

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 - 2011 by Actuate Corporation. All rights reserved. Printed in the United States of America.

Contains information proprietary to:
Actuate Corporation, 2207 Bridgepointe Parkway, 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, 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:

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, Shindig, Struts, Tomcat, Xerces, Xerces2 Java Parser, and Xerces-C++ XML Parser. Bits Per Second, Ltd. and Graphics Server Technologies, L.P.: Graphics Server. Bruno Lowagie and Paulo Soares: iText, licensed under the Mozilla Public License (MPL). Castor (www.castor.org), ExoLab Project (www.exolab.org), and Intalio, Inc. (www.intalio.org): Castor. Codejock Software: Xtreme Toolkit Pro. DataDirect Technologies Corporation: DataDirect JDBC, DataDirect ODBC. 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). Jason Hsueh and Kenton Varda (code.google.com): Protocole Buffer. ImageMagick Studio LLC.: ImageMagick. InfoSoft Global (P) Ltd.: FusionCharts, FusionMaps, FusionWidgets, PowerCharts. Mark Adler and Jean-loup Gailly (www.zlib.net): zlib. 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. KL Group, Inc.: XRT Graph, licensed under XRT for Motif Binary License Agreement. LEAD Technologies, Inc.: LEADTOOLS. 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). Oracle Corporation: Berkeley DB. PostgreSQL Global Development Group: pgAdmin, PostgreSQL, PostgreSQL JDBC driver. 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. Sun Microsystems, Inc.: JAXB, JDK, Jstl. ThimbleWare, Inc.: JMemcached, licensed under the Apache Public License (APL). World Wide Web Consortium (W3C)(MIT, ERCIM, Keio): Flute, JTIty, Simple API for CSS. XFree86 Project, Inc.: (www.xfree86.org): xvfb. Yuri Kanivets (code.google.com): Android Wheel gadget, licensed under the Apache Public License (APL). 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. 111021-2-430342 August 2, 2011

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 Installing BIRT iServer for Windows	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 data- base 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	13
Administering iServer System	15
About Migration and Administration Tools	16
Using JDBC to connect to an Encyclopedia volume database	18
API Compatibility	19
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	45
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	46
About upgrading an iServer with resource groups	46
About the Java Software Development Kit	46
Accessing JAR files for report generation	47
Gathering LDAP information	47
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	53
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
Backing up iServer system and Encyclopedia volume metadata	62
Installing an Encyclopedia volume that uses an alternative database	63

Chapter 4	
Upgrading BIRT iServer	77
Preparing to upgrade BIRT iServer	78
Creating an account with Windows administrator privileges	78
Backing up iServer system and Encyclopedia volume metadata	79
Performing an automatic in-place upgrade	80
About the acserverconfig.xml files saved by the upgrade process	97
Accessing the PostgreSQL Database Server using the pgAdmin utility	98
Performing an in-place upgrade from an earlier version of iServer Release 11	99
Performing a manual side-by-side upgrade	113
Performing a manual side-by-side migration	133
Specifying Squirrel Data Exporter properties	142
Specifying Encyclopedia Data Store Administrator properties	144
Specifying System Data Store Administrator properties	150
Performing operations using System Data Store Administrator utility	152
Creating and Populating a System Schema	152
Performing operations using Encyclopedia Data Store Administrator utility	153
Importing One or More Volumes into a New Schema	153
Importing One or More Volumes into a Populated Schema	154
Exporting All Volumes from a Schema	154
Exporting A Single Volume from a Schema	154
Deleting All Volumes from a Schema	155
Deleting a Single Volume from a Schema	155
Creating a New Volume in an Empty Schema	156
Creating a New Volume in a Populated Schema	156
Creating and Initializing a New Volume in a New Schema	157
Configuring Squirrel Data Exporter and System or Encyclopedia Data Store Administrator properties files	157
Chapter 5	
Installing a BIRT iServer cluster	159
Installing a BIRT iServer cluster node	160
Creating an account with Windows administrator privileges	161
Performing a cluster node installation using the wizard	162
Adding a node to a cluster	175
Chapter 6	
Installing BIRT iServer in a cloud	183
Understanding a BIRT iServer cloud deployment	184
Preparing to deploy BIRT iServer in a cloud environment	184
Deploying an iServer distribution package	186

Specifying AC_SERVER_HOME	189
Specifying AC_DATA_HOME	190

Chapter 7

Installing Information Console 193

Preparing to install Information Console	194
About installing from an FTP download	194
About performing a full installation	194
Installing Information Console on Windows	194
Using the installation wizard	195
Using the WAR file to install	201
General deployment tasks	201
Preparing the server	202
Preparing the WAR file	202
About clusters of servers	204
Avoiding cache conflicts after installing	204
Testing the installation	205

Chapter 8

Installing iServer Integration Technology and Documentation 207

Installing iServer Integration Technology	208
Installing the localization and documentation files	212
About accessing online help	215
Searching PDF manuals using master-index.pdx	216

Part 3

Licensing

Chapter 9

Licensing BIRT iServer 221

Working with licensing	222
Understanding the licensing options	222
About license files	225
About CPU binding	226
Understanding node-key licensing	226
Obtaining the installation license key file	226
Collecting machine information for a node-key license	227
Installing the license key	228
Understanding CPU binding	229
Binding a BIRT iServer to processors on a Windows machine	230
Binding to specific CPUs	230

Binding to multiple-core CPUs	232
Binding an Actuate process to a processor	232
About processors and hyperthreading	233
Checking BIRT iServer bound processors	234
Determining the number of processors an iServer System uses	234
Understanding CPU binding validation while iServer is running	235
Understanding CPU binding validation when an Encyclopedia volume comes online	235
Understanding CPU binding validation when running iServer processes	235
Configuring e-mail for CPU license problems	236

Part 4

Backing Up

Chapter 10

Backing up an Encyclopedia volume	239
Performing an Encyclopedia volume backup	240
Managing the backup and recovery of Encyclopedia volume metadata and data files	240
Using RDBMS and file system backup utilities	241
Avoiding conflict with the file purging process	242
Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database	243
Backing up an Encyclopedia volume using pgAdmin	244
Backing up an Encyclopedia volume using pg_dump	249
Restoring an Encyclopedia volume using pgAdmin	252
Restoring an Encyclopedia volume using pg_restore	259
Backing up and restoring an Encyclopedia volume that uses an Oracle database	261
Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)	261
Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp)	263
Backing up and restoring an Encyclopedia volume that uses a SQL Server database	265
Backing up and restoring an Encyclopedia volume	265
Backing up an Encyclopedia volume using SQL Server Management Studio	266
Restoring an Encyclopedia volume using SQL Server Management Studio	274
Backing up and restoring an Encyclopedia volume that uses a DB2 database	286
Backing up an Encyclopedia volume using DB2 Control Center	287
Restoring an Encyclopedia volume using DB2 Control Center	298
Index	315

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



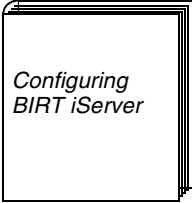
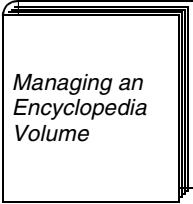
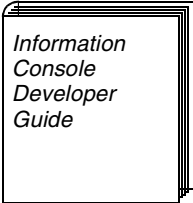
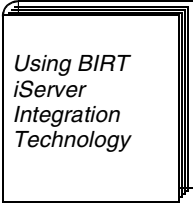
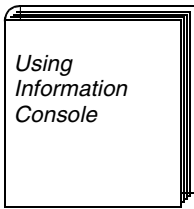
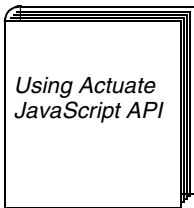
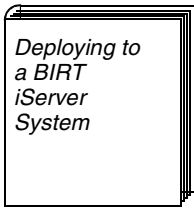
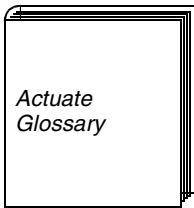

For information about this topic	See the following resource
Installing BIRT iServer for Linux and UNIX	 <p><i>Installing BIRT iServer for Linux and UNIX</i></p>
Installing BIRT iServer for Windows	 <p><i>Installing BIRT iServer for Windows</i></p>
Late-breaking information and documentation updates	Release notes and updated localization, HTML help, and PDF files posted on birt-exchange.com and Actuate Support
<p>Configuring BIRT iServer</p> <p>Use Configuration Console to:</p> <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.

If you request a physical package, install the files using the Online Documentation and Localization Resource Files DVD, which ships as part of your Actuate software package. If you select the typical setup when you install from the DVD, the installation creates the Program Files\Actuate11\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	Dim Text1 As String
File names	Initial capital letter, except where file names are case-sensitive	Detail.roi
Key combination	A + sign between keys means to press both keys at the same time	Ctrl+Shift
Menu items	Capitalized, no bold	File
Submenu items	Separated from the main menu item with a small arrow	File→New
User input or user response	Monospace	M*16*
User input in XML and Java code	Monospace italics	chkjava.exe <i>cab_name.cab</i>

Syntax conventions

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

Table I-3 Syntax conventions

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

Table I-3 Syntax conventions

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

About *Installing BIRT iServer for 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 Actual 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 an 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

1

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 3, 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 an upgrade from Release 11 Service Pack 2 to Service Pack 3
- 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 an iServer system and Encyclopedia volume schemas in the Release 11 environment, see Chapter 10, “Backing up an Encyclopedia volume,” later in this book.

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 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.

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.

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.

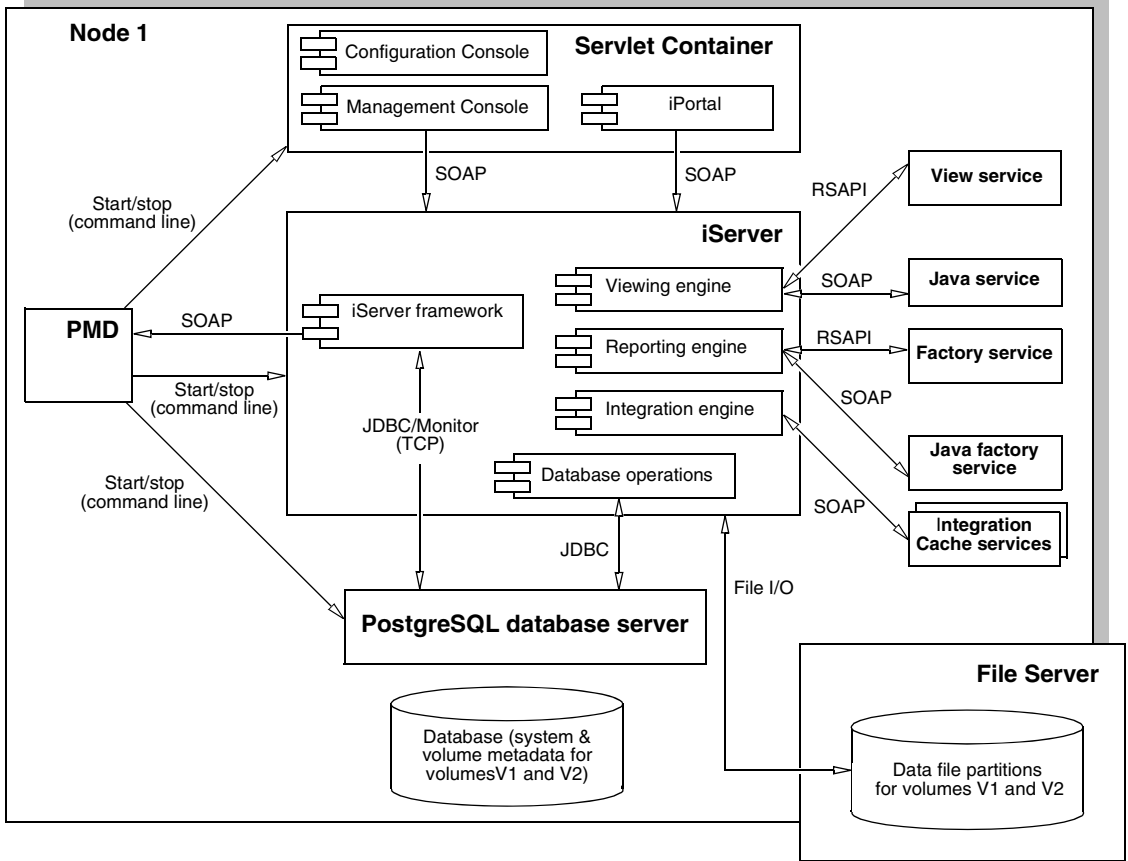


Figure 1-1 iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

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>

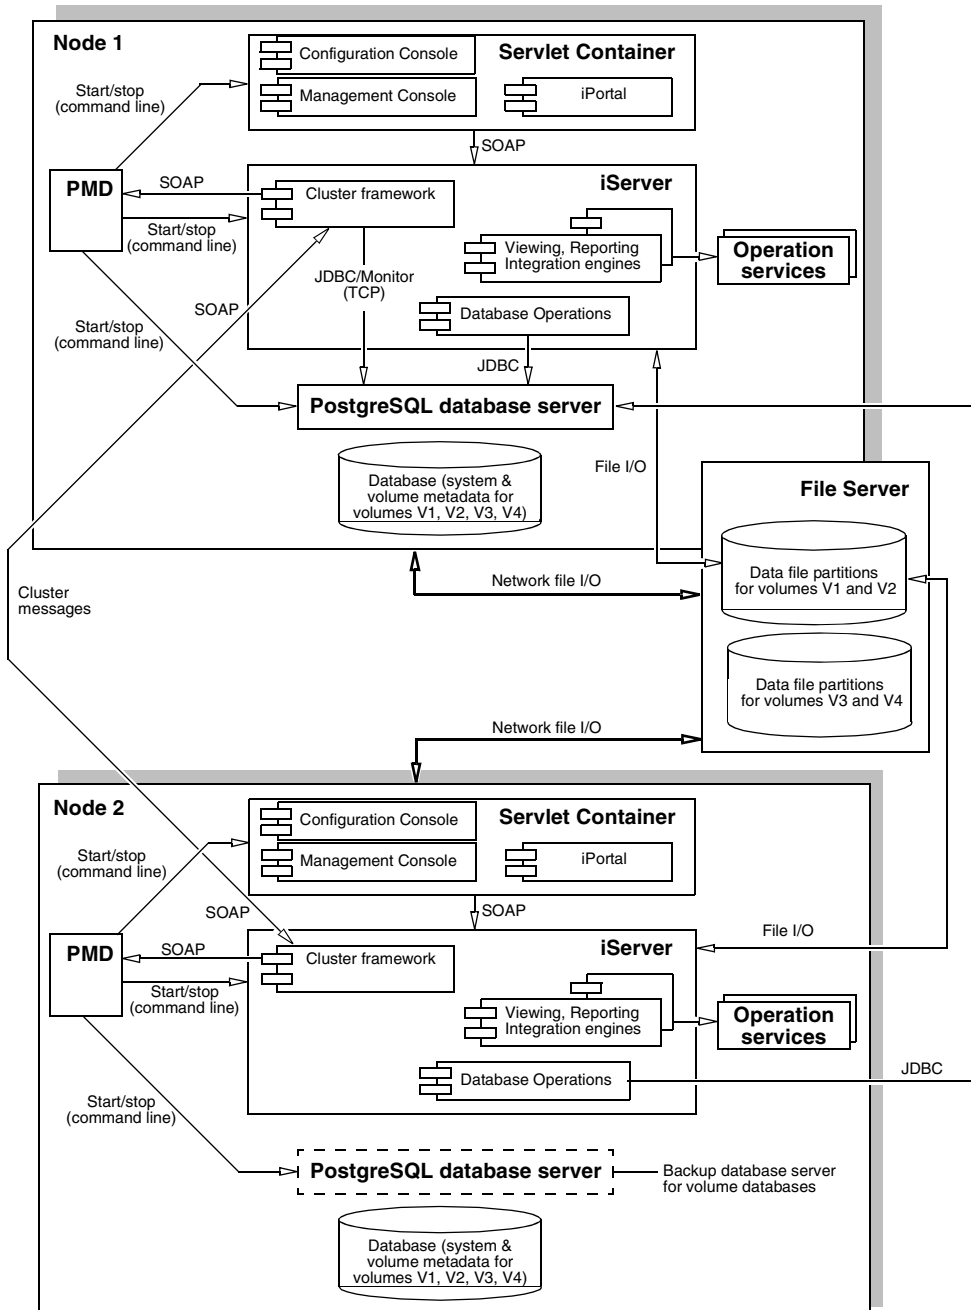


Figure 1-2 iServer RDBMS process architecture for a clustered, two-node, two-volume, OOTB database

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>

Understanding process flow in an iServer cluster

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

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 deployment and configuring the appropriate environment variables in `acpmdconfig.xml`, the administrator launches a deployed iServer image from the command line by passing the necessary arguments or creates a script that executes the commands. 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 `acserverconfig.xml`, located in a shared configuration home directory. Nodes with the same cluster ID, running on the same sub-net, automatically detect and join each other to form the cluster.

- Using Actuate Server Integration Technologies scripts and tools to develop client applications and extend iServer functionality

The Actuate Information Delivery application programming interface (IDAPI) supports integrating and administering iServer using extensible markup language (XML) and the simple object access protocol (SOAP). Using the IDAPI, developers can create applications that perform such tasks as scheduling a custom event, running a 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 the appropriate iServer installation manual, either *Installing BIRT iServer for Windows* or *Installing BIRT iServer for Linux and UNIX*.

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 3 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.

Before 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 this user 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.

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 or install from a product DVD. 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 If you downloaded iServer, run `ActuateBIRTiServerEnterprise.exe`. If you have a DVD or ftp distribution, run `setup.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/Actuate11SP3/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.

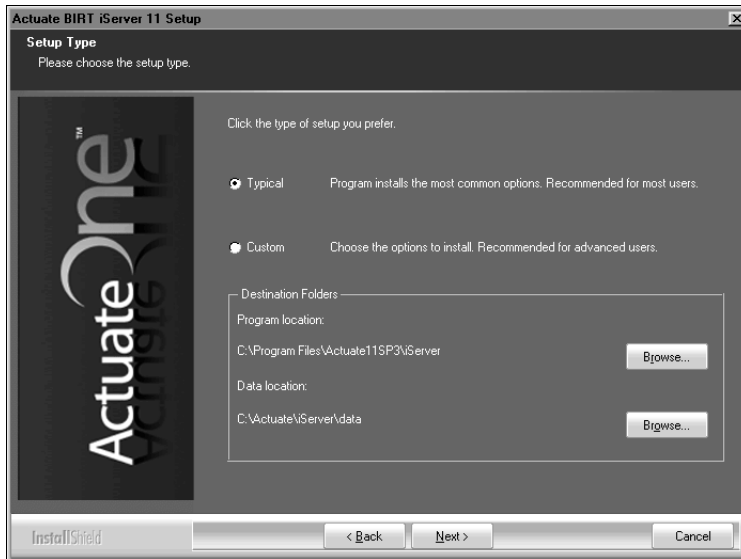


Figure 2-4 Specifying typical or custom setup type

- 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.

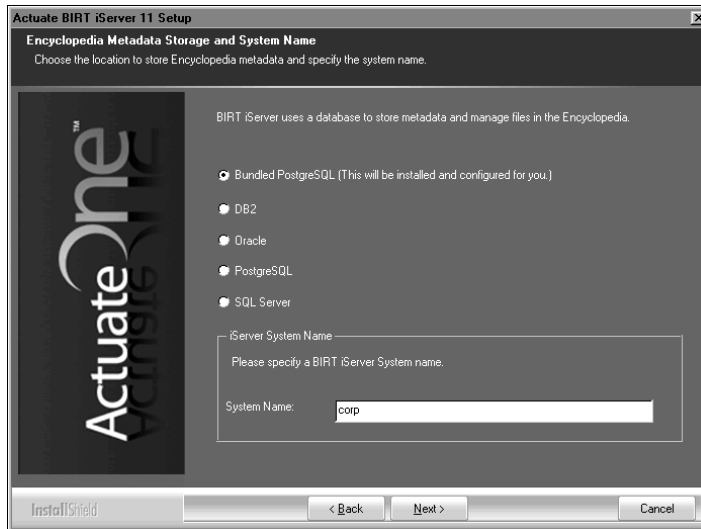


Figure 2-5 Selecting the Encyclopedia volume metadata database to install

- 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.

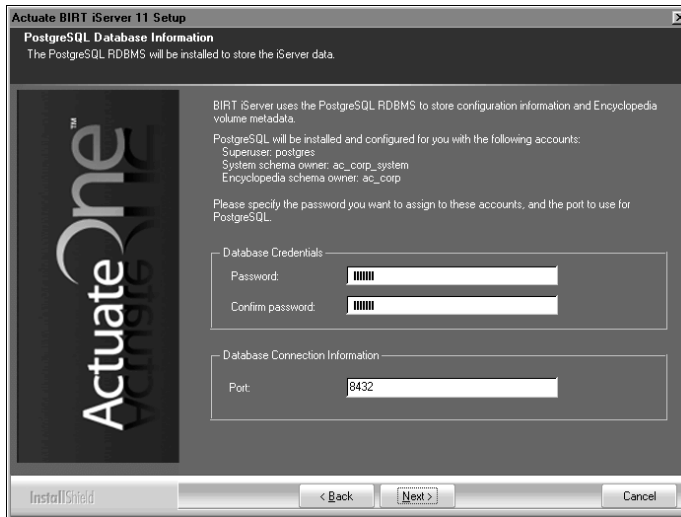


Figure 2-6 Specifying PostgreSQL database information

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 2-7.




Figure 2-7 Setting the Windows local security policy

- 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 2-8. Alternatively, choose Try out the product using the included evaluation license if you do not have a purchased license. Choose Next.

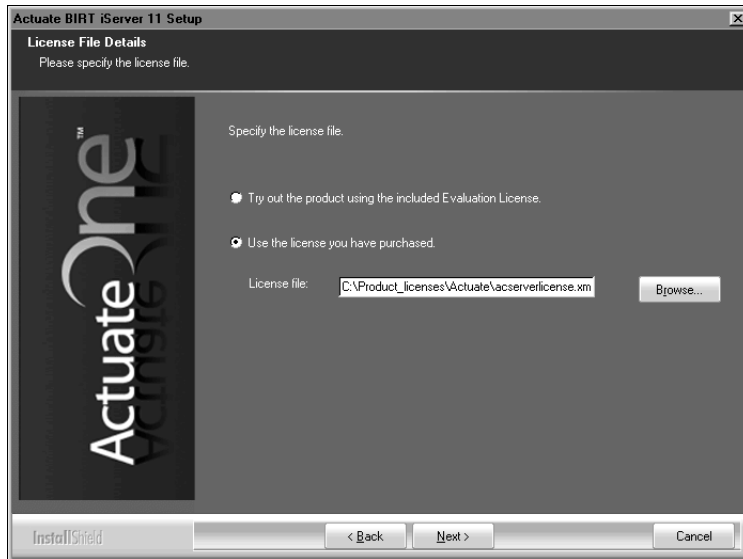


Figure 2-8 Specifying the license file

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 2-9.




Figure 2-9 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 2-10.

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

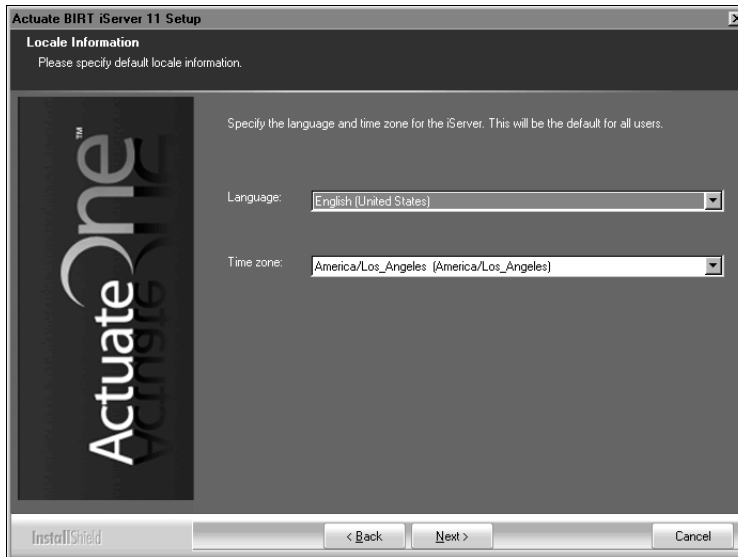


Figure 2-10 Specifying locale information

- 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 2-11. 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 2-11. 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.

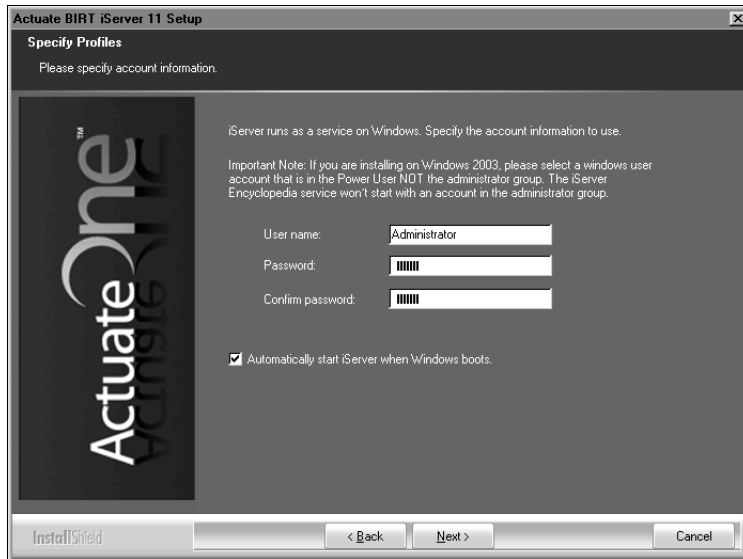


Figure 2-11 Specifying an account for running the iServer service

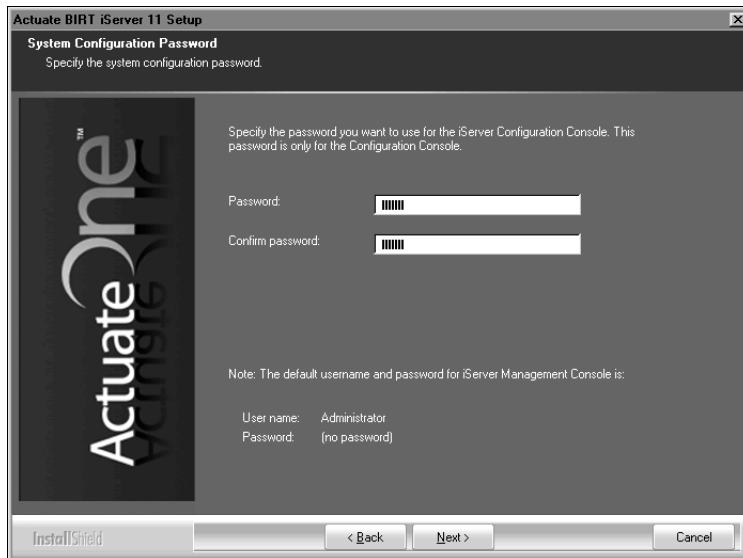


Figure 2-12 Specifying the password for using Configuration Console

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.

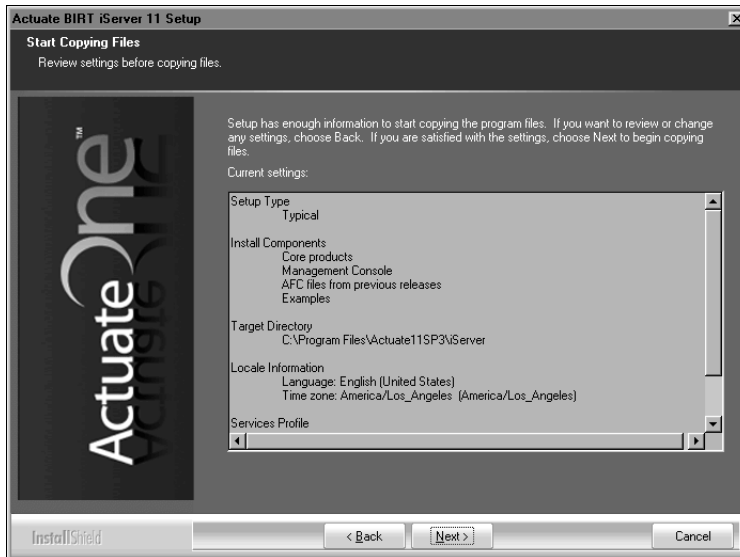


Figure 2-13 Reviewing settings before copying files



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

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

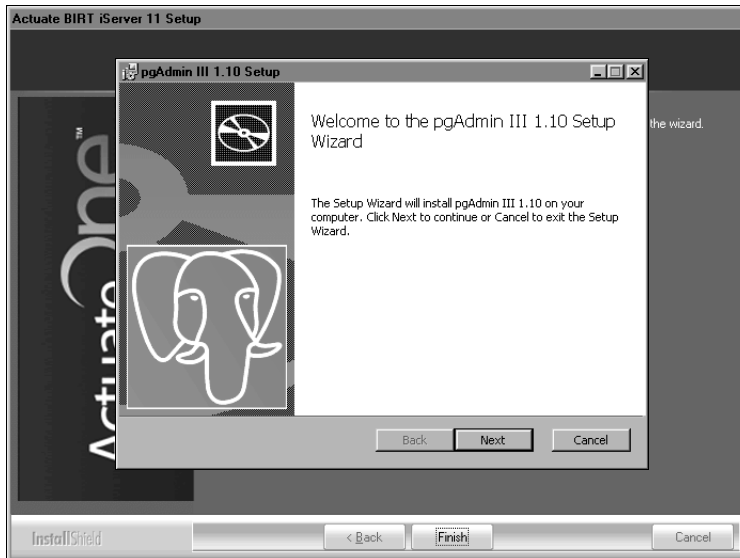


Figure 2-17 Viewing Actuate BIRT iServer and pgAdmin III Setup

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.

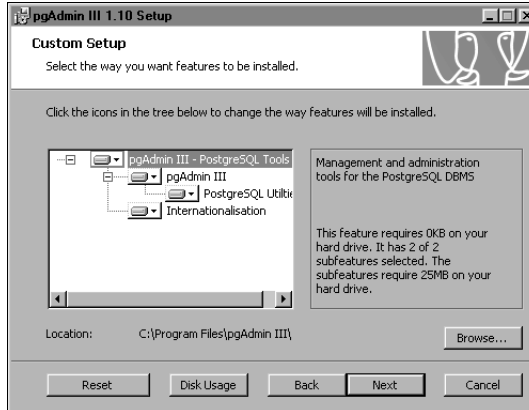


Figure 2-19 Viewing the features to be installed

- 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.

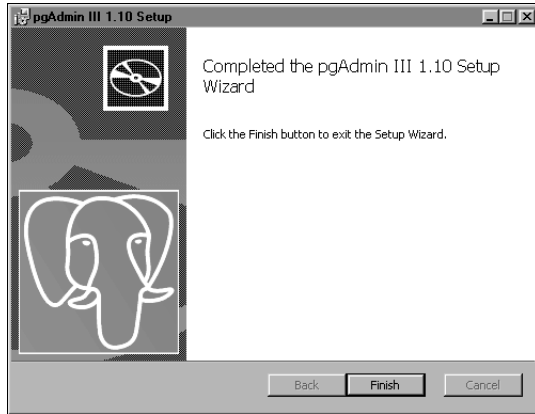


Figure 2-22 Choosing Finish

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:

<http://www.actuate.com>

Alternatively, you can install the online help and manuals in PDF format using the Actuate Localization and Online Documentation media. Figure 2-24 shows the prompt.



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

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.

Configuration Console Log In appears, as shown in Figure 2-26. iServer is ready for use.




Figure 2-26 Configuration Console Log In appears

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 2-27.

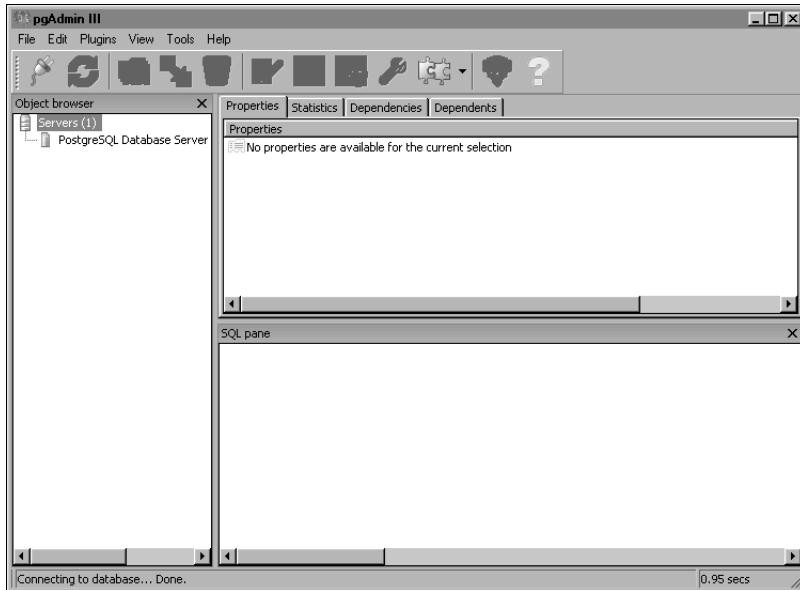


Figure 2-27 Viewing pgAdmin III

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

Table 2-1 New Server Properties

Property	Value
Restore env	Selected
Service	Actuate 11 PostgreSQL for BIRT iServer
Connect now	Selected

New Server Registration—Properties appears, as shown in Figure 2-28.

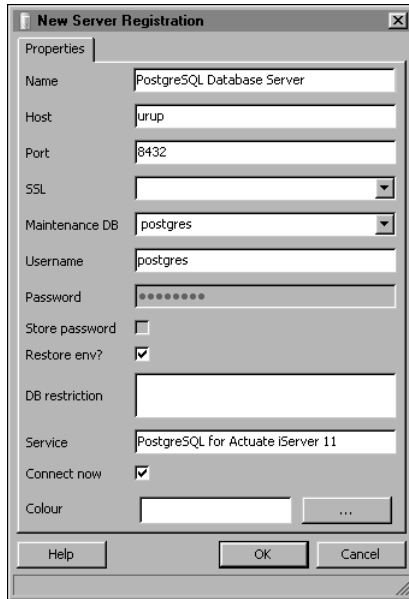


Figure 2-28 Registering a new server

Choose OK.

On pgAdmin III, expand PostgreSQL Database Server, as shown in Figure 2-29.

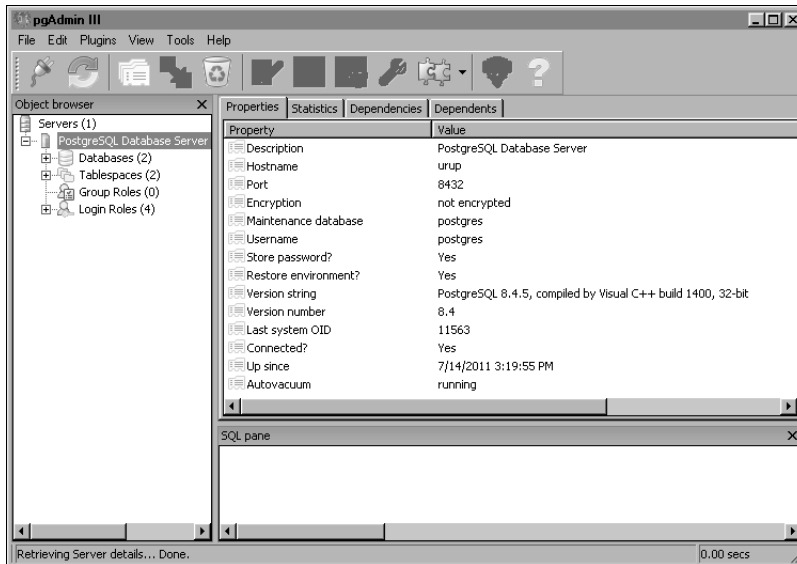


Figure 2-29 Viewing PostgreSQL Database Server properties

- In Object browser, expand Databases—iServer, iServer—Schemas, Tablespaces, and Login Roles to review the server and database installation, as shown in Figure 2-30.

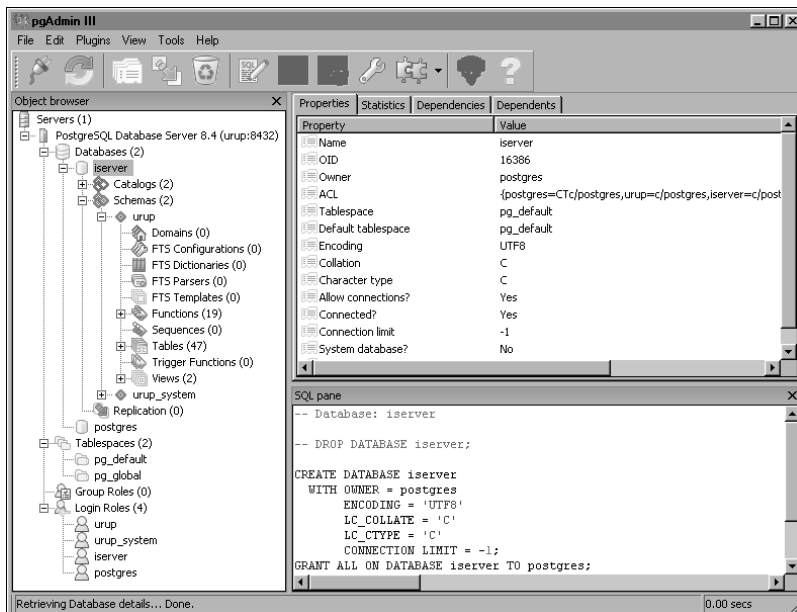


Figure 2-30 Browsing the PostgreSQL Database Server installation

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 3 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 volumes and archive volumes. 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 Encyclopedia volumes by installing Actuate 11 in the same directory, overwriting Actuate 10.

Uninstall BIRT iServer Release 11 from your production staging area and reinstall it in place of your existing Actuate 10 installation. The installer upgrades the Encyclopedia volumes when it starts.

- If you are running both Actuate 10 and Actuate 11, you can migrate volumes to the Actuate 11 location.

Use the Actuate 11 Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate an Actuate 10 Encyclopedia volume to the Actuate 11 Encyclopedia location. Confirm that all partitions for Actuate 11 are in the active state before you use these utilities.

- Install upgraded design and 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 another data store, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance to store Encyclopedia volume metadata. This chapter describes how to install a new instance of BIRT iServer Release 11, using an alternative data store.

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

Creating 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.”

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;
```

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 ANY 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:

```
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 ANY 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 USINGUCA500R1_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;
```

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 or install from a product DVD. 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 Microsoft SQL Server as an alternative data store.

Installing an Encyclopedia volume that uses an alternative database

The following procedures use a pre-existing Microsoft SQL Server 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 If you downloaded iServer, run the ActuateBIRTiServer.exe file. If you have a DVD or ftp distribution, run setup.exe in the iServer directory. 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.

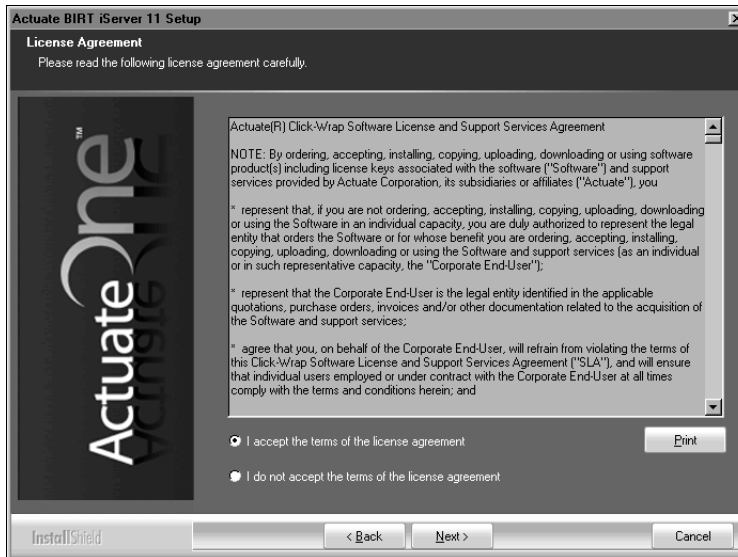


Figure 3-2 Accepting the license agreement

- 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/Actuate11SP3/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:/Actuate11/iServer/data.

Choose Next.

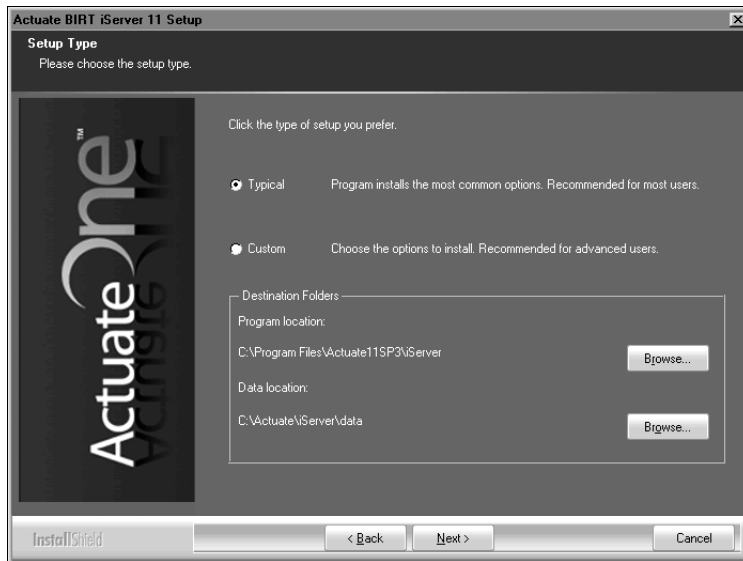


Figure 3-4 Specifying Typical setup type

- 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 SQL Server 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.

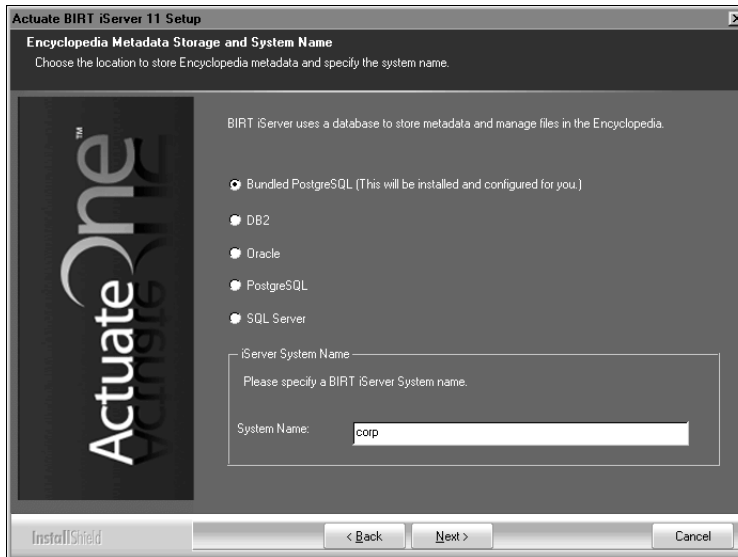


Figure 3-5 Choosing a third-party database

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:

- **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.
- **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. Choose Next.



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

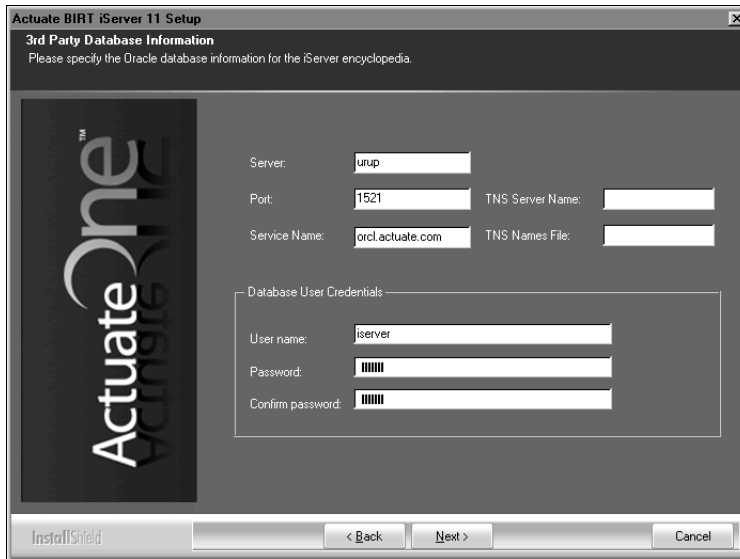


Figure 3-7 Specifying third-party database information for Oracle

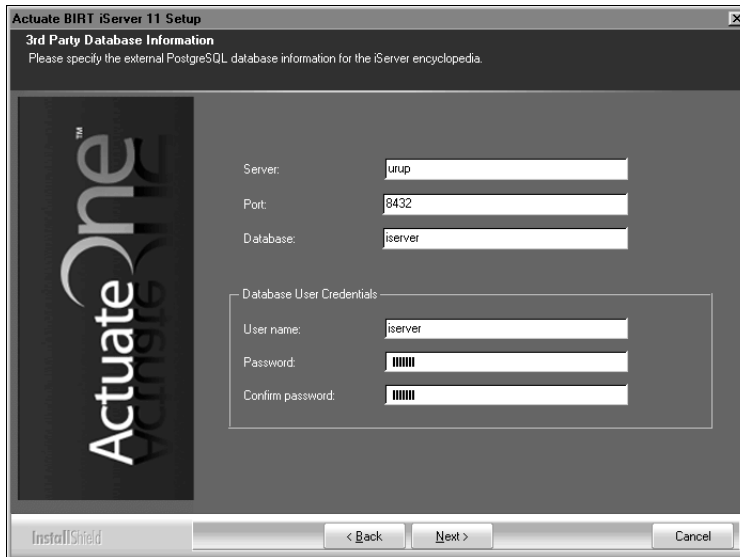


Figure 3-8 Specifying third-party database information for PostgreSQL

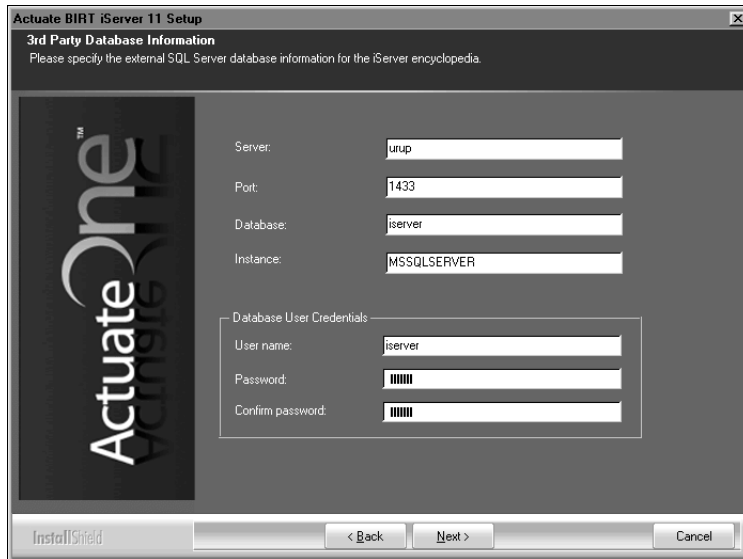


Figure 3-9 Specifying third-party database information for SQL Server

- 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.

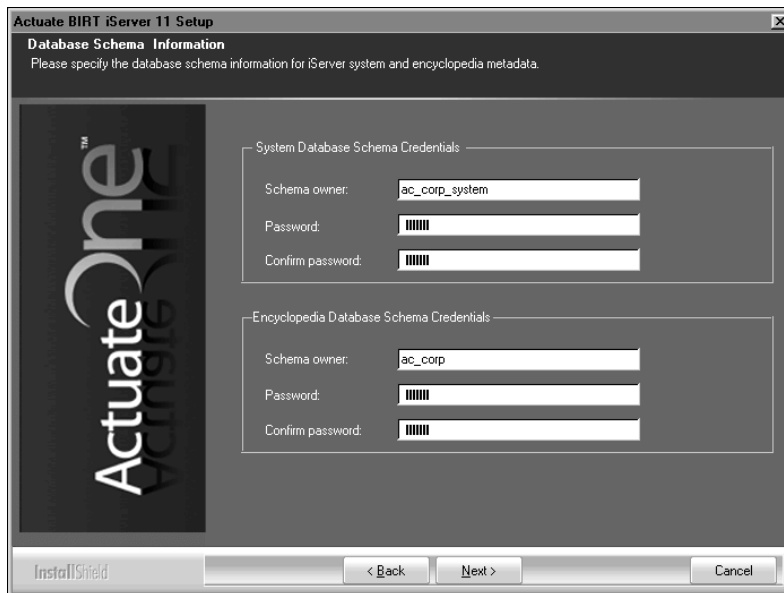


Figure 3-10 Specifying System and Encyclopedia schema passwords

- 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.

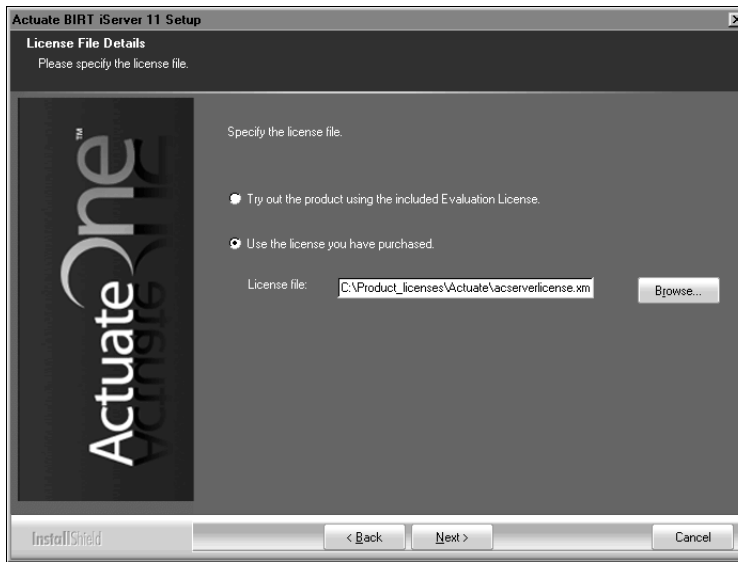


Figure 3-11 Specifying the license file

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.

- 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.

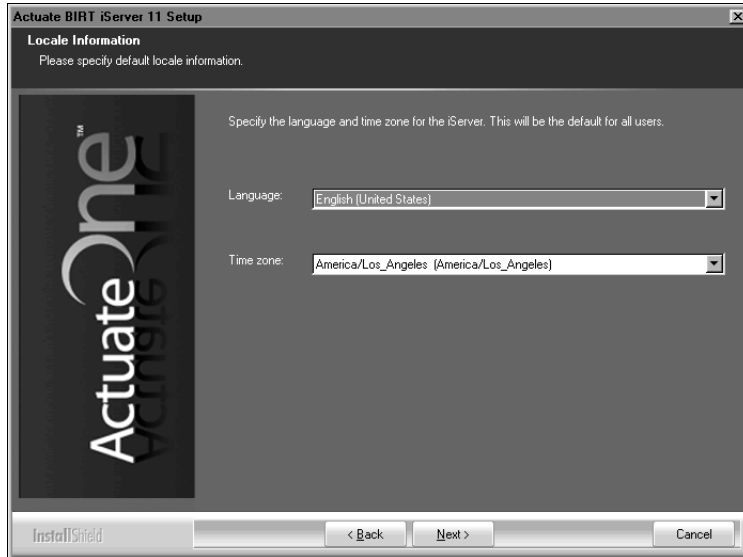


Figure 3-13 Specifying locale information

- 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.

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 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 3-14 Specifying the account for running the iServer service

- 11** In System Configuration Password, type and confirm a password for Configuration Console, as shown in Figure 3-15. 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 3-15 Specifying the password for using Configuration Console

12 In Start Copying Files, review the settings shown in Figure 3-16. Choose Next.

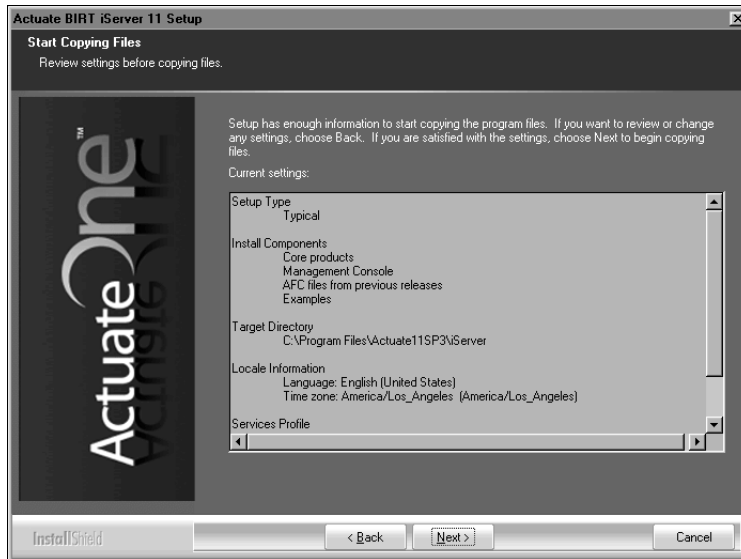


Figure 3-16 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 3-17.



Figure 3-17 Viewing setup status

13 Choose Finish to exit the wizard, as shown in Figure 3-18.



Figure 3-18 Exiting the installation wizard

- 14** The installation program prompts you to install the online help and manuals using the Actuate Localization and Online Documentation media, as shown in Figure 3-19.



Figure 3-19 Viewing the install online help and manuals prompt

The installation program installs shortcuts on the desktop, as shown in Figure 3-20.


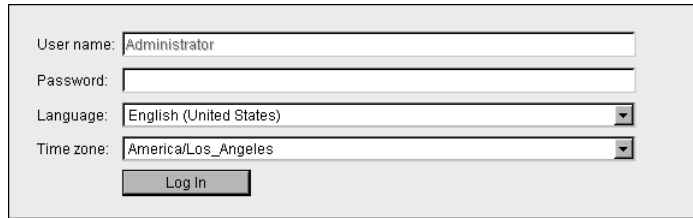


Figure 3-20 Viewing iServer shortcuts on the desktop

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.

Configuration Console Log In appears, as shown in Figure 3-21. iServer is ready to use.



The image shows a login form with the following fields and options:

- User name: Administrator
- Password: (empty field)
- Language: English (United States) (dropdown menu)
- Time zone: America/Los_Angeles (dropdown menu)
- Log In button

Figure 3-21 Configuration Console Log In appears

4

Upgrading BIRT iServer

This chapter discusses the following topics:

- Preparing to upgrade BIRT iServer
- Performing an automatic in-place upgrade
- Performing a manual side-by-side upgrade

Preparing to upgrade 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, such as Release 10 Service Pack 1, 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.

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.”

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 or install from a product DVD. 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

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

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

The automatic option also performs the following operations during installation:

- Creates the system and volume schemas, initializing these schemas with basic configuration information.
- Creates the iServer configuration file, specifying system, volume, and connection information for the default installation

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.

This installation program can encounter a problem overwriting a file linked with a running process. Be sure to stop the iServer service and the NobleNet Portmapper service before proceeding with the upgrade.

How to stop the 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.

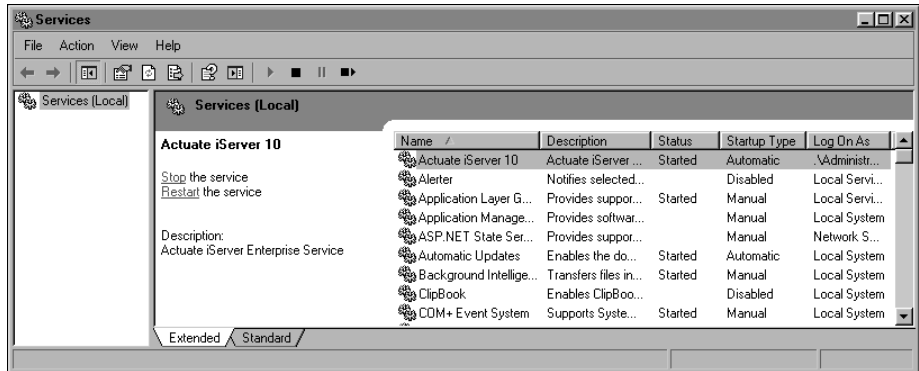


Figure 4-1 Stopping the Actuate iServer 10 service

The service stops, as shown in Figure 4-2.

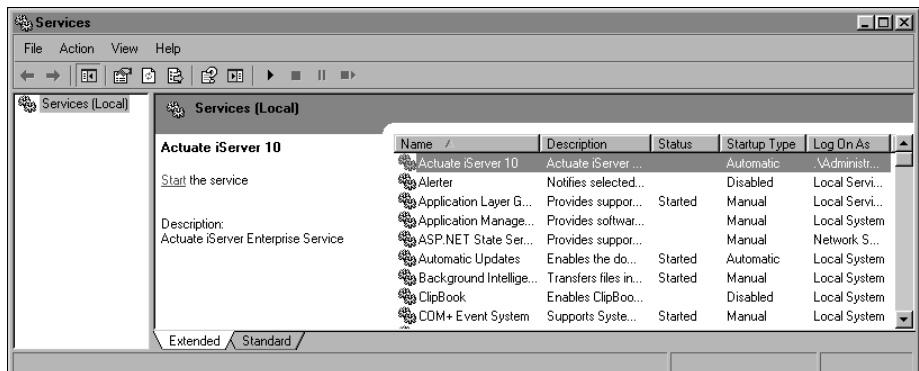


Figure 4-2 Viewing Services after BIRT iServer service stops

How to shut down 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"
```

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 perform an automatic in-place upgrade from an earlier major release to iServer Release 11

- 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 /etc directory
 - acpmdconfig.xml in the /etc directory
 - RSSE code and associated files if you use the Open Security option
- 2 If you downloaded iServer, run the ActuateBIRTiServer.exe file. If you have a DVD or ftp distribution, run setup.exe in the iServer directory. 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.

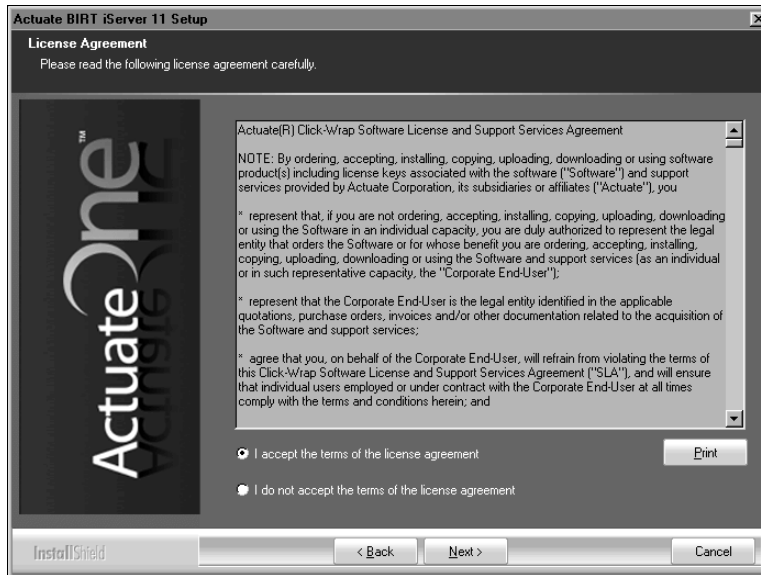


Figure 4-4 Viewing the license agreement

- 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/Actuate11SP3/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`.

- 3 Select Upgrade Existing iServer and Migrate the Encyclopedia Data.
- 4 Check that the default directory in Existing iServer Location matches the location of the old files that you are upgrading, as shown in Figure 4-6. Choose Next.

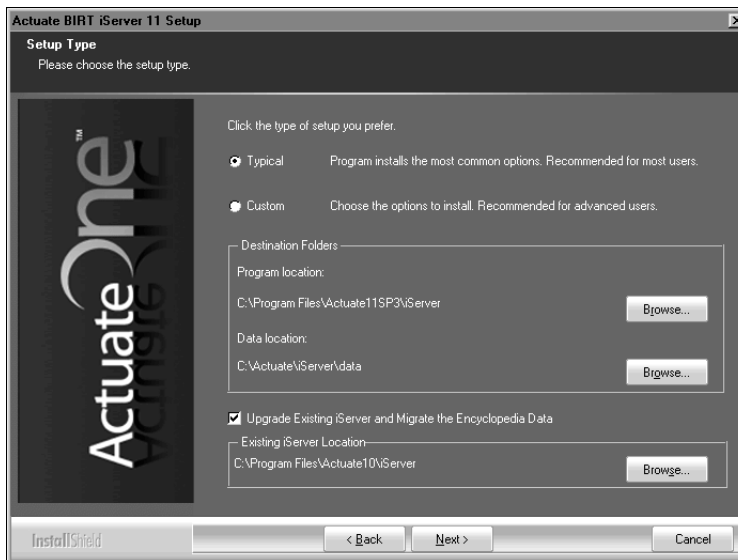


Figure 4-6 Choosing to upgrade iServer and migrate the Encyclopedia

- 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

If you have not shut down the iServer service, a prompt appears, instructing 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

- 7 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.

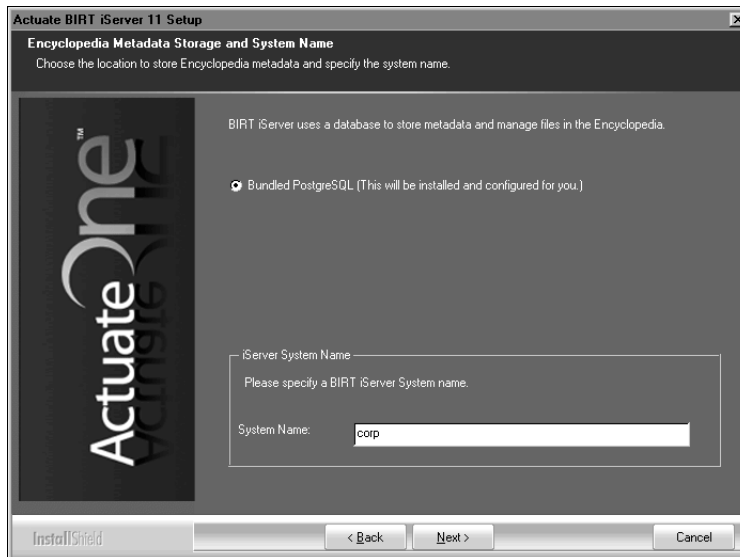


Figure 4-9 Accepting metadata database type

8 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.

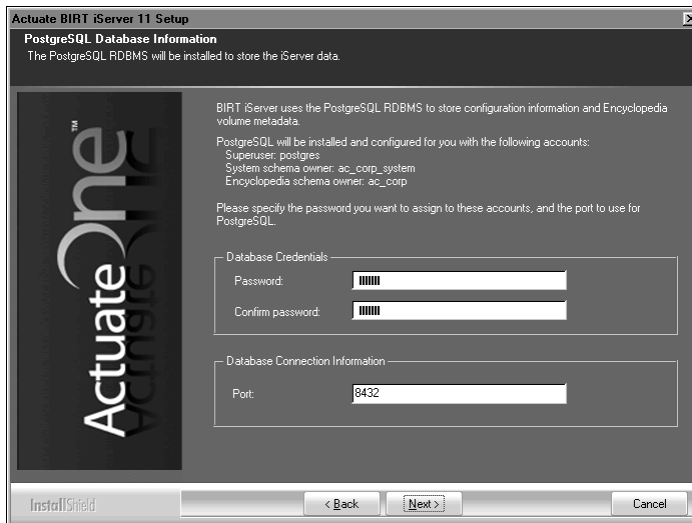


Figure 4-10 Specifying PostgreSQL database information

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 4-11.




Figure 4-11 Setting the Windows local security policy

- 9 Choose Yes to update your license, as shown in Figure 4-12.




Figure 4-12 Updating licensing

- 10 In License File Details, browse to and choose the license file, as shown in Figure 4-13. Choose Next.

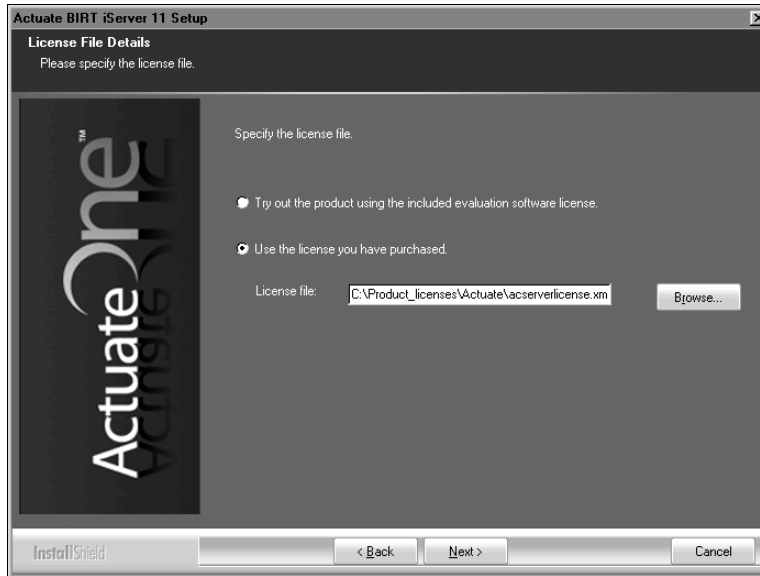


Figure 4-13 Specifying the license file

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 4-14. Choose Yes to continue, then choose Next.




Figure 4-14 Viewing the named user license question

- 11 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 4-15. Alternatively, choose the language and locale settings for your region.

- 12** 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 4-16. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

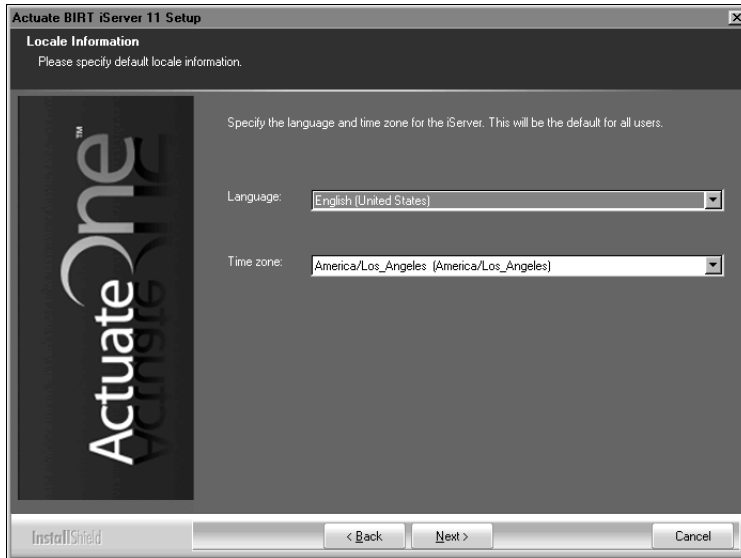


Figure 4-15 Specifying locale information




Figure 4-16 Specifying the account for running the iServer service

- 13** 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.

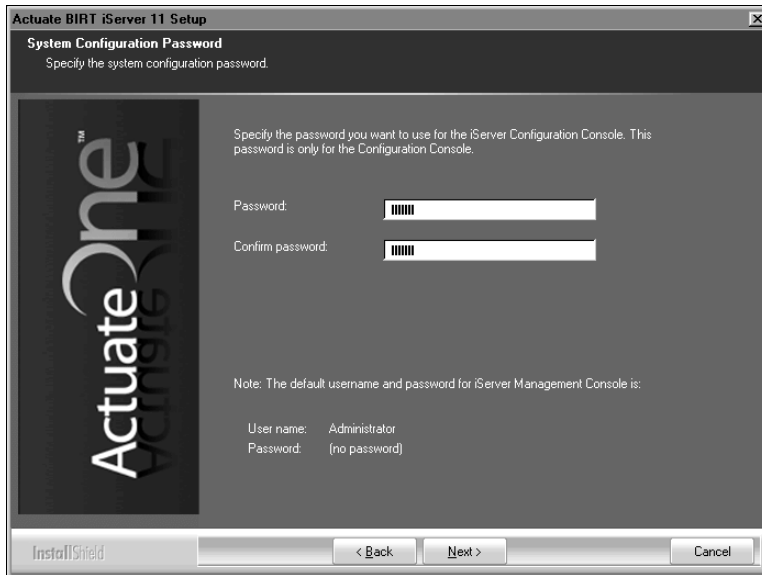


Figure 4-17 Specifying the password for using Configuration Console

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

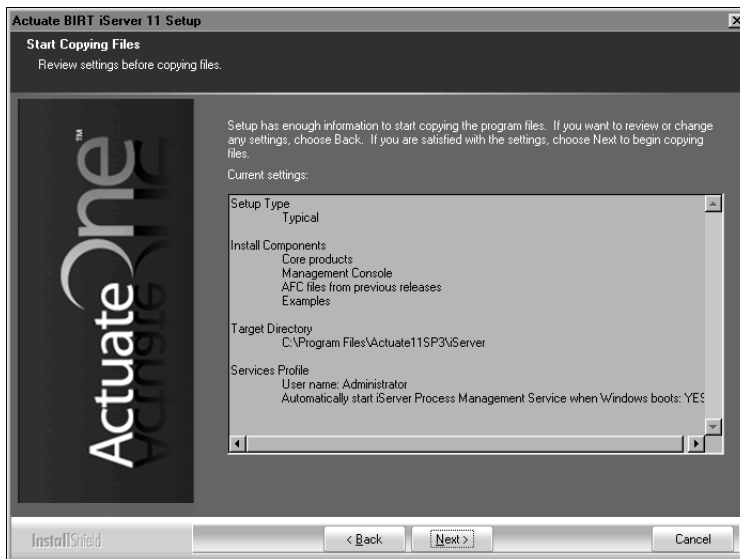


Figure 4-18 Reviewing settings before copying files

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

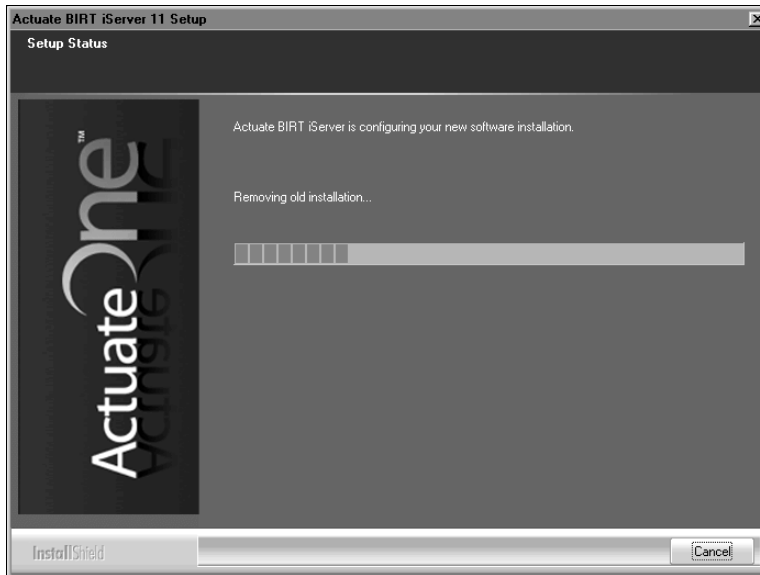


Figure 4-19 Viewing setup status

During installation, a command prompt appears, displaying the run-time commands and messages from the initialization process of the PostgreSQL database system, as shown in Figure 4-20.




Figure 4-20 Initializing the PostgreSQL database system

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.

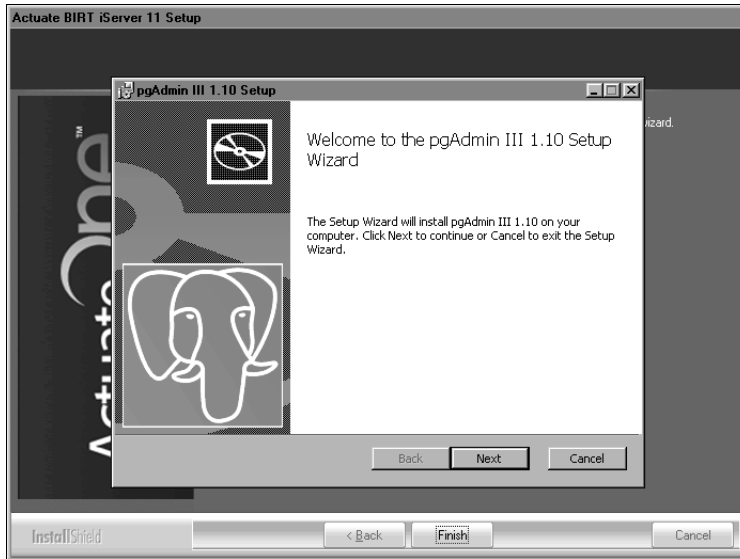


Figure 4-22 Viewing Actuate BIRT iServer and pgAdmin III Setup

15 In pgAdmin III Setup, perform the following tasks:

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




Figure 4-23 Viewing Welcome

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

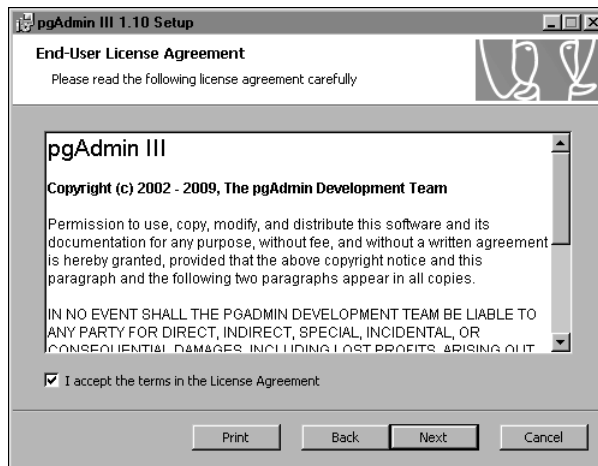


Figure 4-24 Accepting the license agreement

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

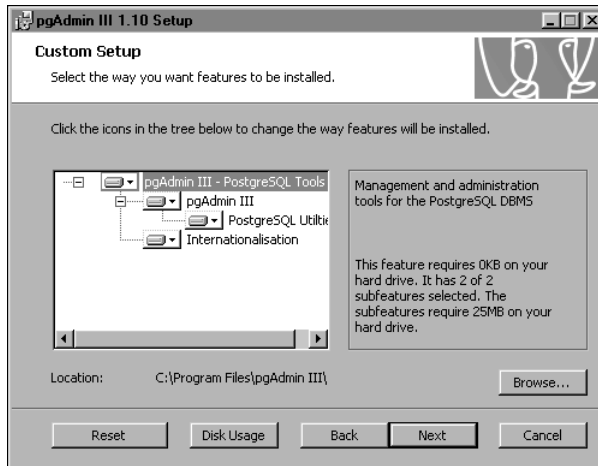


Figure 4-25 Viewing the features to be installed

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




Figure 4-26 Choosing to install pgAdmin III

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

- 16 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

17 The installation program prompts you to install the online help and manuals using the Actuate Localization and Online Documentation media, as shown in Figure 4-30.




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

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.

Configuration Console Log In appears, as shown in Figure 4-32. iServer is ready for use.

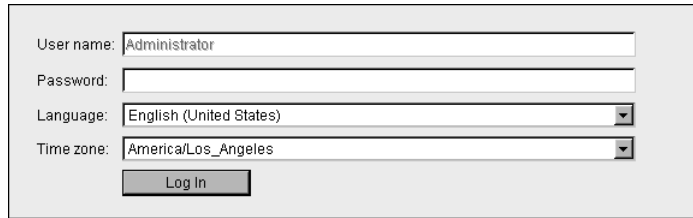
The image shows a login form with the following fields: 'User name:' with the text 'Administrator' entered; 'Password:' with an empty field; 'Language:' with a dropdown menu showing 'English (United States)'; and 'Time zone:' with a dropdown menu showing 'America/Los_Angeles'. Below these fields is a 'Log In' button.

Figure 4-32 Configuration Console Log In appears

About the acserverconfig.xml files saved by the upgrade process

Figure 4-33 shows the date-stamped acserverconfig.xml files saved during an upgrade from iServer Release 10 to Release 11SP3. Figure 4-33 also shows the following changes in the file structure from an upgrade to an earlier iServer Release 11 version:

- The location of acserverconfig.xml is AC_DATA_HOME/config/11SP3 instead of AC_DATA_HOME/config.
- The upgrade to Release 11SP3 installs iServer into a new folder, \Program Files\Actuate11SP3\iServer. This is the Release 11SP3 AC_SERVER_HOME. An upgrade to a previous iServer Release 11 version from an earlier major release installed iServer in the earlier major release AC_SERVER_HOME.

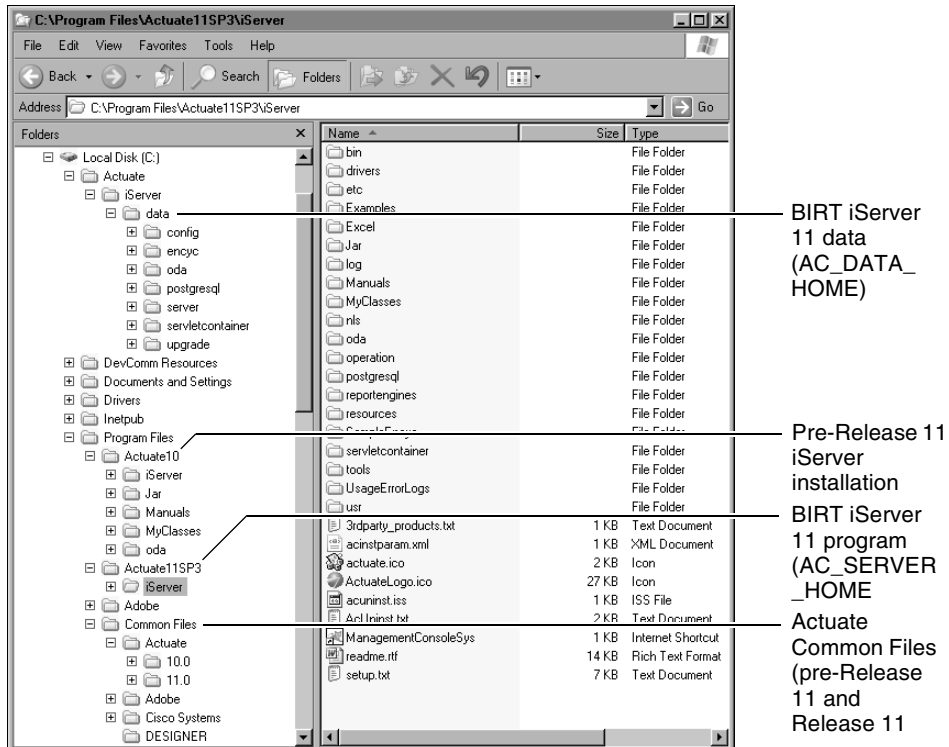


Figure 4-33 Viewing the file system after upgrade to Release 11SP3

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 in-place upgrade from an earlier version of iServer Release 11

When upgrading BIRT iServer from an earlier Release 11 version, you install the new iServer version in the same directory as the earlier version. After completing the install procedure, you run a utility that updates the volume schema, enabling the new iServer version to work with your existing Encyclopedia volume or volumes. You must have write access to the current working directory where you run the utility since the process creates a local temporary file. Then, you bring the volume or volumes online.

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
- 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.

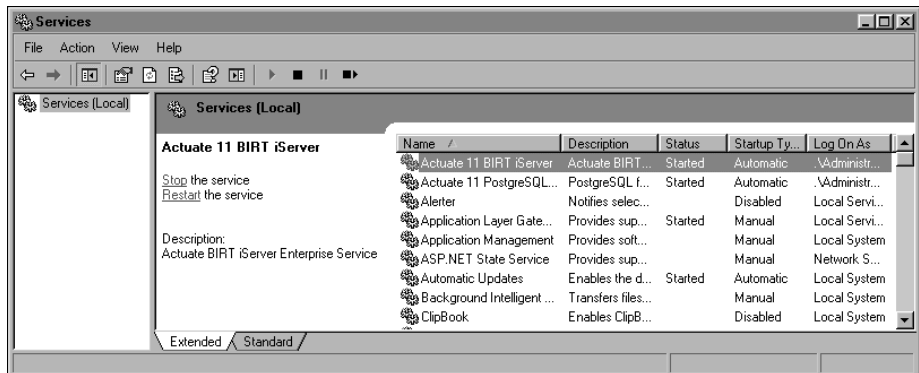


Figure 4-34 Stopping the BIRT iServer service

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

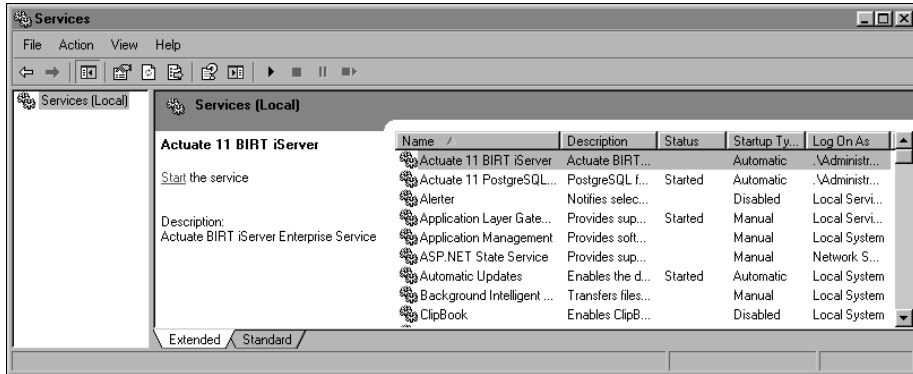


Figure 4-35 Viewing Services after BIRT iServer service stops

How to shut down the Actuate 11 PostgreSQL for BIRT service

- 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.

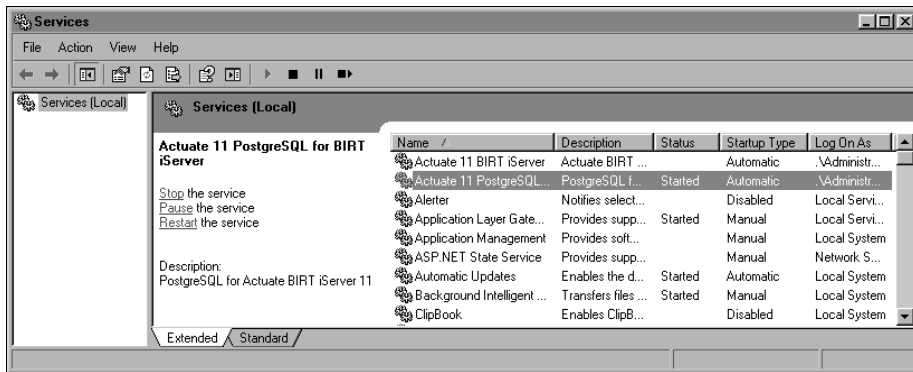


Figure 4-36 Shutting down the PostgreSQL for BIRT iServer service

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

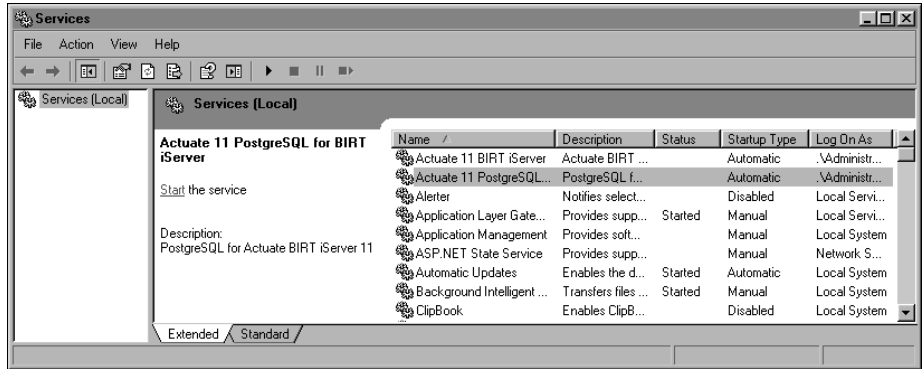


Figure 4-37 Viewing Services after PostgreSQL service stopped

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"
```

How to perform an in-place upgrade from an earlier Release 11 installation

- 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 If you downloaded iServer, run the ActuateBIRTiServer.exe file. If you have a DVD or ftp distribution, run setup.exe in the iServer directory. 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.

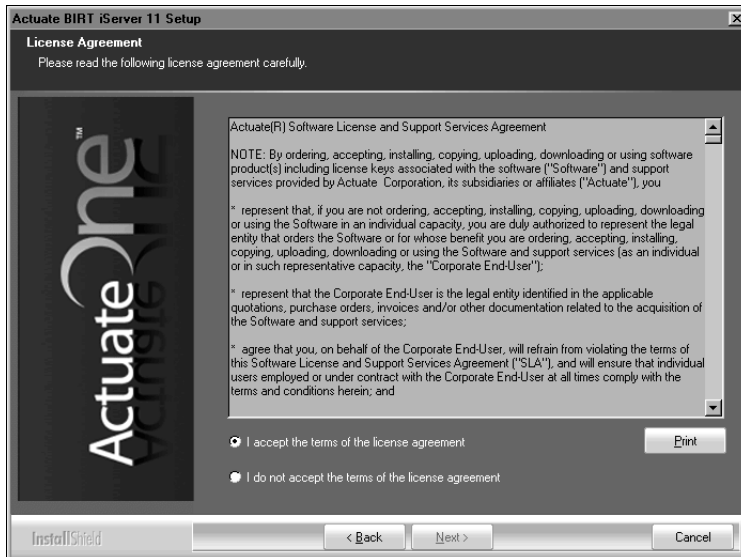


Figure 4-39 Viewing the license agreement

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

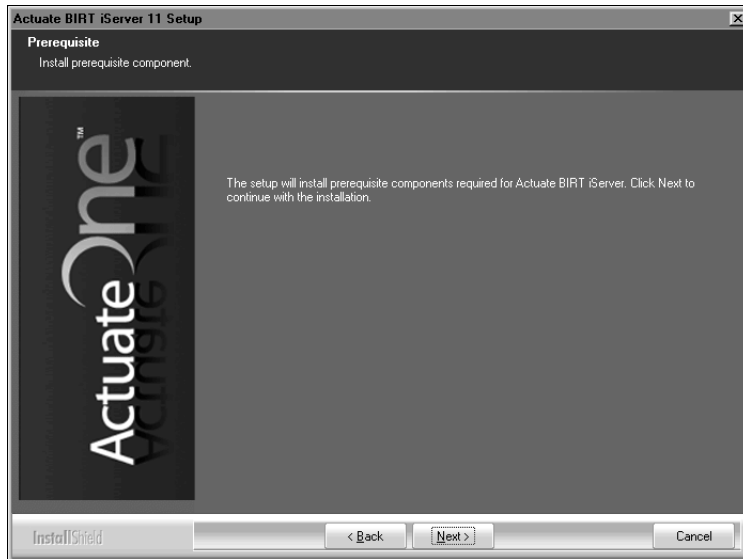


Figure 4-40 Installing Prerequisites

- 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/Actuate11SP3/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.

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

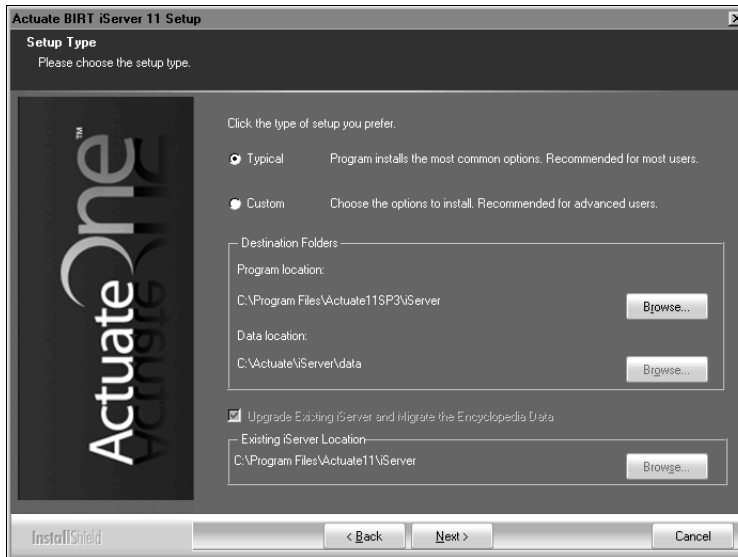


Figure 4-41 Choosing setup type

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



Figure 4-42 Confirming volume migration

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. If you have not also stopped the PostgreSQL for Actuate BIRT iServer 11 service, stop this service too.
- 3 On Setup Type, choose Next.




Figure 4-43 Choosing OK to shut down iServer

- 6 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 . 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

- 7 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.

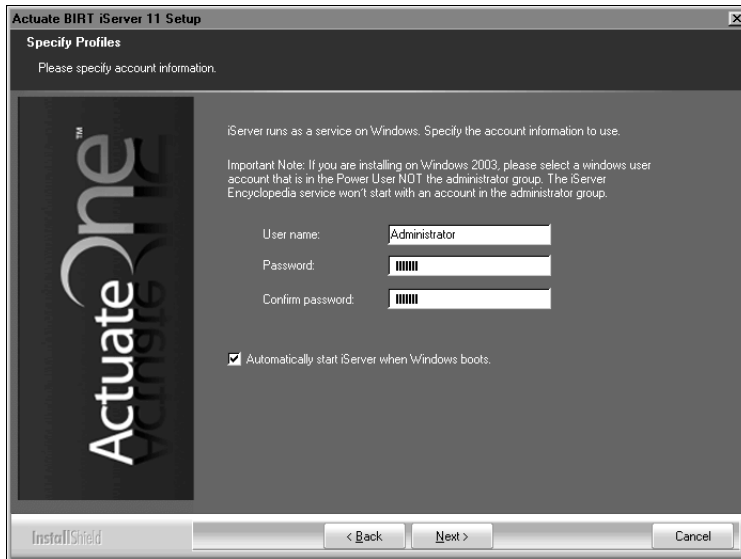


Figure 4-45 Specifying the account for running the iServer service

- 8** 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.

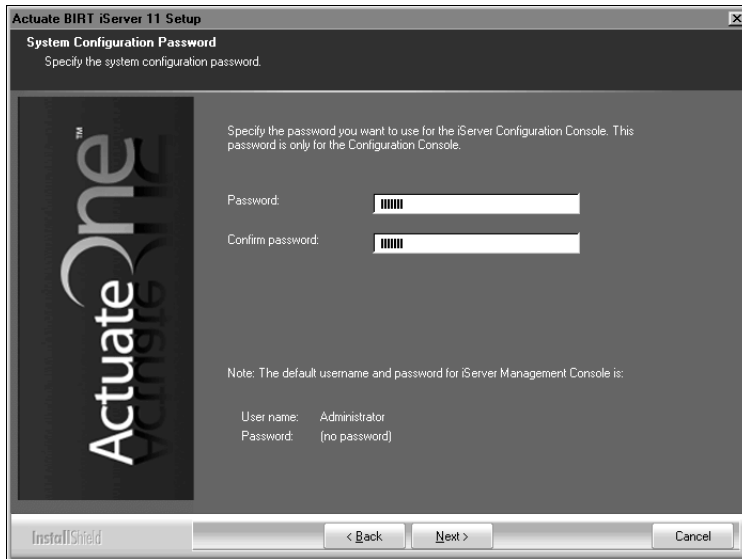


Figure 4-46 Specifying the password for using Configuration Console

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

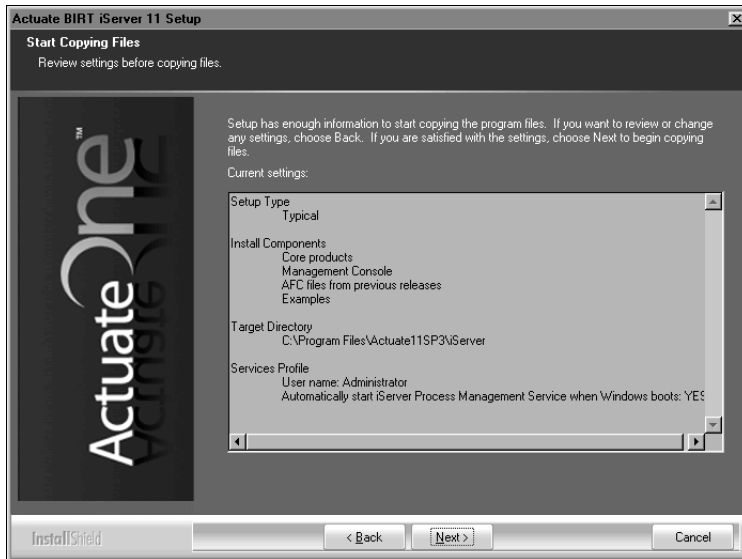


Figure 4-47 Reviewing settings before copying files

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



Figure 4-48 Viewing setup status

A prompt appears, asking whether to install pgAdmin, the administration tool for the PostgreSQL RDBMS, as shown in Figure 4-49. If you do not have pgAdmin installed, choose Yes.




Figure 4-49 Choosing to install the pgAdmin tool

Actuate BIRT iServer Setup appears with pgAdmin III Setup appearing immediately afterward, as shown in Figure 4-50.

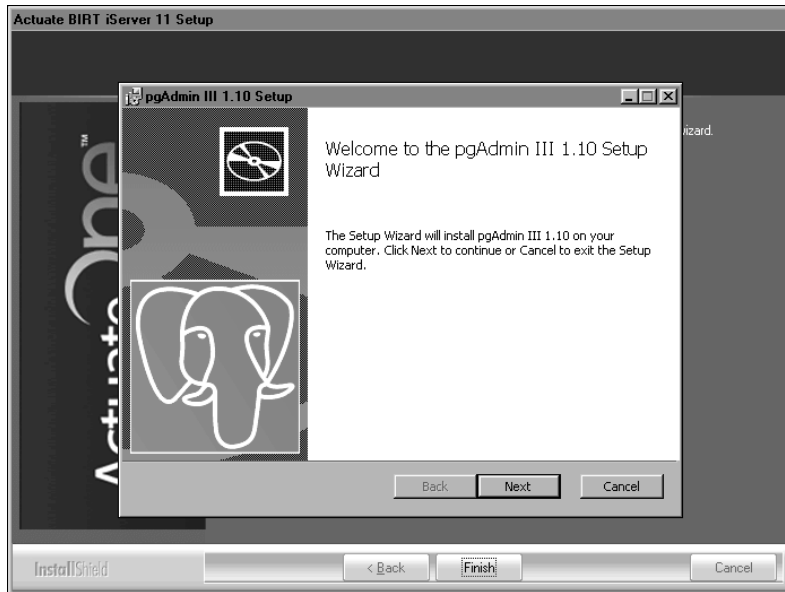


Figure 4-50 Viewing Actuate BIRT iServer and pgAdmin III Setup

10 In pgAdmin III Setup, perform the following tasks:

- 1 In Welcome, shown in Figure 4-51, choose Next.

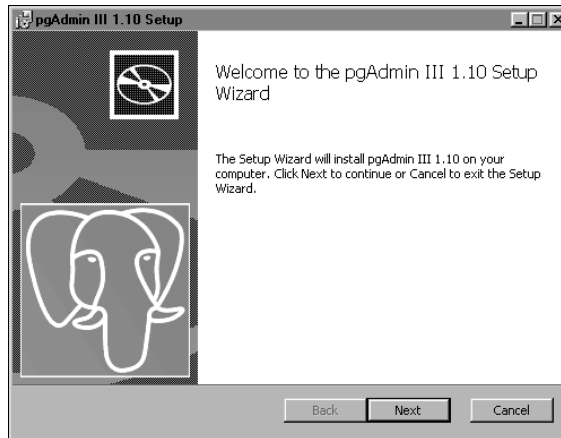


Figure 4-51 Viewing Welcome

- 2 In License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-52. Choose Next.



Figure 4-52 Accepting the license agreement

- 3 In Custom Setup, review the features to be installed, as shown in Figure 4-53. Choose Next.

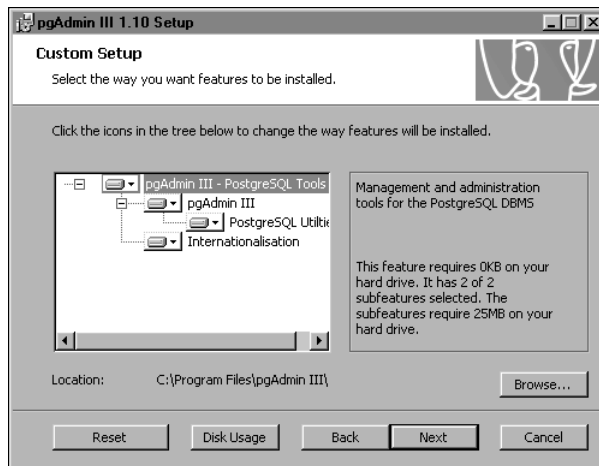


Figure 4-53 Viewing the features to be installed

- 4 In Ready to Install pgAdmin III, shown in Figure 4-54, choose Install.

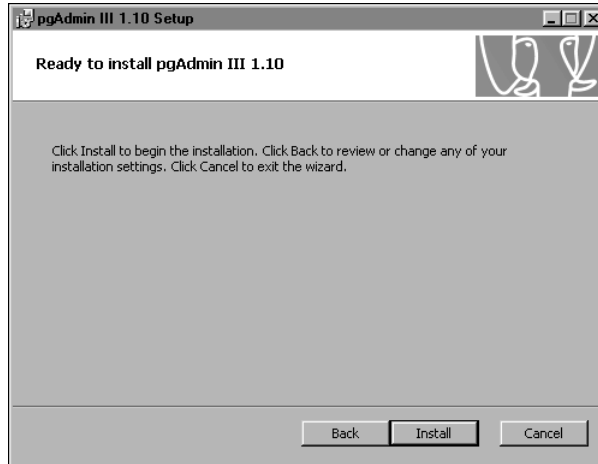


Figure 4-54 Choosing to install pgAdmin II
Installing pgAdmin III appears, as shown in Figure 4-55.




Figure 4-55 Installing pgAdmin III

- 5 When Completed, the pgAdmin III Setup Wizard appears, as shown in Figure 4-56, choose Finish to exit the pgAdmin III Setup Wizard.




Figure 4-56 Choosing Finish

- 11 When the Actuate BIRT iServer 11 setup success message appears, as shown in Figure 4-57, choose Finish to exit this wizard.



Figure 4-57 Exiting the installation wizard

- 12 The installation program prompts you to install the online help and manuals using the Actuate Localization and Online Documentation media, as shown in Figure 4-58.



Figure 4-58 Viewing the install online help and manuals prompt

The installation program installs shortcuts on the desktop, as shown in Figure 4-59.




Figure 4-59 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.

Configuration Console Log In appears, as shown in Figure 4-60. iServer is ready for use.

Figure 4-60 Configuration Console Log In appears

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

This manual option requires the system administrator to export the volume metadata and data from the earlier iServer volume to a temporary workspace, then import these resources into the new Release 11 installation. Actuate provides the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to perform these tasks.

The following procedure describes how to perform a manual side-by-side upgrade of an earlier iServer version to BIRT iServer Release 11.

How to perform a manual side-by-side upgrade to iServer Release 11

- 1 If you downloaded iServer, run the ActuateBIRTiServer.exe file. If you have a DVD or ftp distribution, run setup.exe in the iServer directory. The welcome message appears, as shown in Figure 4-61. Choose Next.



Figure 4-61 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 4-62. Choose Next.

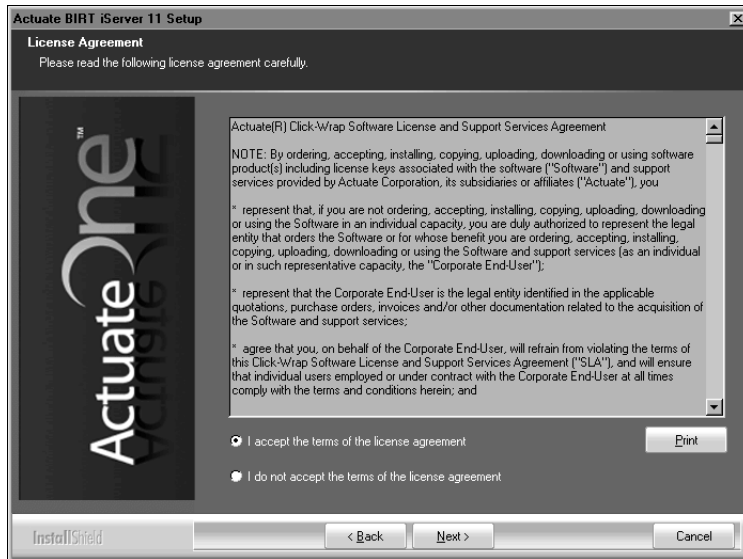


Figure 4-62 Viewing the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-63. Choose Next.

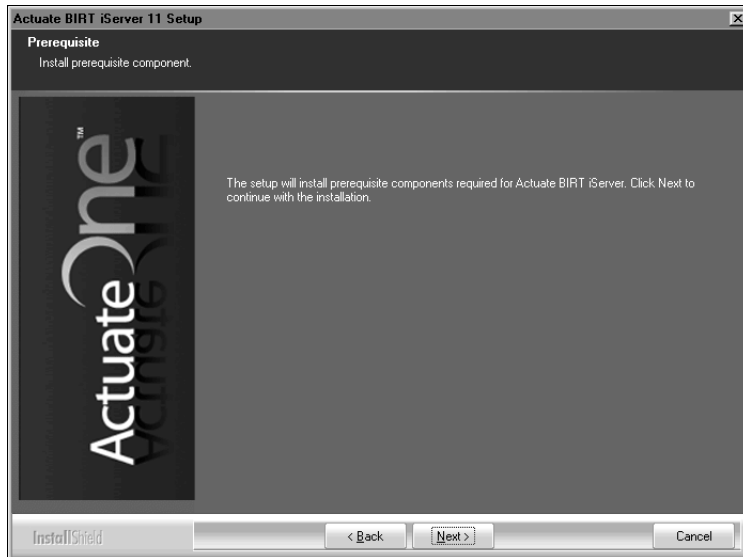


Figure 4-63 Installing Prerequisites

- 4 In Setup Type, select Custom, as shown in Figure 4-64, to install a supported relational database management system (RDBMS) for the Encyclopedia

volume, such as Oracle or PostgreSQL. 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\Actuate11SP3\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.

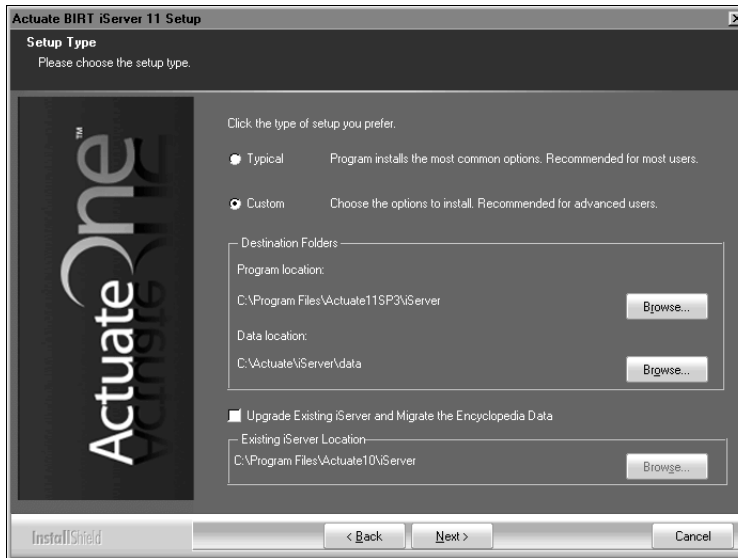


Figure 4-64 Selecting custom setup type

- 5 In Select Features, accept the default, all features selected, as shown in Figure 4-65. The installation program installs all features. Choose Next.

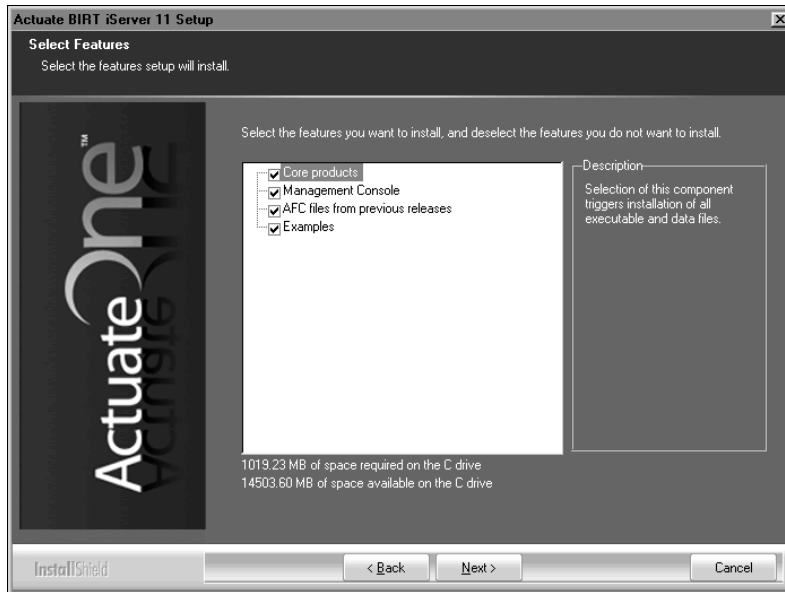


Figure 4-65 Selecting features to install

- 6** In iServer Installation Option, accept the default, Stand-alone, to install iServer as a stand-alone server, as shown in Figure 4-66. Choose Next.
- 7** In Encyclopedia Metadata Storage, select the type of Encyclopedia volume metadata database to install, as shown in Figure 4-67. 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.

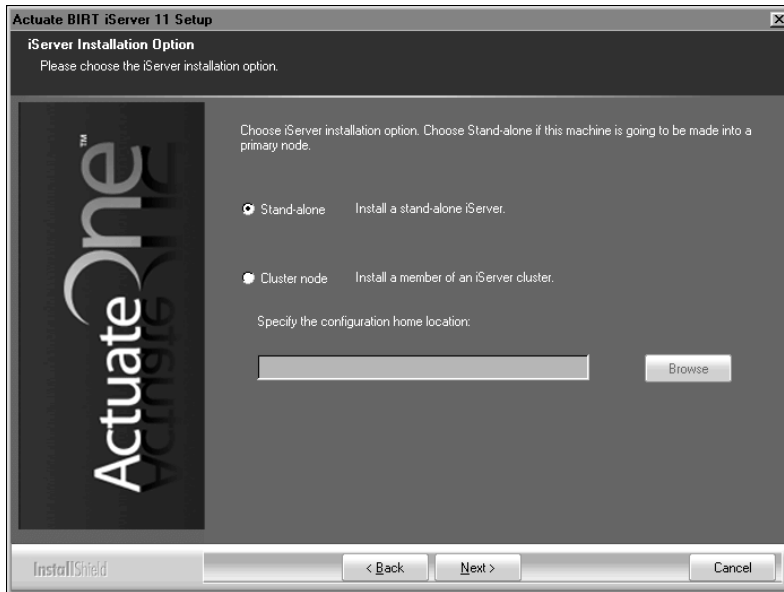


Figure 4-66 Selecting the option to install a Stand-alone iServer

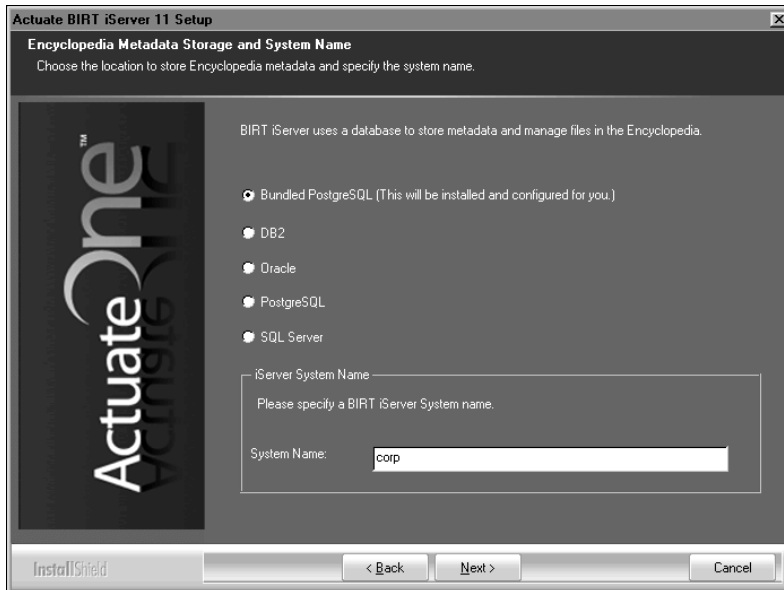


Figure 4-67 Selecting metadata database type

- 8** On PostgreSQL Database Information, in Database Credentials, type and confirm a password, as shown in Figure 4-68. 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.

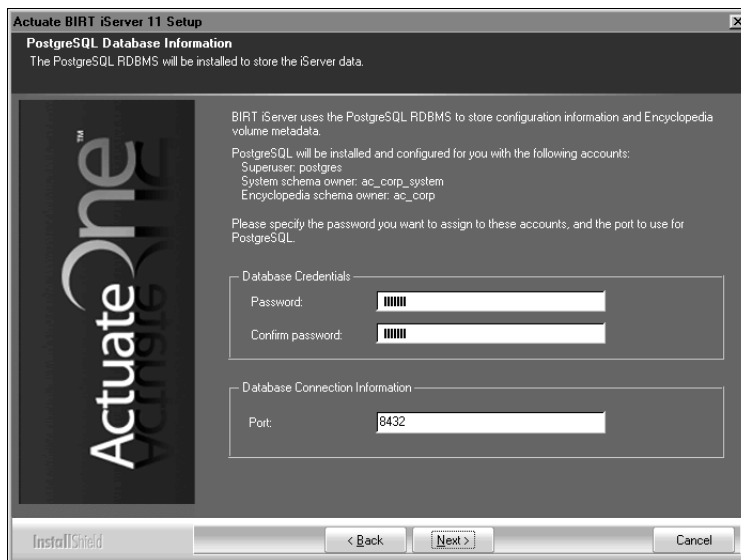


Figure 4-68 Specifying PostgreSQL database information

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 4-69.




Figure 4-69 Setting the Windows local security policy

- 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 4-70.

Alternatively, choose Try out the product using the included evaluation software license if you do not have a purchased license. Choose Next.

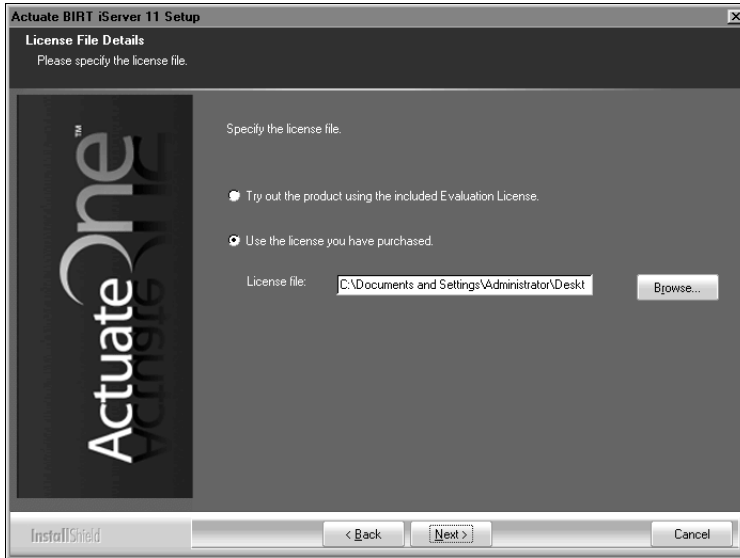


Figure 4-70 Specifying the license file

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 4-71.




Figure 4-71 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 4-72. Alternatively, choose the language and locale settings for your region.

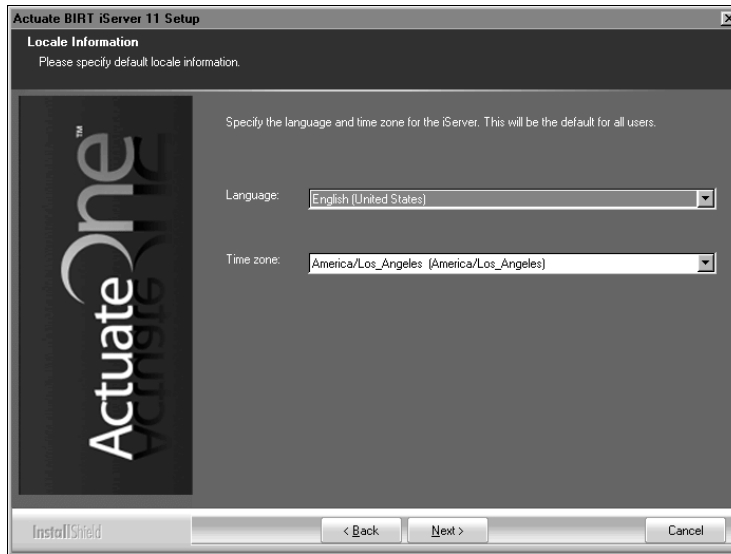


Figure 4-72 Specifying locale information

- 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 4-73. 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 4-73. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

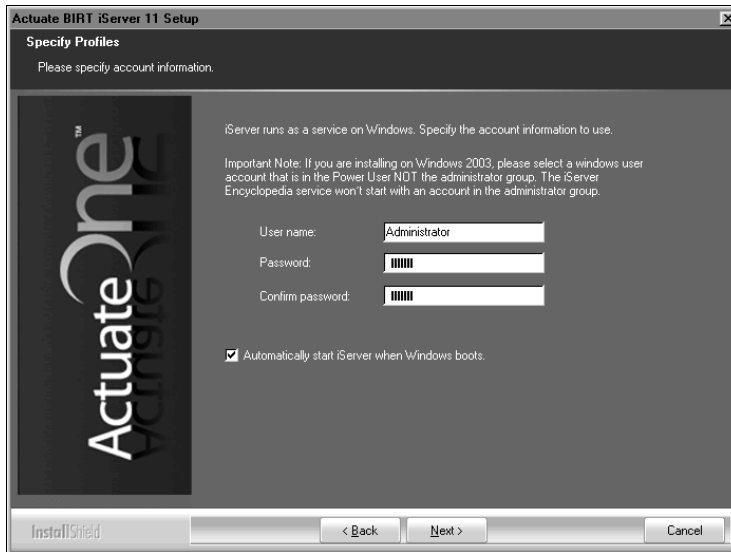


Figure 4-73 Specifying the account for running the iServer service

- 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 4-74. Choose Next.

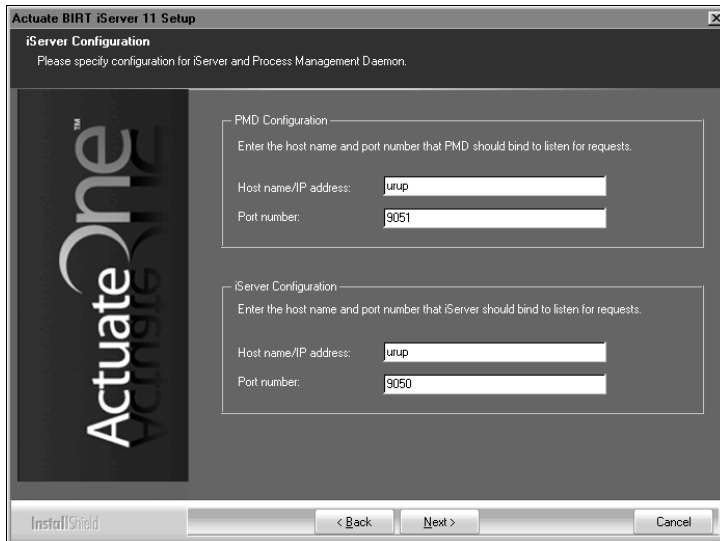


Figure 4-74 Specifying iServer configuration information

- 13** In System Configuration Password, type and confirm a password for Configuration Console, as shown in Figure 4-75. 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.

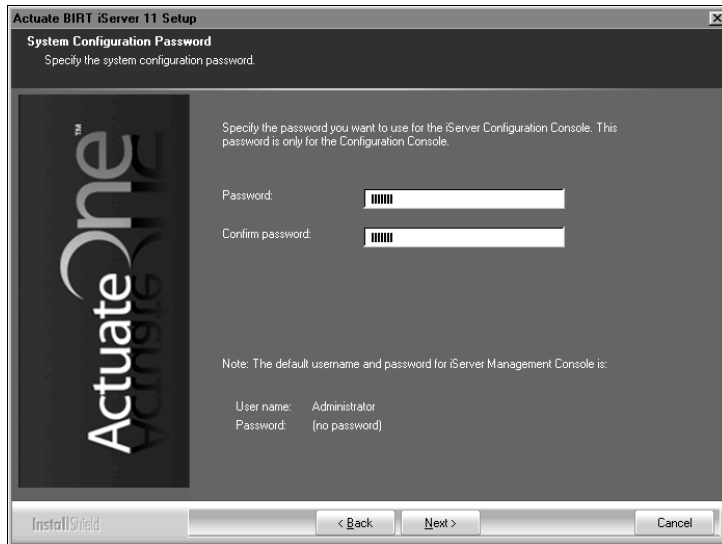


Figure 4-75 Specifying the password for using Configuration Console

- 14** In Volume Information, select Do not use the default volume, as shown in Figure 4-76. The installation process does not create the default Encyclopedia volume. Choose Next.

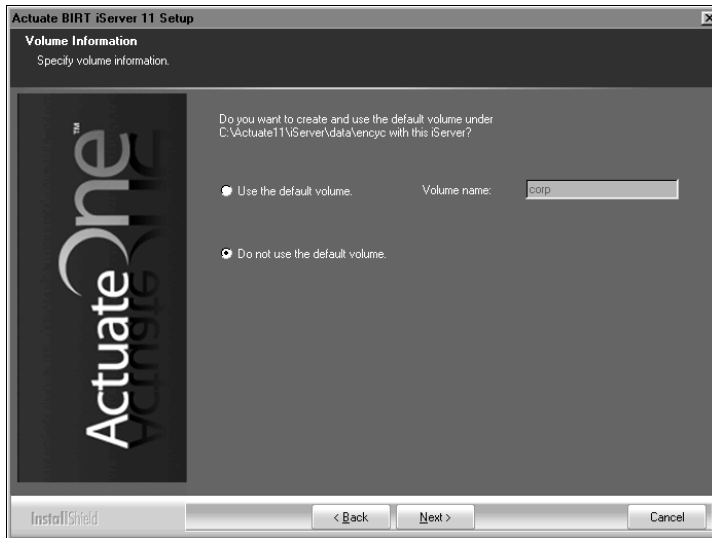


Figure 4-76 Selecting not to install the default volume

- 15** In Management Console Configuration, accept the default Host name and port number values in PMD Configuration and iServer Configuration, as shown in Figure 4-77. Alternatively, specify different values. Choose Next.

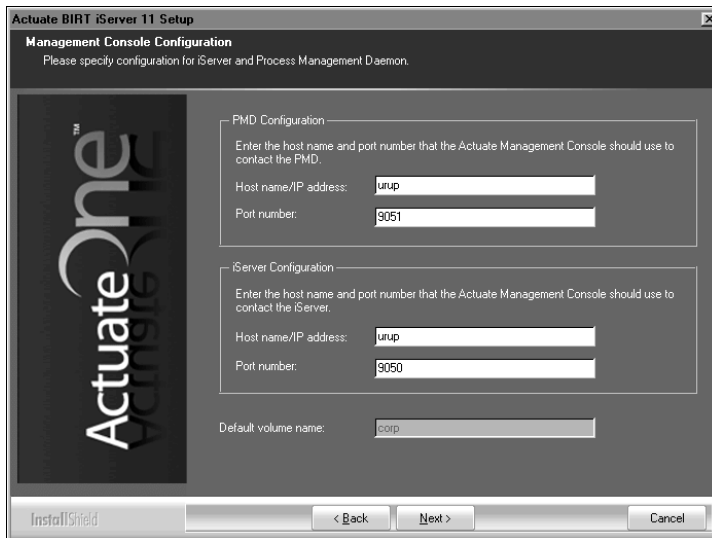


Figure 4-77 Specifying Management Console properties

- 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 4-78. Choose Next.

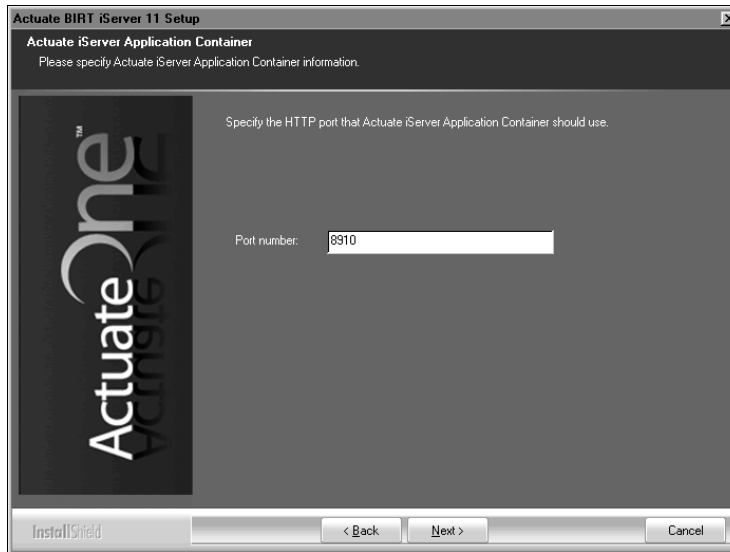


Figure 4-78 Specifying the application container port number

- 17 In Context Path, type the context path for Management Console or accept the default path, /acadmin, as shown in Figure 4-79. Choose Next.



Figure 4-79 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 4-80. Choose Next.



Figure 4-80 Specifying a program folder

19 In Start Copying Files, review the settings shown in Figure 4-81. Choose Next.

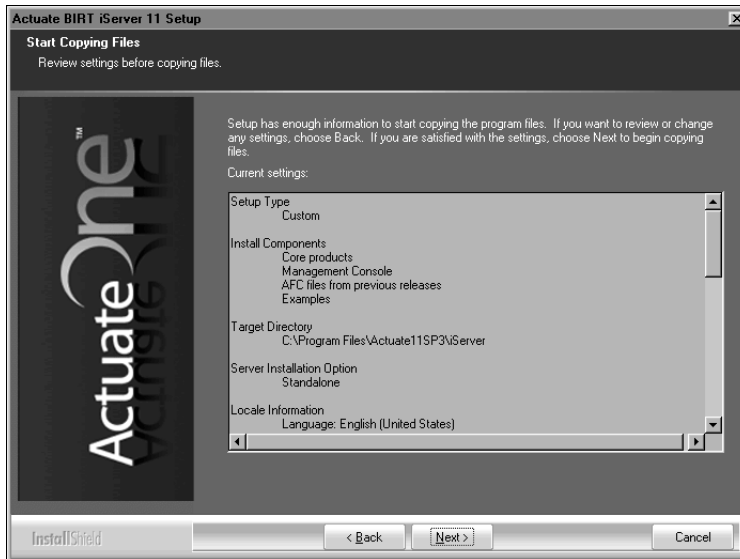


Figure 4-81 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-82.



Figure 4-82 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 4-83.




Figure 4-83 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 4-84. Choose Yes.




Figure 4-84 Choosing to install the pgAdmin tool

Activate BIRT iServer Setup appears, then pgAdmin III Setup appears, as shown in Figure 4-85.



Figure 4-85 Viewing Activate BIRT iServer and pgAdmin III Setup

20 In pgAdmin III Setup, perform the following tasks:

- 1 In Welcome, shown in Figure 4-86, choose Next.

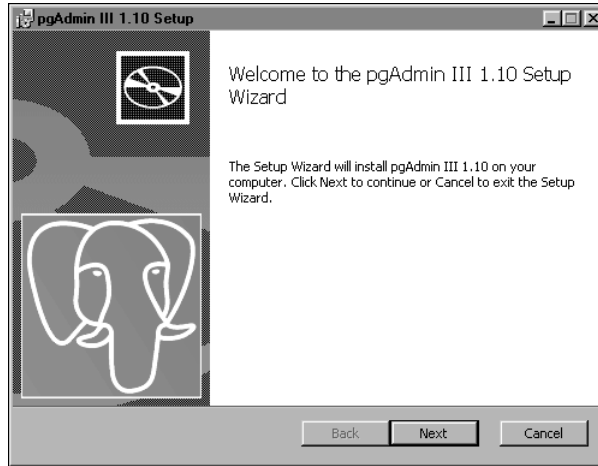


Figure 4-86 Viewing Welcome

- 2 In End-User License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-87. Choose Next.

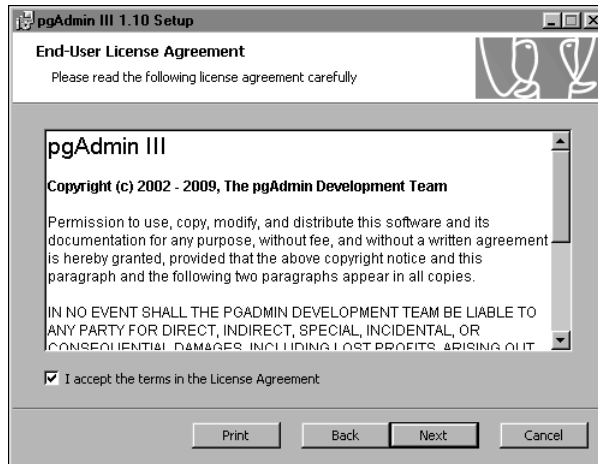


Figure 4-87 Accepting the license agreement

- 3 In Custom Setup, review the features to be installed, as shown in Figure 4-88. Choose Next.




Figure 4-88 Viewing the features to be installed

- 4 In Ready to Install pgAdmin III, shown in Figure 4-89, choose Install.

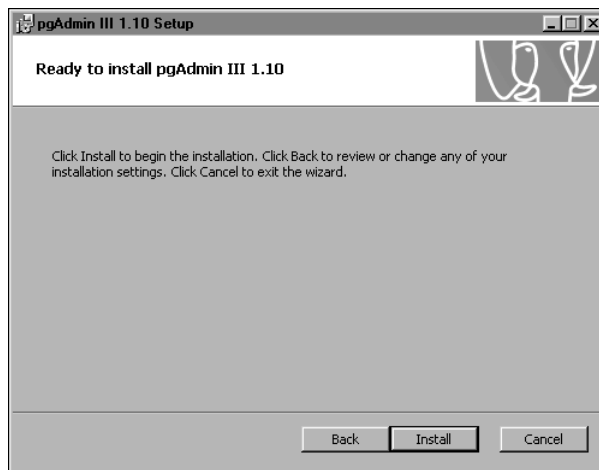


Figure 4-89 Choosing to install pgAdmin II

Installing pgAdmin III appears, as shown in Figure 4-90.




Figure 4-90 Installing pgAdmin III

- 5 When Completed the pgAdmin III Setup Wizard appears, as shown in Figure 4-91, choose Finish to exit the wizard.




Figure 4-91 Choosing Finish

- 21 Choose Finish to exit the wizard, as shown in Figure 4-92.



Figure 4-92 Exiting the installation wizard

- 22** The installation program prompts you to install the online help from the following location:

<http://www.actuate.com>

Alternatively, you can install the online help and manuals in PDF format using the Actuate Localization and Online Documentation media. Figure 4-93 shows the prompt.




Figure 4-93 Viewing the install online help and manuals prompt

The installation program installs shortcuts on the desktop, as shown in Figure 4-94.




Figure 4-94 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

The following section describes how to perform a manual migration of a single Encyclopedia volume in a side-by-side installation. In a side-by-side installation, the administrator installs a new BIRT iServer Release 11 in a path separate from an earlier major release.

The administrator manually copies the contents of the old encyc directory, containing the Encyclopedia volume data, to a new BIRT iServer Release 11 Encyclopedia volume folder. Next, the administrator migrates the Encyclopedia volume schema and metadata from the earlier release to the new BIRT iServer Release 11 in two steps, using the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities. The administrator then creates a new volume in Configuration Console for the migrated volume.

Alternatively, iServer supports a manual upgrade to an Oracle RDBMS for storing Encyclopedia volume metadata. For more information on installing an iServer that uses an Oracle RDBMS, see “Installing an Encyclopedia volume that uses an alternative database,” in Chapter 3, “Installing BIRT iServer using an alternative database.”

The following procedure describes how to migrate an Encyclopedia volume to BIRT iServer Release 11 in a side-by-side installation.

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 iServer release to BIRT iServer Release 11, you must first export this release to a staging area using the Squirrel Data Exporter utility. To export the earlier release, perform the following tasks:

- 1 When the install program completes, the Actuate iServer System Configuration Console Log in appears. The address bar contains the following web address:

```
http://localhost:8910/acadmin/config
```

8910 is the port number that the installer specified for the iServer Application Container to use if you installed with the previous release running in the environment, as shown in Figure 4-78. 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 into Configuration Console as Administrator.
- 3 When you migrate a volume from an earlier release to BIRT iServer Release 11 the install program creates a default partition. You must delete it. If you have already performed this step, skip to Step 4. Otherwise, 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 4-95.

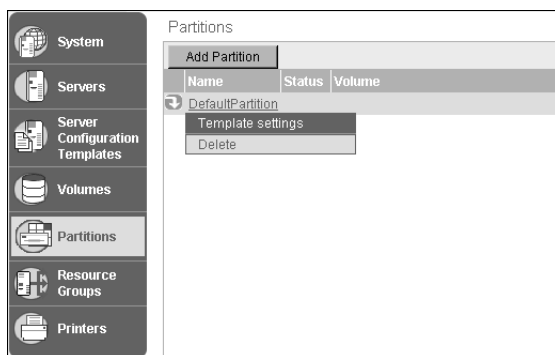


Figure 4-95 Deleting DefaultPartition

- 4 Create a partition on BIRT iServer Release 11 for the volume that you are going to migrate from the previous iServer release. 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 BIRT iServer Release 11 installation.
 - 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 4-96. This path does not need to match the partition path to the same volume on the earlier iServer release. Choose OK.



Figure 4-96 Creating a partition for a migrated volume

- 5 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 10SP1, the location of these folders is AC_SERVER_HOME/encyc. The default AC_SERVER_HOME path in release 10SP1 is /Program Files/Actuate10/iServer.
- 6 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 4-97.

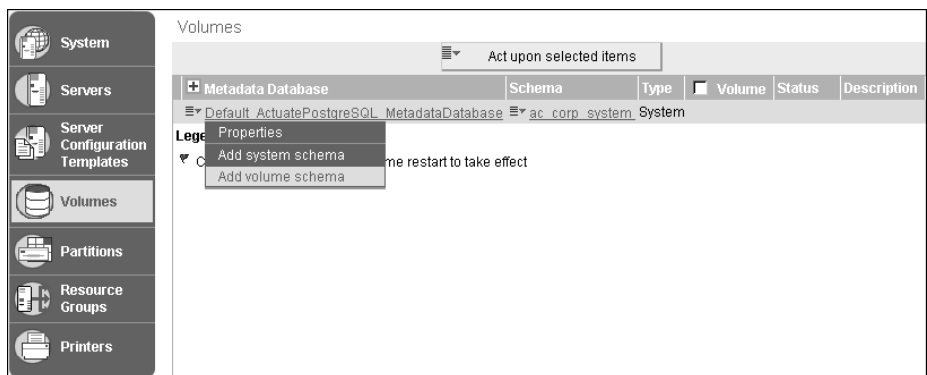


Figure 4-97 Choosing to add a new schema

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 4-68. Choose OK.

Volumes—New Schema appears as shown in Figure 4-98.

Figure 4-98 Creating a new schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-99.

Metadata Database	Schema	Type	Volume	Status	Description
Default_ActuatePostgreSQL_MetadataDatabase	ac_corp	Volume			
	ac_corp_system	System			

Legend
 Changes pending require volume restart to take effect

Figure 4-99 Viewing the new schema

- 7 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 4-100.




Figure 4-100 iServer is offline

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

```
<AC_SERVER_HOME>\bin
```

In this instance, AC_SERVER_HOME refers to the release 11SP3 AC_SERVER_HOME. For example, using the default value for AC_SERVER_HOME, add:

```
C:\Program Files\Actuate11SP3\iServer\bin
```

- 9 Run the Squirrel Data Exporter by performing the following tasks:
 - 1 In Windows Explorer, navigate to release 11SP3 AC_SERVER_HOME\bin.
 - 2 Create a properties file named SquirrelDataExporter.properties. You pass this file to the Squirrel Data Exporter utility when you execute it. Use the following example to create SquirrelDataExporter.properties:

```
AC_SERVER_HOME = C:/Program Files/Actuate11SP3/iServer
SQUIRREL_DATA_HOME =
    C:/Program Files/Actuate10/iServer/ac_corp_partition
SQUIRREL_EXPORT_FOLDER = C:/Actuate11/SquirrelData/ac_corp
SQUIRREL_EXPORT_FORMAT = PostgreSQL
NEW_SCHEMA_NAME = ac_corp
NEW_VOLUME_NAME = corp
```

See Table 4-1 for descriptions of the required Squirrel Data Exporter properties. See Table 4-2 for descriptions of the optional Squirrel Data Exporter properties. See “Configuring Squirrel Data Exporter and System or Encyclopedia Data Store Administrator properties files” for additional notes on property files.

- 10 Open a command prompt and navigate to release 11SP3 AC_SERVER_HOME /bin.

11 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 4-1.

- 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 4-1 `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 is a line-continuation character.

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

To import an Encyclopedia volume from an earlier iServer release to BIRT iServer Release 11 from the staging area, you use the Encyclopedia Data Store Administrator utility. To import the volume, perform the following tasks:

- 1 Run the Encyclopedia Data Store Administrator by performing the following tasks:

- 1 In Release 11 AC_SERVER_HOME\bin, create a properties file named VolumeImport.properties. You pass this file to the Encyclopedia Data Store Administrator utility when you execute it. Use the following example to create VolumeImport.properties:

```
AC_SERVER_HOME = C:/Program Files/Actuate11SP3/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:/Actuate11/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. See “Configuring Squirrel Data Exporter and System or Encyclopedia Data Store Administrator properties files” for additional notes on property files.

- 2 Open a command prompt and navigate to release 11SP3 AC_SERVER_HOME/bin.
- 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 4-2.

- Checks to see if the administrator submitted a property file on the command line when running the script
- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax
- Calls the set_tools_environment.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 4-2 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 batch file, the ^ symbol is a line-continuation character.

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 4-101.




Figure 4-101 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 4-102.

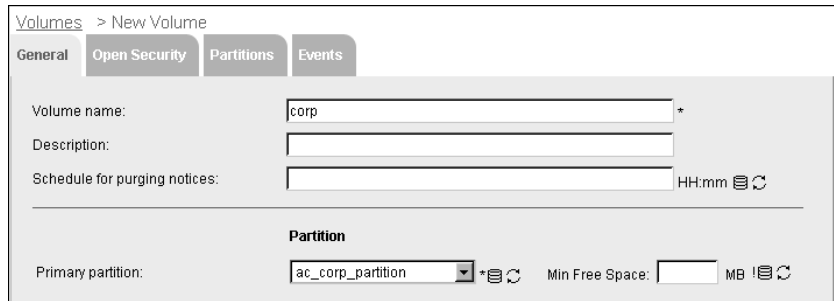


Figure 4-102 Specifying volume and partition name

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 4-103.

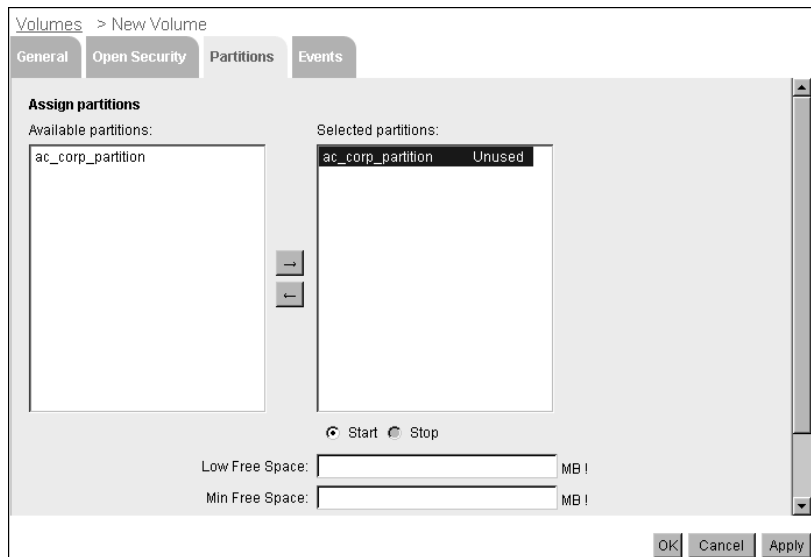


Figure 4-103 Assigning a partition

Choose OK.

- 3 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-104.

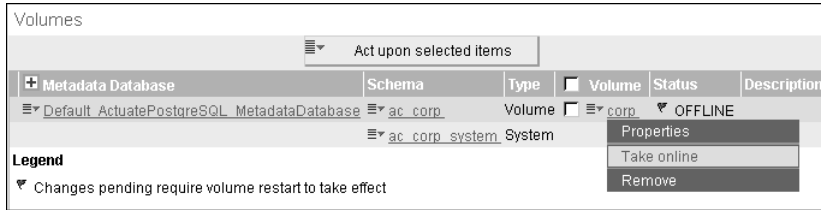


Figure 4-104 Viewing the new volume

The new volume comes online, as shown in Figure 4-105.




Figure 4-105 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.

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, which you specify during the BIRT iServer Release 11 installation, as shown in Figure 4-64.
NEW_SCHEMA_NAME	The 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 <code>[a-z][a-z 0-9]*</code> . Do not use a hyphen.

Table 4-1 Required Squirrel Data Exporter properties

Parameter	Description
NEW_VOLUME_NAME	The name of the migrated volume on BIRT iServer Release 11. This name does not need to match the name this volume had on your earlier iServer release.
SQUIRREL_DATA_HOME	Absolute path to the folder on your earlier iServer release containing the Encyclopedia volume you want to migrate.
SQUIRREL_EXPORT_FOLDER	Absolute path to the folder that Squirrel Data Exporter creates, containing the volume data from SQUIRREL_DATA_HOME, which Encyclopedia Data Store Administrator loads into the new volume on BIRT iServer Release 11.
SQUIRREL_EXPORT_FORMAT	Format of exported data. Specify DB2, Oracle, PostgreSQL, or SQLServer.

Table 4-2 describes the optional Squirrel Data Exporter properties to specify in a SquirrelDataExporter.properties file for performing logging.

Table 4-2 Optional Squirrel Data Exporter properties

Property	Value	Default Value
FILE_LOG_LEVEL	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
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
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.

Specifying Encyclopedia Data Store Administrator properties

Table 4-3 describes the required Encyclopedia Data Store Administrator properties used to configure the environment for a manual migration operation.

Table 4-3 Required Encyclopedia Data Store Administrator properties

Parameter	Description
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 4-64.
DATABASE_TYPE	Type of supported RDBMS that contains the data store: PostgreSQL or Oracle.
LOG_FOLDER	Absolute path to the log folder.
SCHEMA_FILE_NAME	Base name of the file without the extension that contains the meta-schema definition.
SCRIPT_HOME	Absolute path to the root of the folder hierarchy that contains the scripts and the meta-schema definition.

Table 4-4 describes optional Encyclopedia Data Store Administrator properties. The properties used depend on the type of operation performed and the installation environment.

Table 4-4 Optional Encyclopedia Data Store Administrator properties

Parameter	Description	Default Value	Supported Databases
APPLICATION_USER	Application user ID for connecting to the database for normal operations. Must be a legal SQL identifier. Typically iserver. This parameter is required when creating or populating a schema. Restrict the iServer application user name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.		All

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

Parameter	Description	Default Value	Supported Databases
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. Required when USE_SERVER_CONFIG_FILE is true. The schema name can be different from the database 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 Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL
DATA_EXPORT_FOLDER	Absolute path to the folder to which data is exported. This parameter is required if EXPORT_DATA is true. Folder is on the database machine.		All
DATA_EXPORT_FORMAT	Format of exported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE_TYPE}	All
DATA_IMPORT_FOLDER	Absolute path to the folder from which data is imported. This parameter is required if IMPORT_DATA is true. For PostgreSQL, the data must be on the database server.		PostgreSQL
DATA_IMPORT_FORMAT	Format of imported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE_TYPE}	PostgreSQL

(continues)

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

Parameter	Description	Default Value	Supported Databases
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	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
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports PostgreSQL and Oracle.		All
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
<i>(continues)</i>			
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 data from the data store.	False	All

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

Parameter	Description	Default Value	Supported Databases
EXPORT_DATA	Set to true to export filtered data from the data store. This parameter only works when a filter value is available.	False	All
FILE_LOG_LEVEL	The 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
IMPORT_DATA	Set to true to import data into the data store.	False	All
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	The name of the log file. Do not add a file extension. The extension is set to .log. A unique number is appended automatically to the file name to prevent overwriting previous logs.	Encyclopedia DataStore Administrator <number>.log	
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Data Store Administrator tools do not generate large logs.	10,000,000	All
LOG_FOLDER	The full path of folder to write logs.	AC_DATA _HOME/ server/log.	

(continues)

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

Parameter	Description	Default Value	Supported Databases
NEW_SCHEMA_NAME	Name of the new schema to be created, or the name of the target schema for data export. This name is required if {SCHEMA_NAME} is not present. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	{NEW_SCHEMA_NAME}	All
NEW_SCHEMA_PASSWORD	Password of the database superuser. This parameter is required if {NEW_SCHEMA_NAME} is present.	{NEW_SCHEMA_PASSWORD}	All
NEW_VOLUME_NAME	Name of the new volume to be created, or the name of the target schema for data export. Do not use a hyphen in a volume name.	{NEW_VOLUME_NAME}	All
ORACLE_TNS_NAMES_FILE	Absolute path to the Oracle TNS names file. This path can be used instead of {DATABASE_HOST} and {DATABASE_PORT} to generate a JDBC URL.		Oracle
POPULATE_SCHEMA	Set to true to populate the schema with data store elements such as tables, views, stored functions, and procedures. Set to false if the data store elements already exist. Use Configuration Console to populate schema when possible. Schema automatically populates when defined in Configuration Console.	{CREATE_SCHEMA}	All
POPULATE_SCHEMA_PHASE_ONE	Set to true to create the schema objects necessary for data to be loaded. Set to false to prevent this schema object creation. Imports data without building indexes for fast load.	{POPULATE_SCHEMA}	All

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

Parameter	Description	Default Value	Supported Databases
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
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	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema.	{SCHEMA_NAME}	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
SYSTEM_DATABASE_NAME	Name of the system database.		PostgreSQL
TABLESPACE_NAME	Name of the application tablespace.		PostgreSQL
TABLESPACE_LOCATION	Absolute path to the application tablespace folder.		PostgreSQL

(continues)

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

Parameter	Description	Default Value	Supported Databases
TIME_ZONE	Local time zone for installation.		All
USE_SERVER_CONFIG_FILE	Reads connection information from acserverconfig.xml. Password is not picked up from acserverconfig.xml. Default value is false.	False	All
VOLUME_NAME	Name of the volume or target schema for data import.	{VOLUME_NAME}	All

Specifying System Data Store Administrator properties

The system data store is a separate schema in the database that stores iServer metadata. The system schema is a required element for any iServer installation. In a cluster, the nodes share the system schema metadata and use this information to communicate and coordinate processing.

In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this procedure to protect critical system metadata. If a system schema becomes corrupted or is accidentally deleted, the administrator can use the System Data Store Administrator utility to create a new schema. You can define a new system schema in Configuration Console or edit acserverconfig.xml to refer to the current system schema.

Use the System Data Store Administrator utility to perform the following tasks:

- Create or delete a schema
- Populate or depopulate a schema
- Import or export data

Before running the System Data Store Administrator utility, add the following string to the PATH environment variable on your machine:

```
<AC_SERVER_HOME>\bin
```

In this instance, AC_SERVER_HOME refers to the release 11SP3 AC_SERVER_HOME. For example, using the default value for AC_SERVER_HOME, add:

```
C:\Program Files\Actuate11SP3\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 “Creating and Populating a System Schema,” 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 upgrdedbs.bat script, using the following command-line syntax:

```
adminsds systemdatastore.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 System Data Store Administrator utility using the name of the properties file as an argument

Listing 4-3 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 ^ character 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 specify the RDBMS type and JDBC connection details, such as the schema, application user, and superuser names and passwords
- Schema operation properties that specify an action to perform, such as create, populate, or delete a schema
- Logging properties that control messages sent to the console and log files
- Engineering properties used by Actuate Support or Professional Services to assist with diagnosing or resolving specific issues

Refer to the required and optional properties tables in “Specifying Encyclopedia Data Store Administrator properties,” earlier in this chapter for detailed information about these properties.

Performing operations using System Data Store Administrator utility

The System Data Store Administrator utility supports a 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.

Creating and Populating a System Schema

This operation is only supported for PostgreSQL. Creating and populating a schema requires superuser privileges. 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/Actuate11SP3/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SYSTEM_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>
NEW_SCHEMA_PASSWORD = <provide a password>
INITIALIZE_DATA = true
```

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/Actuate11SP3/iServer
DATABASE_TYPE = PostgreSQL
SYSTEM_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>
NEW_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/Actuate11SP3/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. 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/Actuate11SP3/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. 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/Actuate11SP3/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/Actuate11SP3/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/Actuate11SP3/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/Actuate11SP3/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
```

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/Actuate11SP3/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/Actuate11SP3/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SYSTEM_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
NEW_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
```

Configuring Squirrel Data Exporter and System or Encyclopedia Data Store Administrator properties files

When configuring Squirrel Data Exporter and System or Encyclopedia Data Store Administrator 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.

Installing a BIRT iServer cluster

This chapter discusses the following topics:

- Installing a BIRT iServer cluster node
- Performing a cluster node installation using the wizard
- Adding a node to a cluster

Installing a BIRT iServer cluster node

A node is a machine running an BIRT iServer instance. An iServer administrator adds a node to an iServer cluster to 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 them to the process environment variable list. When a node joins a cluster, it configures itself using its template.

After deployment and configuring the appropriate environment variables in `acpmdconfig.xml`, the administrator launches a deployed iServer image from the command line by passing the necessary arguments or creates a script that executes the commands. 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*.

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 cluster node installation using the wizard

When you create a BIRT iServer cluster, you 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 If you downloaded iServer, run `ActuateBIRTiServerEnterprise.exe`. If you have a DVD or ftp distribution, run `setup.exe`. The welcome message appears, as shown in Figure 5-1. Choose Next.



Figure 5-1 Viewing the welcome message

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

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

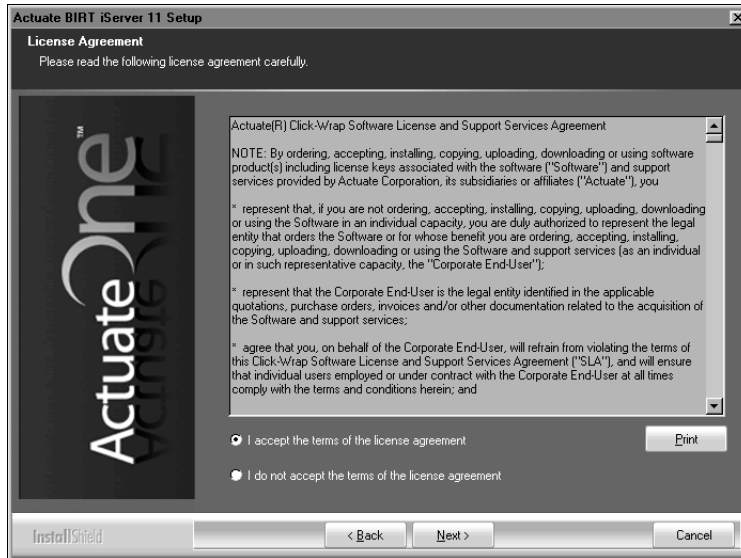


Figure 5-2 Accepting the license agreement

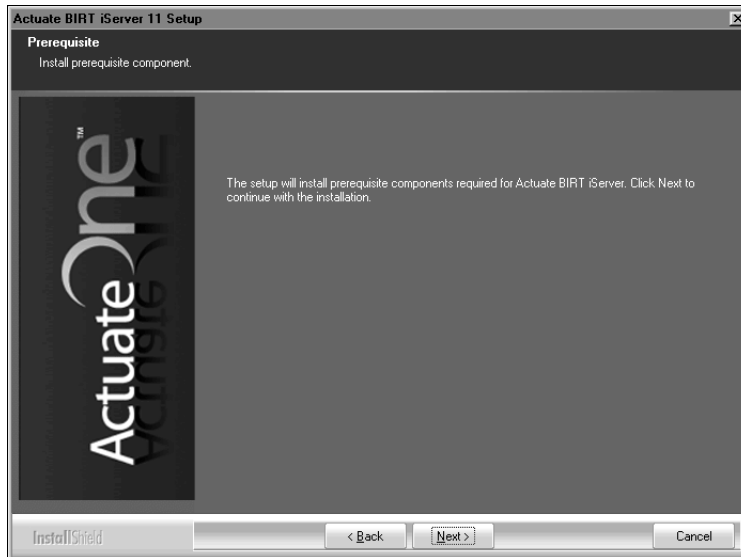


Figure 5-3 Installing Prerequisites

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

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/Actuate11SP3/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 to have its own `AC_SERVER_HOME` and `AC_DATA_HOME` folders. These folders cannot be shared by other nodes in a cluster.

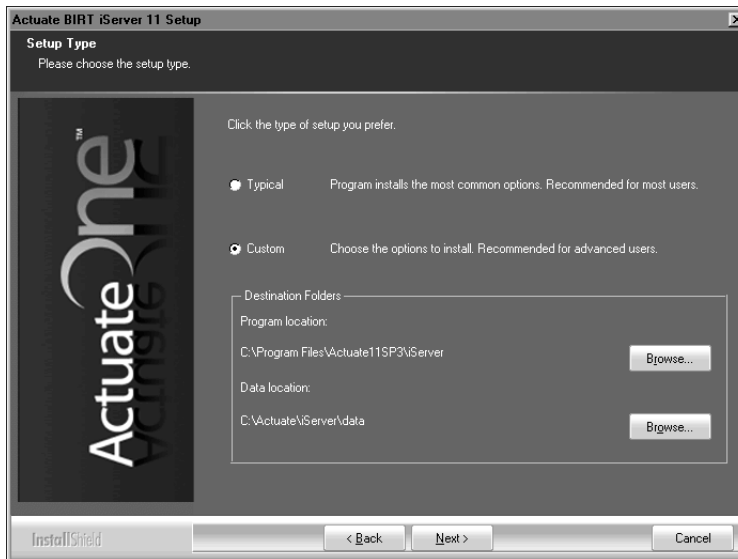


Figure 5-4 Selecting the Custom installation

- 5** In Select Features, select all features—Core products, Management Console, AFC files from previous releases, and Examples, as shown in Figure 5-5. 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 5-6.

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

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

In a Release 11 Service Pack 3 installation, the configuration files are located in AC_DATA_HOME\config\11SP3 by default. If the shared configuration folder is in this default location on a server named urup, then the administrator should specify \\urup\11SP3 as the path.

Note that in a cluster setup, the administrator must edit <ServerFileSystemSetting> in the shared acserverconfig.xml file to point to the shared drive location that contains the Encyclopedia data files. For more information about configuring an iServer node to run in a cluster, see “Adding a node to a cluster,” later in this chapter.

Choose Next.

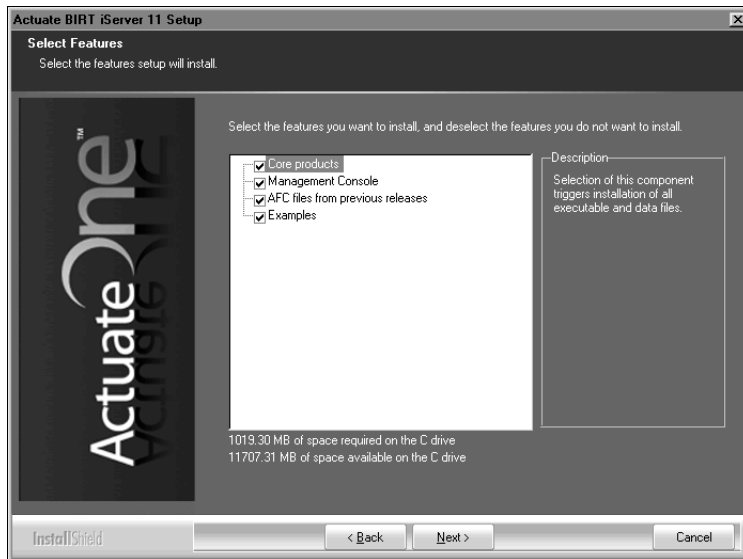


Figure 5-5 Selecting features to install

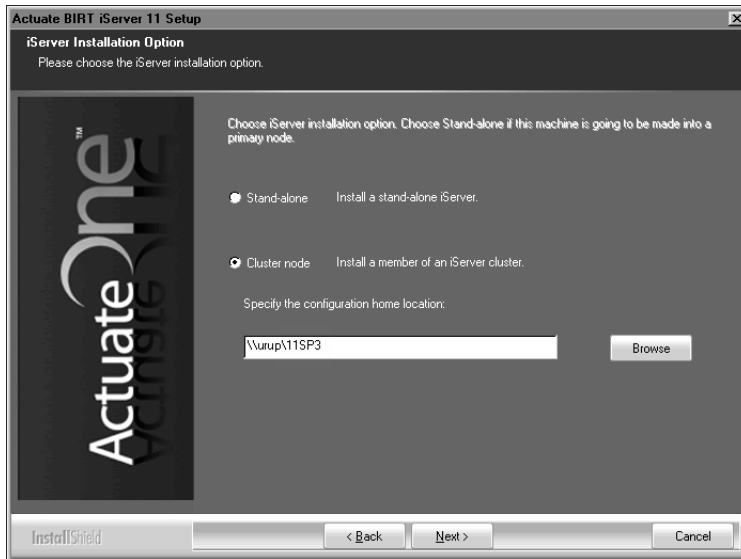


Figure 5-6 Choosing the option to install a node

- 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 5-7. Choose Next. Alternatively, choose Try out the product using the included evaluation software license if you do not have a purchased license.

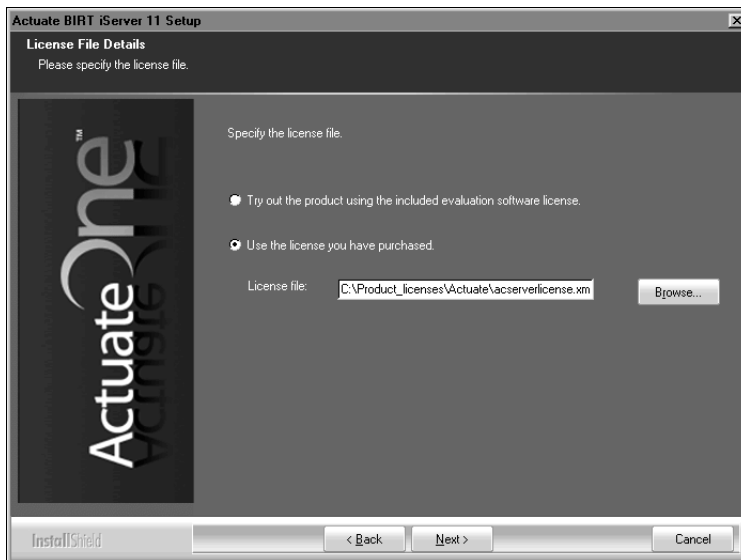


Figure 5-7 Specifying the license file

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 5-8.




Figure 5-8 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 5-9. 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 5-10. 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 BIRT iServer service.

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 5-10. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

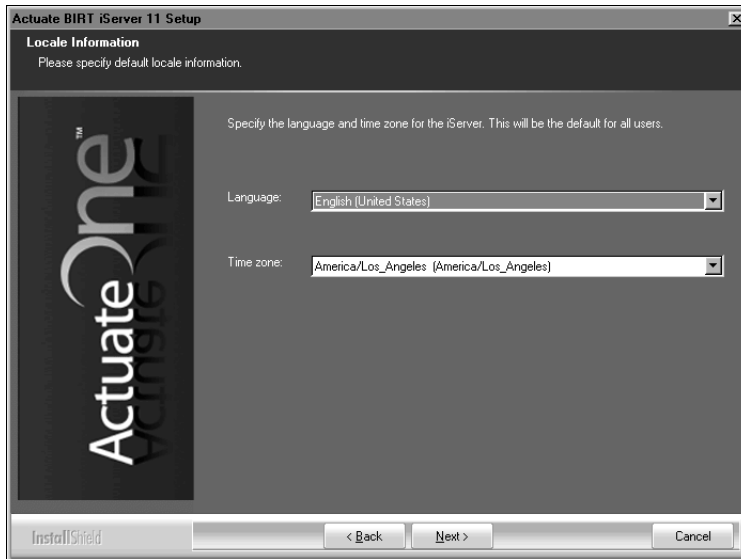


Figure 5-9 Specifying a language

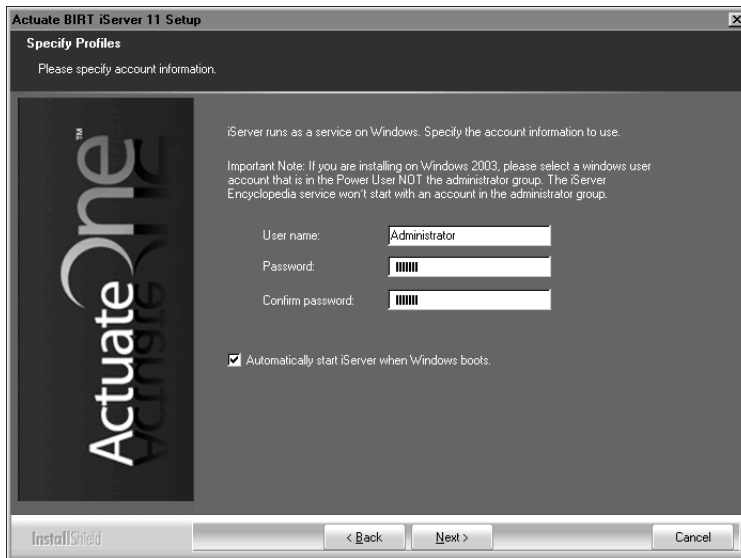


Figure 5-10 Specifying the account for running the iServer service

- 10** In iServer Configuration, shown in Figure 5-11, 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.

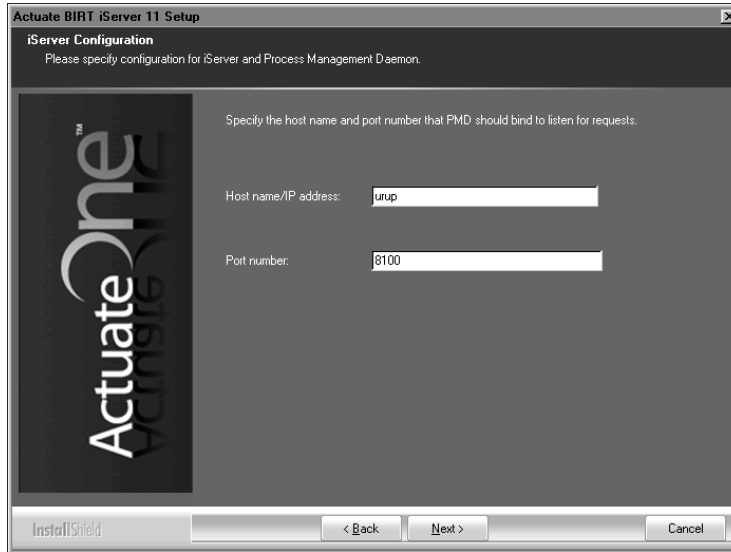


Figure 5-11 Identifying the cluster node and port to bind iServer service

- 11** In System Configuration Password, type a Configuration Console password, as shown in Figure 5-12. Note that the default user name for Management Console is Administrator with no password. Choose Next.

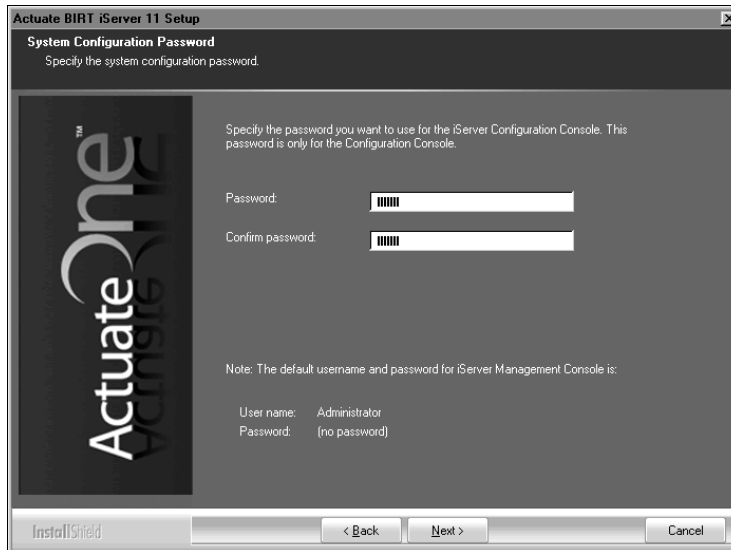


Figure 5-12 Specifying the password for using Configuration Console

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

- 1 Specify the following information, as shown in Figure 5-13. Choose Next:
 - Host name and port number for the following items:
 - PMD (Process Management Daemon) Configuration
 - iServer Configuration
 - Default volume name

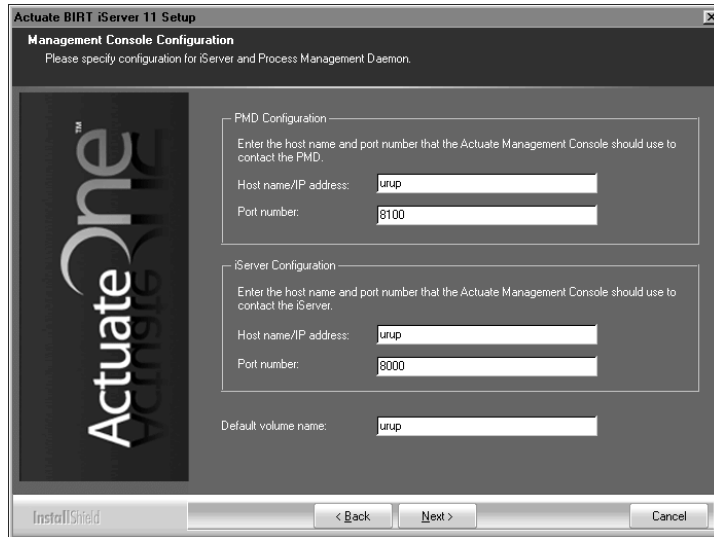


Figure 5-13 Specifying the Management Console Configuration

- 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 5-14. Choose Next.

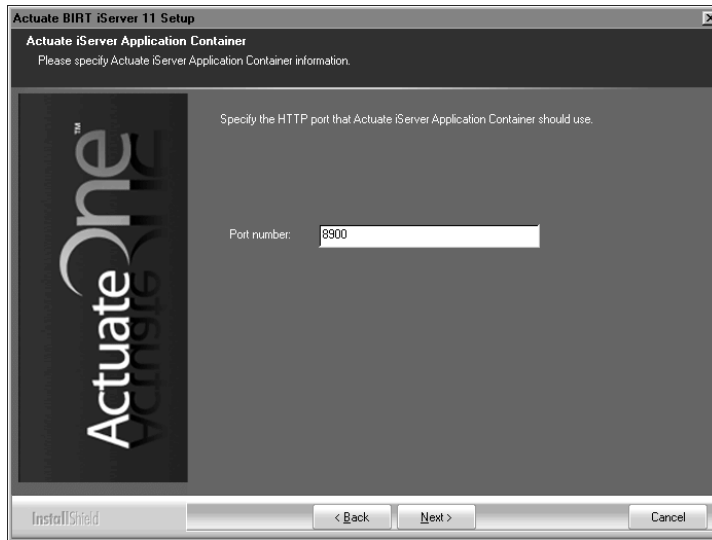


Figure 5-14 Specify the port number for iServer Application Container

- 3 In Context Path, type the context path for Management Console or accept the default path, /acadmin, as shown in Figure 5-15. Choose Next.

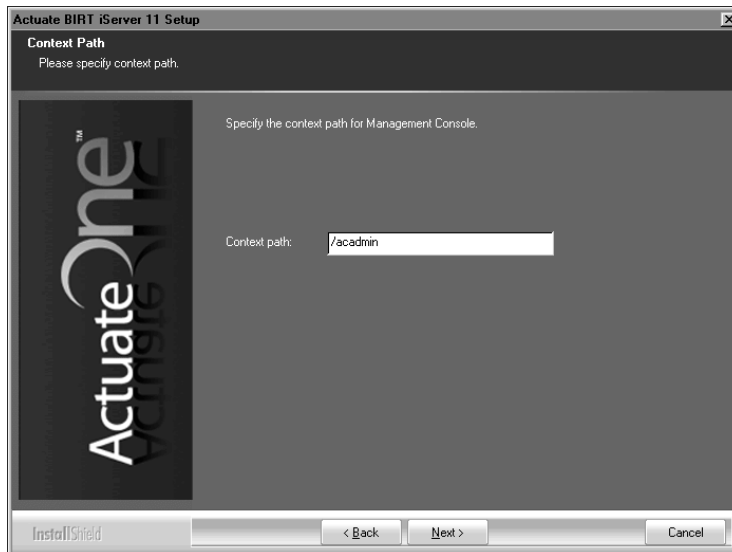


Figure 5-15 Specifying the context path

- 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 5-16. Choose Next.



Figure 5-16 Specifying a program folder

- 13** In Start Copying Files, review the settings shown in Figure 5-17. Choose Next to start copying files.

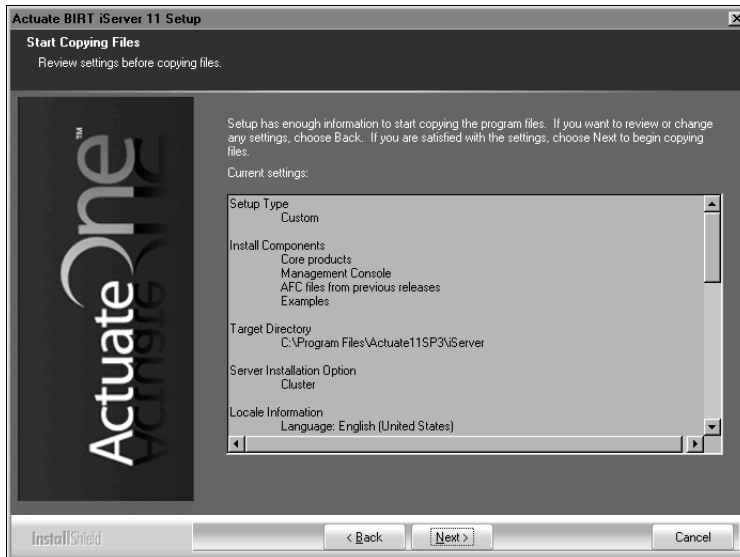


Figure 5-17 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 5-18.



Figure 5-18 Viewing setup status

- 14 When the setup success message appears, as shown in Figure 5-19. Choose Finish to exit the wizard.



Figure 5-19 Exiting the installation wizard

- 15** The installation program prompts you to install the online help from the following location:

<http://www.actuate.com>

Alternatively, you can install the online help and manuals in PDF format using the Actuate Localization and Online Documentation media. Figure 5-20 shows the prompt.




Figure 5-20 Viewing the install online help and manuals prompt

- 16** If you chose to install Management Console in step 4, the installation program installs shortcuts on the desktop, as shown in Figure 5-21.



Figure 5-21 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.
- 17** The installation program automatically opens Configuration Console in a browser, as shown in Figure 5-22, for the administrator to complete the process of adding the node to the cluster. Log into Configuration Console and perform the steps described in “Adding a node to a cluster,” later in this chapter.

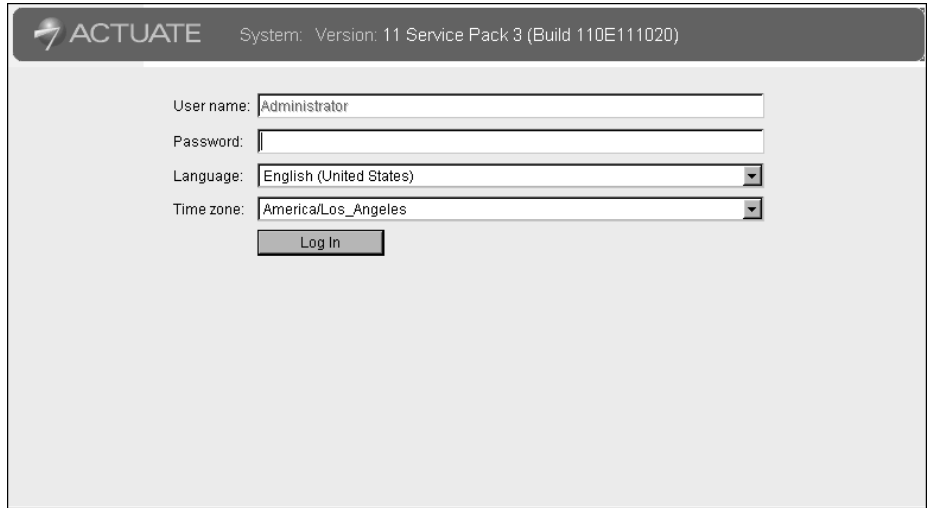


Figure 5-22 Logging in to Configuration Console

Adding a node to a cluster

The Actuate BIRT iServer 11 service on Windows must be running on all cluster nodes. The nodes must communicate with each other over TCP/IP. If necessary, reconfigure your firewall to allow communication between the nodes. Figure 5-23 shows an example of a pre-existing cluster.

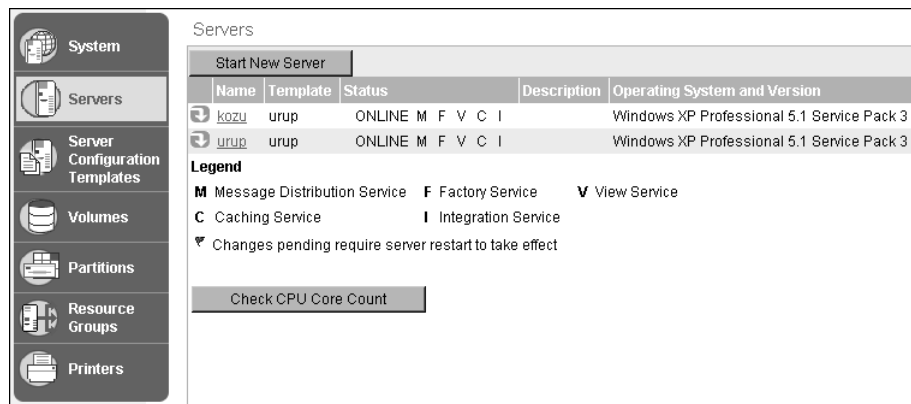


Figure 5-23 Viewing the cluster configuration before adding a node

After installing a node as an iServer instance on a machine, the administrator must still add the node to the cluster. When adding a node to a cluster setup, the administrator must verify that the configuration home directory specified during

the install procedure points to the shared configuration home directory for the cluster and all Encyclopedia volume resources are accessible.

The cluster node must have shared, read-write access to the following system resources:

- `AC_DATA_HOME\config`, including any version-related subfolders on the primary node
In a Release 11 Service Pack 3 installation, the configuration files are located in `AC_DATA_HOME\config\11SP3` by default.
- `AC_DATA_HOME\encyc` or other volumes, including all `file`, `fileType`, `status`, and `tempRov` subfolders
In an iServer installation, where there has been no activity on the system, the `status` or `tempRov` folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes.

The administrator must verify or edit the shared `acpmdconfig.xml` file to contain the following information:

- `<AC_CONFIG_HOME>` to point to the shared configuration home directory for the cluster
- `<AC_TEMPLATE_NAME>` to specify the server template from the available server templates listed in the shared `acserverconfig.xml` file

How to verify and edit `acpmdconfig.xml` file settings

To verify and edit `acpmdconfig.xml` file settings, perform the following tasks:

- 1 Stop the Actuate BIRT iServer 11 service running on the recently installed cluster node.
- 2 Using a text editor, such as Notepad, open `acpmdconfig.xml`, which by default is located in `AC_SERVER_HOME/etc`.
- 3 Verify or edit `<AC_CONFIG_HOME>` to point to the shared configuration home directory for the cluster.

This location is the path that you specified for the configuration home directory during the install procedure, as shown in Figure 5-6. The configuration home location is the shared directory containing `acserverconfig.xml`, `acserverlicence.xml`, and `aconfigowner.lock` files for the cluster. The administrator should specify the location using the Universal Naming Convention (UNC) Format.

In a Release 11 Service Pack 3 installation, the configuration files are located in `AC_DATA_HOME\config\11SP3` by default. If the shared configuration folder is in this default location on a server named `urup`, then the

administrator should specify \\urup\11SP3 as the path, as shown in Listing 5-1.

- 4 Verify or edit <AC_TEMPLATE_NAME>, shown in Listing 5-1, to specify the server template name from the available server templates listed in the shared acserverconfig.xml file, shown in Listing 5-2.
- 5 Save acpmdconfig.xml.

Listing 5-1 acpmdconfig.xml

```
<PMDConfig>
  <!--Actuate system Type -->
  <System>Cluster</System>
  <Mode>Default</Mode>
  <PMDConfigFileVersion>2</PMDConfigFileVersion>
  <!--Daemon SOAP endpoint information -->
  <DaemonSOAPPport>8100</DaemonSOAPPport>
  <!-- Disk Thresholds are in MB -->
  <MinDiskThreshold>100</MinDiskThreshold>
  <LowDiskThreshold>300</LowDiskThreshold>
  <!--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\11SP3</AC_CONFIG_HOME>
    <AC_JRE_HOME>
      C:\Program Files\Common Files\Actuate\11.0\JDK160\jre
    </AC_JRE_HOME><AC_JRE64_HOME/>
    <AC_JAVA_HOME>
      C:\Program Files\Common Files\Actuate\11.0\JDK160
    </AC_JAVA_HOME>
    <AC_ODBC_HOME>
      C:\Program Files\Common Files\Actuate\11.0\odbc
    </AC_ODBC_HOME>
    <AC_SERVER_IP_ADDRESS>urup</AC_SERVER_IP_ADDRESS>
    <AC_SOAP_DISPATCH_ADDRESS>urup</AC_SOAP_DISPATCH_ADDRESS>
    <AC_DOC_BASE>
      http://www.actuate.com/documentation/R11SP3
    </AC_DOC_BASE>
    <AC_ICU_DATA>
      C:\Program Files\Actuate\11SP3\iServer\bin
    </AC_ICU_DATA>
  </Server>
  <!-- Servlet Container information -->
  <ServletContainer>
    <Startup>Auto</Startup>
    <JavaOpts
```

```

        Args="-Xms128m -Xmx512m -XX:MaxPermSize=128m
            -Djava.awt.headless=true"/>
    </ServletContainer>
</PMDCConfig>

```

The administrator must also verify or edit the shared `acserverconfig.xml` file to contain the following information:

- `<ServerFileSystemSetting>` points to the shared drive location that contains the Encyclopedia volume data files
- `server <ConnectionProperty>` specifies the network name of the node that contains the shared Encyclopedia volume database

How to verify and edit `acserverconfig.xml` file settings

To verify and edit `acserverconfig.xml` file settings, perform the following tasks:

- 1 Stop the Actuate BIRT iServer 11 service running on the cluster node that contains the shared configuration home directory.
- 2 Using a text editor, such as Notepad, open the `acserverconfig.xml` file in the configuration home directory.

The configuration home directory is the shared directory on the node that contains the `acserverconfig.xml`, `acserverlicence.xml`, and `acconfigowner.lock` files for the cluster. In a Release 11 Service Pack 3 installation, the configuration files are located in `AC_DATA_HOME\config\11SP3` by default. The location is the path that you specified for the configuration home directory during the install procedure, as shown in Figure 5-6.

- 3 In `<Template>` settings for the node, verify or edit `<ServerFileSystemSettings>` to make sure that it points to the location that contains the Encyclopedia data files, as shown in Listing 5-2.

In Listing 5-2, the Path setting for `DefaultPartition` is `AC_DATA_HOME/encyc`.

- 4 In `<MetadataDatabase>` settings, verify or edit the `<ConnectionProperty>` for the server to make sure that it specifies the network name, not `localhost`, of the node on which the Encyclopedia volume database resides, by performing the following tasks:

- 1 Locate the `<ConnectionProperties>` element under the `<MetadataDatabase>` element.
- 2 In `<ConnectionProperties>` locate:

```

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

```

- 3 Change Value from localhost to the name of the machine on which the Encyclopedia volume database resides, such as urup, as shown in the following code:

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

- 5 Save acpmdconfig.xml.

Listing 5-2 acserverconfig.xml

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<Config>
  ...
  <Templates>
    <Template
      Name="urup"
      PMDPort="8100"
      ActuateBuild="110E111020"
      ActuateVersion="11 Service Pack 3"
      ServerSOAPPort="11100"
      AppContainerPort="8900"
      RequesterRSAPIVolume="corp">
    <ReportingService
      EnableGenerationService="true"/>
    <SOAPDispatchService
      EnableRequestService="true"
      SOAPDispatchSOAPPort="8000"/>
    <ViewingService
      EnableRender="false"
      EnableViewingService="true"/>
    <IntegrationService
      EnableIntegrationService="true"/>
    <CachingService
      EnableCachingService="true"/>
    <ServerFileSystemSettings>
      <ServerFileSystemSetting
        Name="DefaultPartition"
        Path="$AC_DATA_HOME$/encyc"/>
    </ServerFileSystemSettings>
```

Start the Actuate BIRT iServer 11 service on the nodes. The new cluster node will automatically read the settings in the acserverconfig.xml file in the shared configuration home directory to access its template, then join the cluster, as shown in Figure 5-25.

How to start an iServer cluster using Configuration Console

To start iServer using Configuration Console, perform the following tasks:

- 1 On the node containing the configuration home directory for the cluster, log in to Configuration Console and choose Advanced view. Choose Servers, then choose Start New Server.
- 2 On Servers—Start New Server, as shown in Figure 5-24, perform the following tasks:
 - 1 In Server name, type the name of the cluster node.
 - 2 In Host Name or IP Address, type the name or IP address of the cluster node.
 - 3 In iServer Process Manager Port Number, type the Daemon listen port number. The default value for this port is 8100. You specify this port number during the install procedure, as shown in Figure 5-24.
 - 4 In Server template name, choose the name of the template that the cluster node uses.

Choose OK.

Servers > Start New Server

Server name:

Host Name or IP Address: * [refresh]

iServer Process Manager Port Number: * [refresh]

Server template name: * [refresh]

* These fields are required and cannot be left blank

* [refresh] These fields require server restart to take effect

OK Cancel

Figure 5-24 Preparing to start a new server

- 3 Log out of Configuration Console.
- 4 Restart the Actuate BIRT iServer 11 services on the node containing the configuration home directory for the cluster then the new node.
- 5 Log in to Configuration Console and choose Advanced view. Choose Servers from the side menu. The new cluster node automatically reads the `acserverconfig.xml` in the shared configuration home directory to access its template, then joins the cluster, as shown in Figure 5-25.

The screenshot shows a management console with a sidebar on the left containing icons for System, Servers, Server Configuration Templates, Volumes, Partitions, Resource Groups, and Printers. The main window is titled "Servers" and features a "Start New Server" button at the top. Below this is a table with columns for Name, Template, Status, Description, and Operating System and Version. Three servers are listed: kazu, kazu, and urup, all using the "urup" template and running "Windows XP Professional 5.1 Service Pack 3". The Status column contains a sequence of letters (M, F, V, C, I) representing service states. A legend below the table defines these letters: M (Message Distribution Service), F (Factory Service), V (View Service), C (Caching Service), and I (Integration Service). A note states that changes pending require a server restart. A "Check CPU Core Count" button is located at the bottom of the main window.

Name	Template	Status	Description	Operating System and Version
kozu	urup	ONLINE M F V C I		Windows XP Professional 5.1 Service Pack 3
kazu	urup	ONLINE M F V C I		Windows XP Professional 5.1 Service Pack 3
urup	urup	ONLINE M F V C I		Windows XP Professional 5.1 Service Pack 3

Legend
M Message Distribution Service **F** Factory Service **V** View Service
C Caching Service **I** Integration Service
 ⚠ Changes pending require server restart to take effect

Check CPU Core Count

Figure 5-25 Viewing the cluster configuration after adding a node

Installing BIRT iServer in a cloud

This chapter contains the following topics:

- Understanding a BIRT iServer cloud deployment
- Preparing to deploy BIRT iServer in a cloud 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. The administrator deploys the image by unbundling the archive or installing a virtual image on the target machine.
- **Ready-to-launch iServer instance**
Launching an iServer instance requires minimal work. After configuring the appropriate environment variables, the administrator launches the deployed iServer image from the command line by passing the necessary arguments or creates a script that executes 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 each other 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. 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.

For more information on setting up a cluster after performing a cloud deployment of a stand-alone iServer installation, see Chapter 6, “Installing a BIRT iServer cluster,” earlier in this book, and Chapter 9, “Clustering,” in *Configuring BIRT iServer*.

Preparing to deploy 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 or TAR file, or virtual image, unbundling the image on the target machine.

In addition, the administrator must install a supported version of the JAVA SE Development Kit (JDK) 1.6 (32-bit) or earlier. The JDK can be downloaded from <http://www.oracle.com/technetwork/java/javase/downloads/index.html> if it is not already installed on the machine.

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 version 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\Actuate11\iServer
```

The data installs in the following directory:

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

In a default Linux or UNIX setup, the iServer run-time environment installs in the following directory:

```
$HOME/AcServer
```

The data installs in the following directory:

```
$HOME/AcServer/data
```

Deploying an iServer distribution package

To extract the iServer run-time resources, configure the run-time environment, start iServer, and perform the following tasks:

- 1 Create a new folder titled Actuate 11 in a directory location outside of C:\Program Files or C:\Program Files(x86).
- 2 Extract the contents of ActuateBIRTiServer.zip to the folder created in the previous step.

In the iServer Release 11 configuration, the setting for the environment variable, AC_SERVER_HOME, specifies the location of the iServer binaries. The variable AC_JAVA_HOME specifies the location of the JDK.

- 3 To specify the location of AC_SERVER_HOME and AC_JAVA_HOME in the startiServer script, perform the following tasks:
 - 1 Using the text editor, open startiServer script located in the directory where you deployed BIRT iServer 11.
 - 2 Specify the path to the AC_SERVER_HOME and AC_JAVA_HOME on your machine, as shown in Listing 6-1.
 - 3 Save and close the file.

Listing 6-1 startiServer 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

setlocal enabledelayedexpansion

set AC_SERVER_HOME=C:\Actuate11\iServer
set AC_JAVA_HOME=C:\JDK160\bin
...
```

- 4 From the Windows menu, choose Start → Run.
- 5 Open a command prompt by typing cmd, and choose OK.

- 6 Navigate to the folder where you extracted the iServer package.
- 7 Type `startiServer`, and press Enter to execute the script.
- 8 The script determines the settings and paths for the environment variables that iServer uses to locate data and run-time resources.

The script installs the pre-requisite C++ runtime components, sets up the iServer deployment files, creates the iServer Encyclopedia volume, starts the PostgreSQL database system used to store Encyclopedia volume metadata, and finally starts BIRT iServer 11.

The deployment script asks whether to upload the Encyclopedia volume samples, as shown in Figure 6-1. Type 'y' and press Enter. Alternatively, type 'n' to skip this step. If you do not automatically upload the Encyclopedia volume samples in this step, you can upload these objects later.

```

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\bin
Installing pre-requisite C++ runtime components. This will take a few seconds ...
Configuring deployed iServer files ...
The system cannot find the path specified.
Creating the iServer encyclopedia volume. This will take a few seconds ...
Starting PostgreSQL for Actuate ...
Starting Actuate BIRT iServer 11 ...
Waiting a few moments for the iServer to startup and initialize ...
Would you like to upload sample content into the encyclopedia?(y or n):_

```

Figure 6-1 Running startiServer script

- 9 When the script starts iServer, a new command prompt window opens to serve as a console from which iServer can be shut down and restarted, as shown in Figure 6-2.

```

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

```

Figure 6-2 Viewing BIRT iServer 11 console

You can now make modifications to iServer parameters and server templates, create an archive of the iServer folder to replicate, and rapidly re-deploy the newly modified version in a cloud computing environment. For more information about configuring a cluster, see Chapter 6, “Installing a BIRT iServer cluster,” earlier in this book, and Chapter 9, “Clustering,” in *Configuring BIRT iServer*.

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 components:

- BIRT iServer 11 Information Console

Launches the Information Console for viewing report documents

- BIRT iServer 11 Management Console
Launches Management Console for setting up user accounts and scheduling reports

To access Configuration Console to administer iServer after starting the system, open a browser and enter the following URL, as shown in Figure 6-3:

<http://localhost:8900/acadmin/config/>

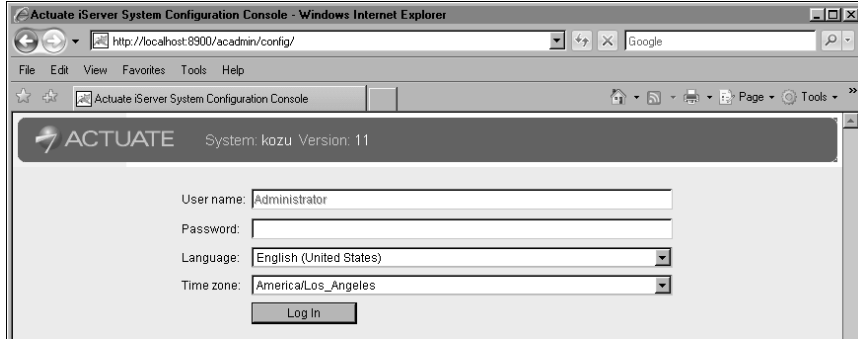


Figure 6-3 Logging into Configuration Console

To access Management Console to administer the Encyclopedia volume after starting the system, use the provided BIRT iServer 11 Management Console shortcut. Alternatively you can open a browser and enter the following URL, as shown in Figure 6-4:

<http://localhost:8900/acadmin/>

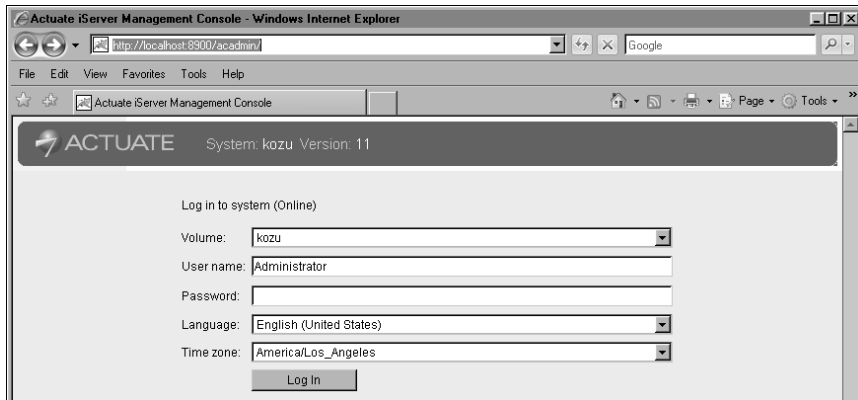


Figure 6-4 Logging into Management Console

How to start and stop BIRT iServer

To start iServer, perform the following tasks:

- 1 Open a command prompt and navigate to the folder where you deployed BIRT iServer.
- 2 Type startiServer and choose Enter to execute the startup script. A new command window opens, entitled Actuate BIRT iServer 11. The first window must remain open, since the PostgreSQL service runs from this window.
- 3 After iServer starts, a browser window to Information Console opens automatically.

To stop iServer, perform the following tasks:

- 1 In the command prompt entitled Actuate BIRT iServer 11, type 's' and choose Enter to shut down iServer. The window closes automatically after iServer shuts down.
- 2 After iServer shuts down, close the main command prompt window.

How to manually upload Encyclopedia samples

You can upload the Encyclopedia samples at any time, by performing the following tasks:

- 1 If BIRT iServer is running, stop it using the steps described in Starting and stopping BIRT iServer.
- 2 Run the startiServer deployment script.
- 3 When the deployment program asks whether you would like to upload sample content into the Encyclopedia volume, type 'y' and press Enter.
- 4 The deployment program uploads the sample content and opens Information Console.

Specifying AC_SERVER_HOME

In iServer Release 11, the location of binary files changes to consolidate these resources in AC_SERVER_HOME to facilitate creating an iServer image for deployment in a cloud environment.

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
- msvcp90.dll
- LTDIS10N.dll
- ltfil10N.DLL
- msucr71.dll
- acicudt18.dll

- ltkrn10N.dll
- msvcr90.dll
- acr7790w.dll
- mfc90u.dll
- msvcr71.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

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.

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 download the DVD image, extract the files, move only some of the files to a new location, and attempt to install from that location.

About performing a full installation

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, you must install the documentation files, which are on the documentation DVD that ships with your Actuate Software.

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 If you downloaded iServer, run the `ActuateInformationConsole.exe` file. If you have a DVD or ftp distribution, run `setup.exe` in the Information Console directory. The welcome message appears, as shown in Figure 7-1. Choose Next.

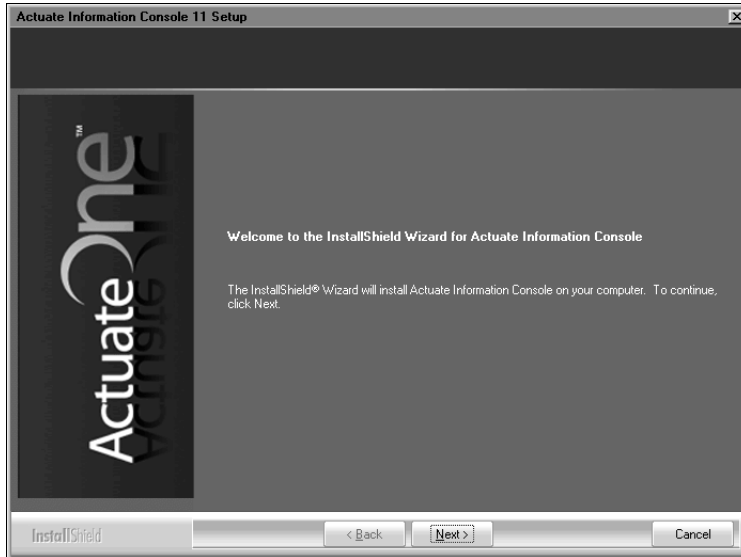


Figure 7-1 Viewing the welcome message

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

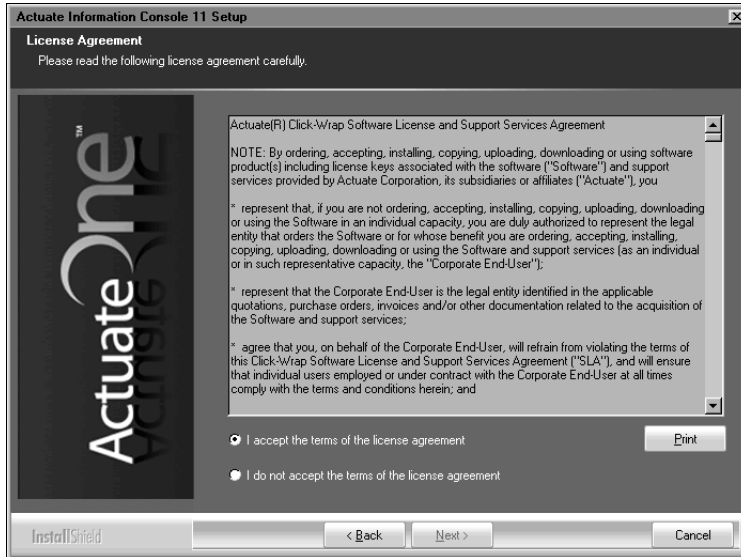


Figure 7-2 Accepting the license agreement

- 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.

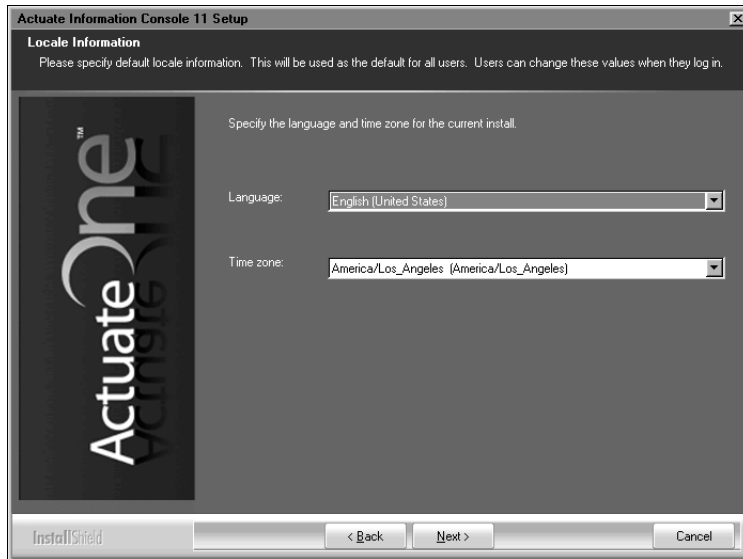


Figure 7-4 Specifying locale information

- 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.

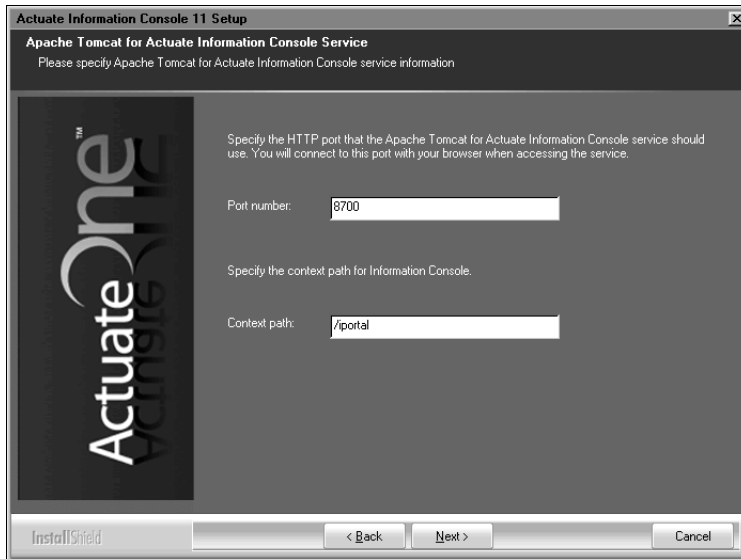


Figure 7-5 Specifying the port number and context path

- 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.

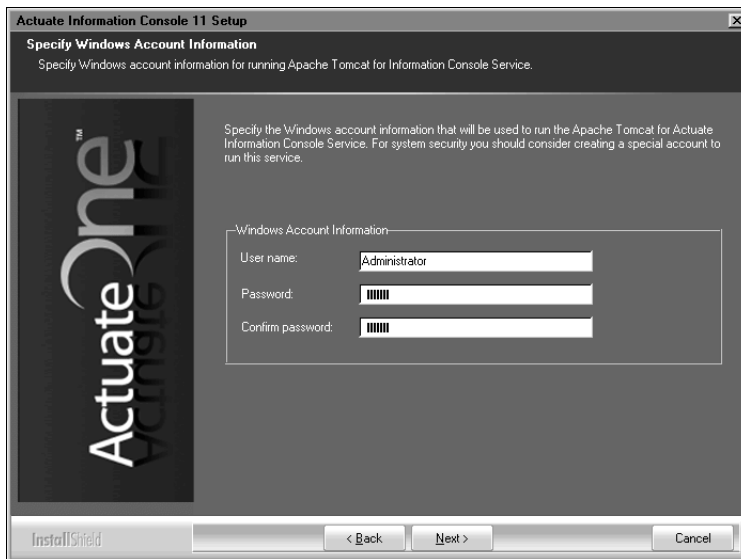


Figure 7-6 Specifying the account running the Information Console service

- 7 In iServer Information, accept the default values for Profile name, Host name, and Port number, as shown in Figure 7-7. 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.

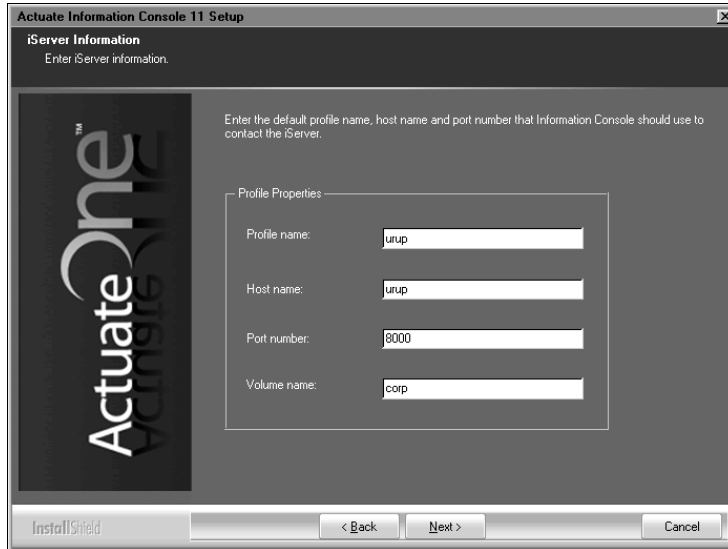


Figure 7-7 Specifying values for iServer profile properties

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

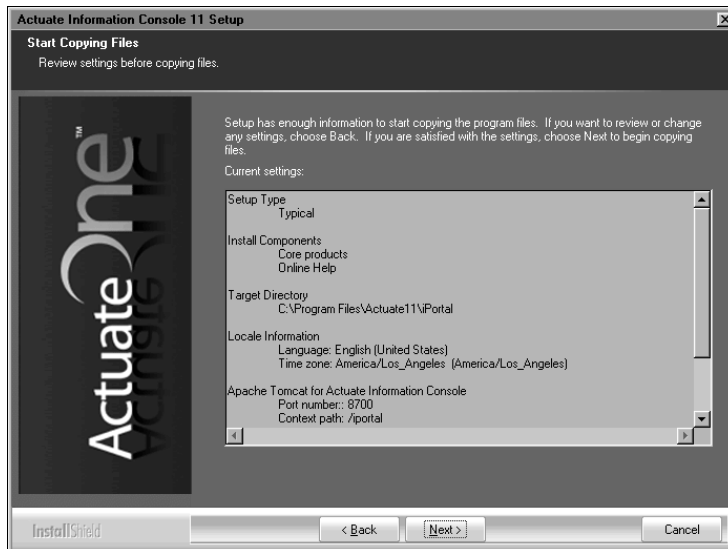


Figure 7-8 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 7-9.

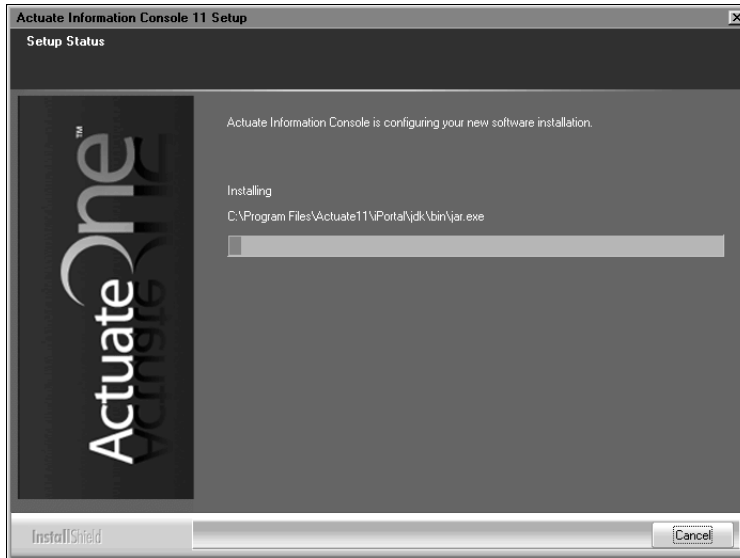


Figure 7-9 Setup Status

- 9 Select I would like to view the ReadMe file, as shown in Figure 7-10, if you want to review this documentation. Choose Finish to exit the wizard.



Figure 7-10 Exiting the installation wizard

If you chose to view the ReadMe file, the installation program opens the document, as shown in Figure 7-11.

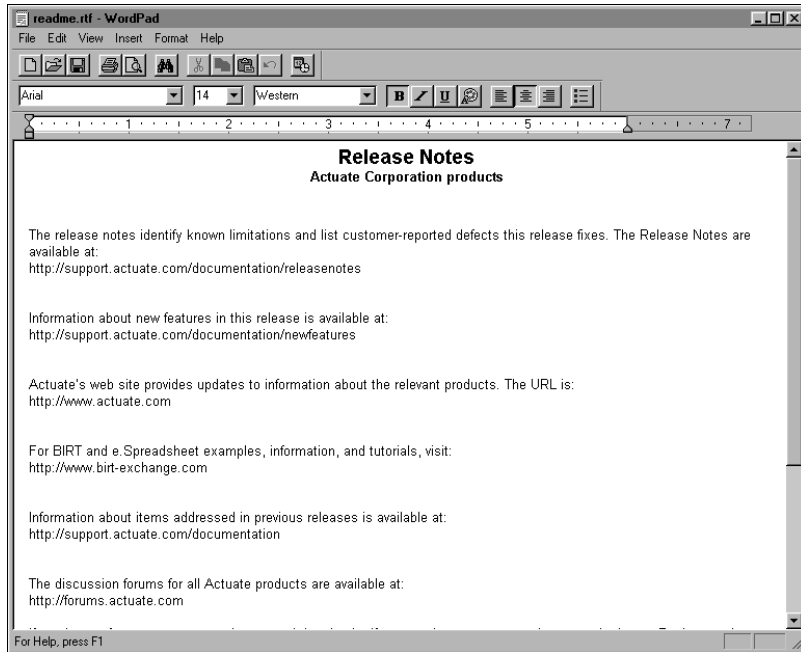


Figure 7-11 The ReadMe file

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 ActuateInformationConsole.war on the installation DVD.

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.

Table 7-1 Information Console configuration parameters (continued)

Parameter name	Description	Action
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** Insert the iServer System installation DVD.
- 2** Create a temporary directory, such as C:/Temp/ic.
If you use an existing directory, ensure that this directory is empty.
- 3** Open a command window and type the following commands, replacing the X: DVD drive letter with a drive letter appropriate to your system:
 - 1** Copy the WAR file to the temporary directory, as shown in the following example:

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

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

```
jar -xf ActuateInformationConsole.war
```

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

- 4 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 3, 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.

- 5 Save and close web.xml.

- 6 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/iportal/getfolderitems.do
?repositoryType=Enterprise&volume=volume1&serverurl=
http://iServer1:8000
```

where

- Actuate1:8900 is the name of your computer and the port you use to access Information Console.
- iportal is the context root for Information Console.
- ? indicates the beginning of a parameter that indicates where to access Information Console files.
- getfolderitems.do is the call to the default Information Console home page.
- repositoryType=Enterprise indicates that this Information Console connects to iServer.
- &volume=volume1&serverurl=http://iServer1:8000 specifies the Encyclopedia volume and 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

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 If you downloaded iServer Integration Technology, run the ActuateiServerIntegrationTechnology.exe file. If you have a DVD or ftp distribution, run setup.exe in the iServer directory. 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.

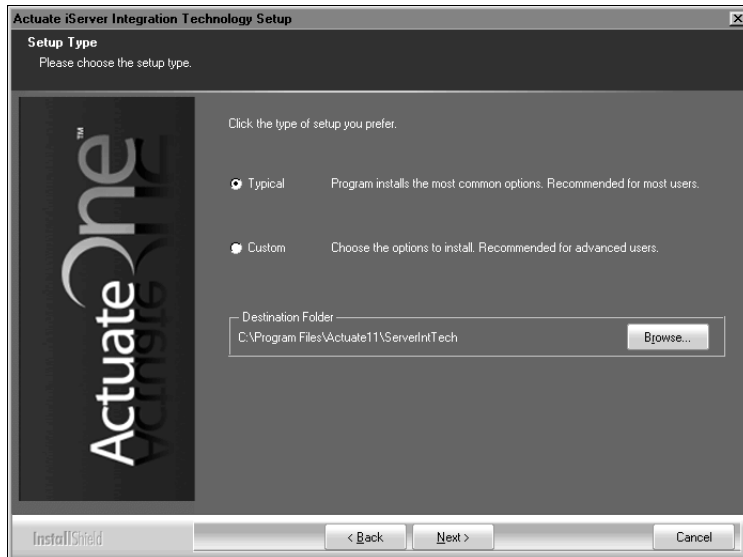


Figure 8-3 Selecting a typical installation

- 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.

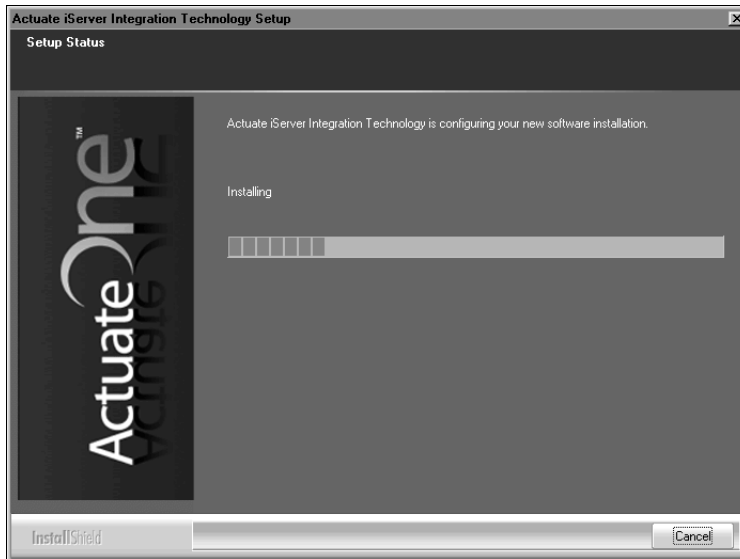


Figure 8-5 Viewing setup status

- 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.

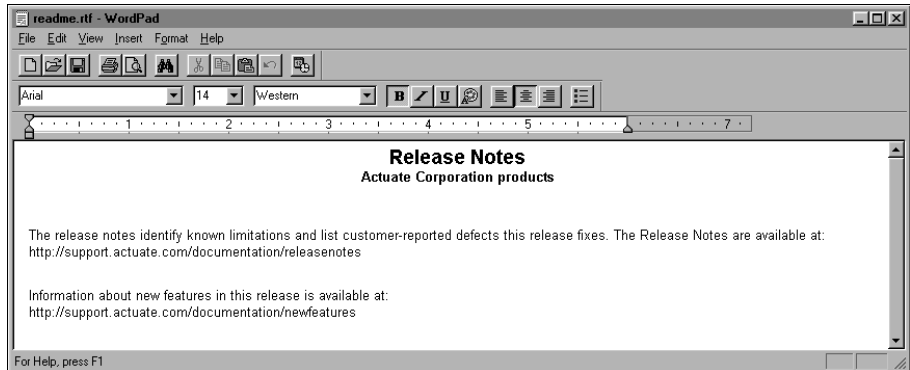


Figure 8-7 Viewing the ReadMe file

- 6 The installation program prompts you to install the online help and manuals using the Actuate Localization and Online Documentation media, as shown in Figure 8-8.



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

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/docupdate11sp1/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 If you downloaded Actuate Localization and Documentation, run the `ActuateLocalizationandOnlineDocumentation.exe` file. If you have a DVD or ftp distribution, run `setup.exe`. The welcome message appears, as shown in Figure 8-9. Choose Next.

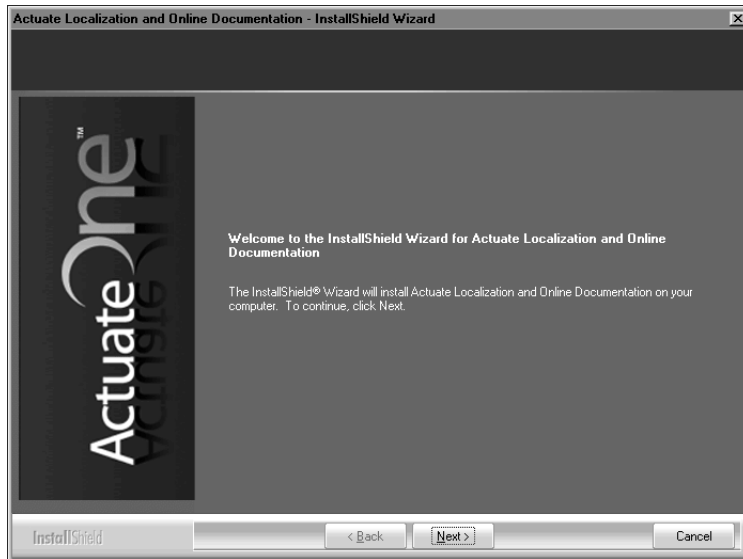


Figure 8-9 Viewing the welcome message

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

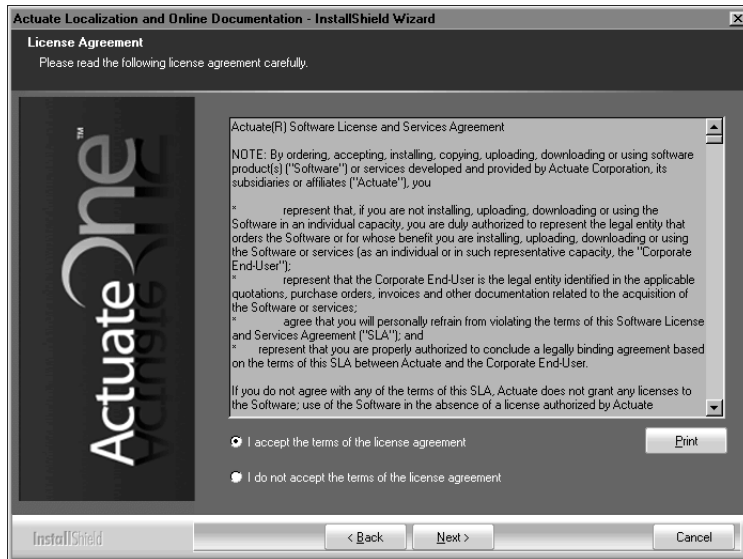


Figure 8-10 Accepting the license agreement

- 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.

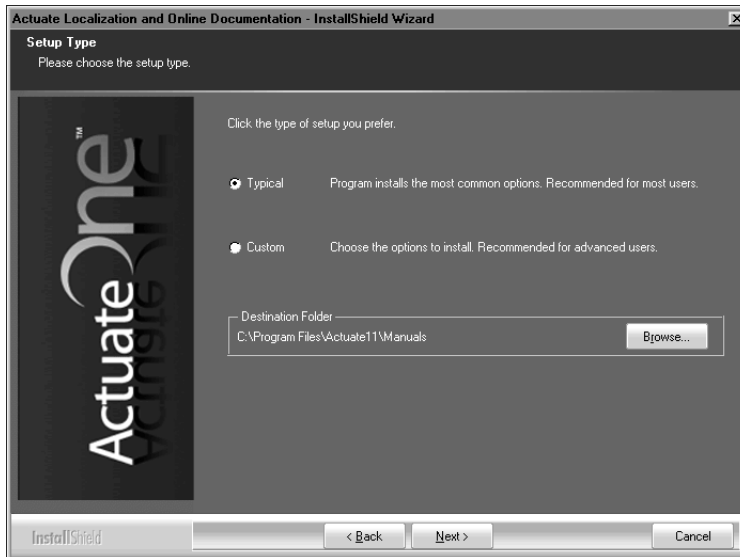


Figure 8-11 Specifying typical or custom setup type

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

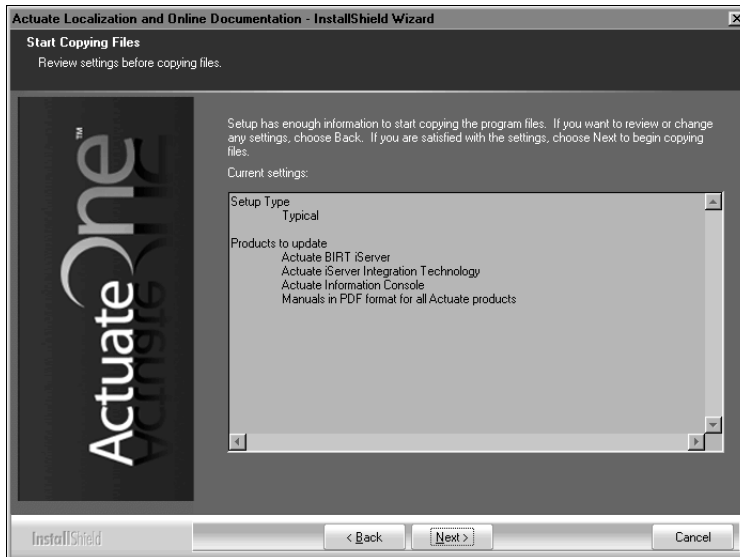


Figure 8-12 Reviewing settings before copying files

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

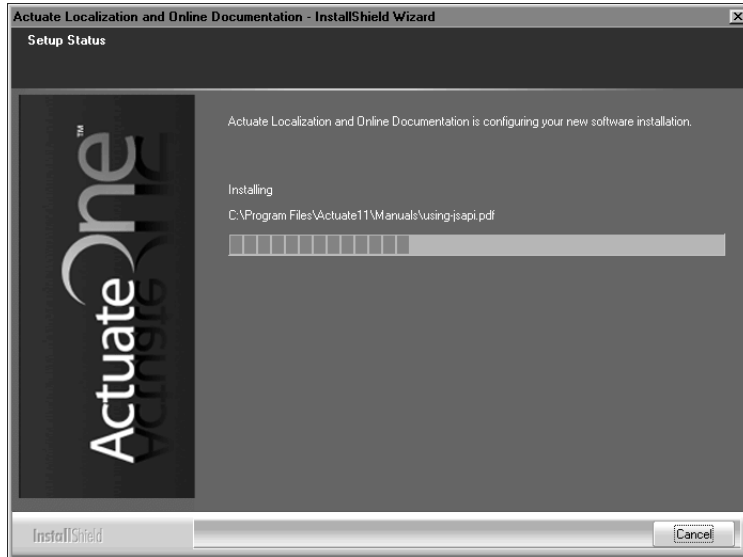


Figure 8-13 Viewing setup status

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.

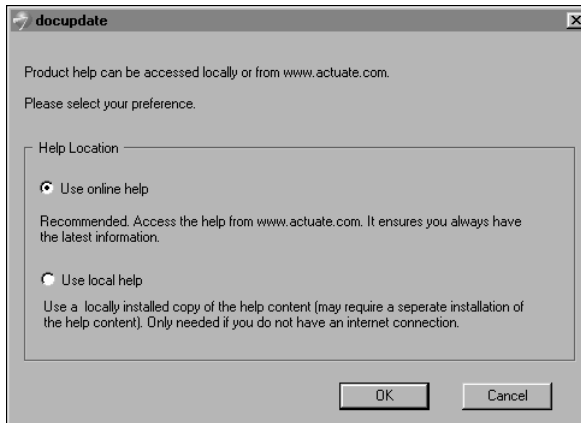


Figure 8-15 Selecting online or local help

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.

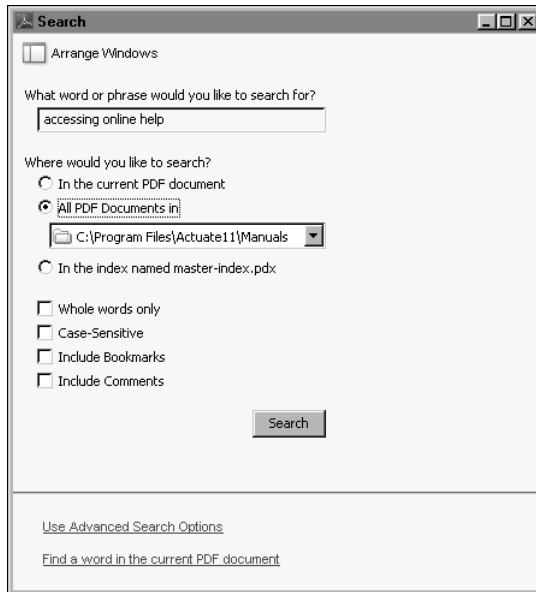


Figure 8-16 Specifying the search

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.

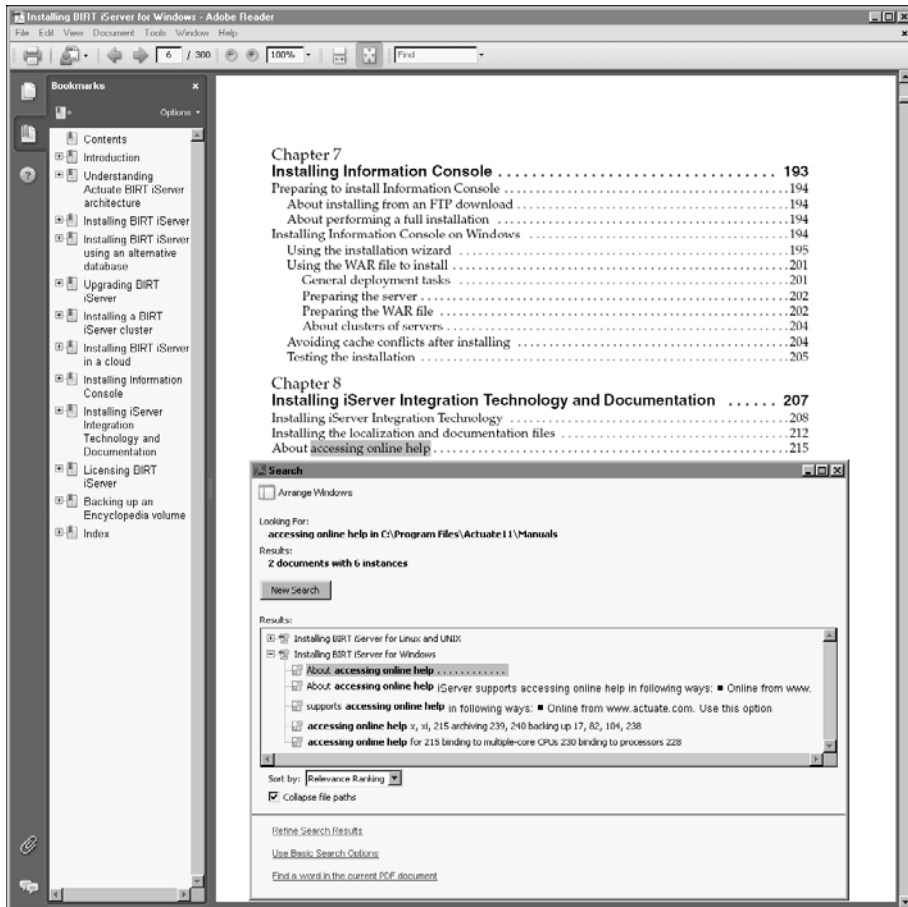


Figure 8-17 Viewing search results

Part **Three**

Licensing

Licensing BIRT iServer

This chapter discusses the following topics:

- Working with licensing
- Understanding the licensing options
- About license files
- Understanding node-key licensing
- Obtaining the installation license key file
- Collecting machine information for a node-key license
- Installing the license key
- Understanding CPU binding

Working with licensing

BIRT iServer System licensing supports running BIRT iServer with sets of features grouped as license options. You enable BIRT iServer System options using either a Named User, Platform, or Work Unit License.

The following list describes the types of licenses:

- **Named User License**

Specifies the maximum number of named users that can use an BIRT iServer System. A named user is a distinct individual who receives content and value from BIRT iServer.

In Release 11, a BIRT iServer administrator must specify the options that a user can access in an Encyclopedia volume. The administrator makes this specification by configuring the user Licensed Options properties in Management Console.

When you license an option, such as BIRT Option, e.Report Option, or BIRT Spreadsheet Option, the license entitles the user to access a single volume in BIRT iServer System. If the user needs additional Encyclopedia volumes for multiple applications, archiving, or other purposes, you must license the Multi-Tenant Option for each additional Encyclopedia volume the user needs to access.
- **Platform License**

Specifies a hybrid model that supports combining Named User Licenses with CPU Licenses for Actuate end-user components and the BIRT iServer System infrastructure. In a multiple-CPU environment, Actuate typically uses the Standard Performance Evaluation Corporation (SPEC) standard benchmark, for measuring machine capacity based on CPU, memory, disk, and network capacity.
- **Work Unit (WU) License**

Specifies iServer features and functionality using an aggregate model. This plan defines each iServer System resource as a work unit.

Understanding the 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
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

(continues)

Table 9-1 BIRT iServer System license options (continued)

Option	Description	Supported releases
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 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
Information Object Caching	Provides the ability to cache data from an information object in a third-party database. This option requires separate licensing of a third party database server for data storage. This option is not available for an information object based on Actuate Basic technology.	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.

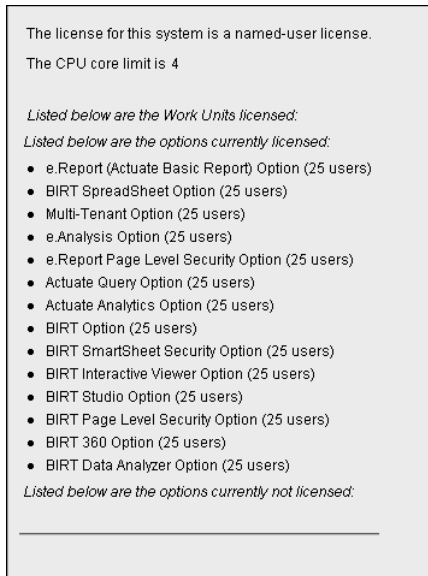


Figure 9-1 iServer License options

About license files

Actuate provides a temporary BIRT iServer license key file to use for the initial installation. The temporary BIRT iServer System license key expires 45 days after installation. A design or document run using a temporary license appears with a watermark when viewed.

After installing BIRT iServer System, the customer must collect information about the system in which BIRT iServer is running and transmit the information to Actuate Licensing. After receiving this information, Actuate Licensing issues a new BIRT iServer System license key file.

This license key 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 ConfigHomeDirectory variable in the acpmdconfig.xml file of each node, and accessible to all nodes in the cluster.

A node key associates an iServer node 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.

When upgrading a cluster node or installing iServer on a new machine, the customer must request a new node-key file and supply the machine ID of the new machine.

About 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.

The following sections provide a detailed description of how the license key process works in Actuate Release 11.

Understanding node-key licensing

Actuate license enforcement for Release 11 requires a single, shared license key for all nodes in a cluster. You receive a temporary license key from Actuate Licensing when you purchase the product.

The name for the BIRT iServer license key file uses the following format:

```
Actuate_iServer_key_XXXXX.xml
```

XXXXX is a unique five-digit number generated by Actuate Licensing when it creates the license key file.

The following sections describe the BIRT iServer System license installation process.

Obtaining the installation license key file

New customers receive an e-mail that contains the license key file information when Actuate processes the order. If you have a problem with a license key file, contact Actuate Licensing at licensing@actuate.com.

To obtain a new file for licensed products, visit the Support web site at:

```
http://support.actuate.com
```

A maintenance customer should have login information for the Actuate Support web site. If you do not have access, please contact Actuate Support at support@actuate.com

If you are not a direct Actuate customer, the partner or distributor who provides the product arranges for your license key file. If you have a problem obtaining your license key file from this source, please contact Actuate Licensing at licensing@actuate.com.

A license key file can be an expiring file that is valid until a specific date. If your license key file is an expiring file, you are reminded 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 key file, or if you have a problem with an expiring file, please contact Actuate Licensing at licensing@actuate.com.

You can easily modify your decisions about Actuate license options by contacting Actuate Licensing. 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.

Collecting machine information for a node-key license

After installing BIRT iServer System using the temporary license key file, you must collect information about the machine running this Actuate software and send it to Actuate Licensing. During the installation process, the InstallShield Wizard prompts you to provide the location of the `Actuate_iServer_key_xxxxx.xml` file. After providing the location of the license key, the InstallShield Wizard issues a prompt similar to the following message:

```
The iServer system license file is locked to the machines that are
used in the iServer system. The following machine id must be
used to request a node key license file from Actuate.
00-13-02-4e-3a-8e00-15-5a-2b-27-0fe0-16-ce-ec-43-52
```

```
Please contact Actuate Licensing (licensing@actuate.com or
http://www.actuate.com/licensing), or your representative, and
request a node key file.
```

```
The machine id required for the node key file can also be
generated by using the acmachineid utility that can be found in
the ACTUATE_HOME\iServer\bin folder.
```

```
Press Enter to continue.
```

The format of the alphanumeric string for the machine ID and location of the node-key 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.

Make a note of the machine ID in the installation prompt and send it to Actuate Licensing. Actuate Licensing processes your request and sends the new license key for BIRT iServer System.

You can also run the BIRT iServer utility, `acmachineid`, from the command line to generate the machine ID information as shown in the following Windows-based example:

```
STATUS:          OK
GEN_VERSION:     11
GEN_BUILD:       10C081031
MACHINEID: 00-13-02-4e-3a-8e00-15-5a-2b-27-0fe0-16-ce-ec-43-52
```

The `acmachineid` utility is located in the `bin` folder of the BIRT iServer installation.

Installing the license key

After installing BIRT iServer System using the temporary license key, the login screen displays two messages.

The following message about expiration of the initial license key always appears on the login screen regardless of the node-key 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 set of license keys from Actuate Licensing appears until you install the new license keys 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 node-key license after you install BIRT iServer System.

How to install the license key file

To update the license key file, perform the following tasks:

- 1 Verify that the format of the `Actuate_iServer_key_XXXXX.xml` license file name is correct.

An Actuate license key is an XML file. Actuate Licensing sends this XML file to you with an appended `.txt` file extension because transmitting a file with an `.xml` extension can cause problems in an e-mail system. You must remove the `.txt` extension from the file name before installing the license key file in the BIRT iServer System. Make sure that the file name contains the unique five-digit number generated by Actuate Licensing.
- 2 Copy the `Actuate_iServer_key_XXXXX.xml` license file to the shared location specified by the `ConfigHomeDirectory` variable in the `acpmdconfig.xml` file of each node in the cluster.
- 3 Log in to Configuration Console. For example, type `http://localhost:8900/admin/config/` in Address of a browser, and use the system configuration password you specified during installation.
- 4 Choose Update License.
- 5 Navigate to the location of the new license key file and select the file. Choose OK.
- 6 Restart any node where the node-key configuration changed.

If you change the machine for a node in a BIRT iServer cluster, you must reapply to Actuate Licensing for a new license file. If you replace the network card on a machine, such as a Windows system, you may be required to obtain a new license file, since the unique identifier for the network card is the source of the machine ID. Use the `acmachineid` utility to obtain the machine ID information and transmit the new information to Actuate Licensing.

Understanding CPU binding

CPU binding can work only 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. 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 9-2.

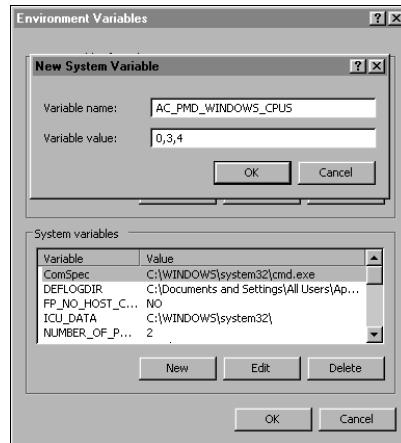


Figure 9-2 Creating the AC_PMD_WINDOWS_CPUS system variable
Choose OK. AC_PMD_WINDOWS_CPUS appears in System variables, as shown in Figure 9-3.

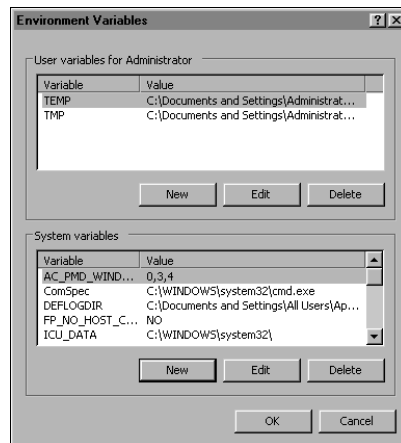


Figure 9-3 Viewing the AC_PMD_WINDOWS_CPUS system variable

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 affect 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 3 environment, the default AC_SERVER_HOME path is:

```
C:\Program Files\Actuate11SP3\
```


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\11SP3
```

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 backup 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 Backup 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.

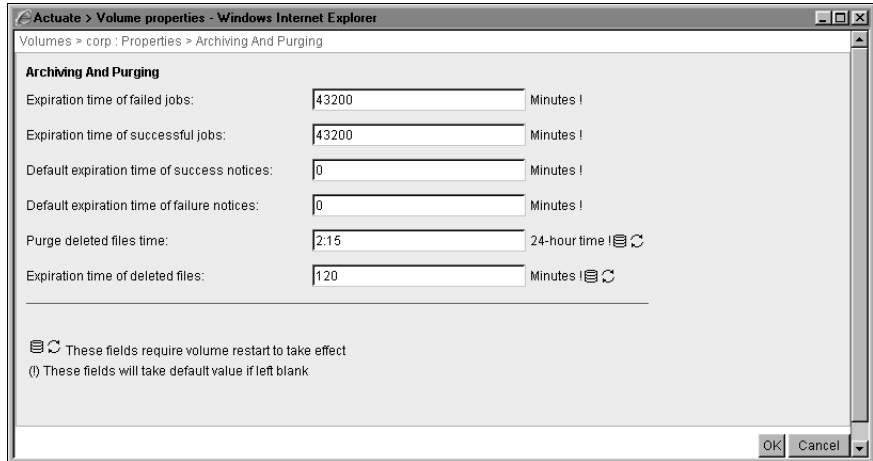


Figure 10-1 Configuring file purging properties

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 backup and restore an Encyclopedia volume that uses the OOTB PostgreSQL database to store the metadata. These demonstrations serve as a detailed reference example. Other RDBMS environments, such as a DB2, Microsoft SQL Server, or Oracle RDBMS, require similar procedures, which are covered in sections later in this chapter.

Backing up an Encyclopedia volume using 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:

```
C:\Actuate\iServer\encyc_backup
```

Figure 10-2 shows the new folder.




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

Backup 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 the 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.

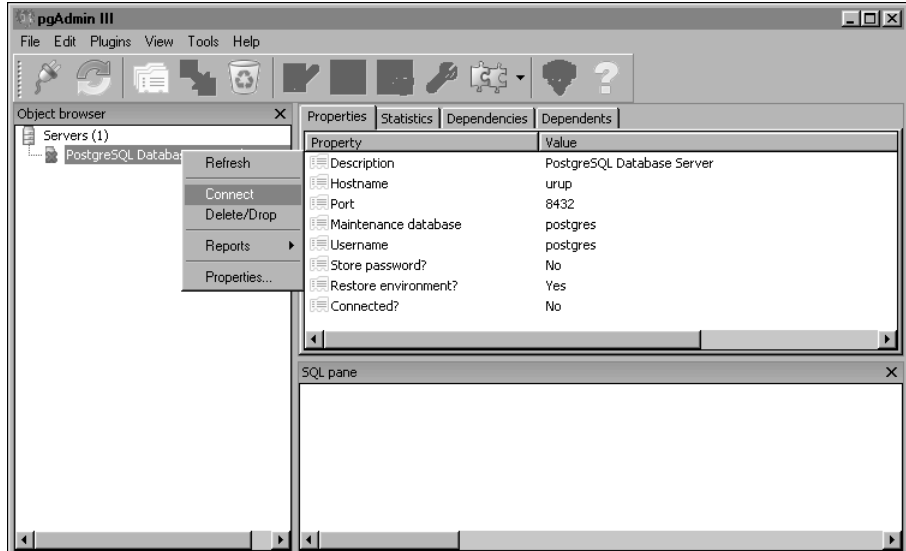


Figure 10-3 Connecting to PostgreSQL Database Server

- 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.

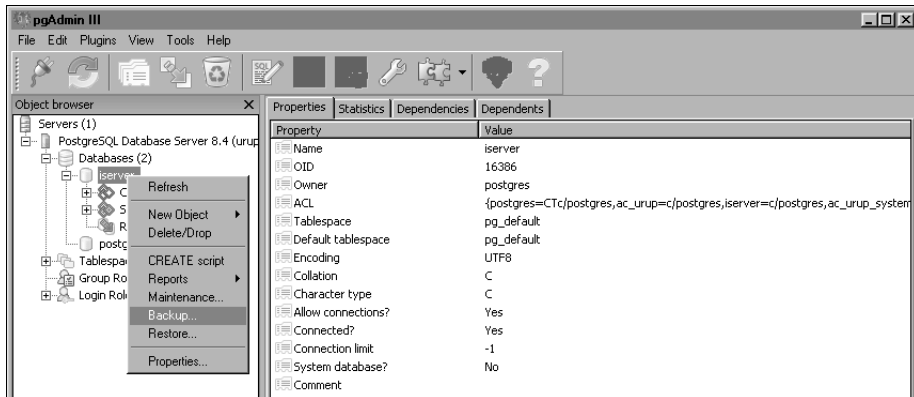


Figure 10-5 Choosing to back up the iserver database

- 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.

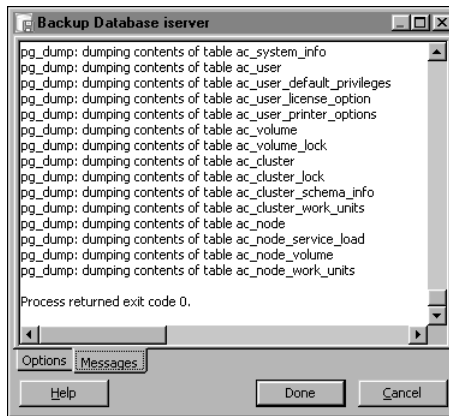


Figure 10-7 Viewing status messages `pg_dump` writes during execution

- 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.

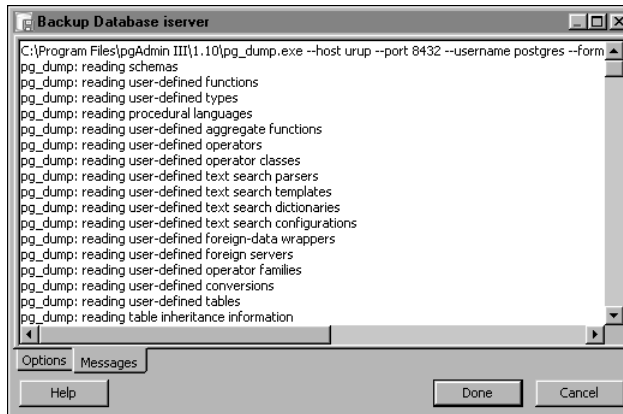


Figure 10-8 Viewing the command that executed `pg_dump`

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 backup 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 backup 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\Actuate11SP3\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 3, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP3
```

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.

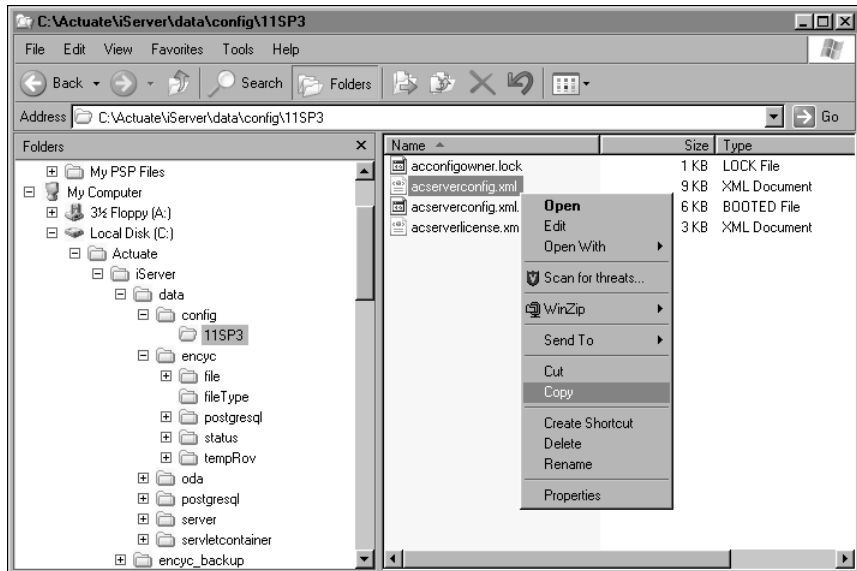


Figure 10-9 Copying acserverconfig.xml

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

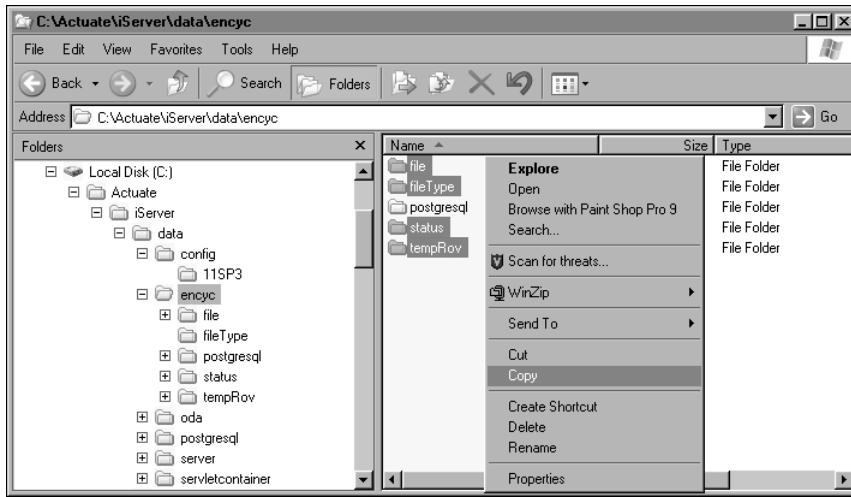


Figure 10-10 Copying the volume data folders

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

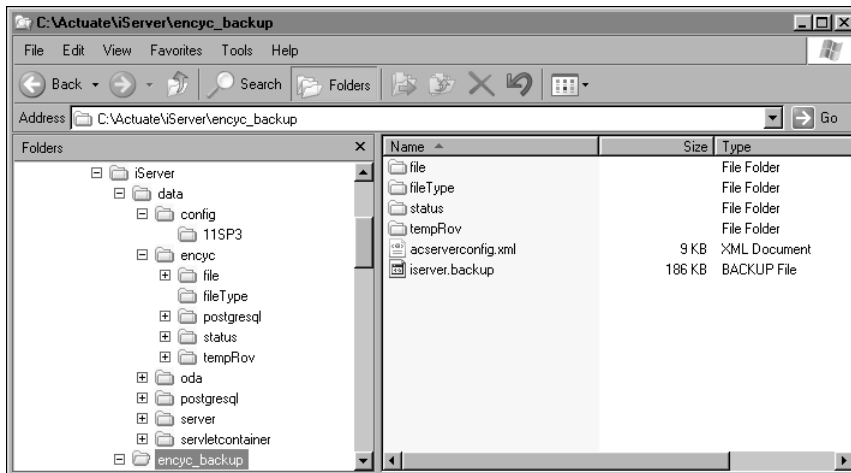


Figure 10-11 Viewing the backed up files

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.

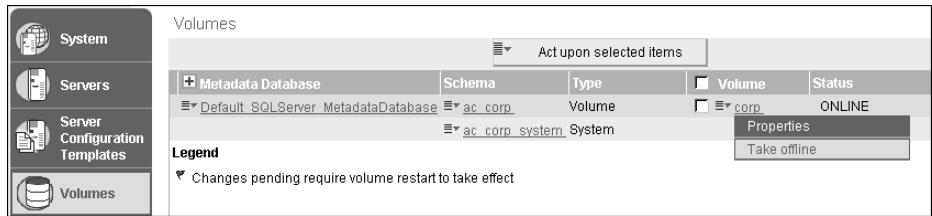


Figure 10-12 Taking the volume offline

How to restore the backed up volume data folders

- 1 In Windows Explorer, navigate to AC_DATA_HOME\config\11SP3.
- 2 Select acserverconfig.xml, right-click, and choose Delete, as shown in Figure 10-13. Confirm the deletion.

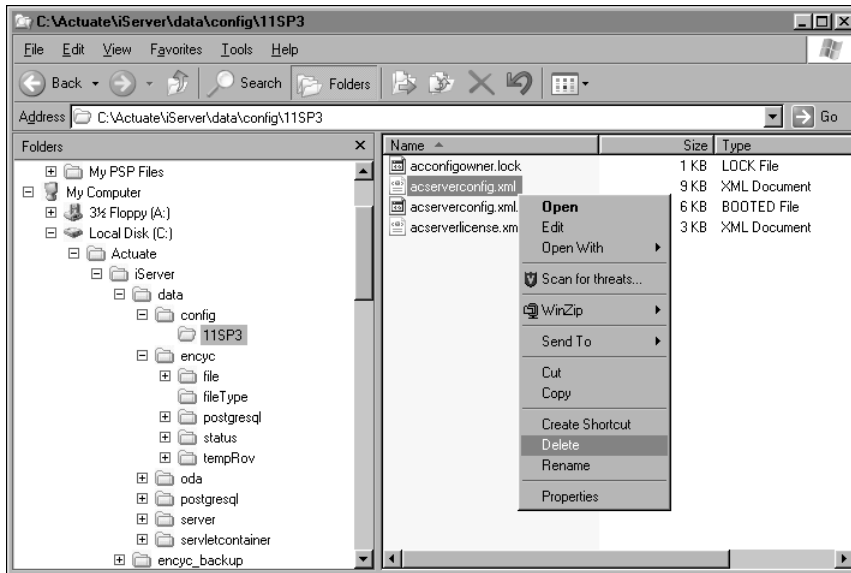


Figure 10-13 Deleting acserverconfig.xml

- 3 In AC_DATA_HOME, open the encyc folder.

In AC_DATA_HOME\encyc, select the file and 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.

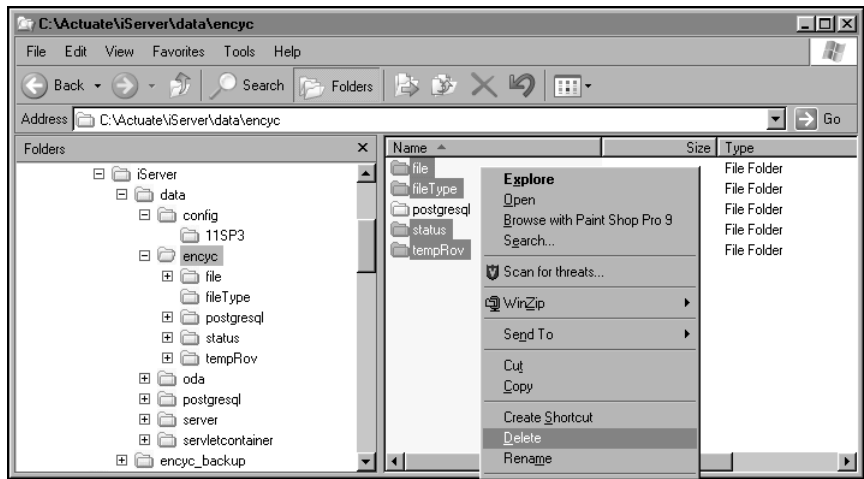


Figure 10-14 Deleting the file, fileType, and status folders from AC_DATA_HOME\encyc

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\11SP3`, as shown in Figure 10-15.

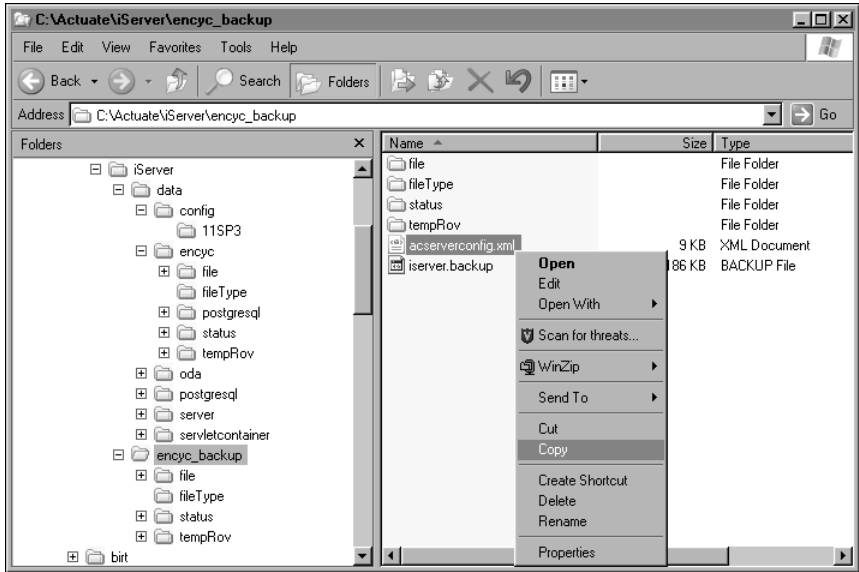


Figure 10-15 Choosing to copy `acserverconfig.xml`

`AC_DATA_HOME\config\11SP3` appears as shown in Figure 10-16.

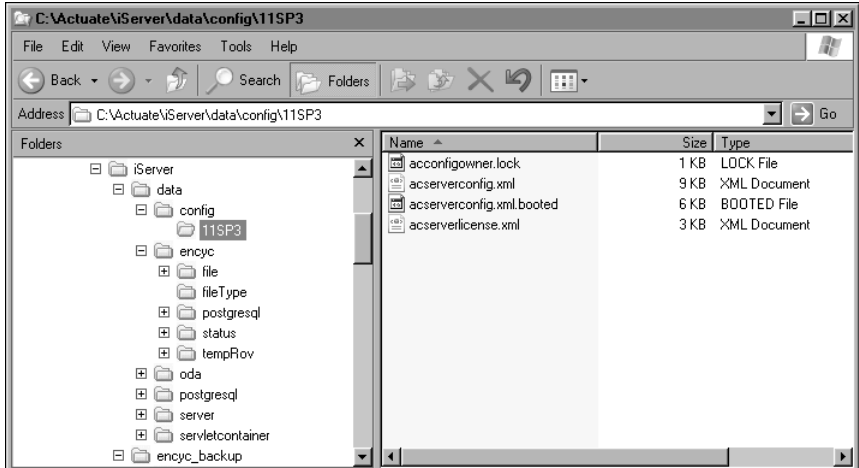


Figure 10-16 Viewing `AC_DATA_HOME\config\11SP3` after restoring `acserverconfig.xml`

- 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, as shown in Figure 10-17.

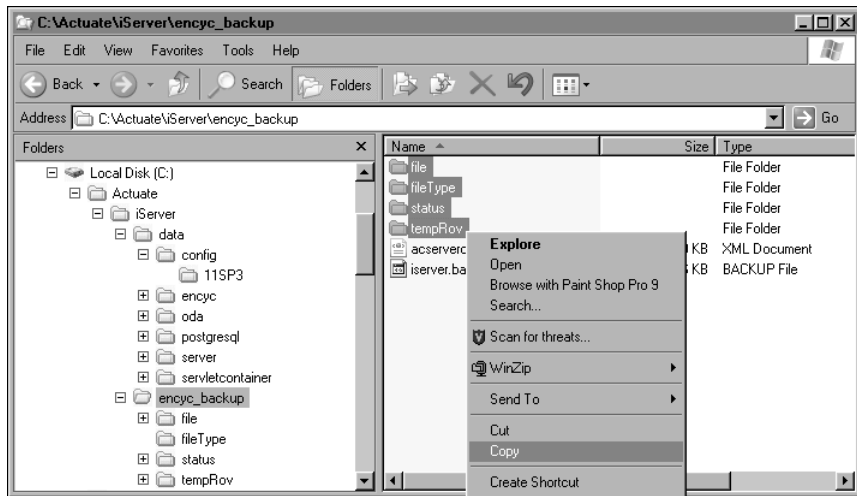


Figure 10-17 Choosing to copy the volume data folders
AC_DATA_HOME\encyc appears as shown in Figure 10-18.

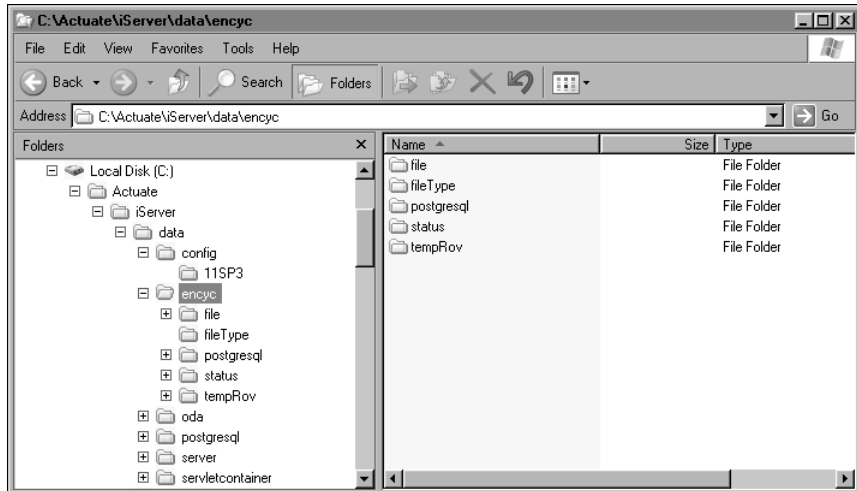


Figure 10-18 AC_DATA_HOME\encyc after copying backed up folders

How to run pg_restore using pgAdmin

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

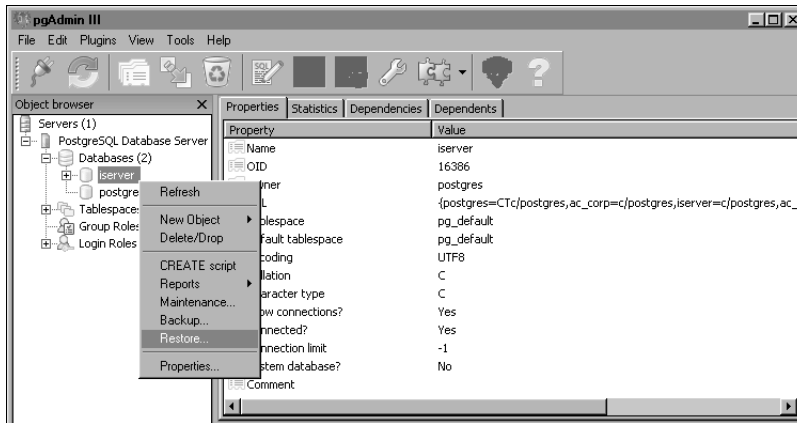


Figure 10-19 Choosing to restore the iserver database from backup

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.

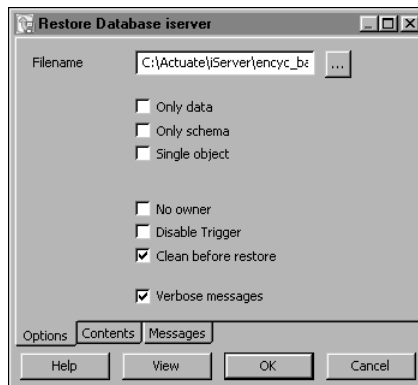


Figure 10-20 Specifying the backup file to restore

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.

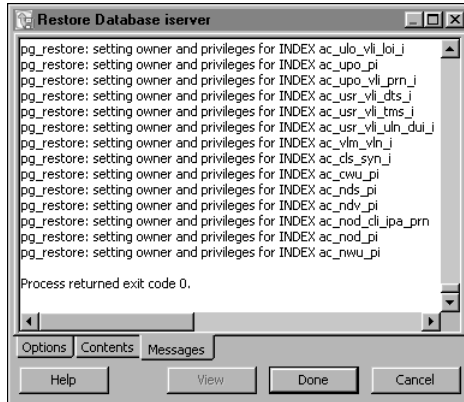


Figure 10-21 Viewing status messages pg_restore writes during execution

- 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.

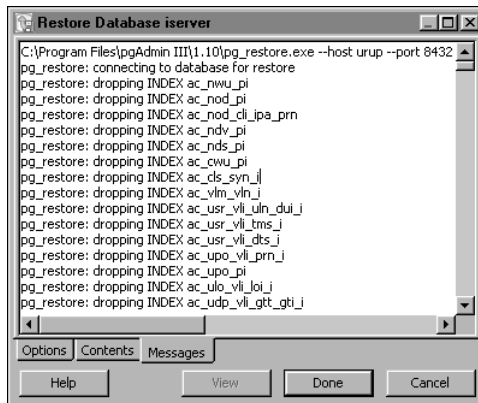


Figure 10-22 Viewing the command that executed pg_restore

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 the Simple view, choose Advanced view. Choose Volumes.
- 3 On Volumes, take the volume online, as shown in Figure 10-23.

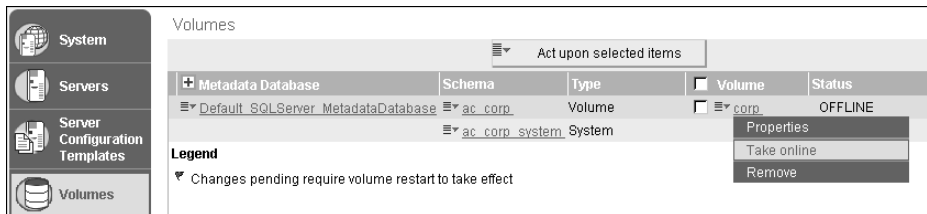


Figure 10-23 Taking the volume online

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
- Backup 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 back up 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 Backup 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`.

- **DUMPFIL**

Specifies the name of the output file, such as `ac_corp_schema.dmp`.

After backing up the Encyclopedia volume metadata, back up the `acserverconfig.xml` file and volume data directories to the backup directory by performing the following tasks.

How to back up the volume data directories

- 1** 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 `acserverconfig.xml` file. In Actuate 11 Service Pack 3, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP3
```

Copy `acserverconfig.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 `acserverconfig.xml` file and volume data directories in `AC_DATA_HOME`
- Copy the backed up `acserverconfig.xml` file and volume data directories from the backup directory to `AC_DATA_HOME`

- Restore the Encyclopedia volume metadata using the Oracle impdp utility
- Take the Encyclopedia volume online

Take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

- 1 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.

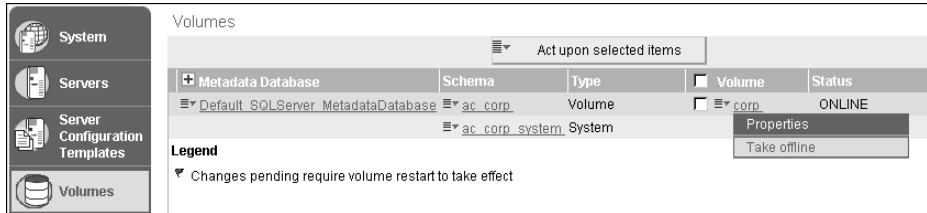


Figure 10-24 Taking the volume offline

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 3, this file is located in the following subfolder:
C:\Actuate\iServer\data\config\11SP3
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\11SP3.
 - 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 back up of the Encyclopedia volume metadata

- 1 In Windows, open a command prompt.
- 2 Restore the Encyclopedia volume metadata using the impdp utility, as shown in the following example:

```
impdp system/password@db_host/dbname.actuate.com
  SCHEMAS=corp
  DIRECTORY=encyc_backup
  DUMPFILE=ac_corp_schema.dmp
```

The Oracle impdp utility runs using arguments similar to the expdp utility.

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

- 1 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 10-25.

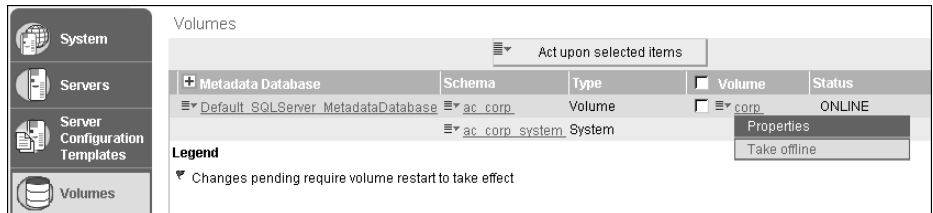


Figure 10-25 Taking the volume online

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.

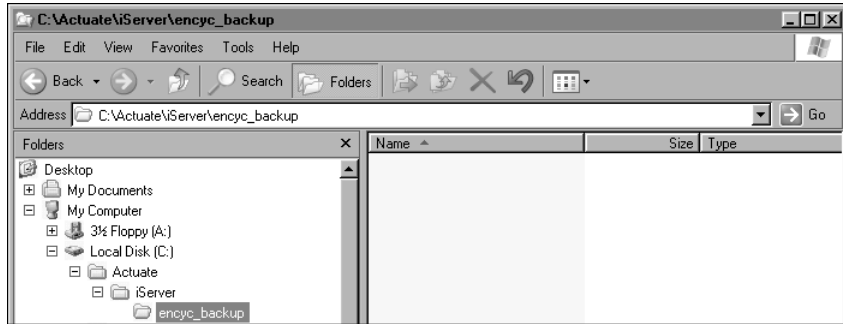


Figure 10-26 Creating a new folder named encyc_backup

Backup 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.

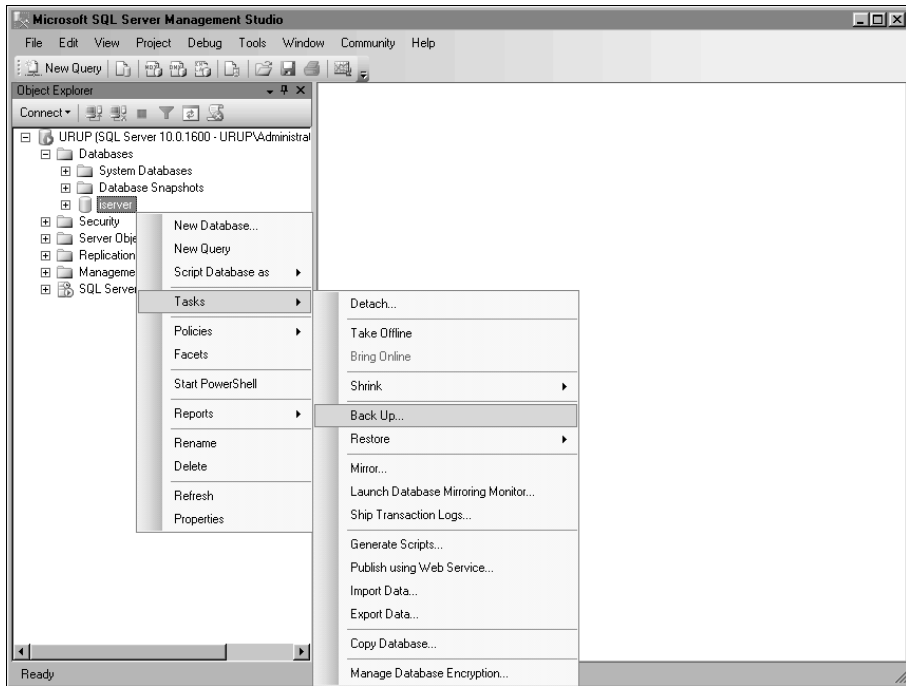


Figure 10-28 Choosing to backup the database named iserver

- 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.

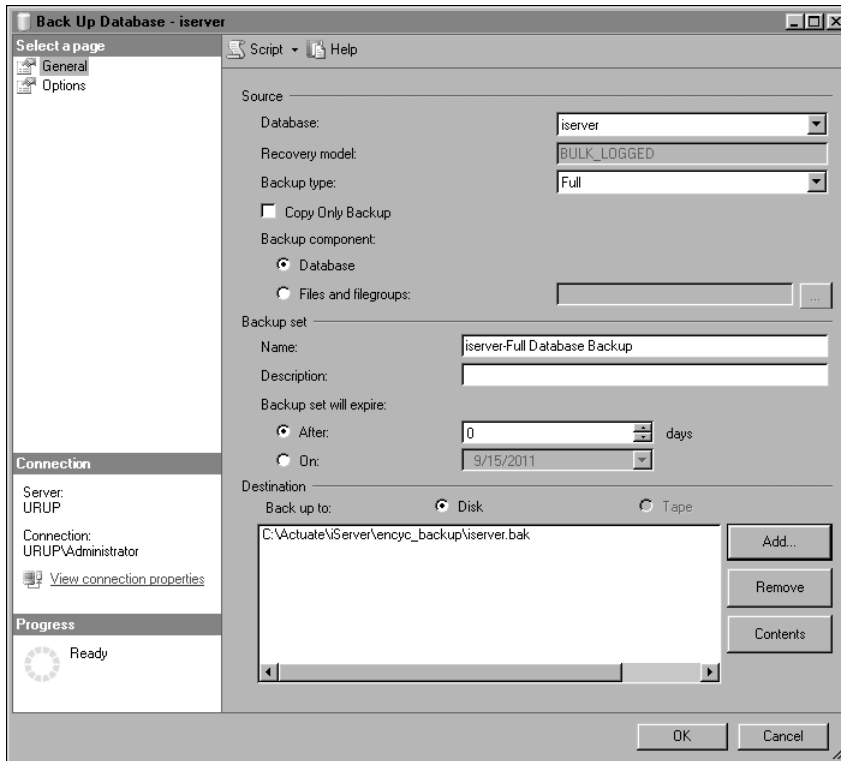


Figure 10-30 Viewing Back Up Database with new backup location

- 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.

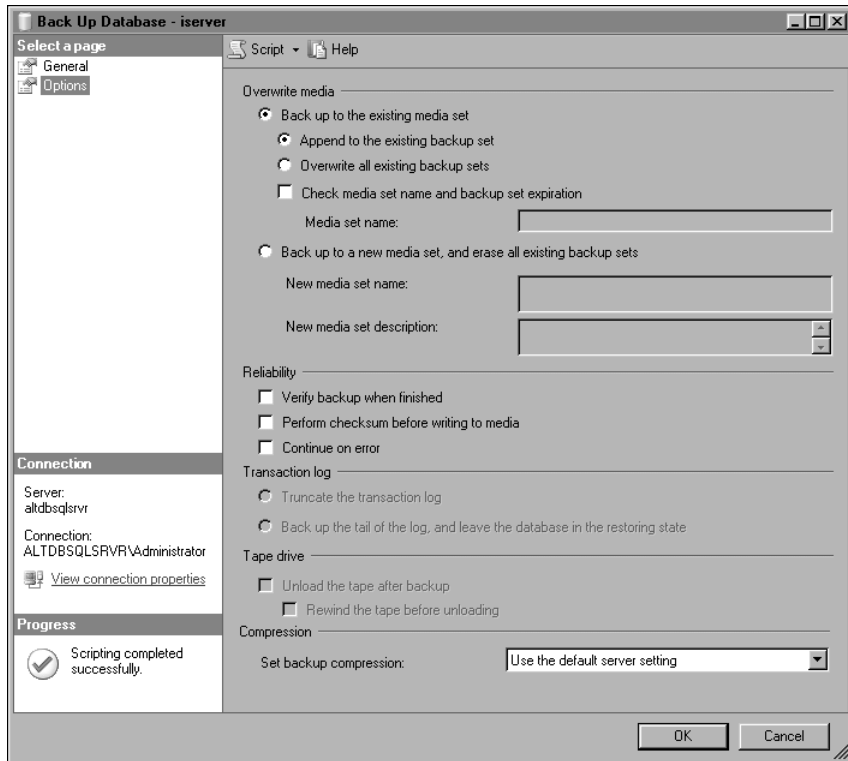


Figure 10-31 Accepting the default values on Back Up Database—Options

- 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 backup 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.




Figure 10-34 Typing the command to backup the metadata

Press enter to execute the command to backup the database. The output appears as shown in Figure 10-35.


```

SQLCMD
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.
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.
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 3, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP3
```

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.

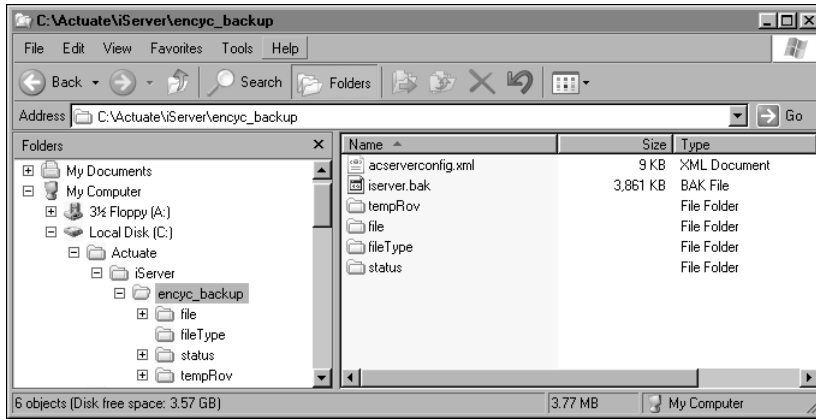


Figure 10-36 Viewing the backed up files

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.

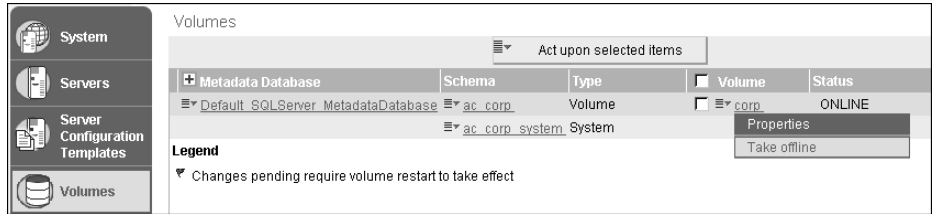


Figure 10-37 Taking the volume offline

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 3, this file is located in the following subfolder:
 C:\Actuate\iServer\data\config\11SP3
 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\11SP3. 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 backup 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.

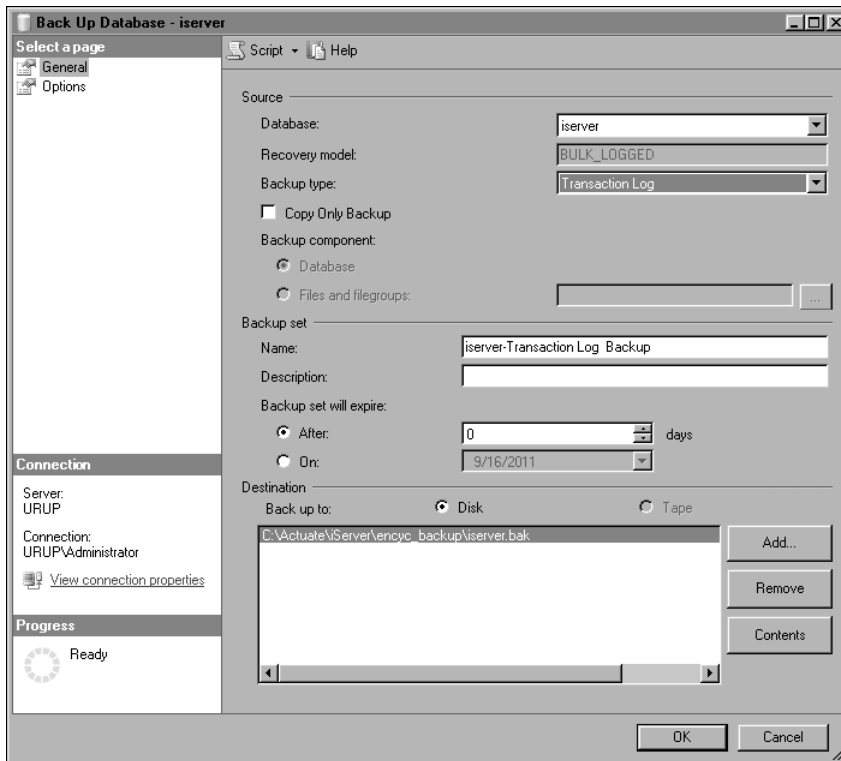


Figure 10-38 Choosing to back up the transaction log

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.

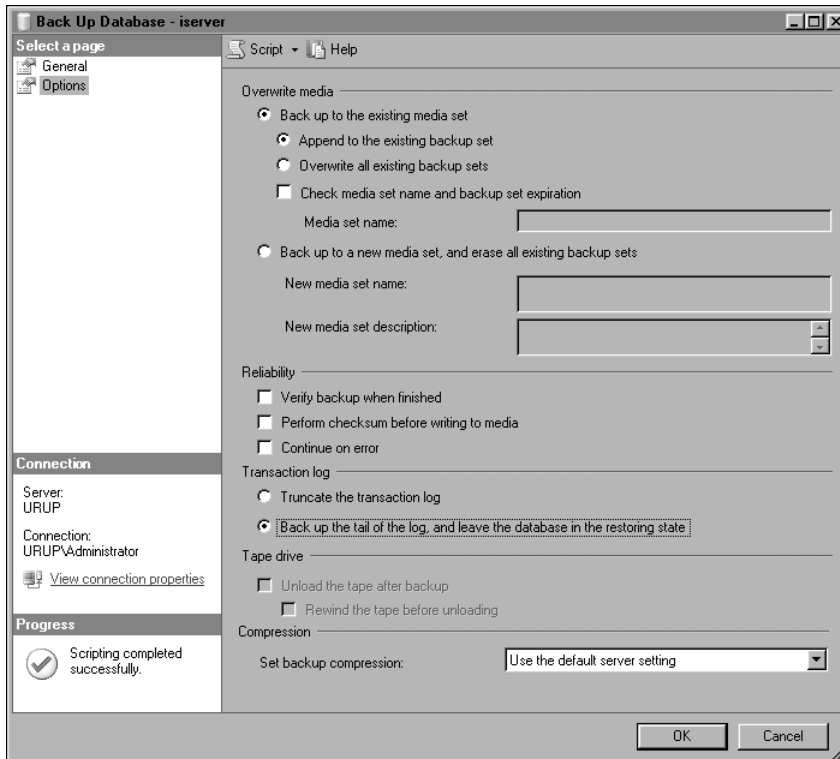


Figure 10-40 Configuring Back Up Database—Options

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 backup 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
```

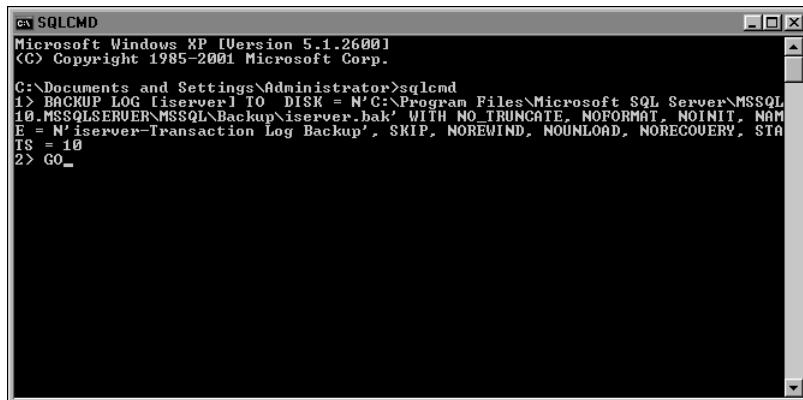
- 3 At the sqlcmd prompt, type the following command and press enter:

```
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, NORECOVERY, STATS = 10
```

At the sqlcmd prompt, type the following command:

```
GO
```

The command prompt appears as shown in Figure 10-42.

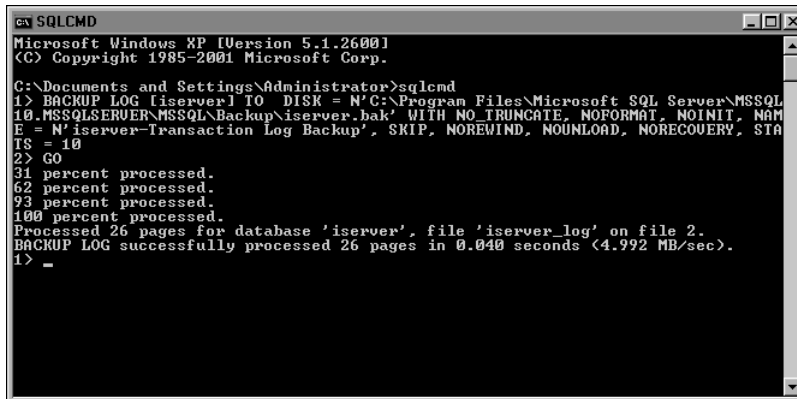


```
SQLCMD
Microsoft Windows XP [Version 5.1.26001
(C) Copyright 1985-2001 Microsoft Corp.

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_
```

Figure 10-42 Typing the command to backup the tail of the log

Press enter to execute the command to backup 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, NAME
 = N'iserver-Transaction Log Backup', SKIP, NOREWIND, NOUNLOAD, NORECOVERY, STAT
IS = 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 back up 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.

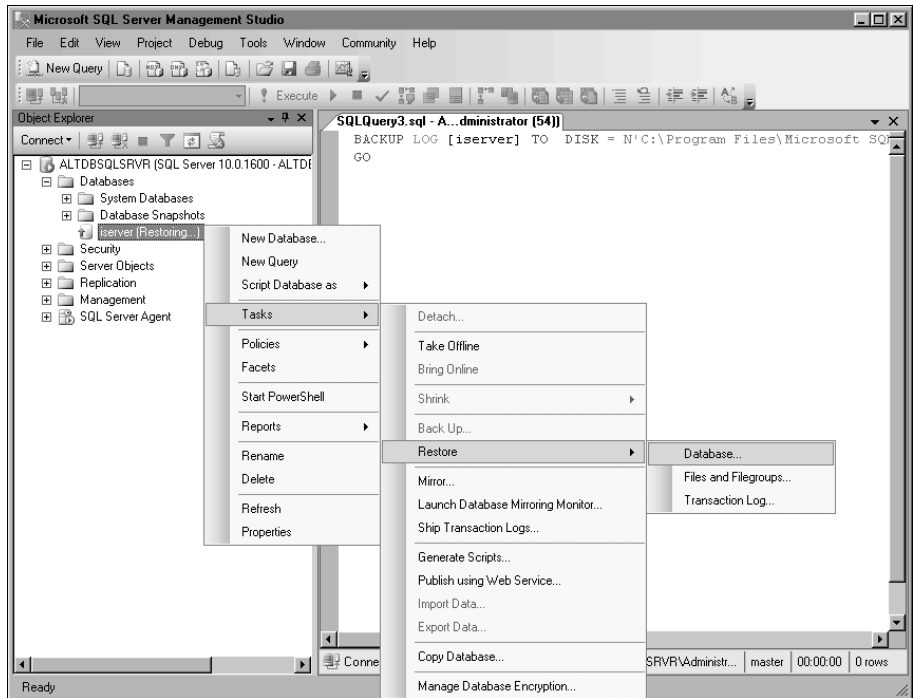


Figure 10-44 Choosing to restore the database

- 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.

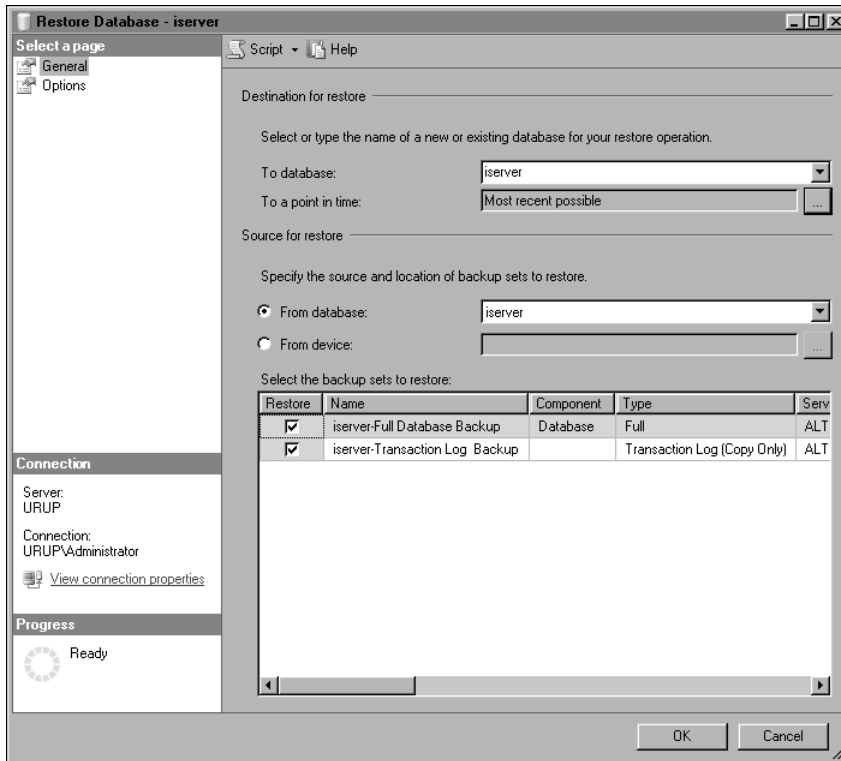


Figure 10-45 Configuring Restore Database—General

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.

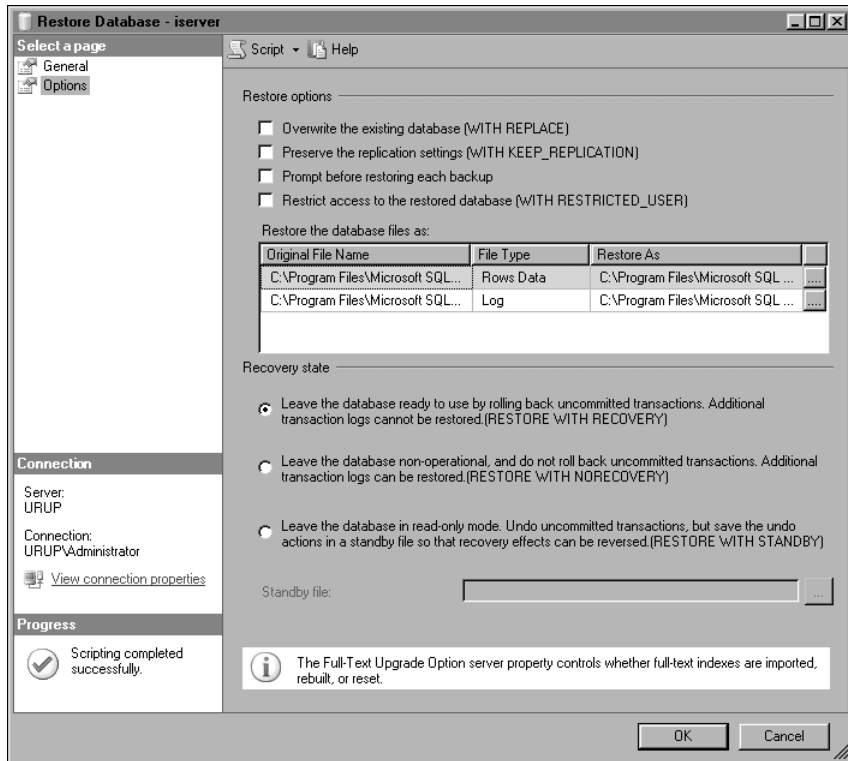


Figure 10-46 Configuring Restore Database—Options

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
```

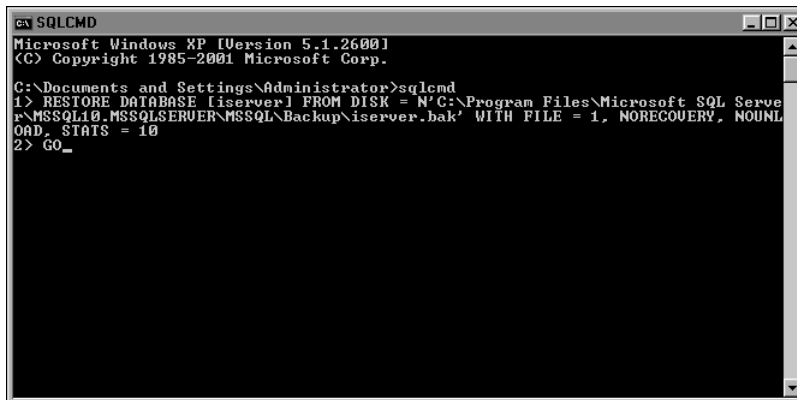
- 3 At the sqlcmd prompt, type the following command and press enter:

```
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

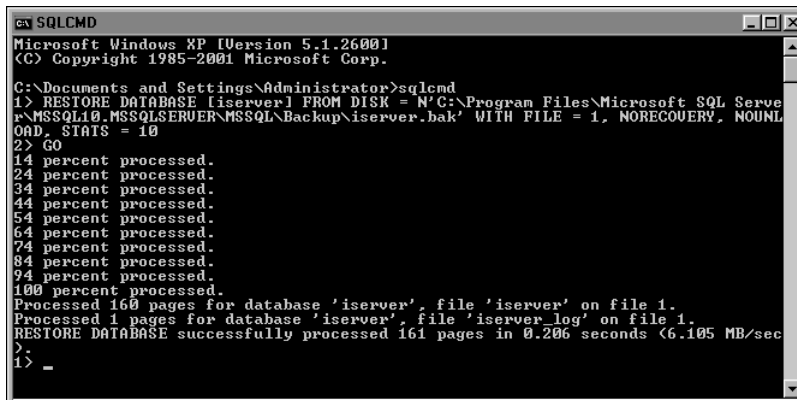


```
SQLCMD
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

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

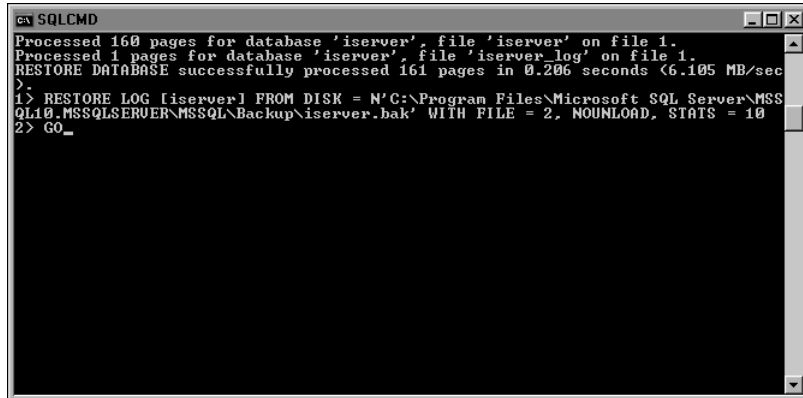
- 4 To restore the tail of the transaction log backup, type the following command at the sqlcmd prompt and press enter:

```
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
```

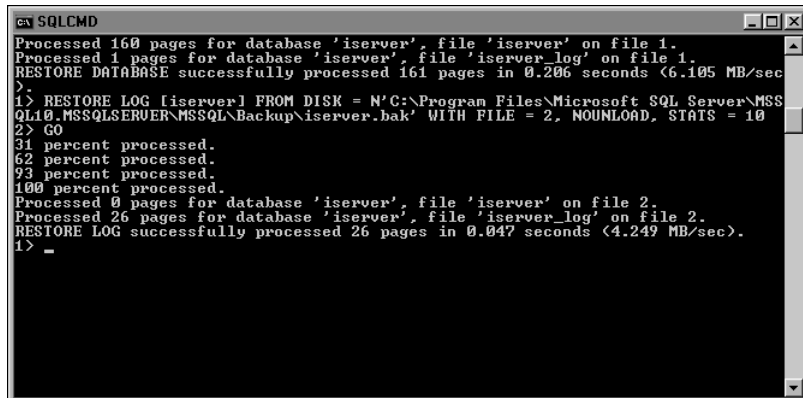
The command prompt appears as shown in Figure 10-50.



```
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.

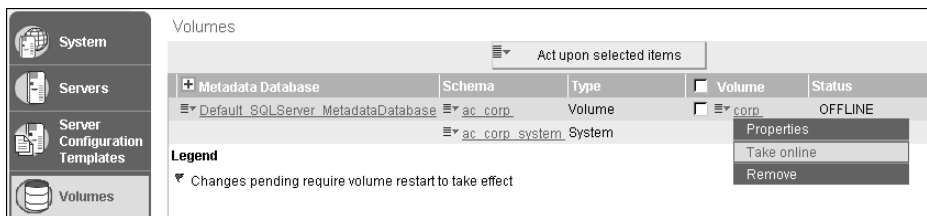


Figure 10-52 Taking the volume online

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

Backup 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.

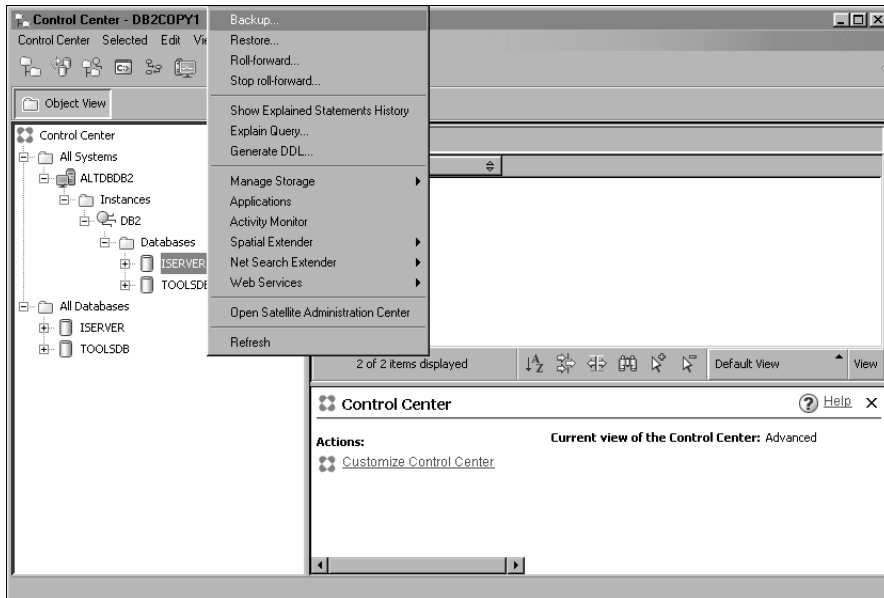


Figure 10-54 Choosing to backup the ISERVER database

- 3 On Backup Wizard—Introduction, select Backup entire database, as shown in Figure 10-55. Choose Next.

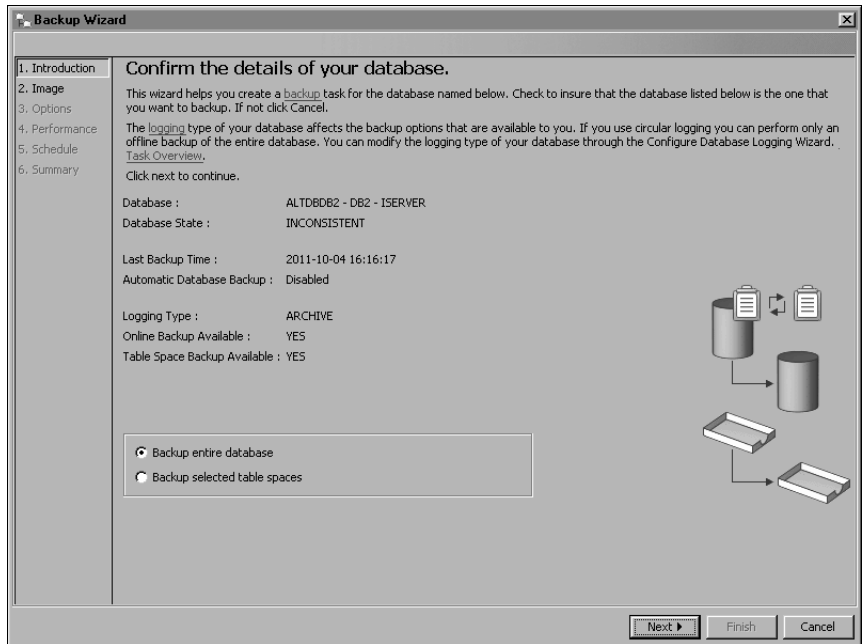


Figure 10-55 Confirming the details of the ISERVER database

- 4 On Backup Wizard—Image, in Media Type, select File System, then select Add, as shown in Figure 10-56.

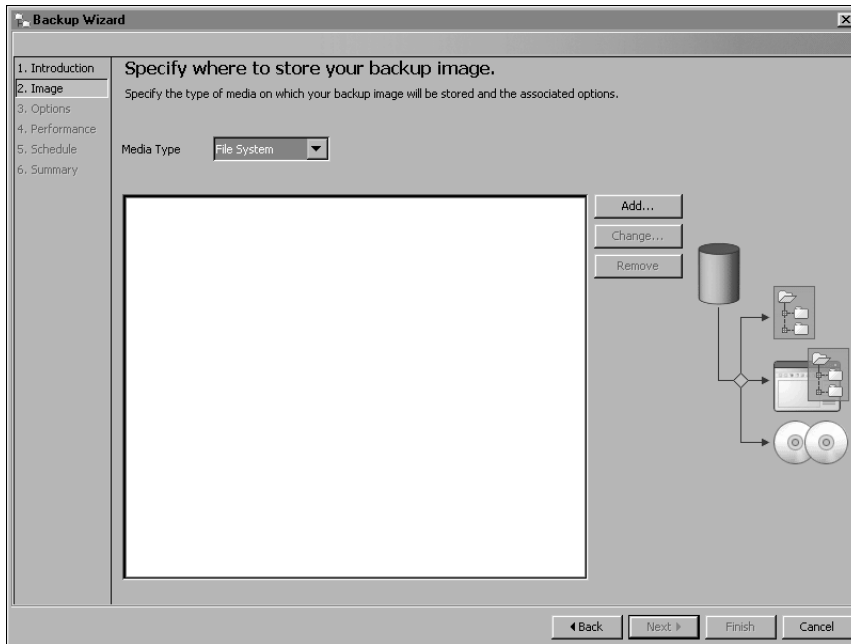


Figure 10-56 Choosing to add a backup location

5 In Path Browser, navigate to the following location, as shown in Figure 10-57.

C:\Actuate\iServer\encyc_backup




Figure 10-57 Navigating the metadata backup location I

Choose OK. Backup Wizard—Image appears as shown in Figure 10-58.

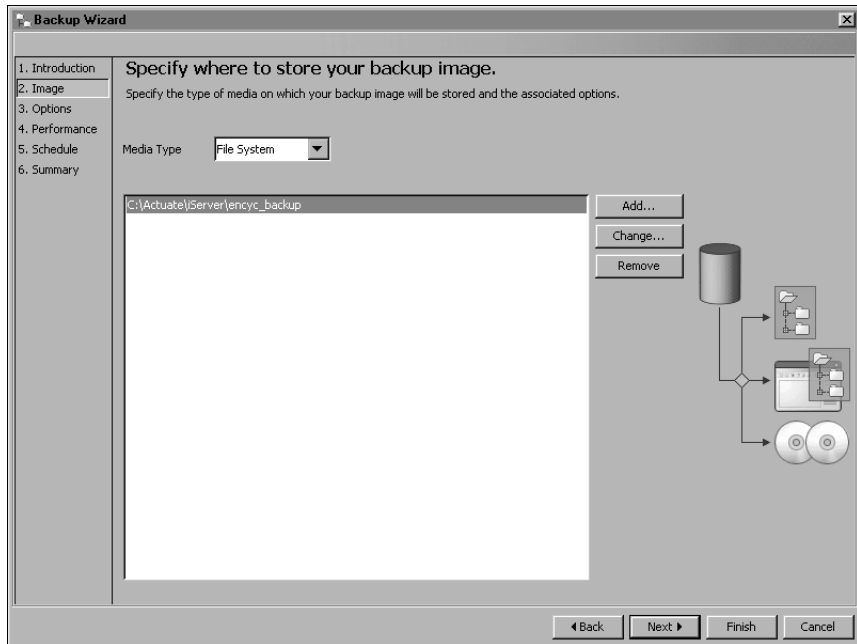


Figure 10-58 Viewing the chosen metadata backup location

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.

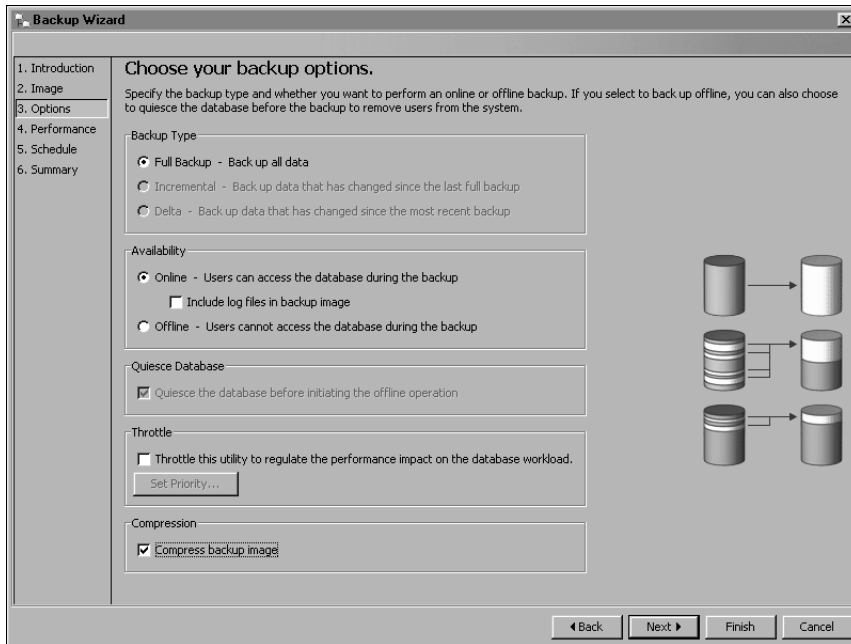


Figure 10-59 Configuring the backup options

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.

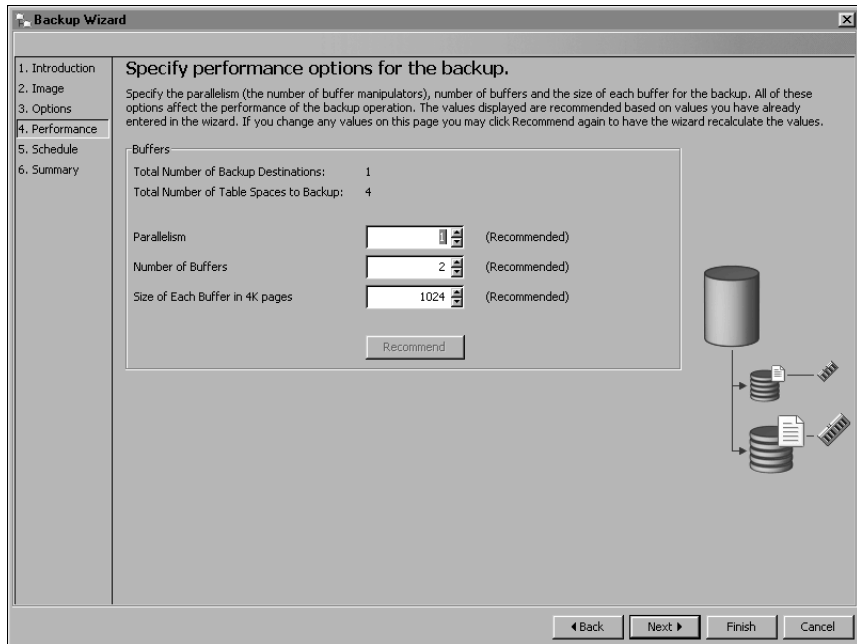


Figure 10-60 Specifying performance options

Choose Next.

- 8 On Backup Wizard—Schedule, accept the default values, as shown in Figure 10-61. Alternatively, specify different values.

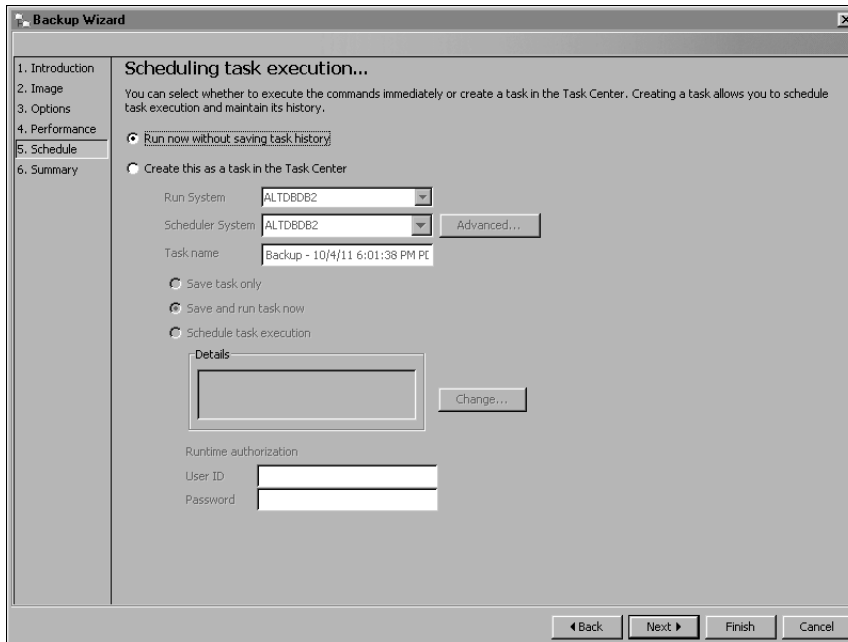


Figure 10-61 Selecting scheduling options

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.

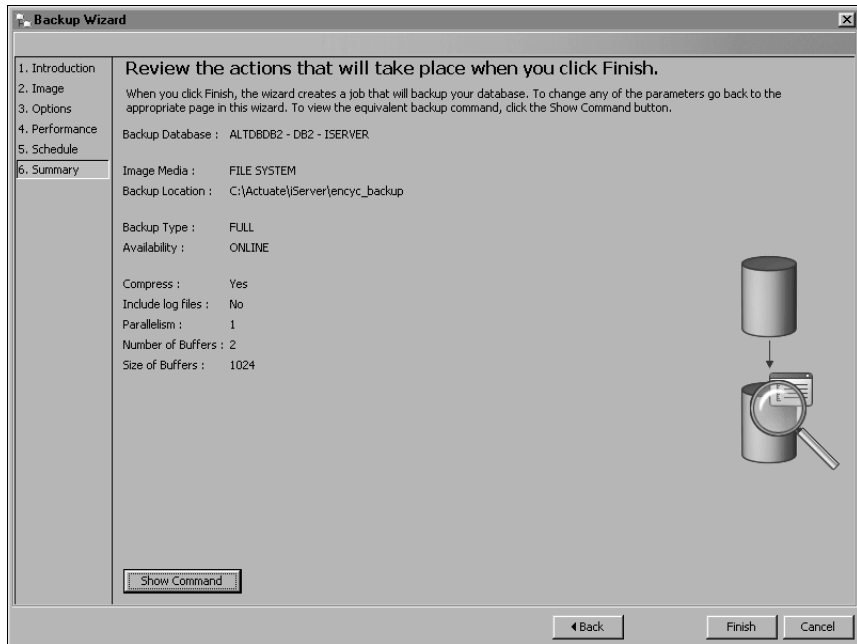


Figure 10-62 Reviewing backup job parameters

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.

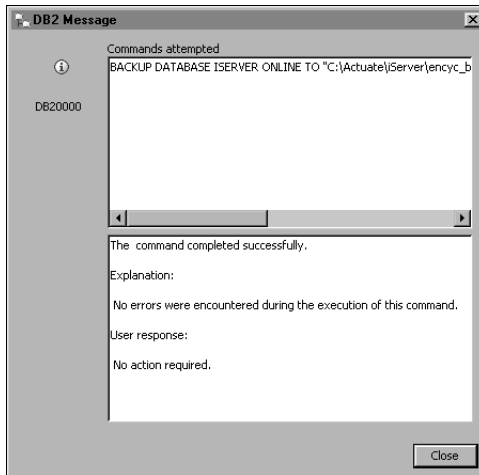


Figure 10-64 Viewing the commands executed and end-of-job messages

Alternatively, you can also backup 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.

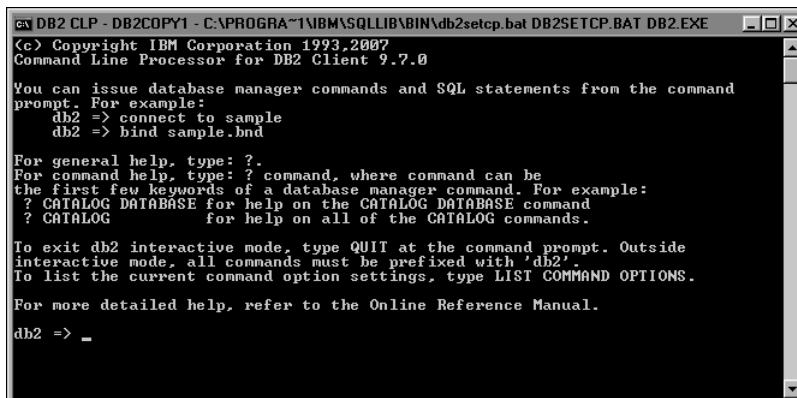
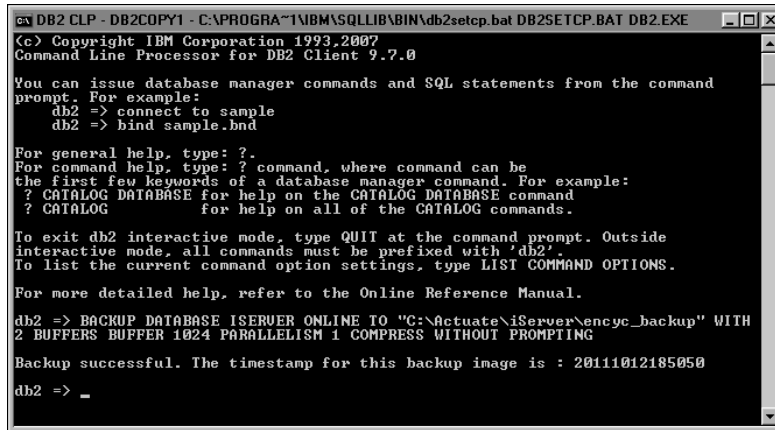


Figure 10-65 Opening the command line processor

- 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
```



The screenshot shows a Windows command prompt window titled "DB2 CLP - DB2COPY1 - C:\PROGRAM FILES\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE". The prompt is "(c) Copyright IBM Corporation 1993.2007 Command Line Processor for DB2 Client 9.7.0". The window contains the following text:

```
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 `acservconfig.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 `acservconfig.xml` file. In Actuate 11 Service Pack 3, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP3
```

Select `acservconfig.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.

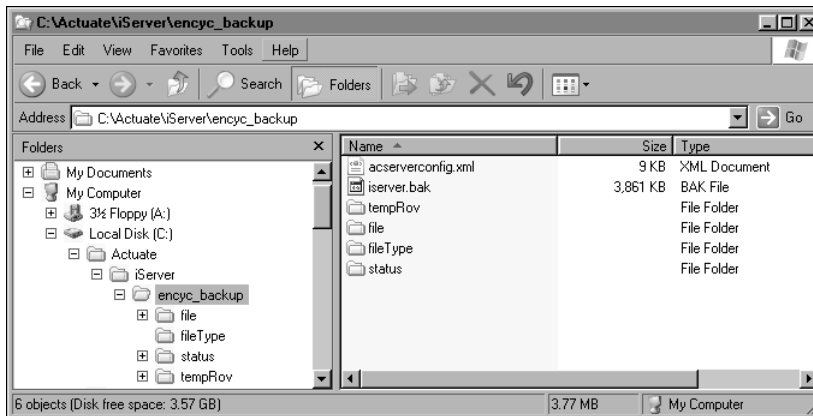


Figure 10-67 Viewing the backed up files

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.

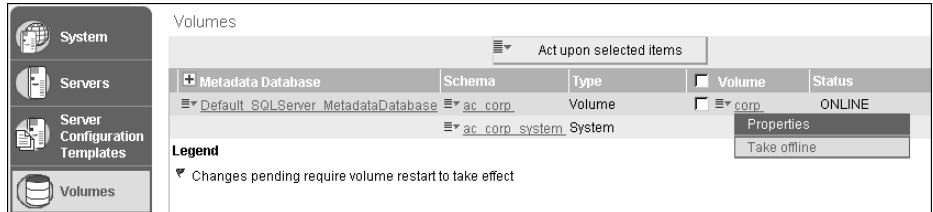


Figure 10-68 Taking the volume offline

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 3, this file is located in the following subfolder:
C:\Actuate\iServer\data\config\11SP3
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.

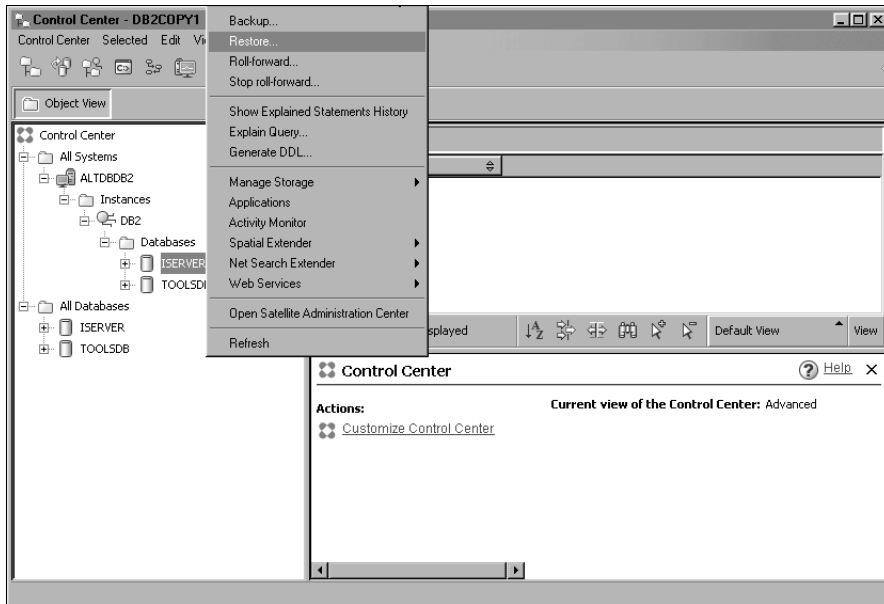


Figure 10-69 Choosing to restore the database

- 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.

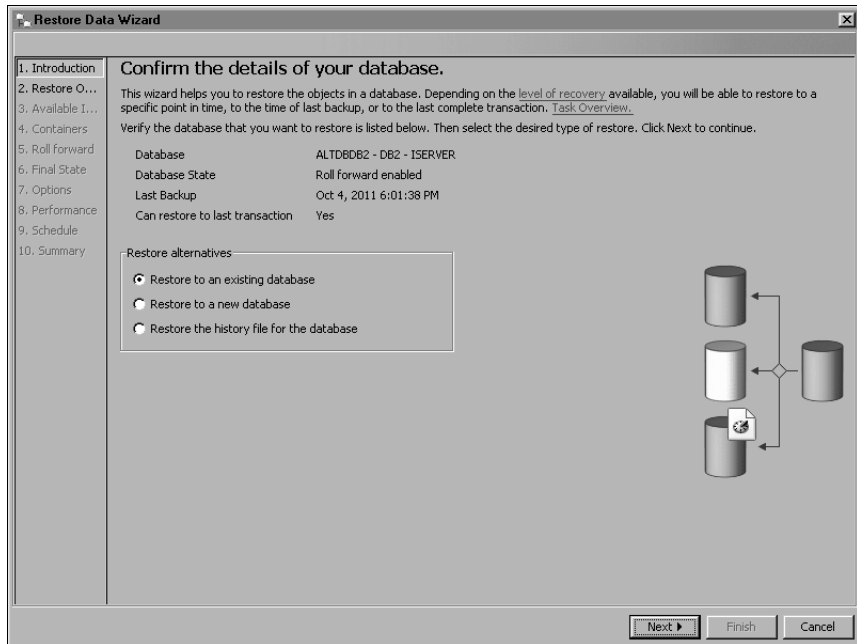


Figure 10-70 Confirming the details of the ISERVER database

- 3 On Restore Data Wizard—Restore Objects, accept the default selection, Restore the entire database, as shown in Figure 10-71. Choose Next.

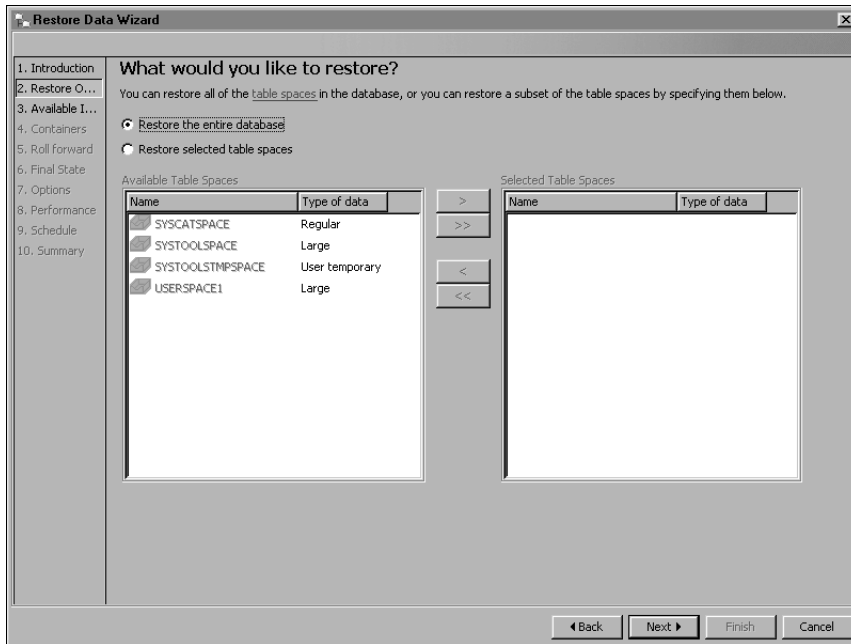


Figure 10-71 Selecting the backup image to restore

- 4** On Restore Data Wizard—Available Images, in Available backup images, select the image you want to restore, as shown in

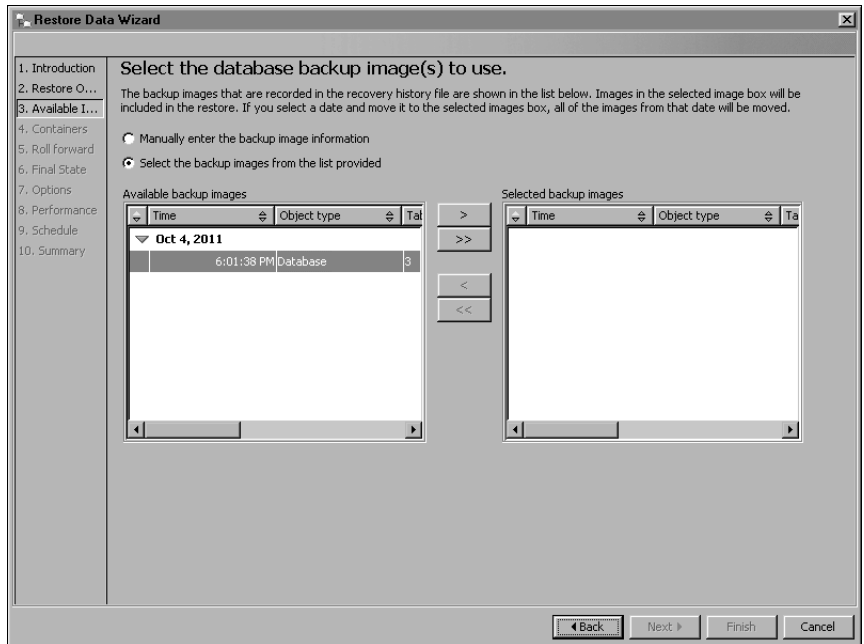


Figure 10-72 Selecting the backup image to restore

Choose the right arrow to move the image to Selected backup images, as shown in Figure 10-73. Choose Next

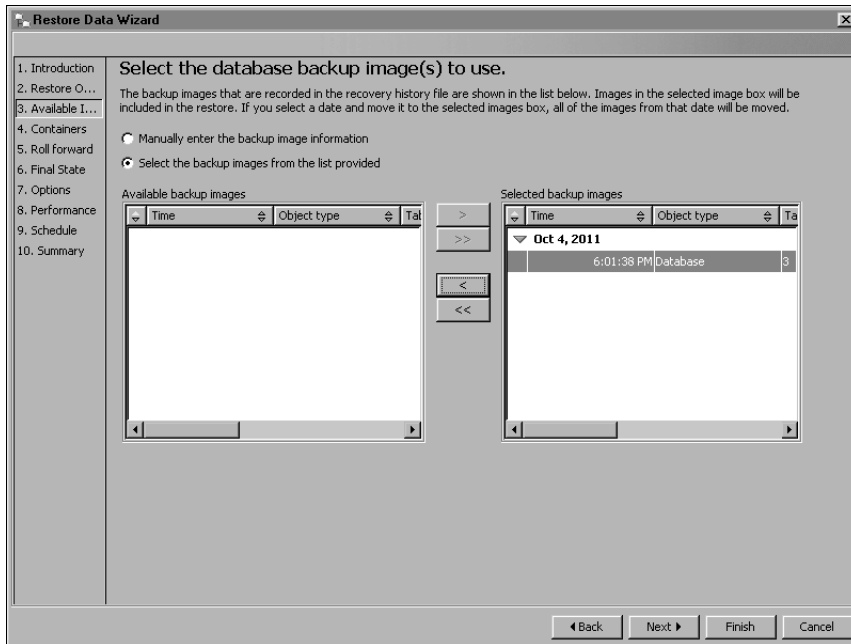


Figure 10-73 Moving the selected image to Selected backup images

- 5** On Restore Data Wizard—Containers, accept the default options, as shown in Figure 10-74. Choose Next.

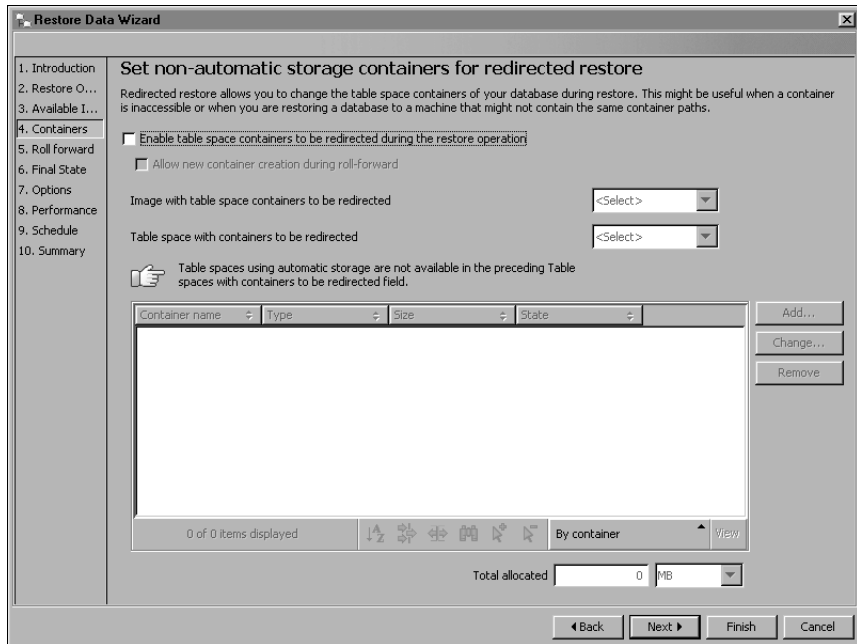


Figure 10-74 Choosing whether to enable redirecting of table space containers

- 6** On Restore Data Wizard—Roll forward, accept the default option, Restore Only, as shown in . Choose Next.

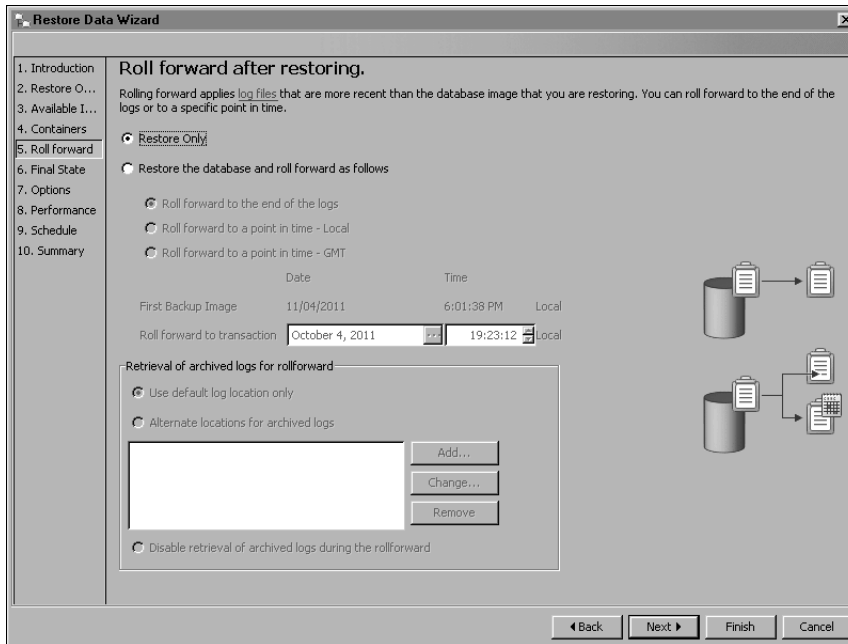


Figure 10-75 Choosing to restore the database backup only

- 7 On Restore Data Wizard—Final State, select Complete the restore and return to the active state, as shown in . Choose Next.

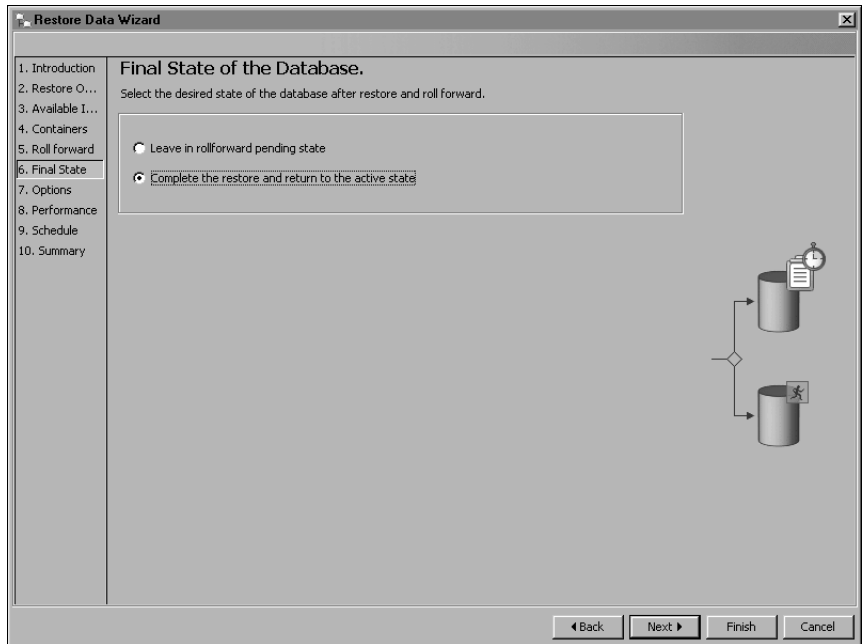


Figure 10-76 Choosing to complete the restore

- 8 On Restore Data Wizard—Options, accept the default options, as shown in Figure 10-77. Choose Next.

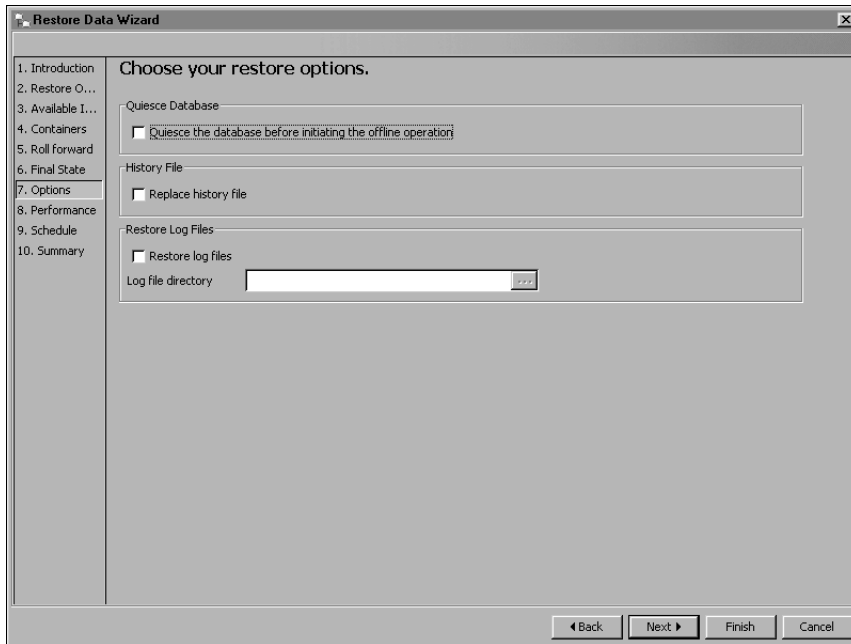


Figure 10-77 Choosing restore options

- 9 On Restore Data Wizard—Performance, accept the default values, as shown in Figure 10-78. Choose Next.

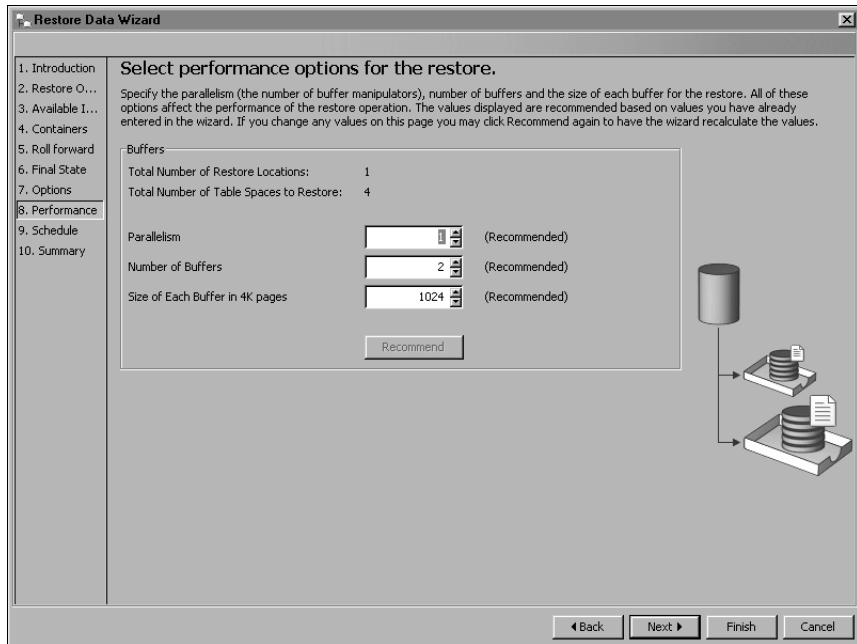


Figure 10-78 Selecting performance options

- 10 On Restore Data Wizard—Schedule, accept the default values, as shown in Figure 10-79. Choose Next.

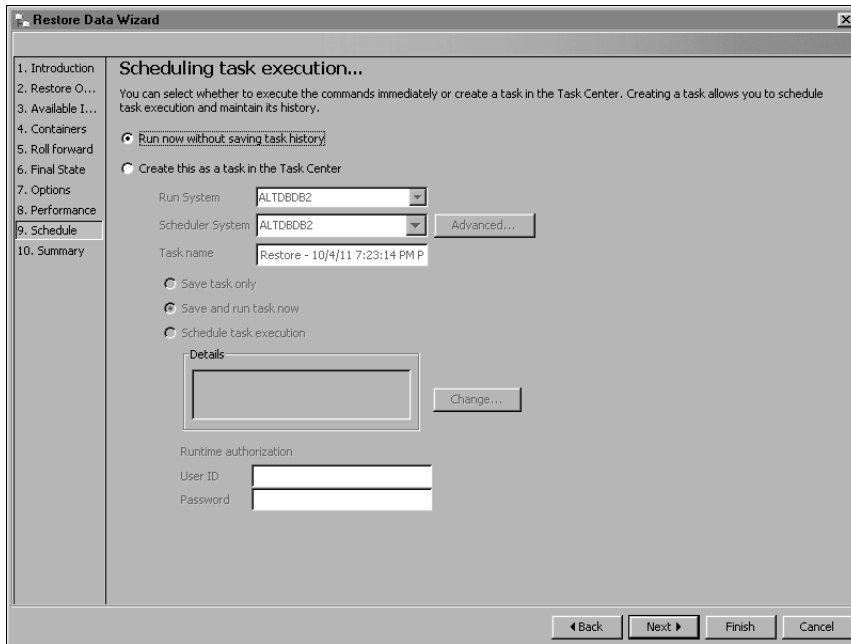


Figure 10-79 Setting backup job scheduling options

- 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.

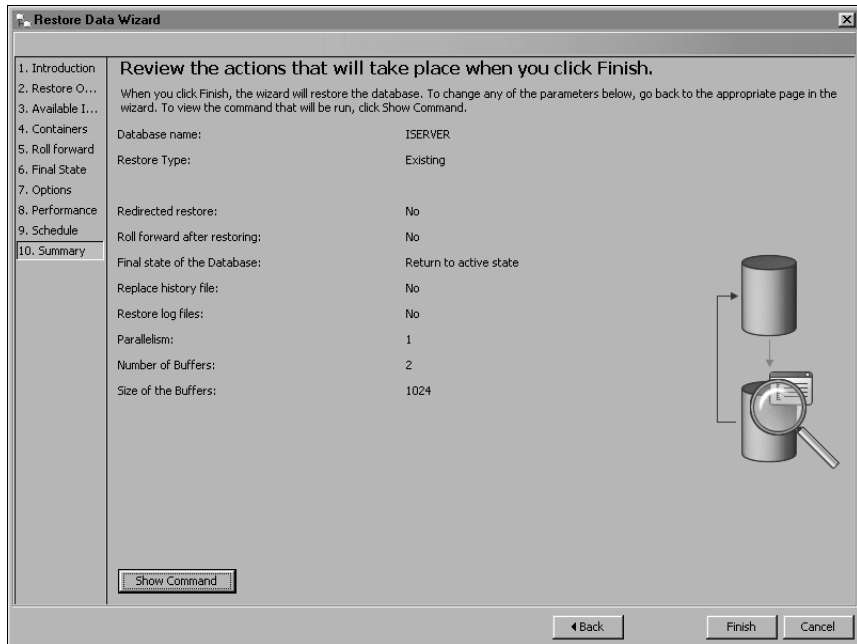


Figure 10-80 Reviewing restore job parameters

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.

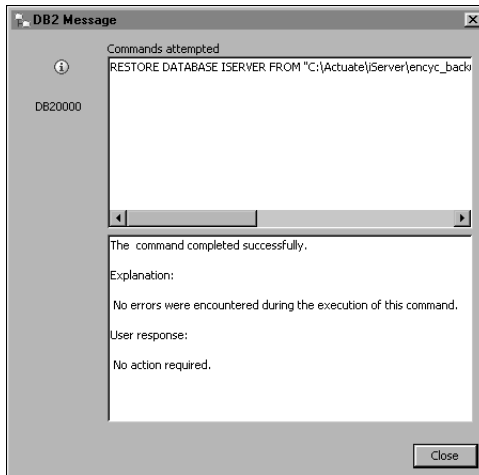


Figure 10-82 Viewing the command executed and end-of-job messages

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.

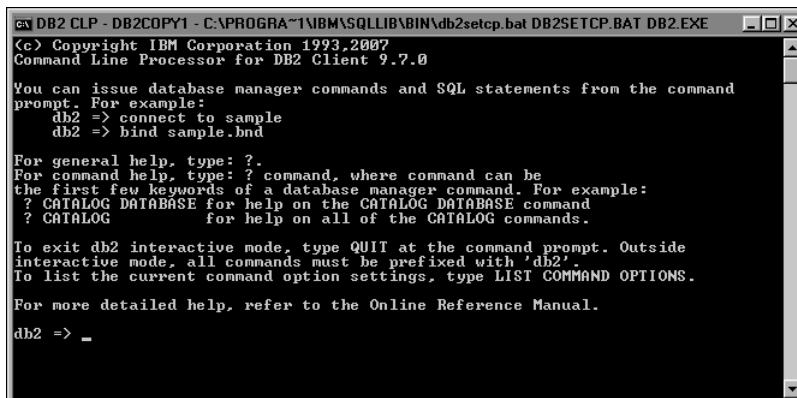


Figure 10-83 Opening the command line processor

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.

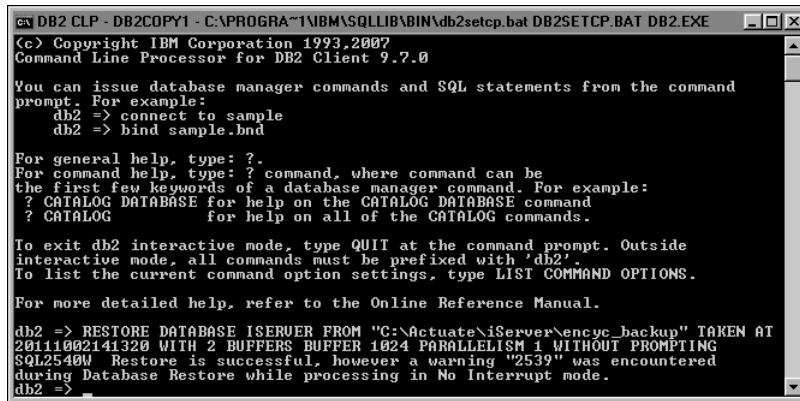


Figure 10-84 Restoring the database from the command line

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.

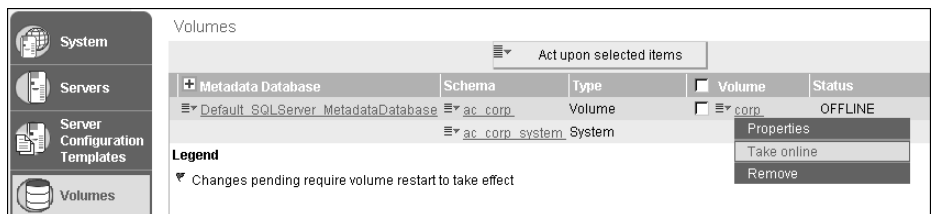


Figure 10-85 Taking the volume online

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_DATA_HOME directory 190
- AC_DATA_HOME variable 29, 116, 190
- AC_JAVA_HOME variable 46, 186
- AC_JRE_HOME variable 46
- AC_JRE64_HOME variable 46
- AC_JVM_HOME variable 46
- AC_PMD_WINDOWS_CPUS variable 231, 232
- AC_SERVER_HOME parameter 142, 144
- AC_SERVER_HOME variable 29, 116, 135, 186, 189
- AC_TEMPLATE_NAME variable 177
- access permissions. *See* privileges
- accessing
 - data 13
 - Encyclopedia system database 18
 - Encyclopedia volumes 47, 222, 224
 - help files x, xi
 - help topics xi
 - iServer features 46, 222
 - iServer System components 41
 - online help 215
 - pgAdmin administration tool 42, 98
 - PostgreSQL servers 41
 - resources 202
 - Update tool 212
 - Visual C++ libraries 185
 - web-based content 223, 224
- accounts
 - cluster configurations and 25
 - connecting to databases and 54, 57, 59, 61
 - installing alternate databases and 53
 - installing cluster nodes and 167
 - installing Information Console and 198
 - iServer 11 service and 34
 - iServer installations and 24, 25
 - Management Console and 34
 - managing 16
 - PostgreSQL database and 31
 - setting up e-mail 236
 - upgrades and 78, 88
 - AcEncycUpgrade utility 18
 - AcExport utility 16
 - AcExtern utility 17
 - AcImport utility 17
 - AcIntern utility 17
 - acmachineid utility 228
 - AcMode utility 18
 - acpmdconfig.xml 13, 160, 176
 - Acrobat Catalog. *See* Adobe Acrobat Catalog
 - acserverconfig.xml 13, 160
 - activity logs 203
 - AcToc utility 17
 - Actuate Analytics Option 223
 - Actuate Basic reports 224
 - Actuate Customer Support xi
 - Actuate licensing options 222, 227
 - Actuate Query Option 223
 - Actuate Support Lifecycle Policy xii
 - actuate.com vii
 - ActuateBIRTiServer.exe 82, 101, 114
 - ActuateBIRTiServer.zip 186
 - ActuateBIRTiServerEnterprise.exe 27, 162
 - ActuateInformationConsole.exe 195
 - ActuateInformationConsole.war 201
 - ActuateiServerIntegrationTechnology.exe 208
 - ActuateLocalizationandOnlineDocumentatio
n.exe 212
 - ActuateOne technology vii
 - AcVerify utility 17
 - Add Partition link 134
 - Add Volume link 140
 - Add volume schema setting 135
- adding
 - administrator accounts 24, 78
 - backup folders 244, 266, 287
 - cluster nodes 13, 160, 175
 - desktop shortcuts 41, 96, 113, 187
 - e-mail accounts 236
 - Encyclopedia volumes 27, 140, 156, 157, 224
 - indexes 5
 - license keys 228

- licensing options 227
 - passwords. *See* passwords
 - PostgreSQL servers 42
 - user accounts 24, 25
 - users 52, 53, 54, 57, 59, 61
 - administration tools 16
 - administrative reports 19
 - administrator accounts 24, 34, 78
 - administrators
 - backing up Encyclopedia and 241
 - customizing third-party databases and 18
 - deploying iServer images and 184
 - installing alternate databases and 52
 - installing iServer and 24, 45
 - managing Encyclopedia and 7
 - managing iServer System and 15
 - migrating to current release and 47, 49
 - obtaining licenses and 222
 - optimizing iServer System and 13, 160
 - preventing data loss and 7, 26
 - storing user information and 5
 - upgrading Encyclopedia and 133
 - upgrading iServer and 78, 79, 113
 - Adobe Acrobat Catalog utility xi
 - aggregation 223
 - Analytics Option 223
 - analyzing
 - data 223
 - search results 224
 - Apache Tomcat servers 195, 197
 - application pages. *See* web pages
 - application programming interfaces (APIs) 19
 - application server cache 204
 - application server configurations 195, 202
 - application servers 195, 201, 204
 - See also* servers
 - APPLICATION_USER_NAME
 - parameter 144
 - APPLICATION_USER_PASSWORD
 - parameter 145
 - applications
 - See also* specific iServer application
 - accessing Encyclopedia and 47
 - backward compatibility with 19
 - creating production environment for 48
 - creating test environments for 48
 - deploying web 204
 - developing 15
 - installing. *See* installation
 - restricting CPU processes for 229
 - running iServer processes and 8
 - upgrades and 82
 - archives (cloud deployments) 6, 184
 - archiving
 - iServer folders 187
 - report files 242
 - Archiving and Purging page 242
 - asynchronous Factory service 11
 - autoarchive file purging 242
 - automated installation option 6
 - automatic rollback 79
 - automatic upgrades 78, 80, 82
- ## B
- backing up
 - data files 241
 - database schemas 7, 26, 249
 - Encyclopedia volumes 240–243, 244, 266, 287
 - folders 250, 273, 297
 - metadata 26, 62, 241
 - PostgreSQL databases 243
 - program files 82, 101
 - report files 17
 - system schemas 79, 150
 - volume databases 17
 - backup utilities 241
 - backward compatibility 19
 - Basic reports 224
 - batch files 138, 140
 - bin directory 189
 - binary files 29, 116, 189
 - BIRT 360 Option 223
 - BIRT Data Analyzer Option 223
 - BIRT Designer Professional 223
 - BIRT Exchange URL 27
 - BIRT Interactive Viewer Option 223
 - BIRT iServer. *See* iServer
 - BIRT open source projects vii
 - BIRT Option 223
 - BIRT Page Level Security option 223
 - BIRT reports. *See* reports

- BIRT Spreadsheet Designer 223
- BIRT Spreadsheet Option 223
- BIRT Studio 202
- BIRT Studio Option 223
- BIRT Viewer 203
- BIRT_RESOURCE_PATH parameter 202
- BIRT_VIEWER_LOG_DIR parameter 203
- birt-exchange.com vii
- browsers. *See* web browsers
- browsing Encyclopedia volume database 41, 98
- bulk load script files 154
- bulk loading utilities 17
- Business Intelligence technology vii

C

- C++ libraries 185
- C++ runtime components 187
- cache (data) 224
- cache (web pages) 204
- cache conflicts 204
- Caching processes 233
- Caching service 9, 11
- CachingService element 233
- changing
 - cluster machines 229
 - CPU binding 234, 236
 - database schemas 98
 - databases 54, 57, 59, 61
 - default locales 202
 - license file names 229
 - licensing options 227
 - network cards 229
 - parameters 187
 - passwords 34
 - port numbers 31
 - time zones 203
 - volume metadata 5
 - volume names 153
- character sets 19
- charts 46, 223
- CLASSPATH variable 47
- clearing web browser cache 204
- client applications 8, 15
- client/server models 13
- cloud deployments 6, 184–191
- cluster IDs 13, 160
- cluster nodes 8, 13, 29, 160, 162, 175
- clusters
 - adding nodes to 13, 160, 175
 - associating with machine IDs 225
 - binding iServer processes to 230
 - changing machines for 229
 - connecting to iServer System and 8
 - creating iServer images and 184
 - deploying web applications and 204
 - determining number of processors for 234
 - exceeding CPU licenses for 235
 - failing 13, 160
 - installing cluster nodes for 162–174
 - installing iServer for 29
 - licensing options for 225, 226
 - load balancing for 14, 160
 - running iServer processes and 13, 14
 - running iServer services and 8
 - setting configuration home location for 164
 - setting port numbers for 168
 - setting up Encyclopedia database for 13, 14
 - setting up user accounts for 25
 - storing information for 184
 - storing volume metadata and 5
- collecting machine information 227
- command-line options 16
- command-line utilities 16, 243
- comments 157
- communications protocols (clusters) 175
- config directory 165, 176
- CONFIG_SCHEMA_NAME parameter 145
- ConfigHomeDirectory variable 225, 229
- Configuration Console
 - accessing documentation for 16
 - adding cluster nodes and 180
 - archiving report files and 242
 - cluster configurations and 14
 - creating Encyclopedia and 140, 156, 157
 - creating system schemas and 135, 150
 - installing 164, 194
 - managing iServer and 188
 - sending notifications and 236
 - setting user names and passwords for 34
 - shutting down iServer and 137

- stand-alone configurations and 9
- updating license keys and 229
- upgrades and 89, 133, 134
- viewing licensing information and 224
- configuration files 82, 157
- configuration home locations 164
- configuration parameters 13, 160, 202
 - See also* parameters
- configuration templates 13, 160
- configurations
 - accessing security database and 47
 - binding CPU processes and 229, 232
 - creating dedicated user accounts and 25
 - deploying Information Console and 195, 201, 203
 - deploying iServer and 186
 - installing cluster nodes and 164, 175, 179, 180
 - installing Information Console and 195, 202
 - installing iServer and 26, 29
 - licensing iServer and 222, 225
 - managing external volume databases and 7
 - master cluster nodes and 8
 - overwriting previous installations and 45
 - purging report files and 242
 - renaming licensing files and 229
 - testing installations and 48
 - upgrading and 82, 157
- configuring
 - application servers 195, 202
 - iServer 9
 - iServer System 6
 - LDAP servers 47
 - service privileges 26
 - system data store 150, 152
 - third-party databases 5, 9, 13, 14
 - upgrade utilities 157
- connection parameters 19
 - See also* parameters
- connection pooling 19, 53
- connections
 - accessing volume database and 18, 53
 - cluster configurations and 8
 - configuring Information Console 203
 - creating user accounts and 54, 57, 59, 61
 - installing cluster nodes and 175
 - running queries and 13
 - upgrading PostgreSQL databases and 86
- context paths 197, 202
- context-sensitive help 212
- copying
 - files 241
 - license files 229
- corrupted schemas 79
- CPU binding 226, 229–236
- CPU binding validation 235, 236
- CPU-based licenses
 - combining named-user licenses with 222
 - determining number of 226, 229
 - exceeding number of 234, 236
 - viewing information about 235
- CPUs
 - binding to multiple core 232
 - deploying iServer over multi-threaded 230
 - determining number of 232, 234
 - hyperthreading and 233
 - licensing options for 222, 229
 - measuring machine capacity for 222
 - restricting processes for 229
 - running encycsrvr11 processes and 234
 - running Information Console and 202
 - viewing maximum number of 235
 - viewing processor IDs for 230
- CREATE_SCHEMA parameter 145
- creating
 - administrator accounts 24, 78
 - backup folders 244, 266, 287
 - cluster nodes 13, 160, 175
 - data cubes 223
 - database schema owners 53, 54, 56, 58, 61
 - database schemas 52, 53, 55, 59, 60, 61, 62, 135
 - e-mail accounts 236
 - Encyclopedia volume partitions 134
 - Encyclopedia volumes 27, 140, 156, 157
 - indexes 5
 - iServer images 184
 - passwords. *See* passwords
 - production environments 48
 - production staging areas 48
 - report designs 223

- system data schemas 150, 152
- system databases 11, 53, 57
- test environments 47
- user accounts 24, 25
- users 52, 53, 54, 57, 59, 61
- credentials. *See* user credentials
- cube reports 223
- cubes 223
- cubeview files 223
- custom configurations 29
- custom installations 29, 64
- Customer Support xi
- customizing
 - Information Console 202, 203
 - third-party databases 5, 18

D

- dashboards vii, 223
- data
 - accessing 13
 - analyzing 223
 - backing up Encyclopedia and 240
 - preventing loss of 7, 26, 62
 - recovering 7
 - storing 185
 - summarizing 223
 - upgrades and 6, 79, 116
- Data Analyzer Option 223
- data cache 224
- data cubes. *See* cubes
- Data Definition Language statements 53
- data directory 29, 116
- Data Exporter 137, 154
- Data Exporter configurations 157
- Data Exporter properties 142, 143
- data files 190, 241
 - See also* files; report files
- data objects 240
- data objects files 5
- data sources 11
- Data Store Administrator
 - See also* System Data Store Administrator
 - configuring 157
 - creating Encyclopedia and 156, 157
 - importing Encyclopedia and 139, 153, 154
 - migrating iServer installations and 6, 17
 - running 153
 - setting properties for 144
- Data Store Upgrader 6, 17
- Data Store Upgrader configurations 157
- data stores 24, 150, 203
- data types 18
- DATA_EXPORT_FOLDER parameter 145
- DATA_EXPORT_FORMAT parameter 145
- DATA_IMPORT_FOLDER parameter 145
- DATA_IMPORT_FORMAT parameter 145
- database administration tool
 - accessing PostgreSQL servers and 41, 98
 - backing up Encyclopedia and 244, 266, 287
 - installing 37
 - restoring Encyclopedia and 252, 274, 298
 - uninstalling 45
 - upgrades and 91, 108
- Database Credentials option 31
- database drivers. *See* drivers
- Database Schema Information page 69
- database schema names 148
- database schema owner 5, 31, 53, 54, 56, 58, 61
- database schemas
 - adding Encyclopedia to 156, 157
 - assigning privileges 53, 54, 56, 57, 58, 61
 - backing up 26, 150, 249
 - browsing 41
 - changing 98
 - creating 52, 53, 55, 59, 60, 61, 62, 135
 - depopulating 146
 - importing volume data in to 153, 154
 - preventing data loss and 7
 - removing Encyclopedia from 155
 - storing volume information and 4, 5
 - upgrades and 79, 87
- database servers. *See* servers
- database superuser 31
- DATABASE_HOST parameter 146
- DATABASE_NAME parameter 146
- DATABASE_PORT parameter 146
- DATABASE_TYPE parameter 144, 146
- databases
 - accessing Encyclopedia and 47
 - adding Encyclopedia and 157
 - adding indexes for 5

- analyzing data and 223
- backing up metadata for 7, 26, 62, 79, 241
- backing up system 17
- caching information objects and 224
- changing 54, 57, 59, 61
- changing volume metadata and 5
- configuring third-party RDBMS 8, 9, 13, 14
- connecting to 54, 57, 59, 61
- creating system 11, 53, 57
- creating system data store and 152
- customizing 5, 18
- importing Encyclopedia and 153, 154
- installing alternate metadata 52, 53, 63
- installing iServer and 24, 27, 29
- managing Encyclopedia and 7
- running iServer clusters and 14
- selecting metadata 65
- setting password for 31
- storing cluster information and 184
- storing user information and 4, 5
- uninstalling 45, 98
- upgrades and 6, 78, 79, 86, 91
- viewing incomplete files in 241
- viewing installation information for 44
- DB2 database administration tool
 - backing up Encyclopedia and 287
 - restoring Encyclopedia and 298
- DB2 databases 52
- DDL statements 53
- default directories. *See* directories
- default locales 33, 202
- default ports. *See* ports
- default time zones 203
- default user names 34
- default values 45
- default volume partitions 134
- DEFAULT_LOCALE parameter 202
- DEFAULT_TIMEZONE parameter 203
- DEFAULT_VOLUME parameter 203
- DELETE_ALL_DATA parameter 146
- DELETE_DATA parameter 146
- deleting
 - Encyclopedia volumes 155
 - installation directories 45, 98
- deploying
 - Information Console 195, 201–204
 - iServer 6, 13, 160, 184, 186, 230
 - new releases 47
 - spreadsheets 223
 - WAR files 201
 - web applications 204
- deployment scripts 185, 186
- deployment tools 202
- DEPOPULATE_SCHEMA parameter 146
- design files 5
- designers vii, 47
- designs 11, 48, 223, 225
- desktop products 48
- desktop shortcuts 41, 96, 113, 187
- destination folder locations 29
- developing applications 15
- DHTML reports 224
- diagnostic fixes 9, 14, 161
- directories
 - configuring cluster nodes and 13, 160
 - customizing WAR files and 203
 - deploying iServer images and 189, 190
 - extracting product files to 194
 - installing cluster nodes and 164
 - installing documentation files and xi
 - installing iServer and 25, 29
 - installing iServer Integration Technology and 208
 - installing JDK files and 46
 - migrating Encyclopedia and 133
 - removing installation 45, 98
 - running multiple releases and 45
 - storing data and 185
 - upgrading iServer and 45, 82, 99, 116
- directory paths 29, 116, 134
- disk partitions 241
 - See also* volume partitions
- disk space 46
- displaying
 - charts 46
 - cube reports 223
 - licensing information 224, 228
 - machine ID information 228
 - processor IDs 230
 - ReadMe files 201, 211
 - reports 223
- DLLs 46, 189
- document files 5, 225

- document generation service 11
- documentation
 - accessing vii–xi
 - administering iServer System and 16
 - deploying WAR files and 201
 - downloading xi
 - installing online xi, 40, 212
 - managing Encyclopedia and 7
 - upgrades and 96, 112, 132
- documentation conventions xii
- documentation update tool 212
- domain servers 25
- downloading
 - documentation files x
 - JAVA SE Development Kit 185
 - product files 27
 - Visual C++ libraries 185
- drivers 11, 14, 18
- DROP_SCHEMA parameter 146
- dual-core CPUs 232
- dynamic link libraries 46, 189

E

- e.Analysis Option 224
- e.Report Designer Professional 224
- e.Report Option 224
- e.reporting server. *See* iServer
- e.reports 185, 224
- Eclipse BIRT open source projects vii
- elastic iServer clustering 13, 160, 184
- e-mail. *See* notifications
- encoding 19
- encyc directory 45, 82, 98, 133
- Encyclopedia Data Store Administrator
 - configuring 157
 - creating Encyclopedia and 156, 157
 - importing Encyclopedia and 139, 153, 154
 - migrating iServer installations and 6, 17
 - running 153
 - setting properties for 144
- Encyclopedia Data Store Upgrader 6, 17
- Encyclopedia Data Store Upgrader configurations 157
- Encyclopedia Metadata Storage page 65, 85
- Encyclopedia processes. *See* encycsrvr11 processes

- Encyclopedia volume databases. *See* PostgreSQL databases; Oracle databases
- Encyclopedia volumes
 - accessing multiple 224
 - accessing objects in 222
 - adding to database schemas 156, 157
 - assigning schema owner for 31
 - backing up 240–243, 244, 266, 287
 - backing up metadata for 26, 62
 - browsing schemas for 41
 - browsing system database for 98
 - connecting to database for 18
 - controlling access to 47
 - CPU binding and 235
 - creating 27
 - creating partitions for 134
 - creating schema owner for 54, 56, 58, 61
 - creating schemas for 55, 59, 60, 61, 62, 135
 - customizing system database for 5, 18
 - deploying iServer and 187
 - exporting 133, 154
 - failover operations for 7, 8
 - importing 138, 153, 154
 - installing alternate database for 52, 53, 63
 - installing Information Console and 195
 - installing licenses and 33
 - installing system database for 4, 6, 31
 - iServer processes and 8
 - licensing and 222
 - localizing 19
 - managing 7
 - migrating data for 17
 - migrating to current release 78, 133
 - naming 140
 - preventing data loss for 7, 26
 - removing from schemas 155
 - renaming 153
 - restoring 243, 252, 260, 266, 274, 286, 298
 - running iServer clusters and 14, 160
 - running system database for 4, 9
 - setting location of 29
 - specifying default 203
 - specifying primary partition for 140
 - starting 9
 - storing user information for 4, 5, 11
 - troubleshooting 241
 - uninstalling databases for 45, 98

- upgrading 49
- upgrading iServer and 6, 17, 79
- uploading sample files for 189
- EncyclopediaUpgrader utility. *See* Encyclopedia Data Store Upgrader
- encycsrvr11 processes 230, 234, 235
- environment variables 13, 46, 160, 230
- environments 201
- error logging reports 19
- escape characters 157
- evaluation copies 27
- evaluation licenses 32, 120
- example files 189
- example reports 19
- Excel formats vii
- Excel spreadsheets 223
- executable files 46
- Expiration time of deleted files property 242
- expired licenses 227, 228
- EXPORT_ALL_DATA parameter 146
- EXPORT_DATA parameter 147
- exporting Encyclopedia volumes 133, 154
- external data sources 11
- external security systems 29
- extracting product files 194

F

- Factory processes 46, 232, 235
- Factory service 9, 11, 46
- failover operations 7, 8
- features xi, 46, 222
- file folder 135
- file I/O operations 11
- file names 226
- file purging properties 242
- file system backup utilities 241
- file systems 5, 11
- FILE_LOG_LEVEL parameter 143, 147
- files
 - accessing online help x, xi, 215
 - archiving 242
 - backing up 17, 82, 101, 241
 - configuring upgrade utilities and 157
 - copying 241
 - deploying iServer and 184, 186, 190
 - downloading documentation x

- downloading product 27, 194
- installing iServer and 26, 27
- installing JDK software and 46
- installing online help 212
- managing 7
- migrating Encyclopedia volumes and 135
- obtaining license key 226, 228
- obtaining temporary license 225
- overwriting product 6, 45
- purging 242
- saving 82, 101
- selecting license 32
- specifying location of 29, 116
- storing report-specific 5
- transmitting license keys and 229
- updating 212
- upgrading iServer and 48, 79, 80
- uploading sample 189
- filetype folder 135
- firewalls 175, 195
- fixes 9, 14, 161
- Flash gadgets vii
- folders
 - archiving 187
 - backing up 250, 273, 297
 - creating 244, 266, 287
 - deploying iServer and 186
 - installing iServer and 29
 - installing Management Console and 171
 - migrating Encyclopedia and 135
 - restoring 253
 - shared configurations and 165, 176
 - specifying program 125
 - specifying resource 202
 - upgrades and 135
- formats vii
- Forrester Wave Open Source Business Intelligence report vii
- ftp distributions
 - cluster installations 162
 - documentation xi, 212
 - Information Console 194, 195
 - iServer 27
 - iServer Integration Technology 208
 - iServer System 194
 - iServer upgrades 82, 101, 114
- full installations 194

G

- gadgets vii
- generating
 - charts 46
 - machine ID information 228
 - reports 47
 - temporary documents 11
- getJDBCMajorVersion method 18
- graphs. *See* charts

H

- heap size 202
- help 212
 - See also* documentation
- help files x, xi, 212, 215
- help topics xi
- hostnames
 - Oracle databases 66
 - PostgreSQL databases 67
- How 253
- HTML documentation vii, x
- hyperthreading 233

I

- I/O operations 11
 - See also* input; output
- icons (console applications) 125
- ICU libraries 25
- ICU_DATA environment variable 25
- IDAPI applications 10, 16
- IMPORT_DATA parameter 147
- importing Encyclopedia volumes 138, 153, 154
- incomplete files 241
- indexed searches 5
- Information Console
 - assigning CPUs to 202
 - caching web pages for 204
 - cluster configurations and 14
 - connecting to 203
 - customizing 202, 203
 - deploying 195, 201–204
 - installing 194, 195, 201
 - logging activities for 203
 - logging in to 41, 205
 - preventing cache conflicts for 204
 - setting application server port for 195
 - setting context paths for 197, 202
 - setting default Encyclopedia for 203
 - stand-alone configurations and 9
 - starting 205
 - testing installations for 205
- Information Delivery API 16
 - See also* IDAPI applications
- information delivery solutions vii
- Information Object Caching Option 224
- information object files 5
- information objects 223, 224
- INITIALIZE_DATA parameter 147
- in-place upgrades 78, 80, 82, 101
- installation
 - alternative volume databases and 52, 53, 63
 - C++ runtime components 187
 - cache conflicts and 204
 - cluster nodes 29, 162–174
 - Configuration Console 164, 194
 - desktop products 48
 - disk space and 46
 - domain servers and 25
 - Encyclopedia volume database 4, 5, 6
 - Information Console 194, 195, 201
 - iServer 6, 24, 27–41, 52, 64
 - iServer Integration Technology 194, 208–211
 - iServer System 6, 194
 - iServer System components 29
 - Java Runtime Environment 195
 - Java Software Development Kit 46, 195
 - license keys 228
 - licensing prompts and 227, 228
 - Management Console 169, 194
 - online documentation xi, 40, 212
 - Open Security application 47
 - pgAdmin administration tool 37
 - previous releases and 45
 - testing 47, 205
 - upgrades and 45, 78, 80
- installation directories 29, 45, 98
- installation guides 16
- installation options 29, 64
- installation prerequisites 24, 52, 161, 194

- installation wizard 195
 - Integration processes 232
 - Integration service 9, 11
 - IntegrationService element 232
 - Interactive Viewer 223
 - Interactive Viewer Option 223
 - international character sets 19
 - international customer licensing 227
 - iportal context path 197
 - iServer
 - accessing functionality 222
 - accessing online help for 215
 - binding to multiple-core CPUs 232
 - binding to processors 230
 - changing CPU binding and 234
 - checking bound processors for 234–235
 - configuring 9
 - controlling access to 29
 - deploying 6, 13, 160, 184, 186, 230
 - downloading evaluation copies for 27
 - getting machine ID for 227
 - installing 6, 24, 27–41, 52, 64
 - installing as cluster node 162–174
 - optimizing performance for 14, 161
 - running 14, 25, 160
 - setting port number for 31
 - setting up user accounts for 24, 25, 26
 - setting up volume database and 5, 9
 - shutting down 137, 189
 - starting 188
 - storing volume information and 4, 5
 - testing new releases for 47–49
 - upgrades and 6, 45
 - upgrading 78, 80, 99, 113
 - validating CPU binding and 235
 - viewing licensing information for 224, 228
 - iServer 11 service 34, 121
 - iServer directory 29
 - iServer distribution package 186
 - iServer images 184, 189
 - iServer Information page 199
 - iServer Integration Technology 16, 194, 208–211
 - iServer processes
 - binding to CPUs 226, 229–236
 - binding to Windows systems 229
 - domain servers and 25
 - multiple-core CPUs and 232
 - overwriting files and 80
 - running 8, 13, 160, 229
 - starting 230
 - upgrading and 80
 - verifying CPU binding for 232
 - iServer services
 - See also* specific service
 - cluster configurations and 8
 - running 8
 - shutting down 80, 99
 - specifying user accounts for 24
 - iServer servlet container 8
 - iServer System
 - administering 15
 - changing CPU binding and 236
 - cluster connections and 8
 - configuring 6
 - deploying new releases and 47
 - downloading files for 27, 194
 - installation options for 6, 29, 64
 - installation requirements for 194
 - licensing options for 222, 227
 - obtaining license keys for 228
 - obtaining licensing information for 225
 - optimizing 13, 160
 - running multiple releases and 45
 - upgrading 79
 - iServer System components 28, 29, 41
 - iServer System names 30
- ## J
- J2EE installations 201
 - JAR files 47
 - Java Factory service. *See* Factory service
 - Java heap size 202
 - Java Object Interface 46
 - Java Runtime Environment (JRE) 195
 - JAVA SE Development Kit 185
 - Java Server Pages. *See* JSPs
 - Java Software Development Kit (JDK) 46, 195
 - Java Virtual Machines. *See* JVM
 - JDBC drivers 11, 18
 - jdbcCompliant method 18
 - jobs 10
 - JSPs 204

JVM libraries 46

K

keys. *See* license keys

L

language settings 33
launching iServer images 184
LDAP configuration file 47
LDAP security database 47
LDAP servers 29, 47
libraries 25, 46, 185, 189
License File Details page 32, 87
license files 32, 225
license key file locations 227
license key file names 226, 229
license key files 227
license key information 226
license key installations 228
license key updates 228
license keys 225, 226, 229, 230
licensed CPUs 235
 See also CPU-based licenses
licenses 87, 120, 222
licensing information 224, 225
licensing options 222, 227
links (documentation) xi
Linux systems 16
load balancing (clusters) 14, 160
local help 215
Local Security Settings dialog 26
Local Users and Groups dialog 25
Locale Information page 33, 87
Localemap.xml 202
locales
 encoding character sets for 19
 obtaining licenses for 227
 selecting 33
 setting default 33, 202
 updating resource files for 212
 upgrades and 87
localhost parameter 203
Localization and Online Documentation
 Update tool 212
log files 29, 116, 203
Log on as a service privilege message 32, 86

LOG_FILE_COUNT parameter 143, 147
LOG_FILE_LOCATION parameter 203
LOG_FILE_NAME parameter 147
LOG_FILE_SIZE parameter 143, 147
LOG_FOLDER parameter 143, 144, 147

logging levels

 Data Store Administrator 147

 Squirrel Data Exporter 143

logical CPUs 232

logical processors 230, 233

losing data 62

M

machine IDs 225, 227, 229

machine information 227

machine names 45

maintenance licenses 226

Management Console

 accessing documentation for 16

 cluster configurations and 14

 installing 169, 194

 licensing and 222

 managing Encyclopedia and 188

 setting application container port for 170

 setting user names for 34

 stand-alone configurations and 9

 starting 41

 upgrades and 89

manual in-place upgrades 78, 80, 82

manual iServer system upgrades 6

manual side-by-side migrations 133

manual side-by-side upgrades 78, 113

manuals. *See* documentation

Manuals directory xi

mapping user information 47

master cluster nodes 8

master index file xi

message routing (requests) 14

metadata

 autoarchiving and 242

 backing up 26, 62, 241

 changing 5

 defined 240

 preventing loss of 7, 26

 storing 52, 150

 storing volume information and 4, 11

- upgrades and 6, 79, 85
- metadata database 5, 52, 65, 241
 - See also* PostgreSQL databases; Oracle databases
- Microsoft Excel spreadsheets 223
- Microsoft SQL Server databases 52
- Microsoft Windows systems. *See* Windows systems
- migration 6, 16, 45, 47, 49, 78
- migration options 133
- migration utilities 49
- missing product files 194
- missing system schemas 79
- multicast configurations 8
- multidimensional data analysis 223
- multiple-core CPU binding 232
- multiple Encyclopedia volumes 224
- multiple product installations 45
- multiple upgrade installations 46
- Multi-Tenant Option 224
- multi-threaded CPUs 230

N

- Named User Platform 222
- named-user licenses 33, 87, 222
- naming conventions 53, 66
- network administrators. *See* administrators
- network cards 227, 229
- Network File Systems (NFS) 11
- networked environments 13, 160, 194
- New Schema page 135
- New Server Registration dialog 42
- NEW_SCHEMA_NAME parameter 142, 148
- NEW_SCHEMA_PASSWORD parameter 148
- NEW_VOLUME_NAME parameter 143, 148
- newinformationconsole.war 204
- NobleNet Portmapper processes 80, 81, 101
- node keys 225
- node-key file locations 227
- node-key licensing 225, 226, 227, 228
- nodes. *See* cluster nodes
- notifications 235, 236

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
 - downloading xi
 - installing xi, 40, 212
 - managing Encyclopedia and 7
 - upgrades and 96, 112, 132
- Online Documentation and Localization
 - resource files xi
- online documentation conventions xii
- online documentation update tool 212
- online Encyclopedia volumes 9
- online help 215
 - See also* online documentation
- Open Security applications 29, 47, 82
- open source projects vii
- opening
 - Information Console 41, 205
 - Management Console 41
 - pgAdmin administration tool 42
- operating systems 19, 229
- optimizing iServer System 13, 160
- options (installation) 29, 64
- options (licensing) 222, 227
- Oracle database application users 57, 61
- Oracle database port 66
- Oracle database schema owners 56, 61
- Oracle database schemas 55
- Oracle databases
 - accessing documentation for 7
 - installing Encyclopedia and 63
 - installing iServer and 24, 52
 - storing user information and 5
 - upgrades and 78
- ORACLE_TNS_NAMES_FILE parameter 148
- output 11
- output formats vii
- overwriting product files 6, 45

P

- Page Level Security Option 224
- page-level security 223, 224

- parameters
 - changing 187
 - configuring clusters and 13, 160
 - creating Encyclopedia volumes and 142, 143
 - deploying Information Console and 202
 - encoding volume data and 19
 - setting Data Store Administrator 144
- partition names 134
- partition paths 134
- partitions (disk) 241
- partitions (volume) 134, 140
- Partitions page 134, 141
- passwords
 - changing 34
 - installing cluster nodes and 169
 - installing Information Console and 198
 - installing iServer and 27
 - setting database 31
 - starting console applications and 34
 - starting iServer service and 34
- patches 9, 14, 161
- paths 29, 116, 134
- PDF documentation vii, x, xi, 212
- performance
 - cluster configurations and 14, 161
 - CPU binding and 232
 - file input/output operations and 11
 - iServer processes and 9
 - upgrades and 46
- permanent licenses 227
- permissions. *See* privileges
- pg_dump command line options 249
- pg_dump utility 243, 245, 250
- pg_restore command-line options 260
- pg_restore utility 243, 257, 260
- pgAdmin database administration tool
 - accessing PostgreSQL servers and 41, 98
 - backing up Encyclopedia and 244
 - installing 37
 - restoring Encyclopedia and 252
 - uninstalling 45
 - upgrades and 91, 108
- pgAdmin III Setup page 109
- pgAdmin III Setup Wizard 111
- platform licenses 222
- POPULATE_SCHEMA parameter 148
- POPULATE_SCHEMA_PHASE_ONE
 - parameter 148
- POPULATE_SCHEMA_PHASE_TWO
 - parameter 149
- portinst.exe 185
- portmapper components 185
- ports
 - application servers 195
 - cloud deployments and 185
 - cluster nodes 168
 - Information Console 203
 - iServer 31
 - multiple product releases and 45
 - Oracle databases 66
 - PostgreSQL databases 67
 - upgrades and 45
- postgres database directory 45, 98
- postgres system database 53, 57
- PostgreSQL Database Connection
 - Information 86
- PostgreSQL Database Information page 31, 86
- PostgreSQL database tools 243
- PostgreSQL databases
 - accessing documentation for 7
 - adding Encyclopedia and 157
 - backing up 17, 243
 - browsing schema definitions in 41
 - changing schemas for 98
 - cluster configurations and 13, 14
 - creating alternate 53, 57
 - creating schema owners for 54, 58
 - file I/O operations and 11
 - importing Encyclopedia and 153, 154
 - installing administration tool for 37
 - installing as alternate 52, 53, 57
 - installing iServer and 24, 27
 - setting password for 31
 - setting port number for 67
 - stand-alone configurations and 9
 - storing user information and 4, 5
 - uninstalling 45, 98
 - upgrades and 78, 86, 91
- postgresql directory 251
- PostgreSQL servers 5, 13, 41, 98
- PostgreSQL superusers 31
- Power Users accounts 24

- prerequisites (installation) 24, 52, 161
- previous releases 6, 45, 79, 99
- printers 14, 24
- privileges
 - creating system schemas and 152
 - creating user accounts and 24
 - installing alternate databases and 53, 54, 57, 59, 61
 - installing Information Console and 194
 - logging in to iServer and 26
 - upgrades and 99
- Process Management Daemon
 - CPU binding and 230, 232
 - running cluster configurations and 13, 14, 160
 - running iServer processes and 8, 10
 - starting encycsrvr processes and 234, 235
- Processor Affinity 232
- processor IDs 230, 232
- ProcessorAffinity element 232
- processors. *See* CPUs
- product files 6, 26, 45, 82, 194, 212
- product information xii
- product updates xi, 212
- production environments 48
- production staging areas 48
- profiles (user) 34
- program files. *See* product files
- program folders 125
- properties files 157, 202
- Properties page (New Server Registration) 43
- publishing reports 223
- Purge deleted files time property 242
- purging report files 242

Q

- queries 11, 13, 54, 57, 59, 61
- Query Option (licensing) 223

R

- RDBMS backup utilities 241
- RDBMS databases 4, 5, 14, 64, 241
 - See also* third-party databases
- RDBMS tools 7, 8, 17
- ReadMe files 201, 211
- rebinding encycsrvr11 processes 235

- recovering data 7
- refreshes 204
- relational databases. *See* databases
- release notes xi
- remote procedure calls 10
- remote servers 194
- removing
 - Encyclopedia volumes 155
 - installation directories 45, 98
- renaming
 - Encyclopedia volumes 153
 - license files 229
- report design files 5
- report designers vii, 47
- report designs 11, 48, 223, 225
- report document files 5, 225
- report document generation service 11
- report files
 - archiving 242
 - backing up 17, 82, 101, 241
 - managing 7
 - purging 242
 - storing 5
 - upgrading iServer and 6, 48, 79, 80
 - uploading sample 189
- report object executable files 46
- Report Server Security Extension API 82
- report templates 202
- report viewers 223
- reporting applications. *See* applications
- reporting server. *See* iServer
- reporting services. *See* specific iServer service
- ReportingService element 232
- reports
 - accessing sample 19
 - displaying 223
 - generating 47
 - publishing 223
 - running 185, 224
- repositories 184
 - See also* Encyclopedia volumes
- requests
 - installing cluster nodes and 168
 - remote procedure calls and 10
 - running cluster configurations and 14, 160
- resource file updates 212
- resource folders 202

- resource groups 46
- resources
 - accessing 202
 - creating cluster nodes and 13, 14, 160
 - deploying iServer and 186, 189
 - installing Information Console and 194
 - installing iServer and 24
 - obtaining licenses and 222
- restoring Encyclopedia volumes 243, 252, 260, 266, 274, 286, 298
- restricting iServer processes 229
- rollbacks 79
- RSSE code 82
- running
 - Data Store Administrator 139
 - encycsrvr11 processes 230, 234, 235
 - iServer 25
 - iServer processes 25, 226, 229
 - iServer services 8, 24
 - jobs 10
 - multiple product releases 45
 - pg_dump utility 245, 250
 - pg_restore utility 257, 260
 - pgAdmin administration tool 41, 98
 - queries 11, 13, 54, 57, 59, 61
 - report designs 223
 - reports 185, 224
 - spreadsheet reports 223
 - Squirrel Data Exporter 137

S

- sample files 189
- sample reports 19
- saving previous releases 82, 101
- scalability vii
- schema names 148
- schema owner 5, 31, 53, 54, 56, 58, 61
- SCHEMA_FILE_NAME parameter 144
- SCHEMA_NAME parameter 149
- SCHEMA_PASSWORD parameter 149
- schemas
 - adding Encyclopedia to 156, 157
 - assigning privileges 53, 54, 56, 57, 58, 61
 - backing up 26, 150, 249
 - browsing 41
 - changing 98

- configuring system data store 150, 152
- creating 52, 53, 55, 59, 60, 61, 62, 135
- depopulating 146
- importing volume data in to 153, 154
- preventing data loss and 7
- removing Encyclopedia from 155
- storing volume information and 4, 5
- upgrades and 79, 87
- SCRIPT_HOME parameter 144
- scripts 4, 53, 185
- search results 224
- searching online documentation xi
- security vii, 224
- security database 47
- security systems 29
- server templates 13, 160, 187
- SERVER_DEFAULT parameter 203
- servers
 - See also* iServer
 - accessing PostgreSQL 41, 98
 - changing cluster machines and 229
 - configuring as cluster node 180
 - configuring as clusters 13, 14, 164
 - controlling Encyclopedia access and 47
 - deploying Information Console to 201, 202
 - exceeding CPU licenses for 235
 - installing iServer for external 29
 - integrating Actuate products with vii, 202
 - preventing cache conflicts for 204
 - running as clusters 180
 - running iServer processes and 25
 - setting port numbers for 195
 - shutting down iServer and 9
 - storing volume metadata and 5
- service names 66
- service privileges 26
- services. *See* specific iServer service
- ServIntTech directory 208
- servlet container 8
- Setup Type page 29, 64
- setup.exe 27, 82
- shared database schemas 156
- shared licenses 225, 226
- shared repository 184
- sharing metadata databases 5
- shortcuts 41, 96, 113, 187
- shutting down. *See* stopping

- side-by-side migrations (Encyclopedia) 133
- side-by-side upgrades 78, 113
- SIDs (system identifiers) 66
- single-point node failure 13, 160
- SmartSheet Security Option 223
- snapshots 241
- SOAP messages 10
- SOAP processor 10
- Specify Profiles page 34, 88, 105
- Specify Windows Account Information page 198
- spreadsheet reports 223
- SQL queries 11, 13, 54, 57, 59, 61
- SQL scripts 4, 53
- SQL Server database administration tool
 - backing up Encyclopedia and 266
 - restoring Encyclopedia and 274
- SQL Server databases 52
- Squirrel Data Exporter 6, 17, 137, 154
- Squirrel Data Exporter configurations 157
- Squirrel Data Exporter properties 142, 143
- Squirrel database 4
- SQUIRREL_DATA_HOME parameter 143
- SQUIRREL_EXPORT_FOLDER parameter 142, 143
- SQUIRREL_EXPORT_FORMAT parameter 143
- stand-alone Encyclopedia database configurations 9
- starting
 - Data Store Administrator 139
 - Encyclopedia volumes 9
 - encycsrvr11 processes 234
 - Information Console 41, 205
 - iServer 188
 - iServer 11 service 34, 121
 - iServer processes 230
 - Management Console 41
 - pg_dump utility 245, 250
 - pg_restore utility 257, 260
 - pgAdmin administration tool 42
 - Squirrel Data Exporter 137
- startiServer script 186
- status folder (encyc) 135
- stopping
 - iServer 137, 189
 - iServer processes 80

- iServer services 80, 99
- NobleNet Portmapper processes 81, 101
- summarizing data 223
- SUPERUSER parameter 149
- superuser password 27, 31
- SUPERUSER_PASSWORD parameter 149
- superusers 31
- Supported Products Matrix xii, 201
- Switch Help Location command 215
- synchronous Factory service 11
- syntax conventions (documentation) xii
- system administrator privileges 194
- system administrators. *See* administrators
- System Configuration Password page 34, 89, 106
- System Data Store Administrator 79, 150, 152
- system databases 53, 57
 - See also* PostgreSQL databases; Oracle databases
- system identifiers 66
- system metadata 26, 62
- system names 30, 85
- system schema owner 54, 55, 56, 58, 61
- system schema passwords 31
- system schemas 55, 59, 61, 79
- system variables 230
- SYSTEM_DATABASE_NAME parameter 149
- SystemDataStoreAdministrator class 152

T

- table of contents (documentation) xi
- tables 5
- TABLESPACE_LOCATION parameter 149
- TABLESPACE_NAME parameter 149
- Taking volume offline 253, 274
- Taking volume online 286
- Taking volume online 260, 313
- TAR files 184
- technical support xi
- TEMP_FOLDER_LOCATION parameter 203
- templates 187, 202
- temporary directories 203
- temporary documents 11
- temporary files 203
- temporary license keys 228

- temporary licenses 225, 227
- testing
 - desktop products 48
 - Information Console installations 205
 - new releases 47–49
- text files 229
- 3rd-Party Database Information page 66
- third-party database tools 243
- third-party databases
 - caching information objects and 224
 - creating indexes for 5
 - customizing 5, 18
 - installing 52, 53
 - installing iServer and 24
 - preventing data loss and 7, 26, 62
 - running iServer clusters and 8, 13, 14
 - selecting 65
 - storing cluster information and 184
 - uninstalling 98
 - upgrades and 79
- third-party deployment tools 202
- threads 202
- time zones 33, 203
- TIME_ZONE parameter 150
- TimeZones.xml 203
- TNS services 66
- Tomcat servers 195, 197
- transient data stores 203
- transient files 203
- TRANSIENT_STORE_PATH parameter 203
- Transparent Network Substrate services 66
- transport protocol (requests) 10, 14, 160
- troubleshooting 241
- typographic conventions
 - (documentation) xii

U

- UCS2 conversions 19
- uninstalling
 - pgAdmin administration tool 45
 - pgAdmin utility 98
 - system databases 98
 - third-party databases 45
- UNIX systems 16
- Update tool 212
- updates (documentation) viii, xi
- updates (product) xi, 212
- updating license key files 228
- upgrade options 78, 80, 114
- upgrade utilities 157
- Upgrader utility 6, 17
- Upgrader utility configurations 157
- upgrades
 - creating backups and 79
 - customizing volume database and 5
 - determining if needed 47
 - disk space and 46
 - Encyclopedia volumes 49, 133
 - installing product files and 79
 - iServer 78, 80, 99, 113
 - iServer processes and 80
 - licensing options and 226
 - migrating volume data and 17
 - overwriting previous releases and 6
 - pgAdmin administration tool 91, 108, 127
 - PostgreSQL database schemas and 98
 - previous releases and 45
 - rolling back 79
- uploading sample files 189
- URLs
 - Actuate product information xii
 - Actuate technical support xi
 - Configuration Console 188
 - documentation updates xi
 - evaluation copies 27
 - Information Console 203, 205
 - JAVA SE Development Kit 185
 - licensed products 226
 - Localization and Online Documentation
 - Update tool 212
 - Management Console 188
 - Network File Systems 11
 - PDF documentation 212
 - release notes xi
 - Supported Products Matrix 201
- usage reports 19
- USE_SERVER_CONFIG_FILE parameter 150
- user accounts
 - cluster configurations and 25
 - connecting to databases and 54, 57, 59, 61
 - installing alternate databases and 53
 - installing Information Console and 198
 - iServer installations and 24, 25

- managing 16
- upgrades and 78, 88
- user credentials 5
- user information 4, 47
- user names 34, 53
- user profiles 34
- User Rights Assignments dialog 26
- user tracking reports 19
- users 33, 52, 54, 57, 59, 61, 222

V

- validating CPU binding 235
- verifying CPU binding 232
- View processes 232, 235
- View service 9, 11, 46
- viewers 223
- viewing
 - charts 46
 - cube reports 223
 - licensing information 224, 228
 - machine ID information 228
 - processor IDs 230
 - ReadMe files 201, 211
 - reports 223
- ViewingService element 232
- Vista operating systems 24
 - See also* Windows systems
- Visual C++ libraries 185
- volume administrators. *See* administrators
- volume data 240
- volume databases. *See* PostgreSQL databases;
Oracle databases
- volume metadata database 241
- volume names 140, 153, 195
- volume partition names 134
- volume partition paths 134
- volume partitions 134, 140
- VOLUME_NAME parameter 150
- volumes. *See* Encyclopedia volumes

W

- WAR files 195, 201, 202, 203
- watermarks 225
- web administrator privileges 194
- web administrators. *See* administrators
- web applications 204

- web archive files. *See* WAR files
- web browsers 204
- web pages 204, 223, 224
- web servers 202

- See also* servers

Windows systems

- binding iServer processes to 229, 230–233
- collecting machine information for 227
- configuring firewalls for 175, 195
- creating iServer images for 185
- distributing iServer System over 6
- installation prerequisites for 24, 52, 161, 194
- installing cluster nodes for 162–174
- installing Information Console for 194, 195, 201, 205
- installing iServer for 27
- installing iServer Integration Technology for 208
- installing JDK files for 46, 195
- installing JRE files for 195
- installing localization and documentation files for 212
- running iServer processes on 25, 80
- setting classpaths for 47
- setting default locale for 33
- setting up user accounts for 24, 25, 26
- starting iServer service for 34, 121
- testing new releases for 47–49
- upgrading Encyclopedia volumes for 137
- upgrading iServer for 45, 78, 80, 99, 113

Work Unit Licenses 222

X

- XML files 229

Z

- ZIP files 184