**Figure 4-72** Entering the port number that PostgreSQL uses

- 20** Press Enter to select the default locale, which is English, as shown in Figure 4-73. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.




**Figure 4-73** Specifying a locale

- 21** Press Enter to select the default time zone, which is America/Los\_Angeles, as shown in Figure 4-74. Alternatively, select another time zone from the numbered list.




**Figure 4-74** Specifying a time zone

- 22** To evaluate the product using the included evaluation software license press Enter, as shown in Figure 4-75. Alternatively, type 2, then type the path to a purchased license file.




**Figure 4-75** Specifying license type

- 23** Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 4-76. Alternatively, type a different IP address.




**Figure 4-76** Specifying the hostname that Management Console uses to contact the PMD

- 24** Press Enter to accept the default port number where Process Management Daemon (PMD) listens for requests, as shown in Figure 4-77. Alternatively, type a different port number and press Enter.



**Figure 4-77** Specifying the port number on which the PMD listens


- 25** Press Enter to accept the default hostname, the name of the machine on which iServer runs, as shown in Figure 4-78. Alternatively, type a different hostname or IP address, then press Enter.



**Figure 4-78** Specifying the machine on which the iServer runs




- 26** Press Enter to accept the default port number where iServer listens to requests, as shown in Figure 4-79. Alternatively, type a different port number and press Enter.




**Figure 4-79** Specifying the port number on which the iServer listens

- 27** Specify the iServer administrator password, as shown in Figure 4-80. You use this password to log in to the iServer Configuration Console.




**Figure 4-80** Specifying the iServer administrator password

- 28** Re-enter the password of the iServer administrator, as shown in Figure 4-81. You use this password to log in to Configuration Console.




**Figure 4-81** Re-entering the iServer administrator password

- 29** Choose n for no, to prevent iServer from using a volume name, as shown in Figure 4-82. In a side-by-side upgrade, you migrate your volumes from your earlier iServer release.




**Figure 4-82** Choosing not to use a volume name

- 30** Press Enter to accept the default option to start iServer automatically, as shown in Figure 4-83. Alternatively, type n for no.




**Figure 4-83** Specifying whether to start iServer automatically

- 31 Press Enter to accept the default option to not integrate LDAP with iServer, as shown in Figure 4-84. Alternatively, you can edit the setting.




**Figure 4-84** Specifying whether to integrate LDAP with iServer

- 32 Press Enter to accept the default option to not use any database drivers/clients, as shown in Figure 4-85. Alternatively, type y for yes, and specify the database drivers/clients you want to use.




**Figure 4-85** Specifying whether to use database drivers/clients


- 33 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-86.
- 34 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 4-87. Alternatively, type a different IP address.
- 35 Press Enter to accept the default port number on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 4-88. Alternatively, type a different port number.



**Figure 4-86** Specifying what kind of X-Server to use, if any




**Figure 4-87** Specifying the hostname that Management Console uses to contact the PMD




**Figure 4-88** Specifying the port number for the PMD to listen for requests from Management Console

- 36** Press Enter to accept the hostname or enter the IP address of the machine that Management Console uses to contact iServer, as shown in Figure 4-89. Alternatively, type a different IP address.




**Figure 4-89** Specifying the hostname that Management Console uses to contact iServer

- 37 Press Enter to accept the default port number on which iServer listens for requests from Management Console, as shown in Figure 4-90. Alternatively, type a different port number.




**Figure 4-90** Specifying the port number on which iServer listens for requests from Management Console

- 38 Press Enter to accept the name of the default Encyclopedia volume to use with Management Console, as shown in Figure 4-91. Alternatively, type a different name for the Encyclopedia volume.



**Figure 4-91** Specifying the name of the default Encyclopedia volume


- 39 Press Enter to accept the default name, acadmin, for the HTTP server context root for Management Console configuration, as shown in Figure 4-92. Alternatively, type a different name.



**Figure 4-92** Specifying the name of the HTTP server context root

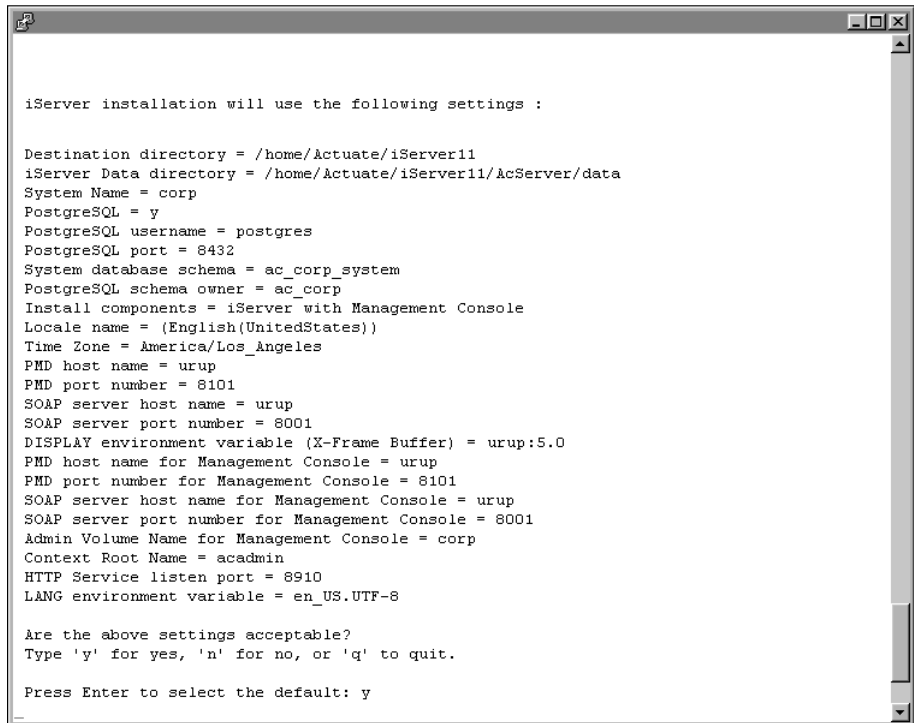
- 40 Press Enter to accept the default HTTP port number on which the application container listens to requests, as shown in Figure 4-93. Alternatively, choose a different port.

You connect to the port from your browser when accessing various features of iServer.




**Figure 4-93** Specifying the application container listening port number

- 41 Review the settings, as shown in Figure 4-94, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.



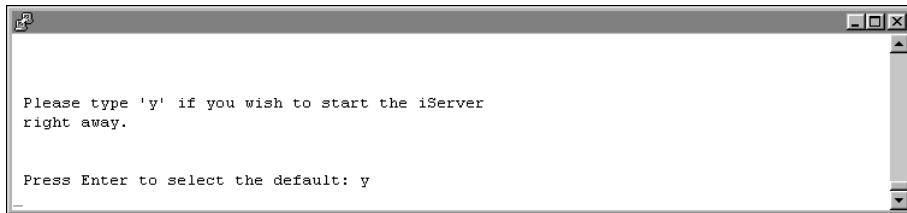
**Figure 4-94** Reviewing the installation settings

- 42 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-95.



**Figure 4-95** Viewing iServer installation progress

- 43 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 4-96.



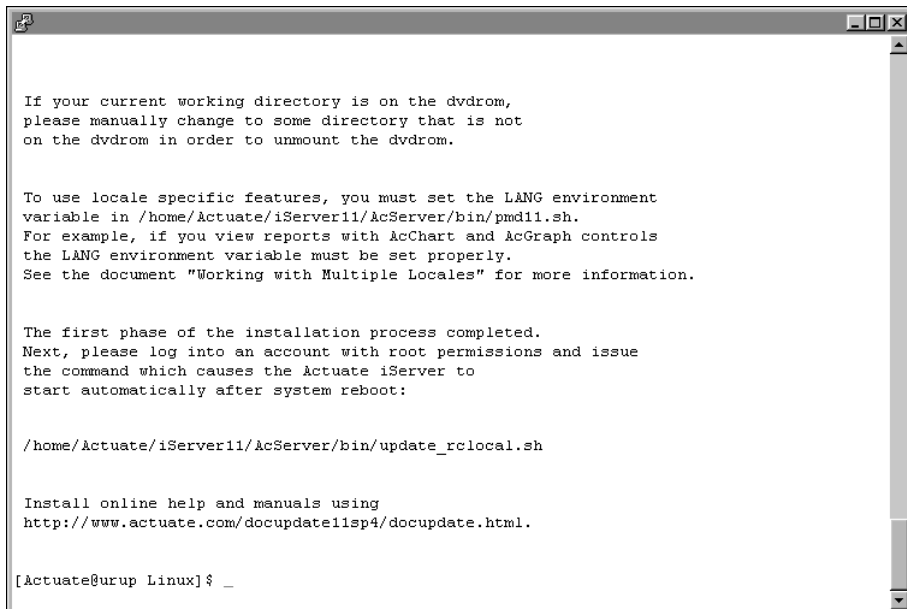
```

Please type 'y' if you wish to start the iServer
right away.

Press Enter to select the default: y
_
```

**Figure 4-96** Specifying whether to start iServer

- 44 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-97.



```

If your current working directory is on the dvdrom,
please manually change to some directory that is not
on the dvdrom in order to unmount the dvdrom.

To use locale specific features, you must set the LANG environment
variable in /home/Actuate/iServer11/AcServer/bin/pmd11.sh.
For example, if you view reports with AcChart and AcGraph controls
the LANG environment variable must be set properly.
See the document "Working with Multiple Locales" for more information.

The first phase of the installation process completed.
Next, please log into an account with root permissions and issue
the command which causes the Actuate iServer to
start automatically after system reboot:

/home/Actuate/iServer11/AcServer/bin/update_rclocal.sh

Install online help and manuals using
http://www.actuate.com/docupdate11sp4/docupdate.html.

[Actuate@urup Linux]$ _
```

**Figure 4-97** Viewing information about localization, logging in, and installing online help

## Performing a manual side-by-side migration

In a side-by-side installation, the administrator installs a new BIRT iServer Release 11 in a path separate from the earlier release. Next, the administrator manually copies the Encyclopedia volume data from the earlier release to the new iServer Release 11 location.

The administrator migrates the Encyclopedia volume schema and metadata from the earlier release to the new BIRT iServer Release 11 in two steps, using the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities. Finally, the administrator then creates a new volume in Configuration Console for the migrated volume.

The following procedure describes how to migrate an Encyclopedia volume to BIRT iServer Release 11 in a side-by-side installation using the required utilities.

### How to perform a manual export of an Encyclopedia volume in a side-by-side installation

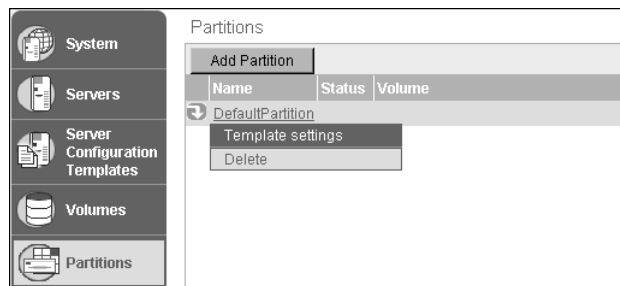
To migrate an Encyclopedia volume from an earlier iServer release to iServer Release 11, you must first export this release to a staging area using the Squirrel Data Exporter utility. To export the earlier release, perform the following tasks:

- 1 In a web browser, type:

```
http://localhost:8910/acadmin/config
```

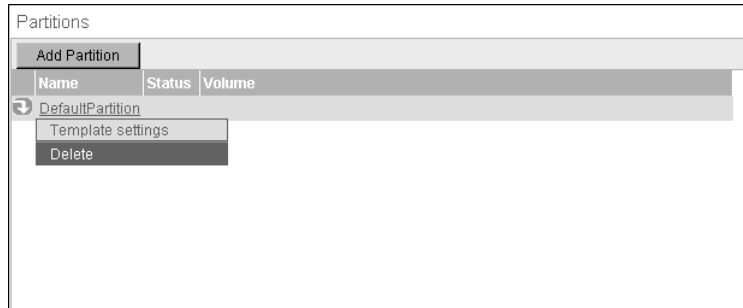
8910 is the port number that the installer specified for the iServer Application Container if you installed with the previous release running on the same machine, as shown in Figure 4-93. If you installed the new release with the previous release shut down, you may have to edit the port specified for the new iServer Application Container to manage potential port conflict if you want to run both releases simultaneously on the same machine.

- 2 Log in to Configuration Console as Administrator.
- 3 When you perform the manual side-by-side upgrade install procedure, the install program creates a default partition named DefaultPartition. If your earlier iServer release does not contain a volume that uses a partition having this name, delete DefaultPartition. If you have already done this, skip to step 5. Otherwise, delete DefaultPartition by performing the following tasks:
  - 1 Choose Advanced view.
  - 2 From the side menu, choose Partitions. On Partitions, point to the arrow next to DefaultPartition and choose Delete, as shown in Figure 4-98.



**Figure 4-98** Deleting DefaultPartition

- 4 If your earlier iServer release contains a volume that uses a partition named DefaultPartition, use DefaultPartition for that volume in Release 11 by performing the following tasks:
  - 1 Choose Advanced view.
  - 2 From the side menu, choose Partitions. On Partitions, point to the icon next to DefaultPartition and choose Template settings, as shown in Figure 4-99.



**Figure 4-99** Choosing Template Settings

- 3 On Template Settings, accept the path to the iServer Release 11 volume folder appearing in Partition Path, as shown in Figure 4-100. \$AC\_DATA\_HOME\$ points to the location of the iServer data, specified during the BIRT iServer Release 11 installation, as shown in Figure 4-62. By default, this location is:

AC\_SERVER\_HOME/data

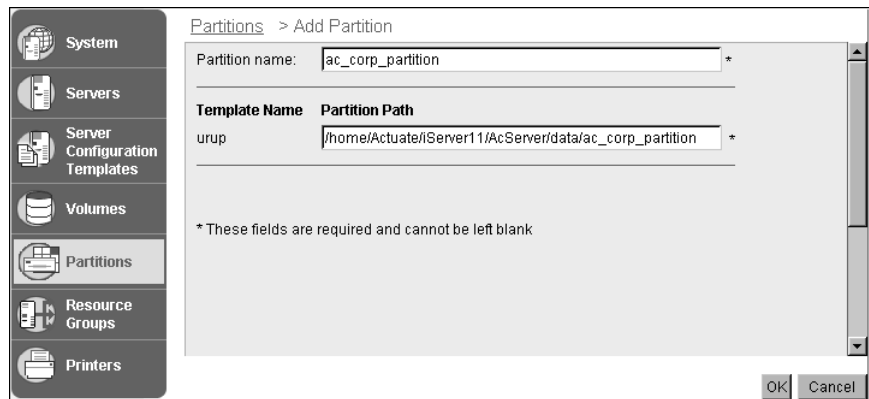
Alternatively, create a new folder for DefaultPartition in AC\_DATA\_HOME. Then, type the name of that path in Partition Path. Choose OK.



**Figure 4-100** Viewing Partition Path



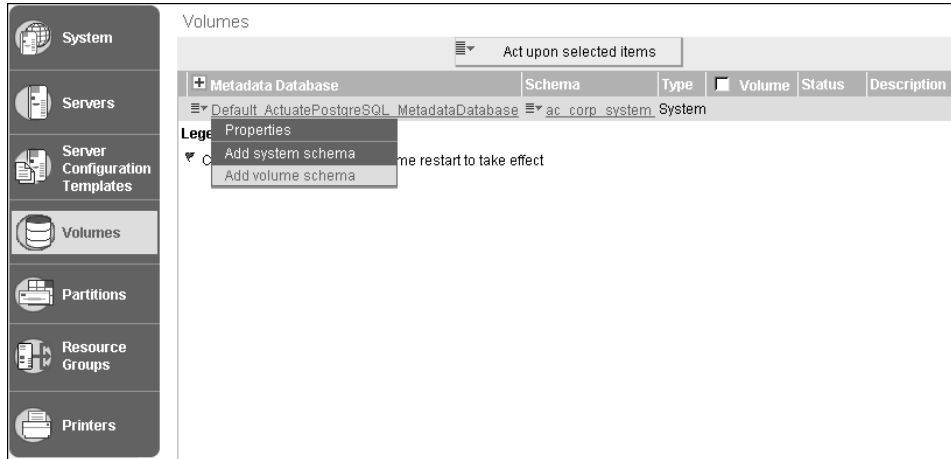
- 5 If the volume you are going to migrate from the previous iServer release does not use DefaultPartition, create a partition on iServer Release 11 for the volume. To create a new partition, perform the following tasks:
  - 1 In the operating system, create a folder for a new partition in AC\_DATA\_HOME, which points to the location of the iServer data specified during the manual side-by-side upgrade procedure.
  - 2 Log in to Configuration Console.
  - 3 Choose Advanced view.
  - 4 From the side menu, choose Partitions, then choose Add Partition.
  - 5 In Partition name, type the name of the earlier iServer release partition that specifies the path to the volume that you want to migrate. For example, if the name of the Encyclopedia volume partition on the earlier iServer release is ac\_corp\_partition, type that name in Partition name.
  - 6 In Partition Path, type the path to the iServer Release 11 Encyclopedia volume folder, as shown in Figure 4-101. This path does not need to match the partition path to the same volume on the earlier iServer release. Choose OK.



**Figure 4-101** Creating a partition for a migrated volume

- 6 Copy the file, filetype, and status folders from the old Encyclopedia volume folder to the iServer Release 11 Encyclopedia volume folder. For example, if you want to migrate the default volume from Release 10 Service Pack 1, the location of these folders is AC\_SERVER\_HOME/encyc. The default AC\_SERVER\_HOME path in Release 10 Service Pack 1 is /home/Actuate/AcServer.
- 7 In this step, create a new schema to use to create a new volume. In Advanced view of Configuration Console, from the side menu, choose Volumes.

On Volumes, point to the icon next to Default ActuatePostgreSQL MetadataDatabase and choose Add volume schema, as shown in Figure 4-102.



**Figure 4-102** Choosing to add a new schema

On Volumes—New Schema, perform the following tasks:

- 1 In Schema name, type a name for the schema. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern `[a-z][a-z 0-9]*`. Do not use a hyphen.
- 2 In Database schema name, type a name for the Database schema. The name must be less than 30 characters. Observe the same naming restrictions for this schema as the volume schema name.
- 3 In Database schema password, type a new password.
- 4 In Database schema password confirm, re-type the new password.
- 5 In Database superuser, type the database superuser name. For the PostgreSQL RDBMS that installs with iServer by default, the PostgreSQL superuser name is `postgres`.
- 6 In Database superuser password, type the superuser password. This password is the same password that the installer provides when installing iServer. Choose OK.

Volumes—New Schema appears as shown in Figure 4-103.

Volumes > New Volume Schema

**Schema**

Metadata Database: Default\_ActuatePostgreSQL\_MetadataDatabase

Schema type: Volume

Schema name:  \*

Database schema name:  \*

Database schema password:  \*

Database schema password confirm:

---

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

Database superuser:  \*

Database superuser password:  \*

\* These fields are required and cannot be left blank

**Figure 4-103** Creating a new schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-104.

Volumes

Act upon selected items

Metadata Database	Schema	Type	Volume	Status	Description
Default_ActuatePostgreSQL_MetadataDatabase	ac_corp	Volume	<input type="checkbox"/>		
	ac_corp_system	System			

**Legend**

Changes pending require volume restart to take effect

**Figure 4-104** Viewing the new schema

**8** Shut down the earlier iServer release by performing the following tasks:


- 1 In a web browser type:

`http://localhost:8900/acadmin/config`

- 2 Log in to Configuration Console as Administrator.
- 3 Choose Advanced view.

- 4 From the side menu, choose System. In System—Status, choose Stop to shut down iServer.

Make sure the earlier iServer is offline before performing the next operation, as shown in Figure 4-105.



**Figure 4-105** iServer is offline

- 9 Edit the PATH environment variable on your machine to contain the following string:

```
<AC_SERVER_HOME>/bin
```

where AC\_SERVER\_HOME refers to the Release 11 Service Pack 4 installation path. For example, using the default value for AC\_SERVER\_HOME, add:

```
/home/Actuate/iServer11/AcServer/bin
```

- 10 Run the Squirrel Data Exporter by performing the following tasks:

- 1 Navigate to AC\_SERVER\_HOME/bin.
- 2 Create a properties file named SquirrelDataExporter.properties. You pass this file to the Squirrel Data Exporter utility when you execute it. Use the following example to create SquirrelDataExporter.properties:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
SQUIRREL_DATA_HOME = /home/Actuate/AcServer/ac_corp_partition
SQUIRREL_EXPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
SQUIRREL_EXPORT_FORMAT = PostgreSQL
NEW_SCHEMA_NAME = ac_corp
NEW_VOLUME_NAME = corp
```

See Table 4-1 for descriptions of the required Squirrel Data Exporter properties. See Table 4-2 for descriptions of the optional Squirrel Data Exporter properties. See “Specifying Squirrel Data Exporter properties,” later in this chapter, for additional notes on property files.

- 11 Run the export\_squirrel\_data.sh file using the following command line syntax:

```
sh ./export_squirrel_data.sh SquirrelDataExporter.properties
```

Alternatively, you can execute the utility by running the upgrdedes.sh script, using the following command-line syntax:

```
sh ./exportsd.sh SquirrelDataExporter.properties
```

The batch file performs the following operations, as shown in Listing 4-2.

- Checks to see if the administrator submitted a property file on the command line when running the script.
- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.

- Calls the `set_tools_environment.sh` script, which sets the environment variables.
- Executes the Squirrel Data Exporter utility using the name of the properties file as an argument.

**Listing 4-2**     `export_squirrel_data.sh`

---

```
#!/bin/sh

if [ "x$1" = "x" ]; then
    echo "Usage: export_squirrel_data.sh <properties file name>"
    exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

# Export Squirrel data
java com.actuate.iserver.encyclopedia.datastore.admin
    .SquirrelDataExporter "$PROPERTY_FILE"
```

**How to import an Encyclopedia volume in a side-by-side installation**

To import an Encyclopedia volume from an earlier iServer release to BIRT iServer Release 11 from the staging area, you use the Encyclopedia Data Store Administrator utility. To import the volume, perform the following tasks:

- 1 Run the Encyclopedia Data Store Administrator by performing the following tasks:
  - 1 In Release 11 `AC_SERVER_HOME/bin`, create a properties file named `VolumeImport.properties`. You pass this file to the Encyclopedia Data Store Administrator utility when you execute it. Use the following example to create `VolumeImport.properties`:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
```

See Table 4-3 for descriptions of the required Encyclopedia Data Store Administrator properties. See Table 4-4 for descriptions of the optional Encyclopedia Data Store Administrator properties. For additional notes on property files, see “Specifying Encyclopedia Data Store Administrator properties,” later in this chapter.

- 2 Run the `administrate_encyclopedia_data_store.sh` file using the following command line syntax:

```
sh ./administrate_encyclopedia_data_store.sh
VolumeImport .properties
```

The batch file performs the following operations, as shown in Listing 4-3:

- Checks to see if the administrator submitted a property file on the command line when running the script.
- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.
- Calls the `set_tools_environment.sh` script, which sets the environment variables.
- Executes the Encyclopedia Data Store Administrator utility using the name of the properties file as an argument.

**Listing 4-3** `administrate_encyclopedia_data_store.sh`

---


```
#!/bin/sh

if [ "$1" = "x" ]; then
    echo "Usage: administrate_encyclopedia_data_store.sh
    <properties file name>"
    exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

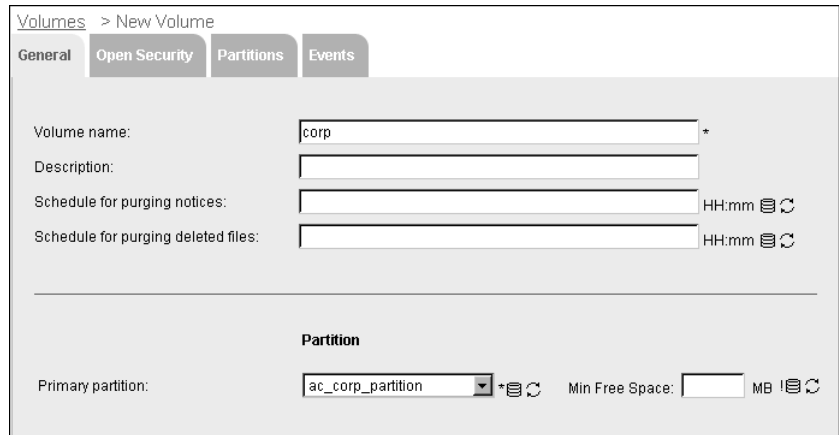
# Administrate data store
java com.actuate.iserver.encyclopedia.datastore.admin
.EncyclopediaDataStoreAdministrator "$PROPERTY_FILE"
```

- 2 Create a new volume for the migrated volume by performing the following tasks:
  - 1 Log in to Configuration Console and choose Advanced view.
  - 2 From the side menu, choose Volumes.
  - 3 On Volumes, point to the icon next to a schema and choose Add Volume, as shown in Figure 4-106.



**Figure 4-106** Choosing to add a volume

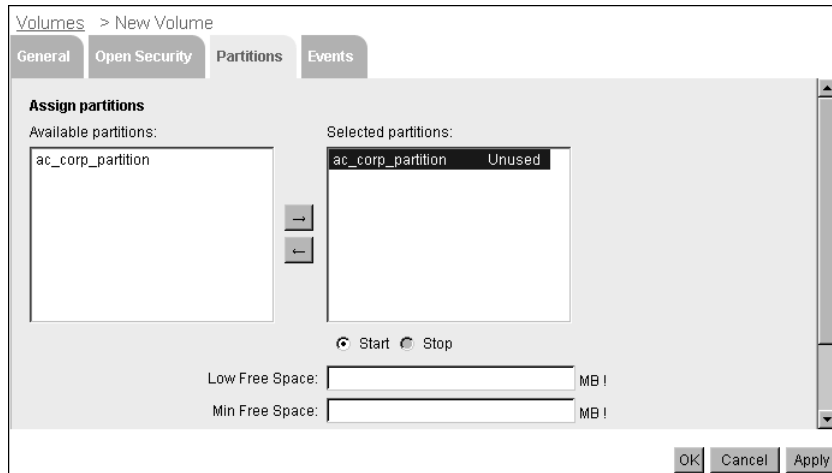
- 4 On New Volume—General, perform the following tasks:
  - 1 In Volume name, type the name you specified for NEW\_VOLUME\_NAME in SquirrelDataExporter.properties.
  - 2 In Primary partition, specify the partition you created previously, as shown in Figure 4-107. If the migrated volume uses DefaultPartition, specify DefaultPartition in Primary partition.



**Figure 4-107** Specifying volume and partition name

Choose Partitions.


- 5 On New Volume—Partitions, perform the following tasks:
  - 1 In Available partitions, select the partition that you specified in the previous step, then move it to Selected by choosing the right arrow.
  - 2 In Selected partitions, select the partition. Choose Start, as shown in Figure 4-108.



**Figure 4-108** Assigning a partition


Choose OK.

- 3 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-109.



**Figure 4-109** Viewing the new volume

The new volume comes online, as shown in Figure 4-110.



**Figure 4-110** Viewing the new volume after it comes online

- 4 Log in to Management Console. In Files and Folders, the data from your previous release appears.



- 5 If you are satisfied that the contents of the migrated volume are correct and complete, delete the folder specified in SQUIRREL\_EXPORT\_FOLDER property in the SquirrelDataExporter.properties file.

---

## Working with iServer utilities

The following sections provide information on how to specify properties and perform operations using the following utilities:

- Squirrel Data Exporter
- Encyclopedia Data Store Administrator
- System Data Store Administrator
- Encyclopedia Data Store Upgrader

When configuring Squirrel Data Exporter, Encyclopedia Data Store Administrator, System Data Store Administrator, or Encyclopedia Data Store Upgrader properties files, observe the following rules:

- Use / in path settings. \ is an escape character.
- Use Unicode property values. Save the property file in UTF-8 format, including the UTF-8 byte order mark (BOM).
- Use # at the start of a line to add a comment or selectively comment out properties.

When reading the property files, the utilities remove leading and trailing whitespace automatically, but preserve embedded whitespace.

## Working with Squirrel Data Exporter

Use Squirrel Data Exporter to configure the environment for a manual export operation and perform the export operation. The utility creates a folder to contain the exported Encyclopedia volume metadata files. Squirrel Data Exporter also creates auxiliary script files to use with RDBMS bulk-load utilities when loading Encyclopedia volume metadata into the database. For more information on these auxiliary script files, see “Using the generated bulk-load script files,” later in this chapter.

## Specifying Squirrel Data Exporter properties

Table 4-1 describes the required Squirrel Data Exporter properties used to configure the environment for a manual export operation.

**Table 4-1** Required Squirrel Data Exporter properties

Parameter	Description
AC_SERVER_HOME	Points to the location of the iServer binaries specified during the BIRT iServer Release 11 installation.
NEW_SCHEMA_NAME	Name of the target schema into which the Squirrel Data Exporter loads data. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z0-9]*. Do not use a hyphen.
NEW_VOLUME_NAME	Name of the migrated volume on BIRT iServer Release 11. This name does not need to match the name this volume had on your earlier iServer release.
SQUIRREL_DATA_HOME	Absolute path to the folder on your earlier iServer release containing the Encyclopedia volume you want to migrate.
SQUIRREL_EXPORT_FOLDER	Absolute path to the folder that Squirrel Data Exporter creates, containing the volume data from SQUIRREL_DATA_HOME, which Encyclopedia Data Store Administrator loads into the new volume on BIRT iServer Release 11.
SQUIRREL_EXPORT_FORMAT	Format of exported data. Specify DB2, Oracle, PostgreSQL, or SQLServer.

Table 4-2 describes the optional Squirrel Data Exporter properties to specify in a SquirrelDataExporter.properties file for performing logging.

**Table 4-2** Optional Squirrel Data Exporter properties

Property	Value	Default Value
FILE_LOG_LEVEL	Minimum logging level for messages sent to the log file. This parameter only works for AcLogger. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Data Store Administrator tools do not generate large logs.	5

**Table 4-2** Optional Squirrel Data Exporter properties

Property	Value	Default Value
LOG_FILE_SIZE	Maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Data Store Administrator tools do not generate large logs.	10,000,000
LOG_FOLDER	Absolute path to the log folder.	AC_DATA_HOME/server/log.

## Using the generated bulk-load script files

The Squirrel Data Exporter utility creates the folder specified by the SQUIRREL\_EXPORT\_FOLDER property. This folder contains the Encyclopedia volume metadata files exported from the volume. Squirrel Data Exporter also creates auxiliary script files to use with RDBMS bulk-load utilities when loading Encyclopedia volume metadata into the database.

The SQUIRREL\_EXPORT\_FORMAT property specifies the database type. The database type determines the type of file the utility creates to contain the metadata as well as what configuration files to create to support loading that data into the database using the standard RDBMS data-loading utility.

The SQUIRREL\_EXPORT\_FORMAT property supports the following values:

- Oracle

The Squirrel Data Exporter utility creates a .dat file for each metadata set that SQL\*Loader loads into a single database table, such as user or security role. For every .dat file, the utility creates a .ctl file containing the information that SQL\*Loader requires to load the contents of the associated .dat file, such as the name of the table to load, the character in the .dat file used to separate each data field, and the column names in the table.

The utility also creates the following files for use by the SQL\*Loader:

- ac\_oracle\_load.bat  
Contains a list of sqlldr commands, one for each .ctl and .dat file pair.
- ac\_oracle\_load.par  
Contains the parameters passed in each invocation of SQL\*Loader. This file is passed to SQL\*Loader in each sqlldr command that ac\_oracle\_load.bat contains. Modify this file to contain the appropriate server, user, password, port, and database information before running an ac\_oracle\_load script.
- ac\_oracle\_load.sh  
Contains the UNIX equivalent of the batch file containing the list of sqlldr commands.

Before running `ac_oracle_load.bat` or `ac_oracle_load.sh`, the database administrator must copy the `SQUIRREL_EXPORT_FOLDER` to the location of the Oracle SQL\*Loader utility.

- **DB2**

The Squirrel Data Exporter utility creates a `.csv` file for each type of metadata that DB2 loads into a single database table, such as user or security role.

The utility also creates the following files for use by DB2:

- `ac_db2_load.bat`  
Contains a single `db2cmd` invocation that runs `ac_db2_load.db2`.
- `ac_db2_load.db2`  
Contains one `LOAD` command for each `.csv` file. Each `LOAD` command includes all the information DB2 needs to load a single table. Modify this file to contain the appropriate server, user, password, port, and database information before running an `ac_db2_load` script.
- `ac_db2_load.sh`  
Contains the UNIX equivalent of `ac_db2_load.bat`.

Before running `ac_db2_load.bat` or `ac_db2_load.sh`, the database administrator must copy the `SQUIRREL_EXPORT_FOLDER` to the machine hosting DB2.

- **SQL Server**

The Squirrel Data Exporter utility creates a `.dat` file for each type of metadata that SQL Server loads into a single database table, such as user or security role.

The utility also creates the following files for use by SQL Server:

- `ac_sql_server_load.bat`  
Contains a single `sqlcmd` invocation that runs `ac_sql_server_load.sql`. Modify this file to contain the appropriate server, user, password, port, and database information before running this batch file.
- `ac_sql_server_load.sql`  
Contains one `BULK INSERT` command for each `.dat` file. This file also contains all the information SQL Server needs to load each table.

Before running `ac_sql_server_load.sql`, the database administrator must copy the `SQUIRREL_EXPORT_FOLDER` to the machine hosting SQL Server.

- **PostgreSQL**

The Squirrel Data Exporter utility creates a `.csv` file for each type of metadata that PostgreSQL loads into a single database table, such as user or security role.

The utility also creates the following files for use by PostgreSQL:

- `ac_postgresql_client_load.bat`  
Contains a single `psql` command invocation that runs `ac_postgresql_client_load.psql`. Modify this file to contain the appropriate server, user, password, port, and database information before running this batch file.
- `ac_postgresql_client_load.psql`  
Contains one copy command for each `.csv` file.
- `ac_postgresql_client_load.sh`  
Contains the UNIX equivalent of `ac_postgresql_client_load.bat`. Modify this file to contain the appropriate server, user, password, port, and database information before running this script.

Before running `ac_postgresql_client_load.bat` or `ac_postgresql_client_load.sh`, the database administrator must copy the `SQUIRREL_EXPORT_FOLDER` to the machine hosting PostgreSQL.

## Working with Encyclopedia Data Store Administrator

Use Encyclopedia Data Store Administrator to import or export volume metadata and create or delete a volume schema.

### Specifying Encyclopedia Data Store Administrator properties

Table 4-3 describes the required Encyclopedia Data Store Administrator properties used to configure the environment for a manual migration operation.

**Table 4-3** Required Encyclopedia Data Store Administrator properties

Parameter	Description
<code>AC_SERVER_HOME</code>	Points to the location of the iServer binaries specified during the BIRT iServer Release 11 installation.
<code>DATABASE_TYPE</code>	Type of supported RDBMS that contains the data store. Specify DB2, Oracle, or PostgreSQL.
<code>LOG_FOLDER</code>	Absolute path to the log folder.
<code>SCHEMA_FILE_NAME</code>	Base name of the file without the extension that contains the schema definition.
<code>SCRIPT_HOME</code>	Absolute path to the root of the folder hierarchy that contains the scripts and the schema definition.

Table 4-4 describes optional Encyclopedia Data Store Administrator properties. The properties used depend on the type of operation performed and the installation environment.

**Table 4-4** Optional Encyclopedia Data Store Administrator properties

Parameter	Description	Default value	Supported databases
APPLICATION_USER	Application user ID for connecting to the database for normal operations. Must be a legal SQL identifier. Typically iserver. This parameter is required when creating or populating a schema. Restrict the iServer application user name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.		All
APPLICATION_USER_PASSWORD	Application user password for connecting to the database for normal operations. Required when creating a schema. The password is not encrypted.		All
CONFIG_SCHEMA_NAME	Specifies the schema definition in acserverconfig.xml. Use this property when any of the following properties is required but is not specified in the properties file: <ul style="list-style-type: none"> <li>■ APPLICATION_USER</li> <li>■ DATABASE_HOST</li> <li>■ DATABASE_INSTANCE</li> <li>■ DATABASE_NAME</li> <li>■ DATABASE_PORT</li> <li>■ DATABASE_TYPE</li> <li>■ ORACLE_TNS_NAMES_FILE</li> <li>■ SCHEMA_NAME</li> </ul>	False	All

**Table 4-4** Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
CREATE_SCHEMA	Set to true to create a new schema, false to use an existing schema. Specifies whether to create a new Encyclopedia schema. PostgreSQL only, not supported for DB2 or Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL
DATA_EXPORT_FOLDER	Absolute path to the folder to which data is exported. This parameter is required if EXPORT_DATA is true. Folder is on the database machine.		All
DATA_EXPORT_FORMAT	Format of exported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE_TYPE}	All
DATA_IMPORT_FOLDER	Absolute path to the folder from which data is imported. This parameter is required if IMPORT_DATA is true. For PostgreSQL, the data must be on the database server.		All
DATA_IMPORT_FORMAT	Format of imported data. Specify DB2, Oracle, or PostgreSQL.	{DATABASE_TYPE}	All
DATABASE_HOST	Hostname or IP address of the machine hosting the database. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
DATABASE_NAME	Database name for PostgreSQL. For Oracle not using TNS, the database service name. For Oracle using TNS, the TNS net service name.		All
DATABASE_PORT	Port that the database server uses. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All

*(continues)*

**Table 4-4** Optional Encyclopedia Data Store Administrator properties (continued)

<b>Parameter</b>	<b>Description</b>	<b>Default value</b>	<b>Supported databases</b>
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports DB2, Oracle, and PostgreSQL in Linux/UNIX.		All
DEFAULT_DATABASE_NAME	Used by the superuser to connect to the default database in order to create the iServer application database. Required when the application database does not exist. PostgreSQL only.		PostgreSQL
DELETE_ALL_DATA	Set to true to delete all data from the data store. Does not delete data from other tables in the schema.	False	All
DELETE_DATA	Set to true to delete filtered data from the data store. This parameter only works when a filter value is available.	False	All
DEPOPULATE_SCHEMA	Set to true to delete data store elements such as tables, views, stored functions, and procedures from the schema. Does not remove any other objects in the schema.	False	All
DROP_SCHEMA	Set to true to delete the schema.	False	PostgreSQL
EXPORT_ALL_DATA	Set to true to export all volumes from a schema.	False	All
EXPORT_DATA	Set to true to export a single volume from a schema. Use VOLUME_NAME to specify the name of the volume to export. Optionally, use NEW_VOLUME_NAME to specify a new name for the exported volume.	False	All



**Table 4-4** Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
FILE_LOG_LEVEL	Minimum logging level for messages sent to the log file. This parameter only works for AcLogger, and is a standard JUL Level name. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts to perform operations instead of performing the operations directly	false	All
IMPORT_DATA	Set to true to import data into the data store.	False	PostgreSQL
INITIALIZE_DATA	Set to true to initialize the data in the data store, using the data initialization script.	False	All
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Data Store Administrator tools do not generate large logs.	5	All
LOG_FILE_NAME	Name of the log file. Do not add a file extension. The extension is set to .log. A unique number is appended automatically to the file name to prevent overwriting previous logs.	Encyclopedia DataStore Administrator <number>.log	
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Data Store Administrator tools do not generate large logs.	10,000,000	All
LOG_FOLDER	Full path of folder to write logs.	AC_DATA _HOME /server/log.	

*(continues)*

**Table 4-4** Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
NEW_SCHEMA_NAME	Name of the new schema to be created, or the name of the target schema for data export. This name is required if {SCHEMA_NAME} is not present. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	NEW_SCHEMA_NAME	All
NEW_VOLUME_NAME	Name of the new volume to be created, or the name of the target schema for data export. Do not use a hyphen in a volume name.	NEW_VOLUME_NAME	All
ORACLE_TNS_NAMES_FILE	Absolute path to the Oracle TNS names file. This path can be used instead of {DATABASE_HOST} and {DATABASE_PORT} to generate a JDBC URL.		Oracle
POPULATE_SCHEMA	Set to true to populate the schema with data store elements such as tables, views, stored functions, and procedures. Set to false if the data store elements already exist. Use Configuration Console to populate schema when possible. Schema automatically populates when defined in Configuration Console.	CREATE_SCHEMA	All
POPULATE_SCHEMA_PHASE_ONE	Set to true to create the schema objects necessary for data to be loaded. Set to false to prevent this schema object creation. Imports data without building indexes for fast load.	POPULATE_SCHEMA	All
POPULATE_SCHEMA_PHASE_TWO	Set to true to create schema objects phase one does not create. Set to false to prevent this schema object creation. Builds indexes and other structures in database.	POPULATE_SCHEMA	All

**Table 4-4** Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
PROMPT_FOR_PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it as a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	True	All
SCHEMA_NAME	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	NEW_SCHEMA_NAME	All
SCHEMA_PASSWORD	Password for the schema owner. Required if PROMPT_FOR_PASSWORDS is false.		All
SUPERUSER	User ID of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL. Typically postgres in a PostgreSQL database.		PostgreSQL
SUPERUSER_PASSWORD	Password of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL.		PostgreSQL
TABLESPACE_NAME	Name of the application tablespace.		PostgreSQL
TABLESPACE_LOCATION	Absolute path to the application tablespace folder.		PostgreSQL
TIME_ZONE	Local time zone for installation.		All
VOLUME_NAME	Name of the volume or target schema for data import.	VOLUME_NAME	All

## Performing operations using Encyclopedia Data Store Administrator utility

The Encyclopedia Data Store Administrator utility supports a wide range of import and export operations in the BIRT iServer environment. The following sections describe a mix of operational scenarios that an administrator can perform using this utility.

### Importing One or More Volumes into a New Schema

This operation is only supported for PostgreSQL. `NEW_SCHEMA_NAME` does not need to match the schema name from which the Squirrel Data Exporter utility exported the data. Do not change the volume names, since these names are defined in the exported data.

Creating a schema requires superuser privileges. The Encyclopedia Data Store Administrator utility creates the database and users if these items do not already exist.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {SQUIRREL_EXPORT_FOLDER}
```

### Importing One or More Volumes into a Populated Schema

This operation is only supported for PostgreSQL. Typically, you perform this operation to import an additional volume into a shared schema that already contains one or more volumes.

`SCHEMA_NAME` does not need to match the schema name from which the Squirrel Data Exporter utility exported the data. Do not change the volume names, because these names are defined in the exported data. Importing data into PostgreSQL requires superuser privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {SQUIRREL_EXPORT_FOLDER}
```

### Exporting All Volumes from a Schema

Squirrel Data Exporter generates bulk load script files in the same folder as the exported data. For more information on using the generated bulk-load script files, see, “Using the generated bulk-load script files,” earlier in this chapter.

You can omit `NEW_SCHEMA_NAME` if it is the same as `SCHEMA_NAME`.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
EXPORT_ALL_DATA = true
NEW_SCHEMA_NAME = <provide a name>
DATA_EXPORT_FOLDER = home/Projects/DataStores/Data
DATA_EXPORT_FORMAT = Oracle
```

### Exporting a Single Volume from a Schema

Squirrel Data Exporter generates bulk load script files in the same folder as the exported data. For more information on using the generated bulk-load script files, see, “Using the generated bulk-load script files,” earlier in this chapter.

You can omit `NEW_SCHEMA_NAME` if it is the same as `SCHEMA_NAME`. You can omit `NEW_VOLUME_NAME` if it is the same as `VOLUME_NAME`.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
```

```
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
EXPORT_DATA = true
VOLUME_NAME = <provide a name>
NEW_SCHEMA_NAME = <provide a name>
NEW_VOLUME_NAME = <provide a name>
DATA_EXPORT_FOLDER = home/Projects/DataStores/Data
DATA_EXPORT_FORMAT = Oracle
```

## Deleting All Volumes from a Schema

Perform this operation to remove all volumes from the schema that SCHEMA\_NAME specifies.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
DELETE_ALL_DATA = true
```

## Deleting a Single Volume from a Schema

Perform this operation to remove the volume that VOLUME\_NAME specifies from the schema that SCHEMA\_NAME specifies.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
DELETE_DATA = true
VOLUME_NAME = <provide a name>
```

## Creating a New Volume in an Empty Schema

iServer performs this operation when you create a new volume in Configuration Console and bring the volume online. When creating a new volume using Encyclopedia Data Store Administrator, set INITIALIZE\_DATA to true and specify a name for NEW\_VOLUME\_NAME.

Setting POPULATE\_SCHEMA to true to have Encyclopedia Data Store Administrator create schema elements, such as tables and indexes in an empty schema. The schema owner and application user must already exist and have appropriate privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

### **Populating an Empty Schema**

iServer performs this operation when you create a new schema in Configuration Console. Before populating an empty schema using Encyclopedia Data Store Administrator, create the schema owner and application user if they do not exist and give them the appropriate privileges.

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
```

### **Creating a New Volume in a Populated Schema**

iServer performs this operation when you create a new volume in Configuration Console and bring the volume online. Typically, you perform this operation to add a volume to a shared schema that already contains one or more volumes.

When creating a new volume using Encyclopedia Data Store Administrator, set INITIALIZE\_DATA to true and specify a name for NEW\_VOLUME\_NAME. The schema owner and application user must already exist and have appropriate privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

## Creating a New Volume in a New Schema

This operation is only supported for PostgreSQL. iServer typically performs this operation when you create a new volume in Configuration Console and bring the volume online.

Creating a schema requires superuser privileges. Encyclopedia Data Store Administrator creates the database and users if these items do not already exist.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

## Working with System Data Store Administrator

Use the System Data Store Administrator utility to perform the following tasks:

- Create or delete a schema.
- Populate or depopulate a schema.



## Specifying System Data Store Administrator properties

The system data store is a separate schema in the database that stores iServer metadata. The system schema is a required element for any iServer installation. In a cluster, the nodes share the system schema metadata and use this information to communicate and coordinate processing.

In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this procedure to protect critical system metadata. If a system schema becomes corrupted or is accidentally deleted, the administrator can use the System Data Store Administrator utility to create a new schema. You can define a new system schema in Configuration Console or edit `acserverconfig.xml` to refer to the current system schema.

Use the System Data Store Administrator utility to perform the following tasks:

- Create or delete a schema.
- Populate or depopulate a schema.
- Import or export data.

Before running the System Data Store Administrator utility, edit the `PATH` environment variable on your machine to contain the following string:

```
<AC_SERVER_HOME>\bin
```

where `AC_SERVER_HOME` refers to the Release 11 Service Pack 4 installation path. For example, using the default value for `AC_SERVER_HOME`, add:

```
/home/Actuate/iServer11/AcServer/bin
```

To run the System Data Store Administrator utility, perform the following tasks:

- 1 Navigate to `AC_SERVER_HOME/bin`.
- 2 Create a properties file containing the property definitions necessary to perform an operation, such as creating and populating a system schema, described in “Performing operations using System Data Store Administrator utility,” later in this chapter. You pass this file to the System Data Store utility when you execute it.
- 3 Run the `administrate_system_data_store.sh` file using the following command line syntax:

```
sh ./administrate_system_data_store.sh systemdatastore
.properties
```

Alternatively, you can execute the utility by running the `upgrdsds.sh` script, using the following command-line syntax:

```
sh ./adminsds systemdatastore.properties
```

The batch file performs the following operations, as shown in Listing 4-4:

- Checks to see if the administrator submitted a property file on the command line when running the script.
- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.
- Calls the `set_tools_environment.sh` script, which sets the environment variables.
- Executes the System Data Store Administrator utility using the name of the properties file as an argument.

**Listing 4-4**      `administrate_system_data_store.sh`

---

```
#!/bin/sh

if [ "x$1" = "x" ]; then
    echo "Usage: administrate_system_data_store.sh <properties
    file name>"
    exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

# Administrate data store
java com.actuate.iserver.system.datastore.admin
    .SystemDataStoreAdministrator "$PROPERTY_FILE"
```

The `SystemDataStoreAdministrator` class has the same parent class as the Encyclopedia Data Store Administrator and uses the same property settings. System Data Store Administrator properties include the following categories:

- Properties that specify details of the iServer installation environment, such as `AC_SERVER_HOME` and `AC_DATA_HOME`
- Database properties that specify the RDBMS type and JDBC connection details, such as the schema, application user, and superuser names and passwords
- Schema operation properties that specify an action to perform, such as create, populate, or delete a schema
- Logging properties that control messages sent to the console and log files
- Engineering properties used by Actuate Support or Professional Services to assist with diagnosing or resolving specific issues

Refer to the required and optional properties tables in “Specifying Encyclopedia Data Store Administrator properties,” earlier in this chapter, for detailed information about these properties.

## Performing operations using System Data Store Administrator utility

This operation is only supported for PostgreSQL. Creating and populating a schema requires superuser privileges. iServer performs this operation automatically when you create a new system schema in Configuration Console. The System Data Store Administrator utility creates the database and users if these items do not already exist.

Configure the properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
INITIALIZE_DATA = true
```

## Working Encyclopedia Data Store Upgrader

Use Encyclopedia Data Store Upgrader to upgrade an Encyclopedia volume manually from an earlier Release 11 installation, such as Service Pack 3.

### Specifying Encyclopedia Data Store Upgrader properties

Table 4-5 describes the required Encyclopedia Data Store Upgrader properties to specify in the `upgrade_encyclopedia_data_store.bat` or other properties file for a manual upgrade operation.

**Table 4-5** Required Encyclopedia Data Store Upgrader properties

Parameter	Description
AC_SERVER_HOME	Points to the location of the iServer binaries, which you specify during the BIRT iServer Release 11 installation.
APPLICATION_USER	User ID used to connect to the database for normal operations.

*(continues)*

**Table 4-5** Required Encyclopedia Data Store Upgrader properties (continued)

Parameter	Description
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports PostgreSQL and Oracle.
LOG_FOLDER	Absolute path to the log folder.
SCHEMA_FILE_NAME	Base name of the file without the file extension that contains the meta-schema definition.
SCHEMA_NAME	Name of the target schema which the Encyclopedia Data Store Upgrader updates. Required if NEW_SCHEMA_NAME is not present. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.
SCRIPT_HOME	Absolute path to the root of the folder hierarchy containing scripts and the meta-schema definition.

Table 4-6 describes the optional Encyclopedia Data Store Upgrader properties to specify in the upgrade\_encyclopedia\_data\_store.bat or other properties file.

**Table 4-6** Optional Encyclopedia Data Store Upgrader properties

Parameter	Description	Default value	Supported databases
CONFIG_SCHEMA_NAME	Specifies the schema definition in acserverconfig.xml. The schema name can be different from the database schema name.	False	All
DATABASE_HOST	Hostname or IP address of the machine hosting the database. This value is required for PostgreSQL. Required for Oracle if not using TNS.	False	All
DATABASE_INSTANCE	RDBMS instance that manages the database.		SQL Server
DATABASE_NAME	Database name for PostgreSQL. For Oracle not using TNS, the database service name. For Oracle using TNS, the TNS net service name.		All

**Table 4-6** Optional Encyclopedia Data Store Upgrader properties (continued)

Parameter	Description	Default value	Supported databases
DATABASE_PORT	Port that the database server uses. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
FILE_LOG_LEVEL	The minimum logging level for messages sent to the log file. This parameter only works for AcLogger. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts to perform operations instead of performing the operations directly	false	All
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Encyclopedia Data Store Upgrader tool does not generate large logs.	5	All
LOG_FILE_NAME	Name of the log file. Do not add a file extension. The extension is set to .log. Standard JUL.FileHandler placeholders are supported. A unique number is appended automatically to the file name to prevent overwriting previous logs.		All
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Encyclopedia Data Store Upgrader tool does not generate large logs.	10,000,000	All
LOG_FOLDER	The full path of folder to write logs.	AC_DATA _HOME /server/log.	All

*(continues)*

**Table 4-6** Optional Encyclopedia Data Store Upgrader properties (continued)

<b>Parameter</b>	<b>Description</b>	<b>Default value</b>	<b>Supported databases</b>
ORACLE_TNS_NAMES_FILE	Absolute path of the Oracle TNS names file to use instead of DATABASE_HOST and DATABASE_PORT to generate a JDBC URL.		Oracle
PROMPT_FOR_PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it as a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	True	All
SCHEMA_PASSWORD	Password for the schema owner. Required if PROMPT_FOR_PASSWORDS is false.		All

# Installing a BIRT iServer cluster

This chapter discusses the following topics:

- Installing a BIRT iServer cluster node
- Preparing to install an iServer cluster
- Performing a cluster installation using the wizard
- Adding a node to a cluster
- Finding the BIRT iServer home directory
- About the Java Development Kit

---

## Installing a BIRT iServer cluster node

A node is a machine running an BIRT iServer instance. An iServer administrator adds a node to an iServer cluster to improve availability and throughput and scale the cluster installation to necessary processing requirements.

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

- Perform an automated, custom installation, using the wizard-driven installation program.
- Perform a manual installation or cloud deployment, using a prepared image of an installed iServer run-time environment.

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

- Shared configuration home and Encyclopedia volume directories
- Cluster resources, such as printers, database systems, and disk storage systems

It is the responsibility of the administrator performing the installation to make sure that all network sharing settings conform to the security policies in force for the environment.

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 the settings to the process environment variable list. When a node joins a cluster, it configures itself using its designated template.

After installation and configuring the appropriate environment variables in `acpmdconfig.xml`, the administrator launches the installed iServer image from the command line by passing the necessary arguments or creates 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 iServer clustering.

The cluster communicates across the network using standard HTTP/IP addressing. The cluster automatically detects the on-off status of any node. Single-point node failure does not affect the availability of other nodes.

One or more nodes in the cluster manage the request message routing. The Process Management Daemons (PMDs) located on each node coordinate processing among available iServer services based on message type to balance load across the nodes.



iServer instances running on multiple machines maintain iServer system and Encyclopedia volume metadata in databases and control access to shared volume data. The volume data can be on machines that are not running iServer, but must be shared and accessible to each iServer instance. It is the responsibility of the administrator performing the installation to make sure that all network sharing settings conform to the security policies in force for the environment.

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

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

This operational model lends itself well to grid, cloud, and other data-center types of deployments. For more information about the pre-packaged Actuate cloud computing deployment option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book. For more information about administering an installed iServer cluster, see Chapter 9, “Clustering,” in *Configuring BIRT iServer*.

---

## Preparing to install an iServer cluster

When you create a BIRT iServer cluster, you must install and run all cluster nodes using the same administrative user account.

### Creating an administrative user account

Before installing iServer, create a user account with the privileges to access the relevant files and directories. Like other Linux and UNIX processes, the processes that perform BIRT iServer tasks run under a specific account.

Actuate recommends creating a dedicated user account for installing and running iServer. Having a dedicated user account isolates iServer-specific issues and events on a machine, making it easier to administer the environment. If you exercise the same control over the operating system account for BIRT iServer that your site exercises for other system administrator and root accounts, you can maintain the same level of security.

Installation of the iServer under the root account is not recommended since the PostgreSQL server must start and be maintained under an unprivileged user ID to prevent system security compromise. If installed under the root account, the default installation will be unable to set up the PostgreSQL schema and Actuate sample Encyclopedia.

## Installing X-frame buffer

Xvfb is an X-Windows server that has neither a graphics card nor a physical graphics display. iServer uses the X-Server for font-rendering information and to generate graphics in documents. Normally, an X-Server requires a graphics card and physical graphics display on the BIRT iServer machine, but you can use Xvfb in place of these components.

Actuate distributes Xvfb for the Sun and AIX operating systems. The installation and configuration of Xvfb are BIRT iServer installation options in these environments.

The Xvfb software installed with iServer includes Type 1 fonts. Actuate maps these fonts to Microsoft Windows fonts for consistent graphics rendering on the various platforms.

The Xvfb software requires:

- X libraries installed on the iServer machine  
If you choose to install Xvfb, the installation script searches for the required libraries and displays a message if the install script cannot find the required libraries on the machine.
- Variables set to the path of the Xvfb libraries
  - XFBDISPLAY variable in start\_srvr.sh.
  - display\_value in pmd11.sh.

To view and print the reports from iServer, you need to set these variables only if you install Xvfb software yourself.

The DISPLAY environment variable specifies the X-Windows server used by the BIRT iServer machine. For example, if the BIRT iServer machine is running X Windows, it sets DISPLAY to the local machine:

```
# setenv DISPLAY :0.0
```

If you use a separate machine as the X-Windows server, specify the machine name in the environment variable DISPLAY. The following example sets DISPLAY to use an X-Windows server on a machine named urup:

```
# setenv DISPLAY urup:0.0
```

The original source code for Xvfb is included as a component of X11R6, but not in earlier X-Window system releases.

---

## Performing a cluster installation using the wizard

The following section describes how to install an iServer Release 11 cluster node in the Linux or UNIX operating system using a Linux system as the example.

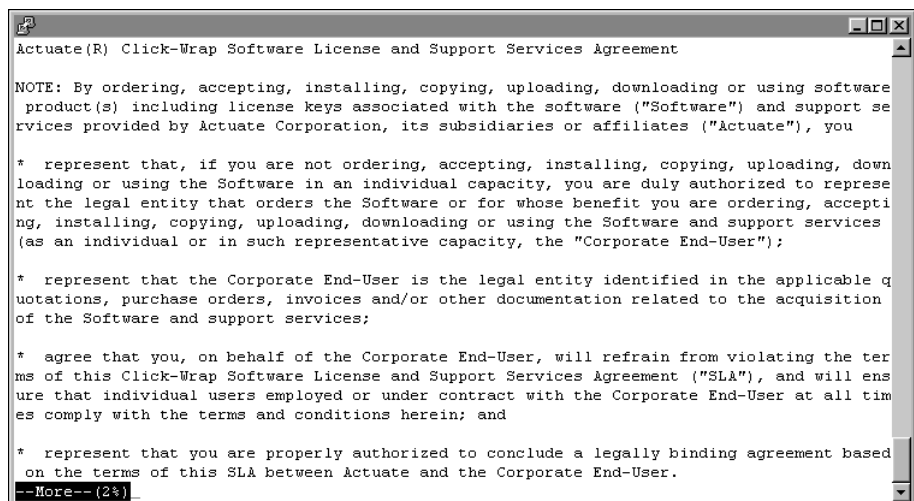
### How to install a cluster node in Linux

- 1 Download the required files from the FTP software distribution site. Extract the files.
- 2 To install the server files, execute the `isinstall` script:

```
sh ./isinstall.sh
```


The script displays a series of prompts. Respond to the prompts as described in the following procedures.

- 3 The license agreement appears, as shown in Figure 5-1.



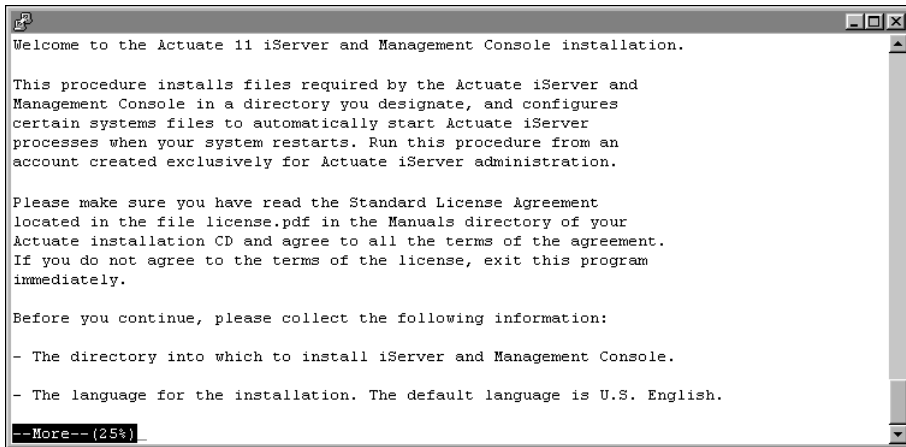
**Figure 5-1** Reviewing the license agreement

- 4 Read the license agreement and press Enter to continue the installation. At the prompt, type `y` for yes if you accept the licensing terms, as shown in Figure 5-2.



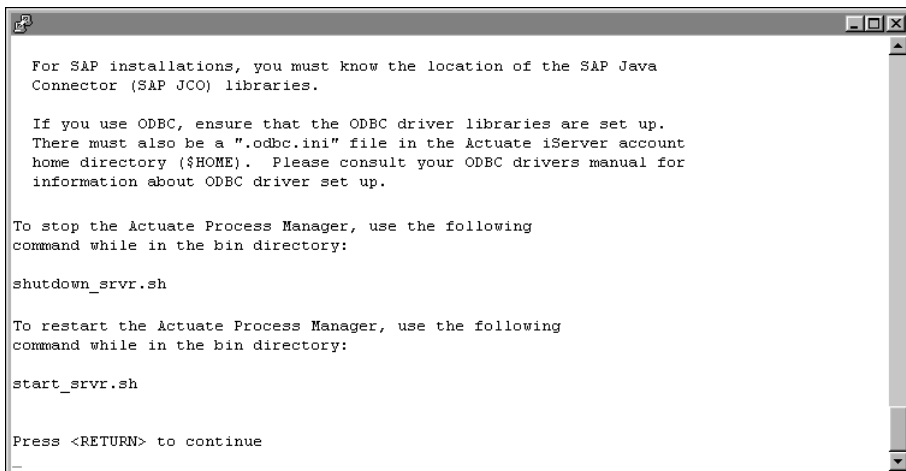
**Figure 5-2** Accepting the licensing terms

- 5 The introduction to the installation appears, as shown in Figure 5-3.



**Figure 5-3** Reviewing the introductory information

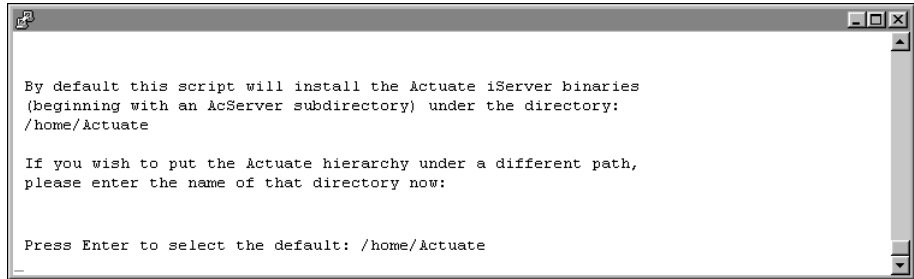
- 6 Press Enter after reviewing the introductory information, as shown in Figure 5-4.



**Figure 5-4** Finishing the review of introductory information

- 7 Press Enter to accept the default location for the installation binaries, as shown in Figure 5-5. Alternatively, type a different directory and press Enter.

The installation program creates the AcServer directory in your chosen location and installs the files.

A terminal window with a title bar containing standard window controls (minimize, maximize, close). The text inside the terminal reads: "By default this script will install the Actuate iServer binaries (beginning with an AcServer subdirectory) under the directory: /home/Actuate". Below this, it says: "If you wish to put the Actuate hierarchy under a different path, please enter the name of that directory now:". At the bottom, it prompts: "Press Enter to select the default: /home/Actuate".

```
By default this script will install the Actuate iServer binaries
(beginning with an AcServer subdirectory) under the directory:
/home/Actuate

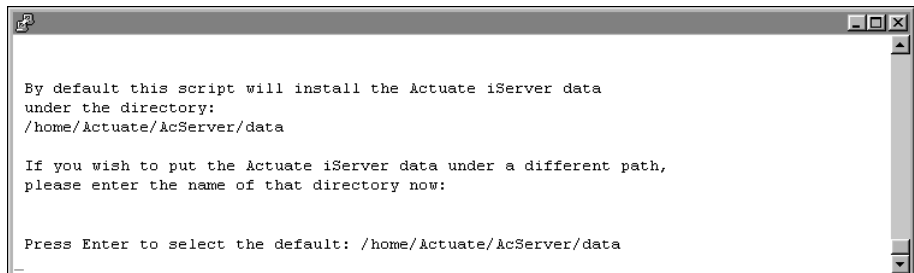
If you wish to put the Actuate hierarchy under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate
```

**Figure 5-5** Specifying the installation directory

- 8 Press Enter to accept the default installation directory for data, `AC_SERVER_HOME/data`, for the iServer data as shown in Figure 5-6. Alternatively, choose a different directory.

iServer uses this data location to store the iServer Encyclopedia volume data, including PostgreSQL metadata, logs, and other files. The default path is `AC_SERVER_HOME/data`, which is referred to in the iServer documentation by the environment variable `AC_DATA_HOME`.

A terminal window with a title bar containing standard window controls. The text inside the terminal reads: "By default this script will install the Actuate iServer data under the directory: /home/Actuate/AcServer/data". Below this, it says: "If you wish to put the Actuate iServer data under a different path, please enter the name of that directory now:". At the bottom, it prompts: "Press Enter to select the default: /home/Actuate/AcServer/data".

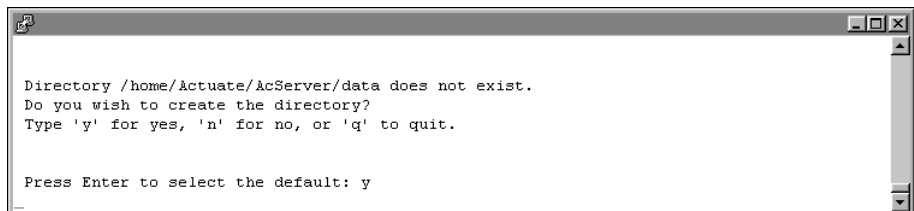
```
By default this script will install the Actuate iServer data
under the directory:
/home/Actuate/AcServer/data

If you wish to put the Actuate iServer data under a different path,
please enter the name of that directory now:

Press Enter to select the default: /home/Actuate/AcServer/data
```

**Figure 5-6** Specifying the data installation directory

- 9 Press Enter to accept the default option of creating the directory for data, as shown in Figure 5-7. Alternatively, type `n` for no, or `q` to quit, and press Enter.


A terminal window with a title bar containing standard window controls. The text inside the terminal reads: "Directory /home/Actuate/AcServer/data does not exist. Do you wish to create the directory? Type 'y' for yes, 'n' for no, or 'q' to quit.". At the bottom, it prompts: "Press Enter to select the default: y".

```
Directory /home/Actuate/AcServer/data does not exist.
Do you wish to create the directory?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```


**Figure 5-7** Creating the `AC_DATA_HOME` directory

- 10 The installer copies prerequisite files to the destination directory, as shown in Figure 5-8. After copying the prerequisite files, the installation continues.



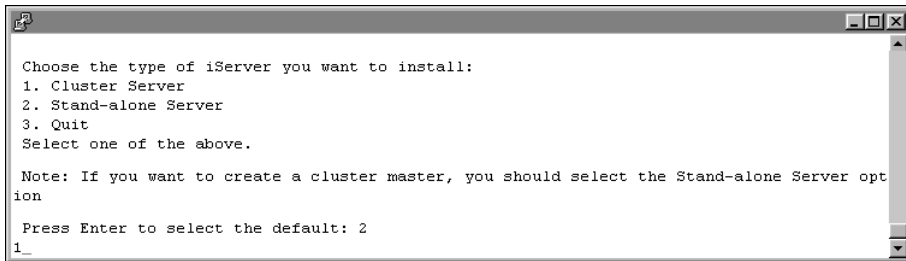
**Figure 5-8** Copying prerequisite files

- 11 Press Enter to choose the default option, Server with Management Console, as shown in Figure 5-9. A cluster node must have access to Configuration Console. Configuration Console installs with Management Console.



**Figure 5-9** Choosing the components to install


- 12 Type 1 and press Enter to select Cluster Server for installation, as shown in Figure 5-10. Alternatively, choose a different type of iServer to install.



**Figure 5-10** Specifying the type of iServer to install


- 13 Specify the path to the configuration home location, and press Enter, as shown in Figure 5-11.

The configuration home location is the shared directory of the cluster containing the files, `acserverconfig.xml`, `acserverlicense.xml`, and `acconfigowner.lock`. In a Release 11 Service Pack 4 installation, the configuration files are located in `AC_DATA_HOME/config/11SP4` by default. For more information about configuring network sharing, see “Adding a node to a cluster,” later in this chapter.



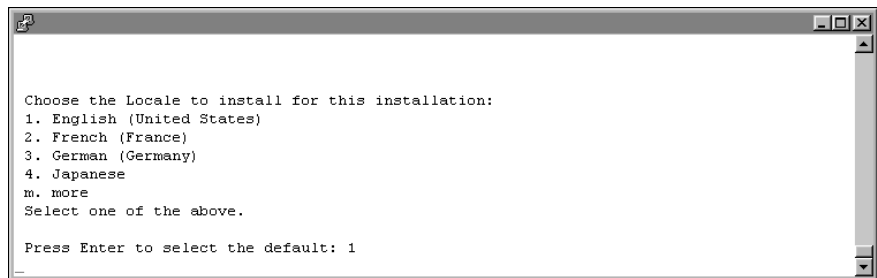
**Figure 5-11** Specifying the configuration home location

- 14 If the specified location for the configuration home directory does not yet exist, you are prompted to create the directory. Press Enter to accept the default option which creates the directory, as shown in Figure 5-12. Alternatively, press n for no, or q to quit.



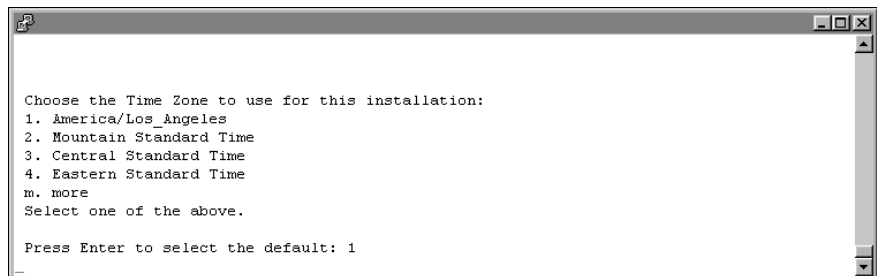
**Figure 5-12** Creating the configuration home location

- 15 If you chose to install Server with Management Console instead of Server only, perform the following steps:
  - 1 Press Enter to select the default locale, which is English (United States), as shown in Figure 5-13. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.




**Figure 5-13** Specifying a locale

- 2 Press Enter to select the default time zone, which is America/Los\_Angeles as shown in Figure 5-14. Alternatively, select another time zone from the numbered list.




**Figure 5-14** Specifying a time zone

- 16 To evaluate the product using the included evaluation software license, press Enter, as shown in Figure 5-15. Alternatively, type 2, then type the path to the license file that you purchased.




**Figure 5-15** Specifying license type

- 17 Press Enter to select the default host name, the name of your machine, where the Process Management Daemon (PMD) runs, as shown in Figure 5-16. Alternatively, type a different IP address or hostname.




**Figure 5-16** Specifying the machine on which the PMD runs

- 18 Press Enter to accept the default port number where the Process Management Daemon (PMD) listens for requests, as shown in Figure 5-17. Alternatively, type a different port number.



**Figure 5-17** Specifying the port number on which the PMD listens


- 19 Specify the administrator password, as shown in Figure 5-18. You use this password to log in to the iServer Configuration Console.



**Figure 5-18** Specifying the administrator password




- 20 Re-enter the password of the administrator, as shown in Figure 5-19.



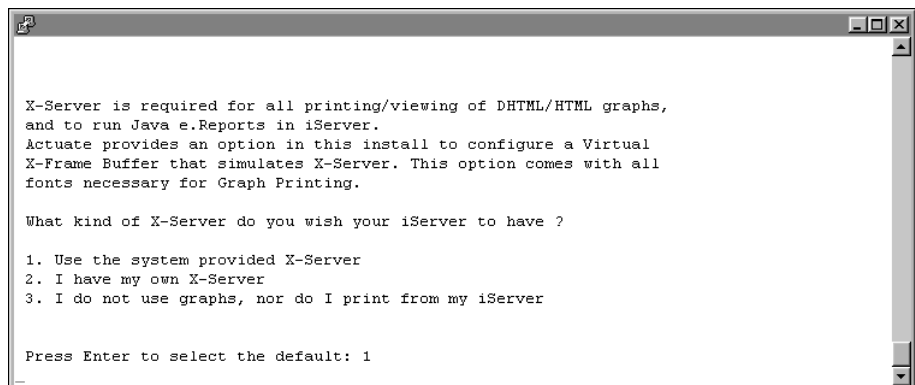
**Figure 5-19** Re-entering the administrator password

- 21 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 5-20. Alternatively type y for yes, specify the database drivers/clients you wish to use and press Enter.



**Figure 5-20** Specifying whether to use database drivers/clients


- 22 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 5-21. For more information about installing X-server, see "Installing X-frame buffer," earlier in this chapter.



**Figure 5-21** Specifying what kind of X-Server to use, if any


**23** If you chose to install Server with Management Console instead of Server only, perform the following steps:

- 1 Press Enter to accept the default hostname, the name of your machine, that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 5-22. Alternatively, type a different IP address.




**Figure 5-22** Specifying the machine host name that Management Console uses to contact the PMD

- 2 Press Enter to accept the default port number, 8100, on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 5-23. Alternatively, type a different port number.




**Figure 5-23** Specifying the port number for the PMD to listen for requests from Management Console

- 3 Press Enter to accept the default hostname, the name of your machine, as shown in Figure 5-24. Alternatively, type a different IP address.




**Figure 5-24** Specifying the name Management Console uses to contact iServer

- 4 Press Enter to accept the default port number, 8000, as shown in Figure 5-25. Alternatively, type a different port number.




**Figure 5-25** Specifying the port number for listening for requests from Management Console

- 5 Press Enter to accept the default name, the name of your machine, for the Encyclopedia volume to use with Management Console, as shown in Figure 5-26. Alternatively, type a different name for the Encyclopedia volume.



**Figure 5-26** Specifying the name of the default Encyclopedia volume


- 6 Press Enter to accept the default name, acadmin, for the HTTP server context root, as shown in Figure 5-27. Alternatively, type a different name.



**Figure 5-27** Specifying the name of the HTTP server context root

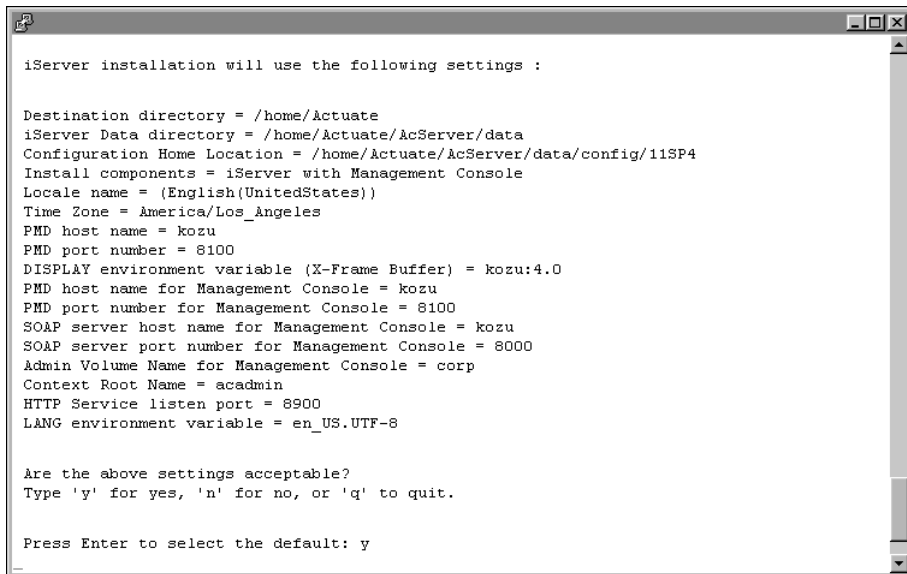
- 24 Press Enter to accept the default port number, 8900, on which the application container listens for requests, as shown in Figure 5-28. Alternatively, choose a different port.

You connect to the port from your browser when accessing various features of iServer.



**Figure 5-28** Specifying the application container listening port number

- 25 Review the settings, as shown in Figure 5-29, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.



```
iServer installation will use the following settings :

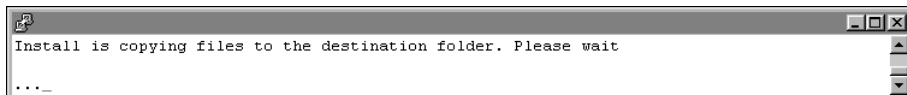
Destination directory = /home/Actuate
iServer Data directory = /home/Actuate/ActServer/data
Configuration Home Location = /home/Actuate/ActServer/data/config/11SP4
Install components = iServer with Management Console
Locale name = (English(UnitedStates))
Time Zone = America/Los_Angeles
PMD host name = kozu
PMD port number = 8100
DISPLAY environment variable (X-Frame Buffer) = kozu:4.0
PMD host name for Management Console = kozu
PMD port number for Management Console = 8100
SOAP server host name for Management Console = kozu
SOAP server port number for Management Console = 8000
Admin Volume Name for Management Console = corp
Context Root Name = acadmin
HTTP Service listen port = 8900
LANG environment variable = en_US.UTF-8

Are the above settings acceptable?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 5-29** Reviewing settings for a Server with Management Console install

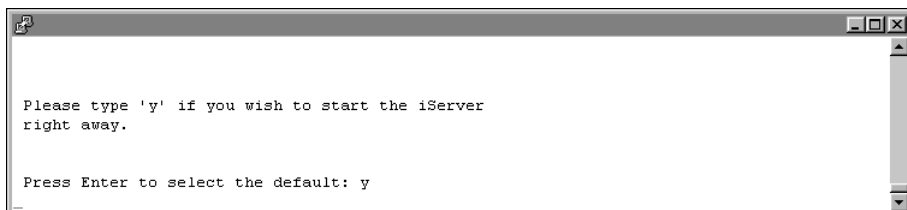
- 26 The installation program installs iServer, and displays an indicator showing how the installation is progressing, as shown in Figure 5-30.



```
Install is copying files to the destination folder. Please wait
..._
```

**Figure 5-30** Copying iServer files to your destination folder

- 27 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start the Process Management Daemon (PMD), as shown in Figure 5-31.

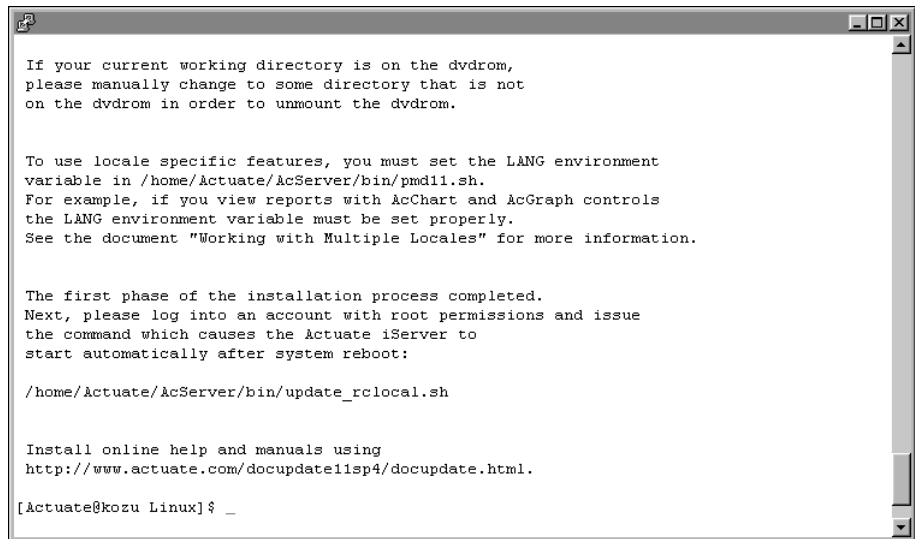


```
Please type 'y' if you wish to start the iServer
right away.

Press Enter to select the default: y
```

**Figure 5-31** Specifying whether to start iServer

- 28 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 5-32.



```
If your current working directory is on the dvdrom,
please manually change to some directory that is not
on the dvdrom in order to unmount the dvdrom.

To use locale specific features, you must set the LANG environment
variable in /home/Actuate/AcServer/bin/pmd11.sh.
For example, if you view reports with AcChart and AcGraph controls
the LANG environment variable must be set properly.
See the document "Working with Multiple Locales" for more information.

The first phase of the installation process completed.
Next, please log into an account with root permissions and issue
the command which causes the Actuate iServer to
start automatically after system reboot:

/home/Actuate/AcServer/bin/update_rclocal.sh

Install online help and manuals using
http://www.actuate.com/docupdate11sp4/docupdate.html.

[Actuate@kozu Linux]$ _
```

**Figure 5-32** Viewing information about localization, logging in, and installing online help

---

## Adding a node to a cluster

After installing a node on a machine, the administrator must still configure sharing and add the node to the cluster. When adding a node to a cluster setup, the administrator must verify that the configuration home directory specified during the install procedure points to the shared configuration home directory and all Encyclopedia volume resources are accessible.

The following section refers to the machine containing the shared configuration directory as node1 and the cluster node accessing these shared resources as node2. The following example assumes that both the configuration folder and Encyclopedia volume folders are located on node1, although in a more complex installation, these configuration and volume resources may reside in another network location.

Before performing a cluster node installation, the Administrator performs the following tasks:

- On node1, the Administrator shares the configuration folder and any Encyclopedia volume folders that a cluster node accesses.

- On node2, the Administrator:
  - Creates folders on which to mount the node1 shared folders
  - Creates a mapping between the node1 and node2 shared folders
  - Mounts the node1 shared folders on the node2 machine

It is the responsibility of the administrator performing the installation to make sure that all settings conform to the security policies in force for the environment.

The following instructions provide a basic reference example of the operations required to configure folder sharing in a Linux environment that supports using the Network File System (NFS), a common, standard, distributed file system protocol.

### **How to share the configuration and Encyclopedia volume files and folders**

In a default iServer Release 11 Service Pack 4 installation, a cluster node requires shared, read-write access to the following system resources:

- AC\_DATA\_HOME/config/11SP4  
In a Release 11 Service Pack 4 installation, the configuration files are located in AC\_DATA\_HOME/config/11SP4.
- AC\_DATA\_HOME/encyc or other volumes, including all file, fileType, status, and tempRov subfolders  
In an iServer installation, where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes.

To give a cluster node read-write access to these files and folders, perform the following tasks:

- 1 Log in to node1 as the root user.
- 2 Add the following entries to the /etc/exports file:
 

```
/home/actuate/AcServer/data/config/11SP4
  *(rw,fsid=1,no_root_squash)
/home/actuate/AcServer/data/encyc
  *(rw,fsid=2,no_root_squash)
```
- 3 Start the NFS server processes by executing the following command:
 

```
service nfs restart
```
- 4 Log in to node2 as the actuate user.
- 5 Create the following directory paths:
 

```
/home/actuate/AcServer/data/config/11SP4
/home/actuate/AcServer/data/encyc
```

6 Log off node2.

7 Log in to node2 as the root user.

8 Add the following entries to the `/etc/fstab` file:

```
<node1 hostname>:/home/actuate/AcServer/data/config/11SP4
/home/actuate/AcServer/data/config/11SP4 nfs nfsvers=3 0 0
<node1 hostname>:/home/actuate/AcServer/data/encyc /home/
actuate/AcServer/data/encyc nfs nfsvers=3 0 0
```

9 Mount the shared folders on node1 by executing the following commands:

```
mount /home/actuate/AcServer/data/config/11SP4
mount /home/actuate/AcServer/data/encyc
```

The administrator must also verify or edit the shared `acpmdconfig.xml` file to contain the following information:

- `<AC_CONFIG_HOME>` to point to the shared configuration home directory for the cluster
- `<AC_TEMPLATE_NAME>` to specify the server template from the available server templates listed in the shared `acserverconfig.xml` file

#### **How to verify and edit `acpmdconfig.xml` file settings**

To verify and edit `acpmdconfig.xml` file settings, perform the following tasks:

1 Shut down the recently installed cluster node.

2 Using a text editor, open `acpmdconfig.xml`, which by default is located in `AC_SERVER_HOME/etc`.

3 Verify or edit `<AC_CONFIG_HOME>` to point to the shared configuration home directory for the cluster, as shown in the following code:

```
<AC_CONFIG_HOME>/home/actuate/AcServer/data/config/11SP4
</AC_CONFIG_HOME>
```

This location is the path that you specified for the configuration home directory during the install procedure.

4 Verify or edit `<AC_TEMPLATE_NAME>` to specify the server template name from the available server templates listed in the shared `acserverconfig.xml` file, as shown in the following code:

```
<AC_TEMPLATE_NAME>urup</AC_TEMPLATE_NAME>
```

In the example, `urup` is server template name

5 Save `acpmdconfig.xml`.

The administrator must also verify or edit the shared `acserverconfig.xml` file to contain the following information:

- `<ServerFileSystemSetting>` points to the shared drive location that contains the Encyclopedia volume data files
- `server <ConnectionProperty>` specifies the network name of the node that contains the shared Encyclopedia volume database

### How to verify and edit `acserverconfig.xml` file settings

To verify and edit `acserverconfig.xml` file settings, perform the following tasks:

- 1 Stop the Actuate BIRT iServer 11 service running on the node that contains the shared configuration home directory.
- 2 Using a text editor, open the `acserverconfig.xml` file in the configuration home directory.

In a Release 11 Service Pack 4 installation, the configuration files are located in `AC_DATA_HOME/config/11SP4` by default. The location is the path that you specified for the configuration home directory during the install procedure.

- 3 In `<Template>` settings for the node, verify or edit `<ServerFileSystemSettings>` to make sure the path `<ServerFileSystemSetting>` points to the location that contains the Encyclopedia data files, by performing the following tasks:

- 1 Locate the `<ServerFileSystemSettings>` element under the `<Template>` element.

- 2 In `<ServerFileSystemSetting>`, locate:

```
<ServerFileSystemSettings>
  <ServerFileSystemSetting
    Name="DefaultPartition"
    Path="$AC_DATA_HOME$/encyc"/>
</ServerFileSystemSettings>
```

- 3 Change Path from the `AC_DATA_HOME` variable notation to the full path specification, as shown in the following code:

```
<ServerFileSystemSettings>
  <ServerFileSystemSetting
    Name="DefaultPartition"
    Path="/home/actuate/AcServer/data/encyc"/>
</ServerFileSystemSettings>
```

The Path setting for DefaultPartition is `/home/actuate/AcServer/data/encyc`. Do not use the `AC_DATA_HOME` variable notation.

- 4 In `<MetadataDatabase>` settings, verify or edit the `<ConnectionProperty>` for the server to make sure that it specifies the network name, not localhost, of the node on which the Encyclopedia volume database resides, by performing the following tasks:



- 1 Locate the <ConnectionProperties> element under the <MetadataDatabase> element.

- 2 In <ConnectionProperties> locate:

```
<ConnectionProperty  
  Name="server"  
  Value="localhost"/>
```

- 3 Change Value from localhost to the name of the machine on which the Encyclopedia volume database resides, such as urup, as shown in the following code:

```
<ConnectionProperty  
  Name="server"  
  Value="urup"/>
```

- 5 Save acserverconfig.xml.

Start Actuate BIRT iServer 11 on each cluster node. The new cluster node will automatically read the settings in the acserverconfig.xml file in the shared configuration directory to access its template, then join the cluster.

#### **How to start an iServer cluster using Configuration Console**

To start iServer using Configuration Console manually, perform the following tasks:

- 1 On the node containing the configuration home directory for the cluster, log in to Configuration Console and choose Advanced view. Choose Servers, then choose Start New Server.
- 2 On Servers—Start New Server, as shown in Figure 5-33, perform the following tasks:
  - 1 In Server name, type the name of the cluster node.
  - 2 In Host Name or IP Address, type the name or IP address of the cluster node.
  - 3 In iServer Process Manager Port Number, type the Daemon listen port number. The default value for this port is 8100. You specify this port number during the install procedure.
  - 4 In Server template name, choose the name of the template that the cluster node uses.Choose OK.

Servers > Start New Server

Server name:

Host Name or IP Address:  \*

iServer Process Manager Port Number:  \*

Server template name:  \*

\*These fields are required and cannot be left blank  
 These fields require server restart to take effect

OK Cancel

**Figure 5-33** Preparing to start a new server

- 3 Log out of Configuration Console.
- 4 Restart the Actuate BIRT iServer 11 services on the node containing the configuration home directory for the cluster then the new node.
- 5 Log in to Configuration Console and choose Advanced view. Choose Servers from the side menu. The new cluster node automatically reads the `acserverconfig.xml` in the shared configuration home directory to access its template, then joins the cluster.

---

## Finding the BIRT iServer home directory

The environment variable for the iServer home directory is `AC_SERVER_HOME`. The iServer installation program sets the variable to the path of your iServer login environment.

### How to find the home directory for BIRT iServer on a Linux or UNIX system

If you use the C Shell on a Linux or UNIX system, you can look in your `.cshrc` file to see the value of `AC_SERVER_HOME`:

```
$ setenv AC_SERVER_HOME /usr/local/AcServer
```

If an iServer process is running on the system, you can also use the following `ps` command piped to the `grep` command to find the `pmd11` executable, which runs from the iServer home `bin` directory:

```
# ps -ef | grep pmd
```

---

## About the Java Development Kit

The BIRT iServer installation routine installs the Java SDK files under the directory specified in the environment variable `AC_SERVER_HOME`:

```
AC_SERVER_HOME/jdk160
```

Some operating systems require an operating system upgrade or patch to use JRE 6.0. For information about requirements for your operating system, see your operating system documentation. Also, see the Actuate Support Lifecycle Policy and Supported Products Matrix on the Actuate Support web site. You can access the Support site at the following URL:

<http://support.actuate.com/documentation/spm>

To use a different JDK with iServer, change the files in the installation directory or change the values of the following Linux and UNIX environment variables:

- `AC_JAVA_HOME`
- `AC_JVM_HOME`
- `AC_JRE_HOME` in the PMD startup script, `pmd11.sh`
- `AC_JRE64_HOME`

Using an earlier release of JDK can cause some Actuate features to fail or to work improperly. For example, using an earlier release of JDK can cause Actuate products to display Actuate charts incorrectly.



# Installing BIRT iServer in a cloud

This chapter contains the following topics:

- Understanding a BIRT iServer cloud deployment
- Deploying BIRT iServer in a cloud environment
- Setting up iServer to use the out-of-the-box (OOTB) PostgreSQL database
- Setting up iServer to use an alternative database
- Setting up iServer only
- Setting up iServer to join an existing cluster
- Understanding the cloud deployment environment

---

## Understanding a BIRT iServer cloud deployment

A cloud deployment in the BIRT iServer environment has the following features:

- Easy-to-prepare, stateless iServer image  
This image of an installed iServer run-time environment does not require modification during installation or the life time of the instance. The administrator can create a customized image by generating an archive of an installed iServer run-time environment. Alternatively, an out-of-the-box (OOTB) image is available as a separate iServer distribution package for Windows and Linux. The administrator deploys the image by unbundling the archive and running an installation script or installing a virtual image on the target machine.
- Ready-to-launch iServer instance  
Launching an iServer instance requires minimal work. After installing and configuring the appropriate environment variables, the administrator launches the deployed iServer image from the command line using a script to execute the commands.
- Elastic iServer clustering  
The use of a ready-to-launch iServer image simplifies iServer cluster installation and management. Nodes with the same cluster ID, running on the same sub-net, automatically detect and join to form the cluster. The cluster automatically detects the on-off status of any node. Single-point node failure does not affect the availability of the other nodes.

In Release 11 and earlier, iServer used multicasting to broadcast event information and synchronize operations in a cluster. Some cloud computing environments do not support multicasting. Starting in Release 11 Service Pack 1 and later, iServer uses the third-party RDBMS as a shared repository for storing cluster information. This enhancement replaces multicasting as a way of managing the cluster environment.

For more information on setting up a cluster after performing a cloud deployment of a stand-alone iServer installation, see “Setting up iServer to join an existing cluster,” later in this chapter; Chapter 5, “Installing a BIRT iServer cluster,” earlier in this book; and Chapter 9, “Clustering,” in *Configuring BIRT iServer*. For more information on iServer architecture, see “Understanding Actuate BIRT iServer architecture,” earlier in this book.

---

## Deploying BIRT iServer in a cloud environment

In an Actuate cloud deployment, a stateless iServer image contains only the run-time environment. The administrator typically transfers the image of the

iServer run-time environment using a compressed archive, such as a TAR file, or virtual image, unbundling the image on the target machine.

In addition, the administrator must install a supported version of the JAVA SE Development Kit (JDK) 1.6 or earlier. If not already installed on the machine, the JDK can be downloaded from the following location:

<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

Set the `JDK_HOME` environment variable to point to the location of the JDK.

If you have an earlier version of BIRT iServer installed on your machine, such as Release 10 Service Pack 1, you can continue to run the earlier version, but not simultaneously with the new version if the earlier version uses the same default ports. You must shut down the earlier version during the deployment process. The earlier version must remain shut down when the newly installed iServer is running. Reconfigure the ports for one of the versions to run both versions at the same time.

In an iServer Release 11 configuration, Actuate recommends storing iServer data in a directory located outside the iServer run-time environment. In a default Linux setup performed using the wizard-based install program, the iServer run-time environment installs in the following directory:

```
$HOME/AcServer
```

The data installs in the following directory:

```
$HOME/AcServer/data
```

In a cloud deployment, Actuate recommends installing in an alternative directory. For example, in Linux, install the run-time environment in the following directory:

```
/home/Actuate/Actuate11/AcServer
```

The data installs in the following directory:

```
/home/Actuate/Actuate11/AcServer/data
```

The environment variable `AC_SERVER_HOME` points to the directory containing the run-time environment. The environment variable `AC_DATA_HOME` points to the directory containing the iServer data.

## Deploying an iServer distribution package

In deploying an iServer distribution package, the administrator performs the following tasks:

- Extracts the contents of the iServer distribution package
- Installs a supported Java Development Kit (JDK)
- Runs the iServer setup script, installing iServer using an evaluation license

## How to extract the contents of the iServer distribution package

To extract the iServer run-time resources and configure the setup script, perform the following tasks:

- 1 Create a new folder in a location outside of \$HOME/AcServer, such as \$HOME/Actuate11.
- 2 Extract the contents of ActuateBIRTiServer.tar.gz to the folder created in the previous step.
- 3 In the iServer Release 11 configuration, the setting for the environment variable, AC\_SERVER\_HOME, specifies the location of the iServer program files. The variable, AC\_JAVA\_HOME, specifies the location of the Java Development Kit (JDK.)

The setupiServer.sh script automatically detects the location of AC\_SERVER\_HOME and AC\_JAVA\_HOME in most cases. If the script is unable to locate these resources in the environment, the script provides an error message. In this case, you can edit the script and set the values for these properties manually.

Specify the location of AC\_SERVER\_HOME and AC\_JAVA\_HOME in the setupiServer script by performing the following tasks:

- 1 Using a text editor, open the script, setupiServer.sh, located in the folder where you extracted BIRT iServer 11.
- 2 Specify the paths to AC\_SERVER\_HOME and AC\_JAVA\_HOME, as shown in Listing 6-1.
- 3 Save and close the file.

### Listing 6-1 setupiServer script

---

```
#!/bin/sh
# Script that sets up the iServer and starts the iServer

#Title Actuate BIRT iServer 11 setup and start script for
  evaluation

AC_SERVER_HOME=/home/Actuate/Actuate11/AcServer
AC_JAVA_HOME=/home/JDK160
...
```

## Running the setup script

The setup script provides the following stand-alone and cluster installation options:



- Stand-alone
  - Set up iServer to use the out-of-the-box (OOTB) PostgreSQL database. Sets up iServer and an embedded out-of-the-box (OOTB) PostgreSQL database for storing iServer system and Encyclopedia volume metadata.
  - Set up iServer to use an alternative database. Sets up iServer and an external database, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL database, for storing iServer system and Encyclopedia volume metadata. Requires superuser or database system administrator access to the external database.
  - Set up iServer only. Sets up only the iServer program files. This option requires logging in to Configuration Console after installing iServer to set up a connection to the database used for storing iServer system and Encyclopedia volume metadata. Requires the database administrator to run SQL Data Definition Language (DDL) scripts to create a database and the following schema owner and application user accounts with appropriate privileges:
    - iServer system schema owner
    - Encyclopedia volume schema owner
    - iserver application user

For more information about creating the iServer system and Encyclopedia volume schemas and iserver user in an alternative database, see Chapter 3, “Installing BIRT iServer using an alternative database.”
- Cluster
  - Set up iServer to join an existing cluster. Sets up an iServer node on an existing cluster to add more capacity. The setup prompts the administrator for the location of the shared configuration files.

#### How to run the setup script

- 1 Navigate to the folder where you extracted the iServer package.
- 2 Execute `setupiServer.sh` by entering the following command:
 

```
sh ./setupiServer.sh
```
- 3 The script prompts you to choose one of the following stand-alone or cluster options in setting up iServer, as shown in Figure 6-1:
  - Stand-alone
    - 1. Set up iServer to use the out-of-the-box (OOTB) PostgreSQL database.

- ❑ 2. Set up iServer to use an alternative database.  
Follow instructions under iServer\tools\install\readme before selecting this choice.
- ❑ 3. Set up iServer only.  
Use this option to setup iServer only and add an external database connection later using Configuration Console.
- Cluster:
  - ❑ 4. Set up iServer to join an existing cluster.
  - ❑ 5. Abort to terminate the operation.

```
[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setupiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?

Standalone:
 1. Set up iServer to use the out-of-the-box PostgreSQL database.
 2. Set up iServer to use an alternative database.
    Follow instructions under AcServer/tools/install/readme before selecting this choice
 .
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connection later
using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.

 5. Abort to terminate this operation.

Default 1 : _
```

**Figure 6-1** Choosing setup type

The following sections describe how to perform an installation for each option.

---

## Setting up iServer to use the out-of-the-box (OOTB) PostgreSQL database

The following section describes how to set up iServer to use the out-of-the-box (OOTB) PostgreSQL database.

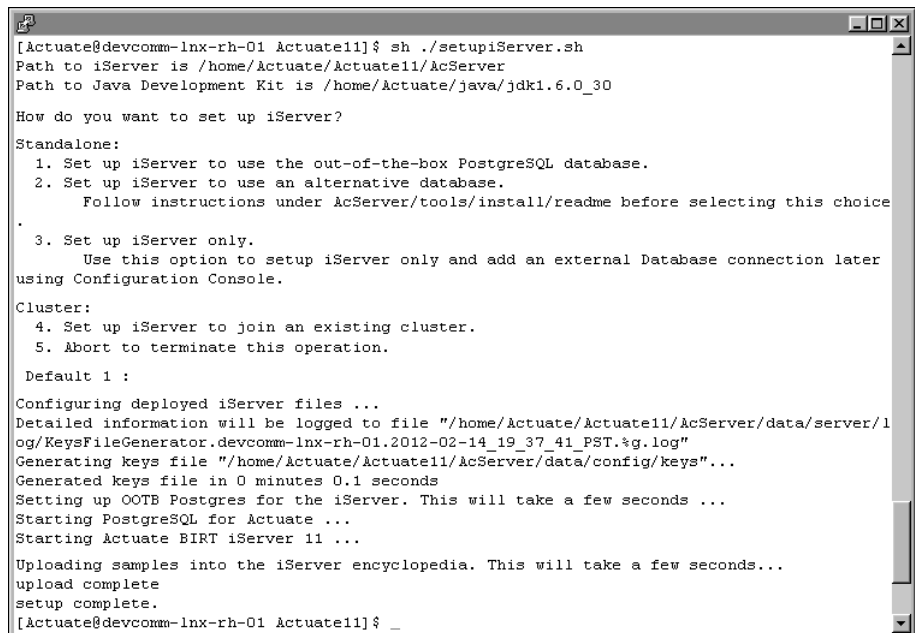
### How to set up iServer to use the out-of-the-box (OOTB) PostgreSQL database

After extracting the contents of ActuateBIRTiServer.tar.gz and running the setup script as described in “How to extract the contents of the iServer distribution package,” and “How to run the setup script,” earlier in this chapter, perform the following tasks:

- 1 In the `setUpiServer` menu, press `Enter` to choose default option 1, Set up iServer to use the out-of-the-box (OOTB) PostgreSQL database, as shown in Figure 6-1.

The script performs the following tasks, as shown in Figure 6-2:

- Sets up the iServer deployment files, including log and security keys files
- Installs and starts the OOTB PostgreSQL database system used to store Encyclopedia volume metadata
- Starts BIRT iServer 11
- Creates the Encyclopedia volume
- Uploads the Encyclopedia volume sample content



```
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setUpiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?
Standalone:
  1. Set up iServer to use the out-of-the-box PostgreSQL database.
  2. Set up iServer to use an alternative database.
     Follow instructions under AcServer/tools/install/readme before selecting this choice
  3. Set up iServer only.
     Use this option to setup iServer only and add an external Database connection later
     using Configuration Console.
Cluster:
  4. Set up iServer to join an existing cluster.
  5. Abort to terminate this operation.

Default 1 :
Configuring deployed iServer files ...
Detailed information will be logged to file "/home/Actuate/Actuate11/AcServer/data/server/1
og/KeysFileGenerator.devcomm-lnx-rh-01.2012-02-14_19_37_41_PST.%g.log"
Generating keys file "/home/Actuate/Actuate11/AcServer/data/config/keys"...
Generated keys file in 0 minutes 0.1 seconds
Setting up OOTB Postgres for the iServer. This will take a few seconds ...
Starting PostgreSQL for Actuate ...
Starting Actuate BIRT iServer 11 ...

Uploading samples into the iServer encyclopedia. This will take a few seconds...
upload complete
setup complete.
[Actuate@devcomm-lnx-rh-01 Actuate11]$ _
```

**Figure 6-2** Completing the iServer setup

## Accessing Information, Management, and Configuration Consoles

After the script finishes running, open a browser to log in to the following BIRT iServer 11 consoles to perform user and administrator tasks:

- **Information Console**  
Perform tasks such as accessing folders and viewing designs and documents.

To access Information Console, open a browser manually and enter the following URL, as shown in Figure 6-3:

`http://localhost:8900/iportal/`

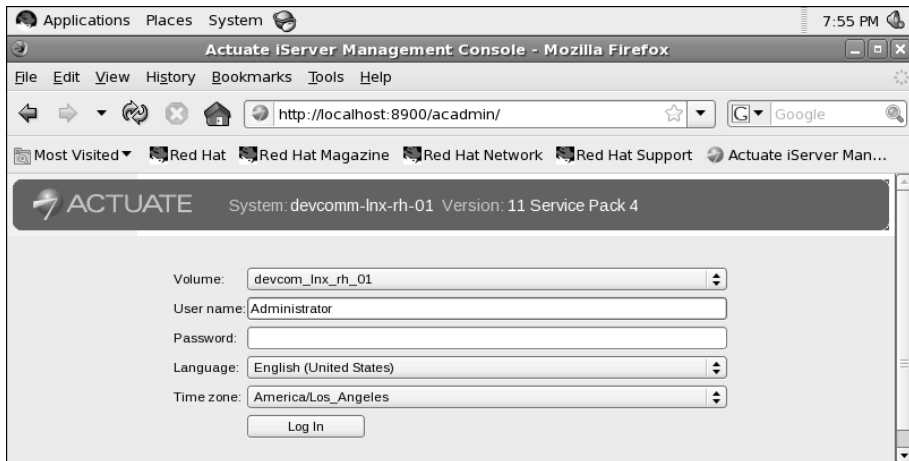


**Figure 6-3** Viewing Welcome to Actuate Information Console

- Management Console  
Set up user accounts and schedule or run a design.

To access Management Console, open a browser manually and enter the following URL, as shown in Figure 6-4:

`http://localhost:8900/acadmin/`

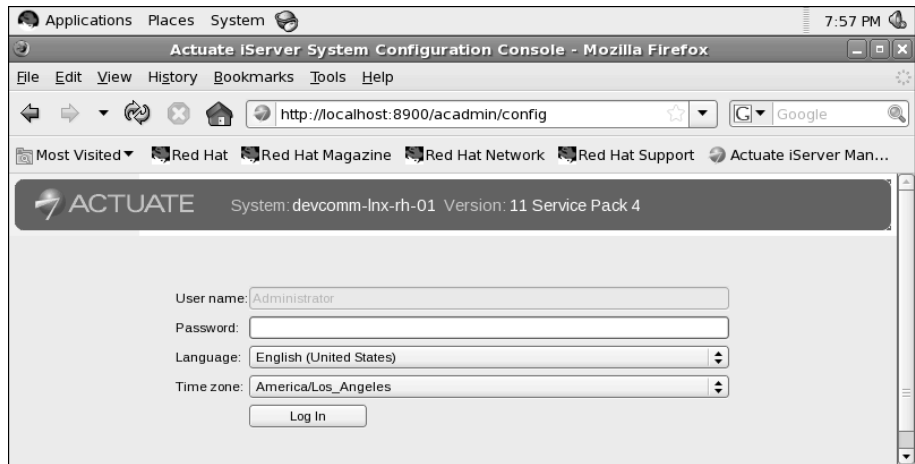


**Figure 6-4** Logging in to Management Console

- Configuration Console
  - Perform administrative operations, such as the following tasks:
    - Add an Encyclopedia volume.
    - Connect to a database.
    - Make modifications to iServer parameters and server templates.
    - Update the license.

To access Configuration Console for administering iServer, open a browser manually and enter the following URL, as shown in Figure 6-5:

`http://localhost:8900/acadmin/config/`



**Figure 6-5** Logging in to Configuration Console

When starting PostgreSQL for Actuate iServer and Actuate BIRT iServer, the best practice is to start PostgreSQL then iServer. When stopping these programs, stop iServer then PostgreSQL, if necessary.

It is not necessary to shut down the database when starting and stopping iServer. iServer starts faster with the database already running and available in the background, which is particularly important in cluster and other high-performance operations.

The following sections describe how to perform these operations in the recommended order.

## Stopping and starting iServer and PostgreSQL processes

After installing iServer, an administrator can stop and start iServer and PostgreSQL processes from a command prompt using scripts provided in the installation.

### How to stop and start BIRT iServer

To stop iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./shutdown_srvr.sh
```

To start iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./startsrvr.sh
```

### How to stop and start PostgreSQL for Actuate iServer

To stop PostgreSQL, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin
- 2 Type the following command and press Enter:

```
sh ./stoppostgresql.sh
```

To restart PostgreSQL for Actuate iServer, perform the following tasks:

- 1 Open a command prompt and navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter:

```
sh ./startpostgresql.sh
```

---

## Setting up iServer to use an alternative database

The following section describes how to install and configure the iServer deployment bundle to use an alternative database.

### How to prepare the installation environment

After extracting the contents of ActuateBIRTiServer.tar.gz as described in “How to extract the contents of the iServer distribution package,” earlier in this chapter, prepare the installation environment before running the setup script by performing the following tasks:

- 1 In a command prompt, navigate to AC\_SERVER\_HOME\tools\install. For example:

```
/home/Actuate/Actuate11/AcServer/tools/install
```

- 2 Using a text editor, open the readme.txt file that contains the preliminary setup instructions, as shown in Listing 6-2.

**Listing 6-2** The readme.txt file

---

1. Install database client software onto the iServer node.
  2. Configure the database client to access the database instance to which you want to install the encyclopedia. In some cases, you may need to set the correct environment variables in the command window from which you run the script. For example, to run the scripts on Unix with DB2 as the database, you will need to configure environment variables such as DB2INSTANCE, INSTHOME, PATH and LD\_LIBRARY\_PATH to make db2 commands accessible.
  3. Create a "lib" folder under \$AC\_SERVER\_HOME/tools/install. Copy JDBC driver jar from database client to this "lib" folder.
    - For Oracle database copy ojdbc14.jar
    - For DB2 database, copy db2jcc.jar
    - For SQL Server database, copy sqljdbc4.jar
    - For PostgreSQL database, copy postgresql-8.4-703.jdbc4.jar
  4. Edit the install.properties file to add database connection properties and other required properties.
  5. The script is called by other scripts, for example startiServer\_Non\_EmbeddedDB.bat. It can also be executed manually as below:

```
ant -f install.xml install
```
- 3 Following the instructions in readme.txt, perform the following tasks:
    - 1 Install an alternative database, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL server.
    - 2 Run a SQL script containing the appropriate Data Definition Language (DDL) statements to create the iserver database. For example, when creating the database in a PostgreSQL server, run the following DDL commands:
      - 1 Connect to the PostgreSQL system database as a user with full administrator privileges, typically named postgres, and execute the following SQL commands to create a database named iserver:

```
CREATE DATABASE iserver
  WITH OWNER = "postgres"
  TEMPLATE = template0 ENCODING = 'UTF-8';
REVOKE ALL ON DATABASE iserver FROM PUBLIC;
```

- 2 In the iserver database, create the plpgsql procedural language by executing the following SQL command:

```
CREATE LANGUAGE plpgsql;
```

Plpgsql is a superset of PostgreSQL that supports advanced programming features, such as variables, conditional expressions, iterative constructs, and events. If the language is already installed, an error message appears. If so, ignore the message.

When the setupiServer script runs, as described in “How to set up iServer to use an alternative database,” later in this section, the script creates the following schema owner and application user accounts with appropriate privileges:

- ❑ iServer system schema owner
- ❑ Encyclopedia volume schema owner
- ❑ iserver application user

For more information about creating the iServer system and Encyclopedia volume schemas and iserver user in an alternative database, see Chapter 3, “Installing BIRT iServer using an alternative database.”

- 3 Create a lib folder in AC\_SERVER\_HOME/tools/install, and copy the JDBC driver JAR file specified for the database to the lib folder.
- 4 In a text editor, open the install.properties file, and specify all required and any necessary optional settings, as shown in Listing 6-3.

The install.properties file requires settings for the following properties:

- ❑ AC\_SERVER\_HOME  
iServer home folder, such as /home/Actuate/Actuate11/AcServer.
- ❑ SYSTEM\_NAME  
Set up automatically by the installation script. Do not change this property. Restrict system, schema, and the iServer application user names to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]\*. Do not use a hyphen.
- ❑ DEFAULT\_DATABASE\_PASSWORD  
Default database server administrator (DBA) or superuser password if the user does not specify a password at the script prompt.
- ❑ Database-specific properties:
  - ❑ jdbc.dbtype  
Server type, such as DB2, Microsoft SQL Server, Oracle, or PostgreSQL.



- ❑ jdbc.serverName  
Database server name.
- ❑ jdbc.portNumber  
Database server port number.
- ❑ jdbc.databaseName  
Database name, such as iserver.
- ❑ dba.name  
Database administrator (DBA) name.
- ❑ dba.password  
Database administrator (DBA) password. The installation script prompts for the password, so this value can be left at the default setting.

Listing 6-3 specifies the required property settings for a PostgreSQL database as an example.

**Listing 6-3** The install.properties file

---

```
#Please specify the follow required properties:
  AC_SERVER_HOME, AC_DATA_HOME and AC_CONFIG_HOME
#for the iServer install. Do not use backslash ("\") in the
  path. Always use forward slash ("/")
#on both Windows, Linux, or Unix.

#REQUIRED
#AC_SERVER_HOME, for example: /home/Actuate/Actuate11/
  AcServer
AC_SERVER_HOME=/home/Actuate/Actuate11/AcServer

#OPTIONAL
#AC_DATA_HOME, for example: ${AC_SERVER_HOME}/data

#OPTIONAL
#AC_CONFIG_HOME, for example: ${AC_SERVER_HOME}/data/config

#REQUIRED
#Set up by the script automatically. User should not change
  it.
SYSTEM_NAME=#AC_SYSTEM_NAME#

#REQUIRED
#Database password that Installer asked for, which will be
  applied to all DB related passwords, if user doesn't
  specify.
```

```

DEFAULT_DATABASE_PASSWORD=xxxxxxx

#OPTIONAL
#APPLICATION_USER_PASSWORD=${DEFAULT_DATABASE_PASSWORD}

#OPTIONAL
#SYSTEM_SCHEMA_PASSWORD=${DEFAULT_DATABASE_PASSWORD}

#OPTIONAL
#VOLUME_SCHEMA_PASSWORD=${DEFAULT_DATABASE_PASSWORD}

#
#Specify database specific properties. The database types
#supported are Microsoft SQL Server, Oracle, DB2, and
#PostgreSQL.
...
jdbc.dbtype=PostgreSQL
jdbc.serverName=localhost
jdbc.portNumber=8432
jdbc.databaseName=iserver

#define the target database
dba.name=postgres
dba.password=xxxxxxx
...

```

### How to set up iServer to use an alternative database

After preparing the installation environment, run the setup script as described in “How to run the setup script” earlier in this chapter, and perform the following installation tasks:

- 1 In the setupiServer menu, choose option 2, Set up iServer to use an alternative database.

The script performs the following tasks, as shown in Figure 6-6:

- Sets up the iServer deployment files, including log and security keys files
- Sets up and starts BIRT iServer 11
- Creates the Encyclopedia volume
- Uploads the Encyclopedia volume sample content

```
[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setupiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?

Standalone:
  1. Set up iServer to use the out-of-the-box PostgreSQL database.
  2. Set up iServer to use an alternative database.
     Follow instructions under AcServer/tools/install/readme before selecting this choice
  3. Set up iServer only.
     Use this option to setup iServer only and add an external Database connection later
     using Configuration Console.

Cluster:
  4. Set up iServer to join an existing cluster.
  5. Abort to terminate this operation.

Default 1 : 2_
```

**Figure 6-6** Completing the iServer alternative database setup

After the script finishes running, open a browser manually and enter the following URL, as shown in Figure 6-7:

<http://localhost:8900/iportal/>

Log in to Information Console to perform tasks such as accessing folders and viewing designs and documents.



**Figure 6-7** Viewing Welcome to Actuate Information Console

For more information about accessing iServer consoles, see “Accessing Information, Management, and Configuration Consoles,” earlier in this chapter.

---

## Setting up iServer only

The following section describes how to install and configure iServer only and add the external database connections later using Configuration Console. In this option, the administrator performs the following installation tasks:

- Sets up iServer run-time environment only using the cloud deployment package
- Configures an external database to store iServer system and Encyclopedia volume metadata
- Uses Configuration Console to configure a database connection by specifying the properties for the following items:
  - Metadata database
  - System and Encyclopedia volume schemas and iserver application user
  - Data partition and Encyclopedia volume

## Setting up iServer only using the cloud deployment package

After extracting the contents of `ActuateBIRTiServer.tar.gz` and running the setup script as described in “How to extract the contents of the iServer distribution package,” and “How to run the setup script,” earlier in this chapter, perform the following tasks.

### How to set up iServer only

- 1 In the `setupiServer` menu, choose option 3, Set up iServer only.

The script performs the following tasks, as shown in Figure 6-8:

- Sets up the iServer deployment files, including log and security keys files
- Sets up and starts BIRT iServer 11

```

[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setupiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?

Standalone:
  1. Set up iServer to use the out-of-the-box PostgreSQL database.
  2. Set up iServer to use an alternative database.
     Follow instructions under AcServer/tools/install/readme before selecting this choice
  3. Set up iServer only.
     Use this option to setup iServer only and add an external Database connection later
     using Configuration Console.

Cluster:
  4. Set up iServer to join an existing cluster.
  5. Abort to terminate this operation.

Default 1 : 3

Configuring deployed iServer files ...
Detailed information will be logged to file "/home/Actuate/Actuate11/AcServer/data/server/1
og/KeysFileGenerator.devcomm-lnx-rh-01.2012-02-22_17_51_21_PST.%g.log"
Generating keys file "/home/Actuate/Actuate11/AcServer/data/config/keys"...
Generated keys file in 0 minutes 0.0 seconds
Setting up iServer. This will take a few seconds ...
Starting Actuate BIRT iServer 11 ...

setup complete.
[Actuate@devcomm-lnx-rh-01 Actuate11]$ _

```

**Figure 6-8** Completing the iServer only setup

2 After installing option 3, stop and restart iServer by performing the following tasks:

- 1 Navigate to AC\_SERVER\_HOME/bin.
- 2 Type the following command and press Enter, as shown in Figure 6-9:

```
sh ./startsrvr.sh
```

```

[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ cd AcServer/bin
[Actuate@devcomm-lnx-rh-01 bin]$ sh ./startsrvr.sh

Starting startsrvr at Wed 02/22/2012 18:07:45

AC_SERVER_HOME=/home/Actuate/Actuate11/AcServer
Path to the Java Development Kit is /home/Actuate/java/jdk1.6.0_30
Calling CheckJavaVersion to validate required java version
Checking java version: Required minimum java version is 1.6.0
Checking java version: Current java version is 1.6.0_30
Checking java version: 1.6.0_30 is okay to run the Actuate 11 iServer
Version written to /home/Actuate/Actuate11/AcServer/javaversion.txt file: 1.6.0_30
Status written to /home/Actuate/Actuate11/AcServer/javaversion.txt file: ok
Completed startsrvr
[Actuate@devcomm-lnx-rh-01 bin]$ _

```

**Figure 6-9** Starting Actuate BIRT iServer 11

## Configuring an external database

The database administrator must install an external database then run SQL scripts containing the Data Definition Language (DDL) statements to create the following database objects with appropriate privileges:

- Metadata database
- iServer system schema and user role
- Encyclopedia volume schema and user role
- iserver application user role

### How to add an external database

The following section describes how to add an external database, schemas, and user roles using PostgreSQL server as an example. Listing 6-4 shows an example of a SQL script containing Data Definition Language (DDL) statements that create these objects with appropriate privileges in a PostgreSQL database.

#### Listing 6-4 PostgreSQL SQL Data Definition Language (DDL) script

---

```
# Run in postgres database

CREATE DATABASE iserver
  WITH OWNER = "postgres"
  TEMPLATE = template0 ENCODING = 'UTF-8';
REVOKE ALL ON DATABASE iserver FROM PUBLIC;

CREATE ROLE ac_corp_system LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO ac_corp_system;

CREATE ROLE ac_corp LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO ac_corp;

CREATE ROLE iserver LOGIN PASSWORD 'password';
GRANT CONNECT ON DATABASE iserver TO iserver;

# Run in iserver database

CREATE LANGUAGE plpgsql;

CREATE SCHEMA ac_corp_system AUTHORIZATION ac_corp_system;
GRANT USAGE ON SCHEMA ac_corp_system TO iserver;

CREATE SCHEMA ac_corp AUTHORIZATION ac_corp;
GRANT USAGE ON SCHEMA ac_corp TO iserver;
```

For more information about creating a metadata database, iServer system and Encyclopedia volume schemas, and iServer user or role in a supported alternative database, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL database, see Chapter 3, “Installing BIRT iServer using an alternative database.”

## Using Configuration Console to configure the database connection

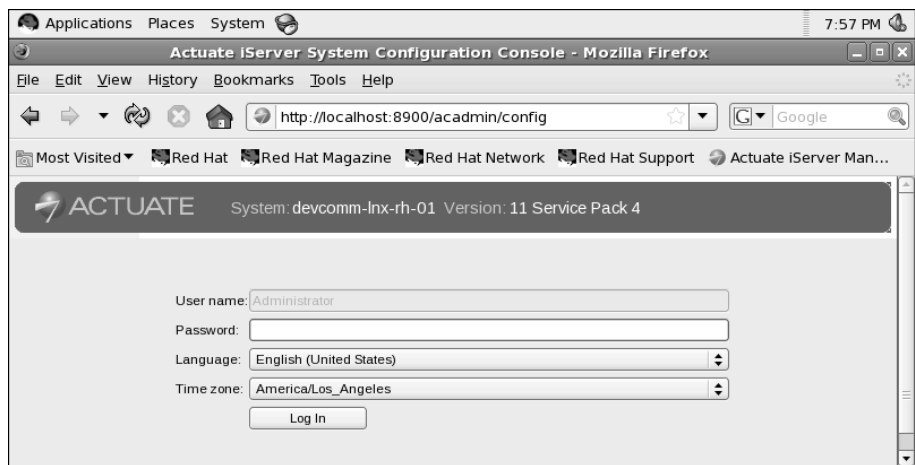
After configuring the database and completing the iServer only installation, the iServer administrator must use Configuration Console to complete the iServer system setup by specifying the database connection properties for the following items:

- Metadata database
- iServer System schema
- Encyclopedia volume schema

To create these items, perform the following tasks:

- 1 Log in to Configuration Console by opening a browser and entering the following URL, as shown in Figure 6-10. The administrator does not have a password yet, so leave Password blank.

`http://localhost:8900/acadmin/config/`

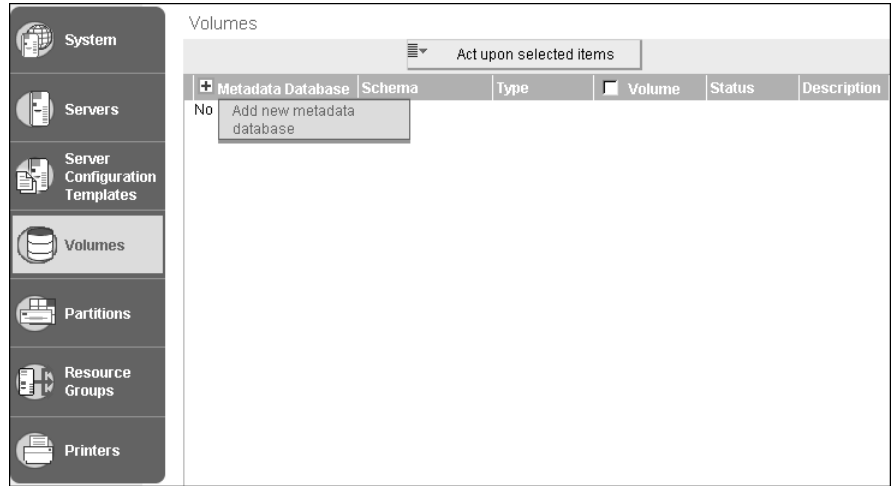


**Figure 6-10** Logging in to Configuration Console

- 2 Choose Advanced view, then choose Volumes.

## How to specify a new metadata database

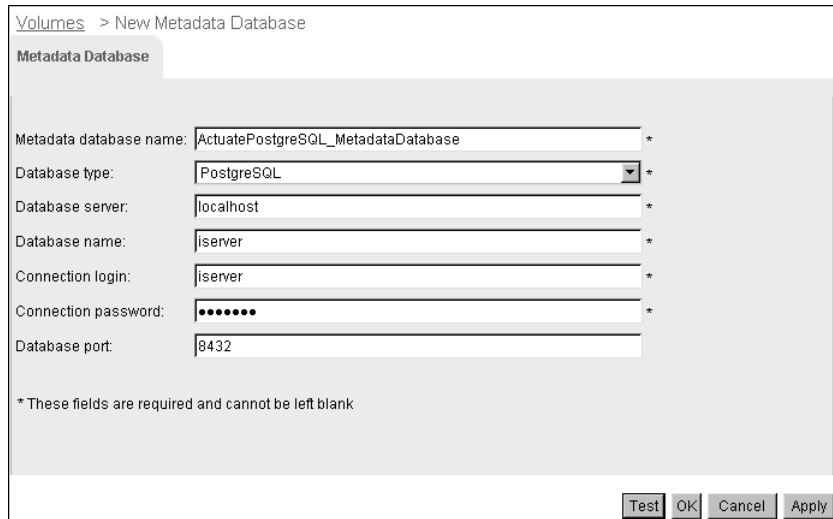
- 1 To specify a new metadata database, perform the following tasks:
  - 1 Point to the icon next to Metadata Database and choose Add new metadata database, as shown in Figure 6-11.



**Figure 6-11** Adding a new metadata database

- 2 On New Metadata Database, perform the following tasks, as shown in Figure 6-12:
  - 1 In Metadata database name, type a name for the metadata database, such as ActuatePostgreSQL\_MetadataDatabase.
  - 2 In Database type, select the type of database connection to create, such as PostgreSQL.
  - 3 In Database server, type the host name of the machine containing the database, such as localhost or the actual machine name if the database resides on a remote system.
  - 4 In Database name, type the name for the database, such as iserver.
  - 5 In Connection login, type the database application user name, such as iserver.
  - 6 In Connection password, type the database application user password.
  - 7 In Database port, specify a port number, such as 8432.






**Figure 6-12** Adding new metadata database properties

8 Choose Test.


The message, *Connection tested successfully!*, appears, if the connection properties are correct and there are no problems communicating with the database system, as shown in Figure 6-13. Choose OK.



**Figure 6-13** Viewing Connection tested successfully! message

On New Metadata Database, choose OK.


9 On Volumes, the ActuatePostgreSQL\_MetadataDatabase appears, as shown in Figure 6-14.



**Figure 6-14** Viewing the metadata database

## How to specify a new system schema

- 1 To specify a new system schema, perform the following tasks:
  - 1 On Volumes, point to the icon next to the metadata database and choose Add system schema, as shown in Figure 6-15.



**Figure 6-15** Choosing Add system schema

- 2 On New System Schema, perform the following tasks, as shown in Figure 6-16:
  - 1 In Schema name, type a name for the new schema, such as ac\_corp\_system. The name must be less than 30 characters.
  - 2 In Schema owner name, type the schema owner name, such as ac\_corp\_system.
  - 3 Type and confirm a password for the Schema owner.

The screenshot shows a 'New System Schema' configuration dialog. At the top, it says 'Volumes > New System Schema'. Below that is a 'Schema' tab. A warning message reads: 'This new schema will replace existing BIRT iServer system schema. Please restart your BIRT iServer system immediately.' The form contains the following fields:


- Metadata Database: ActuatePostgreSQL\_MetadataDatabase
- Schema type: System
- Schema name: ac\_corp\_system \*
- Database schema name: ac\_corp\_system \*
- Database schema password: [masked] \*

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

**Figure 6-16** Adding a new system schema

- 4 Choose Test. If successful, choose OK.

A message stating that the new schema will replace the existing BIRT iServer system schema and requesting to restart your system immediately appears, as shown in Figure 6-17.




**Figure 6-17** Viewing message to restart iServer system

Choose OK.

- 5 Stop and start iServer as described in “Stopping and starting iServer and PostgreSQL processes,” earlier in this chapter.

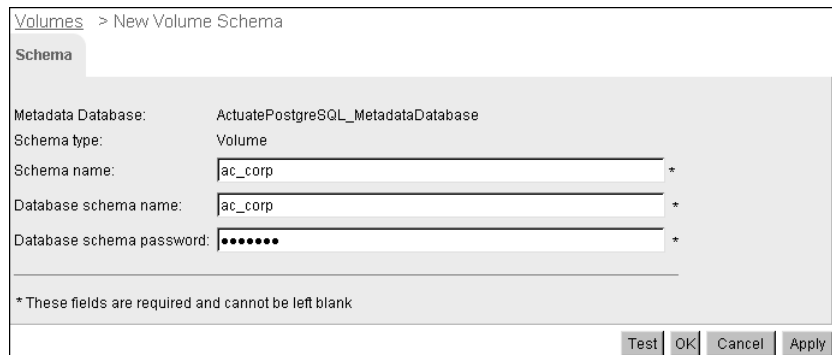
### How to specify a new Encyclopedia volume schema

- 1 To specify a new Encyclopedia volume schema, perform the following tasks:
  - 1 On Volumes, point to the icon next to the metadata database and choose Add volume schema. The metadata database is ActuatePostgreSQL\_MetadataDatabase, as shown in Figure 6-18.



**Figure 6-18** Choosing Add volume schema


- 2 On New Volume Schema, as shown in Figure 6-19, perform the following tasks:
  - 1 In Schema name, type a name for the new schema, such as ac\_corp. The name must be less than 30 characters.
  - 2 In Schema owner name, type the schema owner name, such as ac\_corp.
  - 3 Type and confirm a password for the Schema owner.



**Figure 6-19** Adding a new Encyclopedia volume schema

- 4 Choose Test.

The message, Connection tested successfully!, appears, if the connection properties are correct and there are no problems communicating with the database system, as shown in Figure 6-20. Choose OK.



**Figure 6-20** Viewing Connection tested successfully! message  
On New Volume Schema, choose OK.

## Using Configuration Console to configure the data partition and Encyclopedia volume

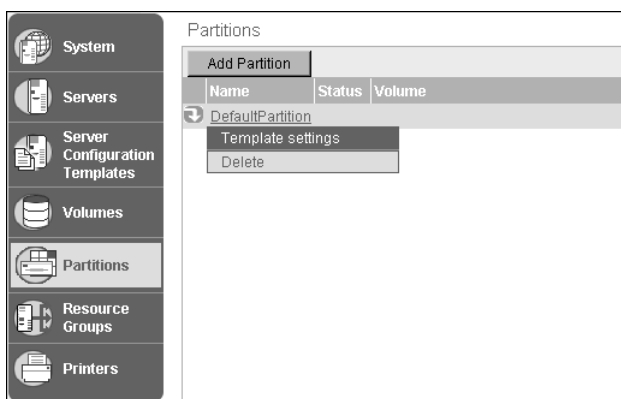
After configuring the database connection, the administrator must create the following items to bring an Encyclopedia volume online:

- Data partition
- Encyclopedia volume

### How to specify a new data partition


In Configuration Console, use the default data partition that installs with iServer or specify a new partition, then create the Encyclopedia volume and bind it to the partition.

- 1 To delete the default partition, and specify a new data partition, perform the following tasks:
  - 1 Choose Advanced view.
  - 2 From the side menu, choose Partitions. On Partitions, point to the arrow next to DefaultPartition and choose Delete, as shown in Figure 6-21.



**Figure 6-21** Deleting DefaultPartition

Choose OK to confirm deleting DefaultPartition, as shown in Figure 6-22.




**Figure 6-22** Confirming to delete the DefaultPartition

- 2 To specify a new data partition, perform the following tasks:
  - 1 Create a directory for the partition on the physical drive of the machine or storage device that iServer can access. For example, create a directory in the following path:  

```
/home/Actuate/Actuate11/AcServer/data/ac_corp_partition
```
  - 2 From Advanced view of Configuration Console, choose Partitions.
  - 3 In Partitions, choose Add partition.
  - 4 In Partition name, specify a name. For example, name the partition `ac_corp_partition`.
  - 5 In Partition Path, specify the fully qualified path to the partition directory, as shown in Figure 6-23.


Choose OK.



**Figure 6-23** Adding a new data partition

- 6 In Partitions, choose the new partition, `ac_corp_partition`, from the list of partitions.
- 7 In Template Settings, choose Test, as shown in Figure 6-24.

Choose OK.



**Figure 6-24** Testing a new data partition

If the test succeeds, the test successful message appears, as shown in message in Figure 6-25. Choose OK.




**Figure 6-25** Viewing test successful message

If the test fails, check that the directory named in the partition path exists.

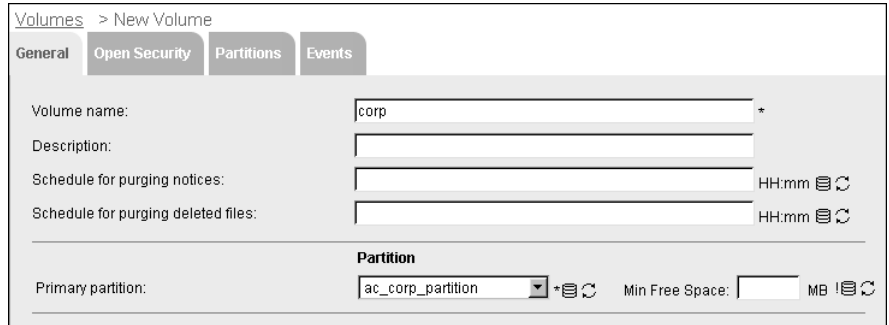
### How to specify a new Encyclopedia volume

- 1 To specify a new Encyclopedia volume, perform the following tasks:
  - 1 On Volumes, point to the icon next to an Encyclopedia volume schema and choose Add Volume, as shown in Figure 6-26.



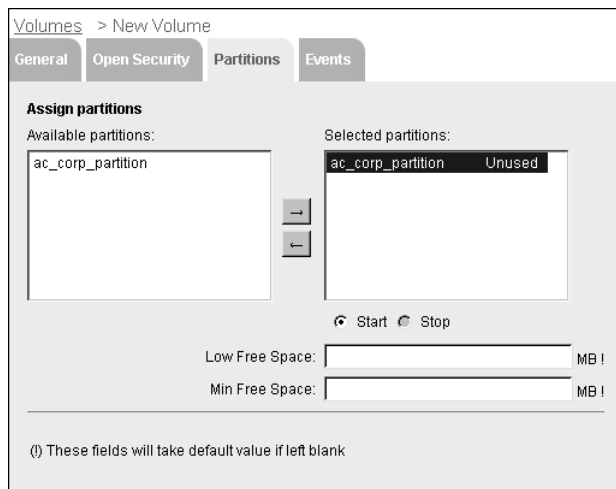
**Figure 6-26** Adding a new volume

- 2 On New Volume—General, perform the following steps:
  - 1 Type a name for the new volume. For example, type corp.
  - 2 In Primary partition, select an unassigned partition. For example, accept ac\_corp\_partition, as shown in Figure 6-27.



**Figure 6-27** Specifying general volume properties


- 3 On New Volume, choose Partitions, and start the partition for the new Encyclopedia volume by performing the following steps:
  - 1 In Available partitions, select a partition, then move it to Selected by choosing the right arrow.
  - 2 In Selected partitions, select the partition. Choose Start, as shown in Figure 6-28.



**Figure 6-28** Starting the partition

Choose OK.


- 4 In Volumes, point to the arrow next to the new volume name, and choose Take online, as shown in Figure 6-29.



**Figure 6-29** Taking a volume online

- In Volumes, check that the status of the new volume changes to ONLINE, as shown in Figure 6-30.

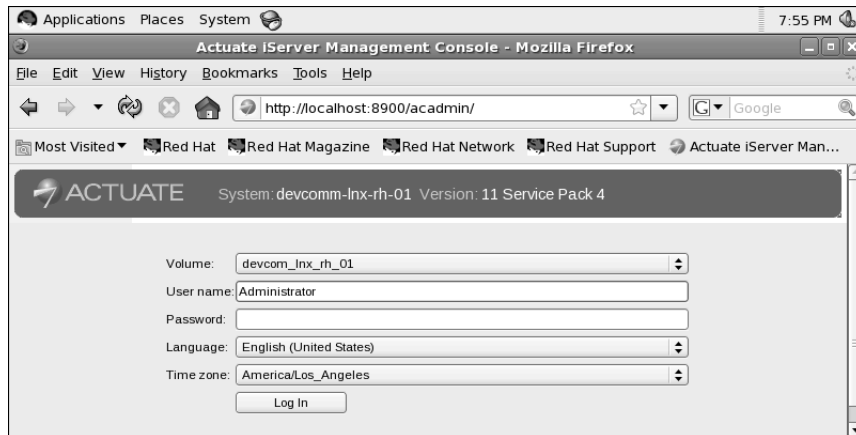
If the volume does not go online, check for insufficient free disk space for the partition and consider configuring the free space threshold.



**Figure 6-30** Viewing the online volume

- Log in to Management Console to inspect the new volume by opening a browser manually and entering the following URL, as shown in Figure 6-31:

`http://localhost:8900/acadmin/`



**Figure 6-31** Logging in to Management Console



- 7 In Files and Folders, the default Encyclopedia volume appears with an empty Resources folder, as shown in Figure 6-32.



Figure 6-32 Viewing new volume contents

---

## Setting up iServer to join an existing cluster

This section assumes the administrator has already created a stand-alone iServer instance on one computer, as described in “How to set up iServer to use the out-of-the-box (OOTB) PostgreSQL database,” earlier in this chapter. The stand-alone iServer instance is the machine that contains the shared configuration directory, which all nodes in the cluster access.

Before performing the cluster node installation, the administrator must also set up network sharing on the configuration folder and any required Encyclopedia volume folders, as described in “Adding a node to a cluster,” in Chapter 5, “Installing a BIRT iServer cluster.” It is the responsibility of the administrator performing the installation to make sure that all network sharing settings conform to the security policies in force for the environment.

This section refers to the machine containing the shared configuration directory as node1 and the cluster node accessing these shared resources as node2. Before performing a cluster node installation, the Administrator performs the following tasks:

- On node1:
  - Turns off the private-network firewall
  - Obtains the machine host name and IP address
  - Tests the network accessibility of the machine
  - Sets the partition path for DefaultPartition in Configuration Console
  - Configures directory sharing and firewall settings in the network environment for the shared configuration and Encyclopedia volume folders
- On node2:
  - Turns off the private-network firewall
  - Obtains the machine host name and IP address
  - Tests the network accessibility of the machine

The following section describes how to install and configure the iServer deployment bundle to join an existing cluster.

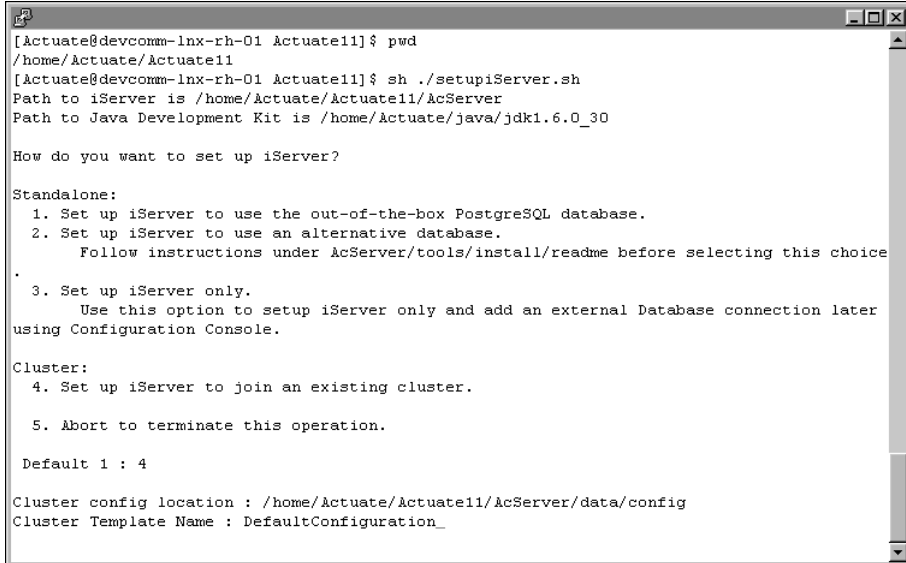
### How to set up iServer to join an existing cluster

After extracting the contents of ActuateBIRTiServer.tar.gz and running the setup script as described in “How to extract the contents of the iServer distribution package,” and “How to run the setup script,” earlier in this chapter, perform the following tasks:

- 1 In the setupiServer menu, choose option 4, Set up iServer to join an existing cluster.
- 2 At the Cluster config location prompt, specify the configuration home location, as shown in Figure 6-33. For example, type:

```
/home/Actuate/Actuate11/AcServer/data/config
```

Press Enter.



```
[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setupiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?

Standalone:
 1. Set up iServer to use the out-of-the-box PostgreSQL database.
 2. Set up iServer to use an alternative database.
    Follow instructions under AcServer/tools/install/readme before selecting this choice
 .
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connection later
    using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.
 5. Abort to terminate this operation.

Default 1 : 4

Cluster config location : /home/Actuate/Actuate11/AcServer/data/config
Cluster Template Name : DefaultConfiguration_
```

**Figure 6-33** Specifying the cluster option, configuration home location, and server template name

In an iServer cluster, the configuration home location, AC\_CONFIG\_HOME, is the shared directory that contains the acserverconfig.xml, acserverlicense.xml, and other related files.

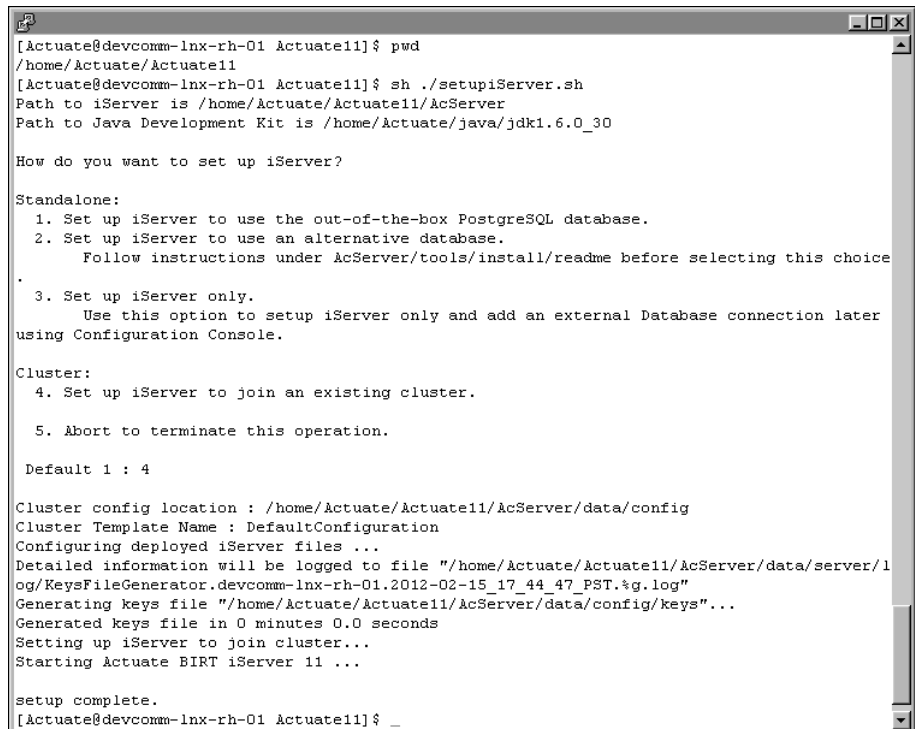
For Cluster Template Name, specify a server template name from the available server templates listed in the shared acserverconfig.xml file, as shown in Figure 6-33. For example, type:

```
DefaultConfiguration
```

The default template name in the shared `acserverconfig.xml` in a cloud-based stand-alone iServer installation is `DefaultConfiguration`. Press Enter.

The script performs the following tasks, as shown in Figure 6-34:

- Configures iServer deployment files
- Writes setup and installation information to log files
- Generates security key files
- Sets up iServer to join the cluster
- Starts BIRT iServer 11



```
[Actuate@devcomm-lnx-rh-01 Actuate11]$ pwd
/home/Actuate/Actuate11
[Actuate@devcomm-lnx-rh-01 Actuate11]$ sh ./setupiServer.sh
Path to iServer is /home/Actuate/Actuate11/AcServer
Path to Java Development Kit is /home/Actuate/java/jdk1.6.0_30

How do you want to set up iServer?

Standalone:
 1. Set up iServer to use the out-of-the-box PostgreSQL database.
 2. Set up iServer to use an alternative database.
    Follow instructions under AcServer/tools/install/readme before selecting this choice
.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connection later
using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.

 5. Abort to terminate this operation.

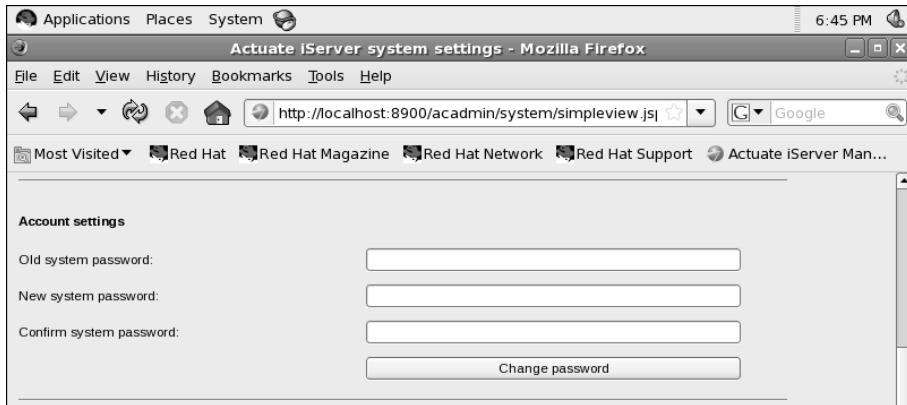
Default 1 : 4

Cluster config location : /home/Actuate/Actuate11/AcServer/data/config
Cluster Template Name : DefaultConfiguration
Configuring deployed iServer files ...
Detailed information will be logged to file "/home/Actuate/Actuate11/AcServer/data/server/1
og/KeysFileGenerator.devcomm-lnx-rh-01.2012-02-15_17_44_47_PST.%g.log"
Generating keys file "/home/Actuate/Actuate11/AcServer/data/config/keys"...
Generated keys file in 0 minutes 0.0 seconds
Setting up iServer to join cluster...
Starting Actuate BIRT iServer 11 ...

setup complete.
[Actuate@devcomm-lnx-rh-01 Actuate11]$ _
```

**Figure 6-34** Completing the iServer cluster node setup

- 3 Log in to Configuration Console. In Simple view, scroll down to Account Settings. In New system password, type a new password. In Confirm system password, type the new password again. Then, choose Change password, as shown in Figure 6-35.



**Figure 6-35** Creating a new Configuration Console password

- 4 Scroll to the top of Simple view and choose Advanced view. In Advanced view, choose Servers. The new node appears in the list of servers.

For more information about configuring a cluster, see Chapter 5, “Installing a BIRT iServer cluster,” earlier in this book, and Chapter 9, “Clustering,” in *Configuring BIRT iServer*.

---

## Understanding the cloud deployment environment

In iServer Release 11, the location of program files changes to consolidate these resources in `AC_SERVER_HOME` to facilitate creating an iServer image for deployment in a cloud environment.

iServer Release 11 introduces a new environment variable, `AC_DATA_HOME`, for specifying the location of iServer data files. `AC_DATA_HOME` separates iServer data from the run-time binaries. This change facilitates deployment in a cloud environment.

### Specifying `AC_SERVER_HOME`

The following binaries, previously installed in `AC_SERVER_HOME\operation` by iServer earlier versions, are now installed `AC_SERVER_HOME\bin` directory:

- `viewsrv11`
- `fctsrvr11`
- `fctcmd11`

The following directories, previously installed under Actuate product home, the parent directory of AC\_SERVER\_HOME, are now installed under AC\_SERVER\_HOME directory:

- \$AC\_PRODUCT\_HOME/jar
- \$AC\_PRODUCT\_HOME/oda

## **Specifying AC\_DATA\_HOME**

In an iServer Release 11 product installation, the following data files move to the AC\_DATA\_HOME directory:

- AC\_SERVER\_HOME/log
- ACTUATE\_HOME/oda/ais/log
- AC\_SERVER\_HOME/postgresql/log
- AC\_SERVER\_HOME/server/encyc
- AC\_SERVER\_HOME/server/encyc/postgresql
- AC\_SERVER\_HOME/tmp
- AC\_SERVER\_HOME/etc/acserverconfig.xml and acserverlicense.xml

Other XML configuration files, which are read-only, remain in AC\_SERVER\_HOME/etc.



# 7

## Installing Information Console

This chapter discusses the following topics:

- Before you begin
- Installing Information Console on Linux and UNIX

---

## Before you begin

Before installing Information Console, you must prepare the operating system environment to ensure that you have all the necessary software and configuration resources in place. This section describes how to prepare prior to an installation.

### About installing from an FTP download

If you download an Actuate product from the Actuate FTP software distribution site, keep all the files together in their original relative locations when you extract them. The installation program verifies that all necessary files are present before beginning the installation process.

If any files are missing, the installation program exits. Files can be missing if you download the DVD image, extract the files, move only some of the files to a new location, and attempt to install from that location.

### About performing a full installation

For Actuate 11, if you perform a full installation for iServer and Information Console, install BIRT iServer System products in the following order:

- BIRT iServer and Management Console
- Information Console
- BIRT iServer Integration Technology

To access online documentation, such as the online help and PDF files of the product manuals, install the documentation files from the following location:

<http://www.actuate.com>

---

## Installing Information Console on Linux and UNIX

This section describes how to install Information Console for Linux and UNIX. Before you begin the installation process, ensure that you have Actuate administrator, system administrator, and web administrator privileges. If you download an Actuate product for Linux and UNIX from the Actuate FTP software distribution site, keep all the files together in their original relative locations after you extract them.

Information Console installation requires the following information:

- Installation directory  
\$HOME, the account's home directory, is the default installation directory.



- Port used by the Apache Tomcat Information Console service  
If you are using a firewall, ensure that the firewall passes the port number you select. The default port is 8900.
- Value of the SPINLOOPTIME environment variable  
If you use AIX in a multiple-CPU environment, Information Console installation sets value of the SPINLOOPTIME environment variable to 2000. For more information about SPINLOOPTIME, see your AIX documentation.
- Encyclopedia volume name that Information Console accesses  
The default is the current machine.

You can install Information Console in following ways:

- Use an installation script.  
The installation script configures Information Console, creates shortcuts, and extracts and installs all necessary files. Use this option for automated configuration.
- Deploy a WAR file to an Application Server.  
Deploying directly requires that you configure Information Console for your application server. Use this option if your application server supports configuration of an application from a WAR file.

## Using the script to install


Complete the steps in the following section to install Information Console using the installation script.

### How to install using a script

- 1 Download the required files from the FTP software distribution site. Extract the files.
- 2 To install the Information Console files, type:
 


```
sh ./infoconsoleinstall.sh
```

The script displays a series of prompts. Respond to the prompts as described in the following procedures.
- 3 The license agreement appears, as shown in Figure 7-1.



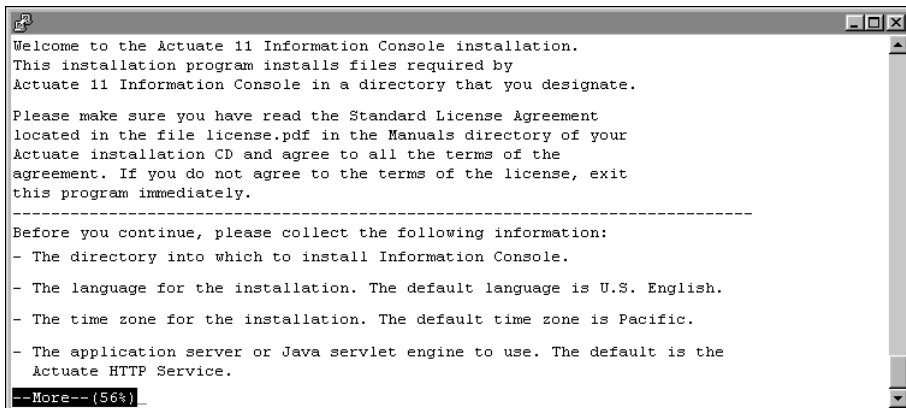
**Figure 7-1** The license agreement

- 4 Read the license agreement and press Enter to continue installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 7-2.



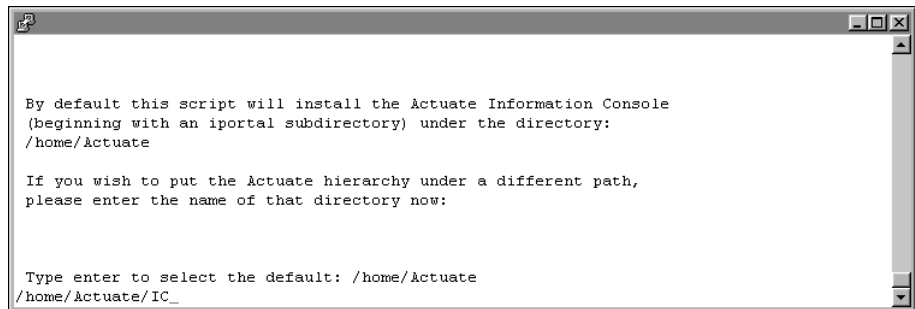
**Figure 7-2** Specifying whether you accept the license agreement

- 5 The introduction to the installation appears, as shown in Figure 7-3. Review the information, then press Enter to continue.



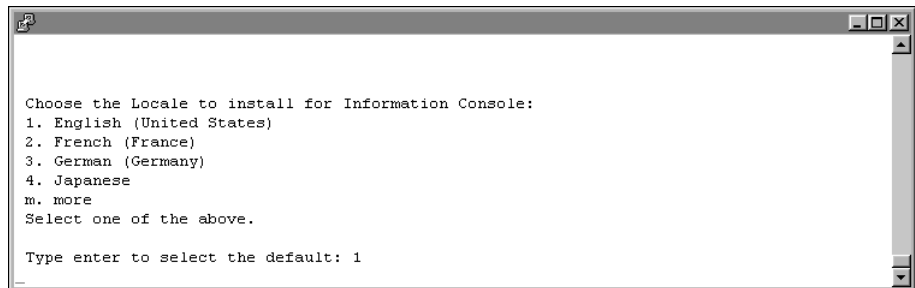
**Figure 7-3** Viewing the introduction

- 6 Type the path for the Information Console installation, for example `/home/Actuate/IC`, as shown in Figure 7-4. Alternatively, press Enter to accept the default directory.



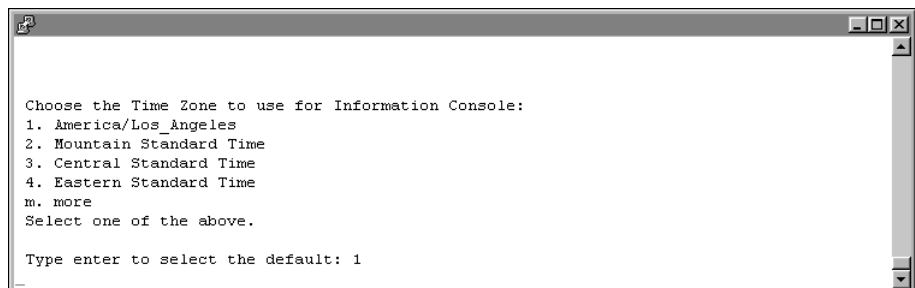
**Figure 7-4** Specifying the Information Console install directory

- 7 Press Enter to select the default locale, which is English, as shown in Figure 7-5. Alternatively, select a different locale. If you do not see the locale for your region, type `m` for more and press Enter.



**Figure 7-5** Specifying a locale

- 8 Press Enter to select the default time zone, which is `America/Los_Angeles` as shown in Figure 7-6. Alternatively, select another time zone from the numbered list.




**Figure 7-6** Specifying a time zone

- 9 Type a default profile name that you will use in Information Console, as shown in Figure 7-7.




**Figure 7-7** Specifying the default profile name

- 10 Type the IP address or host name of the machine where iServer runs, or accept the default, your machine name, as shown in Figure 7-8.




**Figure 7-8** Specifying the machine on which the iServer runs

- 11 Type the number of the port where iServer listens for requests, or accept the default, 8000, as shown in Figure 7-9.



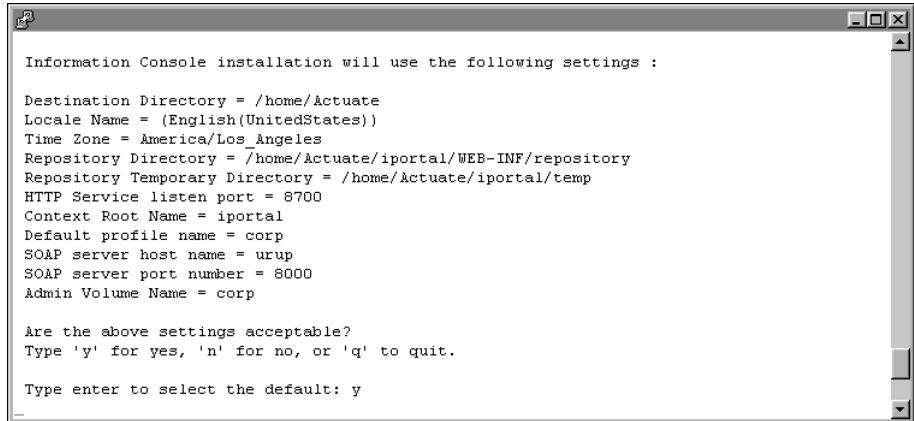
**Figure 7-9** Specifying the port number on which iServer listens

- 12 Type the Encyclopedia volume name that you want to use, or accept the default, your machine name, as shown in Figure 7-10.



**Figure 7-10** Specifying the Encyclopedia volume name

- 13 The installation program displays the settings that you specified during the install process. Review these settings, as shown in Figure 7-11, then specify whether you accept them. Press Enter to accept the default option, y for yes. Alternatively, type n for no, or type q to quit.



```
Information Console installation will use the following settings :

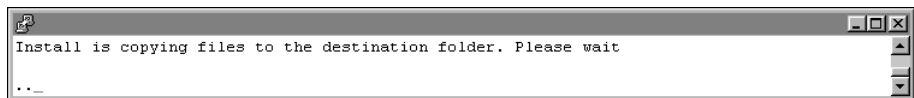
Destination Directory = /home/Actuate
Locale Name = (English(UnitedStates))
Time Zone = America/Los_Angeles
Repository Directory = /home/Actuate/iportal/WEB-INF/repository
Repository Temporary Directory = /home/Actuate/iportal/temp
HTTP Service listen port = 8700
Context Root Name = iportal
Default profile name = corp
SOAP server host name = urup
SOAP server port number = 8000
Admin Volume Name = corp

Are the above settings acceptable?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Type enter to select the default: y
```

**Figure 7-11** Reviewing your settings

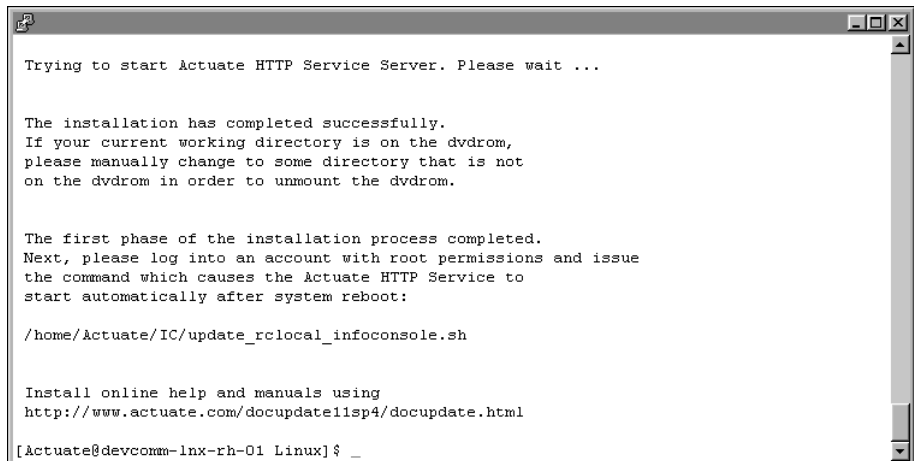
- 14 The installation program installs Information Console, and displays an indicator showing the progress of the installation, as shown in Figure 7-12.



```
Install is copying files to the destination folder. Please wait
.._
```

**Figure 7-12** The install program copies files to your destination folder

- 15 When the installation completes, review the information, as shown in Figure 7-13. Issuing the command: `sh ./update_rlocal_infoconsole.sh` starts the Information Console service at system startup.



```
Trying to start Actuate HTTP Service Server. Please wait ...

The installation has completed successfully.
If your current working directory is on the dvdrom,
please manually change to some directory that is not
on the dvdrom in order to unmount the dvdrom.

The first phase of the installation process completed.
Next, please log into an account with root permissions and issue
the command which causes the Actuate HTTP Service to
start automatically after system reboot:

/home/Actuate/IC/update_rlocal_infoconsole.sh

Install online help and manuals using
http://www.actuate.com/docupdate11sp4/docupdate.html

[Actuate@devcomm-lnx-rh-01 Linux]$ _
```

**Figure 7-13** Typing the command to start the service

## Using the WAR file to install

If Actuate supports your application server, you can deploy Information Console as a WAR (web archive) file. See the application server documentation about how to deploy a WAR file. For information about application servers on which Actuate supports deployment of Information Console, see the Supported Products Matrix for this release on the Actuate Support site at the following URL:

<http://support.actuate.com/documentation/spm>

Customize Information Console for your local environment, if necessary, before beginning deployment. To customize these applications for your local environment, follow the steps in “Preparing the WAR file,” later in this chapter, and deploy the customized Information Console WAR file to the application server instead of the `ActuateInformationConsole.war` file that ships with Information Console.

### General deployment tasks

You must configure the Information Console WAR file and application server to integrate them. You must also verify that you have a standard J2EE installation.

To integrate the application server with Actuate, complete the following tasks:

- Configure the server for best performance with Actuate products.
- Configure the Information Console WAR to integrate with the application server as described in “Preparing the WAR file,” later in this chapter.
- Deploy Information Console to the application server.
  - If the application server has deployment tools, use those tools to integrate the application server with Actuate products. For more information about deployment tools, see the application server documentation.
  - If the application server does not have deployment tools, add the Actuate context root to the application server, typically by modifying the application server configuration file.

### Preparing the server

Actuate recommends the following configuration for best performance:

- Use at least a two-CPU machine for Information Console.
- If iServer uses more than four CPUs, configure one CPU for Information Console for every two CPUs that iServer uses.

Set the following configuration values on the application or web server:

- On a machine with 1 GB of physical memory, set the Java heap size to a value between 256 MB and 512 MB.
- Set the number of threads to a value between 50 and 75.

## Preparing the WAR file

You can use a WAR file to deploy Information Console to many supported servers. Table 7-1 describes the Information Console configuration parameters to review and update before deployment.

**Table 7-1** Information Console configuration parameters

Parameter name	Description	Action
BIRT_RESOURCE_PATH	The location of the standard templates and properties files that BIRT Studio uses. This location can be in a WAR file or on a disk.	If you specify a location on disk, copy the contents of the Information Console resources folder to this physical location on the file system.
DEFAULT_LOCALE	The default locale is en_US. You can leave this value unchanged. A user can select the locale at login.	If you change the locale, select the new locale from the locales in /WEB-INF/Localemap.xml.
DEFAULT_TIMEZONE	The default time zone is Pacific Standard Time (PST). You can leave this value unchanged. A user can select a time zone at login.	If you change the time zone, select the new time zone from the time zones in the TimeZones.xml file, in the WEB-INF directory.
SERVER_DEFAULT	This value specifies the iServer URL to which the Information Console application connects if you do not specify a server URL. The default value is http://localhost:8000	Update this value to the machine and port of the server. Change localhost to the iServer machine name or IP address. Change 8000 to the iServer port number.
DEFAULT_VOLUME	This value specifies the default Encyclopedia volume for Information Console. If you do not specify a volume in an Information Console URL, the JSP application attempts to log in to this volume. The default value is localhost.	Update this value to the name of an Encyclopedia volume.
BIRT_VIEWER_LOG_DIR	The location of the files that log BIRT viewer activity.	You must create this physical location on the file system.
LOG_FILE_LOCATION	The location of the files that log Information Console activity.	You must create this physical location on the file system.
TEMP_FOLDER_LOCATION	The location where Information Console creates temporary files.	You must create this physical location on the file system.
TRANSIENT_STORE_PATH	The location where Information Console creates temporary files.	You must create this physical location on the file system.

## How to customize the WAR file

The following steps describe the generic procedure for customizing an Information Console WAR file:

- 1 Create a temporary directory, such as `/home/Actuate/ic_temp`.  
If you use an existing directory, ensure that this directory is empty.
- 2 Copy the appropriate Information Console WAR file for your environment, such as `TOMCAT_ActuateInformationConsole.war`, to the temporary directory, as shown in the following example:

```
cp TOMCAT_ActuateInformationConsole.war /home/Actuate/ic_temp
```

- 3 Decompress the WAR file, as shown in the following example:

```
jar -xf TOMCAT_ActuateInformationConsole.war
```

The Information Console files appear in the temporary directory.

- 4 Using a text editor that accepts UTF-8 encoding, edit `web.xml` to configure Information Console for your application server.

If you used the temporary path in step 2, the file location is `/home/Actuate/ic_temp/WEB-INF/Web.xml`. Refer to Table 7-1 for a list of entries to modify in `web.xml`.

- 5 Save and close `web.xml`.

- 6 Type the following command:

```
jar -cf ../newinformationconsole.war *
```

This command creates `newinformationconsole.war` in the `/ic_temp` directory. This new WAR file for Information Console contains the modified configuration values.

Use `newinformationconsole.war` to deploy Information Console to your application server.

## About clusters of servers

If your application server supports clustering, see your application server documentation for more information about setting up clusters and deploying web applications such as Information Console.

## Avoiding cache conflicts after installing

Information Console uses Java Server Page (JSP) technology. Application servers and browsers cache pages. A browser can use a cached copy of a page instead of the new page. After you install, using a cached copy of some pages can lead to errors or missing functionality.



To avoid this problem, clear the application server cache after you install Information Console. With some application servers, you must restart the application server. For more information about the necessary steps for clearing the cache, see the application server documentation.

If possible, also clear the browser cache to ensure that the browser does not use an old copy of the page from its cache. Alternatively, you can refresh recently visited pages or clear the browser cache if a problem occurs. For more information about clearing the browser's cache and refreshing a page, see the browser documentation.

## Testing the installation

Complete the steps in the following section to test the Information Console installation.

### How to test the installation

- 1 Start the application server, if necessary.
- 2 Open a web browser.
- 3 Type the URL for the Information Console home page.

You can use a URL similar to the following example:

```
http://Actuate1:8900/iportal/getfolderitems.do
    ?repositoryType=Enterprise&volume=volume1
    &serverurl=http://iServer1:8000
```

where

- Actuate1:8900 is the name of your computer and the port you use to access Information Console.
- iportal is the context root for Information Console.
- ? indicates the beginning of a parameter that indicates where to access Information Console files.
- getfolderitems.do is the call to the default Information Console home page.
- repositoryType=Enterprise indicates that this Information Console connects to iServer.
- &volume=volume1&serverurl=http://iServer1:8000 specifies the Encyclopedia volume and URL to the BIRT iServer.

The Information Console login page appears.

On the Information Console login page:

- For enterprise mode Information Console, in Volume, select an Encyclopedia Volume.

- In User name, type your login name.
  - In Password, type your password.
  - In Language, select a locale.
  - In Time zone, select a time zone.
- 4** Choose Log In.

# 8

## **Installing iServer Integration Technology and Documentation**

This chapter discusses the following topics:

- Installing BIRT iServer Integration Technology
- Installing the localization and documentation files

---

## Installing BIRT iServer Integration Technology

This section describes how to install BIRT iServer Integration Technology for Linux and UNIX. If you download an Actuate product for UNIX or Linux from the Actuate FTP software distribution site, keep all the files together in their original relative locations after you extract them.

### How to install

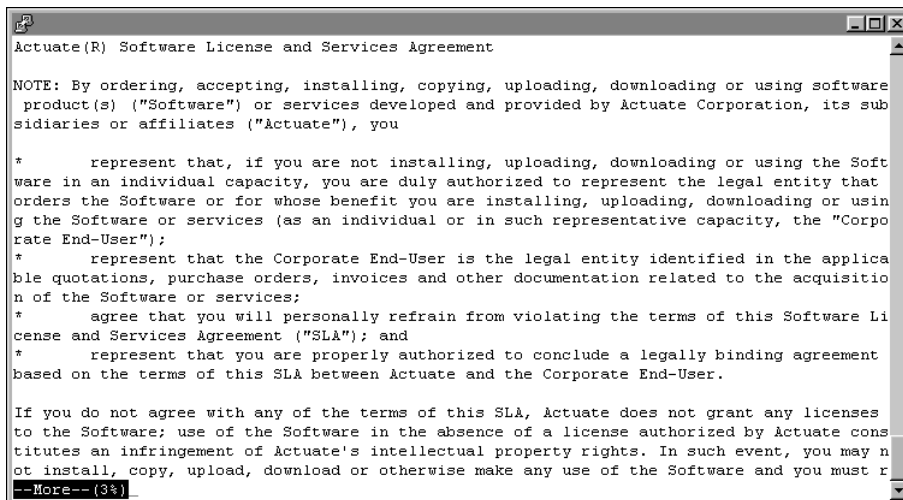
In a default installation, BIRT iServer Integration Technology installs in \$HOME /ServerIntTech. To install BIRT iServer Integration Technology, perform the following steps:

- 1 Download the required files from the FTP software distribution site. Extract the files.
- 2 To install the server files, execute the isitinstall script:

```
sh ./isitinstall.sh
```


The script displays a number of prompts. Respond to the prompts as described in the following procedure.

- 3 The license agreement appears, as shown in Figure 8-1.



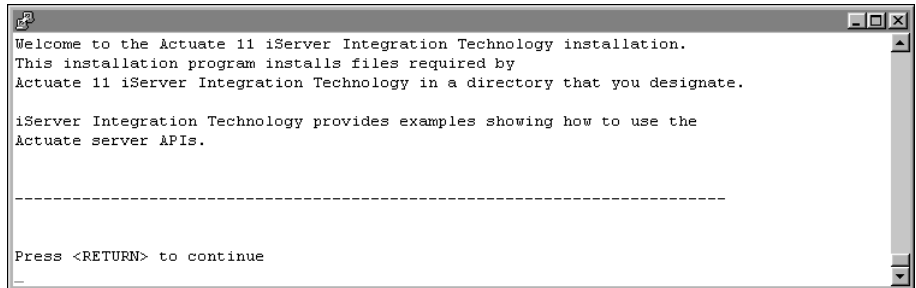
**Figure 8-1** Reviewing the license agreement

- 4 Read the license agreement and press Enter to continue installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 8-2.



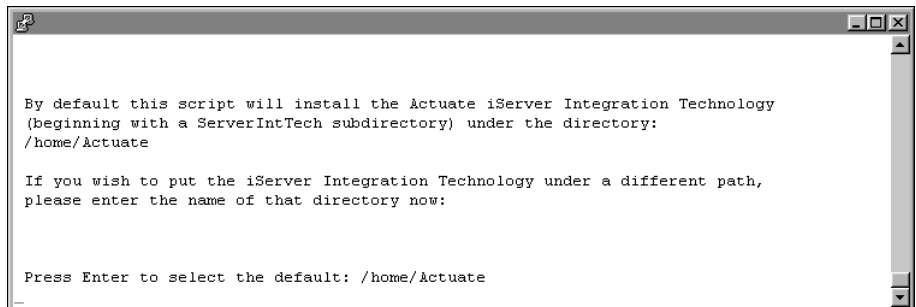
**Figure 8-2** Accepting the license agreement

- 5 The introduction to the installation appears, as shown in Figure 8-3. Press Enter after reviewing the introductory information.



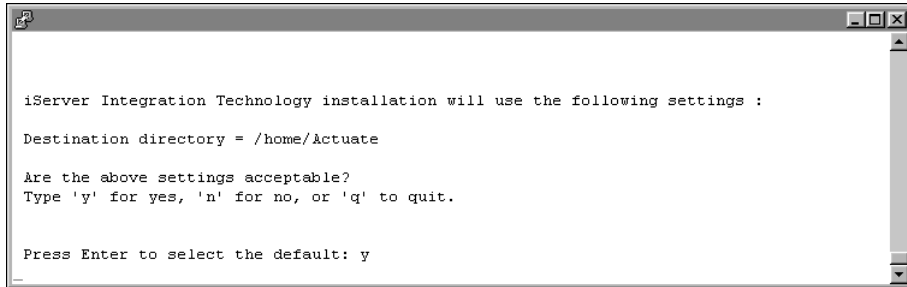
**Figure 8-3** Reviewing introductory information

- 6 Press Enter to accept the default location for installation, \$HOME /ServerIntTech as shown in Figure 8-4. Alternatively, type a different directory and press Enter.



**Figure 8-4** Specifying the installation directory

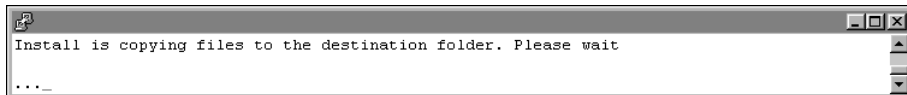
- 7 Review the settings, as shown in Figure 8-5, then specify whether to accept the settings. Press Enter to accept the default, y for yes. Alternatively type n for no, or q to quit.



```
iServer Integration Technology installation will use the following settings :  
Destination directory = /home/Actuate  
  
Are the above settings acceptable?  
Type 'y' for yes, 'n' for no, or 'q' to quit.  
  
Press Enter to select the default: y
```

**Figure 8-5** Reviewing settings before copying files

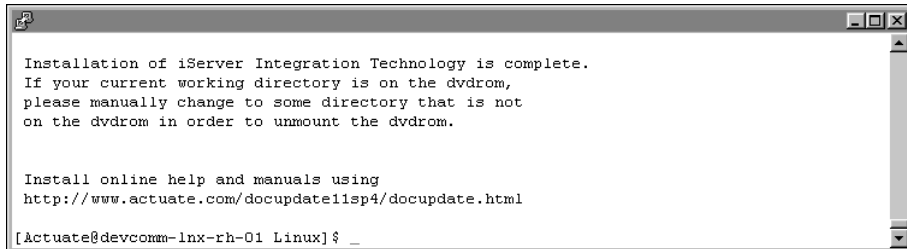
- 8 The installation program installs iServer Integration Technology, and displays an indicator showing how the installation is progressing, as shown in Figure 8-6.



```
Install is copying files to the destination folder. Please wait  
..._
```

**Figure 8-6** Copying files to your destination folder

- 9 When the installation program finishes, it provides additional information about changing the working directory and installing online help and manuals, as shown in Figure 8-7.



```
Installation of iServer Integration Technology is complete.  
If your current working directory is on the dvdrom,  
please manually change to some directory that is not  
on the dvdrom in order to unmount the dvdrom.  
  
Install online help and manuals using  
http://www.actuate.com/docupdate11sp4/docupdate.html  
[&actuate@devcomm-lnx-rh-01 Linux]$ _
```

**Figure 8-7** Viewing information about changing the working directory and installing online help

---

## Installing the localization and documentation files

The information in the printed manuals is available as Adobe Acrobat PDF files and as a context-sensitive help system for Actuate products. After installing Actuate products, install the localization and documentation files.

Both localization and documentation resource file updates can become available between releases. The Actuate Localization and Online Documentation Update tool provides replacements and additional files for PDF documentation,

context-sensitive help, and localization of installed Actuate products. The tool is available from the following Actuate web site location:

<http://support.actuate.com/documentation/releasenotes>

Updates to documentation in PDF form are available at the following Actuate web site locations:

<http://support.actuate.com/documentation>

<http://www.actuate.com/docupdate11sp4/docupdate.html>

If you do not see an update tool for your release on the Support site, no updates exist for the release.

### How to install the localization and documentation files

To install the iServer localization and documentation files, perform the following steps:

- 1 After downloading the tar file, decompress it using the tar command:

```
tar -xvf FILE_NAME.tar
```

- 2 To begin the installation, move to the newly decompressed directory and execute the helpinstall script:

```
sh ./helpinstall.sh
```

The script displays a series of prompts. Respond to the prompts as described in the following procedures.

- 3 The license agreement appears, as shown in Figure 8-8.

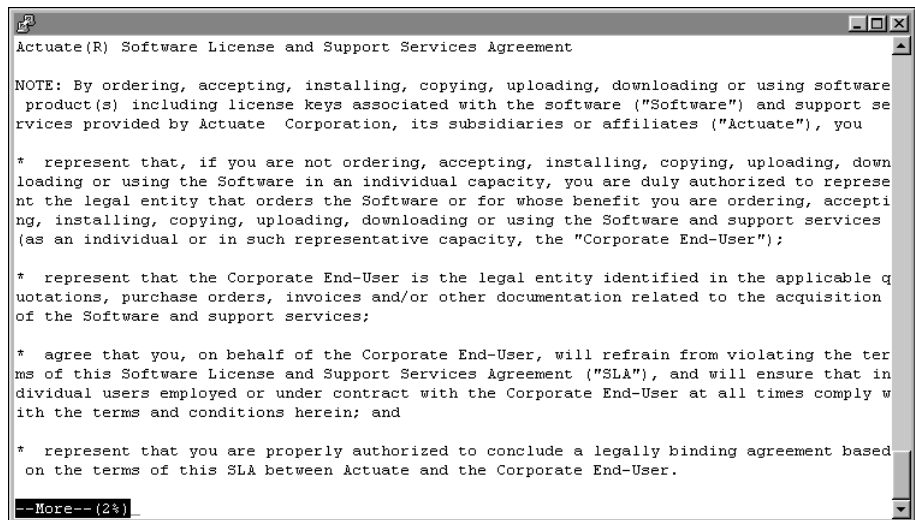
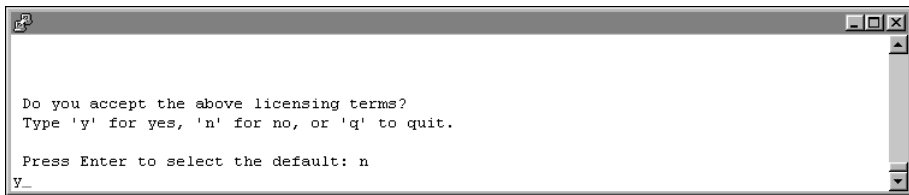


Figure 8-8 Reviewing the license agreement

- 4 Read the license agreement and press Enter to continue installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 8-9.

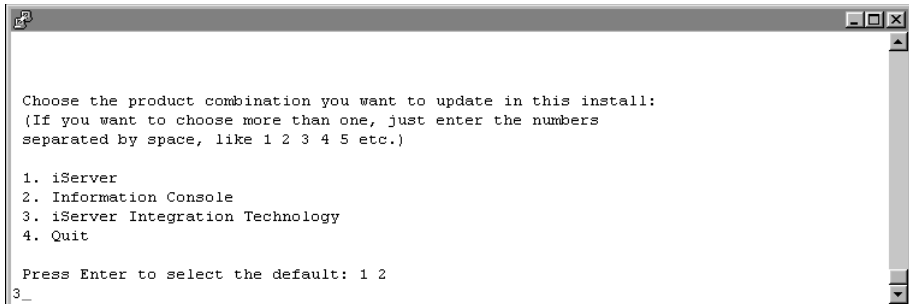


```
Do you accept the above licensing terms?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: n
y_
```

**Figure 8-9** Accepting the license agreement

- 5 Choose the products that you wish to update in this install, as shown in Figure 8-10. If you want to choose more than one, just enter the numbers separated by a space. For example, type 1 2 3 to select all products.



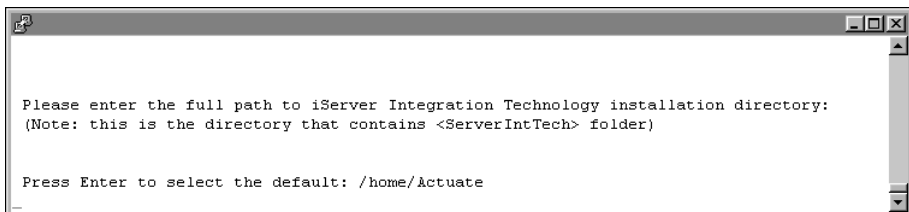
```
Choose the product combination you want to update in this install:
(If you want to choose more than one, just enter the numbers
separated by space, like 1 2 3 4 5 etc.)

1. iServer
2. Information Console
3. iServer Integration Technology
4. Quit

Press Enter to select the default: 1 2
3_
```

**Figure 8-10** Selecting a product combination

- 6 Enter the full path of the product installation directory that you chose to update, as shown in Figure 8-11. If you chose to update more than one product, you are asked for the full path of every product directory in your selection.



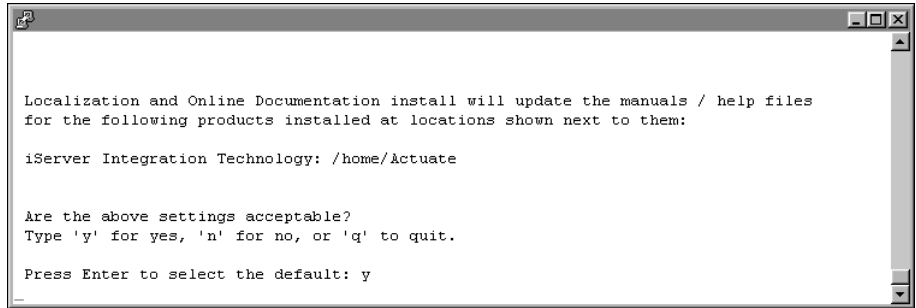
```
Please enter the full path to iServer Integration Technology installation directory:
(Note: this is the directory that contains <ServerIntTech> folder)

Press Enter to select the default: /home/Actuate
_
```

**Figure 8-11** Specifying a product directory

- 7 Review the settings, as shown in Figure 8-12, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.





```
Localization and Online Documentation install will update the manuals / help files
for the following products installed at locations shown next to them:

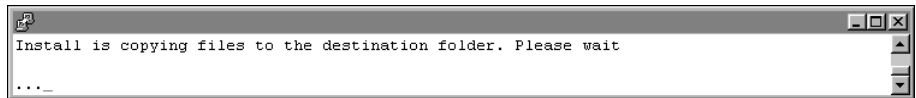
iServer Integration Technology: /home/Actuate

Are the above settings acceptable?
Type 'y' for yes, 'n' for no, or 'q' to quit.

Press Enter to select the default: y
```

**Figure 8-12** Reviewing settings before copying files

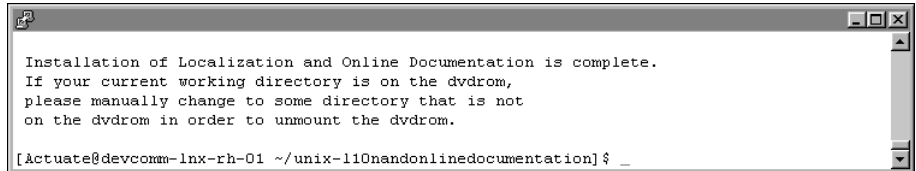
- 8 The installation program starts and displays an indicator showing how the installation is progressing, as shown in Figure 8-13.



```
Install is copying files to the destination folder. Please wait
..._
```

**Figure 8-13** Copying files to your destination folder

- 9 When the installation program finishes, it provides additional information about changing the working directory, as shown in Figure 8-14.



```
Installation of Localization and Online Documentation is complete.
If your current working directory is on the dvdrom,
please manually change to some directory that is not
on the dvdrom in order to unmount the dvdrom.

[Actuate@devcomm-lnx-rh-01 ~/unix-110nandonlinedocumentation]$ _
```

**Figure 8-14** Viewing information about changing the working directory



# Part **Three**

---

**Licensing**



# Licensing BIRT iServer

This chapter discusses the following topics:

- Understanding licensing types
- Understanding licensing options
- Installing Actuate BIRT iServer System license files
- Understanding CPU binding

---

## Understanding licensing types

BIRT iServer System licensing supports running BIRT iServer with sets of features grouped as license options. You enable BIRT iServer System options using one or more of the following types of license models:

- **Named User**

Specifies the maximum number of named users that can use an BIRT iServer System. A named user is a distinct individual who receives content and value from BIRT iServer.

In Release 11, a BIRT iServer administrator must specify the options that a user can access in an Encyclopedia volume. The administrator makes this specification by configuring the user Licensed Options properties in Management Console.

When you license an option, such as BIRT Option, e.Report Option, or BIRT Spreadsheet Option, the license entitles the user to access a single volume in BIRT iServer System. If the user needs additional Encyclopedia volumes for multiple applications, archiving, or other purposes, you must license the Multi-Tenant Option for each additional Encyclopedia volume the user needs to access.

- **CPU Core**

Specifies the maximum number of CPUs that BIRT iServer System can use. Any number of users can access the licensed options on the system provided adequate licensing and capacity exists.

- **Instance**

A BIRT onDemand 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**

An annual payment option that permits the use of the licensed software and 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 can no longer be used.

- **Software as a Service (SaaS)**

Some products are offered as a Software as a Service (SaaS) option, providing customers with a convenient solution without incurring the acquisition and management costs of hardware and traditional 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.
- Work Unit (WU) License  
Specifies iServer features and functionality using an aggregate model. This plan defines each iServer 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.  
  
In a CPU Core and Work Unit licensing, Actuate currently uses the Standard Performance Evaluation Corporation (SPEC) standard benchmark for measuring machine capacity based on CPU, memory, disk, and network capacity.

---

## Understanding licensing options

Table 9-1 lists and describes BIRT iServer System license options. BIRT iServer System options are separately licensed products. Some license options require other options to be licensed before their functionality is available to users. Table 9-1 also describes these prerequisites.

**Table 9-1** BIRT iServer System license options

Option	Description	Supported releases
Actuate Analytics	Allows building a cube and displaying a cube report for the purpose of multidimensional analysis. The Analytic Option enables BIRT iServer to extract data from a database and build a compressed OLAP file. When you analyze the cube, you can aggregate or categorize data, summarize data, and create graphs based on data. You can save and share views of the analysis you perform in the Encyclopedia volume.	10, 11
Actuate Query	Supports retrieving information using an information object.	10, 11
BIRT	Allows a user to publish and run a BIRT design using BIRT iServer. This option is a requirement for BIRT Page Level Security Option.	10, 11
BIRT 360	Allows a user to create, execute, and view dashboard files.	11

*(continues)*

**Table 9-1** BIRT iServer System license options (continued)

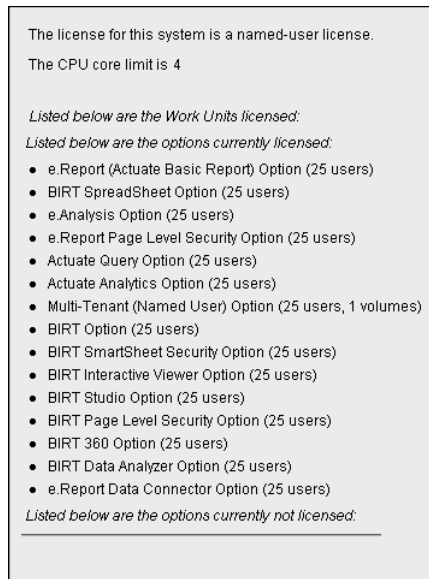
<b>Option</b>	<b>Description</b>	<b>Supported releases</b>
BIRT Data Analyzer	Allows a user to create, view, and modify cubeview files.	11
BIRT Interactive Viewer	Allows a user who has the BIRT Option to use BIRT Interactive Viewer to view and interact with a BIRT document.	10, 11
BIRT Page Level Security	Controls access to structured content available on the web. This option works for reports created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or security role.	10, 11
BIRT SmartSheet Security	Controls access to structured content available on the web. This option works for reports created using BIRT Spreadsheet Designer and requires BIRT Spreadsheet Option. Access privileges are based on user name or security role.	10, 11
BIRT Spreadsheet	Allows a user to deploy and run a spreadsheet built using BIRT Spreadsheet Designer. This tool enables customers to save reports as richly formatted Excel spreadsheets and manage them in an Encyclopedia volume.	10, 11
BIRT Studio	Allows a user to create a BIRT design and to run it in BIRT iServer. BIRT Studio Option supports access to an information object on BIRT iServer System.	10, 11
e.Analysis	Supports analysis of search results from an Actuate Basic report written in dynamic hypertext markup language (DHTML). This tool is available as an additional purchase with BIRT iServer and requires the e.Report (Actuate Basic Report) Option.	10, 11
e.Report (Actuate Basic Report)	Allows a user to deploy and run an e.report built using Actuate e.Report Designer Professional on an Encyclopedia volume. This option is a requirement for e.Analysis Option and e.Report Page Level Security Option.	10, 11
e.Report Data Connector	Allows a BIRT design to access data that an Actuate Report Document (.roi) file contains. This option works for designs created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or security role.	11



**Table 9-1** BIRT iServer System license options (continued)

Option	Description	Supported releases
e.Report Page Level Security	Controls access to structured content available on the web. This option works for reports created using Actuate e.Report Designer Professional and requires the e.Report (Actuate Basic Report) Option. Access privileges are based on user name or security role.	10, 11
Multi-Tenant	Allows a BIRT iServer System user to access more than one Encyclopedia volume. This option is available with an Unlimited User CPU License.	11

To determine the license options installed on iServer, log in to Configuration Console, and choose Show License. The license options appear, as shown in Figure 9-1.



**Figure 9-1** iServer License options

## Installing Actuate BIRT iServer System license files

Actuate provides a license file to use when installing Actuate BIRT iServer System. New customers receive an e-mail containing a temporary BIRT iServer license file to use for the initial installation after Actuate processes the order. The temporary BIRT iServer System license expires 45 days after installation.

Actuate license enforcement for Release 11 requires a single, shared license for all nodes in a cluster. A design or document run using a temporary license appears with a watermark when viewed.

The name for the BIRT iServer license file uses the following format:

```
Actuate_iServer_key_XXXXXXX.xml
```

XXXXXXX is a unique seven-digit number generated by Actuate Licensing when it creates the license file.

Actuate BIRT iServer System customers perform an initial installation using a temporary license. After installing BIRT iServer System using the temporary license, the login screen displays two messages.

The following message about expiration of the initial license always appears on the login screen regardless of the license status:

Reminder

```
Your BIRT iServer license expires in [the number of days] days,
on [the specified date]. When the current license expires, the
iServer 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 iServers in your BIRT iServer System are in
violation of the node locked BIRT iServer license. After the
grace period expires, the iServers that violate the node locked
BIRT iServer 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 iServer nodes that are in
violation. Please restart the iServers on the nodes after
updating the license key file.
```

You have 45 days to apply for and install the license file after you install BIRT iServer System.

After installing Actuate BIRT iServer System, the installation informs a customer requiring a license to obtain the machine ID information on which Actuate BIRT iServer 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 iServer System license file.

## About the license file

This license file specifies the available iServer 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 an iServer node in a cluster with the machine ID. The node-key licensing mechanism restricts the iServer 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 remains valid until a specific date. If your license file is about to expire, the system reminds you that the file expires on a certain date when you log in to the Configuration or Management Consoles. 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 iServer 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 iServer System using a temporary license file, such as an evaluation license, you must collect information about the machines running Actuate BIRT iServer 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 iServer system license file is locked to the machines that are
used in the iServer system. The following machine id must be
used to request a node key license file from Actuate:
```

```
IORRHEHs6S5UCsEtrdVu6jOixmzvFY3BbOqXLiwsWQGDceJmKYYaEu0j18lQxjM
sYCxnka3hVkdZFGwkmQMxb+hgKaz4om2vLUcS0ocYTA7Ta6VTMavLFQo7bEjRyr
olwxAKu0Vr4NA6o8uWCzjGZXX8KrvViSUoROj70hWOY=
```

Please contact Actuate Licensing (licensing@actuate.com or <http://www.actuate.com/licensing>), or your representative, and request a node locked iServer system license.

The machine id required for the node locked iServer system license can also be generated by using the acmachineid utility that can be found in the ACTUATE\_HOME\AcServer\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.

After installing iServer, 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 iServer System.

### How to use the acmachineid utility

Use the acmachineid utility to obtain the machine ID information by performing the following tasks:

- 1 Navigate to AC\_SERVER\_HOME\bin.
- 2 Type the following command and press Enter:

```
./acmachineid
```

The utility provides output in the following format:

```
STATUS:          OK
GEN_VERSION:     11 Service Pack 4
GEN_BUILD:       XxxBuild NumberXxx
MACHINEID:      I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67GbO8
                 VjcFs6pcuAgbtDaZauSbFFa2mRejwVJc7ZjKfMEVl1suXglM
                 KmZLiwTlykwJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9f
                 JUMYzZORKk23jrxaswUDsgKsvlc1A6q8UbmrrAYHD8Ggtptui
                 AmxWt4xjEM6rqlmsNEW/4ViMC0KDBkSn
```

Send Actuate Licensing the output of the acmachineid utility.

### How to obtain a license file

To obtain a new license file for a licensed product or receive assistance with a license already issued to you, perform the following tasks:

- 1 Using a browser, go to the Actuate Support web site at the following location:  
<http://support.actuate.com>
- 2 Choose Downloads/Requests→License Key Request.
- 3 Enter the required contact details and license key request information.
- 4 Choose Submit.

A maintenance customer should have login information for the Actuate e.Support web site. If you do not have access, please contact Actuate Support at support@actuate.com. You can also contact Actuate Customer Care at customercare@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.

### **Updating the Actuate BIRT iServer System license file**

After performing an installation of Actuate BIRT iServer 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 Actuate BIRT iServer System using a license file and for obtaining a license file.
- Actuate\_iServer\_key\_XXXXXXX.xml  
Actuate BIRT iServer System license.

An Actuate license file is an XML file. Actuate Licensing sends this XML file inside of a TXT file because transmitting a file with an .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 Configuration Console. For example, type `http://localhost:8900/acadmin/config/` in Address of a browser, and use the system configuration password that you specified during installation.
- 3 Choose Update License. If you do not see Update License, choose Simple view in the upper right corner.
- 4 On Actuate iServer update license, choose Browse to navigate to the location where you extracted the contents of the ZIP file. Select the Actuate BIRT iServer System license file and choose OK to apply the license.

If iServer 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 iServer license file  
location and restart the iServer system.
```

If this message appears, perform the following tasks:

- 1 Stop iServer system by choosing Stop system.

- 2 Copy the new license file to the shared location specified by the <AC\_CONFIG\_HOME> attribute of the <Server> element in the acpmdconfig.xml file. The <AC\_CONFIG\_HOME> attribute in the acpmdconfig.xml files for all nodes in a cluster point to this shared location.
  - 3 Delete the old acserverlicense.xml file.
  - 4 Rename the new license file to acserverlicense.xml.
  - 5 Start iServer System.
- 5 Restart any node where the node-key configuration changed.

If you change the machine on which you installed Actuate BIRT iServer, 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 Actuate Management Console, contact Actuate Licensing and provide the current Actuate iServer System license file with the output from the machine ID utility.

The Actuate\_iServer\_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 9-1 shows the node key information the license contains, obtained from the acmachineid output you submitted to Actuate Licensing.

**Listing 9-1** Viewing license node key information

---

```
<NodeKeys>
  <NodeKey
    MachineId="E0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67Gb08
VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXg1MKmZLiwtLykDa/
wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORRk23jrxaswUDig
KsvlclA6q8UbmrrAYHD8GgtptuiAmxWt4xjEM6rqlmsNEW/4Vjm40Kx1kSv"
    ServerName="W7CLSTRNODE1" />
  <NodeKey
    MachineId="I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67Gb08
VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXg1MKmZLiwtLykDa/
wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORRk23jrxaswUDsg
KsvlclA6q8UbmrrAYHD8GgtptuiAmxWt4xjEM6rqlmsNEW/4ViMC0KDBkSn"
    ServerName="W7CLSTRNODE2" />
</NodeKeys>
```

## About modifying a license

If you decide later to license additional iServer options, the existing license file becomes invalid. You must install a new license file.

Contact Actuate Licensing for the new license file. If you are an Actuate international customer, please be aware that the e-mail message sent to Actuate goes to Actuate headquarters, and we route your request to a team in the appropriate country.

---

## Understanding CPU binding

BIRT iServer 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 CPUs. If you bind the BIRT iServer System to a subset of CPUs, only those CPUs count toward the total number of licensed CPUs. The CPU limit in the license file applies to all CPUs 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 iServer processes to a specific set of processors on a machine that runs a Windows, Sun Solaris, or HP-UX 11i operating system. The default configuration does not bind BIRT iServer to a set of processors. In the default configuration, all processors on an BIRT iServer machine count toward the maximum number of licensed CPUs. For more information about performing CPU binding on a Windows machine, see Chapter 7, “Licensing BIRT iServer,” in *Installing BIRT iServer for Windows*.

To bind BIRT iServer to a set of processors, bind the Actuate Process Management Daemon (PMD) to the processors. The Actuate PMD starts all BIRT iServer processes. The processes inherit the binding from the Actuate PMD.

In a cluster, BIRT iServer counts only the processors on nodes that join the cluster and run the `encycsrvr11` process. An `encycsrvr11` process runs when a node is online. BIRT iServer counts the number of processors on a machine when the first `encycsrvr11` process starts.

When deploying BIRT iServer on a machine with multi-threaded CPUs that use logical processors, the customer receives a license based on the number of physical processors in the system. To accommodate the use of logical processors, the customer receives a license key that specifies two or four times the number of physical processors.

This section contains the following topics:

- Binding BIRT iServer to processors on a Sun Solaris machine
- Binding BIRT iServer to processors on an HP-UX 11i machine
- Checking BIRT iServer bound processors
- Configuring e-mail for CPU license problems

## Binding BIRT iServer to processors on a Sun Solaris machine

You can perform single processor binding and processor set binding on a Solaris machine with the following types of binding:

- With single processor binding, you bind a process to a single processor using the `processor_bind()` system call or the `pbind` command.
- With processor set binding, you bind a process to a group of processors on a Solaris machine. If you bind a process to multiple processor sets, the sets must not overlap.

To create a processor set, use the `pset_create()` system call or `psrset` command. Then, you can bind a process to the set using the `pset_bind()` system call or the `psrset` command. The `psrset` command restricts other processes from running on the processors to which you bind a BIRT iServer process.

For more information about processor binding and the processor binding commands, refer to your Solaris documentation.

### Binding to specific CPUs

On Solaris, you bind a CPU to BIRT iServer using one of the following commands:

- `pbind`  
Binds the BIRT iServer to a single CPU. Does not prevent other applications from using the CPU.

For example, in an 8-CPU Solaris server, you can use `pbind` to bind the BIRT iServer to CPU 2. An Oracle database running on the same server that is not bound to a CPU can impact BIRT iServer performance. While BIRT iServer can use only the processing power of CPU 2, the Oracle database can use all CPUs, including CPU 2. BIRT iServer cannot fully utilize CPU 2 if the Oracle database obstructs access.

- `psrset`  
Binds the BIRT iServer to one or more CPUs or cores. Prevents other applications from using the CPUs or cores.

For example, in a 4-CPU Solaris server, you can use `psrset` to bind the BIRT iServer to CPUs 1 and 2. WebLogic Application Server runs on the same server and is restricted to run on CPUs 3 and 4. BIRT iServer and WebLogic cannot obstruct access to the CPUs assigned to the other process.

To bind to a single CPU, shut down BIRT iServer and determine the process ID of the Actuate Process Management Daemon (PMD) using the `ps -e` command. Type



the `pbind` command to bind the PMD process to a subset of CPUs using the following syntax:

```
pbind -b processor Actuate_pid
```

where

- `Actuate_pid` is the process ID of the PMD, as reported by `ps -e`.
- `processor` is the processor number as reported by `/usr/sbin/psrinfo`.

The following example shows binding to a specific CPU on Solaris:

```
pbind -b 0 8209 # Binds process 8209 to CPU 0
```

## Binding to multiple CPUs

You use Solaris processor sets to bind BIRT iServer to multiple CPUs. Solaris processor sets are non-overlapping groups of processors.

To bind to multiple CPUs, shut down BIRT iServer and determine the process ID of the PMD using the `ps -e` command. Create a processor set using the `psrset` command, and display its ID as follows:

```
psrset -c [processor_id...]
```

where

`processor_id` is the individual or multiple processor numbers.

Bind a processor set to BIRT iServer as follows:

```
psrset -b processor_set_id Actuate_pid
```

where

- `processor_set_id` is the ID returned by the `psrset -c` command.
- `Actuate_pid` is the process ID of the PMD, as reported by `ps -e`.

## Binding to multiple-core CPUs

CPU binding is done at the operating system level, which means that BIRT iServer can bind to any logical CPU. For example, the UltraSPARC T1 processor has eight cores and four threads per core, which is a total of 32 logical CPUs to the operating system.

BIRT iServer System can bind to any logical CPU to the granularity of a thread, not just to the core on a T1 system, which is the same as binding to one physical CPU on a 32-CPU system. The commands to bind to a thread or logical CPU on a T1 system are the same as binding to a physical CPU as shown in the previous examples.

Binding on different logical CPUs can have different effects in terms of Actuate system throughput due to scalability factors across cores and threads. For

example, on a T1 system, binding to logical CPU 0, 4, 8, and 12, which belong to four separate cores, has better overall throughput than binding to logical CPU 0, 1, 2, and 3, which belong to same core.

## **Binding BIRT iServer to processors on an HP-UX 11i machine**

You can perform processor set binding on an HP-UX 11i machine. The software for creating a processor set runs only on HP-UX 11i or later. Before you use processor sets, you must install the HP-UX 11i June 2004 or later Quality Pack and download the software to create processor sets from the HP web site. The software to create processor sets is not installed with HP-UX 11i.

To create a processor set, use the `pset_create()` system call or `psrset` command. Then, you can bind a process to the set using the `pset_bind()` system call or the `psrset` command.

Like Solaris processor sets, HP-UX 11i processor sets are non-overlapping groups of processors. You can download HP-UX 11i processor sets at no charge from the following location:

<http://www.software.hp.com>

The HP UX 11i `psrset` utility controls the management of processor sets. Processor sets allow you to isolate a subset of processors for use by specific threads and processes. Processes in a set have equal access to CPU cycles on their cores through the HP-UX standard scheduler.

To bind to one or more CPUs, use the `pbind` or `psrset` commands, as described for Solaris. For more information about processor binding and the processor binding commands, see your HP-UX 11i documentation.

## **Checking BIRT iServer bound processors**

BIRT iServer performs the following bound processor checks:

- The number of processors a cluster uses
- The set of bound processors

## **Determining the number of processors an iServer System uses**

When the PMD starts the first `encysrvr11` process on a machine, the PMD determines the number of processors to which BIRT iServer is bound and stores the list of bound processors.

If you change the processor binding, BIRT iServer does not recognize the changes until you shut down all `encysrvr11` processes on the machine and restart one of the `encysrvr11` 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 iServer manages an Encyclopedia volume.
- Machine B has eight processors with BIRT iServer bound to five processors. There is no encycsrvr11 process running on the machine, only the PMD.

The cluster counts four processors, the processors on machine A. If you start an encycsrvr11 process on machine B, BIRT iServer 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 PMD on machine B to six processors, the change has no effect until you shut down all the running encycsrvr11 processes on machine B and restart an encycsrvr11 process on machine B.

After you stop the encycsrvr11 processes and restart an encycsrvr11 process on machine B, BIRT iServer 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 CPUs exceeds the number of CPUs your license permits, BIRT iServer does not start and returns an error message to Configuration Console.

## Understanding CPU binding validation while iServer is running

When BIRT iServer is running, each encycsrvr11 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 iServer writes a message with the processor information to the log file. The message contains the maximum number of processors the BIRT iServer license file permits and the following information:
  - Current and original number of bound processors
  - Current and original list of bound processors
- If configured, BIRT iServer sends an e-mail message to the administrator. The message states that the BIRT iServer 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 iServer sends to the log file.

You must rebind the encycsrvr11 process to the same processors to which it was originally bound.

During the next hour, any attempt to use the encycsrvr11 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 iServer in the cluster shuts down and writes an error to its log file.

## Understanding CPU binding validation when an Encyclopedia volume comes online

BIRT iServer uses a separate `encycsrvr11` process to manage each Encyclopedia volume on a machine. When you take an Encyclopedia volume online, the PMD starts an `encycsrvr11` process.

When the PMD starts an `encycsrvr11` process, the PMD compares the list of processors to which the `encycsrvr11` process is bound to the original list of processors to which the PMD is bound. If the lists differ:

- The `encycsrvr11` process writes an error to its log file and shuts down.
- BIRT iServer 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 iServer 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 iServer System can send e-mail messages to an administrator if a change in processor binding violates the maximum number of licensed CPUs for BIRT iServer System. To send e-mail about a CPU license problem, set up BIRT iServer System by completing the following tasks in this order:

- 1 Configure every BIRT iServer node to send e-mail.
- 2 Specify the administrator e-mail address for BIRT iServer 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 Configuration Console, and choosing System—Properties—Advanced—Cluster Operation—Administrative.

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 iServer System. Restarting applies the changes after you set or change the e-mail address.

# Part **Four**

---

**Backing Up**



# 10

## Backing up an Encyclopedia volume

This chapter discusses the following topics:

- Performing an Encyclopedia volume backup
- Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database
- Backing up and restoring an Encyclopedia volume that uses an Oracle database
- Backing up and restoring an Encyclopedia volume that uses a DB2 database

---

## Performing an Encyclopedia volume backup

When performing a volume backup, it is important to note that there are two types of data:

- **Metadata**  
Information about iServer system and Encyclopedia volume settings and data objects stored in third-party relational database management system (RDBMS) schemas
- **Data**  
iServer system and Encyclopedia volume data objects, such as designs, documents, and information objects, stored as files on disk partitions, and the `acserverconfig.xml` file containing iServer configuration information

The administrator must back up all Encyclopedia volume metadata and data to ensure the recoverability of a volume in the event of failure. In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this operation to protect critical system metadata. The administrator can restore a corrupted or missing system schema using the System Data Store Administrator utility. For more information on this utility, see “Specifying System Data Store Administrator properties,” in Chapter 4, “Upgrading BIRT iServer.”

The third-party database that contains Actuate Encyclopedia metadata is a critical component of Actuate iServer System. An Actuate system administrator must take all necessary precautions to ensure that this database is properly backed up and available to safeguard Encyclopedia 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 Encyclopedia volume metadata and data files

A complete Encyclopedia volume backup must include the following items:

- A database backup of the Encyclopedia volume schema containing the metadata
- A copy of the folders from all Encyclopedia volume disk partitions containing file data
- A copy of the `acserverconfig.xml` file containing iServer configuration information

In the Linux/UNIX Actuate 11 Service Pack 4 environment, the default `AC_SERVER_HOME` path is:

```
/home/Actuate/AcServer/
```



The default AC\_DATA\_HOME path is:

```
/home/Actuate/AcServer/data/
```

The default Encyclopedia volume path is:

```
/home/Actuate/AcServer/data/encyc
```

The default acserverconfig.xml file path is:

```
/home/Actuate/AcServer/data/config/11SP4
```

Back up the Encyclopedia volume metadata in the RDBMS at the same time that you back up the disk partition data files. 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 Encyclopedia volume metadata backup on the RDBMS must be done before the backup of the data on the disk partitions. 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 Encyclopedia volume configuration, including the following items and information:

- A database backup of the Encyclopedia volume schema containing the metadata
- The name of the Encyclopedia volume schema and user that iServer uses to connect to the RDBMS
- A copy of the acserverconfig.xml file containing iServer configuration information
- A copy of the iServer logs

## Using RDBMS and file system backup utilities

The administrator must perform the Encyclopedia volume 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 iServer is using the database. PostgreSQL, Oracle, and DB2 also provide graphical administration tools in addition to command-line tools. For more information on using these RDBMS tools to back up and restore an Encyclopedia volume, see the related sections, later in this chapter.

### **How to perform an Encyclopedia volume backup**

To back up an Encyclopedia volume, perform the following tasks:

- 1 Make sure that the autoarchive file purging process is not running.
- 2 Make an online backup of the volume schema using the tools provided by the RDBMS.
- 3 Back up the volume data files using the tools available in the operating system environment.

### **Avoiding conflict with the autoarchive file purging process**

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 Volumes—Properties—Advanced—Archiving And Purging, the administrator can specify when the file purging process runs.

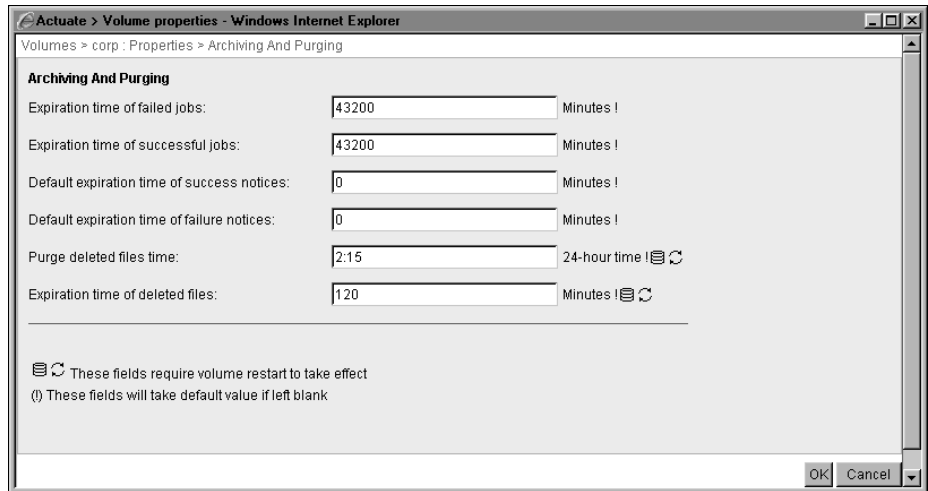
#### **How to configure Archiving And Purging**

To configure the autoarchive file purging process, perform the following tasks:

- 1 From the Advanced view of Configuration Console, choose Volumes.
- 2 On Volumes, point to the icon next to a volume name and choose Properties. In Volumes—Properties, choose Advanced. In Advanced, choose Archiving And Purging.
- 3 In Archiving And Purging, configure the following time-related file purging properties to times that do not conflict with the time when the backup operation runs, as shown in Figure 10-1:
  - 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.

Choose OK.

For information on other aspects of archiving, see Chapter 12 “Archiving files,” in *Configuring BIRT iServer*.



**Figure 10-1** Configuring file purging properties

---

## Backing up and restoring an Encyclopedia volume 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 an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the pgAdmin graphical administration tool or the `pg_dump` PostgreSQL command-line utility
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

Note that a backup of a PostgreSQL database is not portable across all operating systems.

To restore an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the pgAdmin graphical administration tool or the `pg_restore` PostgreSQL command-line utility
- Restores Encyclopedia volume data and configuration files using operating system copy commands

The following sections describe how to back up and restore an Encyclopedia volume that uses the OOTB PostgreSQL database to store the metadata. These demonstrations serve as a detailed reference example. Other RDBMS environments, such as a DB2, Microsoft SQL Server, or Oracle RDBMS, require similar procedures, which are covered in sections later in this chapter.

## Backing up an Encyclopedia volume using pg\_dump

To back up an Encyclopedia volume using the `pg_dump` utility, perform the following tasks:

- Create a folder to contain the metadata and volume data backup files
- Back up Encyclopedia volume metadata using the PostgreSQL `pg_dump` utility
- Back up the `acserverconfig.xml` file and volume data folders to the backup folder

The following example shows a typical `pg_dump` command used to export the contents of an Encyclopedia volume schema to a backup file:

```
pg_dump -F c -n ac_corp -f ac_corp_schema.dmp -h dbhost  
-p 8432 -U postgres dbname
```

This `pg_dump` command example uses the following arguments:

- **F**  
Specifies the output format. The value `c` is an abbreviation for custom, which creates a compressed archive that can be used as input to `pg_restore`.
- **n**  
Specifies the schema. Use multiple `-n` arguments to specify a list. Use wildcard notation to specify a character pattern, such as `ac_*`, to specify all volume names that start with the prefix `ac_`. If `-n` is not specified, `pg_dump` exports all non-system schemas.
- **f**  
Specifies the output file, such as `ac_corp_schema.dmp`.
- **h**  
Specifies the host name of the machine where the PostgreSQL server is running, such as `dbhost`.
- **p**  
Specifies the port where the server listens for connection requests.
- **U**  
Specifies the user name for the connection to the PostgreSQL server, such as `postgres`.

- **dbname**

Replace this string in the example with the database name, such as `actuate_db`.

Re-run the command to back up each Encyclopedia volume schema to a separate archive. To run multiple volume schema backups 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/8.4/static/libpq-pgpass.html>

Create a folder to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

In a Linux environment, create a folder to contain the metadata and volume data backup files by performing the following tasks.


#### **How to create a new backup folder**

- 1 Navigate to your home folder, which by default is:

```
/home/Actuate
```

- 2 Create the following new folder, as shown in Figure 10-2:

```
/home/Actuate/encyc_backup
```



**Figure 10-2** Creating a backup folder

Back up Encyclopedia volume metadata using `pg_dump` by performing the following tasks.

#### **How to run `pg_dump`**

- 1 Navigate to the following location:

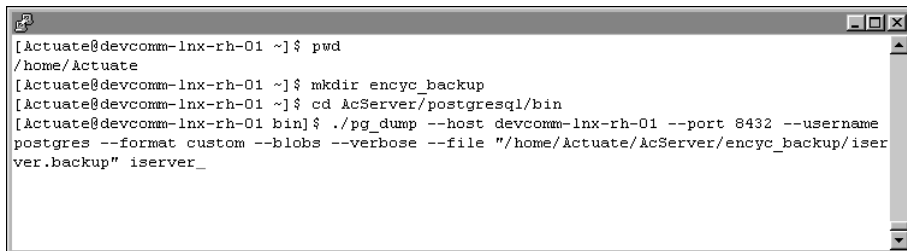
```
/home/Actuate/AcServer/postgresql/bin
```

- 2 Execute the following command. Substitute your machine name for devcomm-lnx-rh-01 in this example:

```
./pg_dump --host devcomm-lnx-rh-01 --port 8432 --username
postgres --format custom --blobs --verbose --file "/home
/Actuate/AcServer/encyc_backup/iserver.backup" iserver
```

This operation backs up the entire iserver database. If the -n argument specifying a specific schema or list of schemas is not specified, pg\_dump exports all non-system schemas. Alternatively, you can back up only one Encyclopedia volume schema, such as ac\_corp, by using the -n argument to specify a particular schema.

- 3 The command line appears as shown in Figure 10-3.



```
[Actuate@devcomm-lnx-rh-01 ~]$ pwd
/home/actuate
[Actuate@devcomm-lnx-rh-01 ~]$ mkdir encyc_backup
[Actuate@devcomm-lnx-rh-01 ~]$ cd AcServer/postgresql/bin
[Actuate@devcomm-lnx-rh-01 bin]$ ./pg_dump --host devcomm-lnx-rh-01 --port 8432 --username
postgres --format custom --blobs --verbose --file "/home/Actuate/AcServer/encyc_backup/iserver.backup" iserver_
```

**Figure 10-3** Entering the command to execute pg\_dump

- 4 Type the postgres superuser password. The administrator specified this password during the iServer installation procedure.

pg\_dump executes, writing status messages to the command prompt.

After backing up the Encyclopedia volume 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 Navigate to AC\_DATA\_HOME, which is the location of the iServer data. You specified this location on Setup Type during the install procedure. The default path for AC\_DATA\_HOME is:

```
/home/Actuate/AcServer/data
```

- 2 In AC\_DATA\_HOME, navigate to the config folder that contains acserverconfig.xml file. In Actuate Release 11 Service Pack 4, the acserverconfig.xml file is located in the config/11SP4 subfolder.

Copy acserverconfig.xml to the following backup location, as shown in Figure 10-4:

```
/home/Actuate/encyc_backup
```

```
[Actuate@devcomm-lnx-rh-01 bin]$ pwd
/home/Actuate/AcServer/bin
[Actuate@devcomm-lnx-rh-01 bin]$ cd ..
[Actuate@devcomm-lnx-rh-01 AcServer]$ cd data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock  acserverconfig.xml  acserverconfig.xml.booted  acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ cp acserverconfig.xml /home/Actuate/encyc_backup_
[Actuate@devcomm-lnx-rh-01 config]$
```

**Figure 10-4** Copying acserverconfig.xml to the backup location

**3** Navigate to AC\_SERVER-HOME/encyc.

Copy the file, fileType, status, and tempRov folders to the following backup location, as shown in Figure 10-5:

/home/Actuate/encyc\_backup

```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file  fileType  postgresql  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ cp -R file fileType status tempRov /home/Actuate/encyc_b
ackup
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-5** Copying the volume data folders to the backup location

In a backup taken immediately after an iServer installation, where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes. If these folders are not present in the environment, simply back up the file and fileType folders.

Do not back up the postgresql folder in an Encyclopedia volume backup operation. The postgres folder contains data, such as log files, from the OOTB PostgreSQL RDBMS installation, which remains active. Inadvertently including these files in an iServer backup, then accidentally overwriting the files with a stale version in a restore operation can cause problems in the PostgreSQL RDBMS installation.

The contents of the backup folder appear as shown in Figure 10-6.

```
[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ pwd
/home/Actuate
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml  file  fileType  iserver.backup  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _
```

**Figure 10-6** Viewing the contents of the backup folder

## Restoring an Encyclopedia volume using pg\_restore

To restore a backed up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia volume offline
- Delete the acserverconfig.xml and volume data folders in AC\_DATA\_HOME
- Copy the backed up acserverconfig.xml file and volume data folders from the backup folder to AC\_DATA\_HOME
- Restore the Encyclopedia volume metadata using the PostgreSQL pg\_restore utility
- Take the Encyclopedia volume online

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 an Encyclopedia volume schema:

```
pg_restore -h mydbhost -p 8432 -U postgres -d db_name
           ac_corp_schema.dmp
```

Run pg\_restore by performing the following tasks.


### How to take the Encyclopedia volume offline

- 1 In a web browser type:

```
http://localhost:8900/acadmin/config
```

Log into Configuration Console as Administrator.

- 2 On Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume offline, as shown in Figure 10-7.



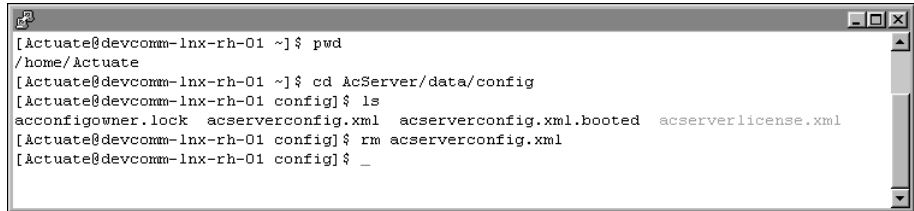
**Figure 10-7** Taking the volume offline

### How to restore the backed up volume data folders

- 1 Navigate to AC\_DATA\_HOME/config/11SP4.

Delete acserverconfig.xml, as shown in Figure 10-8.





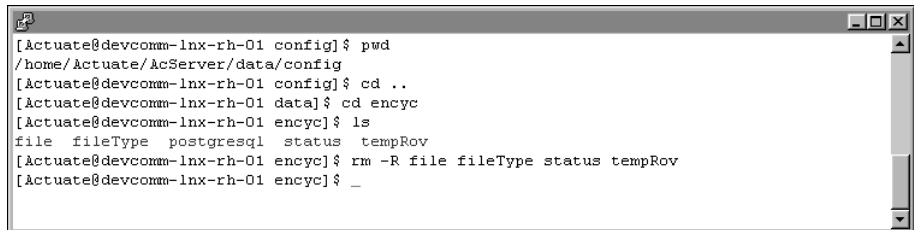
```
[Actuate@devcomm-lnx-rh-01 ~]$ pwd
/home/Actuate
[Actuate@devcomm-lnx-rh-01 ~]$ cd AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock  acserverconfig.xml  acserverconfig.xml.booted  acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ rm acserverconfig.xml
[Actuate@devcomm-lnx-rh-01 config]$ _
```

**Figure 10-8** Deleting acserverconfig.xml

- 2 In AC\_DATA\_HOME, open the encyc folder.

In AC\_DATA\_HOME/encyc, delete the file, fileType, status, and tempRov folders, as shown in Figure 10-9.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov folders may not exist. Be sure to not delete the postgresql folder.



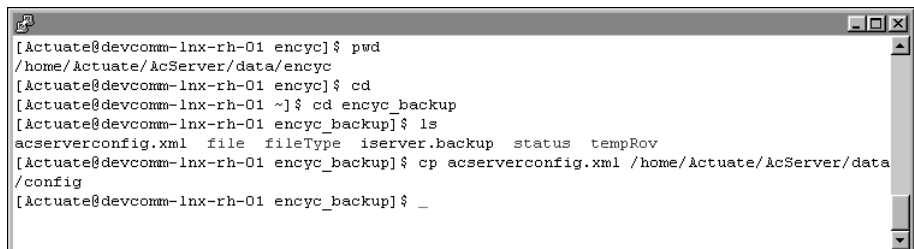
```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file  fileType  postgresql  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ rm -R file fileType status tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-9** Deleting the file, filetype, status, and tempRov folders

- 3 Navigate to the following location:

/home/Actuate/encyc\_backup

Copy acserverconfig.xml to AC\_DATA\_HOME/config/11SP4, as shown in Figure 10-10.



```
[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml  file  fileType  iserver.backup  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp acserverconfig.xml /home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _
```

**Figure 10-10** Copying acserverconfig.xml to AC\_DATA\_HOME/config/11SP4

The contents of AC\_DATA\_HOME/config/11SP4 appear as shown in Figure 10-11.

```
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ pwd
/home/Actuate/encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock  acserverconfig.xml  acserverconfig.xml.booted  acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ _
```

**Figure 10-11** Viewing the contents of AC\_DATA\_HOME/config/11SP4

- 4 Copy the file, fileType, status, and tempRov folders to AC\_DATA\_HOME /encyc, as shown in Figure 10-12.

```
[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml  file  fileType  iserver.backup  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp acserverconfig.xml /home/Actuate/AcServer/data/
config
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp -R file fileType status tempRov /home/Actuate/
AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _
```

**Figure 10-12** Copying the volume data folders to AC\_DATA\_HOME/encyc  
The contents of AC\_DATA\_HOME/encyc appear as shown in Figure 10-13.

```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file  fileType  postgresql  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-13** Viewing the contents of AC\_DATA\_HOME/encyc

### How to run pg\_restore

- 1 Navigate to the following location:  
`/home/Actuate/AcServer/postgresql/bin`
- 2 Execute the following command. Substitute your machine name for devcomm-lnx-rh-01 in this example:  
`./pg_restore --host devcomm-lnx-rh-01 --port 8432 --username postgres --dbname iserver --clean --verbose "/home/Actuate/AcServer/encyc_backup/iserver.backup"`
- 3 The command line appears as shown in Figure 10-14.

```


[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd ../../
[Actuate@devcomm-lnx-rh-01 AcServer]$ cd postgresql/bin
[Actuate@devcomm-lnx-rh-01 bin]$ ./pg_restore --host devcomm-lnx-rh-01 --port 8432 --userna
me postgres --dbname iserver --clean --verbose "/home/Actuate/AcServer/encyc_backup/iserver
.backup"_

```

**Figure 10-14** Entering the command to execute pg\_restore

### How to take the Encyclopedia volume online

- 1 In a web browser type:  
`http://localhost:8900/acadmin/config`  
Log in to Configuration Console as Administrator.
- 2 On Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume online, as shown in Figure 10-15.



**Figure 10-15** Taking the volume online

More information about backing up and restoring an Encyclopedia volume schema using the PostgreSQL pg\_dump and pg\_restore utilities is available at the following location:

<http://www.postgresql.org/docs/8.4/static/backup.html>

## Backing up and restoring an Encyclopedia volume that uses an Oracle database

Oracle provides the Data Pump Export (expdp) and Import (impdp) command-line utilities to back up and restore a database. Oracle Data Pump utilities enable high-speed uploading and downloading of data and metadata in an Oracle RDBMS. Unlike the PostgreSQL backup and restore utilities, the Oracle Data Pump utilities run on the server, not the client, for greater efficiency.

To back up an Encyclopedia volume stored in an Oracle RDBMS, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) command-line utility
- Backs up Encyclopedia volume data and configuration file using operating system copy commands

To restore an Encyclopedia volume in an Oracle RDBMS, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility
- Restores Encyclopedia volume data and configuration file using operating system copy commands

The following sections provide more information on how to perform these backup and restore operations.

## Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)

Oracle Data Pump utilities require the database administrator to map a database dump directory to a physical directory in the file system. The Oracle RDBMS writes to and reads from this directory when performing export and import operations.

Create a directory to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

To back up an Encyclopedia volume metadata using expdp, perform the following tasks:

- Create a directory to contain the metadata and volume data backup files using sqlplus.
- Back up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility.
- Back up the acserverconfig.xml file and volume data directories to the backup folder.

### How to create a backup of the Encyclopedia volume metadata

- 1 In Linux, open a command window.
- 2 Use sqlplus, an Oracle SQL editing tool, to log in as the system administrator, as shown in the following example:

```
sqlplus system/password@db_host/dbname.actuate.com
```

The example specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name.

- 3 In sqlplus, create a backup directory to contain the Encyclopedia volume metadata, as shown in the following example:

```
SQL> CREATE DIRECTORY encyc_backup
      AS '/home/Actuate/encyc_backup' ;
SQL> exit
```

Exit sqlplus.

- 4 Back up the Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility, as shown in the following example:

```
expdp system/password@db_host/dbname.actuate.com
      SCHEMAS=corp
      DIRECTORY=encyc_backup
      DUMPFILE=ac_corp_schema.dmp
```

The expdp example uses the following arguments:

- system/password@db\_host/dbname.actuate.com  
Specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name
- SCHEMAS  
Specifies the Encyclopedia volume schema to export, such as corp
- DIRECTORY  
Specifies the directory for writing the database dump, such as encyc\_backup, created by the system administrator in sqlplus
- DUMPFILE  
Specifies the name of the output file, such as ac\_corp\_schema.dmp

After backing up the Encyclopedia volume 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 directories

- 1 Navigate to AC\_DATA\_HOME, the location of the iServer data. The administrator specified this location on Setup Type during the install. The default path for AC\_DATA\_HOME is:

```
/home/Actuate/AcServer/data
```

- 2 In AC\_DATA\_HOME, navigate to the config folder. In Actuate Release 11 Service Pack 4, the acserverconfig.xml file is located in the config/11SP4 subfolder.

Copy `acserverconfig.xml` to the following backup location:

```
/home/Actuate/encyc_backup
```

- 3 Navigate to `AC_DATA_HOME/encyc`, then copy the file, `fileType`, `status`, and `tempRov` directories to the following backup location:

```
/home/Actuate/encyc_backup
```

In a backup taken immediately after an iServer installation where there has been no activity on the system, the `status` or `tempRov` directories may not exist. These directories contain information about job details and completion notices and do not appear until a job executes. If these directories are not present in the environment, simply back up the file and `fileType` directories.

## Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp)

To restore a backed up Encyclopedia volume using Oracle Data Pump Import (impdp), perform the following tasks:

- Take the Encyclopedia volume offline.
- Delete the `acserverconfig.xml` file and volume data directories in `AC_DATA_HOME`.
- Copy the backed up `acserverconfig.xml` file and volume data directories from the backup directory to `AC_DATA_HOME`.
- Restore the Encyclopedia volume metadata using the Oracle impdp utility.
- Take the Encyclopedia volume online.

Take the Encyclopedia volume offline by performing the following tasks.


### How to take the Encyclopedia volume offline

- 1 In a web browser type:

```
http://localhost:8900/acadmin/config
```

Log in to Configuration Console as Administrator.

- 2 On Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume offline, as shown in Figure 10-16.



**Figure 10-16** Taking the volume offline

### How to restore the backed up volume data directories

- 1 In Linux, open a command window.
- 2 Navigate to AC\_DATA\_HOME/config/11SP4 directory and delete the acserverconfig.xml file.
- 3 Navigate to AC\_DATA\_HOME/encyc directory and delete the file, fileType, status, and tempRov directories.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov directories may not exist.

- 4 Navigate to the following backup directory location:

/home/Actuate/encyc\_backup

- 5 From the backup directory location, perform the following tasks:
  - 1 Copy acserverconfig.xml to AC\_DATA\_HOME/config/11SP4.
  - 2 Copy the file, filetype, status, and tempROV directories to AC\_DATA\_HOME/encyc.

Restore the Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility, by performing the following tasks.

### How to restore a backup of the Encyclopedia volume metadata

- 1 Restore the Encyclopedia volume metadata using the impdp utility, as shown in the following example:

```
impdp system/password@db_host/dbname.actuate.com
  SCHEMAS=corp
  DIRECTORY=encyc_backup
  DUMPFILE=ac_corp_schema.dmp
```

The Oracle impdp utility runs using arguments similar to the expdp utility.

Take the Encyclopedia volume online by performing the following tasks.

## How to take the Encyclopedia volume online


1 In a web browser type:

<http://localhost:8900/acadmin/config>

Log in to Configuration Console as Administrator.

2 On Simple view, choose Advanced view. Choose Volumes.

3 On Volumes, take the volume online, as shown in Figure 10-17.



**Figure 10-17** Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the Oracle Data Pump utilities, go to the following location:

[http://download.oracle.com/docs/cd/B19306\\_01/server.102/b14215/dp\\_overview.htm](http://download.oracle.com/docs/cd/B19306_01/server.102/b14215/dp_overview.htm)

---

## Backing up and restoring an Encyclopedia volume that uses a DB2 database

To back up an Encyclopedia volume in the DB2 RDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using DB2 commands
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

To restore an Encyclopedia volume in the DB2 RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume data and configuration files using operating system copy commands
- Restores Encyclopedia volume metadata using DB2 commands

The following sections show how to perform these backup and restore operations.



## Backing up an Encyclopedia volume

To back up an Encyclopedia volume, perform the following tasks:

- Create a folder to contain the backup files.
- Back up Encyclopedia volume metadata.
- Back up the `acsserverconfig.xml` file and volume data folders to the backup folder.

Create a folder to contain the volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

In a Windows environment, create a folder to contain the volume data backup files by performing the following tasks.

### How to create a new backup folder

- 1 Navigate to your home folder, which by default is:

```
/home/Actuate
```

- 2 Create the following new folder:

```
/home/Actuate/encyc_backup
```

Back up Encyclopedia volume metadata by performing the following tasks.

### How to create a backup of the Encyclopedia volume metadata

- 1 Add the path to `db2profile` to the `PATH` variable on your machine.
- 2 Execute the following command to perform an online backup of the entire database and compress the backup image:

```
BACKUP DATABASE ISERVER ONLINE  
  TO "/home/Actuate/encyc_backup"  
  WITH 2 BUFFERS BUFFER 1024 PARALLELISM 1  
  COMPRESS WITHOUT PROMPTING
```

After backing up the Encyclopedia volume metadata, back up the `acsserverconfig.xml` file and volume data directories to the backup directory by performing the following tasks.

### How to back up the volume data folders

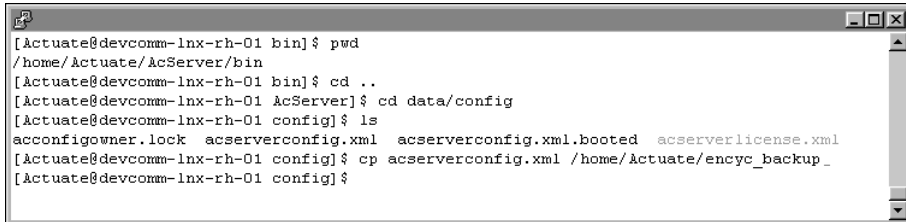
- 1 Navigate to `AC_DATA_HOME`, which is the location of the iServer data. The administrator specified this location on Setup Type during the install procedure. The default path for `AC_DATA_HOME` is:

```
/home/Actuate/AcServer/data
```

- 2 In AC\_DATA\_HOME, navigate to the config directory. In Actuate Release 11 Service Pack 4, the acserverconfig.xml file is located in the config/11SP4 subfolder.

Copy acserverconfig.xml to the following backup location, as shown in Figure 10-18:

```
/home/Actuate/encyc_backup
```



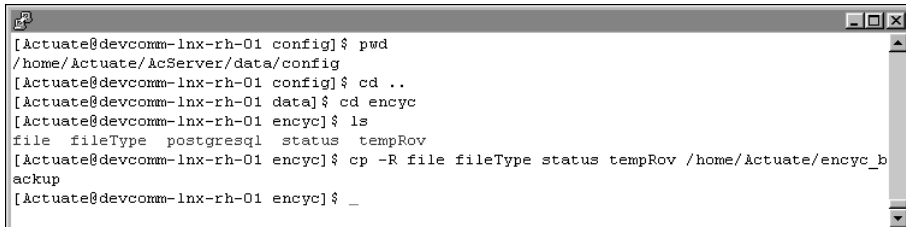
```
[Actuate@devcomm-lnx-rh-01 bin]$ pwd
/home/Actuate/AcServer/bin
[Actuate@devcomm-lnx-rh-01 bin]$ cd ..
[Actuate@devcomm-lnx-rh-01 AcServer]$ cd data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock  acserverconfig.xml  acserverconfig.xml.booted  acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ cp acserverconfig.xml /home/Actuate/encyc_backup_
[Actuate@devcomm-lnx-rh-01 config]$
```

**Figure 10-18** Copying acserverconfig.xml to the backup location

- 3 Navigate to AC\_SERVER-HOME/encyc.

Copy the file, fileType, status, and tempRov folders to the following backup location, as shown in Figure 10-19:

```
/home/Actuate/encyc_backup
```



```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file  fileType  postgresql  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ cp -R file fileType status tempRov /home/Actuate/encyc_b
ackup
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-19** Copying the volume data folders to the backup location

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes. If these folders are not present in the environment, simply back up the file and fileType folders.

The contents of the backup folder appear as shown in Figure 10-20.

```

[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ pwd
/home/Actuate
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml file fileType iserver.backup status tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _

```

**Figure 10-20** Viewing the contents of the backup folder

## Restoring an Encyclopedia volume

To restore a backed up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia volume offline.
- Delete the acserverconfig.xml file and volume data folders in AC\_DATA\_HOME.
- Copy the backed up acserverconfig.xml file and volume data folders from the backup folder to AC\_DATA\_HOME.
- Restore the Encyclopedia volume metadata.
- Take the Encyclopedia volume online.

Take the Encyclopedia volume offline by performing the following tasks.


### How to take the Encyclopedia volume offline

- 1 In a web browser type:

`http://localhost:8900/acadmin/config`

Log in to Configuration Console as Administrator.

- 2 On Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume offline, as shown in Figure 10-21.

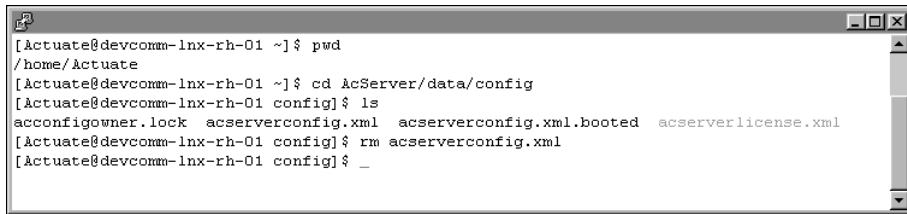


**Figure 10-21** Taking the volume offline

### How to restore the backed up volume data folders

- 1 Navigate to AC\_DATA\_HOME/config/11SP4. In Actuate Release 11 Service Pack 4, the acserverconfig.xml file is located in the config/11SP4 subfolder.

Delete `acserverconfig.xml`, as shown in Figure 10-22.



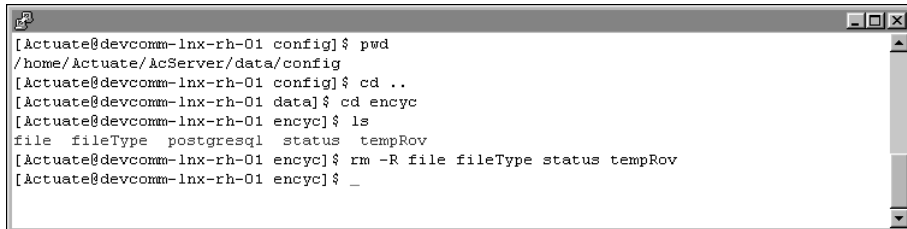
```
[Actuate@devcomm-lnx-rh-01 ~]$ pwd
/home/Actuate
[Actuate@devcomm-lnx-rh-01 ~]$ cd AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock  acserverconfig.xml  acserverconfig.xml.booted  acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ rm acserverconfig.xml
[Actuate@devcomm-lnx-rh-01 config]$ _
```

**Figure 10-22** Deleting `acserverconfig.xml`

- 2 In `AC_DATA_HOME`, open the `encyc` folder.

In `AC_DATA_HOME/encyc`, delete the file, `fileType`, `status`, and `tempRov` folders, as shown in Figure 10-23.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the `status` or `tempRov` folders may not exist.



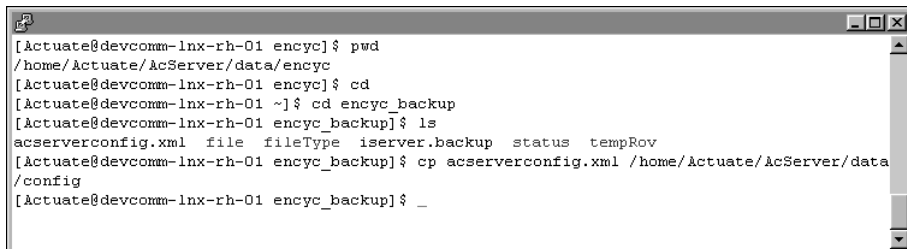
```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file  fileType  postgresql  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ rm -R file fileType status tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-23** Deleting the file, `filetype`, `status`, and `tempRov` folders

- 3 Navigate to the following location:

`/home/Actuate/encyc_backup`

Copy `acserverconfig.xml` to `AC_DATA_HOME/config/11SP4`, as shown in Figure 10-24.



```
[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml  file  fileType  iserver.backup  status  tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp acserverconfig.xml /home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _
```

**Figure 10-24** Copying `acserverconfig.xml` to `AC_DATA_HOME/config`

- 4 Copy the file, `fileType`, `status`, and `tempRov` folders to `AC_DATA_HOME/encyc`, as shown in Figure 10-25.

```
[Actuate@devcomm-lnx-rh-01 encyc]$ pwd
/home/Actuate/AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ ls
acserverconfig.xml file fileType iserver.backup status tempRov
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp acserverconfig.xml /home/Actuate/AcServer/data
/config
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cp -R file fileType status tempRov /home/Actuate/
AcServer/data/encyc
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ _
```

**Figure 10-25** Copying the volume data folders to AC\_DATA\_HOME/encyc  
The contents of AC\_DATA\_HOME/config appear as shown in Figure 10-26.

```
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ pwd
/home/Actuate/encyc_backup
[Actuate@devcomm-lnx-rh-01 encyc_backup]$ cd
[Actuate@devcomm-lnx-rh-01 ~]$ cd AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ ls
acconfigowner.lock acserverconfig.xml acserverconfig.xml.booted acserverlicense.xml
[Actuate@devcomm-lnx-rh-01 config]$ _
```

**Figure 10-26** Viewing the contents of AC\_DATA\_HOME/config  
The contents of AC\_DATA\_HOME/encyc appear as shown in Figure 10-27.

```
[Actuate@devcomm-lnx-rh-01 config]$ pwd
/home/Actuate/AcServer/data/config
[Actuate@devcomm-lnx-rh-01 config]$ cd ..
[Actuate@devcomm-lnx-rh-01 data]$ cd encyc
[Actuate@devcomm-lnx-rh-01 encyc]$ ls
file fileType postgresql status tempRov
[Actuate@devcomm-lnx-rh-01 encyc]$ _
```

**Figure 10-27** Viewing the contents of AC\_DATA\_HOME/encyc

### How to restore a backup of the Encyclopedia volume metadata

- 1 Add the path to db2profile to your PATH variable.
- 2 Execute the following command to perform an online backup of the entire database and compress the backup image:

```
db2 RESTORE DATABASE ISERVER FROM "C:\Actuate\iServer\
encyc_backup" TAKEN AT 20111004180138 WITH 2 BUFFERS BUFFER
1024 PARALLELISM 1 WITHOUT ROLLING FORWARD WITHOUT PROMPTING;
```

where


```
20111004180138
```

is the time stamp of the backup image.

Take the Encyclopedia volume online by performing the following tasks.

### How to take the Encyclopedia volume online

- 1 In a web browser type:  
`http://localhost:8900/acadmin/config`  
Log in to Configuration Console as Administrator.
- 2 On Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume online, as shown in Figure 10-28.



**Figure 10-28** Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the DB2 database utilities, go to the following location:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-0910db2incrementalbackup/index.html>

# Index

## A

- AC\_CONFIG\_HOME attribute 235, 238
- AC\_CONFIG\_HOME directory 202
- AC\_DATA\_HOME variable 205
- AC\_JAVA\_HOME variable 45, 171, 176
- AC\_JRE\_HOME variable 45, 171
- AC\_JRE64\_HOME variable 45, 171
- AC\_JVM\_HOME variable 45, 171
- AC\_SERVER\_HOME parameter 130, 133, 147
- AC\_SERVER\_HOME variable 170, 176, 184, 204
- access permissions. *See* privileges
- accessing
  - configuration parameters 13
  - data 13, 232
  - database schemas 4, 24
  - Encyclopedia volumes 45, 230, 233
  - help files xi
  - help topics xi
  - iServer features 45, 230
  - license files 13
  - online documentation 208
  - RDBMS database documentation 7
  - resources 215
  - sample volumes 24
  - volume databases 18
  - web-based content 232, 233
- accounts
  - creating dedicated iServer 24, 50
  - deploying iServer and 203
  - installing cluster nodes and 153
  - installing DB2 databases and 58
  - installing Information Console and 208
  - installing iServer and 27
  - licensing iServer and 244
  - managing 16
  - upgrades and 80
- AcEncycUpgrade utility 18
- AcExport utility 16
- AcExtern utility 17
- AcImport utility 16
- AcIntern utility 17
- AcMode utility 17
- acpmdconfig.xml 13, 152, 167
- Acrobat Catalog. *See* Adobe Acrobat Catalog
- AcServer directory 156
- acserverconfig.xml 13, 152, 168
- acserverlicense.xml 13
- activity logs 215
- AcToc utility 17
- Actuate Analytics Option 231
- Actuate Basic reports 232
- Actuate Customer Support xi
- Actuate Distribution xi
- Actuate ftp site 25
- Actuate licensing options 230, 231, 238
- Actuate product information 171
- Actuate Query Option 231
- Actuate Support Lifecycle Policy xii
- actuate.com vii
- ActuateBIRTiServer.tar.gz 176
- ActuateInformationConsole.war 214
- ActuateOne technology vii
- AcVerify utility 17
- adding
  - backup folders 253, 265
  - cluster nodes 13, 152, 165
  - database schemas 27, 50, 194, 195
  - e-mail accounts 244
  - Encyclopedia volumes 26, 198
  - external databases 190
  - indexes 5
  - JDBC drivers 18
  - license keys 234
  - licensing options 238
  - metadata databases 192
  - partitions 197
  - passwords. *See* passwords
  - processor sets 240, 241, 242
  - resource groups 44
  - schema owners 50, 194, 195
  - system databases 51, 58
  - user accounts 24, 50, 80
  - volume partitions 121
- administration tools 16

- administrative reports 19
  - administrator accounts 27
    - See also* accounts
  - administrators
    - assigning to PostgreSQL database 32
    - backing up database schemas and 24
    - backing up Encyclopedia and 248, 249, 253
    - configuring iServer clusters and 165, 168
    - customizing database connections and 18
    - deploying iServer and 174, 175
    - installing alternate databases and 50, 51, 53, 57
    - installing DB2 databases and 58, 59
    - installing iServer and 24, 43, 46
    - installing SQL Server databases and 55
    - managing Encyclopedia and 5, 6
    - managing iServer System and 15
    - migrating to current release and 46, 119
    - obtaining licenses and 230
    - optimizing iServer System and 13, 152
    - preventing data loss and 7
    - setting passwords for 34
    - sharing metadata databases and 5
    - testing new releases and 46, 47
    - upgrading Encyclopedia and 118
    - upgrading iServer and 80, 81, 93
  - Adobe Acrobat Catalog utility xi
  - aggregation 231
  - AIX servers 26, 154, 209
    - See also* UNIX systems
  - alternative databases. *See* metadata databases; third-party databases
  - Analytics Option 231
  - analyzing
    - data 231
    - search results 232
  - Apache Tomcat service 209
  - application container port 163
  - application pages. *See* web pages
  - application programming interfaces 15, 18
  - application server cache 217
  - application servers 214, 216
    - See also* servers
  - APPLICATION\_USER parameter 134, 147
  - APPLICATION\_USER\_PASSWORD parameter 134
  - applications
    - accessing Encyclopedia and 45
    - backward compatibility with 18
    - creating production environment for 47
    - creating test environments for 46
    - developing 15, 16
    - installing Open Security 45
    - restricting processes for 239, 240
    - running iServer processes and 8, 10
    - upgrading and 46, 83
  - archives (cloud deployments) 6, 174
  - Archiving and Purging page 250
  - archiving report files 250
  - asynchronous Factory service 11
  - autoarchive file purging 250
  - auto-login scripts (pgpass) 253
  - automated installation option 6
  - automated installation programs 16
  - automatic upgrades 80, 82, 83
- ## B
- backing up
    - data files 249
    - Encyclopedia volumes 81, 248, 249, 250
    - folders 254, 265
    - metadata 249, 250
    - PostgreSQL databases 251
    - report files 17, 81, 83
    - system metadata 24, 62
    - system schemas 24, 81, 145
    - volume databases 17
    - volume metadata 24, 62, 81
    - volume schemas 24, 253
  - backup operations 7, 16, 81
  - backup utilities 17, 249
  - backward compatibility 18
  - Basic reports 232
  - binary files 29, 85, 204
  - BIRT 360 Option 231
  - BIRT Data Analyzer Option 232
  - BIRT Designer Professional 232
  - BIRT Exchange 25
  - BIRT Exchange URL 25
  - BIRT Interactive Viewer Option 232
  - BIRT iServer 16
    - See also* iServer



- BIRT iServer System names 31
- BIRT open source projects vii
- BIRT Option 231
- BIRT Page Level Security option 232
- BIRT report designers 46
- BIRT reports 19, 231
  - See also* reports
- BIRT SmartSheet Security Option 232
- BIRT Spreadsheet Designer 232
- BIRT Spreadsheet Option 232
- BIRT Studio 215
- BIRT Studio Option 232
- BIRT viewer 215
- BIRT\_RESOURCE\_PATH parameter 215
- BIRT\_VIEWER\_LOG\_DIR parameter 215
- birt-exchange.com vii
- browsers. *See* web browsers
- Business Intelligence technology vii

## C

- cache (web pages) 216
- cache conflicts 216
- Caching service 9, 11
- changing
  - cluster machines 238
  - CPU binding 242, 244
  - database encoding 19
  - default locales 215
  - IP addresses 33, 34
  - license file names 237
  - licensing options 238
  - locales 33
  - metadata databases 4, 5
  - network cards 238
  - port numbers 34
  - superuser names 32
  - time zones 33, 215
  - volume names 35, 140
- character encoding 19, 58
- character sets 19
- charts 45, 171, 231
- CLASSPATH variable 45
- clearing web browser cache 217
- client applications 8
- client/server models 13
- cloud deployment environments 204
  - cloud deployments 6, 174–205
  - Cluster config location prompt 202
  - cluster IDs 13, 152
  - cluster nodes 8, 11, 13, 152, 165
    - See also* clusters
  - Cluster Server for installation option 158
  - cluster template names 202
  - clusters
    - accessing Configuration Console and 158
    - adding nodes to 13, 152, 165
    - associating with machine IDs 235
    - binding iServer processes to 239, 243
    - changing machines for 238
    - configuring 13, 14, 167, 168, 169
    - connecting to iServer System and 8
    - creating iServer images and 174
    - creating user accounts for 153
    - deploying Information Console to 216
    - deploying iServer and 177, 202
    - determining number of processors for 243
    - distributing requests among 8, 14
    - exceeding CPU licenses for 244
    - failing 152
    - failover operations and 7
    - installing iServer as 153, 155
    - installing Xvfb software for 154
    - licensing 239
    - licensing options for 233, 235
    - load balancing for 14, 152
    - managing volume metadata and 5
    - running iServer services and 8
    - running on X Windows servers 154
    - setting up RDBMS databases for 11
    - starting 169
    - storing information for 174
    - viewing nodes in 204
- collation 55, 58
- collecting machine information 235
- command line utilities 16
- comments 129
- CONFIG\_SCHEMA\_NAME parameter 134, 148
- Configuration Console
  - adding database schemas and 194, 195
  - adding metadata databases and 192
  - adding partitions and 196
  - administering iServer and 42, 78, 181

- Configuration Console (*continued*)
  - archiving report files and 250
  - configuring BIRT iServer and 16
  - connecting to databases and 191
  - creating Encyclopedia and 142, 143, 144, 198
  - deploying to clusters and 203
  - installing 158
  - licensing iServer and 233
  - logging in to 42, 181
  - migrating Encyclopedia and 119
  - running iServer clusters and 13, 158, 169
  - running iServer processes and 8
  - running stand-alone servers and 9
  - sending notifications and 244
  - setting passwords for 34, 89, 113
  - updating license keys and 237
  - upgrades and 101
- configuration files 27, 83, 158, 205
- configuration home location 158
- configuration parameters
  - See also* parameters
  - accessing 13, 152
  - Data Store Administrator 133
  - Squirrel Data Exporter 129, 130
- configuration templates 13, 152, 167, 169
- configurations
  - binding iServer processes and 239
  - connecting to databases and 18
  - creating cluster nodes and 8
  - customizing database encoding and 19
  - deploying product files and 209
  - deploying WAR files and 214, 215, 216
  - failover operations and 7
  - installing cluster nodes and 13, 158, 165, 167, 168, 169
  - installing Information Console and 214
  - installing iServer and 27
  - installing X frame buffer and 26, 154
  - licensing iServer and 230, 235
  - overwriting previous installations and 43
  - purging report files and 250
  - renaming licensing files and 238
  - running Open Security applications and 45
  - running RDBMS databases and 9, 11
  - storing metadata for 4
  - testing installations and 46
  - upgrades and 83, 129
- configuring
  - clusters 13, 14, 167, 168, 169
  - external databases 190
  - iServer System 6
  - JDBC drivers 18
  - system data store 145
  - upgrade utilities 129
  - X Windows servers 25, 26, 154
- connection pooling 18, 51
- ConnectionProperty parameter 168
- connections
  - accessing multiple volume schemas and 51
  - accessing volume databases and 18
  - changing encoding schemes and 19
  - cluster configurations and 8, 13
  - creating DB2 schemas and 59
  - creating Oracle schemas and 54
  - creating PostgreSQL schemas and 52
  - creating SQL Server schemas and 56
  - customizing database 18
  - deploying iServer and 188
  - running Information Console and 215
  - setting up database 191
  - testing 193
- context roots 38, 116, 163, 214
- context-sensitive help 222
- conventions (documentation) xii
- copying
  - license files 238
- copying files 249
- corrupted schemas 81
- CPU binding 239–244
- CPU binding validation 243, 244
- CPU-based licenses 230, 239, 243, 244
- CPUs
  - AIX servers and multiple 209
  - configuring X Windows servers for 25, 26, 154
  - deploying iServer over multi-threaded 239
  - determining number of 242
  - licensing and 239
  - restricting processes for 239, 240
  - running Information Console and 214
  - viewing maximum number of 243
- CREATE\_SCHEMA parameter 135

- creating
  - backup folders 253, 265
  - cluster nodes 13
  - data cubes 231
  - database schema owners 50, 194, 195
  - database schemas 27, 50, 194, 195
  - databases 11, 51, 58, 192
  - e-mail accounts 244
  - Encyclopedia volume partitions 119
  - Encyclopedia volumes 26, 142, 144, 198
  - indexes 5
  - iServer images 174
  - partitions 197
  - passwords. *See* passwords
  - processor sets 240, 241, 242
  - production environments 47
  - production staging areas 46
  - report designs 232
  - resource groups 44
  - test environments 46
  - user accounts 24, 50, 80
- credentials. *See* user credentials
- cube reports 231
- cubes 231
- cubeview files 232
- custom applications 46
- Customer Support xi
- customizing
  - database connections 18
  - database encoding 19
  - Information Console 214
  - metadata databases 5
  - WAR files 216

## D

- Daemon listen port 169
- dashboards vii, 231
- data
  - accessing 13, 232
  - analyzing 231
  - backing up 81, 248
  - caching 11
  - deleting 136
  - importing 137
  - preventing loss of 7, 24
  - recovering 6, 16, 81

- setting default location for 29
  - sharing 14
  - storing 175
  - upgrades and 6, 48
- Data Analyzer Option 232
- Data Connector Option 232
- data cubes. *See* cubes
- data directory 29
- Data Exporter
  - configuring 129
  - migrating iServer installations and 6
  - migrating volume metadata and 16
  - running 124
  - setting properties for 129
- data files 29, 204, 249
- data objects 5, 248
- data partitions 196
- data retrieval operations 11
- data sources 11
- data store 145
- Data Store Administrator
  - See also* System Data Store Administrator
  - configuring 129
  - migrating iServer installations and 6
  - migrating volume metadata and 16
  - running 140
  - upgrading iServer and 133, 140
- Data Store Upgrader
  - configuring 129
  - migrating volume metadata and 17
  - running 102
  - setting properties for 147, 148
  - upgrading iServer and 6
- data stores
  - See also* database schemas
- data types 18, 58
- DATA\_EXPORT\_FOLDER parameter 135
- DATA\_EXPORT\_FORMAT parameter 135
- DATA\_IMPORT\_FOLDER parameter 135
- DATA\_IMPORT\_FORMAT parameter 135
- database administration tool 251
- database clients 36, 161
- database collation 55, 58
- database drivers. *See* drivers
- database objects 190
- database schema names 31

- database schema owners
  - creating 50, 194, 195
  - installing DB2 databases and 58
  - installing iServer and 5
  - installing Oracle databases and 53, 54
  - installing PostgreSQL databases and 51, 52
  - installing SQL Server databases and 55, 56
- database schemas
  - See also* system schemas; volume schemas
  - accessing 4
  - assigning privileges 51
  - backing up 24, 253
  - creating 27, 50, 194, 195
  - customizing metadata databases and 5
  - deploying iServer and 177
  - initializing 27
  - installing DB2 databases and 59
  - installing external databases and 190
  - installing Oracle databases and 53, 54
  - installing PostgreSQL databases and 52, 53
  - installing SQL Server databases and 56, 57
  - naming 51, 194, 195
  - populating 147
  - preventing data loss and 7
  - storing metadata and 4
  - upgrades and 81
  - user accounts and 24
  - user-defined functions and 61
- database servers. *See* servers
- database types 184
- DATABASE\_HOST parameter 135, 148
- DATABASE\_NAME parameter 135, 148
- DATABASE\_PORT parameter 135, 149
- DATABASE\_TYPE parameter 133, 136, 148
- dbName property 185
- databases
  - See also* specific type
  - accessing documentation for 7
  - accessing Encyclopedia and 45
  - analyzing data and 231
  - backing up 17
  - changing metadata tables and 5
  - committing resources to 27
  - configuring external 190
  - configuring failover operations for 7
  - configuring stand-alone RDBMS 9
  - connecting to 18, 191
  - copying files in 249
  - CPU binding and 240
  - creating metadata 11, 192
  - creating system 51, 58
  - customizing 5
  - deploying iServer and 177, 182–188
  - installing alternate metadata 63
  - installing iServer and 24, 27, 50
  - localizing 19
  - managing Encyclopedia and 5, 6
  - running iServer clusters and 11, 14
  - setting passwords for 32
  - setting port number for 192
  - setting up iServer environments for 50
  - specifying administrator for 32
  - specifying default locations for 29
  - storing metadata and 4
  - upgrades and 80, 93
  - viewing incomplete files in 249
- DB2 databases
  - accessing documentation for 7
  - connecting to 59
  - creating schema owners for 58
  - creating user accounts for 58
  - extending SQL functionality 59–61
  - installing 58, 59
  - managing Encyclopedia and 5
  - naming schemas for 59
  - running volume backup and restore
    - operations with 264, 265, 267
  - setting up iServer environments for 57–59
  - storing metadata and 4
  - taking Encyclopedia offline for 267
  - taking Encyclopedia online for 270
- dbtype property 184
- decompressing localization and
  - documentation files 223
- default cluster template 203
- default directories. *See* directories
- default Encyclopedia volume 38
- default hostname 34
- default installation directory 208
- default locale 33, 159, 215
- default ports. *See* ports
- default program paths 29
- default time zone 33, 159, 215
- default values 43

- default volume name 35
- default volume partition 196
- default volume partitions 119
- DEFAULT\_DATABASE\_PASSWORD parameter 184
- DEFAULT\_DATABASE\_NAME parameter 136
- DEFAULT\_LOCALE parameter 215
- DEFAULT\_TIMEZONE parameter 215
- DEFAULT\_VOLUME parameter 215
- DELETE\_ALL\_DATA parameter 136
- DELETE\_DATA parameter 136
- deleting
  - data store data 136
  - data store elements 136
  - Encyclopedia volumes 142
  - partitions 196
  - product files. *See* uninstalling
- deploying
  - Information Console 214, 215, 216
  - iServer 6, 13, 14, 174–205, 239
  - new releases 46
  - spreadsheets 232
- deployment tools 214
- DÉPOPULATE\_SCHEMA parameter 136
- design files 5
- designers vii, 46
- designs 11, 46, 231, 232, 234
- desktop products 46
- developers 15, 51
- developing applications 15, 16
- DHTML formats 232
- diagnostic fixes 9, 14, 153
- digit wildcard character 59
- directories
  - backing up Encyclopedia and 255
  - backing up volume data 254, 261, 265
  - creating cluster nodes and 13
  - customizing WAR files and 216
  - deploying iServer and 175, 176, 205
  - extracting product files to 208
  - installing cluster nodes and 152, 156, 158, 165
  - installing Information Console and 208, 209, 211
  - installing iServer and 29
  - installing iServer Integration Technology and 220
  - installing JDK files and 44, 171
  - installing upgrades and 83
  - locating home 170
  - mapping Oracle database dump 260
  - migrating Encyclopedia and 118, 119
  - restoring 256, 263, 267
  - running multiple releases and 44
  - upgrading iServer and 43, 83
- directory paths 29, 85, 121, 215, 248
- disk partitions 249
- disk space 44, 200
- DISPLAY environment variable 25
- display\_value variable 25
- displaying
  - charts 45
  - cluster nodes 204
  - cube reports 231
  - images 25
  - licensing information 233, 234
  - process IDs 240, 241
  - reports 11, 25, 232
- DLLs 45
- document files 5, 234
- document generation service 11
- documentation
  - accessing vii–xi, 208
  - administering iServer System and 16
  - conventions for xii
  - deploying Information Console and 214
  - downloading xi
  - installing 222
  - managing Encyclopedia and 7
- documentation updates 222
- downloading
  - documentation files xi
  - HP-UX processor sets 242
  - JAVA SE Development Kit 175
  - product files 25, 208
  - report files 11
- drivers
  - cluster installations and 14, 161
  - database connections and 11
  - Encyclopedia connections and 18
  - iServer installations and 36
- DROP\_SCHEMA parameter 136

dump directories 260  
dynamic link libraries 45

## E

- e.Analysis Option 232
- e.Report Data Connector Option 232
- e.Report Designer Professional 232, 233
- e.Report Option 232
- e.reports 232
- Eclipse BIRT open source projects vii
- elastic iServer clustering 13, 152, 174
- e-mail. *See* notifications
- encoding 19, 58
- encyc directory 83, 119
- Encyclopedia Data Store Administrator
  - configuring 129
  - migrating iServer installations and 6
  - migrating volume metadata and 16
  - running 140
  - upgrading iServer and 133, 140
- Encyclopedia Data Store Upgrader
  - configuring 129
  - migrating volume metadata and 17
  - running 102
  - setting properties for 147, 148
  - upgrading iServer and 6
- Encyclopedia processes. *See* encycsrvr11 processes
- Encyclopedia volumes
  - accessing multiple 233
  - accessing objects in 230
  - adding partitions for 121
  - adding to database schemas 55
  - adding to existing schemas 51, 53, 57
  - backing up 81, 248, 249, 250
  - connecting to database for 18
  - controlling access to 45
  - CPU binding and 244
  - creating default 26
  - creating new 142, 143, 144, 198
  - creating partitions for 119
  - creating schema database for 53, 57, 59
  - creating schema owners for 52
  - creating schemas for 195
  - creating user accounts and 24
  - customizing metadata databases for 5
  - deleting 142
  - downloading files from 11
  - exporting 141
  - importing 125, 140
  - installing Information Console and 209, 212
  - installing metadata database for 4, 5, 24, 50, 63
  - installing sample 19
  - licensing and 230
  - localizing databases for 19
  - managing 5, 6
  - migrating metadata for 16, 119
  - migrating to current release 47, 80, 101, 118
  - naming 35, 198
  - overwriting product files and 47
  - preserving metadata for 24
  - preventing data loss for 7, 24
  - renaming 35, 140
  - restoring. *See* restoring Encyclopedia volumes
  - running iServer clusters and 14, 153
  - running iServer processes and 8
  - running multiple schemas on 51
  - specifying default 38, 215
  - specifying system name for 31
  - starting 9, 199
  - storing metadata for 4, 11
  - taking partitions online for 199
  - troubleshooting 249
  - upgrading iServer and 6, 48, 81
- encycsrvr11 processes 239, 242, 243, 244
- environment variables 13, 25, 45, 152, 170, 171
- environments 208
- error logging applications 15
- error logging reports 19
- error messages 243
- errors 216
- escape characters 129
- evaluation copies 25, 33, 160
- events 8, 15
- example reports 19
- example volume 19, 24
- Excel formats vii
- Excel spreadsheets 232

- executable files 45
- expdp command line options 261
- expdp command line utility 260, 261
- Expiration time of deleted files property 250
- expired licenses 234, 235
- export utilities 140
- EXPORT\_ALL\_DATA parameter 136
- EXPORT\_DATA parameter 136
- exporting
  - Encyclopedia volumes 141
- external data sources 11
- external database connections 188
- external databases 190
- external security system 45
- external user information 15
- extracting product files 208

## F

- Factory processes 44, 244
- Factory service 8, 11, 44
- failover operations 7, 81
- features xi, 45, 171, 230
- file I/O operations 11
- file names 234
- file paths. *See* directory paths
- file purging properties 250
- file system backup utilities 249
- file system backups 17
- file systems 5, 11
- FILE\_LOG\_LEVEL parameter 130, 137, 149
- files
  - accessing license 13
  - accessing online help xi
  - archiving 250
  - backing up 17, 81, 83, 249
  - configuring upgrade utilities and 129
  - copying 249
  - deploying iServer and 174, 177, 204
  - downloading documentation xi
  - downloading product 25, 208
  - downloading report 11
  - installing iServer and 25
  - installing JDK software and 44
  - installing online help 222, 223
  - obtaining license key 234
  - overwriting product 6, 43, 47

- purging 250
- specifying default locations for 29
- storing report-specific 5
- transmitting license keys and 237
- updating 222
- upgrading iServer and 44, 46
- firewalls 209
- fixes 9, 14, 153
- Flash gadgets vii
- folders
  - See also* directories
  - backing up 254, 265
  - creating backup 253, 265
  - deploying iServer and 176
  - restoring 256, 263, 267
- fonts 25
- formats vii
- Forrester Wave Open Source Business Intelligence report vii
- ftp distributions xi, 155, 208
- ftp site (Actuate) 25
- full installations 208

## G

- gadgets vii
- generating
  - charts 45
  - images 25
  - reports 11, 45
  - temporary documents 11
- getJDBCMajorVersion method 18
- graphical window manager 26
- graphics cards 25
- graphics displays 25
- graphs. *See* charts
- grep command 170

## H

- heap size 214
- help 208, 222
  - See also* documentation
- help files xi
- help topics xi
- helpinstall script 223
- home directory 170, 208
- host machines 27

- hostnames 33, 34
- hosts 215
- HP-UX servers
  - See also* UNIX systems
  - CPU binding and 242
  - installing Xvfb software for 26
- HTML documentation vii, xi
- HTTP port 38, 116
- HTTP server context 38
- HTTP server context root 116

**I**

- I/O operations 11
- IBM DB2 databases. *See* DB2 databases
- IDAPI applications 10, 15, 16
- images 25
- impdp command line utility 262
- import utilities 140
- IMPORT\_DATA parameter 137
- importing
  - data 137
  - Encyclopedia volumes 125, 140
- incomplete files 249
- indexed searches 5
- infoconsoleinstall.sh 209
- Information Console
  - assigning CPUs to 214
  - caching conflicts and 216
  - configuring 216
  - connecting to 215
  - customizing 214
  - deploying 209, 214, 215, 216
  - installing 208, 209
  - logging activities for 215
  - logging in to 40, 78, 179, 217
  - running clusters and 13
  - running iServer processes and 8
  - running stand-alone servers and 9
  - setting context root for 214
  - setting default Encyclopedia for 215
  - setting IP address or host name for 212
  - starting 217
  - testing installations for 217
- Information Console service 213
- Information Delivery API 15, 16
- information delivery solutions vii
- information object cache 11
- information object files 5
- information objects 231
- INITIALIZE\_DATA parameter 137
- in-place upgrades 80, 82, 83, 100
- install.properties file 184
- installation
  - alternate metadata databases 50, 63
  - cache conflicts and 216
  - cloud deployments and 175
  - Configuration Console 158
  - DB2 databases 58, 59
  - desktop products 46
  - disk space and 44
  - documentation files 223
  - external databases 190
  - Information Console 208
  - iServer 6, 24, 26–40, 50
  - iServer clusters 153, 155
  - iServer Integration Technology 220–222
  - iServer System 6, 208
  - iServer System components 30
  - Java Software Development Kit 44, 171
  - license keys 234
  - licensing prompts and 235, 236
  - Management Console 159, 162
  - metadata databases 4, 5
  - online documentation 222, 223
  - Open Security application 45
  - previous releases and 44, 46
  - sample Encyclopedia volume 19
  - sample reports 19
  - SQL Server databases 55
  - testing 46, 217
  - upgrades and 43, 93
  - X frame buffer 26, 154
  - X Windows servers 36
  - Xvfb software 25, 26, 154
- installation directories 29
- installation guides 16
- installation prerequisites 24, 50, 208
- installation programs 16
- installation requirements 208
- installation scripts 27, 94, 209, 220, 223
- Integration service 9, 11
- Interactive Viewer 232
- international character sets 19



- international customer licensing 239
  - IP addresses 33, 34
  - iportal context root 217
  - iServer
    - accessing functionality 230
    - binding CPUs to 214
    - changing CPU binding and 242, 244
    - checking bound processors for 242–244
    - configuring system data store for 145
    - configuring X Windows servers for 25, 26, 154
    - creating data objects and 5
    - deploying 6, 13, 14, 174–205, 239
    - downloading evaluation copies for 25
    - initializing 27
    - installing as cluster 153, 155
    - installing stand alone 6, 24, 26–40, 50
    - integrating LDAP servers with 36
    - locating home directory for 170
    - optimizing performance for 14, 153
    - running 14, 153
    - setting machine names for 34
    - setting port number for 34
    - setting up user accounts for 24, 50, 80
    - starting 9, 35, 43, 113, 182, 189
    - stopping 43, 182
    - testing new releases for 46–48
    - upgrading 6, 43, 80, 82, 93, 100
    - viewing licensing information for 233, 234
  - iServer distribution package 175
  - iServer distribution setup script
    - alternate databases and 182, 184, 186
    - deploying to clusters 202, 203
    - environment variables and 176
    - installation options and 176
    - Postgre databases and 179
    - running 177
    - setting up iServer only 188
  - iServer images 174
  - iServer Integration Technology 16, 220–222
  - iServer Process Manager Port Number setting 169
  - iServer processes
    - binding to CPUs 239–244
    - cluster configurations and 153
    - creating resource groups and 44
    - running 8, 152
    - upgrading and 83
  - iServer services 8
    - See also* specific iServer service
  - iServer servlet container 8
  - iServer System
    - administering 15
    - backup operations and 253
    - configuring 6
    - deploying new releases and 46
    - downloading files for 25, 208
    - installation options for 6
    - licensing options for 230, 231, 238
    - optimizing 13, 152
    - preventing data loss for 7, 24
    - running multiple releases and 44, 46
    - storing metadata for. *See* system databases; system metadata
    - third-party databases supported 4
    - upgrading 80, 81, 82
  - iServer System components 30, 208, 224
  - iServer System names 31
  - isinstall script 27, 83
  - isitinstall script 220
- ## J
- J2EE installations 214
  - JAR files 45
  - Java developer guide 16
  - Java Factory service. *See* Factory service
  - Java heap 214
  - Java Object Interface 45
  - Java Report Server Security Extension 16
  - Java Runtime Environment 18, 171
  - JAVA SE Development Kit 175
  - Java Server Pages. *See* JSPs
  - Java Software Development Kit 44, 171
  - JDBC drivers 11, 18
  - jdbcCompliant method 18
  - JDK files 44, 171
  - jobs 10, 166
  - JRE environment 18, 171
  - JSPs 216
  - JVM libraries 45
- ## K
- keys. *See* license keys

## L

- LDAP configuration file 45
- LDAP security database 45
- LDAP servers 36, 45
- libraries 25, 26, 45
- libstdc++ library 26
- license file names 237
- license files 13, 33
- license key file locations 235
- license key file names 234
- license key files 235
- license key installations 234
- license keys 233, 237, 239
- licensed CPUs 239, 243, 244
- licenses 230, 233
- licensing options 230, 231, 238
- licensing policies 16
- links (documentation) xi
- Linux servers
  - backing up Encyclopedia and 253
  - deploying iServer to 175
  - downloading product files for 208
  - getting home directory for 170
  - installation prerequisites for 24, 50
  - installing cluster nodes for 153, 155
  - installing Information Console for 208, 209, 214
  - installing iServer for 26, 27
  - installing iServer Integration Technology for 220
  - installing JDK files for 44, 171
  - installing Xvfb software for 25, 26
  - setting classpaths for 45
  - setting default locale for 33
  - setting run level for 26
  - setting up user accounts for 24, 50, 80
  - testing new releases for 46–48
  - upgrading iServer for 43, 80, 82
- load balancing (clusters) 14, 152
- Localemap.xml 215
- locales
  - changing 33
  - encoding character sets for 19
  - obtaining licenses for 239
  - selecting 33
  - setting default 33, 159, 215

- updating resource files for 222
- localhost parameter 215
- Localization and Online Documentation
  - Update tool 222
- log files 29, 215, 243
- LOG\_FILE\_COUNT parameter 130, 137, 149
- LOG\_FILE\_LOCATION parameter 215
- LOG\_FILE\_NAME parameter 137, 149
- LOG\_FILE\_SIZE parameter 131, 137, 149
- LOG\_FOLDER parameter 131, 133, 137, 148, 149
- logging applications 15
- logging in to
  - Configuration Console 42, 181
  - Information Console 40, 78, 179, 217
  - Management Console 41, 78, 180
- logging levels
  - Data Store Upgrader 149
  - Squirrel Data Exporter 130
- logging operations 130
- losing data 7, 24

## M

- machine IDs 235, 238
- machine information 235
- machine names 33, 34, 44, 209
- Management Console
  - configuring clusters and 13
  - creating Encyclopedia and 200
  - installing 159, 162
  - licensing and 230
  - logging in to 41, 78, 180
  - managing Encyclopedia and 16
  - running iServer processes and 8
  - running stand-alone servers and 9
  - setting context root for 163
  - setting default volume for 38
  - setting host names for 160, 162
  - setting machine names for 33
  - setting passwords for 160
  - setting port number for 38
- manual in-place upgrades 80, 82, 83, 100
- manual iServer system upgrades 6
- manual side-by-side upgrades 80, 82, 104
- manual volume upgrades 147
- manuals. *See* documentation

- mapping fonts 25
- master cluster nodes 8
- master index file xi
- message routing (requests) 10, 13
- metadata
  - See also* system metadata; volume metadata
  - archiving and 250
  - backing up 24, 62, 81, 249, 250
  - changing 5
  - defined 248
  - migrating 16
  - preventing loss of 24, 81
  - restoring 263
  - setting default location for 29
  - storing 4, 11, 145
  - upgrades and 6
- metadata databases
  - See also* third-party databases
  - backing up 17
  - creating 192
  - customizing connections for 18
  - deploying iServer and 177, 182
  - editing 5
  - incomplete files and 249
  - installing alternate 4, 5, 50
  - installing PostgreSQL 27
  - naming 192
  - previous releases and 4
  - selecting alternate 4
  - setting administrator for 32
  - setting port number for 192
  - setting up iServer environments for 50
  - sharing 5
- Microsoft Excel spreadsheets 232
- Microsoft SQL Server databases. *See* SQL Server databases
- migration
  - customizing metadata databases and 5
  - installing new releases and 26, 46
  - performing side-by-side 118
  - upgrading and 6, 43, 80
- migration tools 16, 47
- missing functionality 216
- missing product files 208
- missing system schemas 81
- multicast configurations 8

- multidimensional data analysis 231
- multiple CPU environments 209
- multiple Encyclopedia volumes 233
- multiple product releases 44, 46
- multiple upgrade installations 44
- multiple-core CPU binding 241
- Multi-Tenant Option 233
- multi-threaded CPUs 239
- multi-user mode 26

## N

- name property 185
- named-user licenses 230
- naming
  - data partitions 197
  - database schema owners 51, 52
  - database schemas 51
  - Encyclopedia volumes 35, 198
  - iServer System 31
  - metadata databases 192
  - system schemas 194
  - users 51
  - volume schemas 57, 195
- naming conventions 51
- network cards 238
- Network File Systems (NFS) 11
- networked environments 11, 13, 26, 27, 152
- New Metadata Database page 192
- New System Schema page 194
- New Volume page 198
- New Volume Schema page 195
- NEW\_SCHEMA\_NAME parameter 130, 138
- NEW\_VOLUME\_NAME parameter 130, 138
- newinformationconsole.war 216
- node keys 235
- node-key licensing 233, 234, 235
- notifications 166, 243, 244

## O

- obsolete command line utilities 16
- obsolete features 7
- obsolete product information xii
- online documentation
  - accessing vii–xi, 208
  - administering iServer System and 16
  - conventions for xii

- online documentation (*continued*)
    - deploying Information Console and 214
    - downloading xi
    - installing 222
    - managing Encyclopedia and 7
  - online help. *See* online documentation
  - Open Security applications 45, 83
  - open source projects vii
  - opening
    - Configuration Console 42, 181
    - Information Console 41, 180
    - Management Console 41, 180
  - Openmotif bundle 26
  - operating system user accounts 58
  - operating systems 19, 171, 239
  - optimizing iServer System 13, 152
  - options (licensing) 230, 231, 238
  - Oracle Data Pump Export utility 260–262
  - Oracle Data Pump Import utility 262–264
  - Oracle databases
    - accessing documentation for 7
    - connecting to 54
    - CPU binding and 240
    - creating schema owners for 53, 54
    - installing as schema database 63–77
    - managing Encyclopedia and 5
    - running volume backup and restore operations with 259, 260, 262
    - setting up iServer environments for 53–55
    - storing metadata and 4
    - taking Encyclopedia offline for 262
    - taking Encyclopedia online for 264
    - upgrades and 80, 93
  - ORACLE\_TNS\_NAMES\_FILE parameter
    - 138, 150
  - output 11
  - output formats vii, 11
  - overwriting previous installations 43
  - overwriting product files 6, 43, 47
- P**
- Page Level Security Option 233
  - page-level security 232, 233
  - parameters
    - accessing Information Console files and 217
    - configuring clusters and 13, 152
    - Data Store Administrator 133
    - encoding volume databases and 19
    - Squirrel Data Exporter 129, 130
    - upgrading Encyclopedia and 147, 148
  - partition names 197
  - partition paths 121, 197
  - partitions (disk) 249
  - partitions (volume) 119, 121, 196
  - Partitions page 119, 196
  - password property 185
  - passwords
    - connecting to databases and 192
    - creating schema owners and 194, 195
    - deploying iServer and 203
    - installing iServer and 27
    - installing Management Console and 160
    - logging in to Configuration Console and 35
    - setting administrator 34
    - setting superuser 32
  - patches 9, 14, 153
  - paths 29, 85, 121, 215, 248
  - pbind utility 240, 242
  - PDF documentation vii, xi, 208, 222
  - performance
    - caching data and 11
    - cluster configurations and 14, 153
    - file input/output operations and 11
    - iServer processes and 9
    - running iServer and 42, 181
    - upgrades and 44
  - performance monitoring extensions 15
  - permanent licenses 235
  - permissions. *See* privileges
  - pg\_dump command line options 252
  - pg\_dump utility 251, 253
  - pg\_restore utility 251, 256, 258
  - pgAdmin graphical administration tool 251
  - pgpass files 253
  - plpgsql language 51, 184
  - PMD port 34
    - See also* Process Management Daemon
  - pmd11.sh 45
  - POPULATE\_SCHEMA parameter 138
  - POPULATE\_SCHEMA\_PHASE\_ONE parameter 138

- POPULATE\_SCHEMA\_PHASE\_TWO
  - parameter 138
- port numbers 44
- portNumber property 185
- ports
  - application container 163
  - cloud deployments and 175
  - daemon listening 169
  - HTTP listening 38, 116
  - Information Console 209
  - iServer 34, 215
  - Management Console 38, 162
  - multiple releases and 44
  - PMD listening 34
  - system database 32
- postgres superuser 32
- PostgreSQL command line utilities 251
- PostgreSQL databases
  - accessing documentation for 7
  - adding 11, 183
  - backing up 17, 251
  - backing up data folders for 254, 255
  - cluster configurations and 11
  - connecting to 52
  - creating schema owners for 51, 52
  - creating user accounts and 24
  - deploying iServer and 178, 185
  - file I/O operations and 11
  - installing iServer and 24, 27
  - managing Encyclopedia and 5
  - naming schemas for 52, 53
  - running volume backup and restore
    - operations with 251, 252, 256
  - setting administrator for 32
  - setting default locations for 29
  - setting passwords for 32
  - setting superuser names for 32
  - setting up iServer environments for 51–53
  - shutting down 42, 43, 181, 182
  - stand-alone configurations and 9
  - storing metadata and 4
  - taking Encyclopedia offline for 256
  - taking Encyclopedia online for 259
  - taking online 43, 182
  - upgrades and 80, 93
- postgresql directory 255
- PostgreSQL listening port 32
- PostgreSQL processes 42, 182
- PostgreSQL server clusters 5
- PostgreSQL servers 24, 42, 181
- PostgreSQL SQL language 51
- PostgreSQL superuser name 122
- prerequisites (installation) 24, 50
- previous releases 6, 43, 47, 81, 175
- printers 14
- printing 25
- privileges
  - cluster configurations and 153
  - connecting to DB2 databases and 58, 59
  - connecting to Oracle databases and 53, 54
  - connecting to PostgreSQL databases and 51, 52
  - connecting to SQL Server databases and 55, 56
  - installing cluster nodes and 166
  - installing Information Console and 208
  - running multiple schemas and 51
  - upgrading iServer and 80
- process IDs 240, 241
- Process Management Daemon
  - CPU binding and 239
  - running cluster configurations and 13, 14, 152
  - running iServer processes and 8, 10
  - setting hostnames for 33
  - setting listening port for 34
  - starting encycsrvr processes and 242, 244
- Process Management Daemon processes 240, 241
- Process Manager Port Number setting 169
- processor set binding 240, 242
- processor sets 240, 241, 242
- processor\_bind method 240
- processors. *See* CPUs
- product files 6, 25, 43, 47, 208
- product information xii, 171
- product updates xi, 222
- production environments 47
- production staging areas 46
- profile names (users) 212
- program files. *See* product files
- PROMPT\_FOR\_PASSWORDS parameter 139, 150
- prompts (installation) 27

- properties
  - Encyclopedia Data Store Administrator 133
  - Squirrel Data Exporter 129, 130
  - System Data Store Administrator 146
- properties files 129, 215
- ps command 170
- pset\_bind method 240, 242
- pset\_create method 240, 242
- psrset utility 240, 241, 242
- publishing reports 231
- Purge deleted files time property 250
- purging report files 250

## Q

- queries 11, 13, 59
- Query Option 231

## R

- RDBMS database configurations 9, 11
- RDBMS databases 4, 249
  - See also* third-party databases
- RDBMS documentation 7
- RDBMS environments 4, 14, 51, 249
- RDBMS tools 6, 7, 17
- rebinding encycsrvr11 processes 243
- recovery operations 6, 16, 81
- refreshes (web pages) 217
- relational databases. *See* databases
- release notes xi
- remote procedure calls 10
- removing
  - data store data 136
  - data store elements 136
  - Encyclopedia volumes 142
  - partitions 196
  - product files. *See* uninstalling
- renaming
  - Encyclopedia volumes 35
  - license files 237
  - postgres superuser 32
- rendering reports 25
- replacing previous releases 43
- report design files 5
- report designers vii, 46
- report designs 11, 46, 231, 232, 234

- report document files 5, 234
- report document generation service 11
- Report Encyclopedia. *See* Encyclopedia
- report files
  - archiving 250
  - backing up 17, 81, 249
  - downloading 11
  - purging 250
  - specifying default directories for 29
  - storing 5
  - upgrading iServer and 6, 44, 46
- report object executable files 45
- Report Server Security Extension 15, 16
- report templates 215
- report viewer activity logs 215
- report viewers 232
- reporting applications. *See* applications
- reporting server. *See* iServer
- reporting services. *See* specific iServer service reports
  - displaying 11, 25, 232
  - generating 11, 45
  - installing sample 19
  - printing 25
  - publishing 231
  - rendering 25
  - running 232
  - saving 232
- repositories 174
  - See also* Encyclopedia volumes
- repositoryType parameter 217
- requests
  - distributing 8, 13
  - remote procedure calls and 10
  - running cluster configurations and 14, 152
  - setting listening ports for 32, 34, 38
- resource files 222
- resource groups 44
- resources
  - creating cluster nodes and 13, 14
  - deploying iServer and 176
  - installing cluster nodes and 152, 165
  - installing iServer and 27
  - obtaining licenses and 231
- restoring Encyclopedia volumes
  - DB2 databases and 264, 267
  - PostgreSQL databases and 251, 256

- rollbacks 81
- RSSE applications 15, 16, 83
- run levels 26
- running
  - Data Store Upgrader 102
  - iServer distribution setup script 177
  - iServer processes 8, 239
  - iServer services 8
  - jobs 10
  - multiple product releases 44, 46
  - Open Security applications 45
  - pg\_dump utility 253
  - pg\_restore utility 258
  - PostgreSQL servers 24
  - queries 11, 13
  - report designs 231, 232
  - reports 232
  - spreadsheet reports 232
  - Squirrel Data Exporter 124

## S

- sample reports 19
- sample volume 19, 24
- saving reports 232
- scalability vii
- scheduling iServer installations 46
- schema names 31, 102
- schema owners
  - creating 50, 194, 195
  - installing DB2 databases and 58
  - installing iServer and 5
  - installing Oracle databases and 53, 54
  - installing PostgreSQL databases and 51, 52
  - installing SQL Server databases and 55, 56
- SCHEMA\_FILE\_NAME parameter 133, 148
- SCHEMA\_NAME parameter 139, 140, 148
- SCHEMA\_PASSWORD parameter 139, 150
- schemas
  - See also* system schemas; volume schemas
  - accessing 4
  - assigning privileges 51
  - backing up 24, 253
  - creating 27, 50, 194, 195
  - customizing metadata databases and 5
  - deploying iServer and 177
  - initializing 27
  - installing DB2 databases and 59
  - installing external databases and 190
  - installing Oracle databases and 53, 54
  - installing PostgreSQL databases and 52, 53
  - installing SQL Server databases and 56, 57
  - naming 51, 194, 195
  - populating 147
  - preventing data loss and 7
  - storing metadata and 4
  - upgrades and 81
  - user accounts and 24
  - user-defined functions and 61
- SCRIPT\_HOME parameter 133, 148
- scripts 51, 53, 57, 133, 249
  - See also* installation scripts
- search results 232
- searching online documentation xi
- security vii, 24, 80
- security application 45
- security database 45
- security roles 190
- sending notifications 244
- server context roots 38, 116
- server template names 202
- server templates 13, 152, 167, 169
- SERVER\_DEFAULT parameter 215
- ServerFileSystemSetting parameter 168
- serverName property 185
- servers
  - See also* iServer
  - binding iServer processes to 240, 242
  - caching web pages and 216
  - changing cluster machines and 238
  - configuring as cluster node 167, 168
  - configuring Xvfb 25–26, 154
  - deploying Information Console to 214, 215, 216
  - deploying to clusters and 216
  - exceeding CPU licenses for 243
  - installing Information Console on 209
  - installing stand-alone 30
  - installing X Windows 36
  - integrating Actuate products with vii, 214
  - managing volume metadata and 5
  - preventing cache conflicts for 217
  - running as clusters 13, 14, 153, 155, 169
  - running PostgreSQL 24

- servers (*continued*)
  - shutting down iServer and 9
  - starting PostgreSQL 42, 181
- service packs 248
- services. *See* specific iServer service
- servlet container 8
- setup instructions (cloud deployments) 183
- setuptools script
  - alternate databases and 182, 184, 186
  - deploying to clusters 202, 203
  - environment variables and 176
  - installation options and 176
  - Postgre databases and 179
  - running 177
  - setting up iServer only 188
- shared licenses 233, 235
- shared repository 174
- sharing metadata databases 5
- side-by-side migrations 118
- side-by-side upgrades 80, 82, 104, 125
- single processor binding 240
- single-point media failures 253
- single-point node failure 13, 152
- SmartSheet Security Option 232
- snapshots 249
- SOAP processor 10
- SOAP-based messaging 8, 10
- Solaris systems 19, 240–242
- SPINLOOPTIME variable 209
- spreadsheet reports 232
- SQL languages 51
- SQL queries 11, 13, 59
- SQL scripts 51, 53, 55, 57, 183
  - See also* installation scripts
- SQL Server databases
  - accessing documentation for 7
  - connecting to 56
  - creating schema owners for 55, 56
  - creating users for 56
  - installing 55
  - managing Encyclopedia and 5
  - naming schemas for 57
  - setting up iServer environments for 55–57
  - storing metadata and 4
- sqlplus utility 260
- Squirrel Data Exporter
  - configuring 129
  - migrating iServer installations and 6
  - migrating volume metadata and 16
  - running 124
  - setting properties for 129
- Squirrel database 4
- SQUIRREL\_DATA\_HOME parameter 130
- SQUIRREL\_EXPORT\_FOLDER parameter 130
- SQUIRREL\_EXPORT\_FORMAT parameter 130
- stand alone servers 177
- stand-alone configurations 9, 30
- starting
  - Configuration Console 42, 181
  - Data Store Upgrader 102
  - Encyclopedia volumes 9, 199
  - Information Console 41, 180, 217
  - Information Console service 213
  - iServer 9, 35, 43, 113, 182, 189
  - iServer clusters 169
  - iServer processes 239
  - Management Console 41, 180
  - pg\_dump utility 253
  - pg\_restore utility 258
  - PostgreSQL servers 42, 181
  - Squirrel Data Exporter 124
  - volume partitions 199
- startup scripts 45, 213
- stopping
  - iServer 43, 182
  - iServer processes 83
  - PostgreSQL processes 42, 182
- Sun operating systems 26, 154
- superuser name 32
- SUPERUSER parameter 139
- superuser password 27, 32
- SUPERUSER\_PASSWORD parameter 139
- Support Lifecycle Policy 171
- Supported Products Matrix xii, 171
- synchronous Factory service 11
- syntax conventions (documentation) xii
- system. *See* iServer System
- system administrators. *See* administrators
- system data store 145
- System Data Store Administrator 81, 145
- system databases 4, 5, 51, 58
  - See also* third-party databases



- system failover operations 7
- system metadata
  - backing up 24, 62
  - preventing loss of 81
  - storing 4, 11
- system names 31
- system schema configurations 4
- system schema names 31, 51, 102
- system schema owner 51, 53, 55, 58, 194
- system schemas
  - See also* database schemas
  - backing up 81
  - creating 27, 57, 194
  - naming 52, 59, 194
  - running DB2 databases and 59
  - running PostgreSQL databases and 52
  - upgrading iServer and 145
- SYSTEM\_NAME parameter 184

## T

- table of contents (documentation) xi
- tables 5
- TABLESPACE\_LOCATION parameter 139
- TABLESPACE\_NAME parameter 139
- taking Encyclopedia online 9, 199
- tar command 223
- technical support xi
- TEMP\_FOLDER\_LOCATION parameter 215
- template names 202
- templates 152, 167, 169, 215
- temporary directories 216
- temporary documents 11
- temporary files 215
- temporary license keys 234
- temporary licenses 233, 235
- tempRov directory 166
- testing
  - data partitions 197
  - database connections 193
  - desktop products 46
  - Information Console installations 217
  - new releases 46–48
- text files 237
- third-party database schemas. *See* database schemas

- third-party databases
  - See also* specific database type
  - accessing documentation for 7
  - changing tables in 5
  - configuring stand-alone 9
  - creating indexes for 5
  - customizing connections for 18
  - deploying iServer and 177, 182–188
  - installing iServer and 24, 50
  - managing Encyclopedia and 4, 6
  - preventing data loss for 24
  - running iServer clusters and 8, 11, 14
  - storing cluster information and 174
  - storing metadata and 4
  - upgrades and 80, 81
- third-party deployment tools 214
- third-party RDBMS tools 6
- threads 214
- time zones 33, 159, 215
- TIME\_ZONE parameter 139
- TimeZones.xml 215
- Tomcat service 209
- Transact-SQL scripts 55
- transient files 215
- TRANSIENT\_STORE\_PATH parameter 215
- transport protocol (requests) 10, 13, 152
- troubleshooting 249
- typographic conventions (documentation) xii

## U

- UNIX systems
  - backing up Encyclopedia and 253
  - downloading product files for 208
  - getting home directory for 170
  - installation prerequisites for 24, 50
  - installing cluster nodes for 153, 155
  - installing help system for 223
  - installing Information Console for 208, 209, 214
  - installing iServer for 27
  - installing iServer Integration Technology for 220
  - installing JDK files for 44, 171
  - installing Xvfb software for 25, 26
  - setting classpaths for 45
  - setting default locale for 33

- UNIX systems (*continued*)
    - setting run level for 26
    - setting up user accounts for 24, 50, 80
    - starting Information Console on 217
    - testing new releases for 46–48
    - upgrading iServer for 43, 80, 82, 93, 100
  - unprivileged user IDs 24
  - updates (documentation) viii, xi
  - updates (product) xi, 222
  - upgrade options 80, 82
  - upgrade utilities 129
  - Upgrader utility. *See* Data Store Upgrader
  - upgrades
    - application programming interfaces and 18
    - backing up files for 81, 83
    - customizing metadata databases and 5
    - determining if needed 46
    - disk space and 44
    - Encyclopedia volumes 47, 101
    - installing metadata databases and 4
    - installing product files and 81
    - iServer 43, 80, 82, 93, 100
    - licensing options and 235
    - maintaining version consistency for 25
    - migrating volume metadata and 16
    - overwriting previous releases and 6
    - previous releases and 43
    - rolling back 81
    - running iServer processes and 83
    - user accounts and 80
  - URLs
    - Actuate product information xii, 171
    - Actuate technical support xi
    - Configuration Console 42, 181
    - documentation updates xi
    - evaluation copies 25
    - Information Console 41, 180, 215
    - JAVA SE Development Kit 175
    - Localization and Online Documentation
      - Update tool 223
    - Management Console 41, 180
    - Network File Systems 11
    - PDF documentation 223
    - RDBMS documentation 7
    - release notes xi
  - usage logging applications 15
  - usage reports 19
  - user accounts
    - cluster configurations and 153
    - creating 24, 50
    - installing DB2 databases and 58
    - licensing iServer and 244
    - managing 16
    - upgrades and 80
  - user credentials 5
  - user IDs 24
  - user information 15, 45
  - user names 32, 51
  - user roles 190
  - user tracking reports 19
  - users
    - creating DB2 schemas and 58, 59
    - creating Oracle schemas and 54
    - creating PostgreSQL schemas and 52
    - creating SQL Server schemas and 56
    - licensing and 230
- ## V
- VARCHAR types 58
  - VARGRAPHIC types 58
  - View processes 244
  - View service 9, 11, 44
  - viewer activity logs 215
  - viewers 232
  - viewing
    - charts 45
    - cluster nodes 204
    - cube reports 231
    - images 25
    - licensing information 233, 234
    - process IDs 240, 241
    - reports 11, 25, 232
  - volume administrators. *See* administrators
  - volume data 248
  - volume data directories
    - DB2 environments 265, 267
    - Oracle environments 261, 263
    - PostgreSQL environments 254, 256
  - volume database connections 18
  - volume database encoding 19
  - volume databases 4, 5, 17
    - See also* third-party databases

- volume metadata
  - backing up 24, 62, 249, 260, 265
  - changing 5
  - migrating 16
  - restoring 263, 269
  - storing 4, 11
- volume names 35, 140, 198, 209
- volume partition names 197
- volume partition paths 121
- volume partitions 119, 121, 196
- volume schema configurations 4
- volume schema names 31, 52, 57
- volume schema owner 52, 54, 56, 58, 195
- volume schemas
  - See also* database schemas
  - backing up 253
  - creating 27, 195
  - naming 53, 59, 195
  - running DB2 databases and 59
  - running PostgreSQL databases and 53
  - setting privileges for 51
  - SQL Server databases and 57
- VOLUME\_NAME parameter 139
- volumes. *See* Encyclopedia volumes

## W

- WAR files 209, 214, 215, 216
- watermarks 234

- web administrators. *See* administrators
- web archive files. *See* WAR files
- web browsers 38, 40, 78, 116, 179, 216
- web pages 216, 232, 233
- web servers 214
  - See also* servers
- web services 10
- Windows systems 16, 19
  - collecting machine information for 235
- Work Unit Licenses 231

## X

- X frame buffer 25, 26, 154
- X libraries 25
- X Windows cluster configurations 154
- X Windows servers 25, 26, 36, 161
- XML configuration files. *See* configuration files
- XML files 237
- Xvfb cluster configurations 154
- Xvfb libraries 25
- Xvfb servers 25, 26
- Xvfb software 25, 26, 154
- XVFBDISPLAY variable 25

