

ActuateOne™

One Design
One Server
One User Experience

**Installing BIRT iServer
for Windows**

This documentation has been created for software version 11.0.5.

It is also valid for subsequent software versions as long as no new document version is shipped with the product or is published at <https://knowledge.opentext.com>.

Open Text Corporation

275 Frank Tompa Drive, Waterloo, Ontario, Canada, N2L 0A1

Tel: +1-519-888-7111

Toll Free Canada/USA: 1-800-499-6544 International: +800-4996-5440

Fax: +1-519-888-0677

Support: <https://support.opentext.com>

For more information, visit <https://www.opentext.com>

Copyright © 2017 Actuate. All Rights Reserved.

Trademarks owned by Actuate

“OpenText” is a trademark of Open Text.

Disclaimer

No Warranties and Limitation of Liability

Every effort has been made to ensure the accuracy of the features and techniques presented in this publication. However, Open Text Corporation and its affiliates accept no responsibility and offer no warranty whether expressed or implied, for the accuracy of this publication.

Document No. 170215-2-430342 February 15, 2017

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
Obtaining technical support	xi
About supported and obsolete products	xii
Typographical conventions	xii
Syntax conventions	xii
About <i>Installing BIRT iServer for Windows</i>	xiii

Part 1

Architecture

Chapter 1

Understanding Actuate BIRT iServer architecture	3
Understanding BIRT iServer architecture	4
Using a third-party RDBMS with an Encyclopedia volume	4
Customizing Encyclopedia volume databases	5
Installing and configuring iServer System	6
Managing the backup, recovery, and failover capabilities of the Encyclopedia volume database and data files	6
Managing an iServer cluster	8
Understanding the iServer System process model	8
Understanding process flow in a stand-alone iServer	9
Understanding process flow in an iServer cluster	11
Administering iServer System	15
Using JDBC to connect to an Encyclopedia volume database	16
API Compatibility	17
About international character sets	17
Administrative reports	17
Supported operating systems	17

Part 2

Installing

Chapter 2

Installing BIRT iServer	21
Preparing to install BIRT iServer	22
Creating an account with Windows administrator privileges	22
Configuring the iServer user account	23
Configuring log on as a service privilege	24
Backing up iServer system and Encyclopedia volume metadata	24
Performing a new installation	25
Installing a new instance of BIRT iServer Release 11 Service Pack 5	25
Understanding the iServer installation environment	36
About upgrading or migrating an earlier Release 11 iServer to Release 11 Service Pack 5	36
About running different releases on the same machine	36
About performance and disk space issues	37
About upgrading an iServer with resource groups	37
About the Java Software Development Kit	37
Accessing JAR files for report generation	38
Gathering LDAP information	38
Following best practices	38
Using a test environment	38
Setting up a production staging area	39
Setting up a production environment	40

Chapter 3

Installing BIRT iServer using an alternative database	41
Preparing to install BIRT iServer using an alternative database	42
Creating an account with Windows administrator privileges	42
Creating the iServer system and Encyclopedia volume schemas and iserver user in an alternative database	42
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing PostgreSQL database	43
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing Oracle database	45
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing SQL Server database	47
Creating the system and Encyclopedia volume schemas and iserver user in a pre-existing DB2 database	50
Adding support for the digit wildcard character in iServer when the metadata database is DB2	51
Filtering for a user name returns an empty result if name contains certain special characters	53

Backing up iServer system and Encyclopedia volume metadata	54
Installing an Encyclopedia volume that uses an alternative database	55

Chapter 4

Upgrading BIRT iServer 69

Upgrading BIRT iServer	70
Preparing to upgrade BIRT iServer	70
Creating an account with Windows administrator privileges	71
Backing up iServer system and Encyclopedia volume metadata	71
Performing an automatic in-place upgrade from an earlier Release 11 iServer to Release 11SP4 Fix 5	72
Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services	73
Running the in-place upgrade on an earlier Release 11 iServer	75
About changes to the file structure	88
Performing a manual upgrade of an Encyclopedia volume schema for an earlier Release 11 iServer	89
Performing an automatic in-place upgrade from iServer Release 11SP4 Fix 5 to iServer Release 11SP5	93
Setting PostgreSQL system variables	94
Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services	94
Running the in-place upgrade on iServer 11SP4 Fix 5	97
Performing a side-by-side migration from iServer Release 11SP4 Fix 5 to iServer Release 11SP5	106
Exporting volume metadata from the 11SP4 Fix 5 installation	107
Preparing to run the Data Store Administrator utility	107
Creating the properties file	108
Running the Data Store Administrator utility	110
Importing volume metadata and data to the 11SP5 installation	111
Preparing to import volume metadata and migrating the volume data	112
Importing the volume metadata	115
Creating a new volume	123
Migrating a volume to a database of a different type	125
Creating schemas in the new database	127
Creating a metadata database and schemas in Configuration Console	127
Exporting volume metadata from the source database	134
Preparing to run the Data Store Administrator utility	134
Creating the properties file	135
Running the Data Store Administrator utility	136
Completing the volume migration	136
Importing the volume metadata to the target database	136
Creating a new volume	137
Working with iServer utilities	140
Working with Encyclopedia Data Store Administrator	141

Specifying Encyclopedia Data Store Administrator properties	141
Using the generated bulk-load script files	148
Performing operations using Encyclopedia Data Store Administrator utility	150
Working with System Data Store Administrator	155
Specifying System Data Store Administrator properties	155
Performing operations using the System Data Store Administrator utility	157
Working with Encyclopedia Data Store Upgrader	157
Specifying Encyclopedia Data Store Upgrader properties	157

Chapter 5

Installing a BIRT iServer cluster 161

Installing a BIRT iServer cluster node	162
Preparing the iServer cluster environment	163
Creating an account with Windows administrator privileges	171
Performing a wizard-based cluster node installation	171
Adding a node to a cluster	184

Chapter 6

Installing Information Console 187

Preparing to install Information Console	188
About installing from a download	188
About performing a full installation	188
Installing Information Console on Windows	189
Using the installation wizard	189
Using the WAR file to install	195
General deployment tasks	195
Preparing the server	196
Preparing the WAR file	196
About clusters of servers	198
Avoiding cache conflicts after installing	198
Testing the installation	198

Chapter 7

Installing iServer Integration Technology and documentation 201

Installing iServer Integration Technology	202
Installing the localization and documentation files	206
About accessing online help	209
Searching PDF manuals using master-index.pdx	209

Part 3

Licensing

Chapter 8

Licensing BIRT iServer 215

Understanding licensing types	216
Understanding licensing options	217
Installing Actuate BIRT iServer System license files	219
About the license file	220
Collecting machine information for a license	221
About modifying a license	225
Understanding CPU binding	225
Binding a BIRT iServer to processors on a Windows machine	226
Binding to specific CPUs	226
Binding to multiple-core CPUs	228
Binding an Actuate process to a processor	228
About processors and hyperthreading	229
Checking BIRT iServer bound processors	229
Determining the number of processors an iServer System uses	230
Understanding CPU binding validation while iServer is running	230
Understanding CPU binding validation when an Encyclopedia volume comes online	231
Understanding CPU binding validation when running iServer processes	231
Configuring e-mail for CPU license problems	231

Part 4

Backing up

Chapter 9

Backing up an Encyclopedia volume 235

Performing an Encyclopedia volume backup	236
Managing the backup and recovery of Encyclopedia volume metadata and data files . . .	236
Using RDBMS and file system backup utilities	237
Avoiding conflict with the file purging process	238
Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database	239
Backing up an Encyclopedia volume using pg_dump	240
Restoring an Encyclopedia volume using pg_restore	244
Backing up and restoring an Encyclopedia volume that uses an Oracle database	249
Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)	250
Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp)	252
Backing up and restoring an Encyclopedia volume that uses a SQL Server database	254
Backing up and restoring an Encyclopedia volume	254
Backing up an Encyclopedia volume using SQL Server Management Studio	254
Restoring an Encyclopedia volume using SQL Server Management Studio	262
Backing up and restoring an Encyclopedia volume that uses a DB2 database	273

Backing up an Encyclopedia volume using DB2 Control Center	274
Restoring an Encyclopedia volume using DB2 Control Center	283
Index	293

Understanding ActuateOne

ActuateOne™ includes the Release 11 series 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 documentation includes the materials described in Table I-1. You can obtain HTML and PDF files from developer.actuate.com.

Documentation updates are created in response to customer requirements.

Table I-1 BIRT iServer documentation



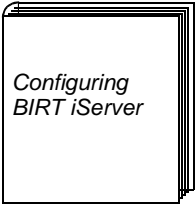
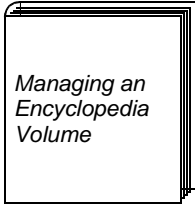
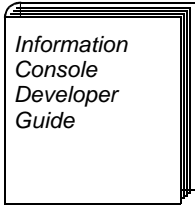
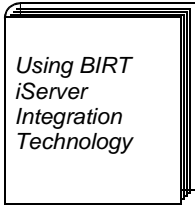
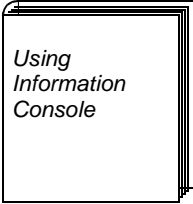
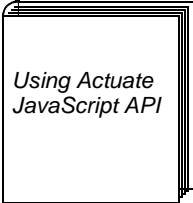
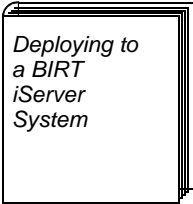
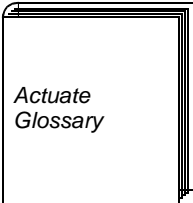

For information about this topic	See the following resource
Documentation updates	Updated HTML help and PDF files posted on developer.actuate.com .
Installing BIRT iServer for Linux and UNIX	
Installing BIRT iServer for Windows	
Configuring BIRT iServer Use Configuration Console to: <ul style="list-style-type: none">■ Add additional Encyclopedia volumes■ Configure clusters of iServers■ Tune iServer services and processes■ Configure e-mail notification■ Review and update license options■ Open ports for iServer use■ Manage iServer printers and resources■ Configure diagnostic logging	

Table I-1 BIRT iServer documentation (continued)

For information about this topic	See the following resource
Managing an Encyclopedia Volume Use Management Console to: <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>
Information Console Developer Guide <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>
Using BIRT iServer Integration Technology <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
Using Information Console <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>
Using Actuate JavaScript API <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>
Deploying to a BIRT iServer System Describes how to deploy designs and information objects to iServer	 <p><i>Deploying to a BIRT iServer System</i></p>
Actuate Glossary Definitions of product terminology	 <p><i>Actuate Glossary</i></p>
Adobe Acrobat Catalog A utility that can search all the documents in the Actuate manuals directory	 <p><i>Adobe Acrobat Catalog</i></p>

Obtaining documentation

Actuate provides technical documentation in PDF and HTML formats. Actuate products access HTML-format documentation from the Actuate web site. If you

do not have web access or prefer to use documentation hosted on your local system, install the documentation from the Online Documentation and Localization Resource Files package. You can download PDF or view HTML versions of the documentation from developer.actuate.com. If you purchase the product, you can also download documentation from OpenText My Support. If you select the typical setup when you install from the downloaded `ActuateLocalizationandOnlineDocumentation.exe`, the installation creates the `ACTUATE_HOME\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 `ACTUATE_HOME\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 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 OpenText My Support at the following URL:

<https://support.opentext.com>

A new user must first register on the site and log in to view the release notes.

Updates to documentation in PDF form are available on OpenText My Support or at the following URL:

<http://www.developer.actuate.com>

Alternatively, after installing Actuate release 11 products, use `Start`→`Programs`→`Actuate 11`→`Update Documentation` to access this URL.

Obtaining technical support

You can contact customer support by submitting a ticket. To submit a ticket, go to the following URL:

<https://support.opentext.com>

About supported and obsolete products

The Actuate Support Lifecycle Policy and Supported Products Matrix are available on OpenText My Support. You can access this site at the following URL:

<https://support.opentext.com>

Typographical conventions

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

Table I-2 Typographical conventions

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

Syntax conventions

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

Table I-3 Syntax conventions

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

Table I-3 Syntax conventions

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

About *Installing BIRT iServer for Windows*

Installing BIRT iServer for Windows includes the following chapters:

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

- *Chapter 5. Installing a BIRT iServer cluster.* Describes how to install a BIRT iServer cluster node in a Windows environment.
- *Chapter 6. Installing Information Console.* Describes how to install Actuate Information Console in Windows.
- *Chapter 7. 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 8. 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 9. Backing up an Encyclopedia volume.* Describes how to back up and restore BIRT iServer Encyclopedia volume metadata and data.

Part One



Architecture

Understanding Actuate BIRT iServer architecture

This chapter contains the following topics:

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

Understanding BIRT iServer architecture

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

Actuate also adapted iServer to support alternative, customizable, third-party database installations. In Release 11 Service Pack 5, Actuate currently supports DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL instance.

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

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

In Release 11, Actuate provides the following installation options:

- Install a new iServer with a PostgreSQL or other supported, alternative, third-party database.
- Upgrade the Actuate proprietary database installation from a previous major release, such as Release 10 Service Pack 1, to a Release 11 version.
- Upgrade an earlier Release 11 version to a newer Release 11 service pack, such as an upgrade from Release 11 Service Pack 4 to Service Pack 5.
- 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 Release 11 iServer system in place, to iServer Release 11 Service Pack 4 Fix 5, automatically migrating one or more Encyclopedia volumes.
Upgrade iServer Release 11 Service Pack 4 Fix 5 in place, to iServer Release 11 Service Pack 5, automatically migrating one or more Encyclopedia volumes.
- Upgrade iServer Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5) to iServer Release 11 Service Pack 5 (11SP5) by installing 11SP5 on a separate machine, then manually migrating Encyclopedia volume metadata and data from the 11SP4 Fix 5 installation to the 11SP5 installation.
After installation, the administrator uses the Encyclopedia Data Store Administrator utility to migrate Encyclopedia volume metadata to the 11SP5 installation. This utility is a Java program 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.

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

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

Since iServer Release 11 no longer uses the internal proprietary Squirrel database to store Encyclopedia volume metadata, the automatic backup, recovery, and

failover features available in earlier iServer releases are now obsolete. For information on the recommended procedures to back up and recover iServer system and Encyclopedia volume schemas in the Release 11 environment, see Chapter 9, “Backing up an Encyclopedia volume,” later in this book.

In Actuate Release 11 and later, there is no concept of volume failover, since each node in a cluster can operate on all the volumes. Configuring system and Encyclopedia volume database failover is the responsibility of the third-party RDBMS administrator. The database administrator must use the facilities available in the RDBMS to configure failover capability.

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

Documentation for a PostgreSQL RDBMS is available at:

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

Documentation for an Oracle RDBMS is available at:

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

Documentation for Microsoft SQL Server RDBMS is available at:

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

Documentation for IBMDB2 RDBMS is available at:

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

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the Encyclopedia volume schemas and all related file data to ensure the recoverability in the event of failure. For more information on backing an iServer installation, see Chapter 9, “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 OpenText Support at the time of installation if you have any questions about the backup, recovery, or failover procedures necessary to protect against the possibility of catastrophic failure.

Managing an iServer cluster

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

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

Understanding the iServer System process model

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

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

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

- **Factory**
Executes requests to generate queries and documents and perform server-side printing.
- **View**
Supports viewing documents in DHTML and other output formats, such as Excel and PDF. Handles requests to download files from an Encyclopedia volume.

- **Integration**
Coordinates the running of information object (IOB) files that extract data from multiple data sources.
- **Caching**
Controls the Actuate Caching process that manages an information object cache and enables caching of data retrieved from data sources.

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

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

Understanding process flow in a stand-alone iServer

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

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

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

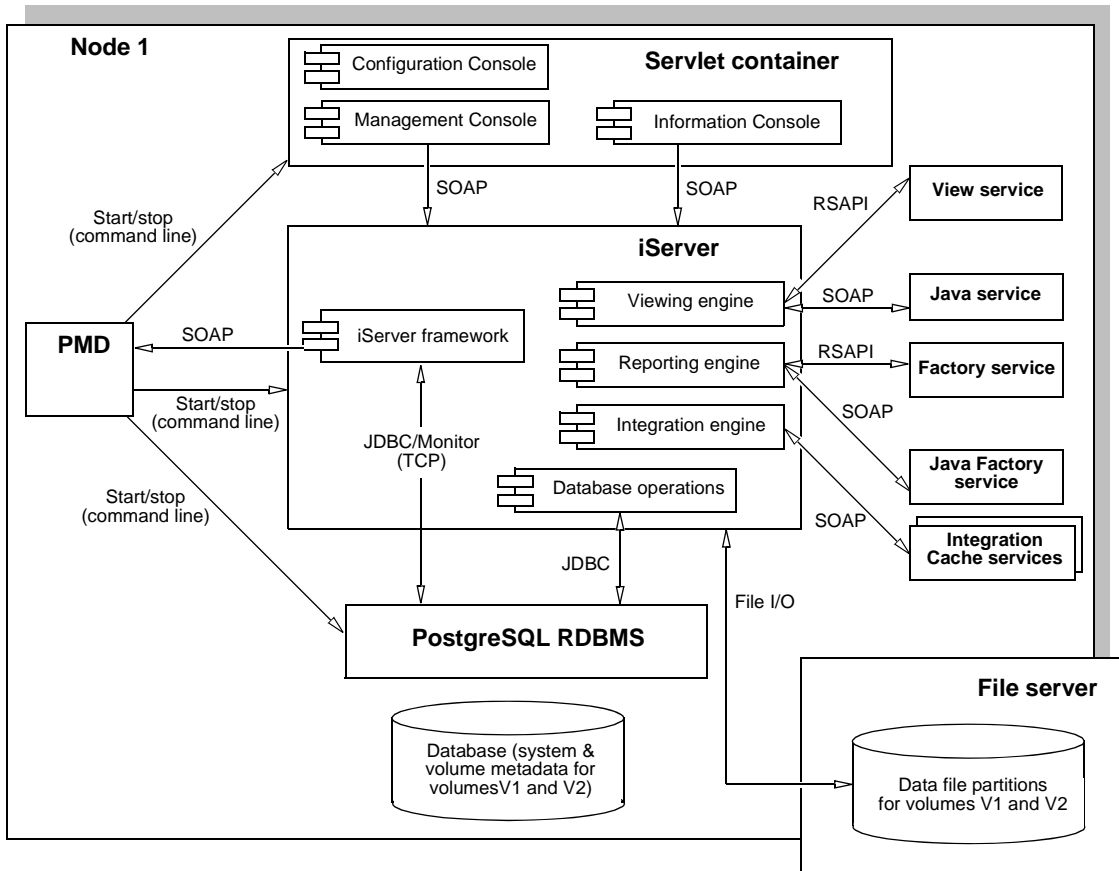


Figure 1-1 iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

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

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

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

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

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

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

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

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

The out-of-the-box (OOTB) iServer PostgreSQL installation configures the Encyclopedia volume database on the local disk to increase the reliability and performance of file input and output (I/O) operations. PostgreSQL discourages creating databases accessed using a Network File Systems (NFS) for these reasons. For more information, see section 17.2.1 Network File Systems at the following URL:

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

Understanding process flow in an iServer cluster

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

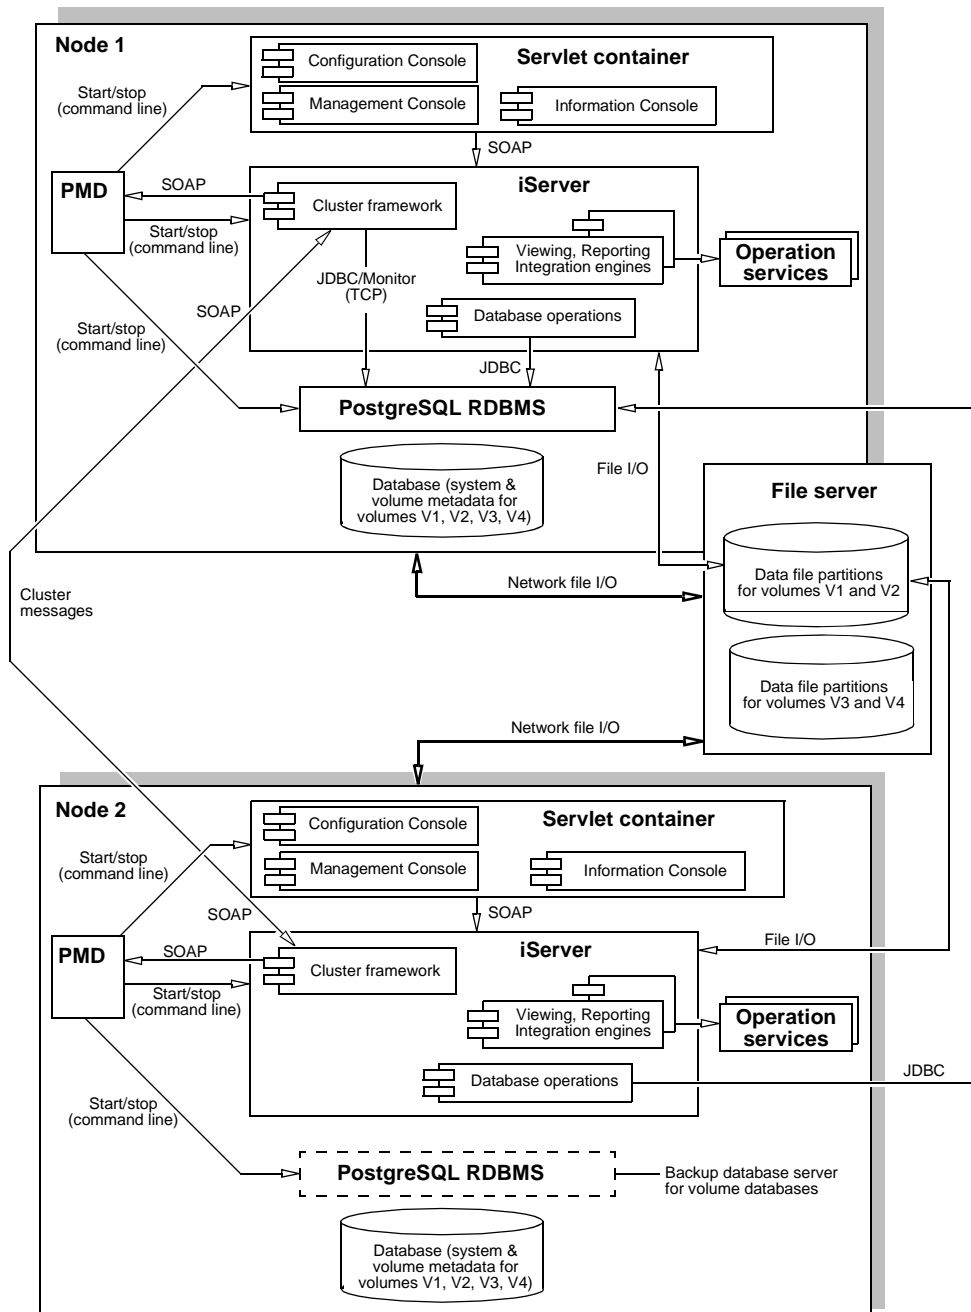


Figure 1-2 iServer RDBMS process architecture for a clustered, two-node, four-volume, OOTB database

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

For more information about the cloud computing deployment option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book. For more information about the cluster installation option, see *Configuring BIRT iServer*.

Administering iServer System

Administering an iServer System includes the following tasks:

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

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

- Scheduling jobs to run designs and generate documents

Each stand-alone iServer and node in an iServer cluster has a job scheduler and dispatcher. A job dispatcher send jobs to the local resource group factories.

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

- Reviewing logs and auditing the information to diagnose system problems
iServer can capture usage and error information in log files to assist an administrator in evaluating resource usage and troubleshoot problems. The usage and error logging applications are open framework applications, which are available as DLLs in Windows and shared libraries in Linux or UNIX.

- Configuring a cluster using automated installation programs and cloud computing base images

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

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

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

LDAP server, and installing and customizing usage and error logging and performance monitoring extensions.

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

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

- *Installing BIRT iServer for Windows or Installing BIRT iServer for Linux and UNIX*
Describes iServer System architecture. Provides detailed instructions on how to use automated installation programs and command-line utilities to install stand-alone iServer and clustered nodes that store Encyclopedia volume metadata in an external, third-party RDBMS, such as DB2, Oracle, PostgreSQL, or SQL Server. Also describes OpenText 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).

Using JDBC to connect to an Encyclopedia volume database

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

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

An administrator, who decides to customize iServer to connect to a database other than the OOTB PostgreSQL database, must ensure that the JDBC driver

returns adequate information about the types on the database. At a minimum, the database must return the common data types, such as integer, floating-point, and character. If the database does not return these common data types, then the database administrator must customize the database mapping framework to specify the types.

The JDBC driver must also support the following features:

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

When using connection pooling, the tracing functionality of the JDBC driver is used to capture the connection pool run-time statistics.

API Compatibility

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

About international character sets

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

Administrative reports

The default iServer Encyclopedia volume contains sample BIRT reports that provide information using the metadata and data extracted from the OOTB database, including job schedule, file, and user tracking and usage and error logging. Installing the sample volume is an option in a custom installation.

Supported operating systems

Actuate BIRT iServer Release 11 Service Pack 4 supports the following operating systems:

- Windows
- Solaris
- Linux

Part Two

Installing

Installing BIRT iServer

This chapter discusses the following topics:

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

Preparing to install BIRT iServer

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

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

Creating an account with Windows administrator privileges

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

The iServer user account must meet the following requirements:

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

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

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

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

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

How to create a Power User account

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

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

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

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

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

Configuring the iServer user account

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

How to configure a user account with administrator privileges

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

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

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

Configuring log on as a service privilege

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

How to configure the log on as a service privilege manually

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

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

Backing up iServer system and Encyclopedia volume metadata

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

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult OpenText 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 9, “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 OpenText My Support:

<https://support.opentext.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 Service Pack 5 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 Service Pack 5

The following section describes how to install a new instance of BIRT iServer Release 11 Service Pack 5 (11SP5) in the Windows operating system.

How to perform a new installation of BIRT iServer Release 11SP5 in Windows

To install iServer, perform the following tasks:

- 1 Download the iServer distribution package. Extract the files. Run the self-extracting executable file, ActuateBIRTiServerEnterprise.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.

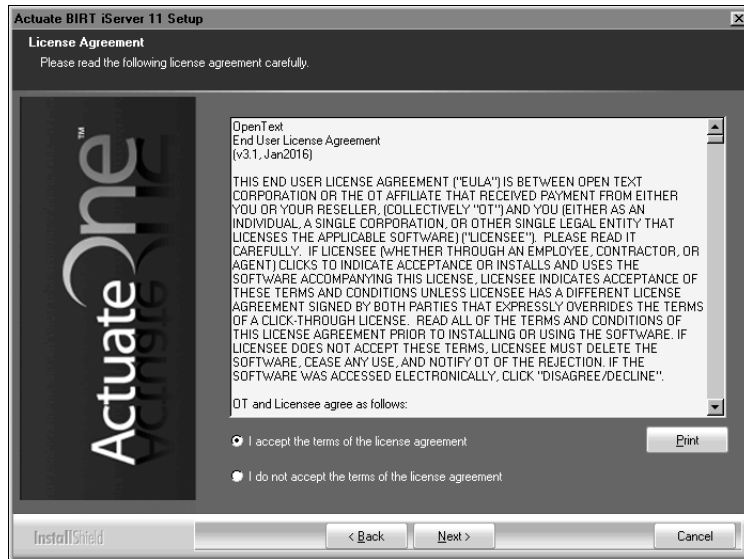


Figure 2-2 Accepting the license agreement

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

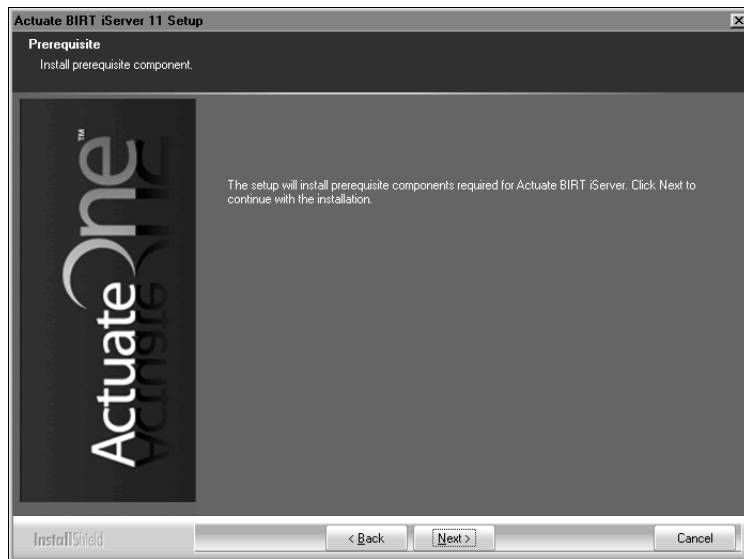


Figure 2-3 Installing Prerequisites

- 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 (x86)\Actuate11SP5\iServer.

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

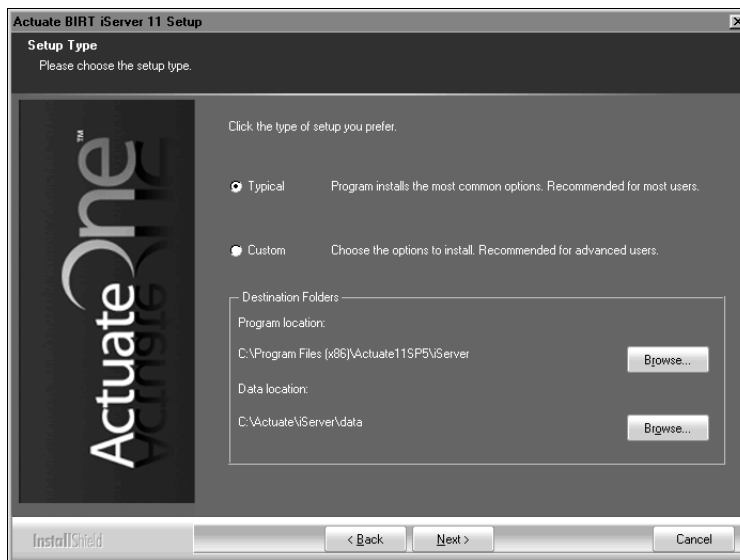


Figure 2-4 Specifying typical or custom setup type

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

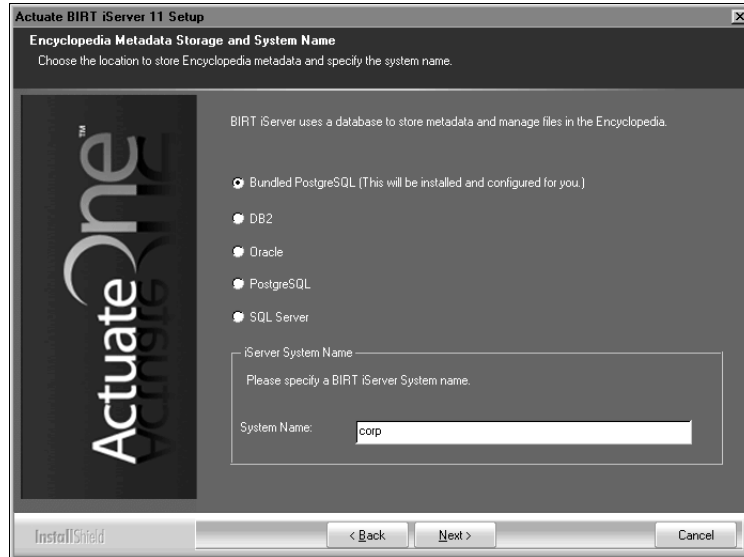


Figure 2-5 Selecting the Encyclopedia volume metadata database to install

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

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

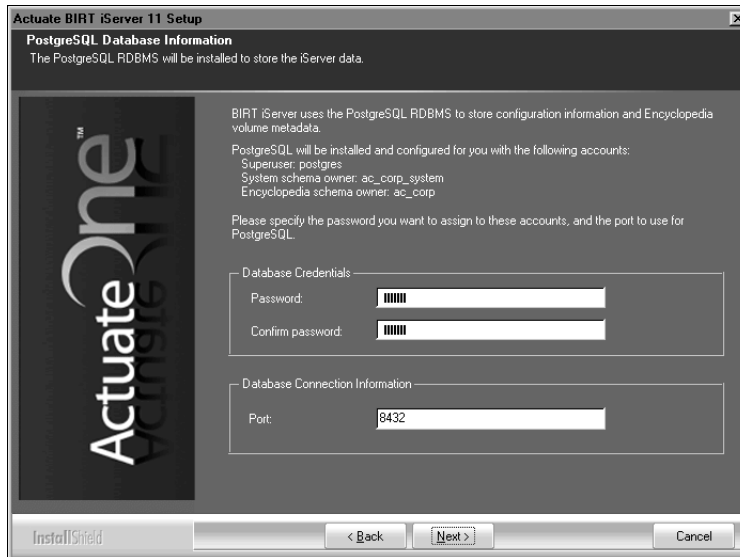


Figure 2-6 Specifying PostgreSQL database information

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

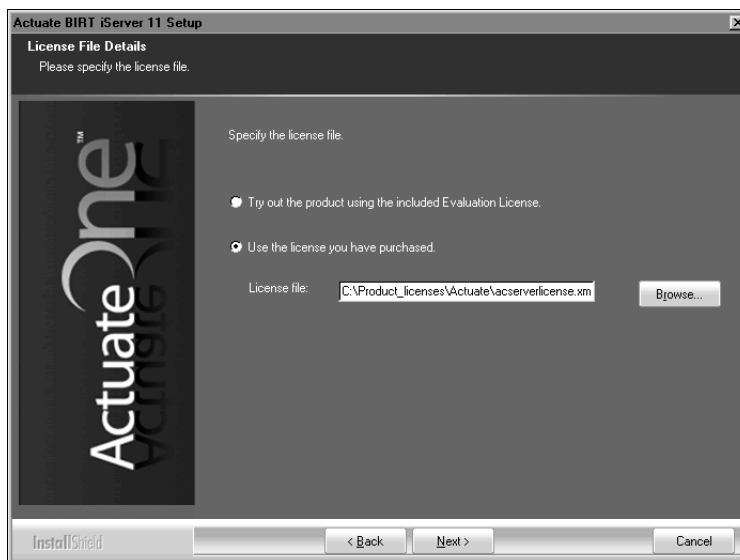


Figure 2-7 Specifying the license file

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




Figure 2-8 Viewing the named-user license question

Choose Yes to continue the installation.

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

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

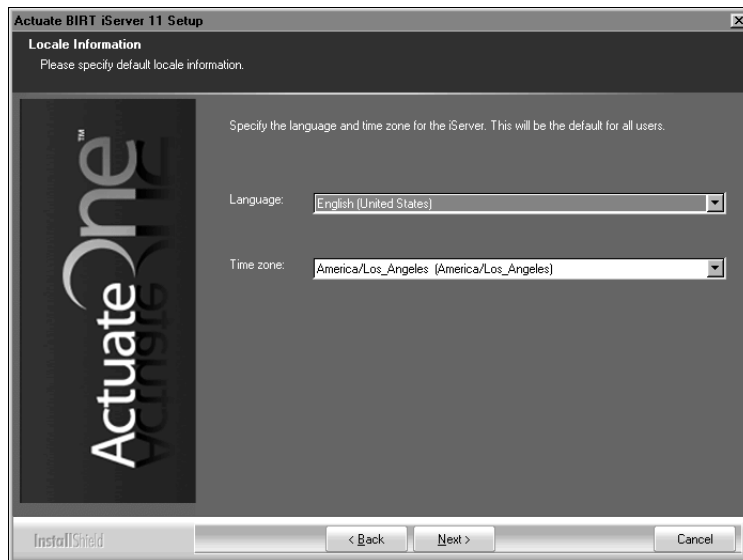


Figure 2-9 Specifying locale information

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

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

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

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

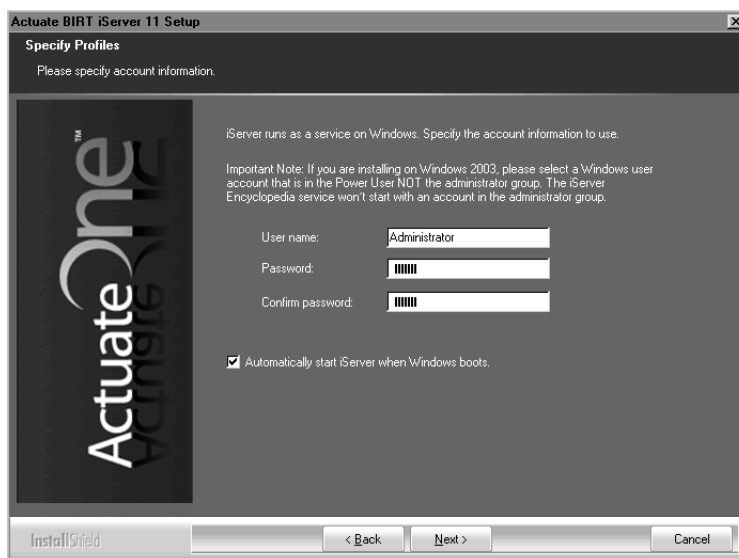


Figure 2-10 Specifying an account for running the iServer service

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




Figure 2-11 Setting the Windows local security policy

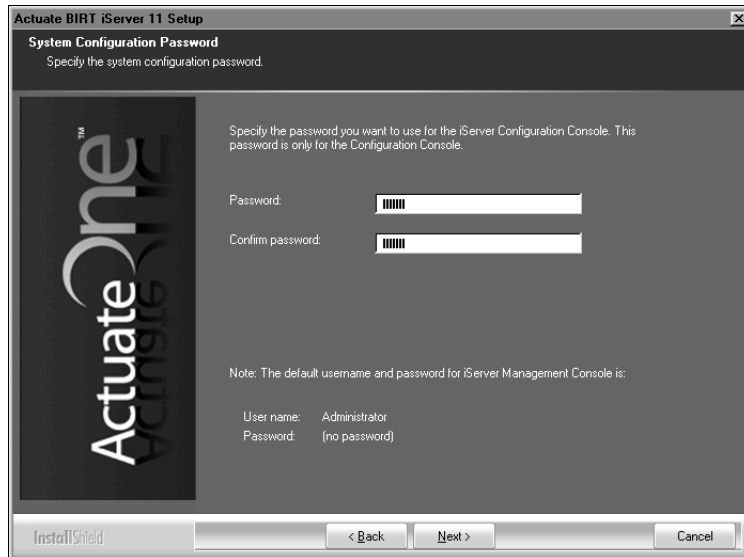


Figure 2-12 Specifying the password for using Configuration Console

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

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

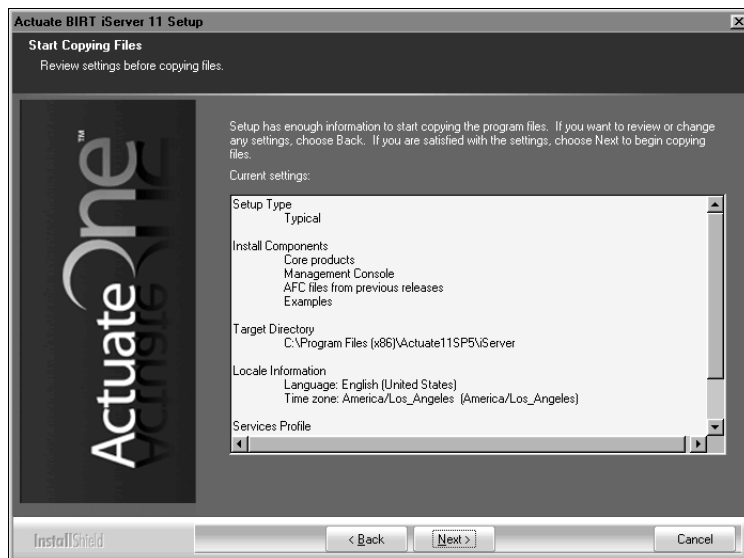


Figure 2-13 Reviewing settings before copying files



Figure 2-14 Viewing setup status

During installation, a command prompt appears, displaying the run-time commands and messages from the initialization process of the PostgreSQL RDBMS and Encyclopedia volume schema. Listing 2-1 shows the contents of this command prompt.

Listing 2-1 Contents of command prompt

```
C:\DOCUME~1\ADMINI~1\LOCALS~1\Temp\{C774ECE4-911A-470F-A9C2-CFD
4FE0CCEEA}\{0EC9D796-3F68-43BE-82D0-B05A7799EE4B}"C:\Program Fi
les (x86)\Actuate11SP5\iServer\postgresql\bin\initDB.exe" -A pa
ssword -U postgres --pwfile="C:\ProgramFiles\Actuate11SP5\iServ
er\pwtemp.txt" -E UTF8 --locale="C" -D "C:\Actuate\iServer\data
\encyc\postgresql\data" 1>>"C:\Actuate\iServer\data\server\log
\initDB.log" 2>&1
```

Choose Finish to exit the wizard, as shown in Figure 2-15.



Figure 2-15 Exiting the installation wizard

- 12** The installation program prompts you to install the online help from the following location, as shown in Figure 2-16:

<http://www.actuate.com>




Figure 2-16 Viewing the install online help and manuals prompt

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

The installation program installs shortcuts on the desktop, as shown in Figure 2-17.




Figure 2-17 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 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 a PostgreSQL database not residing in the AC_DATA_HOME/encyc directory, nor 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 upgrading or migrating an earlier Release 11 iServer to Release 11 Service Pack 5

To upgrade a Release 11 iServer earlier than Release 11 Service Pack 4 Fix 5, (11SP4 Fix 5) you must first upgrade to Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5), then upgrade to Release 11 Service Pack 5 (11SP5). To upgrade from an earlier Release 11 iServer to 11SP4 Fix 5, perform an automatic in-place upgrade, as described in “Performing an automatic in-place upgrade from an earlier Release 11 iServer to Release 11SP4 Fix 5.” To upgrade from 11SP4 Fix 5 to 11SP5, perform an automatic in-place upgrade as described in “Performing an automatic in-place upgrade from iServer Release 11SP4 Fix 5 to iServer Release 11SP5,” or a side-by-side migration, as described in “Performing a side-by-side migration from iServer Release 11SP4 Fix 5 to iServer Release 11SP5.” All the upgrade procedures are in Chapter 4, “Upgrading BIRT iServer.”

The in-place upgrade programs preserve any previous iServer configuration information and reuse the earlier settings. The upgrade programs resolve any differences in default values between releases, ignoring old configuration defaults in favor of new default values. For example, the upgrade programs do not prompt the user for port information and machine name. The programs detect the current port numbers and machine name and keep those settings.

About running different releases on the same machine

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 5 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 Release 11SP5 when the testing is complete. Earlier iServer 11 releases and Release 11SP5 cannot coexist on the same machine.

How and when you upgrade to Release 11SP5 depends on your site configuration and requirements. Complete the following general tasks in this order to determine how to upgrade your site to Release 11SP5:

- Create a test environment for Release 11SP5. The test environment must be 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 Release 11SP5 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 Release 11SP5 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 on a separate machine. You can install all Release 11SP5 products or the Release 11SP5 server products and a subset of the desktop products.

If you plan to test Release 11SP5 desktop products, identify which users to include in the final testing. Developers and users can then confirm that applications perform as expected in the Release 11SP5 production staging environment.

Complete the following general tasks to test Release 11SP5:

- Install BIRT iServer Release 11SP5 software in a production staging area.
- Install Release 11SP5 desktop software on the test user machines.
- Verify that the Release 11SP5 production staging environment works correctly.
- Install the remaining Release 11SP5 desktop products, if you installed a subset earlier.
- Verify that all the Release 11SP5 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 Release 11SP5 environment. Set up the production environment and schedule a date and time to switch from earlier versions to Release 11SP5.

When you switch to Release 11SP5, use the following procedure list as a general guideline:

- Shut down all Actuate servers.
- Back up earlier Actuate Encyclopedia volumes.
- Upgrade existing Encyclopedia volume schemas if necessary. Install upgraded design and document files.
- Start BIRT iServer Release 11SP5.
- Inform users that they can start using Release 11SP5 design tool products.

Installing BIRT iServer using an alternative database

This chapter discusses the following topics:

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

Preparing to install BIRT iServer using an alternative database

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

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

Creating an account with Windows administrator privileges

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

The iServer user account must meet the following requirements:

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

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

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

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

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

Restrict schema and the iServer application user names to alphanumeric and underscore characters with an initial alphabetic character in the pattern `[a-z][a-z 0-9]*`. Do not use a hyphen.

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

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

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

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

Creating a database

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

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

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

```
CREATE LANGUAGE plpgsql;
```

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

Creating the system schema owner

In an iServer installation, the system schema owner must have the same name as the system schema. The system schema owner has all privileges on the schema

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

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

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

Creating the Encyclopedia volume schema owner

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

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

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

Creating the iServer application user

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

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

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

Creating the system schema

The system schema must have the same name as the system schema owner. The following commands create a system schema named `ac_corp_system`, owned by the user, `ac_corp_system`, then grant privileges to use that schema to the application user role named `iserver`. Connect to the `iserver` application database,

not the PostgreSQL system database, as a user with full administrator privileges and execute the following commands:

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

Creating the Encyclopedia volume schema

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

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

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

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

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

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

Creating the system schema owner

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

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

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

```

GRANT CREATE TABLE TO ac_corp_system;
GRANT CREATE VIEW TO ac_corp_system;
GRANT CREATE SEQUENCE TO ac_corp_system;
GRANT CREATE TYPE TO ac_corp_system;
GRANT CREATE PROCEDURE TO ac_corp_system;
GRANT CREATE OPERATOR TO ac_corp_system;
GRANT CREATE TRIGGER TO ac_corp_system;
GRANT CREATE SESSION TO ac_corp_system;
ALTER USER ac_corp_system QUOTA UNLIMITED ON USERS;
COMMIT;

```

Creating the Encyclopedia volume schema owner

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

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

```

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

```

Creating the iServer application user

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

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

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

```
DROP USER iserver CASCADE;
CREATE USER iserver
  IDENTIFIED BY password
  DEFAULT TABLESPACE USERS
  TEMPORARY TABLESPACE TEMP;
GRANT CREATE SESSION TO iserver;
ALTER USER iserver QUOTA UNLIMITED ON USERS;
COMMIT;
```

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

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

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

Creating a database

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

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

Any database created for iServer processing must use a case-insensitive collation, such as SQL_Latin1_General_CP1_CI_AS. The names of case-insensitive collations typically include the letters, CI.

Creating the system schema owner

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

The following commands create a user named ac_corp_system to function as the system schema owner with appropriate privileges to connect to the previously

created iserver database. Connect to the iserver database as a user with full administrator privileges and execute the following SQL commands:

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

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

Creating the Encyclopedia volume schema owner

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

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

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

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

Creating the iServer application user

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

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

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

Creating the system schema

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

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

Creating the Encyclopedia volume schema

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

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

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

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

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

Creating user accounts

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

Creating a database

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

Actuate also requires a DB2 database to use a case-insensitive collation, such as UCA500R1_LEN_S2, which is not the default. DB2 supports this functionality only in DB2 9.5 Fix Pack 1 and later versions.

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

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

Creating the system schema owner

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

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

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

Creating the Encyclopedia volume schema owner

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

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

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

Creating the iServer application user

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

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

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

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

Creating the system schema

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

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

Creating the Encyclopedia volume schema

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

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

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

In a stand-alone iServer installation that uses DB2 to store Encyclopedia volume metadata, the pound sign ('#') in iServer is treated as a single alphanumeric

wildcard character instead of a digit wildcard character. This limitation exists because DB2 does not support regular expressions in SQL syntax.

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

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

The solution consists of the following parts:

- Create a Java implementation that performs the regular-expression testing functionality. In the article, IBM provides the sample java code for such an implementation. See the Implementation section in the article. IBM also provides the pre-built java package, `db2_regex.jar`, for download. See the Download section in the article.
- Create a UDF to access the external Java method. Use the following sample implementation as an example:

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

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

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

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

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

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

```
<MatchOpMapper SingleMatch="_"
    GreedyMatch="%"
    DigitMatch="_"
    EscapeTemplate="@ $"
    AdditionalSpecialChars="@ ">
    <FunctionMappings>
        <FunctionMapping FunctionName="MATCH">
            $P0 LIKE $P1 ESCAPE '@'
        </FunctionMapping>
    </FunctionMappings>
</MatchOpMapper>
```

to the following:

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

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

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

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

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

To resolve this problem, perform the following tasks:

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

```
CREATE DATABASE iserver
WITH OWNER = "postgres"
TEMPLATE = template0 ENCODING = 'UTF-8'
LC_COLLATE = 'English, United States'
LC_CTYPE = 'English, United States'
CONNECTION LIMIT = -1;
```

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

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

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

Backing up iServer system and Encyclopedia volume metadata

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

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult OpenText 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 9, “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 OpenText My Support:

<https://support.opentext.com>

OpenText also supports the cloud deployment of BIRT iServer using a ready-to-launch iServer image. For more information about this installation option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book. For information about the new Release 11 BIRT iServer System architecture, see Chapter 1, “Understanding Actuate BIRT iServer architecture,” earlier in this book.

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

Installing an Encyclopedia volume that uses an alternative database

The following procedures use a pre-existing PostgreSQL database and schema as an example. During the iServer installation, the administrator provides the schema owner and database user credentials. The iServer installation program creates the necessary volume database structures, then loads the metadata.

How to install an Encyclopedia volume that uses an alternative database

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



Figure 3-1 Viewing the welcome message

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



Figure 3-2 Accepting the license agreement

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

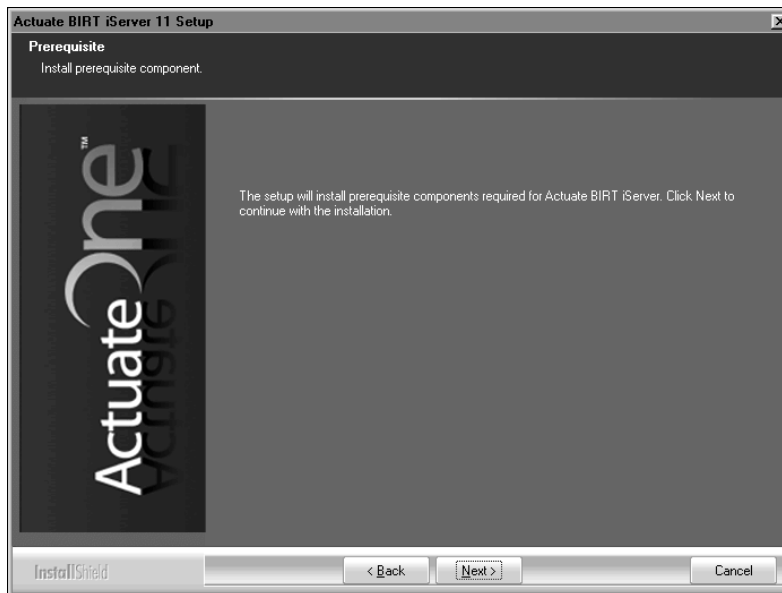


Figure 3-3 Installing prerequisites

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

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

iServer uses the Program location to resolve the paths to all the binaries that it launches. The environment variable, AC_SERVER_HOME, points to the location of the iServer binaries. The default path for the program location is C:/Program Files (x86)/Actuate11SP5/iServer.

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

Choose Next.

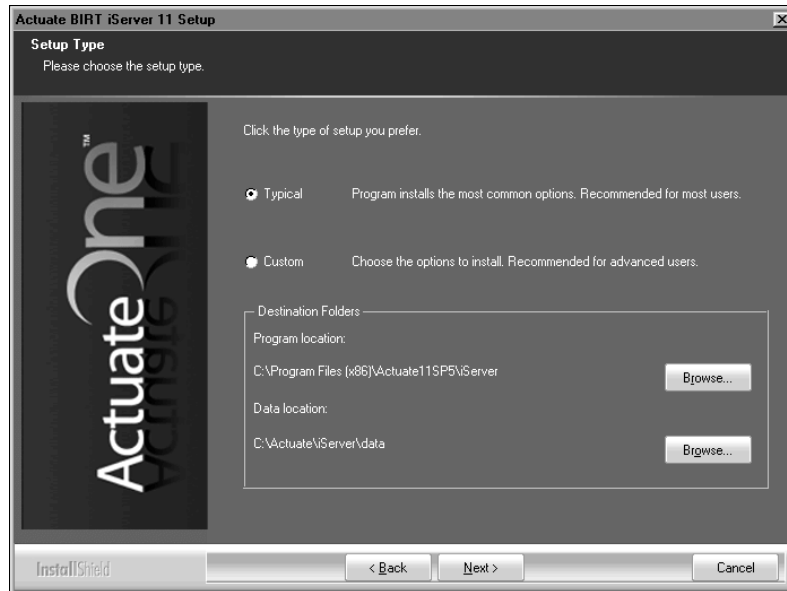


Figure 3-4 Specifying Typical setup type

- 5 In Encyclopedia Metadata Storage, select DB2, Oracle, PostgreSQL, or SQL Server to use as a pre-existing third-party database for storing Encyclopedia volume metadata. This example uses the PostgreSQL RDBMS, as shown in Figure 3-5. In System Name, type a name for the BIRT iServer system. Restrict the name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.

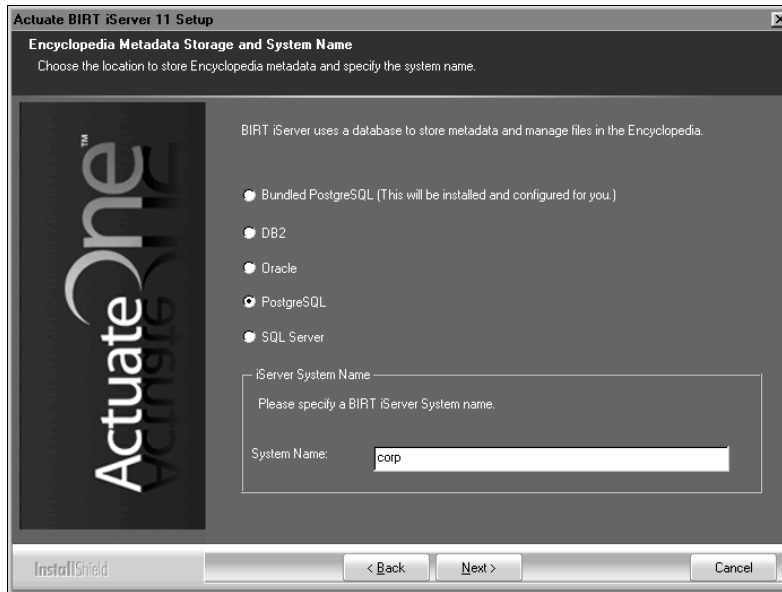


Figure 3-5 Choosing a third-party database

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

- **DB2**

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

- **Oracle**

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

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

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

Actuate BIRT iServer 11 Setup

3rd Party Database Information
Please specify the external DB2 database information for the iServer encyclopedia.

Server:

Database:

Port:

IANAAppCodePage:

Database User Credentials

User name:

Password:

Confirm password:

InstallShield

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

Actuate BIRT iServer 11 Setup

3rd Party Database Information
Please specify the Oracle database information for the iServer encyclopedia.

Server:

Port:

TNS Server Name:

Service Name:

TNS Names File:

Database User Credentials

User name:

Password:

Confirm password:

InstallShield

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

- PostgreSQL

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

- SQL Server

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

Actuate BIRT iServer 11 Setup

3rd Party Database Information
Please specify the external PostgreSQL database information for the iServer encyclopedia.

Server:

Port:

Database:

Database User Credentials

User name:

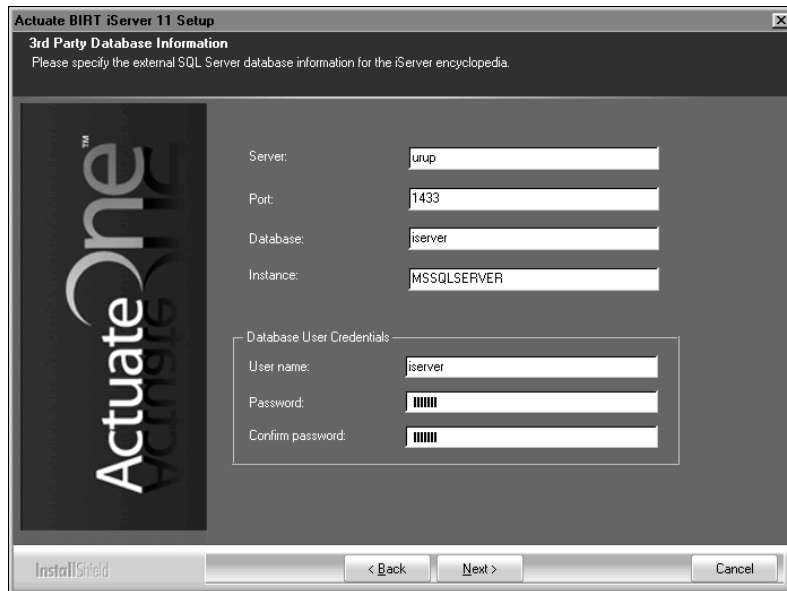
Password:

Confirm password:

InstallShield

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

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



Actuate BIRT iServer 11 Setup

3rd Party Database Information
Please specify the external SQL Server database information for the iServer encyclopedia.

Server:

Port:

Database:

Instance:

Database User Credentials

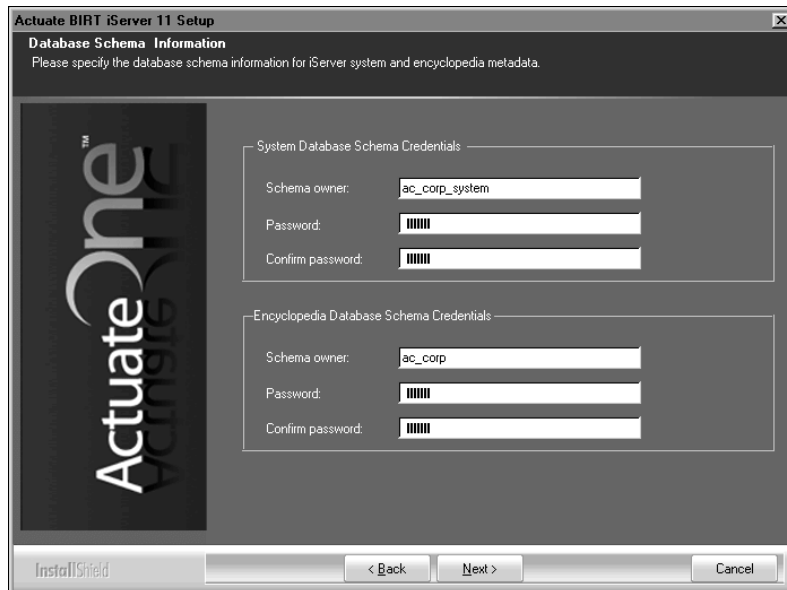
User name:

Password:

Confirm password:

InstallShield

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



Actuate BIRT iServer 11 Setup

Database Schema Information
Please specify the database schema information for iServer system and encyclopedia metadata.

System Database Schema Credentials

Schema owner:

Password:

Confirm password:

Encyclopedia Database Schema Credentials

Schema owner:

Password:

Confirm password:

InstallShield

Figure 3-10 Specifying System and Encyclopedia schema passwords

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

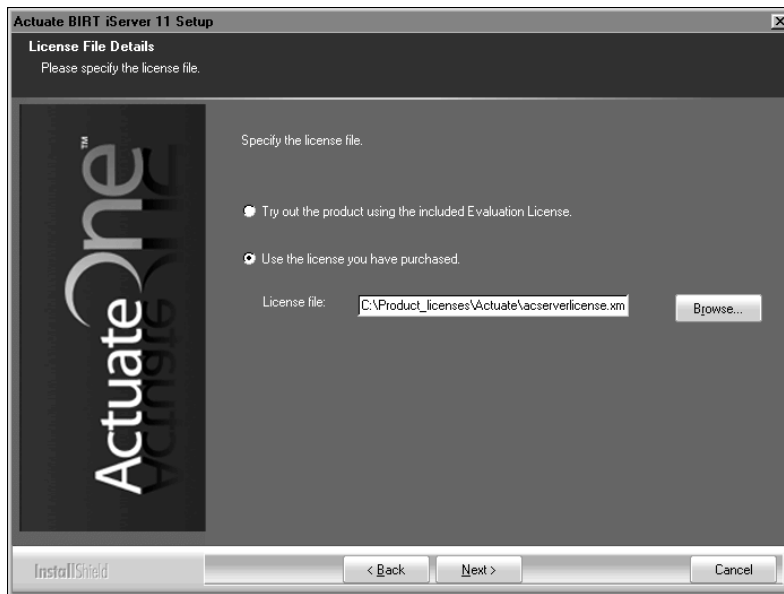


Figure 3-11 Specifying the license file

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




Figure 3-12 Viewing the named user license question

Choose Yes to continue the installation.

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

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

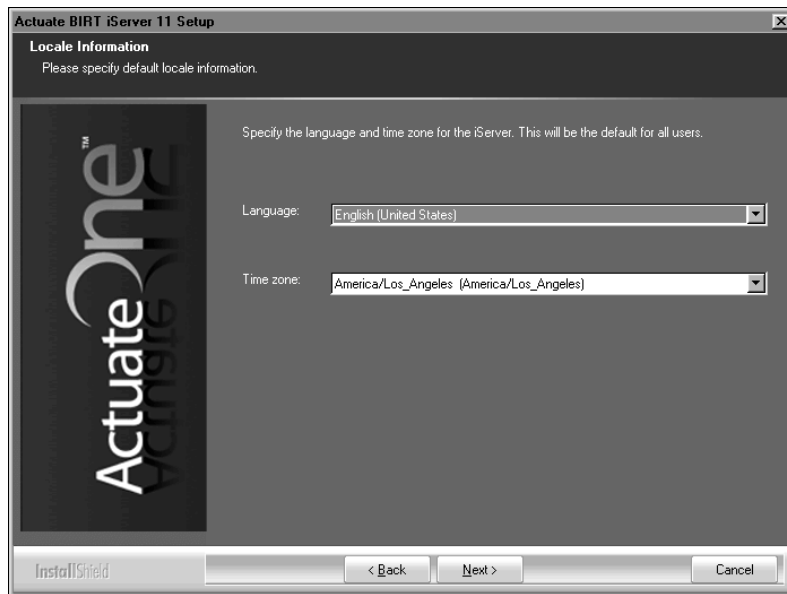


Figure 3-13 Specifying locale information

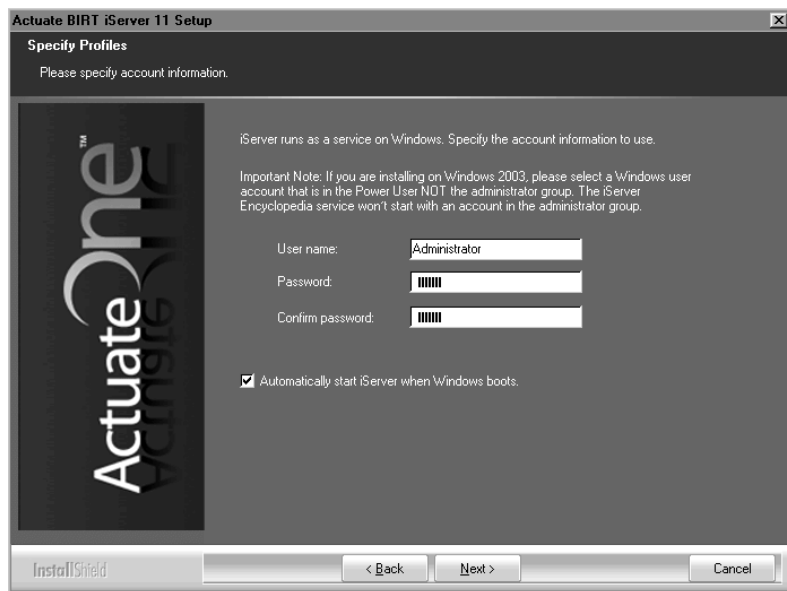


Figure 3-14 Specifying the account for running the iServer service

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




Figure 3-15 Setting the Windows local security policy

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

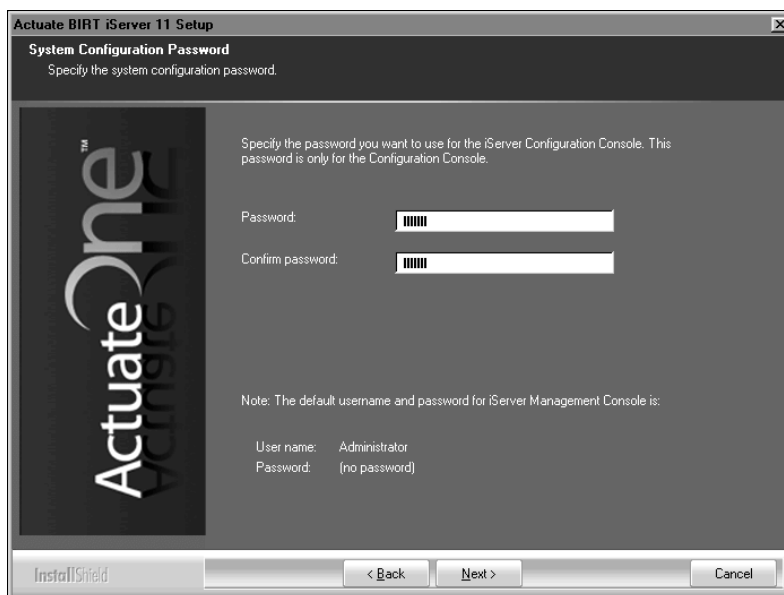


Figure 3-16 Specifying the password for using Configuration Console

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



Figure 3-17 Reviewing settings before copying files



Figure 3-18 Viewing setup status

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



Figure 3-19 Exiting the installation wizard

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

<http://www.actuate.com>




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

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

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




Figure 3-21 Viewing iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

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

- iServer 11

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

Upgrading BIRT iServer

This chapter discusses the following topics:

- Upgrading BIRT iServer
- Preparing to upgrade BIRT iServer
- Performing an automatic in-place upgrade from an earlier Release 11 iServer to Release 11SP4 Fix 5
- Performing an automatic in-place upgrade from iServer Release 11SP4 Fix 5 to iServer Release 11SP5
- Performing a side-by-side migration from iServer Release 11SP4 Fix 5 to iServer Release 11SP5
- Working with iServer utilities

Upgrading BIRT iServer

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

- **Automatic in-place upgrade**
Upgrades an earlier Release 11 iServer system, such as Release 11 Service Pack 3, to iServer Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5), in place. Also upgrades iServer 11SP4 Fix 5 to iServer Release 11 Service Pack 5 (11SP5) in place. Automatically migrates one or more existing Encyclopedia volumes during the installation process. To upgrade from a Release 11 iServer system earlier than 11SP4 Fix 5 to 11SP5, you must first upgrade to 11SP4 Fix 5, then upgrade from 11SP4 Fix 5 to 11SP5.
- **Manual side-by-side migration**
Migrates an iServer Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5) system to an iServer Release 11 Service Pack 5 (11SP5) system. The administrator installs the 11SP5 system on a separate machine from the machine hosting 11SP4 Fix 5. After installation, the administrator uses the Data Store Administrator utility to migrate the volumes manually from the 11SP4 Fix 5 installation to the 11SP5 installation. The side-by-side migration procedure also supports migrating volumes for which the volume metadata resides in a pre-existing data store, such as DB2, Microsoft SQL Server, Oracle, or a pre-existing PostgreSQL database.

For upgrading from an 11SP4 Fix 5 system to an 11SP5 system, OpenText recommends the side-by-side, manual migration procedure as a best practice since this process does not install in the same location as the previous installation. For more information about manually migrating in a side-by-side installation scenario, see “Performing a side-by-side migration from iServer Release 11SP4 Fix 5 to iServer Release 11SP5,” later in this chapter.

Preparing to upgrade BIRT iServer

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

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

Creating an account with Windows administrator privileges

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

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

The iServer user account must meet the following requirements:

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

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

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

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

Backing up iServer system and Encyclopedia volume metadata

When upgrading iServer using an automatic in-place upgrade procedure, 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 OpenText 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 9, “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 OpenText My Support:

<https://support.opentext.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 Service Pack 5 as an upgrade using the available installation options.

Performing an automatic in-place upgrade from an earlier Release 11 iServer to Release 11SP4 Fix 5

When upgrading in-place from an earlier Release 11 iServer, such as Service Pack 2 or 3, to iServer Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5), the administrator upgrades iServer using the automatic upgrade process. The automatic upgrade program performs the following tasks:

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

- Brings the volume or volumes online

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

The following sections describe how to perform an automatic in-place upgrade of an earlier minor BIRT iServer release, such as Release 11 Service Pack 3, to Release 11 Service Pack 4 Fix 5.

Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services

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

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

How to stop the Actuate 11 BIRT iServer service

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

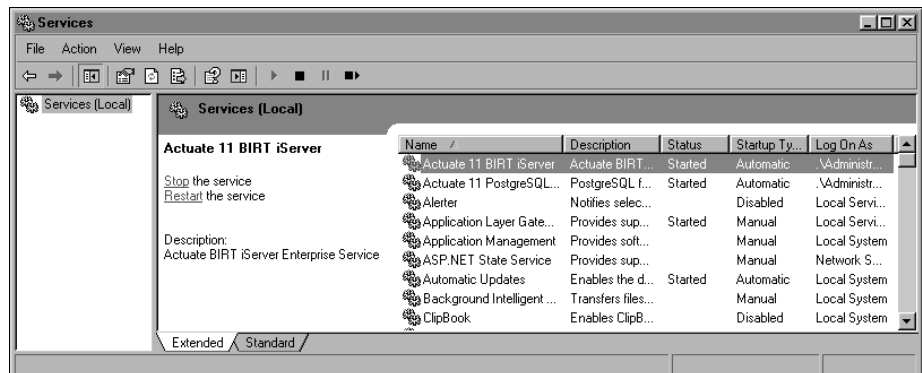


Figure 4-1 Stopping the BIRT iServer service

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

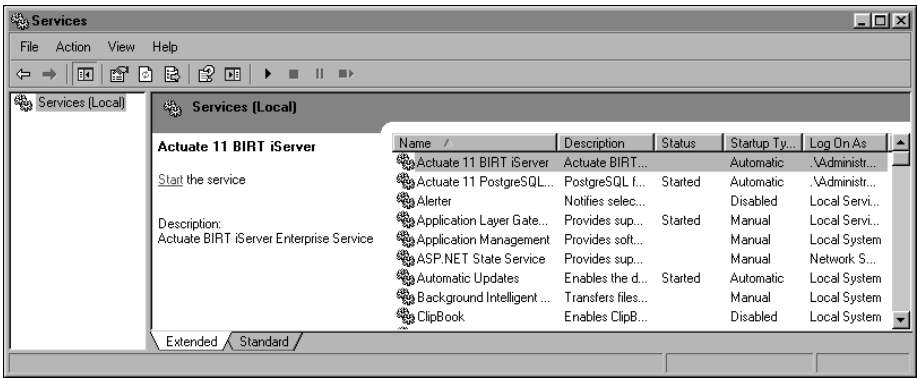


Figure 4-2 Viewing Services after BIRT iServer service stops

How to shut down the Actuate 11 PostgreSQL for BIRT service

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

- 1 Choose Start→Programs→Administrative Tools→Services.
On Control Panel, choose Administrative Tools.
On Administrative Tools, choose Services.
- 2 On Services, select Actuate 11 PostgreSQL for BIRT Service, as shown in Figure 4-3.

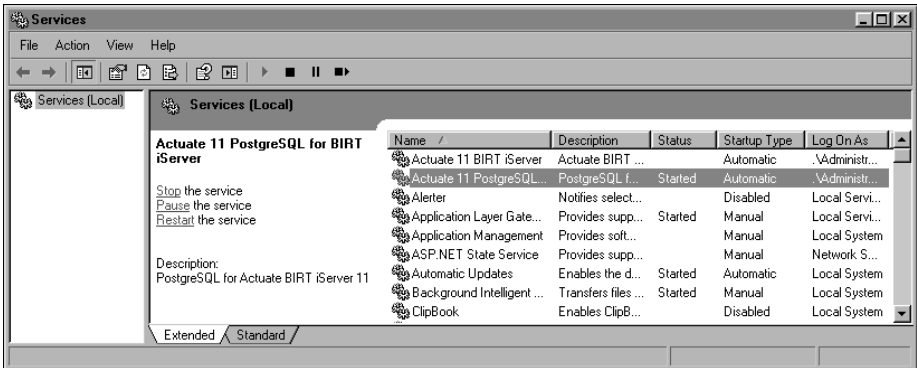


Figure 4-3 Shutting down the PostgreSQL for BIRT iServer service

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

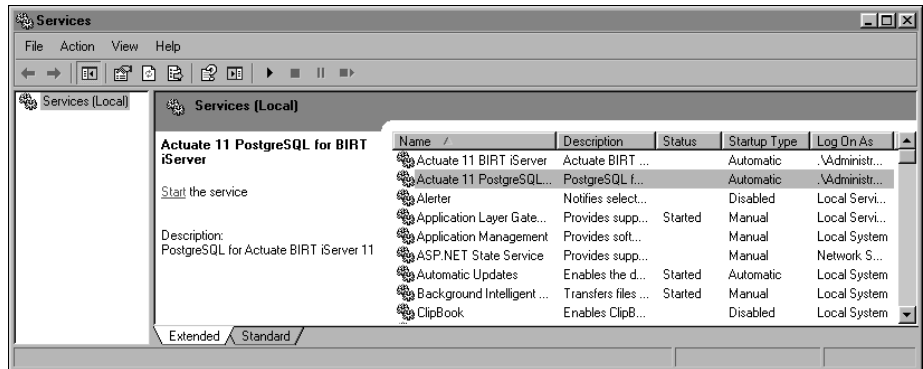


Figure 4-4 Viewing Services after PostgreSQL service stopped

How to shut down the NobleNet Portmapper

Stop NobleNet Portmapper by performing the following tasks:

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

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

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

Running the in-place upgrade on an earlier Release 11 iServer

The following procedure describes step-by-step how to perform an automatic upgrade in place from an earlier Release 11 iServer, such as Service Pack 2 or 3, to Release 11 Service Pack 4 Fix 5.

How to run the in-place upgrade on an earlier Release 11 iServer

- 1 Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
 - encyc directories from all nodes
 - acserverconfig.xml in the AC_DATA_HOME/config directory
 - acpmdconfig.xml in the AC_SERVER_HOME/etc directory
 - RSSE code and associated files if you use the Open Security option
- 2 Download the iServer distribution package. Extract the files. Run the self-extracting executable file, ActuateBIRTiServerEnterprise.exe. The welcome message appears, as shown in Figure 4-5. Choose Next.



Figure 4-5 Viewing the welcome message

- 3 Read and accept the license agreement, as shown in Figure 4-6. Choose Next.



Figure 4-6 Viewing the license agreement

- 4 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-7. Choose Next.

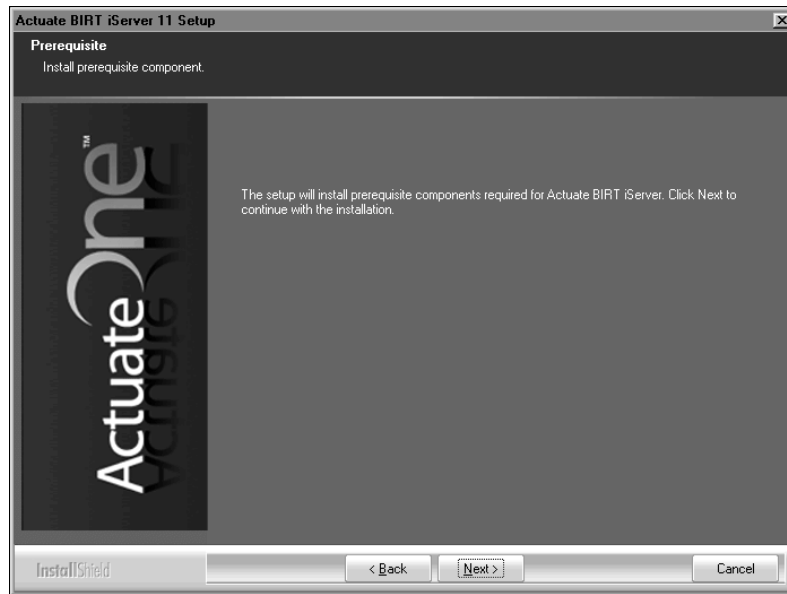


Figure 4-7 Installing Prerequisites

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

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

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

The installer detects the previous Release 11 installation and selects Upgrade Existing iServer and Migrate the Encyclopedia Data. The installer lists the Existing iServer Location, as shown in Figure 4-8.

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

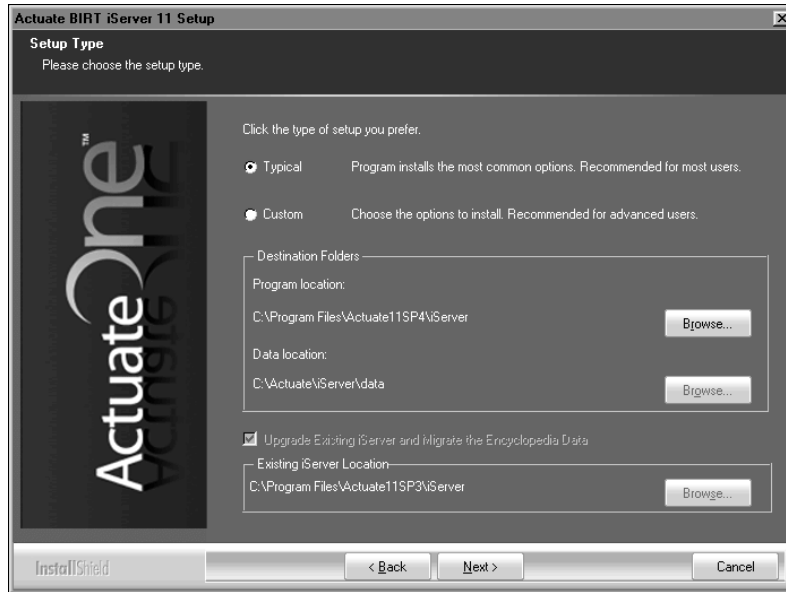


Figure 4-8 Choosing setup type

A message appears stating that the install program is migrating a previous version of the Encyclopedia volume, which might take a long time, as shown in Figure 4-9. Choose OK.



Figure 4-9 Confirming volume migration

- 6 If you have not shut down the BIRT iServer service, a prompt appears, instructing you to shut down the service, as shown in Figure 4-10. Perform the following tasks:
 - 1 Choose OK on the prompt shown in Figure 4-10.
 - 2 Stop the Actuate 11 BIRT iServer service.
 - 3 On Setup Type, choose Next.




Figure 4-10 Choosing OK to shut down iServer

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

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

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

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




Figure 4-11 Choosing OK to shut down the PostgreSQL service

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

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

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 4-12. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

- 9 In System Configuration Password, type a password for Configuration Console, as shown in Figure 4-13. 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.

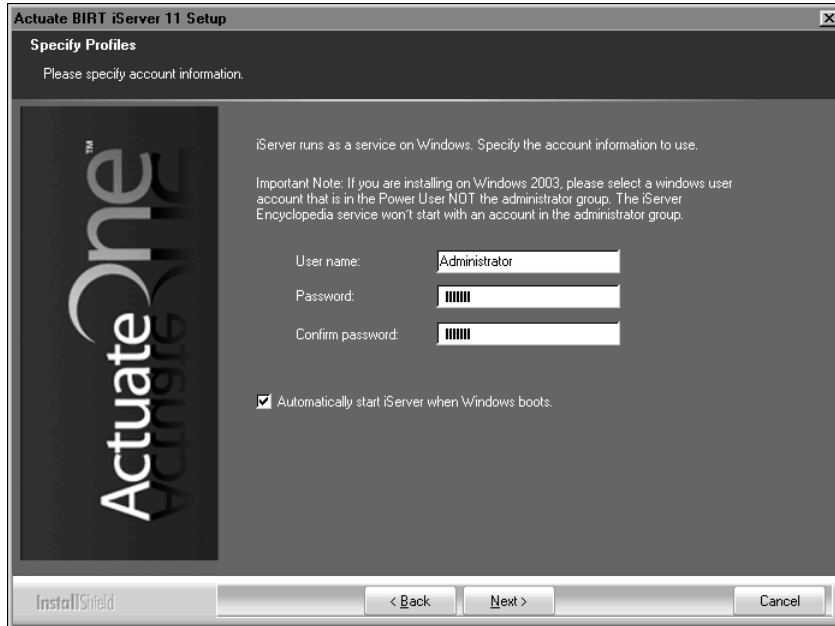


Figure 4-12 Specifying the account for running the iServer service

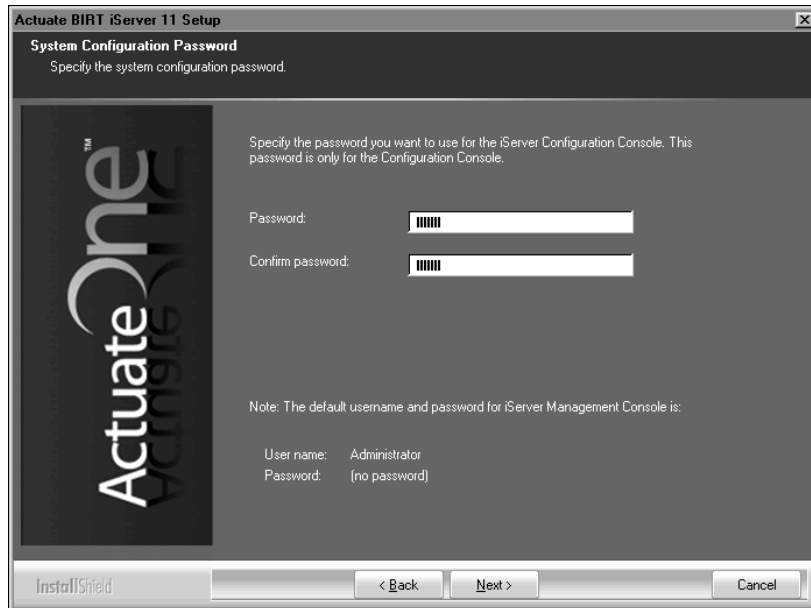


Figure 4-13 Specifying the password for using Configuration Console

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



Figure 4-14 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-15.



Figure 4-15 Viewing setup status

During installation, a command prompt appears, displaying messages regarding installation run-time processing, as shown in Figure 4-16.

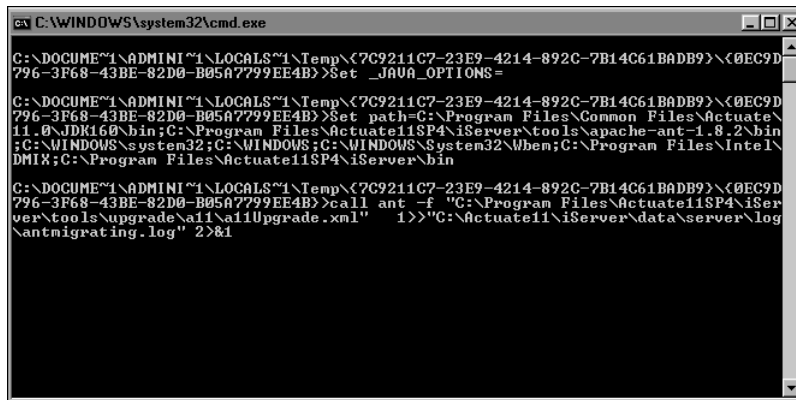


Figure 4-16 Viewing run-time processing

How to install pgAdmin for an installation that uses PostgreSQL RDBMS

If you are upgrading an iServer installation that uses an alternative RDBMS, such as DB2, Oracle, or Microsoft SQL Server, skip this section and go to “How to finish installing Actuate BIRT iServer 11,” later in this chapter.

For an installation that uses the PostgreSQL RDBMS, a prompt appears, asking whether to install pgAdmin, the administration tool for the PostgreSQL RDBMS, as shown in Figure 4-17. To install pgAdmin, perform the following tasks:

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




Figure 4-17 Choosing to install the pgAdmin tool

Actuate BIRT iServer Setup appears with pgAdmin III Setup appearing immediately afterward, as shown in Figure 4-18.



Figure 4-18 Viewing Actuate BIRT iServer and pgAdmin III Setup

- 2 In pgAdmin III Setup, perform the following tasks:
 - 1 In Welcome, shown in Figure 4-19, choose Next.

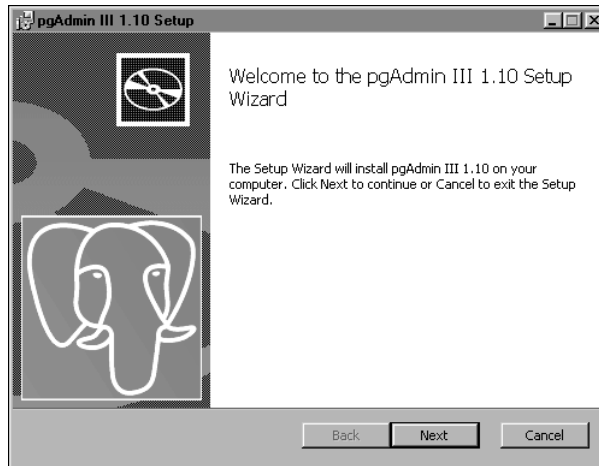


Figure 4-19 Viewing Welcome

- 2 In License Agreement, select I accept the terms in the License Agreement, as shown in Figure 4-20. Choose Next.



Figure 4-20 Accepting the license agreement

- 3 In Custom Setup, review the features to be installed, as shown in Figure 4-21. Choose Next.

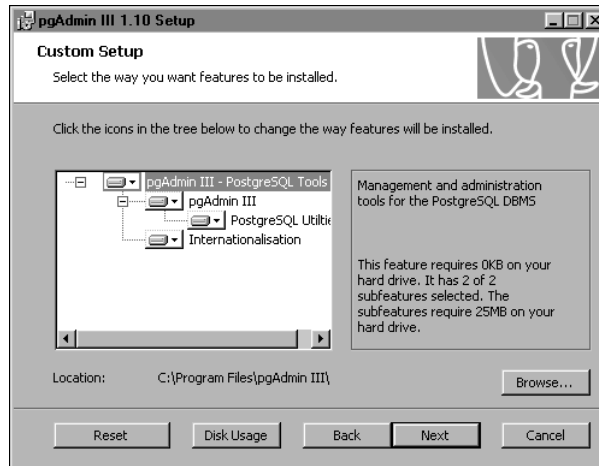


Figure 4-21 Viewing the features to be installed

- 4 In Ready to Install pgAdmin III, shown in Figure 4-22, choose Install.




Figure 4-22 Choosing to install pgAdmin II

Installing pgAdmin III appears, as shown in Figure 4-23.




Figure 4-23 Installing pgAdmin III

- 5 When Completed, the pgAdmin III Setup Wizard appears, as shown in Figure 4-24, choose Finish to exit the pgAdmin III Setup Wizard.




Figure 4-24 Choosing Finish

How to finish installing Actuate BIRT iServer 11

- 1 When the Actuate BIRT iServer 11 setup success message appears, as shown in Figure 4-25, choose Finish to exit this wizard.



Figure 4-25 Exiting the installation wizard

- 2 The installation program prompts you to install the online help from the following location, as shown in Figure 4-26:

<http://www.actuate.com>




Figure 4-26 Viewing the install online help and manuals prompt

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

The installation program installs shortcuts on the desktop, as shown in Figure 4-27.




Figure 4-27 iServer shortcuts installed on the desktop

These shortcuts provide access to the following iServer components:

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

■ iServer 11

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

Choose Start→Programs→Actuate 11→BIRT iServer Configuration Console to open Configuration Console. Configuration Console Log In appears, as shown in Figure 4-28. iServer is ready for use.

User name: Administrator

Password:

Language: English (United States)

Time zone: America/Los_Angeles

Log In

Figure 4-28 Configuration Console Log In appears

About changes to the file structure

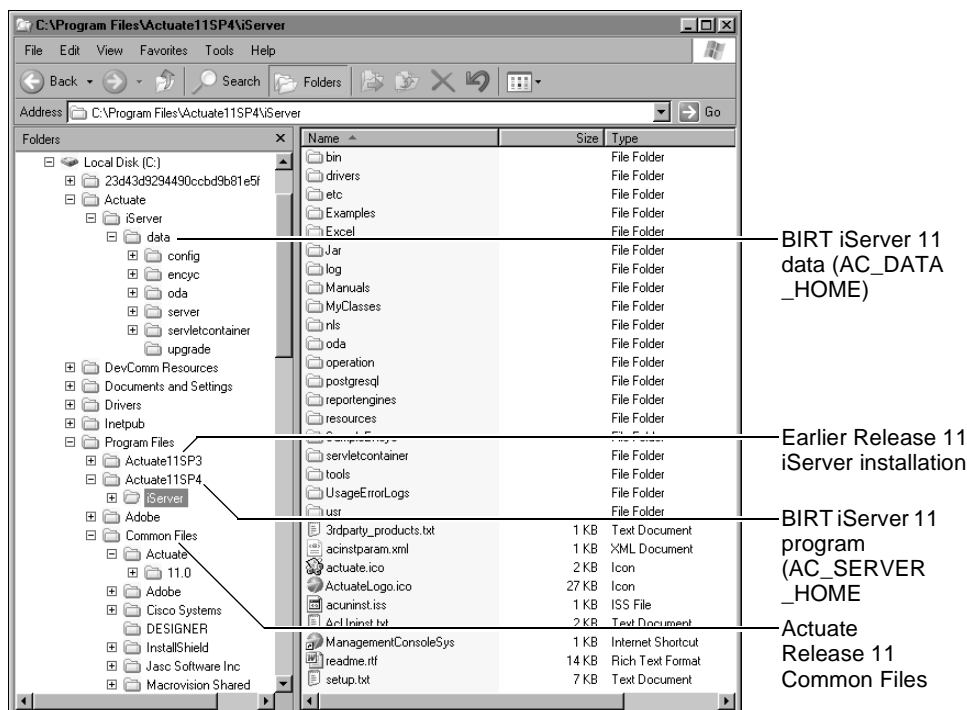


Figure 4-29 Viewing the file system after upgrade to Release 11SP4

Figure 4-29 shows the results of the changes to the upgraded installation file structure:

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

Performing a manual upgrade of an Encyclopedia volume schema for an earlier Release 11 iServer

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

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

How to perform a manual, in-place upgrade of an Encyclopedia volume for an earlier Release 11 iServer

- 1 Choose Start→Programs→Actuate 11→iServer Configuration Console. Log in to Configuration Console as Administrator.

Choose Advanced view. Then, from the side menu, choose Volumes.

On Volumes, point to the icon next to the volume and choose Take offline, if necessary, as shown in Figure 4-30.




Figure 4-30 Choosing to take volume offline

Choose OK to accept the default grace period to allow current transactions on the volume to complete, as shown in Figure 4-31.




Figure 4-31 Choosing to take volume offline

- 2 On Volumes, point to the icon next to the volume schema and choose Properties, as shown in Figure 4-32.




Figure 4-32 Choosing to view volume schema properties

- 3 On Properties, make note of the value that appears for Schema name. You must specify the schema name, not the database schema name, in a later step. Schema name is `ac_corp`, as shown in Figure 4-33.

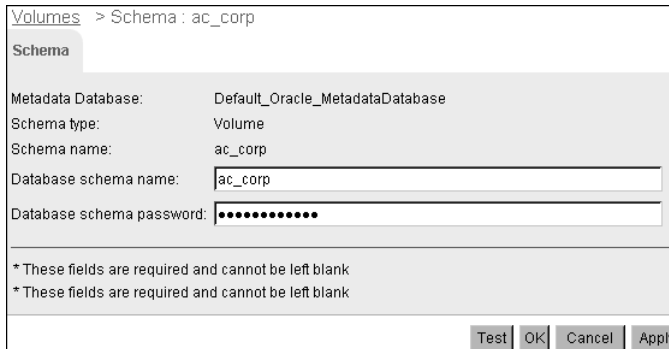


Figure 4-33 Making note of Schema name

- 4 Run the Encyclopedia Data Store Upgrader utility by performing the following tasks:
 - 1 Add the following string to the PATH environment variable on your machine:

`<AC_SERVER_HOME>\bin`

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

`C:\Program Files (x86)\Actuate11SP5\iServer\bin`

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

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

where <schema name> is the Encyclopedia volume schema name or <property file name> is the name of a pre-existing file containing the environment variable settings the utility requires. Type the schema password when prompted for it.

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

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

The batch file performs the following operations, as shown in Listing 4-1:

- ❑ Checks to see if the administrator submitted a property file or a schema name on the command line when running the script.
- ❑ If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.
- ❑ Calls the set_tools_environment.bat script, which sets the environment variables.
- ❑ If the property file exists, the script executes the Encyclopedia Data Store Upgrader utility using the name of the property file as an argument.
- ❑ If the property file does not exist, the script creates a property file, upgrade_encyclopedia_data_store.properties, which contains the environment variable settings.

Listing 4-1 upgrade_encyclopedia_data_store.bat

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

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

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

:: Create property file
```

```

SET PROPERTY_FILE=upgrade_encyclopedia_data_store.properties
ECHO AC_SERVER_HOME = %AC_SERVER_HOME%> "%PROPERTY_FILE%"
ECHO AC_DATA_HOME = %AC_DATA_HOME%>> "%PROPERTY_FILE%"
ECHO USE_SERVER_CONFIG_FILE = true>> "%PROPERTY_FILE%"
ECHO CONFIG_SCHEMA_NAME = %1>> "%PROPERTY_FILE%"

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

GOTO :END

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

:END

```

In the batch file, the ^ symbol is a line-continuation character. You must run Encyclopedia Data Store Upgrader separately for every schema that you want to upgrade when upgrading from an earlier Actuate 11 version.

See Table 4-3 for descriptions of the required Encyclopedia Data Store Upgrader properties later in this chapter. See Table 4-4 for descriptions of the optional Encyclopedia Data Store Upgrader properties.

- 4 Type the password for the schema, as shown in Figure 4-34.




Figure 4-34 Typing the schema password

After the utility runs, close the command prompt.

- 5 In the Advanced view of Configuration Console, choose Volumes from the side menu.

Point to the icon next to the volume name and choose Take online, as shown in Figure 4-35.

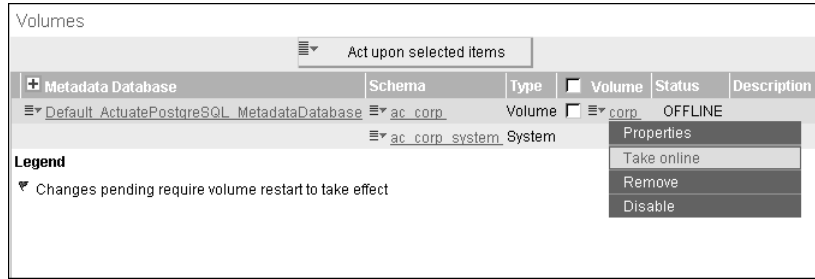


Figure 4-35 Taking the volume online

The volume comes online, as shown in Figure 4-36.

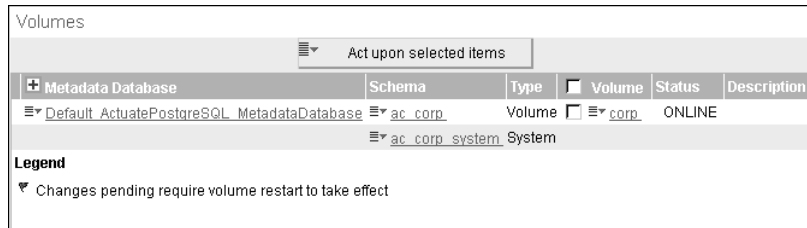


Figure 4-36 Viewing the status of the volume after it comes online

For more information about setting Encyclopedia Data Store Upgrader utility properties, see “Working with iServer utilities,” later in this chapter.

Performing an automatic in-place upgrade from iServer Release 11SP4 Fix 5 to iServer Release 11SP5

When upgrading in-place from iServer Release 11 Service Pack 4 (11SP4 Fix 5) to iServer Release 11 Service Pack 5 (11SP5), the administrator can upgrade iServer using the automatic upgrade process. The automatic upgrade program performs the following tasks:

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

The upgrade from 11SP4 Fix 5 automatically updates the metadata in the installed RDBMS. In 11SP5, it is not necessary to perform a manual in-place upgrade of an Encyclopedia volume schema for any supported RDBMS after performing the automatic, wizard-based, iServer system upgrade.

The following sections describe how to perform an automatic in-place upgrade of iServer Release 11 Service Pack 4 Fix 5, to Release 11 Service Pack 5.

Setting PostgreSQL system variables

If PostgreSQL system variables exist on your machine, set them to the following values before upgrading iServer:

- PGHOST=localhost
- PGPORT=8432
- PGUSER=postgres
- PGDATABASE=iserver

You can reset these variables to their previous values when the upgrade is complete. If you fail to set these variables on a Windows system before upgrading iServer, iServer may not launch properly.

Stopping the Actuate 11 BIRT iServer and NobleNet Portmapper services

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

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

How to stop the Actuate 11 BIRT iServer service

- 1 Choose Start→Settings→Control Panel.
On Control Panel, choose Administrative Tools.
On Administrative Tools, choose Services.
- 2 On Services, select Actuate 11 BIRT iServer Service, as shown in Figure 4-37.
Then, choose Stop the service.

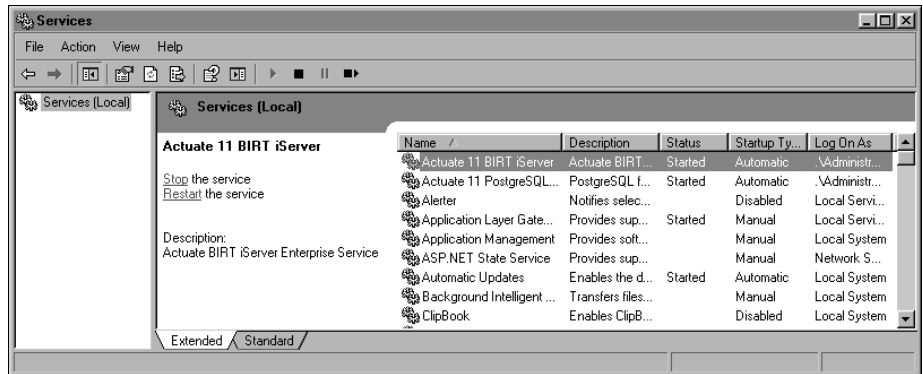


Figure 4-37 Stopping the BIRT iServer service

The service stops, as shown in Figure 4-38.

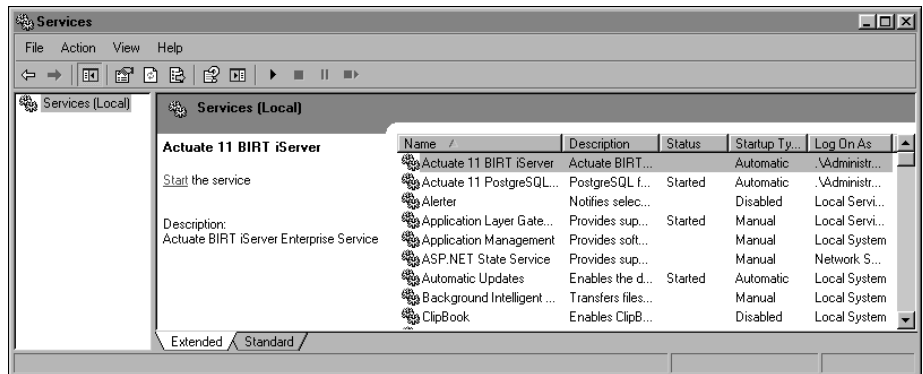


Figure 4-38 Viewing Services after BIRT iServer service stops

How to shut down the Actuate 11 PostgreSQL for BIRT service

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

- 1 Choose Start→Programs→Administrative Tools→Services.
On Control Panel, choose Administrative Tools.
On Administrative Tools, choose Services.
- 2 On Services, select Actuate 11 PostgreSQL for BIRT iServer, as shown in Figure 4-39.

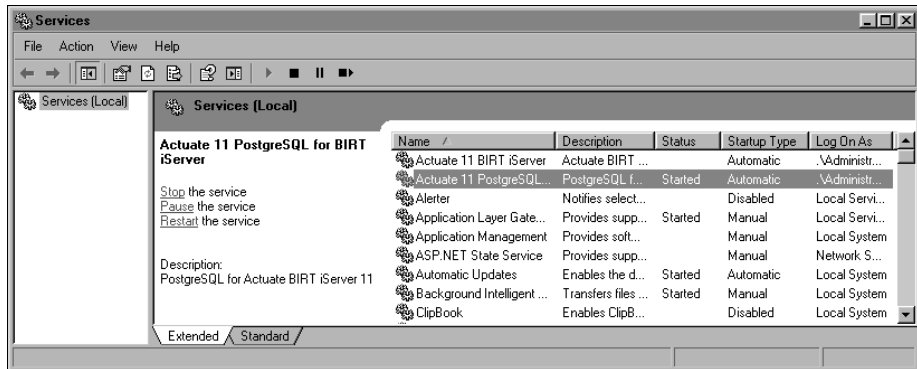


Figure 4-39 Shutting down the PostgreSQL for BIRT iServer service
The service stops, as shown in Figure 4-40.

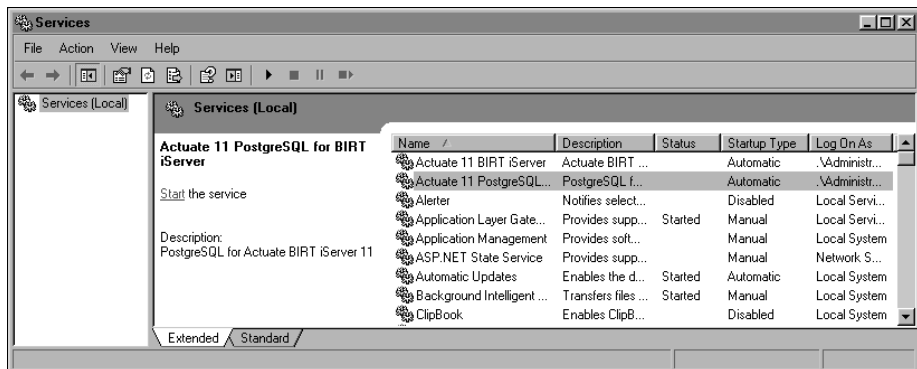


Figure 4-40 Viewing Services after PostgreSQL service stopped

How to shut down the NobleNet Portmapper

Stop NobleNet Portmapper by performing the following tasks:

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

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

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


Running the in-place upgrade on iServer 11SP4 Fix 5

The following procedure describes step-by-step how to perform an automatic upgrade in place from iServer Release 11 Service Pack 4 Fix 5 to Release 11 Service Pack 5.

How to run the in-place upgrade on iServer 11SP4 Fix 5

- 1 Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
 - encyc directories from all nodes
 - acserverconfig.xml in the AC_DATA_HOME/config directory
 - acpmdconfig.xml in the AC_SERVER_HOME/etc directory
 - RSSE code and associated files if you use the Open Security option
- 2 Download the iServer distribution package. Extract the files. Run the self-extracting executable file, ActuateBIRTiServerEnterprise.exe. The welcome message appears, as shown in Figure 4-41. Choose Next.



Figure 4-41 Viewing the welcome message

- 3 Read and accept the license agreement, as shown in Figure 4-42. Choose Next.



Figure 4-42 Viewing the license agreement

- 4 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 4-43. Choose Next.



Figure 4-43 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 (x86)/Actuate11SP5/iServer`.

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

The installer detects the previous Release 11 installation and selects Upgrade Existing iServer and Migrate the Encyclopedia Data. The installer lists the Existing iServer Location, as shown in Figure 4-44.

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

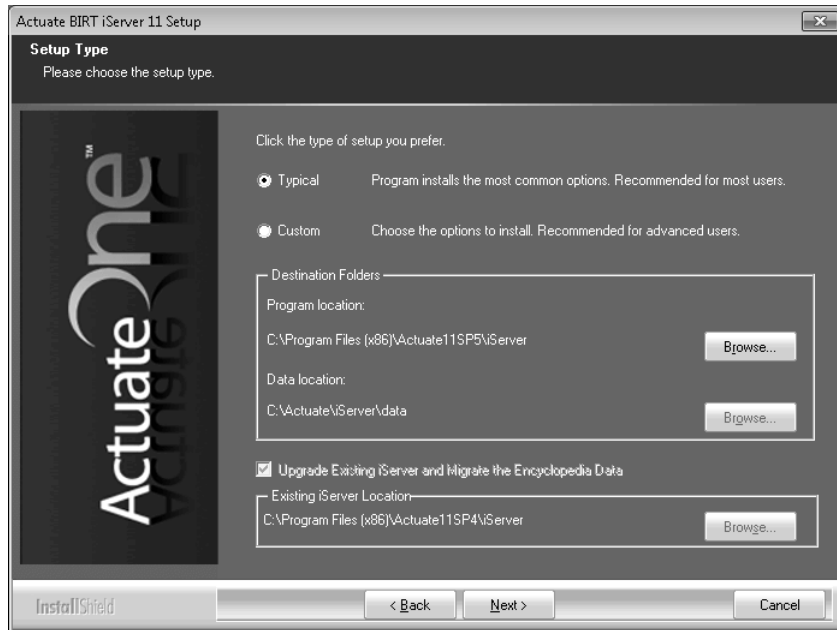


Figure 4-44 Choosing setup type

A message appears stating that the install program is migrating a previous version of the Encyclopedia volume, which might take a long time, as shown in Figure 4-45. Choose OK.

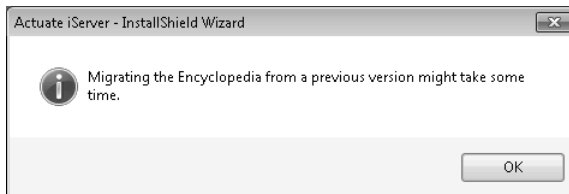


Figure 4-45 Confirming volume migration

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

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

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

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

- 1** Choose OK on the prompt.
 - 2** Stop the Actuate 11 PostgreSQL for BIRT iServer service.
 - 3** On Setup Type, choose Next.
- 8** In Specify Profiles, type the user name and password for the account used to start the Actuate BIRT iServer 11 service. The account must be a member of the Administrators group. Actuate recommends that you limit access to this account for security reasons.

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

Accept Automatically start the Actuate BIRT iServer 11 service when Windows boots, as shown in Figure 4-46. If you deselect this option, you must start the service manually from Windows Services. Choose Next.

- 9** In System Configuration Password, type a password for Configuration Console, as shown in Figure 4-47. 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.

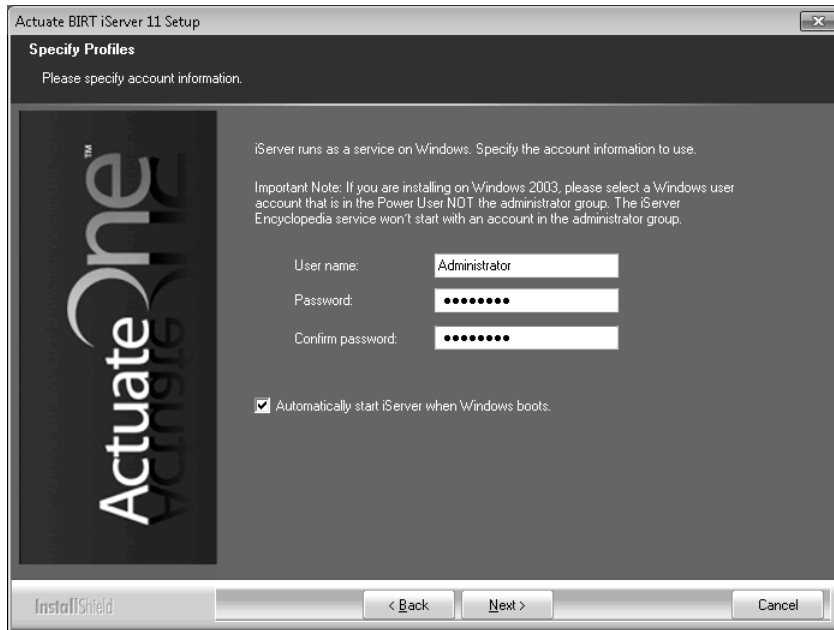


Figure 4-46 Specifying the account for running the iServer service

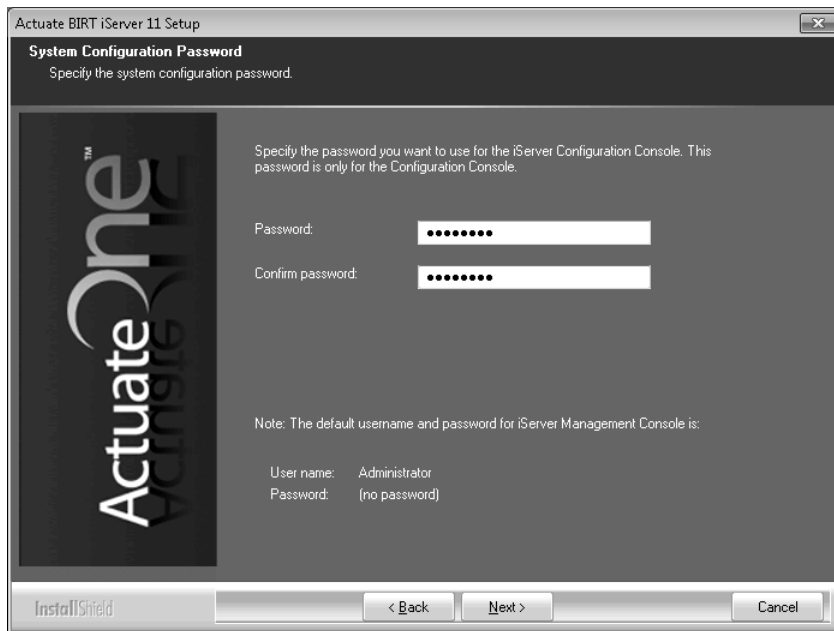


Figure 4-47 Specifying the password for using Configuration Console

10 In Start Copying Files, review the settings shown in Figure 4-48. Choose Next.

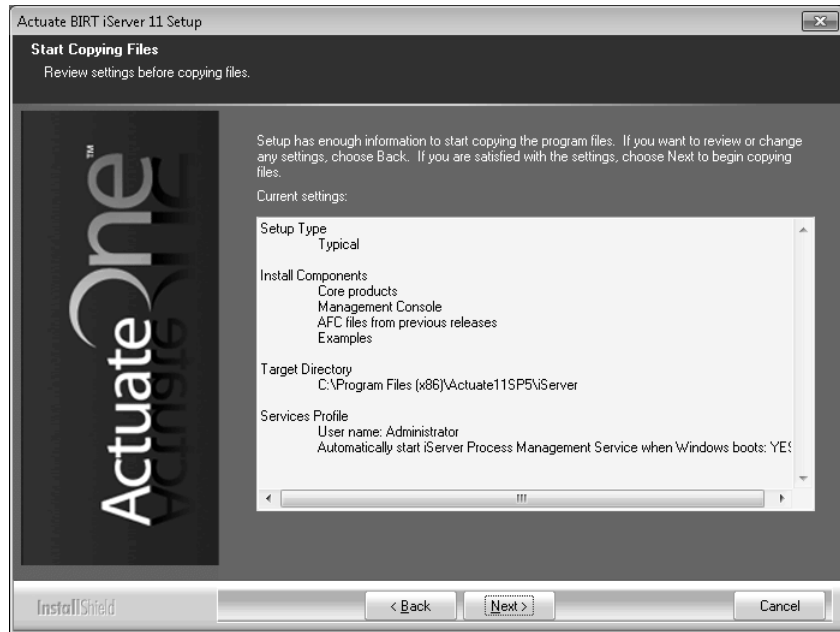


Figure 4-48 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 4-49.



Figure 4-49 Viewing setup status

During installation, a command prompt appears, displaying messages regarding installation run-time processing, as shown in Figure 4-50.

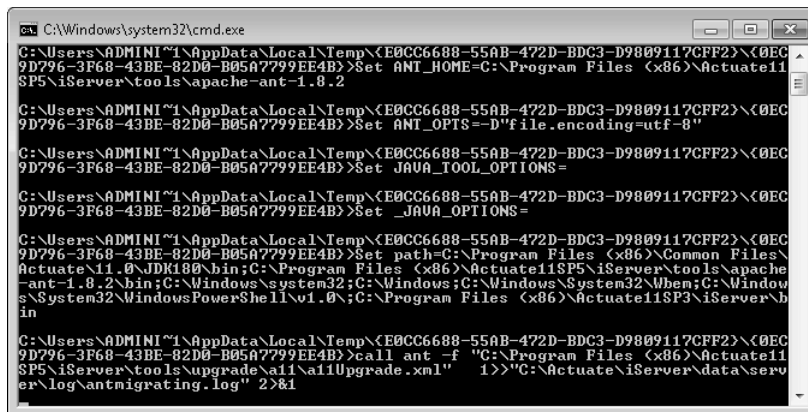


Figure 4-50 Viewing run-time processing

How to finish installing Actuate BIRT iServer 11

- 1 When the Actuate BIRT iServer 11 setup success message appears, as shown in Figure 4-51, choose Finish to exit this wizard.



Figure 4-51 Exiting the installation wizard

- 2 The installation program prompts you to install the online help from the following location, as shown in Figure 4-52:

<http://www.actuate.com>




Figure 4-52 Viewing the install online help and manuals prompt

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

The installation program installs shortcuts on the desktop, as shown in Figure 4-53.




Figure 4-53 iServer shortcuts installed on the desktop

These shortcuts provide access to the following iServer components:

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

Choose Start ▶ Programs ▶ Actuate 11 ▶ BIRT iServer Configuration Console to open Configuration Console. Configuration Console Log In appears, as shown in Figure 4-54. iServer is ready for use.

Figure 4-54 Configuration Console Log In appears

Performing a side-by-side migration from iServer Release 11SP4 Fix 5 to iServer Release 11SP5

To migrate from iServer Release 11 Service Pack 4 Fix 5 (11SP4 Fix 5) to Release 11 Service Pack 5 (11SP5), you perform a manual side-by-side migration. At the highest level, a side-by-side migration consists of exporting iServer volume data and metadata from the 11SP4 Fix 5 release and importing the data and metadata to the 11SP5 release. You install 11SP5 on a separate machine from where 11SP4 Fix 5 is installed. The advantage that a side-by-side migration offers is that it does not disturb your 11SP4 Fix 5 installation. The side-by-side migration supports continuing to use the 11SP4 Fix 5 installation after migrating to 11SP5 or repeating the migration procedure if necessary.

Conceptually, the process consists of the following steps:

- Run a utility to export the iServer volume metadata from the 11SP4 Fix 5 installation. You run the utility once for each volume schema that your 11SP4 Fix 5 installation contains.

- Install 11SP5 on a separate machine from where the 11SP4 Fix 5 iServer is installed.
- Copy the folder or folders containing the exported iServer volume metadata files from the machine where the 11SP4 Fix 5 iServer is installed, to the machine hosting the metadata database for 11SP5. If you are using the PostgreSQL database that installs with iServer, the machine hosting the metadata database for 11SP5 is always the same machine that hosts 11SP5.
- Copy the volume data from the 11SP4 Fix 5 installation to the 11SP5 installation.
- Run a utility on the machine hosting the metadata database for 11SP5 to import the iServer volume metadata. You run the utility once, for each volume schema that your 11SP4 Fix 5 installation contains.
- Re-create the migrated volumes using the 11SP5 Configuration Console.

Exporting volume metadata from the 11SP4 Fix 5 installation

Exporting volume metadata consists of the following operations:

- Preparing to run the Data Store Administrator utility
 - Shut down the 11SP4 Fix 5 iServer
 - Add the path of the folder containing the utility to the PATH environment variable
- Creating a properties file that you pass to the utility when you run it
- Running the Data Store Administrator utility to export the metadata

Preparing to run the Data Store Administrator utility

Perform the procedure in this section to prepare to run the Data Store Administrator utility to export the volume metadata.

How to prepare to run the Data Store Administrator utility

- 1 Shut down 11SP4 Fix 5 by performing the following tasks:
 - 1 Choose Start>Programs>Actuate 11>BIRT 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 4-55.



Figure 4-55 iServer is offline

- 2 Edit the PATH environment variable on your machine to contain the following string:

```
<AC_SERVER_HOME>\bin
```

where AC_SERVER_HOME refers to the 11SP4 Fix 5 installation path. For example, using the default value for AC_SERVER_HOME, add:

```
C:\Program Files (x86)\Actuate11SP4\iServer\bin
```

Creating the properties file

On the machine where 11SP4 Fix 5 is installed, you create a properties file, containing name-value pairs that the Encyclopedia Data Store Administrator utility uses when you run it. If you are exporting metadata for more than one volume schema, a good practice is to create a separate properties file for each schema. Properties this file contains include:

- Type of database the utility is running against.
- Name of the database containing the volume schema or schemas.
- Name of the machine hosting the database.
- Name of the schema you are running the export utility against. You run the export utility once for each volume schema that your 11SP4 Fix 5 installation contains.
- The path of the folder that the utility creates to contain the exported metadata. Specify a separate folder for each volume schema you run against.

The following questions are common when creating a properties file for exporting volume metadata:

- How to export the metadata for all volumes in a schema?
A: Specify `EXPORT_ALL_DATA = true`
- How to export the data for only a single volume in a schema?
A: Specify `EXPORT_DATA = true` and `VOLUME_NAME = <name of the volume>` instead of specifying `EXPORT_ALL_DATA = true`
- How to use the same database server for 11SP5 that 11SP4 Fix 5 uses, when using a pre-existing database server for volume metadata?
A: Specify `SCHEMA_NAME = <name of the volume schema in 11SP4 Fix 5>` and specify `NEW_SCHEMA_NAME = <new name>`. When you perform the

import operation, you create a new volume schema in the database and a new volume schema in the 11SP5 Configuration Console. In both these operations, you name the schema <new name>.

For the complete list of properties that the Data Store Administrator uses, see “Specifying Encyclopedia Data Store Administrator properties,” later in this chapter.

Use the example properties files in this section to create a properties file named VolumeExport.properties in the AC_SERVER_HOME\bin folder, for the type of database you are using.

Setting properties for export when using the PostgreSQL database that installs with iServer

Use the following example to create VolumeExport.properties:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP4/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp
DATA_EXPORT_FORMAT = PostgreSQL
```

Setting properties for export when using a pre-existing PostgreSQL database

Use the following example to create VolumeExport.properties:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP4/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = <machine name of metadata database host>
DATABASE_PORT = 5432
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp
DATA_EXPORT_FORMAT = PostgreSQL
```

Setting properties for export when using a pre-existing Oracle database

Use the following example to create VolumeExport.properties:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP4/iServer
DATABASE_TYPE = Oracle
```

```

DATABASE_NAME = iserver
DATABASE_HOST = <machine name of metadata database host>
DATABASE_PORT = 1521
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp
DATA_EXPORT_FORMAT = Oracle

```

Setting properties for export when using a pre-existing SQL Server database

Use the following example to create VolumeExport.properties:

```

AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP4/iServer
DATABASE_TYPE = SQLServer
DATABASE_NAME = iserver
DATABASE_HOST = <machine name of metadata database host>
DATABASE_INSTANCE = MSSQLSERVER
DATABASE_PORT = 1433
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp
DATA_EXPORT_FORMAT = SQLServer

```

Setting properties for export when using a pre-existing DB2 database

Use the following example to create VolumeExport.properties:

```

AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP4/iServer
DATABASE_TYPE = DB2
DATABASE_NAME = iserver
DATABASE_HOST = <machine name of metadata database host>
DATABASE_PORT = 50000
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp
DATA_EXPORT_FORMAT = SQLServer

```

For more information on Encyclopedia Data Store Administrator properties, see “Specifying Encyclopedia Data Store Administrator properties,” later in this chapter.

Running the Data Store Administrator utility

The utility creates the folder that the DATA_EXPORT_FOLDER property specifies, containing the metadata for the schema that the SCHEMA_NAME

property specifies. For each table in the volume schema that you run the utility against, the utility creates a file in this folder.

How to run the Data Store Administrator utility

- 1 Open a command prompt and navigate to AC_SERVER_HOME\bin in the 11SP4 Fix 5 installation.
- 2 Run the administrate_encyclopedia_data_store.bat file using the following command line syntax:

```
administrate_encyclopedia_data_store VolumeExport.properties
```

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

```
admineds VolumeExport.properties
```

Importing volume metadata and data to the 11SP5 installation

First, you install 11SP5 on a different machine from where 11SP4 Fix 5 is installed. Next, you copy the folder or folders containing the exported volume metadata from the machine hosting 11SP4 Fix 5 to the machine hosting the metadata database for 11SP5. If you are using the PostgreSQL database that installs with iServer for volume metadata, the machine hosting the metadata database for 11SP5 is always the same machine that hosts 11SP5. Then, you migrate the volume or volumes that the 11SP4 Fix 5 installation contains to the 11SP5 installation.

When migrating a volume from 11SP4 Fix 5 to 11SP5, you create a folder on the machine hosting 11SP5 for the volume data. Next, you create a partition in Configuration Console that points to that folder. Then, you copy the folders containing the volume data from the 11SP4 Fix 5 volume data folder to the 11SP5 volume data folder.

The iServer install program for the release 11 series creates a default volume and a default partition. If you are migrating the default volume, you do not need to create a new folder and partition in 11SP5 because they already exist. The folder for the default volume data is AC_DATA_HOME\encyc, and the partition for the default volume is named DefaultPartition.

After copying the volume data folders from the 11SP4 Fix 5 installation to the 11SP5 installation, you remove the default volume and the default volume schema in the 11SP5 installation using Configuration Console.

Then, you create a new volume schema for each schema for which you exported volume metadata. If you are using the PostgreSQL database that installs with 11SP5 to contain volume metadata, you create the schema or schemas in Configuration Console. If you are using a pre-existing database for metadata, you create the schema or schemas in both the database and in Configuration Console.

If you are using the default PostgreSQL database that installs with iServer, you next create a properties file for each schema for which you exported volume metadata. Then, you run the Data Store Administrator utility to import the volume metadata for a schema, passing the properties file in the command that executes the utility. The properties file includes properties that specify:

- The type of database the utility is running against.
- The name of the schema you created in 11SP5 for which you are importing volume metadata.
- The name of the folder containing the exported volume metadata.

If you are using a pre-existing database for metadata, you execute a bulk-load script specific to the database type you are using for volume metadata to import the metadata. The Data Store Administrator utility creates the bulk-load script and any necessary supporting scripts when you run the utility to export metadata, writing the scripts to the folder containing the exported volume metadata files. Depending on the database type, you edit one or two scripts to provide information needed at runtime, such as database user name and password, before executing the script to import the metadata. For more information on bulk-load scripts, see “Using the generated bulk-load script files,” later in this chapter.

Finally, for each volume you migrated from the 11SP4 Fix 5 installation, you create a volume in the 11SP5 Configuration Console and bring the volume online.

Preparing to import volume metadata and migrating the volume data

Perform the following procedure to prepare to import volume metadata and migrate the volume data.

How to prepare to import volume metadata and migrate volume metadata to 11SP5

- 1 Install 11SP5 on a different machine from the machine that hosts 11SP4 Fix 5. For more information, see “Installing a new instance of BIRT iServer Release 11 Service Pack 5,” in Chapter 2, “Installing BIRT iServer.”
- 2 Copy the folder or folders containing the exported volume metadata files that the Encyclopedia Data Store Administrator utility created from the 11SP4 Fix 5 installation to the machine hosting the metadata database for 11SP5. If you are using the PostgreSQL database that installs with iServer, the machine hosting the metadata database for 11SP5 is always the same machine that hosts 11SP5.
- 3 For any volume you are migrating that does not use DefaultPartition, perform the following tasks on the machine that is hosting 11SP5:
 - 1 Create a folder for a new partition in AC_DATA_HOME. This folder will contain the data for the volume you are migrating. AC_DATA_HOME points to the location of the iServer data specified during the iServer install

procedure, as shown in Figure 2-4. The default path for AC_DATA_HOME is:

C:\Actuate\iServer\data

- 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 11SP4 Fix 5 partition that specifies the path to the volume in 11SP4 Fix 5 that you are migrating. For example, if the name of the Encyclopedia volume partition in 11SP4 Fix 5 is sales_partition, type that name in Partition name.
- 6 In Partition Path, type the path to the 11SP5 volume folder that you created in step 2.1, as shown in Figure 4-56. This path does not need to match the partition path to the same volume in 11SP4 Fix 5. Choose OK.

Partitions > Add Partition

Partition name: sales_partition *

Template Name	Partition Path
urup	C:\Actuate\iServer\data\sales_volume *

* These fields are required and cannot be left blank

OK Cancel

Figure 4-56 Creating a partition for a migrated volume

- 7 Copy the file, filetype, and if they exist, the status and tempRov folders from the 11SP4 Fix 5 volume folder to the 11SP5 volume folder that you created in step 3.1.
- 4 If the 11SP4 Fix 5 installation contains a volume that uses a partition named DefaultPartition, use DefaultPartition for that volume in 11SP5 by performing the following tasks:
 - 1 On the machine hosting 11SP5, use Windows Explorer to navigate to the AC_DATA_HOME\encyc folder. This folder contains three folders, file, fileType, and postgresql. Delete the file and fileType folders only. Do not delete the postgresql folder. The default path for AC_DATA_HOME is:
C:\Actuate\iServer\data
 - 2 In the 11SP4 Fix 5 installation, log in to Configuration Console.

- 3 Choose Advanced view.
- 4 From the side menu, choose Partitions. On Partitions, point to the icon next to DefaultPartition and choose Template settings, as shown in Figure 4-57.

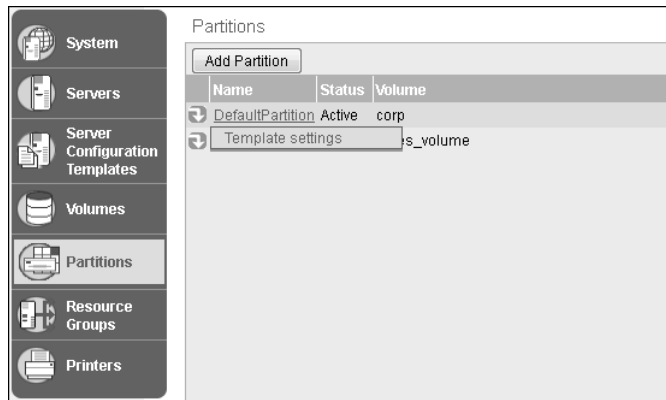


Figure 4-57 Choosing Template settings

- 5 On Template Settings, take note of the path to the volume folder appearing in Partition Path. In a typical iServer Release 11 installation, this path is \$AC_DATA_HOME\$\encyc.
- 6 Using Windows Explorer, navigate to the location appearing in Partition Path.
- 7 Copy the file, filetype, and if they exist, the status and tempRov folders from the 11SP4 Fix 5 AC_DATA_HOME\encyc folder to the 11SP5 AC_DATA_HOME\encyc folder.
- 5 In the 11SP5 installation, remove the default volume and the default volume schema by performing the following tasks:
 - 1 In the 11SP5 installation, log in to Configuration Console.
 - 2 Choose Advanced view.
 - 3 From the side menu, choose Volumes. On Volumes, point to the icon next to the default volume and choose Take offline, as shown in Figure 4-58. The name of the default volume is the name you specified during the 11SP5 install procedure, as shown in Figure 2-5.




Figure 4-58 Taking the default volume offline

- 4 On Volumes, remove the default volume by pointing to the icon next to the default volume and choosing Remove.
- 5 On Volumes, remove the volume schema by pointing to the icon next to the volume schema and choosing Remove, as shown in Figure 4-59.




Figure 4-59 Removing the volume schema in 11SP5

Importing the volume metadata

To import the volume metadata, you perform the following operations:

- Create the necessary volume schema or schemas. You create a new volume schema for each volume schema for which you exported volume metadata.
- Run a utility to import the metadata into the metadata database. You run the utility once, for each volume schema.

This section contains a procedure that describes these tasks for each of the database types that iServer supports for containing iServer volume metadata.

Importing volume metadata when using the PostgreSQL database that installs with iServer

Perform the following procedure if you are using the default PostgreSQL database that installs with iServer to contain volume metadata.

How to import volume metadata when using the PostgreSQL database that installs with iServer

- 1 Create a new volume schema for one or more volumes that you are migrating by pointing to the icon next to Default ActuatePostgreSQL MetadataDatabase and choosing Add volume schema, as shown in Figure 4-60.

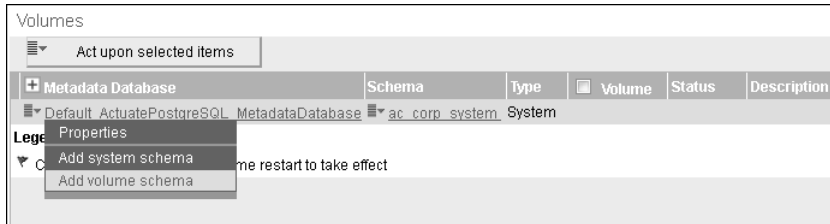


Figure 4-60 Choosing to create a new volume schema

On Volumes—New Schema, perform the following tasks:

- 1 In Schema name, type a name for the schema. The name must be different than the name of the schema you removed. 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 2-6. Choose OK.

Volumes—New Schema appears as shown in Figure 4-61.

Volumes > New Volume Schema

Schema

Metadata Database: Default_ActuatePostgreSQL_MetadataDatabase

Schema type: Volume

Schema name: ac_corp_11sp5 *

Database schema name: ac_corp_11sp5 *

Database schema password: *

Database schema password confirm: *

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

Database superuser: postgres *

Database superuser password: *

* These fields are required and cannot be left blank

OK Cancel Apply

Figure 4-61 Creating a new volume schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-62.

Volumes

Act upon selected items

Metadata Database	Schema	Type	Volume	Status	Description
Default_ActuatePostgreSQL_MetadataDatabase	ac_corp_11sp5	Volume			
	ac_corp_system	System			

Legend

Changes pending require volume restart to take effect

Figure 4-62 Viewing the new schema

- 2 To import the volume metadata, run the Encyclopedia Data Store Administrator by performing the following tasks:
 - 1 In 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 (x86)/Actuate11SP5/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
```

```

DATABASE_PORT = 8432
SUPERUSER = postgres
SUPERUSER_PASSWORD = <superuser password>
SCHEMA_NAME = ac_corp_11sp5
SCHEMA_PASSWORD = <your schema password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = C:/Actuate/11SP4_data/ac_corp

```

For more information on Encyclopedia Data Store Administrator properties, see “Specifying Encyclopedia Data Store Administrator properties,” later in this chapter.

- 2 Open a command prompt and navigate to AC_SERVER_HOME\bin in the 11SP4 Fix 5 installation.
- 3 Run the administrate_encyclopedia_data_store.bat file using the following command line syntax:

```
administrate_encyclopedia_data_store VolumeImport.properties
```

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

```
admineds VolumeImport.properties
```

Importing volume metadata when using a pre-existing PostgreSQL database

When you run the Data Store Administrator utility to export volume metadata, the utility creates a bulk-load script for the database type that the DATA_EXPORT_FORMAT property in VolumeExport.properties specifies. The utility includes the bulk-load script in the folder to which the utility writes the files containing the exported volume metadata. The bulk-load script for importing volume metadata to a PostgreSQL database is ac_postgresql_client_load.bat. You edit this script before executing it to specify information needed at runtime.

Perform the following procedure if you are using a pre-existing PostgreSQL database to contain volume metadata.

How to import volume metadata when using a pre-existing PostgreSQL database

- 1 In the PostgreSQL database, create a new schema for each of the volume schemas for which you exported metadata. Give each schema the same name that you specified for the schema in VolumeExport.properties.
- 2 In Configuration Console—Advanced View—Volumes, create a volume schema for each schema you created in PostgreSQL. Give each schema the same name that you specified for the schema in VolumeExport.properties. For a description of this task, see step 1 in “How to import volume metadata when using the PostgreSQL database that installs with iServer,” earlier in this chapter.

- 3 Navigate to the location of the folder containing the exported volume metadata files. Using a text editor, open the `ac_postgresql_client_load.bat` file. The following example shows the information you must provide appearing between the brackets:

```
"<psql>" --quiet --host <localhost> --port <8432> --dbname  
      <server> --username <postgres> --file  
      ac_postgresql_client_load.psql
```

where:

- `<psql>` is the full path of `psql`, the PostgreSQL interactive terminal.
- `<localhost>` is the name of the machine hosting PostgreSQL. Same as the `DATABASE_HOST` specified in `VolumeExport.properties`.
- `<8432>` is the number of the port on which PostgreSQL listens. Same as the `DATABASE_PORT` specified in `VolumeExport.properties`.
- `<server>` is the name of the database containing the volume schema or schemas. By default, the value is `iserver`.
- `<postgres>` is the user id of a user having the necessary privileges to insert rows into the volume schema in the PostgreSQL database. By default, the value is `postgres`.

Save and exit the file.

- 4 Open a command prompt and navigate to the location of the folder containing the exported volume metadata files. Execute the `ac_postgresql_client_load.bat` script.

Importing volume metadata when using a pre-existing Oracle database

When you run the Data Store Administrator utility to export volume metadata, the utility creates a bulk-load script for the database type that the `DATA_EXPORT_FORMAT` property in `VolumeExport.properties` specifies. The utility includes the bulk-load script in the folder to which the utility writes the files containing the exported volume metadata. The bulk-load script for importing volume metadata to an Oracle database is `ac_oracle_load.bat`. The Data Store Administrator utility also creates a file containing parameters that `ac_oracle_load.bat` uses, named `ac_oracle_load.par`. You edit this file to specify parameter values before executing `ac_oracle_load.bat`.

Perform the following procedure if you are using a pre-existing Oracle database to contain volume metadata.

How to import volume metadata when using a pre-existing Oracle database

- 1 In the Oracle database, create a new schema for each of the volume schemas for which you exported metadata. Give each schema the same name that you specified for the schema in `VolumeExport.properties`.

- 2 In Configuration Console—Advanced View—Volumes, create a volume schema for each schema you created in Oracle. Give each schema the same name that you specified for the schema in VolumeExport.properties. For a description of this task, see step 1 in “How to import volume metadata when using the PostgreSQL database that installs with iServer,” earlier in this chapter.
- 3 Navigate to the location of the folder containing the exported volume metadata files. Using a text editor, open the ac_oracle_load.par file. The following example shows the information you must provide appearing between the brackets:

```
USERID=<userid>/<password>@<instance>  
SILENT=HEADER  
DIRECT=TRUE  
COLUMNARRAYROWS=1000  
READSIZE=4194304  
STREAMSIZE=4194304  
BINDSIZE=4194304
```

where:

- <userid> is typically the volume schema owner. Alternatively, it can be the user id of a user having the necessary privileges to bulk load data into the volume schema in the database.
- <password> is the password of the user that <userid> specifies
- <instance> specifies the database instance to which to connect. This can be a net service name (TNS alias name) or connect descriptor.

Save and exit the file.

- 4 Open a command prompt and navigate to the location of the folder containing the exported volume metadata files. Execute the ac_oracle_load.bat script.

Importing volume metadata when using a pre-existing SQL Server database

When you run the Data Store Administrator utility to export volume metadata, the utility creates a bulk-load script for the database type that the DATA_EXPORT_FORMAT property in VolumeExport.properties specifies. The utility includes the bulk-load script in the folder to which the utility writes the files containing the exported volume metadata. The bulk-load script for importing volume metadata to a SQL Server database is ac_sql_server_load.bat. The Data Store Administrator utility also creates a file which ac_sql_server_load.bat calls, ac_sql_server_load.sql. You edit both files to provide configuration information before executing ac_sql_server_load.bat.

Perform the following procedure if you are using a pre-existing SQL Server database to contain volume metadata.

How to import volume metadata when using a pre-existing SQL Server database

- 1 In the SQL Server database, create a new schema for each of the volume schemas for which you exported metadata. Give each schema the same name that you specified for the schema in VolumeExport.properties.
- 2 In Configuration Console—Advanced View—Volumes, create a volume schema for each schema you created in SQL Server. Give each schema the same name that you specified for the schema in VolumeExport.properties. For a description of this task, see step 1 in “How to import volume metadata when using the PostgreSQL database that installs with iServer,” earlier in this chapter.
- 3 Navigate to the location of the folder containing the exported volume metadata files. Using a text editor, open the ac_sql_server_load.bat file. The following example shows information that you must provide appearing between the brackets.

```
@ECHO OFF
DEL ac_*.bad
DEL ac_*.bad.Error.Txt
sqlcmd -S <server> -U <user> -P <password> -d <database> -i
    ac_sql_server_load.sql
<bindsize>
```

where:

- <server> is the connection string for the SQL Server instance. The DATABASE_INSTANCE property in VolumeExport.properties specifies the name of the SQL Server instance.
- <user> is the user id of the database administrator, or of a user having the necessary privileges to bulk load data into the volume schema in the SQL Server database.
- <password> is the password of the user that <user> specifies.
- <database> is iserver.
- <bindsize> is BINDSIZE=4194304

Save and exit the file.

- 4 Using a text editor, open the ac_sql_server_load.sql file. The following example shows information that you must provide appearing between the brackets.

```
-- SQL Server bulk load for Actuate iServer
DECLARE @v_data_folder NVARCHAR(1000)
= '<location of exported metadata files>';
DECLARE @v_log_folder NVARCHAR(1000)
= '<location of exported metadata files>';
...
```

where <location of exported metadata files> is the full path of the folder containing the exported metadata files, on the machine hosting the metadata database for 11SP5. Note that this path must appear between single quotes.

Save and exit the file.

- 5 Open a command prompt and navigate to the location of the folder containing the exported volume metadata files. Execute the `ac_sql_server_load.bat` script.

Importing volume metadata when using a pre-existing DB2 database

When you run the Data Store Administrator utility to export volume metadata, the utility creates a bulk-load script for the database type that the `DATA_EXPORT_FORMAT` property in `VolumeExport.properties` specifies. The utility includes the bulk-load script in the folder to which the utility writes the files containing the exported volume metadata. The bulk-load script for importing volume metadata to a DB2 database is `ac_db2_load.bat`. The Data Store Administrator utility also creates a file containing configuration information, `ac_db2_load.db2`. You edit this file to provide configuration information before executing `ac_db2_load.bat`.

Perform the following procedure if you are using a pre-existing DB2 database to contain volume metadata.

How to import volume metadata when using a pre-existing DB2 database

- 1 In the DB2 database, create a new schema for each of the volume schemas for which you exported metadata. Give each schema the same name that you specified for the schema in `VolumeExport.properties`.
- 2 In Configuration Console—Advanced View—Volumes, create a volume schema for each schema you created in DB2. Give each schema the same name that you specified for the schema in `VolumeExport.properties`. For a description of this task, see step 1 in “How to import volume metadata when using the PostgreSQL database that installs with iServer,” earlier in this chapter.
- 3 Navigate to the location of the folder containing the exported volume metadata files. Using a text editor, open the `ac_db2_load.db2` file and make the following changes:

```
CONNECT TO <database> USER <userid> USING <password>
```

where:

- <database> is the name of the database containing the volume schema or schemas. By default, the name of this database is `iserver`.
- <userid> is the user id of a user having the necessary privileges to bulk load data into the volume schema in the SQL Server database. In most cases, this is the volume schema owner. Typically, the volume schema owner has the same name as the volume schema.

- <password> is the password of the user that <userid> specifies.

Save and exit the file.

- 4 Open a command prompt and navigate to the location of the folder containing the exported volume metadata files. Execute the ac_db2_load.bat script.

Creating a new volume

For each migrated volume, create a new volume in 11SP5. The new volume must have the same name that the migrated volume had in 11SP4 Fix 5, and the partition that the new volume uses must also have the same name that the partition for the migrated volume had in 11SP4 Fix 5. For example, for a volume named corp, which uses a partition named DefaultPartition in 11SP4 Fix 5, the corresponding volume you create in 11SP5 must be named corp and must use a partition named DefaultPartition.

How to create a new volume

- 1 In the Advanced view of Configuration Console, on Volumes, point to the icon next to a volume schema and choose Add Volume, as shown in Figure 4-63.




Figure 4-63 Choosing to add a volume

- 2 On New Volume—General, perform the following tasks:
 - 1 In Volume name, type the name of the volume.
 - 2 In Primary partition, specify the partition you created in “How to prepare to import volume metadata and migrate volume metadata to 11SP5,” earlier in this chapter, if the partition for this volume is not DefaultPartition. If the migrated volume uses DefaultPartition, specify DefaultPartition in Primary partition, as shown in Figure 4-64.

Volumes > New Volume

General Open Security Partitions Events

Volume name: corp *

Description:

Schedule for purging notices: HH:mm

Schedule for purging deleted files: HH:mm

Partition

Primary partition: DefaultPartition * Min Free Space: MB

Figure 4-64 Specifying volume and partition name

Choose Partitions.

- 3 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 4-65.

Volumes > New Volume

General Open Security Partitions Events

Assign partitions

Available partitions:

DefaultPartition

sales_partition

Selected partitions:

DefaultPartition Unused

Start Stop

Low Free Space: MB

Min Free Space: MB

OK Cancel Apply

Figure 4-65 Assigning a partition

Choose OK.

- 4 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-66.




Figure 4-66 Viewing the new volume

The new volume comes online, as shown in Figure 4-67.




Figure 4-67 Viewing the new volume after it comes online

- 5 Log in to Management Console. In Files and Folders, the data from the previous release appears.
- 6 If you are satisfied that the contents of the migrated volume or volumes are correct and complete, delete the folder containing the exported volume metadata files.

Migrating a volume to a database of a different type

An iServer volume consists of two types of data:

- **Metadata**
iServer system information, and information about objects a volume contains, such as users, roles, jobs, and files. Volume metadata is stored in relational database management system (RDBMS) schemas.
- **Data**
Data objects, such as designs and documents, and the acserverconfig.xml file containing iServer configuration information. Volume data is stored in files on disk partitions in the file system.

To migrate a volume to a different database, you migrate the volume metadata. It is not necessary to move the volume data.

You can migrate volume metadata to a database of a different type than the database type in which the metadata currently resides. iServer installs with an

embedded PostgreSQL RDBMS for containing volume metadata. If your enterprise uses Oracle to contain all other data, for example, you may wish to migrate iServer volume metadata to Oracle as well.

This section contains a procedure which demonstrates how to migrate volume metadata from the out-of-the-box (OOTB) PostgreSQL database that installs with iServer, to an Oracle database. You can use this procedure as a model for migrating volume metadata to any of the databases that iServer supports.

To migrate a volume to a different database, you perform the following operations:

- Create a system schema and one or more volume schemas in the database to which you are migrating. You create a volume schema in the target database for each volume schema that your iServer installation contains.
- Create a new metadata database and system schema in the iServer Configuration Console.
- Remove the original system schema in Configuration Console.
- Remove any existing volumes and volume schemas in Configuration Console.
- Create a new volume schema in Configuration Console for each volume schema you removed.
- Run a utility to export the volume metadata. You run the utility once, for each volume schema that your iServer installation contains. For each volume schema you run against, the utility creates a folder containing the volume metadata for the schema.
- Copy the folder or folders containing exported iServer volume metadata from the machine where iServer is installed, to the machine hosting the metadata database.
- Run a utility on the machine hosting the metadata database to import the iServer volume metadata. You run the utility once, for each volume schema that your iServer installation contains.
- Re-create the migrated volumes in Configuration Console.

Before beginning the volume migration procedure, make a backup copy of the iServer configuration file, `acserverconfig.xml`. In iServer release 11 Service Pack 5 for example, the location of this file is `AC_DATA_HOME\config\11SP5`, where `AC_DATA_HOME` represents the following path by default:

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

The `acserverconfig.xml` file contains information about the current environment, including metadata database name, schema and volume names. This information changes during the volume migration procedure.

Creating schemas in the new database

The first step in migrating one or more volumes to a different database is to create the necessary schemas in the target database, the database to which you are migrating.

How to create the schemas in the target database

- 1 Create a new system schema in the target database. The name of this schema must be different than the name of the current system schema in the source database. To see the name of the current system schema, log in to Configuration Console, choose Advanced view, then from the side menu, choose Volumes. The current system schema appears as the only schema having System as the type. The current system schema is `ac_corp_system`, as shown in Figure 4-68.




Figure 4-68 Viewing current system schema name, `ac_corp_system`

This example procedure uses the name `ac_corp_system_oracle` for the new system schema in the target database. For information on creating schemas in a pre-existing database, see Chapter 3, “Installing BIRT iServer using an alternative database.”

- 2 For each volume schema that your iServer installation contains, create a new volume schema in the target database using the same name. In this example procedure, the iServer installation has one volume schema, `ac_corp`, which is the name the procedure uses for the current volume schema in the source database, and for the new volume schema in the target database.

Creating a metadata database and schemas in Configuration Console

In this section, you use Configuration Console to perform the following tasks in iServer:

- Create a new metadata database
- Create a new system schema:
 - Create a new system schema in the new metadata database
 - Restart iServer

- Remove the system schema from the old metadata database
- Create a new volume schema or schemas:
 - Remove volumes and volume schemas from the old metadata database
 - Remove the old metadata database
 - Create a new volume schema or schemas in the new metadata database.

How to create a new metadata database in Configuration Console

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Volumes.
- 2 Point to the icon next to the Metadata Database heading and choose Add new metadata database, as shown in Figure 4-69.

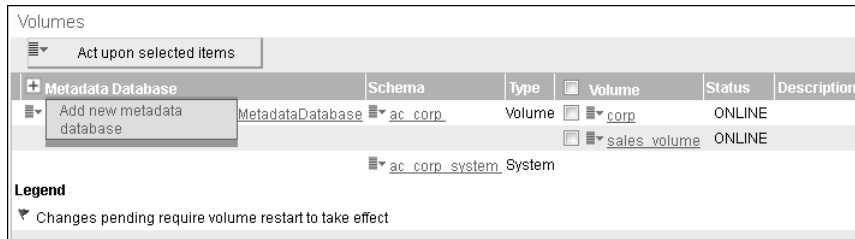


Figure 4-69 Choosing to create a metadata database

- 3 On New Metadata Database, as shown in Figure 4-70, perform the following tasks:
 - 1 In Metadata database name, type a name for the metadata database.
 - 2 In Database type, select the type of database to create.
 - 3 In Connection login, type the database user name.
 - 4 In Connection password, type the database user password.
 - 5 In Database server, type the host name of the machine hosting the database. You do not need to provide a value for this property if you have provided values for Database TNS server name and Database TNS names file.
 - 6 In Database port, specify the port number on which the database listens, or accept the default value. You do not need to provide a value for this property if you have provided values for Database TNS server name and Database TNS names file.
 - 7 In Database service name, type the service name with which to connect to the database. You do not need to provide a value for this property if you have provided values for Database TNS server name and Database TNS names file.

- 8 In Database TNS server name, type a TNS server name. This is typically the Oracle System ID (SID). You do not need to provide a value for this property if you have provided values for Database server, Database port, and Database service name.
- 9 In Database TNS names file, type the full path to a TNS names file. The TNS names file must reside on the same machine as the iServer installation. You do not need to provide a value for this property if you have provided values for Database server, Database port, and Database service name.
- 10 Choose Test to test the database connection. If the connection test is successful, choose OK.

Volumes > New Metadata Database

Metadata Database

Metadata database name: Oracle_MetadataDatabase *

Database type: Oracle *

Connection login: acTest *

Connection password: ***** *

Database server: DB2-Win

Database port: 1521

Database service name: ORA11G64.actuate.com

Database TNS server name:

Database TNS names file:

* These fields are required and cannot be left blank

Test OK Cancel Apply

Figure 4-70 Specifying values for a new Oracle metadata database

How to create a new system schema in Configuration Console

- 1 Point to the icon next to the new Oracle metadata database and choose Add system schema, as shown in Figure 4-71.

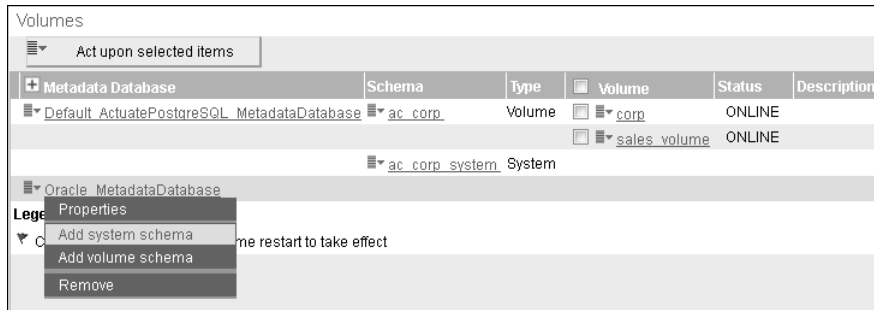


Figure 4-71 Choosing to create a new system schema

- 2 On New System Schema, as shown in Figure 4-72, perform the following tasks:
 - 1 In Schema name, type a name for the system schema. The name must be different than the name of the existing system schema.
 - 2 In Database schema name, type the name of the system schema you created in the database.
 - 3 In Database schema password, type the password for the system schema you created in the database. Then, type the password again in Database schema password confirm.
 - 4 Choose Test. If the connection test is successful, choose OK.

The 'New System Schema' dialog box contains the following information:

Schema

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

Metadata Database: Oracle_MetadataDatabase

Schema type: System

Schema name: ac_corp_system_oracle *

Database schema name: ac_corp_system_oracle *

Database schema password: *

* These fields are required and cannot be left blank

Buttons: Test, OK, Cancel, Apply

Figure 4-72 Creating a new system schema

The following message appears:

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

In the message box, choose OK.

- 3 Restart iServer by performing the following steps:
 - 1 From the side menu, choose System.
 - 2 In System—Status, choose Stop to shut down iServer. Wait for the following message to appear:

System is currently offline
 - 3 Choose Start System.
- 4 When the iServer system is back online, choose Volumes from the side menu.
- 5 Remove the system schema belonging to the old metadata database by pointing to the icon next to the schema name and choosing Remove, as shown in Figure 4-73.

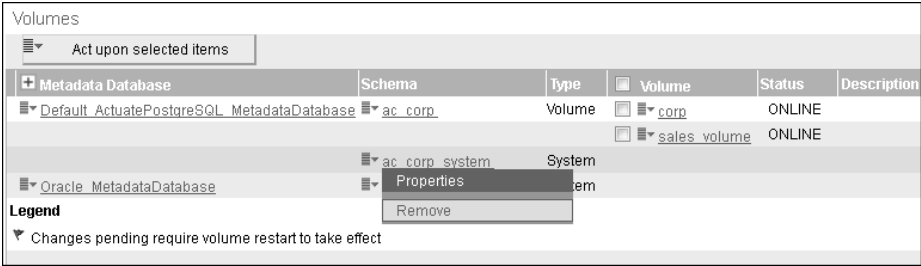


Figure 4-73 Removing the system schema

Confirm that you want to disconnect this schema from the system.

How to create a new volume schema in Configuration Console

- 1 Remove the volume or volumes and the volume schema from the old metadata database by performing the following tasks:
 - 1 Record the names of the volumes you are removing, and the partition that each volume uses. For any volume you remove in this step, you must create a new volume in a later step that has the same name and that uses the same partition. Choose Partitions from the side menu. Each partition and volume pair appear, as shown in Figure 4-74.




Figure 4-74 Viewing partition and volume pairs

- 2 Choose Volumes from the side menu. Point to the icon next to a volume and choose Take offline, as shown in Figure 4-75.




Figure 4-75 Taking a volume offline

- 3 Point to the icon next to the volume again and choose Remove, as shown in Figure 4-76.




Figure 4-76 Removing a volume

- 4 Repeat steps 1.2 and 1.3 for each volume.
- 5 Remove the volume schema by pointing to the icon next to the volume schema and choosing Remove, as shown in Figure 4-77.




Figure 4-77 Removing a volume schema

Confirm that you want to disconnect the schema from the system.

- 2 Remove the old Default_ActuatePostgreSQL_MetadataDatabase metadata database by pointing to the icon next to the metadata database name and choosing Remove, as shown in Figure 4-78.




Figure 4-78 Removing a metadata database

Confirm that you want to disconnect the metadata database from the system.

- 3 On Volumes, create a new volume schema by pointing to the icon next to the new Oracle_MetadataDatabase metadata database and choosing Add volume schema, as shown in Figure 4-79.




Figure 4-79 Choosing to create a volume schema

- 4 On Volumes—New Schema, perform the following tasks:
 - 1 In Schema name, type the name of the volume schema you created in the Oracle database.
 - 2 In Database schema name, type the name of the volume schema you created in the Oracle database.
 - 3 In Database schema password, type the password for the volume schema you created in the Oracle database.
 - 4 Choose Test. If the connection test is successful, choose OK.

Volumes—New Schema appears as shown in Figure 4-80.

Volumes > New Volume Schema

Schema

Metadata Database: Oracle_MetadataDatabase

Schema type: Volume

Schema name: ac_corp *

Database schema name: ac_corp *

Database schema password: •••••• *

* These fields are required and cannot be left blank

Test OK Cancel Apply

Figure 4-80 Creating a new volume schema

Volumes appears as shown in Figure 4-81.

Volumes

Act upon selected items

Metadata Database	Schema	Type	Volume	Status	Description
Oracle_MetadataDatabase	ac_corp	Volume			
	ac_corp_system_oracle	System			

Legend

Changes pending require volume restart to take effect

Figure 4-81 Viewing the new volume schema on Volumes

Exporting volume metadata from the source database

Exporting volume metadata consists of the following operations:

- Preparing to run the Data Store Administrator utility
 - Shut down iServer
 - Add the path of the folder containing the utility to the PATH environment variable
- Creating a properties file that you pass to the utility when you run it
- Running the Data Store Administrator utility against the source database to export the metadata

Preparing to run the Data Store Administrator utility

Perform the procedure in this section to prepare to run the Data Store Administrator utility to export the volume metadata.

How to prepare to run the Data Store Administrator utility

- 1 Shut down iServer by performing the following steps:
 - 1 In Configuration Console—Advanced view, choose System from the side menu.
 - 2 In System—Status, choose Stop to shut down iServer. Wait for the following message to appear:
System is currently offline
- 2 Edit the PATH environment variable on your machine to contain the following string:

```
<AC_SERVER_HOME>\bin
```

where AC_SERVER_HOME refers to the iServer installation path. For example, using the default value for AC_SERVER_HOME, in iServer Release 11 Service Pack 5, add:

```
C:\Program Files (x86)\Actuate11SP5\iServer\bin
```

Creating the properties file

On the machine where iServer is installed, you create a properties file, containing name-value pairs that the Encyclopedia Data Store Administrator utility uses when you run it. If you are exporting metadata for more than one volume schema, a good practice is to create a separate properties file for each schema. Properties this file contains include:

- Type of database the utility is running against.
- Name of the database containing the volume schema or schemas.
- Name of the machine hosting the database.
- Name of the schema you are running the export utility against. You run the export utility once for each volume schema that your iServer installation contains.
- Type of database to which you are migrating volume metadata.
- The path of the folder that the utility creates to contain the exported metadata. Specify a separate folder for each volume schema you run against.

For the complete list of properties that the Data Store Administrator uses, see “Specifying Encyclopedia Data Store Administrator properties,” later in this chapter.

How to create the properties file

Use the following example to create a properties file named VolumeExport.properties in the AC_SERVER_HOME\bin folder.

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/iServer
```

```

DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = ac_corp
SCHEMA_PASSWORD = <your volume schema password>
EXPORT_ALL_DATA = true
DATA_EXPORT_FOLDER = C:/Actuate/11SP5_data/ac_corp
DATA_EXPORT_FORMAT = Oracle

```

Running the Data Store Administrator utility

The utility creates the folder that the `DATA_EXPORT_FOLDER` property specifies, containing the metadata for the schema that the `SCHEMA_NAME` property specifies. For each table in the volume schema that you run the utility against, the utility creates a file in this folder.

How to run the Data Store Administrator utility

- 1 Open a command prompt and navigate to `AC_SERVER_HOME\bin`.
- 2 Run the `administrat_encyclopedia_data_store.bat` file using the following command line syntax:

```
administrat_encyclopedia_data_store VolumeExport.properties
```

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

```
admineds VolumeExport.properties
```

Completing the volume migration

After exporting the volume metadata from the source database, completing the volume migration consists of the following operations:

- Importing the volume metadata to the target database
- Creating the new volume or volumes in the new metadata database in Configuration Console

Importing the volume metadata to the target database

Perform the following procedure to import the volume metadata to the target database.

How to import volume metadata to the target database

- 1 Copy the folder or folders containing the exported volume metadata files that the Encyclopedia Data Store Administrator utility created from the iServer installation to the machine hosting the database to which you want to migrate the volume metadata.

- 2 Edit the PATH environment variable on the machine hosting the target database to contain the path specifying the location of the folder containing the exported volume metadata files. If you have multiple folders, add the path of each folder to the PATH variable.
- 3 Navigate to the location of the folder containing the exported volume metadata files. Using a text editor, open the ac_oracle_load.par file. The following example shows the information you must provide appearing between the brackets:

```
USERID=<userid>/<password>@<instance>  
SILENT=HEADER  
DIRECT=TRUE  
COLUMNARRAYROWS=1000  
READSIZE=4194304  
STREAMSIZE=4194304  
BINDSIZE=4194304
```

where:

- <userid> is typically the volume schema owner. Alternatively, it can be the user id of a user having the necessary privileges to bulk load data into the volume schema in the database.
- <password> is the password of the user that <userid> specifies
- <instance> specifies the database instance to which to connect. This can be a net service name (TNS alias name) or connect descriptor.

Save and exit the file.

- 4 Open a command prompt and navigate to the location of the folder containing the exported volume metadata files. Execute the ac_oracle_load.bat script.
- 5 Verify that the schema populated successfully.
- 6 On the machine containing the iServer installation, in Configuration Console, on System—Status, choose Start system to start iServer.

Creating a new volume

For each migrated volume, you create a new volume in the new metadata database in Configuration Console—Volumes. The new volume must have the same name that the volume had in the old metadata database, and use the same partition. For example, for a volume in the old metadata database named corp, which used a partition named DefaultPartition, the corresponding new volume you create in the new metadata database must also be named corp and must also use the partition named DefaultPartition.

How to create a new volume

- 1 In the Advanced view of Configuration Console, on Volumes, point to the icon next to a volume schema and choose Add Volume, as shown in Figure 4-82.




Figure 4-82 Choosing to add a volume

- 2 On New Volume—General, perform the following tasks:
 - 1 In Volume name, type the name of the volume. Use the same name that the volume had in the old metadata database.
 - 2 In Primary partition, select the same partition that the volume used in the old metadata database. For example, for the volume named corp, select DefaultPartition, as shown in Figure 4-83.

Volumes > New Volume

General Open Security Partitions Events

Volume name: corp *

Description:

Schedule for purging notices: HH:mm

Schedule for purging deleted files: HH:mm

Partition

Primary partition: DefaultPartition * Min Free Space: MB

Volume archive service provider

Use archive service:

Metadata database and schema

Metadata database name: Oracle_MetadataDatabase

Database schema name: ac_corp

Figure 4-83 Specifying volume and partition name

- 3 If you are not using an RSSE application, choose Partitions and skip this step. If you are using an RSSE application, choose Open Security, and select Enable as a web service, as shown in Figure 4-84. Selecting this option enables the RSSE application to map volume objects such as users, roles, and groups, to their corresponding objects in the external security source. For more information, see *Configuring BIRT iServer*.




Figure 4-84 Selecting Enable as web service if using an RSSE application
Choose Partitions.

- 4 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 4-85.



Figure 4-85 Assigning a partition
Choose OK.

- 5 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-86.




Figure 4-86 Viewing the new volume

The new volume comes online.

- 6 Repeat steps 1 through 4 for each migrated volume. The online volumes appear as shown in Figure 4-87.




Figure 4-87 Viewing the new volume after it comes online

- 7 Log in to Management Console. In Files and Folders, the data in the migrated volume appears.
- 8 If you are satisfied that the contents of the migrated volume or volumes are correct and complete, delete the folder containing the exported volume metadata files.

Working with iServer utilities

The following sections provide information on how to specify properties and perform operations using the following utilities:

- Encyclopedia Data Store Administrator
- System Data Store Administrator
- Encyclopedia Data Store Upgrader

When configuring Encyclopedia Data Store Administrator, System Data Store Administrator, or Encyclopedia Data Store Upgrader properties files, observe the following rules:

- Use / in path settings. \ is an escape character.

- Use Unicode property values. Save the property file in UTF-8 format, including the UTF-8 byte order mark (BOM).
- Use # at the start of a line to add a comment or selectively comment out properties.

When reading the property files, the utilities remove leading and trailing whitespace automatically, but preserve embedded whitespace.

Working with Encyclopedia Data Store Administrator

Use Encyclopedia Data Store Administrator to import or export volume metadata and create or delete a volume schema.

Specifying Encyclopedia Data Store Administrator properties

Table 4-1 describes the required Encyclopedia Data Store Administrator properties used to configure the environment for a manual import operation.

Table 4-1 Required Encyclopedia Data Store Administrator properties

Parameter	Description
AC_SERVER_HOME	Points to the location of the iServer binaries, specified during the BIRT iServer Release 11 installation.
DATABASE_TYPE	Type of supported RDBMS that contains the data store. Specify DB2, Oracle, PostgreSQL, or SQLServer.
LOG_FOLDER	Absolute path to the log folder.
SCHEMA_FILE_NAME	Base name of the file without the extension that contains the schema definition.
SCRIPT_HOME	Absolute path to the root of the folder hierarchy that contains the scripts and the schema definition.

Table 4-2 describes optional Encyclopedia Data Store Administrator properties. The properties used depend on the type of operation performed and the installation environment.

Table 4-2 Optional Encyclopedia Data Store Administrator properties

Parameter	Description	Default value	Supported databases
APPLICATION_USER	Application user ID for connecting to the database for normal operations. Must be a legal SQL identifier. Typically iserver. This parameter is required when creating or populating a schema. Restrict the iServer application user name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.		All
APPLICATION_USER_PASSWORD	Application user password for connecting to the database for normal operations. Required when creating a schema. The password is not encrypted.		All
CONFIG_SCHEMA_NAME	Specifies the schema definition in acserverconfig.xml. Use this property when any of the following properties is required but is not specified in the properties file: <ul style="list-style-type: none">■ APPLICATION_USER■ DATABASE_HOST■ DATABASE_INSTANCE■ DATABASE_NAME■ DATABASE_PORT■ DATABASE_TYPE■ ORACLE_TNS_NAMES_FILE■ SCHEMA_NAME	False	All

Table 4-2 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
CREATE_SCHEMA	Set to true to create a new schema, false to use an existing schema. Specifies whether to create a new Encyclopedia schema. PostgreSQL only, not supported for DB2, Microsoft SQL Server, or Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL
DATA_EXPORT_FOLDER	Absolute path to the folder to which data is exported. This parameter is required if EXPORT_DATA is true. Folder is on the database machine.		All
DATA_EXPORT_FORMAT	Format of exported data. Specify DB2, Oracle, PostgreSQL, or SQL Server.	{DATABASE_TYPE}	All
DATA_IMPORT_FOLDER	Absolute path to the folder from which data is imported. This parameter is required if IMPORT_DATA is true. For PostgreSQL, the data must be on the database server.		All
DATA_IMPORT_FORMAT	Format of imported data. Specify DB2, SQL Server, Oracle, or PostgreSQL.	{DATABASE_TYPE}	All
DATABASE_HOST	Hostname or IP address of the machine hosting the database. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
DATABASE_INSTANCE	Specifies the SQL Server database instance. The default value is the name of the database instance on the host machine and port. The default name of the SQL Server 2008 database instance is MSSQLSERVER.		SQL Server

(continues)

Table 4-2 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
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 RDBMS that contains the data store. Actuate Release 11 currently supports DB2, PostgreSQL, Microsoft SQL Server, and Oracle.		All
DEFAULT_DATABASE_NAME	Used by the superuser to connect to the default database in order to create the iServer application database. Required when the application database does not exist. PostgreSQL only.		PostgreSQL
DELETE_ALL_DATA	Set to true to delete all data from the data store. Does not delete data from other tables in the schema.	False	All
DELETE_DATA	Set to true to delete filtered data from the data store. This parameter only works when a filter value is available.	False	All
DEPOPULATE_SCHEMA	Set to true to delete data store elements such as tables, views, stored functions, and procedures from the schema. Does not remove any other objects in the schema.	False	All
DROP_SCHEMA	Set to true to delete the schema.	False	PostgreSQL
EXPORT_ALL_DATA	Set to true to export all volumes from a schema.	False	All
EXPORT_DATA	Set to true to export a single volume from a schema. Use VOLUME_NAME to specify the name of the volume to export. Optionally, use NEW_VOLUME_NAME to specify a new name for the exported volume.	False	All

Table 4-2 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
FILE_LOG_LEVEL	Minimum logging level for messages sent to the log file. This parameter only works for AcLogger, and is a standard JUL Level name. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts to perform operations instead of performing the operations directly	False	All
IMPORT_DATA	Set to true to import data into the data store.	False	PostgreSQL
INITIALIZE_DATA	Set to true to initialize the data in the data store, using the data initialization script.	False	All
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Data Store Administrator tools do not generate large logs.	5	All
LOG_FILE_NAME	Name of the log file. Do not add a file extension. The extension is set to .log. A unique number appends automatically to the file name to prevent overwriting previous logs.	Encyclopedia DataStore Administrator <number>.log	
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Data Store Administrator tools do not generate large logs.	10,000,000	All
LOG_FOLDER	Full path of folder to write logs.	AC_DATA _HOME /server/log.	

(continues)

Table 4-2 Optional Encyclopedia Data Store Administrator properties (continued)

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

Table 4-2 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default value	Supported databases
PROMPT_FOR_PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it as a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	True	All
SCHEMA_NAME	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	SCHEMA_NAME	All
SCHEMA_PASSWORD	Password for the schema owner. Required if PROMPT_FOR_PASSWORDS is false.		All
SUPERUSER	User ID of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL. Typically postgres in a PostgreSQL database.		PostgreSQL
SUPERUSER_PASSWORD	Password of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL.		PostgreSQL
TABLESPACE_NAME	Name of the application tablespace.		PostgreSQL
TABLESPACE_LOCATION	Absolute path to the application tablespace folder.		PostgreSQL
TIME_ZONE	Local time zone for installation.		All
VOLUME_NAME	Name of the volume or target schema for data import.	VOLUME_NAME	All

Using the generated bulk-load script files

When exporting volume metadata, the Data Store Administrator utility creates the folder specified by the `DATA_EXPORT_FOLDER` property. This folder contains the Encyclopedia volume metadata files exported from the volume. Data Store Administrator also creates auxiliary script files to use with RDBMS bulk-load utilities when loading Encyclopedia volume metadata into the database.

The `DATA_EXPORT_FORMAT` property specifies the database type. The database type determines the type of file the utility creates to contain the metadata as well as what files to create to support loading that data into the database using the standard RDBMS data-loading utility.

The `DATA_EXPORT_FORMAT` property supports the following values

- Oracle

The Data Store Administrator utility creates a `.dat` file for each metadata set that SQL*Loader loads into a single database table, such as user or security role. For every `.dat` file, the utility creates a `.ctl` file containing the information that SQL*Loader requires to load the contents of the associated `.dat` file, such as the name of the table to load, the character in the `.dat` file used to separate each data field, and the column names in the table.

The utility also creates the following files for use by the SQL*Loader:

- `ac_oracle_load.bat`

Contains a list of `sqlldr` commands, one for each `.dtl` and `.ctl` file pair.

- `ac_oracle_load.par`

Contains the parameters passed in each invocation of SQL*Loader. This file is passed to SQL*Loader in each `sqlldr` command that `ac_oracle_load.bat` contains. Modify this file to contain the appropriate server, user, password, port, and database information before running an `ac_oracle_load` script.

- `ac_oracle_load.sh`

Contains the UNIX equivalent of the batch file containing the list of `sqlldr` commands.

Before running `ac_oracle_load.bat` or `ac_oracle_load.sh`, the database administrator must copy the folder that `DATA_EXPORT_FOLDER` specifies to the machine hosting Oracle.

- DB2

The Data Store Administrator utility creates a `.csv` file for each type of metadata that DB2 loads into a single database table, such as user or security role.

The utility also creates the following files for use by DB2:

- `ac_db2_load.bat`

Contains a single `db2cmd` invocation that runs `ac_db2_load.db2`.

- `ac_db2_load.db2`
Contains one LOAD command for each .csv file. Each LOAD command includes all the information DB2 needs to load a single table. Modify this file to contain the appropriate server, user, password, port, and database information before running an `ac_db2_load` script.
- `ac_db2_load.sh`
Contains the UNIX equivalent of `ac_db2_load.bat`.

Before running `ac_db2_load.bat` or `ac_db2_load.sh`, the database administrator must copy the folder that `DATA_EXPORT_FOLDER` specifies to the machine hosting DB2.

- **SQL Server**

The Data Store Administrator utility creates a .dat file for each type of metadata that SQL Server loads into a single database table, such as user or security role.

The utility also creates the following files for use by SQL Server:

- `ac_sql_server_load.bat`
Contains a single `sqlcmd` invocation that runs `ac_sql_server_load.sql`. Modify this file to contain the appropriate server, user, password, port, and database information before running this batch file.
- `ac_sql_server_load.sql`
Contains one BULK INSERT command for each .dat file. This file also contains all the information SQL Server needs to load each table.

Before running `ac_sql_server_load.sql`, the database administrator must copy the folder that `DATA_EXPORT_FOLDER` specifies to the machine hosting SQL Server.

- **PostgreSQL**

The Data Store Administrator utility creates a .csv file for each type of metadata that PostgreSQL loads into a single database table, such as user or security role.

The utility also creates the following files for use by PostgreSQL:

- `ac_postgresql_client_load.bat`
Contains a single `psql` command invocation that runs `ac_postgresql_client_load.psql`. Modify this file to contain the appropriate server, user, password, port, and database information before running this batch file.
- `ac_postgresql_client_load.psql`
Contains one copy command for each .csv file.

- `ac_postgresql_client_load.sh`
Contains the UNIX equivalent of `ac_postgresql_client_load.bat`. Modify this file to contain the appropriate server, user, password, port, and database information before running this script.

Before running `ac_postgresql_client_load.bat` or `ac_postgresql_client_load.sh`, the database administrator must copy the folder that `DATA_EXPORT_FOLDER` specifies to the machine hosting PostgreSQL.

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 Data Store Administrator 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 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 (x86)/Actuate11SP5/iServer
DATABASE_TYPE = PostgreSQL
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {DATA_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 Data Store Administrator 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 (x86)/Actuate11SP5/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 = {DATA_EXPORT_FOLDER}
```

Exporting all volumes from a schema

Data Store Administrator generates bulk-load script files in the same folder as the exported data. For more information on using the generated bulk-load script files, see “Using the generated bulk-load script files,” earlier in this book.

You can omit NEW_SCHEMA_NAME if it is the same as SCHEMA_NAME.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/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

Data Store Administrator generates bulk load script files in the same folder as the exported data. For more information on using the generated bulk-load script files, see “Using the generated bulk-load script files,” earlier in this book.

You can omit NEW_SCHEMA_NAME if it is the same as SCHEMA_NAME. You can omit NEW_VOLUME_NAME if it is the same as VOLUME_NAME.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/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 (x86)/Actuate11SP5/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 (x86)/Actuate11SP5/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 (x86)/Actuate11SP5/iServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Populating an empty schema

iServer performs this operation when you create a new schema in Configuration Console. Before populating an empty schema using Encyclopedia Data Store Administrator, create the schema owner and application user if they do not exist and give them the appropriate privileges.

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/iServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
```

Creating a new volume in a populated schema

iServer performs this operation when you create a new volume in Configuration Console and bring the volume online. Typically, you perform this operation to add a volume to a shared schema that already contains one or more volumes.

When creating a new volume using Encyclopedia Data Store Administrator, set INITIALIZE_DATA to true and specify a name for NEW_VOLUME_NAME. The schema owner and application user must already exist and have appropriate privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/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 (x86)/Actuate11SP5/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Working with System Data Store Administrator

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

- Create or delete a schema.
- Populate or depopulate a schema.

Specifying System Data Store Administrator properties

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

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

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

```
<AC_SERVER_HOME>\bin
```

where `AC_SERVER_HOME` refers to the Release 11 Service Pack 5 installation path. For example, using the default value for `AC_SERVER_HOME`, add:

```
C:\Program Files (x86)\Actuate11SP5\iServer\bin
```

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

- 1 In Windows Explorer, navigate to `AC_SERVER_HOME\bin`.
- 2 Create a properties file containing the property definitions necessary to perform an operation, such as creating and populating a system schema, described in “Performing operations using the System Data Store Administrator utility,” later in this book. You pass this file to the System Data Store Administrator utility when you execute it.
- 3 Open a command prompt and navigate to `AC_SERVER_HOME/bin`.
- 4 Run the `adminstrate_system_data_store.bat` file using the following command line syntax:

```
adminstrate_system_data_store systemdatastore.properties
```

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

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

Listing 4-2 `administrate_system_data_store.bat`

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

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

:: Administrate data store
java com.actuate.iserver.system.datastore.admin
    .SystemDataStoreAdministrator %1

GOTO :END

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

:END
```

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

The `SystemDataStoreAdministrator` class has the same parent class as the `Encyclopedia Data Store Administrator` and uses the same property settings. For more information about these properties, see Table 4-1.

System Data Store Administrator properties include the following categories:

- Properties that specify details of the iServer installation environment, such as `AC_SERVER_HOME` and `AC_DATA_HOME`
- Database properties that specify the RDBMS type and JDBC connection details, such as the schema, application user, superuser, and passwords
- Schema operation properties that specify an action to perform, such as create, populate, or delete a schema
- Logging properties that control messages sent to the console and log files

- Engineering properties used by OpenText 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 the System Data Store Administrator utility

Creating and populating a new system data store schema is only supported for PostgreSQL. Creating and populating a schema requires superuser privileges. iServer performs this operation automatically when you create a new system schema in Configuration Console. The System Data Store Administrator utility creates the database and users if these items do not already exist.

Configure the properties as shown in the following example:

```
AC_SERVER_HOME = C:/Program Files (x86)/Actuate11SP5/iServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
INITIALIZE_DATA = true
```

Working with Encyclopedia Data Store Upgrader

Use Encyclopedia Data Store Upgrader to upgrade an Encyclopedia volume manually from an earlier Release 11 installation, such as Service Pack 3.

Specifying Encyclopedia Data Store Upgrader properties

Table 4-3 describes the required Encyclopedia Data Store Upgrader properties to specify in the `upgrade_encyclopedia_data_store.bat` or other properties file for a manual upgrade operation.

Table 4-3 Required Encyclopedia Data Store Upgrader properties

Parameter	Description
AC_SERVER_HOME	Points to the location of the iServer binaries, which you specify during the BIRT iServer Release 11 installation.
APPLICATION_USER	User ID used to connect to the database for normal operations.
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports PostgreSQL and Oracle.
LOG_FOLDER	Absolute path to the log folder.
SCHEMA_FILE_NAME	Base name of the file without the file extension that contains the meta-schema definition.
SCHEMA_NAME	Name of the target schema which the Encyclopedia Data Store Upgrader updates. Required if NEW_SCHEMA_NAME is not present. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.
SCRIPT_HOME	Absolute path to the root of the folder hierarchy containing scripts and the meta-schema definition.

Table 4-4 describes the optional Encyclopedia Data Store Upgrader properties to specify in the upgrade_encyclopedia_data_store.bat or other properties file.

Table 4-4 Optional Encyclopedia Data Store Upgrader properties

Parameter	Description	Default value	Supported databases
CONFIG_SCHEMA_NAME	Specifies the schema definition in acserverconfig.xml. The schema name can be different from the database schema name.	False	All
DATABASE_HOST	Hostname or IP address of the machine hosting the database. This value is required for PostgreSQL. Required for Oracle if not using TNS.	False	All
DATABASE_INSTANCE	RDBMS instance that manages the database.		SQL Server

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

Parameter	Description	Default value	Supported databases
DATABASE_NAME	Database name for PostgreSQL. For Oracle not using TNS, the database service name. For Oracle using TNS, the TNS net service name.		All
DATABASE_PORT	Port that the database server uses. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
FILE_LOG_LEVEL	The minimum logging level for messages sent to the log file. This parameter only works for AcLogger. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts to perform operations instead of performing the operations directly	false	All
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Encyclopedia Data Store Upgrader tool does not generate large logs.	5	All
LOG_FILE_NAME	Name of the log file. Do not add a file extension. The extension is set to .log. Standard JUL.FileHandler placeholders are supported. A unique number is appended automatically to the file name to prevent overwriting previous logs.		All
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Encyclopedia Data Store Upgrader tool does not generate large logs.	10,000,000	All

(continues)

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

Parameter	Description	Default value	Supported databases
LOG_FOLDER	The full path of folder to write logs.	AC_DATA_HOME /server/log.	All
ORACLE_TNS_NAMES_FILE	Absolute path of the Oracle TNS names file to use instead of DATABASE_HOST and DATABASE_PORT to generate a JDBC URL.		Oracle
PROMPT_FOR_PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it as a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	True	All
SCHEMA_PASSWORD	Password for the schema owner. Required if PROMPT_FOR_PASSWORDS is false.		All

5

Installing a BIRT iServer cluster

This chapter discusses the following topics:

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

Installing a BIRT iServer cluster node

A node is a machine running a BIRT iServer instance. An iServer administrator adds a node to an iServer cluster to improve availability and throughput and scale the cluster installation to necessary processing requirements.

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

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

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

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

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

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

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

The cluster communicates across the network using standard HTTP/IP addressing. The cluster automatically detects the on-off status of any node. Single-point node failure does not affect the availability of other nodes.

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

iServer instances running on multiple machines maintain iServer system and Encyclopedia volume metadata in databases and control access to shared volume

data. The volume data can be on machines that are not running iServer, but must be shared and accessible to each iServer instance.

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

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

This operational model lends itself well to grid, cloud, and other data-center types of deployments. For more information about the cloud computing deployment option, see Chapter 6, “Installing BIRT iServer in a cloud,” later in this book. For more information about administering an installed iServer cluster, see *Configuring BIRT iServer*.

Preparing the iServer cluster environment

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

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

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

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

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

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

How to share the encyc and config\11SP5 folders

In an iServer Release 11 Service Pack 5 installation, cluster nodes must have read-write access to AC_DATA_HOME\encyc and AC_DATA_HOME\config\11SP5 folders on node1. To give a cluster node read-write access to the folders and files in the encyc folder and to the files in the config\11SP5 folder, perform the following tasks on node1:

- 1 Using Windows Explorer on node1, right-click the folder, AC_DATA_HOME\encyc. Choose Properties. On encyc Properties, choose Sharing. Sharing appears as shown in Figure 5-1.

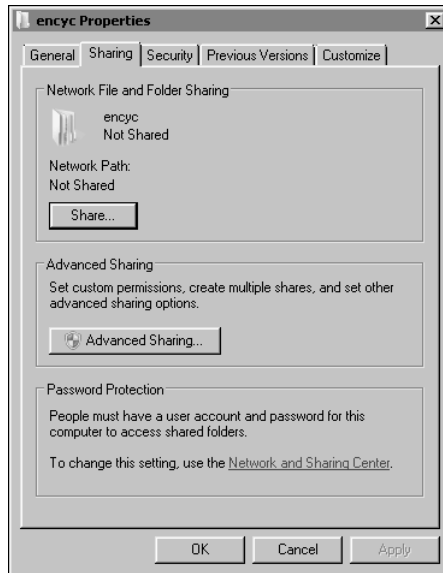


Figure 5-1 Sharing the Encyclopedia volume

On encyc Properties, choose Advanced Sharing.

- 2 On Advanced Sharing, select Share this folder, as shown in Figure 5-2.

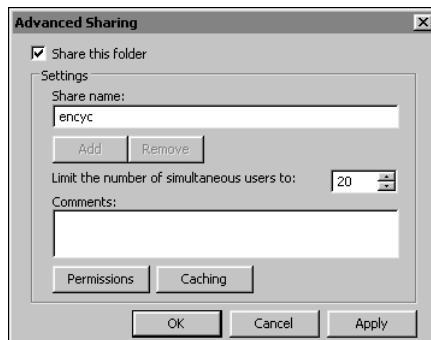


Figure 5-2 Selecting Share this folder

On Advanced Sharing, choose Permissions.

- 3 On Permissions for encyc—Share Permissions, select the user who installed BIRT iServer. Select Change and Read to allow this user these permissions, as shown in Figure 5-3. Make sure that all settings conform to the security policies in force for the environment.

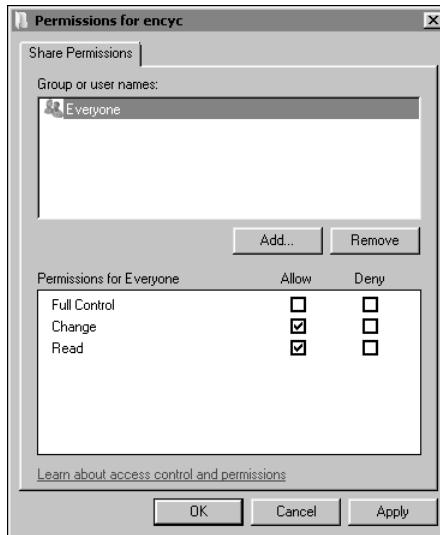


Figure 5-3 Selecting Change and Read permission

Choose OK.

On Advanced Sharing, choose OK.

On encyc Properties, choose Close.

- 4 Repeat steps 1 through 3 for the \config\11SP5 folder. Make sure that all settings conform to the security policies in force for the environment.

Close Windows Explorer.

How to turn off the Windows firewall

Perform the following steps on node1 and node2:

- 1 Choose Start→Control Panel→System and Security→Windows Firewall.
- 2 On Windows Firewall, choose Turn Windows Firewall on or off. Make sure that all settings conform to the security policies in force for the environment.
- 3 For example, on Customize Settings, in Home or work (private) network location settings, choose Turn off Windows Firewall, as shown in Figure 5-4.

Choose OK.

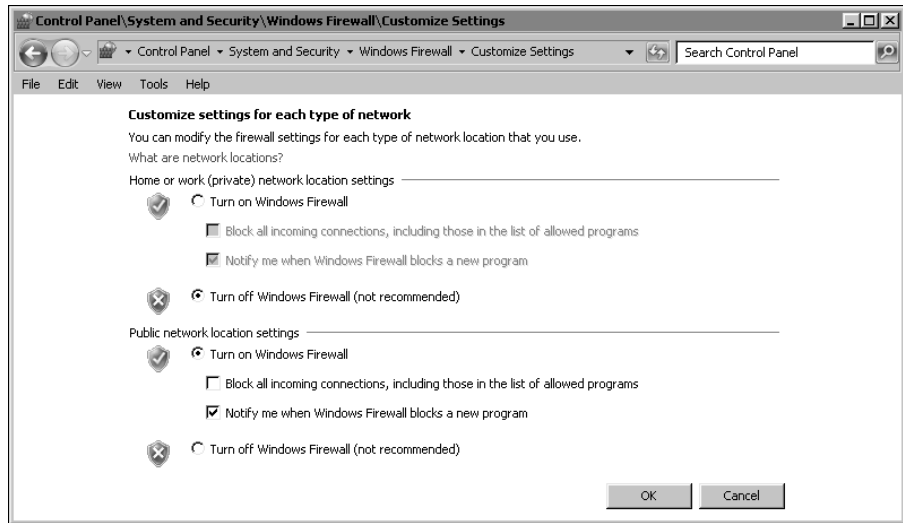


Figure 5-4 Turning off the Home or Work network location firewall
Close Windows Firewall.

How to display a computer's IP address

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

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

```
ipconfig /all
```

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

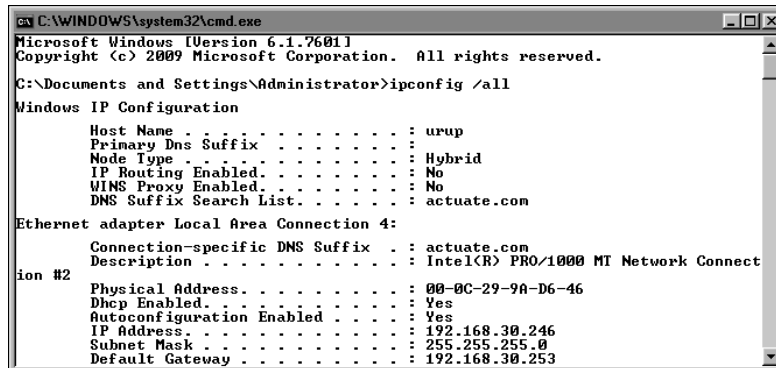


Figure 5-5 Displaying the host name

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

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

iServer	Host name	IP address
Node1	urup	192.168.30.246
Node2	kozu	192.168.30.233

How to test the connection between computers

Perform the following steps on both computers:

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

```
ping kozu
```

Press Enter.

If your computer reaches the other computer, Command Prompt displays a series of replies, as shown in Figure 5-6.

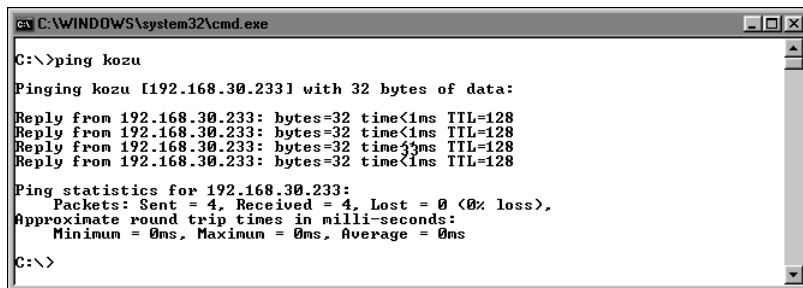


Figure 5-6 Receiving a reply to a ping command

- 2 Close Command Prompt.

How to prepare node1 for clustering

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

- 1 Log in to Configuration Console and choose Advanced view. From the side menu, choose Volumes. On Volumes, point to the icon next to the default volume, corp in this example, and choose Take offline, as shown in Figure 5-7.

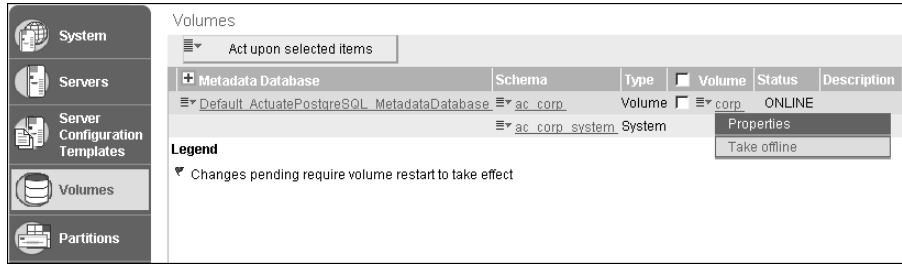


Figure 5-7 Taking a volume offline

In Volume Offline Grace Period, choose OK.

From the side menu, choose Partitions.

- 2 On Partitions, point to the icon next to DefaultPartition and choose Template settings, as shown in Figure 5-8.




Figure 5-8 Choosing Template settings

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

\\urup\encyc

Template Settings appears as shown in Figure 5-9.

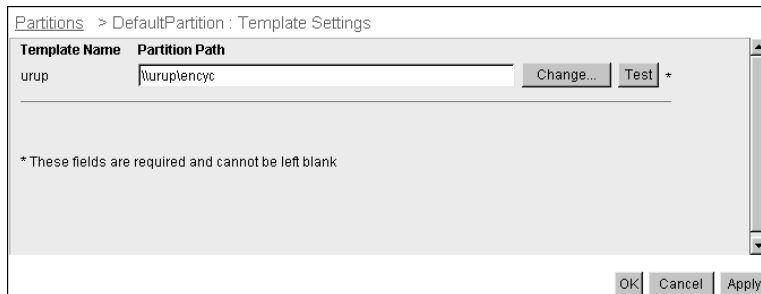


Figure 5-9 Typing the partition path

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

A message appears, stating that the test was successful, as shown in Figure 5-10. Choose OK.



Figure 5-10 Verifying a successful partition path test

On Template Settings, choose OK.

From the side menu, choose System.

- 3 On System—Status, choose Stop to stop the system, as shown in Figure 5-11.




Figure 5-11 Choosing to stop the system on System Status

- 4 Log out of Configuration Console.

- 5 Choose Start→Settings→Control Panel.

In Control Panel, choose Administrative Tools—Services.

On Services, stop the Actuate 11 BIRT iServer service.

- 6 Open Windows Explorer, and navigate to AC_DATA_HOME/config/11SP5. Make a backup copy of acserverconfig.xml, then open acserverconfig.xml and perform the following tasks:

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

- 2 Under <ConnectionProperties>, locate:

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

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

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

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

Creating an account with Windows administrator privileges

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

The iServer user account must meet the following requirements:

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

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

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

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

Performing a wizard-based cluster node installation

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

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

How to install a cluster node in Windows

- 1 Download the iServer distribution package. Extract the files. Run the self-extracting executable file, ActuateBIRTiServer.exe. The welcome message appears, as shown in Figure 5-12. Choose Next.



Figure 5-12 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 5-13. Choose Next.



Figure 5-13 Accepting the license agreement

- 3 The setup installs the prerequisite components that BIRT iServer requires, as shown in Figure 5-14. Choose Next.



Figure 5-14 Installing Prerequisites

- 4 In Setup Type, select Custom, as shown in Figure 5-15.

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 (x86)/Actuate11SP5/iServer`.

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

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

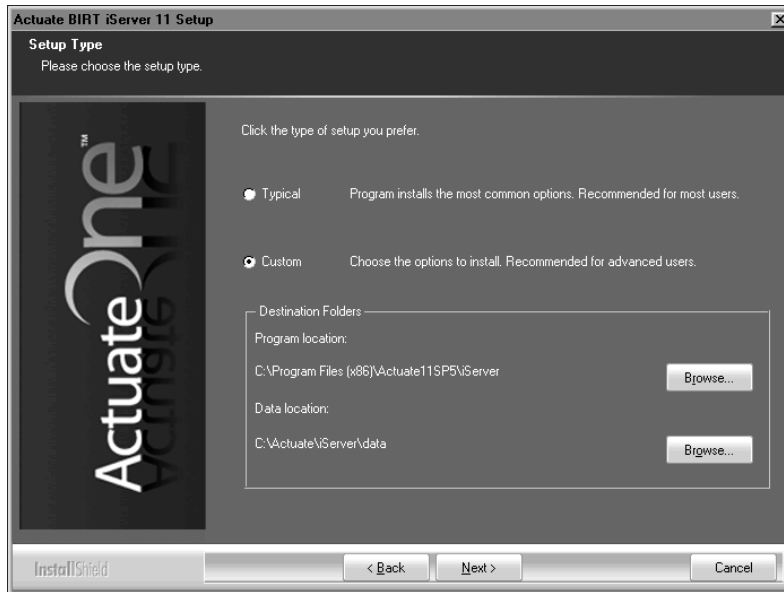


Figure 5-15 Selecting the custom installation

- 5 In Select Features, select all features—Core products, Management Console, AFC files from previous releases, and Examples, as shown in Figure 5-16. 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 5-17.

The configuration home location is the shared directory containing the `acserverconfig.xml`, `acserverlicense.xml`, and `acconfigowner.lock` files for the cluster. The administrator should specify the location using the Universal Naming Convention (UNC) Format.

In a Release 11 Service Pack 5 installation, the configuration files are located in `AC_DATA_HOME\config\11SP5` by default. If the shared configuration folder is in this default location on a server named `urup`, then the administrator should specify `\\urup\11SP5` as the path.

Choose Next.

- 7 In License File Details, select Use the license that you purchased. Choose Browse then navigate to and choose the license file, as shown in Figure 5-18. Choose Next. Alternatively, choose Try out the product using the included evaluation software license if you do not have a purchased license.

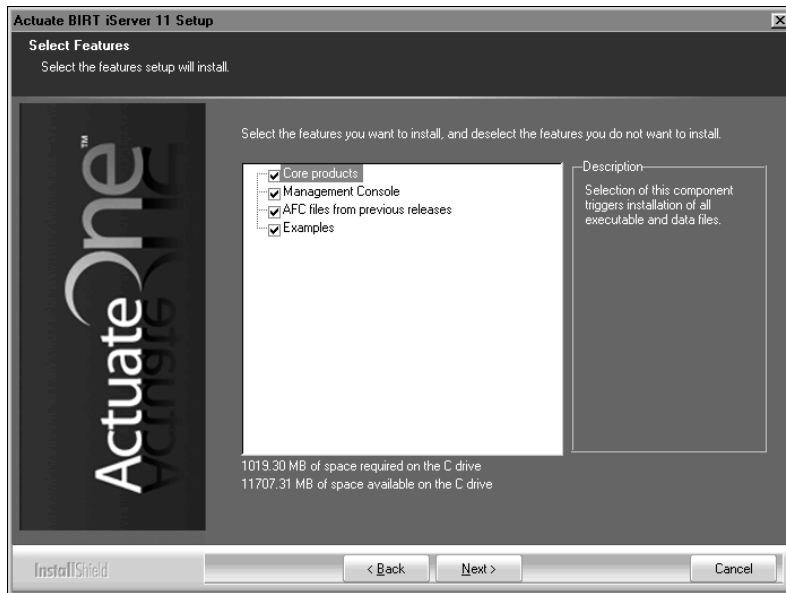


Figure 5-16 Selecting features to install

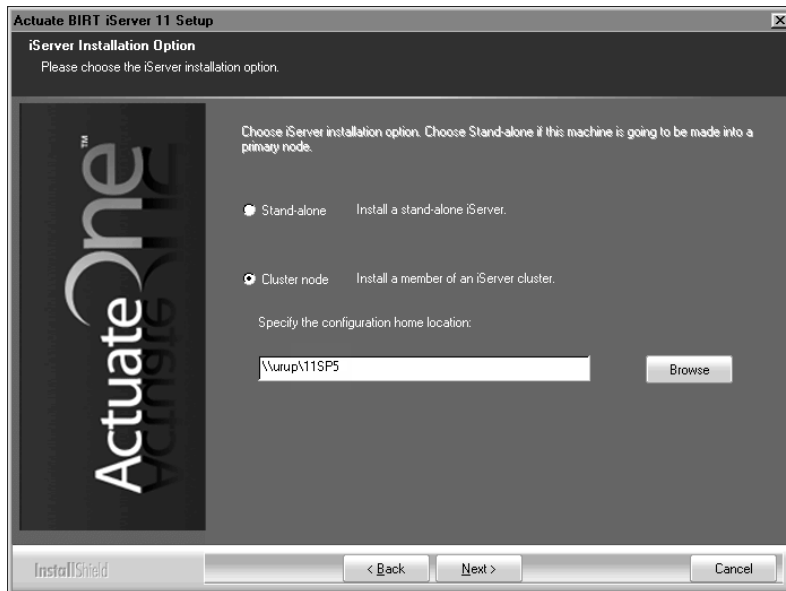


Figure 5-17 Choosing the option to install a node

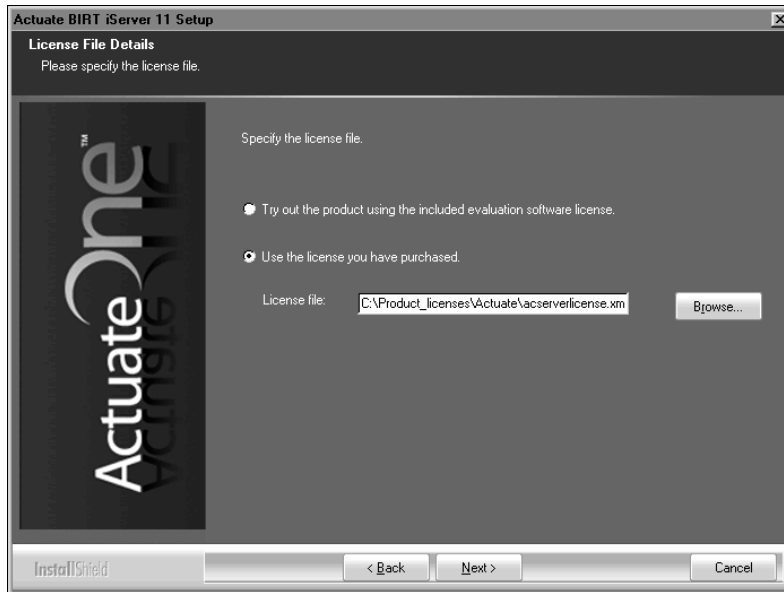


Figure 5-18 Specifying the license file

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:

You are using a named user license.

Please make sure that the internal or external Encyclopedia volume does not exceed the number of registered users authorized by the license.

Please contact OpenText Support or support@opentext.com for further assistance.

Click Yes to continue or No to abort the installation.

Choose Yes to continue the installation.

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

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

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



Figure 5-19 Specifying a language

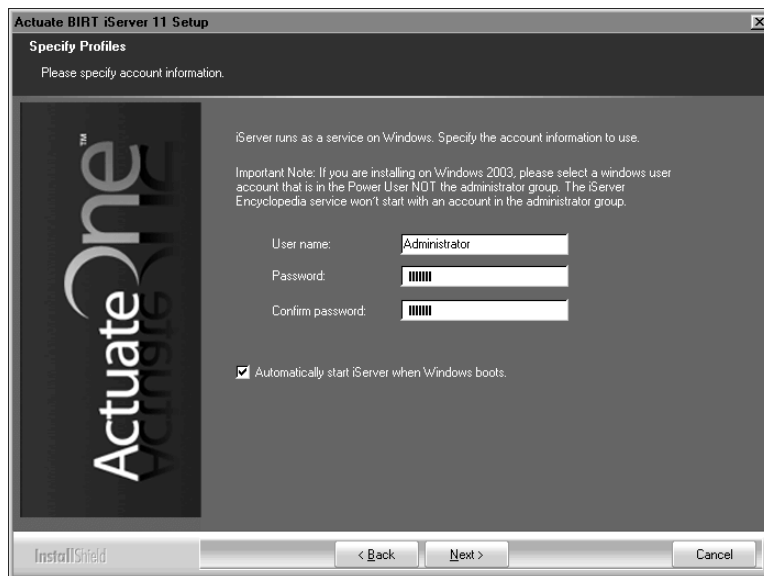


Figure 5-20 Specifying the account for running the iServer service

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




Figure 5-21 Setting the Windows local security policy

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

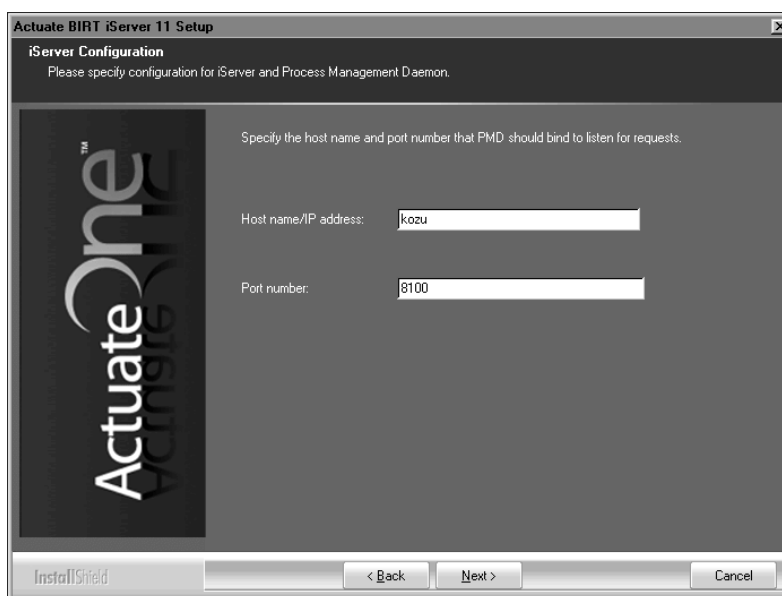


Figure 5-22 Identifying the cluster node and port to bind iServer service

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

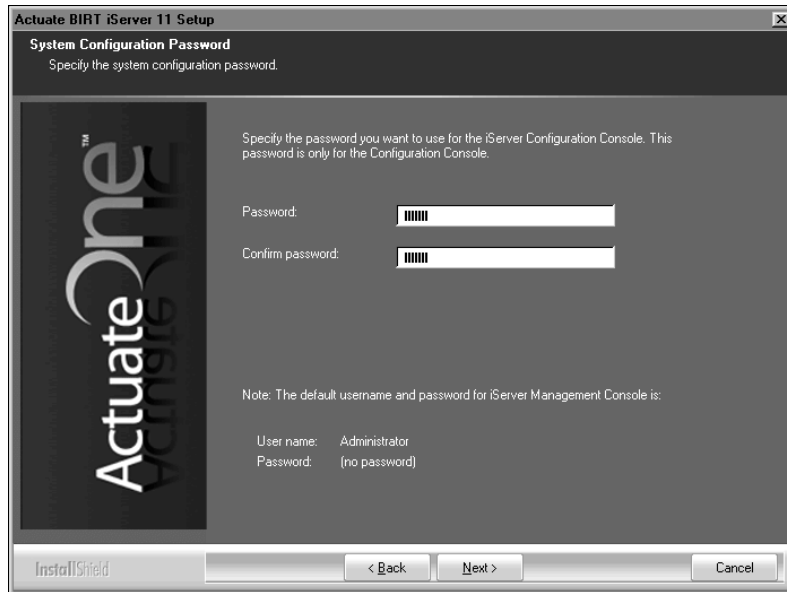


Figure 5-23 Specifying the password for using Configuration Console

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

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

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

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

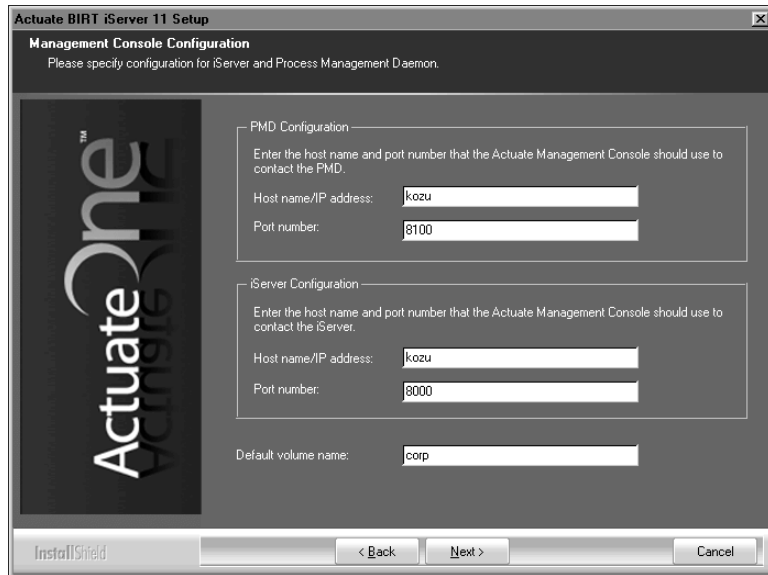


Figure 5-24 Specifying the Management Console Configuration

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

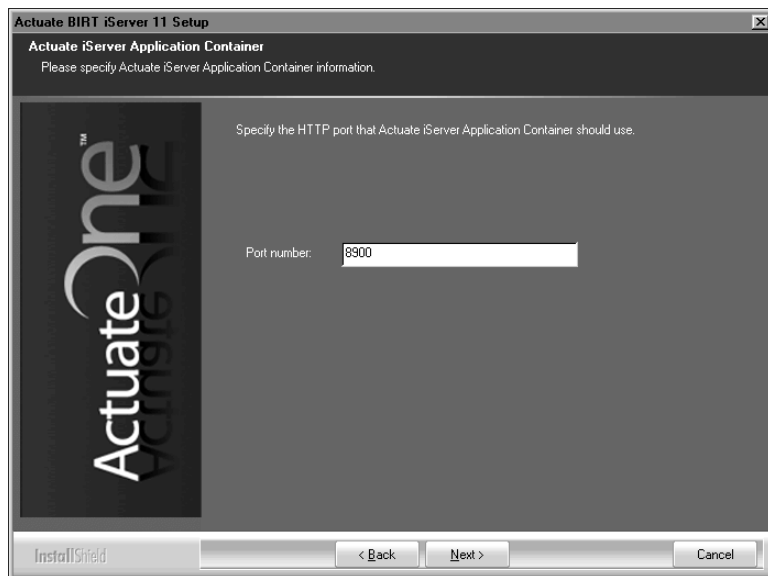


Figure 5-25 Specify the port number for iServer Application Container

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



Figure 5-26 Specifying the context path

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

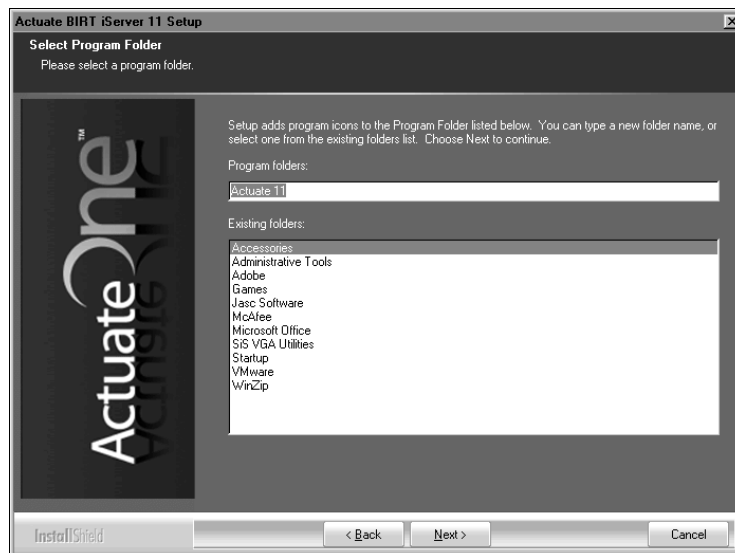


Figure 5-27 Specifying a program folder

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

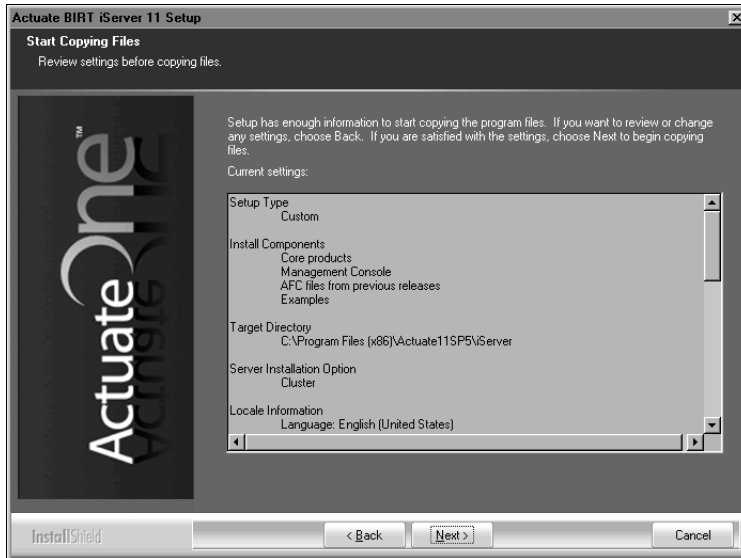


Figure 5-28 Reviewing settings before copying files

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



Figure 5-29 Viewing setup status

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



Figure 5-30 Exiting the installation wizard

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

<http://www.actuate.com>




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

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

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




Figure 5-32 Viewing BIRT iServer shortcuts on the desktop

These shortcuts provide access to the following iServer components:

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

Adding a node to a cluster

This section assumes the Administrator has completed the following tasks:

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

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

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

How to add a newly installed node to a cluster

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

- 1 Choose Start→Settings→Control Panel.

In Control Panel, choose Administrative Tools—Services.

On Services, stop the Actuate 11 BIRT iServer service.

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

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

Listing 5-1 Viewing the shared acserverconfig.xml

```
<Templates>
  <Template
    Name="urup"
    PMDPort="8100"
    ActuateBuild="110F161206"
    ActuateVersion="11 Service Pack 5"
    ServerSOAPPort="11100"
    AppContainerPort="8900"
    RequesterRSAPIVolume="corp">
    ...
  </Template>
</Templates>
```

- 2 Verify that <AC_CONFIG_HOME> points to the shared configuration home directory for the cluster. This is the path you specified for the configuration home location during the install procedure, as shown in Figure 5-17.

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

Listing 5-2 Modifying acpmdconfig.xml

```
<!--Server information -->
<Server>
  <Startup>Manual</Startup>
  <AC_TEMPLATE_NAME>urup</AC_TEMPLATE_NAME>
  <AC_DATA_HOME>C:\Actuate\iServer\data</AC_DATA_HOME>
  <AC_CONFIG_HOME>\\urup\11SP5</AC_CONFIG_HOME>
  ...
</Server>
<!-- Servlet Container information -->
```

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

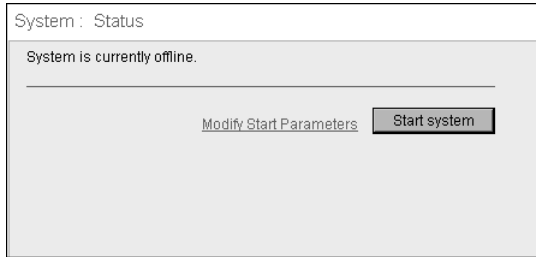


Figure 5-33 Starting the server

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

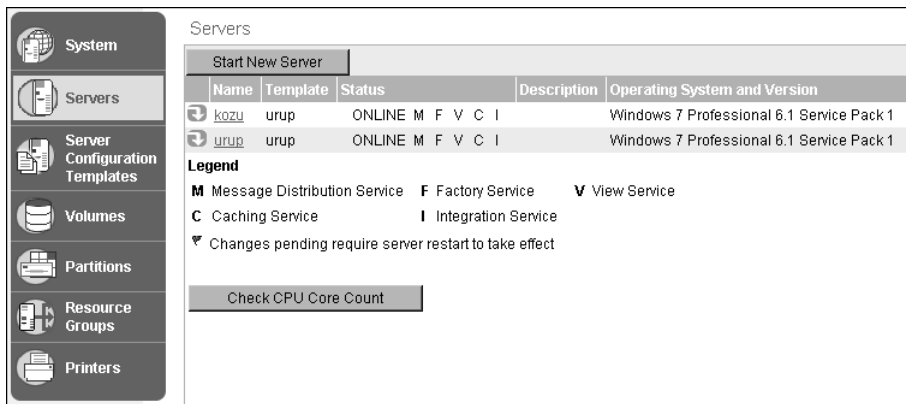


Figure 5-34 Viewing a cluster composed of urup and kozu

Installing 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 a download

If you download an OpenText product from OpenText My Support, keep all the files together in their original relative locations when you extract them. The installation program verifies that all necessary files are present before installing Information Console.

If any files are missing, the installation program exits. Files can be missing if you extract the files, move only some of the files to a new location, and attempt to install from that location.

About performing a full installation

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

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

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

To access online documentation, such as the online help and PDF files of the product manuals, install the documentation files from the following location:

<https://support.opentext.com>

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

Installing Information Console on Windows

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

Information Console installation requires the following information:

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

You can install Information Console in two ways:

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

Using the installation wizard

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

How to install using the wizard

- 1 Download the Information Console distribution package. Extract the files. Run the self-extracting executable file, `ActuateInformationConsole.exe`. The welcome message appears, as shown in Figure 6-1. Choose Next.



Figure 6-1 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 6-2. Choose Next.



Figure 6-2 Accepting the license agreement

- 3 In Setup Type, select Typical, as shown in Figure 6-3. Choose Next.

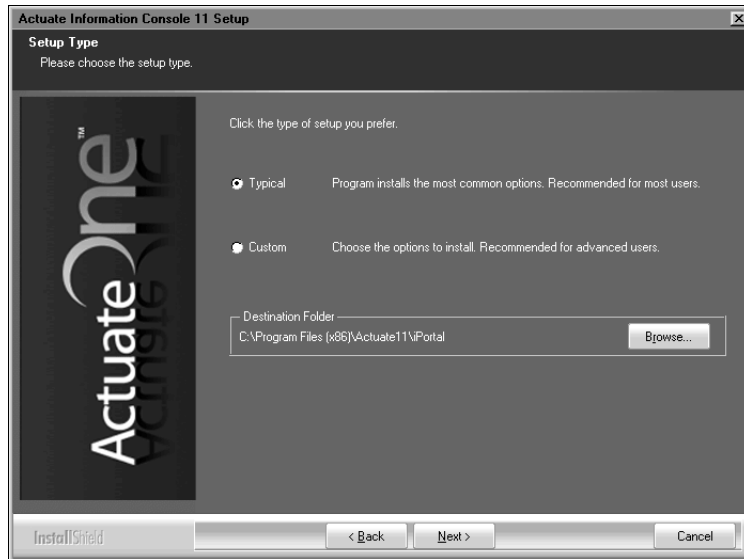


Figure 6-3 Specifying the typical or custom setup type

- 4 In Locale Information, choose Next to accept the default language and time zone, as shown in Figure 6-4, or specify these settings for your region.

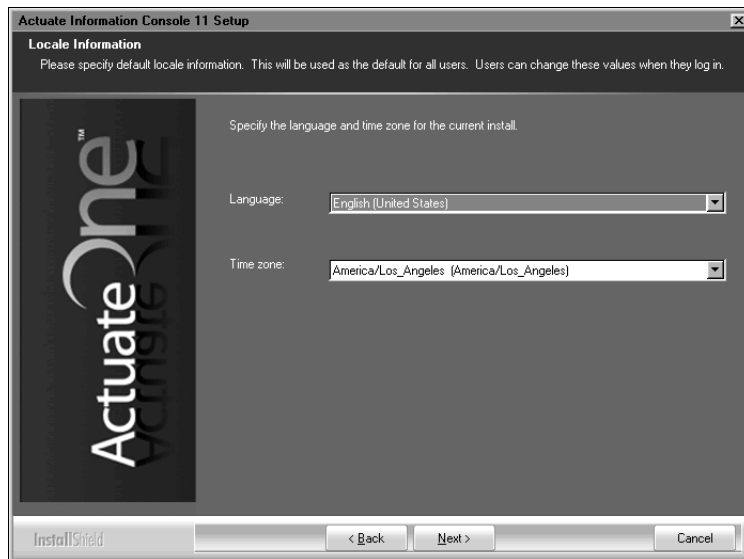


Figure 6-4 Specifying locale information

- 5 In Apache Tomcat for Actuate Information Console Service, accept the default port, 8700, and context path, /iportal, as shown in Figure 6-5. Choose Next.

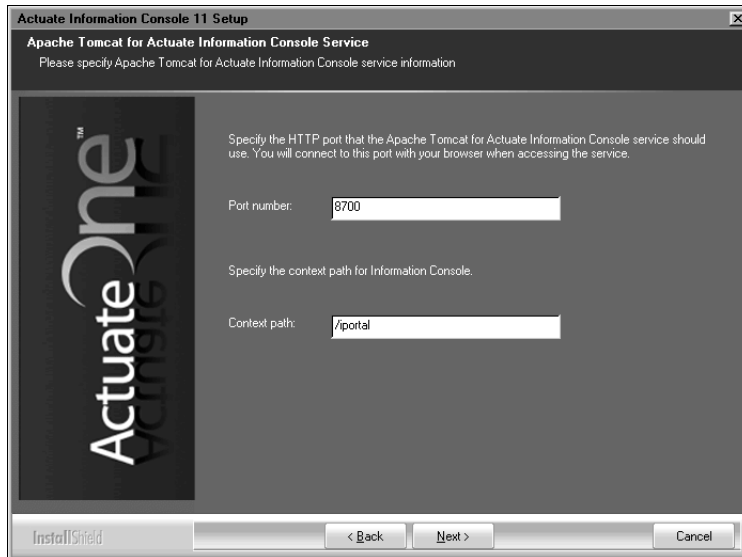


Figure 6-5 Specifying the port number and context path

- 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 6-6. Choose Next.

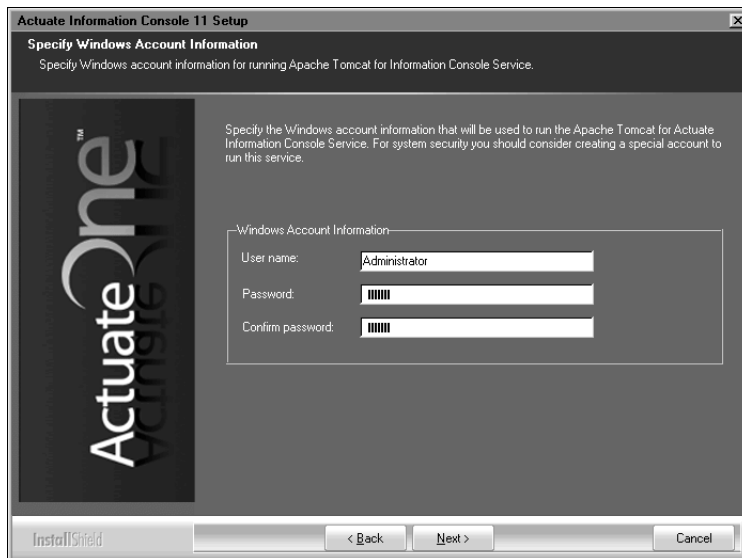


Figure 6-6 Specifying the account running the Information Console service

If prompted to add Log on as a service privilege, choose Yes, as shown in Figure 6-7.




Figure 6-7 Setting the Windows local security policy

- 7 In iServer Information, accept the default values for Profile name, Host name, and Port number, as shown in Figure 6-8. Alternatively, type a different value for one or more of these properties. In Volume name, type a name for the default Encyclopedia volume. Choose Next.

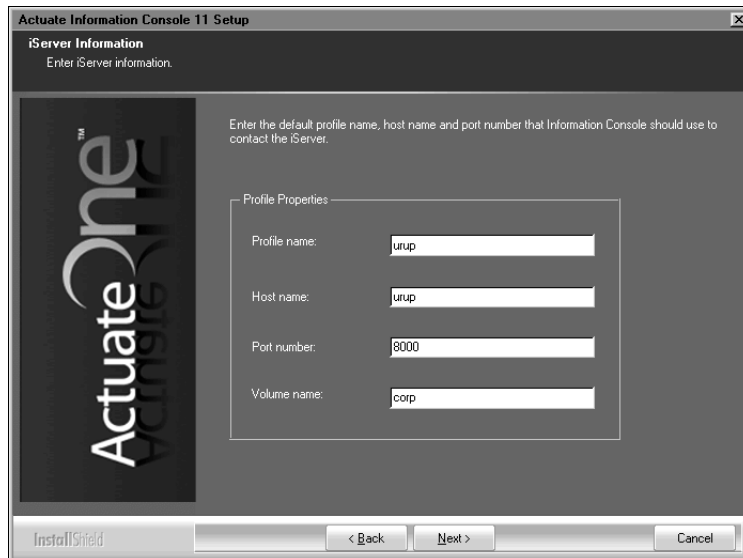


Figure 6-8 Specifying values for iServer profile properties

- 8 In Start Copying Files, review the settings shown in Figure 6-9. Choose Next.



Figure 6-9 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 6-10.



Figure 6-10 Setup Status

- 9 Select I would like to view the ReadMe file, as shown in Figure 6-11, if you want to review this documentation. Choose Finish to exit the wizard.

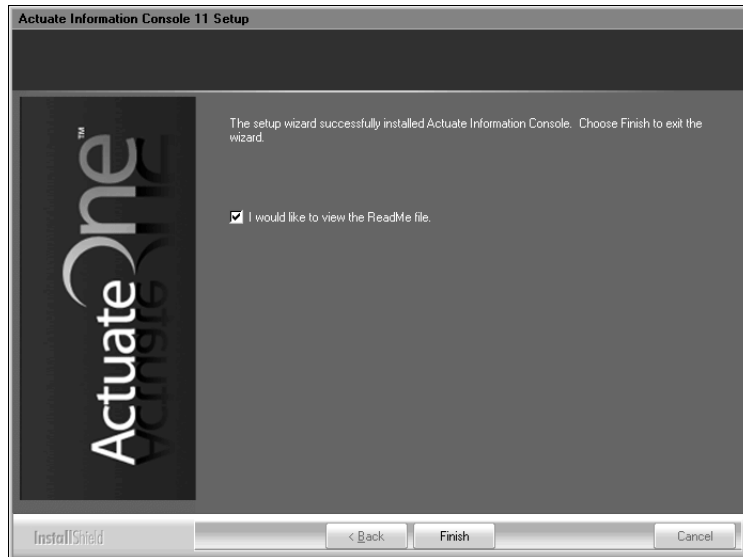


Figure 6-11 Exiting the installation wizard

If you chose to view the ReadMe file, the installation program opens the document. The Readme file contains the locations of online resources such as release notes, documentation, and access to Customer Support.

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 OpenText My Support:

<https://support.opentext.com>

Customize Information Console for your local environment, if necessary, before beginning deployment. To customize the application for your local environment, follow the steps in “Preparing the WAR file,” later in this section. Deploy the customized Information Console WAR file to the application server instead of the ActuateInformationConsole.war file that ships with Information Console.

General deployment tasks

You must configure the Information Console WAR and application server for integration. You must also verify that you have a standard J2EE installation.

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

- Configure the application server to operate with Actuate products.

- Configure the Information Console WAR to integrate with the application server as described in “Preparing the WAR file,” later in this chapter.
- Deploy Information Console to the application server.
 - If the application server has deployment tools, use those tools to integrate the application server with Actuate products. For more information about deployment tools, see 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 file to deploy Information Console to many supported servers. Table 6-1 describes the Information Console configuration parameters to review and update before deployment.

Table 6-1 Information Console configuration parameters

Parameter name	Description	Action
BIRT _RESOURCE _PATH	The location of the standard templates and properties files that BIRT Studio uses. This location can be in a WAR file or on a disk.	If you specify a location on disk, copy the contents of the Information Console resources folder to this physical location on the file system.
DEFAULT _LOCALE	The default locale is en_US. You can leave this value unchanged. A user can select a locale at login.	If you change the locale, select the new locale from the locales in /WEB-INF/Localemap.xml.
DEFAULT _TIMEZONE	The default time zone is Pacific Standard Time (PST). You can leave this value unchanged. A user can select a time zone at login.	If you change the time zone, select the new time zone from the time zones in the TimeZones.xml file, in the WEB-INF directory.

Table 6-1 Information Console configuration parameters

Parameter name	Description	Action
SERVER_DEFAULT	This value specifies the iServer URL to which the Information Console application connects if you do not specify a server URL. The default value is http://localhost:8000	Update this value to the machine and port of the server. Change localhost to the iServer machine name or IP address. Change 8000 to the iServer port number.
DEFAULT_VOLUME	This value specifies the default Encyclopedia volume for Information Console. If you do not specify a volume in an Information Console URL, the JSP application attempts to log in to this volume. The default value is localhost.	Update this value to the name of an Encyclopedia volume.
BIRT_VIEWER_LOG_DIR	The location of the files that log Actuate BIRT viewer activity.	You must create this physical location on the file system.
LOG_FILE_LOCATION	The location of the files that log Information Console activity.	You must create this physical location on the file system.
TEMP_FOLDER_LOCATION	The location where Information Console creates temporary files.	You must create this physical location on the file system.
TRANSIENT_STORE_PATH	The location where Information Console creates temporary files.	You must create this physical location on the file system.

How to customize the WAR file

The following steps describe the general procedure for customizing an Information Console WAR file:

- 1 Create a temporary directory, such as C:/Temp/ic.
If you use an existing directory, ensure that this directory is empty.
- 2 Open a command window and type the following commands, replacing the X: drive letter with a drive letter appropriate to your system:
 - 1 Copy the appropriate Information Console WAR file for your environment, such as TOMCAT_ActuateInformationConsole.war, to the temporary directory, as shown in the following example:


```
cd C:/Temp/ic
copy X:/TOMCAT_ActuateInformationConsole.war .
```
 - 2 Decompress the file, as shown in the following example:


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

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

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

If you used the temporary path in step 2, the file location is C:/Temp/ic/WEB-INF/Web.xml. Refer to Table 6-1 for a list of entries to modify in web.xml.

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

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

This command creates newinformationconsole.war in the /temp directory. This new WAR file for Information Console contains the modified configuration values.

Use newinformationconsole.war to deploy Information Console to your application server.

About clusters of servers

If the application server supports clustering, see the application server documentation for more information about setting up clusters and deploying web applications such as Information Console.

Avoiding cache conflicts after installing

Information Console uses Java Server Page (JSP) technology. Application servers and browsers cache pages. A browser can use a cached copy of a page instead of the new page. After you install, using a cached copy of some pages can lead to errors or missing functionality.

To avoid this problem, clear the application server cache after you install Information Console. With some application servers, you must restart the application server. For more information about the necessary steps for clearing the cache, see the application server documentation.

If possible, also clear the browser cache to ensure that the browser does not use an old copy of the page from its cache. Alternatively, you can refresh recently visited pages or clear the browser cache if a problem occurs. For more information about clearing the browser's cache and refreshing a page, see the browser documentation.

Testing the installation

Complete the steps in the following section to test the Information Console installation.

How to test the installation

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

```
http://Actuate1:8700/iportal/getfolderitems.do
?repositoryType=Enterprise&volume=volume1&serverurl=
http://iServer1:8000
```

where

- Actuate1:8700 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 Volume Profile, Language, and Time zone.
- 4 Choose Log In.

Installing iServer Integration Technology and documentation

This chapter discusses the following topics:

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

Installing iServer Integration Technology

This section describes how to install iServer Integration Technology for Windows. In a default installation, iServer Integration Technology installs in C:\Program Files (x86)\Actuate11\ServIntTech.

How to install

To install iServer Integration Technology, perform the following tasks:

- 1 Download the iServer Integration Technology distribution package. Extract the files. Run the self-extracting executable file, ActuateiServerIntegrationTechnology.exe. The welcome message appears, as shown in Figure 7-1. Choose Next.



Figure 7-1 Viewing the welcome message

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



Figure 7-2 Accepting the license agreement

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

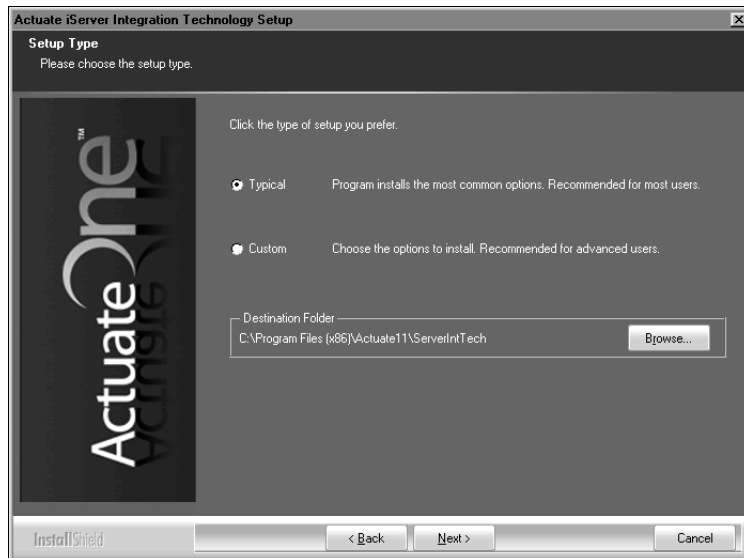


Figure 7-3 Selecting a typical installation

- 4 In Start Copying Files, review the settings shown in Figure 7-4. Choose Next.



Figure 7-4 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 7-5.

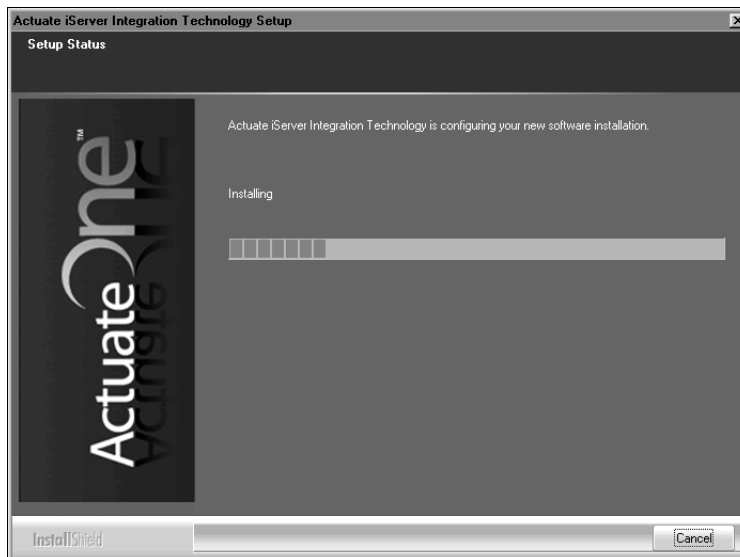


Figure 7-5 Viewing Setup Status

- 5 When the setup success message appears, select I would like to view the ReadMe file, as shown in Figure 7-6, if you want to review this documentation. Choose Finish to exit the wizard.



Figure 7-6 Exiting the installation wizard

If you chose to view the ReadMe file, the installation process opens the document. The Readme file contains the locations of online resources such as release notes, documentation, and access to Customer Support.

- 6 The installation program prompts you to install the online help from the following location, as shown in Figure 7-7:

<http://www.actuate.com>




Figure 7-7 Viewing the install online help and manuals prompt

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

Installing the localization and documentation files

The information in the 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 OpenText My Support:

<https://support.opentext.com>

If you do not see an update tool for your release on the Support site, no updates exist for the release.

How to install the localization and documentation files

To install the iServer localization and documentation files, perform the following tasks:

- 1 Download the Actuate Localization and Documentation distribution package. Extract the files. Run the self-extracting executable file, `ActuateLocalizationandOnlineDocumentation.exe`. The welcome message appears, as shown in Figure 7-8. Choose Next.



Figure 7-8 Viewing the welcome message

- 2 Read and accept the license agreement, as shown in Figure 7-9. Choose Next.



Figure 7-9 Accepting the license agreement

- 3 In Setup Type, select Typical, as shown in Figure 7-10, or select Custom to specify a limited set of localization and documentation files. Choose Next.

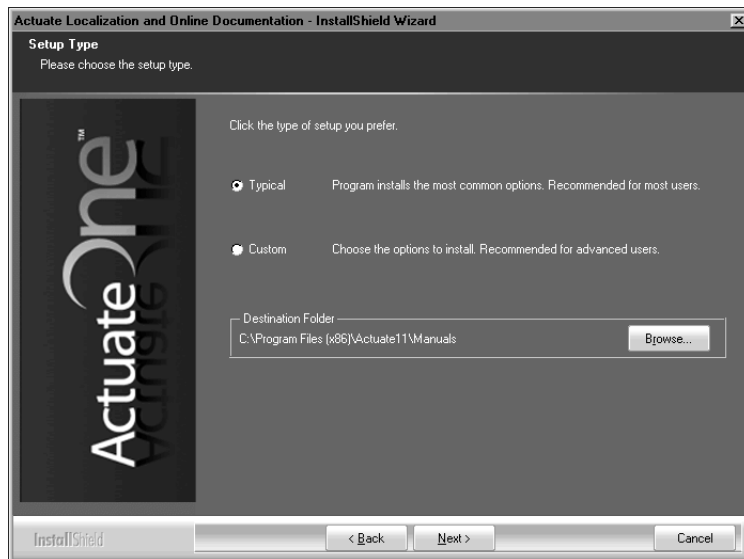


Figure 7-10 Specifying typical or custom setup type

- 4 In Start Copying Files, review the settings shown in Figure 7-11. Choose Next.



Figure 7-11 Reviewing settings before copying files

Setup Status displays an indicator showing how the installation is progressing, as shown in Figure 7-12.

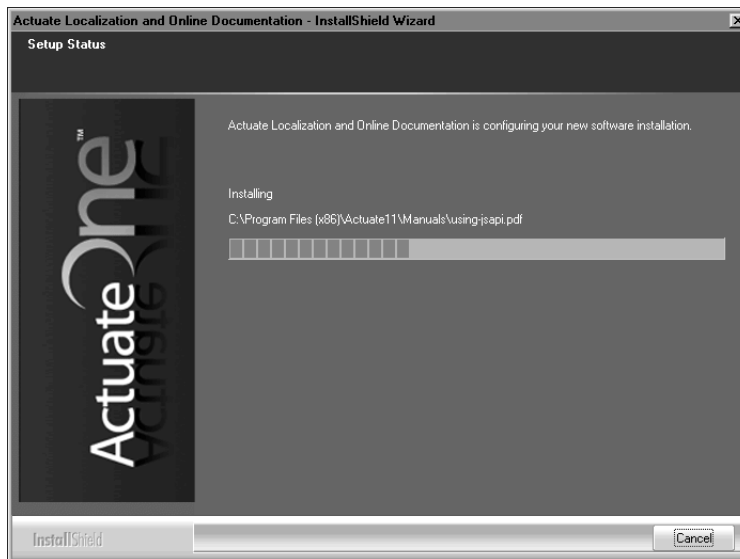


Figure 7-12 Viewing Setup Status

Setup completed successfully message appears, as shown in Figure 7-13. Choose OK.




Figure 7-13 Viewing successful setup message

About accessing online help

iServer supports accessing online help in following ways:

- Online from developer.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 7-14.

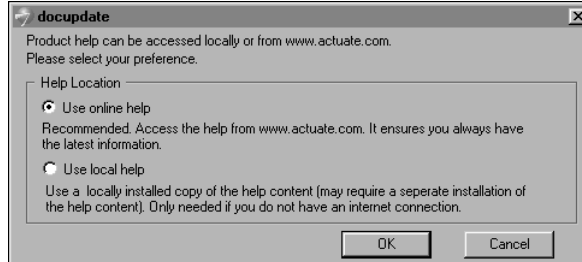


Figure 7-14 Selecting online or local help

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 7-15.

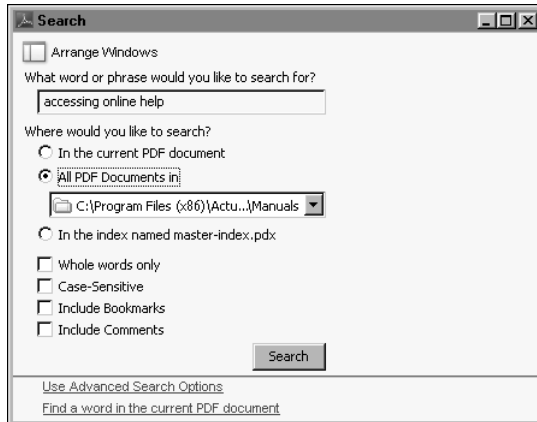


Figure 7-15 Specifying the search

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 7-16.

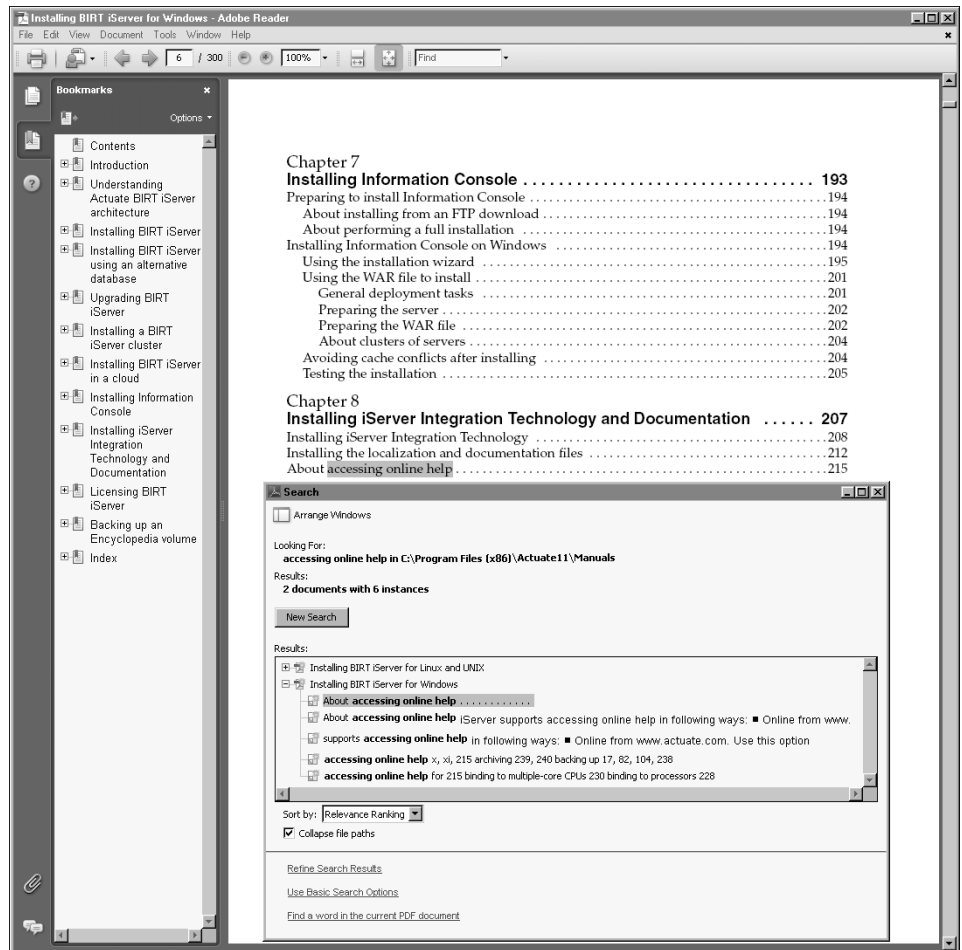


Figure 7-16 Viewing search results

Part Three

Licensing

Licensing BIRT iServer

This chapter discusses the following topics:

- Understanding licensing types
- Understanding licensing options
- Installing Actuate BIRT iServer System license files
- Understanding CPU binding

Understanding licensing types

BIRT iServer System licensing supports running BIRT iServer with sets of features grouped as license options. You enable BIRT iServer System options using one or more of the following types of license models:

- **Named User**

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

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

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

- **CPU Core**

Specifies the maximum number of CPUs that BIRT iServer System can use. Any number of users can access the licensed options on the system provided adequate licensing and capacity exists.

- **Instance**

A BIRT onDemand licensing option that provides a pre-packaged amount of dedicated capacity for a customer application. With instance licensing, the customer does not need to count named users. Multiple instances can be combined to meet capacity needs.

- **Subscription**

An annual payment option that permits the use of the licensed software and includes maintenance. Offered with some of the other licensing models.

A subscription license is not a perpetual license. Once the subscription term expires, the software can no longer be used.

- **Software as a Service (SaaS)**

Some products are offered as a Software as a Service (SaaS) option, providing customers with a convenient solution without incurring the acquisition and management costs of hardware and traditional licenses.

- **Packages**
Some options are offered as packages to customers for convenience and cost-saving benefits. These packages can be used in conjunction with individually selected options.
 - **Work Unit (WU) License**
Specifies iServer features and functionality using an aggregate model. This plan defines each iServer System resource as a work unit.

Similar to CPU Core licensing, but defined at a more granular level. With Work Unit Licensing, the customer can license just the precise amount of capacity needed for application requirements. Any number of users can access the licensed options provided sufficient capacity has been purchased.
- In a CPU Core and Work Unit licensing, Actuate currently uses the Standard Performance Evaluation Corporation (SPEC) standard benchmark for measuring machine capacity based on CPU, memory, disk, and network capacity.

Understanding licensing options

Table 8-1 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 8-1 also describes these prerequisites.

Table 8-1 BIRT iServer System license options

Option	Description	Supported releases
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

(continues)

Table 8-1 BIRT iServer System license options (continued)

Option	Description	Supported releases
BIRT Page Level Security	Controls access to structured content available on the web. This option works for reports created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or security role.	10, 11
BIRT SmartSheet Security	Controls access to structured content available on the web. This option works for reports created using BIRT Spreadsheet Designer and requires BIRT Spreadsheet Option. Access privileges are based on user name or security role.	10, 11
BIRT Spreadsheet	Allows a user to deploy and run a spreadsheet built using BIRT Spreadsheet Designer. This tool enables customers to save reports as richly formatted Excel spreadsheets and manage them in an Encyclopedia volume.	10, 11
BIRT Studio	Allows a user to create a BIRT design and to run it in BIRT iServer. BIRT Studio Option supports access to an information object on BIRT iServer System.	10, 11
e.Analysis	Supports analysis of search results from an Actuate Basic report written in dynamic hypertext markup language (DHTML). This tool is available as an additional purchase with BIRT iServer and requires the e.Report (Actuate Basic Report) Option.	10, 11
e.Report (Actuate Basic Report)	Allows a user to deploy and run an e.report built using Actuate e.Report Designer Professional on an Encyclopedia volume. This option is a requirement for e.Analysis Option and e.Report Page Level Security Option.	10, 11
e.Report Data Connector	Allows a BIRT design to access data that an Actuate Report Document (.roi) file contains. This option works for designs created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or security role.	11
e.Report Page Level Security	Controls access to structured content available on the web. This option works for reports created using Actuate e.Report Designer Professional and requires the e.Report (Actuate Basic Report) Option. Access privileges are based on user name or security role.	10, 11
Multi-Tenant	Allows a BIRT iServer System user to access more than one Encyclopedia volume. This option is available with a Unlimited User CPU License.	10, 11

To determine the license options installed on iServer, log in to Configuration Console, and choose Show License. The license options appear, as shown in Figure 8-1.

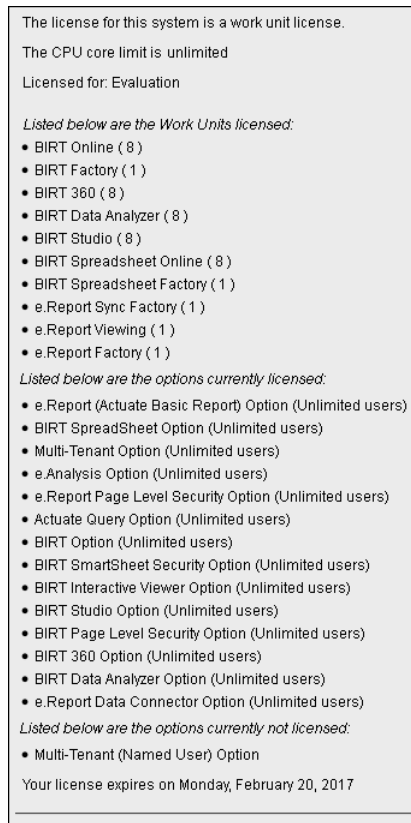


Figure 8-1 iServer License options

Installing Actuate BIRT iServer System license files

OpenText provides a license file to use when installing Actuate BIRT iServer System. New customers receive an e-mail containing a temporary BIRT iServer license file to use for the initial installation after OpenText processes the order. The temporary BIRT iServer System license expires 45 days after installation.

OpenText license enforcement for Release 11 requires a single, shared license for all nodes in a cluster. A design or document run using a temporary license appears with a watermark when viewed.

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

```
Actuate_iServer_key_XXXXXXX.xml
```

XXXXXXX is a unique seven-digit number generated by OpenText Licensing when it creates the license file.

Actuate BIRT iServer System customers perform an initial installation using a temporary license. After installing BIRT iServer System using the temporary license, the login screen displays two messages.

The following message about expiration of the initial license always appears on the login screen regardless of the license status:

Reminder

```
Your BIRT iServer license expires in [the number of days] days,  
on [the specified date]. When the current license expires, the  
iServer will shut down and require a new license to restart.  
Please contact Actuate to purchase a new license.
```

The following message about how to obtain the second license file from OpenText Licensing appears until you install the new license issued by OpenText 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  
OpenText Licensing (https://support.opentext.com), or your  
representative, and request a new license file for the iServer  
nodes that are in violation. Please restart the iServers on the  
nodes after updating the license key file.
```

You have 45 days to apply for and install the license file after you install BIRT iServer System.

After installing Actuate BIRT iServer System, the installation informs a customer requiring a license to obtain the machine ID information on which Actuate BIRT iServer is running and transmit this information to OpenText Licensing. The machine ID is displayed in the reminder message. You can also use the utility, `acmachineid`, to obtain the machine ID. For information on how to use the `acmachineid` utility, see “How to use the `acmachineid` utility,” later in this chapter.

After receiving the machine ID information, OpenText Licensing issues a new Actuate BIRT iServer System license file.

About the license file

This license file specifies the available iServer license options and node-key information for the cluster nodes. This license file must be in a shared location,

specified by the <AC_CONFIG_HOME> attribute of the <Server> element in the apcmdconfig.xml file of each node, and accessible to all nodes in the cluster.

A node key associates an iServer node in a cluster with the machine ID. The node-key licensing mechanism restricts the iServer node installation to that machine.

On startup, each node in the cluster checks the shared license file, verifies the installed options, and determines whether its node key, which is generated at run time, matches the license information. If the node key matches, the node joins the cluster. Otherwise, it shuts down with an error if the node-lock-violation grace period has been exceeded.

A license file remains valid until a specific date. If your license file is about to expire, the system reminds you that the file expires on a certain date when you log in to the Configuration or Management Consoles. Reminders also appear in the system log file. To arrange for a permanent license file, or if you have a problem with an expiring file, please contact OpenText Licensing at support@opentext.com.

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

Collecting machine information for a license

After installing BIRT iServer System using a temporary license file, such as an evaluation license, you must collect information about the machines running Actuate BIRT iServer software and send it to OpenText Licensing. During the installation process, the InstallShield Wizard prompts you to provide the location of the license file. After providing the location of the license file, the InstallShield Wizard issues a prompt similar to the following message:

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:

```
IORRHEHs6S5UCsEtrdVu6jOixmzvFY3BbOqXLIswQGDceJmKYYaEu0j18lQxjM  
sYCxnka3hVkdZFGwkmQMxb+hgKaz4om2vLUcS0ocYTA7Ta6VTMavLFQo7bEjRyr  
olwxAKu0Vr4NA6o8uWCzjGZXX8KrkjViSUoROj70hWOY=
```

Please contact OpenText Licensing (<https://support.opentext.com>), or your representative, and request a node locked iServer system license.

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

The format of the alphanumeric string for the machine ID and location of the license file are different depending on the operating system. On a Windows

system, the unique identifier for the network card is the source of the machine ID. You must have the network card enabled on the BIRT iServer machine to obtain the machine ID.

After installing iServer, you must run the utility, `acmachineid`, from the command line to generate the machine ID information. Copy the machine ID in the command prompt to a file or e-mail message and send it to OpenText Licensing. OpenText Licensing processes your request and sends the new license file for BIRT iServer System.

How to use the `acmachineid` utility

Use the `acmachineid` utility to obtain the machine ID information by performing the following tasks:

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

```
acmachineid
```

The utility provides output in the following format:

```
C:\Program Files (x86)\Actuate11SP5\iServer\bin>acmachineid
STATUS:          OK
GEN_VERSION:     11 Service Pack 5
GEN_BUILD:       XxxBuild NumberXxx
SERVERNAME:      <hostname>
MACHINEID:
                I0RREHs0Jk6tu0o8AbCrVL61x7kDpL1QKwS2t1W7qM67GbO8
                VjcFs6pcuAgbtDaZauSbFFa2mRejwVJc7ZjKfMEV11suXglM
                KmZLiwtLykwJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9f
                JUMYzZORKk23jrxaswUDsgKsvlc1A6q8UbmrrAYHD8Ggtptui
                AmxWt4xjEM6rq1msNEW/4ViMC0KDBkSn
```

Figure 8-2 shows the output as it appears in the command prompt.




Figure 8-2 Executing `acmachineid`

Send OpenText Licensing the output of the `acmachineid` utility.

How to obtain a license file

To register as a new user and obtain a new license file for a product, go to OpenText My Support at the following location:

<https://support.opentext.com> (Log in required)

Choose License Keys. Then, submit a ticket.

If you are a maintenance customer, you should have login information for OpenText My Support. If you do not have access, please contact your OpenText sales representative.

If you are not a direct OpenText customer, contact the partner or distributor who supplies the product for the license file. If you have a problem obtaining a license file from this source, please contact your OpenText sales representative.

Updating the Actuate BIRT iServer System license file

After performing an installation of Actuate BIRT iServer System and transmitting the required machine ID information to obtain a license, OpenText sends an e-mail containing an attached .txt (TXT) file. Replace the .txt extension with a .zip (ZIP) extension and open the file. This ZIP file contains the following files:

- `readme.txt`
Instructions for installing Actuate BIRT iServer System using a license file and for obtaining a license file
- `Actuate_iServer_key_XXXXXXX.xml`
Actuate BIRT iServer System license

An Actuate license file is an XML file. OpenText Licensing sends this XML file inside of a TXT file because transmitting a file with an .xml extension can cause problems in an e-mail system.

How to install the license file

To install the license file, perform the following steps:

- 1 Extract the contents of the ZIP file to a location on your local file system.
- 2 Log in to Configuration Console. For example, type `http://localhost:8900/acadmin/config/` in Address of a browser, and use the system configuration password that you specified during installation.
- 3 Choose Update License. If you do not see Update License, choose Simple view in the upper right corner.
- 4 On Actuate iServer update license, choose Browse to navigate to the location where you extracted the contents of the ZIP file. Select the Actuate BIRT iServer System license file and choose OK to apply the license.

If iServer requires a system restart to update the license file, the following message appears:

The license file cannot be applied without a server restart.
Please copy the license file to the iServer license file location and restart the iServer system.

If this message appears, perform the following tasks:

- 1 Stop iServer system by choosing Stop system.
 - 2 Copy the new license file to the shared location specified by the <AC_CONFIG_HOME> attribute of the <Server> element in the acpmdconfig.xml file. The <AC_CONFIG_HOME> attribute in the acpmdconfig.xml files for all nodes in a cluster point to this shared location.
 - 3 Delete the old acserverlicense.xml file.
 - 4 Rename the new license file to acserverlicense.xml.
 - 5 Start iServer System.
- 5 Restart any node where the node-key configuration changed.

If you change the machine on which you installed Actuate BIRT iServer, you must re-apply to OpenText Licensing for a new license file. If you replace the network card on some machines, such as a Windows system, you may have to obtain a new license file since the unique identifier for the network card may be the source of the machine ID. If you have a license file installed and a reminder message appears when logging into Actuate Management Console, contact OpenText Licensing and provide the current Actuate iServer System license file with the output from the machine ID utility.

Actuate_iServer_key_XXXXX.xml will contain the node key information for the stand-alone machine or all machines in a cluster. There is no separate node license file for each machine.

Listing 8-1 shows the node key information the license contains, obtained from the acmachineid output you submitted to OpenText Licensing.

Listing 8-1 Viewing license node key information

```
<NodeKeys>
  <NodeKey
    MachineId="E0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67GbO8
    VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwtLykDa/
    wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzzZORkK23jrxasWUDig
    Ksvlc1A6q8UbmrrAYHD8GgtptuiAmxWt4xjEM6rqlmsNEW/4Vjm40Kx1kSv"
    ServerName="W7CLSTRNODE1"/>
  <NodeKey
    MachineId="I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67GbO8
    VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEV11suXglMKmZLiwtLykDa/
```



```
wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORKk23jrxaswUDsg  
Ksvlc1A6q8UbmrrAYHD8GgtpuiAmxWt4xjEM6rqlmsNEW/4ViMC0KDBkSn"  
ServerName="W7CLSTRNODE2" />  
</NodeKeys>
```

About modifying a license

If you decide later to license additional iServer options, the existing license file becomes invalid. You must install a new license file.

Contact OpenText Licensing for the new license file.

Understanding CPU binding

BIRT iServer System supports CPU binding on a machine with an appropriate CPU-based license. CPU binding restricts a process or processes to run on a subset of CPUs. If you bind the BIRT iServer System to a subset of CPUs, only those CPUs count toward the total number of licensed CPUs. The CPU limit in the license file applies to all CPUs for all machines in the cluster. Depending on the operating system and specific system command, you can restrict other processes from running on the processor to which you bind a process.

You can bind BIRT iServer processes to a specific set of processors on a machine that runs a Windows, Sun Solaris, or HP-UX 11i operating system. The default configuration does not bind BIRT iServer to a set of processors. In the default configuration, all processors on a BIRT iServer machine count toward the maximum number of licensed CPUs. For more information about performing CPU binding on a Sun Solaris or HP-UX 11i machine, see *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 8-3.

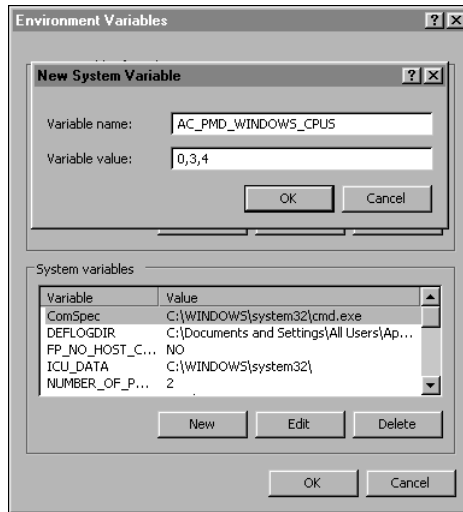


Figure 8-3 Creating the AC_PMD_WINDOWS_CPUS system variable
Choose OK. AC_PMD_WINDOWS_CPUS appears in System variables, as shown in Figure 8-4.

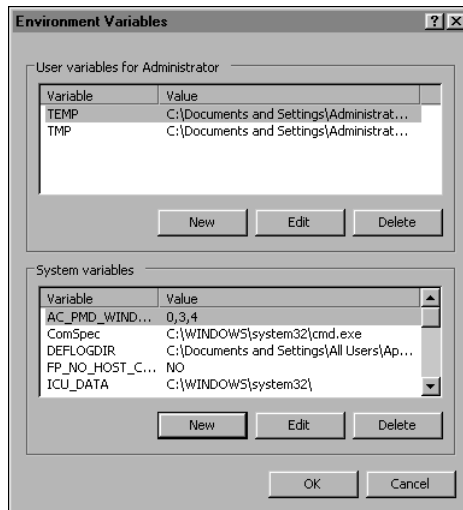


Figure 8-4 Viewing the AC_PMD_WINDOWS_CPUS system variable
On Environment Variables, choose OK. Then, on System Properties, choose OK.

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 OpenText 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 5 environment, the default `AC_SERVER_HOME` path is:

C:\Program Files (x86)\Actuate11SP5\iServer

The default AC_DATA_HOME path is:

C:\Actuate\iServer\data\

The default Encyclopedia volume path is:

C:\Actuate\iServer\data\encyc

The default acserverconfig.xml file path is:

C:\Actuate\iServer\data\config\11SP5

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 OpenText Support to troubleshoot problems, it is best to provide a snapshot of the Encyclopedia volume configuration, including the following items and information:

- A database backup of the Encyclopedia volume schema containing the metadata
- The name of the Encyclopedia volume schema and user that iServer uses to connect to the RDBMS
- A copy of the acserverconfig.xml file containing iServer configuration information
- A copy of the iServer logs

Using RDBMS and file system backup utilities

The administrator must perform the Encyclopedia volume metadata backup using the tools provided or supported by the RDBMS. Copying the physical files of a database at the operating system level while an RDBMS is running does not create a valid backup.

Most RDBMS backup tools can be scripted and run while iServer is using the database. PostgreSQL, Oracle, Microsoft SQL Server, and DB2 also provide graphical administration tools in addition to command-line tools. For more information on using these RDBMS tools to back up and restore an Encyclopedia volume, see the related sections, later in this chapter.

How to perform an Encyclopedia volume backup

To back up an Encyclopedia volume, perform the following tasks:

- 1 Make sure that the autoarchive file purging process is not running.
- 2 Make an online backup of the volume schema using the tools provided by the RDBMS.
- 3 Back up the volume data files using the tools available in the operating system environment.

Avoiding conflict with the file purging process

A metadata backup is consistent with a data backup only if the file purging process that runs during an autoarchive operation does not occur between the time you back up the metadata and the time you back up the data. In Volumes—Properties—Advanced—Archiving And Purging, the administrator can specify when the file purging process runs.

How to configure Archiving And Purging

To configure the autoarchive file purging process, perform the following tasks:

- 1 From the advanced view of Configuration Console, choose Volumes.
- 2 On Volumes, point to the icon next to a volume name and choose Properties. In Volumes—Properties, choose Advanced. In Advanced, choose Archiving And Purging.
- 3 In Archiving And Purging, configure the following time-related file purging properties to times that do not conflict with the time when the backup operation runs, as shown in Figure 9-1:

Actuate > Volume properties - Windows Internet Explorer

Volumes > corp : Properties > Archiving And Purging

Archiving And Purging

Expiration time of failed jobs: 43200 Minutes !

Expiration time of successful jobs: 43200 Minutes !

Default expiration time of success notices: 0 Minutes !

Default expiration time of failure notices: 0 Minutes !

Purge deleted files time: 2:15 24-hour time !

Expiration time of deleted files: 120 Minutes !

These fields require volume restart to take effect

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

OK Cancel

Figure 9-1 Configuring file purging properties

- Purge deleted files time
Specifies the time when the file purging process runs to permanently delete expired files
- Expiration time of deleted files
Specifies the length of time that must elapse before the file purging process permanently deletes an expired file

Choose OK.

For information on other aspects of archiving, see *Configuring BIRT iServer*.

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

PostgreSQL provides the pgAdmin graphical administration tool or the `pg_dump` and `pg_restore` command-line utilities to back up and restore a database. These PostgreSQL utilities run on the client not the server.

To back up an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the pgAdmin graphical administration tool or the `pg_dump` PostgreSQL command-line utility
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

Note that a backup of a PostgreSQL database is not portable across all operating systems.

To restore an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the pgAdmin graphical administration tool or the `pg_restore` PostgreSQL command-line utility
- Restores Encyclopedia volume data and configuration files using operating system copy commands

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

Backing up an Encyclopedia volume using pg_dump

To back up an Encyclopedia volume using the pg_dump utility, perform the following tasks:

- Create a folder to contain the metadata and volume data backup files.
- Back up Encyclopedia volume metadata using the pg_dump utility.
- Back up the acserverconfig.xml file and volume data folders to the backup folder.

Create a folder to contain the metadata, acserverconfig.xml file, 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 at a storage location that is physically separate from the Encyclopedia volume data location.

In a Windows environment, create a folder to contain the metadata and volume data backup files by performing the following tasks.

How to create a new backup folder

- 1 Open Windows Explorer.
- 2 In Windows Explorer, choose File→New→Folder repeatedly to create a set of nested folders in the following location, as shown in Figure 9-2:

C:\Actuate\iServer\encyc_backup




Figure 9-2 Creating a new folder named encyc_backup

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 volumes names that start with the prefix `ac_`. If `-n` is not specified, `pg_dump` exports all non-system schemas.
- **f**
Specifies the output file, such as `ac_corp_schema.dmp`.
- **h**
Specifies the host name of the machine where the PostgreSQL server is running, such as `dbhost`.
- **p**
Specifies the port where the server listens for connection requests.
- **U**
Specifies the user name for the connection to the PostgreSQL server, such as `postgres`.
- **dbname**
Replace this string in the example with the database name, such as `actuate_db`.

Re-run the command to back up each Encyclopedia volume schema to a separate archive. To run multiple volume schema backups using a script, set up auto-login using a `.pgpass` file. The file should contain connection information in the following format:

```
hostname:port:database:username:password
```

More information about setting up a scripted backup using a `.pgpass` file is available at:

```
http://www.postgresql.org/docs/8.4/static/libpq-pgpass.html
```

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

How to run `pg_dump` from a command prompt

- 1 Open a command prompt.
- 2 Navigate to the following location:

```
C:\Program Files (x86)\Actuate11SP5\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 If prompted to enter a password, type the postgres superuser password. The administrator specified this password in PostgreSQL Database Information during the iServer installation procedure.

pg_dump executes, writing status messages to the command prompt.

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

How to back up the volume data folders

- 1 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 5, this file is located in the following subfolder:

C:\Actuate\iServer\data\config\11SP5

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

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.

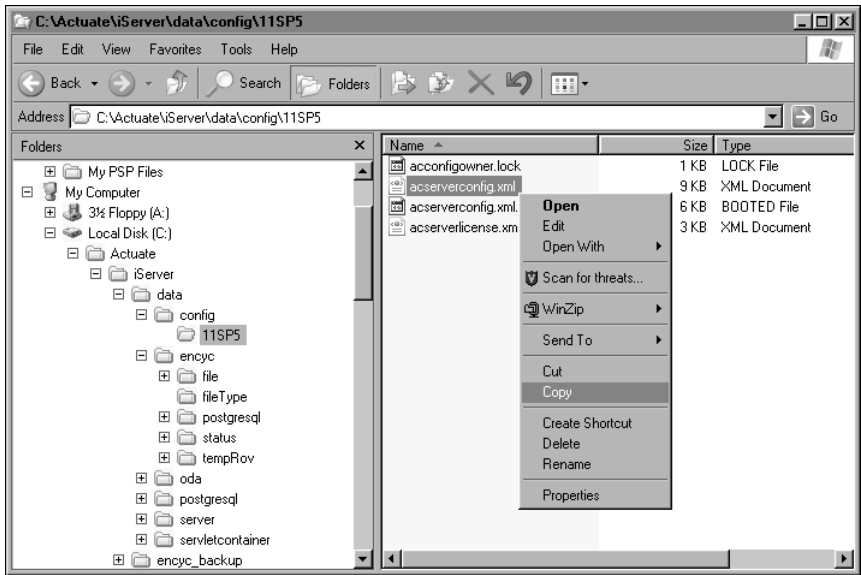


Figure 9-3 Copying acserverconfig.xml

Figure 9-4 shows a copy operation that contains the file, fileType, status, and tempRov folders.

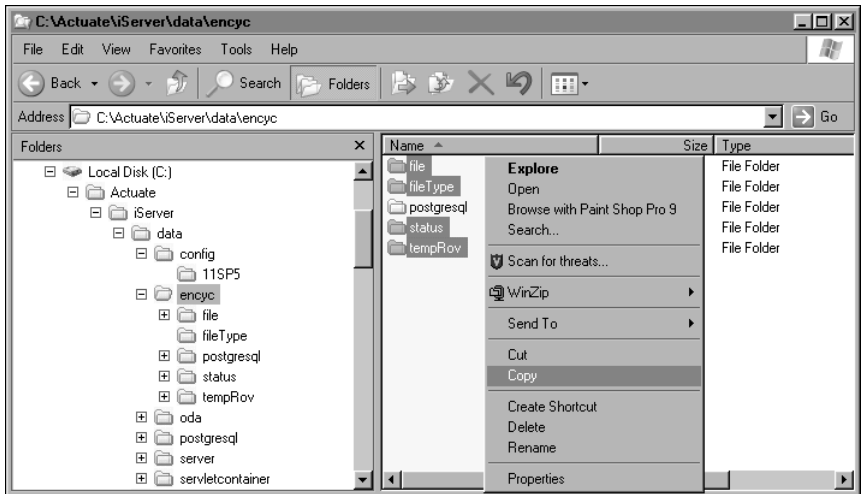


Figure 9-4 Copying the volume data folders

The contents of the backup folder appear as shown in Figure 9-5.

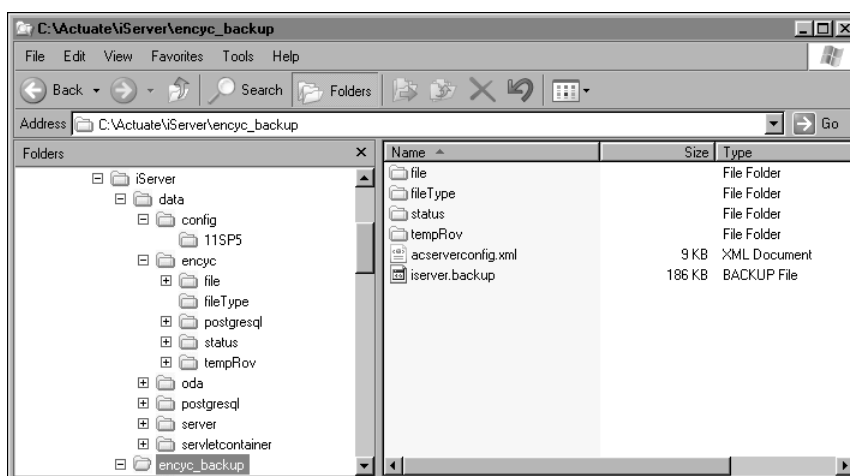


Figure 9-5 Viewing the backed-up files

Restoring an Encyclopedia volume using pg_restore

To restore a backed-up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia volume offline.
- Delete the acserverconfig.xml 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.

Alternatively, the administrator can restore an individual volume by selectively backing up and restoring only the related volume schema and data.

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 9-6.

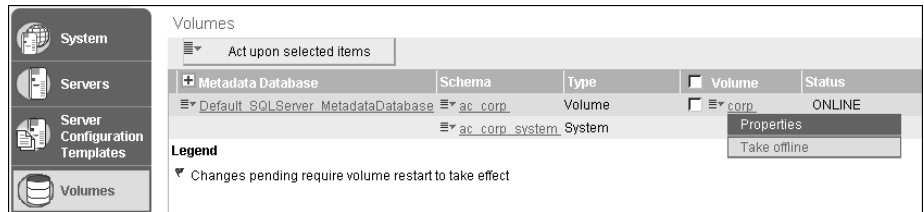


Figure 9-6 Taking the volume offline

How to restore the backed-up volume data folders

- 1 In Windows Explorer, navigate to AC_DATA_HOME\config\11SP5.
- 2 Select acserverconfig.xml, right-click, and choose Delete, as shown in Figure 9-7. Confirm the deletion.

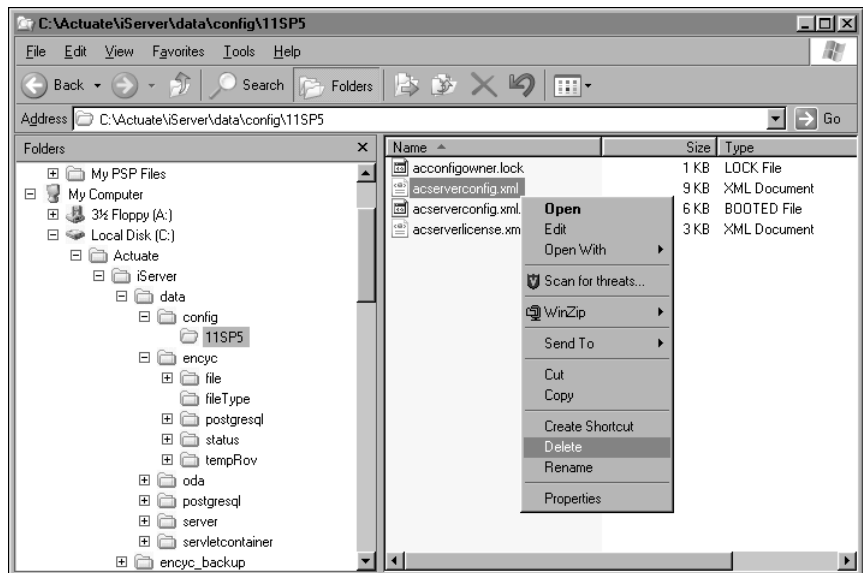


Figure 9-7 Deleting acserverconfig.xml

- 3 In AC_DATA_HOME, open the encyc folder. In AC_DATA_HOME\encyc, select the file, fileType, status, and tempRov folders, right-click, then choose Delete, as shown in Figure 9-8. 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.

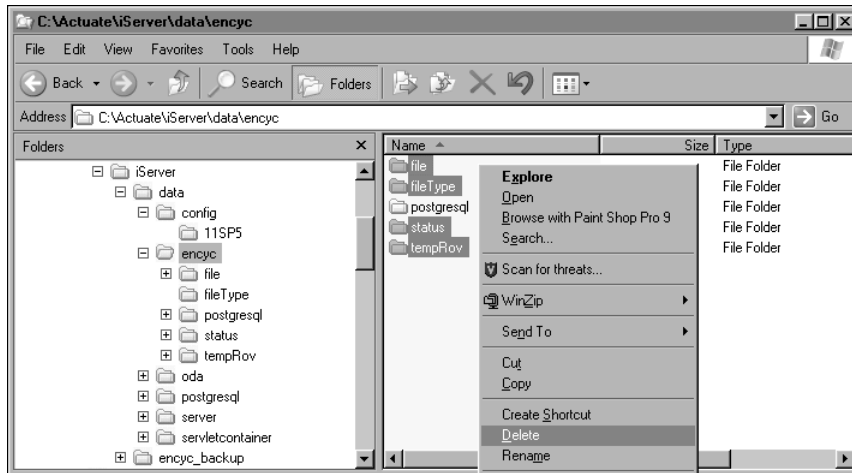


Figure 9-8 Deleting the file, fileType, and status folders from AC_DATA_HOME\encyc

- 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\11SP5, as shown in Figure 9-9.

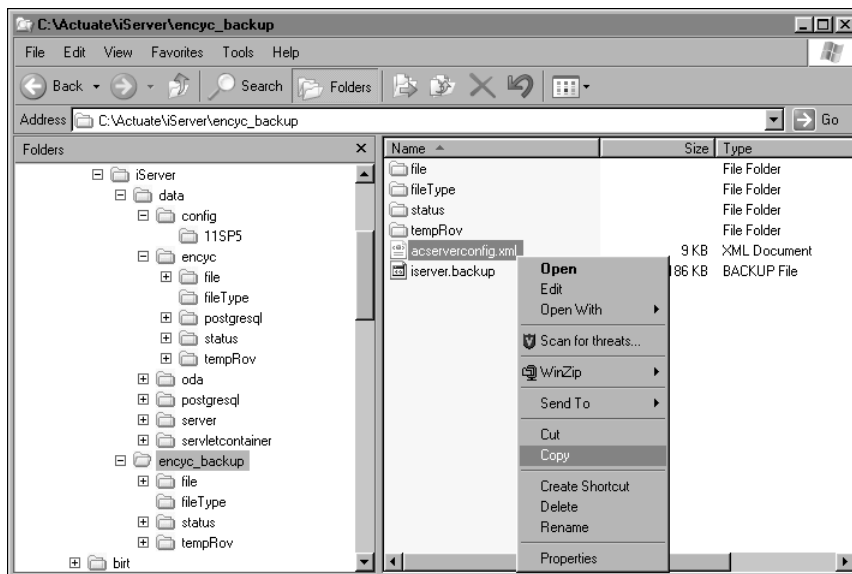


Figure 9-9 Choosing to copy acserverconfig.xml

AC_DATA_HOME\config\11SP5 appears as shown in Figure 9-10.

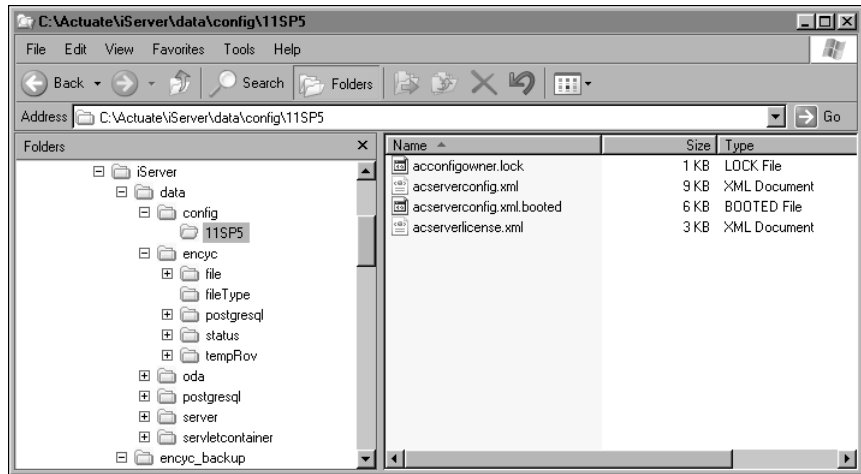


Figure 9-10 Viewing AC_DATA_HOME\config\11SP5 after restoring acserverconfig.xml

- 5 In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempProv folders, right-click, choose Copy, and copy these folders to AC_DATA_HOME\encyc, as shown in Figure 9-11.

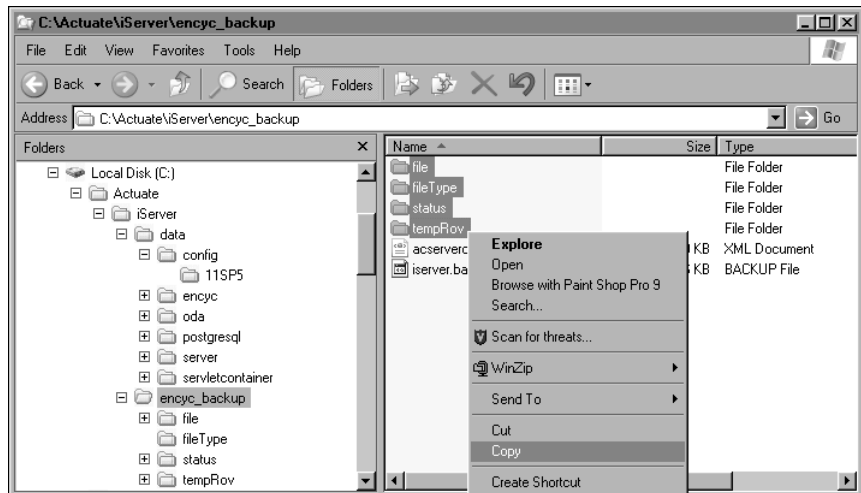


Figure 9-11 Choosing to copy the volume data folders

AC_DATA_HOME\encyc appears as shown in Figure 9-12.

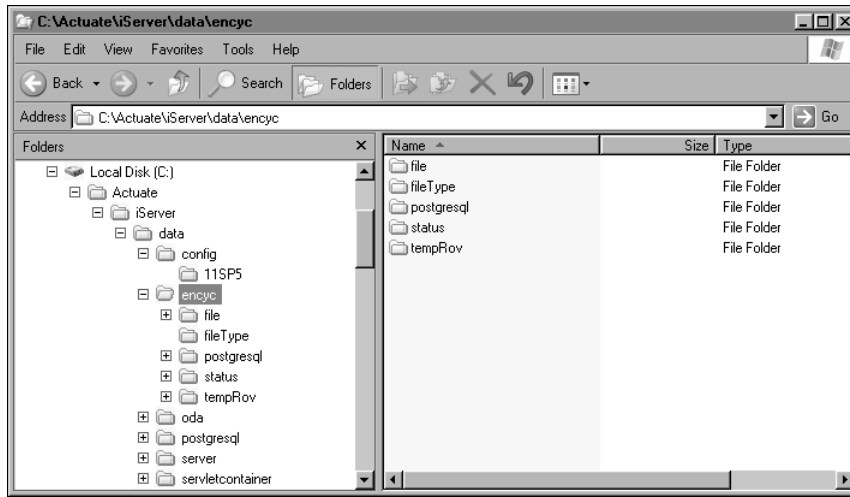


Figure 9-12 AC_DATA_HOME\encyc after copying backed-up folders

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 (x86)\Actuate\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 If prompted, type the postgres superuser password. The administrator specified this password in PostgreSQL Database Information during the iServer installation procedure. Press Enter.

`pg_restore` executes, writing status messages to the command prompt.

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

- 1 Log in to Configuration Console. On the simple view, choose Advanced view. Choose Volumes.
- 2 On Volumes, take the volume online, as shown in Figure 9-13.




Figure 9-13 Taking the volume online

More information about backing up and restoring an Encyclopedia volume schema using the PostgreSQL `pg_dump` and `pg_restore` utilities is available at the following location:

<http://www.postgresql.org/docs/8.4/static/backup.html>

Backing up and restoring an Encyclopedia volume that uses an Oracle database

Oracle provides the Data Pump Export (`expdp`) and Import (`impdp`) command-line utilities to back up and restore a database. Oracle Data Pump utilities enable high-speed uploading and downloading of data and metadata in an Oracle RDBMS. Unlike the PostgreSQL backup and restore utilities, the Oracle Data Pump utilities run on the server, not the client, for greater efficiency.

To back up an Encyclopedia volume stored in an Oracle RDBMS, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the Oracle Data Pump Export (`expdp`) command-line utility
- Backs up Encyclopedia volume data and configuration file using operating system copy commands

To restore an Encyclopedia volume in an Oracle RDBMS, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the Oracle Data Pump Import (`impdp`) command-line utility
- Restores Encyclopedia volume data and configuration file using operating system copy commands

The following sections provide more information on how to perform these backup and restore operations.

Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)

Oracle Data Pump utilities require the database administrator to map a database dump directory to a physical directory in the file system. The Oracle RDBMS writes to and reads from this directory when performing export and import operations.

Create a directory to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

To back up an Encyclopedia volume metadata using expdp, perform the following tasks:

- Create a directory to contain the metadata and volume data backup files using sqlplus.
- Back up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility.
- Back up the acserverconfig.xml file and volume data folders to the backup folder.

How to create a backup of the Encyclopedia volume metadata

- 1 In Windows, open a command prompt.
- 2 Use sqlplus, an Oracle SQL editing tool, to log in as the system administrator, as shown in the following example:

```
sqlplus system/password@db_host/dbname.actuate.com
```

The example specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name.

- 3 In sqlplus, create a backup directory to contain the Encyclopedia volume metadata, as shown in the following example:

```
SQL> CREATE DIRECTORY encyc_backup  
      AS 'C:\Actuate\iServer\encyc_backup';  
SQL> exit
```

Exit sqlplus.

- 4 Back up the Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility, as shown in the following example:

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

The expdp example uses the following arguments:

- **system/password@db_host/dbname.actuate.com**
Specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name
- **SCHEMAS**
Species the Encyclopedia volume schema to export, such as corp
- **DIRECTORY**
Specifies the directory for writing the database dump, such as encyc_backup, created by the system administrator in sqlplus
- **DUMPFILE**
Specifies the name of the output file, such as ac_corp_schema.dmp

After backing up the Encyclopedia volume metadata, back up the 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:\Actuate\iServer\data
```

- 2 In AC_DATA_HOME, navigate to the config folder that contains acserverconfig.xml file. In Actuate 11 Service Pack 5, this file is located in the following subfolder:

```
C:\Actuate\iServer\data\config\11SP5
```

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 9-14.




Figure 9-14 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 5, this file is located in the following subfolder:
`C:\Actuate\iServer\data\config\11SP5`
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\11SP5.
- 2 In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempROV directories, right-click, choose Copy, and copy these directories to AC_DATA_HOME\encyc.

Restore the Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility by performing the following tasks.

How to restore a backup of the Encyclopedia volume metadata

- 1 In Windows, open a command prompt.
- 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 9-15.




Figure 9-15 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 Server RDBMS 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 9-16 shows the new folder.




Figure 9-16 Creating a new folder named encyc_backup

Back up Encyclopedia volume metadata using the SQL Server graphical administration tool, SQL Server Management Studio, by performing the following tasks.

How to create a backup of the Encyclopedia volume metadata using SQL Server Management Studio

- 1 Choose Start→Programs→Microsoft SQL Server 2008→SQL Server Management Studio.
- 2 On Connect to Server, in Server name, type your machine name, such as urup, as shown in Figure 9-17. Choose Connect.

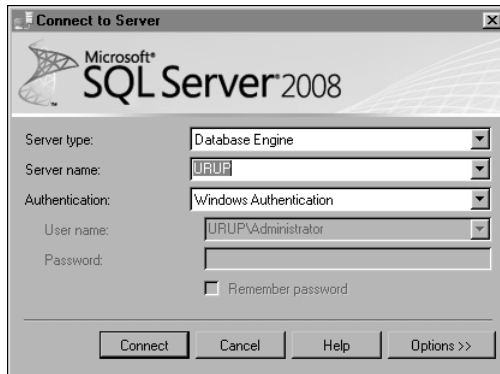


Figure 9-17 Connecting to a machine named urup

- 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 9-18.

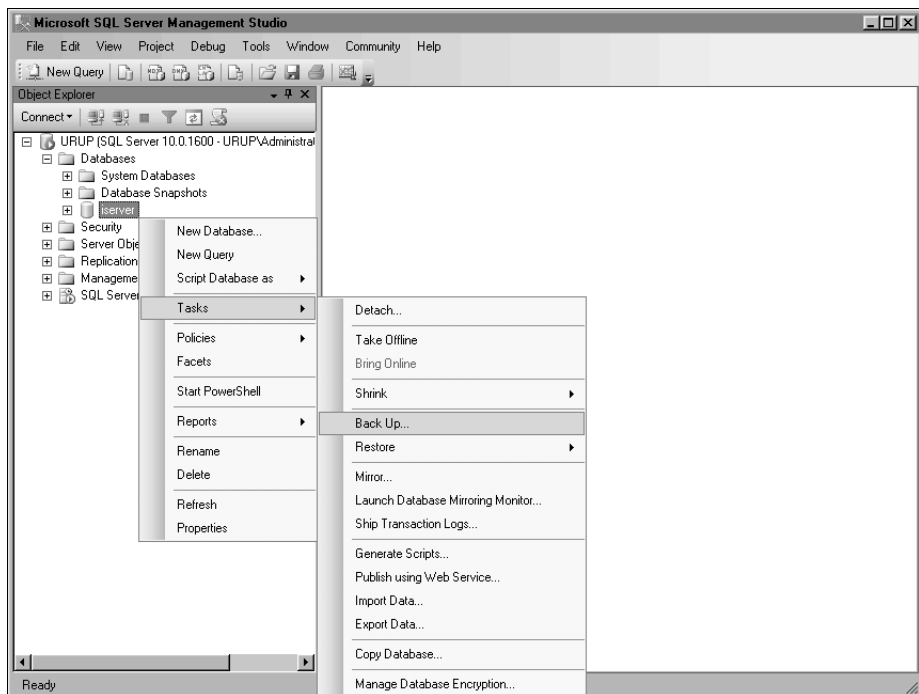


Figure 9-18 Choosing to back up the database named iserver

- 4 On Back Up Database—General, perform the following tasks:
 - 1 In the Source section, in Database, select iserver.
 - 2 In Backup type, select Full.
 - 3 Leave Copy Only Backup unselected.
 - 4 For Backup component, select Database.
 - 5 In the Backup set section, in Name, accept the default value as the name for the backup set. Alternatively, type a different name.
 - 6 In Description, type an optional description for the backup set.
 - 7 For Backup set will expire, in After, accept the default value of 0 to specify that the backup set will not expire and cannot be overwritten. Alternatively, specify a number of days after which the backup set expires and can be overwritten.

Optionally, select On to specify an expiration date for the backup set.

- 8 In the Destination section, perform the following tasks:
 - 1 For Back up to, accept the default value of Disk.
 - 2 Choose Remove to delete the default backup destination.
 - 3 Choose Add.
 - 4 On Select Backup Destination, as shown in Figure 9-19, in File name, type:

C:\Actuate\iServer\encyc_backup\iserver.bak




Figure 9-19 Typing the backup destination

On Select Backup Destination, choose OK. Back Up Database—General appears as shown in Figure 9-20. Choose OK.

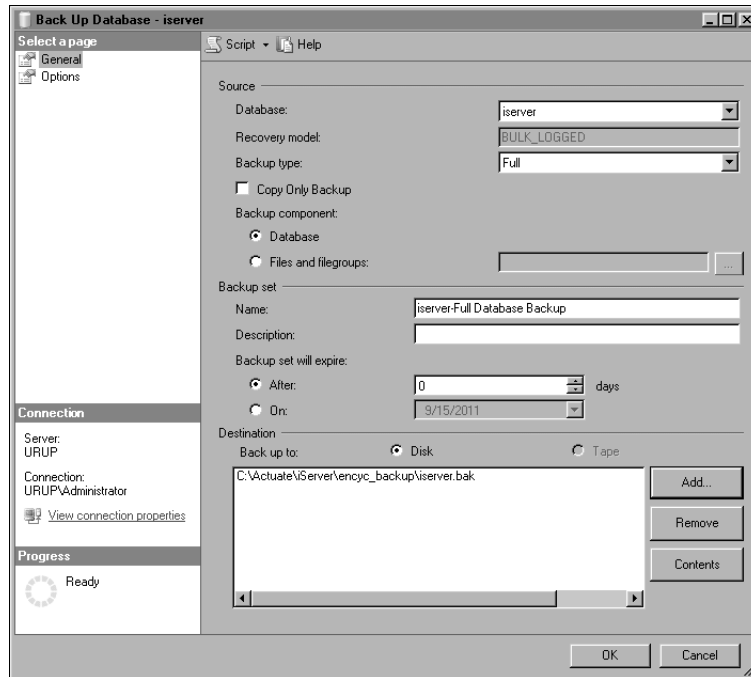


Figure 9-20 Viewing Back Up Database with new backup location

- 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 9-21.

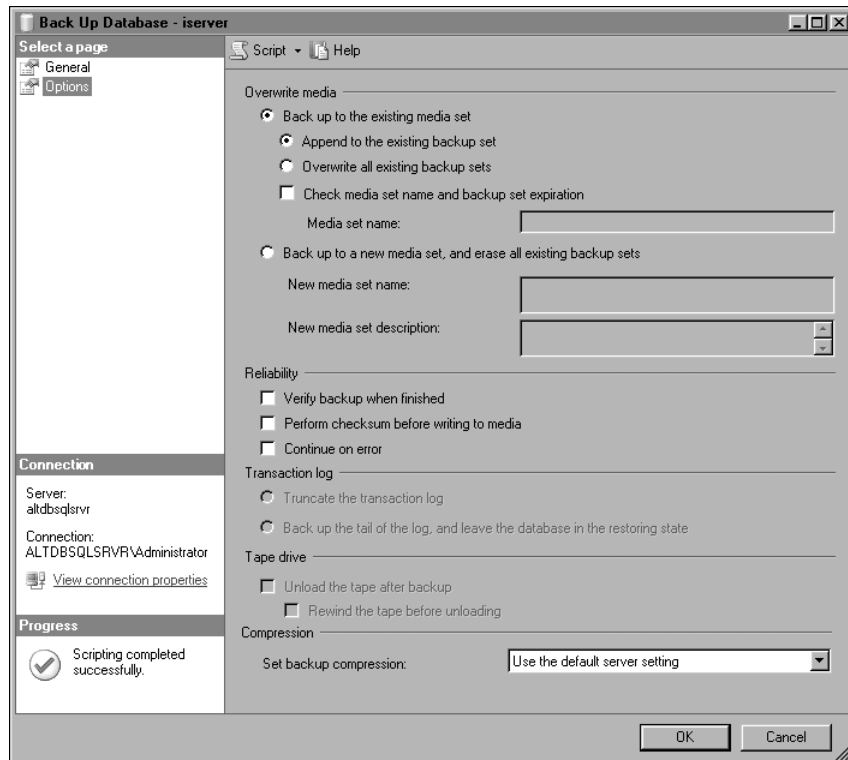


Figure 9-21 Accepting the default values on Back Up Database—Options

- 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 9-22.




Figure 9-22 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 9-23.




Figure 9-23 Acknowledging the successful backup operation

Alternatively, you can also back up the Encyclopedia volume metadata using the `sqlcmd` utility. The following example duplicates the operations performed in the previous section using the graphical SQL Server administration tool, SQL Server Management Studio. You do not need to do both activities.

How to create a backup of the Encyclopedia volume metadata using the `sqlcmd` utility from a command prompt

- 1 Open a command prompt.
- 2 To connect with the default instance of SQL Server running on your machine, type the following command and press Enter:

`sqlcmd`

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

```
BACKUP DATABASE [iserver] TO DISK = N'C:\Program Files
\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\Backup
\iserver.bak' WITH NOFORMAT, NOINIT, NAME = N'iserver-Full
Database Backup', SKIP, NOREWIND, NOUNLOAD, STATS = 10
```

At the `sqlcmd` prompt, type the following command:

`GO`

The command prompt appears as shown in Figure 9-24.

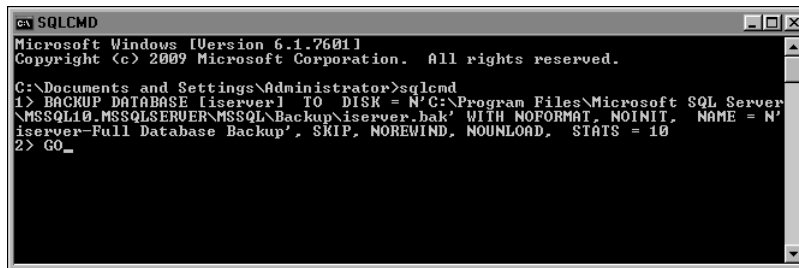
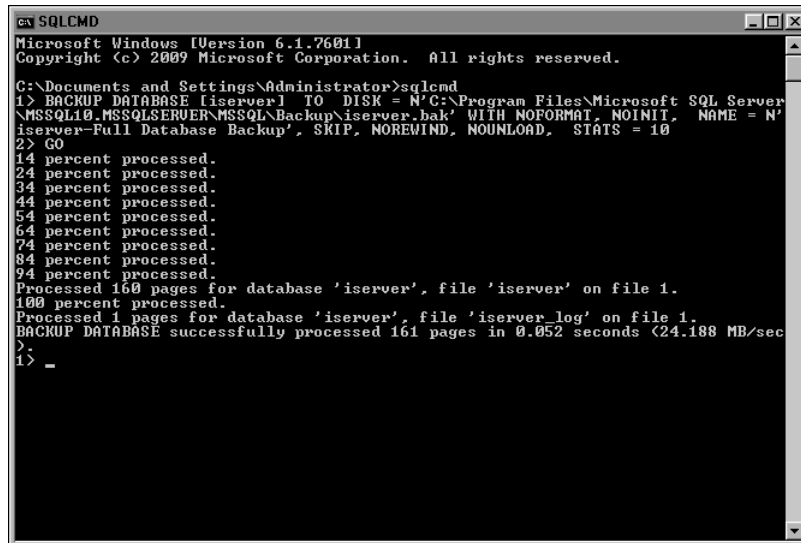


Figure 9-24 Typing the command to back up the metadata

Press Enter to execute the command to back up the database. The output appears as shown in Figure 9-25.



```
SQLCMD
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Documents and Settings\Administrator>sqlcmd
1> 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
2> GO
14 percent processed.
24 percent processed.
34 percent processed.
44 percent processed.
54 percent processed.
64 percent processed.
74 percent processed.
84 percent processed.
94 percent processed.
Processed 160 pages for database 'iserver', file 'iserver' on file 1.
100 percent processed.
Processed 1 pages for database 'iserver', file 'iserver_log' on file 1.
BACKUP DATABASE successfully processed 161 pages in 0.052 seconds (24.188 MB/sec)
>
1> _
```

Figure 9-25 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:\Actuate\iServer\data

- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 5, this file is located in the following subfolder:

C:\Actuate\iServer\data\config\11SP5

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 9-26.

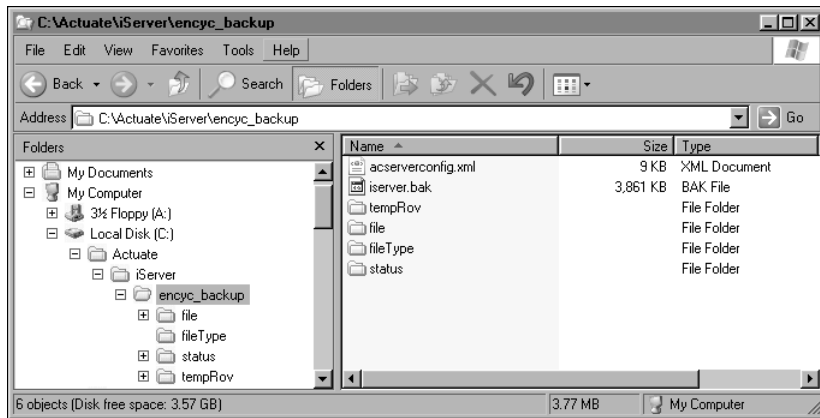


Figure 9-26 Viewing the backed-up files

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 9-27.




Figure 9-27 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 5, this file is located in the following subfolder:
`C:\Actuate\iServer\data\config\11SP5`
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, 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\11SP5. In C:\Actuate\iServer\encyc_backup, select the file, filetype, status, and tempRov folders, right-click, choose Copy, and copy these folders to AC_DATA_HOME\encyc.

Before restoring a full database backup, you must back up the tail of the transaction log. A backup taken of the tail of the transaction log just before a restore operation is called a tail-log backup.

How to back up the tail of the transaction log

- 1 On Microsoft SQL Server Management Studio, in Object Explorer, expand Databases. Right-click the database named iserver and choose Tasks➤Back Up.
- 2 On Back Up Database—General, perform the following tasks:
 - 1 In the Source section, in Database, select iserver.
 - 2 In Backup type, select Transaction Log.
 - 3 Leave Copy Only Backup unselected.
 - 4 In the Backup set section, in Name, accept the default value as the name for the backup set. Alternatively, type a different name.
 - 5 In Description, type an optional description for the backup set.
 - 6 For Backup set will expire, in After, accept the default value of 0 to specify that the backup set will not expire and cannot be overwritten. Alternatively, specify a number of days after which the backup set expires and can be overwritten.

Optionally, select On to specify an expiration date for the backup set.
 - 7 In the Destination section, for Back up to, accept the default value of Disk.
 - 8 Accept the default backup destination.

Back Up Database—General appears as shown in Figure 9-28.

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.

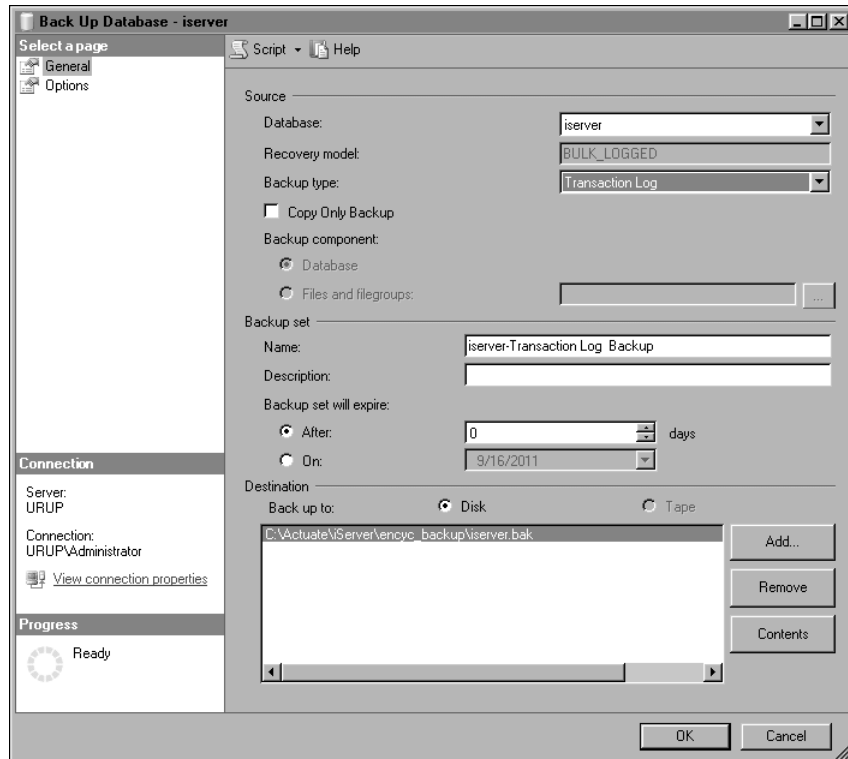


Figure 9-28 Choosing to back up the transaction log

- 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 9-29.




Figure 9-29 Choosing to display Transact-SQL statements

Back Up Database—Options appears as shown in Figure 9-30.

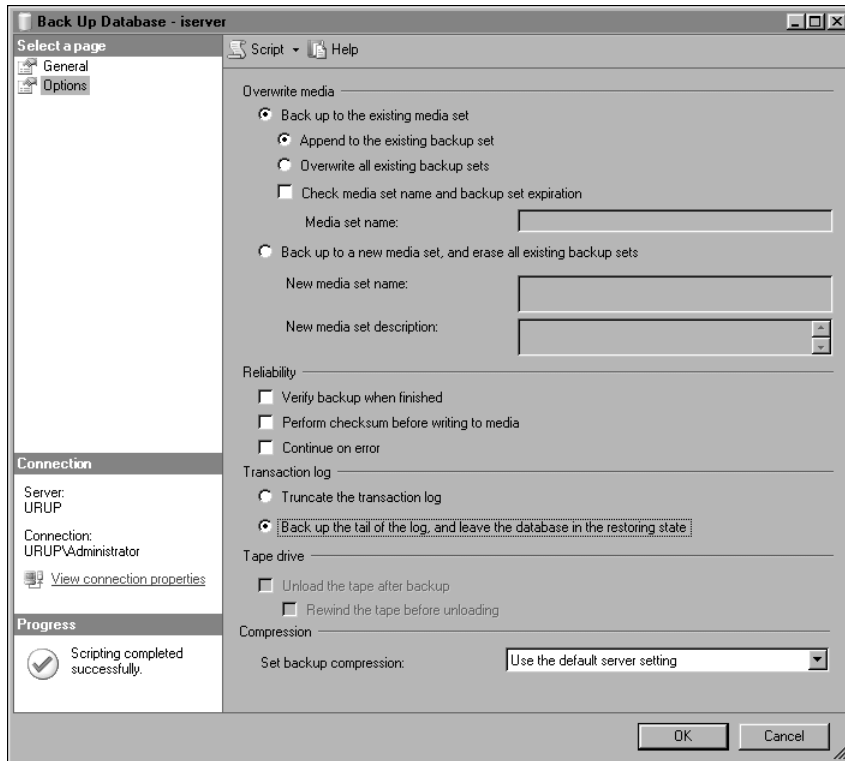


Figure 9-30 Configuring Back Up Database—Options

Choose OK.

- 4 In the dialog box stating that the backup operation completed successfully, choose OK, as shown in Figure 9-31.




Figure 9-31 Acknowledging the successful backup operation

Alternatively, you can also back up the tail of the transaction log using the `sqlcmd` utility. The following example duplicates the operations performed in the previous section using the graphical SQL Server administration tool, SQL Server Management Studio. You do not need to do both activities.

How to back up of the tail of the transaction log using `sqlcmd` from a command prompt

- 1 Open a command prompt.

- 2 To connect with the default instance of SQL Server running on your machine, type the following command and press Enter:

```
sqlcmd
```

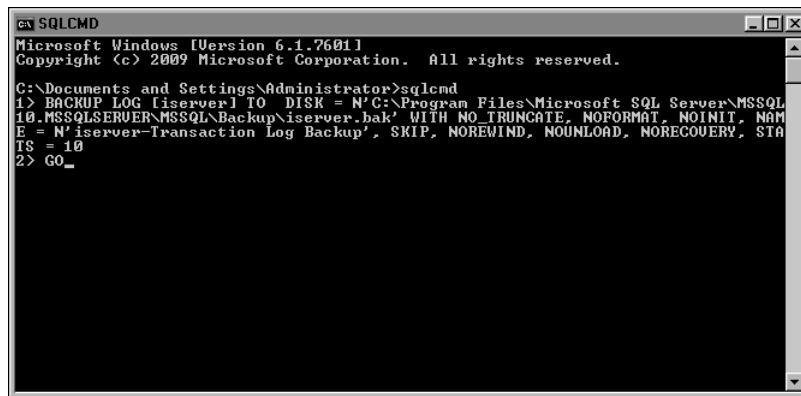
- 3 At the sqlcmd prompt, type the following command and press Enter:

```
BACKUP LOG [iserver] TO DISK = N'C:\Program Files\Microsoft SQL  
Server\MSSQL10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH  
NO_TRUNCATE, NOFORMAT, NOINIT, NAME = N'iserver-Transaction  
Log Backup', SKIP, NOREWIND, NOUNLOAD, NORECOVERY, STA  
TS = 10
```

At the sqlcmd prompt, type the following command:

```
GO
```

The command prompt appears as shown in Figure 9-32.

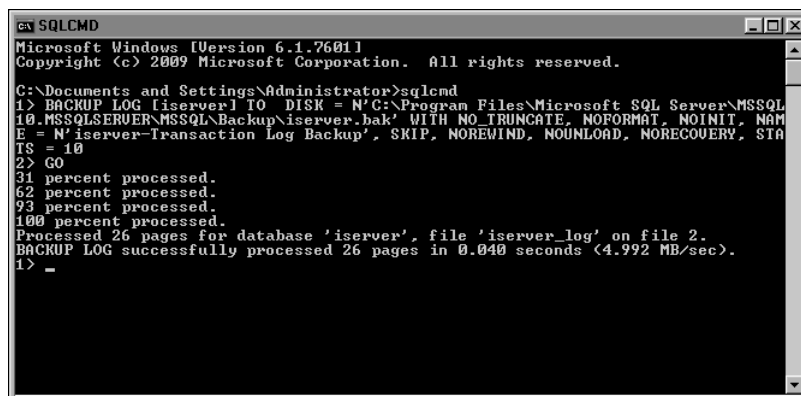


```
SQLCMD
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

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 9-32 Typing the command to back up the tail of the log

Press Enter to execute the command to back up the tail of the transaction log. The output appears as shown in Figure 9-33.



```
SQLCMD
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Documents and Settings\Administrator>sqlcmd
1> BACKUP LOG [iserver] TO DISK = N'C:\Program Files\Microsoft SQL Server\MSSQL
10.MSSQLSERVER\MSSQL\Backup\iserver.bak' WITH NO_TRUNCATE, NOFORMAT, NOINIT, NAM
E = N'iserver-Transaction Log Backup', SKIP, NOREWIND, NOUNLOAD, NORECOVERY, STA
TS = 10
2> GO
31 percent processed.
62 percent processed.
93 percent processed.
100 percent processed.
Processed 26 pages for database 'iserver', file 'iserver_log' on file 2.
BACKUP LOG successfully processed 26 pages in 0.040 seconds (4.992 MB/sec).
1> _
```

Figure 9-33 Executing the command to back up the tail of the log

- 4 To end the sqlcmd session, type the following command at the sqlcmd prompt:

EXIT

When restoring the Encyclopedia volume metadata, you must also restore the tail of the log backup. Restore both of these backups by performing the following tasks.

How to restore a backup of the Encyclopedia volume metadata using SQL Server Management Studio

- 1 On Microsoft SQL Server Management Studio, in Object Explorer, expand Databases. Right-click the database named iserver and choose Tasks→Restore→Database, as shown in Figure 9-34.

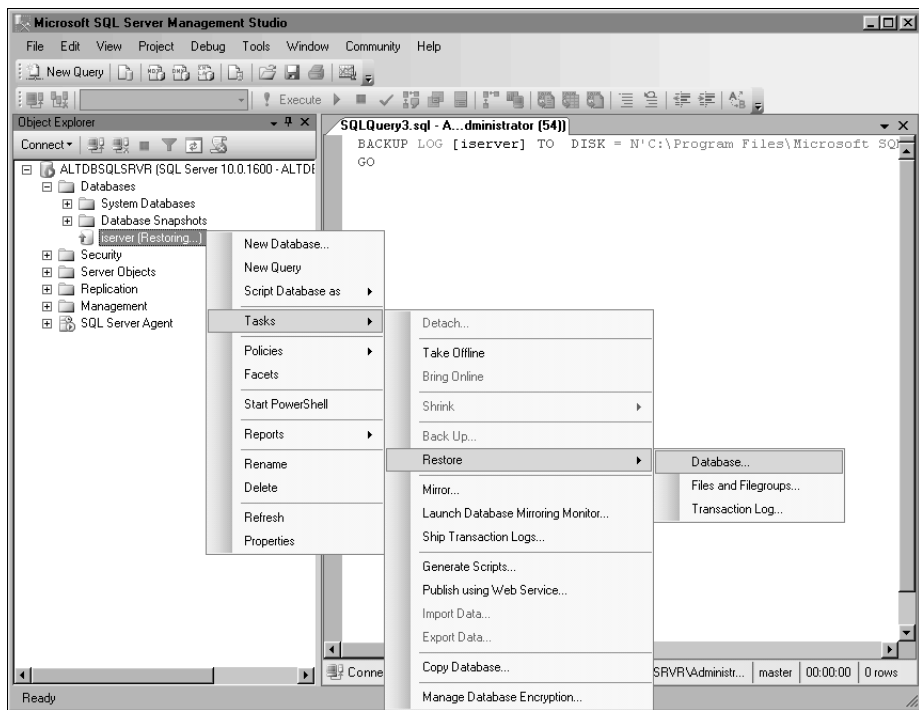


Figure 9-34 Choosing to restore the database

- 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 9-35.

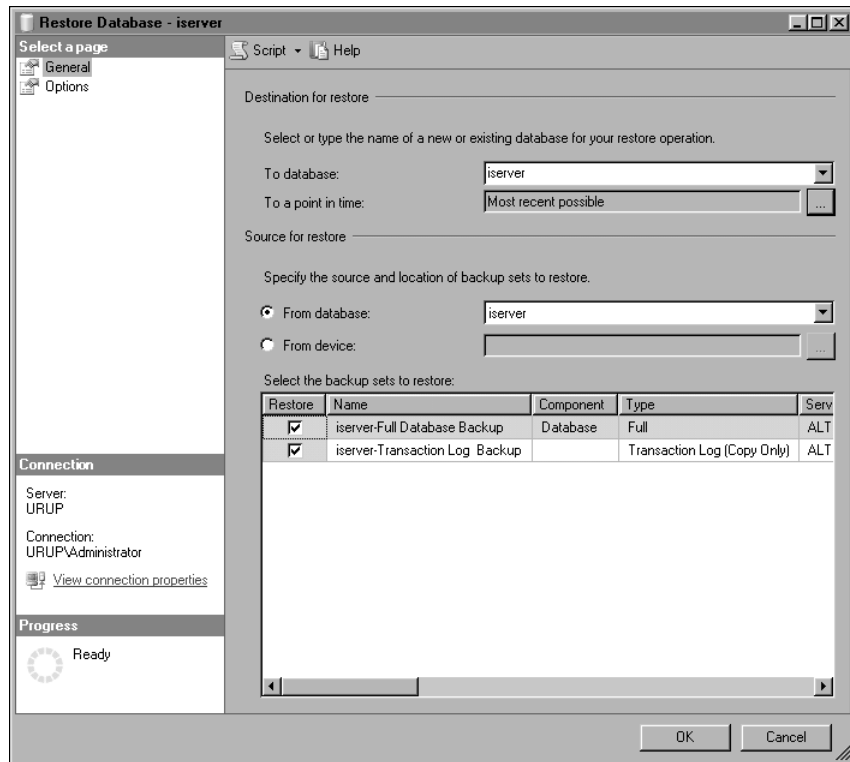


Figure 9-35 Configuring Restore Database—General

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 9-36.

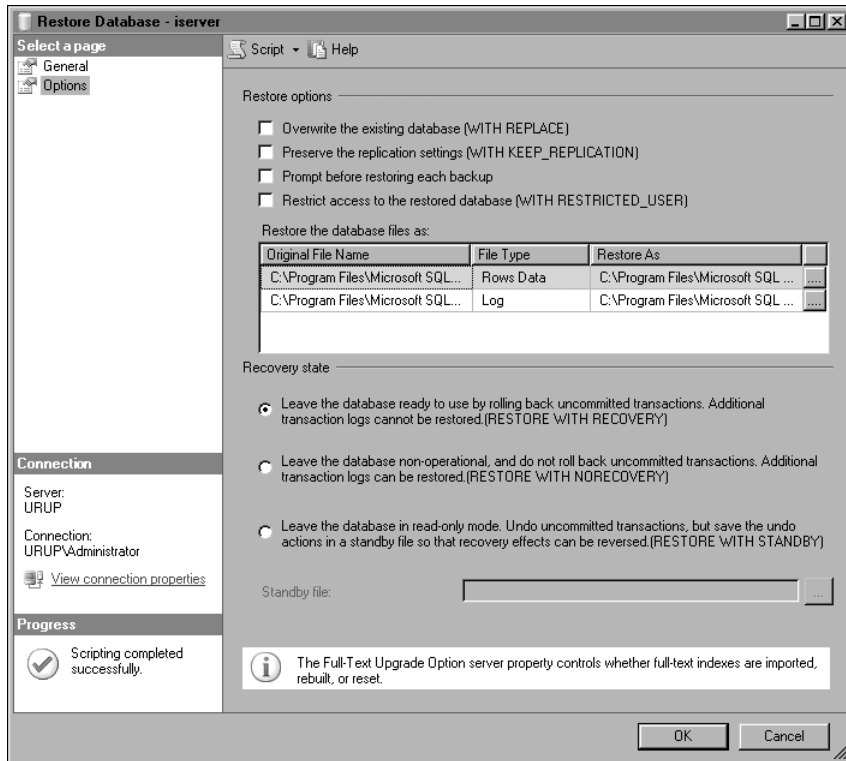


Figure 9-36 Configuring Restore Database—Options

Choose OK.

In the dialog box stating that the restore operation completed successfully, choose OK, as shown in Figure 9-37.




Figure 9-37 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 9-38.




Figure 9-38 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 9-39.




Figure 9-39 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 9-40.




Figure 9-40 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 9-41.




Figure 9-41 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 9-42.




Figure 9-42 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 9-43 shows the new folder.




Figure 9-43 Creating a new folder named encyc_backup

Back up Encyclopedia volume metadata using the DB2 graphical administration tool, Control Center, by performing the following tasks.

How to create a backup of the Encyclopedia volume metadata using DB2 Control Center

- 1 Choose Start→Programs→IBM DB2→DB2COPY1 (Default)→General Administration Tools→Control Center.
- 2 On Control Center, in Object View, navigate to the Databases folder. Select the ISERVER database, right-click, and choose Backup, as shown in Figure 9-44.

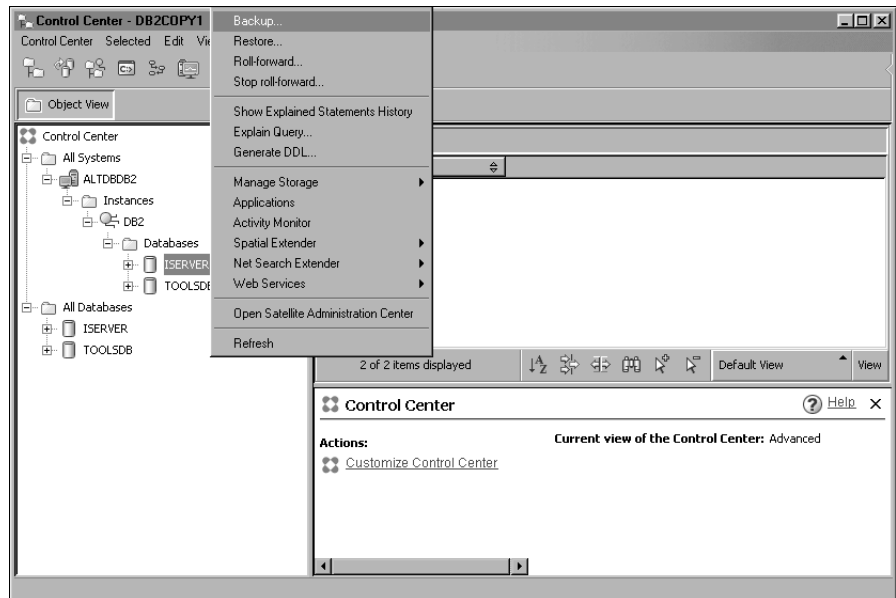


Figure 9-44 Choosing to back up the ISERVER database

- 3 On Backup Wizard—Introduction, select Backup entire database, as shown in Figure 9-45. Choose Next.
- 4 On Backup Wizard—Image, in Media Type, select File System, then select Add, as shown in Figure 9-46.
- 5 In Path Browser, navigate to the following location, as shown in Figure 9-47.
C:\Actuate\iServer\encyc_backup



Figure 9-45 Confirming the details of the ISERVER database

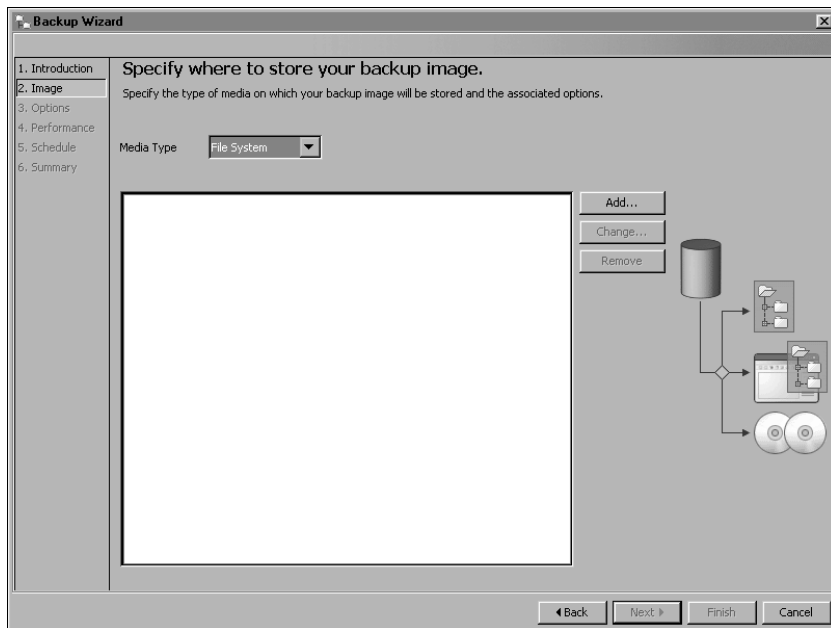


Figure 9-46 Choosing to add a backup location

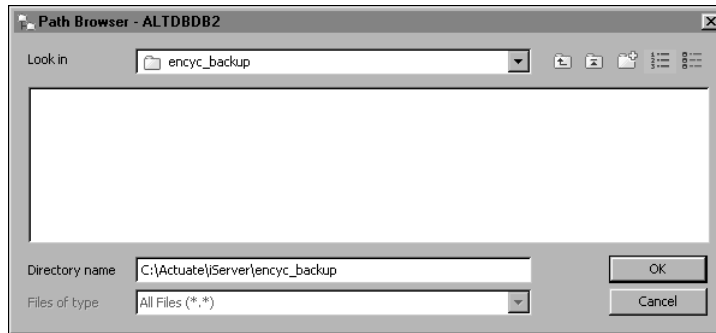


Figure 9-47 Navigating to the metadata backup location
Choose OK. Backup Wizard—Image appears as shown in Figure 9-48.

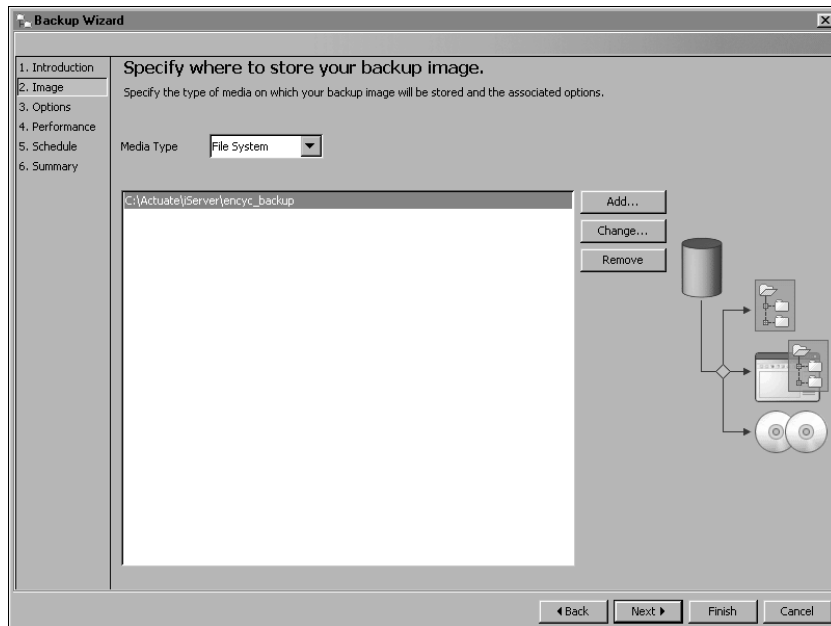


Figure 9-48 Viewing the chosen metadata backup location
Choose Next.

- 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 9-49.

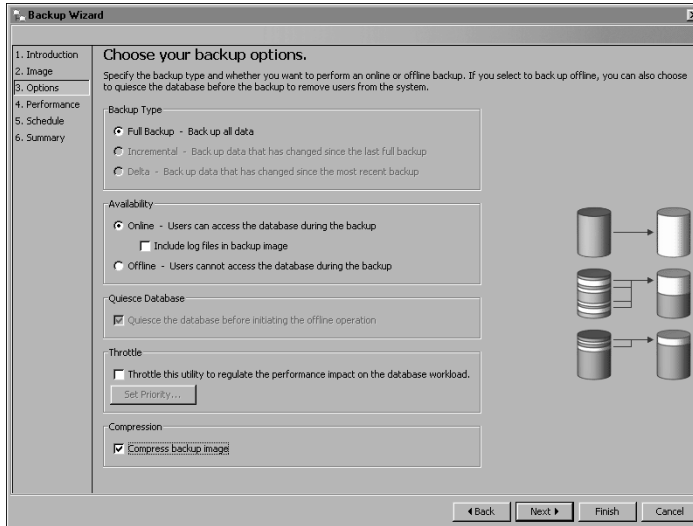


Figure 9-49 Configuring the backup options

Choose Next

- 7 On Backup Wizard—Performance, in the Buffers section, accept the recommended default values, as shown in Figure 9-50. Alternatively, select different values.

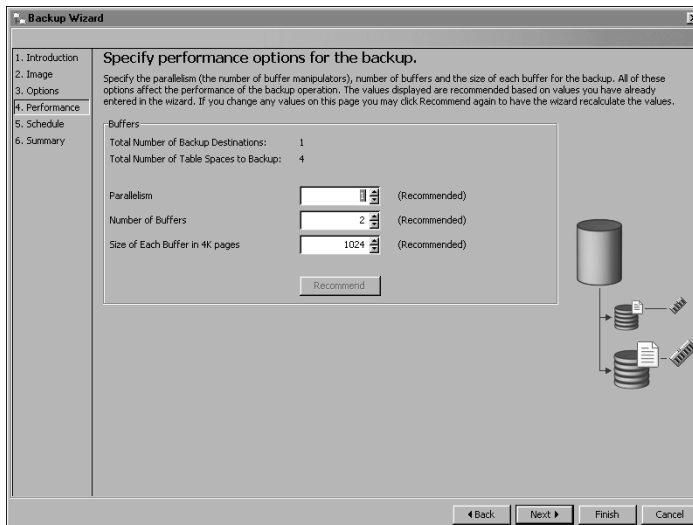


Figure 9-50 Specifying performance options

Choose Next.

- 8 On Backup Wizard—Schedule, accept the default values, as shown in Figure 9-51. Alternatively, specify different values.

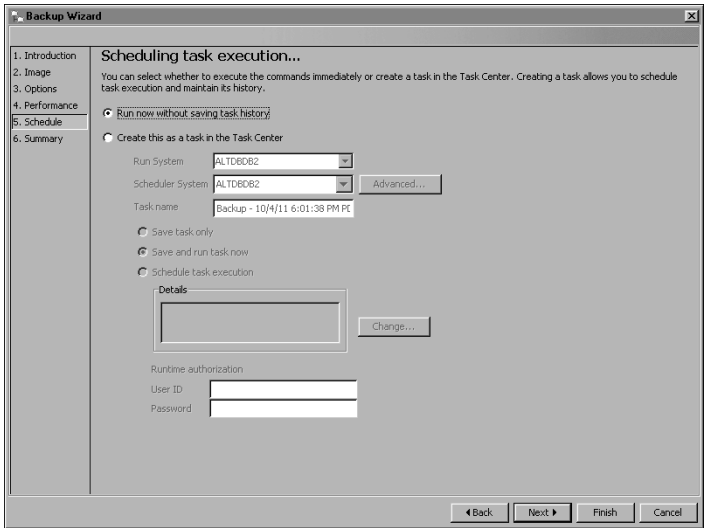


Figure 9-51 Selecting scheduling options

Choose Next.

- 9 On Backup Wizard—Summary, review the backup job parameters, as shown in Figure 9-52. Choose Back to change any backup job parameters, or choose Show Command to view the command that executes the backup.

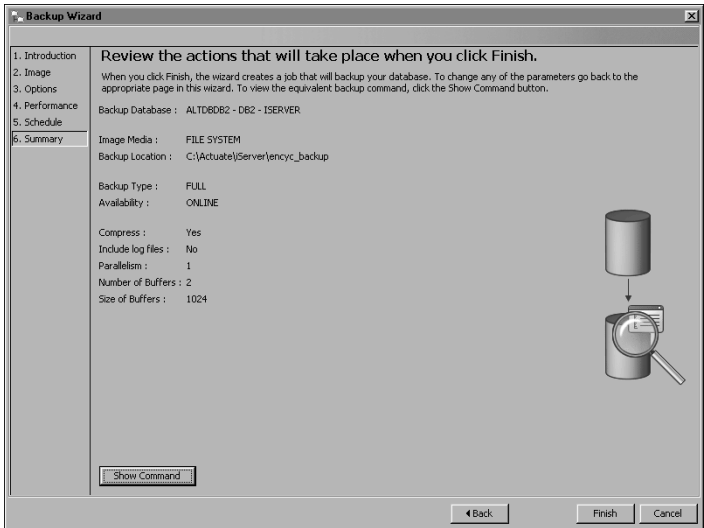


Figure 9-52 Reviewing backup job parameters

Choose Finish to execute the backup. The elapsed time for the backup job displays, as shown in Figure 9-53.




Figure 9-53 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 9-54.

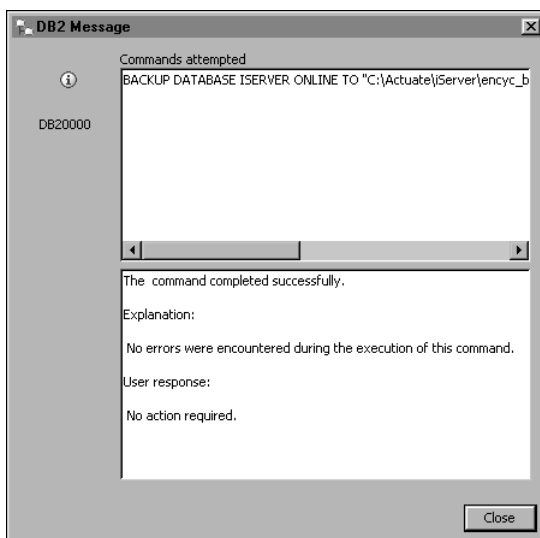


Figure 9-54 Viewing the commands executed and end-of-job messages

Alternatively, you can also back up the Encyclopedia volume metadata using the command line processor utility. The following example duplicates the operations performed in the previous section using the graphical administration tool, DB2 Control Center. You do not need to do both activities.

How to create a backup of the Encyclopedia volume metadata using the DB2 command line processor

- 1 Choose Start→Programs→IBM DB2→DB2COPY1 (Default)→Command Line Tools→Command Line Processor.

The command line processor appears, as shown in Figure 9-55.


```

c:\DB2CLP> DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => _

```

Figure 9-55 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 9-56:

```

BACKUP DATABASE ISERVER ONLINE
  TO "C:\Actuate\iServer\encyc_backup"
  WITH 2 BUFFERS BUFFER 1024 PARALLELISM 1
  COMPRESS WITHOUT PROMPTING

```

```

c:\DB2CLP> DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat DB2SETCP.BAT DB2.EXE
(c) Copyright IBM Corporation 1993,2007
Command Line Processor for DB2 Client 9.7.0

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 => connect to sample
  db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 => BACKUP DATABASE ISERVER ONLINE TO "C:\Actuate\iServer\encyc_backup" WITH
2 BUFFERS BUFFER 1024 PARALLELISM 1 COMPRESS WITHOUT PROMPTING
Backup successful. The timestamp for this backup image is : 20111012185050

db2 => _

```

Figure 9-56 Executing the commands to back up the database

- 3 To end the DB2 command line processor session, type the following command:

```
quit
```

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

How to back up the volume data folders

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

C:\Actuate\iServer\data

- 2 Navigate to the config folder that contains the acserverconfig.xml file. In Actuate 11 Service Pack 5, this file is located in the following subfolder:

C:\Actuate\iServer\data\config\11SP5

Select acserverconfig.xml, right-click, and choose Copy.

Copy this file to the following backup location:

C:\Actuate\iServer\encyc_backup

- 3 Navigate to AC_DATA_HOME\encyc. Select the file, fileType, status, and tempRov directories, right-click, and choose Copy. Copy these directories to the following backup location:

C:\Actuate\iServer\encyc_backup

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov directories may not exist. These directories contain information about job details and completion notices and do not appear until a job executes. If these directories are not present in the environment, simply back up the file and fileType directories.

The contents of the backup folder appear as shown in Figure 9-57.

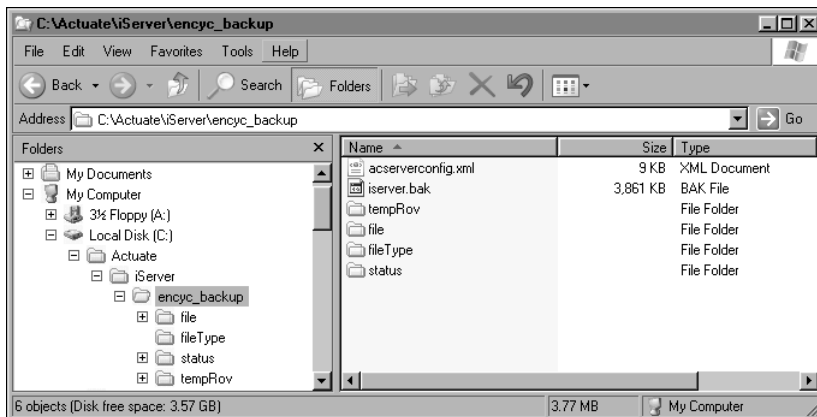


Figure 9-57 Viewing the backed-up files

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 9-58.




Figure 9-58 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 5, this file is located in the following subfolder:
`C:\Actuate\iServer\data\config\11SP5`
Select acserverconfig.xml, right-click, and choose Delete. Confirm the deletion.
- 3 Navigate to AC_DATA_HOME\encyc. Select the file, 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 9-59.

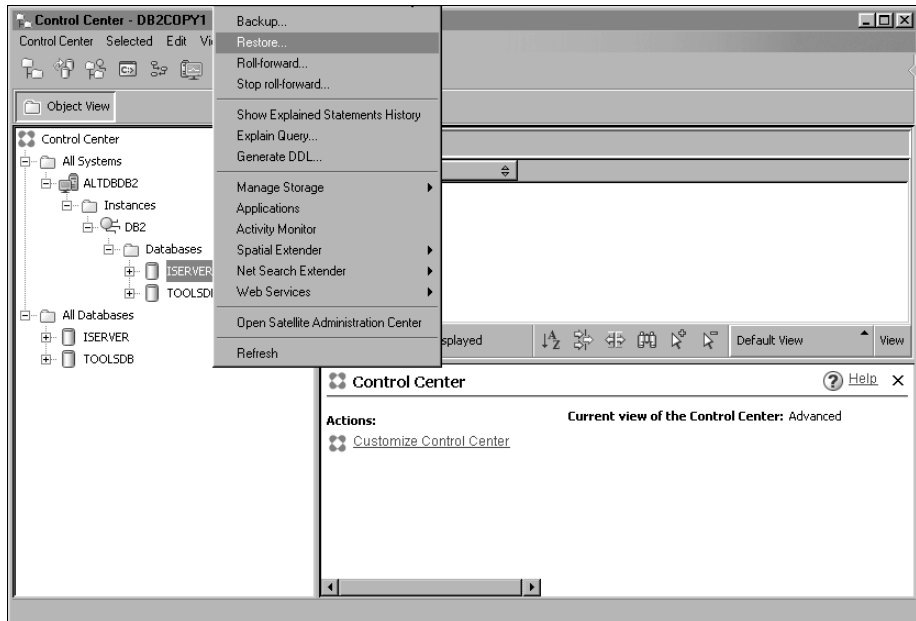


Figure 9-59 Choosing to restore the database

- 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 9-60. Choose Next.

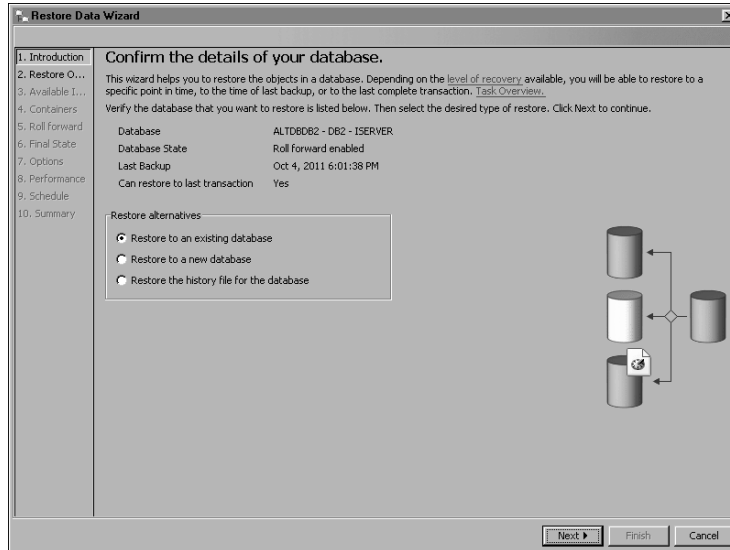


Figure 9-60 Confirming the details of the ISERVER database

- 3 On Restore Data Wizard—Restore Objects, accept the default selection, Restore the entire database, as shown in Figure 9-61. Choose Next.

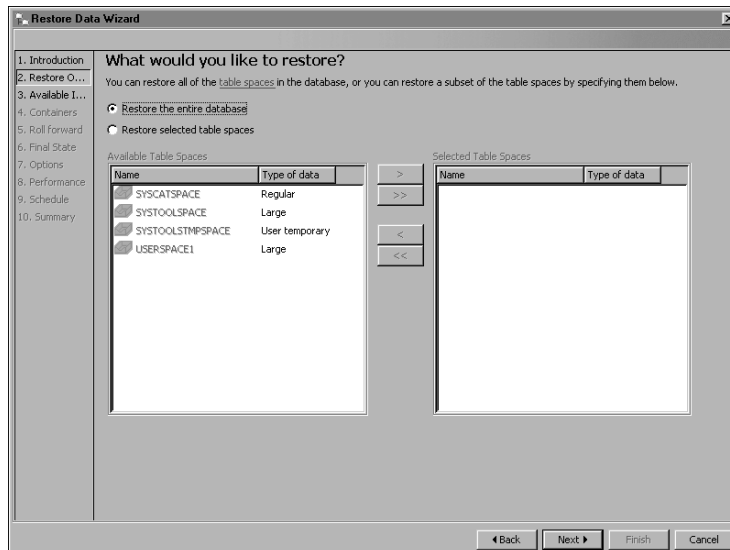


Figure 9-61 Selecting the backup image to restore

- 4 On Restore Data Wizard—Available Images, in Available backup images, select the image you want to restore, as shown in Figure 9-62.

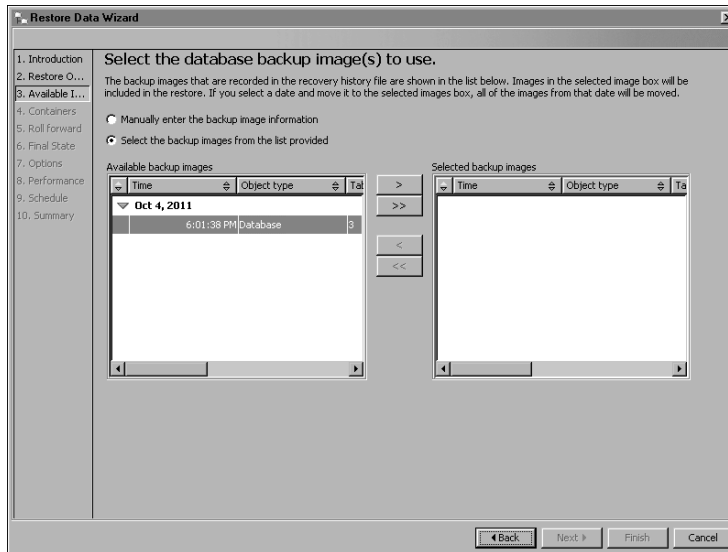


Figure 9-62 Selecting the backup image to restore

Choose the right arrow to move the image to Selected backup images, as shown in Figure 9-63. Choose Next.

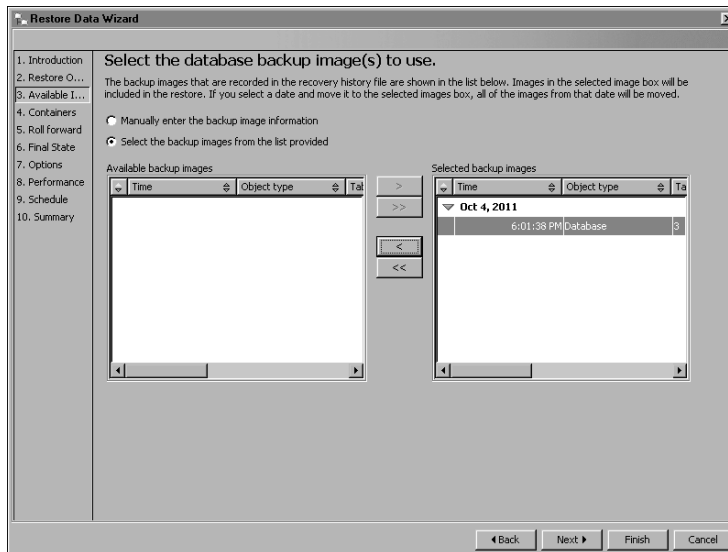


Figure 9-63 Moving the selected image to Selected backup images

- 5 On Restore Data Wizard—Containers, accept the default options, as shown in Figure 9-64. Choose Next.

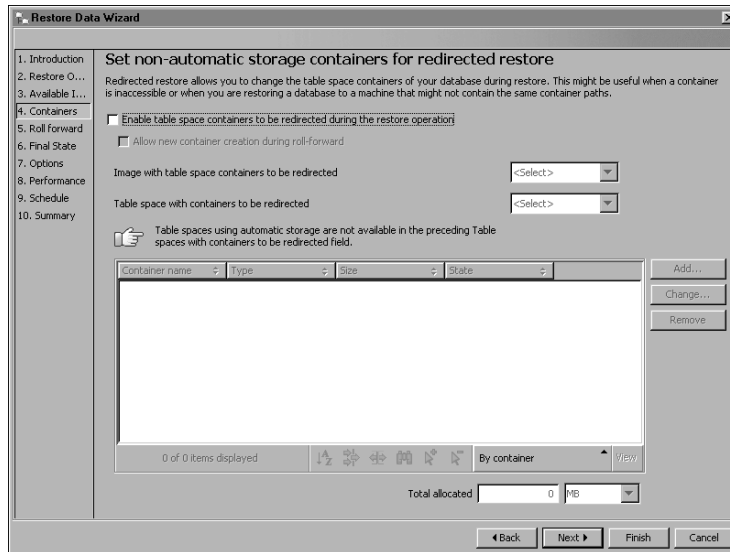


Figure 9-64 Choosing whether to enable redirecting of table space containers

- 6 On Restore Data Wizard—Roll forward, accept the default option, Restore Only, as shown in Figure 9-65. Choose Next.

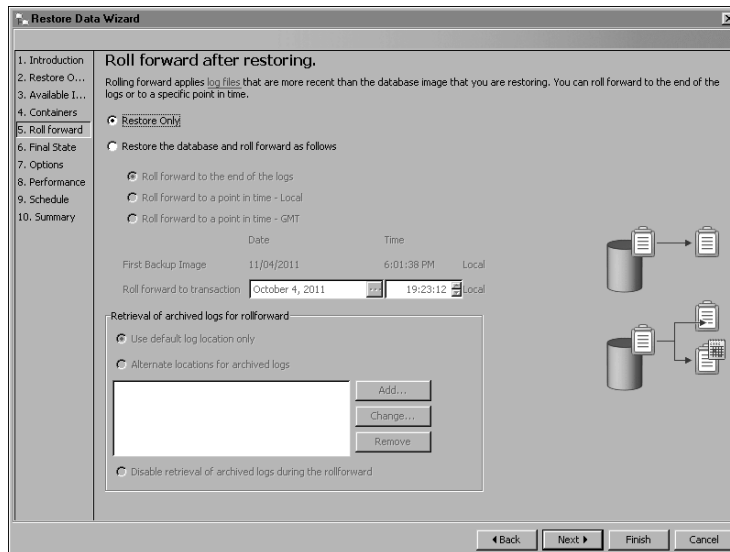


Figure 9-65 Choosing to restore the database backup only

- 7 On Restore Data Wizard—Final State, select Complete the restore and return to the active state, as shown in Figure 9-66. Choose Next.

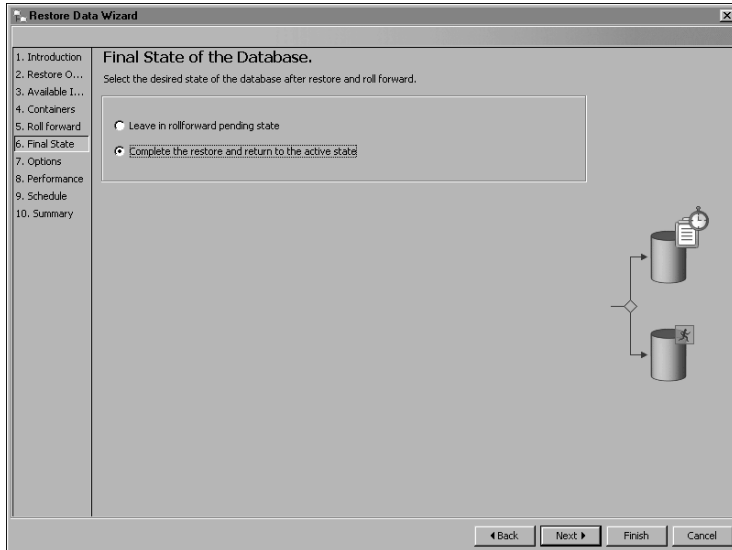


Figure 9-66 Choosing to complete the restore

- 8 On Restore Data Wizard—Options, accept the default options, as shown in Figure 9-67. Choose Next.

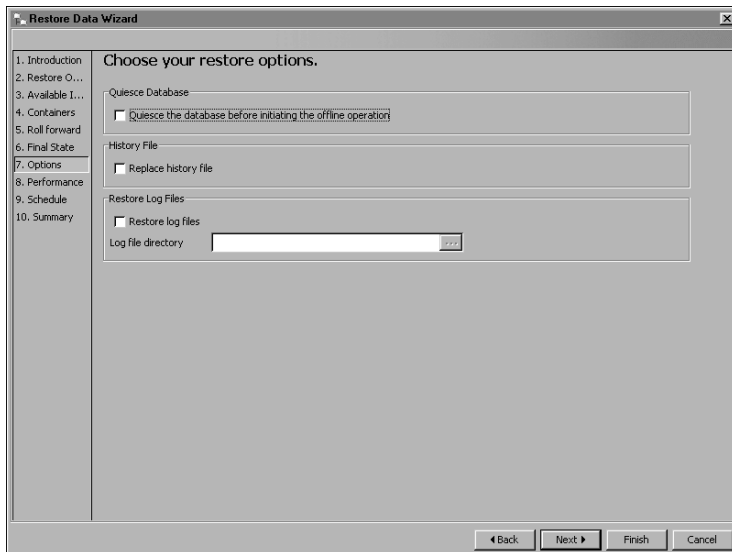


Figure 9-67 Choosing restore options

- 9 On Restore Data Wizard—Performance, accept the default values, as shown in Figure 9-68. Choose Next.

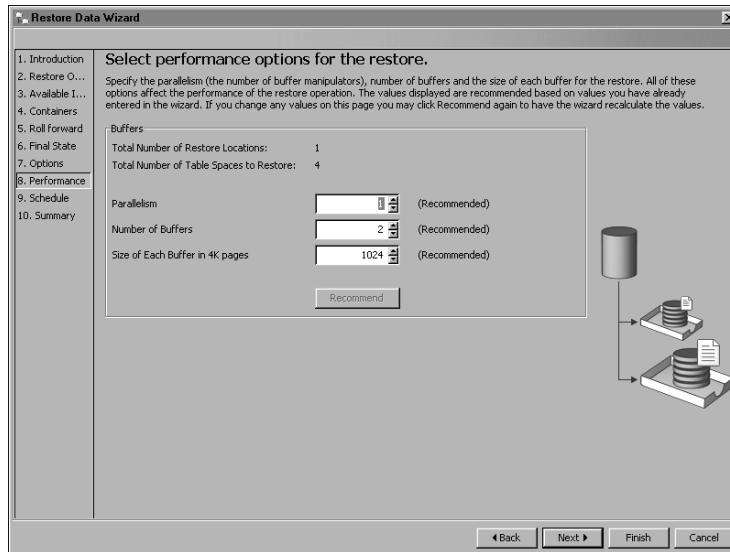


Figure 9-68 Selecting performance options

- 10** On Restore Data Wizard—Schedule, accept the default values, as shown in Figure 9-69. Choose Next.

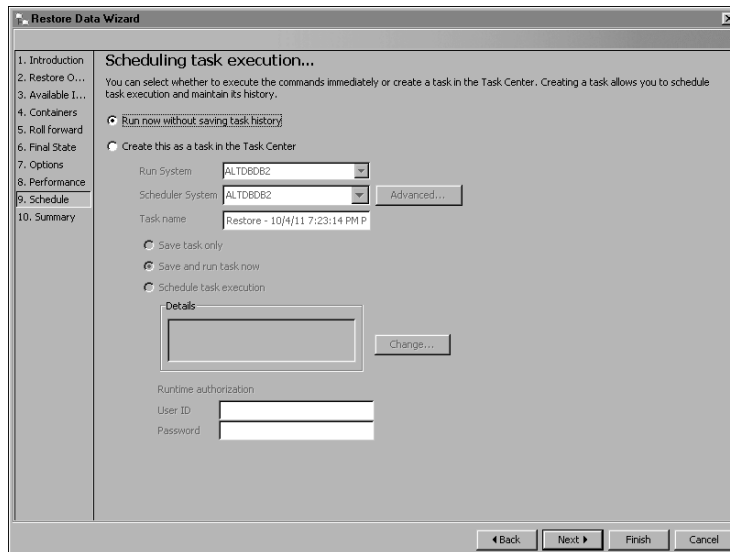


Figure 9-69 Setting backup job scheduling options

- 11** On Restore Data Wizard—Summary, review the restore job parameters, as shown in Figure 9-70. Choose Back to change any restore job parameters, or choose Show Command to view the SQL that executes the restore.

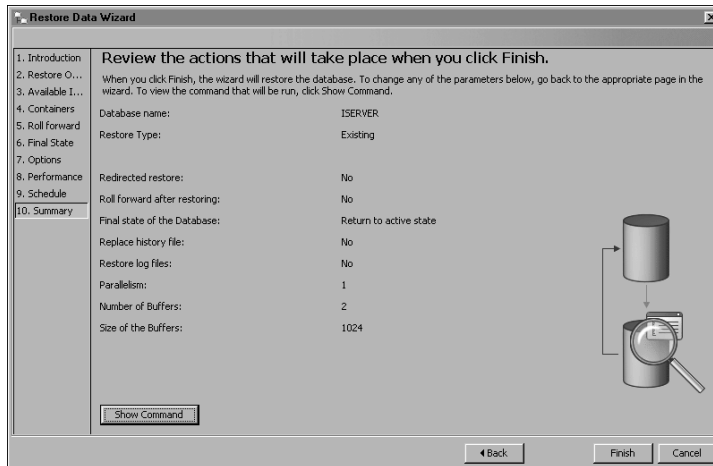


Figure 9-70 Reviewing restore job parameters

Choose Finish to execute the restore. The elapsed time for the restore job displays, as shown in Figure 9-71.




Figure 9-71 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 9-72.

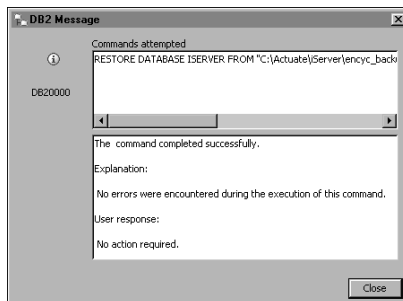


Figure 9-72 Viewing the command executed and end-of-job messages

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 9-73.

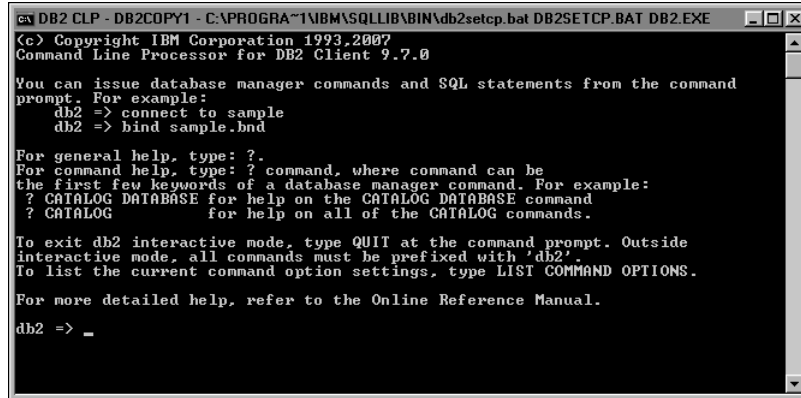


Figure 9-73 Opening the command line processor

- 2 Execute the following command, as shown in Figure 9-74:

```
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.

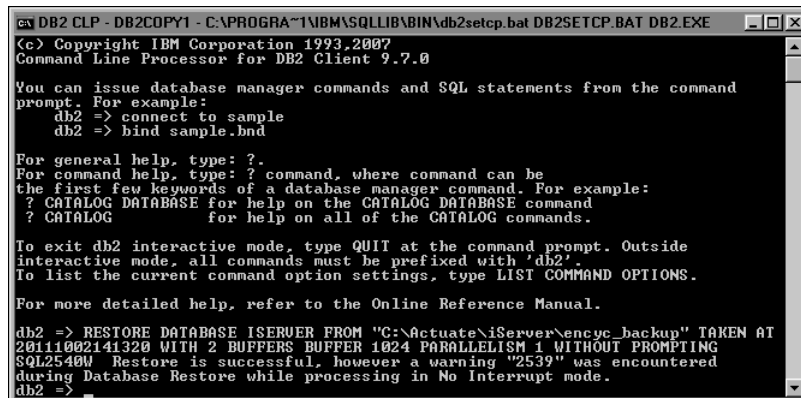


Figure 9-74 Restoring the database from the command line

- 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 9-75.




Figure 9-75 Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the DB2 database utilities, go to the following location:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-0910db2incrementalbackup/index.html>

Index

A

- AC_CONFIG_HOME attribute 221, 224
- AC_JAVA_HOME variable 37
- AC_JRE_HOME variable 37
- AC_JRE64_HOME variable 37
- AC_JVM_HOME variable 37
- AC_PMD_WINDOWS_CPUS variable 226, 228
- AC_SERVER_HOME parameter 141, 158
- AC_TEMPLATE_NAME variable 184
- access permissions. *See* privileges
- accessing
 - configuration parameters 13
 - data 13, 218
 - Encyclopedia volumes 38, 216, 218
 - help topics xi
 - iServer features 37, 216
 - iServer System components 35
 - license files 13
 - metadata database schemas 4
 - online help 209
 - RDBMS database documentation 7
 - report files 22
 - resources 196
 - Update tool 206
 - volume databases 16
 - web-based content 218
- accounts
 - cluster configurations and 23
 - creating user 22, 23, 42
 - installing alternate databases and 50, 64
 - installing cluster nodes and 176
 - installing Information Console and 192
 - managing 16
 - setting passwords for 29, 31
 - setting up e-mail 231
 - starting iServer 11 service and 31
 - starting Management Console and 32
 - upgrades and 71
- acpmdconfig.xml 13, 162
- Acrobat Catalog. *See* Adobe Acrobat Catalog
- acserverconfig.xml 13, 162
- acserverlicense.xml 13
- activity logs 197
- Actuate Basic reports 218
- Actuate ftp site 25
- Actuate licensing options 216, 217, 225
- Actuate Query Option 217
- actuate.com vii
- ActuateInformationConsole.exe 189
- ActuateInformationConsole.war 195
- ActuateiServerIntegrationTechnology.exe 202
- ActuateLocalizationandOnlineDocumentation.exe xi
- ActuateOne technology vii
- Add Partition link 113
- Add Volume link 123, 138
- adding
 - administrator accounts 22, 71
 - backup folders 240, 255, 274
 - cluster nodes 162
 - clusters 168
 - desktop shortcuts 35, 87, 105
 - e-mail accounts 231
 - Encyclopedia volumes 25, 153, 154
 - indexes 5
 - JDBC drivers 16
 - license keys 220
 - licensing options 225
 - metadata databases 11
 - passwords. *See* passwords
 - resource groups 37
 - schema owners 42
 - system databases 43, 50
 - system schemas 25, 42, 155, 157
 - user accounts 23, 42
 - volume schemas 25, 42
- admineds.bat file 111, 118, 136
- administrate_encyclopedia_data_store.bat file 111, 118, 136
- administrate_system_data_store.bat file 155
- administrative reports 17
- administrator accounts 22, 32, 64, 71
- administrators
 - backing up database schemas and 24
 - backing up Encyclopedia and 236, 237

- configuring iServer clusters and 13, 14
- customizing database connections and 16
- installing DB2 databases and 50, 51
- installing iServer and 22, 36, 39, 42
- installing Oracle databases and 45
- installing Postgre databases and 43
- installing SQL Server databases and 47
- managing Encyclopedia and 5, 6
- managing iServer System and 15
- obtaining licenses and 216
- optimizing iServer System and 13, 162
- preventing data loss and 7
- recreating systems schemas and 155, 157
- sharing metadata databases and 5
- testing new releases and 39, 40
- upgrading iServer and 70, 71, 72, 93
- adminsds.bat script 155
- Adobe Acrobat Catalog utility xi
- alternative databases. *See* metadata databases; third-party databases
- analyzing
 - search results 218
- Apache Tomcat servers 189, 191
- application pages. *See* web pages
- application programming interfaces 15, 17
- application server cache 198
- application server configurations 189, 195, 196
- application servers 189, 195, 198
 - See also* servers
- APPLICATION_USER parameter 158
- APPLICATION_USER_NAME parameter 142
- APPLICATION_USER_PASSWORD parameter 142
- applications
 - accessing Encyclopedia and 38
 - backward compatibility with 17
 - creating production environment for 40
 - deploying web 198
 - developing 15, 16
 - installing Open Security 28, 38
 - restricting CPU processes for 225
 - running iServer processes and 8, 10
- archives (cloud deployments) 6
- Archiving and Purging page 238
- archiving report files 238

- asynchronous Factory service 11
- autoarchive file purging 238
- auto-login scripts (pgpass) 241
- automated installation option 6
- automated installation programs 16
- automatic rollback 71
- automatic upgrades 70

B

- backing up
 - data files 237
 - database schemas 24, 241
 - Encyclopedia volumes 236, 237, 238, 239, 249, 254, 273
 - folders 242, 261, 282
 - metadata 24, 54, 237
 - PostgreSQL databases 239
 - program files 75, 97
 - system schemas 7, 72, 155
- backup operations 6, 7, 16
- backup utilities 237
- backward compatibility 17
- Basic reports 218
- batch files
 - Data Store Upgrader 91, 92
 - Squirrel Data Exporter 148, 149
 - System Data Store Administrator 155
- binary files 28
- BIRT 360 Option 217
- BIRT Data Analyzer Option 217
- BIRT Designer Professional 218
- BIRT Interactive Viewer Option 217
- BIRT iServer 16
 - See also* iServer
- BIRT iServer System names 29, 57
- BIRT open source projects vii
- BIRT Option 217
- BIRT Page Level Security option 218
- BIRT reports 17, 218
 - See also* reports
- BIRT SmartSheet Security Option 218
- BIRT Spreadsheet Designer 218
- BIRT Spreadsheet Option 218
- BIRT Studio 196
- BIRT Studio Option 218
- BIRT Viewer 197
- BIRT_RESOURCE_PATH parameter 196

- BIRT_VIEWER_LOG_DIR parameter 197
- birt-exchange.com vii
- browsers. *See* web browsers
- bulk load script files 151
- Business Intelligence technology vii

C

- cache (web pages) 198
- cache conflicts 198
- Caching processes 228
- Caching service 9, 11
- CachingService element 228
- changing
 - cluster machines 224
 - CPU binding 230, 231
 - database encoding 17
 - default locales 196
 - license file names 223
 - licensing options 225
 - metadata databases 4, 5
 - network cards 224
 - passwords 32
 - time zones 196
 - volume names 150
- character encoding 17, 50
- character sets 17
- charts 37
- CLASSPATH variable 38
- clearing web browser cache 198
- client applications 8
- client/server models 13
- cloud deployments 6
- cluster IDs 13, 162
- cluster nodes 162, 171
 - creating shared folders for 164
- clusters
 - adding nodes to 13, 162
 - associating with machine IDs 221
 - binding iServer processes to 225
 - changing machines for 224
 - configuring 13, 14
 - connecting to iServer System and 8
 - creating 168
 - deploying web applications and 198
 - determining number of processors for 230
 - distributing requests among 8, 14
 - exceeding CPU licenses for 231
 - failing 162
 - failover operations and 7
 - getting IP addresses for 167
 - installing cluster nodes for 172–184
 - installing iServer for 28
 - licensing options for 219, 220, 225
 - load balancing for 14, 162
 - managing volume metadata and 5
 - running iServer services and 8
 - setting configuration home location for 174
 - setting port numbers for 178
 - setting up RDBMS databases for 11
 - setting up system data store for 155
 - setting up user accounts for 23
 - testing connections for 168
 - turning off firewalls for 166
- collation 47, 50
- collecting machine information 221
- command line utilities 16, 239
- comments 141
- config directory 174
- CONFIG_SCHEMA_NAME parameter 142, 158
- ConfigHomeDirectory variable 185
- Configuration Console
 - adding alternate databases and 64
 - adding partitions and 113
 - archiving report files and 238
 - configuring BIRT iServer and 16
 - configuring clusters and 13
 - creating Encyclopedia and 153, 154
 - creating system schemas and 155
 - installing 174, 188
 - running iServer processes and 8
 - running stand-alone servers and 9
 - sending notifications and 231
 - setting user names and passwords for 32
 - shutting down iServer and 107
 - updating license keys and 223
 - upgrades and 89
 - viewing licensing information and 219
- configuration files 25, 140
- configuration home locations 174, 185
- configuration parameters 13, 162, 196
 - See also* parameters
- configuration templates 13, 162

- configurations
 - binding CPU processes and 225, 228
 - connecting to databases and 16, 17
 - creating cluster nodes and 8
 - customizing database encoding and 17
 - deploying Information Console and 189, 195, 197
 - failover operations and 7
 - installing cluster nodes and 13, 174
 - installing Information Console and 189, 196
 - installing iServer and 24, 28
 - licensing iServer and 216, 220
 - purging report files and 238
 - renaming licensing files and 224
 - running Open Security applications and 38
 - running RDBMS databases and 11
 - storing metadata for 4
 - testing installations and 39
 - upgrading and 140
- configuring
 - application servers 189, 196
 - clusters 13, 14
 - iServer System 6
 - JDBC drivers 16, 17
 - service privileges 24
 - system data store 155
 - upgrade utilities 140
- connection pooling 17, 43
- connections
 - accessing multiple volume schemas and 43
 - accessing volume databases and 16
 - changing encoding schemes and 17
 - configuring clusters and 8, 13
 - configuring Information Console 197
 - creating DB2 schemas and 51
 - creating Oracle schemas and 46
 - creating PostgreSQL schemas and 44
 - creating SQL Server schemas and 49
 - customizing database 16
 - testing 168
- context paths 191, 196
- context-sensitive help 206
- conventions (documentation) xii
- copying
 - files 237
 - license files 224
- corrupted schemas 72
- CPU binding 225–232
- CPU binding validation 230, 231
- CPU-based licenses 216
 - determining number of 225
 - exceeding number of 230, 231
 - viewing information about 230
- CPUs
 - binding to multiple core 228
 - deploying iServer over multi-threaded 225
 - determining number of 228, 230
 - hyperthreading and 229
 - licensing options for 225
 - restricting processes for 225
 - running encycsrvr11 processes and 230
 - running Information Console and 196
 - viewing maximum number of 230
 - viewing processor IDs for 226
- CREATE_SCHEMA parameter 143
- creating
 - administrator accounts 22, 71
 - backup folders 240, 255, 274
 - cluster nodes 162
 - clusters 168
 - e-mail accounts 231
 - Encyclopedia volumes 25, 153, 154
 - indexes 5
 - metadata databases 11
 - passwords. *See* passwords
 - production environments 40
 - production staging areas 39
 - report designs 218
 - resource groups 37
 - schema owners 42
 - shared folders 164
 - system databases 43, 50
 - system schemas 25, 42, 155, 157
 - test environments 39
 - user accounts 23, 42
 - volume schemas 25, 42
- credentials. *See* user credentials
- cubeview files 217
- current release xi
- custom configurations 28
- custom installations 28
- customer licenses 223
- customizing

- database connections 16
- database encoding 17
- Information Console 196, 197
- metadata databases 5

D

- dashboards vii, 217

- data

- accessing 13, 218
 - backing up Encyclopedia and 236
 - caching 11
 - preventing loss of 7, 24
 - recovering 6, 7, 16
 - setting default location for 28
 - sharing 14
 - upgrades and 71

- Data Analyzer Option 217

- Data Connector Option 218

- data cubes. *See* cubes

- data directory 28

- Data Exporter

- configuring 140
 - migrating iServer installations and 6
 - setting up export environments and 151

- data files 28, 237

- See also* files; report files

- data objects 5, 236

- data retrieval operations 11

- data sources 11

- Data Store Administrator

- See also* System Data Store Administrator
 - configuring 140
 - creating Encyclopedia and 153, 154
 - creating schema elements and 153
 - importing Encyclopedia and 141, 150
 - migrating iServer installations and 6
 - running 141, 150
 - setting properties for 141

- Data Store Upgrader

- configuring 140
 - running 90, 91, 157
 - setting properties for 157, 158

- data stores 155, 197

- See also* database schemas

- data types 17, 50

- DATA_EXPORT_FOLDER parameter 143

- DATA_EXPORT_FORMAT parameter 143

- DATA_IMPORT_FOLDER parameter 143

- DATA_IMPORT_FORMAT parameter 143

- database administration tool

- backing up Encyclopedia and 240–242

- restoring Encyclopedia and 244, 245

- uninstalling 36

- upgrades and 83

- database collation 47, 50

- Database Credentials option 29

- database drivers. *See* drivers

- Database Schema Information page 60

- database schema names 29, 90, 146, 147

- database schema owners

- creating 42

- installing DB2 databases and 50, 51

- installing iServer and 5

- installing Oracle databases and 45, 46

- installing PostgreSQL databases and 43, 44

- installing SQL Server databases and 47, 48

- setting passwords for 29, 147

- database schemas

- accessing 4

- adding Encyclopedia to 153, 154

- assigning privileges 43

- backing up 7, 24, 155, 241

- configuring system data store as 155

- creating 25, 42

- customizing metadata databases and 5

- depopulating 144

- dropping 144

- importing volume data to 150

- initializing 25

- installing DB2 databases and 51

- installing Oracle databases and 45, 46

- installing PostgreSQL databases and 44, 45

- installing SQL Server databases and 48, 49

- naming 29, 43, 147

- populating 146, 153

- preventing data loss and 7

- removing Encyclopedia from 152

- restoring 155, 157

- setting passwords for 147

- storing metadata and 4

- upgrades and 71

- user-defined functions and 53

- database servers. *See* servers

- database superuser 29, 147
- DATABASE_HOST parameter 143, 158
- DATABASE_INSTANCE parameter 143, 158
- DATABASE_NAME parameter 144, 159
- DATABASE_PORT parameter 144, 159
- DATABASE_TYPE parameter 141, 144, 158
- databases
 - accessing documentation for 7
 - accessing Encyclopedia and 38
 - adding data store schemas and 157
 - changing metadata tables and 5
 - configuring failover operations for 7
 - connecting to 16
 - copying files in 237
 - creating metadata 11
 - creating system 43, 50
 - customizing 5
 - importing Encyclopedia and 150
 - installing iServer and 22, 25, 28, 42
 - installing metadata 29, 58
 - localizing 17
 - managing Encyclopedia and 5, 6
 - running iServer clusters and 11, 14
 - selecting alternate 57
 - setting passwords for 29
 - setting up iServer environments for 42
 - specifying default locations for 28
 - storing metadata and 4
 - uninstalling 36
 - upgrades and 70, 71, 72, 93
 - viewing incomplete files in 237
- DB2 command line processor 280
- DB2 Control Center 274–280, 283
- DB2 databases
 - accessing documentation for 7
 - connecting to 51
 - creating schema owners for 50, 51
 - creating user accounts for 50
 - exporting metadata to 148
 - extending SQL functionality 52–53
 - installing 50, 51, 58
 - managing Encyclopedia and 5
 - naming schemas for 51
 - running volume backup and restore
 - operations with 273, 274, 280, 283
 - setting port number for 58
 - setting up iServer environments for 50–51
 - storing metadata and 4
 - taking Encyclopedia offline for 283
 - taking Encyclopedia online for 292
- default directories. *See* directories
- default locale 31, 196
- default paths 28
- default ports. *See* ports
- default time zones 196
- default user names 32
- DEFAULT_DATABASE_NAME parameter 144
- DEFAULT_LOCALE parameter 196
- DEFAULT_TIMEZONE parameter 196
- DEFAULT_VOLUME parameter 197
- DELETE_ALL_DATA parameter 144
- DELETE_DATA parameter 144
- deleting
 - database schemas 144
 - installation directories 36
 - product files. *See* uninstalling
- deploying
 - Information Console 189, 195–198
 - iServer 6, 13, 14, 162, 225
 - new releases 38
 - spreadsheets 218
 - WAR files 195
 - web applications 198
- deployment tools 196
- DEPOPULATE_SCHEMA parameter 144
- design files 5
- designers vii
- designs 11, 39, 217, 218, 219
- desktop shortcuts 35, 87, 105
- destination folder locations 28
- developers 15, 43
- developing applications 15, 16
- DHTML reports 218
- diagnostic fixes 9, 14, 163
- digit wildcard character 51
- directories
 - backing up volume data 242, 251, 261, 282
 - configuring cluster nodes and 162
 - creating cluster nodes and 13
 - customizing WAR files and 197
 - extracting product files to 188
 - file structure changes and 89
 - ICU libraries in 23

- installing cluster nodes and 174
- installing iServer and 28
- installing iServer Integration Technology and 202
- installing JDK files and 37
- mapping Oracle database dump 250
- removing installation 36
- restoring 245, 252, 263, 283
- running multiple releases and 37
- directory paths 28, 113, 236
- disk partitions 237
 - See also* volume partitions
- disk space 37
- displaying
 - charts 37
 - folder properties 164
 - host names 167
 - IP addresses 167
 - licensing information 219, 220
 - processor IDs 226
 - ReadMe files 195, 205
 - reports 11, 217
- DLLs 38
- document files 5, 219
- document generation service 11
- documentation
 - accessing vii–xi
 - administering iServer System and 16
 - conventions for xii
 - deploying WAR files and 195
 - installing online 35, 206
 - managing Encyclopedia and 7
 - obtaining x
 - upgrades and 87, 105
- documentation update tool 206
- domain servers 23
- downloading
 - documentation x
 - product files 25
 - report files 11
- drivers 11, 14, 16, 17
- DROP_SCHEMA parameter 144
- dual-core CPUs 228
- dump directories 250
- dynamic link libraries 38

E

- e.Analysis Option 218
- e.Report Data Connector Option 218
- e.Report Designer Professional 218
- e.Report Option 218
- e.reports 218
- Eclipse BIRT open source projects vii
- elastic iServer clustering 13, 162
- e-mail. *See* notifications
- encoding 17, 50
- encyc directory 36
- Encyclopedia Data Store Administrator
 - configuring 140
 - creating Encyclopedia and 153, 154
 - creating schema elements and 153
 - importing Encyclopedia and 141, 150
 - migrating iServer installations and 6
 - running 141, 150
 - setting properties for 141
- Encyclopedia Data Store Upgrader
 - configuring 140
 - running 90, 91, 157
 - setting properties for 157, 158
- Encyclopedia Metadata Storage page 29, 57
- Encyclopedia processes. *See* encycsrvr11 processes
- Encyclopedia volumes
 - accessing multiple 218
 - accessing objects in 216
 - adding to database schemas 43, 45, 47, 50, 154
 - backing up 236, 237, 238, 239, 249, 254, 273
 - connecting to database for 16
 - controlling access to 38
 - CPU binding and 231
 - creating default 25
 - creating new 153
 - creating schema database for 45, 49, 51
 - creating schema owners for 44
 - customizing metadata databases for 5
 - downloading files from 11
 - exporting 141, 150
 - importing 141, 150
 - installing Information Console and 189

- installing metadata database for 4, 5, 22, 42
- installing sample 17
- installing system database for 29
- licensing and 216
- localizing databases for 17
- managing 5, 6
- migrating to current release 70, 89
- naming 29, 123, 138, 146
- preserving metadata for 24
- preventing data loss and 7, 24
- removing from schemas 152
- renaming 150
- restoring. *See* restoring Encyclopedia volumes
- running iServer clusters and 14, 162
- running iServer processes and 8
- running multiple schemas on 43
- setting location of 28
- specifying default 197
- specifying primary partition for 123
- specifying schema owners for 29
- starting 9, 124, 139
- storing metadata for 4, 11
- troubleshooting 237
- uninstalling databases for 36
- upgrading iServer and 6, 71
- EncyclopediaUpgrader utility. *See* Encyclopedia Data Store Upgrader
- encycsrvr11 processes 225, 230, 231
- environment variables 13, 37, 91, 162, 226
- environments 195
- error logging applications 15
- error logging reports 17
- escape characters 140
- evaluation licenses 30
- events 8, 15
- example reports 17
- example volume 17
- Excel formats vii
- Excel spreadsheets 218
- executable files 37
- expdp command line options 251
- expdp command line utility 250
- Expiration time of deleted files property 239
- expired licenses 220, 221
- EXPORT_ALL_DATA parameter 144

- EXPORT_DATA parameter 144
- exporting Encyclopedia volumes 141, 150
- external data sources 11
- external security systems 28, 38
- external user information 15
- extracting product files 188

F

- Factory processes 37, 228, 231
- Factory service 8, 11, 37
- failover operations 6, 7
- features 37, 216
- file I/O operations 11
- file names 220
- file purging properties 238
- file system backup utilities 237
- file systems 5, 11
- FILE_LOG_LEVEL parameter 145, 159
- files
 - accessing license 13
 - accessing online help 209
 - accessing report 22
 - archiving 238
 - backing up 75, 97, 237
 - configuring upgrade utilities and 140
 - copying 237
 - downloading product 25, 188
 - downloading report 11
 - installing iServer and 25
 - installing JDK software and 37
 - installing online help 206
 - obtaining license key 220
 - overwriting product 6
 - purging 238
 - running processes and 73, 94
 - saving 75, 97, 141
 - selecting license 30
 - sharing 164
 - specifying default locations for 28
 - storing report-specific 5
 - transmitting license keys and 223
 - updating 206
 - upgrading iServer and 37, 39, 72
- firewalls 166, 189
- fixes 9, 14, 163
- Flash gadgets vii

- folders
 - backing up 242, 261, 282
 - creating backup 240, 255, 274
 - installing iServer and 28
 - installing Management Console and 181
 - restoring 245, 252, 263, 283
 - shared configurations and 174
 - sharing 164
 - specifying resource 196
 - viewing properties for 164
- formats vii
- Forrester Wave Open Source Business Intelligence report vii
- ftp distributions
 - Information Console 188, 189
 - iServer Integration Technology 202
 - iServer System 188
- ftp site (Actuate) 25
- full installations 188

G

- gadgets vii
- generating
 - charts 37
 - reports 11, 38
 - temporary documents 11
- getJDBCMinorVersion method 16
- graphs. *See* charts

H

- heap size 196
- help 206
 - See also* documentation
- help files 35, 206, 209
- help topics xi
- host names
 - displaying 167
- HTML documentation vii
- hyperthreading 229

I

- I/O operations 11
 - See also* input; output
- IBM DB2 databases. *See* DB2 databases
- ICU libraries 23
- ICU_DATA environment variable 23

- IDAPI applications 10, 15, 16
- impdp command line utility 252, 253
- IMPORT_DATA parameter 145
- importing Encyclopedia volumes 141, 150
- incomplete files 237
- indexed searches 5
- Information Console
 - assigning CPUs to 196
 - caching web pages for 198
 - configuring clusters and 13
 - connecting to 197
 - customizing 196, 197
 - deploying 189, 195–198
 - installing 188, 189, 195
 - logging activities for 197
 - logging in to 36, 199
 - preventing cache conflicts for 198
 - running iServer processes and 8
 - running stand-alone servers and 9
 - setting application server port for 189
 - setting context paths for 191, 196
 - setting default Encyclopedia for 197
 - starting 199
 - testing installations for 198
- Information Delivery API 15, 16
- information delivery solutions vii
- information object cache 11
- information object files 5
- information objects 217
- INITIALIZE_DATA parameter 145, 153
- in-place upgrades 70, 75, 89, 97
- installation
 - alternate metadata databases 42, 58
 - cache conflicts and 198
 - cluster nodes 162, 172–184
 - Configuration Console 174, 188
 - DB2 databases 50, 51, 58
 - disk space and 37
 - domain servers and 23
 - Information Console 188, 189, 195
 - iServer 6, 22, 25–36, 42, 172
 - iServer Integration Technology 188, 202–205
 - iServer System 6, 188
 - iServer System components 27, 28
 - Java Runtime Environment 189
 - Java Software Development Kit 37, 189

- license keys 220
 - licensing prompts and 221, 222
 - Management Console 179, 188
 - metadata databases 4, 5, 29
 - online documentation 35, 206
 - Open Security application 28, 38
 - PDF documentation 35
 - sample Encyclopedia volume 17
 - sample reports 17
 - SQL Server databases 47, 60
 - testing 38, 198
 - upgrades and 70, 72, 93
 - installation directories 28, 36
 - installation guides 16
 - installation options 28
 - installation prerequisites 22, 42, 171
 - installation programs 16
 - installation wizard 189
 - Integration processes 228
 - Integration service 9, 11
 - IntegrationService element 228
 - Interactive Viewer 217
 - Interactive Viewer Option 217
 - international character sets 17
 - IP addresses
 - displaying 167
 - ipconfig command 167
 - iportal context path 191
 - iServer
 - accessing functionality 216
 - accessing online help for 209
 - binding to multiple-core CPUs 228
 - binding to processors 225, 226
 - changing CPU binding and 230
 - checking bound processors for 229–231
 - creating data objects and 5
 - deploying 6, 13, 14, 162, 225
 - getting machine ID for 222
 - initializing 25
 - installing 6, 22, 25–36, 42, 172
 - optimizing performance for 14, 163
 - running 14, 162
 - setting up user accounts for 22, 23, 24, 42
 - starting 9
 - storing metadata for 4
 - testing new releases for 38–40
 - upgrade options for 6
 - upgrading 36, 70, 72, 89, 93
 - validating CPU binding and 230
 - viewing licensing information for 219, 220
 - iServer 11 service 31
 - iServer directory 28
 - iServer Information page 193
 - iServer Integration Technology 16, 188, 202–205
 - iServer processes
 - binding to CPUs 225–232
 - binding to Windows systems 225
 - creating resource groups and 37
 - domain servers and 23
 - multiple-core CPUs and 228
 - running 8, 162, 225
 - starting 225
 - upgrades and 73, 94
 - verifying CPU binding for 227
 - iServer services
 - See also* specific service
 - running multiple 8
 - setting user profiles and 31
 - shutting down 73, 94
 - specifying user accounts for 22
 - iServer servlet container 8
 - iServer System
 - administering 15
 - changing CPU binding and 231
 - configuring 6
 - creating dedicated account for 23
 - deploying new releases and 38
 - downloading files for 25, 188
 - installation options for 6, 28
 - installation requirements for 188
 - licensing options for 216, 217, 225
 - optimizing 13, 162
 - preventing data loss for 7, 24
 - storing metadata for. *See* system databases; system metadata
 - third-party databases supported 4
 - upgrading 72
 - iServer System components 27, 28, 35
 - iServer System names 29, 57
- ## J
- J2EE installations 195

- JAR files 38
- Java developer guide 16
- Java Factory service. *See* Factory service
- Java heap size 196
- Java Object Interface 38
- Java Report Server Security Extension 16
- Java Runtime Environment 16, 189
- Java Server Pages. *See* JSPs
- Java Software Development Kit 37, 189
- Java Virtual Machines. *See* JVM
- JDBC drivers 11, 16, 17
- jdbcCompliant method 16
- JDK files 37, 189
- jobs 11
- JRE environment 16, 189
- JSPs 198
- JVM libraries 38

K

- keys. *See* license keys

L

- language settings 31
- LDAP configuration file 38
- LDAP security database 38
- LDAP servers 28, 38
- libraries 23, 38
- License File Details page 30
- license file names 223
- license files 13, 30
 - obtaining 223
- license key file locations 221
- license key file names 220
- license key files 221
- license key installations 220
- license keys 219, 223, 225
- licensed CPUs 230
 - See also* CPU-based licenses
- licenses 216
- licensing information 219
- licensing options 16, 216, 217, 225
- licensing support (Actuate) 223
- links (documentation) xi
- Linux systems 16, 17
- load balancing (clusters) 14, 162
- local help 209

- Local Security Settings dialog 24
- Local Users and Groups dialog 23
- Locale Information page 31
- Localemap.xml 196
- locales
 - encoding character sets for 17
 - selecting 31
 - setting default 31, 196
 - updating resource files for 206
- localhost parameter 197
- Localization and Online Documentation
 - Update tool 206
- log files 28, 197
- Log on as a service privilege message 32, 64, 178, 193
- LOG_FILE_COUNT parameter 145, 159
- LOG_FILE_LOCATION parameter 197
- LOG_FILE_NAME parameter 145, 159
- LOG_FILE_SIZE parameter 145, 159
- LOG_FOLDER parameter 141, 145, 158, 160
- logging applications 15
- logging levels
 - Data Store Administrator 145
 - Data Store Upgrader 159
- logical CPUs 228
- logical processors 225, 229
- losing data 7, 24

M

- machine IDs 221, 224
- machine information 221
- maintenance customers 223
- Management Console
 - adding alternate databases and 64
 - configuring clusters and 13
 - installing 179, 188
 - licensing and 216
 - managing Encyclopedia and 16
 - running iServer processes and 8
 - running stand-alone servers and 9
 - setting application container port for 180
 - setting user names for 32
 - starting 35
- manual import operations 141
- manual in-place upgrades 70, 89
- manual side-by-side upgrades 70

- manual volume upgrade operations 157
- manuals. *See* documentation
- master cluster nodes 8
- master index file xi
- message routing (requests) 10, 13
- metadata
 - autoarchiving and 238
 - backing up 24, 54, 237
 - changing 5
 - defined 236
 - preventing loss of 7, 24
 - restoring 253, 268, 270
 - storing 4, 11, 155
 - upgrades and 71
- metadata databases
 - See also* third-party databases
 - changing 4
 - customizing connections for 16
 - editing 5
 - incomplete files and 237
 - installing 4, 5, 29, 42, 58
 - previous releases and 4
 - selecting type 57
 - setting up iServer environments for 42
 - sharing 5
- Microsoft Excel spreadsheets 218
- Microsoft SQL Server databases. *See* SQL Server databases
- Microsoft Windows systems. *See* Windows systems
- migration
 - customizing metadata databases and 5
 - installing new releases and 25, 38
 - upgrading and 6, 36, 70
- missing product files 188
- missing system schemas 72
- multicast configurations 8
- multiple-core CPU binding 228
- multiple Encyclopedia volumes 218
- multiple upgrade installations 37
- Multi-Tenant Option 218
- multi-threaded CPUs 225

N

- named-user licenses 31, 216
- names

- obtaining host 167
- naming
 - Encyclopedia volumes 29, 146
 - iServer System 29, 57
 - schema owners 43, 44
 - system schemas 29, 43
 - users 43
 - volume schemas 43, 45, 49, 51
- naming conventions 43, 57
- network administrators. *See* administrators
- network cards 222, 224
- Network File Systems (NFS) 11
- networked environments 11, 13, 162, 188
- NEW_SCHEMA_NAME parameter 146, 150
- NEW_VOLUME_NAME parameter 146, 153
- newinformationconsole.war 198
- NobleNet Portmapper processes 73, 75, 94, 96
- NobleNet Portmapper service 73, 75, 94, 96
- node keys 221
- node-key licensing 219, 220, 221
- nodes. *See* cluster nodes
- notifications 230, 231

O

- obsolete features 6
- online documentation
 - accessing vii–xi
 - administering iServer System and 16
 - conventions for xii
 - installing 35, 206
 - managing Encyclopedia and 7
 - upgrades and 87, 105
- online documentation update tool 206
- online documentation, downloading xi
- online help 209
 - See also* online documentation
- Open Security applications 28, 38
- open source projects vii
- opening
 - Information Console 36, 199
 - Management Console 35
- operating system user accounts 50
- operating systems 17, 225
- optimizing iServer System 13, 162
- options (installation) 28

- options (licensing) 216, 217, 225
- Oracle Data Pump Export utility 250–252
- Oracle Data Pump Import utility 252–254
- Oracle databases
 - accessing documentation for 7
 - connecting to 46
 - creating schema owners for 45, 46
 - exporting metadata to 148
 - installing 58
 - managing Encyclopedia and 5
 - running volume backup and restore
 - operations with 249, 250, 252
 - setting port number for 58
 - setting up iServer environments for 45–47
 - storing metadata and 4
 - taking Encyclopedia offline for 252
 - taking Encyclopedia online for 253
 - upgrades and 70, 72, 93
- Oracle service names 58
- ORACLE_TNS_NAMES_FILE parameter 146, 160
- output 11
- output formats vii, 11
- overwriting product files 6

P

- Page Level Security Option 218
- page-level security 218
- parameters
 - configuring clusters and 13, 162
 - deploying Information Console and 196
 - encoding volume databases and 17
 - setting Data Store Administrator 141
 - upgrading Encyclopedia and 157, 158
- partition names 113
- partition paths 113
- partitions (disk) 237
- partitions (volume)
 - setting primary 123
- Partitions page 124, 139
- passwords
 - changing 32
 - creating schema owners and 147
 - installing alternate databases and 64
 - installing cluster nodes and 178
 - installing Information Console and 192

- installing iServer and 25
- prompting for schema 147
- setting database 29
- starting console applications and 32
- starting iServer service and 31
- patches 9, 14, 163
- PATH environment variable 108, 135, 155
- paths 28, 113, 236
- PDF documentation vii, x, xi, 35, 206
- performance
 - caching data and 11
 - cluster configurations and 14, 163
 - CPU binding and 228
 - file input/output operations and 11
 - iServer processes and 9
 - upgrades and 37
- performance monitoring extensions 15
- permanent licenses 221
- permissions. *See* privileges
- pg_dump command line options 241
- pg_dump utility 239, 241
- pg_restore command line options 248
- pg_restore utility 239, 248
- pgAdmin database administration tool
 - backing up Encyclopedia and 239–242
 - restoring Encyclopedia and 244, 245
 - uninstalling 36
 - upgrades and 83
- pgAdmin III Setup dialog 83
- pgAdmin III Setup Wizard 86
- pgpass files 241
- ping command 168
- plpgsql language 43
- POPULATE_SCHEMA parameter 146, 153
- POPULATE_SCHEMA_PHASE_ONE
 - parameter 146
- POPULATE_SCHEMA_PHASE_TWO
 - parameter 146
- port numbers
 - See also* ports
- Portmapper service 75, 96
- ports
 - application servers 189
 - cluster nodes 178
 - DB2 databases 58
 - Information Console 197
 - multiple product releases and 37

- Oracle databases 58
- PostgreSQL databases 29, 60
- SQL Server databases 60
- postgres database directory 36
- PostgreSQL command line utilities 239
- PostgreSQL Database Information page 29
- PostgreSQL databases
 - accessing documentation for 7
 - adding 11
 - backing up 239
 - backing up data folders for 242
 - cluster configurations and 11
 - connecting to 44
 - creating Encyclopedia for 154
 - creating schema owners for 43, 44
 - exporting metadata to 149
 - file I/O operations and 11
 - importing Encyclopedia and 150
 - installing iServer and 22, 25, 60
 - managing Encyclopedia and 5
 - naming schemas for 44, 45
 - running volume backup and restore
 - operations with 239, 240, 244
 - setting passwords for 29
 - setting port number for 29, 60
 - setting up iServer environments for 43–45
 - specifying default locations for 28
 - storing metadata and 4
 - taking Encyclopedia offline for 244
 - taking Encyclopedia online for 249
 - uninstalling 36
 - upgrades and 70, 72, 74, 93, 95
- postgresql directory 242
- PostgreSQL for BIRT service 31, 73, 74, 94, 95
- PostgreSQL server clusters 5
- PostgreSQL SQL language 43
- PostgreSQL superusers 29, 147
- Power User accounts 22, 23, 31
- prerequisites (installation) 22, 42, 171
- previous releases 6
- primary volume partition 123
- printers 14, 22
- privileges
 - connecting to DB2 databases and 50, 51
 - connecting to Oracle databases and 45, 46
 - connecting to PostgreSQL databases and 43, 44

- connecting to SQL Server databases and 47, 48, 49
- creating system schemas and 157
- creating user accounts and 22, 23
- installing Information Console and 189
- logging in to iServer and 24
- running multiple schemas and 43
- Process Management Daemon
 - CPU binding and 225, 228
 - running cluster configurations and 13, 14, 162
 - running iServer processes and 8, 10
 - starting encycsrvr processes and 230, 231
- Processor Affinity 227
- processor IDs 226, 228
- ProcessorAffinity element 228
- processors. *See* CPUs
- product files 6, 25, 188
- product updates 206
- production environments 40
- production staging areas 39
- profiles (user) 31
- program files. *See* product files
- PROMPT_FOR_PASSWORDS parameter 147, 160
- properties 141
 - file sharing 164
- properties files 91, 117, 140, 196
- publishing reports 217
- Purge deleted files time property 239
- purging report files 238

Q

- queries 11, 13, 52
- Query Option (licensing) 217

R

- RDBMS backup utilities 237
- RDBMS database configurations 11
- RDBMS databases 4, 237
 - See also* third-party databases
- RDBMS documentation 7
- RDBMS environments 4, 14, 43
- RDBMS tools 6, 7
- ReadMe files 195, 205
- rebinding encycsrvr11 processes 231

- recovery operations 6, 7, 16
- refreshes 198
- registered users 31
- relational databases. *See* databases
- release notes xi
- remote procedure calls 10
- remote servers 188
- removing
 - database schemas 144
 - installation directories 36
 - product files. *See* uninstalling
- renaming
 - Encyclopedia volumes 150
 - license files 223
- report design files 5
- report designers vii
- report designs 11, 39, 217, 218, 219
- report document files 5, 219
- report document generation service 11
- Report Encyclopedia. *See* Encyclopedia
- report files
 - accessing 22
 - archiving 238
 - backing up 75, 97, 237
 - downloading 11
 - purging 238
 - running processes and 73, 94
 - sharing 164
 - specifying default directories for 28
 - storing 5
 - upgrading iServer and 37, 39, 72
- report object executable files 37
- Report Server Security Extension 15, 16
- report templates 196
- report viewers 217
- reporting applications. *See* applications
- reporting server. *See* iServer
- reporting services. *See* specific iServer service
- ReportingService element 228
- reports
 - displaying 11, 217
 - generating 11, 38
 - installing sample 17
 - publishing 217
 - running 218
- repositories
 - See also* Encyclopedia volumes

- requests
 - distributing 8, 13
 - installing cluster nodes and 178
 - remote procedure calls and 10
 - running cluster configurations and 14, 162
- resource file updates 206
- resource folders 196
- resource groups 37
- resources
 - accessing 196
 - creating cluster nodes and 13, 14, 162
 - installing Information Console and 188
 - installing iServer and 22
 - obtaining licenses and 217
- restoring Encyclopedia volumes
 - DB2 databases and 273, 283
 - PostgreSQL databases and 239, 244
 - SQL Server databases and 254, 262
- restricting iServer processes 225
- rollbacks 71
- RSSE applications 15, 16
- running
 - Data Store Upgrader 90, 92
 - encycsrvr11 processes 225, 230, 231
 - iServer processes 8, 23, 225
 - iServer services 8, 22
 - jobs 11
 - Open Security applications 38
 - pg_dump utility 241
 - pg_restore utility 248
 - queries 11, 13
 - report designs 217, 218
 - reports 218
 - spreadsheet reports 218
 - System Data Store Administrator 155

S

- sample reports 17
- sample volume 17
- saving
 - previous releases 75, 97
 - properties files 141
- scalability vii
- scheduling iServer installations 39
- schema names 29, 90, 146, 147
- schema owners

- creating 42
- installing DB2 databases and 50, 51
- installing iServer and 5
- installing Oracle databases and 45, 46
- installing PostgreSQL databases and 43, 44
- installing SQL Server databases and 47, 48
- setting passwords for 29, 147
- SCHEMA_FILE_NAME parameter 141, 158
- SCHEMA_NAME parameter 147, 151, 158
- SCHEMA_PASSWORD parameter 147, 160
- schemas
 - accessing 4
 - adding Encyclopedia to 153, 154
 - assigning privileges 43
 - backing up 7, 24, 155, 241
 - configuring system data store as 155
 - creating 25, 42
 - customizing metadata databases and 5
 - depopulating 144
 - dropping 144
 - importing volume data to 150
 - initializing 25
 - installing DB2 databases and 51
 - installing Oracle databases and 45, 46
 - installing PostgreSQL databases and 44, 45
 - installing SQL Server databases and 48, 49
 - naming 29, 43, 147
 - populating 146, 153
 - preventing data loss and 7
 - removing Encyclopedia from 152
 - restoring 155, 157
 - setting passwords for 147
 - storing metadata and 4
 - upgrades and 71
 - user-defined functions and 53
- SCRIPT_HOME parameter 141, 158
- scripts 237
- search results 218
- searching online documentation xi
- security vii, 31, 218
- security application 38
- security database 38
- security systems 28
- server templates 13, 162
- SERVER_DEFAULT parameter 197
- servers
 - See also* iServer
 - changing cluster machines and 224
 - creating clusters for 168
 - creating shared folders for 164
 - deploying Information Console to 195, 196
 - exceeding CPU licenses for 231
 - getting host names for 167
 - installing iServer for external 28
 - installing stand-alone 28
 - integrating Actuate products with vii, 195
 - managing volume metadata and 5
 - preventing cache conflicts for 198
 - running as clusters 13, 14, 174
 - running iServer processes and 23
 - setting port numbers for 189
 - shutting down iServer and 9
- service names 58
- service packs 236
- service privileges 24
- services. *See* specific iServer service
- ServIntTech directory 202
- servlet container 8
- set_tools_environment.bat file 91
- Setup Type page 28
- shared database schemas 153
- shared folders 164
- shared licenses 219, 221
- sharing metadata databases 5
- shortcuts 35, 87, 105
- shutting down. *See* stopping
- side-by-side upgrades 70
- SIDs (system identifiers) 58
- single-point media failures 250
- single-point node failure 13, 162
- SmartSheet Security Option 218
- snapshots 237
- SOAP processor 10
- SOAP-based messaging 8, 10
- Solaris systems 17
- Specify Profiles page 31, 79, 101
- Specify Windows Account Information page 192
- spreadsheet reports 218
- SQL languages 43
- SQL queries 11, 13, 52
- SQL scripts 43, 45, 47, 50
- SQL Server databases
 - accessing documentation for 7

- backing up transaction logs for 264, 266
 - connecting to 49
 - creating schema owners for 47, 48
 - creating users for 48, 49
 - exporting metadata to 149
 - installing 47, 55–67
 - managing Encyclopedia and 5
 - naming schemas for 49
 - running volume backup and restore operations with 254, 262
 - setting port number for 60
 - setting up iServer environments for 47–49
 - storing metadata and 4
 - taking Encyclopedia offline for 263
 - taking Encyclopedia online for 273
 - SQL Server Management Studio
 - backing up Encyclopedia and 254–262
 - restoring Encyclopedia and 262–273
 - SQL*Loader utility 148
 - sqlcmd utility 260, 270
 - sqlplus utility 250
 - Squirrel Data Exporter
 - configuring 140
 - migrating iServer installations and 6
 - setting up export environments and 151
 - Squirrel database 4
 - stand-alone configurations 28
 - stand-alone servers
 - creating clusters for 168
 - creating shared folders for 164
 - starting
 - Data Store Upgrader 90
 - DB2 command line processor 280
 - DB2 Control Center 275
 - Encyclopedia volumes 9, 124, 139
 - encycsrvr11 processes 230
 - Information Console 36, 199
 - iServer 9
 - iServer 11 service 32
 - iServer processes 225
 - Management Console 35
 - pg_dump utility 241
 - pg_restore utility 248
 - SQL Server Management Studio 255
 - stopping
 - iServer services 73, 94
 - NobleNet Portmapper processes 75, 96
 - Portmapper services 75, 96
 - PostgreSQL for BIRT service 74, 95
 - SUPERUSER parameter 147
 - superuser password 25, 29
 - SUPERUSER_PASSWORD parameter 147
 - superusers 29, 147
 - Supported Products Matrix 195
 - supported products page xii
 - Switch Help Location command 209
 - synchronous Factory service 11
 - syntax conventions (documentation) xii
 - system. *See* iServer System
 - system administrators. *See* administrators
 - System Configuration Password page 32, 80, 101
 - System Data Store Administrator
 - recreating system schemas and 157
 - restoring system schemas and 72
 - running 155
 - setting PATH variable for 155
 - setting properties for 155, 156
 - system data stores 155
 - system databases 4, 5, 43, 50
 - See also* third-party databases
 - system failover operations 7
 - system identifiers 58
 - system metadata
 - backing up 24, 54
 - preventing loss of 7
 - storing 4, 11, 155
 - system names 29, 57
 - system schema names 29
 - system schema owner 43, 45, 47, 50
 - system schema owner accounts 29
 - system schema passwords 29
 - system schemas
 - See also* database schemas
 - backing up 72, 155
 - creating 25, 44, 49, 51
 - restoring 155
 - storing metadata and 4
 - system variables 226
 - SystemDataStoreAdministrator class 156
- ## T
- table of contents (documentation) xi

- tables 5
- TABLESPACE_LOCATION parameter 147
- TABLESPACE_NAME parameter 147
- taking Encyclopedia online 9, 124, 139
- TEMP_FOLDER_LOCATION parameter 197
- templates 196
- temporary directories 197
- temporary documents 11
- temporary files 197
- temporary license keys 220
- temporary licenses 221
- testing
 - connections 168
 - desktop products 39
 - Information Console installations 198
 - new releases 38–40
- text files 223
- 3rd-Party Database Information page 58
- third-party command line utilities 239
- third-party database schemas. *See* database schemas
- third-party databases
 - See also* specific database type
 - accessing documentation for 7
 - changing tables in 5
 - creating indexes for 5
 - customizing connections for 16
 - importing Encyclopedia and 150
 - installing iServer and 22, 42
 - managing Encyclopedia and 4, 6
 - preventing data loss for 24
 - running iServer clusters and 8, 11, 14
 - selecting 57
 - setting port numbers for 58, 60
 - storing metadata and 4
 - uninstalling 36
 - upgrades and 71
- third-party deployment tools 196
- third-party RDBMS tools 6
- threads 196
- time zones 31, 196
- TIME_ZONE parameter 147
- TimeZones.xml 196
- TNS services 58
- Tomcat servers 189, 191
- transaction logs 263, 264, 266
- Transact-SQL scripts 47

- transient data stores 197
- transient files 197
- TRANSIENT_STORE_PATH parameter 197
- Transparent Network Substrate services 58
- transport protocol (requests) 10, 13, 162
- troubleshooting 237
- typographic conventions (documentation) xii

U

- uninstalling
 - pgAdmin administration tool 36
 - third-party databases 36
- UNIX systems 16, 17
- Update tool 206
- updates (documentation) viii
- updates (product) 206
- upgrade options 70
- upgrade utilities 140
- upgrade_encyclopedia_data_store.bat file 91
- upgrade_encyclopedia_data_store.properties file 91
- Upgrader utility
 - configuring 140
 - running 90, 91
 - setting properties for 157, 158
- upgrades
 - application programming interfaces and 17
 - creating backups and 71
 - customizing metadata databases and 5
 - disk space and 37
 - Encyclopedia volumes 89
 - file structure changes and 89
 - installing metadata databases and 4
 - installing product files and 72
 - iServer 36, 70, 72, 89, 93
 - licensing options and 221
 - maintaining version consistency for 25
 - overwriting previous releases and 6
 - pgAdmin administration tool 83
 - rolling back 71
 - running processes and 73, 94
- upgrdeds.bat file 91
- URLs
 - Information Console 197, 199

- Localization and Online Documentation
 - Update tool 206
- Network File Systems 11
- RDBMS documentation 7
- Supported Products Matrix 195
- usage logging applications 15
- usage reports 17
- user accounts
 - cluster configurations and 23
 - creating 23, 42
 - installing DB2 databases and 50
 - installing Information Console and 192
 - installing iServer and 22
 - managing 16
 - starting iServer 11 service and 31
 - upgrades and 71
- user credentials 5
- user information 15, 38
- user names 31, 32, 43
- user profiles 31
- User Rights Assignments dialog 24
- user tracking reports 17
- users
 - creating DB2 schemas and 50, 51
 - creating Oracle schemas and 46
 - creating PostgreSQL schemas and 44
 - creating SQL Server schemas and 48, 49
 - licensing and 31, 216
 - setting passwords for 31

V

- validating CPU binding 230, 231
- VARCHAR types 50
- VARGRAPHIC types 50
- verifying CPU binding 227
- View processes 228, 231
- View service 8, 11, 37
- viewers 217
- viewing
 - charts 37
 - folder properties 164
 - host names 167
 - IP addresses 167
 - licensing information 219, 220
 - processor IDs 226
 - ReadMe files 195, 205

- reports 11, 217
- ViewingService element 228
- Vista operating systems 22
 - See also* Windows systems
- volume administrators. *See* administrators
- volume data 236
- volume data directories
 - DB2 environments 282, 283
 - Oracle environments 251, 252
 - PostgreSQL environments 242, 245
 - SQL Server environments 261, 263
- volume database connections 16
- volume database encoding 17
- volume databases 4, 5
 - See also* third-party databases
- volume metadata
 - backing up 237, 250, 255, 260, 275, 280
 - changing 5
 - restoring 253, 268, 270, 284, 291
 - storing 4, 11
- volume metadata database 29, 57
- volume names 29, 146, 150, 189
- volume partition names 113
- volume partition paths 113
- volume partitions
 - setting primary 123
- volume schema names 29, 44, 45, 49
- volume schema owner 44, 46, 48, 51
- volume schemas
 - See also* database schemas
 - assigning privileges to multiple 43
 - backing up 241
 - creating 25
 - DB2 databases and 51
 - PostgreSQL databases and 45
 - SQL Server databases and 49
 - storing metadata and 4
- VOLUME_NAME parameter 147
- VolumeImport.properties file 117
- volumes. *See* Encyclopedia volumes

W

- WAR files 189, 195, 196, 197
- watermarks 219
- web administrator privileges 189
- web administrators. *See* administrators

- web applications 198
- web archive files. *See* WAR files
- web browsers 198
- web pages 198, 218
- web servers 196
 - See also* servers
- web services 10
- whitespace (property files) 141
- Windows platforms
 - creating shared folders on 164
 - displaying IP addresses for 167
 - turning off firewalls for 166
- Windows systems
 - binding iServer processes to 225, 226–229
 - collecting machine information for 221
 - configuring firewalls for 189
 - distributing iServer System over 6
 - installation prerequisites for 22, 42
 - installing Configuration Console for 188
 - installing Information Console for 188, 189, 195, 198
 - installing iServer clusters for 171, 172–184
 - installing iServer for 26
 - installing iServer Integration Technology for 202
 - installing JDK files for 37, 189
 - installing JRE files for 189
 - installing localization and documentation files for 206
 - installing Management Console for 188
 - running iServer processes and 23, 73, 94
 - setting classpaths for 38
 - setting default locale for 31
 - setting up user accounts for 22, 23, 24, 42
 - starting iServer service for 31
 - testing new releases for 38–40
 - upgrading iServer for 36, 70, 72, 89, 93
- Work Unit Licenses 217

X

- XML files 223