One Design
One Server
One User Experience

Installing BIRT iServer for Linux and UNIX

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

Document No. 170215-2-430343 February 15, 2017

Contents

Part 2
Installing

Chapter 2
Installing BIRT iServer 21

Chapter 3
Installing BIRT iServer using an alternative database 47

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: ■ 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	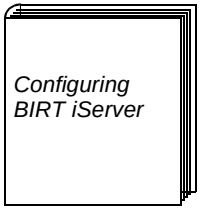

Table I-1 BIRT iServer documentation (continued)

For information about this topic	See the following resource
Managing an Encyclopedia Volume Use Management Console to: ■ 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	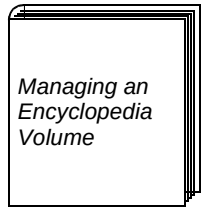
Information Console Developer Guide ■ Overview of Information Console concepts and web applications ■ Using, customizing, and configuring the Deployment Kit ■ Using code components for JSPs, URL parameters, JavaScript files, Java servlets, Java Beans, and security facilities	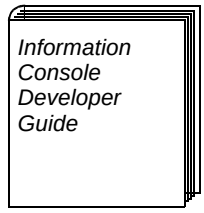
Using BIRT iServer Integration Technology ■ Overview of Actuate web services and SOAP messaging ■ Managing an Encyclopedia volume ■ Developing API applications using Java or .NET ■ Using Java Report Server Security Extension (RSSE) APIs ■ Using logging, performance monitoring, and archiving features ■ Customizing the Actuate software installation process ■ Actuate Information Delivery API operations and data types reference	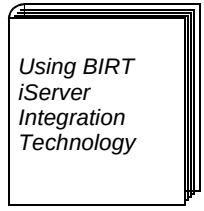

(continues)

Table I-1 BIRT iServer documentation (continued)

For information about this topic	See the following resource
Using Information Console ■ Overview of Information Console concepts and online reporting ■ Accessing and managing files and folders; running designs	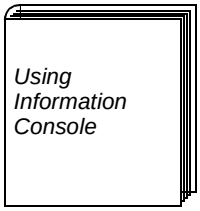
Using Actuate JavaScript API ■ Overview of programming with Actuate JavaScript ■ Creating custom pages using Actuate JavaScript ■ Reference for BIRT JavaScript classes and methods	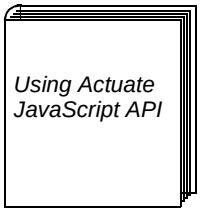
Deploying to a BIRT iServer System Describes how to deploy designs and information objects to iServer	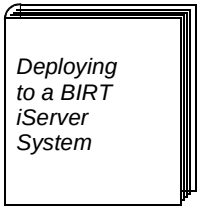
Actuate Glossary Definitions of product terminology	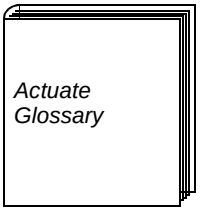
Adobe Acrobat Catalog A utility that can search all the documents in the Actuate manuals directory	

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 information on a particular topic in Actuate documentation.

Obtaining late-breaking information and documentation updates

The release notes contain late-breaking news about Actuate products and features. The release notes are available on 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	`Dim Text1 As String`
File names	Initial capital letter, except where file names are case-sensitive	Detail.roi
Key combination	A + sign between keys means to press both keys at the same time	Ctrl+Shift
Menu items	Capitalized, no bold	File
Submenu items	Separated from the main menu item with a small arrow	File➤New
User input or user response	Monospace	`M*16*`
User input in XML and Java code	Monospace italics	`chkjava.exe` *`cab_name`*`.cab`

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

About *Installing BIRT iServer for Linux and UNIX*

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

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

- *Chapter 3. Installing BIRT iServer using an alternative database.* Describes how to install BIRT iServer using an alternative RDBMS, such as a pre-existing PostgreSQL or Oracle RDBMS, in a Linux and UNIX environment.
- *Chapter 4. Upgrading BIRT iServer.* Describes how to upgrade BIRT iServer in a Linux and UNIX environment.
- *Chapter 5. Installing a BIRT iServer cluster.* Describes how to install a BIRT iServer cluster node in Linux and UNIX.
- *Chapter 6. Installing Information Console.* Describes how to install Actuate Information Console in Linux and UNIX.
- *Chapter 7. Installing iServer Integration Technology and documentation.* Describes how to install BIRT iServer Integration Technology and Documentation in Linux and UNIX.
- *Part 3. Licensing.* Describes the licensing for BIRT iServer.
- *Chapter 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

Chapter

1

Understanding Actuate BIRT iServer architecture

This chapter contains the following topics:

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

Understanding BIRT iServer architecture

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

Actuate also adapted iServer to support alternative, customizable, third-party database installations. In Release 11 Service Pack 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 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."

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

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

Documentation for a PostgreSQL RDBMS is available at:

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

Documentation for an Oracle RDBMS is available at:

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

Documentation for Microsoft SQL Server RDBMS is available at:

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

Documentation for IBMDB2 RDBMS is available at:

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

The third-party database schemas that contain iServer system and Encyclopedia volume metadata are critical components of BIRT iServer System. To guard against data loss, the database administrator must back up the Encyclopedia volume schemas and all related file data to ensure the recoverability in the event of failure. For more information on backing up an iServer installation, see Chapter 9, "Backing up an Encyclopedia volume."

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."

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.

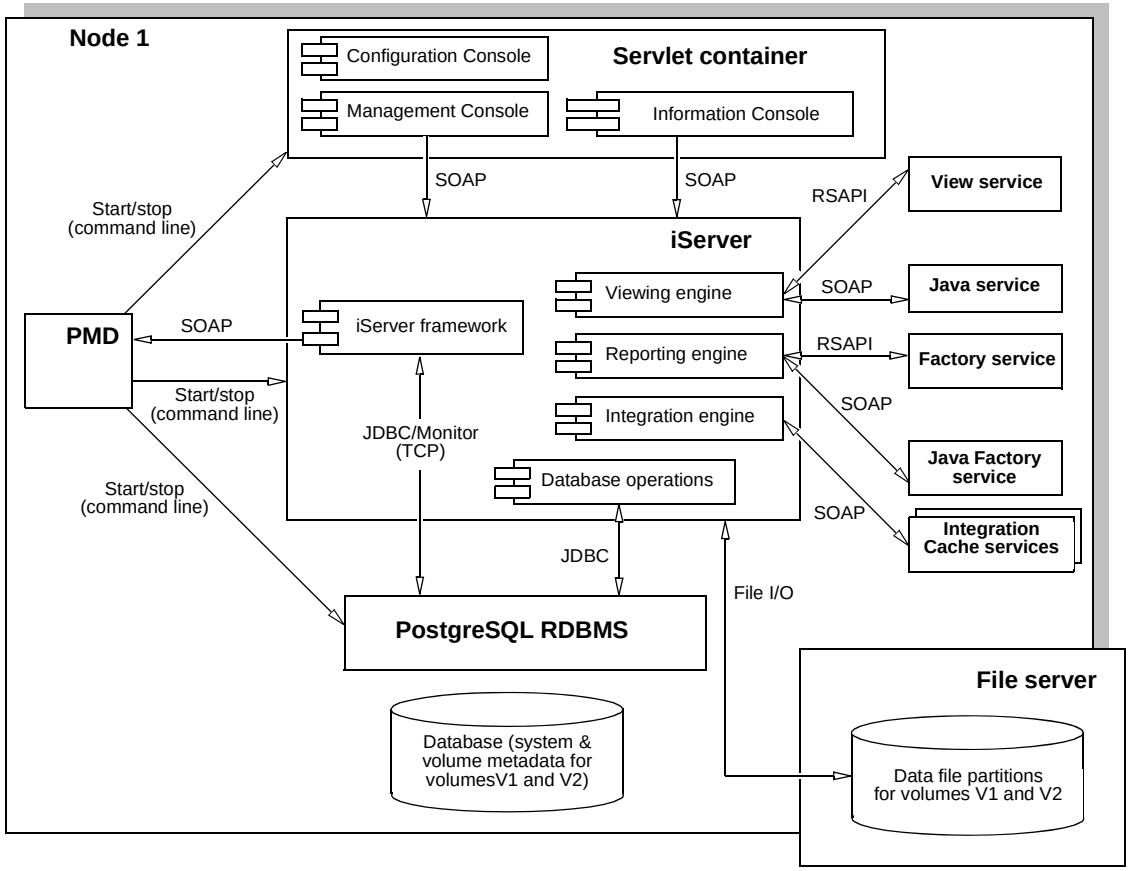

Figure 1-1 iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

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.

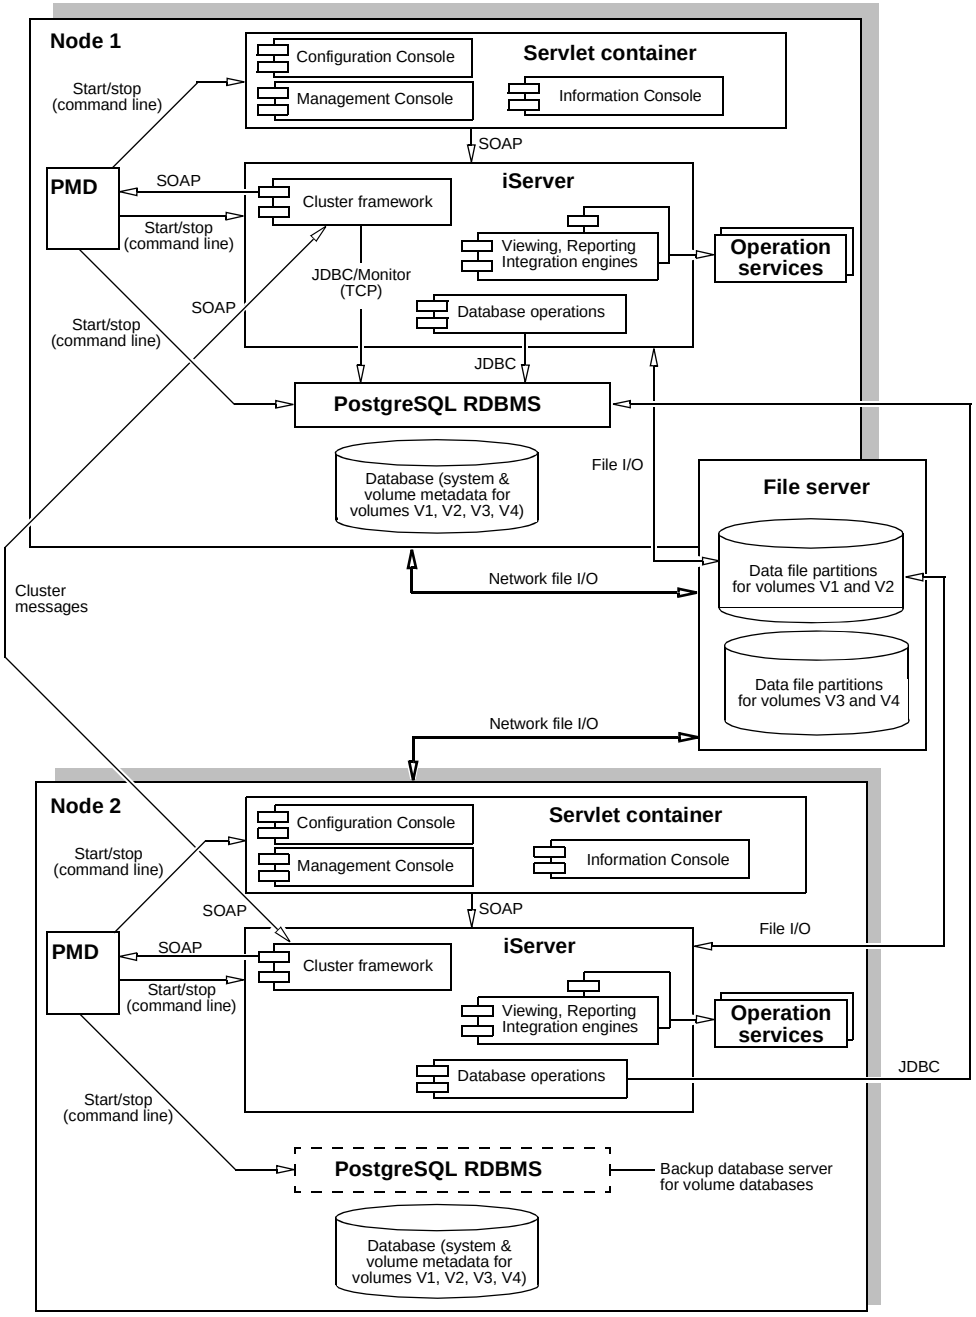

Figure 1-2 iServer RDBMS process architecture for a clustered, two-node, four-volume, OOTB database

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 RBDMS, eliminating internal messages to Encyclopedia services on other nodes.

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

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

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

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

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

For more information about the cluster installation option, see *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 5 supports the following operating systems:

- Windows
- Solaris
- Linux

Part Two

Installing

Chapter

2

Installing BIRT iServer

This chapter discusses the following topics:

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

Preparing to install BIRT iServer

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

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

Creating a dedicated user account for installing and running iServer

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

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

Backing up iServer system and Encyclopedia volume metadata

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

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult 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."

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

If you are a purchasing customer, you can download iServer from OpenText My Support at the following location:

https://support.opentext.com

About X frame buffer

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

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

The Xvfb software uses:

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

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

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

```
# setenv DISPLAY :0.0
```

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

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

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

Installing X frame buffer

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

About libstdc++

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

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

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

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

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

About run levels

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

About Openmotif

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

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

Performing a new installation

Installing a new Release 11 Service Pack 5 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, stand-alone instance of BIRT iServer Release 11 Service Pack 5 (11SP5) in the UNIX and Linux operating system.

How to perform a new installation of BIRT iServer Release 11SP5 in UNIX and Linux

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

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

1. Download the required files. Extract the files.
2. To install the server files, execute the isinstall script:

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

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

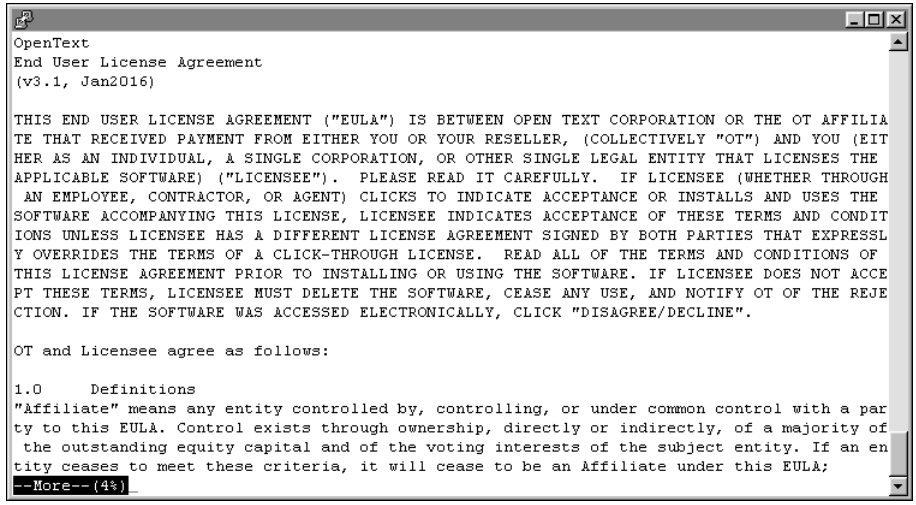

Figure 2-1 Reviewing the license agreement

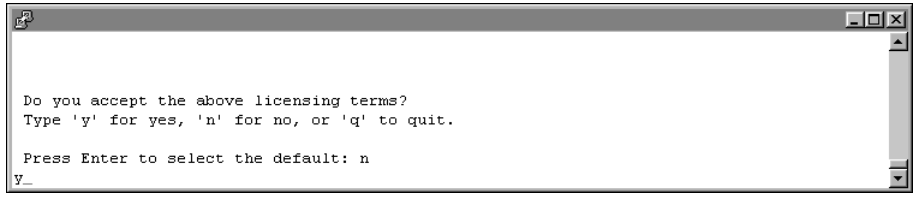

Figure 2-2 Accepting the licensing terms

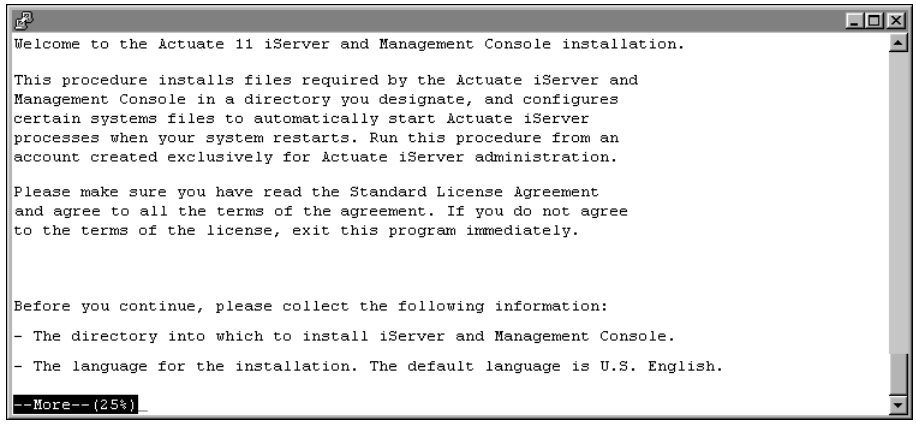

Figure 2-3 Reviewing the introductory information

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

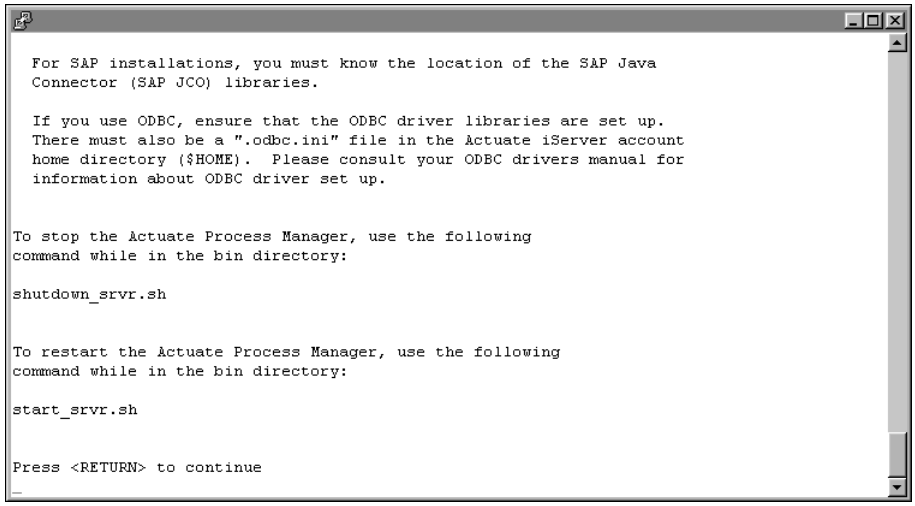

Figure 2-4 Finishing the review of introductory information

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

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

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

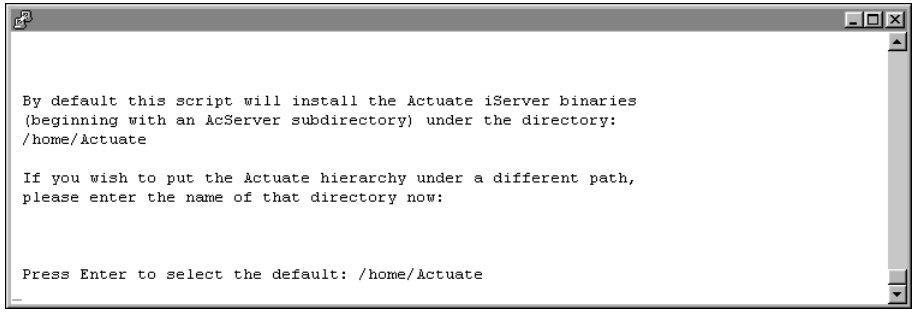

Figure 2-5 Specifying the installation directory

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

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

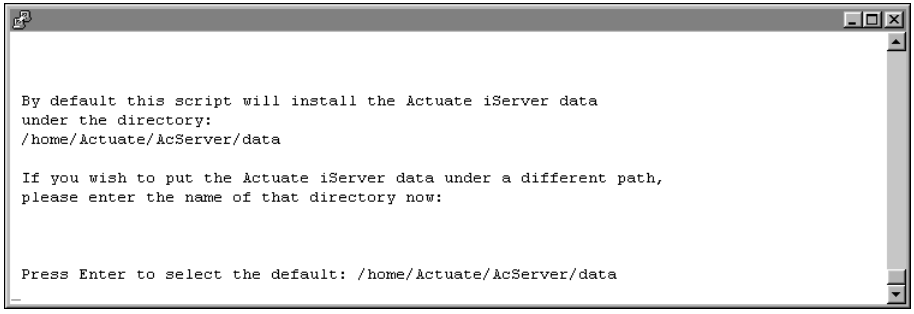

Figure 2-6 Specifying the data installation directory

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

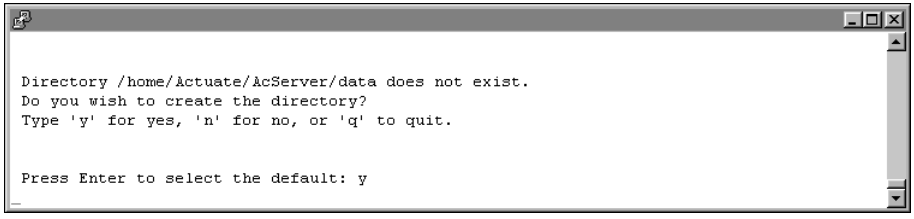

Figure 2-7 Creating the AC_DATA_HOME directory

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

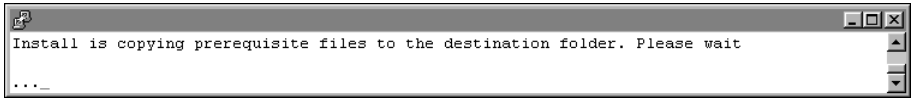

Figure 2-8 Copying prerequisite files

11 Press Enter to accept the default iServer component combination, which includes Management Console, as shown in Figure 2-9. Alternatively, choose a different component combination and press Enter.

12 Press Enter to accept the default stand-alone Server installation, as shown in Figure 2-10. Alternatively, choose a different type of iServer to install. For information on how to install an iServer cluster, see Chapter 5, "Installing a BIRT iServer cluster."

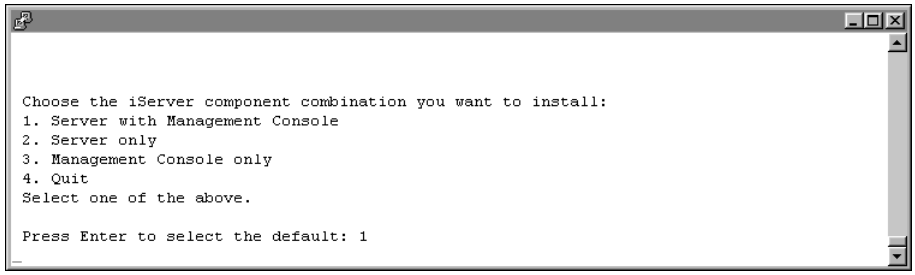

Figure 2-9 Choosing the components to install

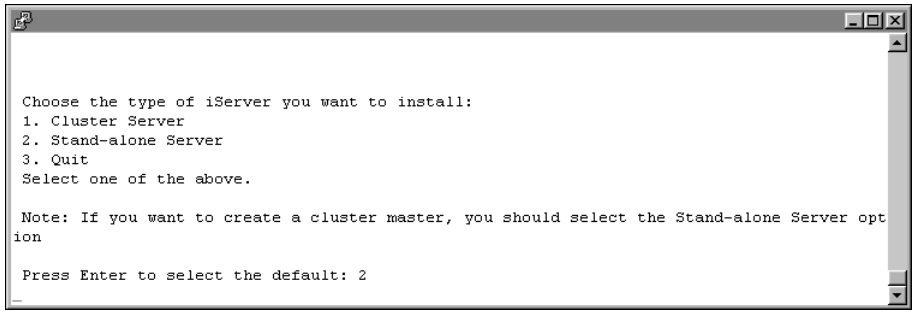

Figure 2-10 Specifying the type of iServer to install

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

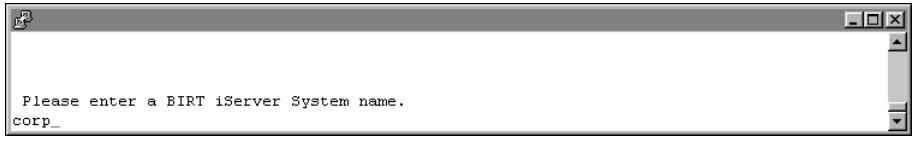

Figure 2-11 Specifying the BIRT iServer System name

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

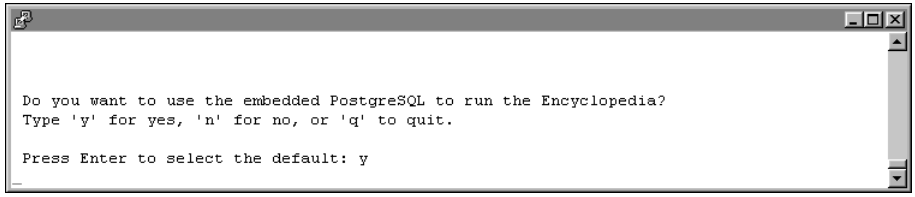

Figure 2-12 Choosing the embedded PostgreSQL

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

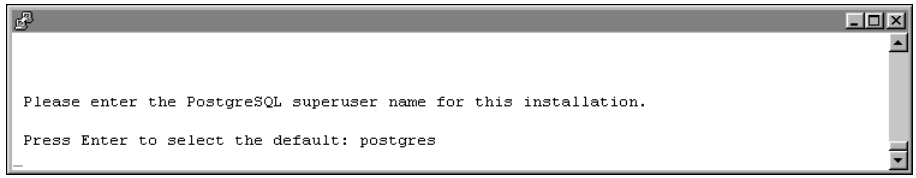

Figure 2-13 Choosing the PostgreSQL superuser name

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

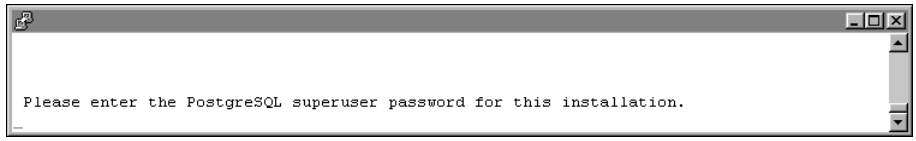

Figure 2-14 Typing the PostgreSQL superuser password

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

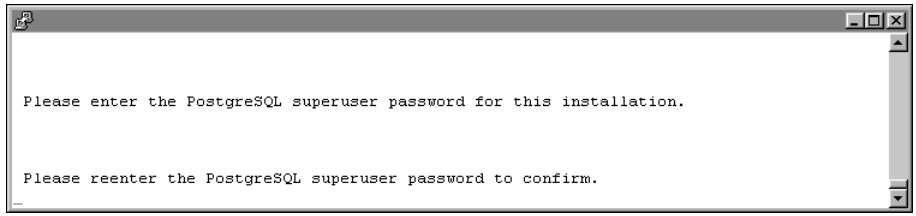

Figure 2-15 Re-entering the PostgreSQL superuser password

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

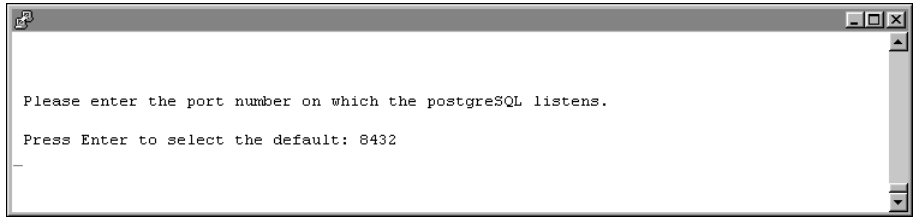

Figure 2-16 Choosing the PostgreSQL port number

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

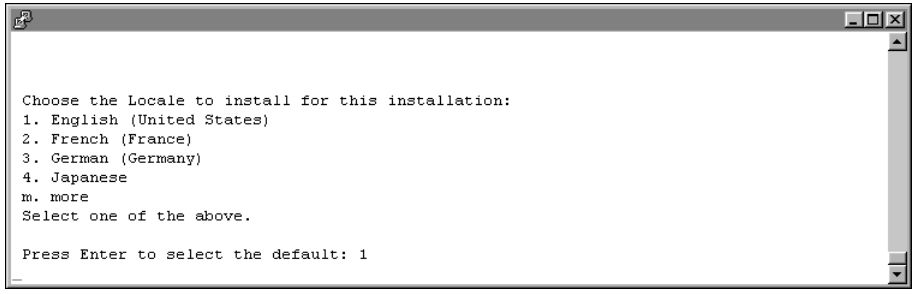

Figure 2-17 Specifying a locale

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

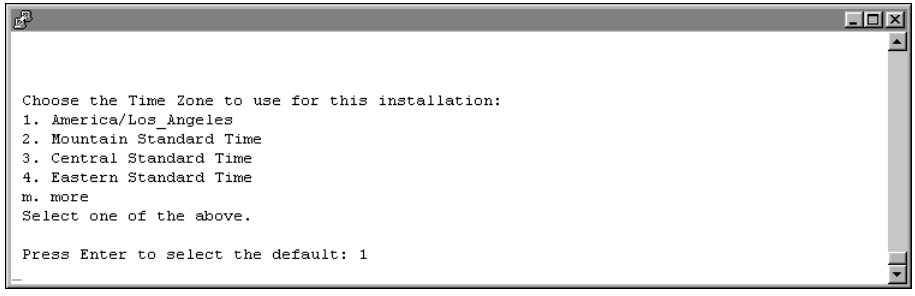

Figure 2-18 Specifying a time zone

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

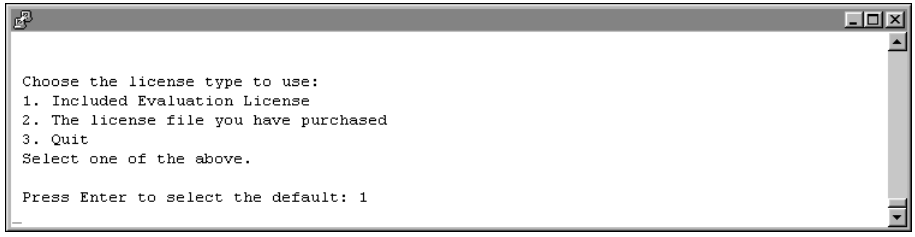

Figure 2-19 Specifying license type

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

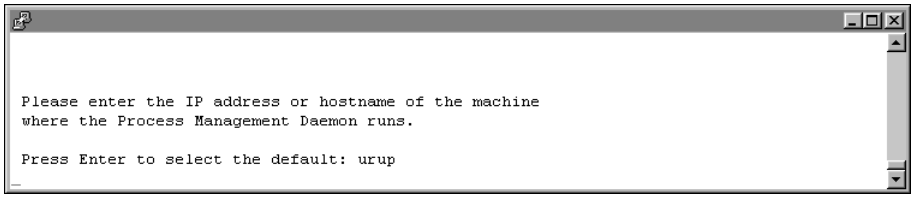

Figure 2-20 Specifying the hostname that Management Console uses to contact the PMD

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

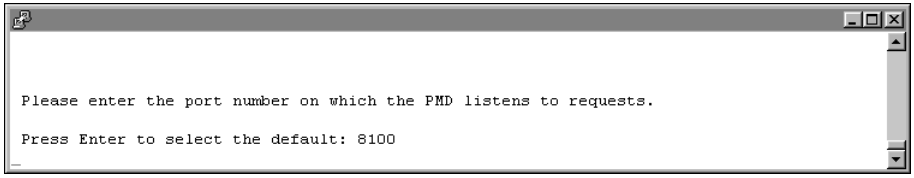

Figure 2-21 Specifying the port number on which the PMD listens

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

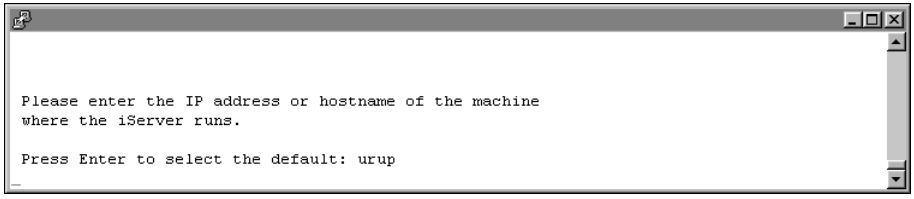

Figure 2-22 Specifying the machine on which the iServer runs

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

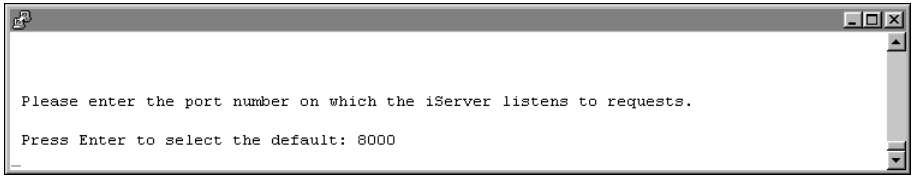

Figure 2-23 Specifying the port number on which the iServer listens

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

Figure 2-24 Specifying the iServer administrator password

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

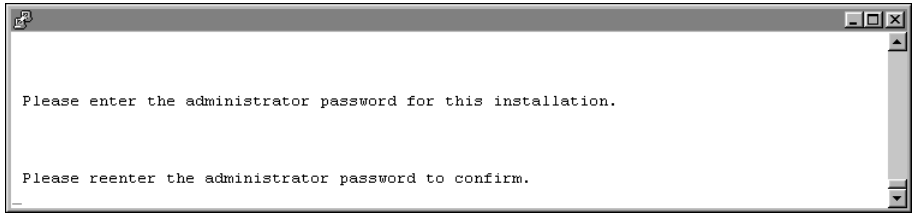

Figure 2-25 Re-entering the iServer administrator password

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

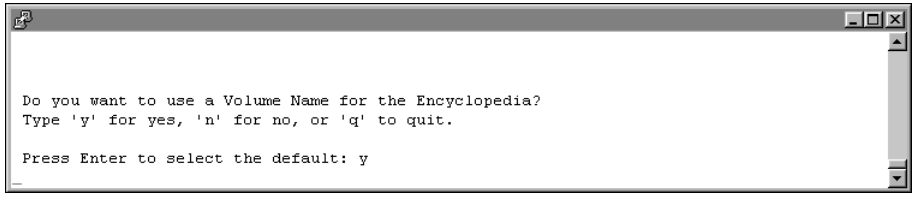

Figure 2-26 Specifying whether to use an Encyclopedia volume name

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

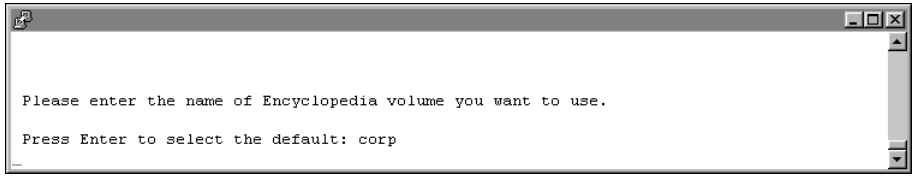

Figure 2-27 Specifying the Encyclopedia volume name

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

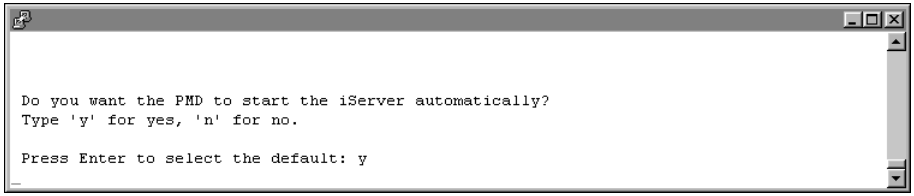

Figure 2-28 Specifying whether to start iServer automatically

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

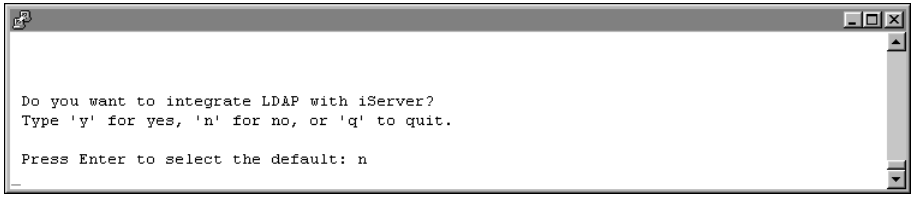

Figure 2-29 Specifying whether to integrate LDAP with iServer

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

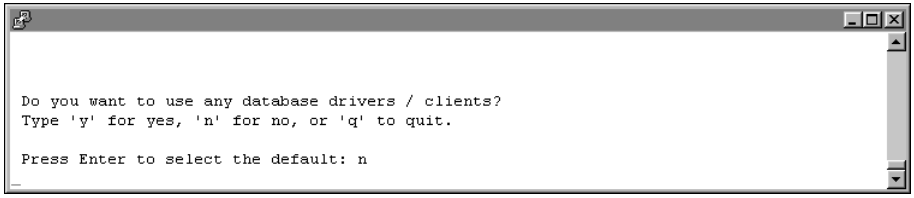

Figure 2-30 Specifying whether to use database drivers/clients

33 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 2-31.

34 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 2-32. Alternatively, type a different IP address.

35 Press Enter to accept the default port number, 8100, on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 2-33. Alternatively, type a different port number.

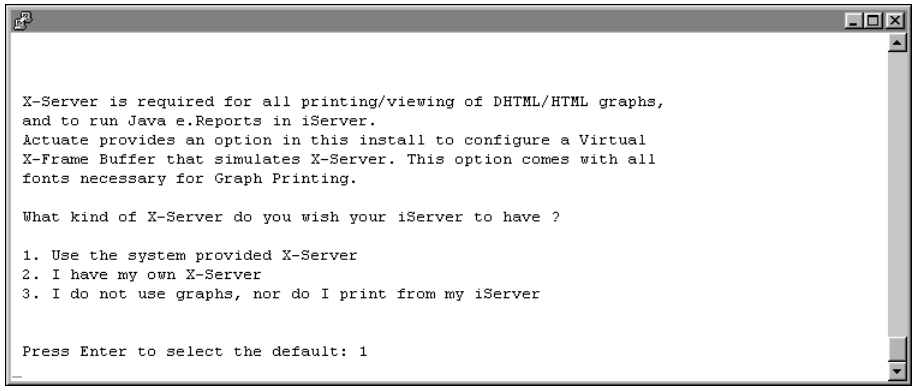

Figure 2-31 Specifying what kind of X-Server to use, if any

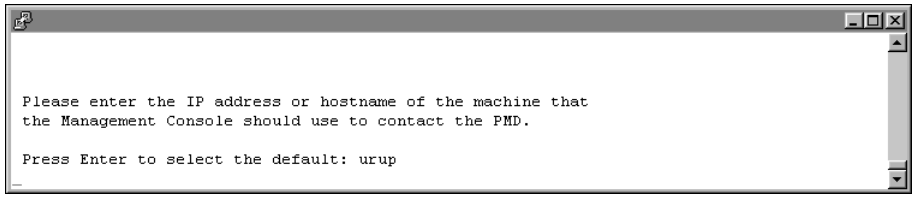

Figure 2-32 Specifying the hostname that Management Console uses to contact the PMD

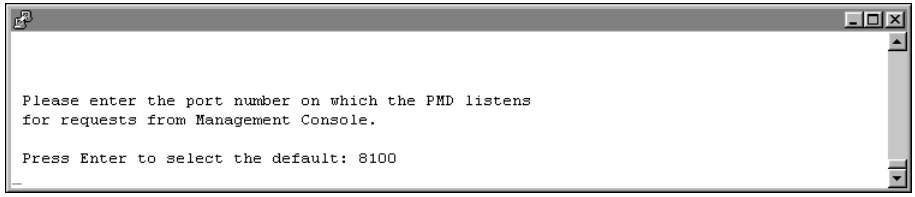

Figure 2-33 Specifying the port number for the PMD to listen for requests from Management Console

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

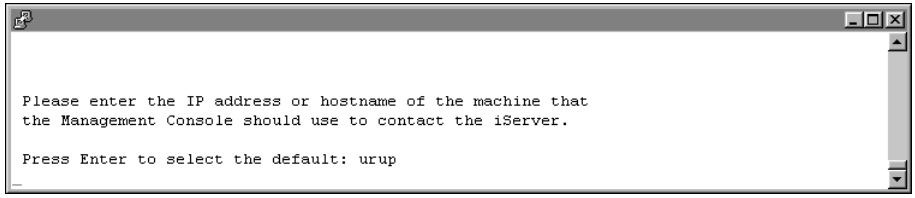

Figure 2-34 Specifying the hostname that Management Console uses to contact iServer

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

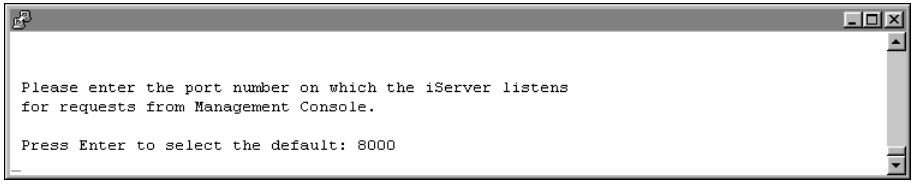

Figure 2-35 Specifying the port number on which iServer listens for requests from Management Console

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

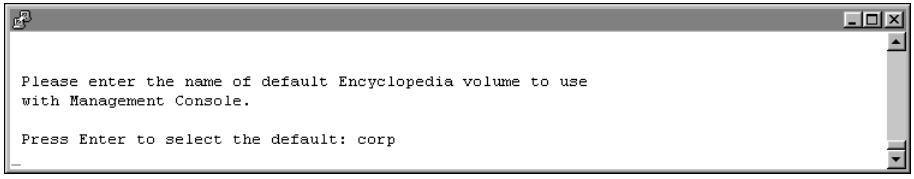

Figure 2-36 Specifying the name of the default Encyclopedia volume

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

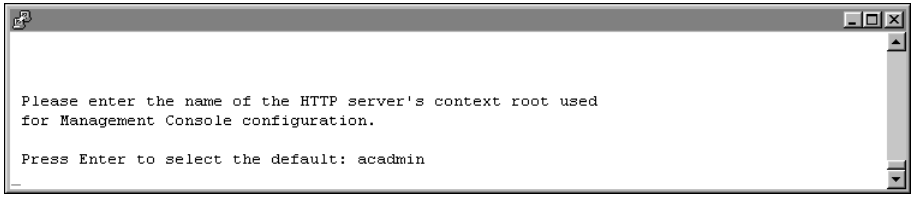

Figure 2-37 Specifying the name of the HTTP server context root

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

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

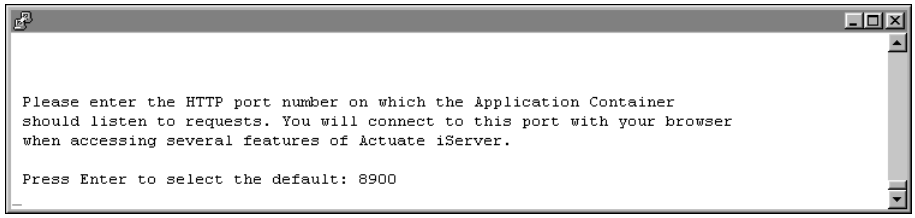

Figure 2-38 Specifying the application container listening port number

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

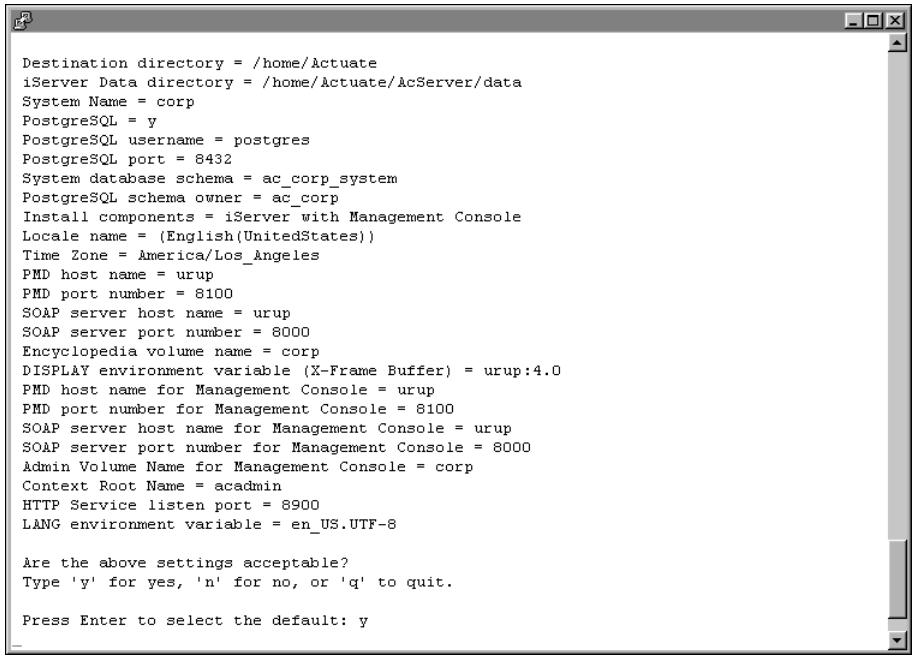

Figure 2-39 Reviewing the installation settings

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

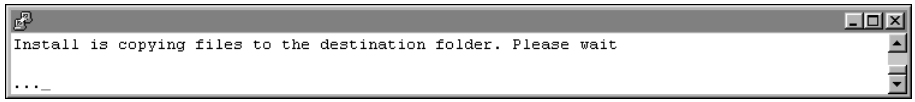

Figure 2-40 Viewing iServer installation progress

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

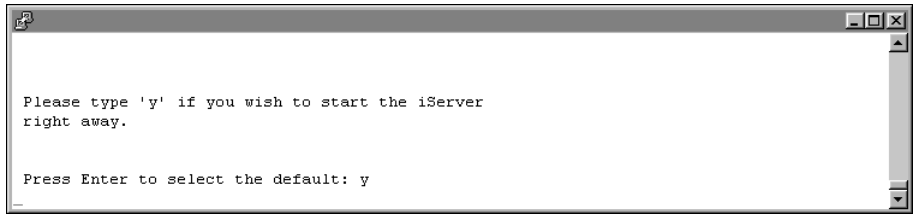

Figure 2-41 Specifying whether to start iServer

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

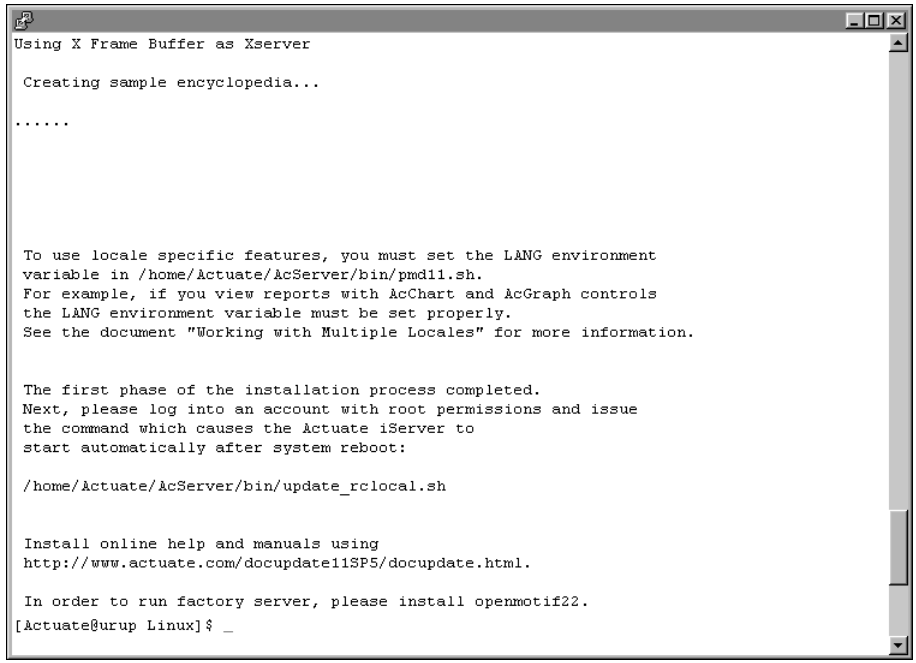

Figure 2-42 Viewing information about localization, logging in, and installing online help

Accessing Information, Management, and Configuration Consoles

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

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

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

```
http://<machine name>:8900/iportal/
```

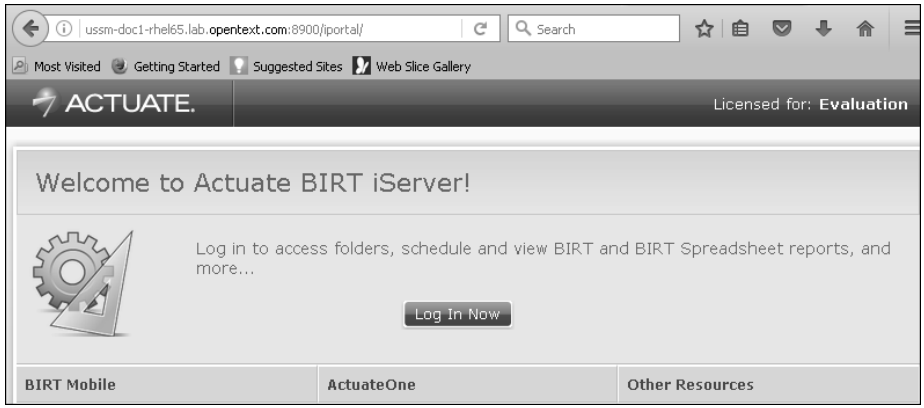

Figure 2-43 Viewing Welcome to Actuate Information Console

- Management Console
  Set up user accounts and schedule or run a design.

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

  ```
  http://<machine name>:8900/acadmin/
  ```

Figure 2-44 Logging in to Management Console

- Configuration Console
  Perform administrative operations, such as the following tasks:
  - Add an Encyclopedia volume.

- Connect to a database.
- Make modifications to iServer parameters and server templates.
- Update the license.

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

```
http://<machine name>:8900/acadmin/config/
```

Figure 2-45 Logging in to Configuration Console

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

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

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

Stopping and starting iServer and PostgreSQL processes

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

How to stop and start BIRT iServer

To stop iServer, perform the following tasks:

1 Open a command prompt and navigate to AC_SERVER_HOME/bin.

2 Type the following command and press Enter:

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

To start iServer, perform the following tasks:

1 Open a command prompt and navigate to AC_SERVER_HOME/bin.

2 Type the following command and press Enter:

```
sh ./start_srvr.sh
```

How to stop and start PostgreSQL for Actuate iServer

To stop PostgreSQL, perform the following tasks:

1 Open a command prompt and navigate to AC_SERVER_HOME/bin

2 Type the following command and press Enter:

```
sh ./stopPostgreSQL.sh
```

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

1 Open a command prompt and navigate to AC_SERVER_HOME/bin.

2 Type the following command and press Enter:

```
sh ./startPostgreSQL.sh
```

Understanding the iServer installation environment

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

About 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:

```
AC_SERVER_HOME/jdk160
```

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

- AC_JAVA_HOME
- AC_JVM_HOME

- AC_JRE_HOME
- AC_JRE64_HOME

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

The following types of Actuate report object executable files use AC_JRE_HOME and AC_JVM_HOME:

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

Accessing JAR files for report generation

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

Gathering LDAP information

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

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

Actuate Open Security uses an LDAP configuration file to map the Encyclopedia volume user information to LDAP object attributes. For more information on Actuate Open Security, see the reference implementations available in Actuate Server Integration Technology.

Following best practices

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

Using a test environment

Set up a test environment then migrate to 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.

Chapter

3

Installing BIRT iServer using an alternative database

This chapter discusses the following topics:

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

Preparing to install BIRT iServer using an alternative database

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

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

Creating a dedicated user account for installing and running BIRT iServer

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

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

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

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

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

Restrict schema and the iServer application user names to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-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 the specific PostgreSQL installation. In the commands, substitute system and schema names appropriate to your environment.

Creating a database

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

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

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

```
CREATE LANGUAGE plpgsql;
```

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

Creating the system schema owner

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

The following commands create a user role named ac_corp_system with appropriate privileges to connect to the previously created iserver database.

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

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

Creating the Encyclopedia volume schema owner

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

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

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

Creating the iServer application user

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

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

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

Creating the system schema

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

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

Creating the Encyclopedia schema

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

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

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

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

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

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

Creating the system schema owner

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

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

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

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

Creating the Encyclopedia volume schema owner

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

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

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

Creating the iServer application user

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

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

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

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

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

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

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

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

Creating a database

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

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

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

Creating the system schema owner

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

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

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

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

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

Creating the Encyclopedia volume schema owner

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

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

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

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

Creating the iServer application user

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

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

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

Creating the system schema

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

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

Creating the Encyclopedia volume schema

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

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

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

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

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

Creating user accounts

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

Creating a database

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

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

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

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

Creating the system schema owner

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

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

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

Creating the Encyclopedia volume schema owner

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

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

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

Creating the iServer application user

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

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

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

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

Creating the system schema

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

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

Creating the Encyclopedia volume schema

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

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

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

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

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

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

The solution consists of the following parts:

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

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

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

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

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

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

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

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

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

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

```
                </FunctionMapping>
          </FunctionMappings>
</MatchOpMapper>
```

to the following:

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

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

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

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

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

To resolve this problem, perform the following tasks:

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

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

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

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

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

Backing up iServer system and Encyclopedia volume metadata

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

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

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 at the following location:

https://support.opentext.com

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

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

Installing an Encyclopedia volume that uses an alternative database

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

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

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

How to install an Encyclopedia volume that uses an alternative database

To install iServer, perform the following steps:

1. Download the required files. Extract the files.
2. To install the server files, execute the following isinstall script:

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

   The script displays a number of prompts. Respond to the prompts as described in the following procedure.
3. The license agreement appears, as shown in Figure 3-1.

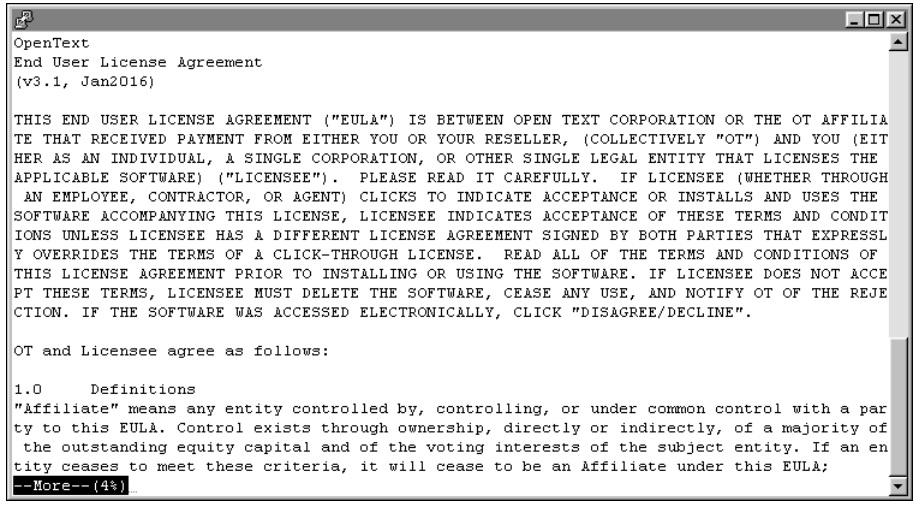

Figure 3-1 Reviewing the license agreement

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

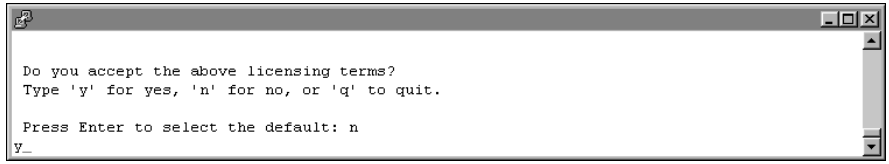

Figure 3-2 Accepting the licensing terms

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

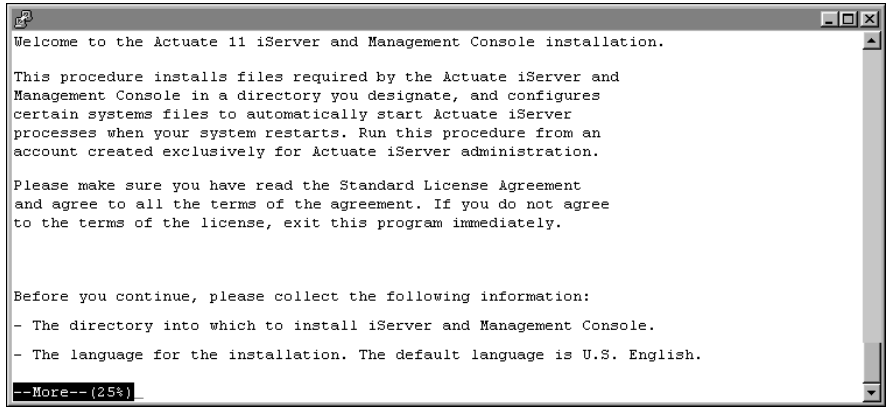

Figure 3-3 Reviewing the introductory information

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

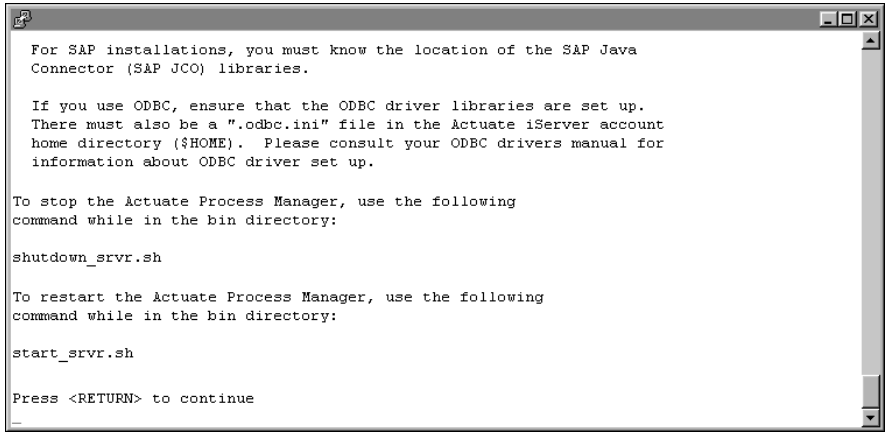

Figure 3-4 Finishing the review of introductory information

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

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

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

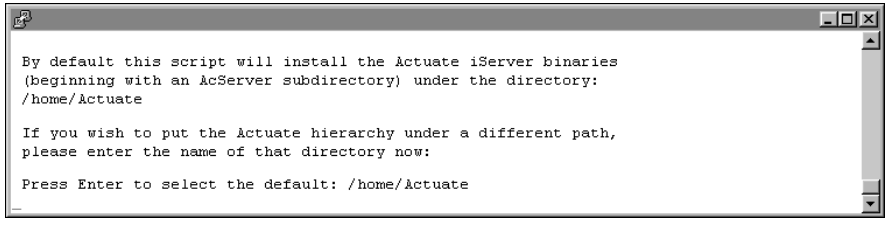

Figure 3-5 Specifying the installation directory

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

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

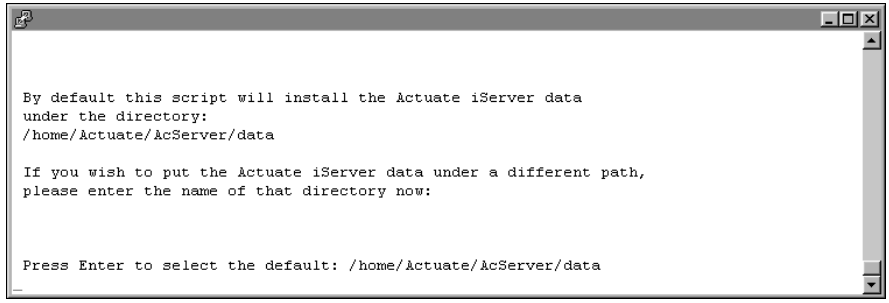

Figure 3-6 Specifying the data installation directory

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

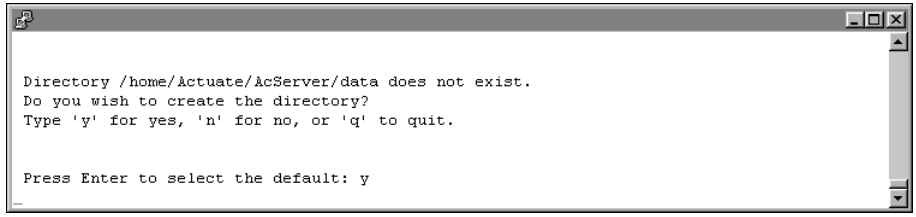

Figure 3-7 Creating the AC_DATA_HOME directory

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

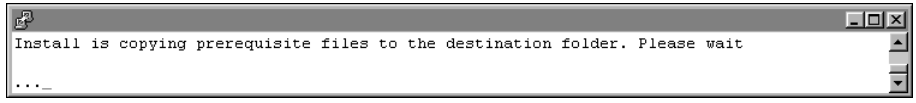

Figure 3-8 Copying prerequisite files

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

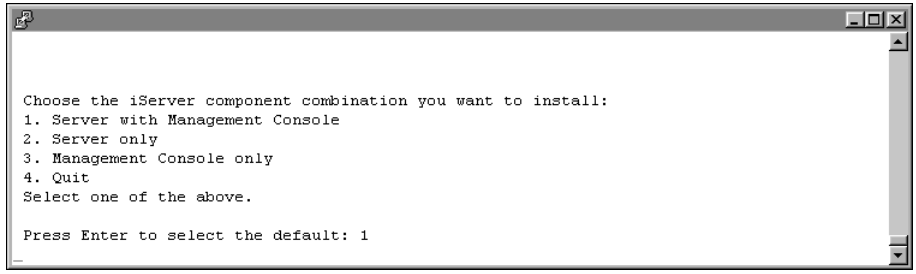

Figure 3-9 Choosing the iServer component combination

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

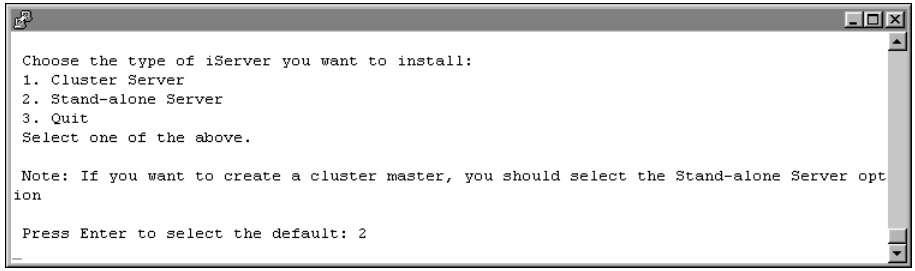

Figure 3-10 Specifying the type of iServer to install

13 Type a name to use for the BIRT iServer System name, as shown in Figure 3-11. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.

Figure 3-11 Specifying the BIRT iServer System name

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

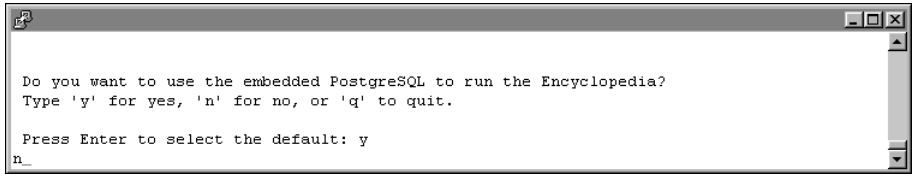

Figure 3-12 Choosing to not use Embedded PostgreSQL

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

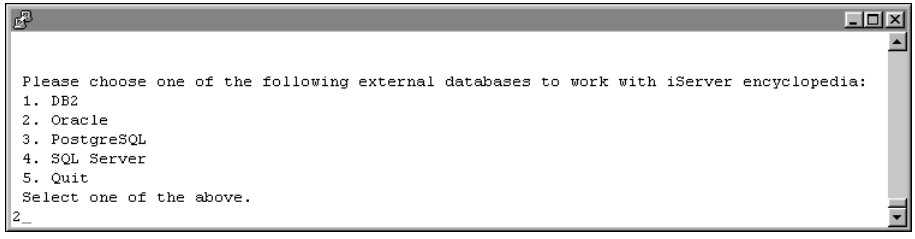

Figure 3-13 Choosing an external database for iServer Encyclopedia

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

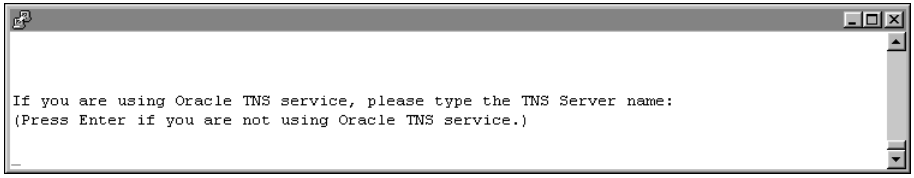

Figure 3-14 Specifying the external Oracle database TNS Server

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

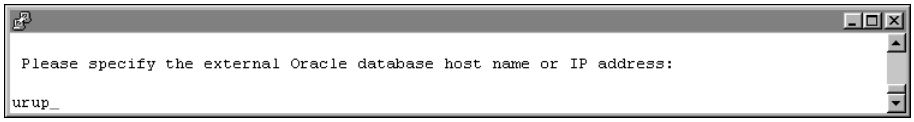

Figure 3-15 Specifying the external Oracle database server name

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

Figure 3-16 Specifying the external Oracle database port

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

Figure 3-17 Specifying the external Oracle database service name

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

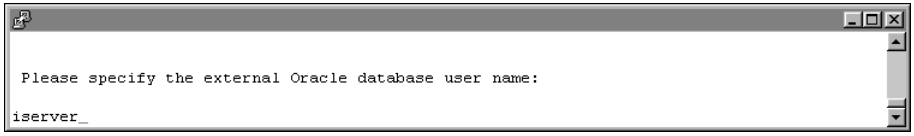

Figure 3-18 Specifying the external Oracle database user name

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

Figure 3-19 Specifying the external Oracle database user password

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

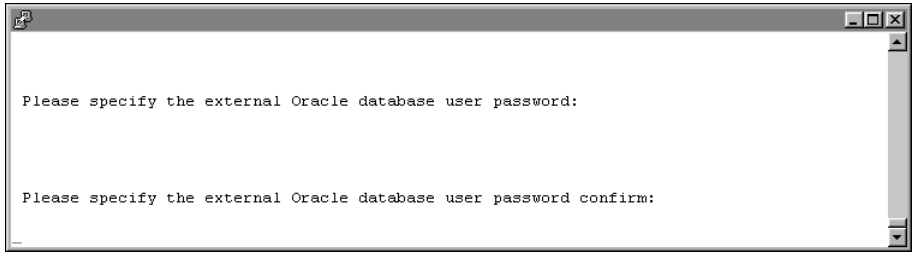

Figure 3-20 Confirming the external Oracle database user password

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

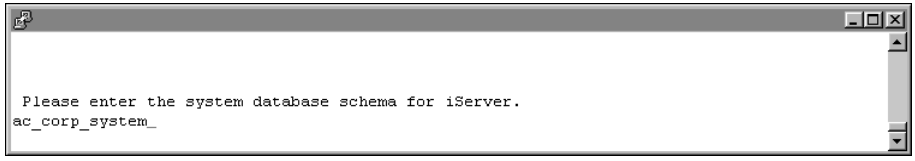

Figure 3-21 Specifying the system database schema

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

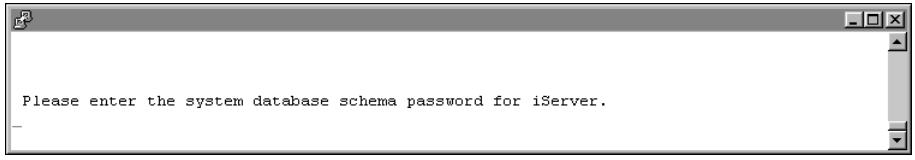

Figure 3-22 Specifying the system database schema password

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

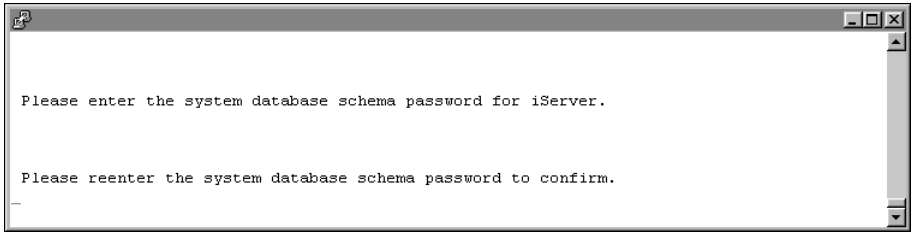

Figure 3-23 Confirming the system database schema password

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

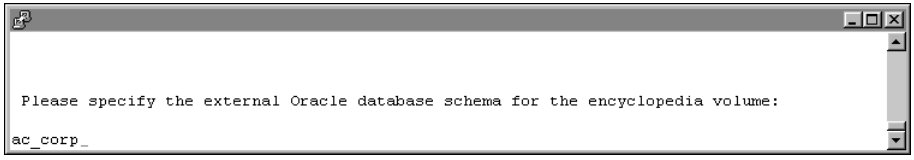

Figure 3-24 Specifying the schema for the Encyclopedia volume

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

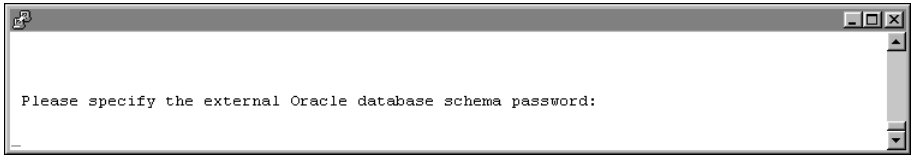

Figure 3-25 Specifying the external Oracle database schema password

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

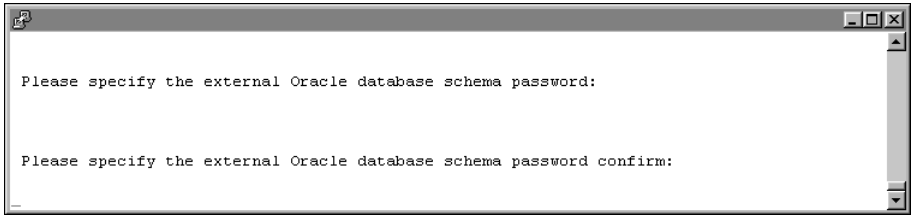

Figure 3-26 Confirm the external Oracle database schema password

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

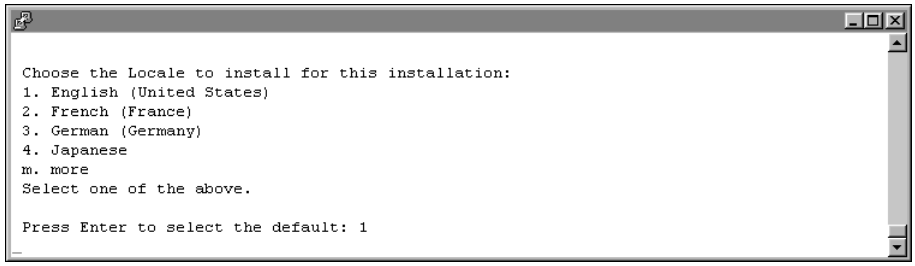

Figure 3-27 Specifying a locale

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

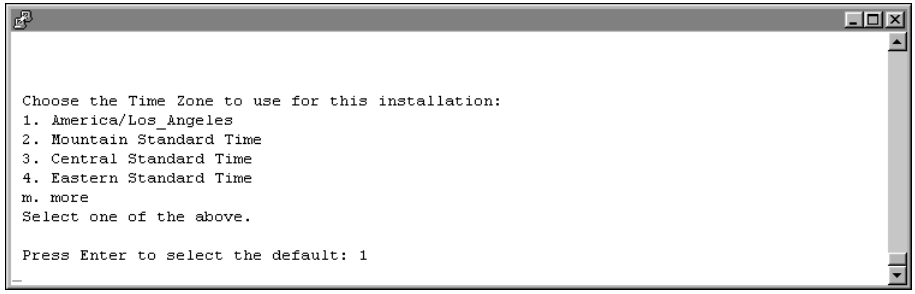

Figure 3-28 Specifying a time zone

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

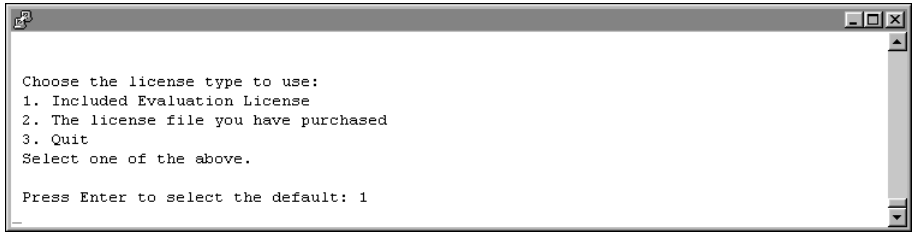

Figure 3-29 Specifying license type

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

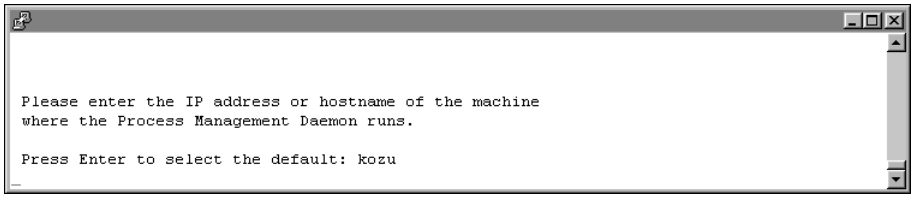

Figure 3-30 Specifying the hostname that Management Console uses to contact the PMD

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

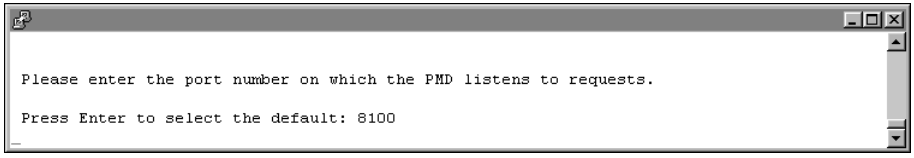

Figure 3-31 Specifying the port number on which the PMD listens

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

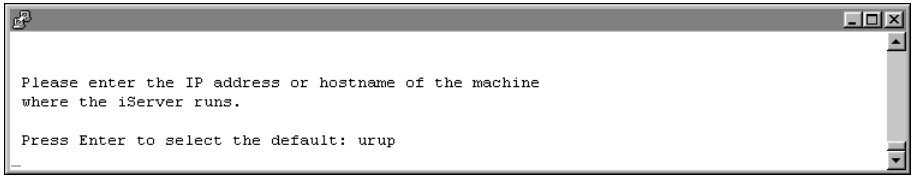

Figure 3-32 Specifying the machine on which the iServer runs

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

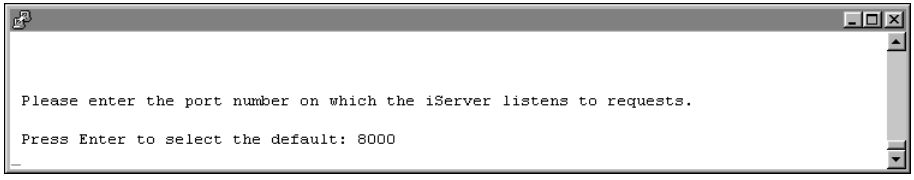

Figure 3-33 Specifying the port number on which the iServer listens

36 Specify the iServer administrator password, as shown in Figure 3-34.

You use this password to log in to the iServer Configuration Console.

Figure 3-34 Specifying the iServer administrator password

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

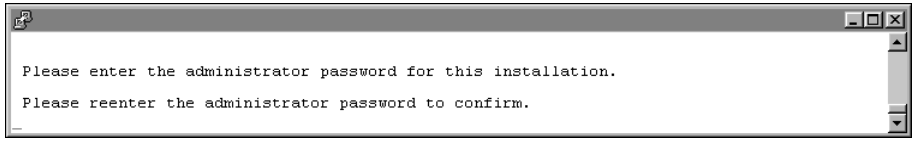

Figure 3-35 Re-entering the iServer administrator password

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

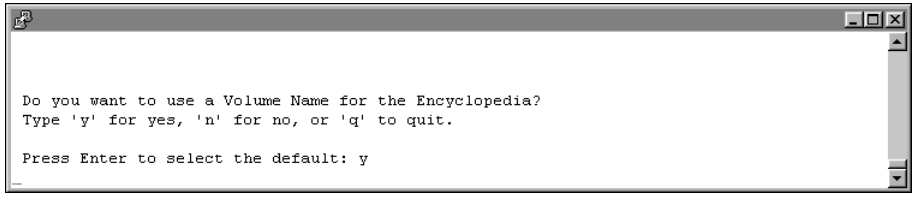

Figure 3-36 Specifying whether to use a volume name

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

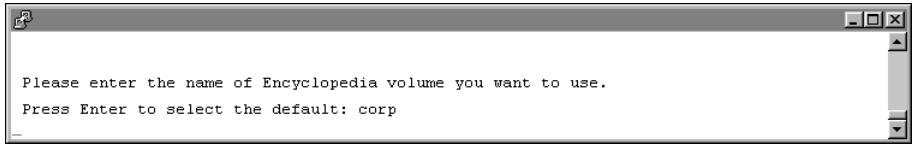

Figure 3-37 Specifying the Encyclopedia volume name

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

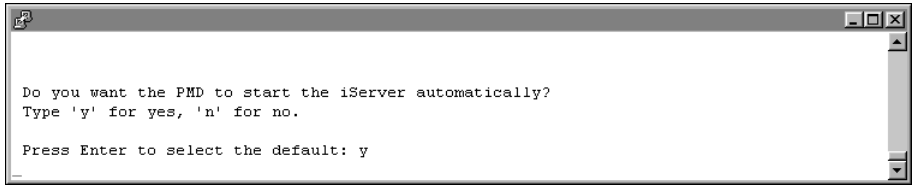

Figure 3-38 Specifying whether to start iServer automatically

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

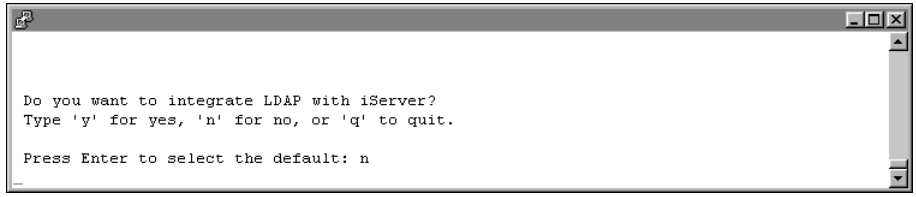

Figure 3-39 Specifying whether to integrate LDAP with iServer

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

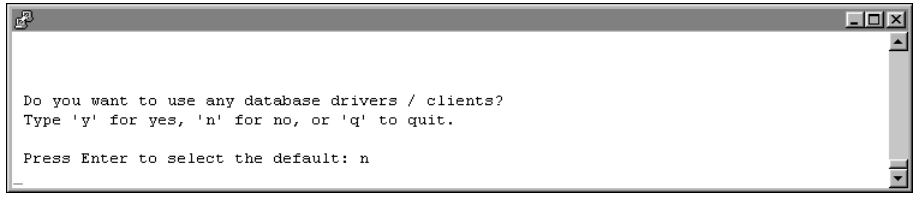

Figure 3-40 Specifying whether to use database drivers/clients

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

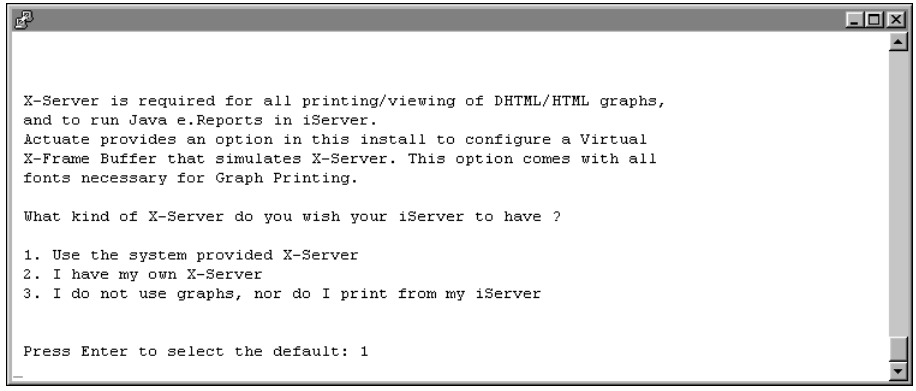

Figure 3-41 Specifying what kind of X-Server to use, if any

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

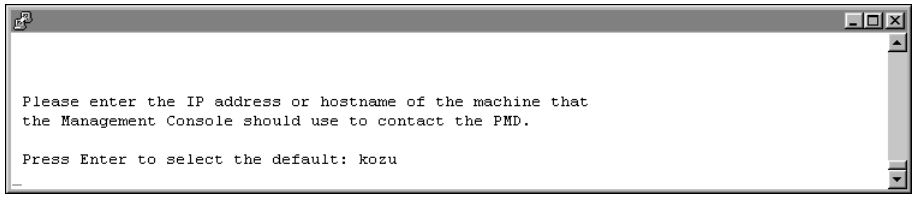

Figure 3-42 Specifying the machine host name that Management Console uses to contact the PMD

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

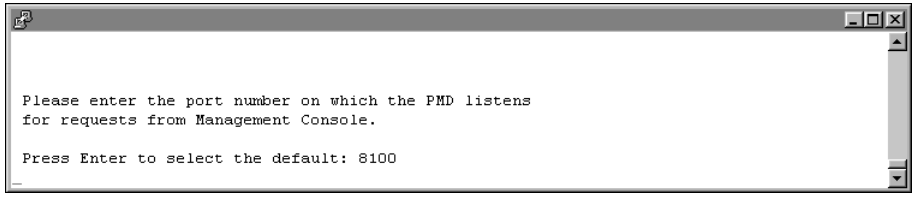

Figure 3-43 Specifying the port number for the PMD to listen for requests from Management Console

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

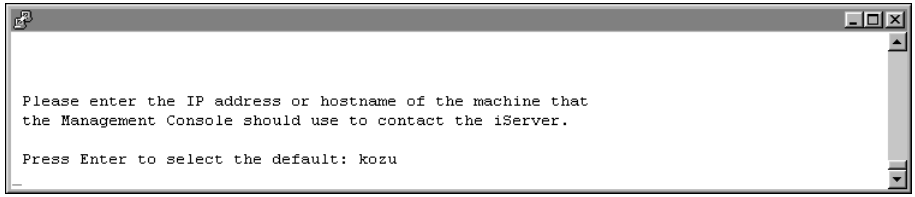

Figure 3-44 Specifying the name Management Console uses to contact iServer

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

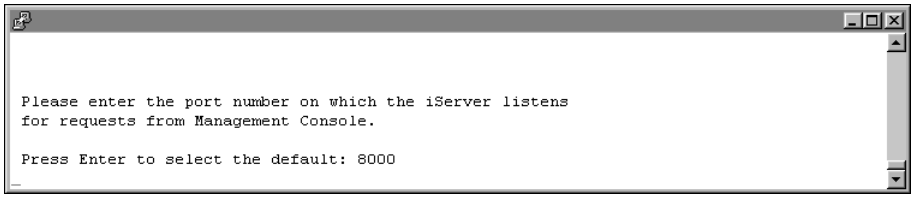

Figure 3-45 Specifying the port number for listening for requests from Management Console

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

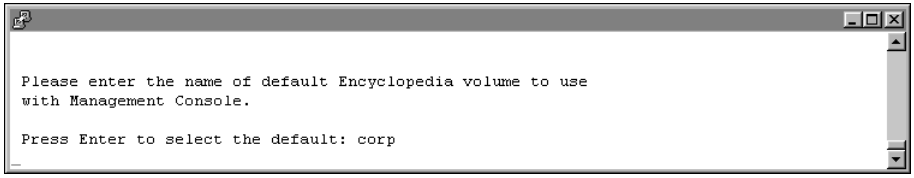

Figure 3-46 Specifying the name of the default Encyclopedia volume

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

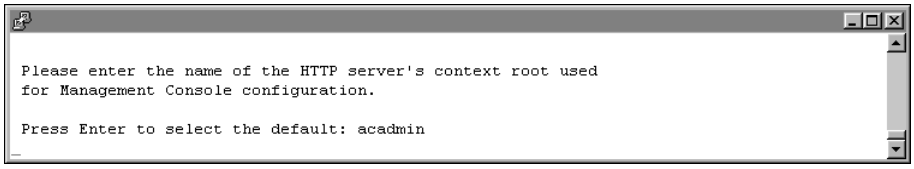

Figure 3-47 Specifying the name of the HTTP server context root

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

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

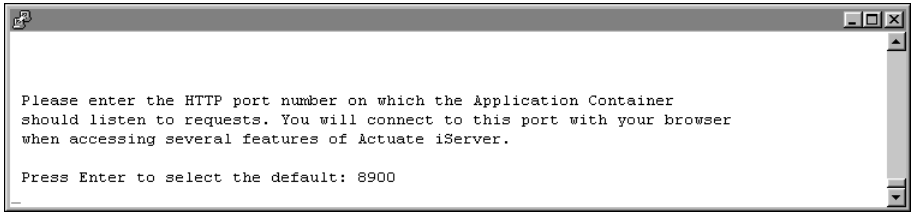

Figure 3-48 Specifying the application container listening port number

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

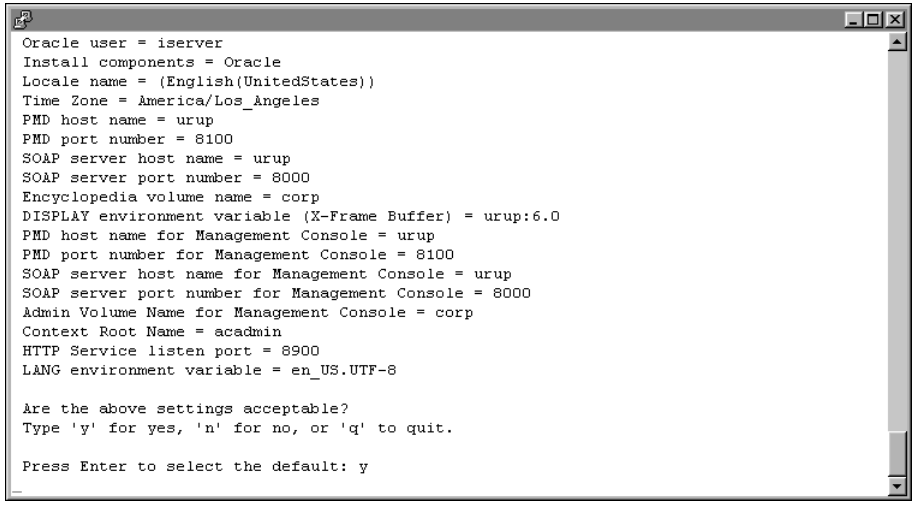

Figure 3-49 Reviewing your settings

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

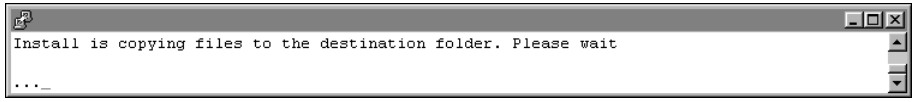

Figure 3-50 Copying iServer files to your destination folder

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

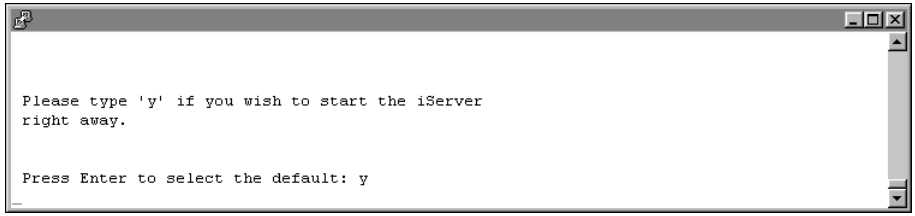

Figure 3-51 Specifying whether to start iServer

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

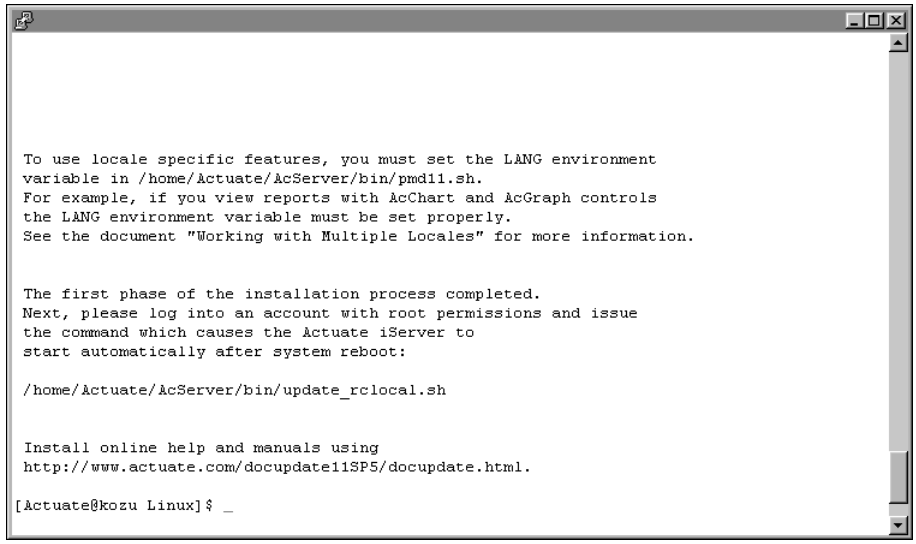

Figure 3-52 Specifying information about localization, logging in, and installing online help

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

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

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

Chapter

4

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.

Preparing to upgrade BIRT iServer

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

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

Creating a dedicated user account for installing and running BIRT iServer

Actuate recommends creating a dedicated user account for installing and running iServer. Like other Linux and UNIX processes, the processes that perform BIRT iServer tasks run under a specific user account. Having a dedicated user account

isolates iServer-specific issues and events on a machine, making it easier to administer the environment.

Before installing iServer, create the user account, making sure to configure it with privileges to access the relevant files and directories. If you exercise the same control over the user account for BIRT iServer as your site exercises for other system administrator and root accounts, you can maintain the same level of security for BIRT iServer.

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

Backing up iServer system and Encyclopedia volume metadata

When upgrading iServer 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 at the following location:

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 Release 11 BIRT iServer System architecture, see Chapter 1, "Understanding Actuate BIRT iServer architecture."

The following sections describe how to install BIRT iServer Release 11 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 in any supported RDBMS after performing the automatic, wizard-based, iServer system upgrade.

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. The installation program can encounter a problem over-writing a file linked with a running process. Be sure to stop all iServer processes before proceeding with the upgrade.

2 Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:

- encyc directories from all nodes
- acserverconfig.xml in the /etc directory
- acpmdconfig.xml in the /etc directory
- RSSE code and associated files if you use the Open Security option

3 Download the required files. Extract the files.

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

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

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

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

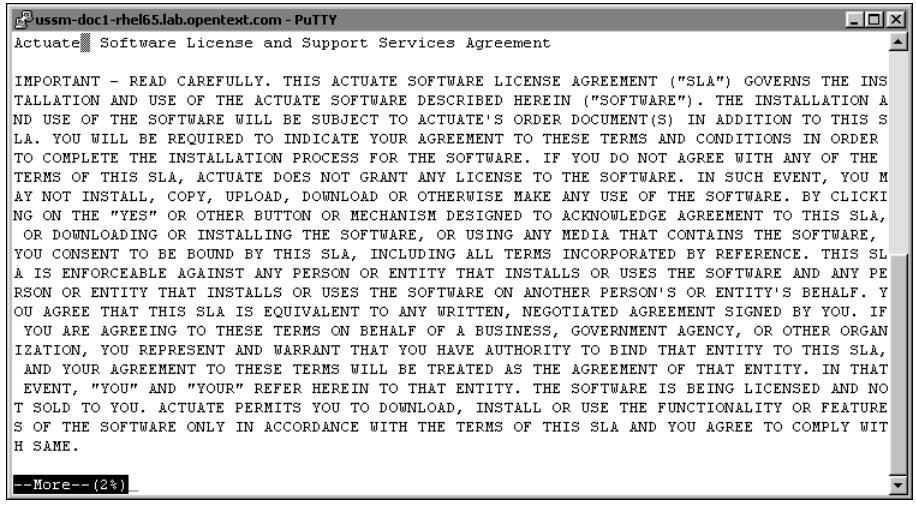

Figure 4-1 Reviewing the license agreement

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

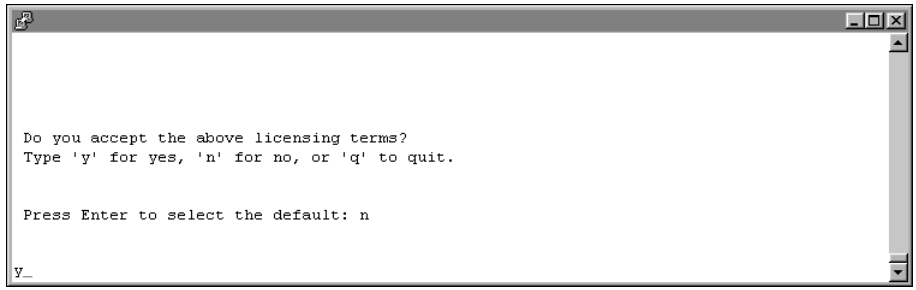

Figure 4-2 Accepting the licensing terms

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

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

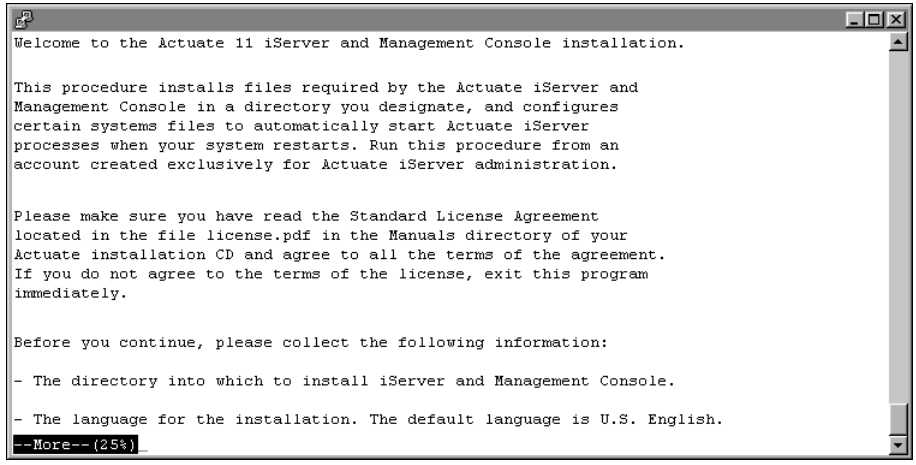

Figure 4-3 Reviewing the introductory information

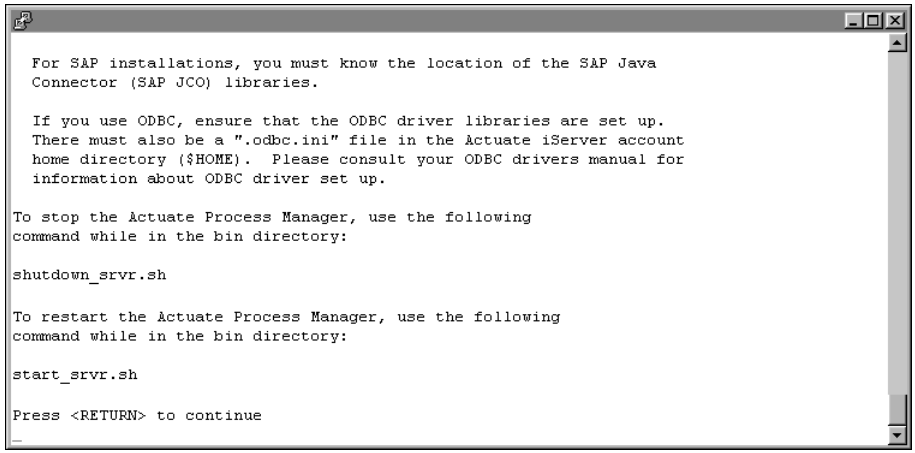

Figure 4-4 Finishing the review of introductory information

9 Press Enter to select the default directory as the location of the earlier Release 11 installation. Alternatively, type the appropriate path to that location and press Enter, as shown in Figure 4-5.

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

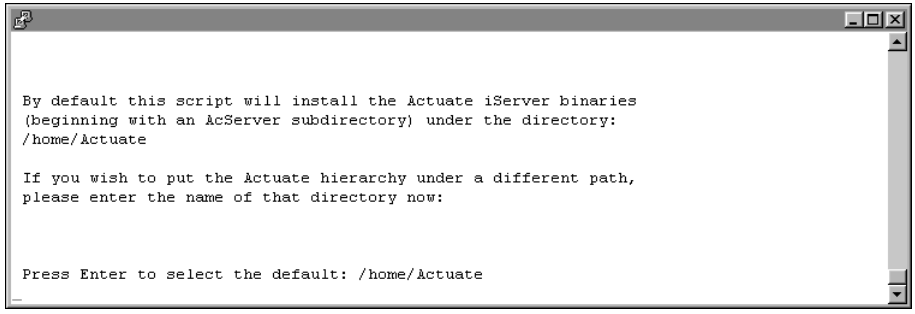

Figure 4-5 Specifying the installation directory

10 The installer detects the previous iServer version, and asks whether to perform an upgrade to the new iServer version, as shown in Figure 4-6. Press Enter to accept the default option of upgrading to the new version. Alternatively, type n for no, or q to quit.

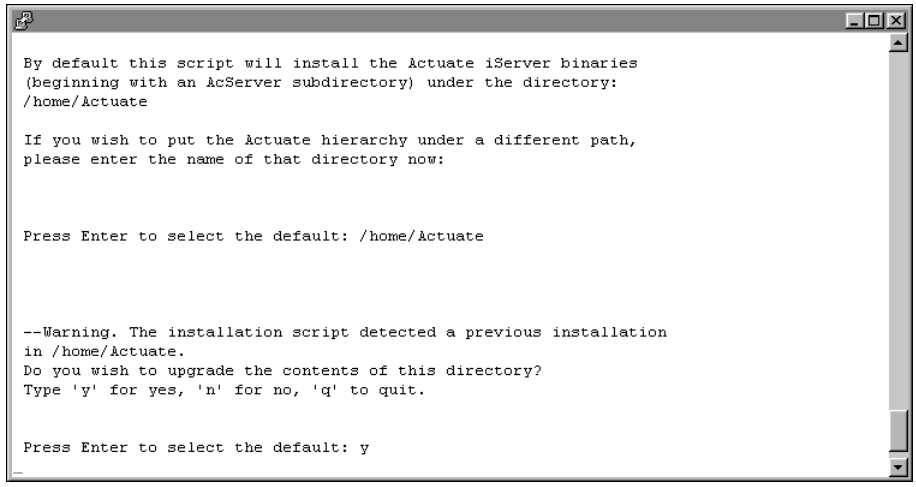

Figure 4-6 Choosing to upgrade iServer

11 If the installer detects that the Actuate servers are currently running, it will try to stop the servers then continue with the installation, as shown in Figure 4-7. This process might take a few minutes.

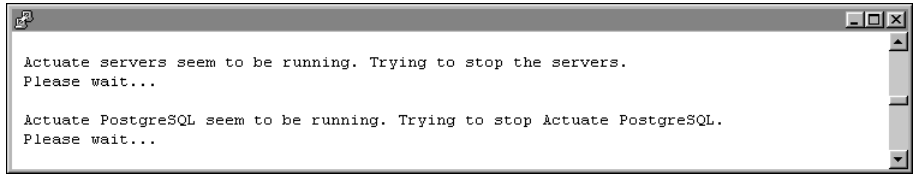

Figure 4-7 Shutting down the servers

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

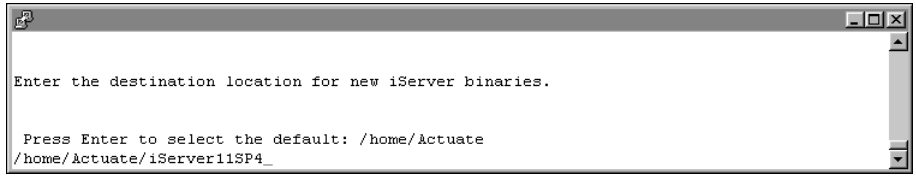

Figure 4-8 Specifying a new location for the iServer binaries

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

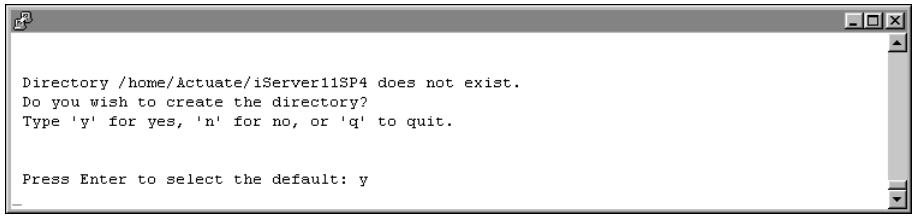

Figure 4-9 Creating the new installation directory

The installer copies prerequisite files to the new installation directory, as shown in Figure 4-10.

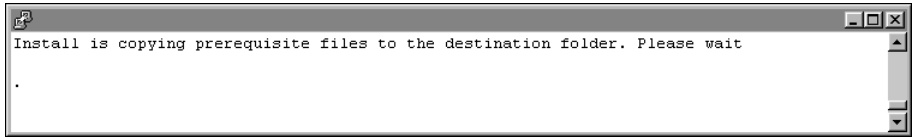

Figure 4-10 Copying prerequisite files to the new installation directory

14 Press Enter to accept the default iServer component combination, as shown in Figure 4-11. Alternatively, choose a different iServer component combination and press Enter.

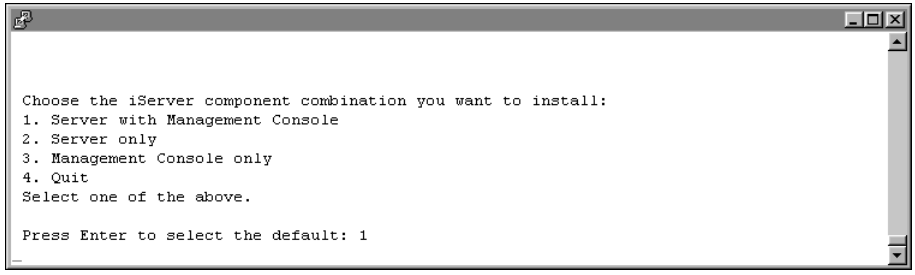

Figure 4-11 Choosing the iServer components to install

15 Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-12.

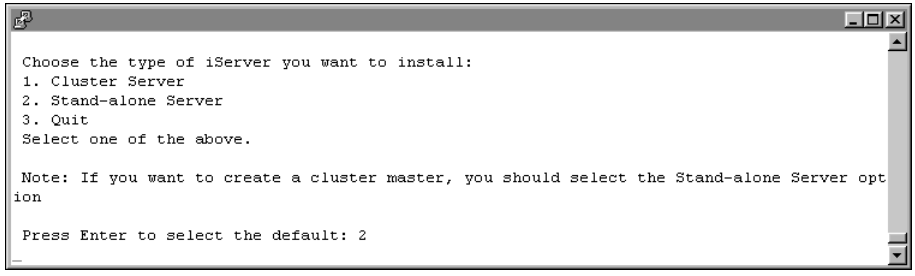

Figure 4-12 Choosing the iServer installation type

16 Specify the iServer administrator password, as shown in Figure 4-13.

You use this password to log in to the iServer Configuration Console.

Figure 4-13 Specifying the iServer administrator password

17 Re-enter the password of the iServer administrator, as shown in Figure 4-14.

You use this password to log in to Configuration Console.

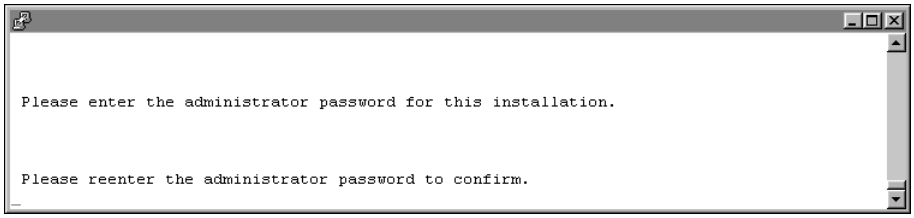

Figure 4-14 Re-entering the iServer administrator password

18 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 4-15. Alternatively, type y for yes, specify the database drivers/clients you want to use, and press Enter.

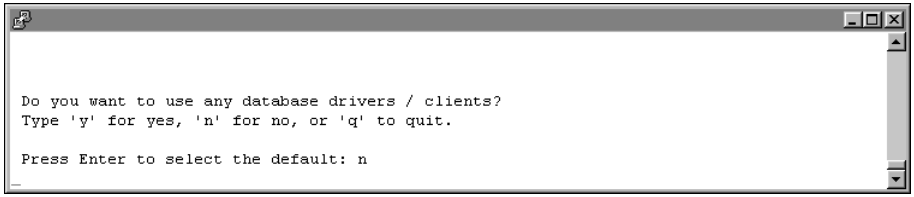

Figure 4-15 Specifying whether to use database drivers/clients

19 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-16.

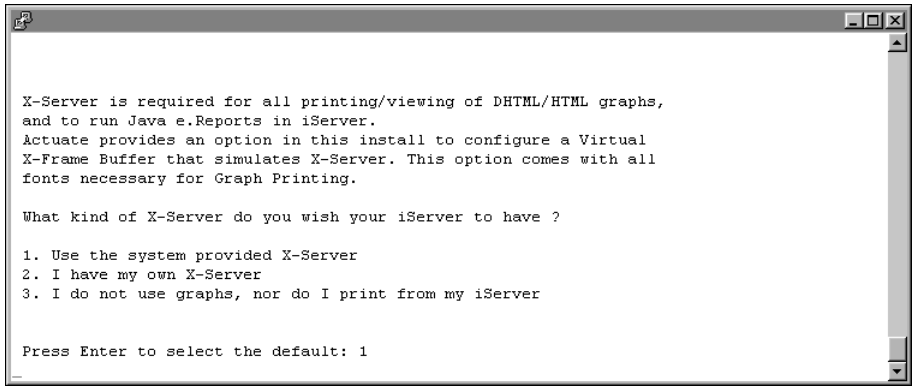

Figure 4-16 Specifying what kind of X-Server to use, if any

20 Review the settings, as shown in Figure 4-17, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.

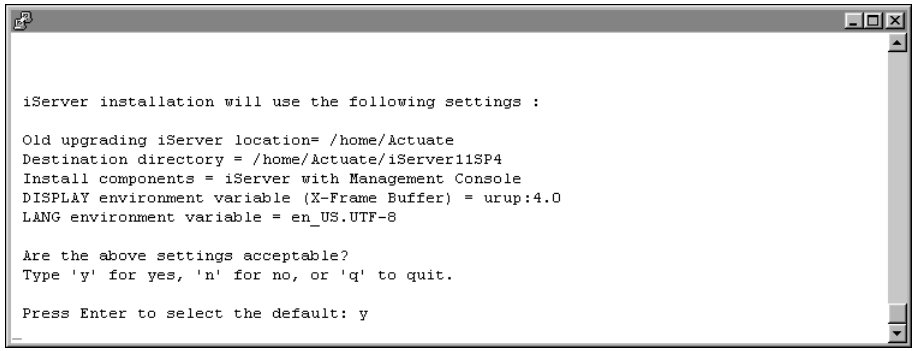

Figure 4-17 Reviewing the installation settings

21 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-18.

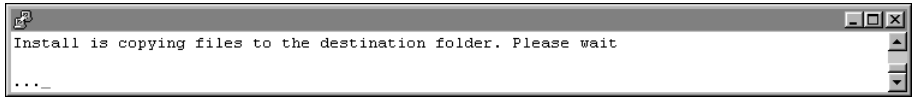

Figure 4-18 Viewing iServer installation progress

22 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 4-19.

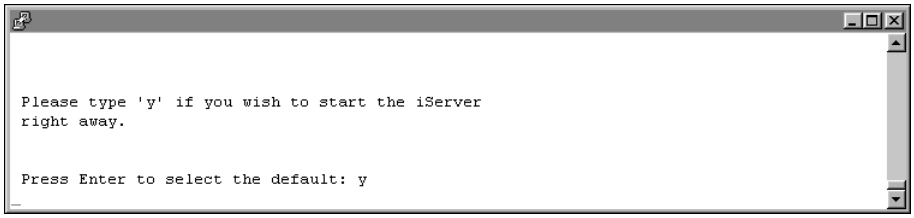

Figure 4-19 Specifying whether to start iServer

23 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-20.

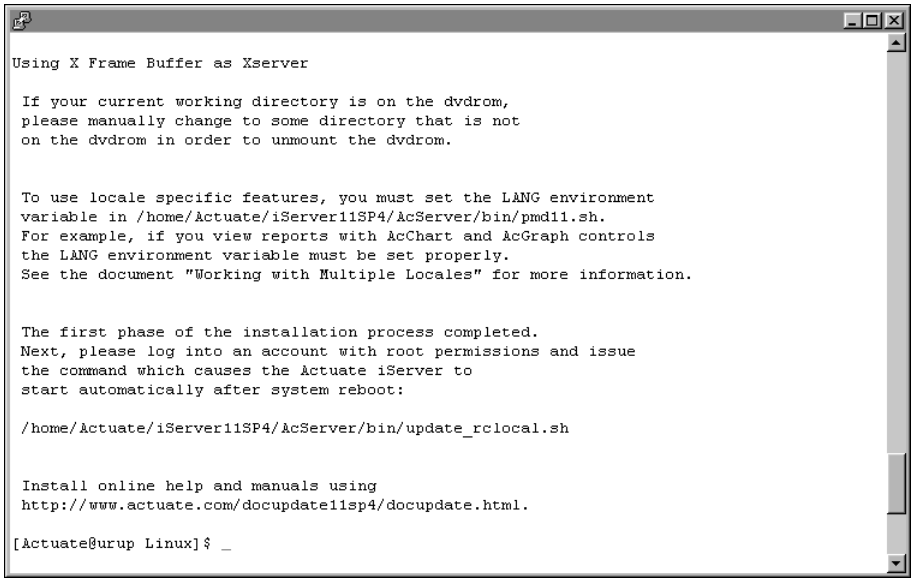

Figure 4-20 Viewing information about localization, logging in, and installing online help

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

Performing a manual upgrade of an Encyclopedia volume schema for an earlier 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 In a web browser, type:

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

2 Log in to the BIRT iServer Release 11 Configuration Console as Administrator.

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

On Volumes, point to the icon next to the volume and choose Take offline, if necessary, as shown in Figure 4-21.

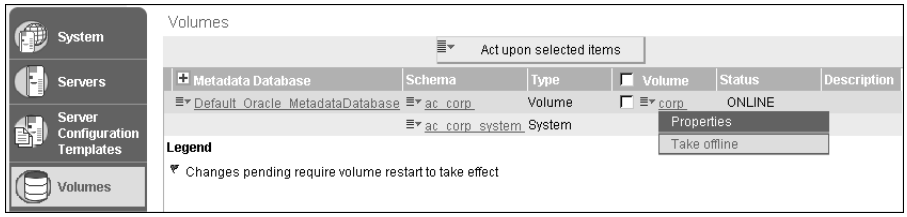

Figure 4-21 Choosing to take volume offline

Choose OK to accept the default grace period to allow current transactions on the volume to complete, as shown in Figure 4-22.

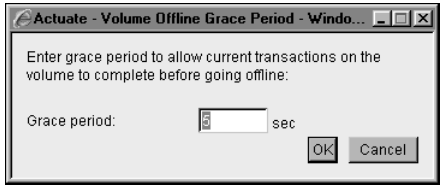

Figure 4-22 Choosing to take volume offline

3 On Volumes, point to the icon next to the volume schema and choose Properties, as shown in Figure 4-23.

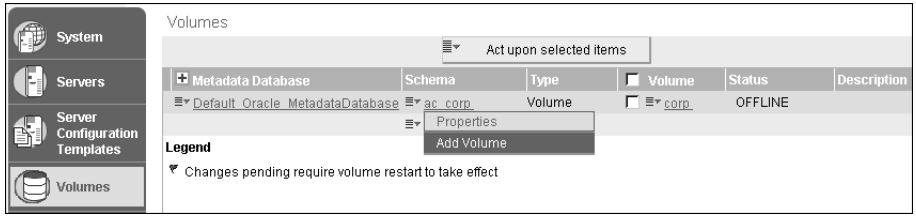

Figure 4-23 Choosing to view volume schema properties

4 On Properties, make note of the value that appears for Schema name. You must specify the schema name, not the database schema name, in a later step. Schema name is ac_corp, as shown in Figure 4-24.

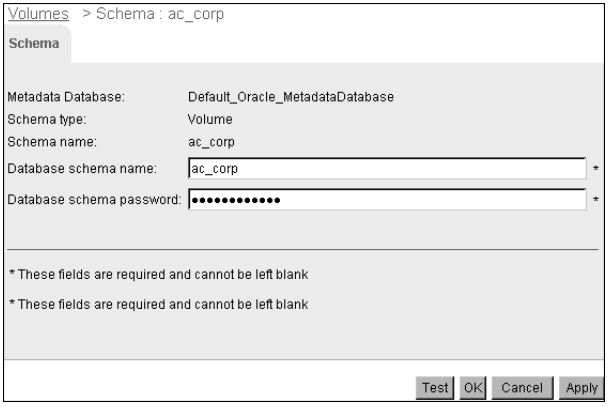

Figure 4-24 Making note of schema name

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

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

```
<AC_SERVER_HOME>/bin
```

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

```
/home/Actuate/AcServer/bin
```

2 Navigate to AC_SERVER_HOME/bin.

3 Run the upgrade_encyclopedia_data_store.sh file using the following command line syntax:

```
sh ./upgrade_encyclopedia_data_store.sh <schema name |
   property file name>
```

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

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

```
sh ./upgrdeds.sh <schema name>
```

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

- Checks to see if the administrator submitted a property file or a schema name on the command line when running the script.

- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.
- Calls the set_tools_environment.sh script, which sets the environment variables.
- If the property file exists, the script executes the EncyclopediaDataStoreUpgrader utility using the name of the property file as an argument.
- If the property file does not exist, the script creates a property file, upgrade_encyclopedia_data_store.properties, which contains the environment variable settings.

Listing 4-1 upgrade_encyclopedia_data_store.sh

```
#!/bin/sh

if [ "x$1" = "x" ]; then
   echo "Usage: upgrade_encyclopedia_data_store.sh <schema
   name | property file name>"
   exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

# Check if argument 1 is a file
if [ -f "$1" ]; then
   # Use property file specified on command line
   PROPERTY_FILE="$1"
else
   # Create property file
   PROPERTY_FILE=upgrade_encyclopedia_data_store.properties
   echo "AC_SERVER_HOME = $AC_SERVER_HOME" > $PROPERTY_FILE
   echo "AC_DATA_HOME = $AC_DATA_HOME" >> $PROPERTY_FILE
   echo "USE_SERVER_CONFIG_FILE = true" >> $PROPERTY_FILE
   echo "CONFIG_SCHEMA_NAME = $1" >> $PROPERTY_FILE
fi

# Upgrade data store
java com.actuate.iserver.encyclopedia.datastore.admin
   .EncyclopediaDataStoreUpgrader "$PROPERTY_FILE"
```

You must run Encyclopedia Data Store Upgrader separately for every schema that you want to upgrade when upgrading from an earlier Actuate 11 version.

See Table 4-3 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.

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

Point to the icon next to the volume name and choose Take online, as shown in Figure 4-25.

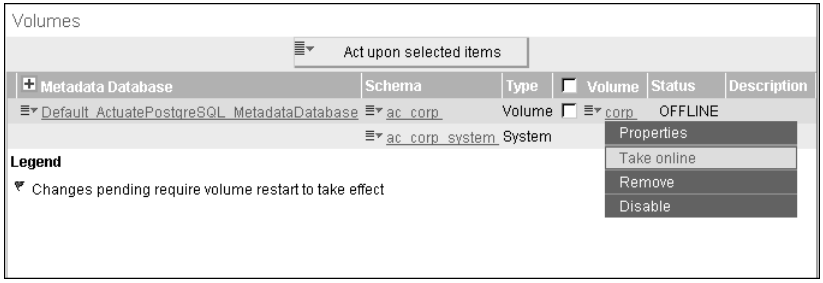

Figure 4-25 Taking the volume online

The volume comes online, as shown in Figure 4-26.

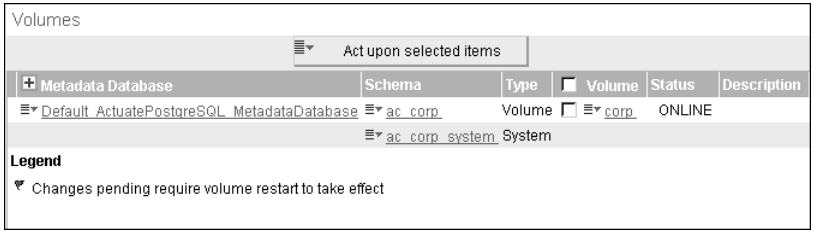

Figure 4-26 Viewing the status of the volume after it comes online

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

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

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

   The script displays a number of prompts. Respond to the prompts as described in the following procedure.
5. The license agreement appears, as shown in Figure 4-27.

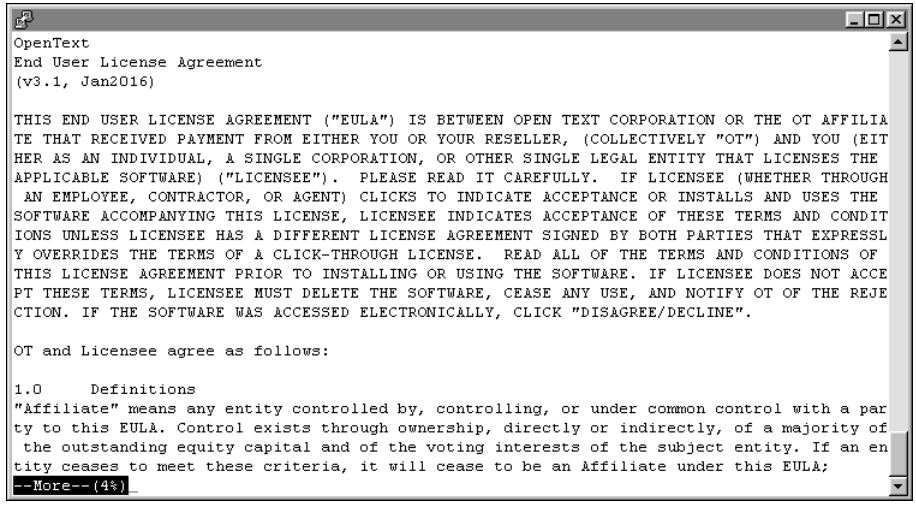

Figure 4-27 Reviewing the license agreement

6 Read the license agreement, then press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 4-28.

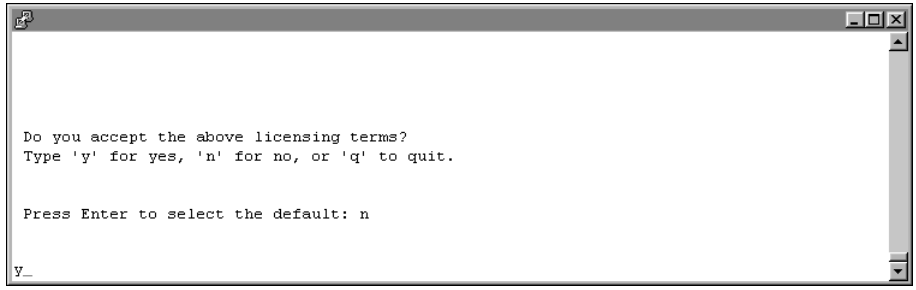

Figure 4-28 Accepting the licensing terms

7 The introduction to the installation appears, as shown in Figure 4-29.

8 Press Enter after reviewing the introductory information, as shown in Figure 4-30.

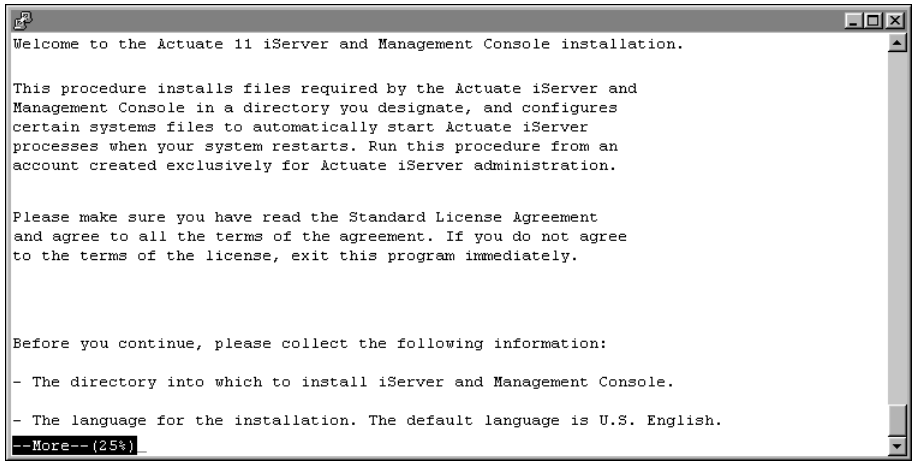

Figure 4-29 Reviewing the introductory information

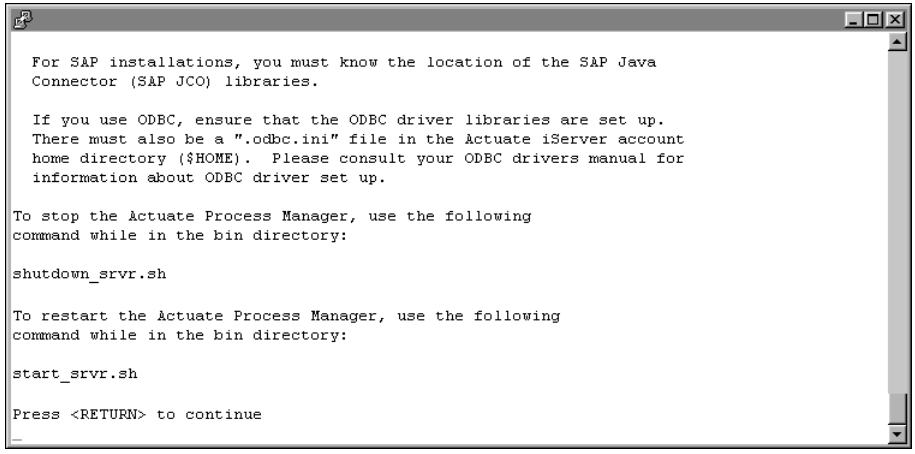

Figure 4-30 Finishing the review of introductory information

9 Press Enter to select the default directory as the location of the 11SP4 Fix 5 installation. Alternatively, type the appropriate path to that location and press Enter, as shown in Figure 4-31.

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

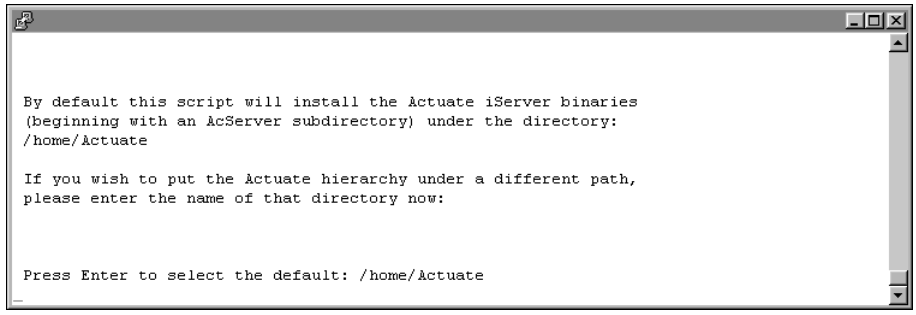

Figure 4-31 Specifying the installation directory

10 The installer detects the previous iServer version, and asks whether to perform an upgrade to the new iServer version, as shown in Figure 4-32. Press Enter to accept the default option of upgrading to the new version. Alternatively, type n for no, or q to quit.

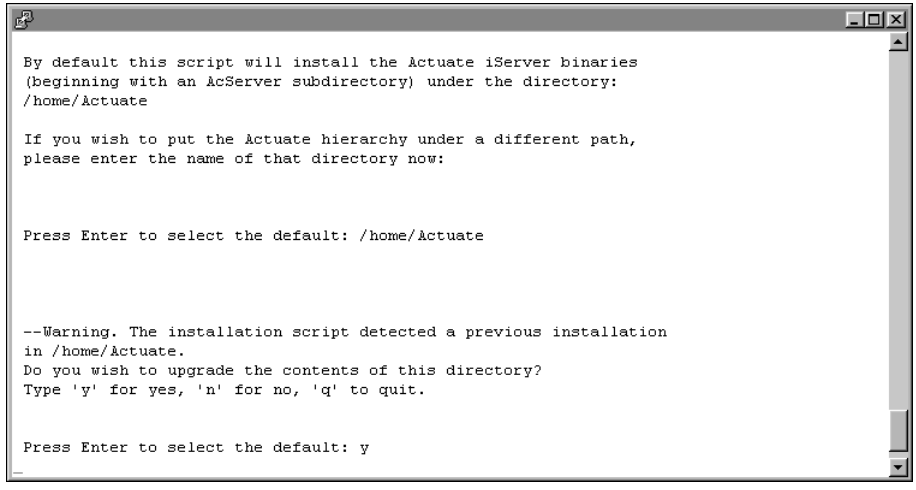

Figure 4-32 Choosing to upgrade iServer

11 If the installer detects that the Actuate servers are currently running, it will try to stop the servers then continue with the installation, as shown in Figure 4-33. This process might take a few minutes.

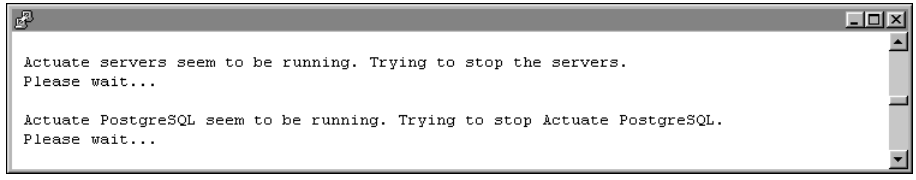

Figure 4-33 Shutting down the servers

12 Type the path to a new directory in which to install the iServer binaries, as shown in Figure 4-34.

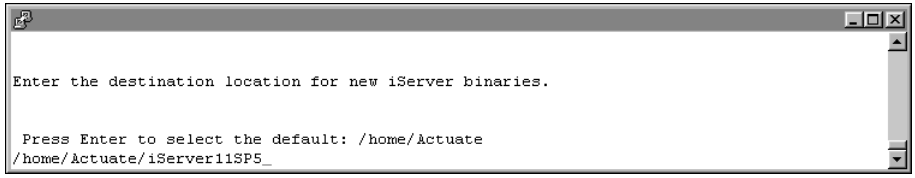

Figure 4-34 Specifying a new location for the iServer binaries

13 Press Enter to accept the default option of creating the folder you specified in the previous step, as shown in Figure 4-35. Alternatively, type n for no, or q to quit, and press Enter.

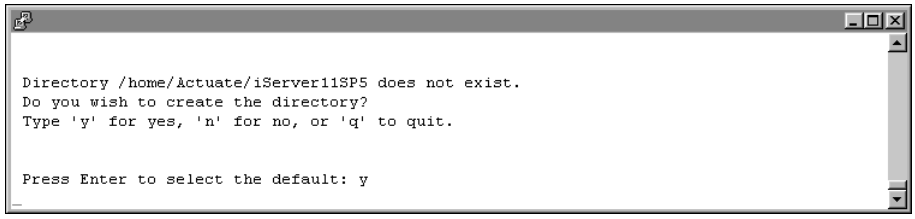

Figure 4-35 Creating the new installation directory

The installer copies prerequisite files to the new installation directory, as shown in Figure 4-36.

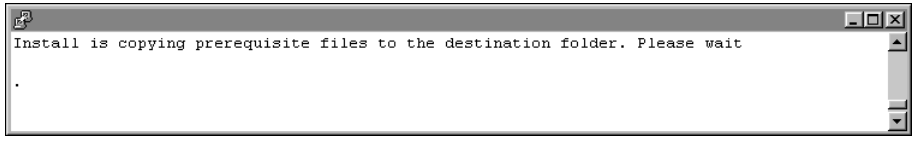

Figure 4-36 Copying prerequisite files to the new installation directory

14 Press Enter to accept the default iServer component combination, as shown in Figure 4-37. Alternatively, choose a different iServer component combination and press Enter.

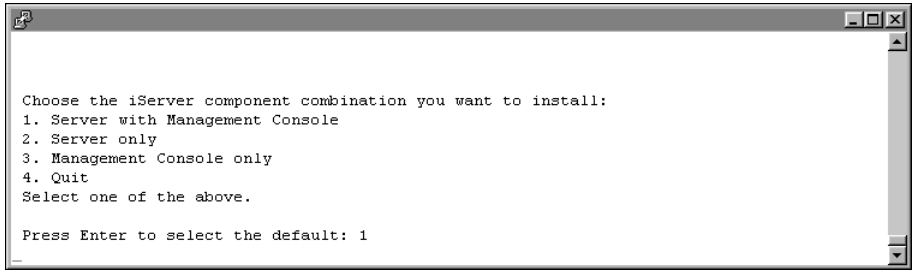

Figure 4-37 Choosing the iServer components to install

15 Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-38.

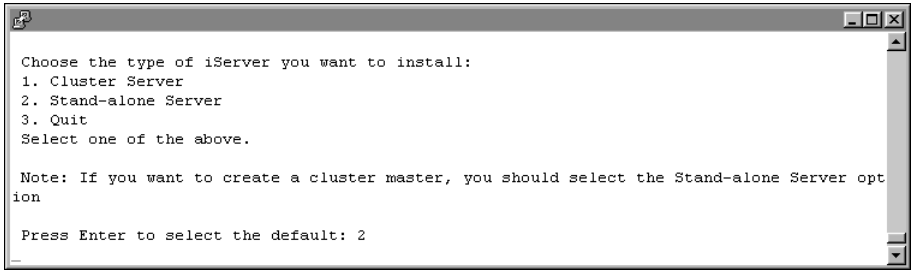

Figure 4-38 Choosing the iServer installation type

16 Specify the iServer administrator password, as shown in Figure 4-39.

You use this password to log in to the iServer Configuration Console.

Figure 4-39 Specifying the iServer administrator password

17 Re-enter the password of the iServer administrator, as shown in Figure 4-40.

You use this password to log in to Configuration Console.

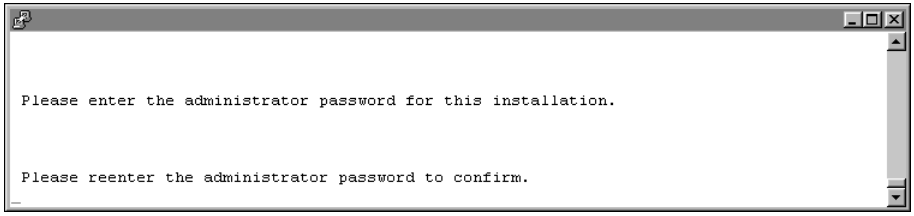

Figure 4-40 Re-entering the iServer administrator password

18 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 4-41. Alternatively, type y for yes, specify the database drivers/clients you want to use, and press Enter.

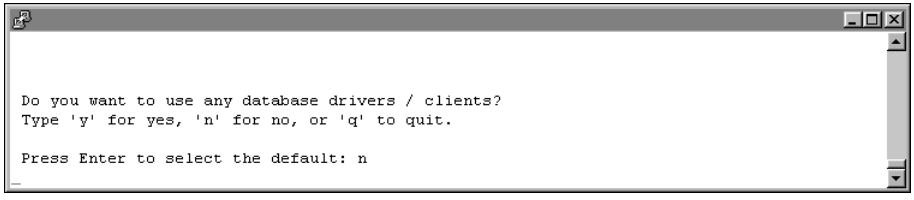

Figure 4-41 Specifying whether to use database drivers/clients

19 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-42.

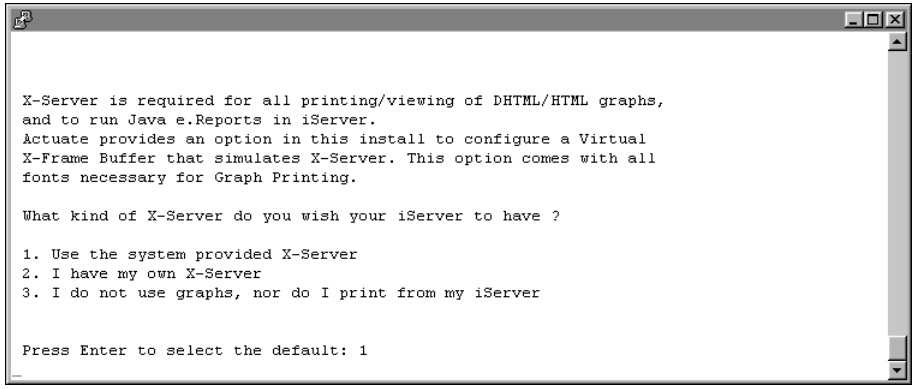

Figure 4-42 Specifying what kind of X-Server to use, if any

20 Review the settings, as shown in Figure 4-43, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.

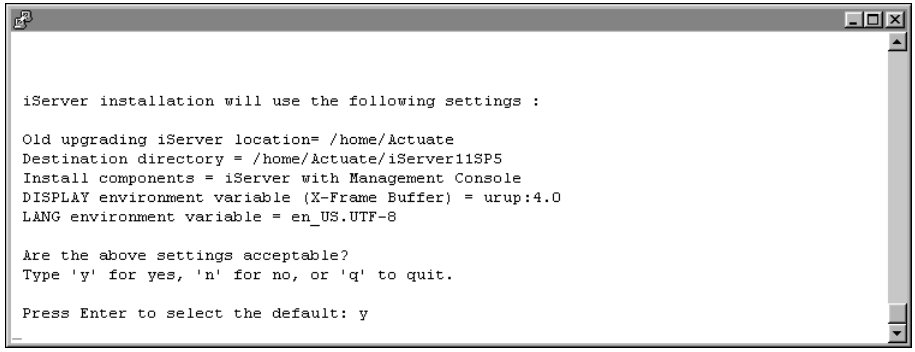

Figure 4-43 Reviewing the installation settings

21 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-44.

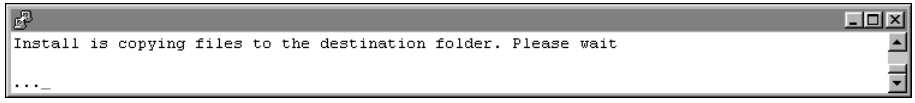

Figure 4-44 Viewing iServer installation progress

22 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 4-45.

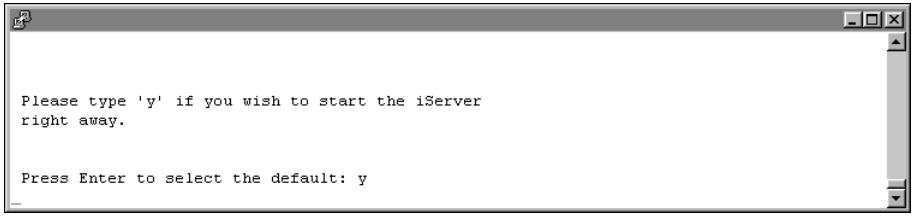

Figure 4-45 Specifying whether to start iServer

23 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-46.

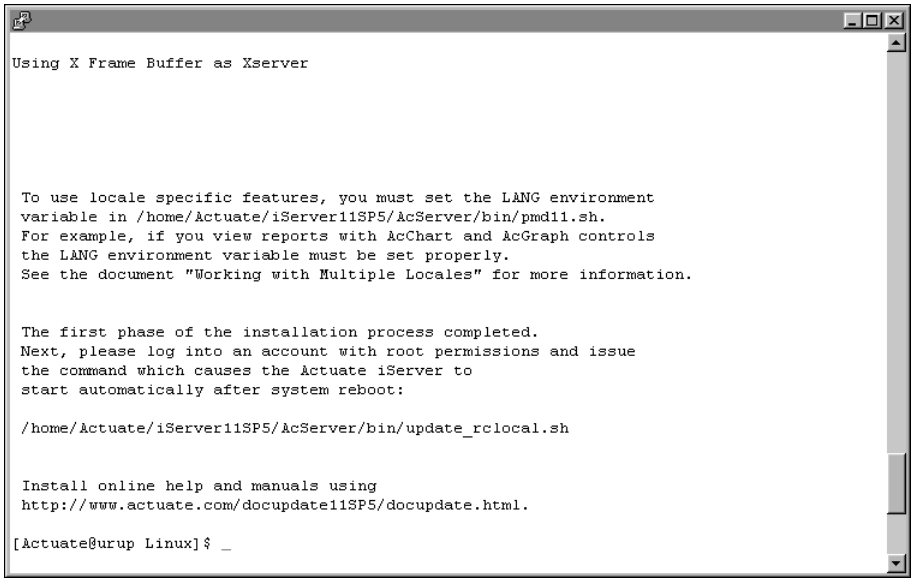

Figure 4-46 Viewing information about localization, logging in, and installing online help

Log in to Management Console. In Files and Folders, the data from the previous release 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 In a web browser type:

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

   2 Log in to Configuration Console as Administrator.

   3 Choose Advanced view.

   4 From the side menu, choose System. In System—Status, choose Stop to shut down iServer.

   Make sure the earlier iServer is offline before performing the next operation, as shown in Figure 4-47.

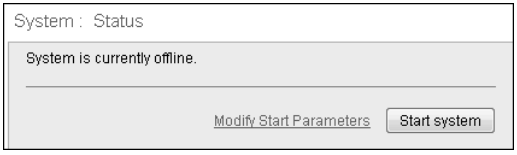

Figure 4-47 iServer is offline

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:

```
/home/Actuate/AcServer/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 = /home/Actuate/AcServer
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 = /home/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 = /home/Actuate/AcServer
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 = /home/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 = /home/Actuate/AcServer
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 = /home/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 = /home/Actuate/AcServer
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 = /home/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 = /home/Actuate/AcServer
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 = /home/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.sh script using the following command line syntax:

   ```
   sh ./administrate_encyclopedia_data_store.sh
      VolumeExport.properties
   ```

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

   ```
   sh ./admineds.sh 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 2-6. The default path for AC_DATA_HOME is:

        ```
        /home/Actuate/AcServer/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 3.1, as shown in Figure 4-48. This path does not need to match the partition path to the same volume in 11SP4 Fix 5. Choose OK.

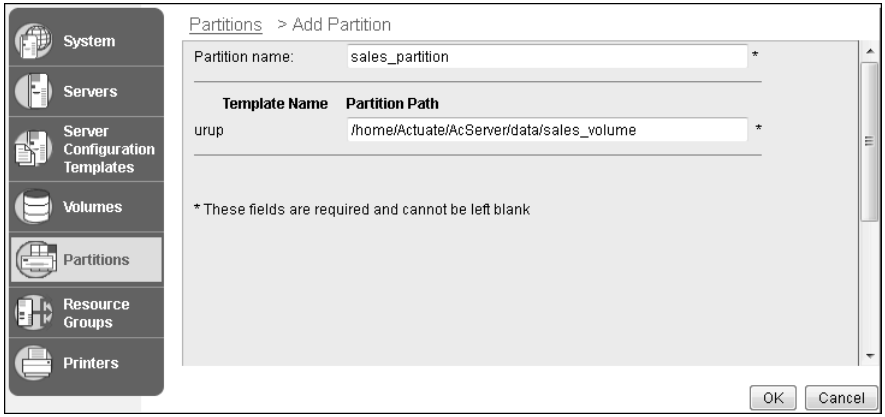

Figure 4-48 Creating a partition for a migrated volume

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, 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:

```
/home/Actuate/AcServer/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 4-49.

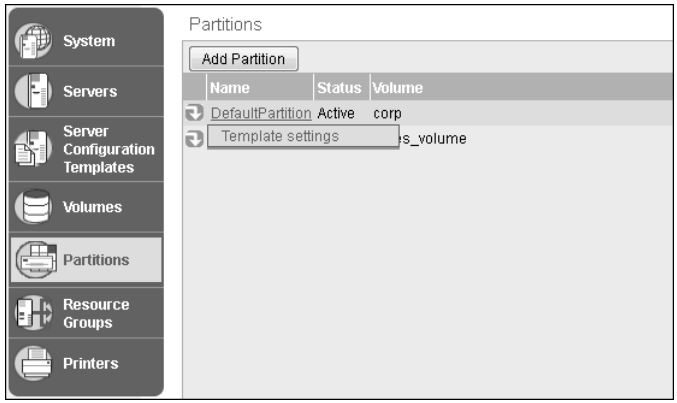

Figure 4-49 Choosing Template settings

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 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 4-50. The name of the default volume is the name you specified during the 11SP5 install procedure, as shown in Figure 2-11.

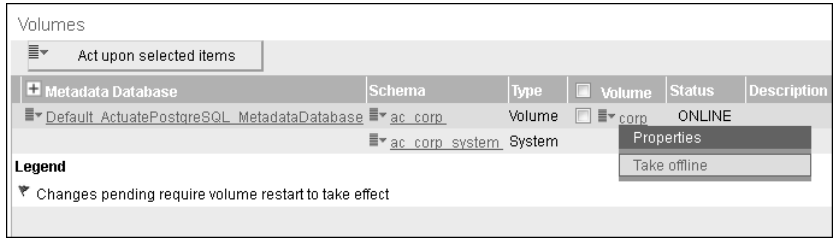

Figure 4-50 Taking the default volume offline

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 4-51.

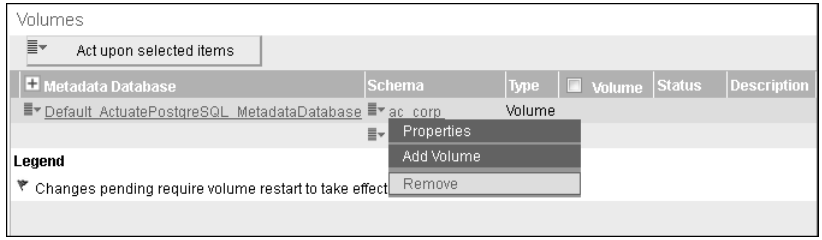

Figure 4-51 Removing the volume schema in 11SP5

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 4-52.

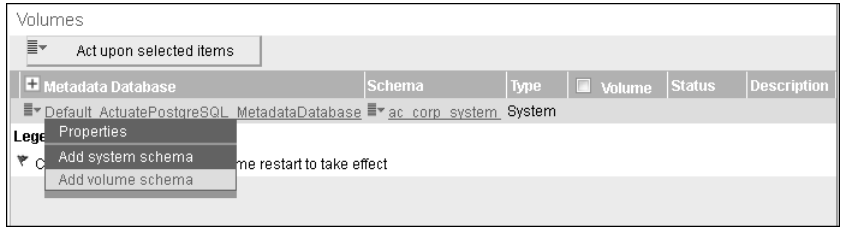

Figure 4-52 Choosing to create a new volume schema

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 2-14. Choose OK.

Volumes—New Schema appears as shown in Figure 4-53.

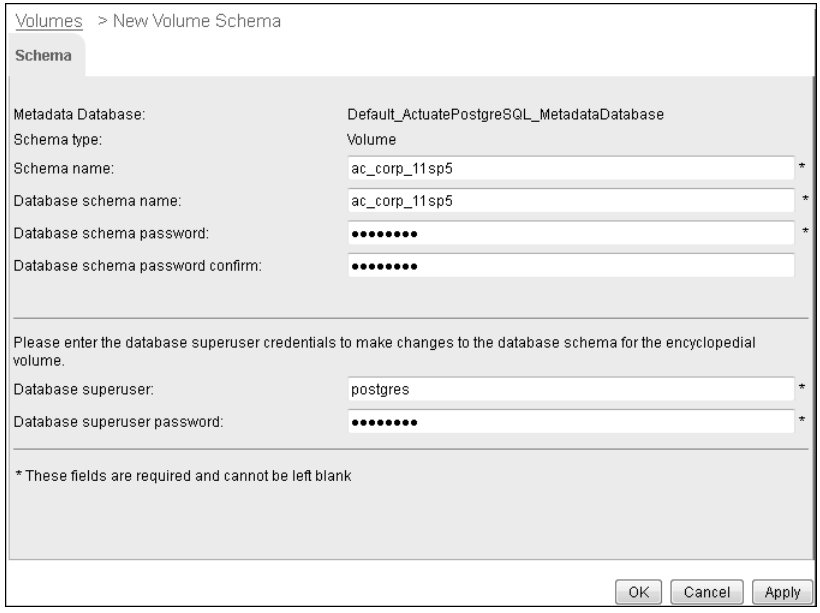

Figure 4-53 Creating a new volume schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-54.

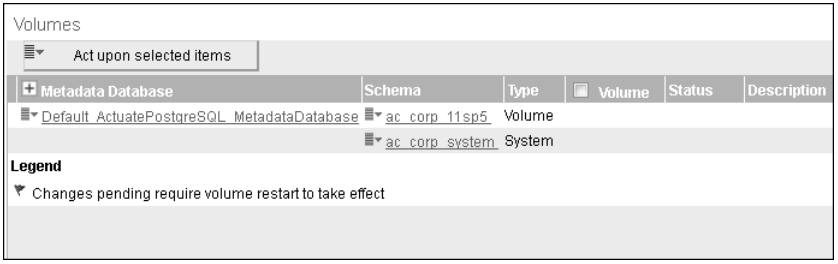

Figure 4-54 Viewing the new schema

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 = /home/Actuate/AcServer
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 = /home/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.sh script using the following command line syntax:

```
sh ./administrate_encyclopedia_data_store.sh
   VolumeImport.properties
```

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

```
sh ./admineds.sh 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.sh. 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.sh 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.sh 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.sh. The Data Store Administrator utility also creates a file containing parameters that ac_oracle_load.sh uses, named ac_oracle_load.par. You edit this file to specify parameter values before executing ac_oracle_load.sh

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.sh 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.sh. The Data Store Administrator utility also creates a file which ac_sql_server_load.sh calls, ac_sql_server_load.sql. You edit both files to provide configuration information before executing ac_sql_server_load.sh.

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.sh 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.sh 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.sh. 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.sh.

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.sh 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 4-55.

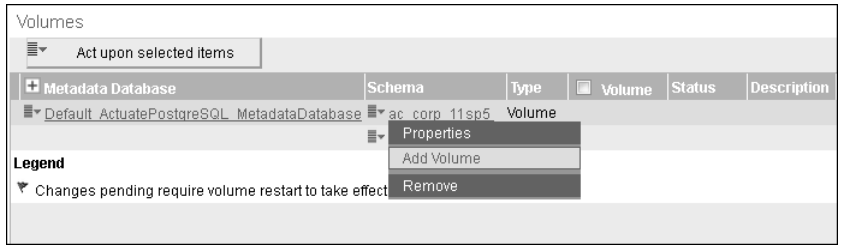

Figure 4-55 Choosing to add a volume

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," if the partition for this volume is not DefaultPartition. If the migrated volume uses DefaultPartition, specify DefaultPartition in Primary partition, as shown in Figure 4-56.

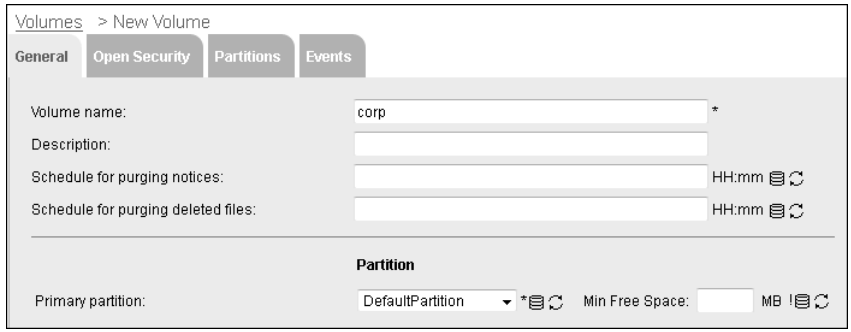

Figure 4-56 Specifying volume and partition name

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 4-57.

Figure 4-57 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 4-58.

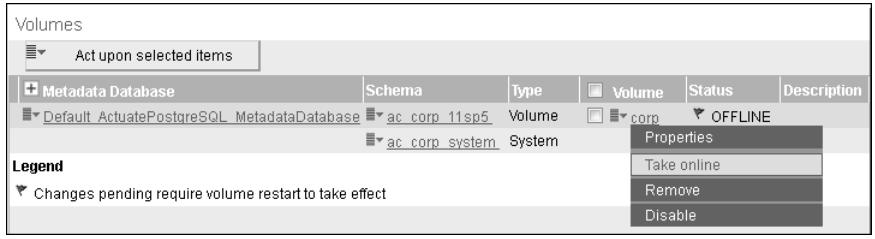

Figure 4-58 Viewing the new volume

The new volume comes online, as shown in Figure 4-59.

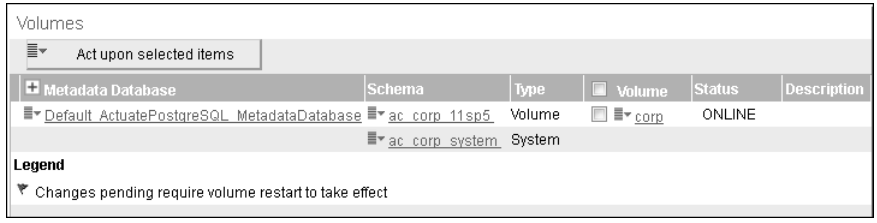

Figure 4-59 Viewing the new volume after it comes online

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 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:

```
/home/Actuate/AcServer/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 4-60.

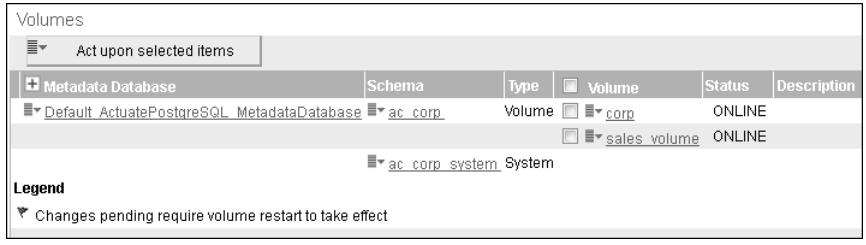

Figure 4-60 Viewing current system schema name, ac_corp_system

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 4-61.

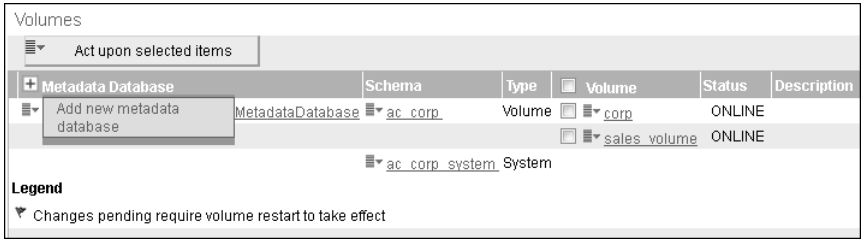

Figure 4-61 Choosing to create a metadata database

3 On New Metadata Database, as shown in Figure 4-62, 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.

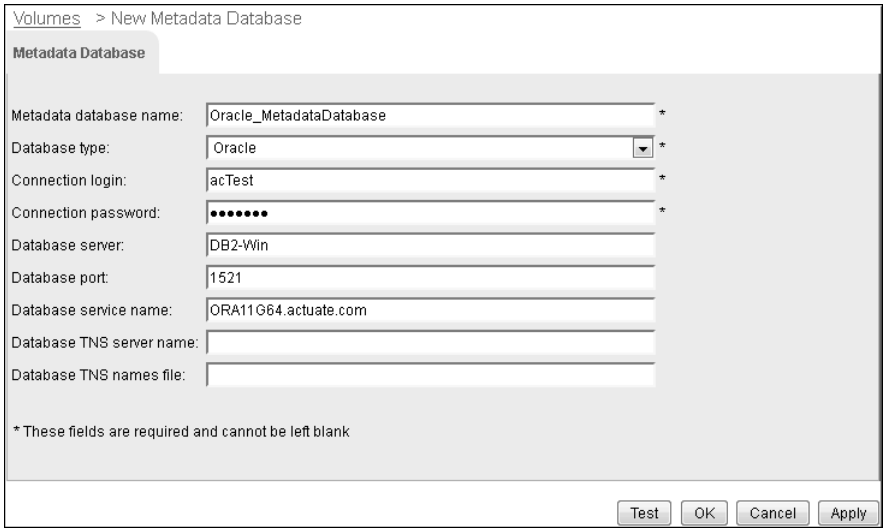

Figure 4-62 Specifying values for a new Oracle metadata database

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 4-63.

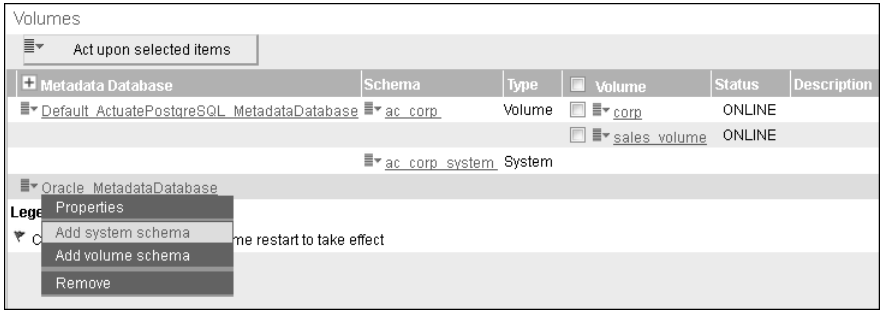

Figure 4-63 Choosing to create a new system schema

2 On New System Schema, as shown in Figure 4-64, 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.

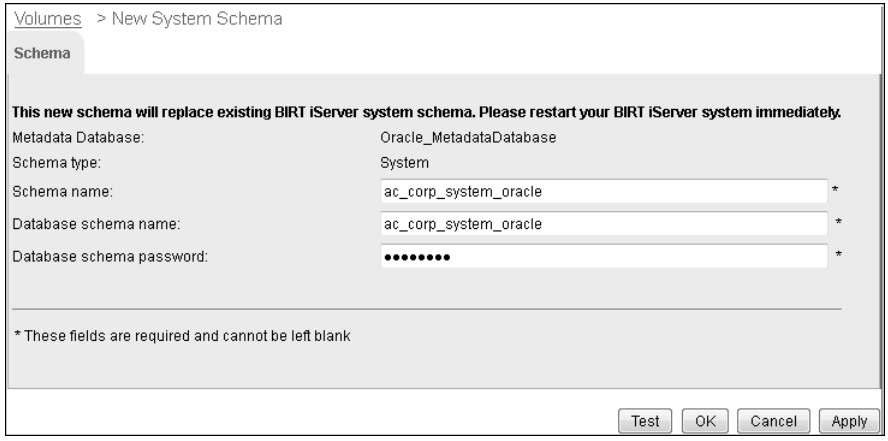

Figure 4-64 Creating a new system schema

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 4-65.

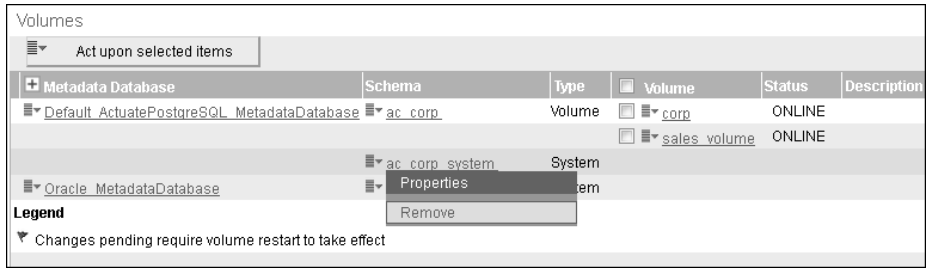

Figure 4-65 Removing the system schema

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 4-66.

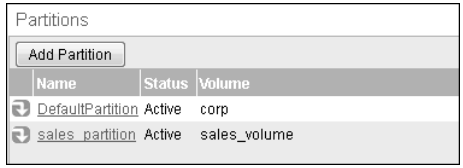

Figure 4-66 Viewing partition and volume pairs

2. Choose Volumes from the side menu. Point to the icon next to a volume and choose Take offline, as shown in Figure 4-67.

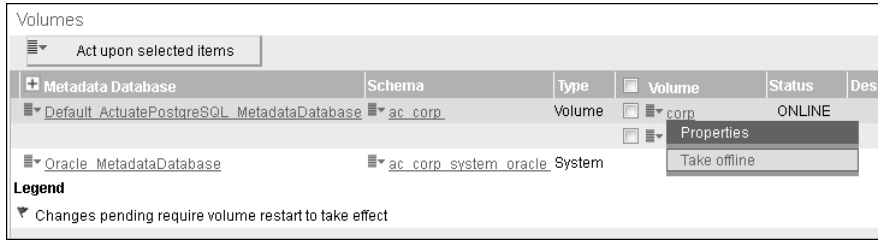

Figure 4-67 Taking a volume offline

3. Point to the icon next to the volume again and choose Remove, as shown in Figure 4-68.

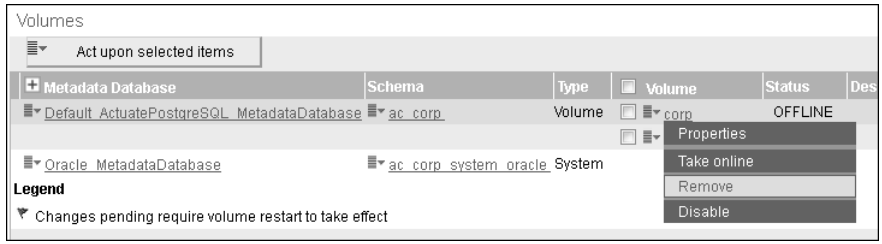

Figure 4-68 Removing a volume

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 4-69.

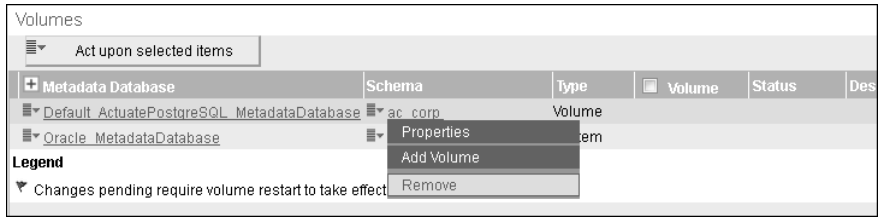

Figure 4-69 Removing a volume schema

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 4-70.

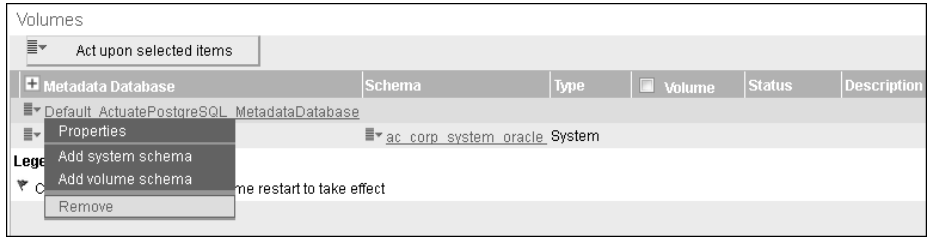

Figure 4-70 Removing a metadata database

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 4-71.

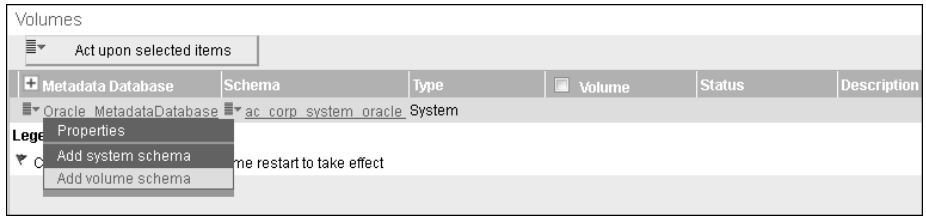

Figure 4-71 Choosing to create a volume schema

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 4-72.

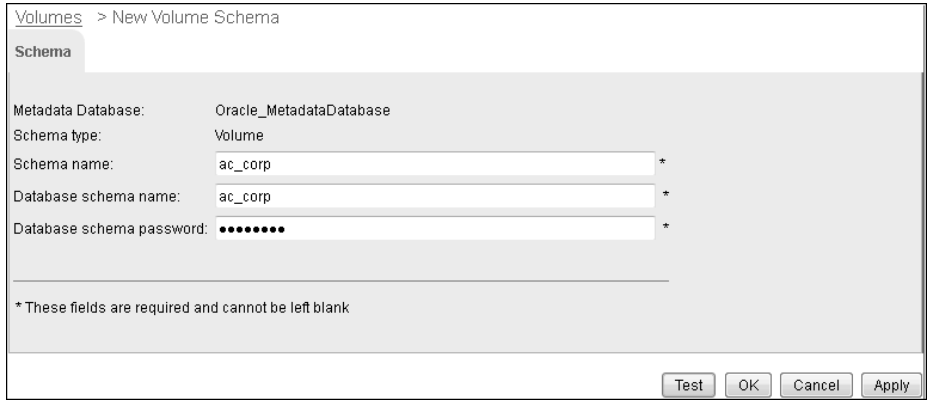

Figure 4-72 Creating a new volume schema

Volumes appears as shown in Figure 4-73.

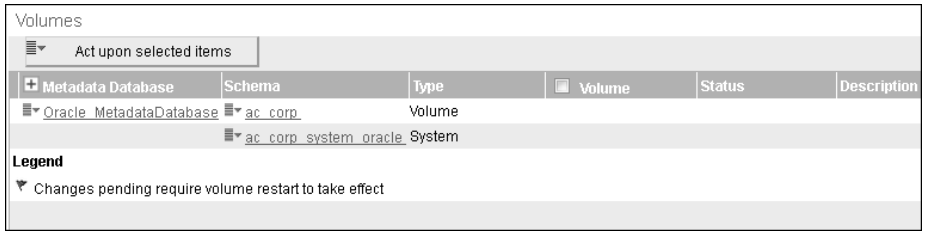

Figure 4-73 Viewing the new volume schema on Volumes

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:

   ```
   /home/Actuate/AcServer/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 = /home/Actuate/AcServer
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 = /home/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 administrate_encyclopedia_data_store.sh file using the following command line syntax:

```
sh ./administrate_encyclopedia_data_store.sh
   VolumeExport.properties
```

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

```
sh ./admineds.sh 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.sh 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 4-74.

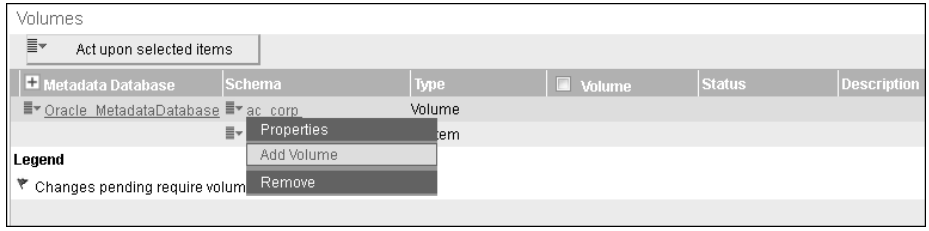

Figure 4-74 Choosing to add a volume

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 4-75.

Figure 4-75 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 4-76. 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*.

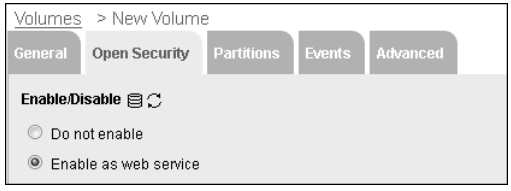

Figure 4-76 Selecting Enable as web service if using an RSSE application

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 4-77.

Figure 4-77 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 4-78.

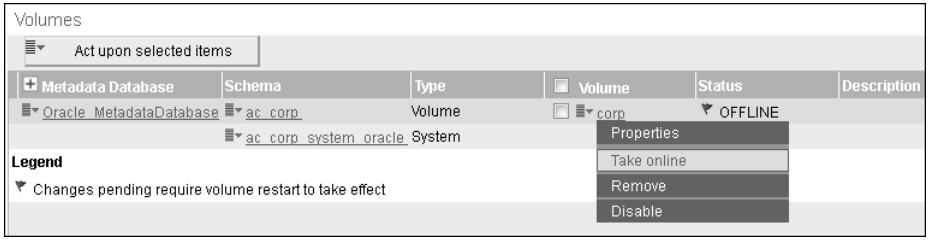

Figure 4-78 Viewing the new volume

The new volume comes online.

6 Repeat steps 1 through 4 for each migrated volume. The online volumes appear as shown in Figure 4-79.

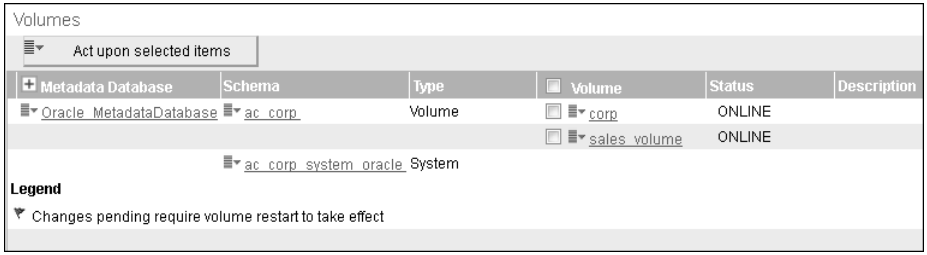

Figure 4-79 Viewing the new volume after it comes online

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 migration 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, or PostgreSQL.
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

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

Parameter	Description	Default value	Supported databases
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: ■ APPLICATION_USER ■ DATABASE_HOST ■ DATABASE_INSTANCE ■ DATABASE_NAME ■ DATABASE_PORT ■ DATABASE_TYPE ■ ORACLE_TNS_NAMES_FILE ■ SCHEMA_NAME	False	All
CREATE_SCHEMA	Set to true to create a new schema, false to use an existing schema. Specifies whether to create a new Encyclopedia schema. PostgreSQL only, not supported for DB2 or Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL
DATA_EXPORT _FOLDER	Absolute path to the folder to which data is exported. This parameter is required if EXPORT_DATA is true. Folder is on the database machine.		All
DATA_EXPORT _FORMAT	Format of exported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE _TYPE}	All
DATA_IMPORT _FOLDER	Absolute path to the folder from which data is imported. This parameter is required if IMPORT_DATA is true. For PostgreSQL, the data must be on the database server.		All

(continues)

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

Parameter	Description	Default value	Supported databases
DATA_IMPORT _FORMAT	Format of imported data. Specify DB2, Oracle, or PostgreSQL.	{DATABASE _TYPE}	All
DATABASE_HOST	Hostname or IP address of the machine hosting the database. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
DATABASE_NAME	Database name for PostgreSQL. For Oracle not using TNS, the database service name. For Oracle using TNS, the TNS net service name.		All
DATABASE_PORT	Port that the database server uses. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports DB2, Oracle, and PostgreSQL in Linux/UNIX.		All
DEFAULT_DATABASE _NAME	Used by the superuser to connect to the default database in order to create the iServer application database. Required when the application database does not exist. PostgreSQL only.		PostgreSQL
DELETE_ALL_DATA	Set to true to delete all data from the data store. Does not delete data from other tables in the schema.	False	All
DELETE_DATA	Set to true to delete filtered data from the data store. This parameter only works when a filter value is available.	False	All

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

Parameter	Description	Default value	Supported databases
DEPOPULATE _SCHEMA	Set to true to delete data store elements such as tables, views, stored functions, and procedures from the schema. Does not remove any other objects in the schema.	False	All
DROP_SCHEMA	Set to true to delete the schema.	False	PostgreSQL
EXPORT_ALL_DATA	Set to true to export all volumes from a schema.	False	All
EXPORT_DATA	Set to true to export a single volume from a schema. Use VOLUME_NAME to specify the name of the volume to export. Optionally, use NEW_VOLUME _NAME to specify a new name for the exported volume.	False	All
FILE_LOG_LEVEL	Minimum logging level for messages sent to the log file. This parameter only works for AcLogger, and is a standard JUL Level name. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
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

(continues)

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

Parameter	Description	Default value	Supported databases
LOG_FILE_NAME	Name of the log file. Do not add a file extension. The extension is set to .log. A unique number is appended automatically to the file name to prevent overwriting previous logs.	Encyclopedia DataStore Administrator <number>.log	
LOG_FILE_SIZE	Approximate maximum size of log files, in byte, before a new log file starts. Must be a valid integer. Data Store Administrator tools do not generate large logs.	10,000,000	All
LOG_FOLDER	Full path of folder to write logs.	AC_DATA _HOME /server/log.	
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

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

Parameter	Description	Default value	Supported databases
POPULATE_SCHEMA	Set to true to populate the schema with data store elements such as tables, views, stored functions, and procedures. Set to false if the data store elements already exist. Use Configuration Console to populate schema when possible. Schema automatically populates when defined in Configuration Console.	CREATE _SCHEMA	All
POPULATE_SCHEMA _PHASE_ONE	Set to true to create the schema objects necessary for data to be loaded. Set to false to prevent this schema object creation. Imports data without building indexes for fast load.	POPULATE _SCHEMA	All
POPULATE_SCHEMA _PHASE_TWO	Set to true to create schema objects phase one does not create. Set to false to prevent this schema object creation. Builds indexes and other structures in database.	POPULATE _SCHEMA	All
PROMPT_FOR _PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it as a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	True	All
SCHEMA_NAME	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	NEW _SCHEMA _NAME	All

(continues)

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

Parameter	Description	Default value	Supported databases
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.sh
  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.sh 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.sh 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.sh
    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.sh.

  Before running ac_db2_load.sh 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.sh
  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.sh
    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.sh. Modify this file to contain the appropriate server, user, password, port, and database information before running this script.

  Before running ac_postgresql_client_load.sh 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 = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {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 = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {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 chapter.

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

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
EXPORT_ALL_DATA = true
NEW_SCHEMA_NAME = <provide a name>
DATA_EXPORT_FOLDER = home/Projects/DataStores/Data
DATA_EXPORT_FORMAT = Oracle
```

Exporting a single volume from a schema

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 chapter.

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

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
```

```
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
EXPORT_DATA = true
VOLUME_NAME = <provide a name>
NEW_SCHEMA_NAME = <provide a name>
NEW_VOLUME_NAME = <provide a name>
DATA_EXPORT_FOLDER = home/Projects/DataStores/Data
DATA_EXPORT_FORMAT = Oracle
```

Deleting All Volumes from a Schema

Perform this operation to remove all volumes from the schema that SCHEMA_NAME specifies.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
DELETE_ALL_DATA = true
```

Deleting a single volume from a schema

Perform this operation to remove the volume that VOLUME_NAME specifies from the schema that SCHEMA_NAME specifies.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
DELETE_DATA = true
VOLUME_NAME = <provide a name>
```

Creating a new Volume in an empty schema

iServer performs this operation when you create a new volume in Configuration Console and bring the volume online. When creating a new volume using Encyclopedia Data Store Administrator, set INITIALIZE_DATA to true and specify a name for NEW_VOLUME_NAME.

Setting POPULATE_SCHEMA to true to have Encyclopedia Data Store Administrator create schema elements, such as tables and indexes in an empty schema. The schema owner and application user must already exist and have appropriate privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Populating an empty schema

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

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
```

Creating a new volume in a populated schema

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

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

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Creating a new volume in a new schema

This operation is only supported for PostgreSQL. iServer typically performs this operation when you create a new volume in Configuration Console and bring the volume online.

Creating a schema requires superuser privileges. Encyclopedia Data Store Administrator creates the database and users if these items do not already exist.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Working with System Data Store Administrator

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

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

Specifying System Data Store Administrator properties

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

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

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

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

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

```
<AC_SERVER_HOME>\bin
```

where AC_SERVER_HOME refers to the Release 11 Service Pack 4 installation path. For example, using the default value for AC_SERVER_HOME, add:

```
/home/Actuate/iServer11/AcServer/bin
```

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

1. Navigate to AC_SERVER_HOME/bin.
2. Create a properties file containing the property definitions necessary to perform an operation, such as creating and populating a system schema, described in "Performing operations using System Data Store Administrator utility," later in this chapter. You pass this file to the System Data Store utility when you execute it.
3. Run the administrate_system_data_store.sh file using the following command line syntax:

   ```
   sh ./administrate_system_data_store.sh systemdatastore
      .properties
   ```

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

   ```
   sh ./adminsds systemdatastore.properties
   ```

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

- Checks to see if the administrator submitted a property file on the command line when running the script.
- If the administrator does not submit an argument, the script echoes a usage statement that describes the command-line syntax.
- Calls the set_tools_environment.sh script, which sets the environment variables.
- Executes the System Data Store Administrator utility using the name of the properties file as an argument.

Listing 4-2 administrate_system_data_store.sh

```
#!/bin/sh

if [ "x$1" = "x" ]; then
   echo "Usage: administrate_system_data_store.sh <properties
   file name>"
   exit 1
fi

# Set up environment variables
. `dirname $0`/set_tools_environment.sh

# Administrate data store
java com.actuate.iserver.system.datastore.admin
   .SystemDataStoreAdministrator "$PROPERTY_FILE"
```

The SystemDataStoreAdministrator class has the same parent class as the Encyclopedia Data Store Administrator and uses the same property settings. System Data Store Administrator properties include the following categories:

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

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

Configure the properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
DEFAULT_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
INITIALIZE_DATA = true
```

Working 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.sh 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.

Table 4-3 Required Encyclopedia Data Store Upgrader properties

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

Table 4-4 describes the optional Encyclopedia Data Store Upgrader properties to specify in the upgrade_encyclopedia_data_store.sh 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
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

(continues)

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

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

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

Parameter	Description	Default value	Supported databases
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

Chapter

5

Installing a BIRT iServer cluster

This chapter discusses the following topics:

- Installing a BIRT iServer cluster node
- Preparing to install an iServer cluster
- Performing a cluster installation using the wizard
- Adding a node to a cluster
- Finding the BIRT iServer home directory
- About the Java Development Kit

Installing a BIRT iServer cluster node

A node is a machine running 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:

- Shared configuration home and Encyclopedia volume directories
- Cluster resources, such as printers, database systems, and disk storage systems

It is the responsibility of the administrator performing the installation to make sure that all network sharing settings conform to the security policies in force for the environment.

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

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

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

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

iServer instances running on multiple machines maintain iServer system and Encyclopedia volume metadata in databases and control access to shared volume data. The volume data can be on machines that are not running iServer, but must be shared and accessible to each iServer instance. It is the responsibility of the administrator performing the installation to make sure that all network sharing settings conform to the security policies in force for the environment.

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

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

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

Preparing to install an iServer cluster

When you create a BIRT iServer cluster, you must install and run all cluster nodes using the same administrative user account.

Creating an administrative user account

Before installing iServer, create a user account with the privileges to access the relevant files and directories. Like other Linux and UNIX processes, the processes that perform BIRT iServer tasks run under a specific account.

Actuate recommends creating a dedicated user account for installing and running iServer. Having a dedicated user account isolates iServer-specific issues and events on a machine, making it easier to administer the environment. If you exercise the same control over the operating system account for BIRT iServer that your site exercises for other system administrator and root accounts, you can maintain the same level of security.

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

Installing X-frame buffer

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

Actuate distributes Xvfb for the Sun and AIX operating systems. The installation and configuration of Xvfb are BIRT iServer installation options in these environments.

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

The Xvfb software requires:

- X libraries installed on the iServer machine

  If you choose to install Xvfb, the installation script searches for the required libraries and displays a message if the install script cannot find the required libraries on the machine.

- Variables set to the path of the Xvfb libraries
  - XVFBDISPLAY variable in start_srvr.sh.
  - display_value in pmd11.sh.

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

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

```
# setenv DISPLAY :0.0
```

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

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

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

Performing a cluster installation using the wizard

The following section describes how to install an iServer Release 11 cluster node in the Linux or UNIX operating system using a Linux system as the example.

How to install a cluster node in Linux

1. Download the required files. Extract the files.
2. To install the server files, execute the isinstall script:

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

   The script displays a series of prompts. Respond to the prompts as described in the following procedures.
3. The license agreement appears, as shown in Figure 5-1.

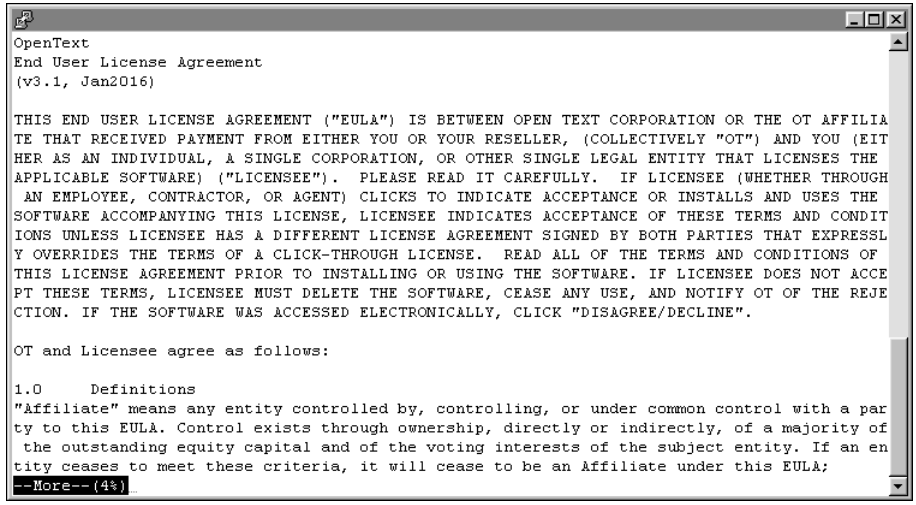

Figure 5-1 Reviewing the license agreement

4. Read the license agreement and press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 5-2.

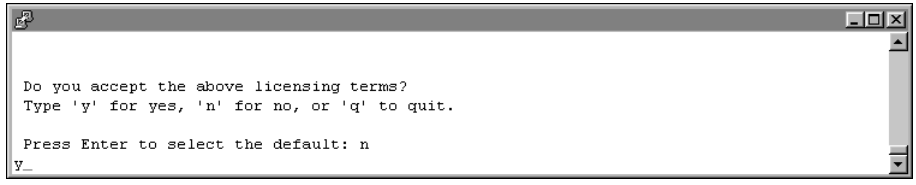

Figure 5-2 Accepting the licensing terms

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

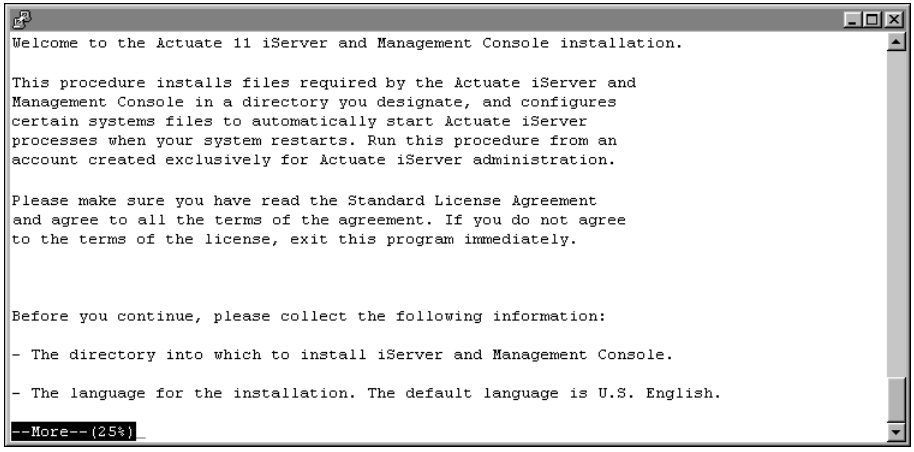

Figure 5-3 Reviewing the introductory information

6 Press Enter after reviewing the introductory information, as shown in Figure 5-4.

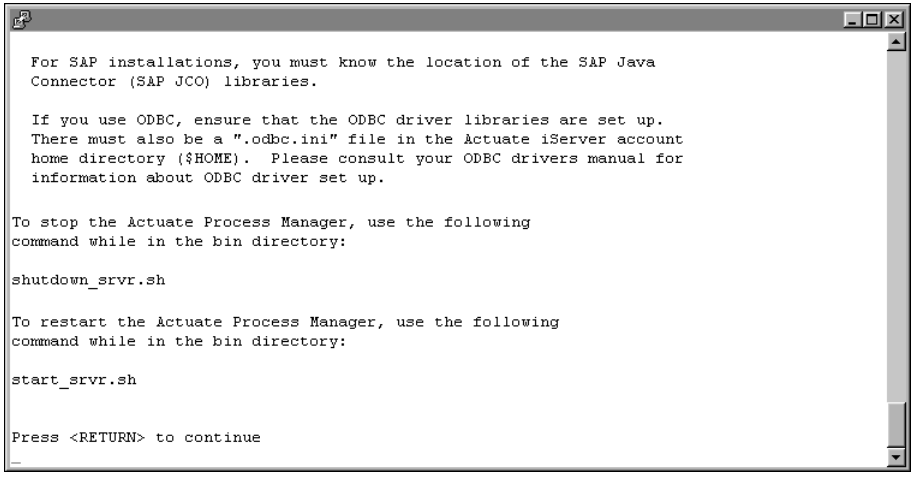

Figure 5-4 Finishing the review of introductory information

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

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

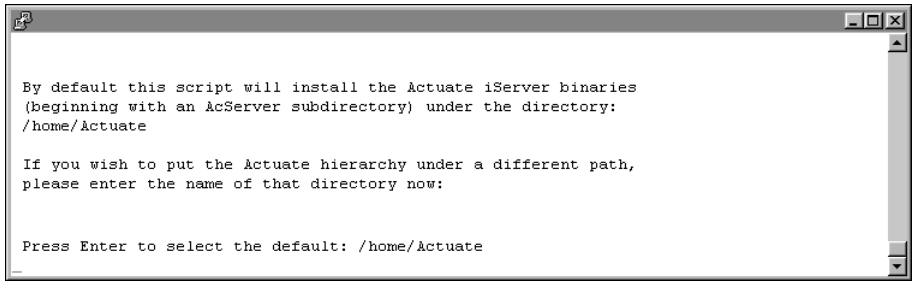

Figure 5-5 Specifying the installation directory

8 Press Enter to accept the default installation directory for data, AC_SERVER_HOME/data, for the iServer data as shown in Figure 5-6. Alternatively, choose a different directory.

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

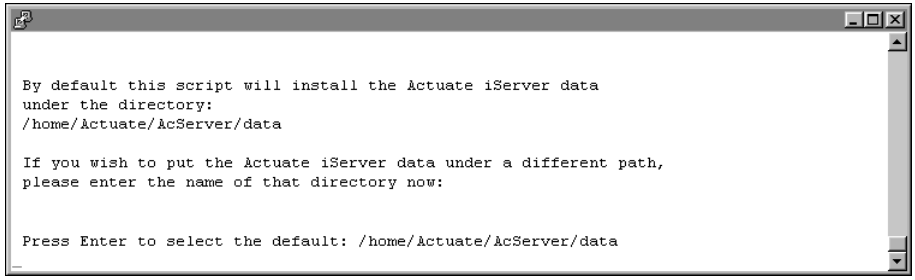

Figure 5-6 Specifying the data installation directory

9 Press Enter to accept the default option of creating the directory for data, as shown in Figure 5-7. Alternatively, type n for no, or q to quit, and press Enter.

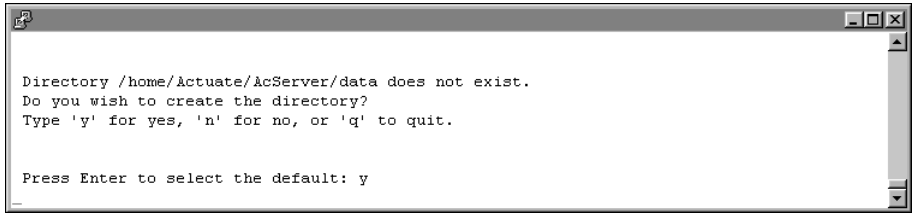

Figure 5-7 Creating the AC_DATA_HOME directory

10 The installer copies prerequisite files to the destination directory, as shown in Figure 5-8. After copying the prerequisite files, the installation continues.

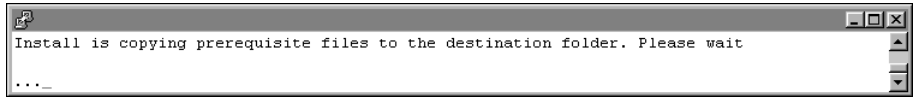

Figure 5-8 Copying prerequisite files

11 Press Enter to choose the default option, Server with Management Console, as shown in Figure 5-9. A cluster node must have access to Configuration Console. Configuration Console installs with Management Console.

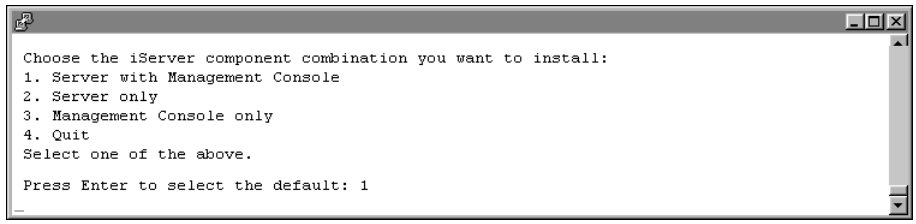

Figure 5-9 Choosing the components to install

12 Type 1 and press Enter to select Cluster Server for installation, as shown in Figure 5-10. Alternatively, choose a different type of iServer to install.

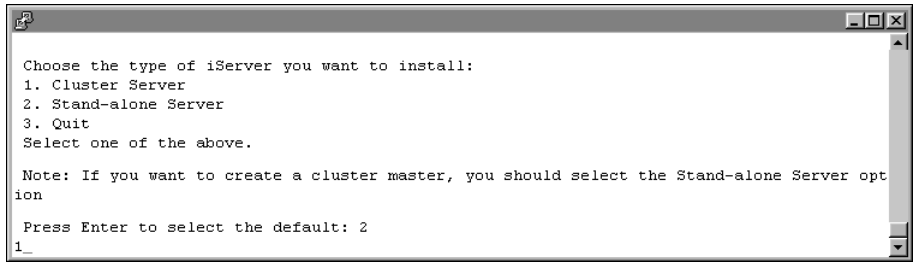

Figure 5-10 Specifying the type of iServer to install

13 Specify the path to the configuration home location, and press Enter, as shown in Figure 5-11.

The configuration home location is the shared directory of the cluster containing the files, acserverconfig.xml, acserverlicense.xml, and acconfigowner.lock.The configuration files are located in AC_DATA_HOME /config/11SP5 by default. For more information about configuring network sharing, see "Adding a node to a cluster," later in this chapter.

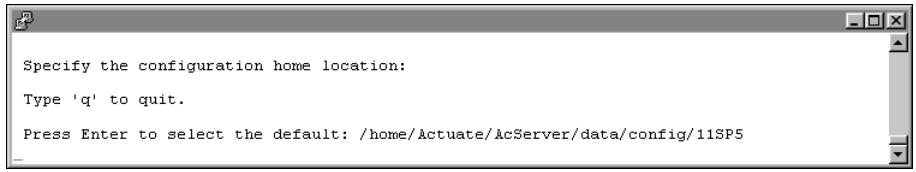

Figure 5-11 Specifying the configuration home location

14 If the specified location for the configuration home directory does not yet exist, you are prompted to create the directory. Press Enter to accept the default option which creates the directory, as shown in Figure 5-12. Alternatively, press n for no, or q to quit.

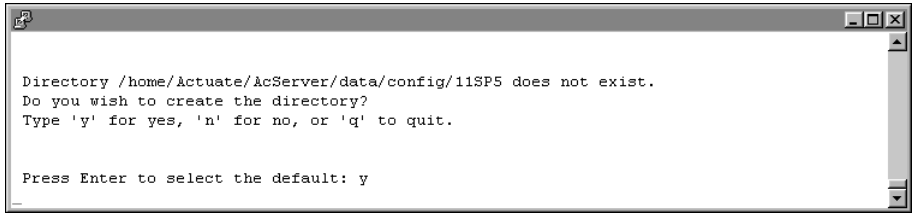

Figure 5-12 Creating the configuration home location

15 If you chose to install Server with Management Console instead of Server only, perform the following steps:

1 Press Enter to select the default locale, which is English (United States), as shown in Figure 5-13. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.

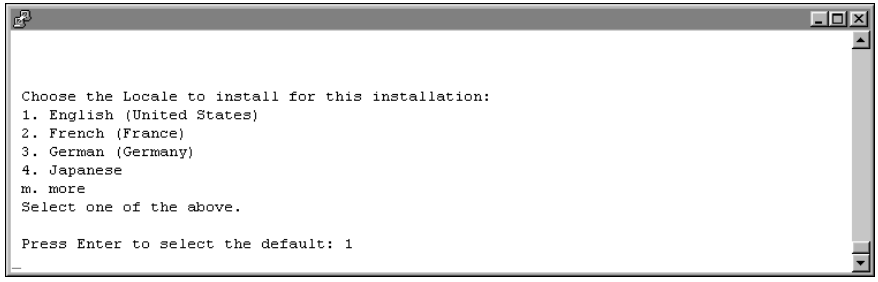

Figure 5-13 Specifying a locale

2 Press Enter to select the default time zone, which is America/Los_Angeles as shown in Figure 5-14. Alternatively, select another time zone from the numbered list.

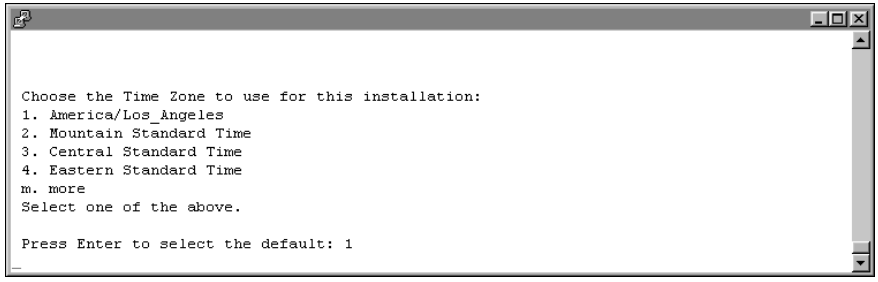

Figure 5-14 Specifying a time zone

16 To evaluate the product using the included evaluation software license, press Enter, as shown in Figure 5-15. Alternatively, type 2, then type the path to the license file that you purchased.

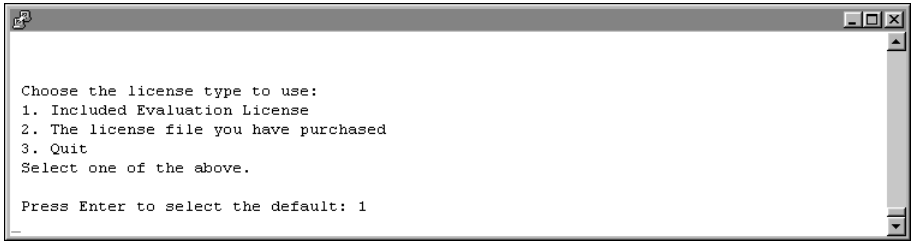

Figure 5-15 Specifying license type

17 Press Enter to select the default host name, the name of your machine, where the Process Management Daemon (PMD) runs, as shown in Figure 5-16. Alternatively, type a different IP address or hostname.

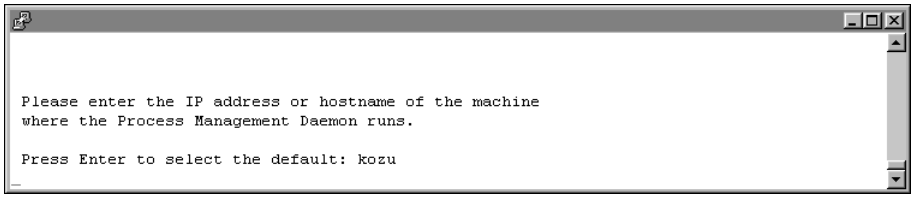

Figure 5-16 Specifying the machine on which the PMD runs

18 Press Enter to accept the default port number where the Process Management Daemon (PMD) listens for requests, as shown in Figure 5-17. Alternatively, type a different port number.

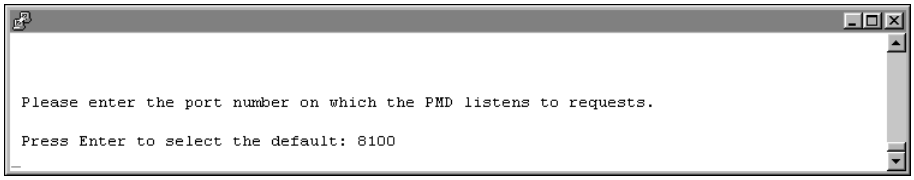

Figure 5-17 Specifying the port number on which the PMD listens

19 Specify the administrator password, as shown in Figure 5-18.

You use this password to log in to the iServer Configuration Console.

Figure 5-18 Specifying the administrator password

20 Re-enter the password of the administrator, as shown in Figure 5-19.

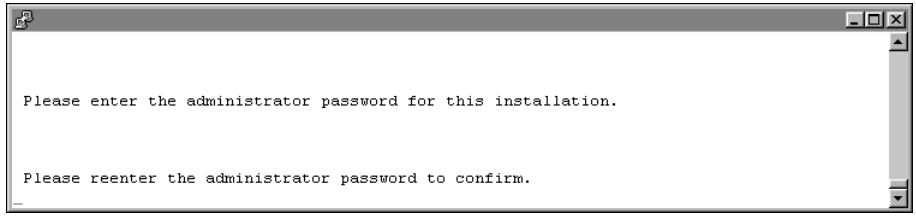

Figure 5-19 Re-entering the administrator password

21 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 5-20. Alternatively type y for yes, specify the database drivers/clients you wish to use and press Enter.

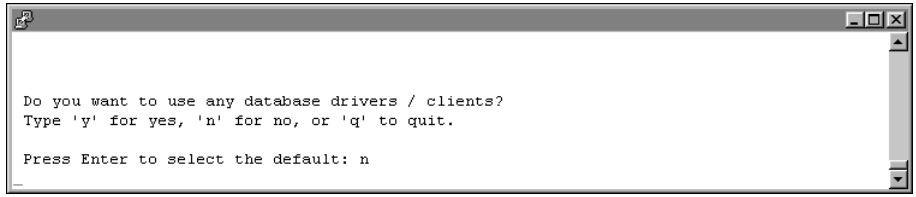

Figure 5-20 Specifying whether to use database drivers/clients

22 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 5-21. For more information about installing X-server, see "Installing X-frame buffer," earlier in this chapter.

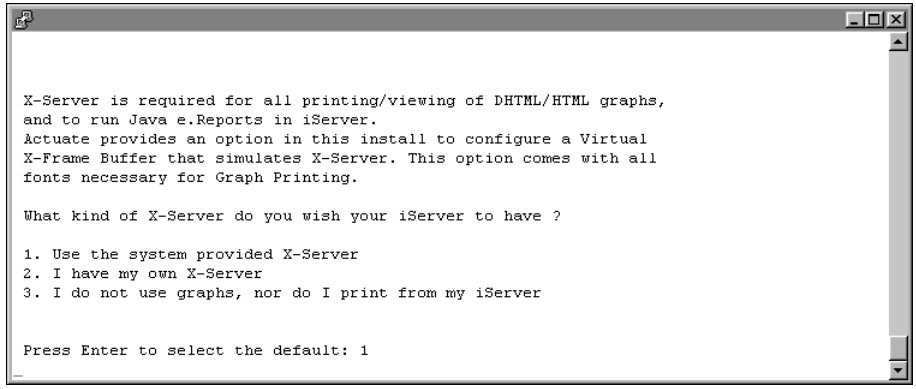

Figure 5-21 Specifying what kind of X-Server to use, if any

23 If you chose to install Server with Management Console instead of Server only, perform the following steps:

1. Press Enter to accept the default hostname, the name of your machine, that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 5-22. Alternatively, type a different IP address.

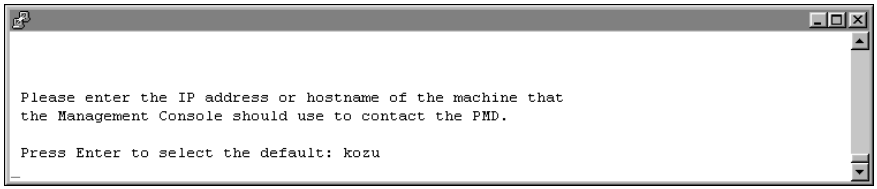

Figure 5-22 Specifying the machine host name that Management Console uses to contact the PMD

2. Press Enter to accept the default port number, 8100, on which the Process Management Daemon (PMD) listens for requests from Management Console, as shown in Figure 5-23. Alternatively, type a different port number.

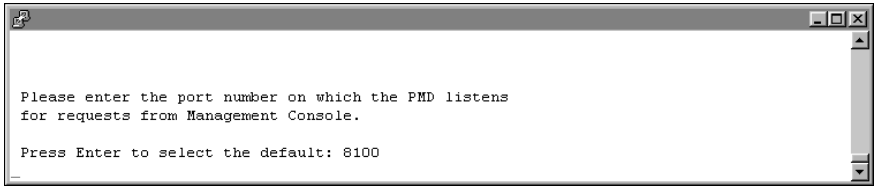

Figure 5-23 Specifying the port number for the PMD to listen for requests from Management Console

3. Press Enter to accept the default hostname, the name of your machine, as shown in Figure 5-24. Alternatively, type a different IP address.

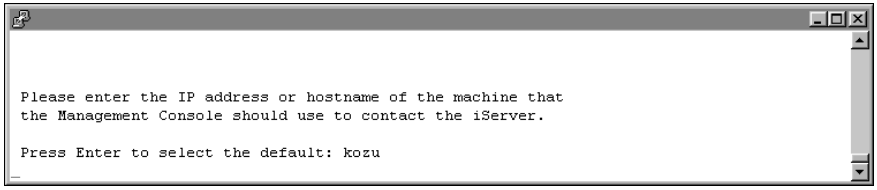

Figure 5-24 Specifying the name Management Console uses to contact iServer

4. Press Enter to accept the default port number, 8000, as shown in Figure 5-25. Alternatively, type a different port number.

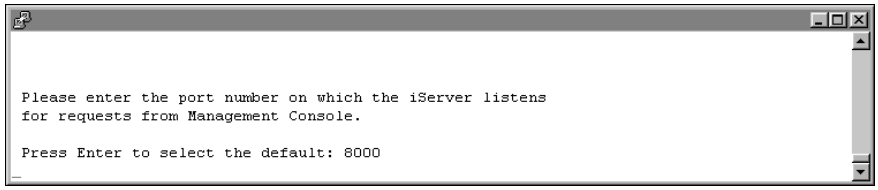

Figure 5-25 Specifying the port number for listening for requests from Management Console

5 Press Enter to accept the default name, the name of your machine, for the Encyclopedia volume to use with Management Console, as shown in Figure 5-26. Alternatively, type a different name for the Encyclopedia volume.

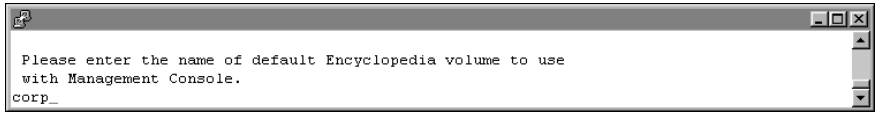

Figure 5-26 Specifying the name of the default Encyclopedia volume

6 Press Enter to accept the default name, acadmin, for the HTTP server context root, as shown in Figure 5-27. Alternatively, type a different name.

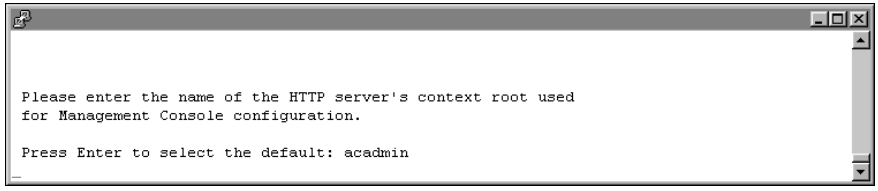

Figure 5-27 Specifying the name of the HTTP server context root

24 Press Enter to accept the default port number, 8900, on which the application container listens for requests, as shown in Figure 5-28. Alternatively, choose a different port.

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

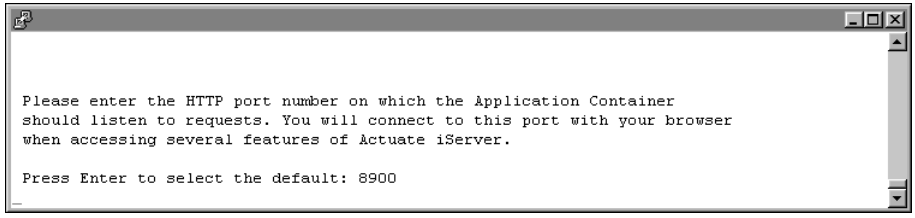

Figure 5-28 Specifying the application container listening port number

25 Review the settings, as shown in Figure 5-29, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.

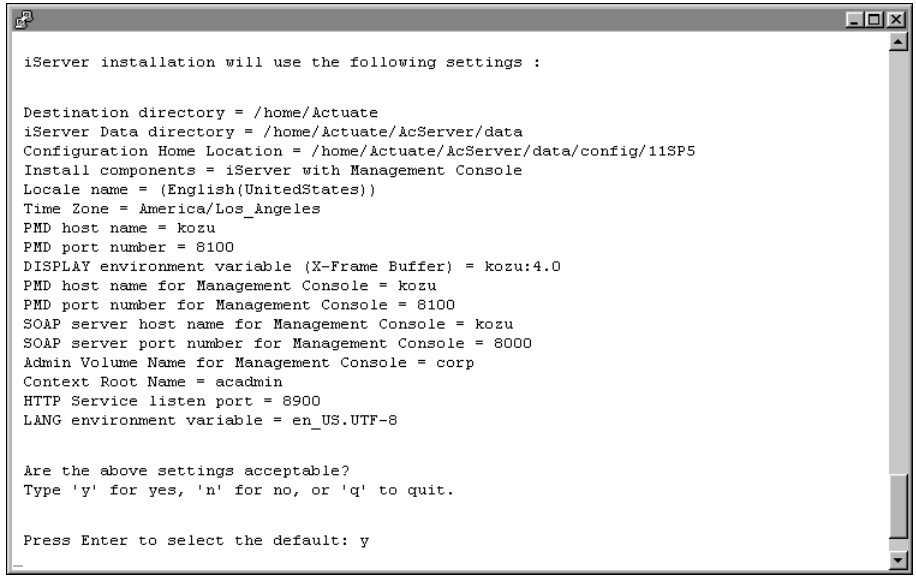

Figure 5-29 Reviewing settings for a Server with Management Console install

26 The installation program installs iServer, and displays an indicator showing how the installation is progressing, as shown in Figure 5-30.

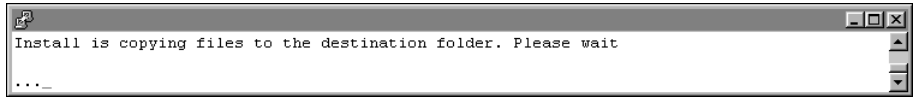

Figure 5-30 Copying iServer files to your destination folder

27 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start the Process Management Daemon (PMD), as shown in Figure 5-31.

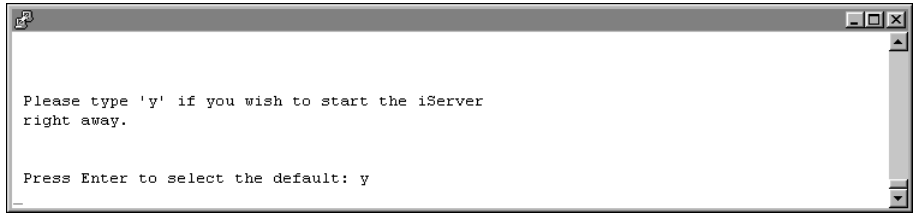

Figure 5-31 Specifying whether to start iServer

28 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 5-32.

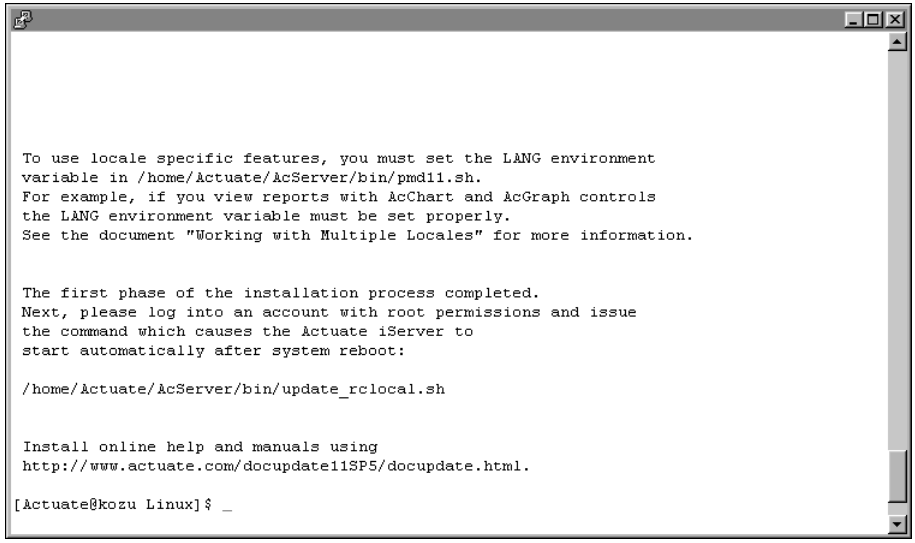

Figure 5-32 Viewing information about localization, logging in, and installing online help

Adding a node to a cluster

After installing a node on a machine, the administrator must still configure sharing and add the node to the cluster. When adding a node to a cluster setup, the administrator must verify that the configuration home directory specified during the install procedure points to the shared configuration home directory and all Encyclopedia volume resources are accessible.

The following section refers to the machine containing the shared configuration directory as node1 and the cluster node accessing these shared resources as node2.The following example assumes that both the configuration folder and Encyclopedia volume folders are located on node1, although in a more complex installation, these configuration and volume resources may reside in another network location.

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

- On node1, the Administrator shares the configuration folder and any Encyclopedia volume folders that a cluster node accesses.

- On node2, the Administrator:
  - Creates folders on which to mount the node1 shared folders
  - Creates a mapping between the node1 and node2 shared folders
  - Mounts the node1 shared folders on the node2 machine

It is the responsibility of the administrator performing the installation to make sure that all settings conform to the security policies in force for the environment.

The following instructions provide a basic reference example of the operations required to configure folder sharing in a Linux environment that supports using the Network File System (NFS), a common, standard, distributed file system protocol.

How to share the configuration and Encyclopedia volume files and folders

In a default iServer Release 11 Service Pack 5 installation, a cluster node requires shared, read-write access to the following system resources:

- AC_DATA_HOME/config/11SP5
  In a Release 11 Service Pack 5 installation, the configuration files are located in AC_DATA_HOME/config/11SP5.
- AC_DATA_HOME/encyc or other volumes, including all file, fileType, status, and tempRov subfolders
  In an iServer installation, where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes.

To give a cluster node read-write access to these files and folders, perform the following tasks:

1. Log in to node1 as the root user.
2. Add the following entries to the /etc/exports file:

```
/home/actuate/AcServer/data/config/11SP5
   *(rw,fsid=1,no_root_squash)
/home/actuate/AcServer/data/encyc
   *(rw,fsid=2,no_root_squash)
```

3. Start the NFS server processes by executing the following command:

```
service nfs restart
```

4. Log in to node2 as the actuate user.
5. Create the following directory paths:

```
/home/actuate/AcServer/data/config/11SP5
/home/actuate/AcServer/data/encyc
```

6 Log off node2.

7 Log in to node2 as the root user.

8 Add the following entries to the /etc/fstab file:

```
<node1 hostname>:/home/actuate/AcServer/data/config/11SP5
   /home/actuate/AcServer/data/config/11SP5 nfs nfsvers=3 0 0
<node1 hostname>:/home/actuate/AcServer/data/encyc /home/
   actuate/AcServer/data/encyc nfs nfsvers=3 0 0
```

9 Mount the shared folders on node1 by executing the following commands:

```
mount /home/actuate/AcServer/data/config/11SP5
mount /home/actuate/AcServer/data/encyc
```

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

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

How to verify and edit acpmdconfig.xml file settings

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

1 Shut down the recently installed cluster node.

2 Using a text editor, open acpmdconfig.xml, which by default is located in AC_SERVER_HOME/etc.

3 Verify or edit <AC_CONFIG_HOME> to point to the shared configuration home directory for the cluster, as shown in the following code:

```
<AC_CONFIG_HOME>/home/actuate/AcServer/data/config/11SP5
   </AC_CONFIG_HOME>
```

This location is the path that you specified for the configuration home directory during the install procedure.

4 Verify or edit <AC_TEMPLATE_NAME> to specify the server template name from the available server templates listed in the shared acserverconfig.xml file, as shown in the following code:

```
<AC_TEMPLATE_NAME>urup</AC_TEMPLATE_NAME>
```

In the example, urup is the server template name

5 Save acpmdconfig.xml.

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

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

How to verify and edit acserverconfig.xml file settings

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

1. Stop Actuate BIRT iServer 11 on the node that contains the shared configuration home directory.
2. Using a text editor, open the acserverconfig.xml file in the configuration home directory.

   In a Release 11 Service Pack 5 installation, the configuration files are located in AC_DATA_HOME/config/11SP5 by default. The location is the path that you specified for the configuration home directory during the install procedure.
3. In <Template> settings for the node, verify or edit <ServerFileSystemSettings> to make sure the path <ServerFileSystemSetting> points to the location that contains the Encyclopedia data files, by performing the following tasks:
   1. Locate the <ServerFileSystemSettings> element under the <Template> element.
   2. In <ServerFileSystemSetting>, locate:

      ```
      <ServerFileSystemSettings>
         <ServerFileSystemSetting
            Name="DefaultPartition"
            Path="$AC_DATA_HOME$/encyc"/>
      </ServerFileSystemSettings>
      ```
   3. Change Path from the AC_DATA_HOME variable notation to the full path specification, as shown in the following code:

      ```
      <ServerFileSystemSettings>
         <ServerFileSystemSetting
            Name="DefaultPartition"
            Path="/home/actuate/AcServer/data/encyc"/>
      </ServerFileSystemSettings>
      ```

   The Path setting for DefaultPartition is /home/actuate/AcServer/data /encyc. Do not use the AC_DATA_HOME variable notation.
4. In <MetadataDatabase> settings, verify or edit the <ConnectionProperty> for the server to make sure that it specifies the network name, not localhost, of the node on which the Encyclopedia volume database resides, by performing the following tasks:

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

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

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

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

5. Save acserverconfig.xml.

Start Actuate BIRT iServer 11 on each cluster node. The new cluster node will automatically read the settings in the acserverconfig.xml file in the shared configuration directory to access its template, then join the cluster.

How to start an iServer cluster using Configuration Console

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

1. On the node containing the configuration home directory for the cluster, log in to Configuration Console and choose Advanced view. Choose Servers, then choose Start New Server.
2. On Servers—Start New Server, as shown in Figure 5-33, perform the following tasks:
   1. In Server name, type the name of the cluster node.
   2. In Host Name or IP Address, type the name or IP address of the cluster node.
   3. In iServer Process Manager Port Number, type the Daemon listen port number. The default value for this port is 8100. You specify this port number during the install procedure.
   4. In Server template name, choose the name of the template that the cluster node uses.

      Choose OK.

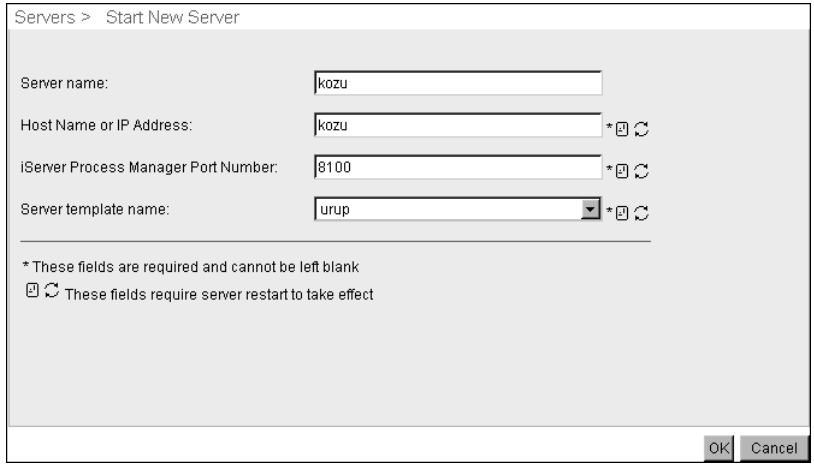

Figure 5-33 Preparing to start a new server

3 Log out of Configuration Console.

4 Restart the Actuate BIRT iServer 11 on the node containing the configuration home directory for the cluster then the new node.

5 Log in to Configuration Console and choose Advanced view. Choose Servers from the side menu. The new cluster node automatically reads the acserverconfig.xml in the shared configuration home directory to access its template, then joins the cluster.

Finding the BIRT iServer home directory

The environment variable for the iServer home directory is AC_SERVER_HOME. The iServer installation program sets the variable to the path of your iServer login environment.

How to find the home directory for BIRT iServer on a Linux or UNIX system

If you use the C Shell on a Linux or UNIX system, you can look in your .cshrc file to see the value of AC_SERVER_HOME:

```
$ setenv AC_SERVER_HOME /usr/local/AcServer
```

If an iServer process is running on the system, you can also use the following ps command piped to the grep command to find the pmd11 executable, which runs from the iServer home bin directory:

```
# ps -ef | grep pmd
```

About the Java Development Kit

The BIRT iServer installation routine installs the Java SDK files under the directory specified in the environment variable AC_SERVER_HOME:

```
AC_SERVER_HOME/jdk160
```

Some operating systems require an operating system upgrade or patch to use JRE 6.0. For information about requirements for your operating system, see your operating system documentation. Also, see the Actuate Support Lifecycle Policy and Supported Products Matrix at the following URL:

https://support.opentext.com

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

- AC_JAVA_HOME
- AC_JVM_HOME
- AC_JRE_HOME in the PMD startup script, pmd11.sh
- AC_JRE64_HOME

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

Chapter

6

Installing Information Console

This chapter discusses the following topics:

- Before you begin
- Installing Information Console on Linux and UNIX

Before you begin

Before installing Information Console, you must prepare the operating system environment to ensure that you have all the necessary software and configuration resources in place. This section describes how to prepare prior to an installation.

About performing a full installation

For Actuate 11, if you perform a full installation for iServer and Information Console, install BIRT iServer System products in the following order:

- BIRT iServer and Management Console
- Information Console
- BIRT iServer Integration Technology

To access online documentation, such as the online help and PDF files of the product manuals, install the documentation files from the following location:

https://support.opentext.com

Installing Information Console on Linux and UNIX

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

Information Console installation requires the following information:

- Installation directory
  $HOME, the account's home directory, is the default installation directory.
- Port used by the Apache Tomcat Information Console service
  If you are using a firewall, ensure that the firewall passes the port number you select. The default port is 8700.
- Value of the SPINLOOPTIME environment variable
  If you use AIX in a multiple-CPU environment, Information Console installation sets value of the SPINLOOPTIME environment variable to 2000. For more information about SPINLOOPTIME, see your AIX documentation.
- Encyclopedia volume name that Information Console accesses
  The default is the current machine.

You can install Information Console in the following ways:

- Use an installation script.

The installation script configures Information Console, creates shortcuts, and extracts and installs all necessary files. Use this option for automated configuration.

- Deploy a WAR file to an Application Server.
  Deploying directly requires that you configure Information Console for your application server. Use this option if your application server supports configuration of an application from a WAR file.

Using the script to install

Complete the steps in the following section to install Information Console using the installation script.

How to install using a script

1. Download the required files. Extract the files.
2. To install the Information Console files, type:

   ```
   sh ./infoconsoleinstall.sh
   ```

   The script displays a series of prompts. Respond to the prompts as described in the following procedures.
3. The license agreement appears, as shown in Figure 6-1.

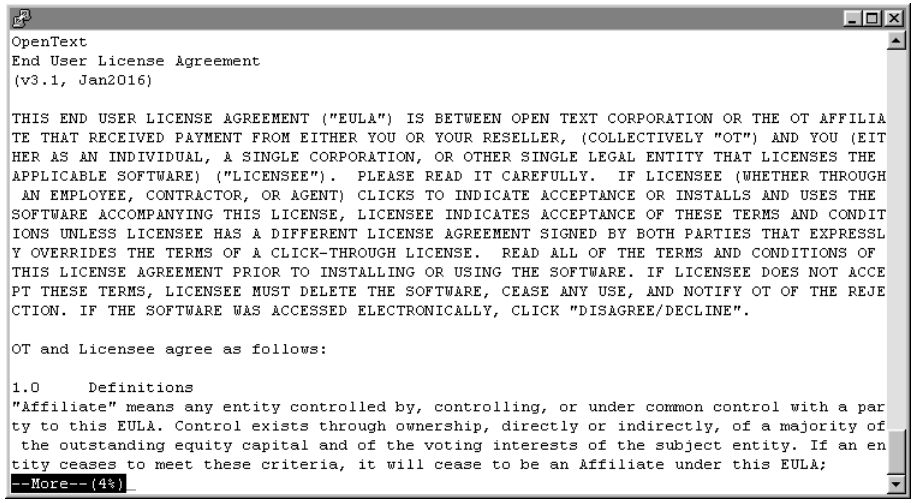

Figure 6-1 The license agreement

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

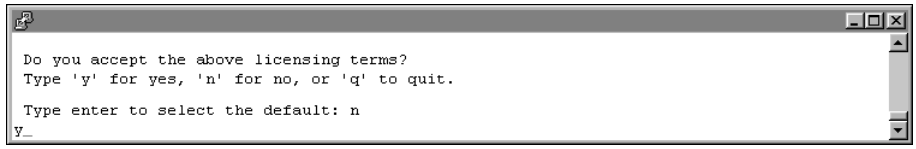

Figure 6-2 Specifying whether you accept the license agreement

5 The introduction to the installation appears, as shown in Figure 6-3. Review the information, then press Enter to continue.

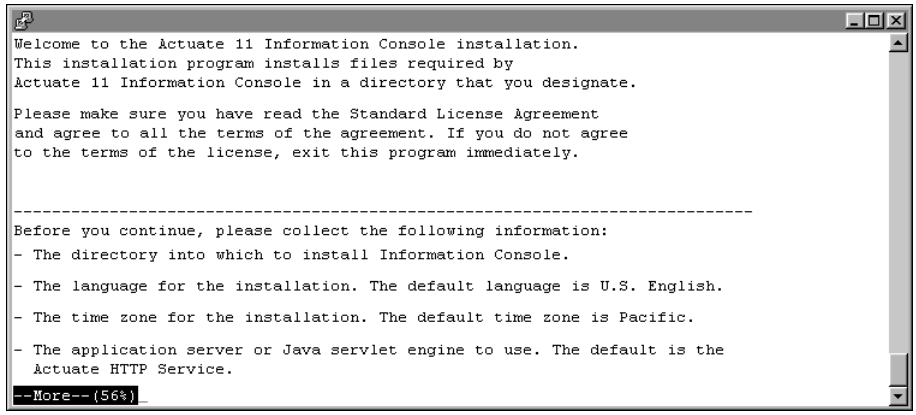

Figure 6-3 Viewing the introduction

6 Type the path for the Information Console installation, for example /home /Actuate/IC, as shown in Figure 6-4. Alternatively, press Enter to accept the default directory.

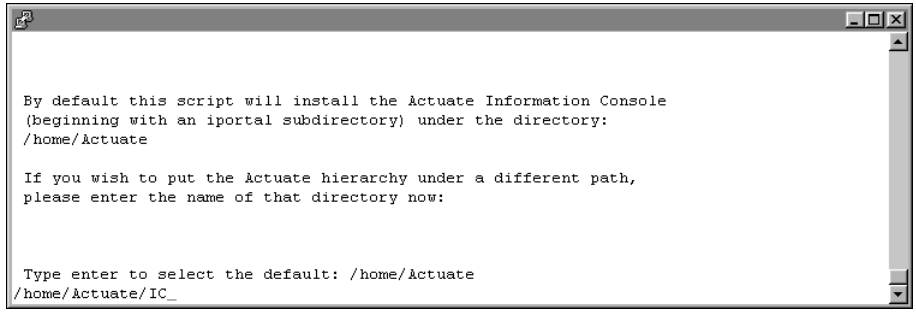

Figure 6-4 Specifying the Information Console install directory

7 Press Enter to select the default locale, which is English, as shown in Figure 6-5. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.

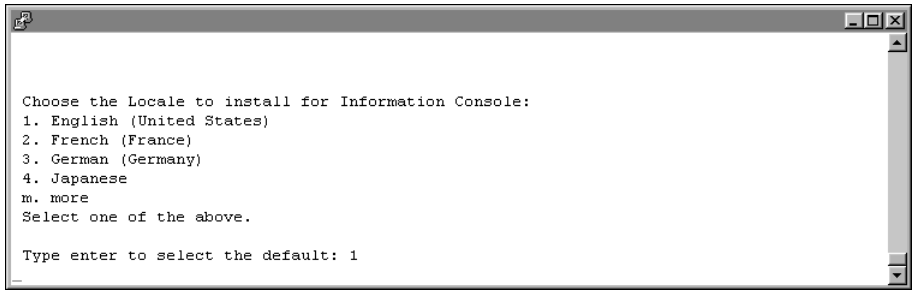

Figure 6-5 Specifying a locale

8 Press Enter to select the default time zone, which is America/Los_Angeles as shown in Figure 6-6. Alternatively, select another time zone from the numbered list.

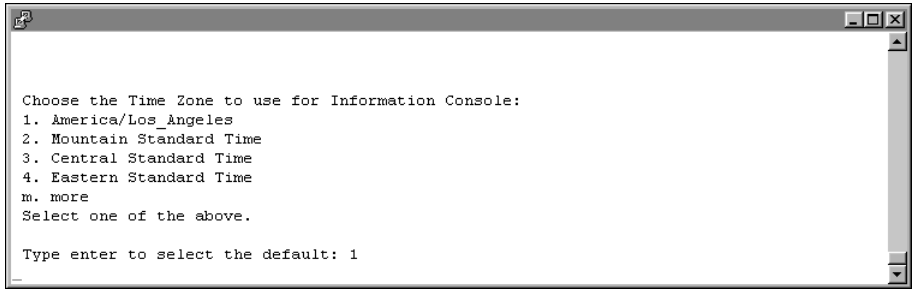

Figure 6-6 Specifying a time zone

9 Type a default profile name that you will use in Information Console, as shown in Figure 6-7.

Figure 6-7 Specifying the default profile name

10 Type the IP address or host name of the machine where iServer runs, or accept the default, your machine name, as shown in Figure 6-8.

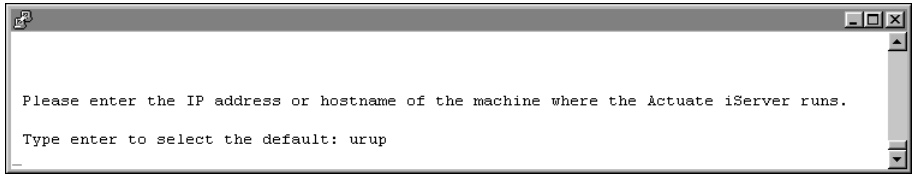

Figure 6-8 Specifying the machine on which the iServer runs

11 Type the number of the port where iServer listens for requests, or accept the default, 8000, as shown in Figure 6-9.

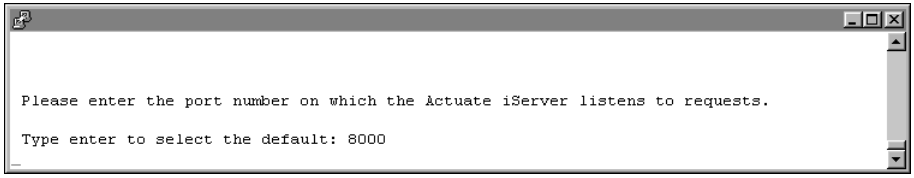

Figure 6-9 Specifying the port number on which iServer listens

12 Type the Encyclopedia volume name that you want to use, or accept the default, your machine name, as shown in Figure 6-10.

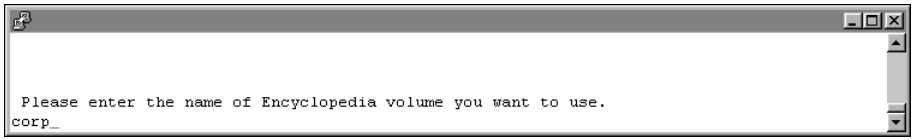

Figure 6-10 Specifying the Encyclopedia volume name

13 The installation program displays the settings that you specified during the install process. Review these settings, as shown in Figure 6-11, then specify whether you accept them. Press Enter to accept the default option, y for yes. Alternatively, type n for no, or type q to quit.

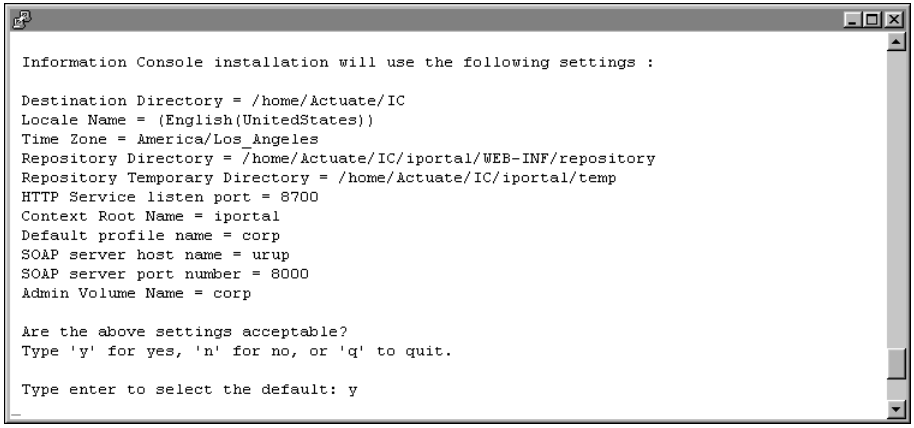

Figure 6-11 Reviewing your settings

14 The installation program installs Information Console, and displays an indicator showing the progress of the installation, as shown in Figure 6-12.

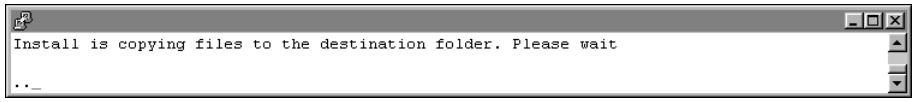

Figure 6-12 The install program copies files to your destination folder

15 When the installation completes, review the information, as shown in Figure 6-13. Issuing the command: sh ./update_rclocal_infoconsole.sh starts the Information Console service at system startup.

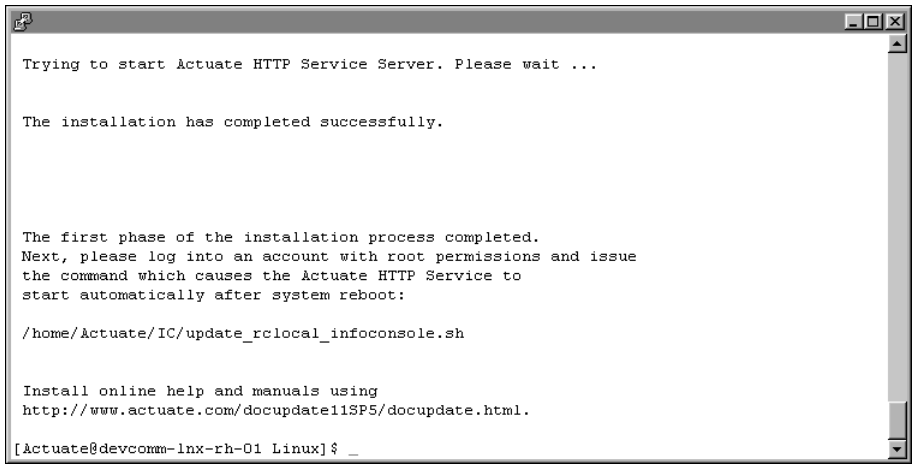

Figure 6-13 Typing the command to start the service

Using the WAR file to install

If Actuate supports your application server, you can deploy Information Console as a WAR (web archive) file. See the application server documentation about how to deploy a WAR file. For information about application servers on which Actuate supports deployment of Information Console, see the Supported Products Matrix for this release at the following URL:

https://support.opentext.com

Customize Information Console for your local environment, if necessary, before beginning deployment. To customize these applications for your local environment, follow the steps in "Preparing the WAR file," later in this chapter, and deploy the customized Information Console WAR file to the application server instead of the ActuateInformationConsole.war file that ships with Information Console.

General deployment tasks

You must configure the Information Console WAR file and application server to integrate them. You must also verify that you have a standard J2EE installation.

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

- Configure the server for best performance with Actuate products.
- Configure the Information Console WAR to integrate with the application server as described in "Preparing the WAR file," later in this chapter.
- Deploy Information Console to the application server.
  - If the application server has deployment tools, use those tools to integrate the application server with Actuate products. For more information about deployment tools, see the application server documentation.
  - If the application server does not have deployment tools, add the Actuate context root to the application server, typically by modifying the application server configuration file.

Preparing the server

Actuate recommends the following configuration for best performance:

- Use at least a two-CPU machine for Information Console.
- If iServer uses more than four CPUs, configure one CPU for Information Console for every two CPUs that iServer uses.

Set the following configuration values on the application or web server:

- On a machine with 1 GB of physical memory, set the Java heap size to a value between 256 MB and 512 MB.
- Set the number of threads to a value between 50 and 75.

Preparing the WAR file

You can use a WAR file to deploy Information Console to many supported servers. Table 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 the locale at login.	If you change the locale, select the new locale from the locales in /WEB-INF/Localemap.xml.

Table 6-1 Information Console configuration parameters

Parameter name	Description	Action
DEFAULT _TIMEZONE	The default time zone is Pacific Standard Time (PST). You can leave this value unchanged.A user can select a time zone at login.	If you change the time zone, select the new time zone from the time zones in the TimeZones.xml. file, in the WEB-INF directory.
SERVER _DEFAULT	This value specifies the iServer URL to which the Information Console application connects if you do not specify a server URL. The default value is http://localhost:8000	Update this value to the machine and port of the server. Change localhost to the iServer machine name or IP address. Change 8000 to the iServer port number.
DEFAULT _VOLUME	This value specifies the default Encyclopedia volume for Information Console. If you do not specify a volume in an Information Console URL, the JSP application attempts to log in to this volume. The default value is localhost.	Update this value to the name of an Encyclopedia volume.
BIRT_VIEWER _LOG_DIR	The location of the files that log BIRT viewer activity.	You must create this physical location on the file system.
LOG_FILE _LOCATION	The location of the files that log Information Console activity.	You must create this physical location on the file system.
TEMP_FOLDER _LOCATION	The location where Information Console creates temporary files.	You must create this physical location on the file system.
TRANSIENT _STORE_PATH	The location where Information Console creates temporary files.	You must create this physical location on the file system.

How to customize the WAR file

The following steps describe the generic procedure for customizing an Information Console WAR file:

1. Create a temporary directory, such as /home/Actuate/ic_temp.

   If you use an existing directory, ensure that this directory is empty.

2. Copy the appropriate Information Console WAR file for your environment, such as TOMCAT_ActuateInformationConsole.war, to the temporary directory, as shown in the following example:

   ```
   cp TOMCAT_ActuateInformationConsole.war /home/Actuate/ic_temp
   ```

3. Decompress the WAR file, as shown in the following example:

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

   The Information Console files appear in the temporary directory.

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

   If you used the temporary path in step 2, the file location is /home/Actuate /ic_temp/WEB-INF/Web.xml. Refer to Table 6-1 for a list of entries to modify in web.xml.

5 Save and close web.xml.

6 Type the following command:

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

   This command creates newinformationconsole.war in the /ic_temp directory. This new WAR file for Information Console contains the modified configuration values.

Use newinformationconsole.war to deploy Information Console to your application server.

About clusters of servers

If your application server supports clustering, see your application server documentation for more information about setting up clusters and deploying web applications such as Information Console.

Avoiding cache conflicts after installing

Information Console uses Java Server Page (JSP) technology. Application servers and browsers cache pages. A browser can use a cached copy of a page instead of the new page. After you install, using a cached copy of some pages can lead to errors or missing functionality.

To avoid this problem, clear the application server cache after you install Information Console. With some application servers, you must restart the application server. For more information about the necessary steps for clearing the cache, see the application server documentation.

If possible, also clear the browser cache to ensure that the browser does not use an old copy of the page from its cache. Alternatively, you can refresh recently visited pages or clear the browser cache if a problem occurs. For more information about clearing the browser's cache and refreshing a page, see the browser documentation.

Testing the installation

Complete the steps in the following section to test the Information Console installation.

How to test the installation

1. Start the application server, if necessary.
2. Open a web browser.
3. Type the URL for the Information Console home page.

   You can use a URL similar to the following example:

   ```
   http://Actuate1:8700/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 URL to the BIRT iServer.

   The Information Console login page appears.

   On the Information Console login page:

   - For enterprise mode Information Console, in Volume Profile, select an Encyclopedia Volume.
   - In User name, type your login name.
   - In Password, type your password.
   - In Language, select a locale.
   - In Time zone, select a time zone.
4. Choose Log In.

Chapter

7

Installing iServer Integration Technology and documentation

This chapter discusses the following topics:

- Installing BIRT iServer Integration Technology
- Installing the localization and documentation files

Installing BIRT iServer Integration Technology

This section describes how to install BIRT iServer Integration Technology for Linux and UNIX.

How to install

In a default installation, BIRT iServer Integration Technology installs in $HOME /ServerIntTech. To install BIRT iServer Integration Technology, perform the following steps:

1. Download the required files. Extract the files.
2. To install the server files, execute the isitinstall script:

   ```
   sh ./isitinstall.sh
   ```

   The script displays a number of prompts. Respond to the prompts as described in the following procedure.
3. The license agreement appears, as shown in Figure 7-1.

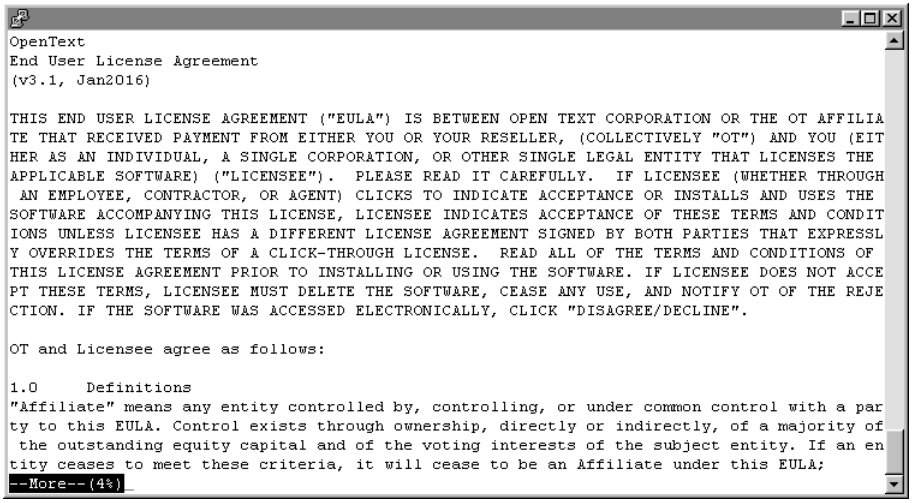

Figure 7-1 Reviewing the license agreement

4. Read the license agreement and press Enter to continue installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 7-2.

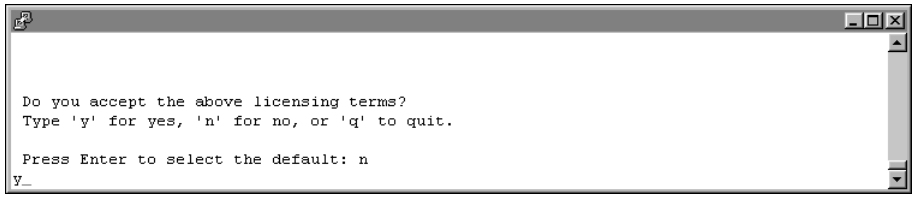

Figure 7-2 Accepting the license agreement

5 The introduction to the installation appears, as shown in Figure 7-3. Press Enter after reviewing the introductory information.

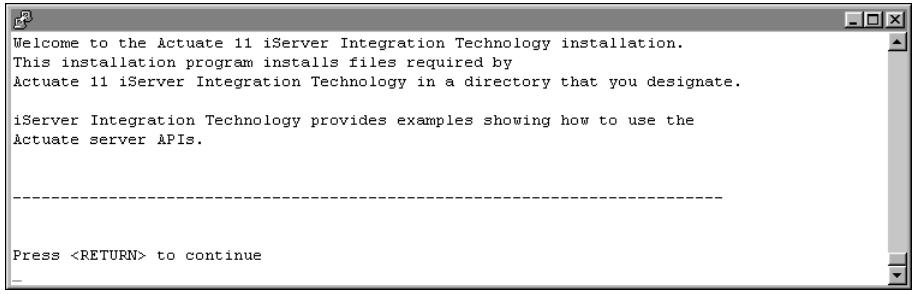

Figure 7-3 Reviewing introductory information

6 Press Enter to accept the default location for installation, $HOME /ServerIntTech as shown in Figure 7-4. Alternatively, type a different directory and press Enter.

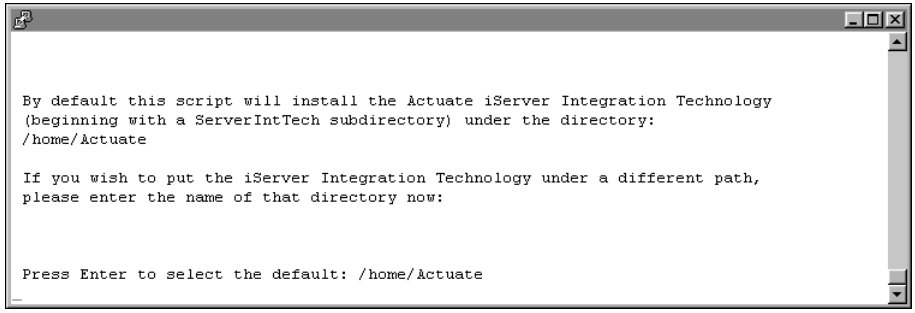

Figure 7-4 Specifying the installation directory

7 Review the settings, as shown in Figure 7-5, then specify whether to accept the settings. Press Enter to accept the default, y for yes. Alternatively type n for no, or q to quit.

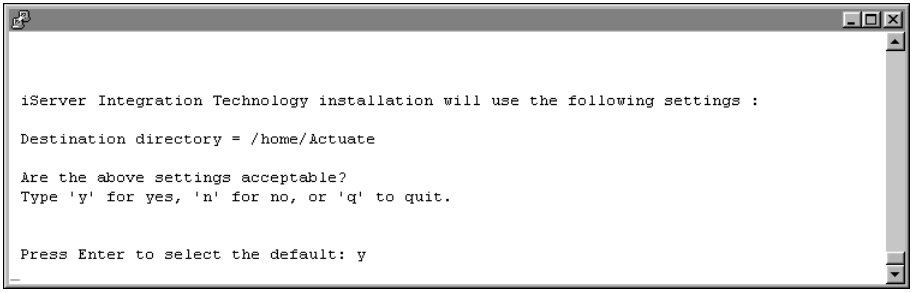

Figure 7-5 Reviewing settings before copying files

8 The installation program installs iServer Integration Technology, and displays an indicator showing how the installation is progressing, as shown in Figure 7-6.

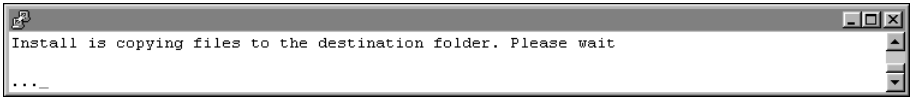

Figure 7-6 Copying files to your destination folder

9 When the installation program finishes, it provides additional information about installing online help and manuals, as shown in Figure 7-7.

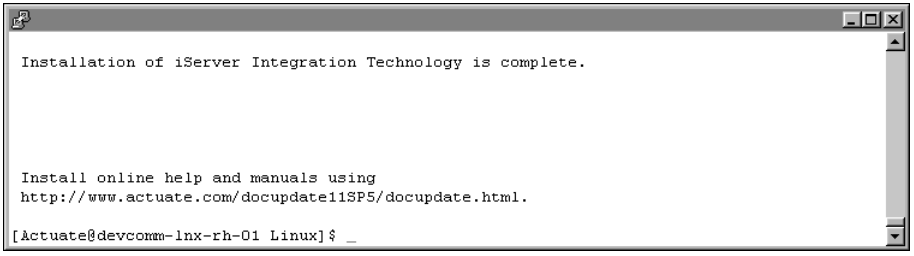

Figure 7-7 Viewing information about installing online help

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 the following location:

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 steps:

1. After downloading the tar file, decompress it using the tar command:

   ```
   tar -xvf FILE_NAME.tar
   ```

2. To begin the installation, move to the newly decompressed directory and execute the helpinstall script:

   ```
   sh ./helpinstall.sh
   ```

   The script displays a series of prompts. Respond to the prompts as described in the following procedures.

3. The license agreement appears, as shown in Figure 7-8.

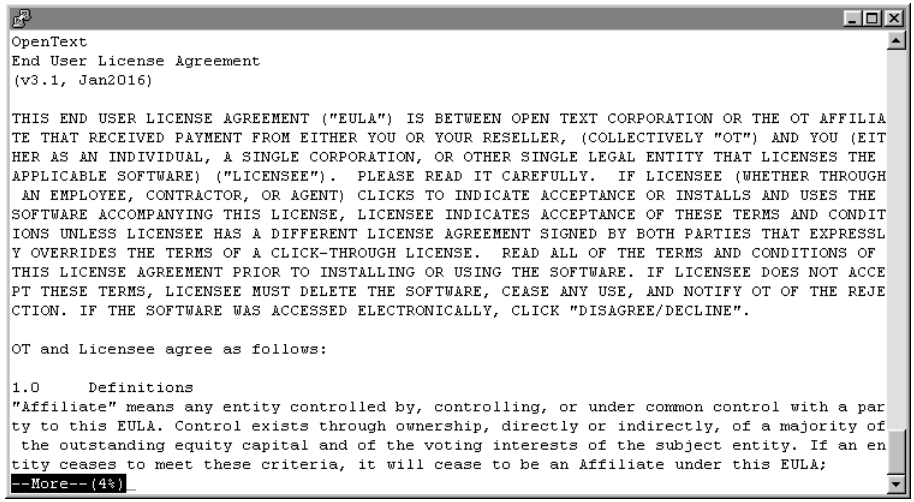

Figure 7-8 Reviewing the license agreement

4. Read the license agreement and press Enter to continue installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 7-9.

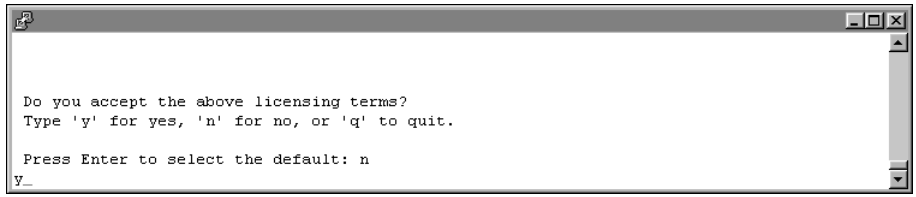

Figure 7-9 Accepting the license agreement

5 Choose the products that you wish to update in this install, as shown in Figure 7-10. If you want to choose more than one, just enter the numbers separated by a space. For example, type 1 2 3 to select all products.

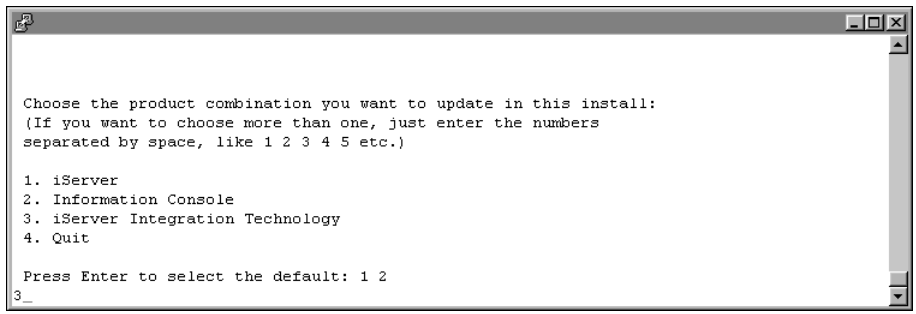

Figure 7-10 Selecting a product combination

6 Enter the full path of the product installation directory that you chose to update, as shown in Figure 7-11. If you chose to update more than one product, you are asked for the full path of every product directory in your selection.

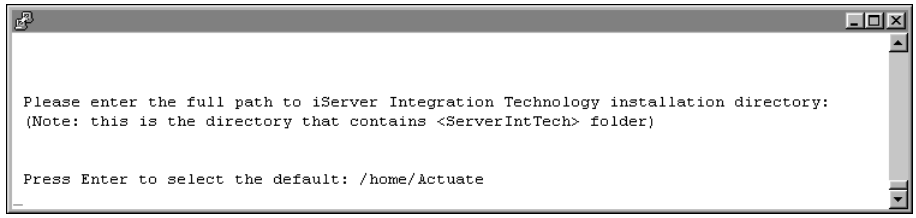

Figure 7-11 Specifying a product directory

7 Review the settings, as shown in Figure 7-12, then specify whether you accept the settings. Press Enter to accept the default, y for yes. Alternatively, type n for no, or q to quit.

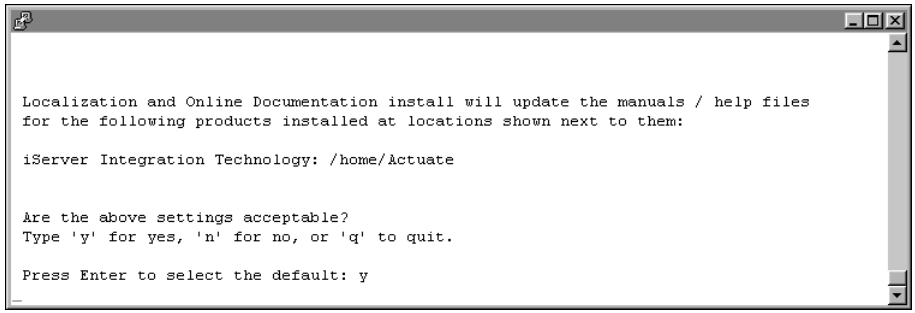

Figure 7-12 Reviewing settings before copying files

8 The installation program starts and displays an indicator showing how the installation is progressing, as shown in Figure 7-13.

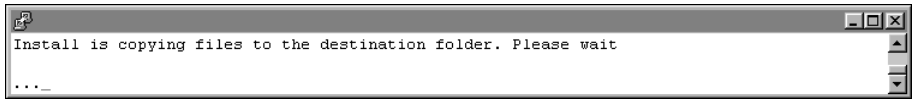

Figure 7-13 Copying files to your destination folder

9 The installation program finishes, as shown in Figure 7-14.

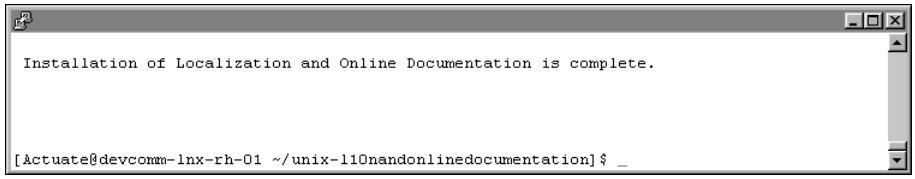

Figure 7-14 The installation program finishes

Part Three

Licensing

Chapter

8

Licensing BIRT iServer

This chapter discusses the following topics:

- Understanding licensing types
- Understanding licensing options
- Installing Actuate BIRT iServer System license files
- Understanding CPU binding

Understanding licensing types

BIRT iServer System licensing supports running BIRT iServer with sets of features grouped as license options. You enable BIRT iServer System options using one or more of the following types of license models:

- Named User

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

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

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

- CPU Core

  Specifies the maximum number of CPUs that BIRT iServer System can use. Any number of users can access the licensed options on the system provided adequate licensing and capacity exists.

- Instance

  A BIRT onDemand licensing option that provides a pre-packaged amount of dedicated capacity for a customer application. With instance licensing, the customer does not need to count named users. Multiple instances can be combined to meet capacity needs.

- Subscription

  An annual payment option that permits the use of the licensed software and includes maintenance. Offered with some of the other licensing models.

  A subscription license is not a perpetual license. Once the subscription term expires, the software can no longer be used.

- Software as a Service (SaaS)

  Some products are offered as a Software as a Service (SaaS) option, providing customers with a convenient solution without incurring the acquisition and management costs of hardware and traditional licenses.

- Packages

  Some options are offered as packages to customers for convenience and cost-saving benefits. These packages can be used in conjunction with individually selected options.

- Work Unit (WU) License

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

  Similar to CPU Core licensing, but defined at a more granular level. With Work Unit Licensing, the customer can license just the precise amount of capacity needed for application requirements. Any number of users can access the licensed options provided sufficient capacity has been purchased.

In a CPU Core and Work Unit licensing, OpenText 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 an Unlimited User CPU License.	11

To determine the license options installed on iServer, log in to Configuration Console, and choose Show License. Figure 8-1 shows the license options that appear for the evaluation license.

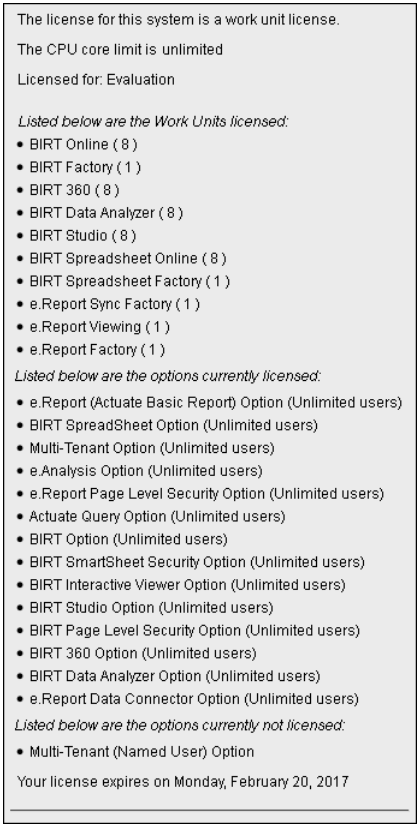

Figure 8-1 iServer License options

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. 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:

```
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 acpmdconfig.xml file of each node, and accessible to all nodes in the cluster.

A node key associates an iServer node in a cluster with the machine ID. The node-key licensing mechanism restricts the iServer node installation to that machine.

On startup, each node in the cluster checks the shared license file, verifies the installed options, and determines whether its node key, which is generated at run time, matches the license information. If the node key matches, the node joins the cluster. Otherwise, it shuts down with an error if the node-lock-violation grace period has been exceeded.

A license file remains valid until a specific date. If your license file is about to expire, the system reminds you that the file expires on a certain date when you log in to the Configuration or Management Consoles. Reminders also appear in the system log file. To arrange for a permanent license file, or if you have a problem with an expiring file, please contact 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 install program prompts you to provide the location of the license file. After providing the location of the license file, the install program issues a prompt similar to the following message:

```
The iServer system license file is locked to the machines that are
   used in the iServer system. The following machine id must be
   used to request a node key license file from Actuate:

   IORRHEHs6S5UCsEtrdVu6jOixmzvFY3BbOqXLiwswQGDceJmKYYaEu0j18lQxjM
   sYCxnka3hVkDZFGwkmQMxb+hgKaz4om2vLUcS0ocYTA7Ta6VTMavLFQo7bEjRyr
   olwxAKu0Vr4NA6o8uWCzjGZXX8KrjViSUoROj70hWOY=

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\AcServer\bin folder.
```

The format of the alphanumeric string for the machine ID and location of the license file are different depending on the operating system.

After installing iServer, you must run the utility, acmachineid, from the command line to generate the machine ID information. Copy the machine ID in the command prompt to a file or e-mail message and send it to 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. Navigate to AC_SERVER_HOME\bin.
2. Type the following command and press Enter:

   ```
   ./acmachineid
   ```

   The utility provides output in the following format:

   ```
   STATUS:          OK
   GEN_VERSION:     11 Service Pack 5
   GEN_BUILD:       XxXBuild NumberXxX
   MACHINEID:
                    I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67GbO8
                    VjcFs6pcuAgbtDaZauSbFFa2mRejwVJc7ZjKfMEVl1suXglM
                    KmZLiwtLykwJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9f
                    JUMYzZORKk23jrxaSwUDsgKsvlc1A6q8UbmrrAYHD8Ggtpui
                    AmxWt4xjEM6rqlmsNEW/4ViMC0KDBkSn
   ```

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 iServer 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://<machine name>: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.

The Actuate_iServer_key_XXXXX.xml will contain the node key information for the stand-alone machine or all machines in a cluster. There is no separate node license file for each machine.

Listing 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
       VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEVl1suXglMKmZLiwtLykDa/
       wJisqMS0EhYe5sCYoKjG+XL2UEnL2GGhLtI9fJUMYzZORKk23jrxaSwUDig
       Ksvlc1A6q8UbmrrAYHD8GgtpuiAmxWt4xjEM6rqlmsNEW/4Vjm40KxlkSv"
       ServerName="W7CLSTRNODE1"/>
   <NodeKey
       MachineId="I0RREHs0Jk6tu0o8AbCrVL61x7kDpLlQKwS2t1W7qM67GbO8
       VjcFs6pcuAgbtZauSbFFa2mRejwVJc7ZjKfMEVl1suXglMKmZLiwtLykDa/
       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 an BIRT iServer machine count toward the maximum number of licensed CPUs. For more information about performing CPU binding on a Windows machine, see Chapter 7, "Licensing BIRT iServer," in *Installing BIRT iServer for Windows*.

To bind BIRT iServer to a set of processors, bind the Actuate Process Management Daemon (PMD) to the processors. The Actuate PMD starts all BIRT iServer processes. The processes inherit the binding from the Actuate PMD.

In a cluster, BIRT iServer counts only the processors on nodes that join the cluster and run the encycsrvr11 process. An encycsrvr11 process runs when a node is online. BIRT iServer counts the number of processors on a machine when the first encycsrvr11 process starts.

When deploying BIRT iServer on a machine with multi-threaded CPUs that use logical processors, the customer receives a license based on the number of physical processors in the system. To accommodate the use of logical processors, the customer receives a license key that specifies two or four times the number of physical processors.

This section contains the following topics:

- Binding BIRT iServer to processors on a Sun Solaris machine
- Binding BIRT iServer to processors on an HP-UX 11i machine
- Checking BIRT iServer bound processors
- Configuring e-mail for CPU license problems

Binding BIRT iServer to processors on a Sun Solaris machine

You can perform single processor binding and processor set binding on a Solaris machine with the following types of binding:

- With single processor binding, you bind a process to a single processor using the processor_bind( ) system call or the pbind command.
- With processor set binding, you bind a process to a group of processors on a Solaris machine. If you bind a process to multiple processor sets, the sets must not overlap.

  To create a processor set, use the pset_create( ) system call or psrset command. Then, you can bind a process to the set using the pset_bind( ) system call or the psrset command. The psrset command restricts other processes from running on the processors to which you bind a BIRT iServer process.

For more information about processor binding and the processor binding commands, refer to your Solaris documentation.

Binding to specific CPUs

On Solaris, you bind a CPU to BIRT iServer using one of the following commands:

- pbind
  Binds the BIRT iServer to a single CPU. Does not prevent other applications from using the CPU.

  For example, in an 8-CPU Solaris server, you can use pbind to bind the BIRT iServer to CPU 2. An Oracle database running on the same server that is not bound to a CPU can impact BIRT iServer performance. While BIRT iServer can use only the processing power of CPU 2, the Oracle database can use all CPUs, including CPU 2. BIRT iServer cannot fully utilize CPU 2 if the Oracle database obstructs access.

- psrset
  Binds the BIRT iServer to one or more CPUs or cores. Prevents other applications from using the CPUs or cores.

  For example, in a 4-CPU Solaris server, you can use psrset to bind the BIRT iServer to CPUs 1 and 2. WebLogic Application Server runs on the same server and is restricted to run on CPUs 3 and 4. BIRT iServer and WebLogic cannot obstruct access to the CPUs assigned to the other process.

To bind to a single CPU, shut down BIRT iServer and determine the process ID of the Actuate Process Management Daemon (PMD) using the ps -e command. Type the pbind command to bind the PMD process to a subset of CPUs using the following syntax:

```
pbind -b processor Actuate_pid
```

where

- Actuate_pid is the process ID of the PMD, as reported by ps -e.
- processor is the processor number as reported by /usr/sbin/psrinfo.

The following example shows binding to a specific CPU on Solaris:

```
pbind -b 0 8209 # Binds process 8209 to CPU 0
```

Binding to multiple CPUs

You use Solaris processor sets to bind BIRT iServer to multiple CPUs. Solaris processor sets are non-overlapping groups of processors.

To bind to multiple CPUs, shut down BIRT iServer and determine the process ID of the PMD using the ps -e command. Create a processor set using the psrset command, and display its ID as follows:

```
psrset -c [processor_id...]
```

where processor_id is the individual or multiple processor numbers.

Bind a processor set to BIRT iServer as follows:

```
psrset -b processor_set_id Actuate_pid
```

where

- processor_set_id is the ID returned by the psrset -c command.
- Actuate_pid is the process ID of the PMD, as reported by ps -e.

Binding to multiple-core CPUs

CPU binding is done at the operating system level, which means that BIRT iServer can bind to any logical CPU. For example, the UltraSPARC T1 processor has eight cores and four threads per core, which is a total of 32 logical CPUs to the operating system.

BIRT iServer System can bind to any logical CPU to the granularity of a thread, not just to the core on a T1 system, which is the same as binding to one physical CPU on a 32-CPU system. The commands to bind to a thread or logical CPU on a T1 system are the same as binding to a physical CPU as shown in the previous examples.

Binding on different logical CPUs can have different effects in terms of Actuate system throughput due to scalability factors across cores and threads. For example, on a T1 system, binding to logical CPU 0, 4, 8, and 12, which belong to four separate cores, has better overall throughput than binding to logical CPU 0, 1, 2, and 3, which belong to same core.

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

You can perform processor set binding on an HP-UX 11i machine. The software for creating a processor set runs only on HP-UX 11i or later. Before you use processor sets, you must install the HP-UX 11i June 2004 or later Quality Pack and download the software to create processor sets from the HP web site. The software to create processor sets is not installed with HP-UX 11i.

To create a processor set, use the pset_create( ) system call or psrset command. Then, you can bind a process to the set using the pset_bind( ) system call or the psrset command.

Like Solaris processor sets, HP-UX 11i processor sets are non-overlapping groups of processors. You can download HP-UX 11i processor sets at no charge from the following location:

`http://www.software.hp.com`

The HP-UX 11i psrset utility controls the management of processor sets. Processor sets allow you to isolate a subset of processors for use by specific threads and processes. Processes in a set have equal access to CPU cycles on their cores through the HP-UX standard scheduler.

To bind to one or more CPUs, use the pbind or psrset commands, as described for Solaris. For more information about processor binding and the processor binding commands, see your HP-UX 11i documentation.

Checking BIRT iServer bound processors

BIRT iServer performs the following bound processor checks:

- The number of processors a cluster uses
- The set of bound processors

Determining the number of processors an iServer System uses

When the PMD starts the first 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 in to 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

Chapter

9

Backing up an Encyclopedia volume

This chapter discusses the following topics:

- Performing an Encyclopedia volume backup
- Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database
- Backing up and restoring an Encyclopedia volume that uses an Oracle database
- Backing up and restoring an Encyclopedia volume that uses a DB2 database

Performing an Encyclopedia volume backup

When performing a volume backup, it is important to note that there are two types of data:

- Metadata

  Information about iServer system and Encyclopedia volume settings and data objects stored in third-party relational database management system (RDBMS) schemas

- Data

  iServer system and Encyclopedia volume data objects, such as designs, documents, and information objects, stored as files on disk partitions, and the acserverconfig.xml file containing iServer configuration information

The administrator must back up all Encyclopedia volume metadata and data to ensure the recoverability of a volume in the event of failure. In Release 11, it is not necessary to back up the iServer system schema, although future versions may require this operation to protect critical system metadata. The administrator can restore a corrupted or missing system schema using the System Data Store Administrator utility. For more information on this utility, see "Specifying System Data Store Administrator properties," in Chapter 4, "Upgrading BIRT iServer."

The third-party database that contains Actuate Encyclopedia metadata is a critical component of Actuate iServer System. An Actuate system administrator must take all necessary precautions to ensure that this database is properly backed up and available to safeguard Encyclopedia volume metadata. Please consult 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 Linux/UNIX Actuate 11 Service Pack 5 environment, the default AC_SERVER_HOME path is:

```
/home/Actuate/AcServer/
```

The default AC_DATA_HOME path is:

```
/home/Actuate/AcServer/data/
```

The default Encyclopedia volume path is:

```
/home/Actuate/AcServer/data/encyc
```

The default acserverconfig.xml file path is:

```
/home/Actuate/AcServer/data/config/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, and DB2 also provide graphical administration tools in addition to command-line tools. For more information on using these RDBMS tools to back up and restore an Encyclopedia volume, see the related sections, later in this chapter.

How to perform an Encyclopedia volume backup

To back up an Encyclopedia volume, perform the following tasks:

1. Make sure that the autoarchive file purging process is not running.
2. Make an online backup of the volume schema using the tools provided by the RDBMS.
3. Back up the volume data files using the tools available in the operating system environment.

Avoiding conflict with the autoarchive file purging process

A metadata backup is consistent with a data backup only if the file purging process that runs during an autoarchive operation does not occur between the time you back up the metadata and the time you back up the data. In Volumes—Properties—Advanced—Archiving And Purging, the administrator can specify when the file purging process runs.

How to configure Archiving And Purging

To configure the autoarchive file purging process, perform the following tasks:

1. From the Advanced view of Configuration Console, choose Volumes.
2. On Volumes, point to the icon next to a volume name and choose Properties.

   In Volumes—Properties, choose Advanced. In Advanced, choose Archiving And Purging.
3. In Archiving And Purging, configure the following time-related file purging properties to times that do not conflict with the time when the backup operation runs, as shown in Figure 9-1:

   - Purge deleted files time

     Specifies the time when the file purging process runs to permanently delete expired files.
   - Expiration time of deleted files

     Specifies the length of time that must elapse before the file purging process permanently deletes an expired file.

   Choose OK.

For information on other aspects of archiving, see *Configuring BIRT iServer*.

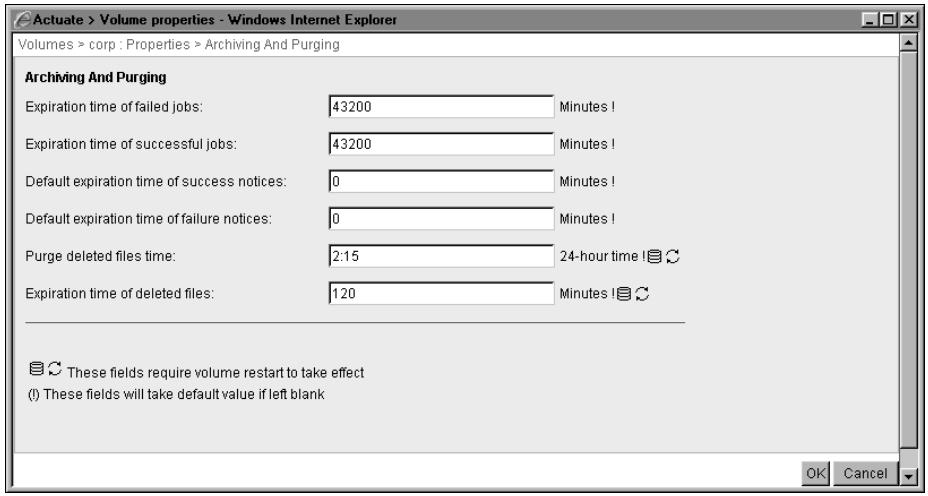

Figure 9-1 Configuring file purging properties

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

PostgreSQL provides the pgAdmin graphical administration tool or the pg_dump and pg_restore command-line utilities to back up and restore a database. These PostgreSQL utilities run on the client not the server.

To back up an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the pgAdmin graphical administration tool or the pg_dump PostgreSQL command-line utility
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

Note that a backup of a PostgreSQL database is not portable across all operating systems.

To restore an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the pgAdmin graphical administration tool or the pg_restore PostgreSQL command-line utility
- Restores Encyclopedia volume data and configuration files using operating system copy commands

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

Backing up an Encyclopedia volume using pg_dump

To back up an Encyclopedia volume using the pg_dump utility, perform the following tasks:

- Create a folder to contain the metadata and volume data backup files
- Back up Encyclopedia volume metadata using the PostgreSQL pg_dump utility
- Back up the acserverconfig.xml file and volume data folders to the backup folder

The following example shows a typical pg_dump command used to export the contents of an Encyclopedia volume schema to a backup file:

```
pg_dump -F c -n ac_corp -f ac_corp_schema.dmp -h dbhost
   -p 8432 -U postgres dbname
```

This pg_dump command example uses the following arguments:

- F

  Specifies the output format. The value c is an abbreviation for custom, which creates a compressed archive that can be used as input to pg_restore.

- n

  Species 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
```

Create a folder to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

In a Linux environment, create a folder to contain the metadata and volume data backup files by performing the following tasks.

How to create a new backup folder

1. Navigate to your home folder, which by default is:

   ```
   /home/Actuate
   ```

2. Create the following new folder, as shown in Figure 9-2:

   ```
   /home/Actuate/encyc_backup
   ```

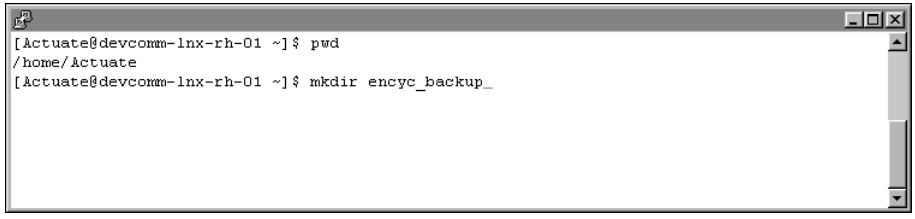

Figure 9-2 Creating a backup folder

Back up Encyclopedia volume metadata using pg_dump by performing the following tasks.

How to run pg_dump

1. Navigate to the following location:

   ```
   /home/Actuate/AcServer/postgresql/bin
   ```

2 Execute the following command. Substitute your machine name for devcomm-lnx-rh-01 in this example:

```
./pg_dump --host devcomm-lnx-rh-01 --port 8432 --username
   postgres --format custom --blobs --verbose --file "/home
   /Actuate/AcServer/encyc_backup/iserver.backup" iserver
```

This operation backs up the entire iserver database. If the -n argument specifying a specific schema or list of schemas is not specified, pg_dump exports all non-system schemas. Alternatively, you can back up only one Encyclopedia volume schema, such as ac_corp, by using the -n argument to specify a particular schema.

3 The command line appears as shown in Figure 9-3.

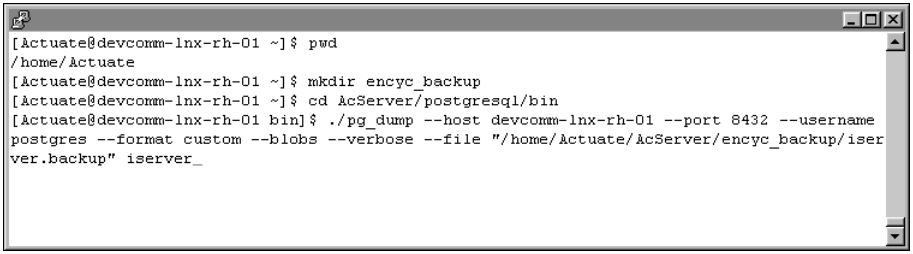

Figure 9-3 Entering the command to execute pg_dump

4 Type the postgres superuser password. The administrator specified this password during the iServer installation procedure.

pg_dump executes, writing status messages to the command prompt.

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

How to back up the volume data folders

1 Navigate to AC_DATA_HOME, which is the location of the iServer data. You specified this location on Setup Type during the install procedure. The default path for AC_DATA_HOME is:

```
/home/Actuate/AcServer/data
```

2 In AC_DATA_HOME, navigate to the config folder that contains acserverconfig.xml file. In Actuate Release 11 Service Pack 5, the acserverconfig.xml file is located in the config/11SP5 subfolder.

Copy acserverconfig.xml to the following backup location, as shown in Figure 9-4:

```
/home/Actuate/encyc_backup
```

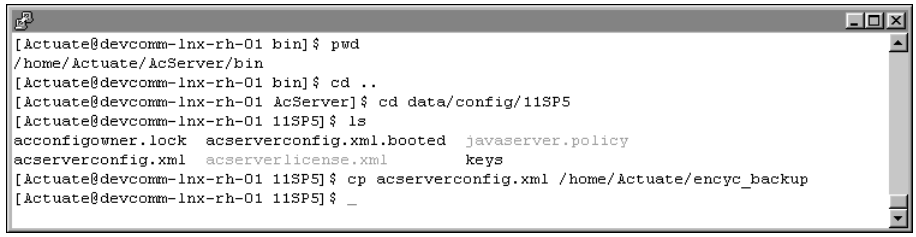

Figure 9-4 Copying acserverconfig.xml to the backup location

3 Navigate to AC_SERVER-HOME/encyc.

Copy the file, fileType, status, and tempRov folders to the following backup location, as shown in Figure 9-5:

```
/home/Actuate/encyc_backup
```

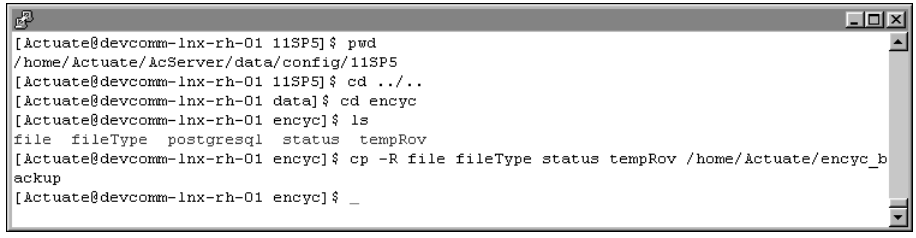

Figure 9-5 Copying the volume data folders to the backup location

In a backup taken immediately after an iServer installation, where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes. If these folders are not present in the environment, simply back up the file and fileType folders.

Do not back up the postgresql folder in an Encyclopedia volume backup operation. The postgres folder contains data, such as log files, from the OOTB PostgreSQL RDBMS installation, which remains active. Inadvertently including these files in an iServer backup, then accidentally overwriting the files with a stale version in a restore operation can cause problems in the PostgreSQL RDBMS installation.

The contents of the backup folder appear as shown in Figure 9-6.

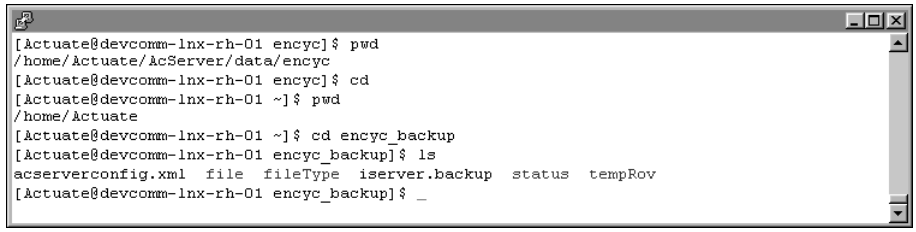

Figure 9-6 Viewing the contents of the backup folder

Restoring an Encyclopedia volume using pg_restore

To restore a backed up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia volume offline
- Delete the acserverconfig.xml and volume data folders in AC_DATA_HOME
- Copy the backed up acserverconfig.xml file and volume data folders from the backup folder to AC_DATA_HOME
- Restore the Encyclopedia volume metadata using the PostgreSQL pg_restore utility
- Take the Encyclopedia volume online

The pg_restore utility runs using arguments similar to the pg_dump utility. The following example shows a typical pg_restore command used to import the contents of a backup file to an Encyclopedia volume schema:

```
pg_restore -h mydbhost -p 8432 -U postgres -d db_name
   ac_corp_schema.dmp
```

Run pg_restore by performing the following tasks.

How to take the Encyclopedia volume offline

1. In a web browser type:

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

   Log into Configuration Console as Administrator.
2. On Simple view, choose Advanced view. Choose Volumes.
3. On Volumes, take the volume offline, as shown in Figure 9-7.

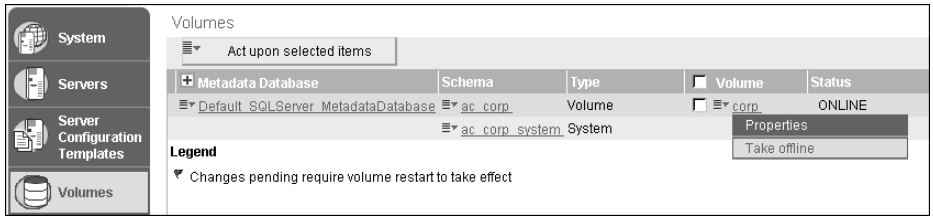

Figure 9-7 Taking the volume offline

How to restore the backed up volume data folders

1 Navigate to AC_DATA_HOME/config/11SP5.

Delete acserverconfig.xml, as shown in Figure 9-8.

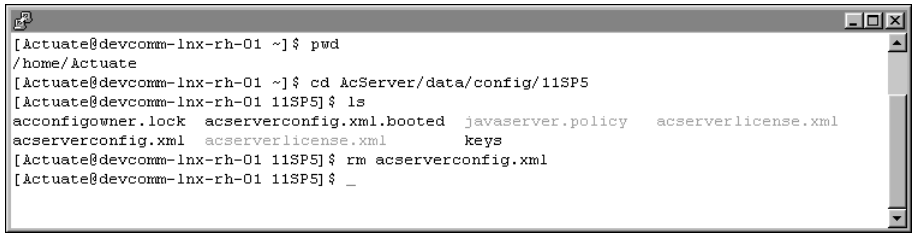

Figure 9-8 Deleting acserverconfig.xml

2 In AC_DATA_HOME, open the encyc folder.

In AC_DATA_HOME/encyc, delete the file, fileType, status, and tempRov folders, as shown in Figure 9-9.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov folders may not exist. Be sure to not delete the postgresql folder.

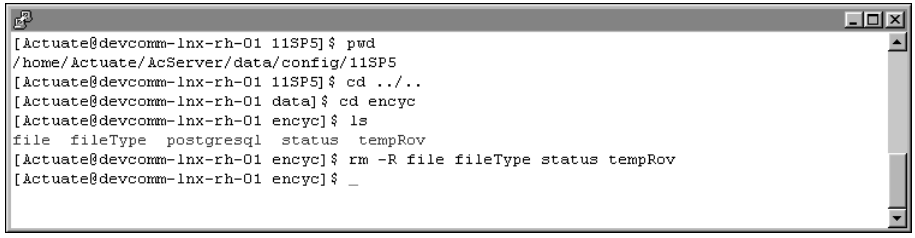

Figure 9-9 Deleting the file, filetype, status, and tempRov folders

3 Navigate to the following location:

`/home/Actuate/encyc_backup`

Copy acserverconfig.xml to AC_DATA_HOME/config/11SP5, as shown in Figure 9-10.

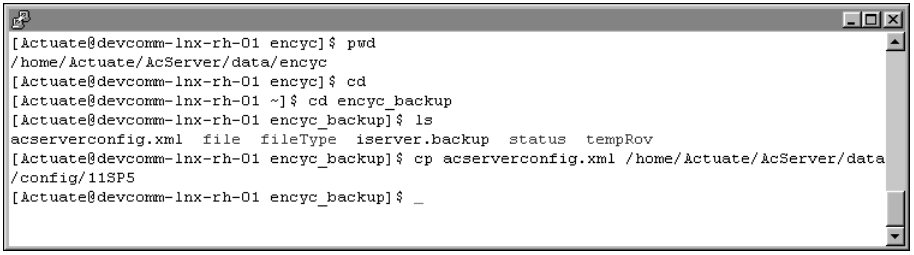

Figure 9-10 Copying acserverconfig.xml to AC_DATA_HOME/config/11SP5

4 Copy the file, fileType, status, and tempRov folders to AC_DATA_HOME /encyc, as shown in Figure 9-11.

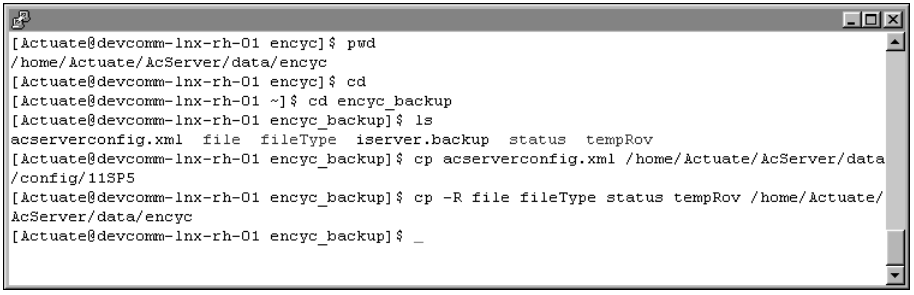

Figure 9-11 Copying the volume data folders to AC_DATA_HOME/encyc

The contents of AC_DATA_HOME/config/11SP5 appear as shown in Figure 9-12.

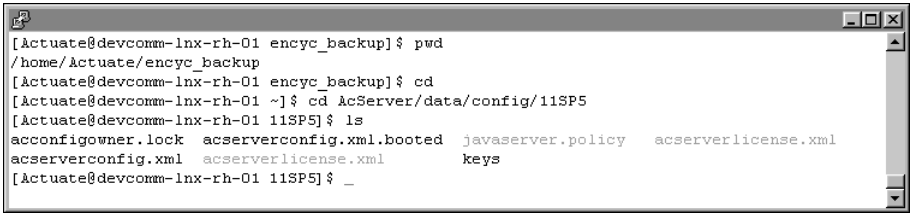

Figure 9-12 Viewing the contents of AC_DATA_HOME/config/11SP5

The contents of AC_DATA_HOME/encyc appear as shown in Figure 9-13.

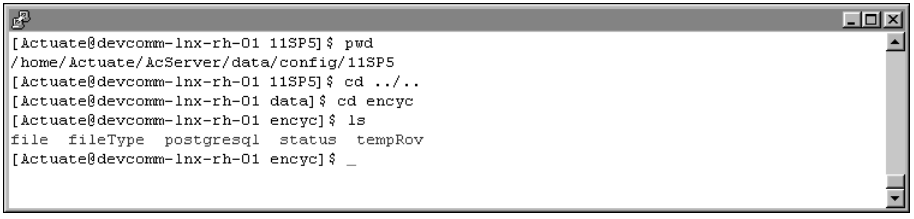

Figure 9-13 Viewing the contents of AC_DATA_HOME/encyc

How to run pg_restore

1 Navigate to the following location:

```
/home/Actuate/AcServer/postgresql/bin
```

2 Execute the following command. Substitute your machine name for devcomm-lnx-rh-01 in this example:

```
./pg_restore --host devcomm-lnx-rh-01 --port 8432 --username
   postgres --dbname iserver --clean --verbose "/home/Actuate
   /AcServer/encyc_backup/iserver.backup"
```

3 The command line appears as shown in Figure 9-14.

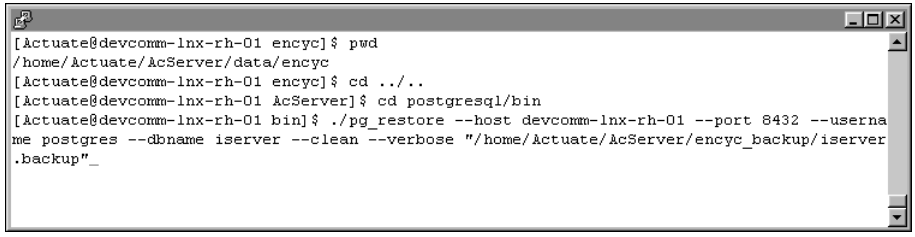

Figure 9-14 Entering the command to execute pg_restore

How to take the Encyclopedia volume online

1 In a web browser type:

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

Log in to Configuration Console as Administrator.

2 On Simple view, choose Advanced view. Choose Volumes.

3 On Volumes, take the volume online, as shown in Figure 9-15.

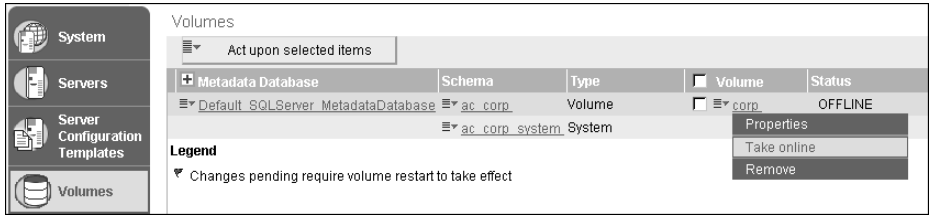

Figure 9-15 Taking the volume online

More information about backing up and restoring an Encyclopedia volume schema using the PostgreSQL pg_dump and pg_restore utilities is available at the following location:

```
http://www.postgresql.org/docs/8.4/static/backup.html
```

Backing up and restoring an Encyclopedia volume that uses an Oracle database

Oracle provides the Data Pump Export (expdp) and Import (impdp) command-line utilities to back up and restore a database. Oracle Data Pump utilities enable high-speed uploading and downloading of data and metadata in an Oracle RDBMS. Unlike the PostgreSQL backup and restore utilities, the Oracle Data Pump utilities run on the server, not the client, for greater efficiency.

To back up an Encyclopedia volume stored in an Oracle RDBMS, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) command-line utility
- Backs up Encyclopedia volume data and configuration file using operating system copy commands

To restore an Encyclopedia volume in an Oracle RDBMS, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility
- Restores Encyclopedia volume data and configuration file using operating system copy commands

The following sections provide more information on how to perform these backup and restore operations.

Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)

Oracle Data Pump utilities require the database administrator to map a database dump directory to a physical directory in the file system. The Oracle RDBMS writes to and reads from this directory when performing export and import operations.

Create a directory to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

To back up an Encyclopedia volume metadata using expdp, perform the following tasks:

- Create a directory to contain the metadata and volume data backup files using sqlplus.

- Back up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility.
- Back up the acserverconfig.xml file and volume data directories to the backup folder.

How to create a backup of the Encyclopedia volume metadata

1. In Linux, open a command window.
2. Use sqlplus, an Oracle SQL editing tool, to log in as the system administrator, as shown in the following example:

   ```
   sqlplus system/password@db_host/dbname.actuate.com
   ```

   The example specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name.
3. In sqlplus, create a backup directory to contain the Encyclopedia volume metadata, as shown in the following example:

   ```
   SQL> CREATE DIRECTORY encyc_backup
          AS '/home/Actuate/encyc_backup';
   SQL> exit
   ```

   Exit sqlplus.
4. Back up the Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility, as shown in the following example:

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

   The expdp example uses the following arguments:

   - system/password@db_host/dbname.actuate.com
     Specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name
   - SCHEMAS
     Specifies the Encyclopedia volume schema to export, such as corp
   - DIRECTORY
     Specifies the directory for writing the database dump, such as encyc_backup, created by the system administrator in sqlplus
   - DUMPFILE
     Specifies the name of the output file, such as ac_corp_schema.dmp

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

How to back up the volume data directories

1. Navigate to AC_DATA_HOME, the location of the iServer data. The administrator specified this location on Setup Type during the install. The default path for AC_DATA_HOME is:

   ```
   /home/Actuate/AcServer/data
   ```

2. In AC_DATA_HOME, navigate to the config folder. In Actuate Release 11 Service Pack 5, the acserverconfig.xml file is located in the config/11SP5 subfolder.

   Copy acserverconfig.xml to the following backup location:

   ```
   /home/Actuate/encyc_backup
   ```

3. Navigate to AC_DATA_HOME/encyc, then copy the file, fileType, status, and tempRov directories to the following backup location:

   ```
   /home/Actuate/encyc_backup
   ```

   In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov directories may not exist. These directories contain information about job details and completion notices and do not appear until a job executes. If these directories are not present in the environment, simply back up the file and fileType directories.

Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp)

To restore a backed up Encyclopedia volume using Oracle Data Pump Import (impdp), perform the following tasks:

- Take the Encyclopedia volume offline.
- Delete the acserverconfig.xml file and volume data directories in AC_DATA_HOME.
- Copy the backed up acserverconfig.xml file and volume data directories from the backup directory to AC_DATA_HOME.
- Restore the Encyclopedia volume metadata using the Oracle impdp utility.
- Take the Encyclopedia volume online.

Take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

1 In a web browser type:

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

Log in to Configuration Console as Administrator.

2 On Simple view, choose Advanced view. Choose Volumes.

3 On Volumes, take the volume offline, as shown in Figure 9-16.

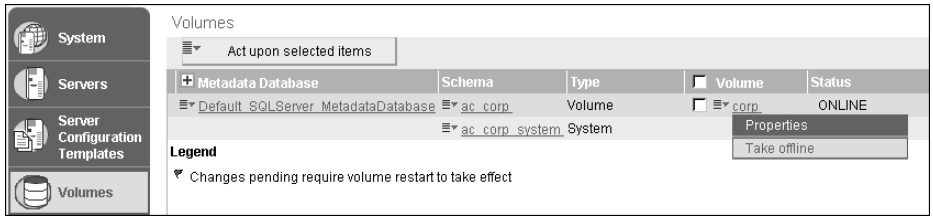

Figure 9-16 Taking the volume offline

How to restore the backed up volume data directories

1 In Linux, open a command window.

2 Navigate to AC_DATA_HOME/config/11SP5 directory and delete the acserverconfig.xml file.

3 Navigate to AC_DATA_HOME/encyc directory and delete the file, fileType, status, and tempRov directories.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov directories may not exist.

4 Navigate to the following backup directory location:

```
/home/Actuate/encyc_backup
```

5 From the backup directory location, perform the following tasks:

1 Copy acserverconfig.xml to AC_DATA_HOME/config/11SP5.

2 Copy the file, filetype, status, and tempRov directories to AC_DATA_HOME/encyc.

Restore the Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility, by performing the following tasks.

How to restore a backup of the Encyclopedia volume metadata

1 Restore the Encyclopedia volume metadata using the impdp utility, as shown in the following example:

```
impdp system/password@db_host/dbname.actuate.com
   SCHEMAS=corp
```

```
DIRECTORY=encyc_backup
DUMPFILE=ac_corp_schema.dmp
```

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

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

1. In a web browser type:

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

   Log in to Configuration Console as Administrator.

2. On Simple view, choose Advanced view. Choose Volumes.

3. On Volumes, take the volume online, as shown in Figure 9-17.

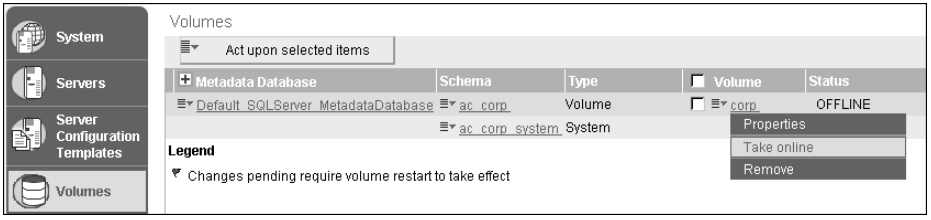

Figure 9-17 Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the Oracle Data Pump utilities, go to the following location:

```
http://download.oracle.com/docs/cd/B19306_01/server.102/b14215
   /dp_overview.htm
```

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

To back up an Encyclopedia volume in the DB2 RDBMS environment, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using DB2 commands
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

To restore an Encyclopedia volume in the DB2 RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume data and configuration files using operating system copy commands
- Restores Encyclopedia volume metadata using DB2 commands

The following sections show how to perform these backup and restore operations.

Backing up an Encyclopedia volume

To back up an Encyclopedia volume, perform the following tasks:

- Create a folder to contain the backup files.
- Back up Encyclopedia volume metadata.
- Back up the 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. Navigate to your home folder, which by default is:

   ```
   /home/Actuate
   ```

2. Create the following new folder:

   ```
   /home/Actuate/encyc_backup
   ```

Back up Encyclopedia volume metadata by performing the following tasks.

How to create a backup of the Encyclopedia volume metadata

1. Add the path to db2profile to the PATH variable on your machine.
2. Execute the following command to perform an online backup of the entire database and compress the backup image:

   ```
   BACKUP DATABASE ISERVER ONLINE
      TO "/home/Actuate/encyc_backup"
      WITH 2 BUFFERS BUFFER 1024 PARALLELISM 1
      COMPRESS WITHOUT PROMPTING
   ```

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

How to back up the volume data folders

1 Navigate to AC_DATA_HOME, which is the location of the iServer data. The administrator specified this location on Setup Type during the install procedure. The default path for AC_DATA_HOME is:

```
/home/Actuate/AcServer/data
```

2 In AC_DATA_HOME, navigate to the config directory. In Actuate Release 11 Service Pack 5, the acserverconfig.xml file is located in the config/11SP5 subfolder.

Copy acserverconfig.xml to the following backup location, as shown in Figure 9-18:

```
/home/Actuate/encyc_backup
```

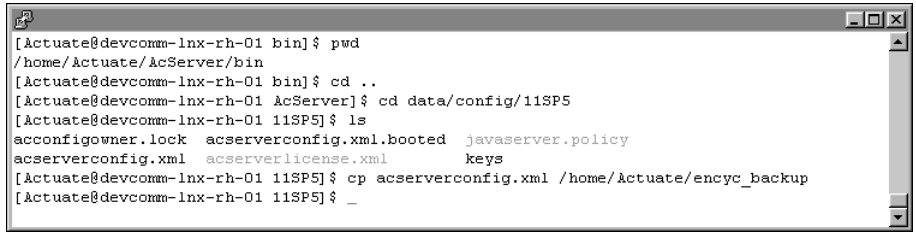

Figure 9-18 Copying acserverconfig.xml to the backup location

3 Navigate to AC_SERVER-HOME/encyc.

Copy the file, fileType, status, and tempRov folders to the following backup location, as shown in Figure 9-19:

```
/home/Actuate/encyc_backup
```

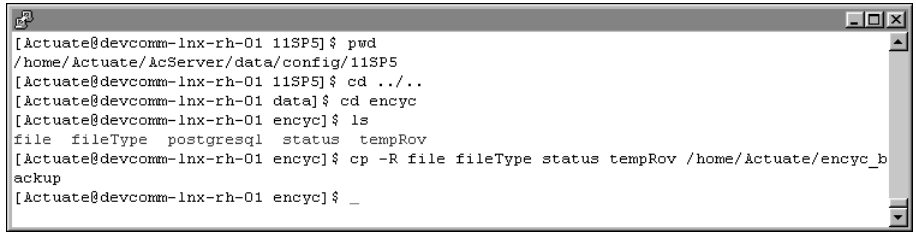

Figure 9-19 Copying the volume data folders to the backup location

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov folders may not exist. These folders contain information about job details and completion notices and do not appear until a job executes. If these folders are not present in the environment, simply back up the file and fileType folders.

The contents of the backup folder appear as shown in Figure 9-20.

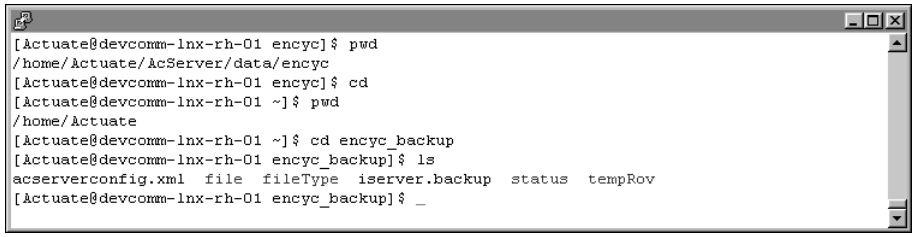

Figure 9-20 Viewing the contents of the backup folder

Restoring an Encyclopedia volume

To restore a backed up Encyclopedia volume, perform the following tasks:

- Take the Encyclopedia volume offline.
- Delete the acserverconfig.xml file and volume data folders in AC_DATA_HOME.
- Copy the backed up acserverconfig.xml file and volume data folders from the backup folder to AC_DATA_HOME.
- Restore the Encyclopedia volume metadata.
- Take the Encyclopedia volume online.

Take the Encyclopedia volume offline by performing the following tasks.

How to take the Encyclopedia volume offline

1 In a web browser type:

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

Log in to Configuration Console as Administrator.

2 On Simple view, choose Advanced view. Choose Volumes.

3 On Volumes, take the volume offline, as shown in Figure 9-21.

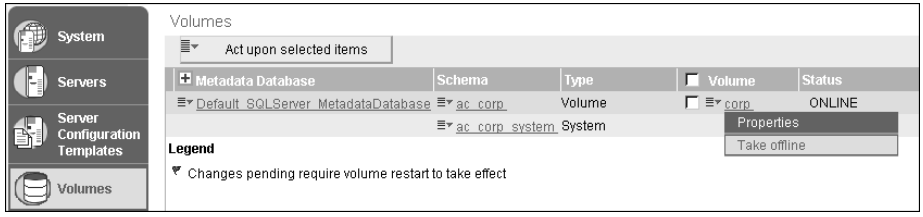

Figure 9-21 Taking the volume offline

How to restore the backed up volume data folders

1 Navigate to AC_DATA_HOME/config/11SP5. In Actuate Release 11 Service Pack 5, the acserverconfig.xml file is located in the config/11SP5 subfolder.

Delete acserverconfig.xml, as shown in Figure 9-22.

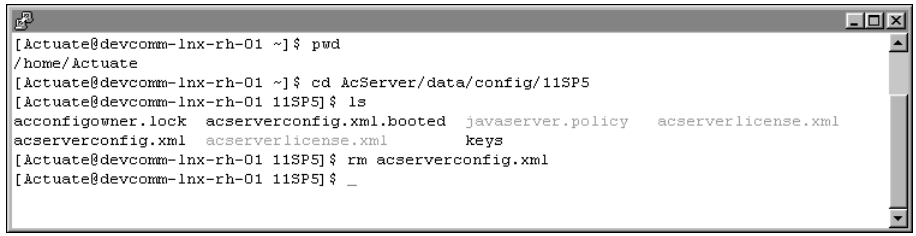

Figure 9-22 Deleting acserverconfig.xml

2 In AC_DATA_HOME, open the encyc folder.

In AC_DATA_HOME/encyc, delete the file, fileType, status, and tempRov folders, as shown in Figure 9-23.

In a backup taken immediately after an iServer installation where there has been no activity on the system, the status or tempRov folders may not exist.

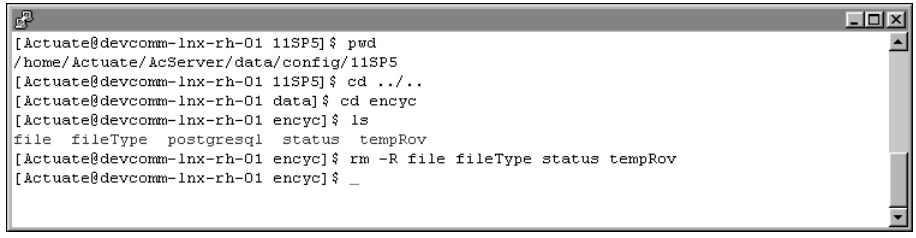

Figure 9-23 Deleting the file, filetype, status, and tempRov folders

3 Navigate to the following location:

`/home/Actuate/encyc_backup`

Copy acserverconfig.xml to AC_DATA_HOME/config/11SP5, as shown in Figure 9-24.

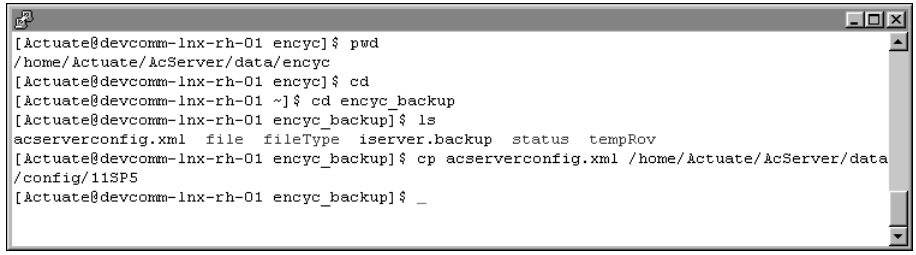

Figure 9-24 Copying acserverconfig.xml to AC_DATA_HOME/config

4 Copy the file, fileType, status, and tempRov folders to AC_DATA_HOME /encyc, as shown in Figure 9-25.

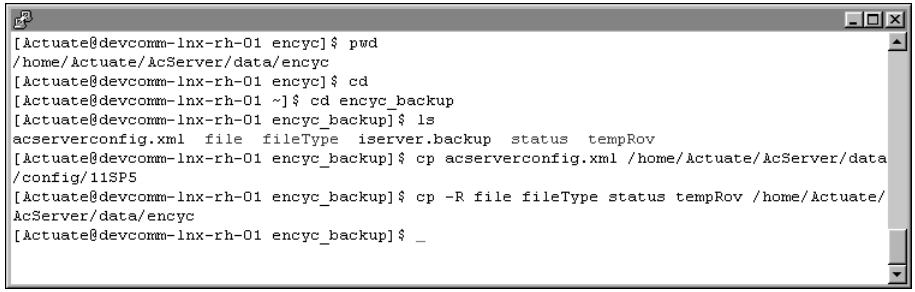

Figure 9-25 Copying the volume data folders to AC_DATA_HOME/encyc

The contents of AC_DATA_HOME/config/11SP5 appear as shown in Figure 9-26.

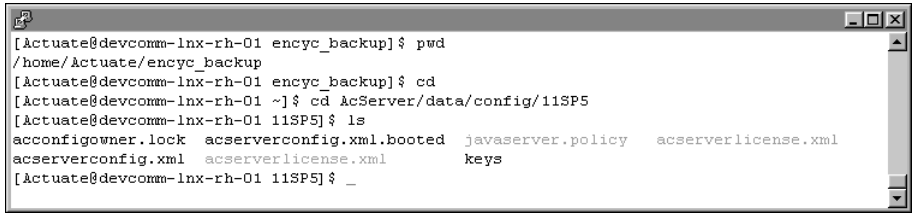

Figure 9-26 Viewing the contents of AC_DATA_HOME/config

The contents of AC_DATA_HOME/encyc appear as shown in Figure 9-27.

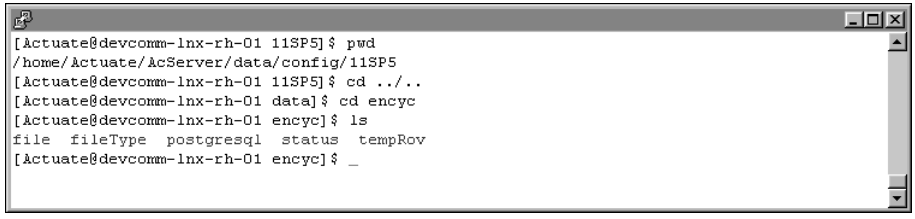

Figure 9-27 Viewing the contents of AC_DATA_HOME/encyc

How to restore a backup of the Encyclopedia volume metadata

1 Add the path to db2profile to your PATH variable.

2 Execute the following command to perform an online backup of the entire database and compress the backup image:

```
db2 RESTORE DATABASE ISERVER FROM "C:\Actuate\iServer
   \encyc_backup" TAKEN AT 20111004180138 WITH 2 BUFFERS BUFFER
   1024 PARALLELISM 1 WITHOUT ROLLING FORWARD WITHOUT PROMPTING;
```

where 20111004180138 is the time stamp of the backup image.

Take the Encyclopedia volume online by performing the following tasks.

How to take the Encyclopedia volume online

1. In a web browser type:

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

   Log in to Configuration Console as Administrator.
2. On Simple view, choose Advanced view. Choose Volumes.
3. On Volumes, take the volume online, as shown in Figure 9-28.

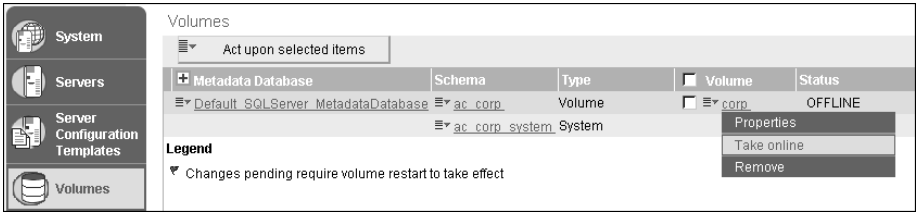

Figure 9-28 Taking the volume online

For more information about backing up and restoring an Encyclopedia volume schema using the DB2 database utilities, go to the following location:

`http://www.ibm.com/developerworks/data/library/techarticle/dm-0910db2incrementalbackup/index.html`

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X