

ActuateOne™

One Design
One Server
One User Experience

**Installing BIRT iServer
for Windows**

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

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): zlib. Adobe Systems Incorporated: Flash Player. Apache Software Foundation (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), ExoLab Project (www.exolab.org), and Intalio, Inc. (www.intalio.org): Castor. Codejock Software: Xtreme Toolkit Pro. Eclipse Foundation, Inc. (www.eclipse.org): Babel, Data Tools Platform (DTP) ODA, Eclipse SDK, Graphics Editor Framework (GEF), Eclipse Modeling Framework (EMF), and Eclipse Web Tools Platform (WTP), licensed under the Eclipse Public License (EPL). Bits Per Second, Ltd. and Graphics Server Technologies, L.P.: Graphics Server. Gargoyle Software Inc.: HtmlUnit, licensed under Apache License Version 2.0. GNU Project: GNU Regular Expression, licensed under the GNU Lesser General Public License (LGPLv3). HighSlide: HighCharts. IDAutomation.com, Inc.: IDAutomation. Jason Hsueth and Kenton Varda (code.google.com): Protocole 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): Ant-Contrib, licensed under Apache License Version 2.0. Matt Ingenthron, Eric D. Lambert, and Dustin Sallings (code.google.com): Spymemcached, licensed under the MIT OSI License. International Components for Unicode (ICU): ICU library. jQuery: jQuery, licensed under the MIT License. Yuri Kanivets (code.google.com): Android Wheel gadget, licensed under the Apache Public License (APL). 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): 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, JTidy, Simple API for CSS. XFree86 Project, Inc.: (www.xfree86.org): xvfb. ZXing authors (code.google.com): ZXing, licensed under the Apache Public License (APL).

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

Document No. 120201-2-430342 October 2, 2012

Contents

Introduction	vii
Understanding ActuateOne	vii
About the BIRT iServer documentation	vii
Obtaining documentation	x
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 Windows</i>	xiii

Part 1

Architecture

Chapter 1

Understanding Actuate BIRT iServer architecture	3
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	7
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 an account with Windows administrator privileges	24
Configuring the iServer user account	25
Configuring log on as a service privilege	26
Backing up iServer system and Encyclopedia volume metadata	26
Performing a new installation	27
Installing a new instance of BIRT iServer Release 11	27
Accessing the PostgreSQL Database Server using the pgAdmin utility	41
Understanding the iServer installation environment	44
About migrating an earlier iServer release to Actuate 11	45
About running different releases on the same machine	45
About performance and disk space issues	45
About upgrading an iServer with resource groups	46
About the Java Software Development Kit	46
Accessing JAR files for report generation	46
Gathering LDAP information	46
Following best practices	47
Using a test environment	47
Setting up a production staging area	48
Setting up a production environment	48

Chapter 3

Installing BIRT iServer using an alternative database	51
Preparing to install BIRT iServer using an alternative database	52
Creating an account with Windows administrator privileges	52
Creating the iServer system and Encyclopedia volume schemas and iserver user in an alternative database	52
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing PostgreSQL database	53
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing Oracle database	55
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing SQL Server database	57
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing DB2 database	60
Adding support for the digit wildcard character in iServer when the metadata database is DB2	61

Filtering for a user name returns an empty result if name contains certain special characters	63
Backing up iServer system and Encyclopedia volume metadata	64
Installing an Encyclopedia volume that uses an alternative database	65

Chapter 4

Upgrading BIRT iServer	81
Upgrading BIRT iServer	82
Preparing to upgrade BIRT iServer	82
Creating an account with Windows administrator privileges	82
Backing up iServer system and Encyclopedia volume metadata	83
Performing an automatic in-place upgrade for an earlier major release	84
Stopping and starting Actuate 10 BIRT iServer and NobleNet Portmapper services	85
Running the in-place upgrade on an earlier major release	87
About changes to the file structure	102
Accessing the PostgreSQL Database Server using the pgAdmin utility	103
Performing an automatic in-place upgrade for an earlier minor release	104
Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services	104
Running the in-place upgrade on an earlier minor release	107
About changes to the file structure	119
Performing a manual upgrade of an Encyclopedia volume schema for an earlier minor release	120
Performing a manual side-by-side upgrade	124
Installing a new Release 11 iServer	125
Performing a manual side-by-side migration	143
Working with iServer utilities	153
Working with Squirrel Data Exporter	153
Specifying Squirrel Data Exporter properties	153
Using the generated bulk-load script files	155
Working with Encyclopedia Data Store Administrator	157
Specifying Encyclopedia Data Store Administrator properties	157
Performing operations using Encyclopedia Data Store Administrator utility	163
Working with System Data Store Administrator	168
Specifying System Data Store Administrator properties	168
Performing operations using System Data Store Administrator utility	170
Working with Encyclopedia Data Store Upgrader	171
Specifying Encyclopedia Data Store Upgrader properties	171

Chapter 5

Installing a BIRT iServer cluster	175
Installing a BIRT iServer cluster node	176
Preparing the iServer cluster environment	177
Creating an account with Windows administrator privileges	184

Performing a wizard-based cluster node installation	184
Adding a node to a cluster	198

Chapter 6

Installing BIRT iServer in a cloud	201
Understanding a BIRT iServer cloud deployment	202
Deploying BIRT iServer in a cloud environment	202
Deploying an iServer distribution package	204
Running the setup script	205
Setting up iServer to use the out-of-the-box (OOTB) PostgreSQL database	207
Accessing Information, Management, and Configuration Consoles	210
Stopping and starting iServer and PostgreSQL processes	212
Setting up iServer to use an alternative database	214
Setting up iServer only	219
Setting up iServer only using the cloud deployment package	219
Configuring an external database	221
Using Configuration Console to configure the database connection	222
Using Configuration Console to configure the data partition and Encyclopedia volume	228
Setting up iServer to join an existing cluster	233
Understanding the cloud deployment environment	243
Specifying AC_SERVER_HOME	243
Specifying AC_DATA_HOME	244

Chapter 7

Installing Information Console	245
Preparing to install Information Console	246
About installing from an FTP download	246
About performing a full installation	246
Installing Information Console on Windows	247
Using the installation wizard	247
Using the WAR file to install	254
General deployment tasks	254
Preparing the server	254
Preparing the WAR file	255
About clusters of servers	256
Avoiding cache conflicts after installing	257
Testing the installation	257

Chapter 8

Installing iServer Integration Technology and Documentation	259
Installing iServer Integration Technology	260
Installing the localization and documentation files	264

About accessing online help	267
Searching PDF manuals using master-index.pdx	268

Part 3

Licensing

Chapter 9

Licensing BIRT iServer	273
Understanding licensing types	274
Understanding licensing options	275
Installing Actuate BIRT iServer System license files	277
About the license file	279
Collecting machine information for a license	279
About modifying a license	283
Understanding CPU binding	283
Binding a BIRT iServer to processors on a Windows machine	284
Binding to specific CPUs	284
Binding to multiple-core CPUs	286
Binding an Actuate process to a processor	286
About processors and hyperthreading	287
Checking BIRT iServer bound processors	288
Determining the number of processors an iServer System uses	288
Understanding CPU binding validation while iServer is running	289
Understanding CPU binding validation when an Encyclopedia volume comes online	289
Understanding CPU binding validation when running iServer processes	290
Configuring e-mail for CPU license problems	290

Part 4

Backing Up

Chapter 10

Backing up an Encyclopedia volume	293
Performing an Encyclopedia volume backup	294
Managing the backup and recovery of Encyclopedia volume metadata and data files ...	294
Using RDBMS and file system backup utilities	295
Avoiding conflict with the file purging process	296
Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database	297
Backing up an Encyclopedia volume using pgAdmin	298
Backing up an Encyclopedia volume using pg_dump	302
Restoring an Encyclopedia volume using pgAdmin	306
Restoring an Encyclopedia volume using pg_restore	311

- Backing up and restoring an Encyclopedia volume that uses an Oracle database 313
 - Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp) 313
 - Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp) 315
- Backing up and restoring an Encyclopedia volume that uses a SQL Server database 317
 - Backing up and restoring an Encyclopedia volume 317
 - Backing up an Encyclopedia volume using SQL Server Management Studio 318
 - Restoring an Encyclopedia volume using SQL Server Management Studio 326
- Backing up and restoring an Encyclopedia volume that uses a DB2 database 337
 - Backing up an Encyclopedia volume using DB2 Control Center 338
 - Restoring an Encyclopedia volume using DB2 Control Center 347
- Index 357**

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 and 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



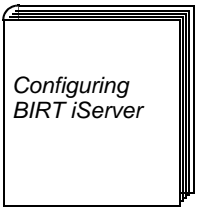
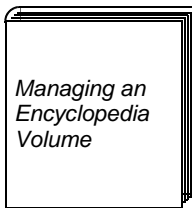
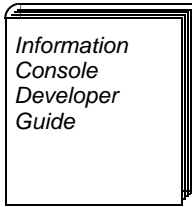
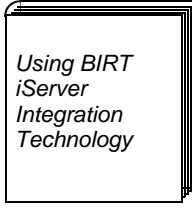
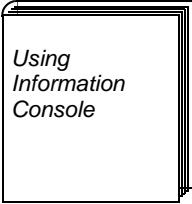
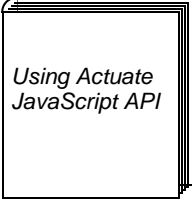

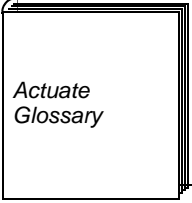
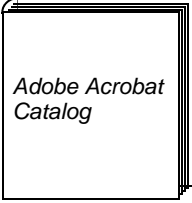
For information about this topic	See the following resource
Installing BIRT iServer for Linux and UNIX	 <p><i>Installing BIRT iServer for Linux and UNIX</i></p>
Installing BIRT iServer for Windows	 <p><i>Installing BIRT iServer for Windows</i></p>
Late-breaking information and documentation updates	Release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

Table I-1 BIRT iServer documentation (continued)

For information about this topic	See the following resource
<p>Managing an Encyclopedia Volume</p> <p>Use Management Console to:</p> <ul style="list-style-type: none">■ Set up user accounts■ Set up channels and notification groups■ Assign security roles■ Manage files and folders■ Schedule, run, and manage reports■ Back up the Encyclopedia volume■ Use Actuate Open Security	 <p><i>Managing an Encyclopedia Volume</i></p>
<p>Information Console Developer Guide</p> <ul style="list-style-type: none">■ Overview of Information Console concepts and web applications■ Using, customizing, and configuring the Deployment Kit■ Using code components for JSPs, URL parameters, JavaScript files, Java servlets, Java Beans, and security facilities	 <p><i>Information Console Developer Guide</i></p>
<p>Using BIRT iServer Integration Technology</p> <ul style="list-style-type: none">■ Overview of Actuate web services and SOAP messaging■ Managing an Encyclopedia volume■ Developing API applications using Java or .NET■ Using Java Report Server Security Extension (RSSE) APIs■ Using logging, performance monitoring, and archiving features■ Customizing the Actuate software installation process■ Actuate Information Delivery API operations and data types reference	 <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">■ Overview of Information Console concepts and online reporting■ Accessing and managing files and folders; running reports	 <p><i>Using Information Console</i></p>
<p>Using Actuate JavaScript API</p> <ul style="list-style-type: none">■ Overview of programming with Actuate JavaScript■ Creating custom pages using Actuate JavaScript■ Reference for Actuate JavaScript classes and methods	 <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. 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 Actuate11SP4/iServer/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 Program Files\Actuate11\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 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}
&or	Java OR operator	int length 4
<>	Argument you must supply	<expression to format>
<>	Delimiter in XML	<xsd:sequence>

About *Installing BIRT iServer for Windows*

Installing BIRT iServer for Windows 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 Windows 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 Windows environment.
- *Chapter 4. Upgrading BIRT iServer*. Describes how to upgrade BIRT iServer in a Windows environment.

- *Chapter 5. Installing a BIRT iServer cluster.* Describes how to install a BIRT iServer cluster node in a Windows environment.
- *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 Windows.
- *Chapter 8. Installing iServer Integration Technology and Documentation.* Describes how to install Actuate iServer Integration Technology and Documentation in a Windows environment.
- *Part 3. Licensing.* Describes how to license 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.
Upgrades 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 installing and upgrading an iServer using the automated installation programs and utilities or installing iServer in a cloud deployment, see the Installing section later in this book. 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.

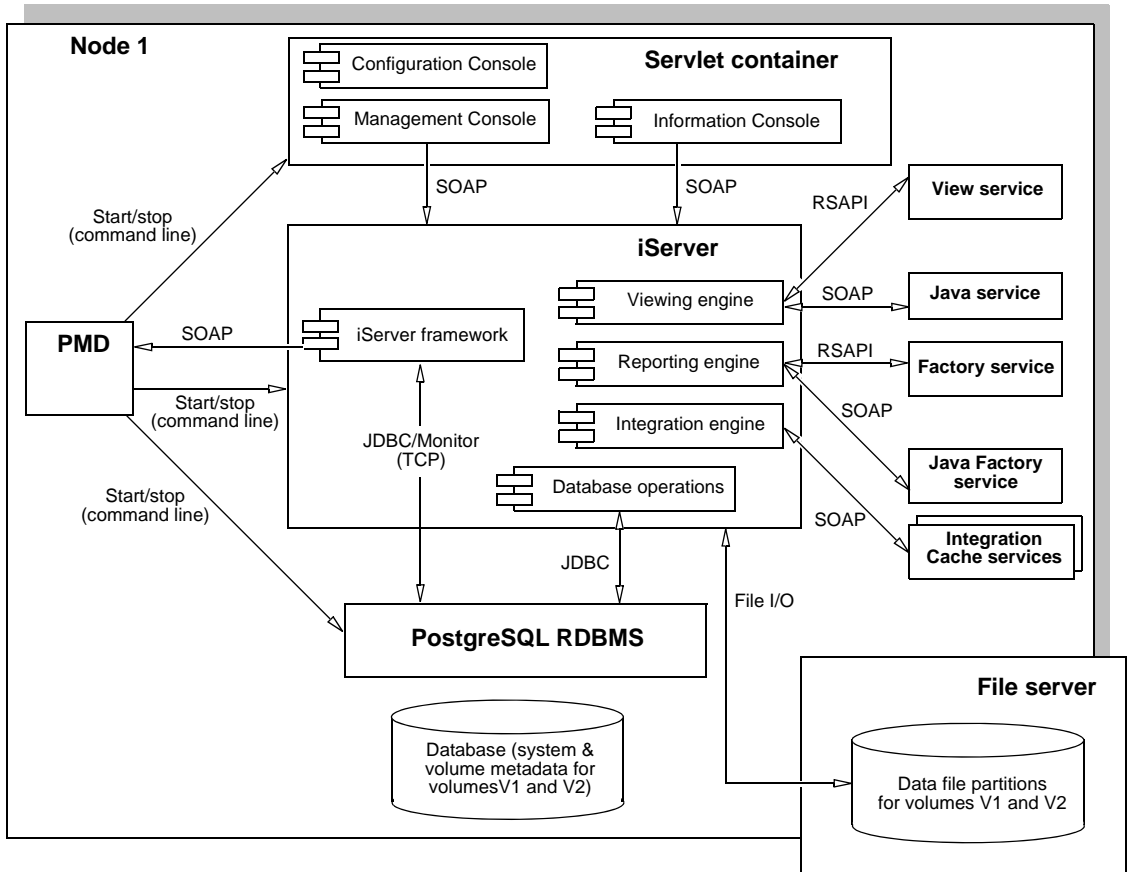


Figure 1-1 iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

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.

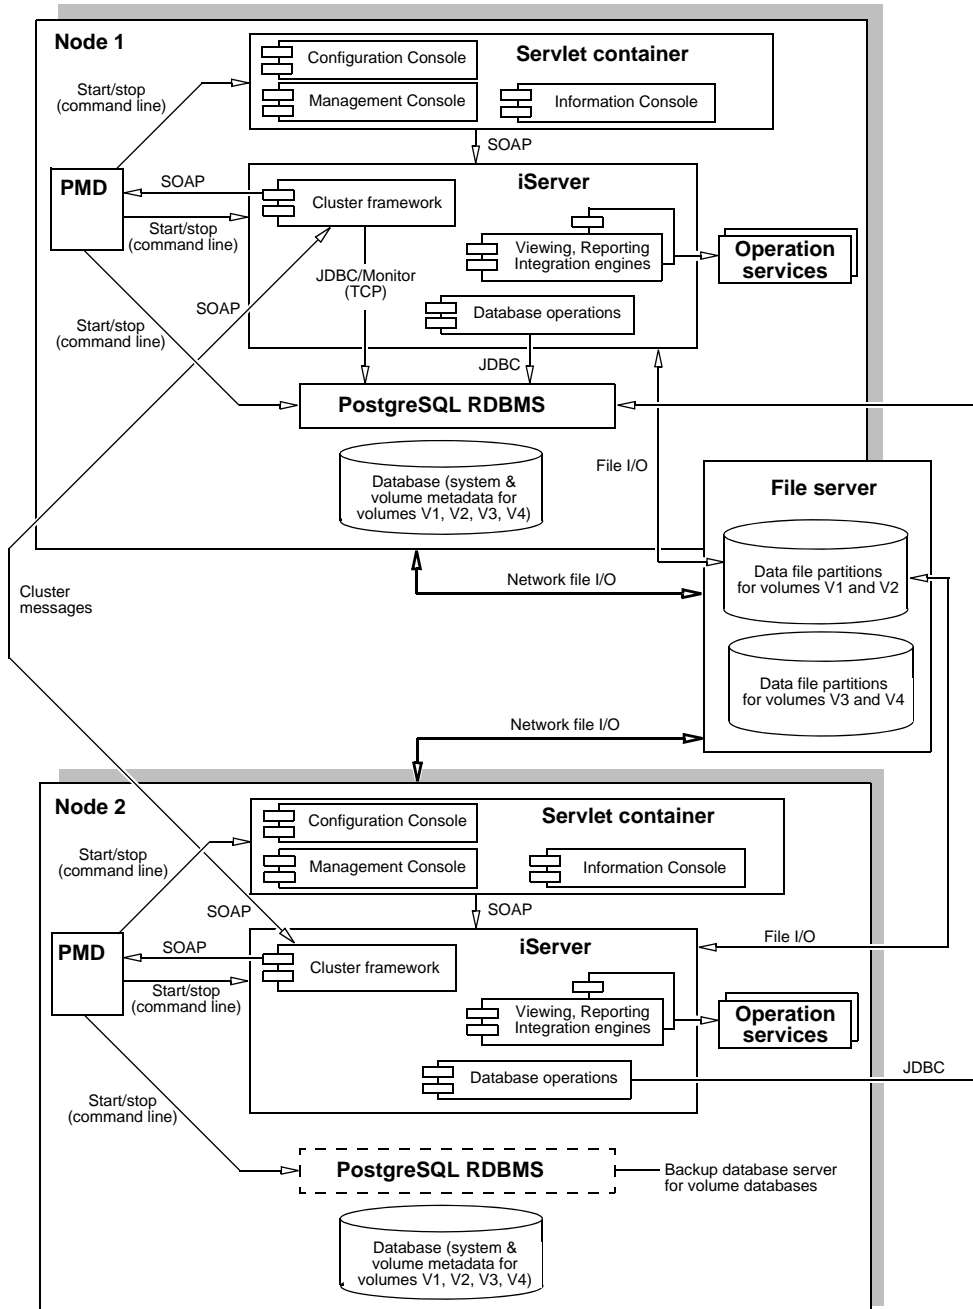


Figure 1-2 iServer RDBMS process architecture for a clustered, two-node, four-volume, OOTB database

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 cloud computing deployment option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book. 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 `acservconfig.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 backup 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 data store, 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 an account with Windows administrator privileges

Before installing iServer, create a Windows user account that is a member of the Administrators group. Use this account when installing and running iServer.

The iServer user account must meet the following requirements:

- Be a member of the Windows Administrators group.
The account must have privileges to access the required software and hardware, such as database servers, printers, and iServer files and folders.
- Have log on as a service privilege.
If the account does not meet this requirement, the iServer installation program prompts you to configure the privilege to run the Windows Actuate iServer service.

On a new Windows Vista installation, the initial user account is not a member of the Administrators group. You must configure this user account to be a member of this group.

In Windows Server 2003, when installing an iServer that uses the OOTB PostgreSQL RDBMS, create a Windows user account that is a member of the Power Users, not the Administrators group. In Windows Server 2003, PostgreSQL cannot run with Administrator privileges.

Perform the iServer installation using an account that has Administrator privileges. During the installation, when prompted to specify the user account that runs the iServer and PostgreSQL services, specify the Power User account and password. Make sure that this user account has permission to access any

required software and hardware, such as database servers, printers, and iServer files and folders.

How to create a Power User account

To create a Power User account, perform the following tasks:

- 1 In Windows, open the Command Prompt and type:
`lusrmgr.msc`
- 2 In Local Users and Groups, choose Users to display the list of users.
- 3 Double-click the user to display the properties.
- 4 In Properties—General, deselect Account is disabled, if necessary.
- 5 In Properties—Member Of, choose Add and perform the following tasks:
 - 1 On Select Groups, in Enter the object names, type:
`Power user`
 - 2 Choose Check Names then choose OK.
- 6 Exit Local Users and Groups

In a Windows installation, verify that the ICU_DATA environment variable is set to the location of the Actuate ICU library. The Actuate installation process sets ICU_DATA to the location of the library on the iServer machine. Change this setting, if necessary, to the location of the library. The following example shows the default path set by the iServer installation process:

```
ICU_DATA= C:\WINNT\system32\
```

If you plan to install iServer processes on a machine controlled by a domain server, install iServer while logged into a user account controlled by the local machine, not the domain server. When you create an iServer cluster, all iServer nodes in the cluster must be installed and run under the same user account.

Configuring the iServer user 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.

How to configure a user account with administrator privileges

To configure a user account with administrator privileges for installing and running iServer, perform the following tasks:

- 1 In Windows, open the Command Prompt and type:
`lusrmgr.msc`
- 2 In Local Users and Groups, choose Users to display the list of users.

- 3 Double-click the user to display the properties.
- 4 In Properties—General, deselect Account is disabled, if necessary.
- 5 In Properties—Member Of, choose Add and perform the following tasks:
 - 1 On Select Groups, in Enter the object names, type:
Administrators
 - 2 Choose Check Names then choose OK.
- 6 Exit Local Users and Groups

Configuring log on as a service privilege

The iServer installation program prompts you to configure log on as a service privilege if the iServer user account does not have this privilege.

How to configure the log on as a service privilege manually

To configure the log on as a service privilege manually, perform the following tasks:

- 1 In Windows Control Panel, open Administrative Tools→Local Security Policy.
- 2 In Local Security Settings, navigate to Security Settings→Local Policies→User Rights Assignments.
- 3 In User Rights Assignments, perform the following tasks:
 - 1 Open Log on as a service Properties. Choose Add User or Group.
 - 2 In Select Users or Groups, add the user name. Choose Check Names then choose OK.
- 4 Exit Local Security Settings.

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 site. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

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

Actuate also supports the cloud deployment of BIRT iServer using a ready-to-launch iServer image. For more information about this installation option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book.

The following sections describe how to install a new BIRT iServer Release 11 using the available installation options.

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 a new instance of BIRT iServer Release 11

The following section describes how to install a new instance of BIRT iServer Release 11 in the Windows operating system.

How to perform a new installation of BIRT iServer Release 11 in Windows

To install iServer, perform the following tasks:

- 1 Download the iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file,

ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 2-1. Choose Next.



Figure 2-1 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 2-2. Choose Next.

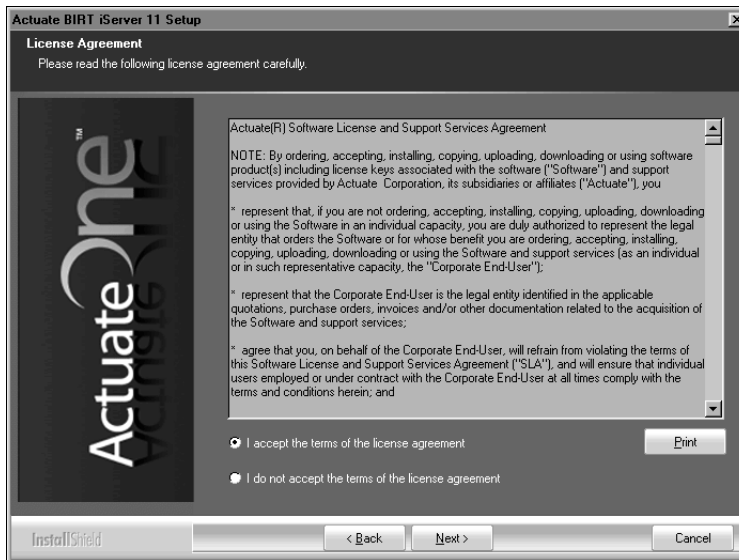


Figure 2-2 Accepting the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 2-3. Choose Next.

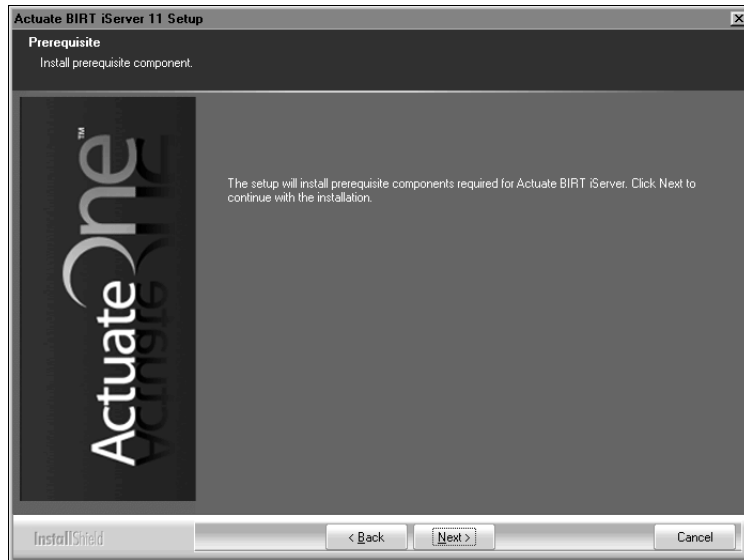


Figure 2-3 Installing Prerequisites

- 4 In Setup Type, select Typical to install the default configuration for a stand-alone iServer, as shown in Figure 2-4. Alternatively, choose Custom for one of the following reasons:
 - To install individual iServer components
 - To install a cluster node
 - To install a custom configuration for a stand-alone iServer and a supported relational database management system (RDBMS), such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance for the Encyclopedia volume metadata database
 - To use an optional Open Security application to control access to iServer using an external system, such as an LDAP server

In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations.

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, `AC_SERVER_HOME`, points to the location of the iServer binaries. The default path for the program location is `C:/Program Files/Actuate11SP4/iServer`.

iServer uses the Data location to store the iServer logs, iServer Encyclopedia, including the PostgreSQL data, and all other run-time data. The environment

variable, `AC_DATA_HOME`, points to the location of the iServer data. The default path for the data location is `C:/Actuate/iServer/data`. Choose Next.

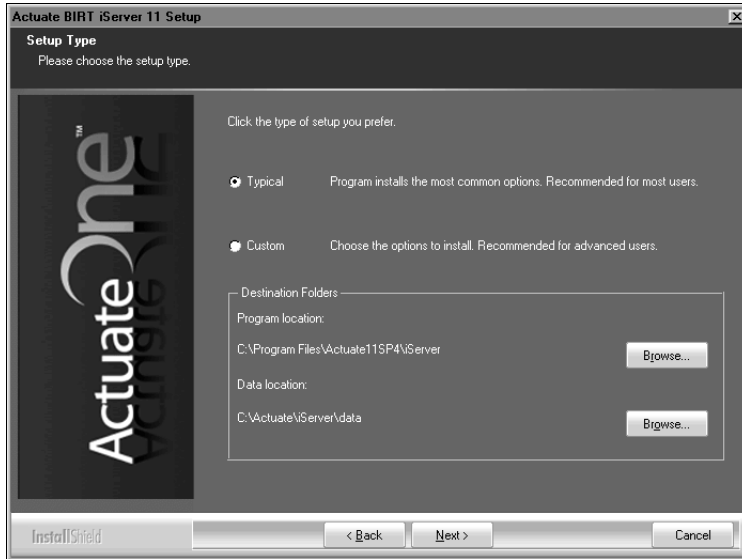


Figure 2-4 Specifying typical or custom setup type

- 5 In Encyclopedia Metadata Storage and System Name, select the type of Encyclopedia volume metadata database to install, as shown in Figure 2-5. This installation example demonstrates installing the bundled OOTB PostgreSQL database. In System Name, type a name for the BIRT iServer System name. 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. Choose Next.

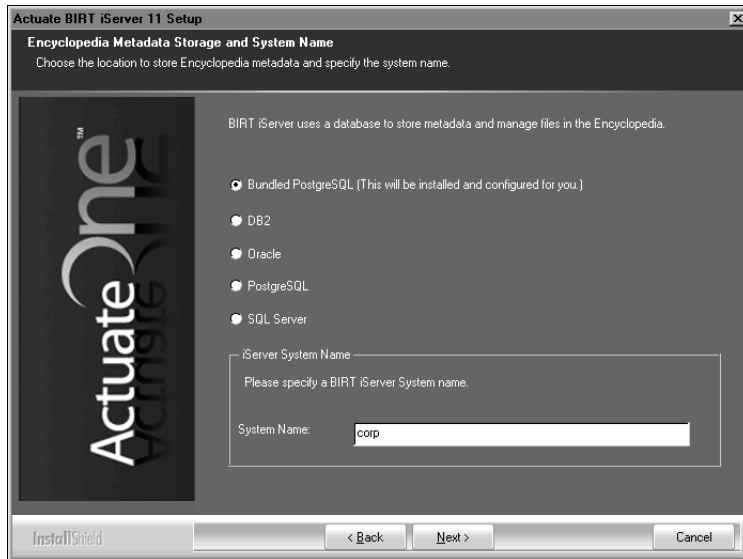


Figure 2-5 Selecting the Encyclopedia volume metadata database to install

- 6 On PostgreSQL Database Information, in Database Credentials, type and confirm a password, as shown in Figure 2-6. iServer creates the following accounts automatically, using this password for each account:
 - `postgres`
The PostgreSQL database superuser. The database superuser administers the PostgreSQL relational database management system (RDBMS).
 - `ac_<BIRT iServer System name>_system`
The System schema owner. iServer creates the iServer system schema and gives it this name. The installation program substitutes <BIRT iServer System name> with the system name you specified in the previous step.
 - `ac_<BIRT iServer System name>`
The Encyclopedia schema owner. iServer creates the Encyclopedia volume schema using `ac_<BIRT iServer System name>`.

In Port, accept the default value of 8432. Alternatively, type a different port. Choose Next.

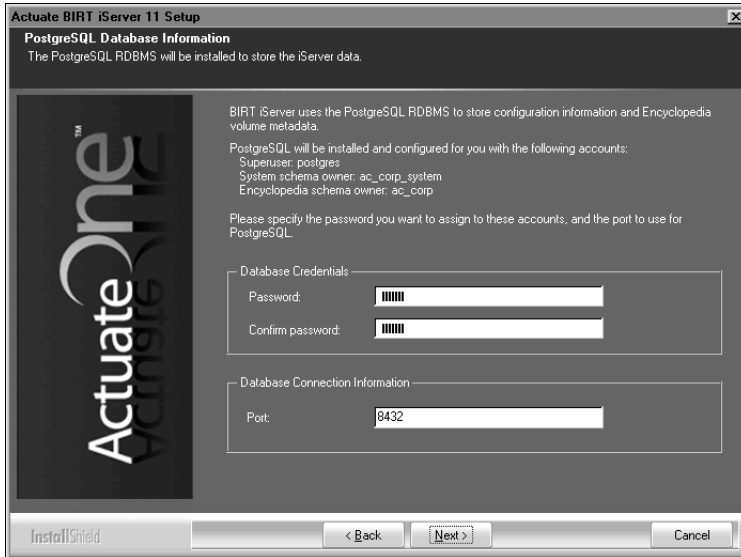


Figure 2-6 Specifying PostgreSQL database information

- 7 In License File Details, select Use the license that you purchased. Choose Browse then navigate to and choose the license file, as shown in Figure 2-7. Alternatively, choose Try out the product using the included evaluation license if you do not have a purchased license. Choose Next.

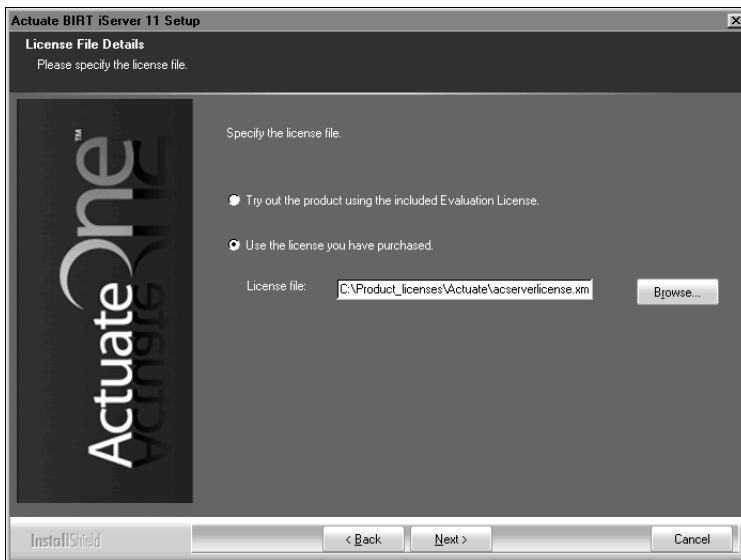


Figure 2-7 Specifying the license file

If installing using a named-user license, a prompt appears advising you to check that the volume does not exceed the number of registered users authorized by the license, as shown in Figure 2-8.




Figure 2-8 Viewing the named-user license question

Choose Yes to continue the installation.

- 8 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 2-9.

Alternatively, choose the language and locale settings for your region.

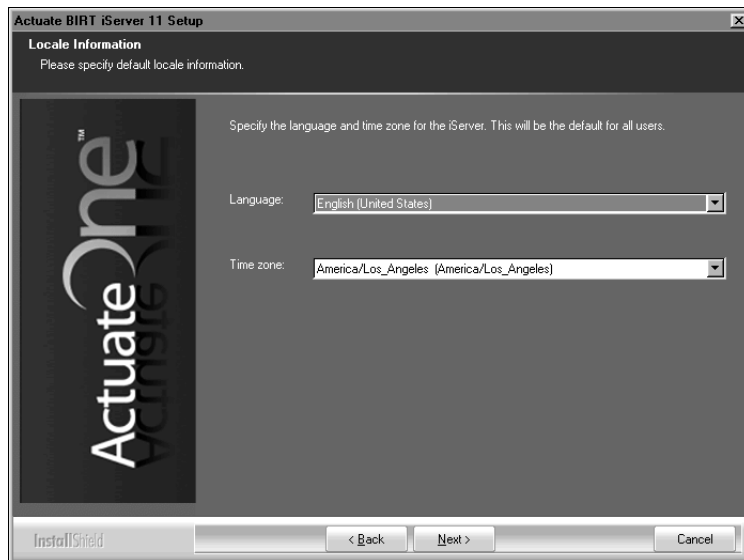


Figure 2-9 Specifying locale information

- 9 In Specify Profiles, type the user name, password, and confirm the password for the account used to start the Actuate iServer 11 service, as shown in Figure 2-10. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

If you are installing on Windows 2003, specify a user account that is in the Power Users not the Administrators group. A user account in the Administrators group cannot start the Actuate 11 PostgreSQL for BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 2-10. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

- 10 In System Configuration Password, type and confirm a password for Configuration Console, as shown in Figure 2-12. For both Configuration Console and Management Console, the default user name is Administrator. The Administrator account for Management Console has no initial password. You can log in to these consoles and change the password settings after installing iServer. Choose Next.

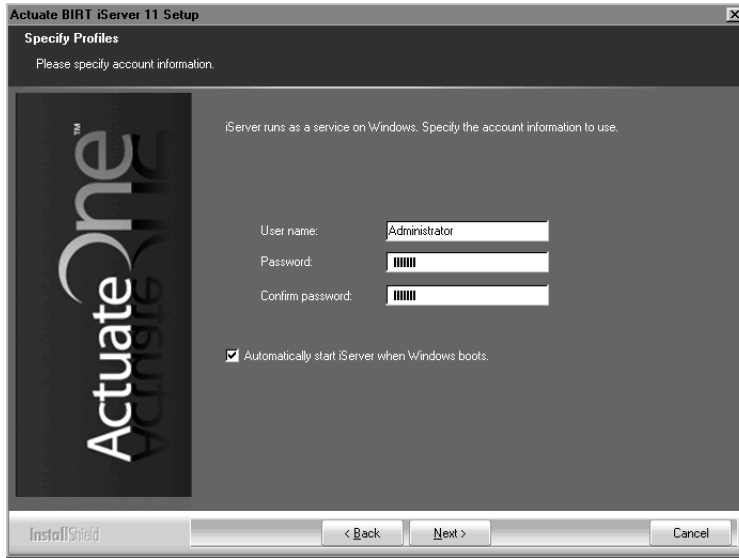


Figure 2-10 Specifying an account for running the iServer service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 2-11.




Figure 2-11 Setting the Windows local security policy

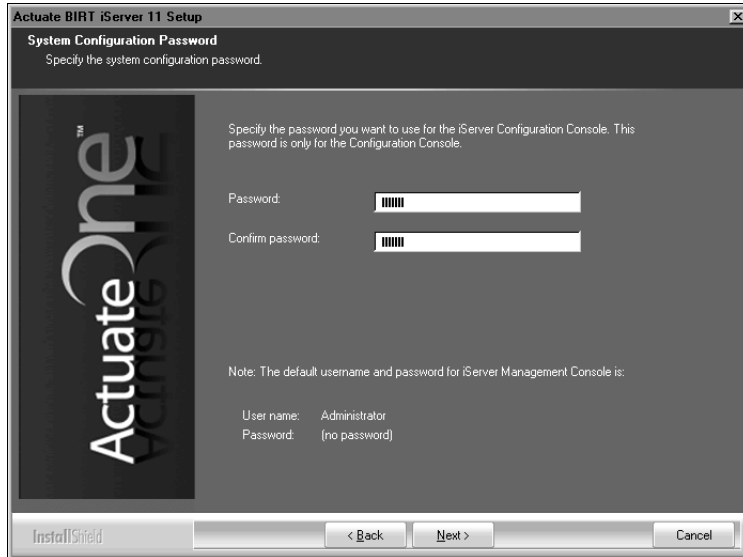


Figure 2-12 Specifying the password for using Configuration Console

11 In Start Copying Files, review the settings shown in Figure 2-13. Choose Next.

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 2-14.

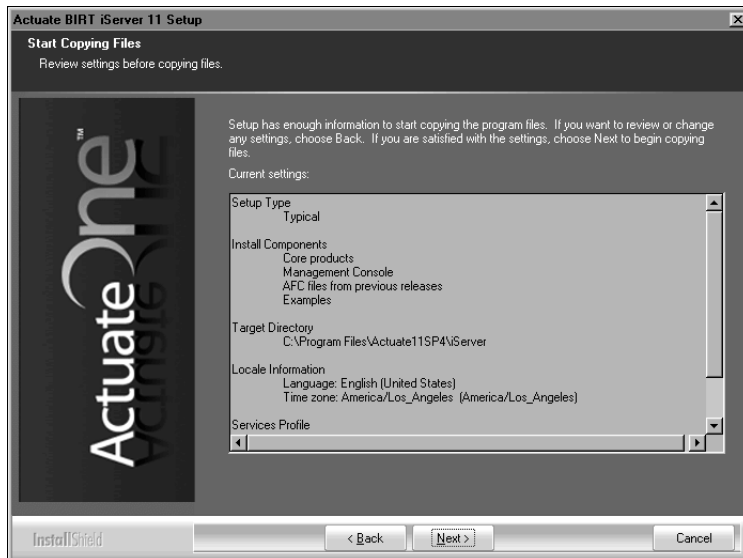


Figure 2-13 Reviewing settings before copying files



Figure 2-14 Viewing setup status

During installation, a command prompt appears, displaying the run-time commands and messages from the initialization process of the PostgreSQL RDBMS and Encyclopedia volume schema, as shown in Figure 2-15.




Figure 2-15 Initializing the PostgreSQL RDBMS

A prompt appears, asking if you want to install the pgAdmin database administration tool for the PostgreSQL RDBMS, as shown in Figure 2-16. Choose Yes.




Figure 2-16 Choosing to install the pgAdmin tool

Activate BIRT iServer Setup appears, then pgAdmin III Setup appears, as shown in Figure 2-17.

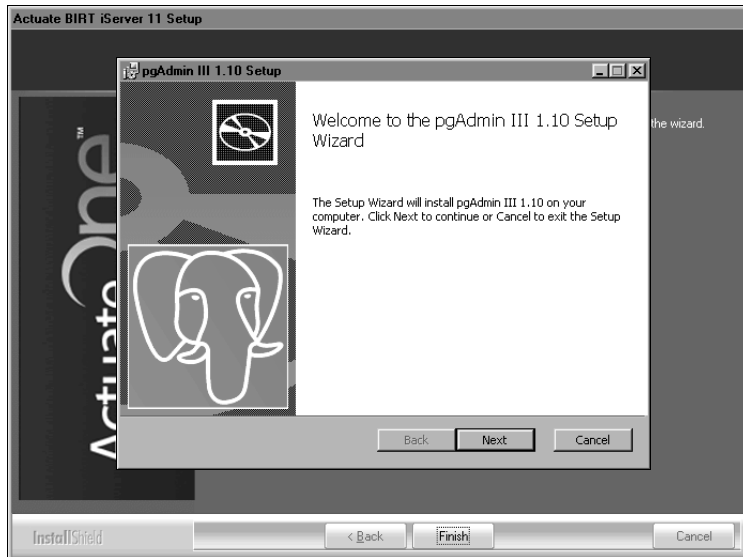


Figure 2-17 Viewing Actuate BIRT iServer and pgAdmin III Setup

- 12 In pgAdmin III Setup, perform the following tasks:
 - 1 In Welcome, shown in Figure 2-17, choose Next.
 - 2 In End-User License Agreement, select I accept the terms in the License Agreement, as shown in Figure 2-18. Choose Next.



Figure 2-18 Accepting the license agreement

- 3 In Custom Setup, review the features to be installed, as shown in Figure 2-19. Choose Next.

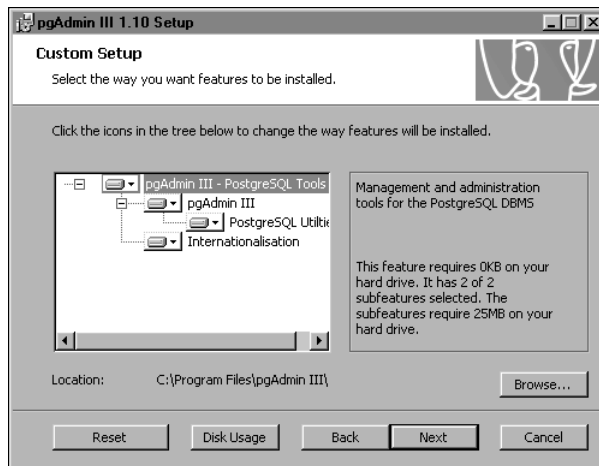


Figure 2-19 Viewing the features to be installed

- 4 In Ready to Install pgAdmin III, shown in Figure 2-20, choose Install.




Figure 2-20 Choosing to install pgAdmin II

Installing pgAdmin III appears, as shown in Figure 2-21.




Figure 2-21 Installing pgAdmin III

- 5 When Completed the pgAdmin III Setup Wizard appears, as shown in Figure 2-22, choose Finish to exit the wizard.




Figure 2-22 Choosing Finish

Choose Finish to exit the wizard, as shown in Figure 2-23.



Figure 2-23 Exiting the installation wizard

- 13** The installation program prompts you to install the online help from the following location, as shown in Figure 2-24:

<http://www.actuate.com>




Figure 2-24 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows, choose Start → Actuate 11 → Update Documentation.

The installation program installs shortcuts on the desktop, as shown in Figure 2-25.




Figure 2-25 Viewing BIRT iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to set up user accounts and run reports.
- BIRT iServer 11
Opens Welcome to Actuate BIRT iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

Accessing the PostgreSQL Database Server using the pgAdmin utility

After migrating the Encyclopedia volume to Release 11, you can optionally open the pgAdmin III utility and access the PostgreSQL Database Server to browse the Encyclopedia volume database. Actuate does not support modifying the BIRT iServer PostgreSQL Database schema. Any changes to the schema made by the customer, such as the addition of an index on a table, must be recreated again manually in any future upgrade.

How to access the PostgreSQL Database Server using the pgAdmin utility

To access the PostgreSQL server, choose Start → Programs → pgAdmin III 1.10 → pgAdmin III.

pgAdmin III appears, showing the PostgreSQL Database Server in the Object browser, as shown in Figure 2-26.

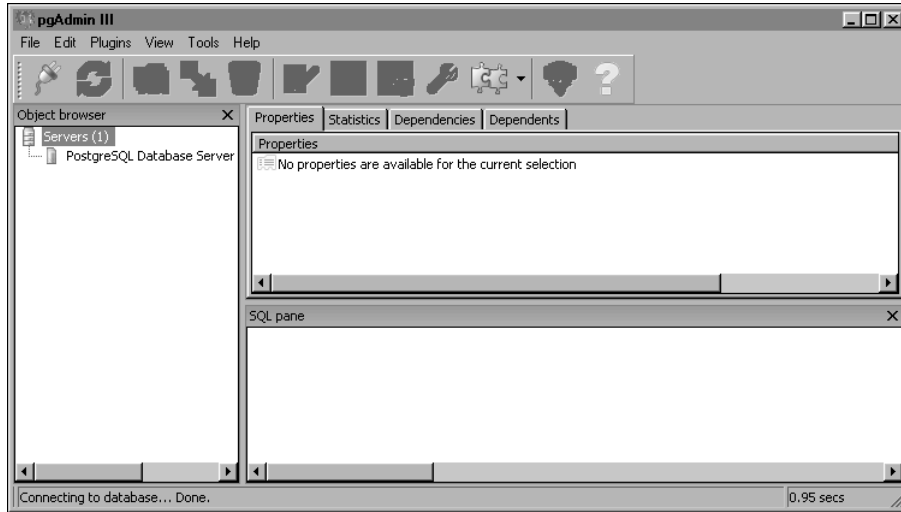


Figure 2-26 Viewing pgAdmin III

If the PostgreSQL Database Server does not appear in the Object browser, you can add the server manually by performing the following tasks:

- 1 In pgAdmin III, choose File→Add Server. In New Server Registration—Properties, type or select the following property values, as listed in Table 2-1.

Table 2-1 New Server Properties

Property	Value
Name	PostgreSQL Database Server
Host	<computer name>
Port	8432
Maintenance DB	postgres
Username	postgres
Password	<your superuser password>
Store password	Not selected
Restore env	Selected
Service	Actuate 11 PostgreSQL for BIRT iServer
Connect now	Selected

New Server Registration—Properties appears, as shown in Figure 2-27.

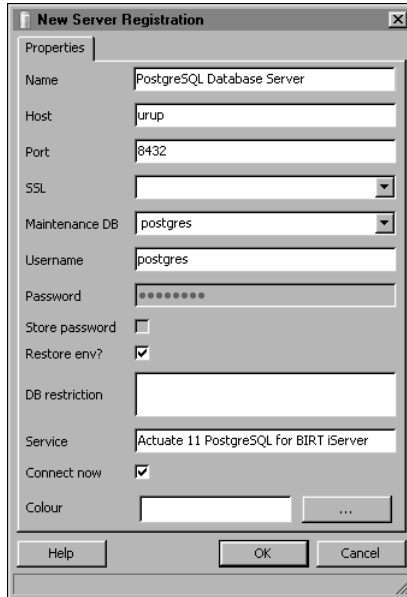


Figure 2-27 Registering a new server

Choose OK. Then, on pgAdmin III, expand PostgreSQL Database Server, as shown in Figure 2-28.

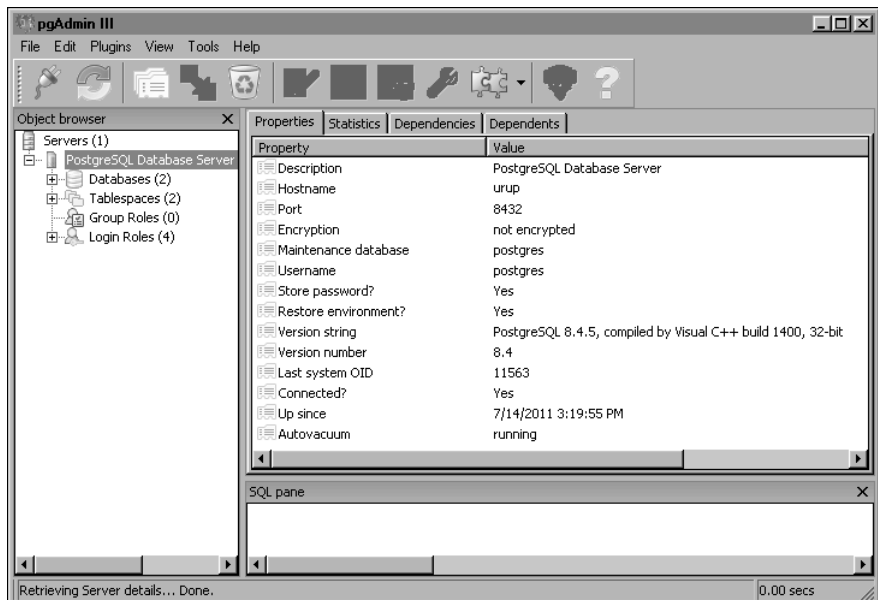


Figure 2-28 Viewing PostgreSQL Database Server properties

- In Object browser, expand Databases—iServer, iServer—Schemas, Tablespaces, and Login Roles to review the server and database installation, as shown in Figure 2-29.

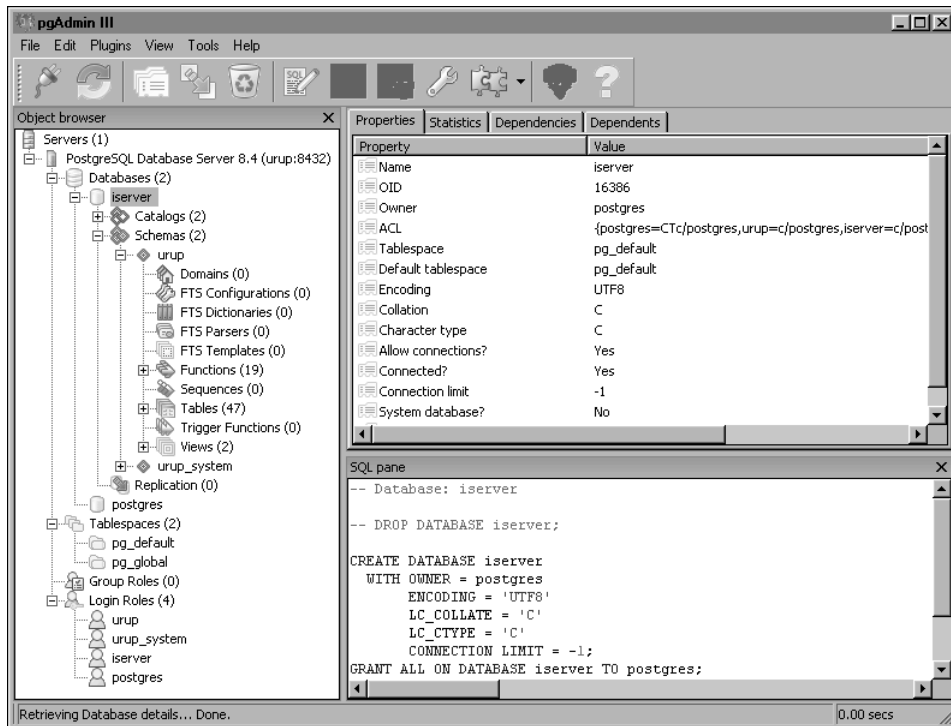


Figure 2-29 Browsing the PostgreSQL Database Server installation

The BIRT iServer uninstall process can optionally remove the iServer installation directory containing the encyc folder, which by default contains the postgres database directory. The uninstall process does not uninstall the pgAdmin III administration tool, a PostgreSQL database not residing in the AC_DATA_HOME/encyc directory, or another third-party database used to store the Encyclopedia volume data. You must uninstall these components separately.

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

A BIRT iServer Release 11 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:

`/Program Files/Common Files/Actuate/11.0/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 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 documents, iServer requires access to jar files in the Jar directory of the iServer installation files. In Windows, include the location of the jar file in the CLASSPATH.

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 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 volume schemas. Install upgraded design and document files.

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

- If you are replacing your Actuate 10 system, upgrade BIRT iServer system by installing Actuate 11 program files in the same directory structure as Actuate 10 and data files in a separate area of the file storage system.

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 document 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 an alternative 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 an account with Windows administrator privileges

Before installing iServer, create a Windows user account that is a member of the Administrators group. Use this account when installing and running iServer.

The iServer user account must meet the following requirements:

- Be a member of the Windows Administrators group.
The account must have privileges to access the required software and hardware, such as database servers, printers, and iServer files and folders.
- Have log on as a service privilege.
If the account does not meet this requirement, the iServer installation program prompts you to configure the privilege to run the Windows Actuate iServer service.

On a new Windows Vista installation, the initial user account is not a member of the Administrators group. You must configure this user account to be a member of this group.

Creating the iServer 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 a specific PostgreSQL installation. In the commands, substitute system and volume 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 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`, 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 the specific Oracle installation. In the commands, substitute system and 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 encodings 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 standalone 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, see 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 site. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

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

Actuate also supports the cloud deployment of BIRT iServer using a ready-to-launch iServer image. For more information about this installation option, see Chapter 6, “Installing BIRT iServer in a cloud,” 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.

The following section describes how to install BIRT iServer Release 11 using PostgreSQL as an alternative data store.

Installing an Encyclopedia volume that uses an alternative database

The following procedures use a pre-existing PostgreSQL database and schema as an example. During the iServer installation, the administrator provides the schema owner and database user credentials. 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

- 1 Download the iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 3-1. Choose Next.



Figure 3-1 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 3-2. Choose Next.

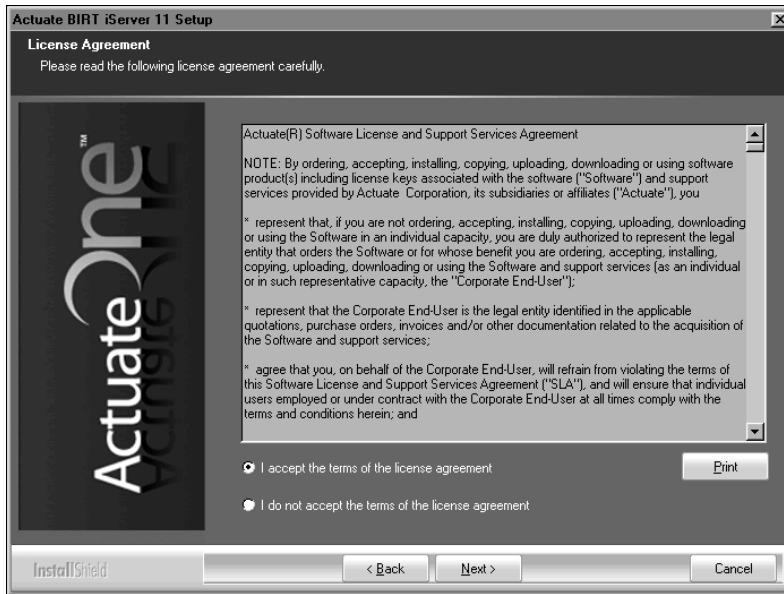


Figure 3-2 Accepting the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 3-3. Choose Next.

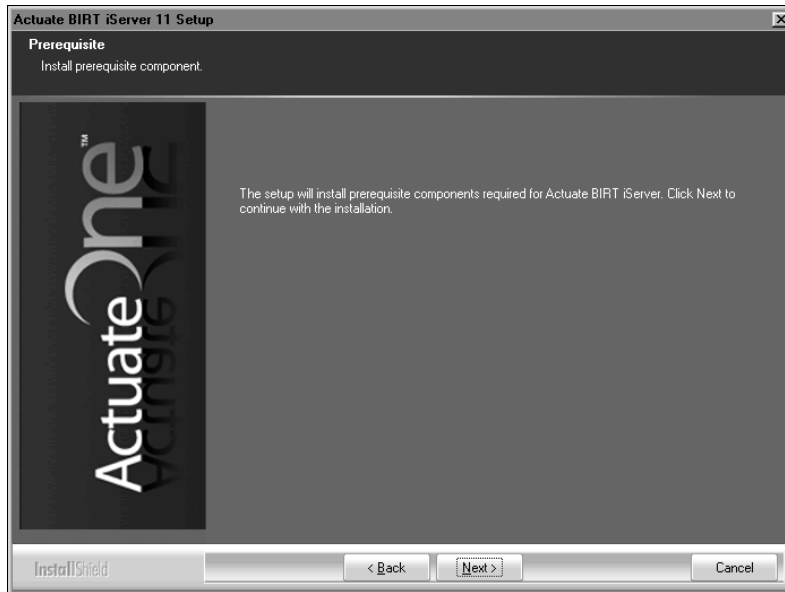


Figure 3-3 Installing prerequisites

- 4 In Setup Type, as shown in Figure 3-4, choose Typical to install a supported relational database management system (RDBMS) for the Encyclopedia volume, such as OOTB PostgreSQL or a pre-existing DB2, Microsoft SQL Server, Oracle or PostgreSQL RDBMS.

In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations.

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, `AC_SERVER_HOME`, points to the location of the iServer binaries. The default path for the program location is `C:/Program Files/Actuate11SP4/iServer`.

iServer uses the Data location to store the iServer logs, Encyclopedia volume data, and other related data. The environment variable, `AC_DATA_HOME`, points to the iServer data location. The default path is `C:/Actuate/iServer/data`.

Choose Next.

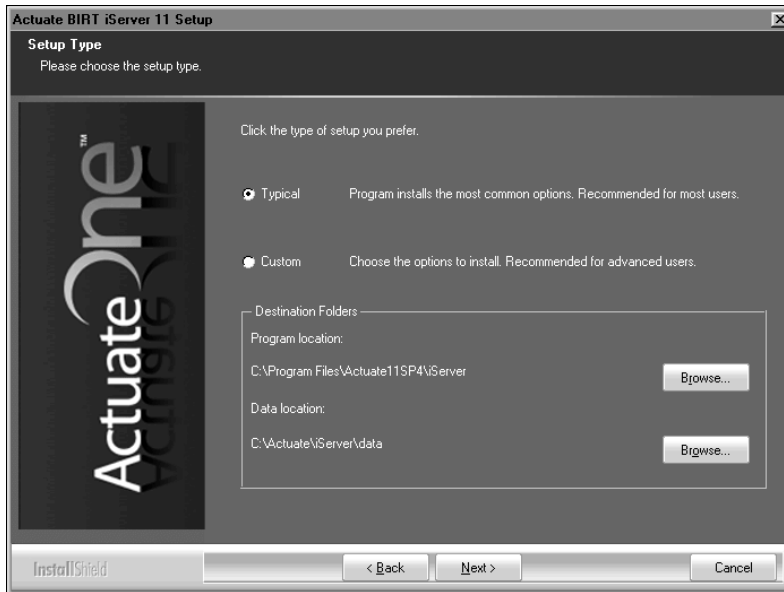


Figure 3-4 Specifying Typical setup type

- 5 In Encyclopedia Metadata Storage, select DB2, Oracle, PostgreSQL, or SQL Server to use as a pre-existing third-party database for storing Encyclopedia volume metadata. This example uses the PostgreSQL RDBMS, as shown in Figure 3-5. In System Name, type a name for the BIRT iServer system. Restrict the 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.

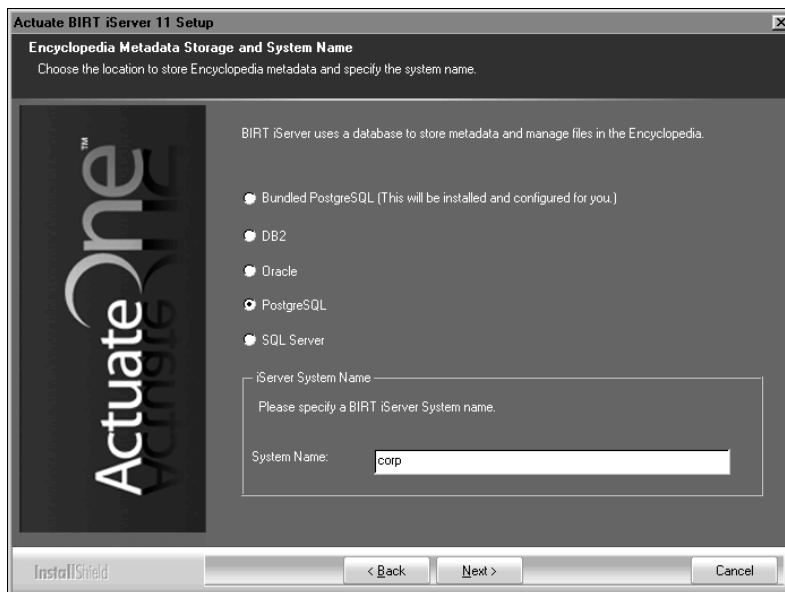


Figure 3-5 Choosing a third-party database

- 6 After you chose DB2, Oracle, PostgreSQL, or SQL Server in the previous step, 3rd Party Database Information appears. Depending on which alternative database you selected, provide the following database information. Then, choose Next:

- **DB2**

For a pre-existing DB2 installation, type the server hostname. In Database, specify the name of the database, such as iserver. Type the port. By default, the port is 50000. You can leave IANAAppCodePage blank. In Database User credentials, specify the iserver application user and a password, as shown in Figure 3-6.

- **Oracle**

For a pre-existing Oracle installation, type the server hostname and port. By default, the port is 1521.

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

In Database User Credentials, specify the iserver application user and a password, as shown in Figure 3-7.

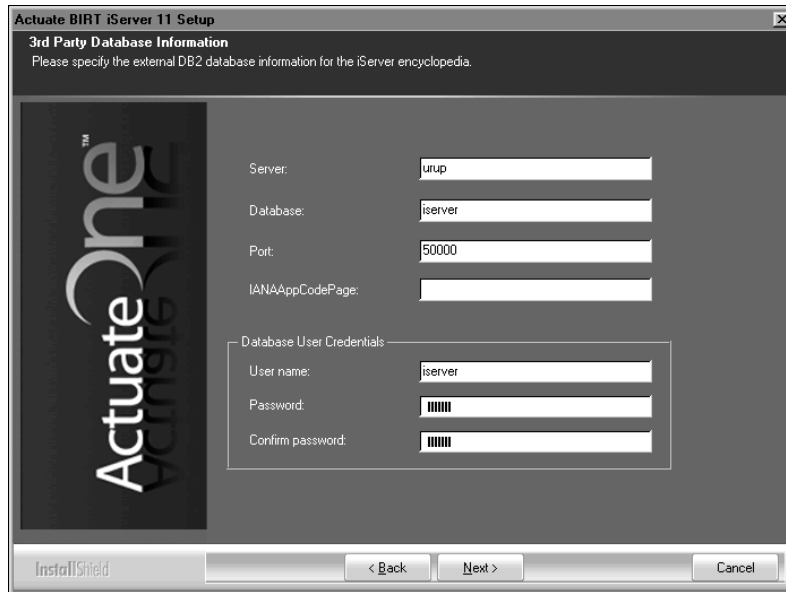


Figure 3-6 Specifying third-party database information for DB2

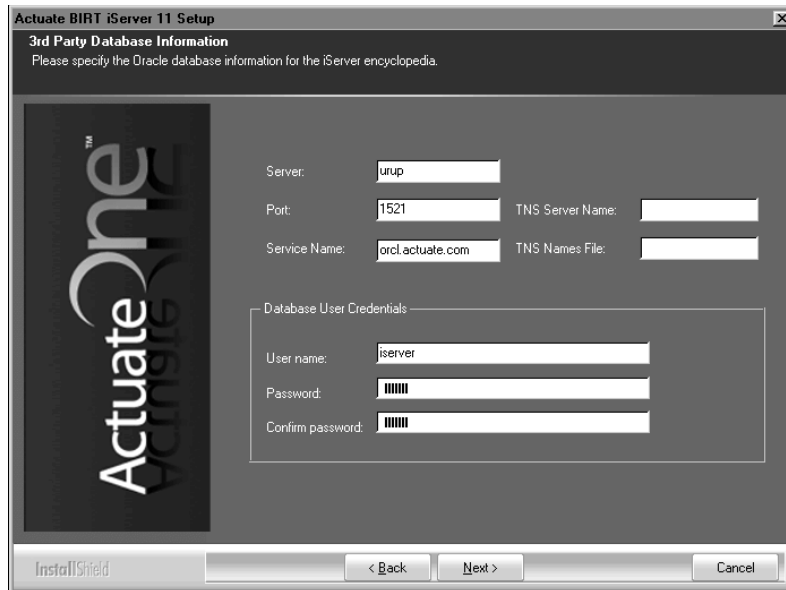


Figure 3-7 Specifying third-party database information for Oracle

- PostgreSQL

For a pre-existing PostgreSQL installation, type the server hostname and port. By default, the port is 5432. In Database, specify the name of the database, such as iserver. In Database User credentials, specify the iserver application user and a password, as shown in Figure 3-8.
- SQL Server

For a pre-existing SQL Server installation, type the server hostname and port. By default, the port is 1433. In Database, specify the name of the database, such as iserver. In Instance, type the SQL Server instance name. In Figure 3-9, Instance contains the SQL Server default instance name. Specify the iserver application user in Database User credentials, as shown in Figure 3-9.



Figure 3-8 Specifying third-party database information for PostgreSQL

- 7 In Database Schema Information, type the system schema owner, password, and confirm the password in System Database Schema Credentials. Type the Encyclopedia database schema owner, password, and confirm the password in Encyclopedia Database Schema Credentials, as shown in Figure 3-10.

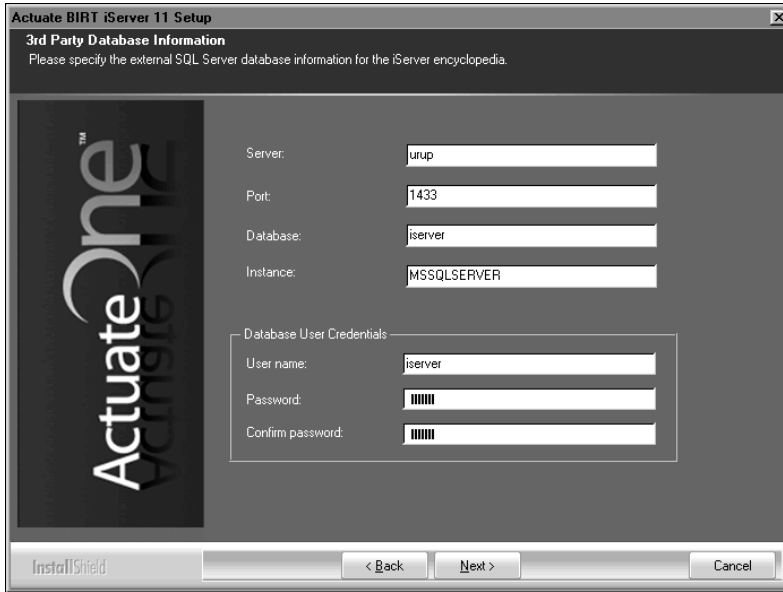


Figure 3-9 Specifying third-party database information for SQL Server

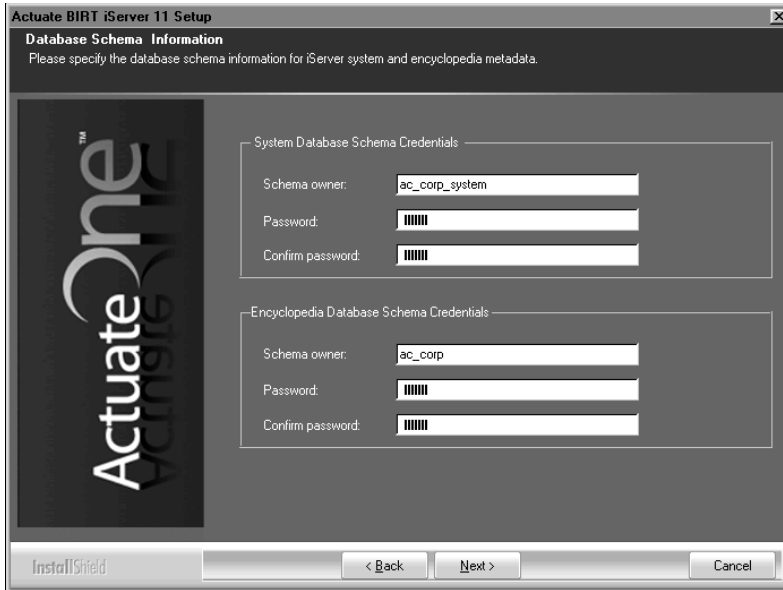


Figure 3-10 Specifying System and Encyclopedia schema passwords

- 8 In License File Details, select Use the license that you purchased. Choose Browse then navigate to and choose the license file, as shown in Figure 3-11. Choose Next.

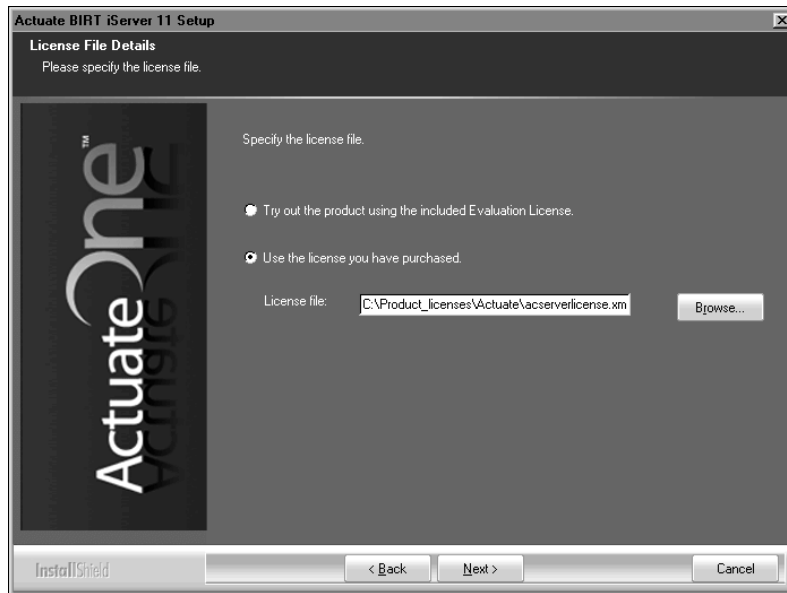


Figure 3-11 Specifying the license file

When installing using a named user license, a prompt appears advising you to check that the volume does not exceed the number of registered users authorized by the license, as shown in Figure 3-12.




Figure 3-12 Viewing the named user license question

Choose Yes to continue the installation.

- 9 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 3-13. Alternatively, choose the language and locale settings for your region.
- 10 In Specify Profiles, type the user name, password, and confirm the password for the account used to start the Actuate BIRT iServer 11 service, as shown in Figure 3-14. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

Accept Automatically start the BIRT iServer 11 service when Windows boots, as shown in Figure 3-14. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

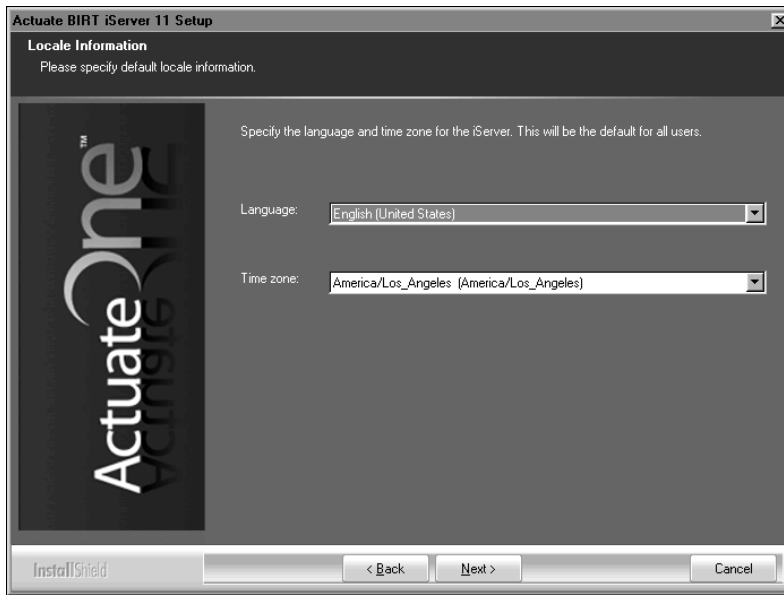


Figure 3-13 Specifying locale information

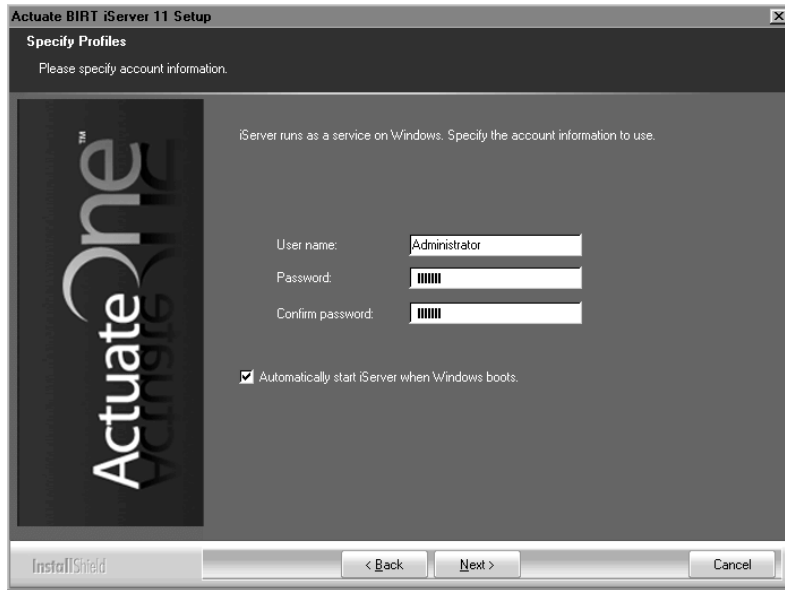


Figure 3-14 Specifying the account for running the iServer service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 3-15.




Figure 3-15 Setting the Windows local security policy

- 11 In System Configuration Password, type and confirm a password for Configuration Console, as shown in Figure 3-16. For both Configuration Console and Management Console, the default user name is Administrator. The Administrator account for Management Console has no password initially. You can log in to these consoles and change the password settings after installing iServer. Choose Next.

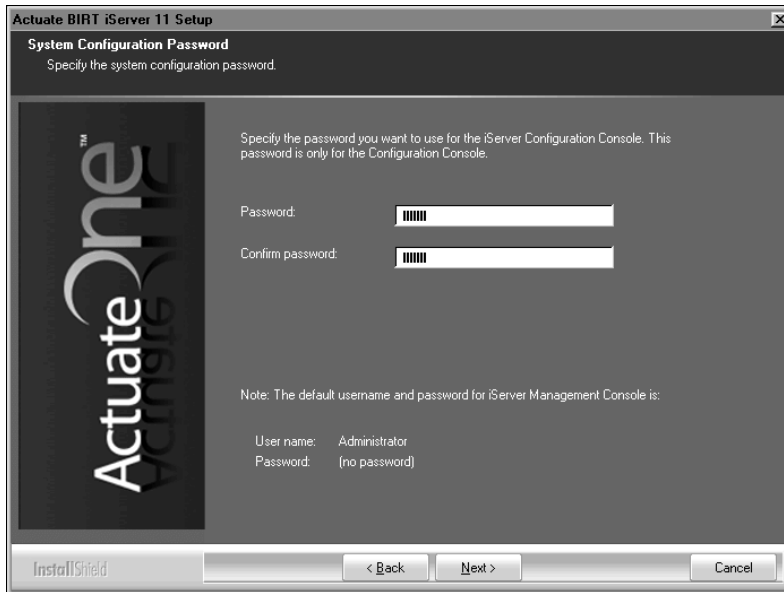


Figure 3-16 Specifying the password for using Configuration Console

- 12** In Start Copying Files, review the settings shown in Figure 3-17. Choose Next. Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 3-18.

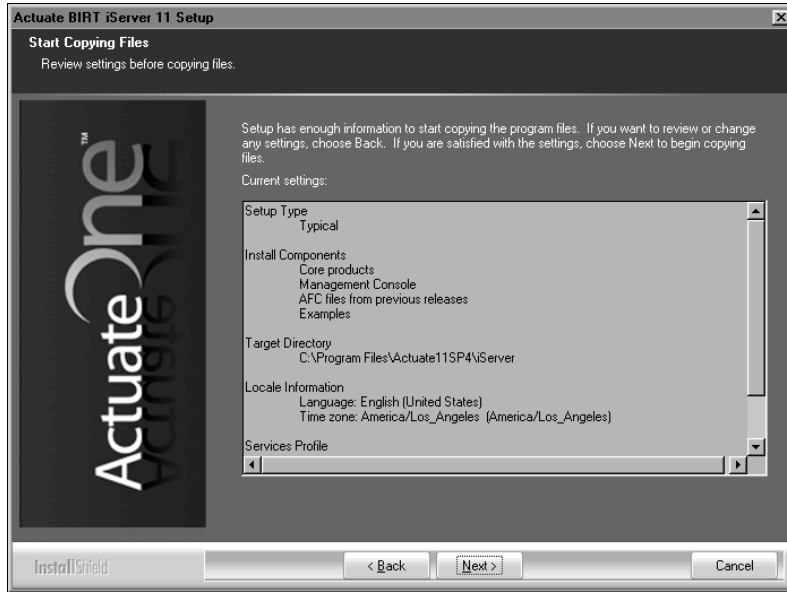


Figure 3-17 Reviewing settings before copying files



Figure 3-18 Viewing setup status

13 Choose Finish to exit the wizard, as shown in Figure 3-19.



Figure 3-19 Exiting the installation wizard

14 The installation program prompts you to install the online help from the following location, as shown in Figure 3-20:

<http://www.actuate.com>




Figure 3-20 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows, choose Start-> Actuate 11->Update Documentation.

The installation program installs shortcuts on the desktop, as shown in Figure 3-21.




Figure 3-21 Viewing iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to set up user accounts and run reports.
- iServer 11
Opens Welcome to Actuate iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

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.

Actuate recommends the side-by-side, manual upgrade procedure as a best practice since this process does not install in the same location as the previous installation. For more information about manually upgrading in a side-by-side installation scenario, see “Performing a manual side-by-side migration,” later in this chapter.

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 administrator privileges for installing and running iServer
- Back up iServer system and Encyclopedia volume metadata

Creating an account with Windows administrator privileges

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.

Before installing iServer, create a Windows user account that is a member of the Administrators group. Use this account when installing and running iServer.

The iServer user account must meet the following requirements:

- Be a member of the Windows Administrators group.
The account must have privileges to access the required software and hardware, such as database servers, printers, and iServer files and folders.
- Have log on as a service privilege.
If the account does not meet this requirement, the iServer installation program prompts you to configure the privilege to run the Windows Actuate iServer service.

On a new Windows Vista installation, the initial user account is not a member of the Administrators group. You must configure this user account to be a member of this group.

When installing iServer in Windows 2003, create a Windows user account that is a member of the Power Users not the Administrators group. Make sure that the Account has permission to access any printers required for printing. Perform the installation using an account that has Administrator privileges. During the installation, when prompted to specify the user account to run the iServer service, specify the Power User account.

For more information about configuring a Power User and iServer account and log on as a service privilege, see “Creating an account with Windows administrator privileges,” in Chapter 2, “Installing BIRT iServer.”

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 site. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

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

Actuate also supports the cloud deployment of BIRT iServer using a ready-to-launch iServer image. For more information about this installation option, see Chapter 6, “Installing BIRT iServer in a cloud,” 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.

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 for 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 upgrade also performs the following operations during installation:

- Creates the system and volume schemas, initializing these schemas 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.

Stopping and starting Actuate 10 BIRT iServer and NobleNet Portmapper services

This installation program can encounter a problem overwriting a file linked with a running process. Before running the installation program, shut down these services in the following order before running the in-place upgrade for the Release 10 version:

- Actuate 10 BIRT iServer service
- NobleNet Portmapper service

If you do not shut down the Actuate 10 BIRT iServer service, the installation program prompts you to do so.

How to stop Actuate iServer 10 service

1 Choose Start→Settings→Control Panel.

On Control Panel, choose Administrative Tools.

On Administrative Tools, choose Services.

- 2 On Services, select Actuate iServer 10 service, as shown in Figure 4-1. Then, choose Stop the service.

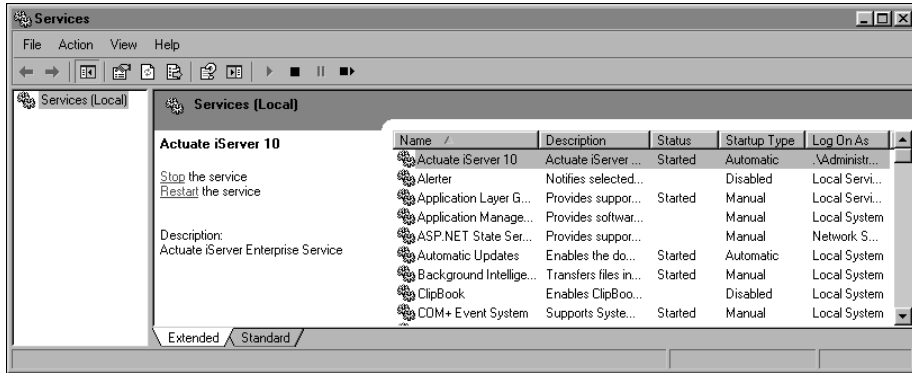


Figure 4-1 Stopping the Actuate iServer 10 service
The service stops, as shown in Figure 4-2.

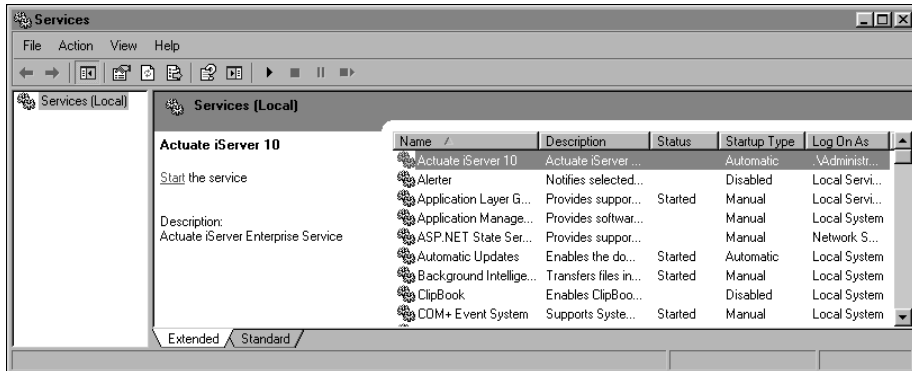


Figure 4-2 Viewing Services after BIRT iServer service stops

How to shut down NobleNet Portmapper

Stop NobleNet Portmapper by performing the following tasks:

- 1 Choose Start>Programs>Administrative Tools>Services.
- 2 In Services, select NobleNet Portmapper for TCP then choose Stop.

Alternatively, open a command prompt and run the following command:

```
net stop "NobleNet Portmapper for TCP"
```

Uninstalling iServer removes the NobleNet Portmapper service. If a co-existing iServer is present in the environment, the administrator can manually start the NobleNet Portmapper service for an installation using a custom batch file.

The batch file starts the NobleNet Portmapper for TCP service for an iServer installation by executing `portserv.exe`, located in `AC_SERVER_HOME/bin`. In a default iServer Release 11 Service Pack 4 installation, the path for `AC_SERVER_HOME` is:

```
C:\Program Files\Actuate11SP4\iServer
```

How to manually start a NobleNet Portmapper service

Create a custom batch file and manually start the NobleNet Portmapper service by performing the following tasks:

- 1 In Windows Explorer, navigate to `AC_SERVER_HOME/bin`.
- 2 In `AC_SERVER_HOME/bin`, using a text editor, create a file containing the commands shown in Listing 4-1.
- 3 Save the file as a `.bat` file, for example, `SwitchPortmapperService.bat`.
- 4 Open a command prompt and navigate to `AC_SERVER_HOME/bin`.
- 5 Execute `SwitchPortmapperService.bat`. In Windows 7, Vista, or other operating system version, you may need to run as administrator to start the service.

Listing 4-1 contains an example of a batch file that manually starts the NobleNet Portmapper service in an iServer Release 11 Service Pack 4 installation.

Listing 4-1 Creating a `.bat` file to point iServer to `portserv.exe`

```
REM Edit AC_SERVER to point to the iServer program files location
SET AC_SERVER=C:\Program Files\Actuate11SP4\iServer
REM Run Service Control(SC) Manager to start portmapper service
sc Stop "NobleNet Portmapper for TCP"
sc Config "NobleNet Portmapper for TCP" binpath= "%AC_SERVER%\bin
\portserv.exe tcp" start= auto
sc Start "NobleNet Portmapper for TCP"
```

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 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
 - `acserversconfig.xml` in the `/etc` directory

- acpmdconfig.xml in the /etc directory
 - RSSE code and associated files if you use the Open Security option
- 2 Download the iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 4-3. Choose Next.



Figure 4-3 Viewing the welcome message

- 3 Read and accept the license agreement, as shown in Figure 4-4. Choose Next.

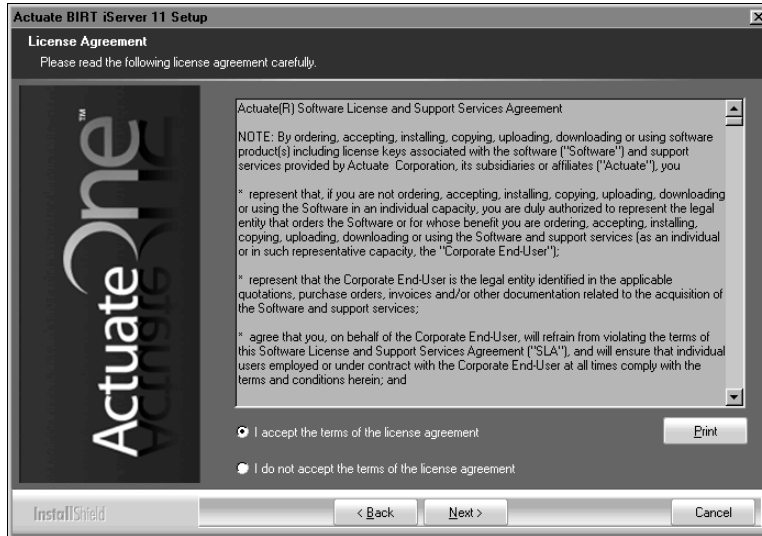


Figure 4-4 Viewing the license agreement

- 4 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-5. Choose Next.

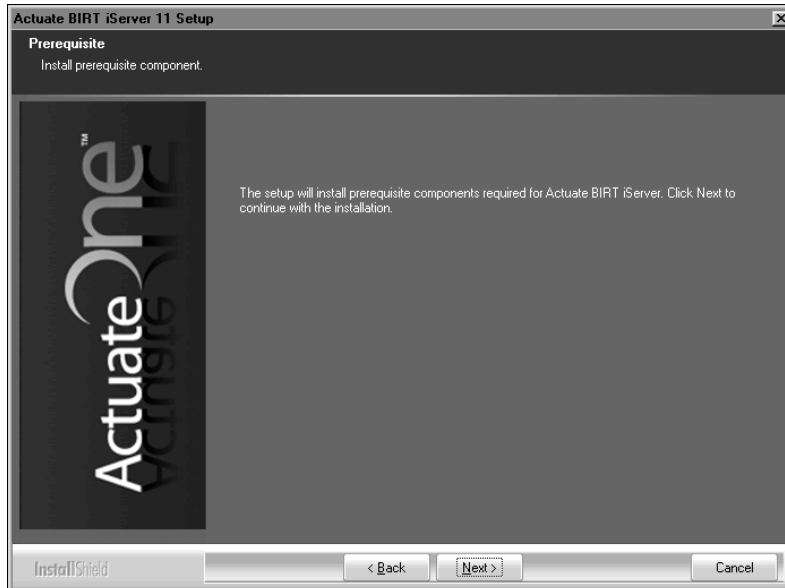


Figure 4-5 Installing Prerequisites

- 5 In Setup Type, perform the following tasks:

- 1 Select Typical setup type.
- 2 In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations.

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, `AC_SERVER_HOME`, points to the location of the iServer binaries. The default path for the program location is `C:/Program Files/Actuate11SP4/iServer`.

iServer uses the Data location to store the iServer logs, Encyclopedia volume files, and run-time data, including PostgreSQL data. The environment variable, `AC_DATA_HOME`, points to the location of the iServer data. When upgrading from an earlier major version of iServer, such as Release 10, the default path for the data location is `C:/Actuate/iServer/data`.

Select Upgrade Existing iServer and Migrate the Encyclopedia Data. The installer lists the Existing iServer Location, as shown in Figure 4-6.

Check that the Existing iServer Location matches the location of the old files that you are upgrading. Choose Next.

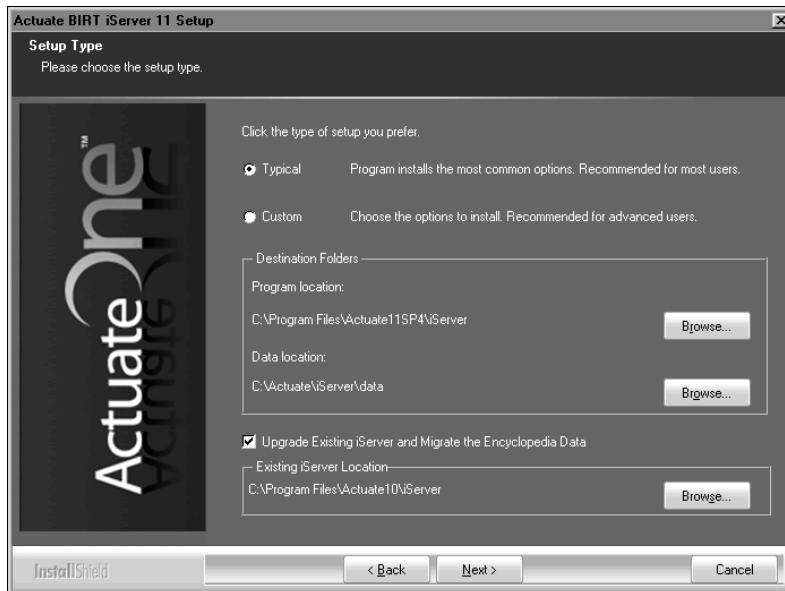


Figure 4-6 Choosing to upgrade iServer and migrate the Encyclopedia

- 6 A message appears stating that the install program is migrating a previous version of the Encyclopedia volume, as shown in Figure 4-7. Choose OK.



Figure 4-7 Confirming volume migration

- 7 If you have not shut down the iServer service, a prompt instructs you to shut down the service, as shown in Figure 4-8. Perform the following tasks:
 - 1 Choose OK on the prompt shown in Figure 4-8.
 - 2 Shut down the Actuate iServer service.
 - 3 On Setup Type, choose Next.




Figure 4-8 Choosing OK to shut down the iServer service

- 8 On Encyclopedia Metadata Storage, accept the default database, Bundled PostgreSQL. In System Name, type a name for the BIRT iServer System name, as shown in Figure 4-9. iServer inserts this name into the Encyclopedia volume schema and the iServer system schema names. Choose Next.

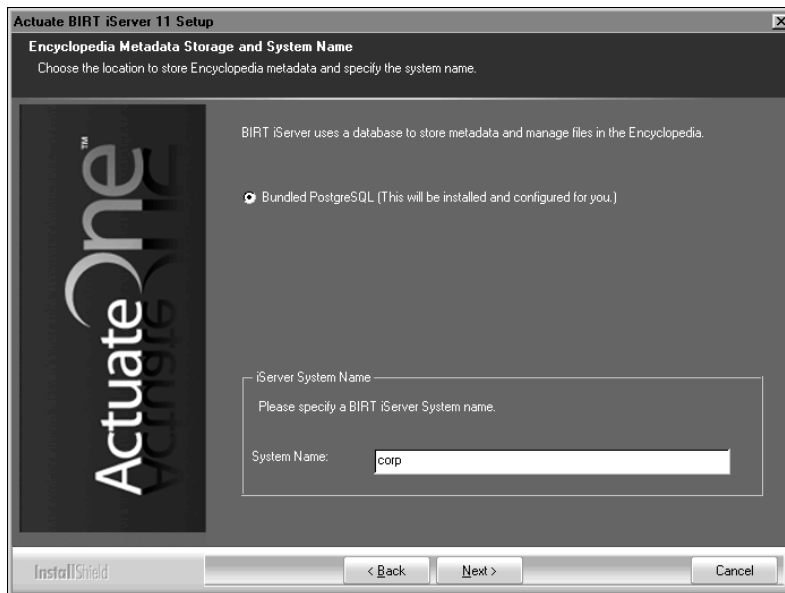


Figure 4-9 Accepting metadata database type

9 On PostgreSQL Database Information, in Database Credentials, type and confirm a password, as shown in Figure 4-10. iServer creates the following accounts automatically, using this password for each account:

- postgres
The PostgreSQL database superuser. The database superuser administers the PostgreSQL relational database management system (RDBMS).
- ac_<BIRT iServer System name>_system
The System schema owner. iServer creates the iServer system schema and gives it this name.
- ac_<BIRT iServer System name>
The Encyclopedia schema owner. iServer creates the Encyclopedia volume schema and gives it this name.

In PostgreSQL Database Connection Information, accept the default port 8432 or type a new port number. Choose Next.

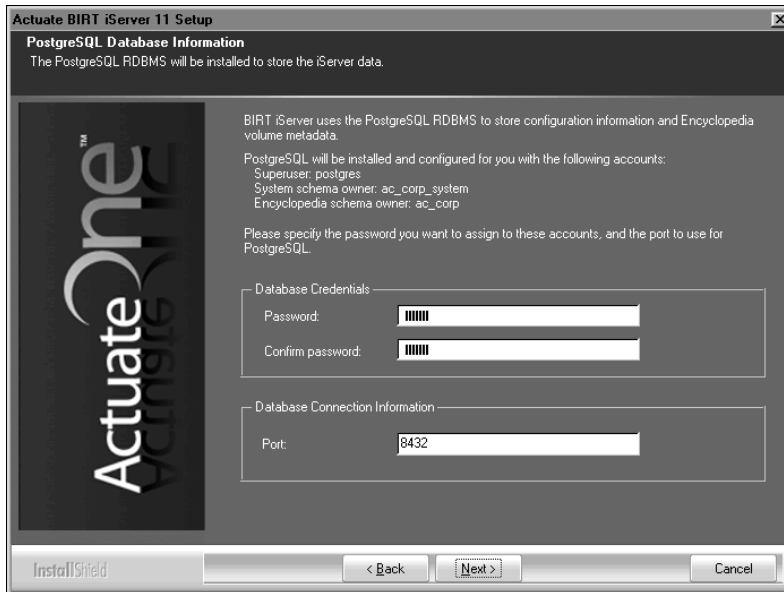


Figure 4-10 Specifying PostgreSQL database information

10 Choose Yes to update your license, as shown in Figure 4-11.




Figure 4-11 Updating licensing

- 11 In License File Details, browse to and choose the license file, as shown in Figure 4-12. Choose Next.

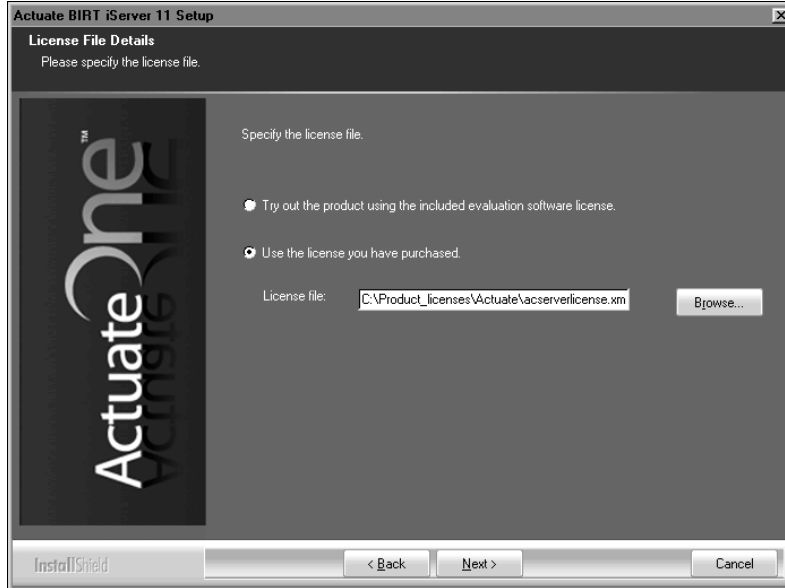


Figure 4-12 Specifying the license file

If installing using a named user license, a prompt appears advising you to check that the volume does not exceed the number of registered users authorized by the license, as shown in Figure 4-13. Choose Yes to continue, then choose Next.




Figure 4-13 Viewing the named user license question

- 12** In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 4-14. Alternatively, choose the language and locale settings for your region.

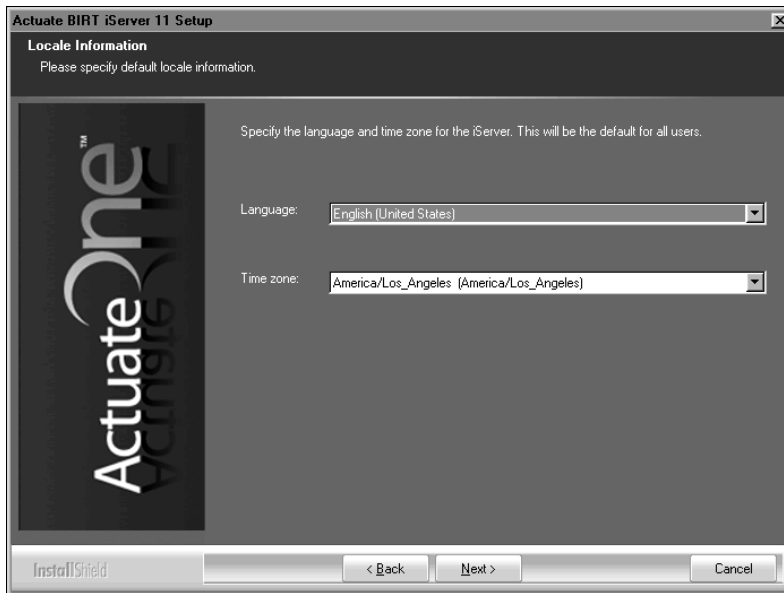


Figure 4-14 Specifying locale information

- 13** In Specify Profiles, type the user name and password for the account used to start the Actuate BIRT iServer 11 service. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

If you are installing on Windows 2003, specify a user account that is in the Power Users not the Administrators group. A user account in the Administrators group cannot start the Actuate 11 BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 4-15. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

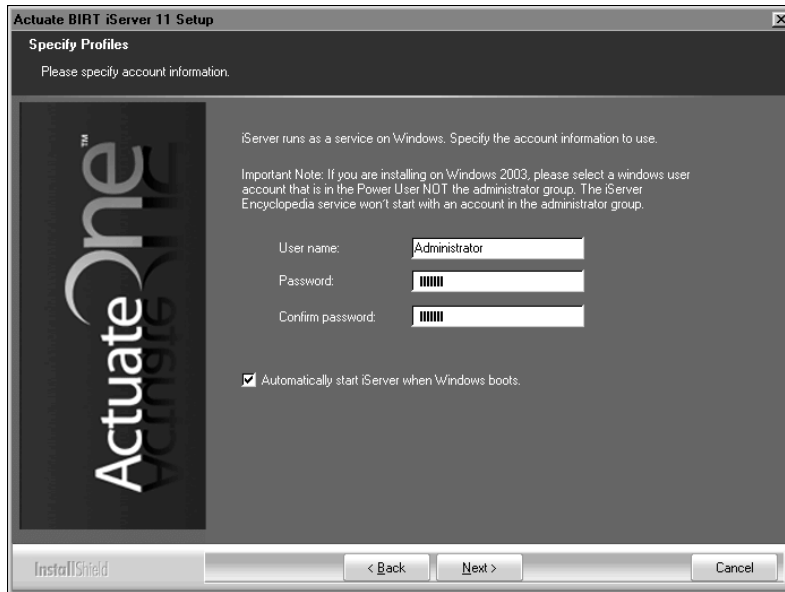


Figure 4-15 Specifying the account for running the iServer service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 4-16.




Figure 4-16 Setting the Windows local security policy

- 14 In System Configuration Password, type a password for Configuration Console, as shown in Figure 4-17. For both Configuration Console and Management Console, the default user name is Administrator. The Administrator account for Management Console has no password initially. You can log in to these consoles and change the password settings after installing iServer. Choose Next.

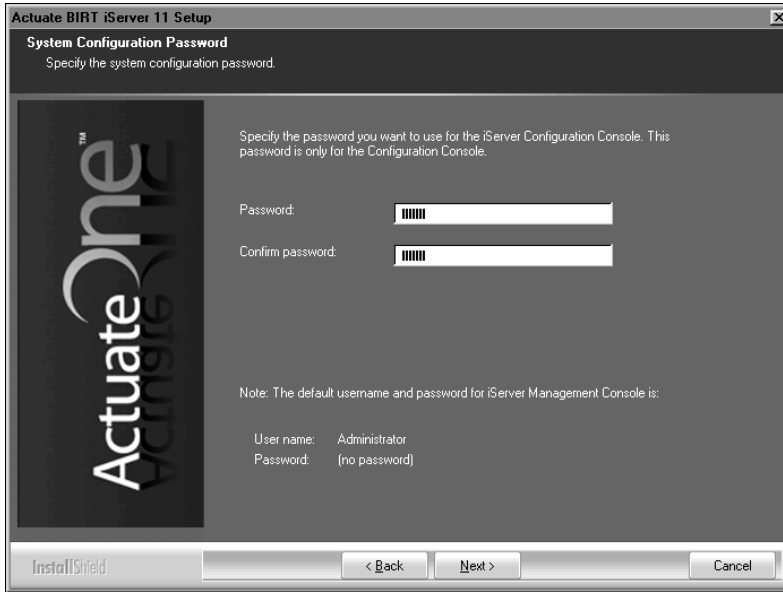


Figure 4-17 Specifying the password for using Configuration Console

15 In Start Copying Files, review the settings shown in Figure 4-18. Choose Next.

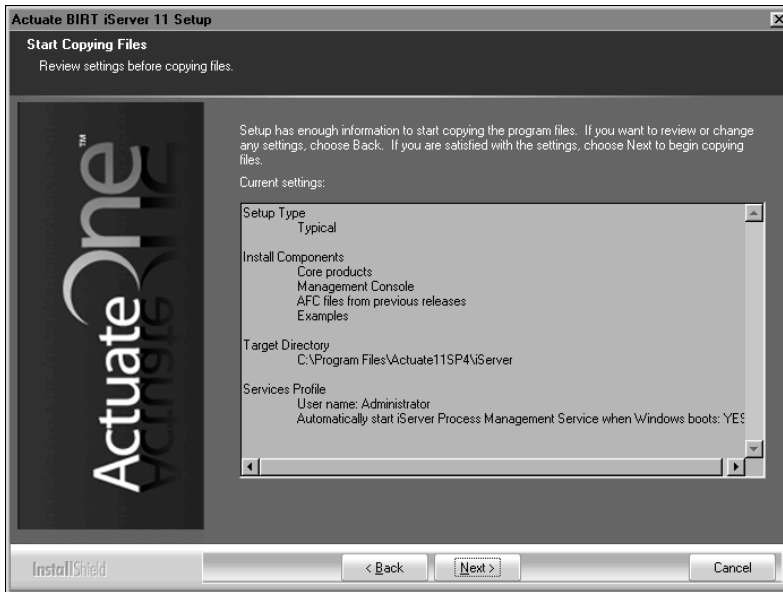


Figure 4-18 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-19.

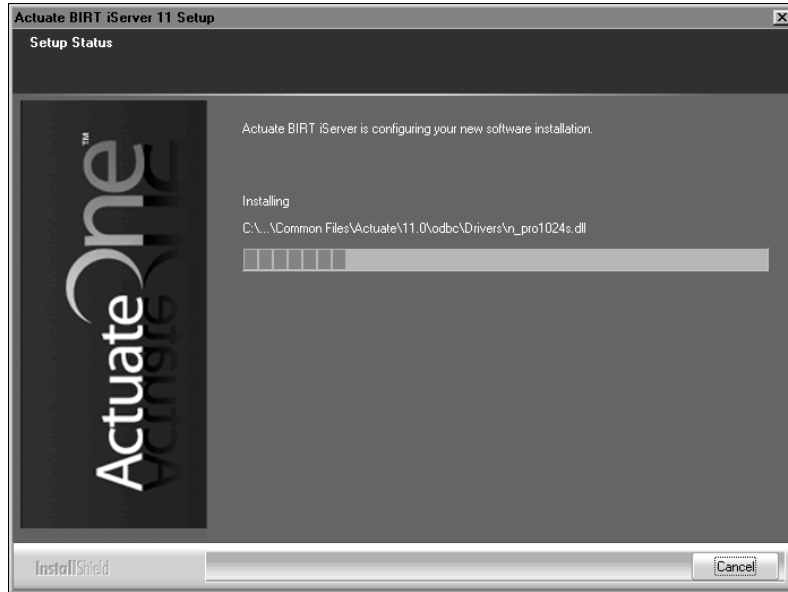


Figure 4-19 Viewing setup status

During installation, a command prompt appears, displaying messages regarding installation run-time processing, as shown in Figure 4-20.




Figure 4-20 Viewing run-time processing

A prompt appears, asking whether to install pgAdmin, the administration tool for the PostgreSQL RDBMS, as shown in Figure 4-21. If you do not have pgAdmin installed, choose Yes.




Figure 4-21 Choosing to install the pgAdmin tool

Actuate BIRT iServer Setup appears with pgAdmin III Setup appearing immediately afterward, as shown in Figure 4-22.

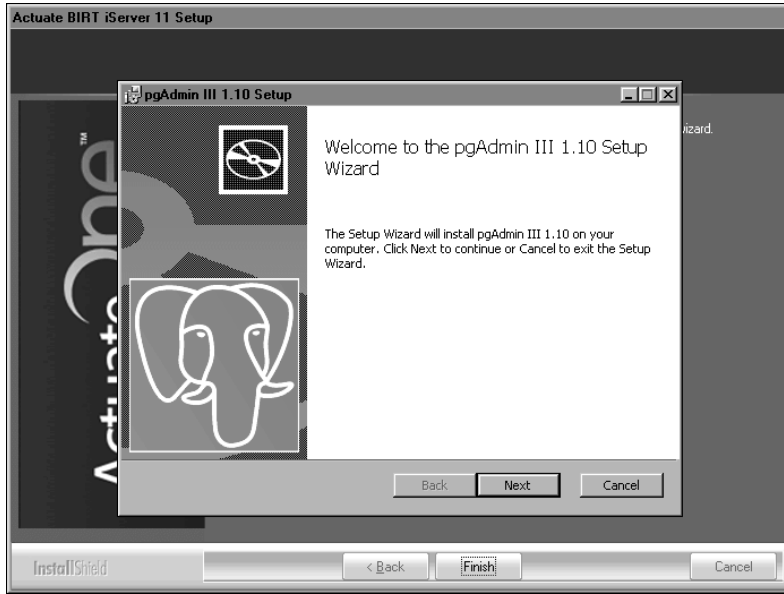


Figure 4-22 Viewing Actuate BIRT iServer and pgAdmin III Setup

16 In pgAdmin III Setup, perform the following tasks:

- 1 In Welcome, shown in Figure 4-23, choose Next.

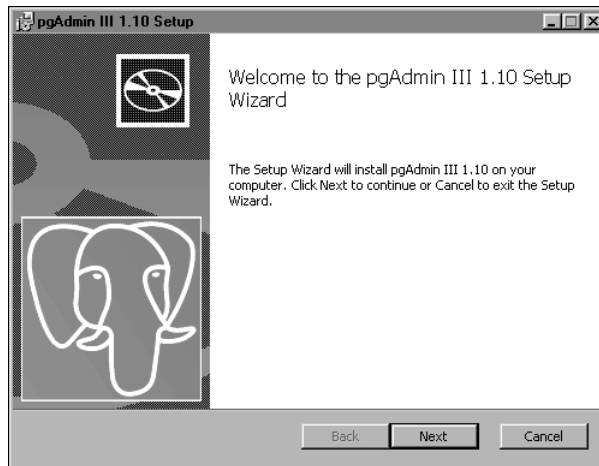


Figure 4-23 Viewing Welcome

- In License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-24. Choose Next.



Figure 4-24 Accepting the license agreement

- In Custom Setup, review the features to be installed, as shown in Figure 4-25. Choose Next.

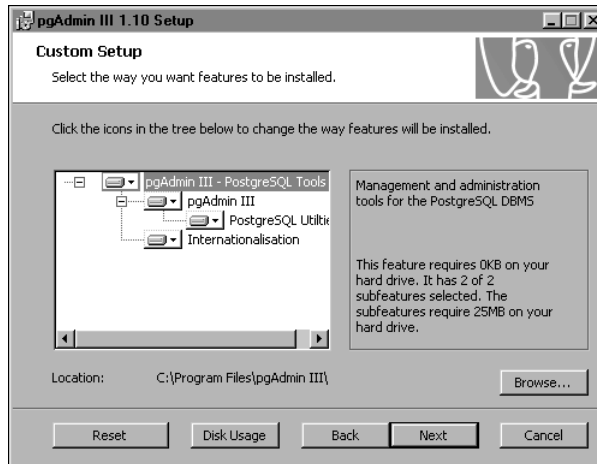


Figure 4-25 Viewing the features to be installed

- In Ready to Install pgAdmin III, shown in Figure 4-26, choose Install.

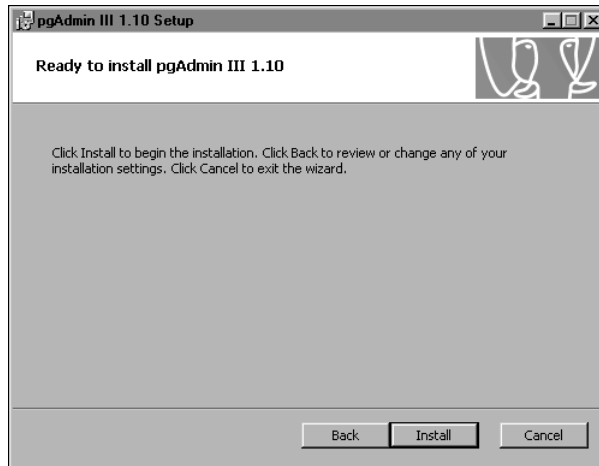


Figure 4-26 Choosing to install pgAdmin II
Installing pgAdmin III appears, as shown in Figure 4-27.




Figure 4-27 Installing pgAdmin III

- 5 When Completed, the pgAdmin III Setup Wizard appears, as shown in Figure 4-28, choose Finish to exit the pgAdmin III Setup Wizard.




Figure 4-28 Choosing Finish

- 17 When the Actuate BIRT iServer 11 setup success message appears, as shown in Figure 4-29, choose Finish to exit this wizard.



Figure 4-29 Exiting the installation wizard

- 18 The installation program prompts you to install the online help from the following location, as shown in Figure 4-30:

<http://www.actuate.com>




Figure 4-30 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows choose Start→ Actuate 11→Update Documentation.

The installation program installs shortcuts on the desktop, as shown in Figure 4-31.



Figure 4-31 iServer shortcuts installed on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to set up user accounts and run reports.
- iServer 11
Opens Welcome to Actuate iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

Choose Start→Programs→Actuate 11→BIRT iServer Configuration Console to open Configuration Console. Configuration Console Log In appears, as shown in Figure 4-32. iServer is ready for use.

Figure 4-32 Configuration Console Log In appears

About changes to the file structure

Figure 4-33 shows the results of the changes to the upgraded installation file structure:

- The location of `acsverconfig.xml` is `AC_DATA_HOME/config/11SP4` instead of `AC_DATA_HOME/config`.

- The upgrade to Release 11 Service Pack 4 requires installing iServer program files into a new location. In this installation, the location is /Program Files /Actuate11SP4/iServer, which becomes the new AC_SERVER_HOME. The pre-release program files remain intact in case the setup needs to be restored.
- The common files location remains /Program Files/Common Files/Actuate \11.0.

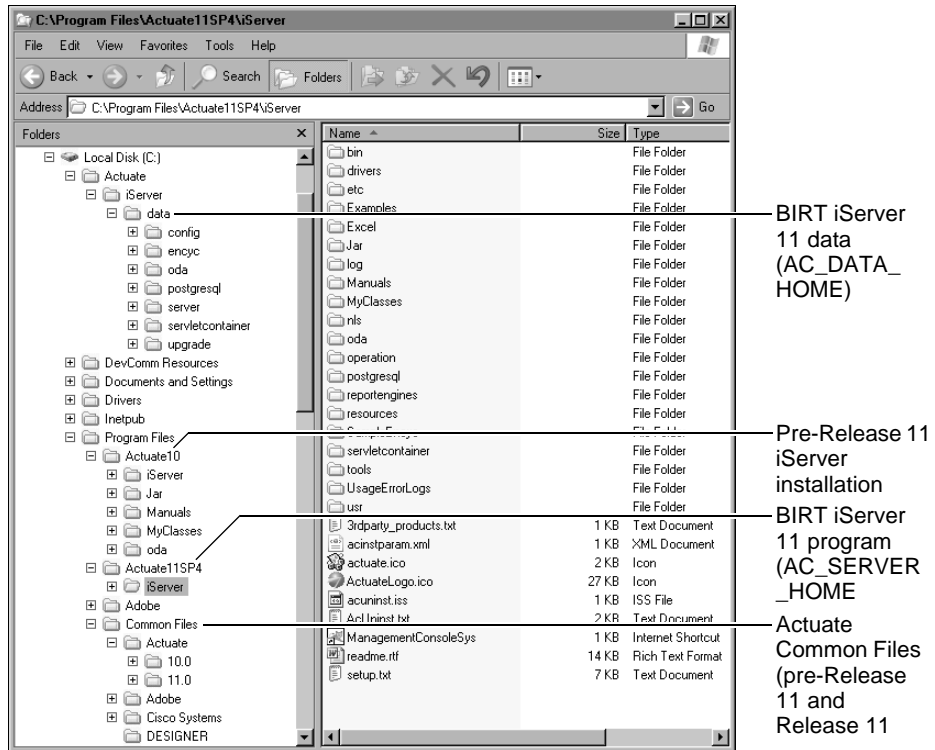


Figure 4-33 Viewing the file system after upgrade to Release 11SP4

Accessing the PostgreSQL Database Server using the pgAdmin utility

After migrating the Encyclopedia volume to Release 11, you can optionally open the pgAdmin III utility and access the PostgreSQL Database Server to browse the Encyclopedia volume database. Actuate does not support modifying the BIRT iServer PostgreSQL Database schema. Any changes to the schema made by the customer, such as the addition of an index on a table, must be recreated again manually in any future upgrade. For more information about using the pgAdmin III utility to access the PostgreSQL Database Server to browse the Encyclopedia volume database, see “Accessing the PostgreSQL Database Server using the

pgAdmin utility,” in Chapter 2, “Installing BIRT iServer.”

The BIRT iServer uninstall process can optionally remove the iServer installation directory containing the encyc folder, which by default contains the postgres database directory. The uninstall process does not uninstall the pgAdmin III administration tool, a PostgreSQL database not residing in the AC_DATA_HOME/encyc directory, or another third-party database used to store the Encyclopedia volume data. You must uninstall these components separately.

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 for any supported RDBMS after performing the automatic, wizard-based, iServer system upgrade.

The following sections describe how to perform an automatic in-place upgrade of an earlier minor BIRT iServer release, such as Release 11 Service Pack 3, to Release 11 Service Pack 4.

Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services

This installation program can encounter a problem overwriting a file linked with a running process. You must shut down the following services in order, before performing the in-place upgrade from an earlier Release 11 version:

- Actuate 11 BIRT iServer service
- Actuate 11 PostgreSQL for BIRT service, if the installation uses the PostgreSQL RDBMS
- NobleNet Portmapper service

How to stop the Actuate 11 BIRT iServer service

- 1 Choose Start→Settings→Control Panel.
On Control Panel, choose Administrative Tools.
On Administrative Tools, choose Services.
- 2 On Services, select Actuate 11 BIRT iServer Service, as shown in Figure 4-34. Then, choose Stop the service.

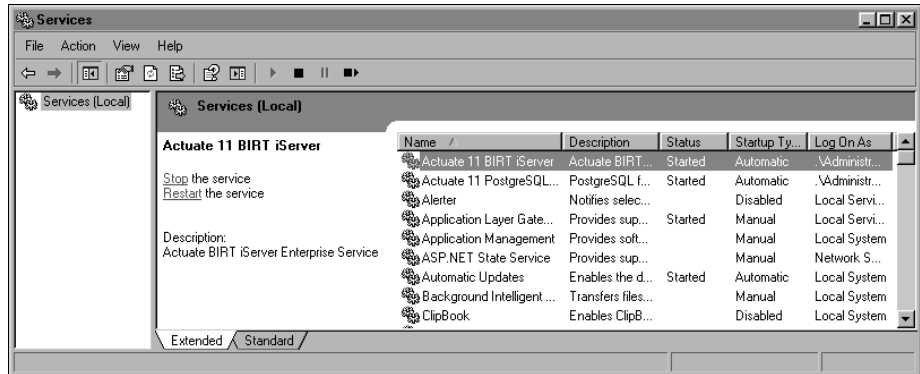


Figure 4-34 Stopping the BIRT iServer service

The service stops, as shown in Figure 4-35.

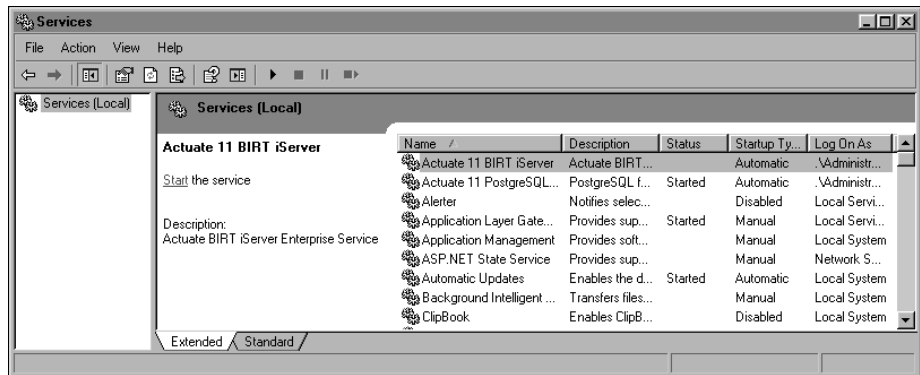


Figure 4-35 Viewing Services after BIRT iServer service stops

How to shut down the Actuate 11 PostgreSQL for BIRT service

If the installation uses the PostgreSQL RDBMS, you must shut down the Actuate 11 PostgreSQL for BIRT service by performing the following tasks:

- 1 Choose Start→Programs→Administrative Tools→Services.
On Control Panel, choose Administrative Tools.

On Administrative Tools, choose Services.

- 2 On Services, select Actuate 11 PostgreSQL for BIRT Service, as shown in Figure 4-36.

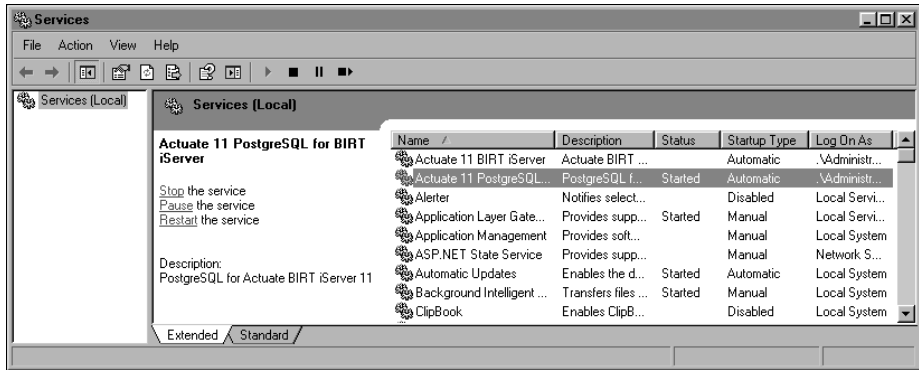


Figure 4-36 Shutting down the PostgreSQL for BIRT iServer service

The service stops, as shown in Figure 4-37.

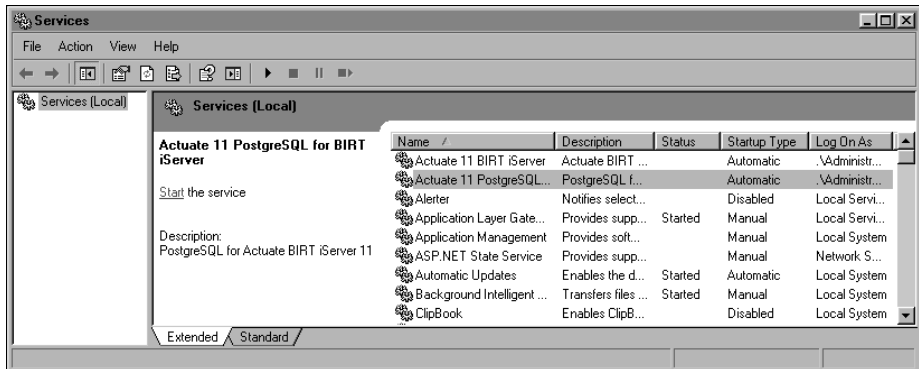


Figure 4-37 Viewing Services after PostgreSQL service stopped

How to shut down the NobleNet Portmapper

Stop NobleNet Portmapper by performing the following tasks:

- 1 Choose Start>Programs>Administrative Tools>Services.
- 2 In Services, select NobleNet Portmapper for TCP then choose Stop.

Alternatively, open a command prompt and run the following command:

```
net stop "NobleNet Portmapper for TCP"
```

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 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 AC_DATA_HOME/config directory
 - acpmdconfig.xml in the AC_SERVER_HOME/etc directory
 - RSSE code and associated files if you use the Open Security option
- 2 Download the iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 4-38. Choose Next.

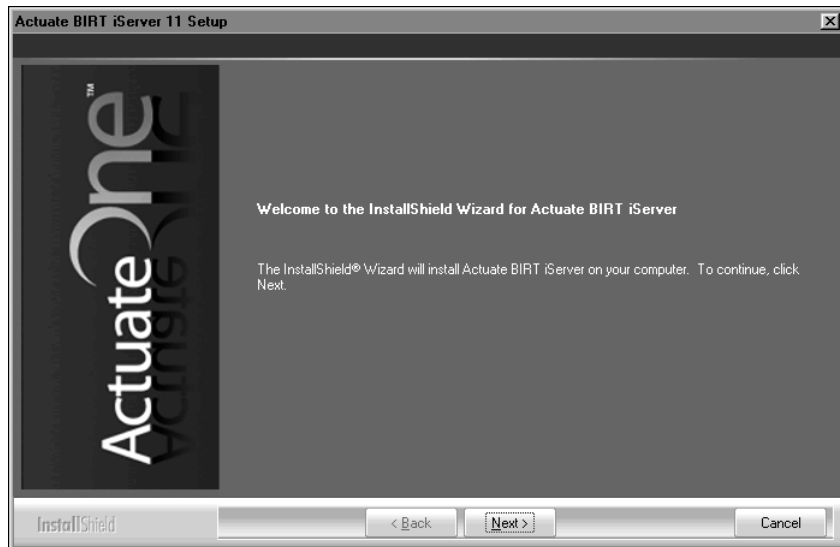


Figure 4-38 Viewing the welcome message

- 3 Read and accept the license agreement, as shown in Figure 4-39. Choose Next.

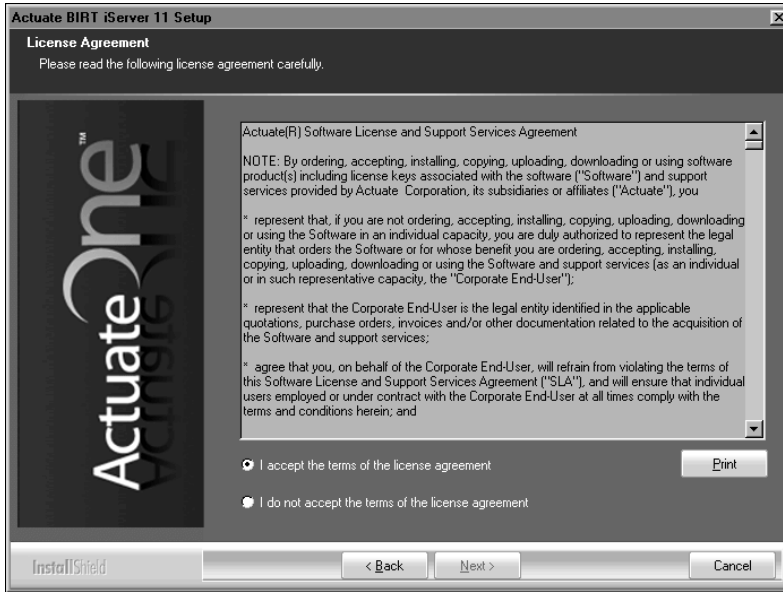


Figure 4-39 Viewing the license agreement

- 4 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-40. Choose Next.

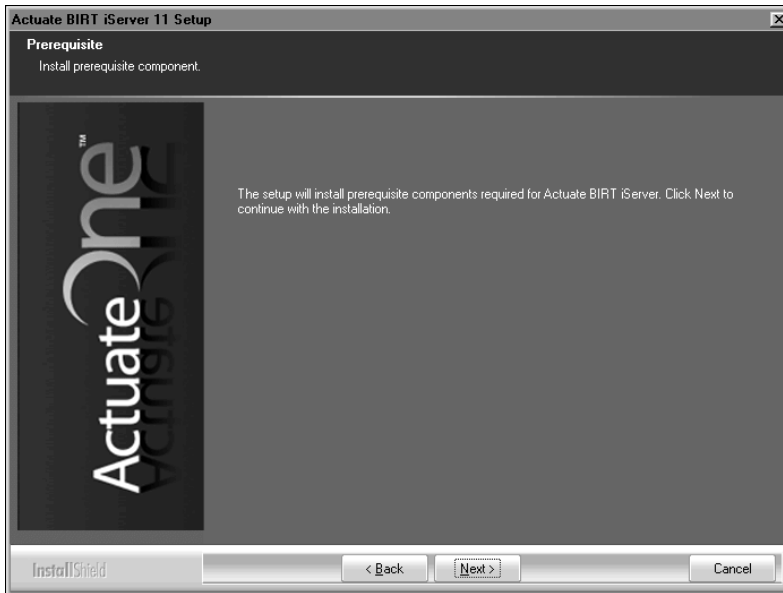


Figure 4-40 Installing Prerequisites

- 5 In Setup Type, perform the following tasks:
 - 1 Select Typical setup type.
 - 2 In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations.

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, AC_SERVER_HOME, points to the location of the iServer binaries. The default path for the program location is C:/Program Files/Actuate11SP4/iServer.

iServer uses the Data location to store the iServer logs, Encyclopedia volume files, and run-time data, including PostgreSQL data. The environment variable, AC_DATA_HOME, points to the location of the iServer data. When upgrading from an iServer Release 11 version earlier than 11SP3, the default path for the data location is C:/Actuate11/iServer/data. In Service Pack 3 and later, the default path for the data location is C:/Actuate/iServer/data. Choose Next.

The installer detects the previous Release 11 installation and selects Upgrade Existing iServer and Migrate the Encyclopedia Data. The installer lists the Existing iServer Location, as shown in Figure 4-41.

Check that the Existing iServer Location matches the location of the old files that you are upgrading. Choose Next.

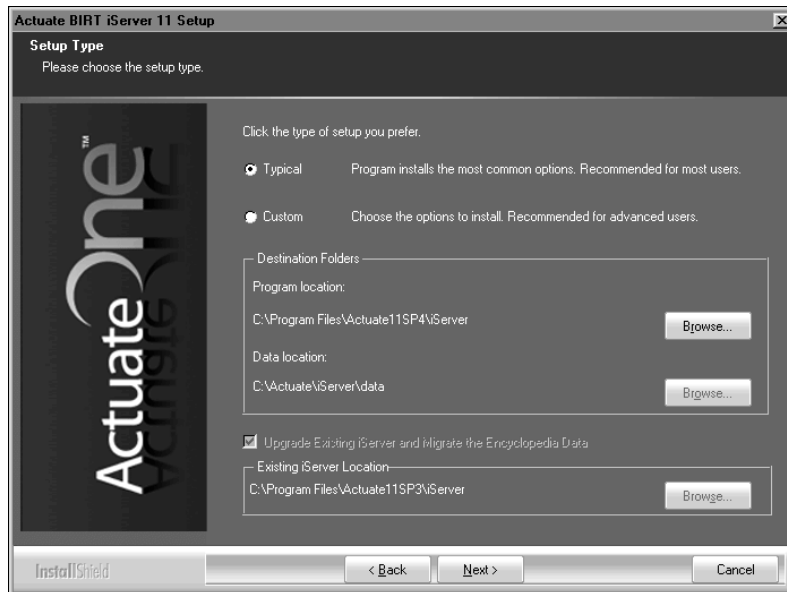


Figure 4-41 Choosing setup type

A message appears stating that the install program is migrating a previous version of the Encyclopedia volume, which might take a long time, as shown in Figure 4-42. Choose OK.



Figure 4-42 Confirming volume migration

- 6 If you have not shut down the BIRT iServer service, a prompt appears, instructing you to shut down the service, as shown in Figure 4-43. Perform the following tasks:
 - 1 Choose OK on the prompt shown in Figure 4-43.
 - 2 Stop the Actuate 11 BIRT iServer service.
 - 3 On Setup Type, choose Next.




Figure 4-43 Choosing OK to shut down iServer

- 7 If the installation uses the PostgreSQL RDBMS, you must shut down the Actuate 11 PostgreSQL for BIRT service before performing an upgrade. If you have not stopped the PostgreSQL for Actuate BIRT iServer 11 service, a prompt appears, instructing you to shut down the service, as shown in Figure 4-44.

If you are upgrading an iServer installation that uses an alternative RDBMS, such as DB2, Oracle, or Microsoft SQL Server, this prompt does not appear.

To shut down an Actuate 11 PostgreSQL for BIRT service, perform the following tasks:

- 1 Choose OK on the prompt shown in Figure 4-44.
- 2 Stop the Actuate 11 PostgreSQL for BIRT iServer service.
- 3 On Setup Type, choose Next.




Figure 4-44 Choosing OK to shut down the PostgreSQL service

- 8 In Specify Profiles, type the user name and password for the account used to start the Actuate BIRT iServer 11 service. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

If you are installing on Windows 2003, specify a user account that is in the Power Users not the Administrators group. A user account in the Administrators group cannot start the Actuate 11 BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 4-45. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

- 9 In System Configuration Password, type a password for Configuration Console, as shown in Figure 4-46. For both Configuration Console and Management Console, the default user name is Administrator. The Administrator account for Management Console has no password initially. You can log in to these consoles and change the password settings after installing iServer. Choose Next.

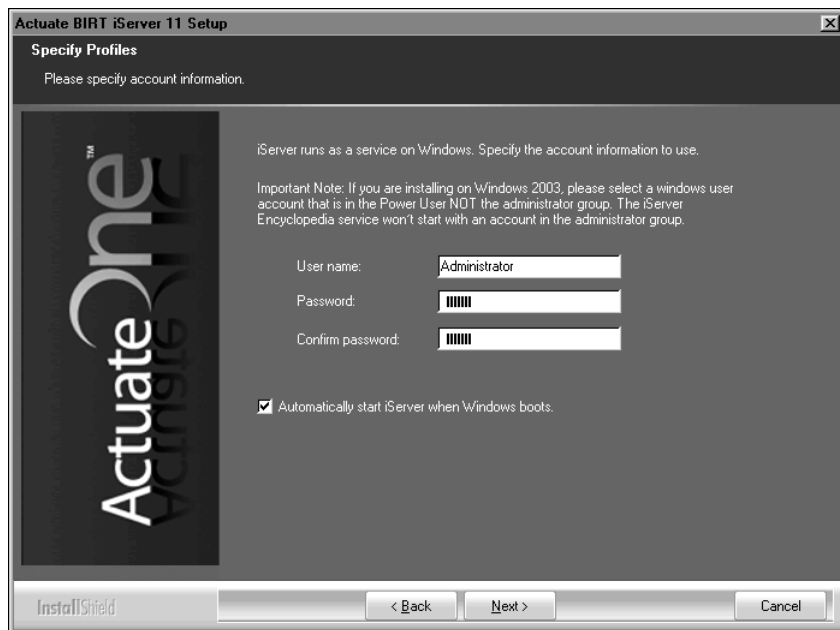


Figure 4-45 Specifying the account for running the iServer service

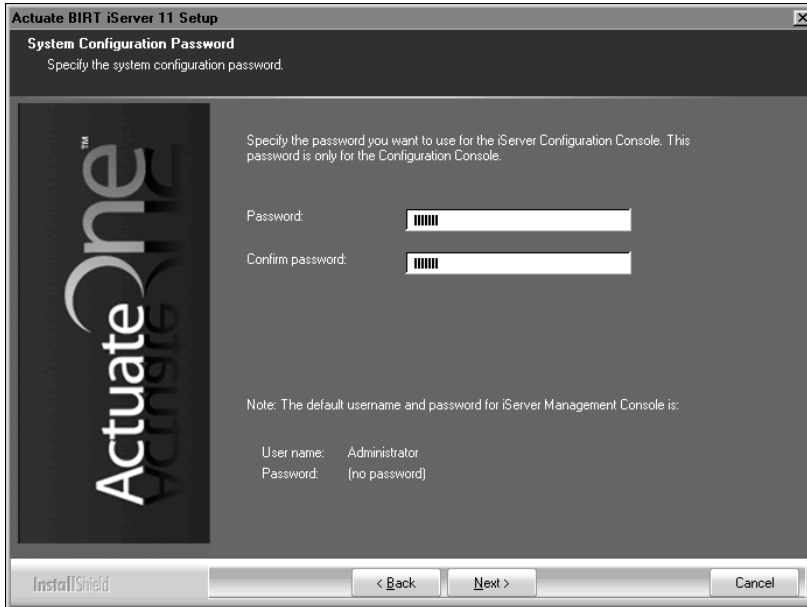


Figure 4-46 Specifying the password for using Configuration Console

10 In Start Copying Files, review the settings shown in Figure 4-47. Choose Next



Figure 4-47 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-48.

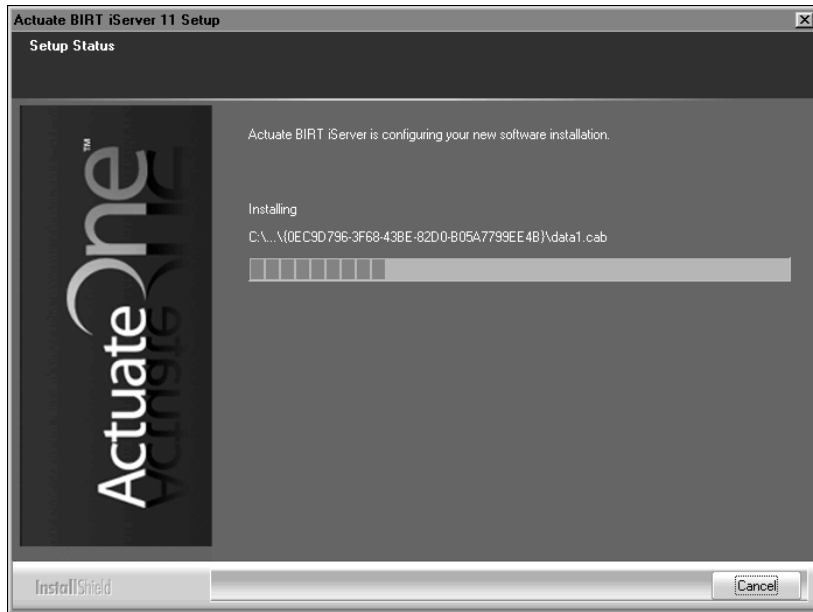


Figure 4-48 Viewing setup status

During installation, a command prompt appears, displaying messages regarding installation run-time processing, as shown in Figure 4-49.




Figure 4-49 Viewing run-time processing

How to install pgAdmin for an installation that uses PostgreSQL RDBMS

If you are upgrading an iServer installation that uses an alternative RDBMS, such as DB2, Oracle, or Microsoft SQL Server, skip this section and go to “How to

finish installing Actuate BIRT iServer 11,” later in this chapter.

For an installation that uses the PostgreSQL RDBMS, a prompt appears, asking whether to install pgAdmin, the administration tool for the PostgreSQL RDBMS, as shown in Figure 4-50. To install pgAdmin, perform the following tasks:

- 1 If you do not have pgAdmin installed, choose Yes.




Figure 4-50 Choosing to install the pgAdmin tool

Actuate BIRT iServer Setup appears with pgAdmin III Setup appearing immediately afterward, as shown in Figure 4-51.



Figure 4-51 Viewing Actuate BIRT iServer and pgAdmin III Setup

- 2 In pgAdmin III Setup, perform the following tasks:
 - 1 In Welcome, shown in Figure 4-52, choose Next.




Figure 4-52 Viewing Welcome

- 2 In License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-53. Choose Next.



Figure 4-53 Accepting the license agreement

- 3 In Custom Setup, review the features to be installed, as shown in Figure 4-54. Choose Next.




Figure 4-54 Viewing the features to be installed

- 4 In Ready to Install pgAdmin III, shown in Figure 4-55, choose Install.

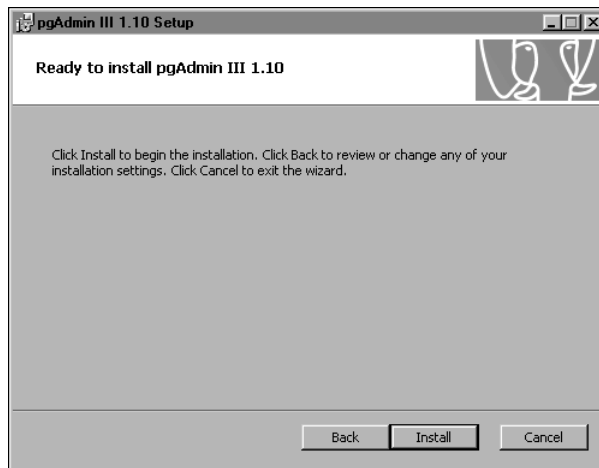


Figure 4-55 Choosing to install pgAdmin II

Installing pgAdmin III appears, as shown in Figure 4-56.




Figure 4-56 Installing pgAdmin III

- 5 When Completed, the pgAdmin III Setup Wizard appears, as shown in Figure 4-57, choose Finish to exit the pgAdmin III Setup Wizard.




Figure 4-57 Choosing Finish

How to finish installing Actuate BIRT iServer 11

- 1 When the Actuate BIRT iServer 11 setup success message appears, as shown in Figure 4-58, choose Finish to exit this wizard.



Figure 4-58 Exiting the installation wizard

- 2 The installation program prompts you to install the online help from the following location, as shown in Figure 4-59:

<http://www.actuate.com>




Figure 4-59 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows choose Start → Actuate 11 → Update Documentation.

The installation program installs shortcuts on the desktop, as shown in Figure 4-60.




Figure 4-60 iServer shortcuts installed on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to set up user accounts and run reports.

- iServer 11

Opens Welcome to Actuate iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

Choose Start→Programs→Actuate 11→BIRT iServer Configuration Console to open Configuration Console. Configuration Console Log In appears, as shown in Figure 4-61. iServer is ready for use.




Figure 4-61 Configuration Console Log In appears

About changes to the file structure

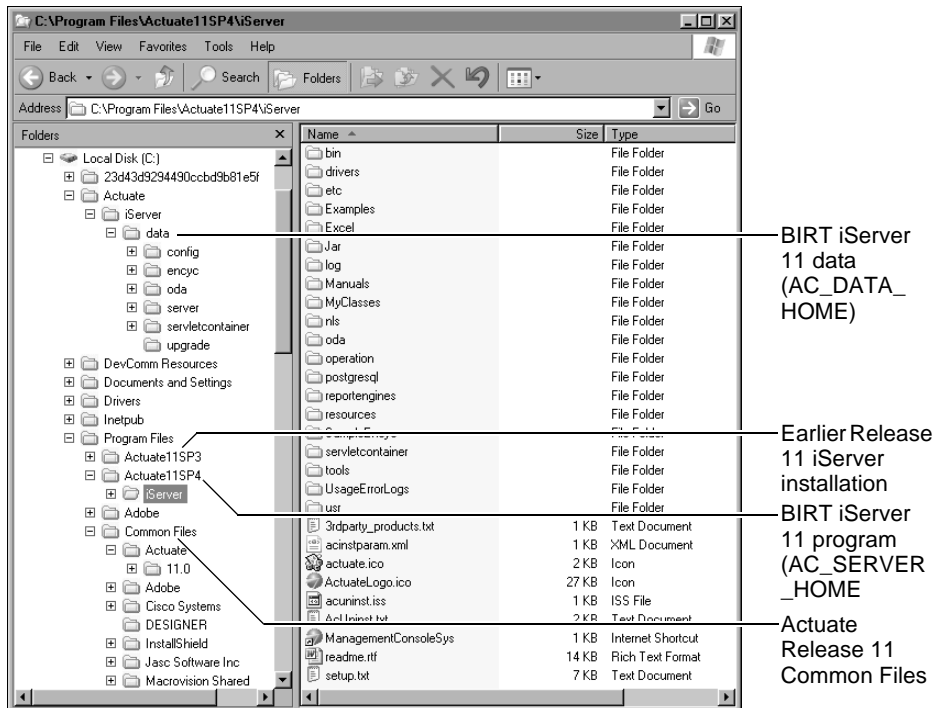


Figure 4-62 Viewing the file system after upgrade to Release 11SP4

Figure 4-62 shows the results of the changes to the upgraded installation file structure:

- The location of acserverconfig.xml is AC_DATA_HOME/config/11SP4. In Release 11 Service Pack 3, the location of acserverconfig.xml is AC_DATA_HOME/config/11SP3.
- The upgrade to Release 11 Service Pack 4 requires installing iServer program files into a new location. In this installation, the location is /Program Files/Actuate11SP4/iServer, which becomes the new AC_SERVER_HOME. The pre-release program files remain intact in case the setup needs to be restored.
- The common files location remains /Program Files/Common Files/Actuate\11.0.

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 Choose Start→Programs→Actuate 11→iServer Configuration Console. Log in to 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 4-63.




Figure 4-63 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 4-64.




Figure 4-64 Choosing to take volume offline

- 2 On Volumes, point to the icon next to the volume schema and choose Properties, as shown in Figure 4-65.




Figure 4-65 Choosing to view volume schema properties

- 3 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 4-66.

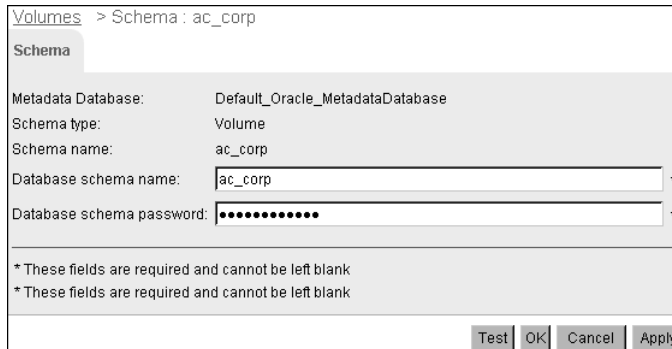


Figure 4-66 Making note of Schema name

- 4 Run the Encyclopedia Data Store Upgrader utility by performing the following tasks:

- 1 Add the following string to the PATH environment variable on your machine:

```
<AC_SERVER_HOME>\bin
```

For example, using the default value for AC_SERVER_HOME, add:

```
C:\Program Files\Actuate11SP4\iServer\bin
```

- 2 Open a command prompt and navigate to AC_SERVER_HOME/bin.
- 3 Run the upgrade_encyclopedia_data_store.bat file using the following command line syntax:

```
upgrade_encyclopedia_data_store <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 upgrdedes.bat script, using the following command-line syntax:

```
upgrdedes <schema name | property file name>
```

The batch file performs the following operations, as shown in Listing 4-2:

- ❑ 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.bat script, which sets the environment variables
- ❑ If the property file exists, the script executes the Encyclopedia Data Store Upgrader 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 4-2 upgrade_encyclopedia_data_store.bat

```
@ECHO OFF
IF %1.==. GOTO :HELP

:: Set environment variables
CALL "%~dp0"set_tools_environment.bat

:: Check if argument 1 is a file
IF EXIST %1 (
    :: Use property file specified on command line
    SET PROPERTY_FILE=%1
    GOTO :UPGRADE
)

:: Create property file
```

```

SET 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%"

:UPGRADE
:: Upgrade data store
java com.actuate.iserver.encyclopedia.datastore.admin
    .EncyclopediaDataStoreUpgrader "%PROPERTY_FILE%"

GOTO :END

:HELP
ECHO Usage: upgrade_encyclopedia_data_store.bat
    ^<schema name ^| property file name^>

:END

```

In the batch file, the ^ symbol is a line-continuation character. 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.

- 4 Type the password for the schema, as shown in Figure 4-67.

```

C:\WINDOWS\system32\cmd.exe - upgrade_encyclopedia_data_store.bat ac_corp
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

G:\Documents and Settings\Administrator>cd C:\Program Files\Actuate11SP4\iServer\bin

C:\Program Files\Actuate11SP4\iServer\bin>upgrade_encyclopedia_data_store.bat ac_corp
Detailed information will be logged to file "C:/Actuate/iServer/data/server/log/EncyclopediaDataStoreUpgrader.xrup.2012-02-01_19_01_59_PST.zg.log"
Schema password: _

```

Figure 4-67 Typing the schema password

After the utility runs, close the command prompt.

- 5 In the 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 4-68.

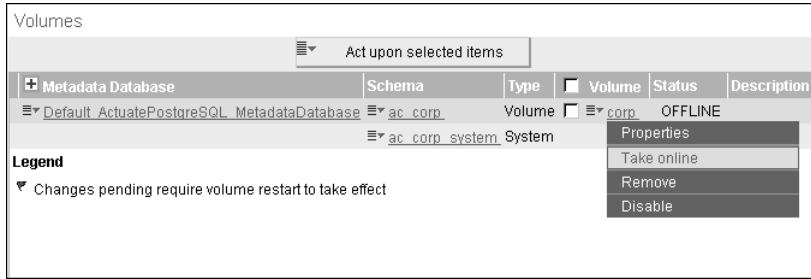


Figure 4-68 Taking the volume online

The volume comes online, as shown in Figure 4-69.




Figure 4-69 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 an earlier major 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

iServer also supports doing a side-by-side upgrade to an alternative RDBMS, such as DB2, Microsoft SQL Server, or Oracle. For more information on installing an iServer that uses one of these RDBMS, see Chapter 3, “Installing BIRT iServer

using an alternative database.”

The side-by-side 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.

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 iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 4-70. Choose Next.



Figure 4-70 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 4-71. Choose Next.

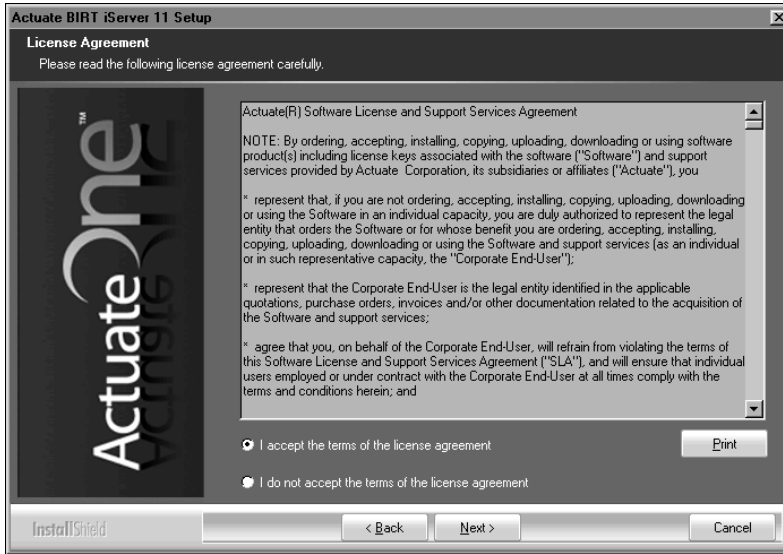


Figure 4-71 Viewing the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-72. Choose Next.

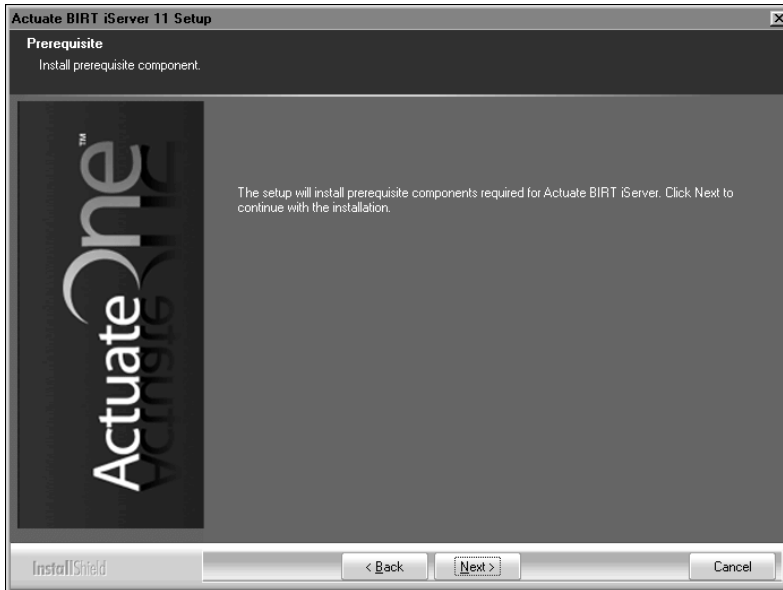


Figure 4-72 Installing Prerequisites

- 4 In Setup Type, select Custom to install a supported relational database management system (RDBMS) for the Encyclopedia volume, such as Oracle or PostgreSQL, as shown in Figure 4-73. In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations.

iServer uses the program location to resolve the paths to all the binaries that it launches. The environment variable, `AC_SERVER_HOME`, points to the location of the iServer program files. The default path for the program location is `C:\Program Files\Actuate11SP4\iServer`.

iServer uses the Data location to store the iServer logs, iServer Encyclopedia, including the PostgreSQL data, and all other run-time data. The environment variable, `AC_DATA_HOME`, points to the location of the iServer data. The default path for the data location is `C:/Actuate/iServer/data`.

Do not select Upgrade Existing iServer and Migrate the Encyclopedia Data. After running the install wizard, you use the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to manually migrate the Encyclopedia volume or volumes from the previous iServer release.

After reviewing this information, choose Next.

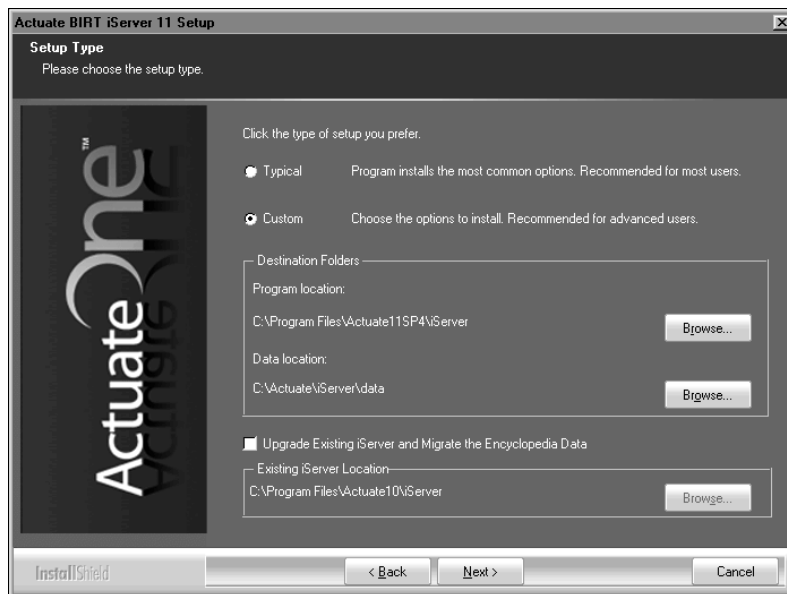


Figure 4-73 Selecting custom setup type

- 5 In Select Features, accept the default, all features selected, as shown in Figure 4-74. The installation program installs all features. Choose Next.

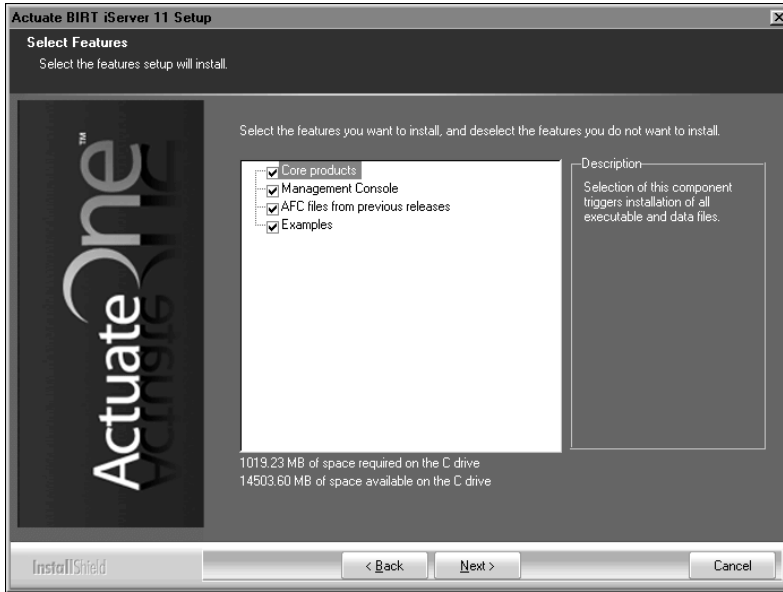


Figure 4-74 Selecting features to install

- 6 In iServer Installation Option, accept the default, Stand-alone, to install iServer as a stand-alone server, as shown in Figure 4-75. Choose Next.

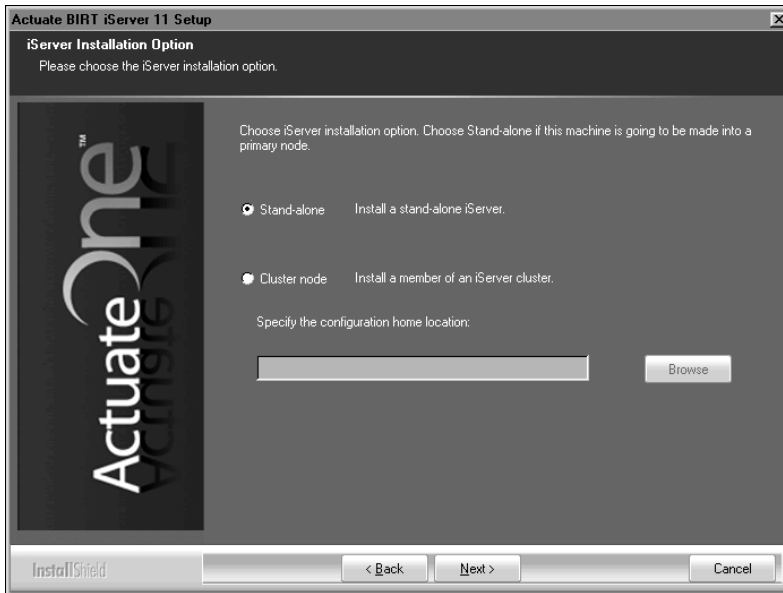


Figure 4-75 Selecting the option to install a Stand-alone iServer

- 7 In Encyclopedia Metadata Storage, select the type of Encyclopedia volume metadata database to install, as shown in Figure 4-76. This installation example demonstrates installing the bundled PostgreSQL database. In System Name, type a name for the BIRT iServer System name. iServer inserts this name into the names of the Encyclopedia volume schema and the iServer system schema. Choose Next.
- 8 On PostgreSQL Database Information, in Database Credentials, type and confirm a password, as shown in Figure 4-77. iServer creates the following accounts automatically, using this password for each account:
 - postgres
The PostgreSQL database superuser. The database superuser administers the PostgreSQL relational database management system (RDBMS).
 - ac_<BIRT iServer System name>_system
The System schema owner. iServer creates the iServer system schema and gives it this name.
 - ac_<BIRT iServer System name>
The Encyclopedia schema owner. iServer creates the Encyclopedia volume schema and gives it this name.

In PostgreSQL Database Connection Information, accept the default port 8432 or type a new port number. Choose Next.

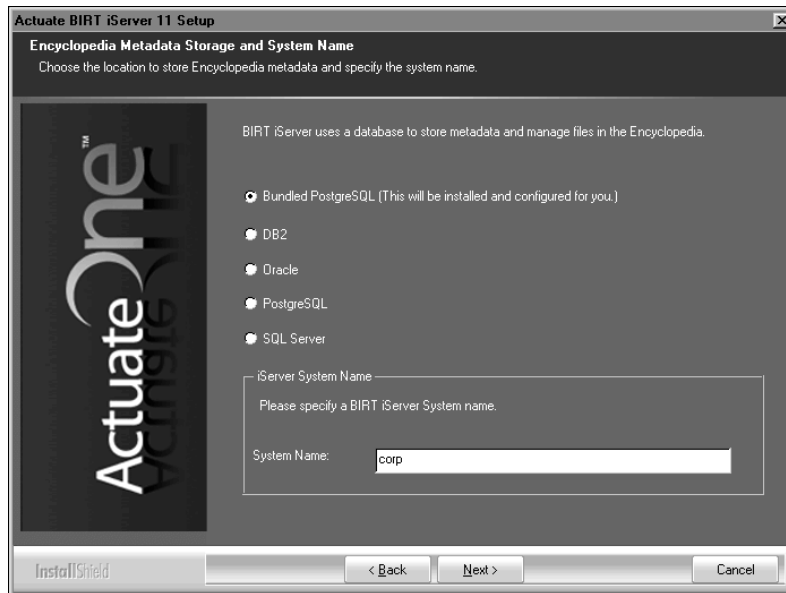


Figure 4-76 Selecting metadata database type

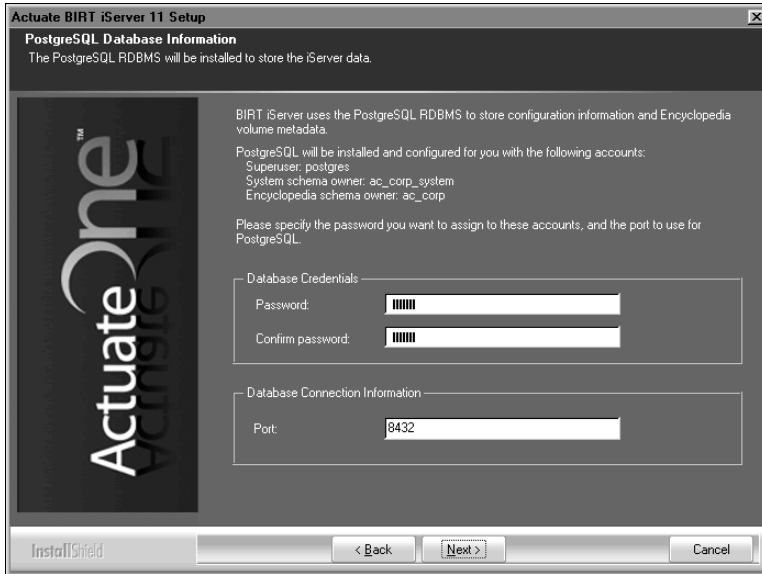


Figure 4-77 Specifying PostgreSQL database information

- 9 In License File Details, select Use the license that you purchased. Choose Browse, then navigate to and choose the license file, as shown in Figure 4-78. Alternatively, choose Try out the product using the included evaluation software license if you do not have a purchased license. Choose Next.

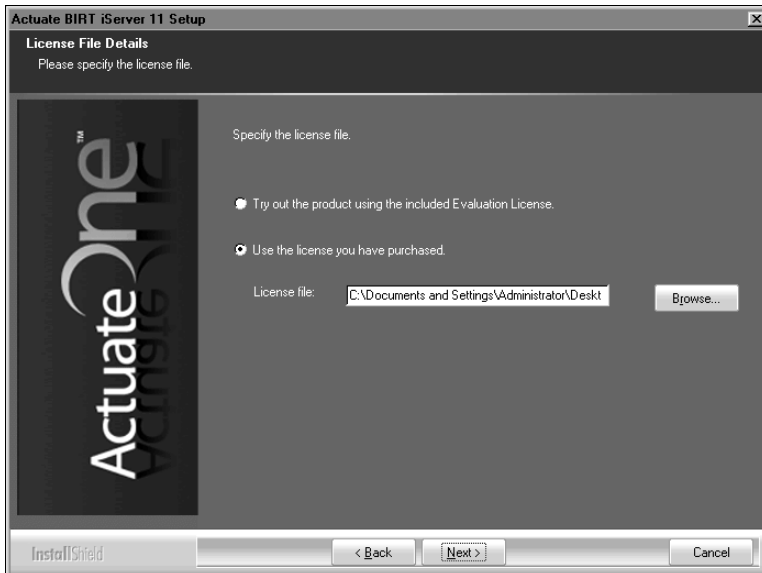


Figure 4-78 Specifying the license file

When installing using a named-user license, a prompt appears advising you to check that the volume does not exceed the number of registered users authorized by the license, as shown in Figure 4-79.




Figure 4-79 Viewing the named-user license question

Choose Yes to continue the installation.

- 10 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 4-80. Alternatively, choose the language and locale settings for your region.

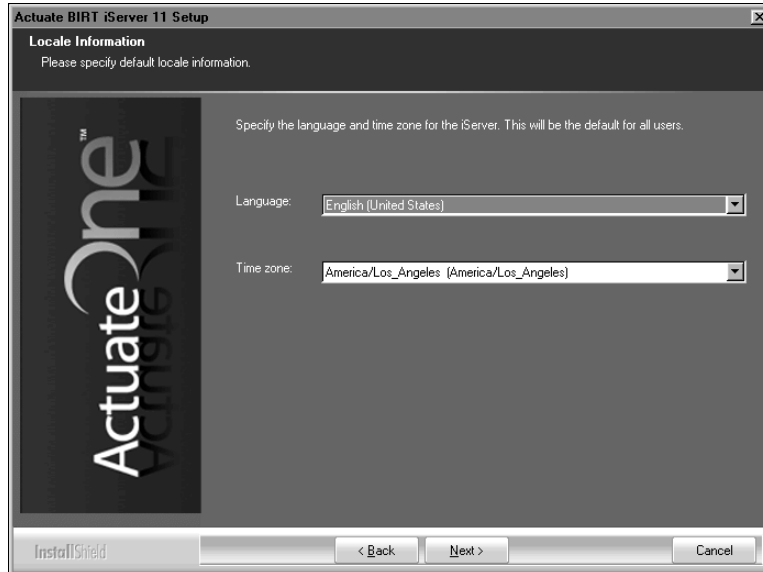


Figure 4-80 Specifying locale information

- 11 In Specify Profiles, type the user name, password, and confirm the password for the account used to start the Actuate iServer 11 service, as shown in Figure 4-81. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

If you are installing on Windows 2003, specify a user account that is in the Power Users not the Administrators group. A user account in the Administrators group cannot start the Actuate 11 BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 4-81. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

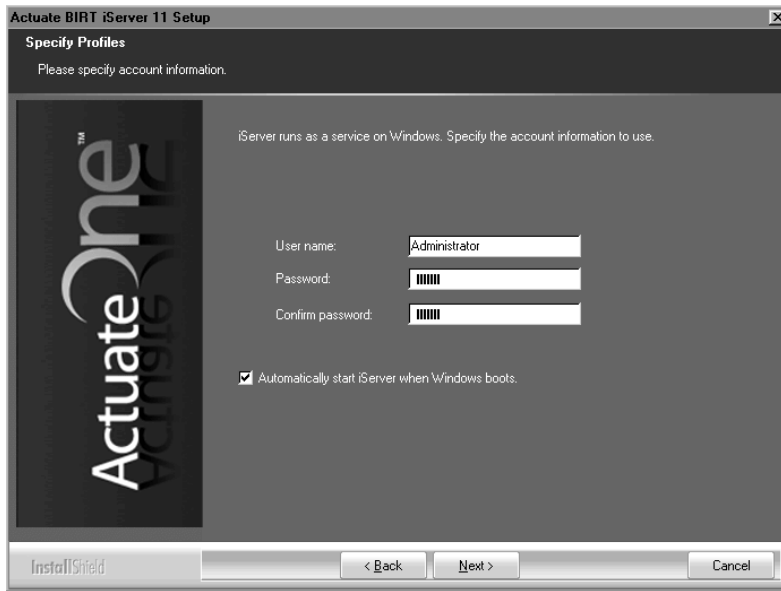


Figure 4-81 Specifying the account for running the iServer service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 4-82.




Figure 4-82 Setting the Windows local security policy

- 12** In iServer Configuration, accept the default values. Alternatively, specify the host names and port numbers where the PMD and iServer bind to listen for requests, as shown in Figure 4-83. Choose Next.
- 13** In System Configuration Password, type and confirm a password for Configuration Console, as shown in Figure 4-84. For both Configuration Console and Management Console, the default user name is Administrator. The Administrator account for Management Console has no initial password. You can log in to these consoles and change the password settings after installing iServer. Choose Next.

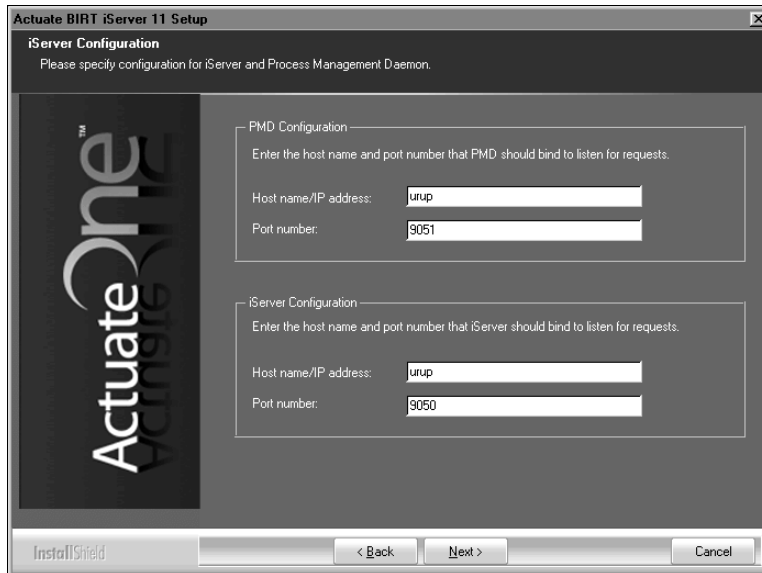


Figure 4-83 Specifying iServer configuration information

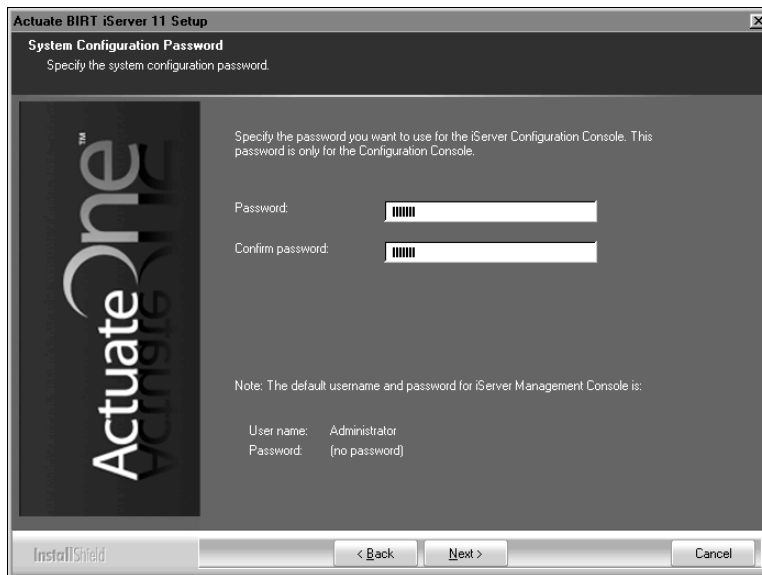


Figure 4-84 Specifying the password for using Configuration Console

- 14** In Volume Information, select Do not use the default volume, as shown in Figure 4-85. The installation process does not create the default Encyclopedia volume. Choose Next.

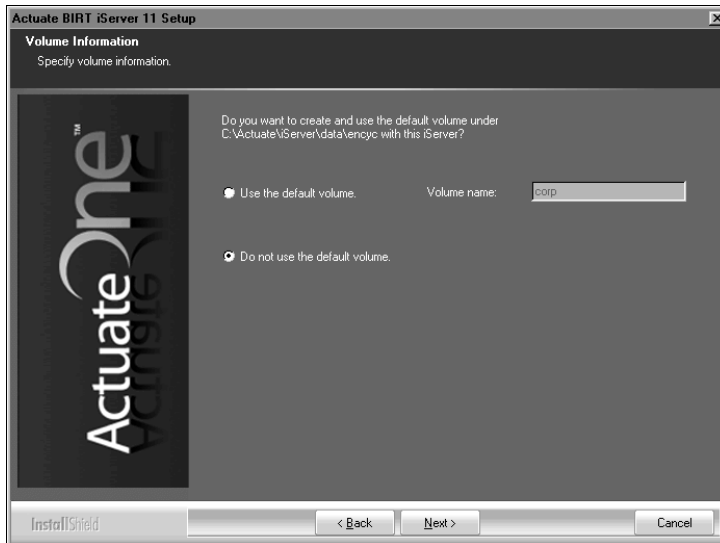


Figure 4-85 Selecting not to install the default volume

- 15 In Management Console Configuration, accept the default Host name and port number values in PMD Configuration and iServer Configuration, as shown in Figure 4-86. Alternatively, specify different values. Choose Next.

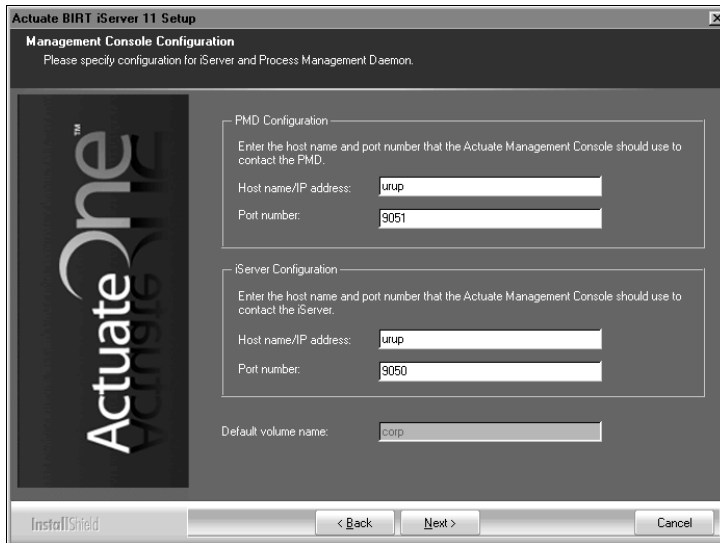


Figure 4-86 Specifying Management Console properties

- 16 In Actuate iServer Application Container, type the port number that the iServer Application Container uses, or accept the default port number, 8910, as

shown in Figure 4-87. Choose Next.

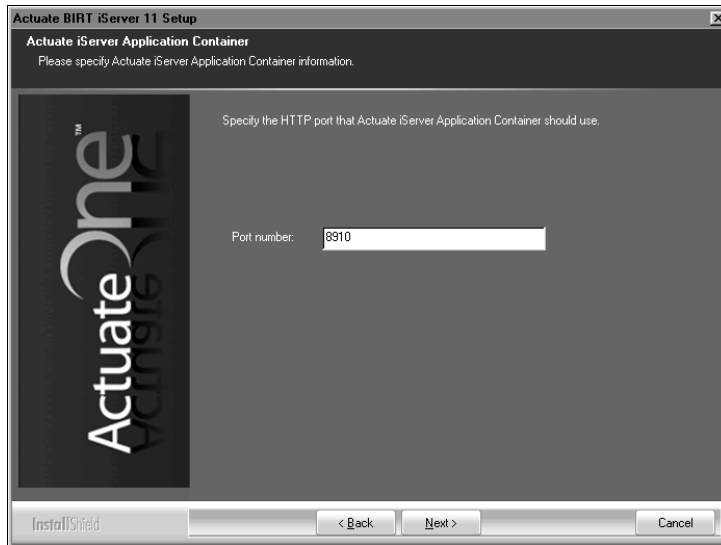


Figure 4-87 Specifying the application container port number

- 17 In Context Path, type the context path for Management Console or accept the default path, /acadmin, as shown in Figure 4-88. Choose Next.

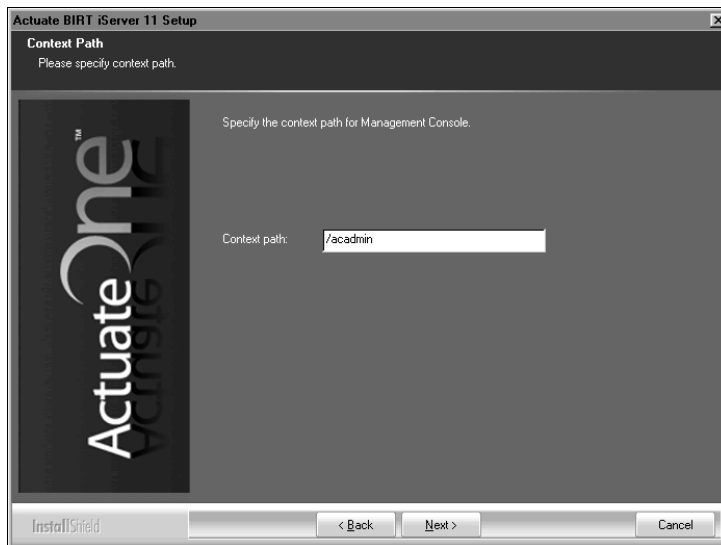


Figure 4-88 Specifying the context path

- 18 In Select Program Folder, specify a folder name containing the iServer icons that are used to launch the consoles, or accept the default name, Actuate 11, as

shown in Figure 4-89. Choose Next.

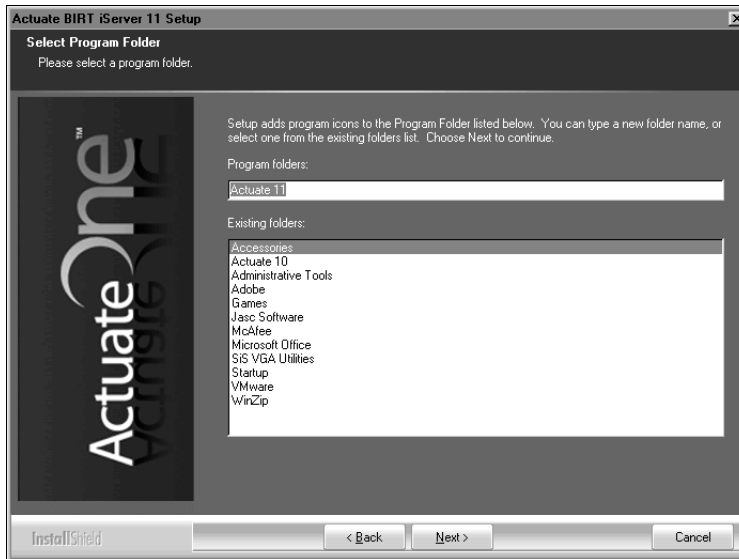


Figure 4-89 Specifying a program folder

19 In Start Copying Files, review the settings shown in Figure 4-90. Choose Next.

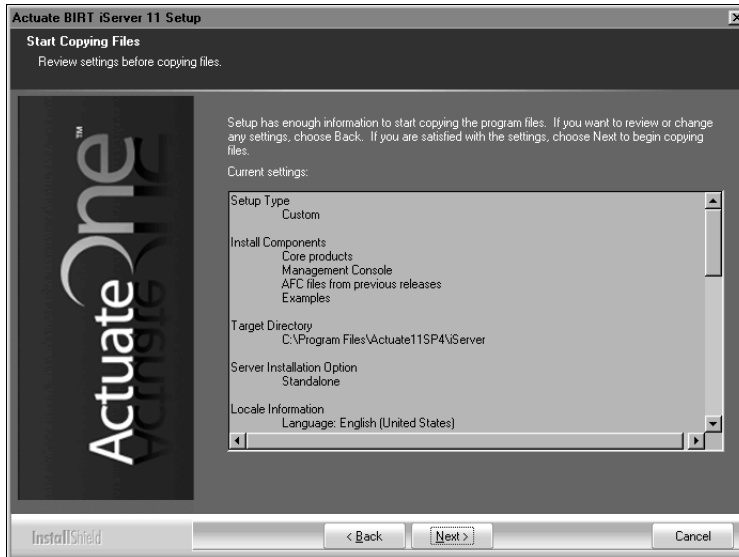


Figure 4-90 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-91.



Figure 4-91 Viewing setup status

During installation, a command prompt appears, displaying the run-time commands and messages from the initialization process of the PostgreSQL RDBMS and Encyclopedia volume schema, as shown in Figure 4-92.




Figure 4-92 Initializing the PostgreSQL RDBMS

A prompt appears, asking if you want to install the pgAdmin database administration tool when installing an iServer that uses the PostgreSQL RDBMS, as shown in Figure 4-93. Choose Yes.




Figure 4-93 Choosing to install the pgAdmin tool

Actuate BIRT iServer Setup appears, then pgAdmin III Setup appears, as shown in Figure 4-94.

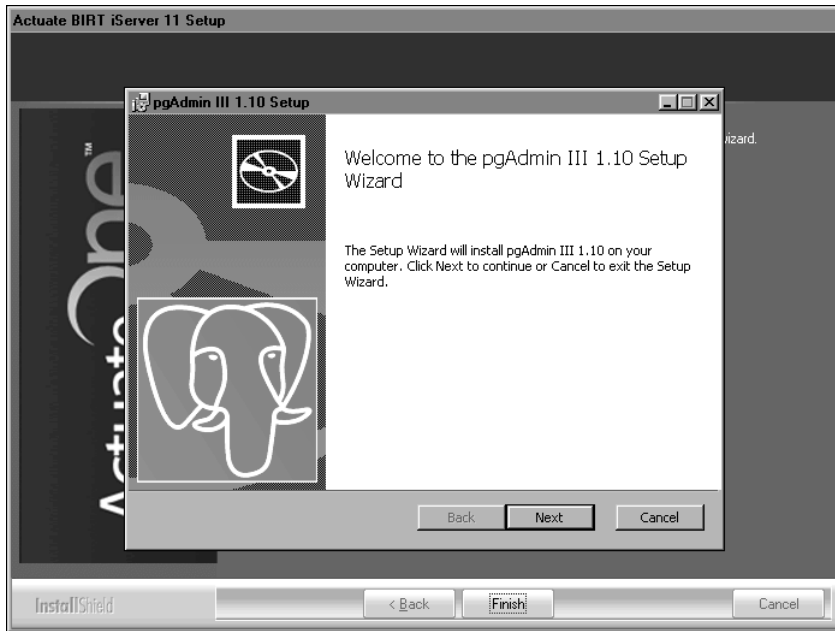


Figure 4-94 Viewing Actuate BIRT iServer and pgAdmin III Setup

20 In pgAdmin III Setup, perform the following tasks:

- 1 In Welcome, shown in Figure 4-95, choose Next.

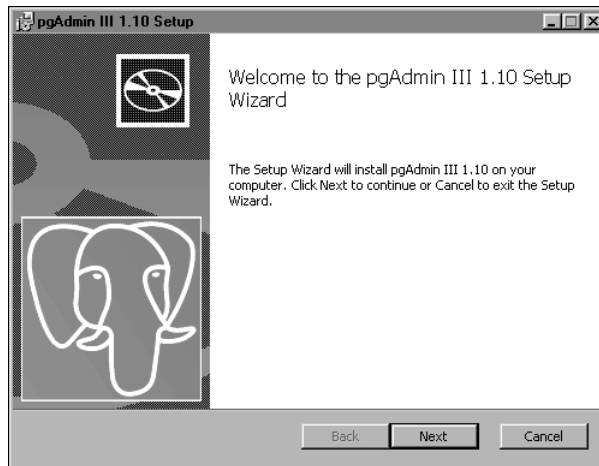


Figure 4-95 Viewing Welcome

- In End-User License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-96. Choose Next.

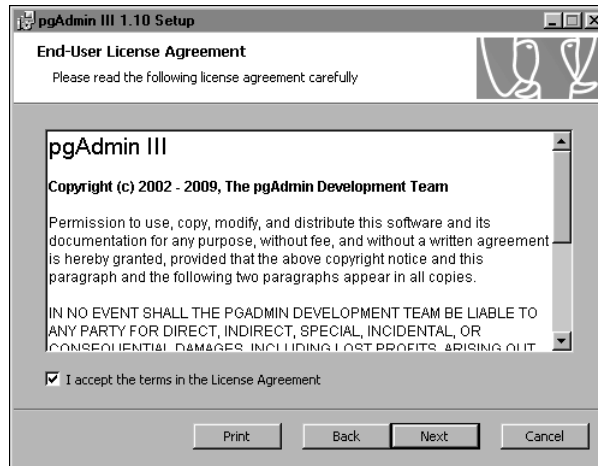


Figure 4-96 Accepting the license agreement

- In Custom Setup, review the features to be installed, as shown in Figure 4-97. Choose Next.

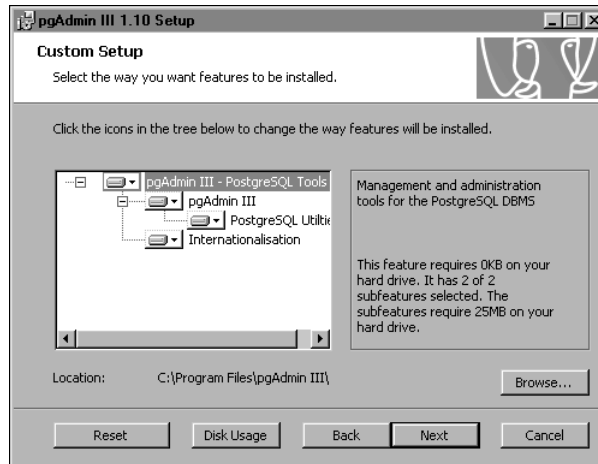


Figure 4-97 Viewing the features to be installed

- In Ready to Install pgAdmin III, shown in Figure 4-98, choose Install.

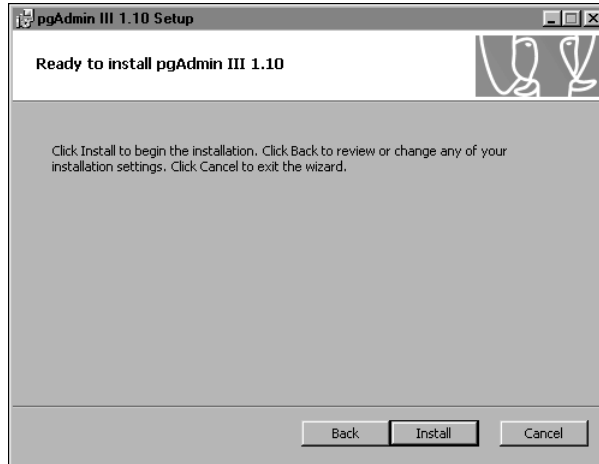


Figure 4-98 Choosing to install pgAdmin II
Installing pgAdmin III appears, as shown in Figure 4-99.




Figure 4-99 Installing pgAdmin III

- 5 When Completed the pgAdmin III Setup Wizard appears, as shown in Figure 4-100, choose Finish to exit the wizard.




Figure 4-100 Choosing Finish

- 21 Choose Finish to exit the wizard, as shown in Figure 4-101.



Figure 4-101 Exiting the installation wizard

- 22 The installation program prompts you to install the online help from the following location, as shown in Figure 4-102:

<http://www.actuate.com>




Figure 4-102 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows choose Start → Actuate 11 → Update Documentation.

The installation program installs shortcuts on the desktop, as shown in Figure 4-103.




Figure 4-103 Viewing BIRT iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to set up user accounts and run reports.
- BIRT iServer 11
Opens Welcome to Actuate BIRT iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

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

Performing a manual side-by-side migration

In the side-by-side installation, the administrator installs a new version of iServer Release 11 in a path separate from an earlier major 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 iServer Release 11 in two steps, using the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities. Finally, the administrator creates a new volume in Configuration Console as the container for the migrated volume.

The following procedure describes how to migrate the Encyclopedia volume to 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 major release to iServer Release 11, you must first export the volume 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 4-87. 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, to delete DefaultPartition, 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 4-104.

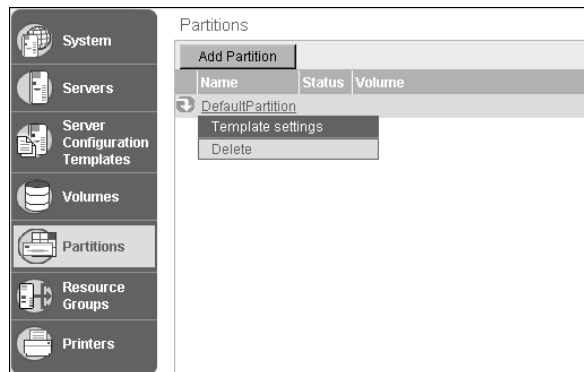


Figure 4-104 Deleting DefaultPartition

- 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 4-105.




Figure 4-105 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 4-106. \$AC_DATA_HOME\$ points to the location of the iServer data, specified during the BIRT iServer Release 11 installation, as shown in Figure 4-106. By default, this location is:

C:\Actuate\iServer\data

Alternatively, create a new folder for DefaultPartition in AC_DATA_HOME. Then, type the name of that path in Partition Path. Choose OK.

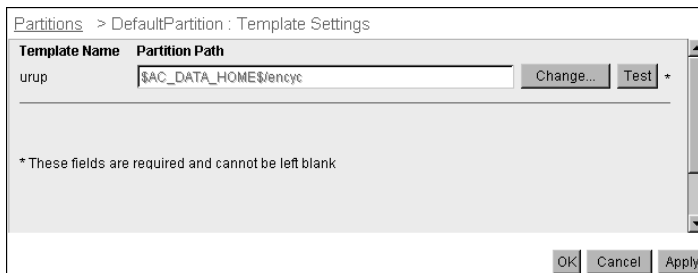


Figure 4-106 Viewing the Partition Path

- 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, as shown in Figure 4-73.
 - 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 4-107. This path does not need to match the partition path to the same volume on the earlier iServer release. Choose OK.



Figure 4-107 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 /Program Files /Actuate10/iServer.
- 7 In this step, you create a new schema that you use to create a new volume. In the 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 4-108.

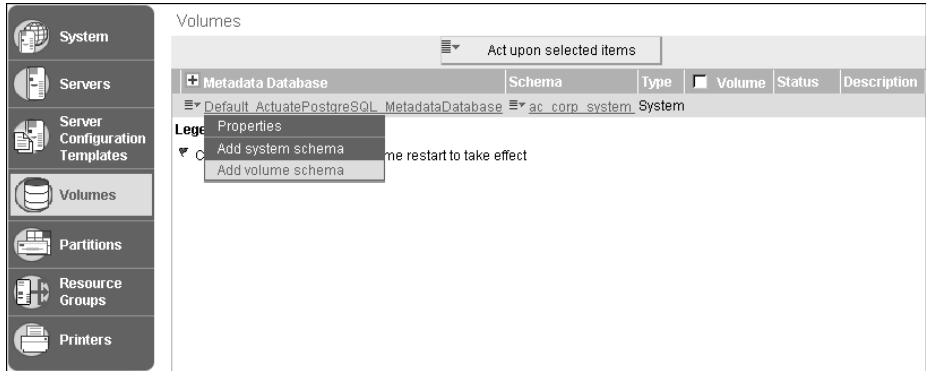


Figure 4-108 Choosing to add a new schema

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, as shown in Figure 4-77. Choose OK.

Volumes—New Schema appears as shown in Figure 4-109.

Volumes > New Volume Schema

Schema

Metadata Database: Default_ActuatePostgreSQL_MetadataDatabase

Schema type: Volume

Schema name: *

Database schema name: *

Database schema password: *

Database schema password confirm:

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

Database superuser: *

Database superuser password: *

* These fields are required and cannot be left blank

OK Cancel Apply

Figure 4-109 Creating a new schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-110.

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 4-110 Viewing the new schema

- 8 Shut down the earlier iServer release by performing the following tasks:
 - 1 Choose Start>Programs><earlier iServer version>>iServer Configuration Console. Log in to Configuration Console as Administrator.
 - 2 Choose Advanced view.
 - 3 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 4-111.




Figure 4-111 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:

```
C:\Program Files\Actuate11SP4\iServer\bin
```

- 10 Run the Squirrel Data Exporter by performing the following tasks:
 - 1 In Windows Explorer, navigate to AC_SERVER_HOME\bin in the Release 11 Service Pack 4 installation.
 - 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 = C:/Program Files/Actuate11SP4/iServer
SQUIRREL_DATA_HOME = C:/Program Files/Actuate10/iServer
/ac_corp_partition
SQUIRREL_EXPORT_FOLDER = C:/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. For additional notes on property files, see “Specifying Squirrel Data Exporter properties.”

- 11 Open a command prompt and navigate to AC_SERVER_HOME/bin in the Release 11 Service Pack 4 installation.
- 12 Run the export_squirrel_data.bat file using the following command line syntax:

```
export_squirrel_data SquirrelDataExporter.properties
```

Alternatively, you can execute the utility by running the exportsd.bat script, using the following command-line syntax:

```
exportsd SquirrelDataExporter.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.bat` script, which sets the environment variables
- Executes the Squirrel Data Exporter utility using the name of the properties file as an argument

Listing 4-3 `export_squirrel_data.bat`

```
@echo off
if %1.==. goto :HELP

:: Set up environment variables
CALL "%~dp0"set_tools_environment.bat

:: Export Squirrel data
java com.actuate.iserver.encyclopedia.datastore.admin
    .SquirrelDataExporter %1

GOTO :END

:HELP
echo Usage: export_squirrel_data.bat ^<properties file name^>

:END
```

In the batch file, the `^` symbol functions as a line-continuation character to concatenate the lines together.

How to import an Encyclopedia volume in a side-by-side installation

To import an Encyclopedia volume from an earlier iServer release to 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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/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.”

- 2 Open a command prompt and navigate to AC_SERVER_HOME/bin in the Release 11 Service Pack 4 installation.
- 3 Run the administrate_encyclopedia_data_store.bat file using the following command line syntax:

```
administrate_encyclopedia_data_store VolumeImport.properties
```

Alternatively, you can execute the utility by running the admineds.bat script, using the following command-line syntax:

```
admineds VolumeImport.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.bat script, which sets the environment variables
- Executes the Encyclopedia Data Store Administrator utility using the name of the properties file as an argument

Listing 4-4 administrate_encyclopedia_data_store.bat

```

@echo off
if %1.==. goto :HELP

:: Set up environment variables
CALL "%~dp0"set_tools_environment.bat

:: Administrate data store
java com.actuate.iserver.encyclopedia.datastore.admin
    .EncyclopediaDataStoreAdministrator %1

```

```

GOTO :END

:HELP
echo Usage: administrate_encyclopedia_data_store.bat
      ^<properties file name^>

:END

```

In the example, the ^ symbol functions as a line-continuation marker to concatenate the lines together.

- 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 volume schema and choose Add Volume, as shown in Figure 4-112.




Figure 4-112 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 4-113. If the migrated volume uses DefaultPartition, specify DefaultPartition in Primary partition.

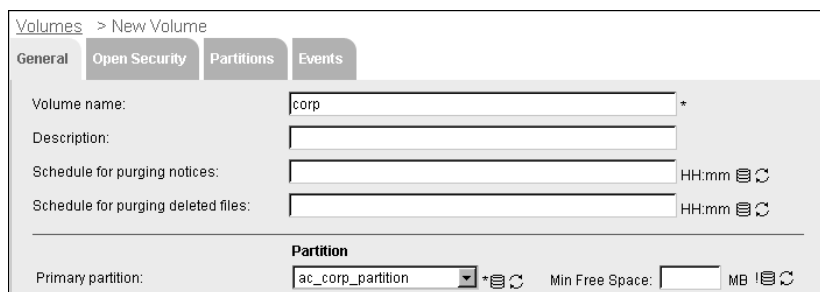


Figure 4-113 Specifying volume and partition name

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 4-114.

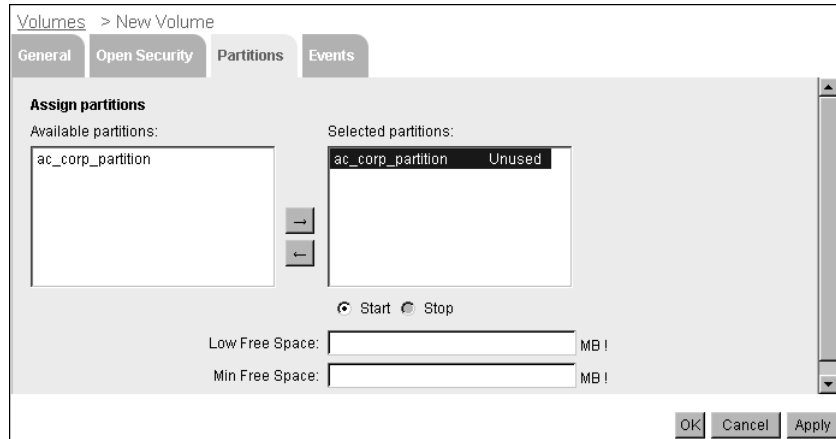


Figure 4-114 Assigning a partition

Choose OK.

- 3 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-115.




Figure 4-115 Viewing the new volume

The new volume comes online, as shown in Figure 4-116.




Figure 4-116 Viewing the new volume after it comes online

- 4 Log in to Management Console. In Files and Folders, the data from the 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.”

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, as shown in Figure 4-73.
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-z 0-9]*. 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 volume name on the earlier iServer release.
SQUIRREL_DATA_HOME	Absolute path to the folder on the 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 .dtl and .ctl 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 import 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, as shown in Figure 4-73.
<code>DATABASE_TYPE</code>	Type of supported RDBMS that contains the data store. Specify DB2, Oracle, PostgreSQL, or SQLServer.
<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-z0-9]*. 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"> ■ APPLICATION_USER ■ DATABASE_HOST ■ DATABASE_INSTANCE ■ DATABASE_NAME ■ DATABASE_PORT ■ DATABASE_TYPE ■ ORACLE_TNS_NAMES_FILE ■ SCHEMA_NAME 	False	All
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, Microsoft SQL Server, or Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL

Table 4-4 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
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 DB2, Oracle, PostgreSQL, or SQL Server.	{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, SQL Server, 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_INSTANCE	Specifies the SQL Server database instance. The default value is the name of the database instance on the host machine and port. The default name of the SQL Server 2008 database instance is MSSQLSERVER.		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
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)

Parameter	Description	Default value	Supported databases
DATABASE_TYPE	Type of RDBMS that contains the data store. Actuate Release 11 currently supports DB2, PostgreSQL, Microsoft SQL Server, and Oracle.		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
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

Table 4-4 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
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 appends 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.	
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

(continues)

Table 4-4 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
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 ancillary structures in database.	POPULATE_SCHEMA	All
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

Table 4-4 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
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.	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 = C:/Program Files/Actuate11SP4/iServer
DATABASE_TYPE = PostgreSQL
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a 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 = C:/Program Files/Actuate11SP4/iServer
DATABASE_TYPE = PostgreSQL
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a 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 book.

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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/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 book.

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 = C:/Program Files/Actuate11SP4/iServer
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 = D:/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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/Program Files/Actuate11SP4/iServer
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 = C:/Program Files/Actuate11SP4/iServer
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 and Initializing a New Volume in a New Schema

This operation is only supported for PostgreSQL. iServer typically performs this operation when you create a new schema and 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 = C:/Program Files/Actuate11SP4/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a 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 system metadata. The system schema is a required element for any iServer installation. In a cluster, the nodes share system schema metadata and use this information for communicating and coordinating 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.

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:

```
C:\Program Files\Actuate11SP4\iServer\bin
```

To run the System Data Store Administrator utility, perform the following tasks:

- 1 In Windows Explorer, 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 book. You pass this file to the System Data Store Administrator utility when you execute it.
- 3 Open a command prompt and navigate to AC_SERVER_HOME/bin.
- 4 Run the administrate_system_data_store.bat file using the following command line syntax:

```
administrate_system_data_store systemdatastore.properties
```

Alternatively, you can execute the utility by running the adminstds.bat script, using the following command-line syntax:

```
adminstds systemdatastore.properties
```

The batch file performs the following operations, as shown in Listing 4-5.

- 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.bat script, which sets the environment variables
- Executes the System Data Store Administrator utility using the name of the properties file as an argument

Listing 4-5 administrate_system_data_store.bat

```
@echo off
if %1.==. goto :HELP

:: Set up environment variables
CALL "%~dp0"set_tools_environment.bat

:: Administrate data store
```

```

java com.actuate.iserver.system.datastore.admin
    .SystemDataStoreAdministrator %1

GOTO :END

:HELP
echo Usage: administrate_system_data_store.bat ^<properties
    file name^>

:END

```

In the example, the ^ symbol functions as a line-continuation marker to concatenate the lines together.

The SystemDataStoreAdministrator class has the same parent class as the Encyclopedia Data Store Administrator and uses the same property settings. For more information about these properties, see Table 4-3.

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, superuser, 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

Creating and populating a new system data store schema 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 = C:/Program Files/Actuate11SP4/iServer
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 with 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, as shown in Figure 4-73.
APPLICATION_USER	User ID used to connect to the database for normal operations.
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 <code>NEW_SCHEMA_NAME</code> is not present. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern <code>[a-z][a-z 0-9]*</code> . 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
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

Table 4-6 Optional Encyclopedia Data Store Upgrader properties

Parameter	Description	Default value	Supported databases
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
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
- Performing a wizard-based cluster node installation
- Adding a node to a cluster

Installing a BIRT iServer cluster node

A node is a machine running a 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:

- 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 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.

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 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 the iServer cluster environment

This section assumes the administrator has already performed a typical, wizard-based installation to create a standalone iServer instance on one computer. This machine contains the shared configuration directory, which all nodes in the cluster access. This section refers to the machine containing the shared configuration directory as node1.

In the section, “Performing a wizard-based cluster node installation,” later in this chapter, the Administrator chooses the custom installation option in the wizard-based install program to create a cluster node. This cluster node, referred to as node2, accesses the shared resources on node1.

Before performing a cluster node installation, the Administrator performs the following tasks:

- On node1:
 - Shares the folders that all cluster nodes access
 - Turns off the Windows firewall
 - Obtains the machine host name and IP address
 - Tests the network accessibility of the machine

- Reconfigures the partition path for the default partition in Configuration Console using UNC format
 - Edits acserverconfig.xml and sets a property to make the metadata database visible to all nodes
- On node2:
- Turns off the Windows firewall.
 - Obtains the machine host name and IP address.
 - Tests the network accessibility of the machine.

The following instructions provide a basic example of the operations required to configure network sharing and firewall settings in the Windows environment. 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.

How to share the encyc and config\11SP4 folders

In an iServer Release 11 Service Pack 4 installation, cluster nodes must have read-write access to AC_DATA_HOME\encyc and AC_DATA_HOME\config\11SP4 folders on node1. To give a cluster node read-write access to the folders and files in the encyc folder and to the files in the config\11SP4 folder, perform the following tasks on node1:

- 1 Using Windows Explorer on node1, right-click the folder, AC_DATA_HOME\encyc. Choose Sharing and Security, as shown in Figure 5-1.

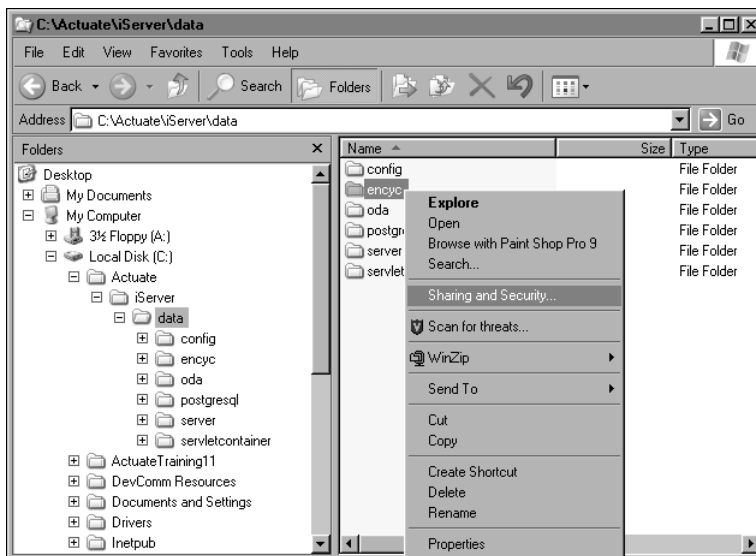


Figure 5-1 Choosing Sharing and Security

2 On encyc Properties—Sharing, select the following options:

- Share this folder on the network
- Allow network users to change my files

Figure 5-2 shows the selected options.

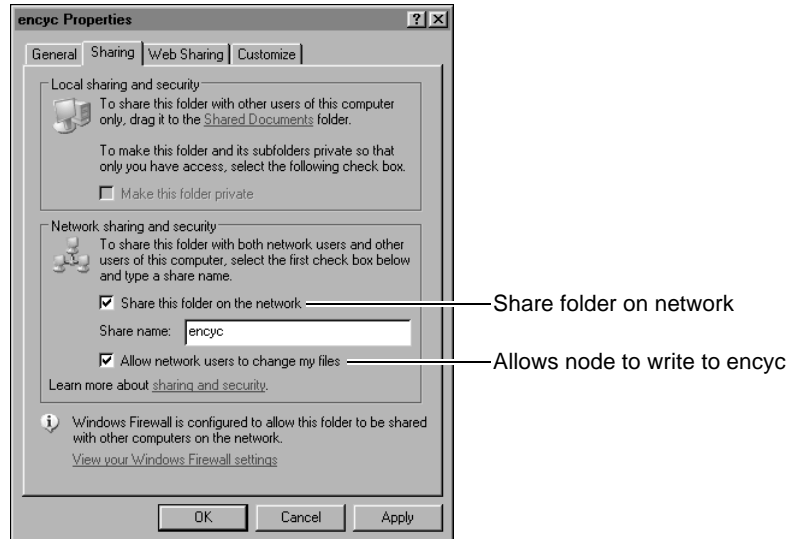


Figure 5-2 Sharing the file folder on the network

Choose OK.

- 3 Repeat steps 1 through 3 for the `\config\11SP4` folder. Make sure that all settings conform to the security policies in force for the environment.
- 4 Close Windows Explorer.

How to turn off the Windows firewall

Perform the following steps on node1 and node2:

- 1 Choose Start → Settings → Control Panel.
- 2 In Control Panel, choose Windows Firewall.
- 3 On Windows Firewall—General, select Off, as shown in Figure 5-3. Make sure that all settings conform to the security policies in force for the environment.

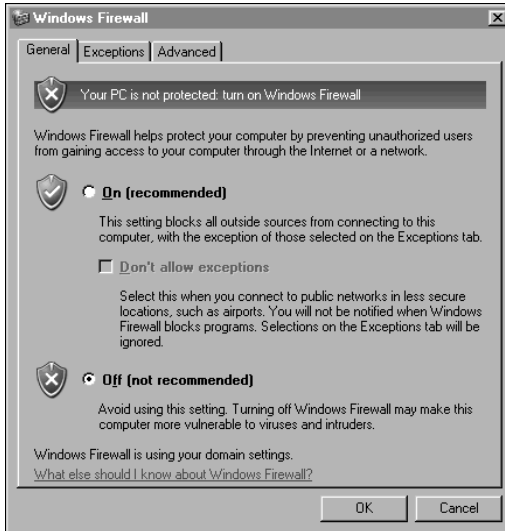


Figure 5-3 Turning off the Windows firewall
Choose OK.

How to display a computer's IP address

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

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

```
ipconfig /all
```

Press Enter. The host name appears, as shown in Figure 5-4. In this example, the host name for node1 is urup.

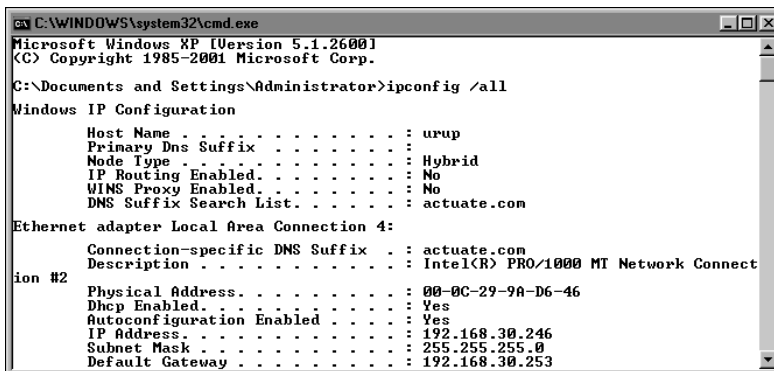


Figure 5-4 Displaying the host name

- 3 Write the host name and IP address in the spaces provided in Table 5-1.

Table 5-1 Host names and IP addresses of computers to be clustered

iServer	Host name	IP address
Node1	urup	192.168.30.246
Node2	kozu	192.168.30.233

How to test the connection between computers

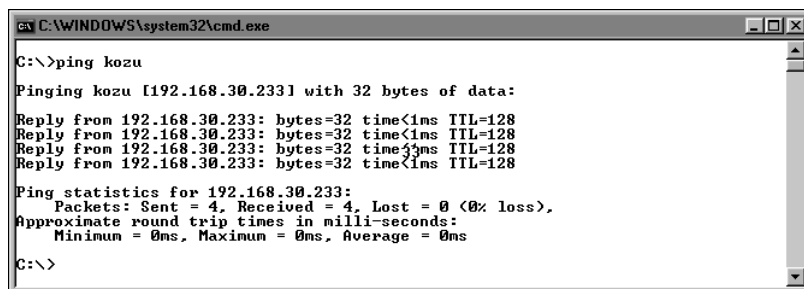
Perform the following steps on both computers:

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

```
ping kozu
```

Press Enter.

If your computer reaches the other computer, Command Prompt displays a series of replies, as shown in Figure 5-5.



```
C:\WINDOWS\system32\cmd.exe
C:\>ping kozu
Pinging kozu [192.168.30.233] with 32 bytes of data:
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Ping statistics for 192.168.30.233:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\>
```

Figure 5-5 Receiving a reply to a ping command

- 2 Close Command Prompt.

How to prepare node1 for clustering

On node1, or `urup` in this example, perform the following tasks:

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Volumes. On Volumes, point to the icon next to the default volume, `corp` in this example, and choose Take offline, as shown in Figure 5-6.

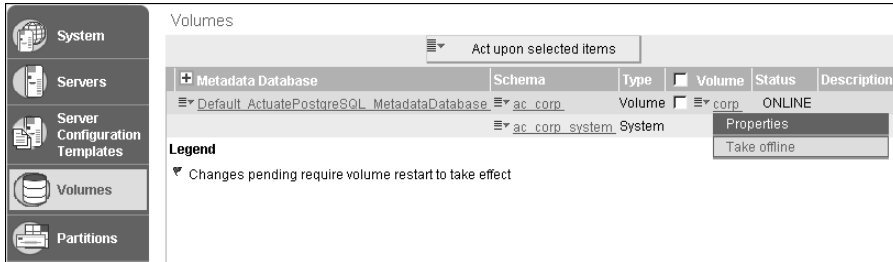


Figure 5-6 Taking a volume offline

In Volume Offline Grace Period, choose OK.

From the side menu, choose Partitions.

- 2 On Partitions, point to the icon next to DefaultPartition and choose Template settings, as shown in Figure 5-7.




Figure 5-7 Choosing Template settings

On Template Settings, choose Change. In Partition Path, type the path to the encyc folder, using UNC format. For example, if the hostname of node1 is urup, type:

\\urup\encyc

Template Settings appears as shown in Figure 5-8.

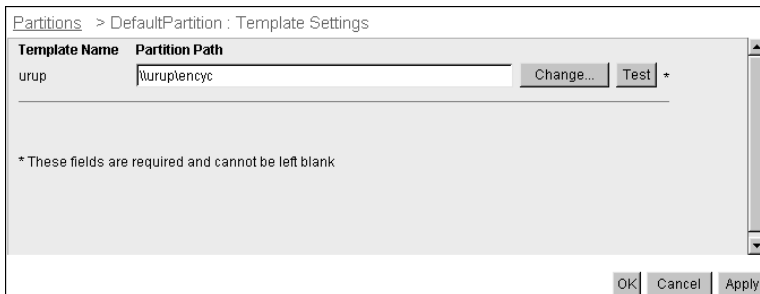


Figure 5-8 Typing the partition path

To verify that iServer can access the encyc folder, choose Test.

A message appears, stating that the test was successful, as shown in Figure 5-9. Choose OK.



Figure 5-9 Verifying a successful partition path test

On Template Settings, choose OK.

From the side menu, choose System.

- 3 On System—Status, choose Stop to stop the system, as shown in Figure 5-10.




Figure 5-10 Choosing to stop the system on System Status

- 4 Log out of Configuration Console.

- 5 Choose Start → Settings → Control Panel.

In Control Panel, choose Administrative Tools—Services.

On Services, stop the Actuate 11 BIRT iServer service.

- 6 Open Windows Explorer, and navigate to AC_DATA_HOME/config/11SP4. Make a backup copy of acserverconfig.xml, then open acserverconfig.xml and perform the following tasks:

- 1 Locate the <ConnectionProperties> element under the <MetadataDatabase> element.

- 2 Under <ConnectionProperties>, locate:

```
<ConnectionProperty
  Name="server"
  Value="localhost"/>
```

- 3 Change Value from localhost to the name of the machine on which the volume resides, in this example “urup”. For urup, the volume is located on localhost, but in a cluster setting the administrator has to specify the machine name, since the volume is not on a localhost from the reference point of the other nodes.

- 4 Save acserverconfig.xml, and exit the file.

- 7 On Services, start the Actuate 11 BIRT iServer service.

Creating an account with Windows administrator privileges

Before installing iServer, create a Windows user account that is a member of the Administrators group. Use this account when installing and running iServer.

The iServer user account must meet the following requirements:

- Be a member of the Windows Administrators group.
The account must have privileges to access the required software and hardware, such as database servers, printers, and iServer files and folders.
- Have log on as a service privilege.
If the account does not meet this requirement, the iServer installation program prompts you to configure the privilege to run the Windows Actuate iServer service.

On a new Windows Vista installation, the initial user account is not a member of the Administrators group. You must configure this user account to be a member of this group.

When installing iServer in Windows 2003, create a Windows user account that is a member of the Power Users not the Administrators group. Make sure that the Account has permission to access any printers required for printing. Perform the installation using an account that has Administrator privileges. During the installation, when prompted to specify the user account to run the iServer service, specify the Power User account.

For more information about configuring a Power User and iServer account and log on as a service privilege, see “Creating an account with Windows administrator privileges,” in Chapter 2, “Installing BIRT iServer.”

Performing a wizard-based cluster node installation

This section assumes the Administrator has already performed the wizard-based Typical install to create a standalone iServer on one computer and has performed the tasks described in the Preparing the iServer cluster environment section, earlier in this chapter. To create a standalone iServer installation, see “Performing a new installation,” in Chapter 2, “Installing BIRT iServer.” To prepare the standalone iServer node for clustering, see , “Preparing the iServer cluster environment,” earlier in this chapter.

When creating a BIRT iServer cluster, the Administrator must install and run all cluster nodes using the same administrative user account. The following section describes how to install an iServer Release 11 cluster node in the Windows operating system using the installation wizard.

How to install a cluster node in Windows

- 1 Download the iServer distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 5-11. Choose Next.



Figure 5-11 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 5-12. Choose Next.

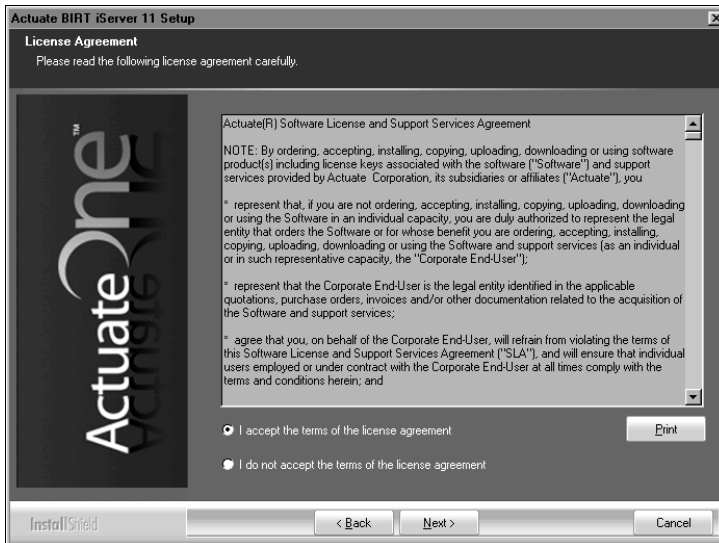


Figure 5-12 Accepting the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 5-13. Choose Next.

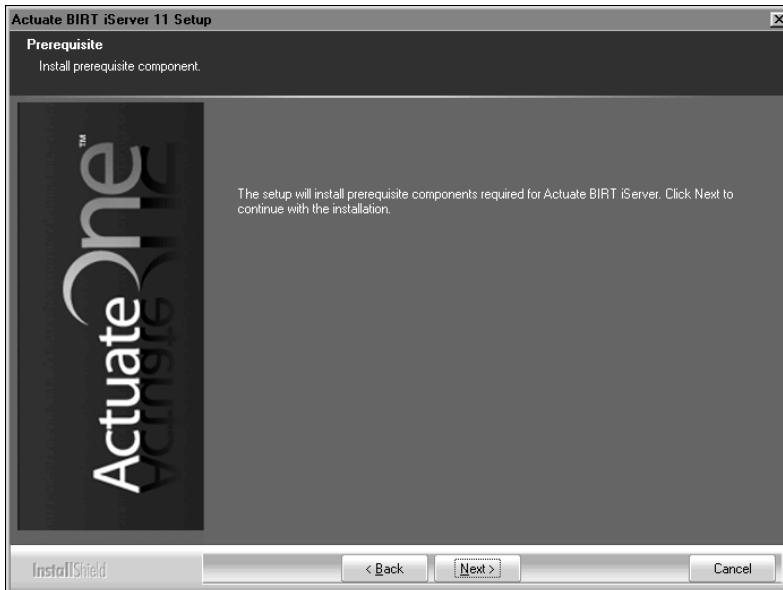


Figure 5-13 Installing Prerequisites

- 4 In Setup Type, select Custom, as shown in Figure 5-14.

In Destination Folder, accept the default or choose a new destination folder for the Program and Data locations. Choose Next.

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, `AC_SERVER_HOME`, points to the location of the iServer binaries. The default path for the program location is `C:/Program Files/Actuate11SP4/iServer`.

iServer uses the Data location to store the iServer logs, iServer encyclopedia, including the PostgreSQL data, and all other run-time data. The environment variable, `AC_DATA_HOME`, points to the location of the iServer data. The default path for the data location is `C:/Actuate/iServer/data`.

Each iServer instance must have its own `AC_SERVER_HOME` and `AC_DATA_HOME` folders. These folders cannot be shared by other nodes in a cluster.

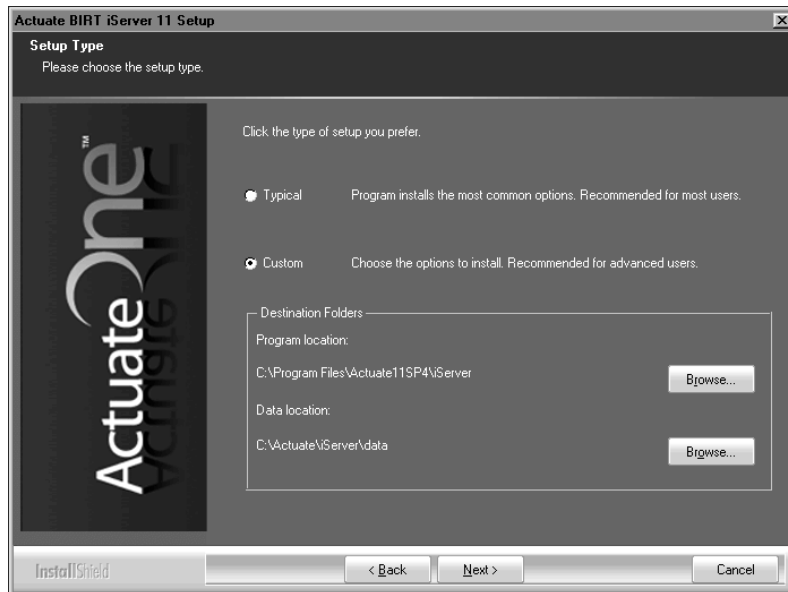


Figure 5-14 Selecting the custom installation

- 5 In Select Features, select all features—Core products, Management Console, AFC files from previous releases, and Examples, as shown in Figure 5-15. A cluster node must have access to Configuration Console. Configuration Console installs with Management Console. Choose Next.
- 6 In iServer Installation Option, select Cluster node to install iServer as a cluster node. Then specify the configuration home location, as shown in Figure 5-16.

The configuration home location is the shared directory containing the `acserverconfig.xml`, `acserverlicense.xml`, and `acconfigowner.lock` files for the

cluster. The administrator should specify the location using the Universal Naming Convention (UNC) Format.

In a Release 11 Service Pack 4 installation, the configuration files are located in AC_DATA_HOME\config\11SP4 by default. If the shared configuration folder is in this default location on a server named urup, then the administrator should specify \\urup\11SP4 as the path.

Choose Next.

- 7 In License File Details, select Use the license that you purchased. Choose Browse then navigate to and choose the license file, as shown in Figure 5-17. Choose Next. Alternatively, choose Try out the product using the included evaluation software license if you do not have a purchased license.

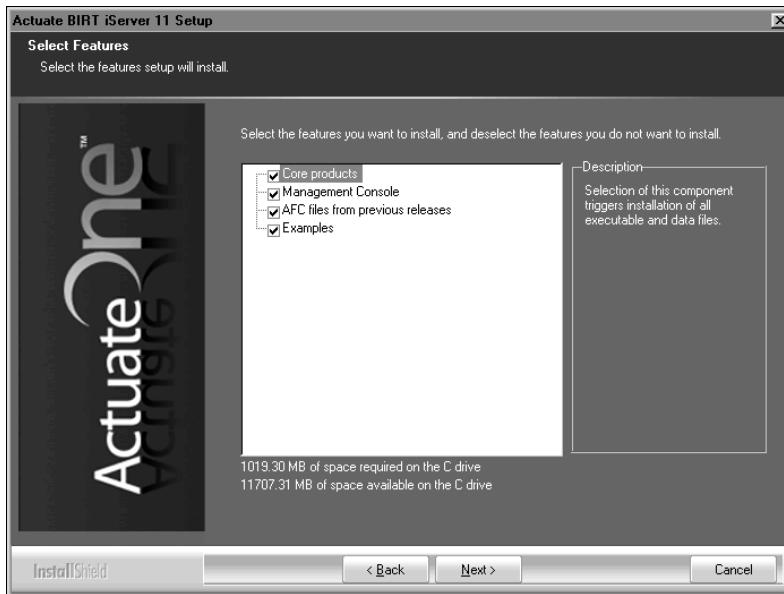


Figure 5-15 Selecting features to install

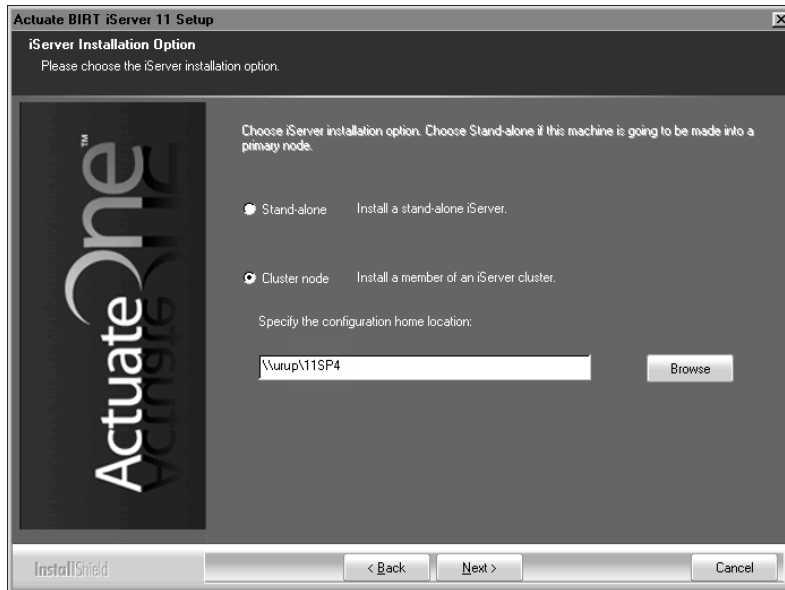


Figure 5-16 Choosing the option to install a node

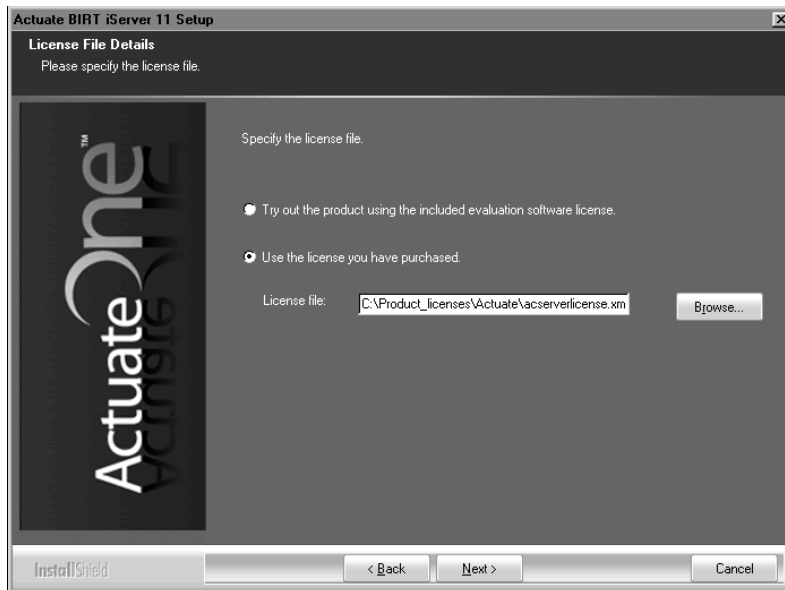


Figure 5-17 Specifying the license file

If installing using a named user license, a prompt appears advising you to check that the volume does not exceed the number of registered users authorized by the license, as shown in Figure 5-18.




Figure 5-18 Viewing the named user license question

Choose Yes to continue the installation.

- 8 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 5-19. Alternatively, choose the language and locale settings for your region.
- 9 In Specify Profiles, type the user name, password, and confirm the password for the account used to start the Actuate BIRT iServer 11 service, as shown in Figure 5-20. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

If you are installing on Windows 2003, specify a user account that is in the Power Users not the Administrators group. A user account in the Administrators group cannot start the Actuate 11 BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 5-20. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

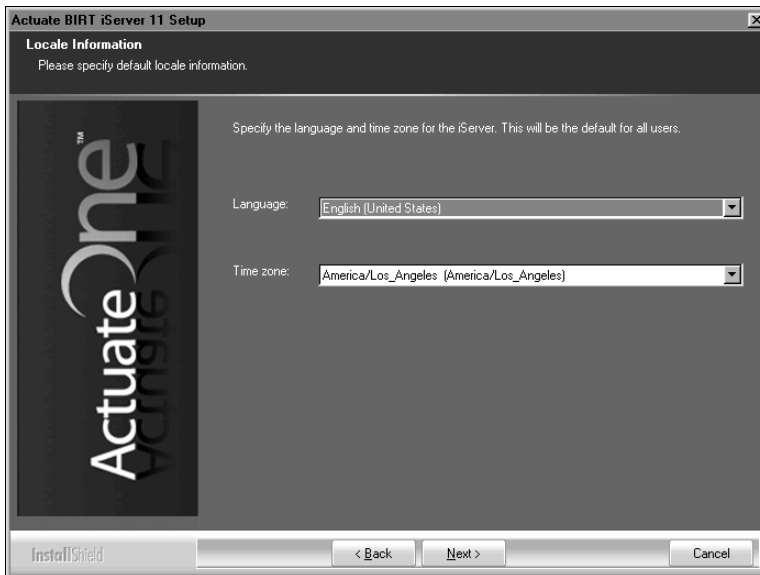


Figure 5-19 Specifying a language

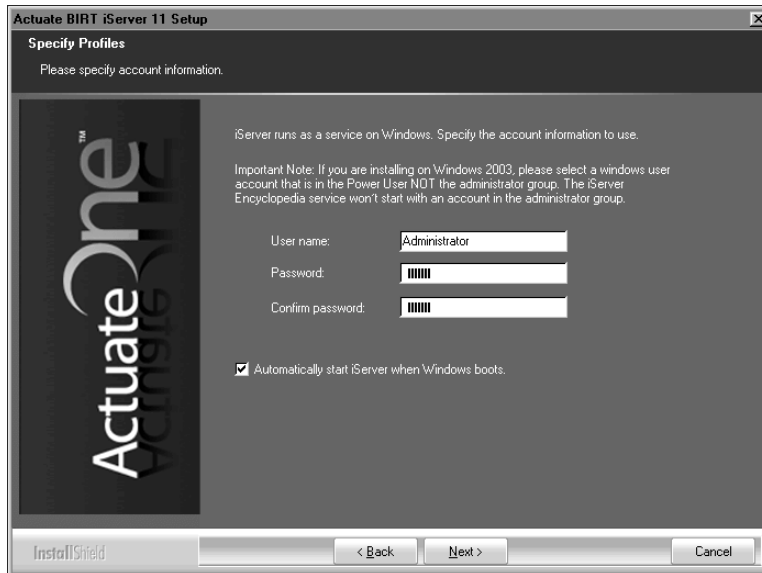


Figure 5-20 Specifying the account for running the iServer service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 5-21.




Figure 5-21 Setting the Windows local security policy

- 10 In iServer Configuration, shown in Figure 5-22, type the cluster node name or IP address for the node. Type a port number or accept the default port number, 8100. The Actuate BIRT iServer 11 service on Windows binds to this port number to listen for requests. Choose Next.

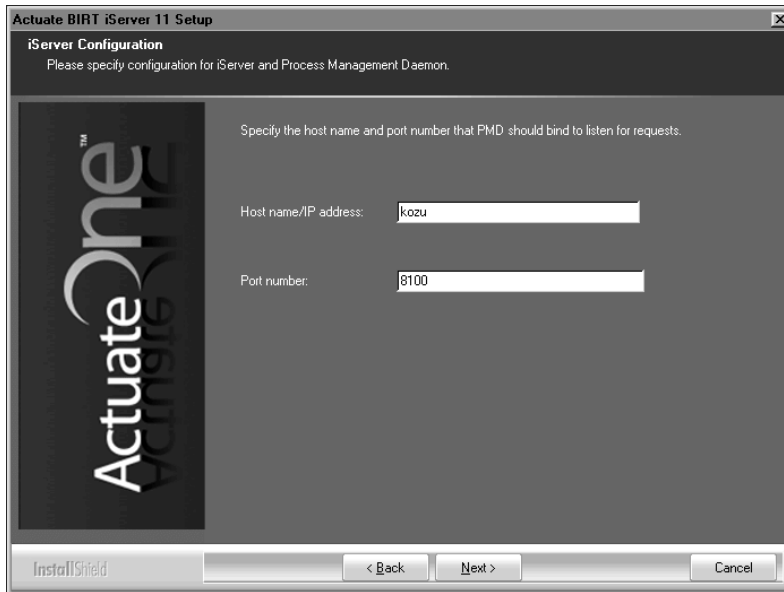


Figure 5-22 Identifying the cluster node and port to bind iServer service

- 11** In System Configuration Password, type a Configuration Console password, as shown in Figure 5-23. Note that the default user name for Management Console is Administrator with no password. Choose Next.

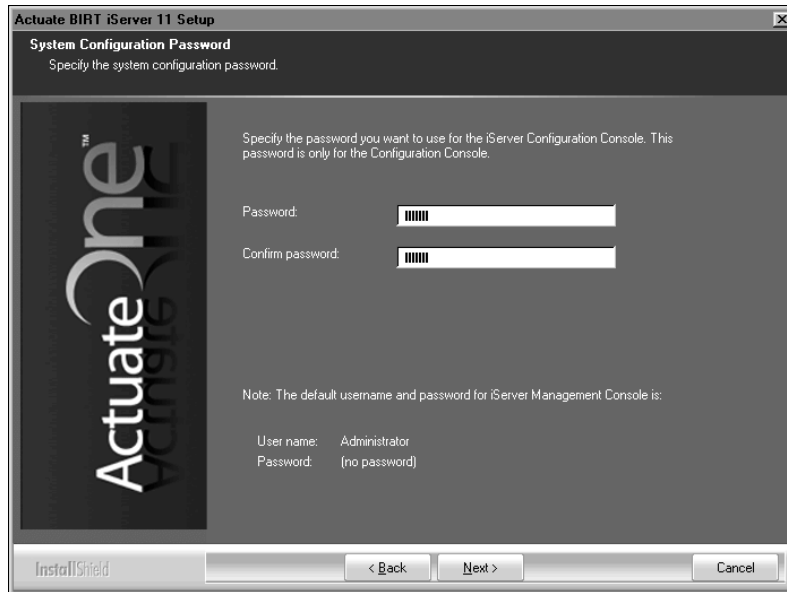


Figure 5-23 Specifying the password for using Configuration Console

12 If you chose to install Management Console in step 5, complete the following steps:

1 Specify the following information, as shown in Figure 5-24. Choose Next:

- Host name and port number for the following items:
 - PMD (Process Management Daemon) Configuration
 - iServer Configuration
- Default volume name

The BIRT iServer System name the Administrator specified on Encyclopedia Metadata Storage and System Name during the wizard-based Typical install to create a standalone iServer, as shown in Figure 2-5 in “Installing a new instance of BIRT iServer Release 11,” in Chapter 2, “Installing BIRT iServer.”

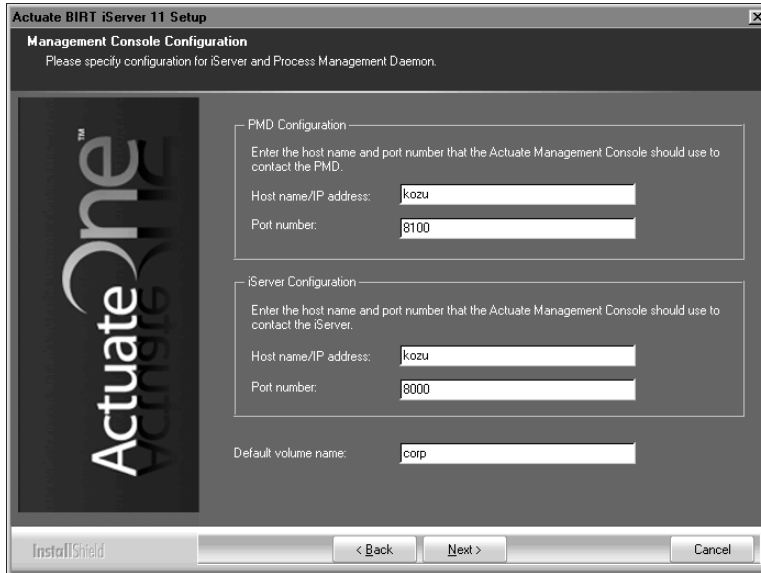


Figure 5-24 Specifying the Management Console Configuration

- 2 In Actuate iServer Application Container, type the port number that the iServer Application Container uses, or accept the default port number, 8900, as shown in Figure 5-25. Choose Next.

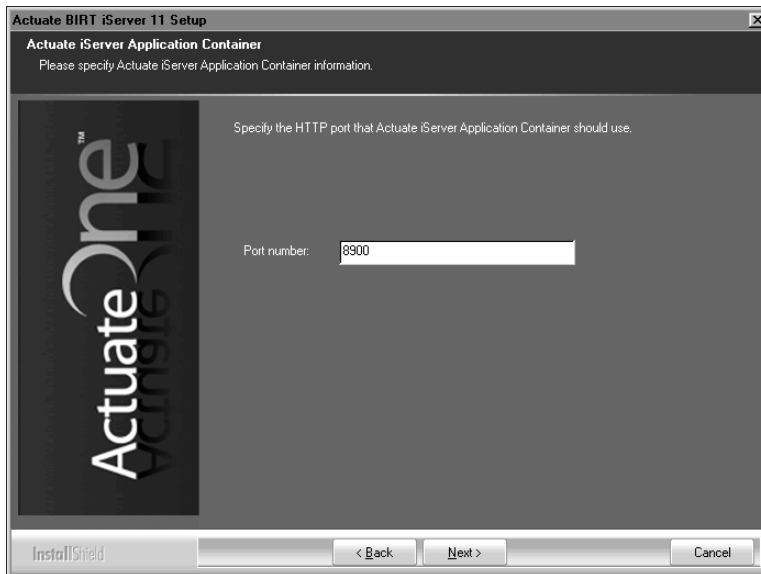


Figure 5-25 Specify the port number for iServer Application Container

- 3 In Context Path, type the context path for Management Console or accept the default path, /acadmin, as shown in Figure 5-26. Choose Next.



Figure 5-26 Specifying the context path

- 4 In Select Program Folder, specify a folder name containing the iServer icons that are used to launch the consoles, or accept the default name, Actuate11, as shown in Figure 5-27. Choose Next.



Figure 5-27 Specifying a program folder

- 13** In Start Copying Files, review the settings shown in Figure 5-28. Choose Next to start copying files.

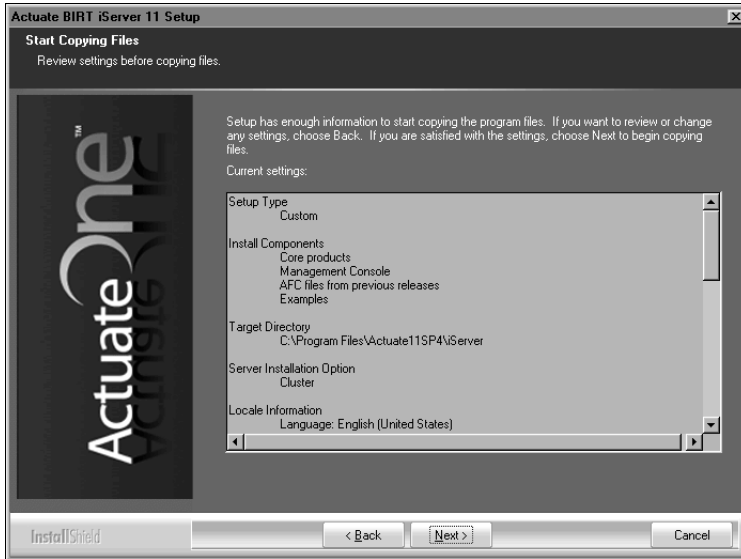


Figure 5-28 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 5-29.



Figure 5-29 Viewing setup status

- 14 When the setup success message appears, as shown in Figure 5-30. Choose Finish to exit the wizard.



Figure 5-30 Exiting the installation wizard

- 15 The installation program prompts you to install the online help from the following location, as shown in Figure 5-31:

<http://www.actuate.com>




Figure 5-31 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows choose Start-> Actuate 11->Update Documentation.

- 16 If you chose to install Management Console in step 4, the installation program installs shortcuts on the desktop, as shown in Figure 5-32.




Figure 5-32 Viewing BIRT iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

- iServer Management Console
Launches Management Console to administer an Encyclopedia volume, set up user accounts, and run reports.
- BIRT iServer 11
Opens Welcome to Actuate BIRT iServer from which you can log in to Information Console to perform tasks, such as accessing folders and viewing reports.

Adding a node to a cluster

This section assumes the Administrator has completed the following tasks:

- Installed iServer as a standalone instance on one machine
- Prepared the node machines for clustering
- Installed iServer as a cluster node instance on a second machine

To add a node to a cluster, the Administrator performs the following tasks:

- Edits `acpmdconfig.xml` on the cluster node to point to the template the cluster node uses
- Starts the iServer system on the cluster node

How to add a newly installed node to a cluster

To add the newly installed node to a cluster, perform the following tasks on the cluster node:

- 1 Choose Start → Settings → Control Panel.

In Control Panel, choose Administrative Tools—Services.

On Services, stop the Actuate 11 BIRT iServer service.

- 2 Open `acpmdconfig.xml`, by default located in `AC_SERVER_HOME/etc`, and perform the following tasks:

- 1 Modify `<AC_TEMPLATE_NAME>` to use the name of the template that you want the node to use. The shared `acserverconfig.xml` contains one template by default. Specify the name that appears in the Name attribute of the `<Template>` element in the shared `acserverconfig.xml`. In this example the server template name is `urup`, as shown in Listing 5-1.

Listing 5-1 Viewing the shared acserverconfig.xml

```
<Templates>
  <Template
    Name="urup"
    PMDPort="8100"
    ActuateBuild="110F120420"
    ActuateVersion="11 Service Pack 4"
    ServerSOAPPort="11100"
    AppContainerPort="8900"
    RequesterRSAPIVolume="corp">
    . . .
  </Template>
</Templates>
```

- 2 Verify that <AC_CONFIG_HOME> points to the shared configuration home directory for the cluster. This is the path you specified for the configuration home location during the install procedure, as shown in Figure 5-16.

acpmdconfig.xml appears as shown in Listing 5-2.

Listing 5-2 Modifying acpmdconfig.xml

```
<!--Server information -->
<Server>
  <Startup>Manual</Startup>
  <AC_TEMPLATE_NAME>urup</AC_TEMPLATE_NAME>
  <AC_DATA_HOME>C:\Actuate\iServer\data</AC_DATA_HOME>
  <AC_CONFIG_HOME>\\urup\11SP4</AC_CONFIG_HOME>
  . . .
</Server>
<!-- Servlet Container information -->
```

- 3 Save acpmdconfig.xml and exit the file.
- 3 Start the Actuate 11 BIRT iServer service on the cluster node.
- 4 On the cluster node, log in to Configuration Console, and start the system, as shown in Figure 5-33. The new cluster node will automatically contact the acserverconfig.xml in the shared configuration home directory to access its template, and then join the cluster.

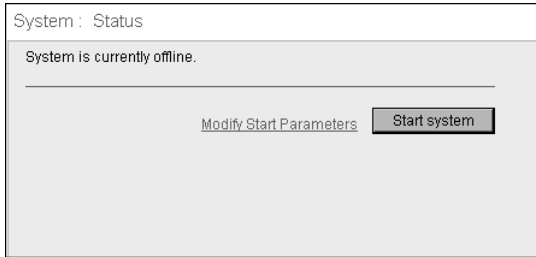


Figure 5-33 Starting the server

- 5 In Advanced view, choose Servers and confirm that the servers are online, and that the Factory and View services are enabled, as shown in Figure 5-34.

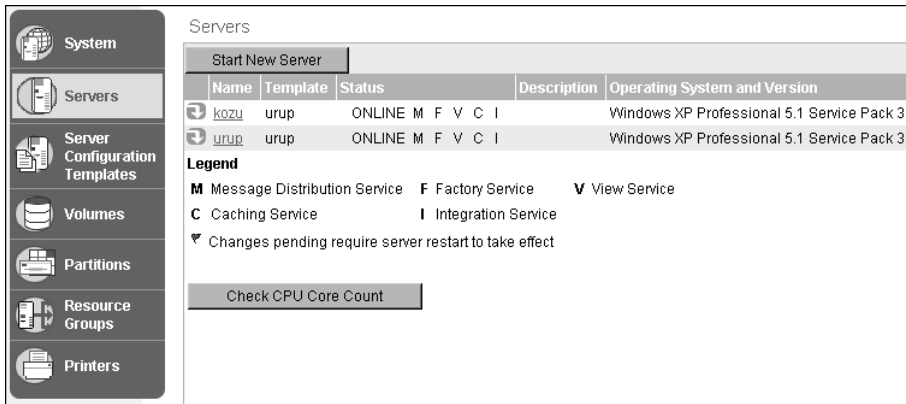


Figure 5-34 Viewing a cluster composed of urup and kozu

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, 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 a 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 ZIP 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 (32-bit) 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>

In Windows, the deployment script automatically installs the following prerequisite **Microsoft Visual C++ Libraries**. These libraries ship with the iServer distribution package in `AC_SERVER_HOME/prerequisites`. You can also download these resources from Microsoft, and manually install them.

In a 32-bit environment:

- Microsoft Visual C++ 2005 SP1 Redistributable Package (x86) - `vc redistrib_vs2005_x86.exe`
- Microsoft Visual C++ 2008 Redistributable Packages (x86) - `vc redistrib_x86.exe`

In a 64-bit environment:

- Microsoft Visual C++ 2005 Redistributable Package (x64) - `vc redistrib_x64.exe`.

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.

The iServer Release 11 distribution package also contains the portmapper components required to run an Actuate e.Report, but the portmapper must be started manually using `iServer/bin/portinst.exe`. If the Actuate e.Reports option is enabled by a license key, iServer looks for the portmapper when it starts. If the portmapper is unavailable, iServer shuts down gracefully after a timeout, writing an appropriate error message to the logs. If the Actuate e.Reports option is not enabled by the license key, an attempt to run an e.Report fails gracefully. iServer does not look for the presence of the portmapper.

In an iServer Release 11 configuration, Actuate recommends storing iServer data in a directory located outside the iServer run-time environment. In a default Windows setup performed using the wizard-based install program, the iServer run-time environment installs in the following directory:

```
C:\Program Files\Actuate11SP4\iServer
```

The data installs in the following directory:

```
C:\Actuate\iServer\data
```

In a cloud deployment, Actuate recommends installing in an alternative directory. For example, in Windows, extract the iServer distribution package files to the following directory:

```
C:\Actuate11
```

Extracting the iServer files installs the run-time environment in the following directory:

```
C:\Actuate11\iServer
```

Running the setup script installs the data in the following directory:

```
C:\Actuate11\iServer\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 `C:\Program Files` or `C:\Program Files(x86)`, such as `C:\Actuate11`.
- 2 Using Winzip or another file extraction tool, extract the contents of `ActuateBIRTiServer.zip` to the folder created in the previous step, as shown in Figure 6-1.

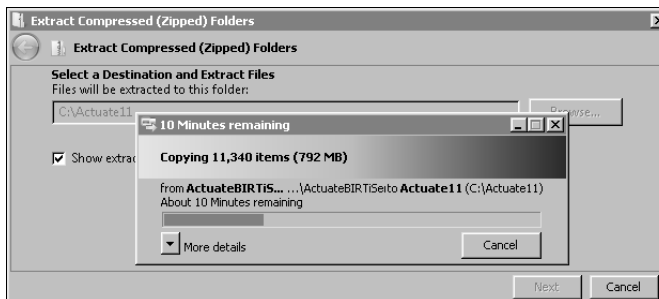


Figure 6-1 Extracting iServer distribution package contents

- 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.bat` script automatically detects the location of `AC_SERVER_HOME` and `AC_JAVA_HOME`. 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 to set the values for these properties manually.

To specify the location of `AC_SERVER_HOME` and `AC_JAVA_HOME` in the `setupiServer.bat` script, perform the following tasks:

- 1 Using a text editor, open the script, `setupiServer.bat`, 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.bat` script

```
@echo off
REM Script that sets up the iServer and starts the iServer

cls
Title Actuate BIRT iServer 11 setup and start script for
  evaluation
...
set AC_SERVER_HOME=C:\Actuate11\iServer
set AC_JAVA_HOME=C:\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 into 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

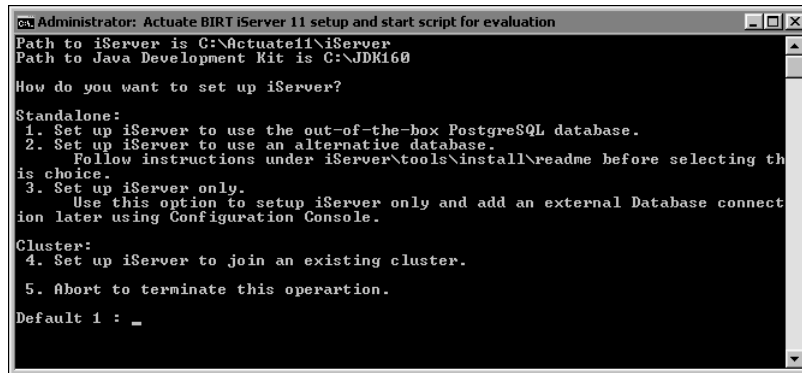
For more information about creating the iServer system and Encyclopedia volume schemas and iserver user in an alternative database, see Chapter 3, “Installing BIRT iServer using an alternative database.”

- 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 From the Windows menu, choose Start→Run.
- 2 Open a command prompt by typing cmd, and choose OK.
- 3 Navigate to the folder where you extracted the iServer package, such as C:\Actuate11.
- 4 Type setupiServer, and press Enter to execute the installation script.
- 5 The script prompts you to choose one of the following stand-alone or cluster options in setting up iServer, as shown in Figure 6-2:
 - Standalone:
 - 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.



```
Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\JDK160
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 iServer\tools\install\readme before selecting th
is choice.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connect
ion later using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.
 5. Abort to terminate this operation.

Default 1 : _
```

Figure 6-2 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.zip 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 6-3.

The script performs the following tasks, as shown in Figure 6-3:

- Installs the prerequisite C++ runtime components
- 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

```
Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\JDK160

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 iServer\tools\install\readme before selecting th
   is choice.
3. Set up iServer only.
   Use this option to setup iServer only and add an external Database connect
   ion later using Configuration Console.

Cluster:
4. Set up iServer to join an existing cluster.
5. Abort to terminate this operation.

Default 1 :

Installing pre-requisite C++ runtime components. This will take a few seconds ..
.
Configuring deployed iServer files ...
Detailed information will be logged to file "C:\Actuate11\iServer\data/server/lo
g/KeysFileGenerator.URUP.2012-06-06_18_43_01_PDI.zg.log"
Generating keys file "C:\Actuate11\iServer\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 ...
  1 file(s) copied.
  1 file(s) moved.
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.

Press any key to close this window
_
```

Figure 6-3 Using the setup script to install option 1

When the script starts Actuate BIRT iServer 11, a command prompt opens to serve as the console from which iServer can be shut down, as shown in Figure 6-4.

```
Actuate BIRT iServer 11
Starting Actuate BIRT iServer 11 on console.

Actuate BIRT iServer 11 console commands:
s - shutdown server
Enter command: _
```

Figure 6-4 Viewing Actuate BIRT iServer 11 console

When the script starts PostgreSQL for Actuate BIRT iServer, a command prompt opens to serve as the console from which PostgreSQL can be shut down, as shown in Figure 6-5.



Figure 6-5 Viewing PostgreSQL for Actuate iServer console

For information about starting and stopping Actuate BIRT iServer 11 from a command prompt, see “How to stop and start Actuate BIRT iServer,” later in this section. For information about starting and stopping PostgreSQL for Actuate iServer, see “How to stop and start PostgreSQL for Actuate iServer,” later in this section.

If a Windows Security Alert appears indicating that the firewall is blocking access to Actuate BIRT iServer programs, perform the following tasks:

- 1 In Allow Actuate BIRT iServer to communicate on these networks, for example, select Private networks, such as my home or work network, then choose Allow access, as shown in Figure 6-6.
- 2 Repeat this step for other Windows Security Alerts, such as Java Platform SE binary.



Figure 6-6 Allowing firewall access to Actuate BIRT iServer

When the script finishes running, press any key to close the command prompt running setupiServer.bat script, as shown in Figure 6-3. If the script exit returns to the command line, type exit to close the command prompt. Be careful to not close the command prompt running the Actuate BIRT iServer 11 console or PostgreSQL by mistake.

Accessing Information, Management, and Configuration Consoles

In Windows, the deployment program installs shortcuts to the folder where you extracted the BIRT iServer deployment package. These shortcuts provide access to the following iServer consoles:

- BIRT iServer 11 Information Console
Launches the Information Console to access folders and view designs and documents
- BIRT iServer 11 Management Console
Launches Management Console to set up user accounts and schedule or run a design

After the script finishes running, a browser opens displaying Welcome to Actuate BIRT iServer, as shown in Figure 6-7.

Log in to Information Console by choosing the shortcut. Alternatively, open a browser manually and enter the following URL, as shown in Figure 6-7:

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

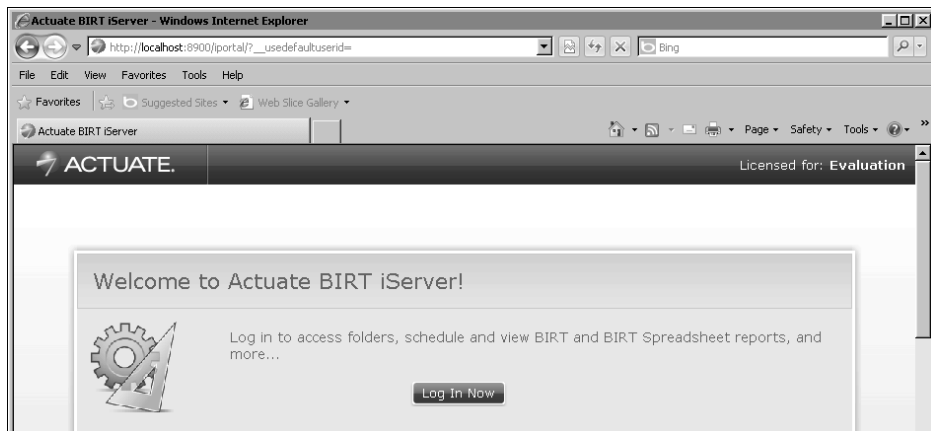


Figure 6-7 Viewing Welcome to Actuate BIRT iServer

Log in to Management Console by choosing the shortcut. Alternatively, open a browser manually and enter the following URL, as shown in Figure 6-8:

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

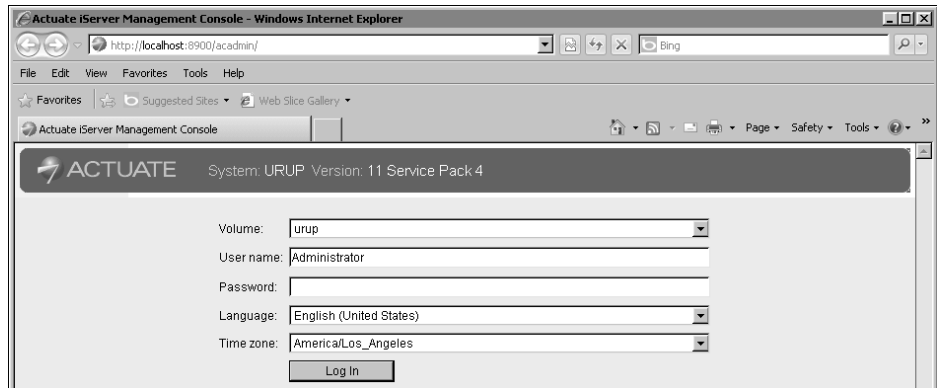


Figure 6-8 Logging in to Management Console

Log in to Configuration Console to 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 6-9:

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

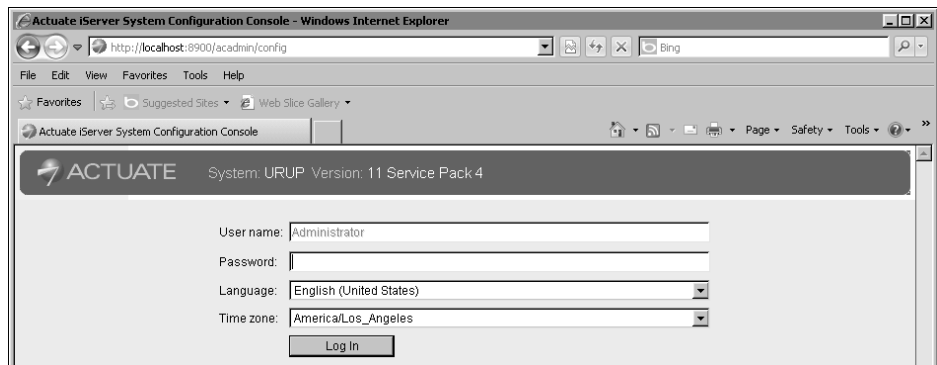


Figure 6-9 Logging in to Configuration Console

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 Actuate BIRT iServer

To stop Actuate BIRT iServer, perform the following tasks:

- 1 From the Actuate BIRT iServer console, type `s` and press Enter.

The message, Shutting down the server appears, as shown in Figure 6-10.

- 2 Wait for the shutdown process to complete.

After shutting down iServer, the command prompt closes automatically.



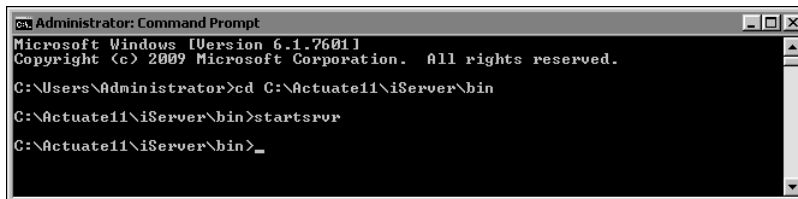
```
Actuate BIRT iServer 11
Starting Actuate BIRT iServer 11 on console.
Actuate BIRT iServer 11 console commands:
s - shutdown server
Enter command: s
Shutting down the server ...
_
```

Figure 6-10 Shutting down Actuate BIRT iServer 11 from the console

To restart 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, as shown in Figure 6-11:

```
startsrvr
```



```
Administrator: Command Prompt
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.
C:\Users\Administrator>cd C:\Actuate11\iServer\bin
C:\Actuate11\iServer\bin>startsrvr
C:\Actuate11\iServer\bin>_
```

Figure 6-11 Starting Actuate BIRT iServer 11

A new Actuate BIRT iServer 11 console opens, as shown earlier in Figure 6-4. This window must remain open, since the iServer service runs from this window.

How to stop and start PostgreSQL for Actuate iServer

To shut down PostgreSQL for Actuate iServer from the console that opens during the installation process, shown in Figure 6-5, type Ctrl-C. Wait for the shutdown process to complete. After shutting down PostgreSQL, the command prompt closes automatically.

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, as shown in Figure 6-12:

```
startpostgresql
```




Figure 6-12 Starting PostgreSQL for Actuate iServer

A new command window opens, entitled PostgreSQL for Actuate iServer, as shown in Figure 6-13. This window must remain open, since the iServer service runs from this window.



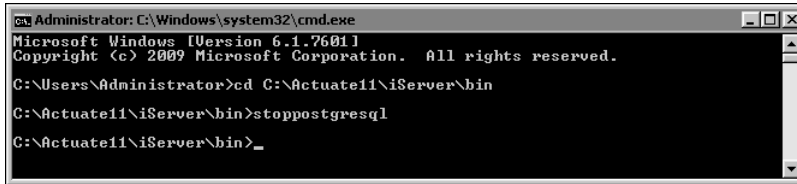
Figure 6-13 Viewing PostgreSQL for Actuate iServer command window

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, as shown in Figure 6-14:

```
stoppostgresql
```

The PostgreSQL for Actuate iServer command window closes.



```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>cd C:\Actuate11\iServer\bin
C:\Actuate11\iServer\bin>stoppostgresql
C:\Actuate11\iServer\bin>_
```

Figure 6-14 Stopping PostgreSQL for Actuate iServer

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.zip 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 Using Windows Explorer or a command prompt, navigate to `AC_SERVER_HOME\tools\install`. For example:
`C:\Actuate11\iServer\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 C:/Actuate11/iServer. Use forward slashes in the path specification.
- 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: D:/ActuateBIRTiServer/iServer
AC_SERVER_HOME=C:/Actuate11/iServer

#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
...

```

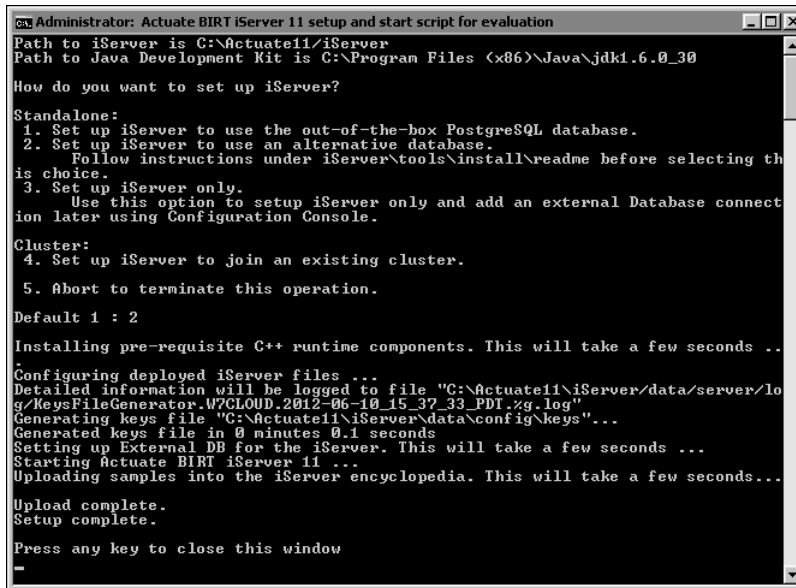
How to set up iServer to use an alternative database

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, as shown in Figure 6-15.

The script performs the following tasks, as shown in Figure 6-15:

- Installs the prerequisite C++ runtime components
- 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



```
Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\Program Files (x86)\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 iServer\tools\install\readme before selecting the
    is choice.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connect
    ion later using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.
 5. Abort to terminate this operation.

Default 1 : 2

Installing pre-requisite C++ runtime components. This will take a few seconds ..
.
Configuring deployed iServer files ...
Detailed information will be logged to file "C:\Actuate11\iServer\data\server\log
g\KeysFileGenerator.W7CLOUD.2012-06-10_15_37_33_PDI.zg.log"
Generating keys file "C:\Actuate11\iServer\data\config\keys"...
Generated keys file in 0 minutes 0.1 seconds
Setting up External DB for the iServer. This will take a few seconds ...
Starting Actuate BIRT iServer 11 ...
Uploading samples into the iServer encyclopedia. This will take a few seconds...

Upload complete.
Setup complete.

Press any key to close this window
-
```

Figure 6-15 Completing the iServer alternative database setup

When the script starts iServer, a new command prompt opens to serve as the Actuate BIRT iServer console from which iServer can be shut down, as shown in Figure 6-16.

When the script finishes running, press any key to close the command prompt running setupiServer.bat script, as shown in Figure 6-15. If the script exit returns to the command line, type exit to close the command prompt. Be careful to not close the command prompt running the Actuate BIRT iServer console by mistake.




Figure 6-16 Viewing Actuate BIRT iServer 11 console

After the script finishes running, a browser opens displaying Welcome to Actuate BIRT iServer, as shown earlier in Figure 6-7. 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 roles
 - Data partition and Encyclopedia volume

Setting up iServer only using the cloud deployment package

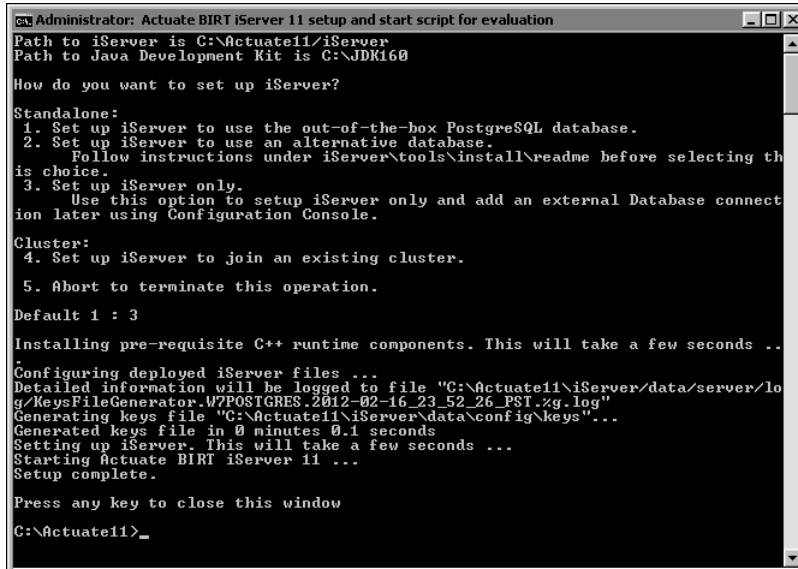
After extracting the contents of ActuateBIRTiServer.zip and running the setup script as described in “How to extract the contents of the iServer distribution package,” and “How to run the setup script,” earlier in this chapter, perform the following tasks.

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 6-17:

- Installs the prerequisite C++ runtime components
- Sets up the iServer deployment files, including log and security keys files
- Sets up and starts BIRT iServer 11



```
Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\JDK160

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 iServer\tools\install\readme before selecting th
is choice.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connect
ion later using Configuration Console.

Cluster:
 4. Set up iServer to join an existing cluster.
 5. Abort to terminate this operation.

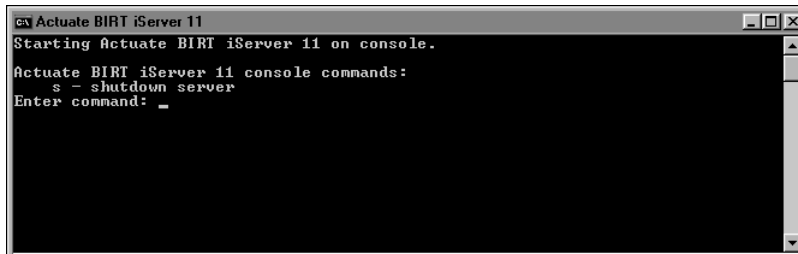
Default 1 : 3

Installing pre-requisite C++ runtime components. This will take a few seconds ..
Configuring deployed iServer files ...
Detailed information will be logged to file "C:\Actuate11\iServer\data\server\log
g\KeyFileGenerator.WPOSTGRES.2012-02-16_23_52_26_PST.zg.log"
Generating keys file "C:\Actuate11\iServer\data\config\keys"...
Generated keys file in 0 minutes 0.1 seconds
Setting up iServer. This will take a few seconds ...
Starting Actuate BIRT iServer 11 ...
Setup complete.

Press any key to close this window
C:\Actuate11>
```

Figure 6-17 Completing the iServer only setup

When the script starts iServer, a new command prompt opens to function as the Actuate BIRT iServer console from which iServer can be shut down, as shown in Figure 6-18.



```
Actuate BIRT iServer 11
Starting Actuate BIRT iServer 11 on console.

Actuate BIRT iServer 11 console commands:
s - shutdown server
Enter command: _
```

Figure 6-18 Viewing Actuate BIRT iServer 11 console

When the script finishes running, a browser opens displaying Welcome to Actuate Information Console, as shown earlier in Figure 6-7, but no Encyclopedia volume is available. You cannot log into Information Console until an Encyclopedia volume exists.

- 2 After the script finishes running, press any key to close the command prompt running setupiServer.bat script, as shown in Figure 6-17. If the script exit returns to the command line, type exit to close the command prompt.
- 3 After installing option 3, stop and restart iServer by performing the following tasks:

- 1 In Actuate BIRT iServer console, type s to shutdown BIRT iServer.
- 2 Open a new command prompt and navigate to AC_SERVER_HOME/bin.
- 3 Type the following command and press Enter, as shown in Figure 6-19:

```
startsrvr
```




Figure 6-19 Starting Actuate BIRT iServer 11

A new Actuate BIRT iServer 11 console opens, as shown earlier in Figure 6-18. This window must remain open, since the iServer service runs from this window. For more information about stopping and starting iServer, see “How to stop and start Actuate BIRT iServer,” earlier in this chapter.

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 roles
- 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 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 6-20. The administrator does not have a password yet, so leave Password blank.

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

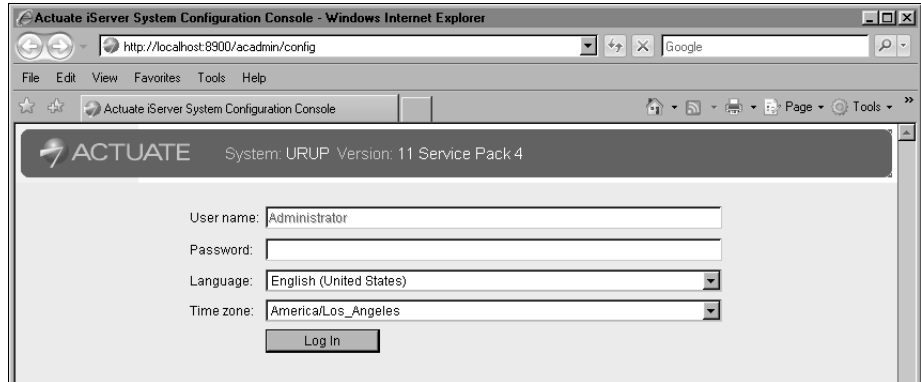


Figure 6-20 Logging in to Configuration Console

- 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 6-21.

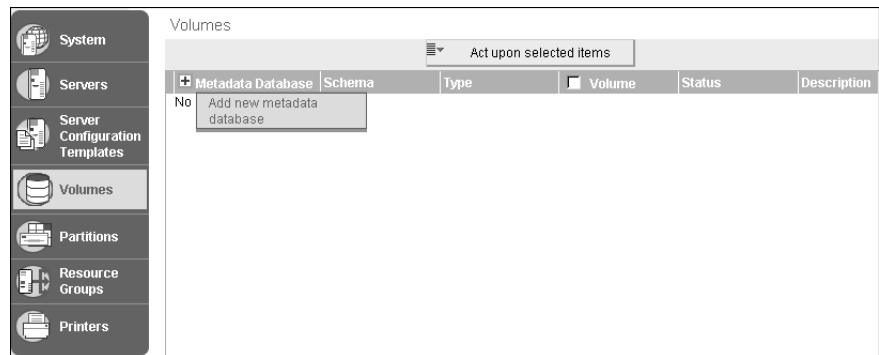


Figure 6-21 Adding a new metadata database

- 2 On New Metadata Database, perform the following tasks, as shown in Figure 6-22:

- 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 6-22 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 6-23. Choose OK.




Figure 6-23 Viewing Connection tested successfully! message

On New Metadata Database, choose OK.

- 9 On Volumes, the ActuatePostgreSQL_MetadataDatabase appears, as shown in Figure 6-24.

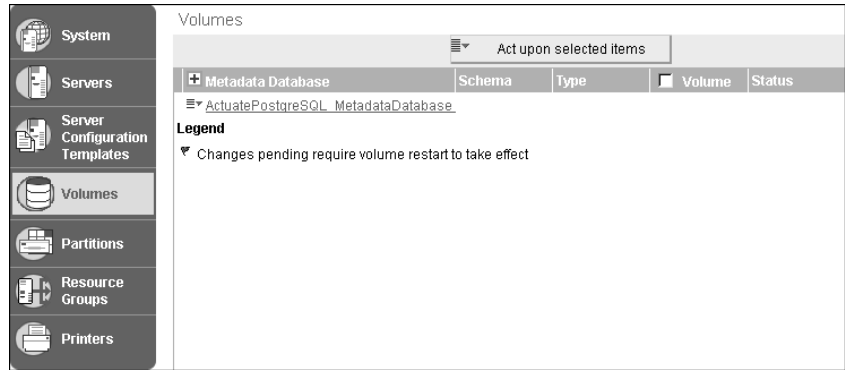


Figure 6-24 Viewing the metadata database

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 6-25.

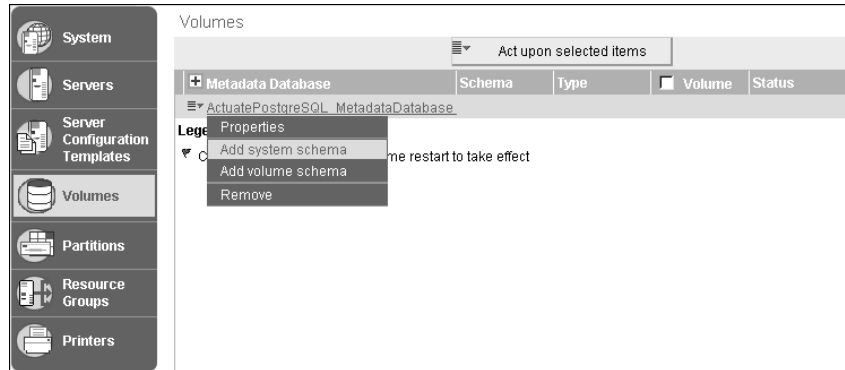


Figure 6-25 Choosing Add system schema

- 2 On New System Schema, perform the following tasks, as shown in Figure 6-26:
 - 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.

Volumes > New System Schema

Schema

This new schema will replace existing BIRT iServer system schema. Please restart your BIRT iServer system immediately.

Metadata Database: ActuatePostgreSQL_MetadataDatabase

Schema type: System

Schema name: ac_corp_system *

Database schema name: ac_corp_system *

Database schema password: •••••• *

* These fields are required and cannot be left blank

Test OK Cancel Apply

Figure 6-26 Adding a new system schema

4 Choose Test. If successful, choose OK.

5 On New System Schema, 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 6-27.




Figure 6-27 Viewing message to restart iServer system

Choose OK.

6 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. Figure 2-13 shows Add volume schema. The metadata database is ActuatePostgreSQL_MetadataDatabase, as shown in Figure 6-28.




Figure 6-28 Choosing Add volume schema

- 2 On New Volume Schema, as shown in Figure 6-29, 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.

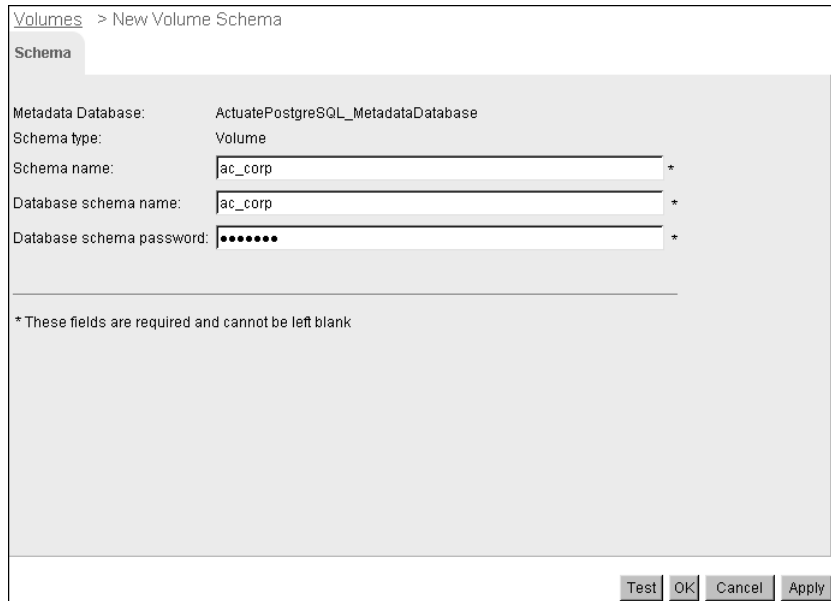


Figure 6-29 Adding a new Encyclopedia volume schema

- 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 6-30. Choose OK.




Figure 6-30 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 6-31.

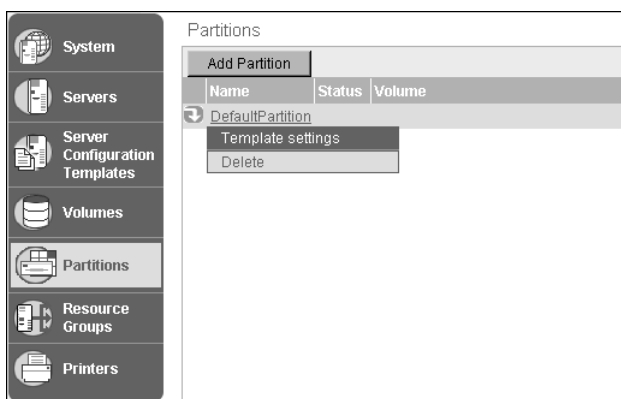


Figure 6-31 Deleting DefaultPartition

Choose OK to confirm the deletion of DefaultPartition, as shown in Figure 6-32.




Figure 6-32 Confirming to delete 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:
`C:\Actuate11\iServer\data\ac_corp_partition`
 - 2 From the 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 6-33.

Choose OK.

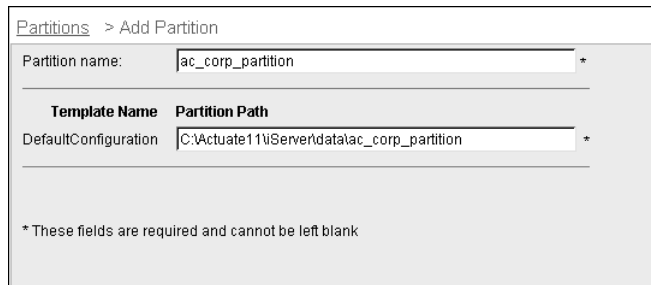


Figure 6-33 Adding a new data partition

- 6 In Partitions, choose the new partition, `ac_corp_partition`, from the list of partitions.
- 7 In Server Settings, choose Test, as shown in Figure 6-34.
Choose OK.



Figure 6-34 Testing a new data partition

If the test succeeds, a success message appears, as shown in message in Figure 6-35. Choose OK.



Figure 6-35 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 6-36.




Figure 6-36 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 6-37.

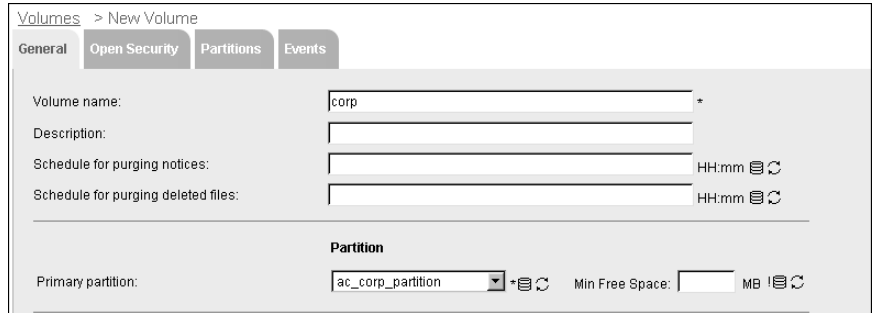


Figure 6-37 Specifying general volume properties

- 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 6-38.

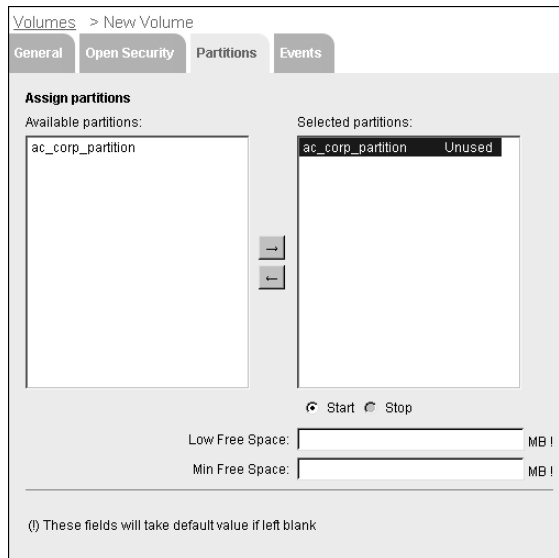


Figure 6-38 Starting the partition

Choose OK.

- 4 In Volumes, point to the arrow next to the new volume name, and choose Take online, as shown in Figure 6-39.




Figure 6-39 Taking a volume online

- In Volumes, check that the status of the new volume changes to ONLINE, as shown in Figure 6-40.

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




Figure 6-40 Viewing the online volume

- Log in to Management Console to inspect the new volume by choosing the shortcut. Alternatively, open a browser manually and enter the following URL, as shown in Figure 6-41:

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

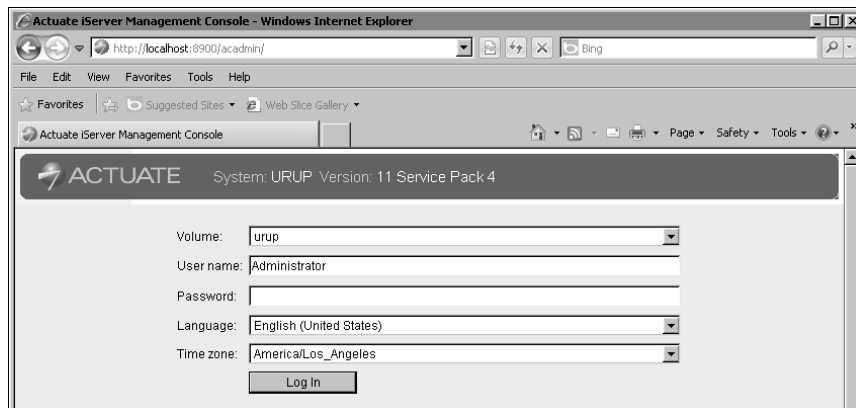


Figure 6-41 Logging in to Management Console

- In Files and Folders, the default Encyclopedia volume appears with an empty Resources folder, as shown in Figure 6-42.



Figure 6-42 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 home directory, which all nodes in the cluster access. This section refers to the machine containing the shared configuration home directory as node1. In “How to set up iServer to join an existing cluster,” later in this chapter, the Administrator installs a cluster node. The cluster node accesses the shared resources on node1. This section refers to the machine on which the administrator installs a cluster node 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 folder sharing and firewall settings in the network environment
- 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 instructions provide a basic example of the operations required to configure network sharing and firewall settings in the Windows environment. 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.

How to share the encyc and config folders

Cluster nodes must have read-write access to AC_DATA_HOME\config and AC_DATA_HOME\encyc on node1. To give a cluster node read-write access to the folders and files in the \encyc folder and to the files in the \config folder, perform the following tasks on node1:

- 1 Using Windows Explorer on node1, right-click the folder, AC_DATA_HOME\encyc. Choose Properties. On encyc Properties, choose Sharing. Sharing appears as shown in Figure 6-43.

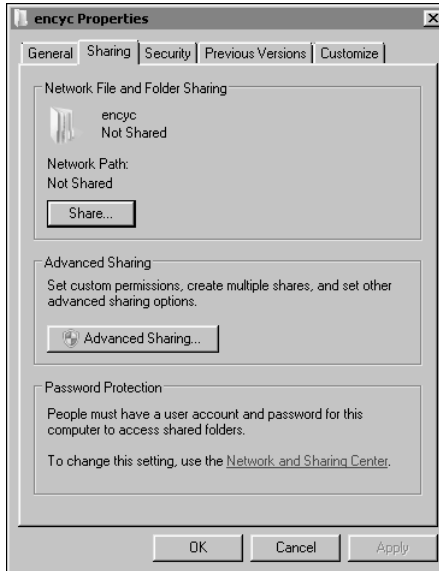


Figure 6-43 Sharing the Encyclopedia volume

On encyc Properties, choose Advanced Sharing.

- 2 On Advanced Sharing, select Share this folder, as shown in Figure 6-44.

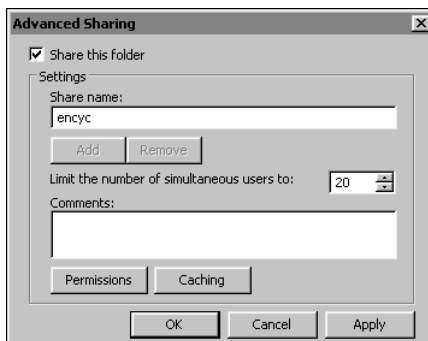


Figure 6-44 Selecting Share this folder

On Advanced Sharing, choose Permissions.

- 3 On Permissions for encyc—Share Permissions, select the user who installed BIRT iServer. Select Change and Read to allow this user these permissions, as shown in Figure 6-45. Make sure that all settings conform to the security policies in force for the environment.

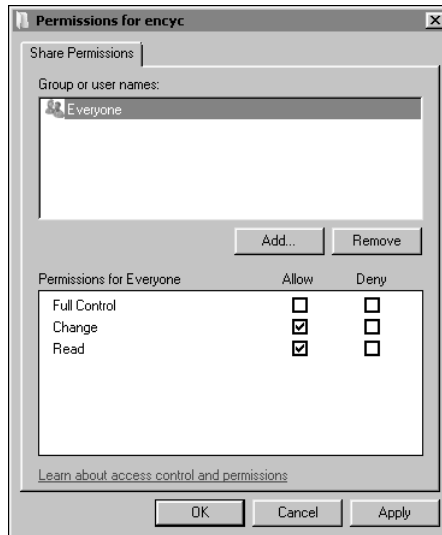


Figure 6-45 Selecting Change and Read permission

Choose OK.

On Advanced Sharing, choose OK.

On encyc Properties, choose Close.

- 4 Repeat steps 1 through 3 for the \config folder. Make sure that all settings conform to the security policies in force for the environment.

Close Windows Explorer.

How to turn off the Windows firewall

Perform the following steps on node1 and node2:

- 1 Choose Start→Control Panel→System and Security→Windows Firewall.
- 2 On Windows Firewall, choose Turn Windows Firewall on or off. Make sure that all settings conform to the security policies in force for the environment.
- 3 For example, on Customize Settings, in Home or work (private) network location settings, choose Turn off Windows Firewall, as shown in Figure 6-46.

Choose OK.

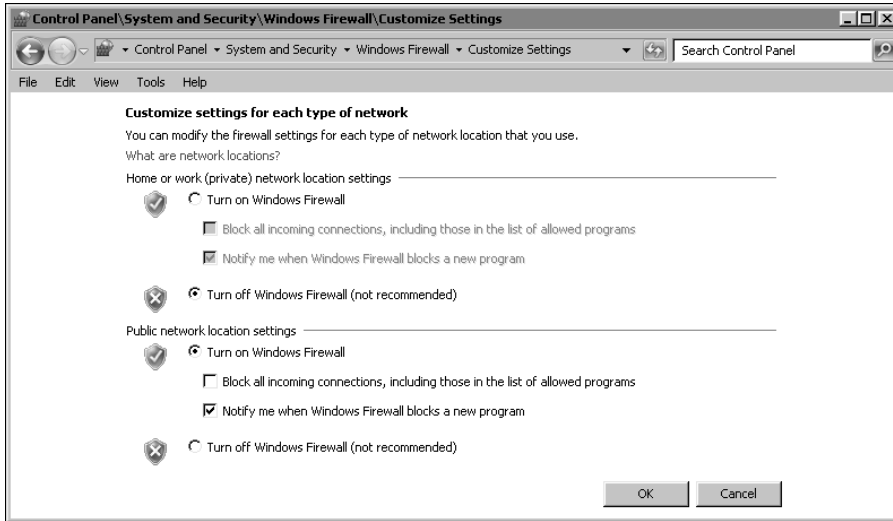


Figure 6-46 Turning off the Home or Work network location firewall
Close Windows Firewall.

How to display a computer's IP address

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

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

```
ipconfig /all
```

Press Enter. The host name appears, as shown in Figure 6-47. In this example, the host name for node1 is urup.

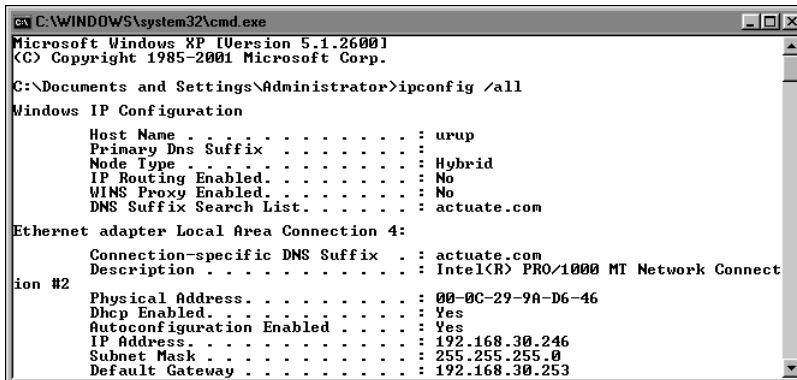


Figure 6-47 Displaying the host name

- 3 Write the host name and IP address in the spaces provided in Table 6-1.

Table 6-1 Host names and IP addresses of computers to be clustered

iServer	Host name	IP address
Node1	urup	192.168.30.246
Node2	kozu	192.168.30.233

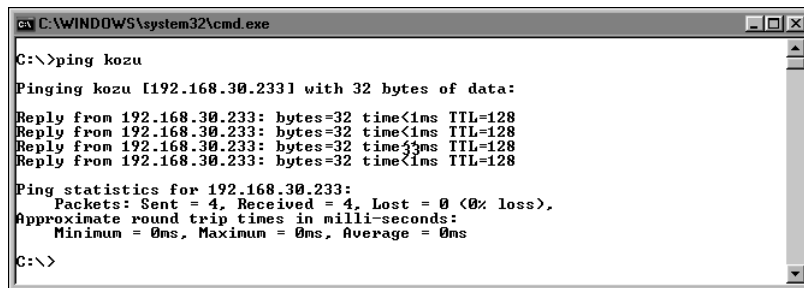
How to test the connection between computers

Perform the following steps on both computers:

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

```
ping kozu
```

Press Enter. If your computer reaches the other computer, Command Prompt displays a series of replies, as shown in Figure 6-48.



```
C:\WINDOWS\system32\cmd.exe
C:\>ping kozu
Pinging kozu [192.168.30.233] with 32 bytes of data:
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Reply from 192.168.30.233: bytes=32 time<1ms TTL=128
Ping statistics for 192.168.30.233:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\>
```

Figure 6-48 Receiving a reply to a ping command

- 2 Close Command Prompt.

How to prepare node1 for clustering

On node1, or `urup` in this example, perform the following tasks:

- 1 Log in to Configuration Console by opening a browser and entering the following URL, as shown in Figure 6-49. The administrator does not have a password yet, so leave Password blank.

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

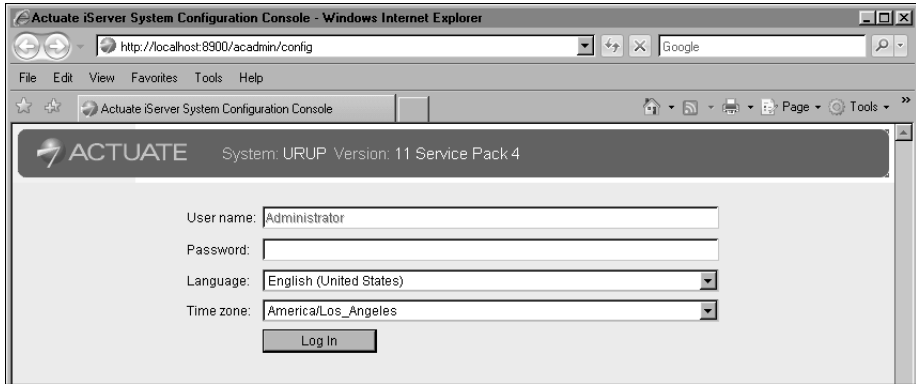


Figure 6-49 Logging in to Configuration Console

- 2 On the simple view, choose Advanced view. From the side menu, choose Volumes. On Volumes, point to the icon next to the default volume, urup in this example, and choose Take offline, as shown in Figure 6-50.

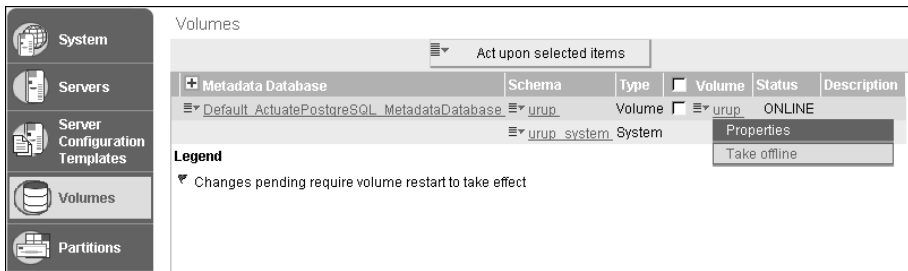


Figure 6-50 Taking a volume offline

In Volume Offline Grace Period, choose OK.

From the side menu, choose Partitions.

- 3 On Partitions, point to the icon next to DefaultPartition and choose Template settings, as shown in Figure 6-51.




Figure 6-51 Choosing Template settings

On Template Settings, choose Change. In Partition Path, type the path to the encyc folder, using UNC format. For example, if the hostname of node1 is urup, type:

```
\\urup\encyc
```

Template Settings appears as shown in Figure 6-52.

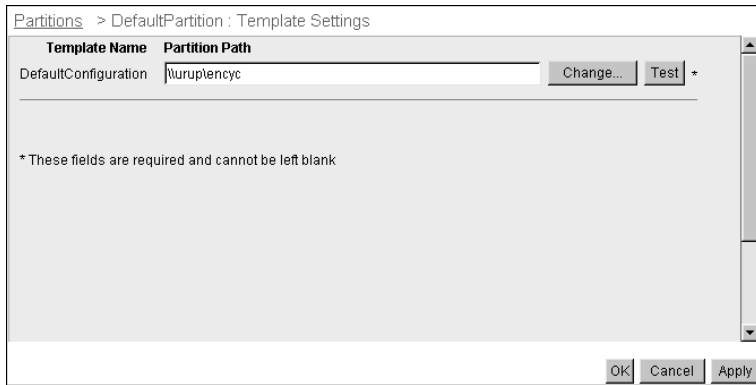


Figure 6-52 Typing the partition path

To verify that iServer can access the encyc folder, choose Test.

A message appears, stating that the test was successful, as shown in Figure 6-53. Choose OK.



Figure 6-53 Verifying a successful partition path test

On Template Settings, choose OK.

4 Log out of Configuration Console.

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.zip 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 using Universal Naming Convention (UNC) Format, as shown in Figure 6-54. For example, type:

```
\\urup\config
```

Press Enter.

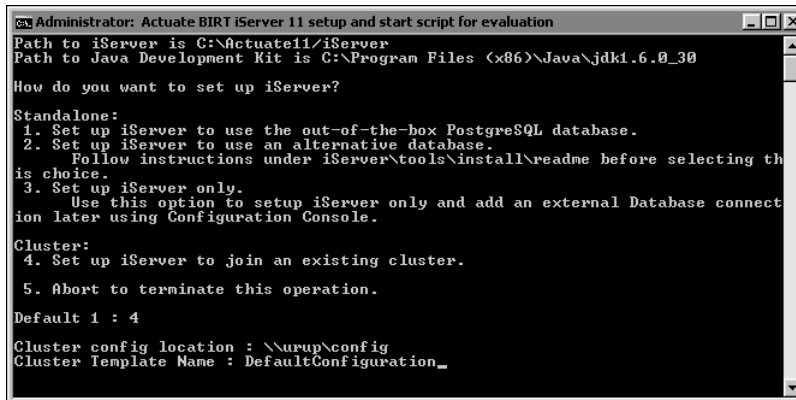
In an iServer cluster, the configuration home location, AC_CONFIG_HOME, is the shared directory that contains the acserverconfig.xml, acserverlicense.xml, and other related files. In a Release 11 Service Pack 4 installation, by default, these files are in AC_DATA_HOME\config. If the shared folder is in this location on a server named urup, then the administrator specifies \\urup\config as the path.

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 6-54. For example, type:

```
DefaultConfiguration
```

The default template name in the shared acserverconfig.xml in a cloud-based stand-alone iServer install is DefaultConfiguration.

Press Enter.



```
Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\Program Files (x86)\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 iServer\tools\install\readme before selecting th
is choice.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connect
ion 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 : \\urup\config
Cluster Template Name : DefaultConfiguration_
```

Figure 6-54 Specifying the cluster option, configuration home location, and server template name

The script performs the following tasks, as shown in Figure 6-55:

- Installs the prerequisite C++ runtime components
- 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
- Sets up and starts BIRT iServer 11


```

Administrator: Actuate BIRT iServer 11 setup and start script for evaluation
Path to iServer is C:\Actuate11\iServer
Path to Java Development Kit is C:\Program Files (x86)\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 iServer\tools\install\readme before selecting th
is choice.
 3. Set up iServer only.
    Use this option to setup iServer only and add an external Database connect
ion 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 : \\usup\config
Cluster Template Name : DefaultConfiguration

Installing pre-requisite C++ runtime components. This will take a few seconds ..
.
Configuring deployed iServer files ...
Detailed information will be logged to file "C:\Actuate11\iServer\data\server\log
g\KeysFileGenerator.KOZU.2012-06-08_18_19_29_PDT.%g.log"
Generating keys file "C:\Actuate11\iServer\data\config\keys"...
Generated keys file in 0 minutes 0.2 seconds
Setting up iServer to join cluster...
Starting Actuate BIRT iServer 11 ...
Setup complete.

Press any key to close this window
=

```

Figure 6-55 Completing the iServer cluster node setup

When the script starts iServer, a new command prompt opens to serve as the Actuate BIRT iServer console from which iServer can be shut down, as shown in Figure 6-56.

```

Administrator: Actuate BIRT iServer 11
Starting Actuate BIRT iServer 11 on console.

Actuate BIRT iServer 11 console commands:
s - shutdown server
Enter command: _

```

Figure 6-56 Viewing Actuate BIRT iServer 11 console

When the script finishes running, press any key to close the command prompt running the setupiServer.bat script, as shown in Figure 6-55. If the script exit returns to the command line, type exit to close the command prompt. Be careful to not close the command prompt running the Actuate BIRT iServer console by mistake.

In Windows, the deployment program installs shortcuts to the folder where you extracted the BIRT iServer deployment package. These shortcuts provide access to the following iServer consoles:

- BIRT iServer 11 Information Console
 - Launches Information Console for viewing report documents

- BIRT iServer 11 Management Console
Launches Management Console for setting up user accounts and scheduling or running a design

After the script finishes running, a browser opens displaying Welcome to Actuate BIRT iServer, as shown earlier in Figure 6-7. Log in to Information Console to perform tasks such as accessing folders and viewing designs and documents.

- 3 Log in to Configuration Console. On the 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 6-57.

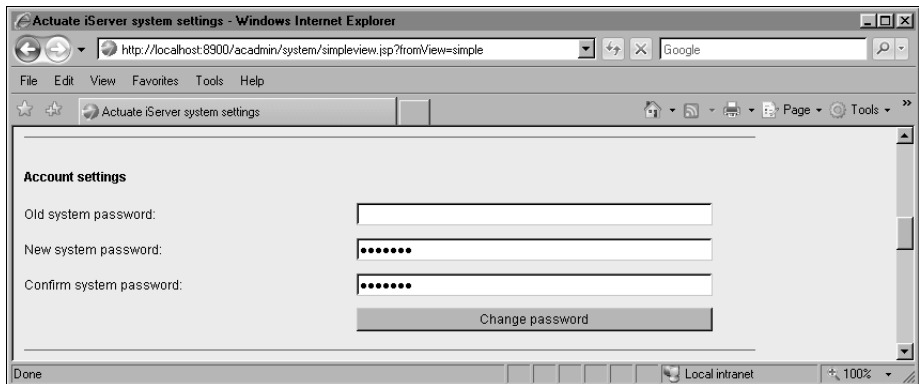


Figure 6-57 Creating a new Configuration Console password

- 4 Scroll to the top of the simple view and choose Advanced view. In Advanced view, choose Servers. In the example, the node named KOZU has joined the cluster, as shown in Figure 6-58.

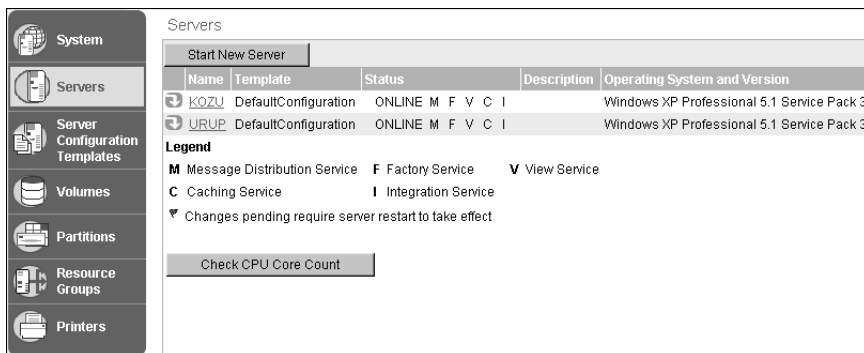


Figure 6-58 Viewing a cluster composed of two nodes

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 DLLs, previously installed in `WINDOWS\system32` by earlier iServer versions, are now installed in `AC_SERVER_HOME\bin` directory:

- `acxerces-c_1_4_90.dll`
- `msvc90.dll`
- `LTDIS10N.dll`
- `ltfil10N.dll`
- `msvcr71.dll`
- `acicudt18.dll`
- `ltkrm10N.dll`
- `msvcr90.dll`
- `acr7790w.dll`
- `mfc90u.dll`
- `msvcrt.dll`
- `acrs11090.dll`

The following binaries, previously installed in `AC_SERVER_HOME\operation` by iServer earlier versions, are now installed `AC_SERVER_HOME\bin` directory:

- `viewsrv11.exe`
- `fctsrvr11.exe`
- `fctcmd11.exe`
- `xmlparse.dll`

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/MyClasses
- \$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:

- Preparing to install Information Console
- Installing Information Console on Windows

Preparing to install Information Console

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 the environment before starting the Information Console installation process.

About installing from an FTP download

If you download an Actuate product from the Actuate FTP 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 installing Information Console.

If any files are missing, the installation program exits. Files can be missing if you 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

In Actuate 11, perform installations of Actuate iServer System products in the following order:

- Actuate BIRT iServer
- Actuate Information Console
- Actuate BIRT iServer Integration Technology

Actuate iServer installs Configuration, Management, and Information Console automatically on the machine where it resides. Typically, a manual Information Console installation is done to provide network access to iServer from a remote machine.

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>

To install the online help and PDF manuals, from Windows choose Start → Actuate 11 → Update Documentation.

Installing Information Console on Windows

This section describes how to install Actuate Information Console Release 11 for Windows. Before you begin the installation process, ensure that you have Actuate administrator, system administrator, and web administrator privileges.

Information Console installation requires the following information:

- The application server and port to use. The default settings are the Apache Tomcat for Actuate Information Console 11 service and port 8700, which the installation program configures. You must configure any other application server. If you are using a firewall, ensure that the firewall allows access to the port number you select.
- The installation wizard installs a Java Development Kit (JDK) distribution and a Java Runtime Environment (JRE) that the Apache Tomcat for Actuate Information Console 11 service uses.
- The name of the Encyclopedia volume that Information Console accesses. The default is the current machine.

You can install Information Console in two ways:

- Using an installation wizard
The installation wizard configures Information Console, creates shortcuts, and extracts and installs all necessary files. The installation wizard configures Apache Tomcat for Actuate Information Console 11 service.
- Deploying a WAR (web archive) file
Deploying directly requires that you configure for the application server. Use this option if your application server supports configuration of an application from a WAR file.

Using the installation wizard

The following section describes how to install Information Console using the installation wizard.

How to install using the wizard

- 1 Download the Information Console distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, `ActuateInformationConsole.exe`. The welcome message appears, as shown in Figure 7-1. Choose Next.



Figure 7-1 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 7-2. Choose Next.

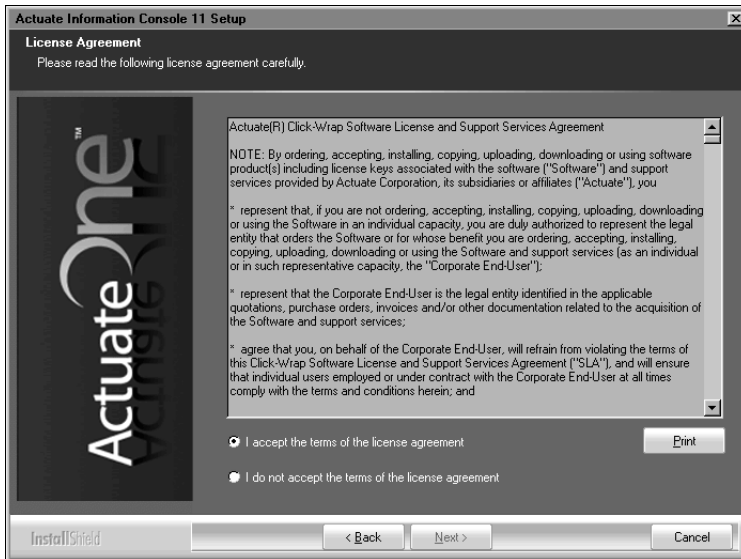


Figure 7-2 Accepting the license agreement

- 3 In Setup Type, select Typical, as shown in Figure 7-3. Choose Next.

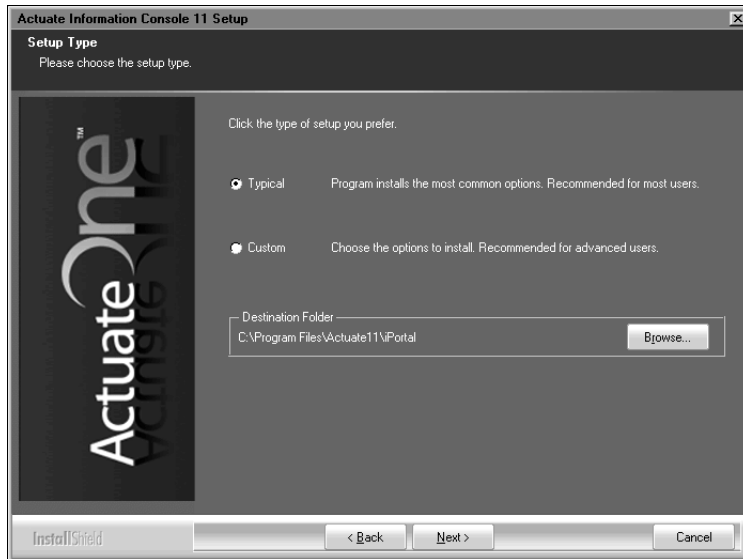


Figure 7-3 Specifying the typical or custom setup type

- 4 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 7-4, or specify these settings for your region.

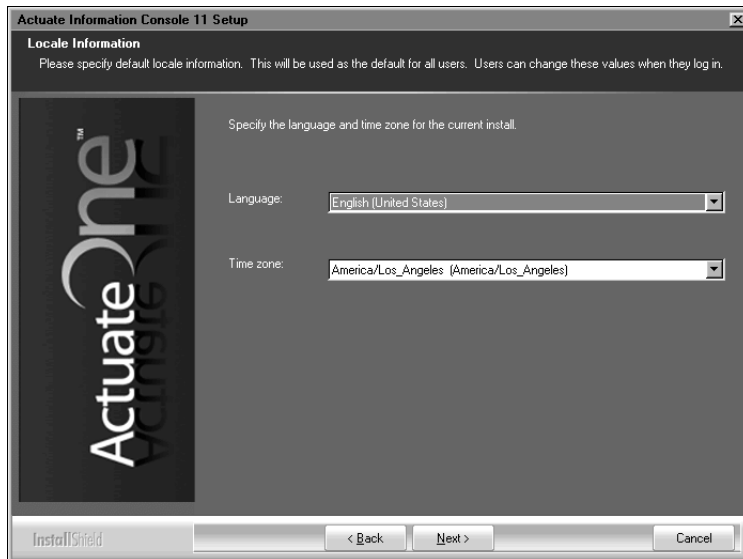


Figure 7-4 Specifying locale information

- 5 In Apache Tomcat for Actuate Information Console Service, accept the default port, 8700, and context path, /iportal, as shown in Figure 7-5. Choose Next.

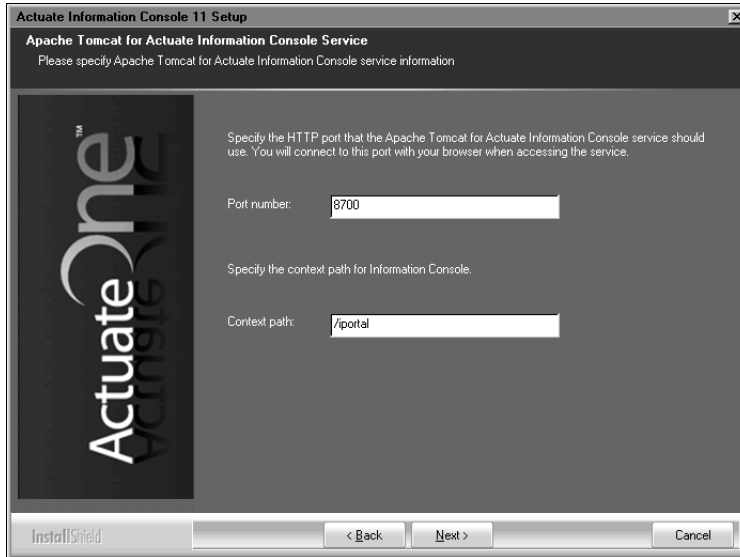


Figure 7-5 Specifying the port number and context path

- 6 In Specify Windows Account Information, type the user name and password for the account used to run the Apache Tomcat for Information Console 11 service, as shown in Figure 7-6. Choose Next.

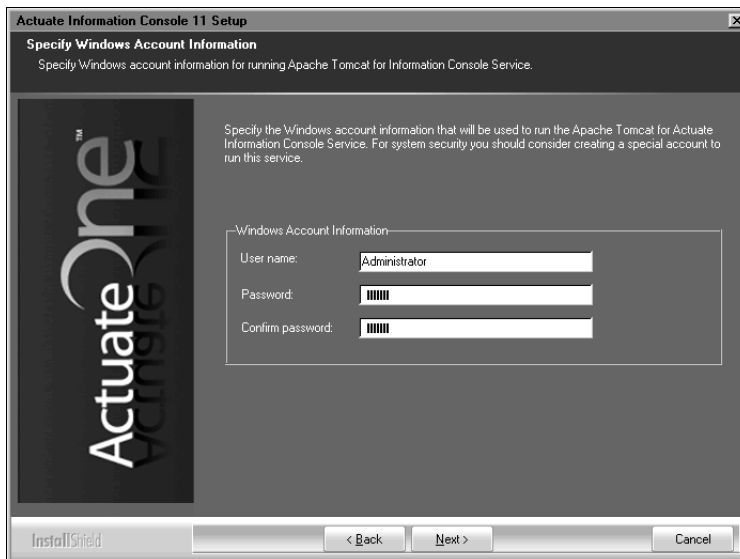


Figure 7-6 Specifying the account running the Information Console service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 7-7.




Figure 7-7 Setting the Windows local security policy

- 7 In iServer Information, accept the default values for Profile name, Host name, and Port number, as shown in Figure 7-8. Alternatively, type a different value for one or more of these properties. In Volume name, type a name for the default Encyclopedia volume. Choose Next.

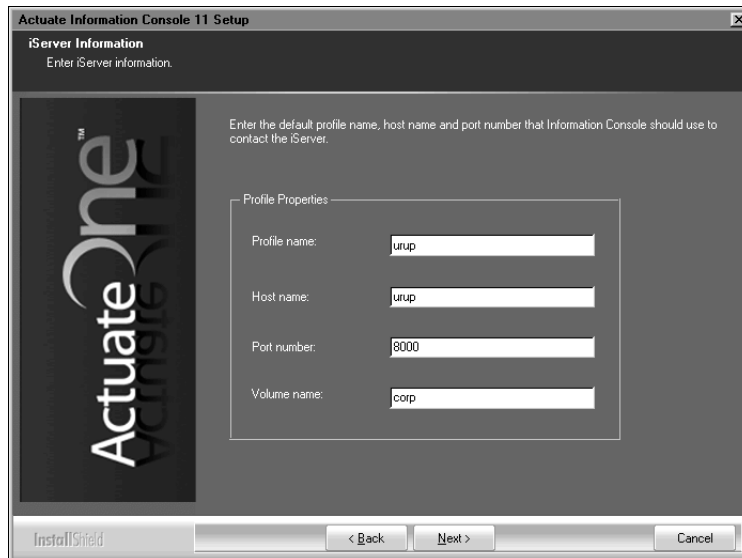


Figure 7-8 Specifying values for iServer profile properties

- 8 In Start Copying Files, review the settings shown in Figure 7-9. Choose Next.



Figure 7-9 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 7-10.

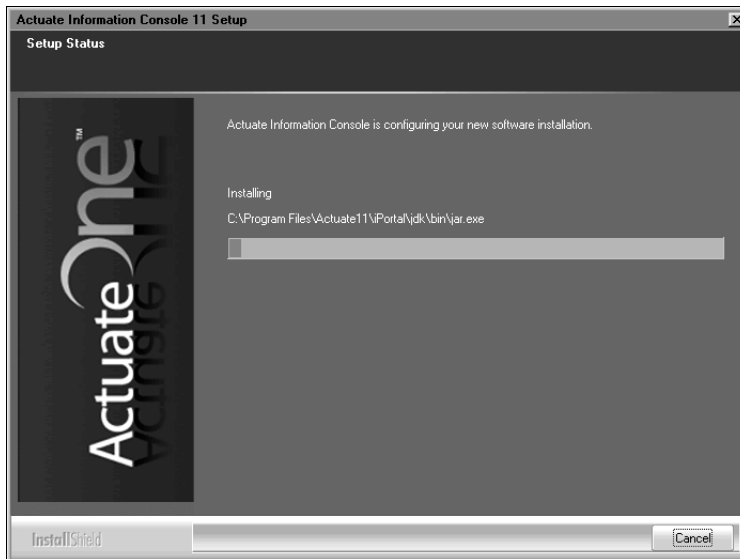


Figure 7-10 Setup Status

- 9 Select I would like to view the ReadMe file, as shown in Figure 7-11, if you want to review this documentation. Choose Finish to exit the wizard.



Figure 7-11 Exiting the installation wizard

If you chose to view the ReadMe file, the installation program opens the document, as shown in Figure 7-12.

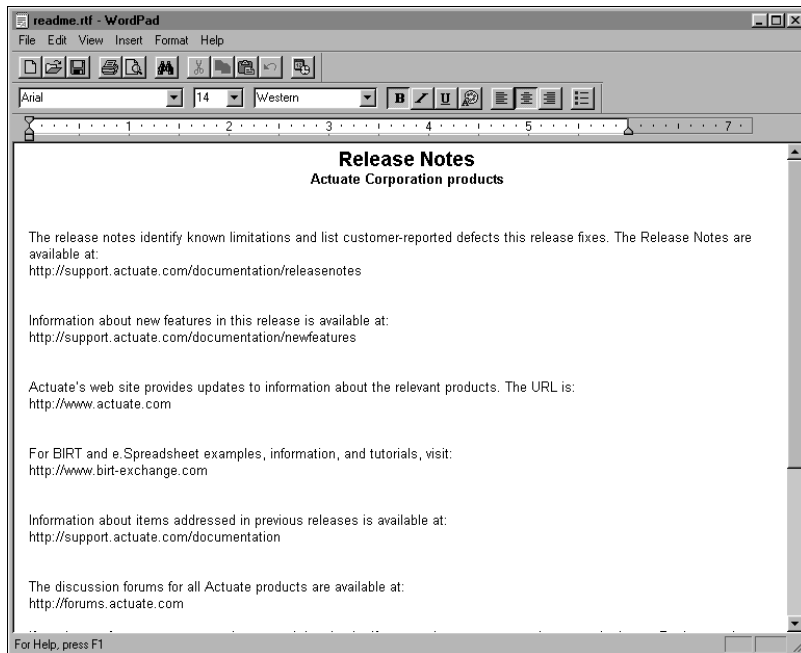


Figure 7-12 The ReadMe file

Using the WAR file to install

If Actuate supports your application server, you can deploy Information Console as a WAR file. See your 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 the application for your local environment, follow the steps in “Preparing the WAR file,” later in this section. 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 and application server for integration. You must also verify that you have a standard J2EE installation.

To integrate your application server with Actuate, complete the following tasks:

- Configure the application server to operate with Actuate products.
- Configure the Information Console WAR to integrate with the application server as described in “Preparing the WAR file.”
- 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 your application server’s documentation.
 - If your application server does not have deployment tools, add the Actuate context root to your application server, typically by modifying the application server's configuration file.

Preparing the server

Actuate recommends the following configuration for best performance:

- Use at least a two-CPU machine for Information Console.
- If Actuate iServer System 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 (web archive) 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 a 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 Actuate 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 general procedure for customizing an Information Console WAR file:

- 1 Create a temporary directory, such as C:/Temp/ic.
If you use an existing directory, ensure that this directory is empty.
- 2 Open a command window and type the following commands, replacing the X: drive letter with a drive letter appropriate to your system:
 - 1 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:

```
cd C:/Temp/ic
copy X:/TOMCAT_ActuateInformationConsole.war .
```

- 2 Decompress the file, as shown in the following example:

```
jar -xf TOMCAT_ActuateInformationConsole.war
```

The Information Console files appear in the temporary directory. Leave the Command window open.

- 3 Using a text editor that accepts UTF-8 encoding, edit web.xml to configure Information Console for the application server.

If you used the temporary path in Step 2, the file location is C:/Temp/ic/WEB-INF/Web.xml. Refer to Table 7-1 for a list of entries to modify in web.xml.

- 4 Save and close web.xml.
- 5 Type the following command:

```
jar -cf ../newinformationconsole.war *
```

This command creates newinformationconsole.war in the /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 the application server supports clustering, see the 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 Apache Tomcat for Actuate Information Console 11 service or the application server and web server, if necessary.
- 2 Start Information Console:
 - If you used the wizard installation, Choose Start→All Programs→Actuate 11→Information Console.
 - If you used the deployment installation, type a URL in your web browser. For example, type a URL similar to the following example:

```
http://Actuate1:8900/portal/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.
- portal 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 the URL to the BIRT iServer.

The Information Console login page appears.

- 3** On the Information Console login page, type a user name and password. Accept the default settings in Encyclopedia Volume, Language, and Time zone.
- 4** Choose Log In.

8

Installing iServer Integration Technology and Documentation

This chapter discusses the following topics:

- Installing iServer Integration Technology
- Installing the localization and documentation files
- About accessing online help
- Searching PDF manuals using master-index.pdx

Installing iServer Integration Technology

This section describes how to install iServer Integration Technology for Windows. In a default installation, iServer Integration Technology installs in C:\Program Files\Actuate11\ServIntTech.

How to install

To install iServer Integration Technology, perform the following tasks:

- 1 Download the iServer Integration Technology distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, ActuateiServerIntegrationTechnology.exe. The welcome message appears, as shown in Figure 8-1. Choose Next.



Figure 8-1 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 8-2. Choose Next.



Figure 8-2 Accepting the license agreement

- 3 In Setup Type, shown in Figure 8-3, select Typical. Choose Next.




Figure 8-3 Selecting a typical installation

- 4 In Start Copying Files, review the settings shown in Figure 8-4. Choose Next.

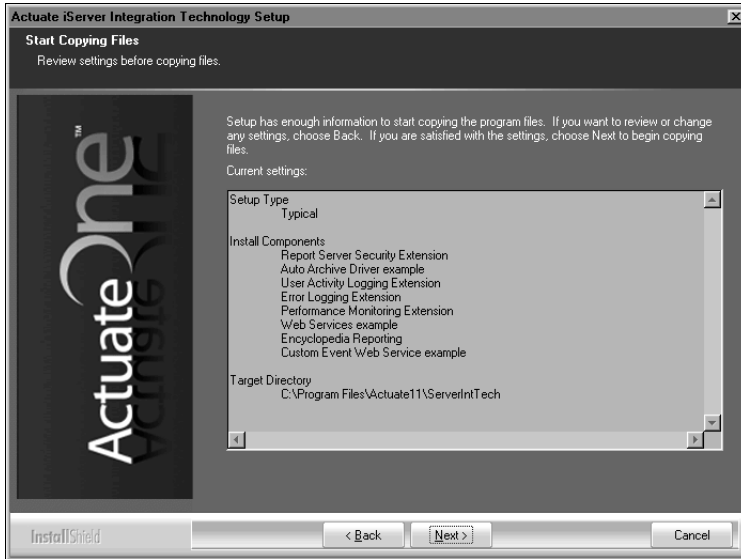


Figure 8-4 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 8-5.

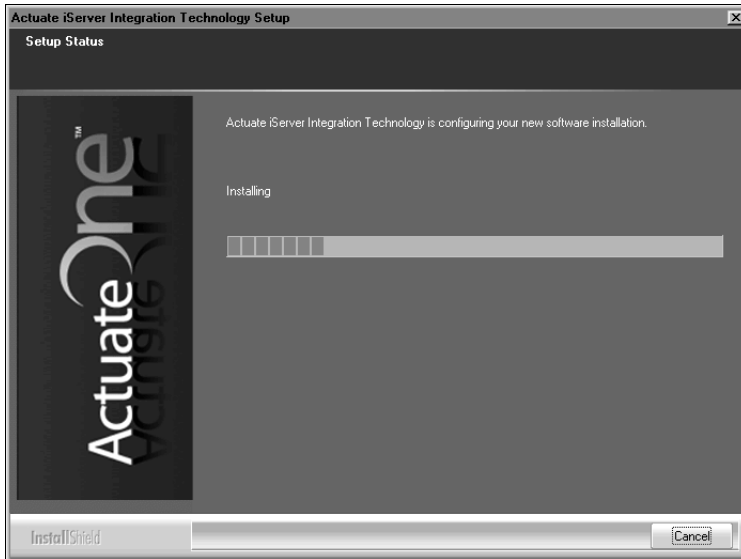


Figure 8-5 Viewing Setup Status

- 5 When the setup success message appears, select I would like to view the ReadMe file, as shown in Figure 8-6, if you want to review this documentation. Choose Finish to exit the wizard.

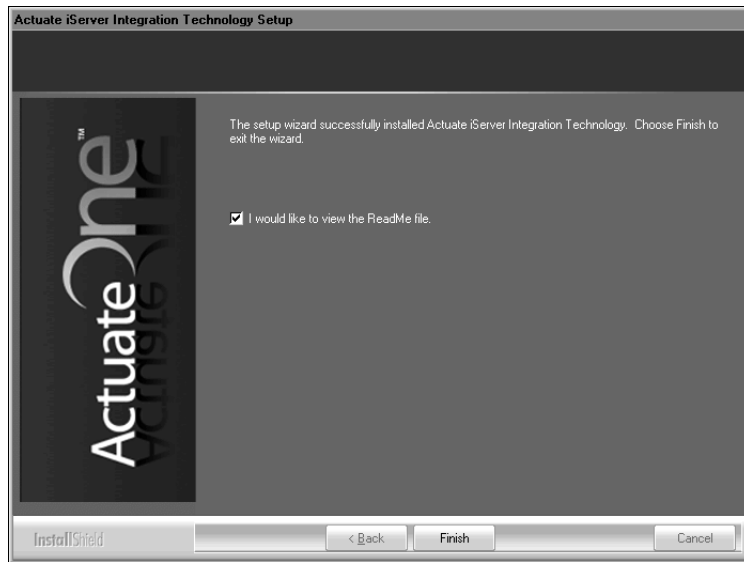


Figure 8-6 Exiting the installation wizard

If you chose to view the ReadMe file, the installation process opens the document, as shown in Figure 8-7.

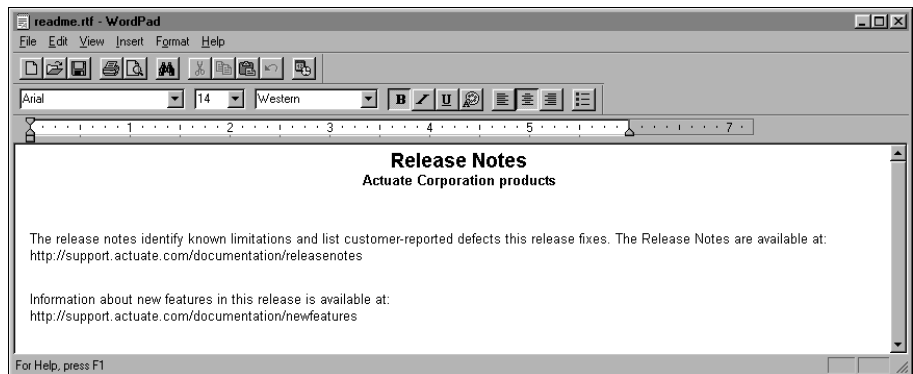


Figure 8-7 Viewing the ReadMe file

- 6 The installation program prompts you to install the online help from the following location, as shown in Figure 8-8:

<http://www.actuate.com>




Figure 8-8 Viewing the install online help and manuals prompt

To install the online help and PDF manuals from this location, in Windows choose Start→ Actuate 11→Update Documentation.

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:

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

Updates to documentation in PDF form are available from the following Actuate web sites:

<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 tasks:

- 1 Download the Actuate Localization and Documentation distribution package from an FTP software distribution site. Extract the files. Run the self-extracting executable file, `ActuateLocalizationandOnlineDocumentation.exe`. The welcome message appears, as shown in Figure 8-9. Choose Next.

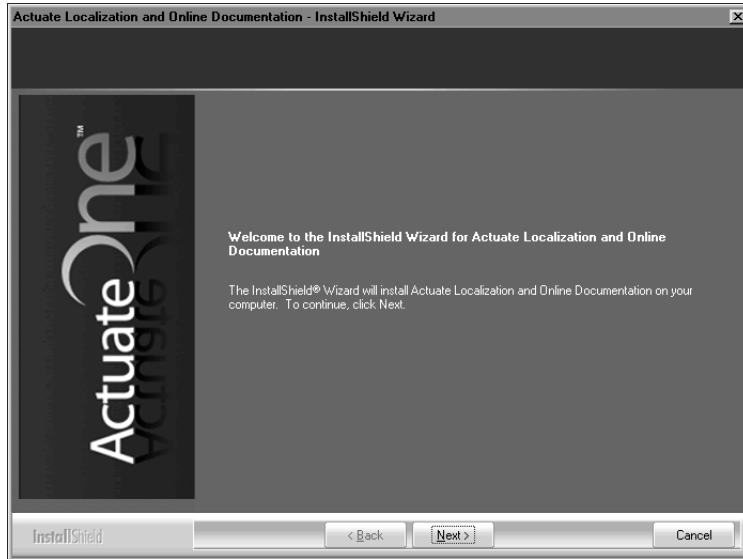


Figure 8-9 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 8-10. Choose Next.

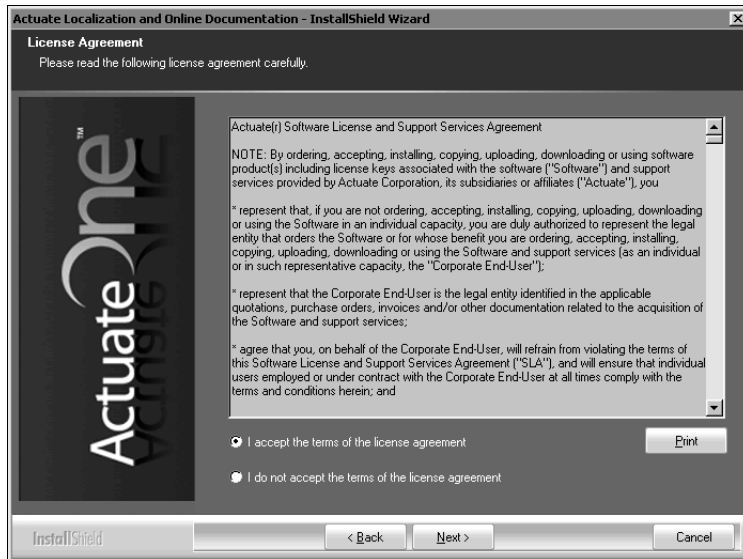


Figure 8-10 Accepting the license agreement

- 3 In Setup Type, select Typical, as shown in Figure 8-11, or select Custom to specify a limited set of localization and documentation files. Choose Next.

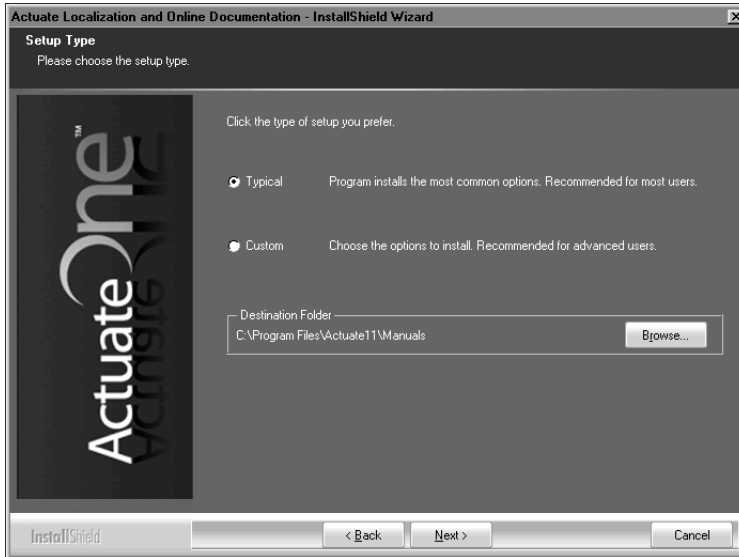


Figure 8-11 Specifying typical or custom setup type

- 4 In Start Copying Files, review the settings shown in Figure 8-12. Choose Next.



Figure 8-12 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 8-13.

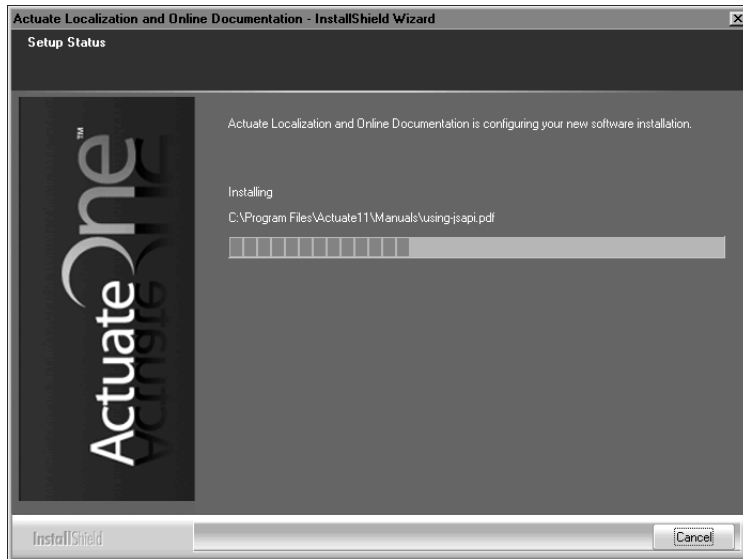


Figure 8-13 Viewing Setup Status

Setup completed successfully message appears, as shown in Figure 8-14. Choose OK.




Figure 8-14 Viewing successful setup message

About accessing online help

iServer supports accessing online help in following ways:

- Online from www.actuate.com
Use this option to ensure that you always have the latest documentation.
- Locally from the installed online localization and documentation files
Use this option if you do not have an internet connection.

How to switch between online help and local help

- 1 Choose Start → Programs → Actuate 11 → Switch Help Location.
- 2 On docupdate, select either Use online help, or Use local help, as shown in Figure 8-15.

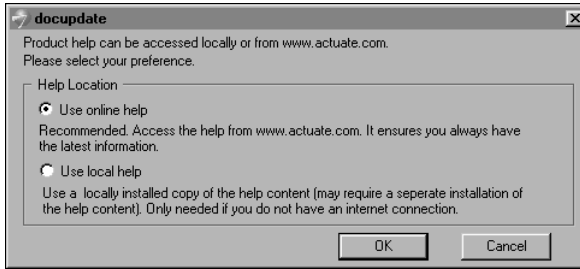


Figure 8-15 Selecting online or local help

Choose OK.

- 3 Restart Actuate BIRT iServer Enterprise Service.

Searching PDF manuals using master-index.pdx

If you install the PDF version of the manuals, you can also use the Actuate Documents Catalog (master-index.pdx) to search for topics across the entire set of books.

How to search the Actuate Documents Catalog

- 1 Navigate to the ACTUATE_HOME\Actuate11\Manuals directory. Open master-index.pdx.
- 2 On Search, in Where would you like to search?, select All PDF documents in, then choose the ACTUATE_HOME\Actuate11\Manuals.
- 3 In What word or phrase would you like to search for?, enter the word or phrase. For example, type accessing online help, as shown in Figure 8-16.

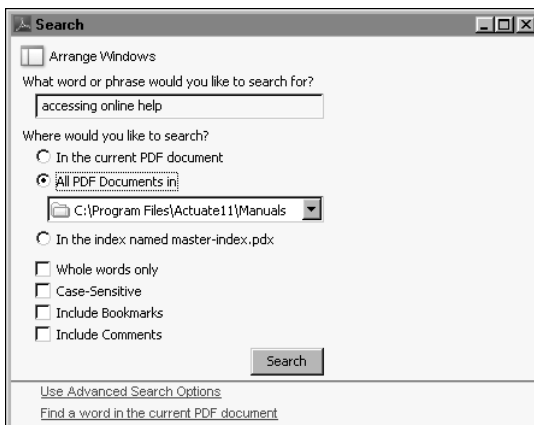


Figure 8-16 Specifying the search

Choose Search.

Search displays all occurrences of the word or phrase in the Actuate Documents Catalog.

- 4 Select an item in the results list to display the documentation in Adobe Reader, as shown in Figure 8-17.

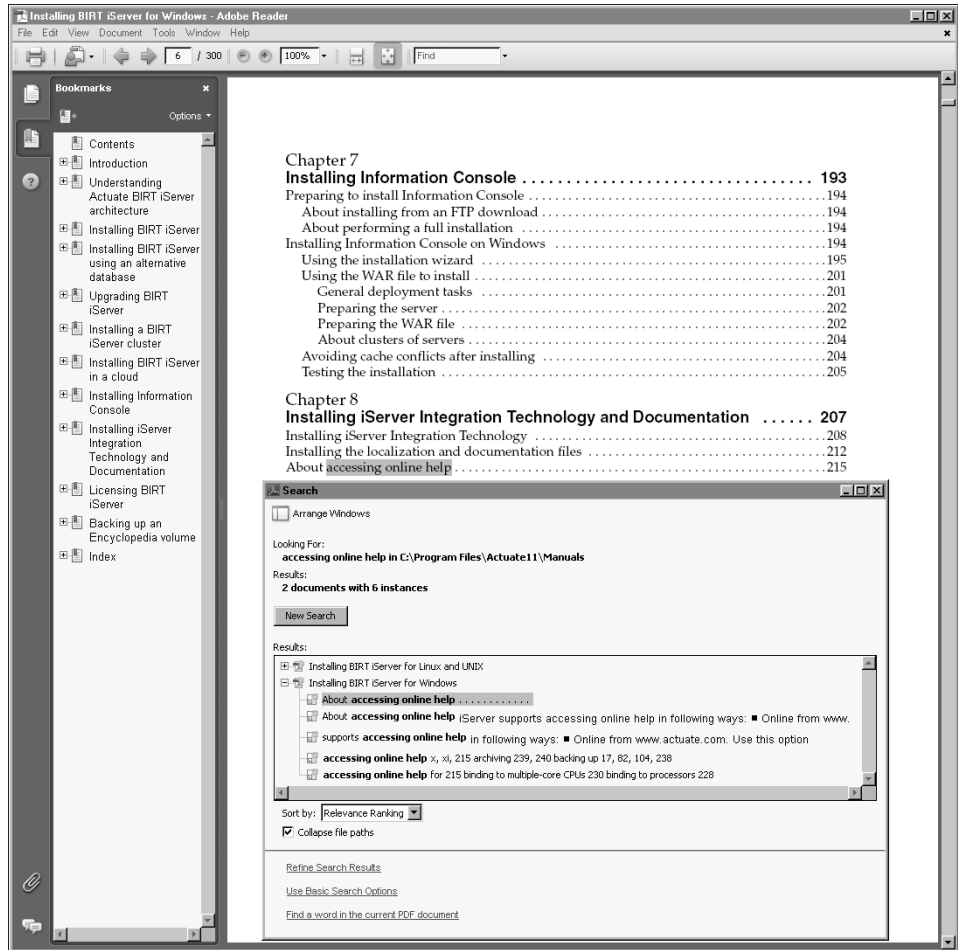


Figure 8-17 Viewing search results

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 a 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)

Option	Description	Supported releases
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 a Unlimited User CPU License.	10, 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.

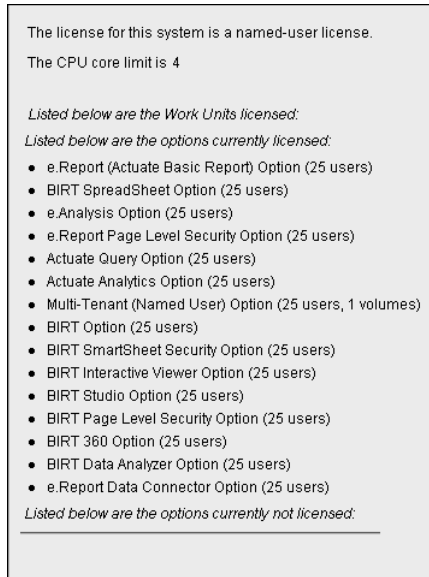


Figure 9-1 iServer License options

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 runtime, 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.

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 InstallShield Wizard prompts you to provide the location of the license file. After providing the location of the license file, the InstallShield Wizard issues a prompt similar to the following message, as shown in Figure 9-2:

```
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
o1wxAKu0Vr4NA6o8uWCzjGZXX8KrvjViSUoROj70hWOY=
```

```
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\iServer\bin folder.
```



Figure 9-2 Viewing iServer temporary license message

The format of the alphanumeric string for the machine ID and location of the license file are different depending on the operating system. On a Windows system, the unique identifier for the network card is the source of the machine ID. You must have the network card enabled on the BIRT iServer machine to obtain the machine ID.

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:

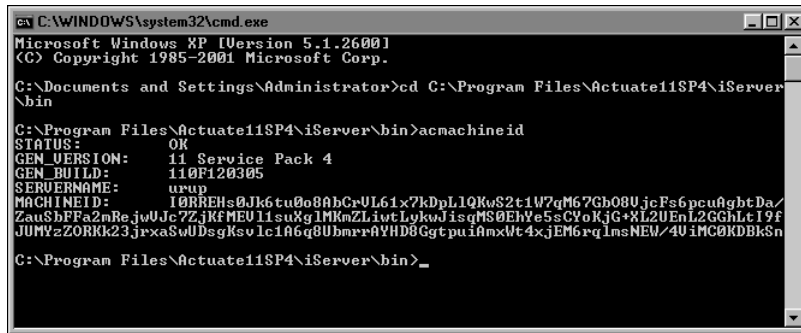
- 1 Open a command prompt and navigate to `AC_SERVER_HOME\bin`.
- 2 Type the following command and press Enter:

```
acmachineid
```

The utility provides output in the following format:

```
C:\Program Files\Actuate11SP4\iServer\bin>acmachineid
STATUS:          OK
GEN_VERSION:     11 Service Pack 4
GEN_BUILD:       XxXBuild NumberXxX
SERVERNAME:      <hostname>
MACHINEID:      IORREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67Gb08
                  VjcF56pcuAgbtDaJzauSbFFa2mRejwVJc7ZjKfMEV11suXglM
                  kmZLiwtLykwJisqMS0EhYe5sCYoKjG+XL2UEEnL2GGHltI9f
                  JUMYzZORkK23jrxaswUDsgKsvlc1A6q8UbmrrAYHD8Ggtptui
                  AmxWt4xjEM6rqlmsNEW/4ViMC0KDBkSn
```


Figure 9-3 shows the output as it appears in the command prompt.



```
C:\WINDOWS\system32\cmd.exe
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Administrator>cd C:\Program Files\Actuate11SP4\iServer\bin

C:\Program Files\Actuate11SP4\iServer\bin>acmachineid
STATUS:          OK
GEN_VERSION:     11 Service Pack 4
GEN_BUILD:       110F120305
SERUENAME:       usup
MACHINEID:       I0RRHs0Jk6tu0o8AbCrUL61x7kDpL1QKwS2t1W7qM67Gb08UjcF6peuAqbtDa/
ZauShFFa2nRejuUJc7ZjKfMEU11su8gIMKmZLiofLykuaIisqMS0EhYe5sCYoKjG+XL2UEnL2GhLcI9f
JUMVzZORRk23jrxasWUDsgKsvlc1A6q8UbmrrAVHD8GgtPuiAmxWt4xjEM6.rq1nsNEW/4U1MC0KDBkSn

C:\Program Files\Actuate11SP4\iServer\bin>
```

Figure 9-3 Executing acmachineid

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.

Actuate_iServer_key_XXXXX.xml will contain the node key information for the standalone 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
    VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwLykDa/
    wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORKk23jrxaswUDig
    Ksvlc1A6q8UbmrrAYHD8GgtpuiaMxWt4xjEM6rqlmsNEW/4Vjm40Kx1kSv"
    ServerName="W7CLSTRNODE1"/>
  <NodeKey
    MachineId="I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67Gb08
    VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwLykDa/
    wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORKk23jrxaswUDsg
    Ksvlc1A6q8UbmrrAYHD8GgtpuiaMxWt4xjEM6rqlmsNEW/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 a BIRT iServer machine count toward the maximum number of licensed CPUs. For more information about performing CPU binding on a Sun Solaris or HP-UX 11i machine, see Chapter 7, “Licensing BIRT iServer,” in *Installing BIRT iServer for Linux and UNIX*.

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 information on the following topics:

- Binding a BIRT iServer to processors on a Windows machine
- Checking BIRT iServer bound processors
- Configuring e-mail for CPU license problems

Binding a BIRT iServer to processors on a Windows machine

You can perform the following types of CPU binding on Windows:

- Binding to specific CPUs
- Binding to multiple-core CPUs
- Binding an Actuate process to a processor

The following sections describe these features.

Binding to specific CPUs

On a multiple-CPU machine running the Windows operating system, the server operating system assigns an ID number to each processor. Windows Task Manager lists the IDs of the available processors. The numbering starts at 0.

How to bind BIRT iServer to a set of processors

To bind BIRT iServer to a set of processors, perform the following steps:

1 Choose Start→Control Panel→System.

On System Properties, choose Advanced. Then select Environment Variables.

2 On Environment Variables, perform the following tasks:

1 In System Variables, choose New.

2 On New System Variable, perform the following tasks:

1 in Variable name, type:

AC_PMD_WINDOWS_CPUS

2 In Variable value, specify the processors to which to bind BIRT iServer by typing a comma-separated list of integers. For example, to bind BIRT iServer to CPU 0, CPU 3, and CPU 4, type the following list:

0, 3, 4

New System Variable looks like Figure 9-4.

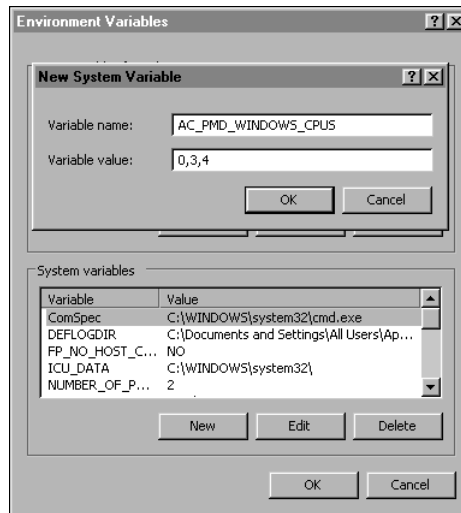


Figure 9-4 Creating the AC_PMD_WINDOWS_CPUS system variable
Choose OK. AC_PMD_WINDOWS_CPUS appears in System variables, as shown in Figure 9-5.

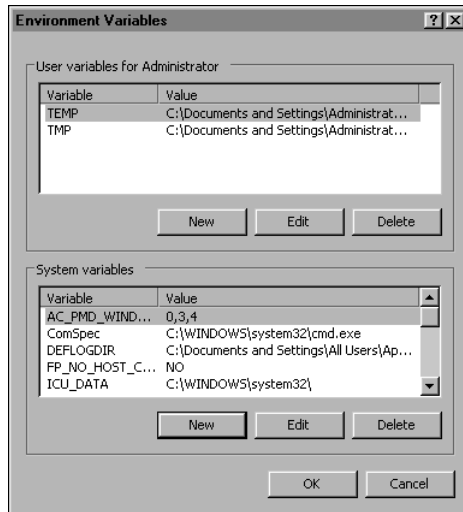


Figure 9-5 Viewing the AC_PMD_WINDOWS_CPUS system variable

On Environment Variables, choose OK. Then, on System Properties, choose OK.

You can verify the CPU binding by checking the Processor Affinity of the BIRT iServer process using Task Manager.

Binding to multiple-core CPUs

You can also perform multiple-core CPU binding, similar to the way you bind to a single CPU, using the AC_PMD_WINDOWS_CPUS environment variable, as described in the previous section. To BIRT iServer, each core appears as a logical CPU.

For example, on a dual-core, 2-CPU system, setting the variable value to 0,1 binds BIRT iServer to both cores on the first CPU. Setting the variable value to 0,2 binds BIRT iServer to one core on each CPU. Setting the variable value to 0 binds BIRT iServer to one core on the first CPU.

Actuate does not recommend restricting BIRT iServer processing on a multiple-core CPU machine to one core for licensing purposes. BIRT iServer System achieves significant performance gains on a multiple-core CPU machine.

For example, BIRT iServer scales nearly perfectly from 1 to 2 cores and gets 50% better throughput on a dual-core system than on a 2-CPU system.

Binding an Actuate process to a processor

If you bind the BIRT iServer PMD to a subset of CPUs on a machine, you can also bind the Factory, View, Integration, and Caching processes to a specific CPU. Under some conditions, binding an Actuate process to a specific CPU can

enhance performance. Binding an Actuate process to a CPU has no effect on the CPU calculations BIRT iServer performs to determine the maximum number of licensed CPUs.

If you bind a process to a CPU, you must bind the CPU to both the BIRT iServer PMD and the process. BIRT iServer writes to the error log and stops the process if you bind a process to a CPU that you do not bind to the PMD.

To bind a BIRT iServer process to CPU processors, use the ProcessorAffinity element in the acserverconfig.xml file for BIRT iServer. List the IDs for the CPUs to which to bind a process as Item subelements in the following ProcessorAffinity elements:

- To bind Factory processes, specify the CPU IDs in the ProcessorAffinity element within the ReportingService element.
- To bind View processes, specify the CPU IDs in the ProcessorAffinity element within the ViewingService element.
- To bind Integration processes, specify the CPU IDs in the ProcessorAffinity element within the IntegrationService element.
- To bind Caching processes, specify the CPU IDs in the ProcessorAffinity element within the CachingService element.

You must also ensure that you bind the specified CPUs to the PMD for the BIRT iServer machine. For example, on a 4-CPU machine, the following ProcessorAffinity example binds View processes to CPU IDs 0 and 2:

```
<ViewingService
  EnableViewingService="true"
  <ProcessorAffinity>
    <Item>0</Item>
    <Item>2</Item>
  </ProcessorAffinity>
/>
```

About processors and hyperthreading

Some Intel processors use hyperthreading, a technology that counts each physical processor as a specific number of logical processors. The operating system and any programs running on the machine see the number of logical processors, not the number of physical processors.

When a machine uses hyperthreading, Windows Task Manager lists the logical processors, not the physical ones. You specify the number of logical processors in the environment variable. When a machine uses hyperthreading, BIRT iServer calculates the number of bound processors by dividing the number of bound logical processors by the number of logical processors for each physical processor. If the result contains a decimal component, BIRT iServer uses the next highest integer. For example, it rounds 4.3 to 5. In the following example, a machine has

four physical processors. With hyperthreading enabled, each physical processor corresponds to two logical processors. The machine has the following logical processors available:

- Physical processor 0 corresponds to logical processors 0 and 1.
- Physical processor 1 corresponds to logical processors 2 and 3.
- Physical processor 2 corresponds to logical processors 4 and 5.
- Physical processor 3 corresponds to logical processors 6 and 7.

If you bind BIRT iServer to the five logical processors 0, 2, 3, 6, and 7, it calculates the number of bound processors as:

$$5/2 = 2.5$$

BIRT iServer rounds this number up to determine that you have three bound processors.

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 `encycsrvr11` 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 `encycsrvr11` processes on the machine and restart one of the `encycsrvr11` 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

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 SQL Server database
- Backing up and restoring an Encyclopedia volume that uses a DB2 database

Performing an Encyclopedia volume backup

When performing a 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 Windows Actuate 11 Service Pack 4 environment, the default AC_SERVER_HOME path is:

```
C:\Program Files\Actuate11SP4\iServer
```

The default AC_DATA_HOME path is:

```
C:\Actuate\iServer\data\
```

The default Encyclopedia volume path is:

```
C:\Actuate\iServer\data\encyc
```

The default acserverconfig.xml file path is:

```
C:\Actuate\iServer\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, Microsoft SQL Server, 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 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

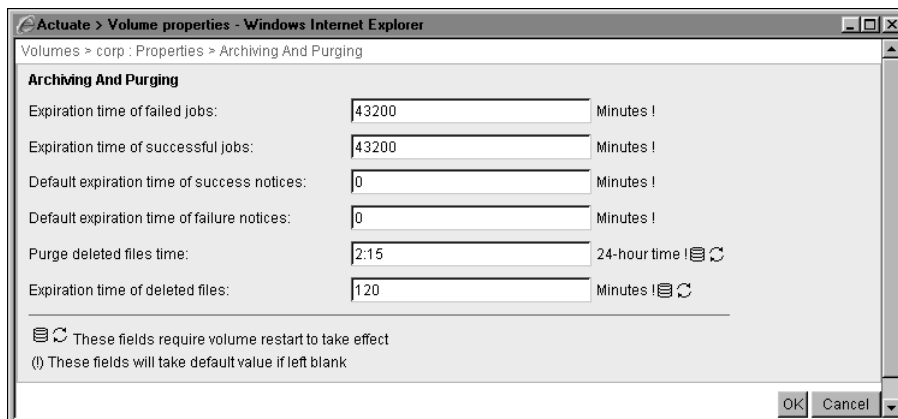


Figure 10-1 Configuring file purging properties

Choose OK.

For information on other aspects of archiving, see Chapter 12 “Archiving files,” in *Configuring BIRT iServer*.

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 examples. 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 pgAdmin

To back up an Encyclopedia volume using the pgAdmin graphical 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 pgAdmin utility.
- Back up the `acserverconfig.xml` file and volume data folders to the backup folder.

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 Windows 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 Open Windows Explorer.
- 2 In Windows Explorer, choose **File**→**New**→**Folder** repeatedly to create a set of nested folders in the following location, as shown in Figure 10-2:

`C:\Actuate\iServer\encyc_backup`




Figure 10-2 Creating a new folder named `encyc_backup`

Back up Encyclopedia volume metadata using the graphical administration tool, pgAdmin, to automatically run `pg_dump` by performing the following tasks.

How to run `pg_dump` using pgAdmin

- 1 In Windows, choose Start→Programs→pgAdmin III→pgAdmin III.
- 2 On pgAdmin III, in Object browser, right-click the PostgreSQL Database Server and choose Connect, as shown in Figure 10-3. If the PostgreSQL Database Server does not appear in Object browser, you can add the server manually. For more information on adding the server manually, see Chapter 2, “Installing BIRT iServer,” earlier in this book.

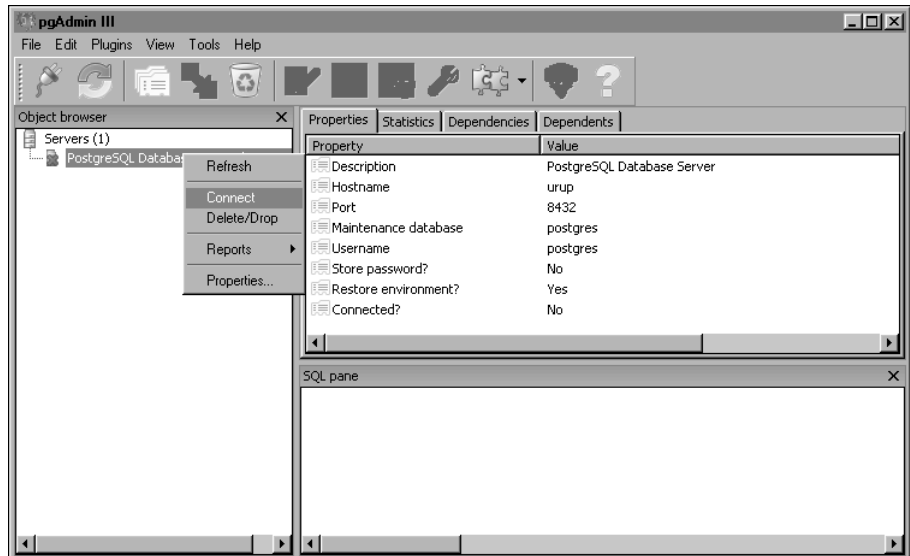


Figure 10-3 Connecting to PostgreSQL Database Server

- 3 In Connect to Server, type the postgres superuser password, as shown in Figure 10-4. You specified this password in PostgreSQL Database Information during the iServer installation.



Figure 10-4 Typing the password to connect to PostgreSQL Database Server

- 4 On pgAdmin III, in Object browser, expand PostgreSQL Database Server, expand Databases, right-click iserver, and choose Backup, as shown in Figure 10-5. This operation backs up the entire iserver database. Alternatively, to back up only one Encyclopedia volume schema, such as ac_corp, right-click the volume, and choose Backup.

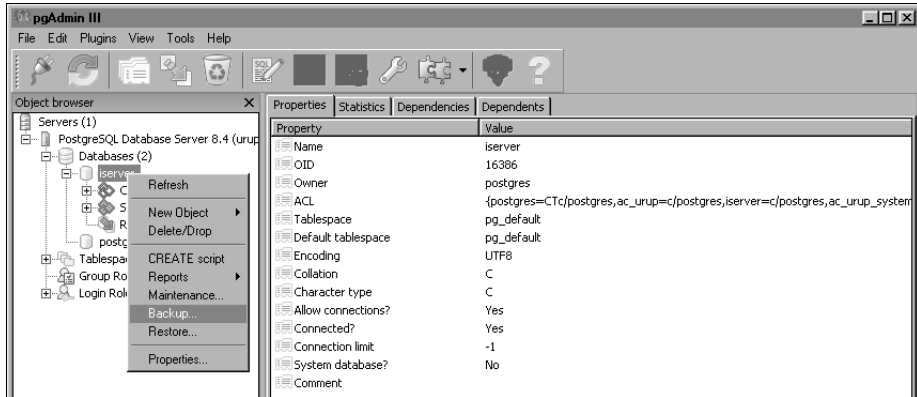


Figure 10-5 Choosing to back up the iserver database

- 5 On Backup Database iserver, perform the following tasks:
 - 1 In Filename, type:
C:\Actuate\iServer\encyc_backup\iserver.backup
 - 2 To execute pg_dump, accept the default option selections, as shown in Figure 10-6, and choose OK.

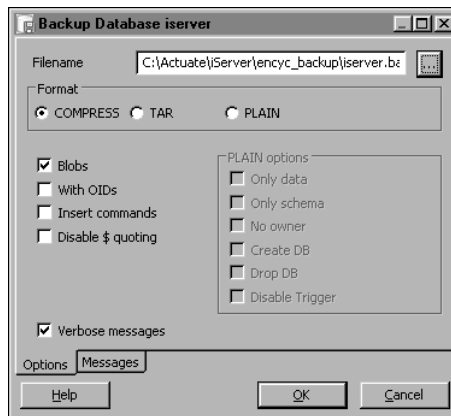


Figure 10-6 Choosing OK to start pg_dump execution

- 3 `pg_dump` executes, writing status messages to BackupDatabase iserver—Messages, as shown in Figure 10-7. Exit code 0 indicates that `pg_dump` ran successfully.

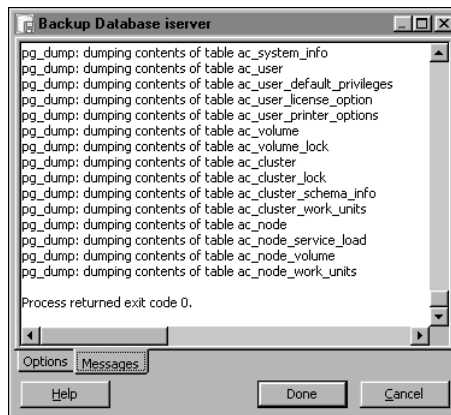


Figure 10-7 Viewing status messages `pg_dump` writes during execution

- 4 Scroll to the top of the output in Backup Database iserver—Messages to see the command that executed `pg_dump`, as shown in Figure 10-8.

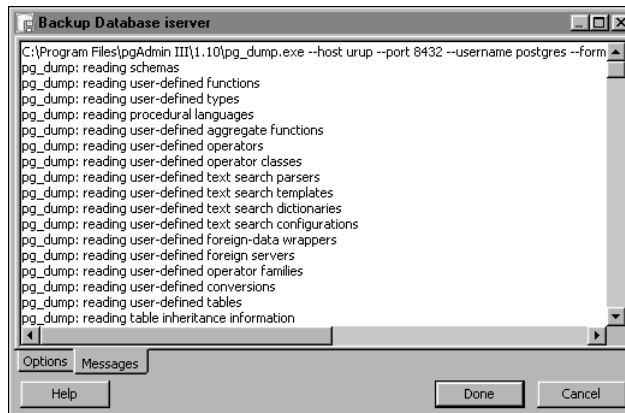


Figure 10-8 Viewing the command that executed `pg_dump`

The complete text of the command is:

```
C:\Program Files\pgAdmin III\1.10\pg_dump.exe --host urup
--port 8432 --username postgres --format custom --blobs
--verbose --file "C:\Actuate\iServer\encyc_backup\
iServer.backup" iServer
```

- 5 On Backup Database iserver, choose Done.

Backing up an Encyclopedia volume using pg_dump

Alternatively, you can back up an Encyclopedia volume schema using the command-line version of `pg_dump`. The following example duplicates the operations performed in the previous section using the graphical PostgreSQL administration tool, pgAdmin. You do not need to do both activities.

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>

Run `pg_dump` from the command line by performing the following tasks.

How to run `pg_dump` from a command prompt

1 Open a command prompt.

2 Navigate to the following location:

```
C:\Program Files\Actuate11SP4\iServer\postgresql\bin
```

3 Execute the following command. Substitute your machine name for `urup` in this example:

```
pg_dump.exe --host urup --port 8432 --username postgres
--format custom --blobs --verbose --file
"C:\Actuate\iServer\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.

4 Type the `postgres` superuser password. The administrator specified this password in PostgreSQL Database Information during the `iServer` installation procedure in Exercise 1.

`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 Open Windows Explorer and navigate to `AC_DATA_HOME`, which is the location of the `iServer` data. You specified this location on Setup Type during the install. The default path for `AC_DATA_HOME` is:

```
C:\Actuate\iServer\data
```

2 In `AC_DATA_HOME`, navigate to the `config` folder that contains `acserverconfig.xml` file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP4
```

3 Select `acserverconfig.xml`, right click, and choose Copy, as shown in Figure 10-9.

Copy this file to the following backup location:

C:\Actuate\iServer\encyc_backup

- 4 In AC_DATA_HOME\encyc, select the file, fileType, status, and tempRov folders, right-click, and choose Copy. Copy these folders to the following backup location:

C:\Actuate\iServer\encyc_backup

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.

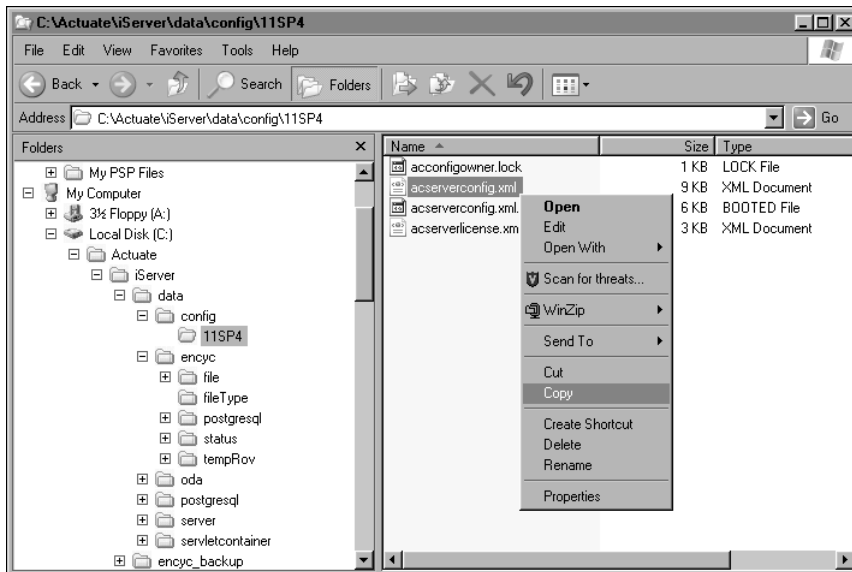


Figure 10-9 Copying acserverconfig.xml

Figure 10-10 shows a copy operation that contains the file, fileType, status, and tempRov folders.

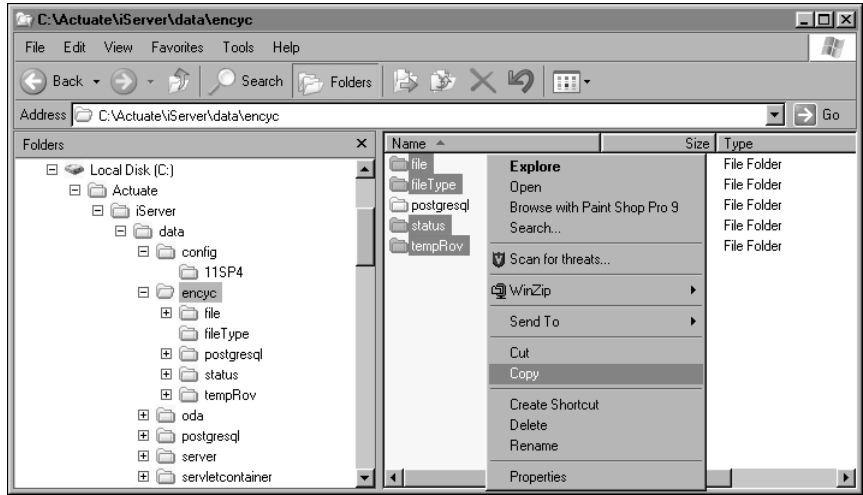


Figure 10-10 Copying the volume data folders

The contents of the backup folder appear as shown in Figure 10-11.

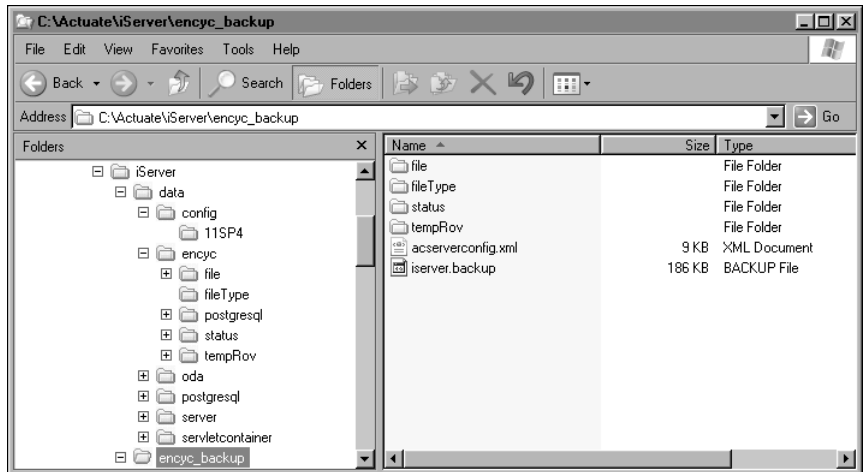


Figure 10-11 Viewing the backed-up files

Restoring an Encyclopedia volume using pgAdmin

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 using the PostgreSQL pg_restore utility.
- Take the Encyclopedia volume online.

To begin the restore operation, take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume offline, as shown in Figure 10-12.




Figure 10-12 Taking the volume offline

How to restore the backed-up volume data folders

- 1 In Windows Explorer, navigate to AC_DATA_HOME\config\11SP4.
- 2 Select acserverconfig.xml, right-click, and choose Delete, as shown in Figure 10-13. Confirm the deletion.

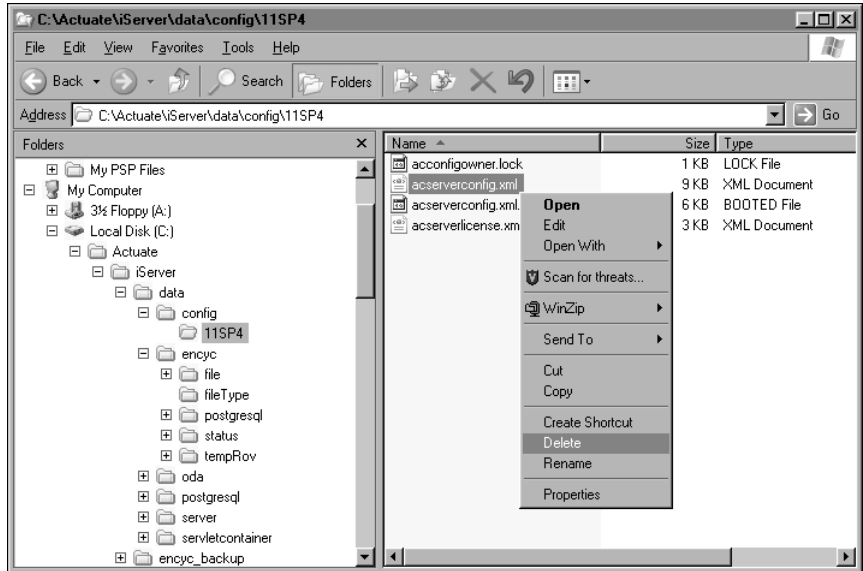


Figure 10-13 Deleting acserverconfig.xml

- 3 In AC_DATA_HOME, open the encyc folder. In AC_DATA_HOME\encyc, select the file, fileType, status, and tempRov folders, right-click, then choose Delete, as shown in Figure 10-14. Confirm the deletion.

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 select and delete the postgresql folder.

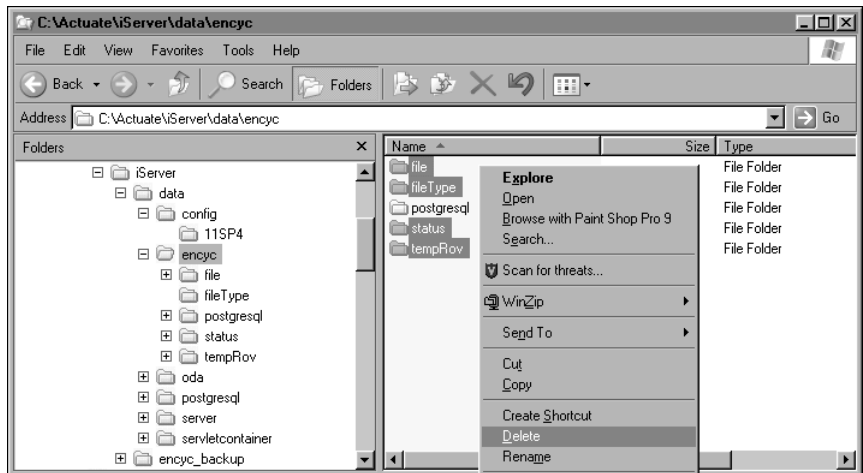


Figure 10-14 Deleting the file, fileType, and status folders from AC_DATA_HOME\encyc

4 In Windows Explorer, navigate to the following location:

C:\Actuate\iServer\encyc_backup

Select acserverconfig.xml, right-click, choose Copy, and copy this file to AC_DATA_HOME\config\11SP4, as shown in Figure 10-15.

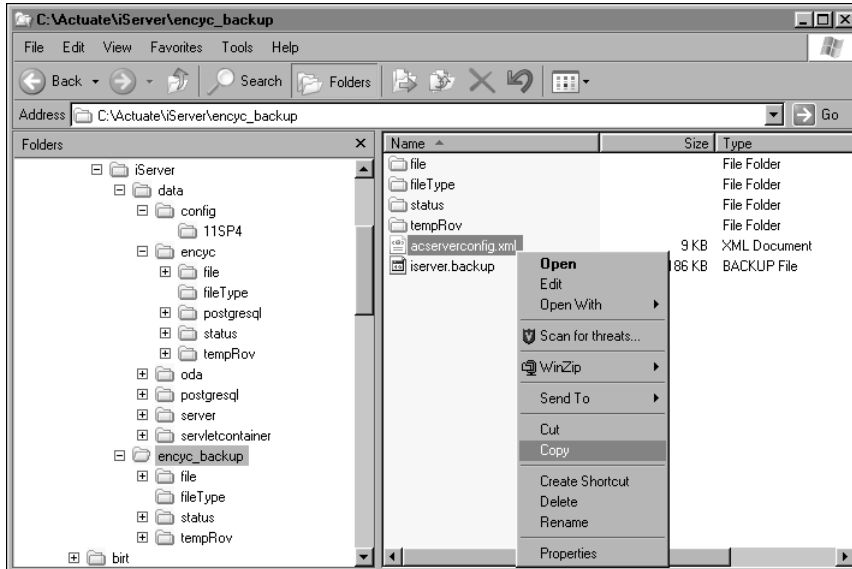


Figure 10-15 Choosing to copy acserverconfig.xml

AC_DATA_HOME\config\11SP4 appears as shown in Figure 10-16.

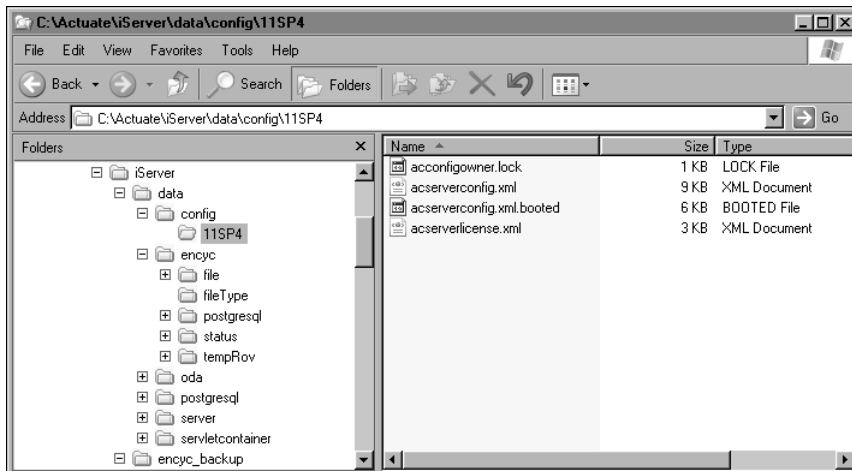


Figure 10-16 Viewing AC_DATA_HOME\config\11SP4 after restoring acserverconfig.xml

- In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempROV folders, right-click, choose Copy, and copy these folders to AC_DATA_HOME\encyc, as shown in Figure 10-17.

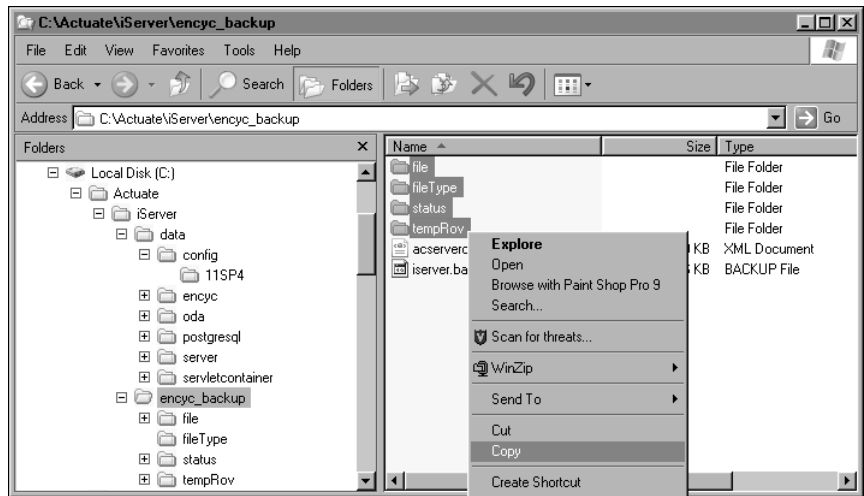


Figure 10-17 Choosing to copy the volume data folders

AC_DATA_HOME\encyc appears as shown in Figure 10-18.

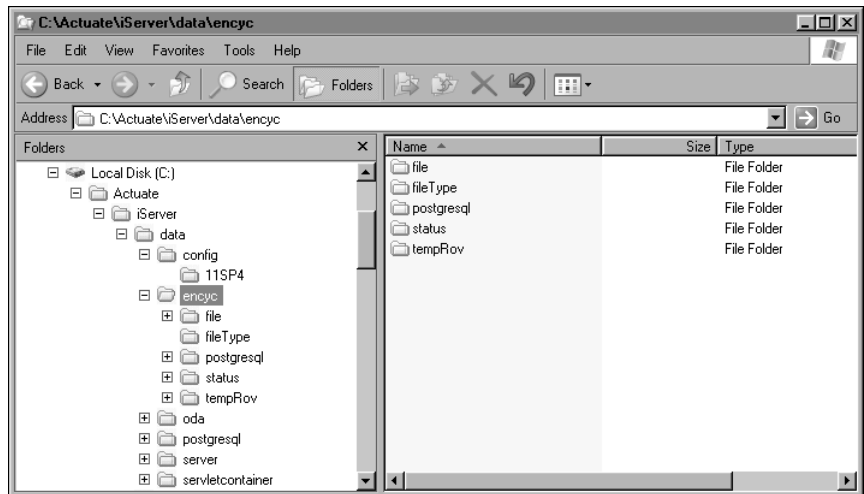


Figure 10-18 AC_DATA_HOME\encyc after copying backed-up folders

How to run pg_restore using pgAdmin

- On pgAdmin III, in Object browser, right-click iserver and choose Restore, as shown in Figure 10-19.

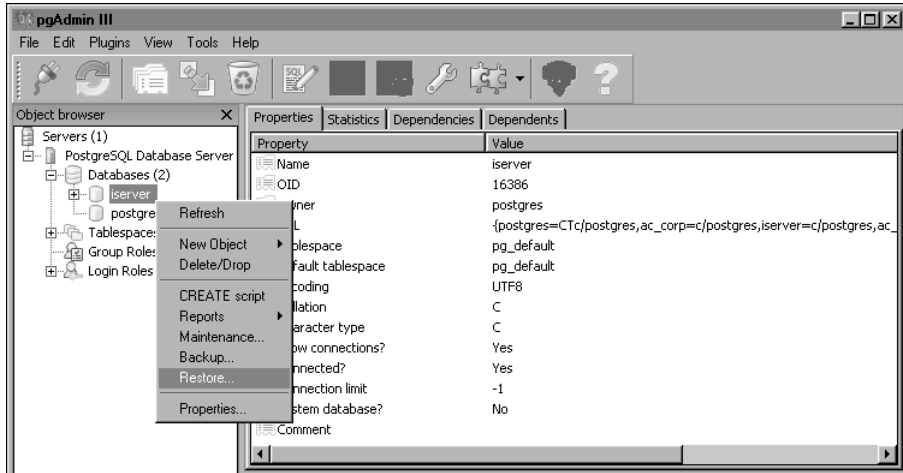


Figure 10-19 Choosing to restore the iserver database from backup



- 2 On Restore Database iserver, perform the following tasks:
 - 1 Choose the ellipse next to Filename. On Select backup filename, navigate to, and select the backup file that pg_dump created. The name of this file is:
C:\Actuate\iServer\encyc_backup\iserver.backup
 - 2 Select Clean before restore.
Restore Database iserver appears, as shown in Figure 10-20.

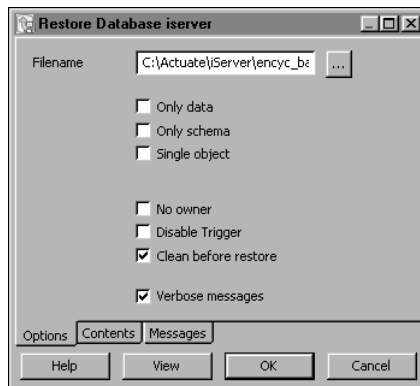


Figure 10-20 Specifying the backup file to restore

Choose OK.

- 3 pg_restore executes, writing status messages to BackupDatabase iserver—Messages, as shown in Figure 10-21. Exit code 0 indicates that pg_restore ran successfully.

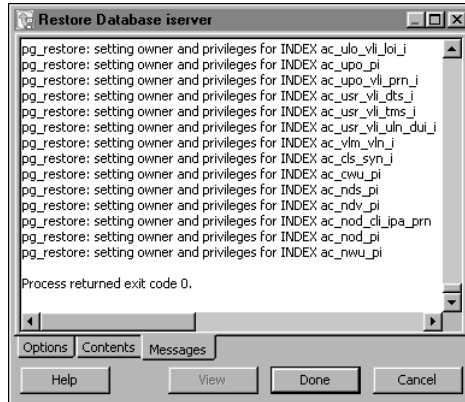


Figure 10-21 Viewing status messages pg_restore writes during execution

- 4 Scroll to the top of the output in Backup Database iserver—Messages to see the command that executed pg_restore, as shown in Figure 10-22.

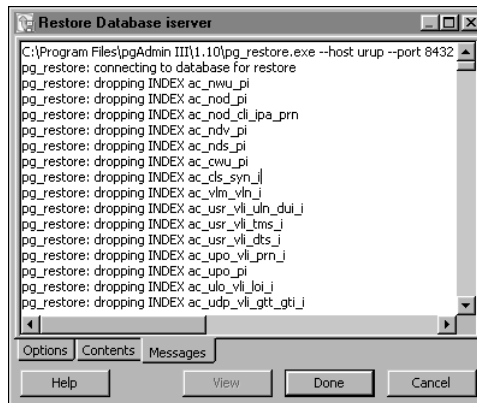


Figure 10-22 Viewing the command that executed pg_restore

The complete text of the command is:

```
C:\Program Files\pgAdmin III\1.10\pg_restore.exe --host urup
--port 8432 --username postgres --dbname iserver --clean -
-verbose "C:\Actuate\iServer\encyc_backup\iserver.backup"
```

- 5 On Backup Database iserver, choose Done.

Restoring an Encyclopedia volume using pg_restore

Alternatively, you can restore an Encyclopedia volume schema using the command-line version of pg_restore. The following example duplicates the restore operations performed in the previous section using the graphical PostgreSQL administration tool, pgAdmin. You do not need to do both activities.

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` from the command line by performing the following tasks.

How to run `pg_restore` from a command prompt

1 Open a command prompt.

2 Navigate to the following location:

```
C:\Program Files\Actuate11\iServer\postgresql\bin
```

3 Enter the following command. Substitute your machine name for `urup` in this example:

```
pg_restore.exe --host urup --port 8432 --username postgres --
             dbname iserver --clean --verbose "C:\Actuate\iServer\
             encyc_backup\iserver.backup"
```

Press Enter.

4 Type the `postgres` superuser password. The administrator specified this password in PostgreSQL Database Information during the `iServer` installation procedure in Exercise 1. Press Enter.

`pg_restore` executes, writing status messages to the command prompt.

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.

2 On Volumes, take the volume online, as shown in Figure 10-23.




Figure 10-23 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 folders to the backup folder.

How to create a backup of the Encyclopedia volume metadata

- 1 In Windows, open a command prompt.
- 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 'C:\Actuate\iServer\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
Species 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 `acsserverconfig.xml` file and volume data directories to the backup directory by performing the following tasks.

How to back up the volume data directories

- 1 Open Windows Explorer and 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:

```
C:\Actuate11\iServer\data
```

- 2 In `AC_DATA_HOME`, navigate to the config folder that contains `acsserverconfig.xml` file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP4
```

Copy `acsserverconfig.xml` to the following backup location:

```
C:\Actuate\iServer\encyc_backup
```

- 3 In `AC_DATA_HOME\encyc`, select the file, `fileType`, `status`, and `tempRov` directories, right-click, and choose Copy to copy these directories to the following backup location:

```
C:\Actuate\iServer\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 `acsserverconfig.xml` file and volume data directories in `AC_DATA_HOME`.
- Copy the backed-up `acsserverconfig.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 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume offline, as shown in Figure 10-24.




Figure 10-24 Taking the volume offline

How to restore the backed-up volume data directories

- 1 In Windows Explorer, navigate to AC_DATA_HOME.
- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:
C:\Actuate\iServer\data\config\11SP4
Select acserverconfig.xml, right-click, and choose Delete. Confirm the deletion.
- 3 In AC_DATA_HOME, navigate to the encyc directory.
In AC_DATA_HOME\encyc, select the file, fileType, status, and tempROV directories, right-click, then choose Delete. Confirm the deletion.
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 In Windows Explorer, navigate to the following backup directory location:
C:\Actuate\iServer\encyc_backup
- 5 From the backup directory location, perform the following tasks:
 - 1 Select acserverconfig.xml, right-click, choose Copy, and copy this file to AC_DATA_HOME\config\11SP4.
 - 2 In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempROV directories, right-click, choose Copy, and copy these 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 In Windows, open a command prompt.

- 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

- Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- On Volumes, take the volume online, as shown in Figure 10-25.




Figure 10-25 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

Backing up and restoring an Encyclopedia volume that uses a SQL Server database

Backing up and restoring an Encyclopedia volume

To back up an Encyclopedia volume in the SQL ServerRDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the SQL Server Management Studio graphical administration tool or the `sqlcmd` command-line utility
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

To restore an Encyclopedia volume in the SQL Server RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume data and configuration files using operating system copy commands
- Backs up the tail of the transaction log using the SQL Server Management Studio graphical administration tool or the sqlcmd command-line utility
- Restores Encyclopedia volume metadata using the SQL Server Management Studio graphical administration tool or the sqlcmd command-line utility

The following sections show how to perform these backup and restore operations.

Backing up an Encyclopedia volume using SQL Server Management Studio

To back up an Encyclopedia volume, perform the following tasks:

- Create a folder to contain the volume data backup files.
- Back up Encyclopedia volume metadata using SQL Server Management Studio.
- Back up the acserverconfig.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 Open Windows Explorer.
- 2 In Windows Explorer, choose File→New→Folder repeatedly to create a set of nested folders in the following location:

C:\Actuate\iServer\encyc_backup

Figure 10-26 shows the new folder.




Figure 10-26 Creating a new folder named `encyc_backup`

Back up Encyclopedia volume metadata using the SQL Server graphical administration tool, SQL Server Management Studio, by performing the following tasks.

How to create a backup of the Encyclopedia volume metadata using SQL Server Management Studio

- 1 Choose Start → Programs → Microsoft SQL Server 2008 → SQL Server Management Studio.
- 2 On Connect to Server, in Server name, type your machine name, such as `urup`, as shown in Figure 10-27. Choose Connect.




Figure 10-27 Connecting to a machine named `urup`

- 3 On Microsoft SQL Server Management Studio, in Object Explorer, expand Databases. Right-click the database named `iserver` and choose Tasks → Back Up, as shown in Figure 10-28.

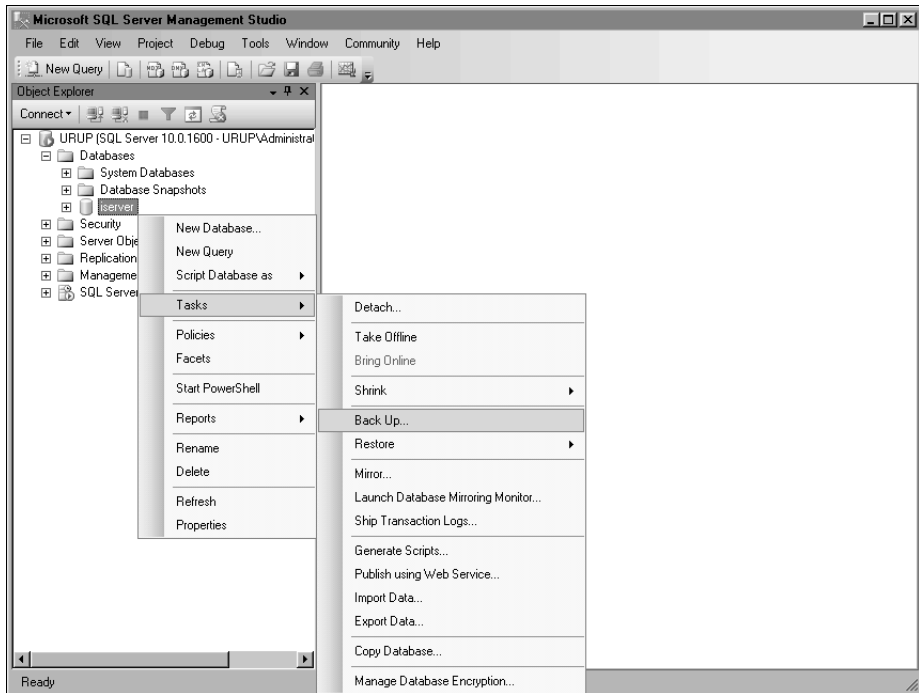


Figure 10-28 Choosing to back up the database named iserver

- 4 On Back Up Database—General, perform the following tasks:
 - 1 In the Source section, in Database, select iserver.
 - 2 In Backup type, select Full.
 - 3 Leave Copy Only Backup unselected.
 - 4 For Backup component, select Database.
 - 5 In the Backup set section, in Name, accept the default value as the name for the backup set. Alternatively, type a different name.
 - 6 In Description, type an optional description for the backup set.
 - 7 For Backup set will expire, in After, accept the default value of 0 to specify that the backup set will not expire and cannot be overwritten. Alternatively, specify a number of days after which the backup set expires and can be overwritten.
Optionally, select On to specify an expiration date for the backup set.
 - 8 In the Destination section, perform the following tasks:
 - 1 For Back up to, accept the default value of Disk.

- 2 Choose Remove to delete the default backup destination.
- 3 Choose Add.
- 4 On Select Backup Destination, as shown in Figure 10-29, in File name, type:

C:\Actuate\iServer\encyc_backup\iserver.bak




Figure 10-29 Typing the backup destination

On Select Backup Destination, choose OK. Back Up Database—General appears as shown in Figure 10-30. Choose OK.

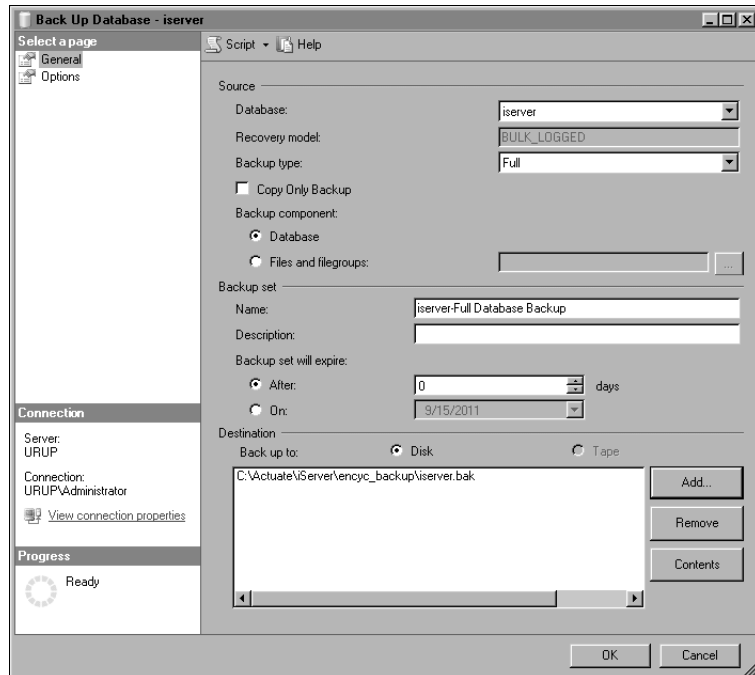


Figure 10-30 Viewing Back Up Database with new backup location

- 5 Select Back Up Database—Options and perform the following tasks:
 - 1 In the Overwrite media section, accept the default option, Back up to the existing media set. For Back up to the existing media set, accept the default option, Append to the existing backup set.
 - 2 In the Reliability section, accept the default value of unselected for all options. Alternatively, select any combination of these options.
 - 3 In the Compression section, for Set backup compression, accept the default value of Use the default server setting. The default value is Do not compress the backup.

Back Up Database—Options appears as shown in Figure 10-31.

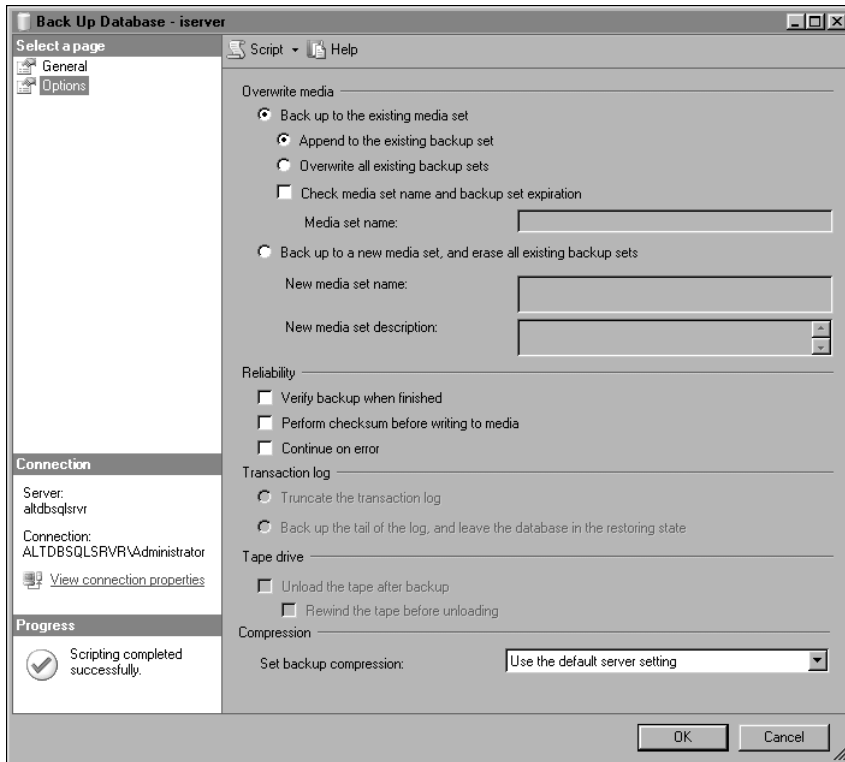


Figure 10-31 Accepting the default values on Back Up Database—Options

- 4 If you want SQL Server Management Studio to generate and display the Transact-SQL statements that the backup operation executes, choose the arrow next to Script then select Script Action to New Query Window, as shown in Figure 10-32.




Figure 10-32 Choosing to display Transact-SQL statements

- 5 Choose OK.
- 6 In the dialog box stating that the backup completed successfully, choose OK, as shown in Figure 10-33.




Figure 10-33 Acknowledging the successful backup operation

Alternatively, you can also back up the Encyclopedia volume metadata using the `sqlcmd` utility. The following example duplicates the operations performed in the previous section using the graphical SQL Server administration tool, SQL Server Management Studio. You do not need to do both activities.

How to create a backup of the Encyclopedia volume metadata using the `sqlcmd` utility from a command prompt

- 1 Open a command prompt.
- 2 To connect with the default instance of SQL Server running on your machine, type the following command and press Enter:

```
sqlcmd
```

- 3 At the `sqlcmd` prompt, type the following command and press Enter:

```
BACKUP DATABASE [iserver] TO DISK = N'C:\Program Files
  \Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\Backup
  \iserver.bak' WITH NOFORMAT, NOINIT, NAME = N'iserver-Full
  Database Backup', SKIP, NOREWIND, NOUNLOAD, STATS = 10
```

At the `sqlcmd` prompt, type the following command:

```
GO
```

The command prompt appears as shown in Figure 10-34.

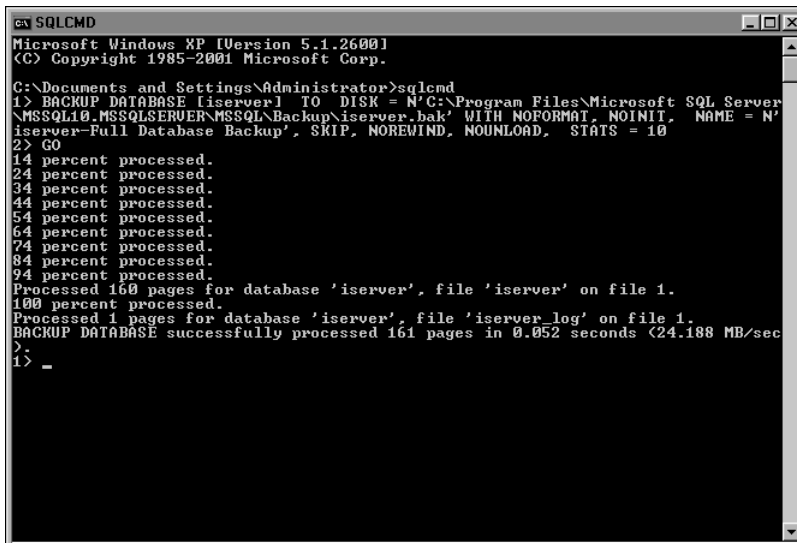


```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Administrator>sqlcmd
1> BACKUP DATABASE [iserver1] TO DISK = N'C:\Program Files\Microsoft SQL Server
\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH NOFORMAT, NOINIT, NAME = N'
iserver-Full Database Backup', SKIP, NOREWIND, NOUNLOAD, STATS = 10
2> GO_
```

Figure 10-34 Typing the command to back up the metadata

Press Enter to execute the command to back up the database. The output appears as shown in Figure 10-35.



```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Administrator>sqlcmd
1> BACKUP DATABASE [iserver1] TO DISK = N'C:\Program Files\Microsoft SQL Server
\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH NOFORMAT, NOINIT, NAME = N'
iserver-Full Database Backup', SKIP, NOREWIND, NOUNLOAD, STATS = 10
2> GO
14 percent processed.
24 percent processed.
34 percent processed.
44 percent processed.
54 percent processed.
64 percent processed.
74 percent processed.
84 percent processed.
94 percent processed.
Processed 160 pages for database 'iserver', file 'iserver' on file 1.
100 percent processed.
Processed 1 pages for database 'iserver', file 'iserver log' on file 1.
BACKUP DATABASE successfully processed 161 pages in 0.052 seconds (24.188 MB/sec)
>
1> _
```

Figure 10-35 Executing the command to back up the metadata

- 4 To end the sqlcmd session, type the following command at the sqlcmd prompt:

EXIT

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 Open Windows Explorer and navigate to AC_DATA_HOME. This is 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:

C:\Actuate11\iServer\data

- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:

C:\Actuate\iServer\data\config\11SP4

Select acserverconfig.xml, right-click, and choose Copy.

Copy this file to the following backup location:

C:\Actuate\iServer\encyc_backup

- 3 In AC_DATA_HOME, navigate to the encyc folder. Select the file, fileType, status, and tempRov directories, right-click, and choose Copy. Copy these directories to the following backup location:

C:\Actuate\iServer\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.

The contents of the backup folder appear as shown in Figure 10-36.

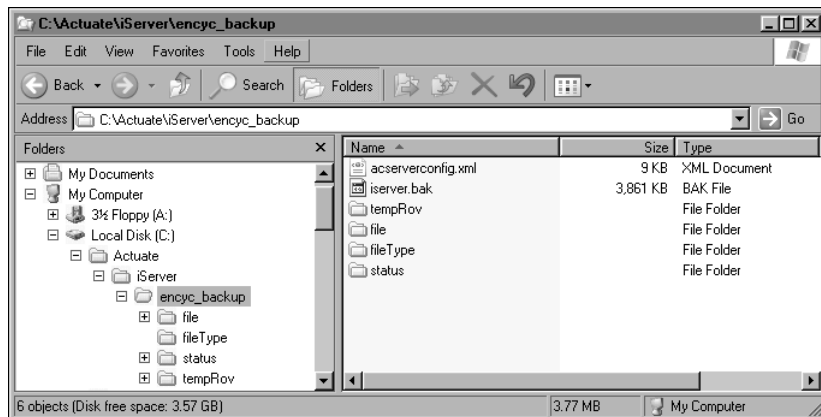


Figure 10-36 Viewing the backed-up files

Restoring an Encyclopedia volume using SQL Server Management Studio

To restore a backed-up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia 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.
- Back up the tail of the transaction log.
- Restore the Encyclopedia volume metadata using SQL Server Management Studio, or execute the appropriate commands using sqlcmd from a command prompt.
- Take the Encyclopedia online.

Take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume offline, as shown in Figure 10-37.




Figure 10-37 Taking the volume offline

How to restore the backed-up volume data folders

- 1 In Windows Explorer, navigate to AC_DATA_HOME.
- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:
C:\Actuate\iServer\data\config\11SP4
Select acserverconfig.xml, right-click, and choose Delete. Confirm the deletion.
- 3 In AC_DATA_HOME, navigate to the encyc folder.
In AC_DATA_HOME\encyc, select the file and fileType, status, and tempRov folders, right-click, then choose Delete. Confirm the deletion.

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.

- 4 In Windows Explorer, navigate to the following location:

C:\Actuate\iServer\encyc_backup

- 5 Right-click acserverconfig.xml, choose Copy, and copy this file to AC_DATA_HOME\config\11SP4. In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempRov folders, right-click, choose Copy, and copy these folders to AC_DATA_HOME\encyc.

Before restoring a full database backup, you must back up the tail of the transaction log. A backup taken of the tail of the transaction log just before a restore operation is called a tail-log backup.

How to back up the tail of the transaction log

- 1 On Microsoft SQL Server Management Studio, in Object Explorer, expand Databases. Right-click the database named iserver and choose Tasks→Back Up.
- 2 On Back Up Database—General, perform the following tasks:
 - 1 In the Source section, in Database, select iserver.
 - 2 In Backup type, select Transaction Log.
 - 3 Leave Copy Only Backup unselected.
 - 4 In the Backup set section, in Name, accept the default value as the name for the backup set. Alternatively, type a different name.
 - 5 In Description, type an optional description for the backup set.
 - 6 For Backup set will expire, in After, accept the default value of 0 to specify that the backup set will not expire and cannot be overwritten. Alternatively, specify a number of days after which the backup set expires and can be overwritten.

Optionally, select On to specify an expiration date for the backup set.
 - 7 In the Destination section, for Back up to, accept the default value of Disk.
 - 8 Accept the default backup destination.

Back Up Database—General appears as shown in Figure 10-38.

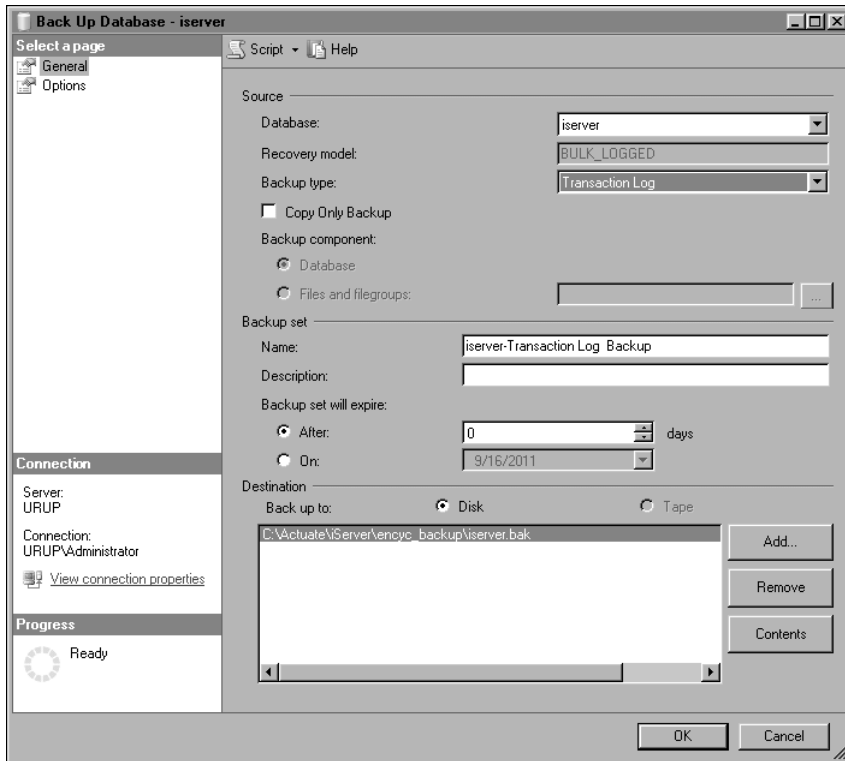


Figure 10-38 Choosing to back up the transaction log

In Select a page, choose Options.

- 3 On Back Up Database—Options, perform the following tasks:
 - 1 In the Overwrite media section, accept the default option, Back up to the existing media set. For Back up to the existing media set, accept the default option, Append to the existing backup set.
 - 2 In the Reliability section, accept the default value of unselected for all options. Alternatively, select any combination of these options.
 - 3 In the Transaction log section, select Back up the tail of the log, and leave the database in the restoring state.
 - 4 In the Compression section, in Set backup compression, accept the default value of Use the default server setting. The default value is Do not compress the backup.
 - 5 If you want SQL Server Management Studio to generate and display the Transact-SQL statements that the backup operation executes, left-click on the arrow next to Script and choose Script Action to New Query Window as shown in Figure 10-39.




Figure 10-39 Choosing to display Transact-SQL statements
Back Up Database—Options appears as shown in Figure 10-40.

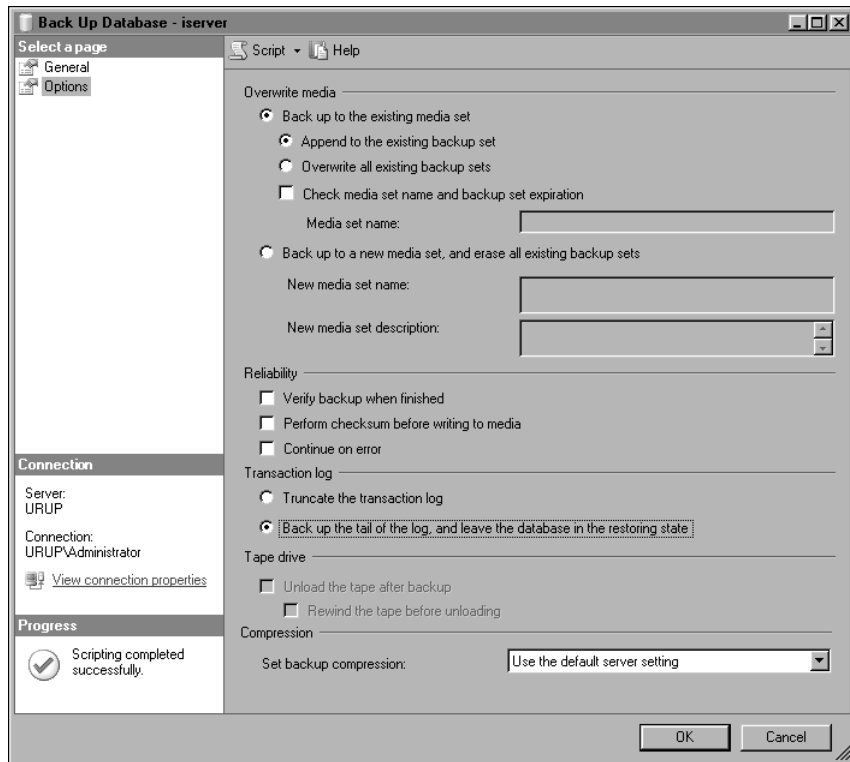


Figure 10-40 Configuring Back Up Database—Options

Choose OK.

- 4 In the dialog box stating that the backup operation completed successfully, choose OK, as shown in Figure 10-41.




Figure 10-41 Acknowledging the successful backup operation

Alternatively, you can also back up the tail of the transaction log using the `sqlcmd` utility. The following example duplicates the operations performed in the previous section using the graphical SQL Server administration tool, SQL Server Management Studio. You do not need to do both activities.

How to back up of the tail of the transaction log using `sqlcmd` from a command prompt

- 1 Open a command prompt.
- 2 To connect with the default instance of SQL Server running on your machine, type the following command and press Enter:

```
sqlcmd
```

- 3 At the `sqlcmd` prompt, type the following command and press Enter:

```
BACKUP LOG [iserver] TO DISK = N'C:\Program Files\Microsoft SQL
Server\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH
NO_TRUNCATE, NOFORMAT, NOINIT, NAME = N'iserver-Transaction
Log Backup', SKIP, NOREWIND, NOUNLOAD, NORECOVER, STATS = 10
```

At the `sqlcmd` prompt, type the following command:

```
GO
```

The command prompt appears as shown in Figure 10-42.


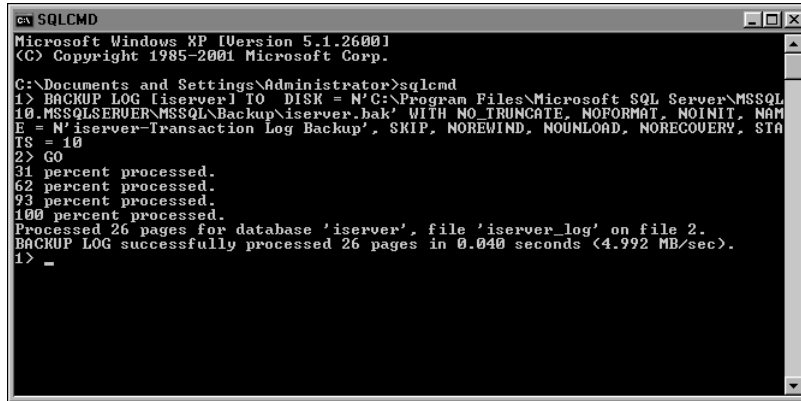


Figure 10-42 Typing the command to back up the tail of the log

Press Enter to execute the command to back up the tail of the transaction log. The output appears as shown in Figure 10-43.



```
SQLCMD
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Administrator>sqlcmd
1> BACKUP LOG [iserver] TO DISK = N'C:\Program Files\Microsoft SQL Server\MSSQL
10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH NO_TRUNCATE, NOFORMAT, NOINIT, NAM
E = N'iserver-Transaction Log Backup', SKIP, NOREWIND, NOUNLOAD, NORECOVERY, STA
TS = 10
2> GO
31 percent processed.
62 percent processed.
93 percent processed.
100 percent processed.
Processed 26 pages for database 'iserver', file 'iserver_log' on file 2.
BACKUP LOG successfully processed 26 pages in 0.040 seconds (4.992 MB/sec).
1> _
```

Figure 10-43 Executing the command to back up the tail of the log

- 4 To end the sqlcmd session, type the following command at the sqlcmd prompt:

```
EXIT
```

When restoring the Encyclopedia volume metadata, you must also restore the tail of the log backup. Restore both of these backups by performing the following tasks.

How to restore a backup of the Encyclopedia volume metadata using SQL Server Management Studio

- 1 On Microsoft SQL Server Management Studio, in Object Explorer, expand Databases. Right-click the database named iserver and choose Tasks>Restore>Database, as shown in Figure 10-44.

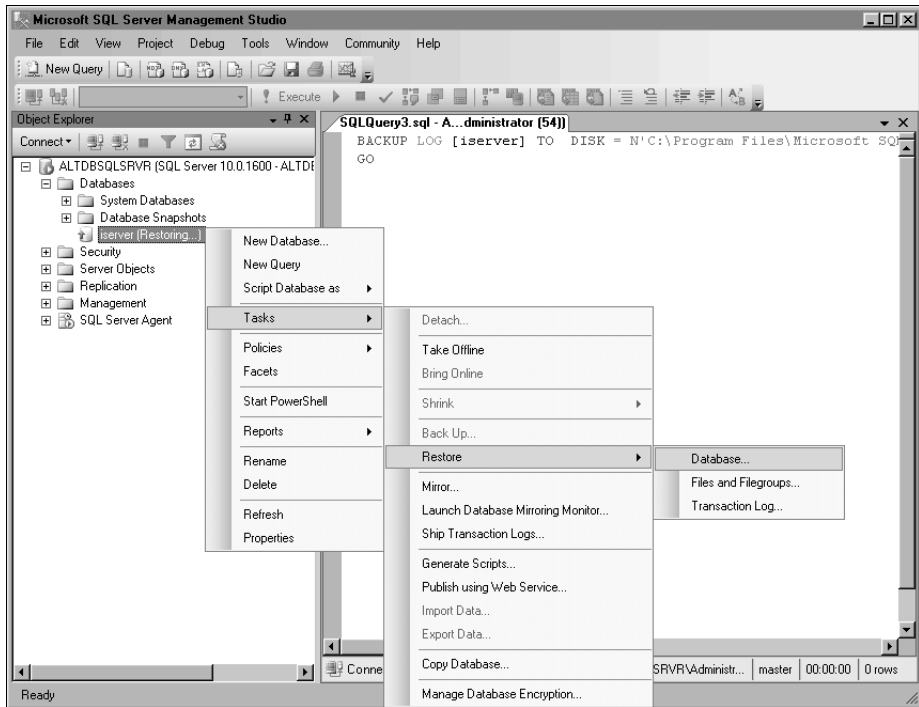


Figure 10-44 Choosing to restore the database

- 2 On Restore Database—General, perform the following tasks:
 - 1 In the Destination for restore section, in To database, select iserver.
 - 2 In To a point in time, accept the default of Most recent possible. Alternatively, click on the ellipsis next to Most recent possible. On Point in time restore, specify a specific date and time to which to restore the backup.
 - 3 In the Source for restore section, in From database, select iserver.
 - 4 In Select the backup sets to restore, select both the Full database backup and the Transaction log backup.

Restore Database—General appears as shown in Figure 10-45.

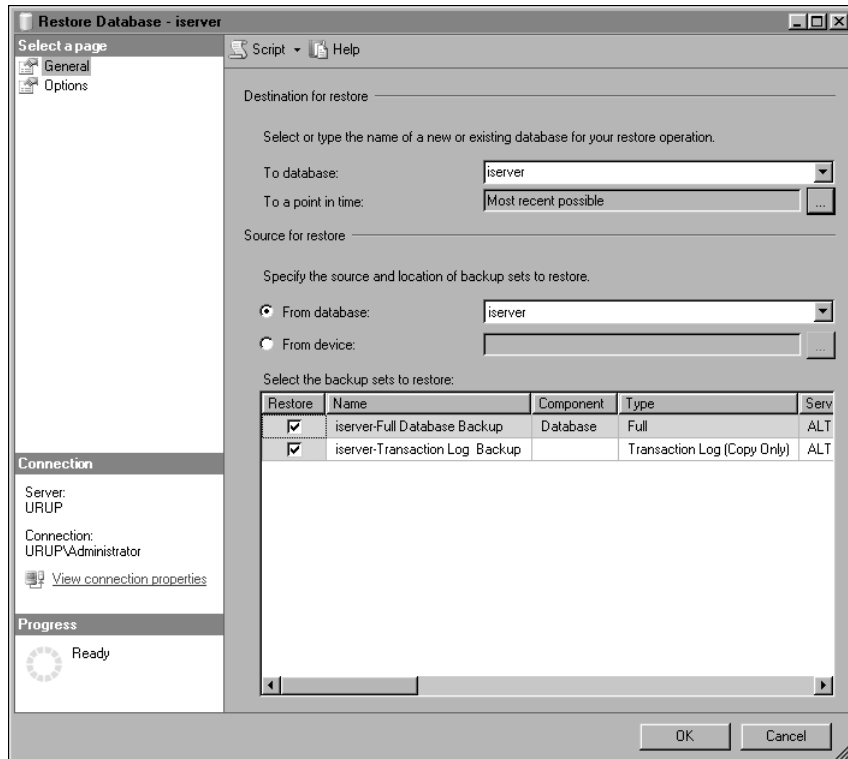


Figure 10-45 Configuring Restore Database—General

In Select a page, choose Options.

- 3 On Restore Database—Options, perform the following tasks:
 - 1 In the Restore options section, accept the default value of unselected for all options. Alternatively, select any combination of these options.
 - 2 In the Recovery state section, select Leave the database ready to use by rolling back uncommitted transactions. Alternatively, select one of the other options.
 - 3 If you want SQL Server Management Studio to generate and display the Transact-SQL statements that the backup operation executes, left-click on the arrow next to Script and choose Script Action to New Query Window.

Restore Database—Options appears as shown in Figure 10-46.

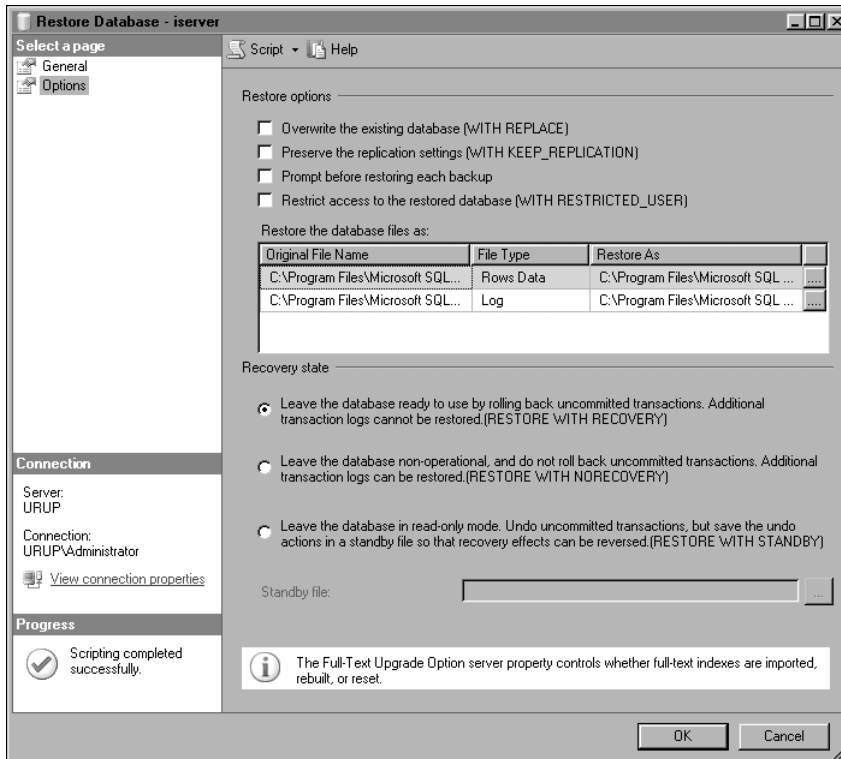


Figure 10-46 Configuring Restore Database—Options

Choose OK.

In the dialog box stating that the restore operation completed successfully, choose OK, as shown in Figure 10-47.




Figure 10-47 Acknowledging the successful restore operation

Alternatively, you can restore an Encyclopedia volume schema using the `sqlcmd` utility. The following example duplicates the operations performed in the previous section using the graphical SQL Server administration tool, SQL Server Management Studio. You do not need to do both activities.

How to restore a backup of the Encyclopedia volume metadata using `sqlcmd` from a command prompt

- 1 Open a command prompt.

- 2 To connect with the default instance of SQL Server running on your machine, type the following command and press Enter:

```
sqlcmd
```

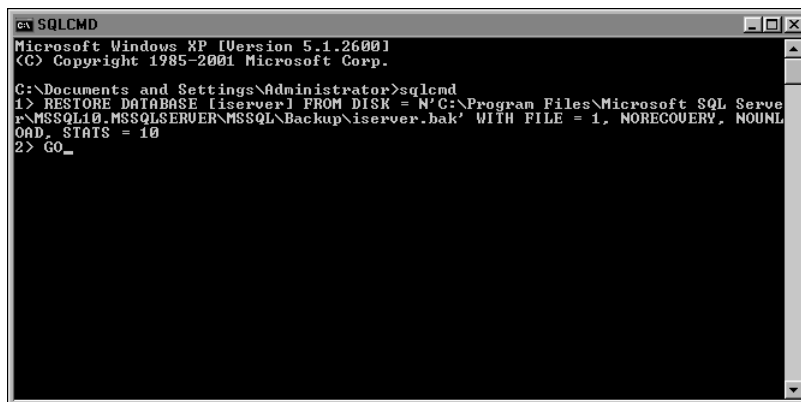
- 3 At the sqlcmd prompt, type the following command and press Enter:

```
RESTORE DATABASE [iserver] FROM DISK = N'C:\Program Files\  
Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\Backup\  
iserver.bak' WITH FILE = 1, NORECOVERY, NOUNLOAD, STATS = 10
```

At the sqlcmd prompt, type the following command:

```
GO
```

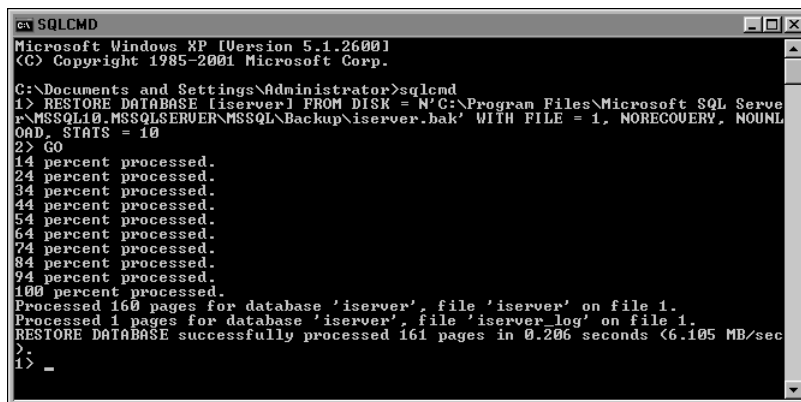
The command prompt appears as shown in Figure 10-48



```
SQLCMD  
Microsoft Windows XP [Version 5.1.2600]  
(C) Copyright 1985-2001 Microsoft Corp.  
  
C:\Documents and Settings\Administrator>sqlcmd  
1> RESTORE DATABASE [iserver] FROM DISK = N'C:\Program Files\Microsoft SQL Serve  
r\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH FILE = 1, NORECOVERY, NOUNL  
OAD, STATS = 10  
2> GO
```

Figure 10-48 Typing the command to restore the backup of the metadata

Press Enter to execute the command to restore the backup of the metadata. The output appears as shown in Figure 10-49.



```
SQLCMD  
Microsoft Windows XP [Version 5.1.2600]  
(C) Copyright 1985-2001 Microsoft Corp.  
  
C:\Documents and Settings\Administrator>sqlcmd  
1> RESTORE DATABASE [iserver] FROM DISK = N'C:\Program Files\Microsoft SQL Serve  
r\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH FILE = 1, NORECOVERY, NOUNL  
OAD, STATS = 10  
2> GO  
14 percent processed.  
24 percent processed.  
34 percent processed.  
44 percent processed.  
54 percent processed.  
64 percent processed.  
74 percent processed.  
84 percent processed.  
94 percent processed.  
100 percent processed.  
Processed 160 pages for database 'iserver', file 'iserver' on file 1.  
Processed 1 pages for database 'iserver', file 'iserver_log' on file 1.  
RESTORE DATABASE successfully processed 161 pages in 0.206 seconds (6.105 MB/sec  
>  
1> _
```

Figure 10-49 Executing the command to restore the backup of the metadata

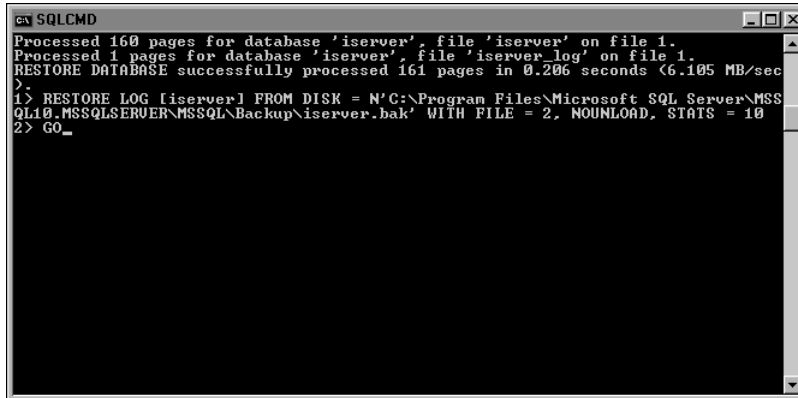
- 4 To restore the tail of the transaction log backup, type the following command at the sqlcmd prompt and press Enter:

```
RESTORE LOG [iserver] FROM DISK = N'C:\Program Files\Microsoft
  SQL Server\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH
  FILE = 2, NOUNLOAD, STATS = 10
```

At the sqlcmd prompt, type the following command:

```
GO
```

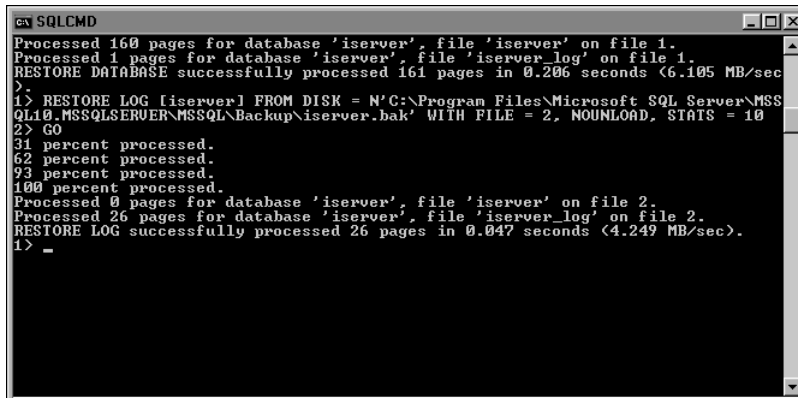
The command prompt appears as shown in Figure 10-50.



```
SQLCMD
Processed 160 pages for database 'iserver', file 'iserver' on file 1.
Processed 1 pages for database 'iserver', file 'iserver_log' on file 1.
RESTORE DATABASE successfully processed 161 pages in 0.206 seconds (6.105 MB/sec)
>
1> RESTORE LOG [iserver] FROM DISK = N'C:\Program Files\Microsoft SQL Server\MSS
QL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH FILE = 2, NOUNLOAD, STATS = 10
2> GO_
```

Figure 10-50 Typing the command to restore the tail of the log backup

Press Enter to execute the command to restore the tail of the log backup. The output appears as shown in Figure 10-51.



```
SQLCMD
Processed 160 pages for database 'iserver', file 'iserver' on file 1.
Processed 1 pages for database 'iserver', file 'iserver_log' on file 1.
RESTORE DATABASE successfully processed 161 pages in 0.206 seconds (6.105 MB/sec)
>
1> RESTORE LOG [iserver] FROM DISK = N'C:\Program Files\Microsoft SQL Server\MSS
QL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH FILE = 2, NOUNLOAD, STATS = 10
2> GO
31 percent processed.
62 percent processed.
93 percent processed.
100 percent processed.
Processed 0 pages for database 'iserver', file 'iserver' on file 2.
Processed 26 pages for database 'iserver', file 'iserver_log' on file 2.
RESTORE LOG successfully processed 26 pages in 0.047 seconds (4.249 MB/sec).
1> _
```

Figure 10-51 Executing the command to restore the tail of the log backup

- 5 To end the sqlcmd session, type the following command at the sqlcmd prompt:

EXIT

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On the simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume online, as shown in Figure 10-52.




Figure 10-52 Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the Microsoft SQL Server database utilities, go to the following location:

<http://msdn.microsoft.com/en-us/library/ms189621.aspx>

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 the DB2 Control Center graphical administration tool or the DB2 command line processor utility
- 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 the SQL Server Control Center graphical administration tool or the DB2 command line processor utility

The following sections show how to perform these backup and restore operations.

Backing up an Encyclopedia volume using DB2 Control Center

To back up an Encyclopedia volume, perform the following tasks:

- Create a folder to contain the backup files.
- Back up Encyclopedia volume metadata using DB2 Control Center.
- Back up the acserverconfig.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 Open Windows Explorer.
- 2 In Windows Explorer, choose File → New → Folder repeatedly to create a set of nested folders in the following location:

C:\Actuate\iServer\encyc_backup

Figure 10-53 shows the new folder.

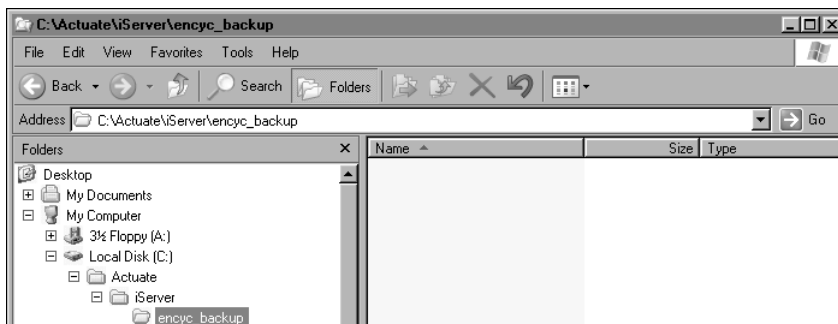


Figure 10-53 Creating a new folder named encyc_backup

Back up Encyclopedia volume metadata using the DB2 graphical administration tool, Control Center, by performing the following tasks.

How to create a backup of the Encyclopedia volume metadata using DB2 Control Center

- 1 Choose Start→Programs→IBM DB2→DB2COPY1 (Default)→General Administration Tools→Control Center.
- 2 On Control Center, in Object View, navigate to the Databases folder. Select the ISERVER database, right-click, and choose Backup, as shown in Figure 10-54.

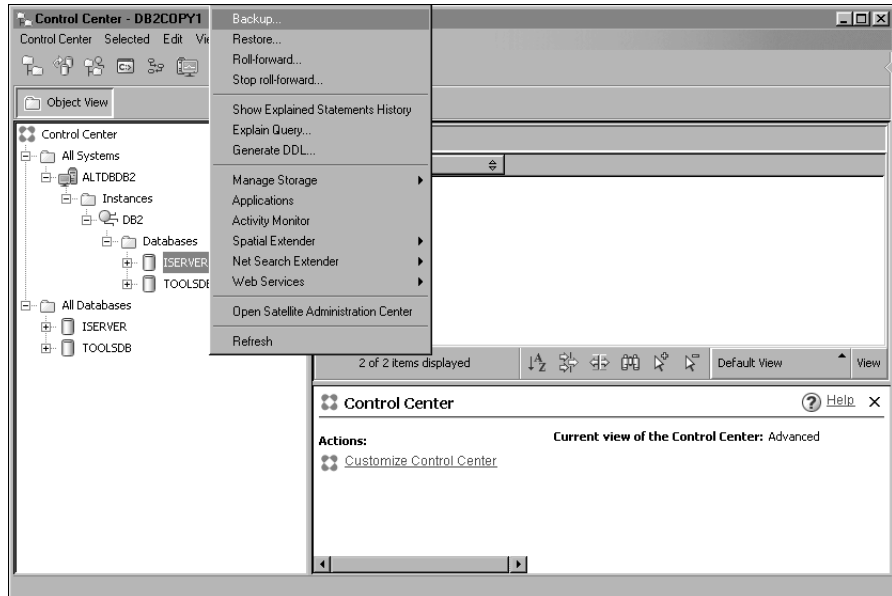


Figure 10-54 Choosing to back up the ISERVER database

- 3 On Backup Wizard—Introduction, select Backup entire database, as shown in Figure 10-55. Choose Next.
- 4 On Backup Wizard—Image, in Media Type, select File System, then select Add, as shown in Figure 10-56.
- 5 In Path Browser, navigate to the following location, as shown in Figure 10-57.
C:\Actuate\iServer\encyc_backup

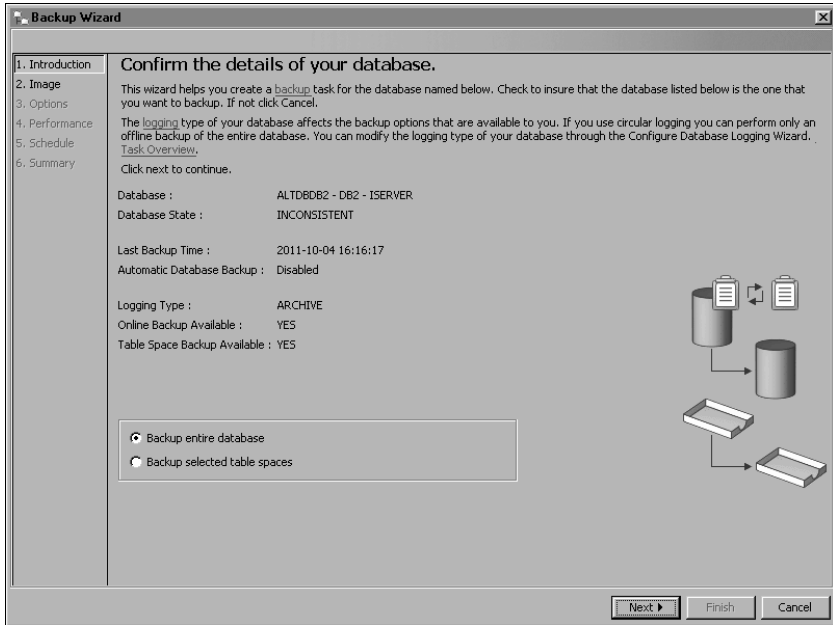


Figure 10-55 Confirming the details of the ISERVER database

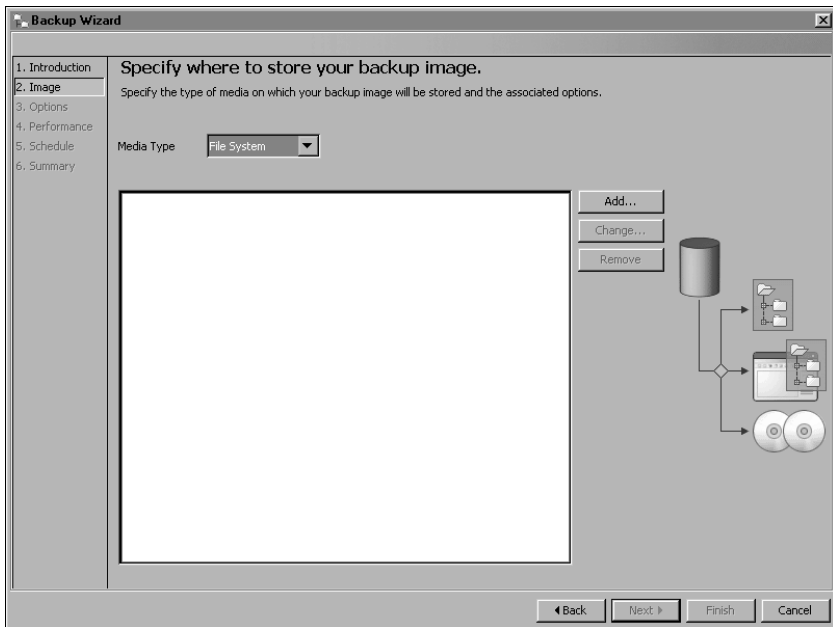


Figure 10-56 Choosing to add a backup location

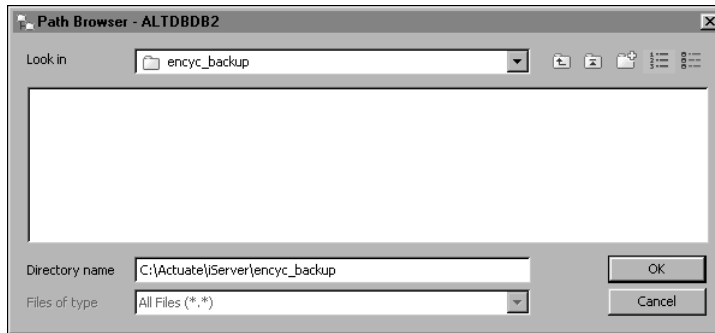


Figure 10-57 Navigating the metadata backup location I
 Choose OK. Backup Wizard—Image appears as shown in Figure 10-58.

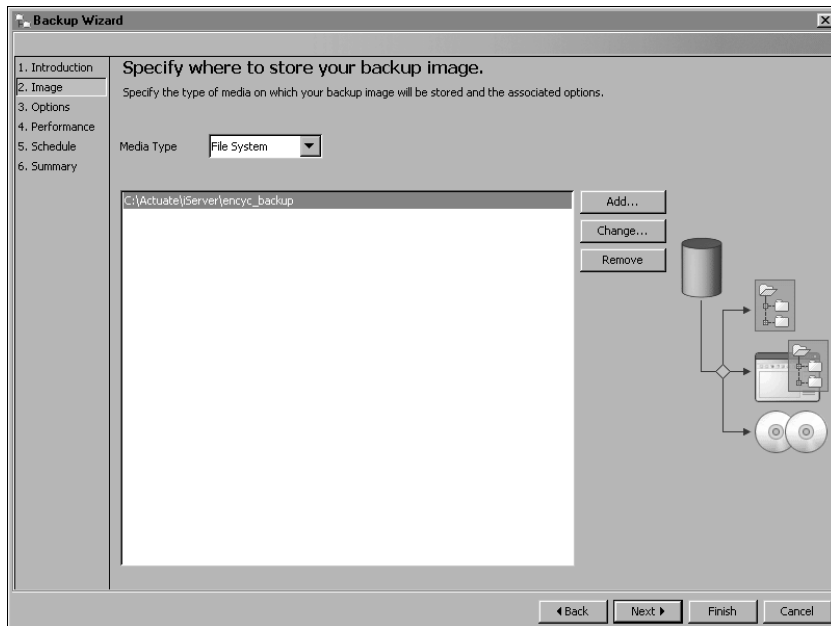


Figure 10-58 Viewing the chosen metadata backup location
 Choose Next.

- 6 On Backup Wizard—Options, perform the following tasks:
 - 1 In the Backup Type section, accept the default option, Full Backup.
 - 2 In the Availability section, select Online.
 - 3 In the Compression section, select Compress backup image.
 - 4 Backup Wizard—Options appears as shown in Figure 10-59.

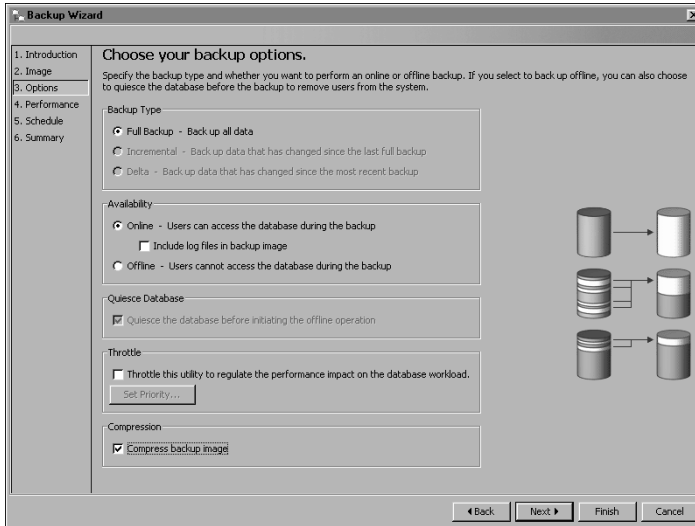


Figure 10-59 Configuring the backup options

Choose Next

- 7 On Backup Wizard—Performance, in the Buffers section, accept the recommended default values, as shown in Figure 10-60. Alternatively, select different values.

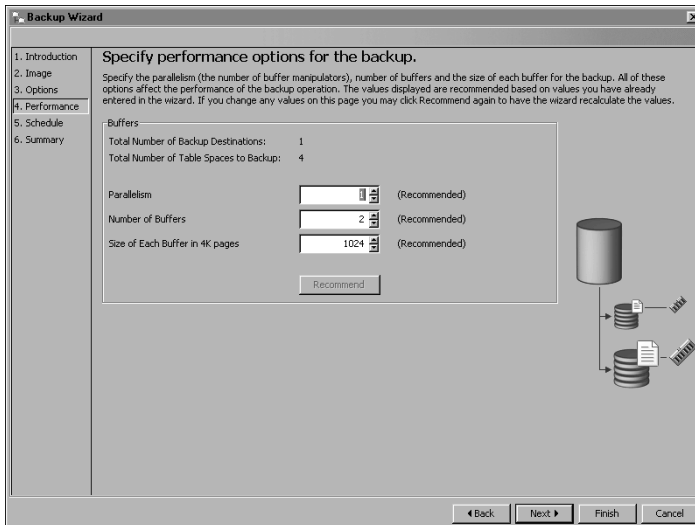


Figure 10-60 Specifying performance options

Choose Next.

- 8 On Backup Wizard—Schedule, accept the default values, as shown in Figure 10-61. Alternatively, specify different values.

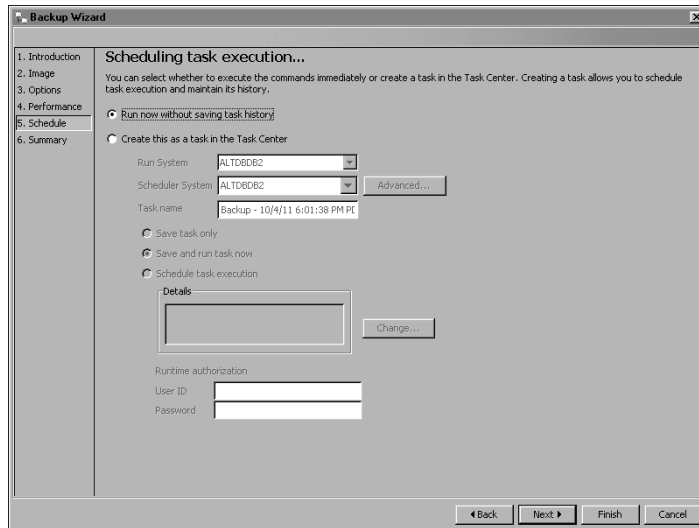


Figure 10-61 Selecting scheduling options

Choose Next.

- 9 On Backup Wizard—Summary, review the backup job parameters, as shown in Figure 10-62. Choose Back to change any backup job parameters, or choose Show Command to view the command that executes the backup.

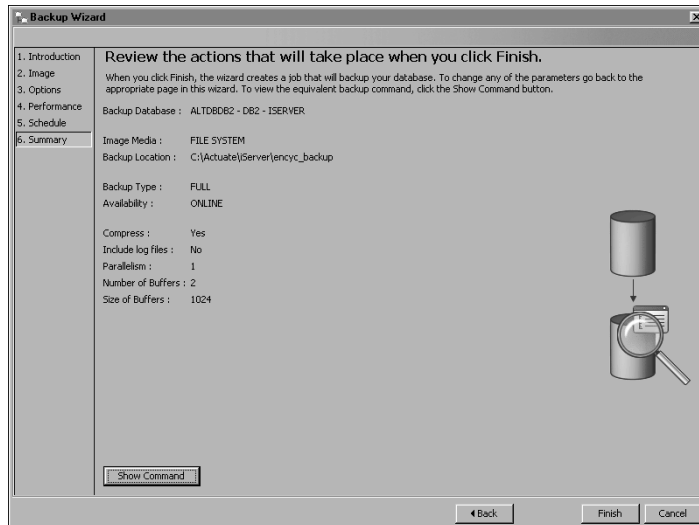


Figure 10-62 Reviewing backup job parameters

Choose Finish to execute the backup. The elapsed time for the backup job displays, as shown in Figure 10-63.



Figure 10-63 Viewing the elapsed time for the backup job

When the backup job completes, DB2 displays the commands that the job executed, and an end-of-job message, as shown in Figure 10-64.

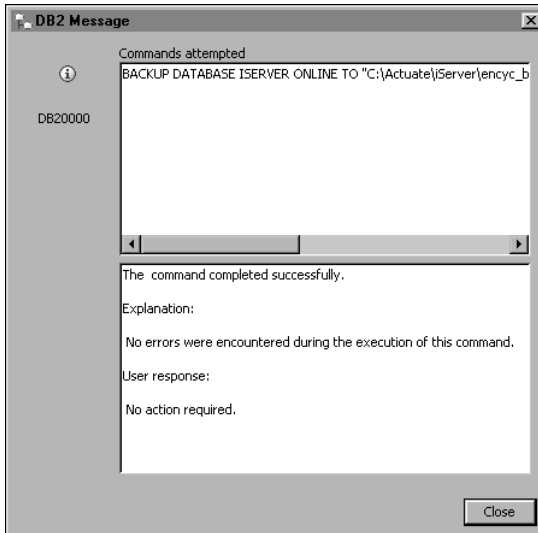


Figure 10-64 Viewing the commands executed and end-of-job messages

Alternatively, you can also back up the Encyclopedia volume metadata using the command line processor utility. The following example duplicates the operations performed in the previous section using the graphical administration tool, DB2 Control Center. You do not need to do both activities.

How to create a backup of the Encyclopedia volume metadata using the DB2 command line processor

- 1 Choose Start → Programs → IBM DB2 → DB2COPY1 (Default) → Command Line Tools → Command Line Processor.

The command line processor appears, as shown in Figure 10-65.


```

c:\DB2 CLP - DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
  ? CATALOG DATABASE for help on the CATALOG DATABASE command
  ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _

```

Figure 10-65 Opening the command line processor

- 2 Execute each of the following commands to perform an online backup of the entire database and compress the backup image, as shown in Figure 10-66:

```

BACKUP DATABASE ISERVER ONLINE
  TO "C:\Actuate\iServer\encyc_backup"
  WITH 2 BUFFERS BUFFER 1024 PARALLELISM 1
  COMPRESS WITHOUT PROMPTING

```

```

c:\DB2 CLP - DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
  ? CATALOG DATABASE for help on the CATALOG DATABASE command
  ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => BACKUP DATABASE ISERVER ONLINE TO "C:\Actuate\iServer\encyc_backup" WITH
2 BUFFERS BUFFER 1024 PARALLELISM 1 COMPRESS WITHOUT PROMPTING
Backup successful. The timestamp for this backup image is : 20111012185050

db2 => _

```

Figure 10-66 Executing the commands to back up the database

- 3 To end the DB2 command line processor session, type the following command:

```
quit
```

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 Open Windows Explorer and navigate to AC_DATA_HOME, which is 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:

C:\Actuate11\iServer\data

- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:

C:\Actuate\iServer\data\config\11SP4

Select acserverconfig.xml, right-click, and choose Copy.

Copy this file to the following backup location:

C:\Actuate\iServer\encyc_backup

- 3 Navigate to AC_DATA_HOME\encyc. Select the file, fileType, status, and tempRov directories, right-click, and choose Copy. Copy these directories to the following backup location:

C:\Actuate\iServer\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.

The contents of the backup folder appear as shown in Figure 10-67.

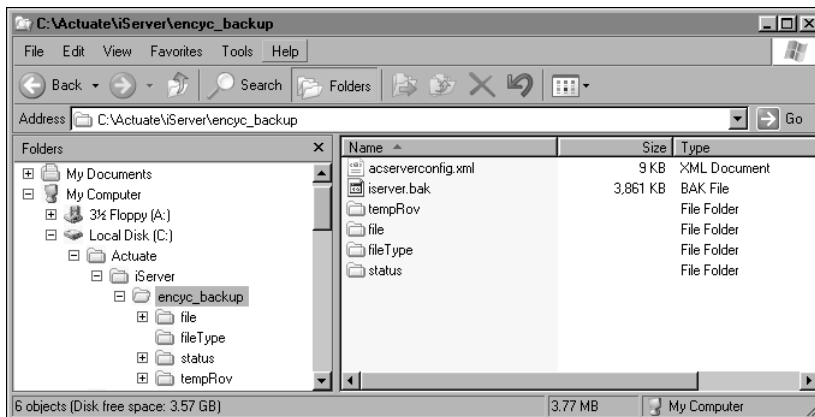


Figure 10-67 Viewing the backed-up files

Restoring an Encyclopedia volume using DB2 Control Center

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 using SQL Server Management Studio, or execute the appropriate commands using sqlcmd from a command prompt.
- Take the Encyclopedia volume online.

Take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume offline, as shown in Figure 10-68.




Figure 10-68 Taking the volume offline

How to restore the backed-up volume data folders

- 1 In Windows Explorer, navigate to AC_DATA_HOME.
- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 4, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP4
```

Select acserverconfig.xml, right-click, and choose Delete. Confirm the deletion.

- 3 Navigate to AC_DATA_HOME\encyc. Select the file and fileType, status, and tempRov folders, right-click, and choose Delete. Confirm the deletion.

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.

- 4 Navigate to the following location:

C:\Actuate\iServer\encyc_backup

Right-click acserverconfig.xml, choose Copy, and copy this file to AC_DATA_HOME\config.

- 5 In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempRov folders, right-click, choose Copy, and copy these folders to AC_DATA_HOME\encyc.

How to restore a back up of the Encyclopedia volume metadata using DB2 Control Center

- 1 On Control Center, in Object View, right-click the database named ISERVER and choose Restore, as shown in Figure 10-69.

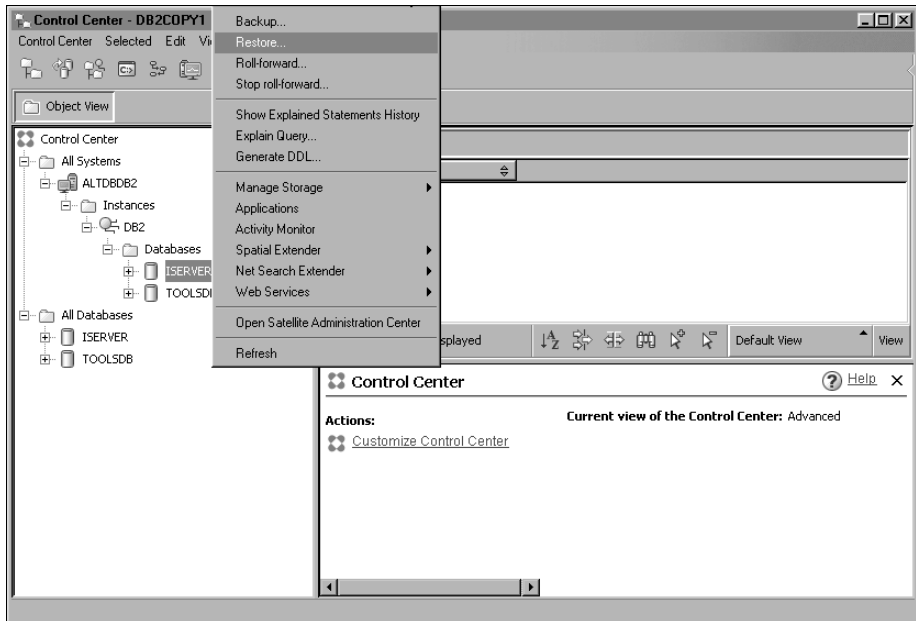


Figure 10-69 Choosing to restore the database

- 2 On Restore Data Wizard—Introduction, in the Restore alternatives section, accept the default value of Restore to an existing database, as shown in Figure 10-70. Choose Next.

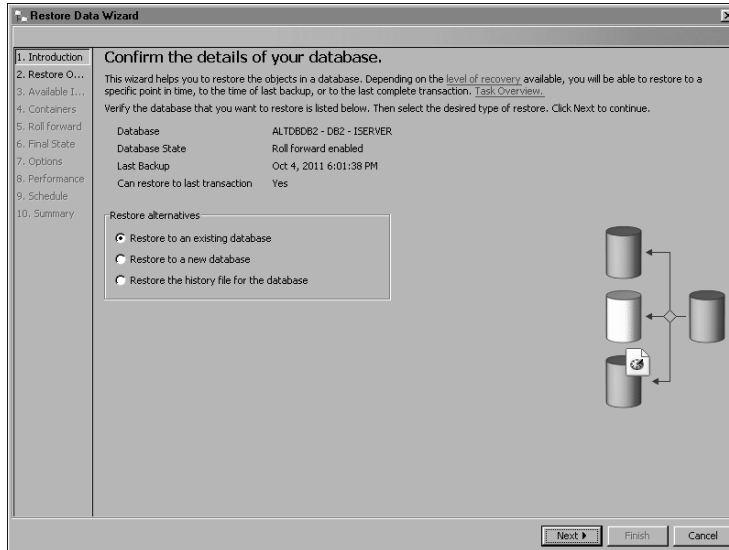


Figure 10-70 Confirming the details of the ISERVER database

- 3 On Restore Data Wizard—Restore Objects, accept the default selection, Restore the entire database, as shown in Figure 10-71. Choose Next.

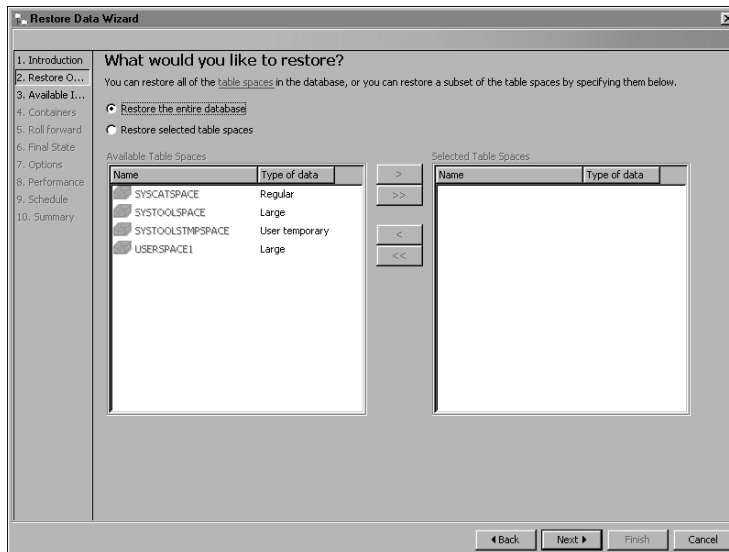


Figure 10-71 Selecting the backup image to restore

- 4 On Restore Data Wizard—Available Images, in Available backup images, select the image you want to restore, as shown in Figure 10-72.

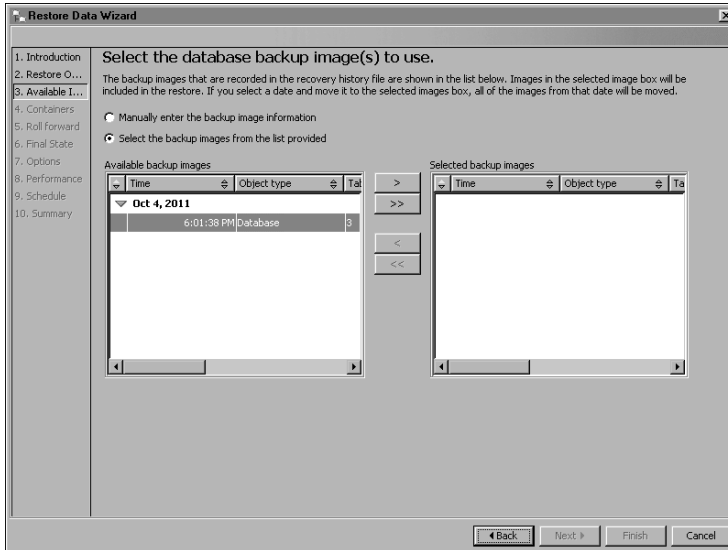


Figure 10-72 Selecting the backup image to restore

Choose the right arrow to move the image to Selected backup images, as shown in Figure 10-73. Choose Next

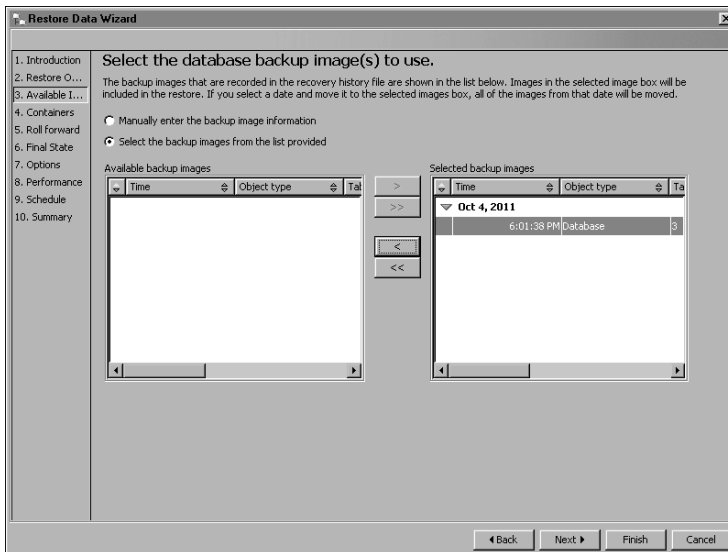


Figure 10-73 Moving the selected image to Selected backup images

- On Restore Data Wizard—Containers, accept the default options, as shown in Figure 10-74. Choose Next.

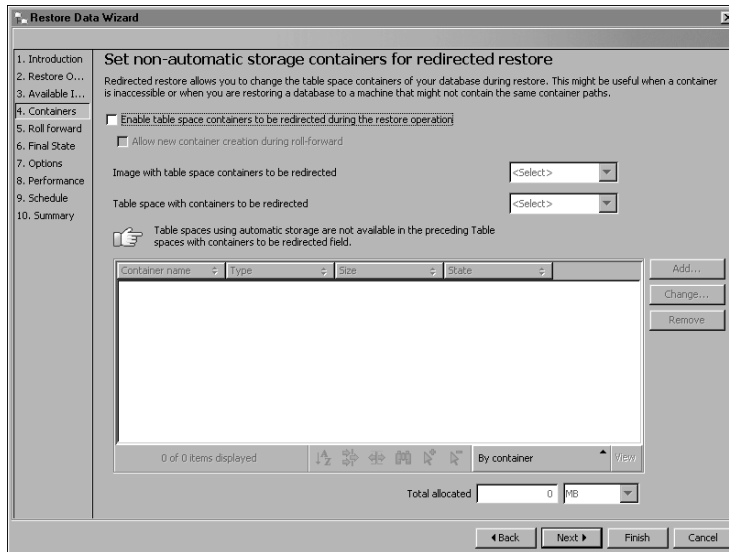


Figure 10-74 Choosing whether to enable redirecting of table space containers

- 6 On Restore Data Wizard—Roll forward, accept the default option, Restore Only, as shown in Figure 10-75. Choose Next.

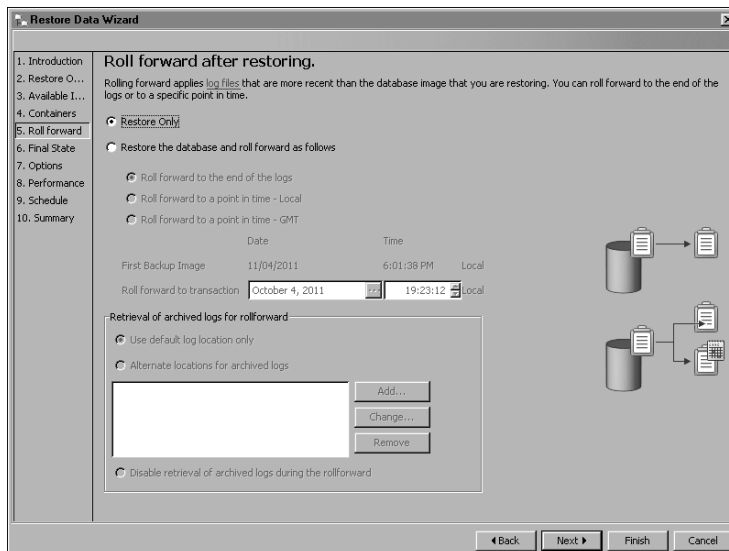


Figure 10-75 Choosing to restore the database backup only

- 7 On Restore Data Wizard—Final State, select Complete the restore and return to the active state, as shown in Figure 10-76. Choose Next.

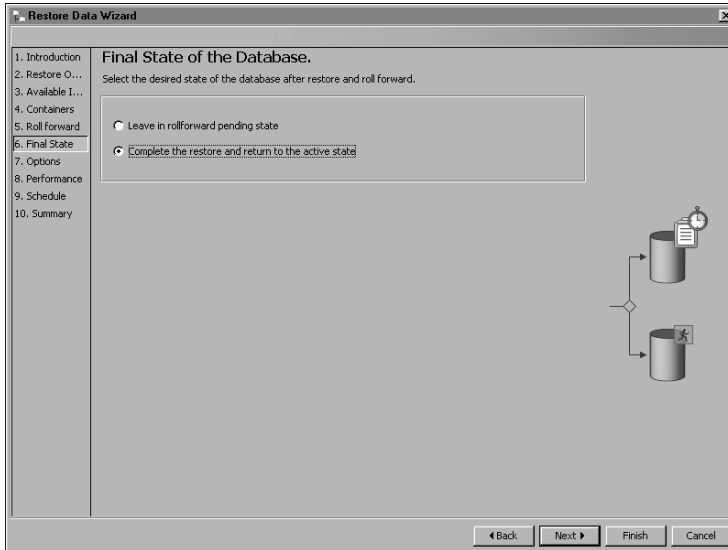


Figure 10-76 Choosing to complete the restore

- 8 On Restore Data Wizard—Options, accept the default options, as shown in Figure 10-77. Choose Next.

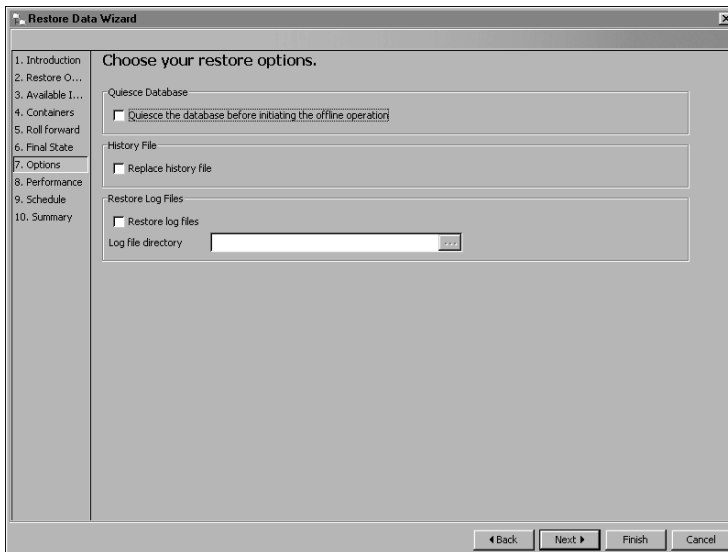


Figure 10-77 Choosing restore options

- 9 On Restore Data Wizard—Performance, accept the default values, as shown in Figure 10-78. Choose Next.

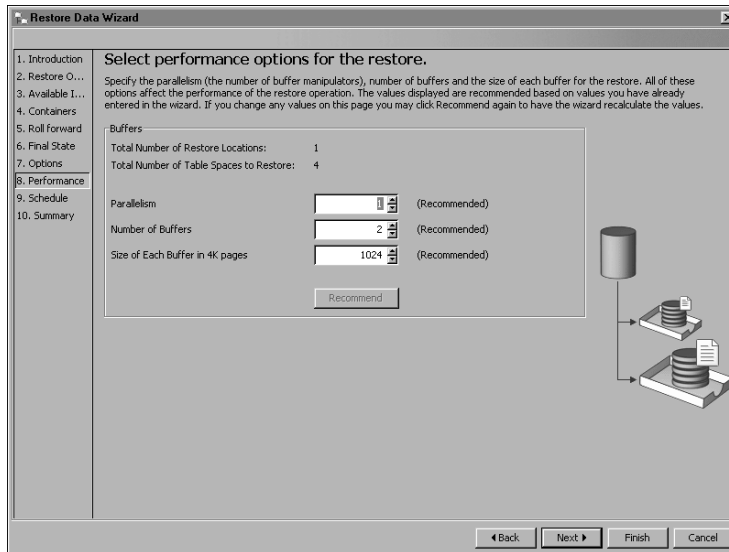


Figure 10-78 Selecting performance options

- 10** On Restore Data Wizard—Schedule, accept the default values, as shown in Figure 10-79. Choose Next.

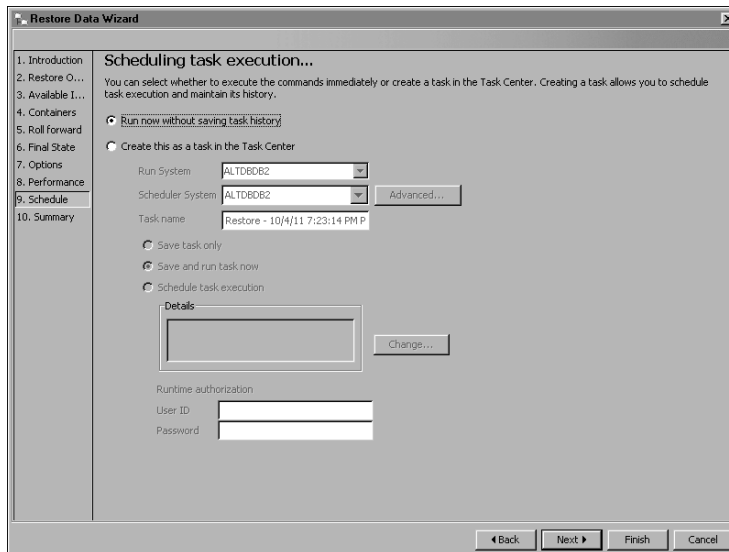


Figure 10-79 Setting backup job scheduling options

- 11** On Restore Data Wizard—Summary, review the restore job parameters, as shown in Figure 10-80. Choose Back to change any restore job parameters, or choose Show Command to view the SQL that executes the restore.

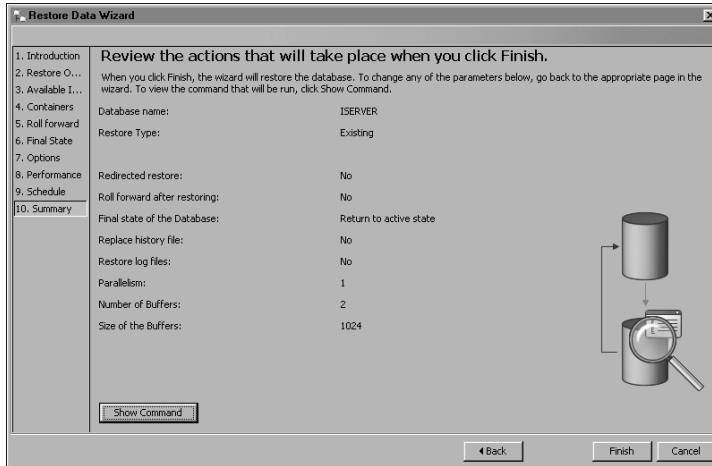


Figure 10-80 Reviewing restore job parameters

Choose Finish to execute the restore. The elapsed time for the restore job displays, as shown in Figure 10-81.



Figure 10-81 Viewing the elapsed time for the restore job

When the job completes, DB2 displays the commands that the job executed, and an end-of-job message, as shown in Figure 10-82.

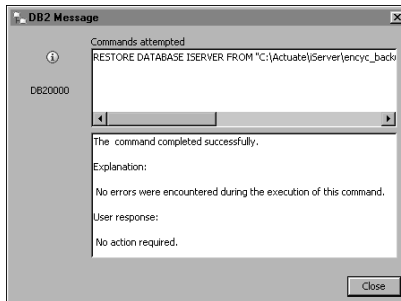


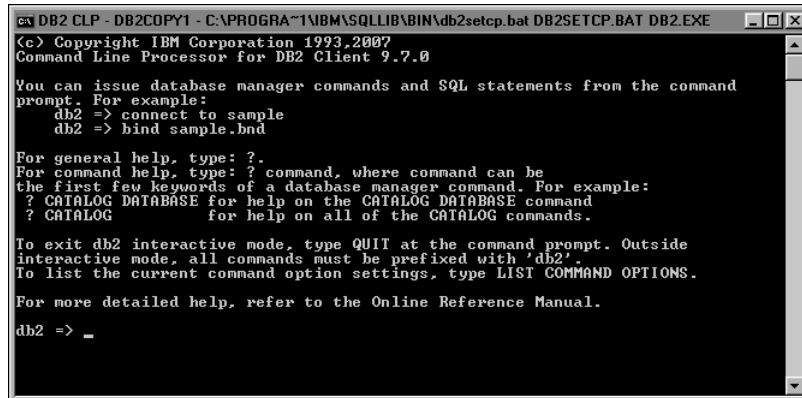
Figure 10-82 Viewing the command executed and end-of-job messages

Alternatively, you can also restore the Encyclopedia volume metadata using the command line processor utility. The following example duplicates the operations performed in the previous section using the graphical administration tool, DB2 Control Center. You do not need to do both activities.

How to restore a backup of the Encyclopedia volume metadata using the DB2 command line processor

- 1 Choose Start→Programs→IBM DB2→DB2COPY1 (Default)→Command Line Tools→Command Line Processor.

The command line processor appears, as shown in Figure 10-83.



```
DB2 CLP - DB2COPY1 - C:\PROGRAMS\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG           for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _
```

Figure 10-83 Opening the command line processor

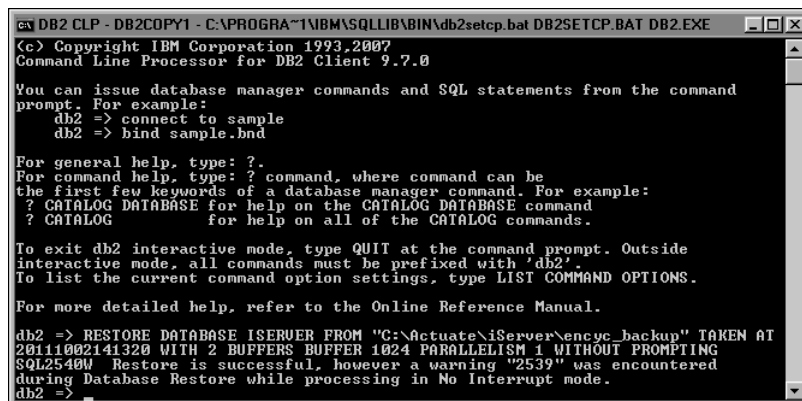
- 2 Execute the following command, as shown in Figure 10-84:

```
RESTORE DATABASE ISERVER FROM "C:\Actuate\iServer\encyc_backup"
TAKEN AT 20111002141320 WITH 2 BUFFERS BUFFER 1024
PARALLELISM 1 WITHOUT PROMPTING;
```

where

20111002141320

is the last node in the name of the iServer backup file, C:\Actuate\iServer\encyc_backup\ISERVER.0.DB2.NODE0000.CATN0000.20111002141320.001.



```
DB2 CLP - DB2COPY1 - C:\PROGRAMS\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG           for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => RESTORE DATABASE ISERVER FROM "C:\Actuate\iServer\encyc_backup" TAKEN AT
20111002141320 WITH 2 BUFFERS BUFFER 1024 PARALLELISM 1 WITHOUT PROMPTING
SQL2540M Restore is successful, however a warning "2539" was encountered
during Database Restore while processing in No Interrupt mode.

db2 => _
```

Figure 10-84 Restoring the database from the command line

- 3 To end the DB2 session, type the following command:

```
quit
```

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume online, as shown in Figure 10-85.




Figure 10-85 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 279, 282
- AC_CONFIG_HOME variable 240
- AC_DATA_HOME directory 244
- AC_DATA_HOME variable 243
- AC_JAVA_HOME variable 46, 205
- AC_JRE_HOME variable 46
- AC_JRE64_HOME variable 46
- AC_JVM_HOME variable 46
- AC_PMD_WINDOWS_CPUS variable 285, 286
- AC_SERVER_HOME parameter 154, 157, 171
- AC_SERVER_HOME variable 146, 205, 216
- AC_TEMPLATE_NAME variable 198
- access permissions. *See* privileges
- accessing
 - configuration parameters 13
 - data 13, 276
 - Encyclopedia volumes 46, 274, 277
 - help files x, xi
 - help topics xi
 - iServer features 46, 274
 - iServer System components 41
 - license files 13
 - metadata database schemas 4
 - online help 267
 - pgAdmin administration tool 41, 103
 - PostgreSQL servers 41
 - RDBMS database documentation 7
 - report files 24
 - resources 255
 - Update tool 264
 - Visual C++ libraries 203
 - volume databases 18
 - web-based content 276, 277
- accounts
 - cluster configurations and 25
 - creating user 24, 25, 52
 - installing alternate databases and 60, 75
 - installing cluster nodes and 190
 - installing Information Console and 250
 - managing 16
 - setting passwords for 31, 33
 - setting up e-mail 290
 - starting iServer 11 service and 33
 - starting Management Console and 34
 - upgrades and 82, 83, 92, 94
- AcEncycUpgrade utility 18
- AcExport utility 16
- AcExtern utility 17
- AcImport utility 16
- AcIntern utility 17
- AcMode utility 17
- acpmdconfig.xml 13, 176
- Acrobat Catalog. *See* Adobe Acrobat Catalog
- acserverconfig.xml 13, 176
- acserverlicense.xml 13
- activity logs 255
- AcToc utility 17
- Actuate Analytics Option 275
- Actuate Basic reports 276
- Actuate Customer Support xi
- Actuate Distribution xi
- Actuate ftp site 27
- Actuate licensing options 274, 275, 283
- Actuate Query Option 275
- Actuate Support Lifecycle Policy xii
- actuate.com vii
- ActuateBIRTiServer.zip 204
- ActuateInformationConsole.exe 247
- ActuateInformationConsole.war 254
- ActuateiServerIntegrationTechnology.exe 260
- ActuateOne technology vii
- AcVerify utility 17
- Add Partition link 145
- Add Volume link 151
- Add volume schema setting 146
- adding
 - administrator accounts 24, 83
 - backup folders 298, 318, 338
 - cluster nodes 176
 - clusters 181, 237
 - desktop shortcuts 41, 102, 118, 210
 - e-mail accounts 290
 - Encyclopedia volumes 27, 151, 166, 168, 230

- adding (*continued*)
 - external databases 221
 - indexes 5
 - JDBC drivers 18
 - license keys 278
 - licensing options 283
 - metadata databases 11, 223
 - partitions 145, 229
 - passwords. *See* passwords
 - PostgreSQL servers 42
 - resource groups 46
 - schema owners 52, 225, 227
 - system databases 53, 60
 - system passwords 242
 - system schemas 27, 52, 168, 170, 225
 - user accounts 25, 52
 - volume schemas 27, 52, 146, 226
- admineds.bat file 150
- administrate_encyclopedia_data_store.bat file 150
- administrate_system_data_store.bat file 169
- administration tools 16
- administrative reports 19
- administrator accounts 24, 34, 75, 83
- administrators
 - backing up database schemas and 26
 - backing up Encyclopedia and 294, 295
 - configuring iServer clusters and 13, 14
 - customizing database connections and 18
 - deploying iServer images and 202
 - installing DB2 databases and 60, 61
 - installing iServer and 24, 44, 47, 52
 - installing Oracle databases and 55
 - installing Postgre databases and 53
 - installing SQL Server databases and 57
 - managing Encyclopedia and 5, 7
 - managing iServer System and 15
 - migrating to current release and 47
 - obtaining licenses and 274
 - optimizing iServer System and 13, 176
 - preventing data loss and 7
 - recreating systems schemas and 168, 170
 - sharing metadata databases and 5
 - testing new releases and 47, 48
 - upgrading Encyclopedia and 143
 - upgrading iServer and 82, 83, 104, 124
- adminsds.bat script 169
- Adobe Acrobat Catalog utility xi
- aggregation 275
- alternative databases. *See* metadata databases; third-party databases
- Analytics Option 275
- analyzing
 - data 275
 - search results 276
- Apache Tomcat servers 247, 249
- application pages. *See* web pages
- application programming interfaces 15, 18
- application server cache 257
- application server configurations 247, 254
- application servers 247, 254, 257
 - See also* servers
- APPLICATION_USER parameter 171
- APPLICATION_USER_NAME parameter 158
- APPLICATION_USER_PASSWORD parameter 158
- applications
 - accessing Encyclopedia and 46
 - backward compatibility with 18
 - creating production environment for 48
 - creating test environments for 47
 - deploying web 256
 - developing 15, 16
 - installing Open Security 29, 46
 - restricting CPU processes for 283
 - running iServer processes and 8, 10
 - upgrading and 47, 88
- archives (cloud deployments) 6, 202
- Archiving and Purging page 296
- archiving report files 296
- asynchronous Factory service 11
- autoarchive file purging 296
- auto-login scripts (pgpass) 302
- automated installation option 6
- automated installation programs 16
- automatic rollback 83
- automatic upgrades 82, 84, 87

B

- backing up
 - data files 295
 - database schemas 26, 302

- Encyclopedia volumes 294, 295, 296, 297, 300, 313, 317, 337
 - folders 303, 325, 346
 - metadata 26, 64, 295
 - PostgreSQL databases 297
 - program files 87, 107
 - report files 17
 - system schemas 7, 83, 168
 - volume databases 17
 - backup operations 7, 8, 16
 - backup utilities 17, 295
 - backward compatibility 18
 - Basic reports 276
 - batch files
 - Data Store Administrator 150
 - Data Store Upgrader 122, 123
 - iServer distribution packages 205
 - Squirrel Data Exporter 148, 149, 155, 156, 157
 - System Data Store Administrator 169
 - bin directory 243
 - binary files 29, 127, 243
 - BIRT 360 Option 275
 - BIRT Data Analyzer Option 276
 - BIRT Designer Professional 276
 - BIRT Exchange 27
 - BIRT Exchange URL 27
 - BIRT Interactive Viewer Option 276
 - BIRT iServer 16
 - See also* iServer
 - BIRT iServer System names 30, 68
 - BIRT open source projects vii
 - BIRT Option 275
 - BIRT Page Level Security option 276
 - BIRT report designers 47
 - BIRT reports 19, 276
 - See also* reports
 - BIRT SmartSheet Security Option 276
 - BIRT Spreadsheet Designer 276
 - BIRT Spreadsheet Option 276
 - BIRT Studio 255
 - BIRT Studio Option 276
 - BIRT Viewer 255
 - BIRT_RESOURCE_PATH parameter 255
 - BIRT_VIEWER_LOG_DIR parameter 255
 - birt-exchange.com vii
 - browsers. *See* web browsers
 - browsing Encyclopedia volume database 41, 103
 - bulk load script files 165
 - bulk-load utilities 153
 - Business Intelligence technology vii
- ## C
- C++ libraries 203
 - C++ runtime components 207, 218, 220, 240
 - cache (web pages) 257
 - cache conflicts 257
 - Caching processes 287
 - Caching service 9, 11
 - CachingService element 287
 - changing
 - cluster machines 282
 - CPU binding 288, 290
 - database encoding 19
 - database schemas 41, 103
 - default locales 255
 - license file names 281
 - licensing options 283
 - metadata databases 4, 5
 - network cards 282
 - passwords 34
 - time zones 255
 - volume names 164
 - character encoding 19, 60
 - character sets 19
 - charts 46, 275
 - CLASSPATH variable 46
 - clearing web browser cache 257
 - client applications 8
 - client/server models 13
 - cloud deployment environments 243
 - cloud deployments 6, 202–244
 - cluster IDs 13, 176
 - cluster nodes 176, 184
 - creating shared folders for 178, 234
 - cluster template names 240
 - clusters
 - adding nodes to 13, 176
 - associating with machine IDs 279
 - binding iServer processes to 284
 - changing machines for 282
 - configuring 13, 14

- clusters (*continued*)
 - connecting to iServer System and 8
 - creating 181, 237
 - creating iServer images and 202
 - deploying iServer and 206, 239
 - deploying web applications and 256
 - determining number of processors for 288
 - distributing requests among 8, 14
 - exceeding CPU licenses for 289
 - failing 176
 - failover operations and 7
 - getting IP addresses for 180, 236
 - installing cluster nodes for 185–198
 - installing iServer for 29
 - licensing options for 277, 279, 283
 - load balancing for 14, 176
 - managing volume metadata and 5
 - running iServer services and 8
 - setting configuration home location for 187
 - setting port numbers for 191
 - setting up RDBMS databases for 11
 - setting up system data store for 168
 - setting up user accounts for 25
 - storing information for 202
 - testing connections for 181, 237
 - turning off firewalls for 179, 235
 - viewing nodes in 242
- collation 57, 60
- collecting machine information 279
- command line utilities 16, 297
- comments 153
- config directory 188
- CONFIG_SCHEMA_NAME parameter 158, 172
- ConfigHomeDirectory variable 199
- Configuration Console
 - adding alternate databases and 75
 - adding database schemas and 146, 225, 226
 - adding metadata databases and 223
 - adding partitions and 145, 228
 - administering iServer and 211
 - archiving report files and 296
 - configuring BIRT iServer and 16
 - configuring clusters and 13
 - connecting to databases and 222
 - creating Encyclopedia and 151, 166, 167, 168, 230
 - creating system schemas and 146, 168
 - installing 187, 246
 - logging in to 211
 - running iServer processes and 8
 - running stand-alone servers and 9
 - sending notifications and 290
 - setting user names and passwords for 34
 - shutting down iServer and 147
 - starting 102
 - updating license keys and 282
 - upgrades and 95, 120, 143
 - viewing licensing information and 277
- configuration files 27, 87, 153, 244
- configuration home locations 187, 199
- configuration parameters 13, 176, 255
 - See also* parameters
- configuration templates 13, 176
- configurations
 - binding CPU processes and 284, 287
 - connecting to databases and 18
 - creating cluster nodes and 8
 - customizing database encoding and 19
 - deploying Information Console and 247, 254, 256
 - failover operations and 7
 - installing cluster nodes and 13, 187
 - installing Information Console and 247, 254
 - installing iServer and 26, 29
 - licensing iServer and 274, 279
 - overwriting previous installations and 45
 - purging report files and 296
 - renaming licensing files and 282
 - running Open Security applications and 47
 - running RDBMS databases and 11
 - storing metadata for 4
 - testing installations and 48
 - upgrading and 87, 153
- configuring
 - application servers 247, 254
 - clusters 13, 14
 - external databases 221
 - iServer System 6
 - JDBC drivers 18
 - service privileges 26

- system data store 168
 - upgrade utilities 153
 - connection pooling 18, 53
 - connections
 - accessing multiple volume schemas and 53
 - accessing volume databases and 18
 - changing encoding schemes and 19
 - configuring clusters and 8, 13
 - configuring Information Console 255
 - creating DB2 schemas and 61
 - creating Oracle schemas and 56
 - creating PostgreSQL schemas and 54
 - creating SQL Server schemas and 59
 - customizing database 18
 - setting up database 222
 - testing 181, 224, 237
 - upgrading PostgreSQL databases and 92
 - context paths 249, 254
 - context-sensitive help 264
 - conventions (documentation) xii
 - copying
 - files 295
 - license files 282
 - corrupted schemas 83
 - CPU binding 283–290
 - CPU binding validation 289, 290
 - CPU-based licenses 274
 - determining number of 284
 - exceeding number of 289, 290
 - viewing information about 289
 - CPUs
 - binding to multiple core 286
 - deploying iServer over multi-threaded 284
 - determining number of 287, 288
 - hyperthreading and 287
 - licensing options for 283
 - restricting processes for 283
 - running encycsrvr11 processes and 288
 - running Information Console and 254
 - viewing maximum number of 289
 - viewing processor IDs for 284
 - CREATE_SCHEMA parameter 158
 - creating
 - administrator accounts 24, 83
 - backup folders 298, 318, 338
 - cluster nodes 176
 - clusters 181, 237
 - data cubes 275
 - e-mail accounts 290
 - Encyclopedia volume partitions 143
 - Encyclopedia volumes 27, 151, 166, 167, 168, 230
 - indexes 5
 - iServer images 202
 - metadata databases 11, 223
 - partitions 145, 229
 - passwords. *See* passwords
 - production environments 48
 - production staging areas 48
 - report designs 276
 - resource groups 46
 - schema owners 52, 225, 227
 - shared folders 178, 234
 - system databases 53, 60
 - system passwords 242
 - system schemas 27, 52, 168, 170, 225
 - test environments 47
 - user accounts 25, 52
 - volume schemas 27, 52, 146, 226
 - credentials. *See* user credentials
 - cube reports 275
 - cubes 275
 - cubeview files 276
 - custom applications 47
 - custom configurations 29
 - custom installations 29
 - Customer Support xi
 - customizing
 - database connections 18
 - database encoding 19
 - Information Console 255, 256
 - metadata databases 5
- ## D
- dashboards vii, 275
 - data
 - accessing 13, 276
 - analyzing 275
 - backing up Encyclopedia and 294
 - caching 11
 - preventing loss of 7, 26
 - recovering 7, 8, 16
 - setting default location for 29

- data (*continued*)
 - sharing 14
 - storing 203
 - upgrades and 6, 49, 83, 127
- Data Analyzer Option 276
- Data Connector Option 276
- data cubes. *See* cubes
- data directory 30, 127
- Data Exporter
 - configuring 153
 - migrating iServer installations and 6
 - migrating volume metadata and 17
 - running 148
 - setting properties for 153
 - setting up export environments and 153, 155, 165
- data files 29, 243, 295
 - See also* files; report files
- data objects 5, 294
- data partitions 228
- data retrieval operations 11
- data sources 11
- Data Store Administrator
 - See also* System Data Store Administrator
 - configuring 153
 - creating Encyclopedia and 166, 167, 168
 - creating schema elements and 166, 167
 - importing Encyclopedia and 149, 157, 164
 - migrating iServer installations and 6
 - migrating volume metadata and 17
 - running 157, 163
 - setting properties for 157
- Data Store Upgrader
 - configuring 153
 - migrating volume metadata and 17
 - running 121, 122, 171
 - setting properties for 171, 172
 - upgrading iServer and 6
- data stores 168, 255
 - See also* database schemas
- data types 18, 60
- DATA_EXPORT_FOLDER parameter 159
- DATA_EXPORT_FORMAT parameter 159
- DATA_IMPORT_FOLDER parameter 159
- DATA_IMPORT_FORMAT parameter 159
- database administration tool
 - accessing PostgreSQL servers and 41, 103
 - backing up Encyclopedia and 298–301
 - installing 36, 37
 - restoring Encyclopedia and 306, 309
 - starting 41
 - uninstalling 44
 - upgrades and 97, 114
- database collation 57, 60
- Database Credentials option 31
- database drivers. *See* drivers
- database objects 221
- Database Schema Information page 71
- database schema names 30, 31, 121, 161, 163
- database schema owners
 - creating 52, 225, 227
 - installing DB2 databases and 60, 61
 - installing iServer and 5
 - installing Oracle databases and 55, 56
 - installing PostgreSQL databases and 53, 54
 - installing SQL Server databases and 57, 58
 - setting passwords for 31, 163
- database schemas
 - accessing 4
 - adding Encyclopedia to 167, 168
 - assigning privileges 53
 - backing up 7, 26, 168, 302
 - browsing 41
 - changing 41, 103
 - configuring system data store as 168
 - creating 27, 52, 146, 225, 226
 - customizing metadata databases and 5
 - deploying iServer and 206
 - depopulating 160
 - dropping 160
 - importing volume data to 164
 - initializing 27
 - installing DB2 databases and 61
 - installing Oracle databases and 55, 56
 - installing PostgreSQL databases and 54, 55
 - installing SQL Server databases and 58, 59
 - naming 31, 53, 163, 225, 227
 - populating 162, 166, 167
 - preventing data loss and 7
 - removing Encyclopedia from 166
 - restoring 168, 170
 - setting passwords for 162
 - storing metadata and 4
 - upgrades and 83, 92

- user-defined functions and 63
- database servers. *See* servers
- database superuser 31, 163
- database types 155, 216
- DATABASE_HOST parameter 159, 172
- DATABASE_INSTANCE parameter 159, 172
- DATABASE_NAME parameter 159, 172
- DATABASE_PORT parameter 159, 172
- DATABASE_TYPE parameter 157, 160, 171
- dbName property 216
- databases
 - accessing documentation for 7
 - accessing Encyclopedia and 46
 - adding data store schemas and 170
 - analyzing data and 275
 - backing up 17
 - changing metadata tables and 5
 - configuring external 221
 - configuring failover operations for 7
 - connecting to 18, 222
 - copying files in 295
 - creating metadata 11, 223
 - creating system 53, 60
 - customizing 5
 - deploying iServer and 205, 214–219
 - exporting metadata to 155
 - importing Encyclopedia and 164
 - installing iServer and 24, 27, 29, 52
 - installing metadata 30, 69
 - localizing 19
 - managing Encyclopedia and 5, 7
 - running iServer clusters and 11, 14
 - selecting alternate 68
 - setting passwords for 31
 - setting port number for 224
 - setting up iServer environments for 52
 - specifying default locations for 29
 - storing metadata and 4
 - uninstalling 44, 104
 - upgrades and 82, 83, 92, 104
 - viewing incomplete files in 295
- DB2 command line processor 344
- DB2 Control Center 338–344, 347
- DB2 databases
 - accessing documentation for 7
 - connecting to 61
 - creating schema owners for 60, 61
 - creating user accounts for 60
 - exporting metadata to 156
 - extending SQL functionality 62–63
 - installing 60, 61, 69
 - managing Encyclopedia and 5
 - naming schemas for 61
 - running volume backup and restore operations with 337, 338, 344, 347
 - setting port number for 69
 - setting up iServer environments for 60–61
 - storing metadata and 4
 - taking Encyclopedia offline for 347
 - taking Encyclopedia online for 356
- dbtype property 216
- default directories. *See* directories
- default locale 33, 255
- default paths 29, 30
- default ports. *See* ports
- default time zones 255
- default user names 34
- default values 45
- default volume partition 143, 144, 145, 228
- DEFAULT_DATABASE_PASSWORD parameter 216
- DEFAULT_DATABASE_NAME parameter 160
- DEFAULT_LOCALE parameter 255
- DEFAULT_TIMEZONE parameter 255
- DEFAULT_VOLUME parameter 255
- DELETE_ALL_DATA parameter 160
- DELETE_DATA parameter 160
- deleting
 - database schemas 160
 - installation directories 44, 104
 - partitions 228
 - product files. *See* uninstalling
- deploying
 - Information Console 247, 254–256
 - iServer 6, 13, 14, 176, 202–237, 284
 - new releases 47
 - spreadsheets 276
 - WAR files 254
 - web applications 256
- deployment scripts 203
- deployment tools 254
- DEPOPULATE_SCHEMA parameter 160
- design files 5

- designers vii, 47
- designs 11, 47, 275, 276, 278
- desktop products 48
- desktop shortcuts 41, 102, 118, 210
- destination folder locations 29
- developers 15, 53
- developing applications 15, 16
- DHTML reports 276
- diagnostic fixes 9, 14, 177
- digit wildcard character 61
- directories
 - backing up volume data 303, 315, 325, 346
 - configuring cluster nodes and 176
 - creating cluster nodes and 13
 - customizing WAR files and 256
 - deploying iServer and 203, 204, 243, 244
 - extracting product files to 246
 - file structure changes and 102, 120
 - ICU libraries in 25
 - installing cluster nodes and 187
 - installing iServer and 29
 - installing iServer Integration Technology and 260
 - installing JDK files and 46
 - mapping Oracle database dump 313
 - migrating Encyclopedia and 143
 - removing installation 44, 104
 - restoring 306, 316, 326, 347
 - running multiple releases and 45
 - upgrading iServer and 45, 87, 127
- directory paths 29, 127, 145, 295
- disk partitions 295
 - See also* volume partitions
- disk space 45, 232
- displaying
 - charts 46
 - cluster nodes 242
 - cube reports 275
 - folder properties 178, 234
 - host names 180, 236
 - IP addresses 180, 236
 - licensing information 277, 278
 - processor IDs 284
 - ReadMe files 253, 263
 - reports 11, 276
- DLLs 46, 243
- document files 5, 278

- document generation service 11
- documentation
 - accessing vii–xi
 - administering iServer System and 16
 - conventions for xii
 - deploying WAR files and 254
 - downloading xi
 - installing online 40, 264
 - managing Encyclopedia and 7
 - upgrades and 101, 118, 142
- documentation update tool 264
- domain servers 25
- downloading
 - documentation files x
 - JAVA SE Development Kit 203
 - product files 27
 - report files 11
 - Visual C++ libraries 203
- drivers 11, 14, 18
- DROP_SCHEMA parameter 160
- dual-core CPUs 286
- dump directories 313
- dynamic link libraries 46, 243

E

- e.Analysis Option 276
- e.Report Data Connector Option 276
- e.Report Designer Professional 276, 277
- e.Report Option 276
- e.reports 203, 276
- Eclipse BIRT open source projects vii
- elastic iServer clustering 13, 176, 202
- e-mail. *See* notifications
- encoding 19, 60
- encyc directory 44, 87, 104
- Encyclopedia Data Store Administrator
 - configuring 153
 - creating Encyclopedia and 166, 167, 168
 - creating schema elements and 166, 167
 - importing Encyclopedia and 149, 157, 164
 - migrating iServer installations and 6
 - migrating volume metadata and 17
 - running 157, 163
 - setting properties for 157
- Encyclopedia Data Store Upgrader
 - configuring 153

- migrating volume metadata and 17
- running 121, 122, 171
- setting properties for 171, 172
- upgrading iServer and 6
- Encyclopedia Metadata Storage page 30, 68, 91
- Encyclopedia processes. *See* encycsrvr11 processes
- Encyclopedia volumes
 - accessing multiple 277
 - accessing objects in 274
 - adding to database schemas 53, 55, 57, 60, 168
 - backing up 294, 295, 296, 297, 300, 313, 317, 337
 - browsing schemas for 41
 - browsing system database for 103
 - connecting to database for 18
 - controlling access to 46
 - CPU binding and 289
 - creating default 27
 - creating new 151, 166, 167, 230
 - creating partitions for 143
 - creating schema database for 55, 59, 61, 146
 - creating schema owners for 54
 - creating schemas for 226
 - customizing metadata databases for 5
 - downloading files from 11
 - exporting 143, 153, 157, 163
 - importing 149, 157, 163
 - installing Information Console and 247
 - installing metadata database for 4, 5, 24, 52
 - installing sample 19
 - installing system database for 30
 - licensing and 274
 - localizing databases for 19
 - managing 5, 7
 - migrating metadata for 17
 - migrating to current release 49, 82, 120, 143
 - naming 30, 151, 162, 230
 - overwriting product files and 48
 - preserving metadata for 26
 - preventing data loss and 7, 26
 - putting partitions online for 231
 - removing from schemas 166
 - renaming 164
 - restoring. *See* restoring Encyclopedia volumes
 - running iServer clusters and 14, 176
 - running iServer processes and 8
 - running multiple schemas on 53
 - setting location of 29
 - specifying default 255
 - specifying primary partition for 151
 - specifying schema owners for 31
 - starting 9, 152, 231, 232
 - storing metadata for 4, 11
 - troubleshooting 295
 - uninstalling databases for 44, 104
 - upgrading iServer and 6, 49, 83
- EncyclopediaUpgrader utility. *See* Encyclopedia Data Store Upgrader
- encycsrvr11 processes 284, 288, 289
- environment variables 13, 46, 122, 176, 285
- environments 254
- error logging applications 15
- error logging reports 19
- escape characters 153
- evaluation copies 27
- evaluation licenses 32, 130
- events 8, 15
- example reports 19
- example volume 19
- Excel formats vii
- Excel spreadsheets 276
- executable files 46
- exit code 0 301
- expdp command line options 314
- expdp command line utility 313, 314
- Expiration time of deleted files property 296
- expired licenses 278, 279
- EXPORT_ALL_DATA parameter 160
- EXPORT_DATA parameter 160
- export_squirrel_data.bat file 148
- exporting Encyclopedia volumes 143, 153, 157, 163
- external data sources 11
- external databases 221
- external security systems 29, 46
- external user information 15
- extracting product files 246

F

- Factory processes 46, 287, 290
- Factory service 9, 11, 46
- failover operations 7, 8
- features xi, 46, 274
- file folder 146
- file I/O operations 11
- file names 278
- file purging properties 296
- file system backup utilities 295
- file system backups 17
- file systems 5, 11
- FILE_LOG_LEVEL parameter 154, 160, 172
- files
 - accessing license 13
 - accessing online help x, 267
 - accessing report 24
 - archiving 296
 - backing up 17, 87, 107, 295
 - configuring upgrade utilities and 153
 - copying 295
 - deploying iServer and 202, 206, 243
 - downloading documentation x
 - downloading product 27, 246
 - downloading report 11
 - installing iServer and 27
 - installing JDK software and 46
 - installing online help 264
 - migrating Encyclopedia volumes and 146
 - obtaining license key 278
 - overwriting product 6, 45, 48
 - purging 296
 - running processes and 85, 104
 - saving 87, 107, 153
 - selecting license 32
 - sharing 178, 234
 - specifying default locations for 29
 - specifying location of binary 127
 - storing report-specific 5
 - transmitting license keys and 282
 - updating 264
 - upgrading iServer and 45, 47, 83
- filetype folder 146
- firewalls 179, 209, 235, 247
- fixes 9, 14, 177
- Flash gadgets vii

folders

- backing up 303, 325, 346
- creating backup 298, 318, 338
- deploying iServer and 204
- installing iServer and 29
- installing Management Console and 195
- migrating Encyclopedia and 146
- restoring 306, 316, 326, 347
- shared configurations and 188
- sharing 178, 234
- specifying program 136
- specifying resource 255
- upgrades and 146
- viewing properties for 178, 234

formats vii

Forrester Wave Open Source Business Intelligence report vii

ftp distributions

- documentation xi
- Information Console 246, 247
- iServer Integration Technology 260
- iServer System 246

ftp site (Actuate) 27

full installations 246

G

- gadgets vii
- generating
 - charts 46
 - reports 11, 46
 - temporary documents 11
- getJDBCMajorVersion method 18
- graphs. *See* charts

H

- heap size 254
- help 264
 - See also* documentation
- help files x, 41, 264, 267
- help topics xi
- host names
 - displaying 180, 236
- HTML documentation vii, x
- hyperthreading 287

I

I/O operations 11

See also input; output

IBM DB2 databases. *See* DB2 databases

icons (console applications) 136

ICU libraries 25

ICU_DATA environment variable 25

IDAPI applications 10, 15, 16

impdp command line utility 315, 317

IMPORT_DATA parameter 161

importing Encyclopedia volumes 149, 157, 163

incomplete files 295

indexed searches 5

Information Console

assigning CPUs to 254

caching web pages for 257

configuring clusters and 13

connecting to 255

customizing 255, 256

deploying 247, 254–256

installing 246, 247, 254

logging activities for 255

logging in to 41, 210, 258

preventing cache conflicts for 257

running iServer processes and 8

running stand-alone servers and 9

setting application server port for 247

setting context paths for 249, 254

setting default Encyclopedia for 255

starting 257

testing installations for 257

Information Delivery API 15, 16

information delivery solutions vii

information object cache 11

information object files 5

information objects 275

INITIALIZE_DATA parameter 161, 166

in-place upgrades 82, 84, 87, 107, 120

install.properties file 215

installation

alternate metadata databases 52, 69

C++ runtime components 207, 218, 220, 240

cache conflicts and 257

cloud deployments and 203, 204

cluster nodes 176, 185–198

Configuration Console 187, 246

DB2 databases 60, 61, 69

desktop products 48

disk space and 45

domain servers and 25

Information Console 246, 247, 254

iServer 6, 24, 27–41, 52, 185

iServer Integration Technology 246, 260–264

iServer System 6, 246

iServer System components 29

Java Runtime Environment 247

Java Software Development Kit 46, 247

license keys 278

licensing prompts and 279, 280

Management Console 193, 246

metadata databases 4, 5, 30

online documentation 40, 264

Open Security application 29, 46

PDF documentation 41

pgAdmin administration tool 36, 37

previous releases and 45, 47

sample Encyclopedia volume 19

sample reports 19

SQL Server databases 57, 71

testing 47, 257

upgrades and 45, 82, 84, 104

installation directories 29, 44, 104

installation guides 16

installation options 29

installation prerequisites 24, 52, 184

installation programs 16

installation wizard 247

Integration processes 287

Integration service 9, 11

IntegrationService element 287

Interactive Viewer 276

Interactive Viewer Option 276

international character sets 19

international customer licensing 283

IP addresses

displaying 180, 236

ipconfig command 180, 236

iportal context path 249

iServer

accessing functionality 274

- iServer (*continued*)
 - accessing online help for 267
 - binding to multiple-core CPUs 286
 - binding to processors 284, 285
 - changing CPU binding and 288
 - checking bound processors for 288–290
 - creating data objects and 5
 - deploying 6, 13, 14, 176, 202–237, 284
 - downloading evaluation copies for 27
 - getting machine ID for 280
 - initializing 27
 - installing 6, 24, 27–41, 52, 185
 - optimizing performance for 14, 177
 - running 14, 176
 - setting up user accounts for 24, 25, 26, 52
 - shutting down 147
 - starting 9, 212, 221
 - stopping 212
 - storing metadata for 4
 - testing new releases for 47–49
 - upgrade options for 6
 - upgrading 45, 82, 84, 104, 120, 124
 - validating CPU binding and 289
 - viewing licensing information for 277, 278
 - iServer 11 service 33, 132
 - iServer Application Container 135, 143
 - iServer directory 29
 - iServer distribution package 204
 - iServer distribution setup script
 - alternate databases and 214, 215, 218
 - desktop shortcuts and 210
 - environment variables and 205
 - installation options and 205
 - Postgre databases and 207
 - running 206
 - setting up clusters 239
 - setting up iServer only 219
 - stopping 209
 - iServer images 202
 - iServer Information page 251
 - iServer Integration Technology 16, 246, 260–264
 - iServer processes
 - binding to CPUs 283–290
 - binding to Windows systems 284
 - creating resource groups and 46
 - domain servers and 25
 - installing iServer and 85
 - multiple-core CPUs and 286
 - running 8, 176, 283
 - starting 284
 - upgrades and 85, 104
 - verifying CPU binding for 286
 - iServer services
 - See also* specific service
 - running multiple 8
 - setting user profiles and 33
 - shutting down 85, 105
 - specifying user accounts for 24
 - iServer servlet container 8
 - iServer System
 - administering 15
 - changing CPU binding and 290
 - configuring 6
 - creating dedicated account for 25
 - deploying new releases and 47
 - downloading files for 27, 246
 - file structure changes and 103
 - firewalls and 209
 - installation options for 6, 29
 - installation requirements for 246
 - licensing options for 274, 275, 283
 - optimizing 13, 176
 - preventing data loss for 7, 26
 - running multiple releases and 45, 47
 - storing metadata for. *See* system databases; system metadata
 - third-party databases supported 4
 - upgrading 84
 - iServer System components 29, 41
 - iServer System names 30, 68
- ## J
- J2EE installations 254
 - JAR files 46
 - Java developer guide 16
 - Java Factory service. *See* Factory service
 - Java heap size 254
 - Java Object Interface 46
 - Java Report Server Security Extension 16
 - Java Runtime Environment 18, 247
 - JAVA SE Development Kit 203
 - Java Server Pages. *See* JSPs

- Java Software Development Kit 46, 247
- Java Virtual Machines. *See* JVM
- JDBC drivers 11, 18
- jdbcCompliant method 18
- JDK files 46, 247
- jobs 10
- JRE environment 18, 247
- JSPs 257
- JVM libraries 46

K

- keys. *See* license keys

L

- language settings 33
- LDAP configuration file 47
- LDAP security database 46
- LDAP servers 29, 46
- libraries 25, 46, 203, 243
- License File Details page 32, 93
- license file names 281
- license files 13, 32
- license key file locations 279
- license key file names 278
- license key files 279
- license key installations 278
- license keys 277, 282, 284
- licensed CPUs 289
 - See also* CPU-based licenses
- licenses 93, 130, 274
- licensing information 277
- licensing options 16, 274, 275, 283
- links (documentation) xi
- Linux systems 16, 19
- load balancing (clusters) 14, 176
- local help 267
- Local Security Settings dialog 26
- Local Users and Groups dialog 25
- Locale Information page 33, 94
- Localemap.xml 255
- locales
 - encoding character sets for 19
 - obtaining licenses for 283
 - selecting 33
 - setting default 33, 255
 - updating resource files for 264

- upgrades and 94
- localhost parameter 255
- Localization and Online Documentation
 - Update tool 264
- log files 29, 127, 154, 255
- Log on as a service privilege message 34, 75, 95, 191, 251
- LOG_FILE_COUNT parameter 154, 161, 172
- LOG_FILE_LOCATION parameter 255
- LOG_FILE_NAME parameter 161, 173
- LOG_FILE_SIZE parameter 155, 161, 173
- LOG_FOLDER parameter 155, 157, 161, 171, 173
- logging applications 15
- logging levels
 - Data Store Administrator 160
 - Data Store Upgrader 172
 - Squirrel Data Exporter 154
- logical CPUs 286
- logical processors 284, 287
- losing data 7, 26

M

- machine IDs 279, 280, 282
- machine information 279
- machine names 45
- Management Console
 - adding alternate databases and 75
 - configuring clusters and 13
 - creating Encyclopedia and 232
 - installing 193, 246
 - licensing and 274
 - logging in to 210
 - managing Encyclopedia and 16
 - running iServer processes and 8
 - running stand-alone servers and 9
 - setting application container port for 194
 - setting user names for 34
 - starting 41
 - upgrades and 95
- manual export operations 153
- manual import operations 157
- manual in-place upgrades 82, 84, 87, 120
- manual iServer system upgrades 6
- manual side-by-side migrations 143
- manual side-by-side upgrades 82, 124

- manual volume upgrade operations 171
- manuals. *See* documentation
- master cluster nodes 8
- master index file xi
- message routing (requests) 10, 13
- metadata
 - autoarchiving and 296
 - backing up 26, 64, 295
 - changing 5
 - defined 294
 - exporting Encyclopedia and 153, 155
 - migrating 17
 - preventing loss of 7, 26
 - restoring 316, 331, 334
 - storing 4, 11, 168
 - upgrades and 6, 83, 91
- metadata databases
 - See also* third-party databases
 - backing up 17
 - changing 4
 - creating 223
 - customizing connections for 18
 - deploying iServer and 205, 214
 - editing 5
 - incomplete files and 295
 - installing 4, 5, 30, 52, 69
 - naming 224
 - previous releases and 4
 - selecting type 68
 - setting port number for 224
 - setting up iServer environments for 52
 - sharing 5
- Microsoft Excel spreadsheets 276
- Microsoft SQL Server databases. *See* SQL Server databases
- Microsoft Windows systems. *See* Windows systems
- migration
 - customizing metadata databases and 5
 - installing new releases and 27, 47
 - upgrading and 6, 45, 82
- migration options 143
- migration tools 16, 49
- missing product files 246
- missing system schemas 83
- multicast configurations 8
- multidimensional data analysis 275

- multiple-core CPU binding 286
- multiple Encyclopedia volumes 277
- multiple product installations 45
- multiple product releases 47
- multiple upgrade installations 45
- Multi-Tenant Option 277
- multi-threaded CPUs 284

N

- name property 216
- named-user licenses 33, 93, 274
- names
 - obtaining host 180, 236
- naming
 - data partitions 229
 - Encyclopedia volumes 30, 162, 230
 - iServer System 30, 68
 - metadata databases 224
 - schema owners 53, 54
 - system schemas 31, 53, 225
 - users 53
 - volume schemas 53, 55, 59, 61, 227
- naming conventions 53, 68
- network administrators. *See* administrators
- network cards 280, 282
- Network File Systems (NFS) 11
- networked environments 11, 13, 176, 246
- New Metadata Database page 223
- New Schema page 146
- New Server Registration dialog 42
- New System Schema page 225
- New Volume page 230
- New Volume Schema page 227
- NEW_SCHEMA_NAME parameter 154, 161, 164
- NEW_VOLUME_NAME parameter 154, 162, 166
- newinformationconsole.war 256
- NobleNet Portmapper processes 104, 106
- NobleNet Portmapper service 85, 86, 104, 106
- node keys 279
- node-key file locations 280
- node-key licensing 277, 278, 279
- nodes. *See* cluster nodes
- notifications 289, 290

O

- Object browser 42
- obsolete command line utilities 16
- obsolete features 7
- obsolete product information xii
- online documentation
 - accessing vii–xi
 - administering iServer System and 16
 - conventions for xii
 - downloading xi
 - installing 40, 264
 - managing Encyclopedia and 7
 - upgrades and 101, 118, 142
- online documentation update tool 264
- online help 267
 - See also* online documentation
- Open Security applications 29, 46, 88
- open source projects vii
- opening
 - Configuration Console 102
 - Information Console 41, 257
 - Management Console 41
 - pgAdmin administration tool 41
- operating system user accounts 60
- operating systems 19, 284
- optimizing iServer System 13, 176
- options (installation) 29
- options (licensing) 274, 275, 283
- Oracle Data Pump Export utility 313–315
- Oracle Data Pump Import utility 315–317
- Oracle databases
 - accessing documentation for 7
 - connecting to 56
 - creating schema owners for 55, 56
 - exporting metadata to 155
 - installing 69
 - managing Encyclopedia and 5
 - running volume backup and restore
 - operations with 313, 315
 - setting port number for 69
 - setting up iServer environments for 55–57
 - storing metadata and 4
 - taking Encyclopedia offline for 316
 - taking Encyclopedia online for 317
 - upgrades and 82, 104
- Oracle service names 69

- ORACLE_TNS_NAMES_FILE parameter
 - 162, 173
- output 11
- output formats vii, 11
- overwriting previous installations 45
- overwriting product files 6, 45, 48

P

- Page Level Security Option 277
- page-level security 276, 277
- parameters
 - configuring clusters and 13, 176
 - deploying Information Console and 255
 - encoding volume databases and 19
 - exporting Encyclopedia volumes and 153
 - setting Data Store Administrator 157
 - upgrading Encyclopedia and 171, 172
- partition names 145, 229
- partition paths 145, 229
- partitions (disk) 295
- partitions (volume)
 - creating 145
 - manual side-by-side upgrades and 143
 - renaming default 144
 - setting default 228
 - setting primary 151, 230
- Partitions page 144, 152, 228
- password property 216
- passwords
 - changing 34
 - connecting to databases and 224
 - creating schema owners and 163, 226, 227
 - creating system schemas and 225, 227
 - installing alternate databases and 75
 - installing cluster nodes and 192
 - installing Information Console and 250
 - installing iServer and 27
 - prompting for schema 162
 - setting database 31
 - setting system 242
 - starting console applications and 34
 - starting iServer service and 33
- patches 9, 14, 177
- PATH environment variable 148, 169
- paths 29, 127, 145, 295
- PDF documentation vii, x, xi, 41, 264

- performance
 - caching data and 11
 - cluster configurations and 14, 177
 - CPU binding and 286
 - file input/output operations and 11
 - iServer processes and 9
 - upgrades and 45
- performance monitoring extensions 15
- permanent licenses 279
- permissions. *See* privileges
- pg_dump command line options 302
- pg_dump utility 297, 299, 302, 303
- pg_restore command line options 312
- pg_restore utility 297, 309, 311, 312
- pgAdmin database administration tool
 - accessing PostgreSQL servers and 41, 103
 - backing up Encyclopedia and 298–301
 - installing 36, 37
 - restoring Encyclopedia and 306, 309
 - starting 41
 - uninstalling 44
 - upgrades and 97, 114
- pgAdmin III Setup dialog 37, 114
- pgAdmin III Setup Wizard 117
- pgpass files 302
- ping command 181, 237
- plpgsql language 53, 215
- POPULATE_SCHEMA parameter 162, 166
- POPULATE_SCHEMA_PHASE_ONE
 - parameter 162
- POPULATE_SCHEMA_PHASE_TWO
 - parameter 162
- port numbers 45
 - See also* ports
- portinst.exe 203
- portmapper components 203
- Portmapper service 85, 86, 106
- portNumber property 216
- ports
 - application servers 247
 - cloud deployments and 203
 - cluster nodes 191
 - DB2 databases 69
 - Information Console 255
 - iServer Application Container 135, 143
 - multiple product releases and 45
 - Oracle databases 69
 - PostgreSQL databases 31, 71
 - SQL Server databases 71
 - upgrades and 45
- postgres database directory 44, 104
- PostgreSQL command line utilities 297
- PostgreSQL Database Connection
 - Information 92
- PostgreSQL Database Information page 31, 92
- PostgreSQL databases
 - accessing documentation for 7
 - adding 11
 - backing up 17, 297
 - backing up data folders for 303, 304
 - browsing schema definitions in 41
 - changing schemas for 103
 - cluster configurations and 11
 - connecting to 54
 - creating Encyclopedia for 168
 - creating schema owners for 53, 54
 - deploying iServer and 207–209, 216
 - exporting metadata to 156
 - file I/O operations and 11
 - importing Encyclopedia and 164
 - installing administration tool for 36, 37
 - installing iServer and 24, 27, 71
 - managing Encyclopedia and 5
 - naming schemas for 54, 55
 - running volume backup and restore
 - operations with 297, 298, 302, 306, 311
 - setting passwords for 31
 - setting port number for 31, 71
 - setting up iServer environments for 53–55
 - shutting down 212, 213
 - specifying default locations for 29
 - starting admin tool for 41
 - storing metadata and 4
 - taking Encyclopedia offline for 306
 - taking Encyclopedia online for 312
 - taking online 213
 - uninstalling 44, 104
 - upgrades and 82, 92, 104, 105
 - viewing installation information for 44
- postgresql directory 304
- PostgreSQL for BIRT service 33, 104, 105
- PostgreSQL processes 212
- PostgreSQL server clusters 5

- PostgreSQL servers 41, 42, 103, 211
- PostgreSQL SQL language 53
- PostgreSQL superusers 31, 163
- Power User accounts 24, 25, 33
- prerequisites (installation) 24, 52, 184
- previous releases 6, 45, 48, 84, 203
- primary volume partition 151
- printers 14, 24
- privileges
 - connecting to DB2 databases and 60, 61
 - connecting to Oracle databases and 55, 56
 - connecting to PostgreSQL databases and 53, 54
 - connecting to SQL Server databases and 57, 58, 59
 - creating system schemas and 170
 - creating user accounts and 24, 25
 - installing Information Console and 247
 - logging in to iServer and 26
 - running multiple schemas and 53
- Process Management Daemon
 - CPU binding and 284, 286
 - running cluster configurations and 13, 14, 176
 - running iServer processes and 8, 10
 - starting encycsrvr processes and 288, 289
- Processor Affinity 286
- processor IDs 284, 287
- ProcessorAffinity element 287
- processors. *See* CPUs
- product files 6, 27, 45, 48, 246
- product information xii
- product updates xi, 264
- production environments 48
- production staging areas 48
- profiles (user) 33
- program files. *See* product files
- program folders 136
- PROMPT_FOR_PASSWORDS parameter 162, 173
- properties 153
 - file sharing 178, 234
- properties files 122, 149, 153, 255
- Properties page (New Server Registration) 42
- publishing reports 275
- Purge deleted files time property 296
- purging report files 296

Q

- queries 11, 13, 62
- Query Option (licensing) 275

R

- RDBMS backup utilities 295
- RDBMS bulk-load utilities 153
- RDBMS database configurations 11
- RDBMS databases 4, 295
 - See also* third-party databases
- RDBMS documentation 7
- RDBMS environments 4, 14, 53
- RDBMS tools 7, 17
- ReadMe files 253, 263
- rebinding encycsrvr11 processes 289
- recovery operations 7, 8, 16
- refreshes 257
- registered users 33
- relational databases. *See* databases
- release notes xi
- remote procedure calls 10
- remote servers 246
- removing
 - database schemas 160
 - installation directories 44, 104
 - partitions 228
 - product files. *See* uninstalling
- renaming
 - Encyclopedia volumes 164
 - license files 281
- replacing previous releases 45
- report design files 5
- report designers vii, 47
- report designs 11, 47, 275, 276, 278
- report document files 5, 278
- report document generation service 11
- Report Encyclopedia. *See* Encyclopedia
- report files
 - accessing 24
 - archiving 296
 - backing up 17, 87, 107, 295
 - downloading 11
 - purging 296
 - running processes and 85, 104
 - sharing 178, 234
 - specifying default directories for 29

- report files (*continued*)
 - storing 5
 - upgrading iServer and 6, 45, 47, 83
- report object executable files 46
- Report Server Security Extension 15, 16
- Report Server Security Extension API 88
- report templates 255
- report viewers 276
- reporting applications. *See* applications
- reporting server. *See* iServer
- reporting services. *See* specific iServer service
- ReportingService element 287
- reports
 - displaying 11, 276
 - generating 11, 46
 - installing sample 19
 - publishing 275
 - running 203, 276
- repositories 202
 - See also* Encyclopedia volumes
- requests
 - distributing 8, 13
 - installing cluster nodes and 191
 - remote procedure calls and 10
 - running cluster configurations and 14, 176
- resource file updates 264
- resource folders 255
- resource groups 46
- resources
 - accessing 255
 - creating cluster nodes and 13, 14, 176
 - deploying iServer and 204
 - installing Information Console and 246
 - installing iServer and 24
 - obtaining licenses and 275
- restoring Encyclopedia volumes
 - DB2 databases and 337, 347
 - PostgreSQL databases and 297, 306
 - SQL Server databases and 318, 326
- restricting iServer processes 283
- rollbacks 83
- RSSE applications 15, 16
- RSSE code 88
- running
 - Data Store Administrator 149
 - Data Store Upgrader 121, 123
 - encycsrvr11 processes 284, 288, 289

- iServer distribution setup script
 - running 206
- iServer processes 8, 25, 283
- iServer services 8, 24
- jobs 10
- multiple product releases 45, 47
- Open Security applications 46
- pg_dump utility 299, 303
- pg_restore utility 309, 312
- pgAdmin administration tool 41, 103
- queries 11, 13
- report designs 275, 276
- reports 203, 276
- spreadsheet reports 276
- Squirrel Data Exporter 148
- System Data Store Administrator 169

S

- sample reports 19
- sample volume 19
- saving
 - previous releases 87, 107
 - properties files 153
- scalability vii
- scheduling iServer installations 47
- schema names 30, 31, 121, 161, 163
- schema owners
 - creating 52, 225, 227
 - installing DB2 databases and 60, 61
 - installing iServer and 5
 - installing Oracle databases and 55, 56
 - installing PostgreSQL databases and 53, 54
 - installing SQL Server databases and 57, 58
 - setting passwords for 31, 163
- SCHEMA_FILE_NAME parameter 157, 171
- SCHEMA_NAME parameter 163, 164, 171
- SCHEMA_PASSWORD parameter 163, 173
- schemas
 - accessing 4
 - adding Encyclopedia to 167, 168
 - assigning privileges 53
 - backing up 7, 26, 168, 302
 - browsing 41
 - changing 41, 103
 - configuring system data store as 168
 - creating 27, 52, 146, 225, 226

- customizing metadata databases and 5
- deploying iServer and 206
- depopulating 160
- dropping 160
- importing volume data to 164
- initializing 27
- installing DB2 databases and 61
- installing Oracle databases and 55, 56
- installing PostgreSQL databases and 54, 55
- installing SQL Server databases and 58, 59
- naming 31, 53, 163, 225, 227
- populating 162, 166, 167
- preventing data loss and 7
- removing Encyclopedia from 166
- restoring 168, 170
- setting passwords for 162
- storing metadata and 4
- upgrades and 83, 92
- user-defined functions and 63
- script files 155
- SCRIPT_HOME parameter 157, 171
- scripts 203, 295
- search results 276
- searching online documentation xi
- security vii, 33, 277
- security application 46
- security database 46
- security systems 29
- server template names 240
- server templates 13, 176
- SERVER_DEFAULT parameter 255
- serverName property 216
- servers
 - See also* iServer
 - accessing PostgreSQL 41, 103
 - adding PostgreSQL 42
 - changing cluster machines and 282
 - creating clusters for 181, 237
 - creating shared folders for 178, 234
 - deploying Information Console to 254, 255
 - exceeding CPU licenses for 289
 - getting host names for 180, 236
 - installing iServer for external 29
 - installing stand-alone 29
 - integrating Actuate products with vii, 254
 - managing volume metadata and 5
 - preventing cache conflicts for 257
 - running as clusters 13, 14, 187
 - running iServer processes and 25
 - setting port numbers for 247
 - shutting down iServer and 9
 - starting PostgreSQL 211
- service names 69
- service packs 103, 295
- service privileges 26
- services. *See* specific iServer service
- ServIntTech directory 260
- servlet container 8
- set_tools_environment.bat file 122
- setup instructions (cloud deployments) 214
- Setup Type page 29
- setupiServer script
 - alternate databases and 214, 215, 218
 - desktop shortcuts and 210
 - environment variables and 205
 - installation options and 205
 - Postgre databases and 207
 - running 206
 - setting up clusters 239
 - setting up iServer only 219
 - stopping 209
- shared database schemas 167
- shared folders 178, 234
- shared licenses 277, 279
- shared repository 202
- sharing metadata databases 5
- shortcuts 41, 102, 118, 210
- shutting down. *See* stopping
- side-by-side migrations (Encyclopedia) 143
- side-by-side upgrades 82, 124
- SIDs (system identifiers) 69
- single-point media failures 313
- single-point node failure 13, 176
- SmartSheet Security Option 276
- snapshots 295
- SOAP processor 10
- SOAP-based messaging 8, 10
- Solaris systems 19
- Specify Profiles page 33, 94, 111
- Specify Windows Account Information page 250
- spreadsheet reports 276
- SQL languages 53
- SQL queries 11, 13, 62

- SQL scripts 53, 55, 57, 60, 221
- SQL Server databases
 - accessing documentation for 7
 - backing up transaction logs for 327, 330
 - connecting to 59
 - creating schema owners for 57, 58
 - creating users for 58, 59
 - exporting metadata to 156
 - installing 57, 65–79
 - managing Encyclopedia and 5
 - naming schemas for 59
 - running volume backup and restore
 - operations with 317, 318, 326
 - setting port number for 71
 - setting up iServer environments for 57–59
 - storing metadata and 4
 - taking Encyclopedia offline for 326
 - taking Encyclopedia online for 337
- SQL Server Management Studio
 - backing up Encyclopedia and 318–325
 - restoring Encyclopedia and 326–337
- SQL*Loader utility 155
- sqlcmd utility 323, 334
- sqlplus utility 314
- Squirrel Data Exporter
 - configuring 153
 - migrating iServer installations and 6
 - migrating volume metadata and 17
 - running 148
 - setting properties for 153
 - setting up export environments and 153, 155, 165
- Squirrel database 4
- SQUIRREL_DATA_HOME parameter 154
- SQUIRREL_EXPORT_FOLDER parameter 153, 154, 155
- SQUIRREL_EXPORT_FORMAT parameter 154
- SquirrelDataExporter.properties file 148
- stand alone servers 205
- stand-alone configurations 29
- stand-alone servers
 - creating clusters for 181, 237
 - creating shared folders for 178, 234
- starting
 - Configuration Console 102
 - Data Store Administrator 149
 - Data Store Upgrader 121
 - DB2 command line processor 344
 - DB2 Control Center 339
 - Encyclopedia volumes 9, 152, 231, 232
 - encycsrvr11 processes 288
 - Information Console 41, 257
 - iServer 9, 212, 221
 - iServer 11 service 34, 132
 - iServer processes 284
 - Management Console 41
 - pg_dump utility 299, 303
 - pg_restore utility 309, 312
 - pgAdmin administration tool 41
 - PostgreSQL servers 211
 - SQL Server Management Studio 319
 - Squirrel Data Exporter 148
 - volume partitions 231
 - status folder (encyc) 146
 - stopping
 - iServer 147, 212
 - iServer services 85, 105
 - NobleNet Portmapper processes 86, 106
 - Portmapper services 106
 - PostgreSQL for BIRT service 105
 - PostgreSQL processes 212
 - SUPERUSER parameter 163
 - superuser password 27, 31
 - SUPERUSER_PASSWORD parameter 163
 - superusers 31, 163
 - Supported Products Matrix xii, 254
 - Switch Help Location command 267
 - synchronous Factory service 11
 - syntax conventions (documentation) xii
 - system. *See* iServer System
 - system administrators. *See* administrators
 - System Configuration Password page 34, 95, 111
 - System Data Store Administrator
 - recreating system schemas and 170
 - restoring system schemas and 83
 - running 168, 169
 - setting PATH variable for 169
 - setting properties for 168, 170
 - system data stores 168
 - system databases 4, 5, 53, 60
 - See also* third-party databases
 - system failover operations 7

- system identifiers 69
- system metadata
 - backing up 26, 64
 - preventing loss of 7
 - storing 4, 11, 168
- system names 30, 68, 91
- system passwords 242
- system schema names 30, 31
- system schema owner 53, 55, 57, 60, 225
- system schema owner accounts 31
- system schema passwords 31
- system schemas
 - See also* database schemas
 - backing up 83, 168
 - creating 27, 54, 59, 61, 225
 - naming 225
 - restoring 168
 - storing metadata and 4
- system variables 285
- SYSTEM_NAME parameter 216
- SystemDataStoreAdministrator class 170

T

- table of contents (documentation) xi
- tables 5
- TABLESPACE_LOCATION parameter 163
- TABLESPACE_NAME parameter 163
- taking Encyclopedia online 9, 152, 231, 232
- taking iServer off line 147
- TAR files 202
- technical support xi
- TEMP_FOLDER_LOCATION parameter 255
- templates 255
- temporary directories 256
- temporary documents 11
- temporary files 255
- temporary license keys 278
- temporary licenses 279
- testing
 - connections 181, 237
 - data partitions 229, 230
 - database connections 224
 - desktop products 48
 - Information Console installations 257
 - new releases 47–49
- text files 282

- 3rd-Party Database Information page 69
- third-party command line utilities 297
- third-party database schemas. *See* database schemas
- third-party databases
 - See also* specific database type
 - accessing documentation for 7
 - changing tables in 5
 - creating indexes for 5
 - customizing connections for 18
 - deploying iServer and 205, 214–219
 - exporting metadata to 155
 - importing Encyclopedia and 164
 - installing iServer and 24, 52
 - managing Encyclopedia and 4, 7
 - preventing data loss for 26
 - running iServer clusters and 8, 11, 14
 - selecting 68
 - setting port numbers for 69, 71
 - storing cluster information and 202
 - storing metadata and 4
 - uninstalling 44, 104
 - upgrades and 83

- third-party deployment tools 254
- third-party RDBMS tools 7
- threads 254
- time zones 33, 255
- TIME_ZONE parameter 163
- TimeZones.xml 255
- TNS services 69
- Tomcat servers 247, 249
- transaction logs 327, 330
- Transact-SQL scripts 57
- transient data stores 255
- transient files 255
- TRANSIENT_STORE_PATH parameter 255
- Transparent Network Substrate services 69
- transport protocol (requests) 10, 13, 176
- troubleshooting 295
- typographic conventions (documentation) xii

U

- uninstalling
 - pgAdmin administration tool 44
 - pgAdmin utility 104
 - system databases 104

- uninstalling (*continued*)
 - third-party databases 44
- UNIX systems 16, 19
- Update tool 264
- updates (documentation) viii, xi
- updates (product) xi, 264
- upgrade options 82, 85, 125
- upgrade utilities 153
- upgrade_encyclopedia_data_store.bat file 122
- upgrade_encyclopedia_data_store.properties file 122
- Upgrader utility
 - configuring 153
 - running 121, 122
 - setting properties for 171, 172
- upgrades
 - application programming interfaces and 18
 - changing database schemas and 41
 - creating backups and 83
 - customizing metadata databases and 5
 - determining if needed 47
 - disk space and 45
 - Encyclopedia volumes 48, 120, 143
 - file structure changes and 102, 120
 - installing metadata databases and 4
 - installing product files and 84
 - iServer 45, 82, 84, 104, 120, 124
 - licensing options and 279
 - maintaining version consistency for 27
 - migrating volume metadata and 17
 - overwriting previous releases and 6
 - pgAdmin administration tool 97, 114, 137
 - PostgreSQL database schemas and 103
 - previous releases and 45
 - rolling back 83
 - running processes and 85, 104
- upgrded.bat file 122
- URLs
 - Actuate product information xii
 - Actuate technical support xi
 - Configuration Console 211, 223
 - documentation updates xi
 - evaluation copies 27
 - Information Console 210, 255, 257
 - JAVA SE Development Kit 203
 - Localization and Online Documentation
 - Update tool 264
 - Management Console 210, 232
 - Network File Systems 11
 - PDF documentation 264
 - RDBMS documentation 7
 - release notes xi
 - Supported Products Matrix 254
- usage logging applications 15
- usage reports 19
- user accounts
 - cluster configurations and 25
 - creating 25, 52
 - installing DB2 databases and 60
 - installing Information Console and 250
 - installing iServer and 24
 - managing 16
 - starting iServer 11 service and 33
 - upgrades and 82, 83, 94
- user credentials 5
- user information 15, 47
- user names 33, 34, 53
- user profiles 33
- User Rights Assignments dialog 26
- user tracking reports 19
- users
 - creating DB2 schemas and 60, 61
 - creating Oracle schemas and 56
 - creating PostgreSQL schemas and 54
 - creating SQL Server schemas and 58, 59
 - licensing and 33, 274
 - setting passwords for 33

V

- validating CPU binding 289, 290
- VARCHAR types 60
- VARGRAPHIC types 60
- verifying CPU binding 286
- View processes 287, 290
- View service 9, 11, 46
- viewers 276
- viewing
 - charts 46
 - cluster nodes 242
 - cube reports 275
 - folder properties 178, 234

- host names 180, 236
- IP addresses 180, 236
- licensing information 277, 278
- processor IDs 284
- ReadMe files 253, 263
- reports 11, 276
- ViewingService element 287
- Vista operating systems 24
 - See also* Windows systems
- Visual C++ libraries 203
- volume administrators. *See* administrators
- volume data 294
- volume data directories
 - DB2 environments 346, 347
 - Oracle environments 315, 316
 - PostgreSQL environments 303, 306
 - SQL Server environments 325, 326
- volume database connections 18
- volume database encoding 19
- volume databases 4, 5, 17
 - See also* third-party databases
- volume metadata
 - backing up 295, 314, 319, 323, 339, 344
 - changing 5
 - exporting Encyclopedia and 153, 155
 - migrating 17
 - restoring 316, 331, 334, 348, 355
 - storing 4, 11
- volume metadata database 30, 68
- volume names 30, 162, 164, 230, 247
- volume partition names 145
- volume partition paths 145
- volume partitions
 - creating 145
 - manual side-by-side upgrades and 143
 - renaming default 144
 - setting default 228
 - setting primary 151, 230
- volume schema names 30, 31, 54, 55, 59
- volume schema owner 54, 56, 58, 61, 227
- volume schemas
 - See also* database schemas
 - assigning privileges to multiple 53
 - backing up 302
 - creating 27, 146, 226
 - DB2 databases and 61
 - naming 227

- PostgreSQL databases and 55
- SQL Server databases and 59
- storing metadata and 4
- VOLUME_NAME parameter 163
- VolumeImport.properties file 149
- volumes. *See* Encyclopedia volumes

W

- WAR files 247, 254, 255, 256
- watermarks 278
- web administrator privileges 247
- web administrators. *See* administrators
- web applications 256
- web archive files. *See* WAR files
- web browsers 257
- web pages 257, 276, 277
- web servers 254
 - See also* servers
- web services 10
- whitespace (property files) 153
- Windows platforms
 - creating shared folders on 178, 234
 - displaying IP addresses for 180, 236
 - turning off firewalls for 179, 235
- Windows Security Alerts 209
- Windows systems
 - binding iServer processes to 284, 284–288
 - collecting machine information for 279
 - configuring firewalls for 247
 - creating iServer images for 203
 - distributing iServer System over 6
 - installation prerequisites for 24, 52
 - installing Configuration Console for 246
 - installing Information Console for 246, 247, 254, 257
 - installing iServer clusters for 184, 185–198
 - installing iServer for 27
 - installing iServer Integration Technology for 260
 - installing JDK files for 46, 247
 - installing JRE files for 247
 - installing localization and documentation files for 264
 - installing Management Console for 246
 - running iServer processes and 25, 85, 104
 - setting classpaths for 46

Windows systems (*continued*)

- setting default locale for 33
- setting up user accounts for 24, 25, 26, 52
- starting iServer service for 33, 132
- testing new releases for 47–49
- upgrading Encyclopedia volumes for 148
- upgrading iServer for 45, 82, 84, 104, 120, 124

Work Unit Licenses 275

X

XML files 282

Z

ZIP files 202