One Design
One Server
One User Experience

Installing BIRT iServer for Linux and UNIX

Contains information proprietary to:
Actuate Corporation, 2207 Bridgepointe Parkway, San Mateo, CA 94404

www.actuate.com
www.birt-exchange.com

Document No. 110812-2-430343 July 8, 2011

Contents

Part 2
Installing

Chapter 2
Installing BIRT iServer 23

Chapter 3
Installing BIRT iServer using an alternative database 47

Understanding ActuateOne

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

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

About the BIRT iServer documentation

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

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

Table I-1 BIRT iServer documentation

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

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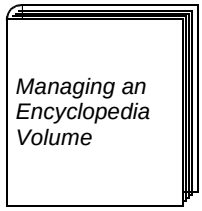
Information Console Developer Guide ■ Overview of Deployment Kit concepts and web applications ■ Using, customizing, and configuring the Deployment Kit ■ Using code components for JSPs, URL parameters, JavaScript files, Java servlets, Java Beans, and security facilities	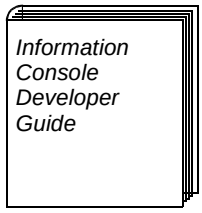
Using BIRT iServer Integration Technology ■ Overview of Actuate web services and SOAP messaging ■ Managing an Encyclopedia volume ■ Developing API applications using Java or .NET ■ Using Java Report Server Security Extension (RSSE) APIs ■ Using logging, performance monitoring, and archiving features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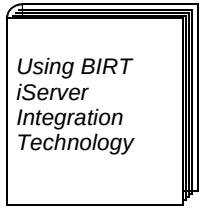

(continues)

Table I-1 BIRT iServer documentation (continued)

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

Obtaining documentation

Actuate provides technical documentation in PDF, HTML, and print formats. You can download PDF or view HTML versions of the documentation from

birt-exchange.com. If you purchase the product, you can also download documentation using ftp as instructed in the e-mail from Actuate Distribution.

If you request a physical package, install the files using the Online Documentation and Localization Resource Files DVD, which ships as part of your Actuate software package. If you accept the default location for the installation directory when you install from the DVD, the installation program creates the /home/Actuate/AcServer/manuals directory.

Using PDF documentation

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

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

Obtaining late-breaking information and documentation updates

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

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

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

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

If you are a new user, you must first register on the site and log in to view the release notes. Birt-exchange.com and actuate.com also provide product update information.

About obtaining technical support

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

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

About supported and obsolete products

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

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

Typographical conventions

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

Table I-2 Typographical conventions

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

Table I-3 Syntax conventions

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

About *Installing BIRT iServer for Linux and UNIX*

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

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

- *Chapter 5. Installing a BIRT iServer cluster.* Describes how to install an 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 Encyclopedia volume information related to users, roles, groups, channels, folders, files, and other objects. 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 is also adapting iServer to support alternative, customizable, third-party database installations. In Release 11, Actuate currently supports PostgreSQL and Oracle.

In Release 11, Actuate provides the following Encyclopedia volume database installation options:

- Install a new iServer with a PostgreSQL or other supported, alternative, third-party database
- Upgrade a former Actuate proprietary database installation to the new Release 11 version
- 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 stores iServer system and Encyclopedia volume metadata in schemas in the third-party RDBMS. Actuate automatically installs these schemas in the OOTB PostgreSQL RDBMS installation. Installation of these schemas in a pre-existing PostgreSQL RDBMS or alternative RDBMS, such as Oracle, 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 PostgreSQL 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:

- Installing 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.

- Upgrading an earlier iServer installation to Release 11 using the installation program to overwrite automatically the earlier installation.

  Upgrades an earlier iServer system in place, such as Release 10 Service Pack 1, automatically migrating one or more Encyclopedia volumes.

- Upgrading an earlier iServer installation to Release 11 using the installation program to install iServer, then manually migrate Encyclopedia volume metadata and data from an earlier to the new installation.

  Upgrades an earlier iServer system in place, such as Release 10 Service Pack 1, without migrating any Encyclopedia volumes. During installation, the administrator chooses to migrate the volumes manually.

  After installation, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the existing iServer installation to either an in-place or a new, side-by-side instance of iServer Release 11 when upgrading from an earlier major release. When upgrading a Release 11 installation to a new Release 11 service pack, the administrator uses the Encyclopedia Data Store Upgrader utility. These utilities are Java programs run from the command line.

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

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

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

Since iServer Release 11 no longer uses the internal proprietary Squirrel database to store Encyclopedia volume metadata, the automatic backup, recovery, and failover features available in earlier iServer releases are now obsolete. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, see Chapter 9, "Backing up an Encyclopedia volume," later in this book.

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

Documentation for a PostgreSQL RDBMS is available at:

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

Documentation for an Oracle RDBMS is available at:

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

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," in Chapter 4, "Upgrading BIRT iServer," later in this book.

An Actuate system administrator must take all necessary precautions to ensure that a database is properly backed up and available to safeguard Encyclopedia volume metadata. Please consult Actuate Support at the time of installation if you have any questions about the backup, recovery, or failover procedures necessary to protect against the possibility of catastrophic failure.

Managing an iServer cluster

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

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

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

RDBMS administrator. The database administrator must use the facilities available in the RDBMS to configure failover capability.

Understanding the iServer System process model

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

- Encyclopedia volume

  Stores metadata in an OOTB (PostgreSQL) or alternative RDBMS and coordinates processing for designs, documents, information objects, and other iServer data objects stored in the file system.

- Process Management Daemon (PMD)

  Distributes service requests among available iServer services and nodes.

- iServer servlet container

  Provides the run-time environment for client applications, such as Actuate Information, Management, and Configuration Consoles. Client applications communicate with iServer System using SOAP-based messaging.

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

- Factory

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

- View

  Supports viewing documents in DHTML and other output formats, such as Excel and PDF. Handles requests to download files from an Encyclopedia volume.

- Integration

  Coordinates the running of information object (IOB) files that extract data from multiple data sources.

- Caching

  Controls the Actuate Caching process that manages an information object cache and enables caching of data retrieved from data sources.

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

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

Understanding process flow in a stand-alone iServer

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

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

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

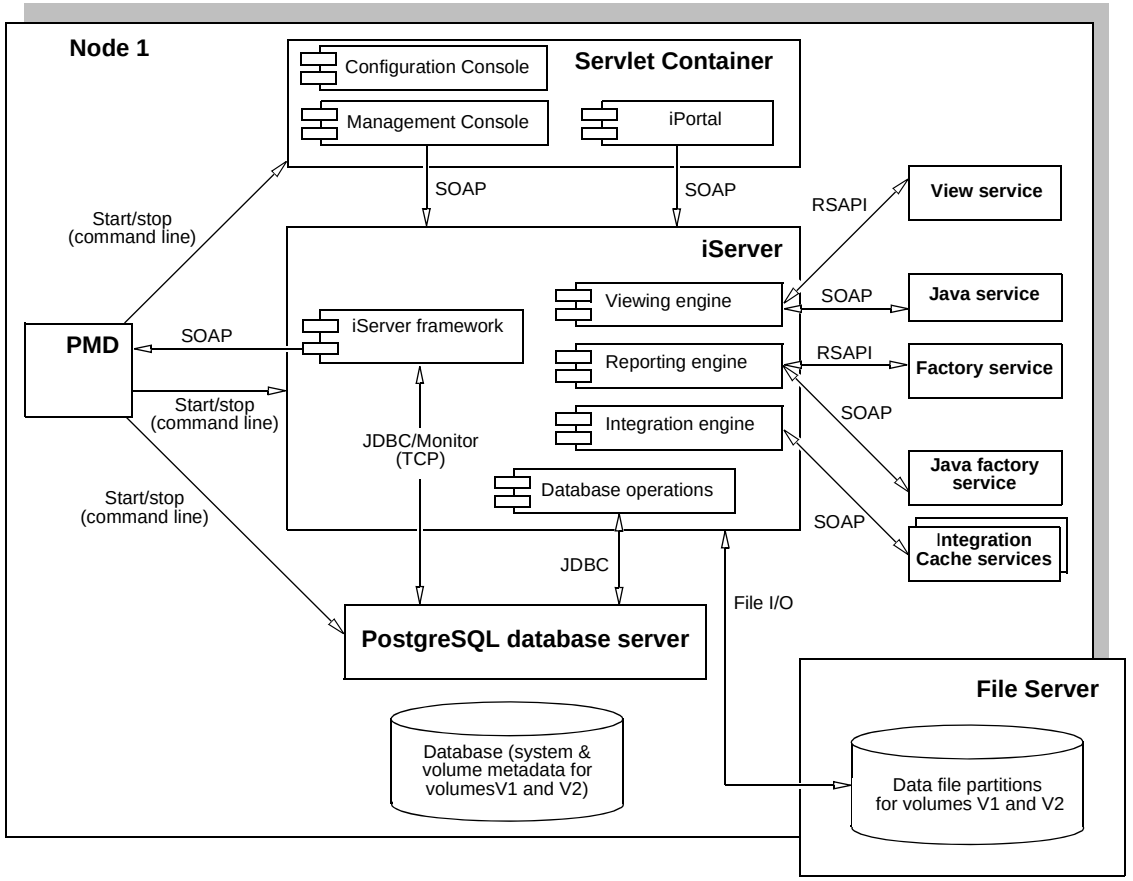

Figure 1-1 iServer RDBMS process architecture for a stand-alone, two-volume, OOTB database

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

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

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

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

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

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

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

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

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

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

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

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

Understanding process flow in an iServer cluster

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

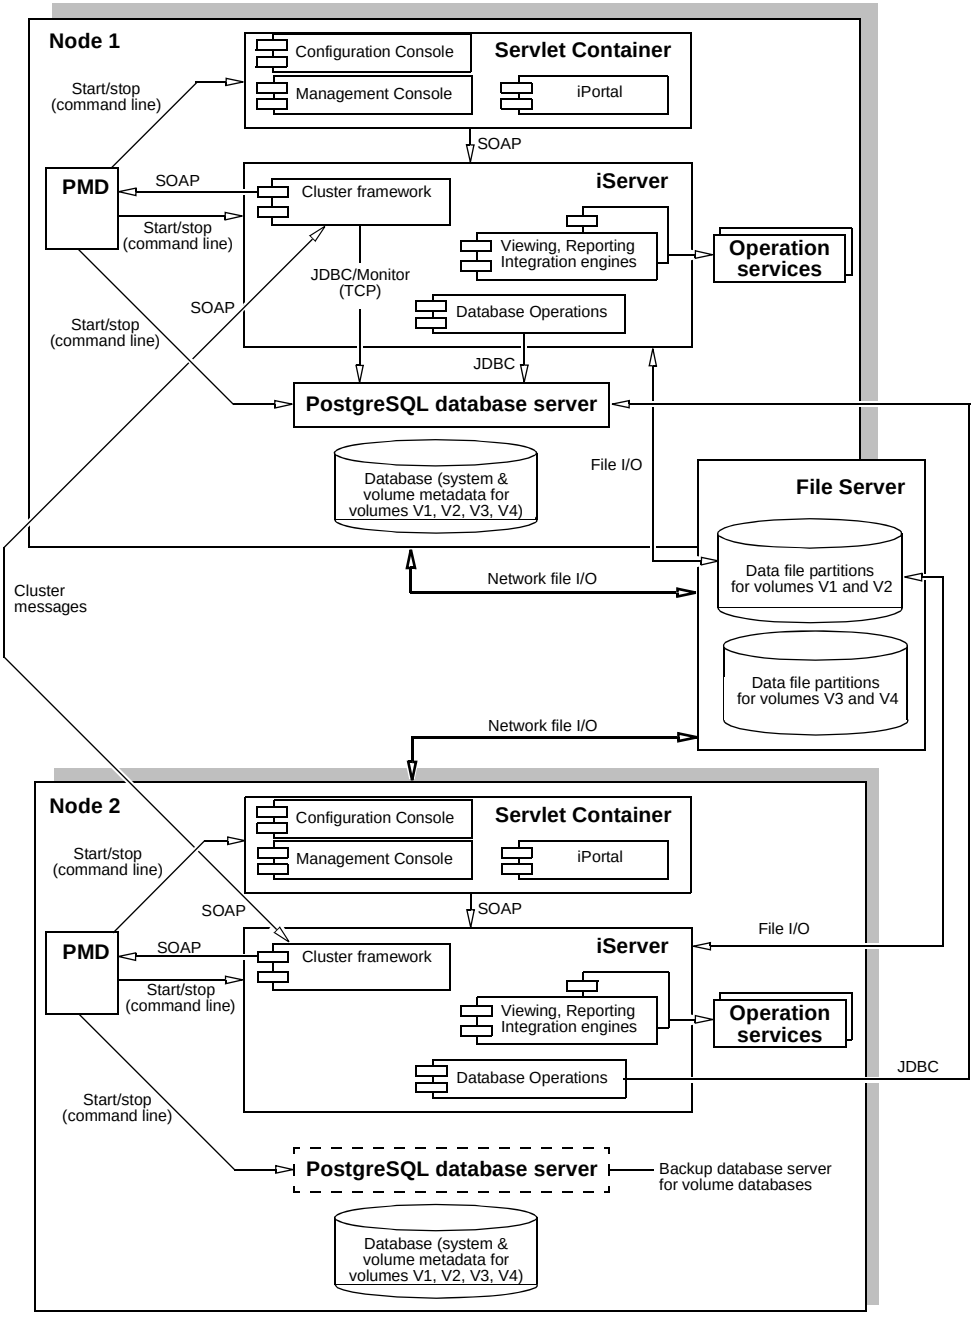

Figure 1-2 iServer RDBMS process architecture for a clustered, two-node, two-volume, OOTB database

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

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

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

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

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

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

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

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

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

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

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

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 PostgreSQL or Oracle. Also describes Actuate licensing policies and procedures and backup and recovery operations.

- *Managing an Encyclopedia Volume*

  Describes how to use Management Console and command-line options to perform tasks such as managing Encyclopedia volume user accounts, assigning privileges, scheduling jobs, and distributing documents.

- *Configuring BIRT iServer*

  Describes how to use Configuration Console to perform tasks such as managing an iServer cluster, adding Encyclopedia volumes to iServer, connecting to databases, updating the license, and configuring iServer properties, such as logging levels, e-mail notification, and printing from iServer.

- *Using BIRT iServer Integration Technology*

  Provides information about application programming using the SOAP-based Actuate Information Delivery API (IDAPI), including a Java developer guide and sections on logging, auto archiving, and using the Java Report Server Security Extension (RSSE).

About Migration and Administration Tools

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

- AcExport

  Formerly used to write a copy of the Encyclopedia volume metadata to a file, so the administrator can import the metadata into another release of iServer.

- AcImport

  Formerly used to populate an Encyclopedia volume with metadata previously written to an exported file.

  In an upgrade from an earlier major release to Release 11, use the Squirrel Data Exporter and the System and Encyclopedia Data Store Administrator utilities that Actuate provides to migrate Encyclopedia volume metadata from the Squirrel database to an alternative RDBMS. Use the Encyclopedia Data Store Upgrader utility to upgrade an earlier Release 11 installation to a newer Release 11 service pack.

  Back up the database using the utilities that the RDBMS provides. For example, PostgreSQL provides the pg_dump and pg_restore utilities and

Oracle provides the expdp and impdp utilities to create and restore a database backup.

Use operating system or other third-party tools to backup and load designs, documents, information objects, and other iServer data objects stored in the file system. For more information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 9, "Backing up an Encyclopedia volume," later in this book.

- AcToc

  Formerly used to list the contents of an export directory. In Release 11, no comparable functionality exists.

- AcVerify

  Formerly used to validate an offline Encyclopedia volume and repair problems.

  In Release 11, use the tools available in the third-party RDBMS containing the Encyclopedia volume metadata to verify data integrity and make repairs.

- AcExtern

  Formerly used to convert Encyclopedia volume user security from internal to external registration to allow administration from another system, such as an LDAP server. In Release 11, use the iServer Integration Technology custom installation program to configure the Report Server Security Extension (RSSE) when implementing external registration.

- AcIntern

  Formerly used to convert Encyclopedia volume user security from external to internal registration from an LDAP or other system to iServer. In Release 11, no comparable functionality exists. For more information on how to install and configure RSSE in the Release 11 environment, refer to Chapter 11, "Configuring iServer security," in *Configuring BIRT iServer* and Chapter 10, "Using Java Report Server Security Extension" in *Using BIRT iServer Integration Technology.*

- AcMode

  Formerly used to put an Encyclopedia volume in and out of online backup mode. In Release 11, a dynamic backup no longer requires putting the system into online backup mode. The administrator performs an Encyclopedia volume metadata backup using the tools provided by the third-party RDBMS, which provides comparable features. The administrator uses standard operating system or other third-party tools to backup the data files.

- AcEncycUpgrade

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

For more information on how to upgrade iServer and Encyclopedia volume to Release 11, refer to the appropriate iServer installation manual, either *Installing BIRT iServer for Windows* or *Installing BIRT iServer for Linux and UNIX*.

Using JDBC to connect to an Encyclopedia volume database

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

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

An administrator, who decides to customize iServer to connect to a database other than the OOTB PostgreSQL database, must ensure that the JDBC driver returns adequate information about the types on the database. At a minimum, the database must return the common data types, such as integer, floating-point, and character. If the database does not return these common data types, then the database administrator must customize the database mapping framework to specify the types.

The JDBC driver must also support the following features:

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

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

API Compatibility

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

About international character sets

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

Administrative reports

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

Supported operating systems

Actuate BIRT iServer Release 11 Service Pack 2 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 Oracle, or a pre-existing PostgreSQL instance to store Encyclopedia volume metadata. This chapter describes how to install a new instance of BIRT iServer Release 11, using the out-of-the-box (OOTB) PostgreSQL RDBMS.

For more information about installing BIRT iServer using an alternative data store, such as Oracle or a pre-existing PostgreSQL database, see Chapter 3, "Installing BIRT iServer using an alternative database," later in this book. For more information about upgrading an existing BIRT iServer installation, see Chapter 4, "Upgrading BIRT iServer," later in this book. For information about the new Release 11 BIRT iServer System architecture, see Chapter 1, "Understanding Actuate BIRT iServer architecture," earlier in this book.

Creating a dedicated user account for installing and running 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.

Before installing iServer, create a user account with 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 user account. Creating a special user account for running Actuate iServer System is the recommended practice. However, you can install the software from an existing account.

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

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

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult Actuate Support at the time of installation if you have any questions about these backup procedures to protect against the possibility of catastrophic failure. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 3, "Configuring an Encyclopedia volume," in *Configuring BIRT iServer.*

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

If you are a purchasing customer, you can download iServer from an Actuate ftp site or install from a product DVD. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

```
http://www.birt-exchange.com
```

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 HP-UX 11i installation

Minimum hardware requirements for HP-UX 11i are a PA-RISC 2.0 processor and 1024 MB of RAM. For more information about HP-UX system requirements, see the Supported Products and Obsolescence Policy on the Actuate web site at the following URL:

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

Working with large files on HP-UX-11i

Actuate works with large files, greater than two gigabytes, on all platforms. You must reconfigure the HP-UX 11i file system as a large file system for this feature to take effect. The disk where the file resides must be a local disk. Contact HP Customer Support for information about how to reconfigure the file system.

Installing Xvfb

Actuate does not provide an Xvfb library for HP-UX 11i. If you choose to use an X frame buffer configured by the BIRT iServer installation instead of using an existing X server, your machine must have the Xvfb binary installed.

How to configure X frame buffer

To configure X frame buffer after the Actuate iServer System installation, perform the following steps:

1. Download and install the Xvfb patch from HP.
2. Create a symbolic link to the X frame buffer under AC_SERVER_HOME /xvfb/bin.
3. Start your Actuate iServer System.

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.

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.

Performing a new installation

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

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

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

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

How to perform a new stand-alone installation of BIRT iServer Release 11 in UNIX and Linux

To reduce network traffic, you can install BIRT iServer on the same host machine as your database. You can install BIRT iServer on a different host machine to dedicate resources to the database.

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

1. Insert the installation DVD, mount the DVD device on your system, and set the working directory to the mounted DVD directory, or download the required 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.

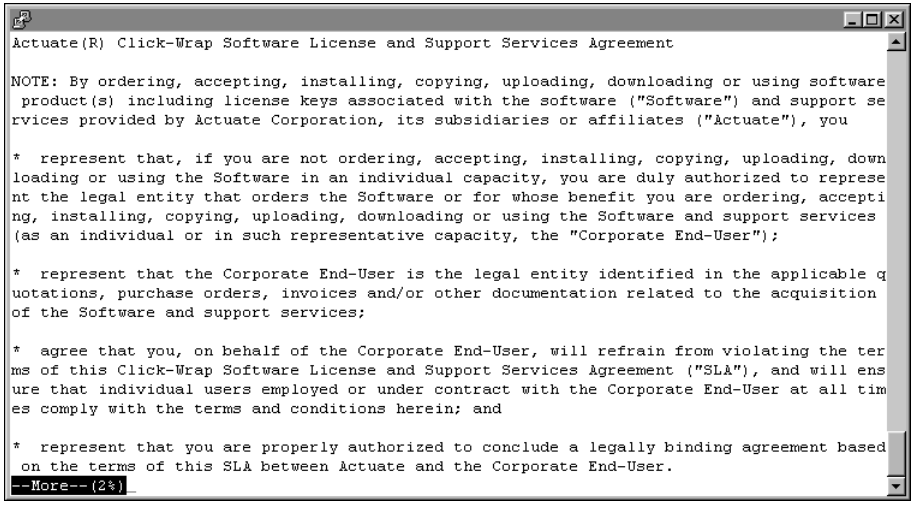

Figure 2-1 Reviewing the license agreement

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.

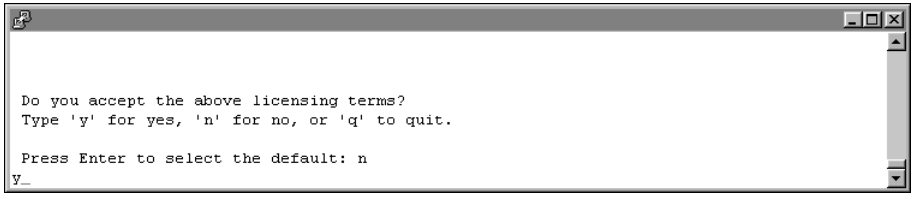

Figure 2-2 Accepting the licensing terms

5. The introduction to the installation appears, as shown in Figure 2-3.

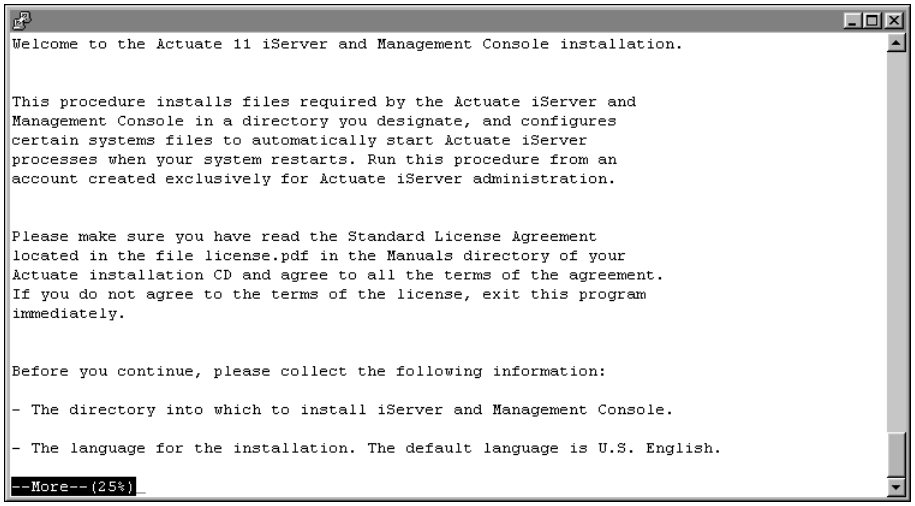

Figure 2-3 Reviewing the introductory information

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

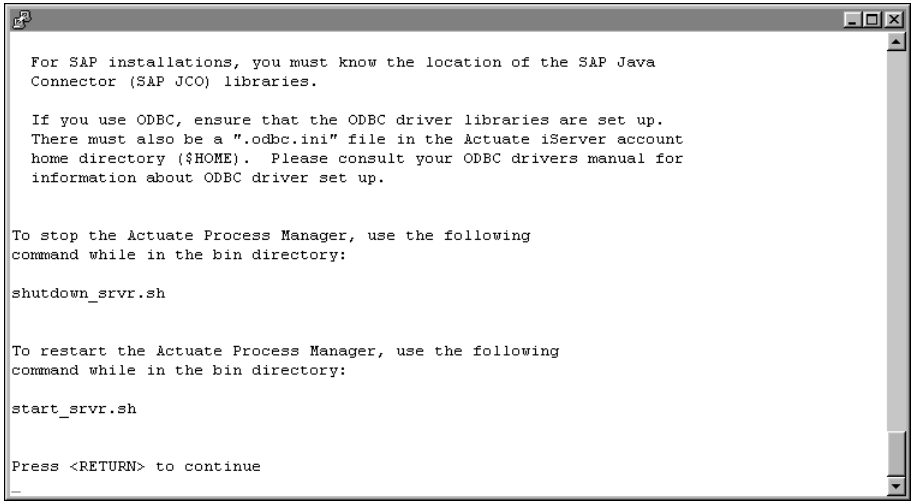

Figure 2-4 Finishing the review of introductory information

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.

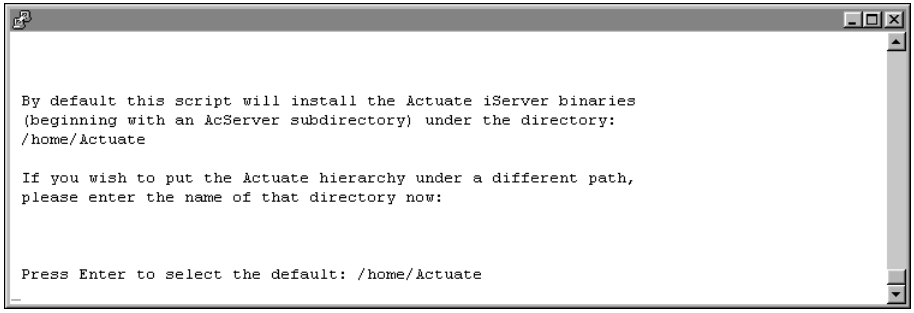

Figure 2-5 Specifying the installation directory

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.

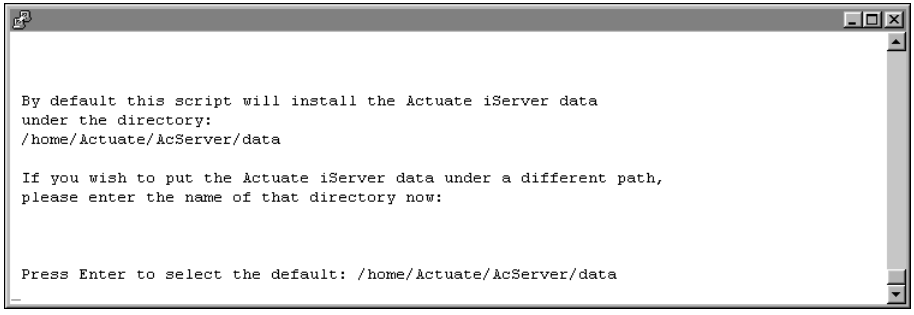

Figure 2-6 Specifying the data installation directory

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.

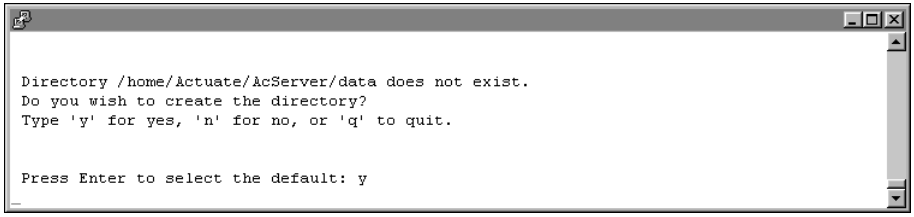

Figure 2-7 Creating the AC_DATA_HOME directory

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.

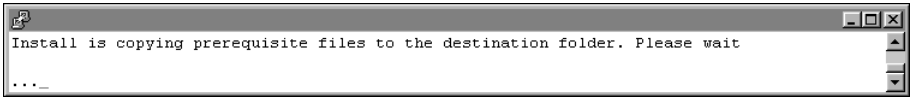

Figure 2-8 Copying prerequisite files

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.

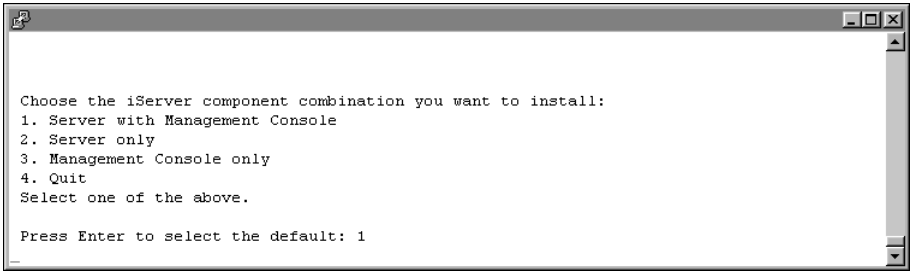

Figure 2-9 Choosing the components to install

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.".

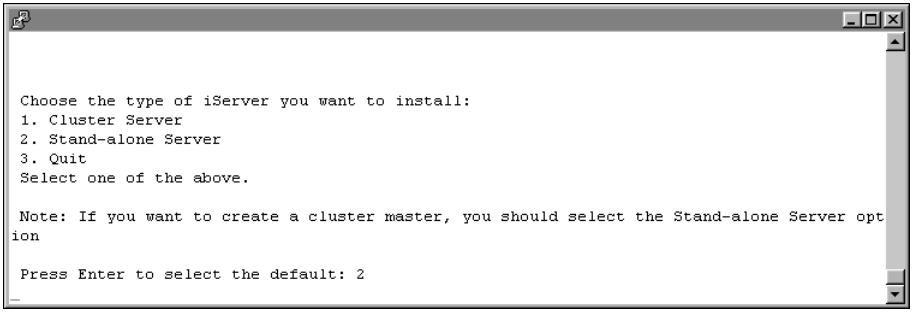

Figure 2-10 Specifying the type of iServer to install

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.

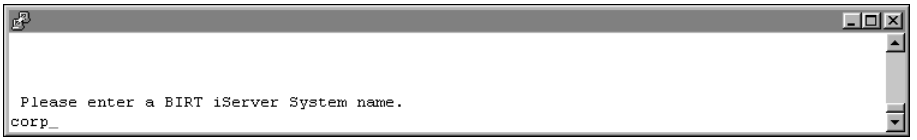

Figure 2-11 Specifying the BIRT iServer System name

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

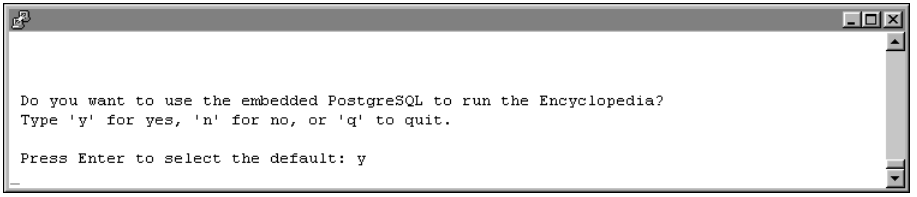

Figure 2-12 Choosing the embedded PostgreSQL

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).

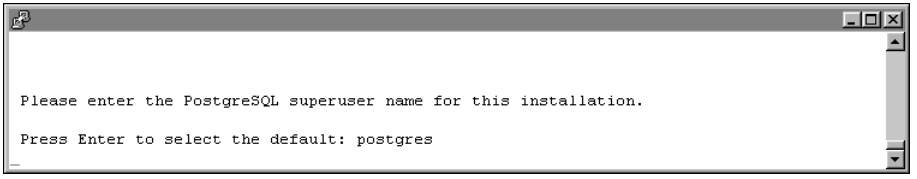

Figure 2-13 Choosing the PostgreSQL superuser name

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.

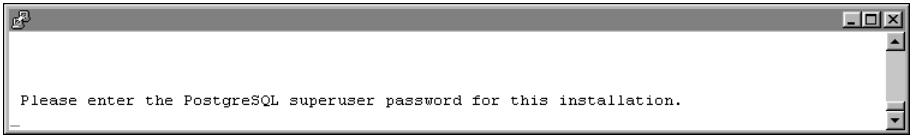

Figure 2-14 Typing the PostgreSQL superuser password

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

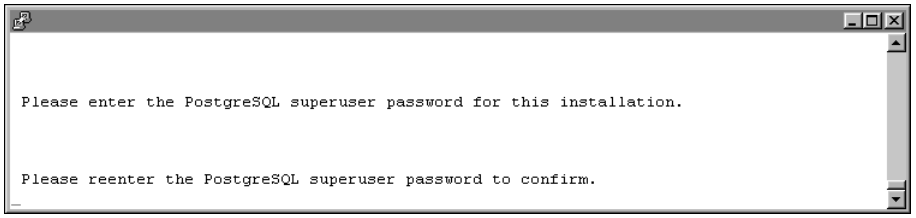

Figure 2-15 Re-entering the PostgreSQL superuser password

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.

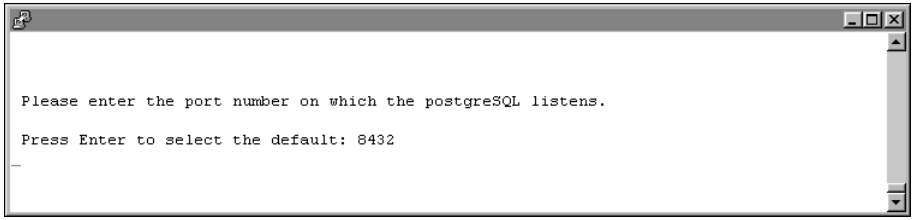

Figure 2-16 Choosing the PostgreSQL port number

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.

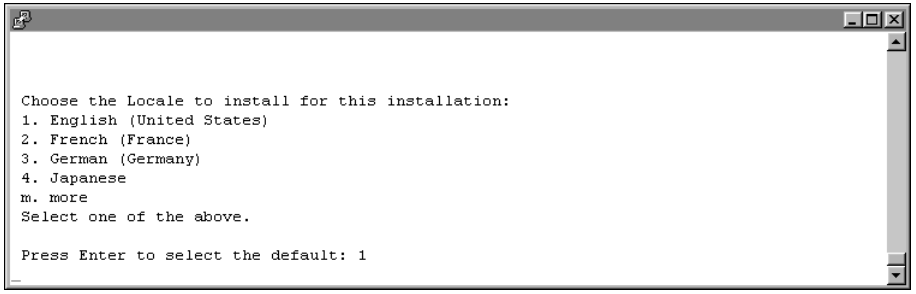

Figure 2-17 Specifying a locale

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.

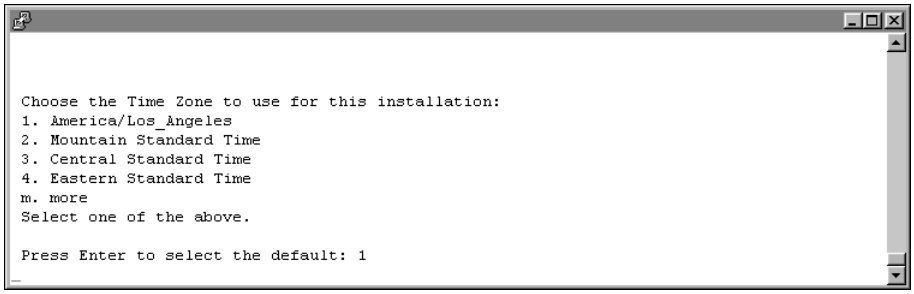

Figure 2-18 Specifying a time zone

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.

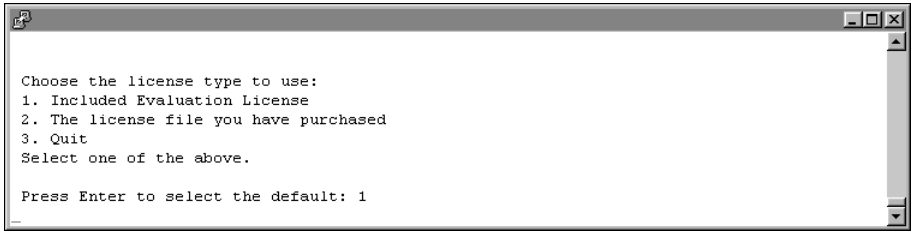

Figure 2-19 Specifying license type

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.

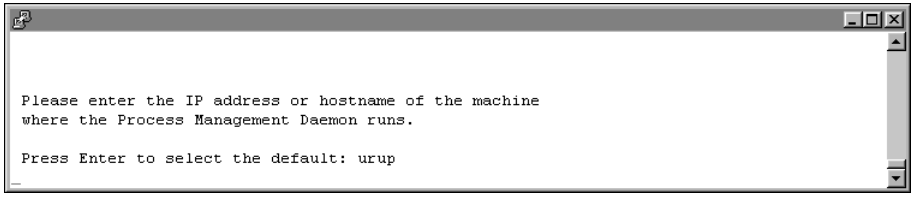

Figure 2-20 Specifying the hostname that Management Console uses to contact the PMD

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.

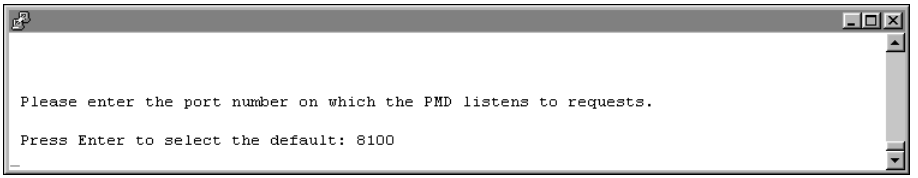

Figure 2-21 Specifying the port number on which the PMD listens

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.

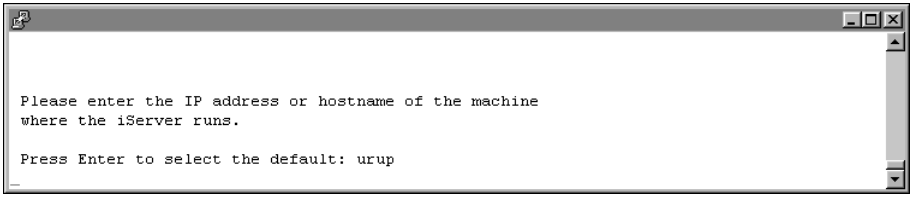

Figure 2-22 Specifying the machine on which the iServer runs

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.

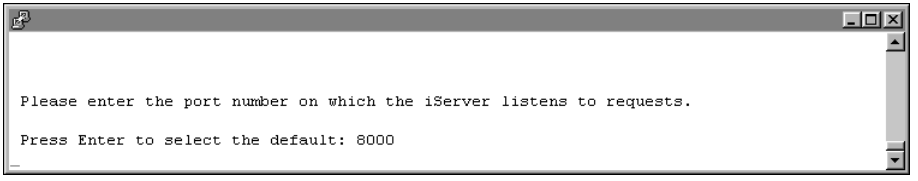

Figure 2-23 Specifying the port number on which the iServer listens

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.

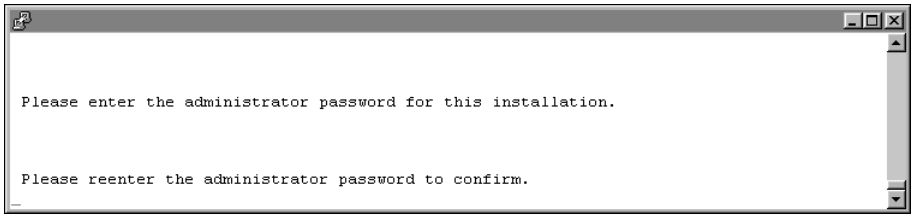

Figure 2-25 Re-entering the iServer administrator password

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.

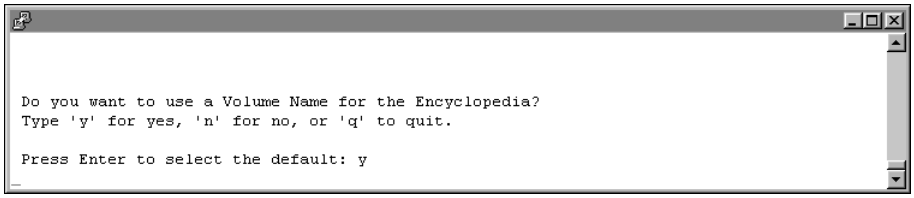

Figure 2-26 Specifying whether to use an Encyclopedia volume name

29 Press Enter to accept the default Encyclopedia volume name, the name of your machine, as shown in Figure 2-27. Alternatively, type a different Encyclopedia volume name.

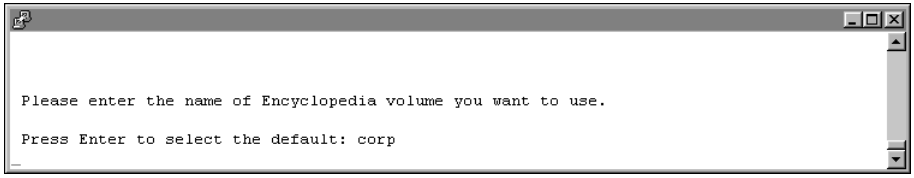

Figure 2-27 Specifying the Encyclopedia volume name

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

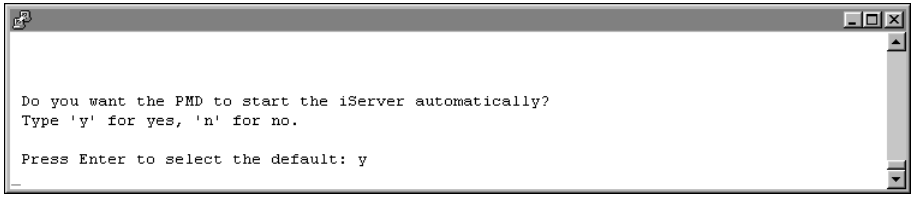

Figure 2-28 Specifying whether to start iServer automatically

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.

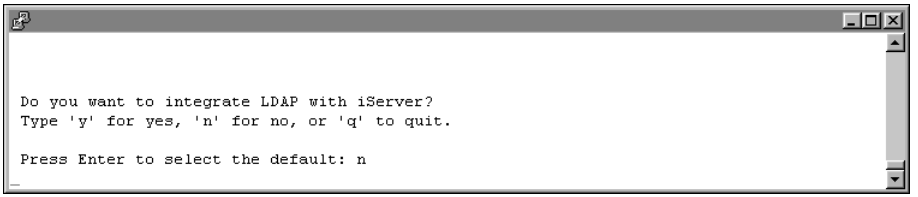

Figure 2-29 Specifying whether to integrate LDAP with iServer

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.

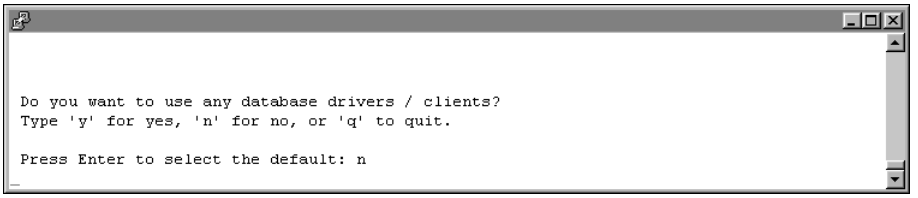

Figure 2-30 Specifying whether to use database drivers/clients

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.

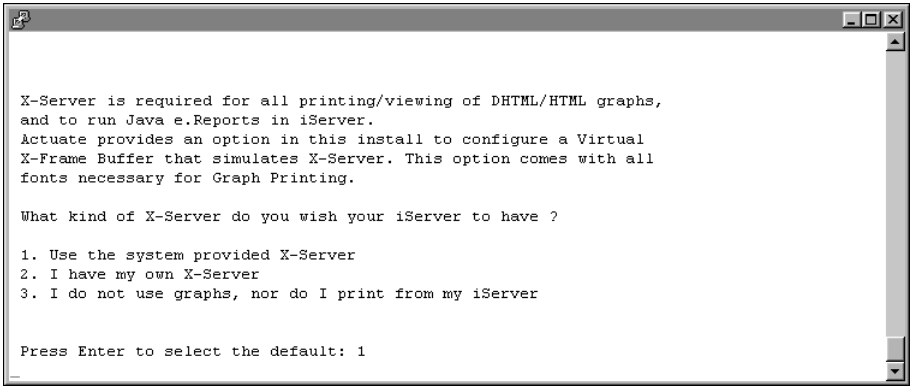

Figure 2-31 Specifying what kind of X-Server to use, if any

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.

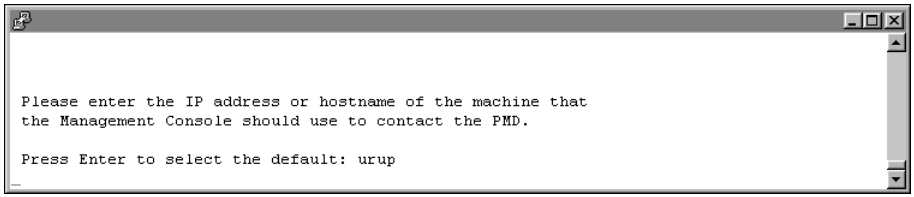

Figure 2-32 Specifying the hostname that Management Console uses to contact the PMD

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.

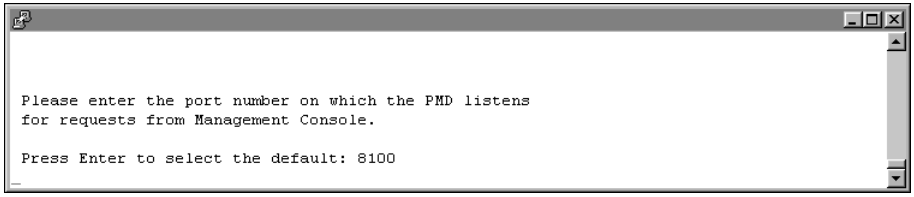

Figure 2-33 Specifying the port number for the PMD to listen for requests from Management Console

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.

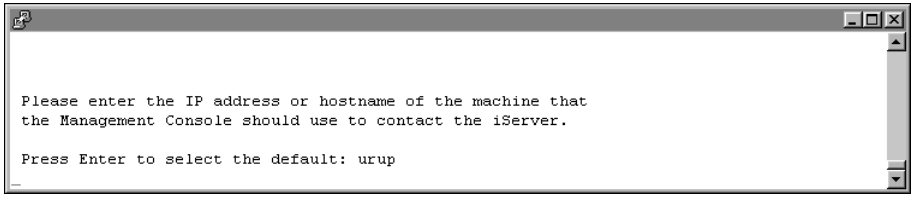

Figure 2-34 Specifying the hostname that Management Console uses to contact iServer

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.

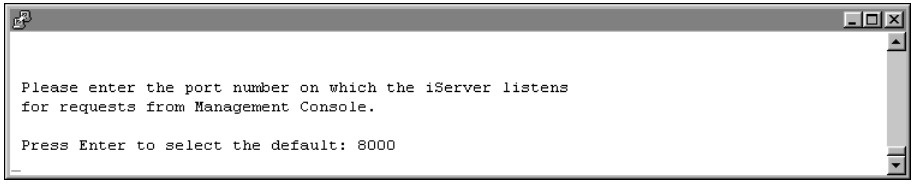

Figure 2-35 Specifying the port number on which iServer listens for requests from Management Console

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.

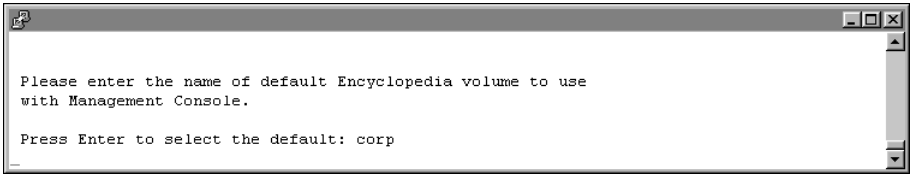

Figure 2-36 Specifying the name of the default Encyclopedia volume

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.

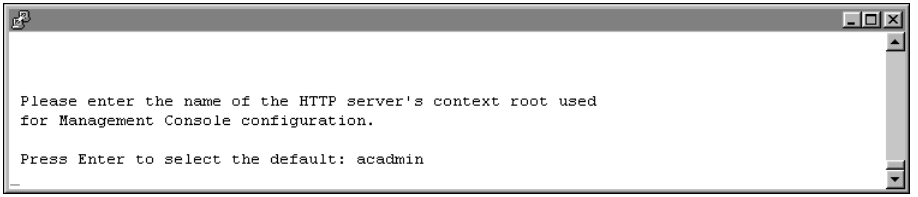

Figure 2-37 Specifying the name of the HTTP server context root

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.

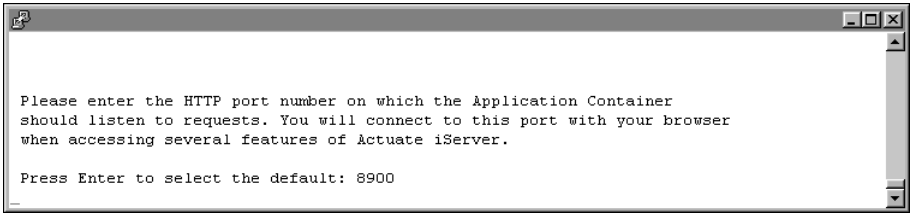

Figure 2-38 Specifying the application container listening port number

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.

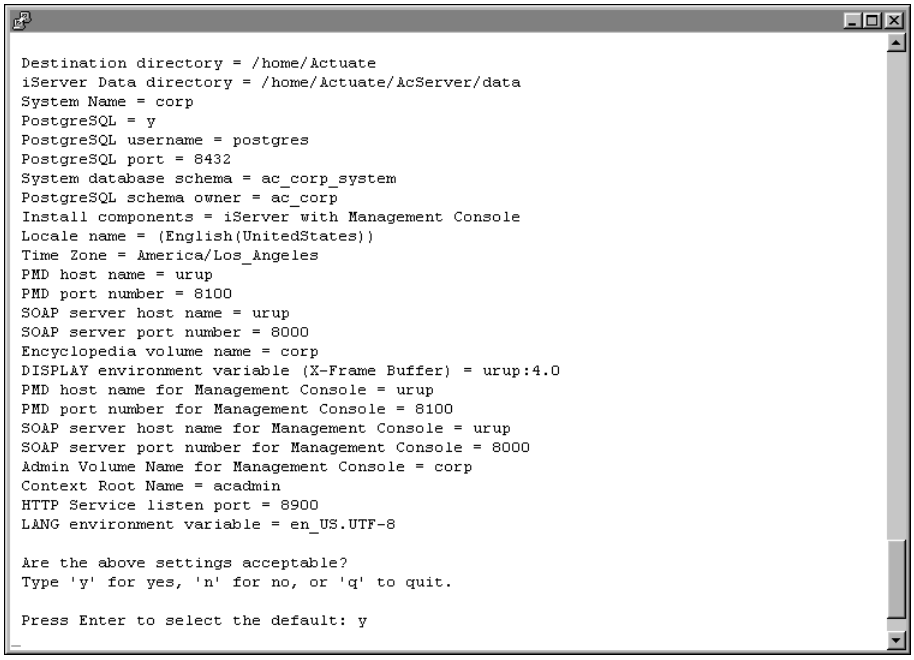

Figure 2-39 Reviewing the installation settings

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

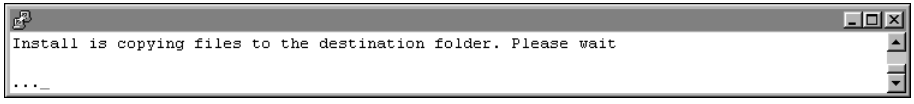

Figure 2-40 Viewing iServer installation progress

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.

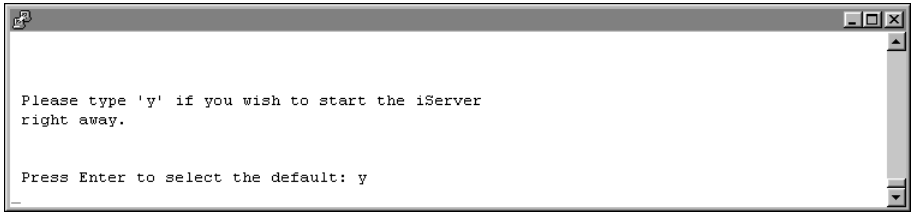

Figure 2-41 Specifying whether to start iServer

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.

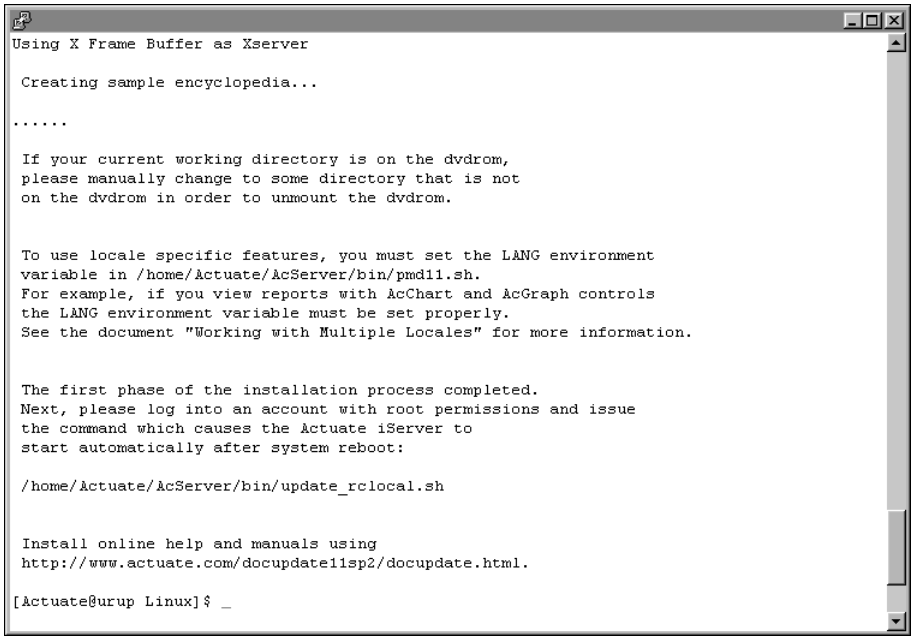

Figure 2-42 Viewing information about localization, logging in, and installing online help

Understanding the iServer installation environment

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

About migrating an earlier iServer release to Actuate 11

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

The upgrade program preserves any previous iServer configuration information and reuses the earlier settings. The installation program resolves any differences in default values between releases, ignoring old configuration defaults in favor of new default values. For example, the installation does not prompt the user for port information and machine name. The installation detects the current port numbers and machine name and keeps those settings.

About running different releases on the same machine

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

Actuate does not support running multiple releases from the same version on a machine. For example, you cannot run Release 10 and Release 10 Service Pack 1 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 it.

Using a test environment

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

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

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

Setting up a production staging area

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

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

Complete the following general tasks to test Actuate 11:

- Install BIRT iServer Release 11 software in a production staging area.
- Install Actuate 11 desktop software on the test user machines. Using separate folders, you can install Actuate 11 desktop software in conjunction with the earlier desktop software. Users can continue to use the existing Actuate software in production while testing the Actuate 11 desktop software.

- Verify that the Actuate 11 production staging environment works correctly.
- Install the remaining Actuate 11 desktop products, if you installed a subset earlier.
- Verify that all the Actuate 11 desktop products work correctly.
- Begin setting up a production environment, described in the following section.

Setting up a production environment

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

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

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

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

  - If you are replacing your Actuate 10 system, upgrade Encyclopedia volumes by installing Actuate 11 in the same directory, overwriting Actuate 10.

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

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

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

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

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 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, such as Oracle or a pre-existing PostgreSQL instance.

For all database systems other than the 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 Encyclopedia volume schema owner and iServer application user credentials. The iServer installation program creates the necessary volume 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.

Before installing iServer, create a user account with 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 user account. Creating a special user account for running Actuate iServer System is the recommended practice. However, you can install the software from an existing account.

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.

Creating the system and Encyclopedia volume schemas 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 the following schema owner and application user accounts with appropriate privileges:

- System schema owner
- 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 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 schema owner and grant privileges in a pre-existing Oracle database. These steps are not necessary when adding an Encyclopedia volume to an existing schema.

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

Creating the system schema owner

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

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

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

(continues)

```
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 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.

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

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

Creating the iServer application user

iServer connects to the database as an application user. The application user requires only the privileges necessary to perform basic SQL data manipulation 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 in a pre-existing PostgreSQL database

The following SQL scripts provide an example of DDL statements that create the schema owner and grant privileges in a pre-existing PostgreSQL database. 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, typically named postgres, as a user with full administrator privileges 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;
```

Connect to the iserver application database, not the system database, as a user with full administrator privileges, and execute the following SQL command:

```
CREATE LANGUAGE plpgsql;
```

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 system schema owner named ac_corp_system, with appropriate privileges to connect to the previously created database named iserver. Connect to the PostgreSQL system database, typically named postgres, 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 schema owner

In an iServer installation, the Encyclopedia schema owner must have the same name as the Encyclopedia schema. The Encyclopedia 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 schema owner named ac_corp with appropriate privileges to connect to a database named iserver. Connect to the PostgreSQL system database, typically named postgres, 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 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 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 and grant privileges to use that schema to the application user named iserver. Connect to the application database, not the 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 schema named ac_corp and grant privileges to use the schema to the application user named iserver. Connect to the application database, not the

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;
```

Backing up iServer system and Encyclopedia volume metadata

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

An iServer system administrator must take all necessary precautions to ensure that the schemas are properly backed up to safeguard the metadata. Please consult Actuate Support at the time of installation if you have any questions about these backup procedures to protect against the possibility of catastrophic failure. For information on the recommended procedures to back up an iServer system and Encyclopedia volume schemas in the Release 11 environment, refer to Chapter 2, "Configuring an Encyclopedia volume," in *Configuring BIRT iServer*.

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

If you are a purchasing customer, you can download iServer from an Actuate ftp site or install from a product DVD. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

```
http://www.birt-exchange.com
```

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

The following sections describe how to install BIRT iServer Release 11 using 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 an example. In Oracle, there is a one-to-one relationship between a user and a schema. A schema is not a separate entity. The following example creates a schema owner named urup and a database user named iserver.

During the iServer installation, the administrator provides the schema owner and database user credentials. The iServer installation program creates the necessary volume database structures, then loads the metadata.

How to install an Encyclopedia volume that uses an alternative database

To install iServer, perform the following steps:

1. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required 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.

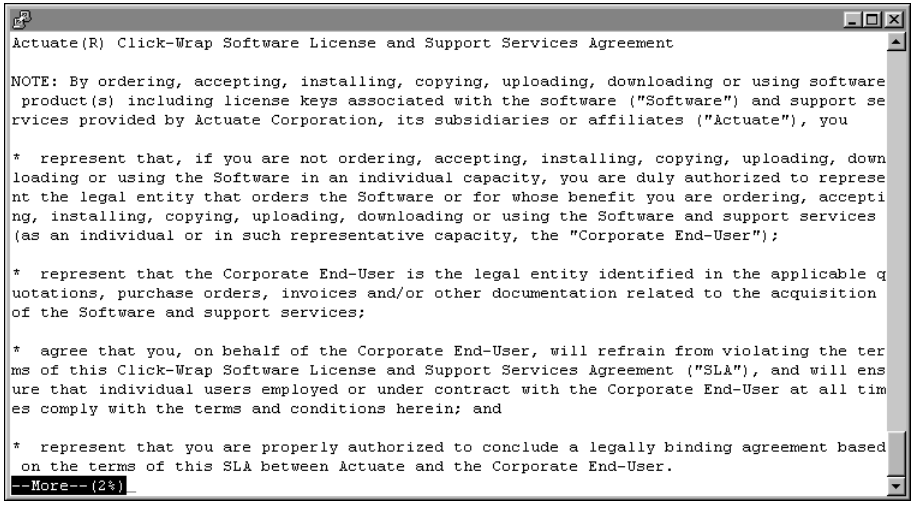

Figure 3-1 Reviewing the license agreement

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.

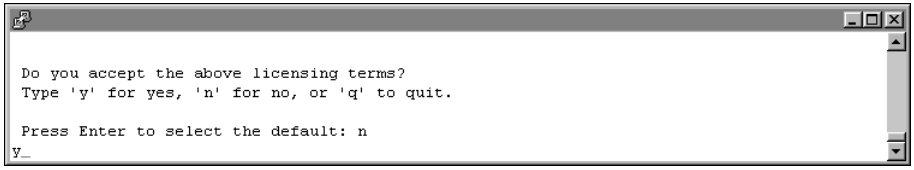

Figure 3-2 Accepting the licensing terms

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

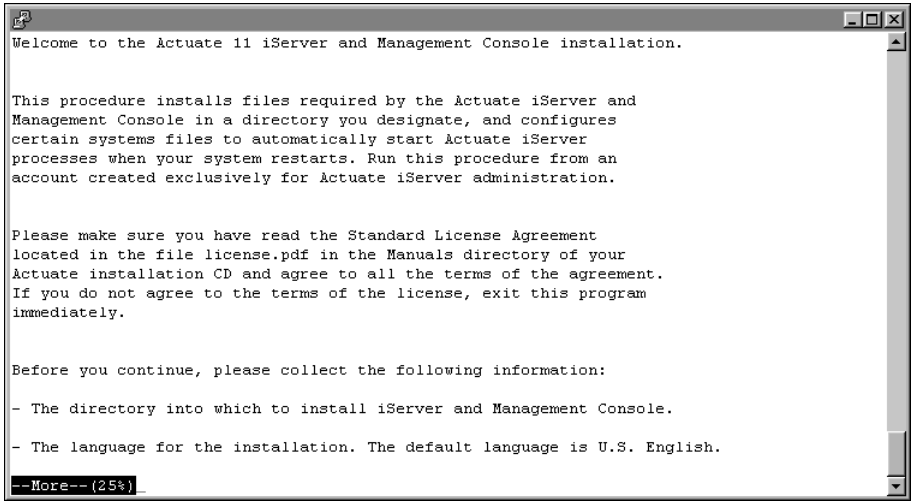

Figure 3-3 Reviewing the introductory information

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

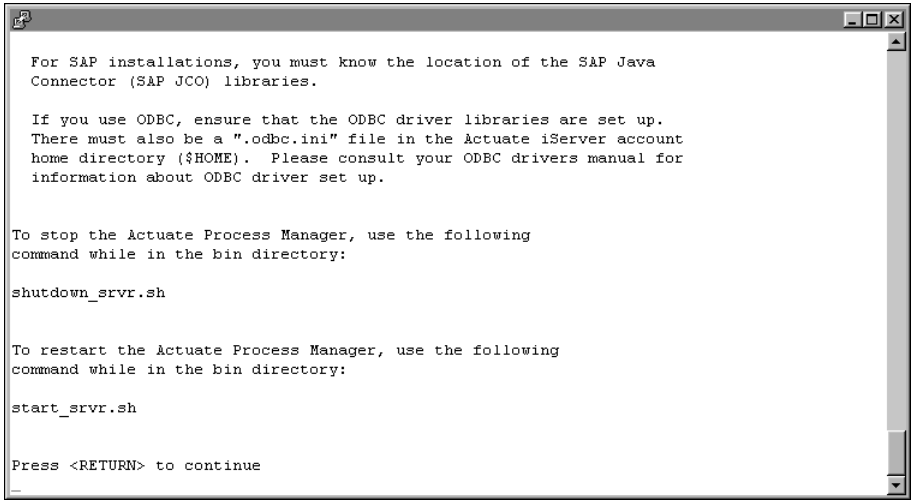

Figure 3-4 Finishing the review of introductory information

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.

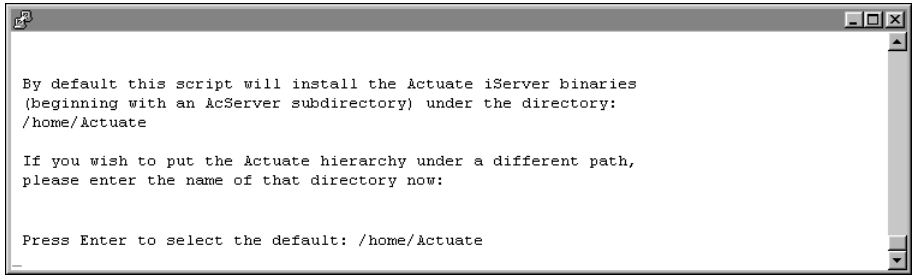

Figure 3-5 Specifying the installation directory

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, 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.

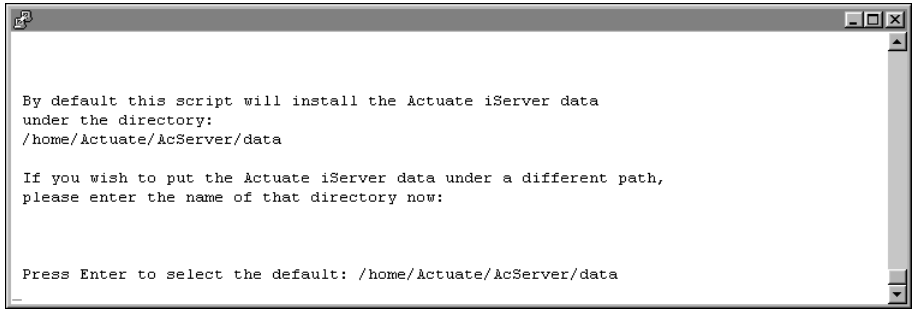

Figure 3-6 Specifying the data installation directory

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.

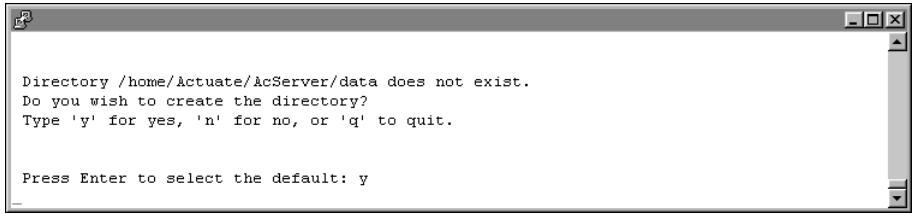

Figure 3-7 Creating the AC_DATA_HOME directory

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

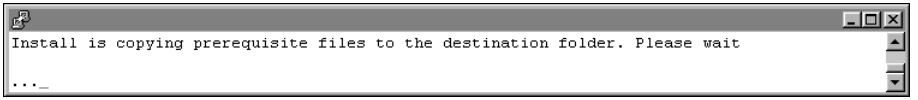

Figure 3-8 Copying prerequisite files

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.

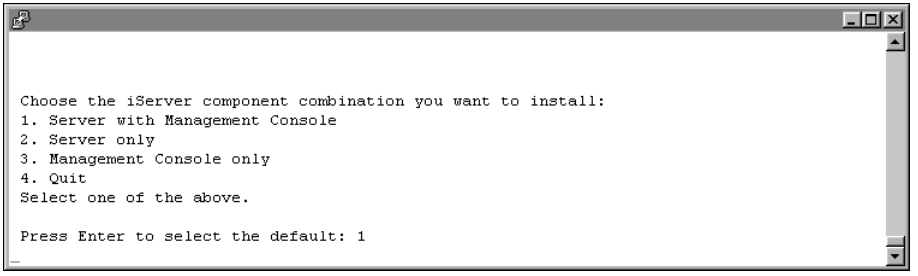

Figure 3-9 Choosing the iServer component combination

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.

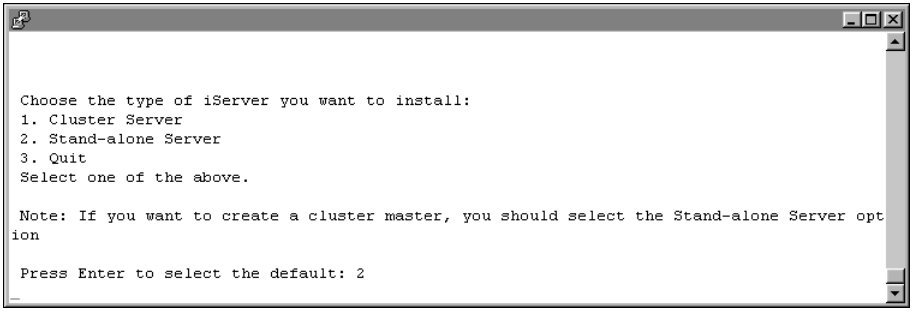

Figure 3-10 Specifying the type of iServer to install

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 to store Encyclopedia volume metadata.

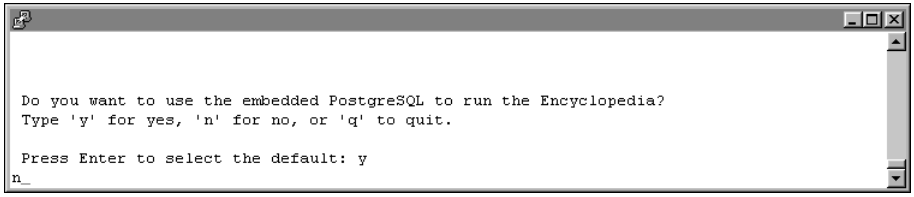

Figure 3-12 Choosing to not use Embedded PostgreSQL

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

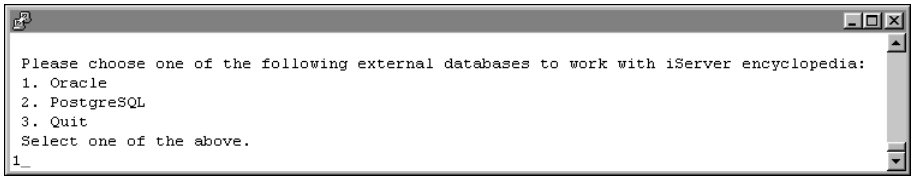

Figure 3-13 Choosing an external database for iServer Encyclopedia

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 server and press Enter, as shown in Figure 3-15.

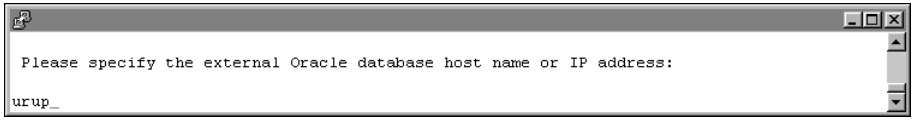

Figure 3-15 Specifying the external Oracle database server name

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.

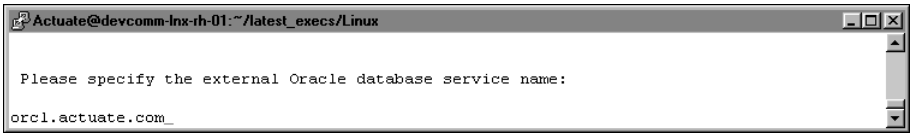

Figure 3-17 Specifying the external Oracle database service name

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

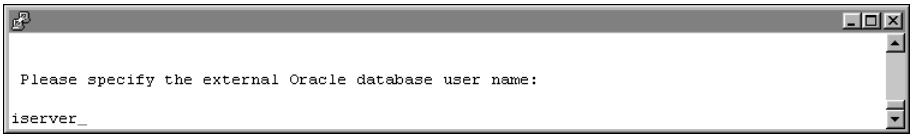

Figure 3-18 Specifying the external Oracle database user name

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

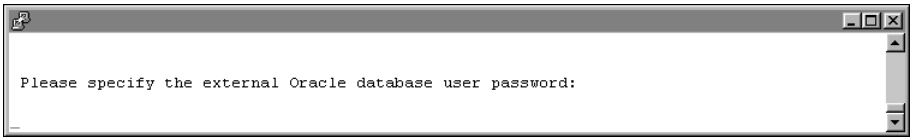

Figure 3-19 Specifying the external Oracle database user password

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

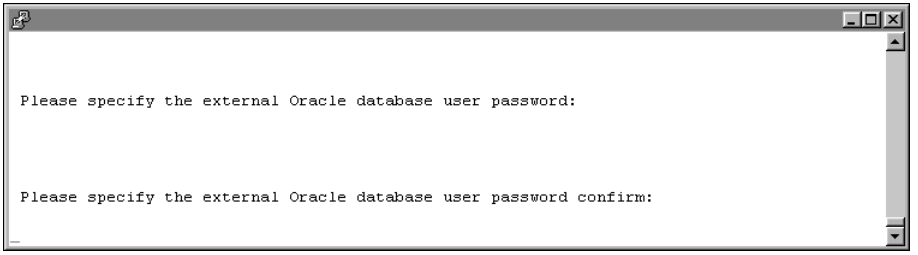

Figure 3-20 Confirming the external Oracle database user password

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

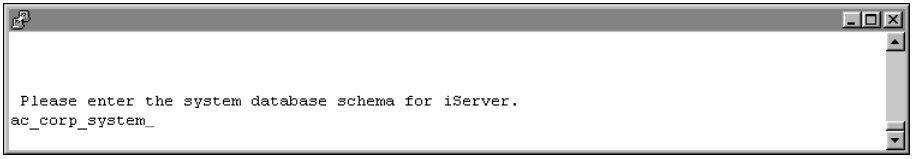

Figure 3-21 Specifying the system database schema

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

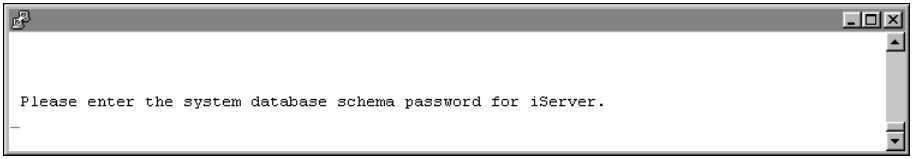

Figure 3-22 Specifying the system database schema password

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

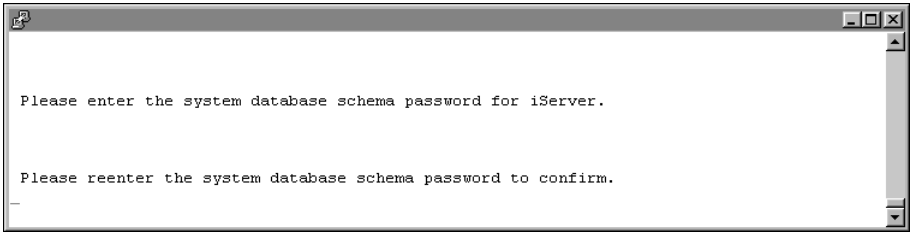

Figure 3-23 Confirming the system database schema password

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

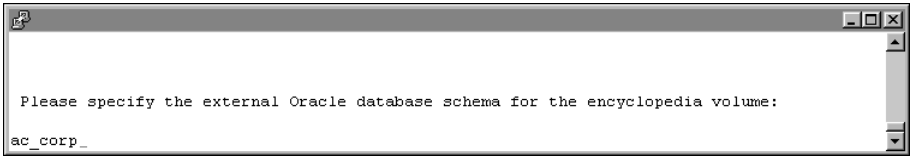

Figure 3-24 Specifying the schema for the Encyclopedia volume

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

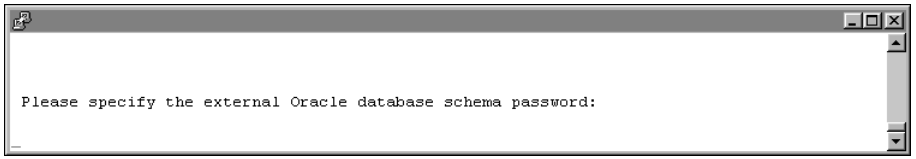

Figure 3-25 Specifying the external Oracle database schema password

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

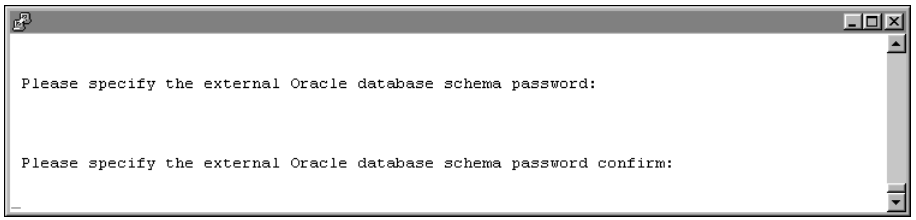

Figure 3-26 Confirm the external Oracle database schema password

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.

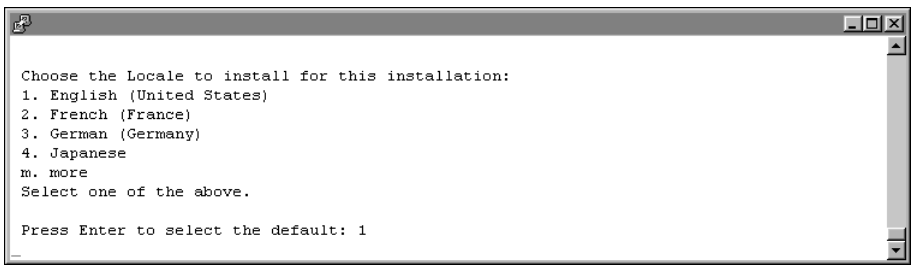

Figure 3-27 Specifying a locale

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.

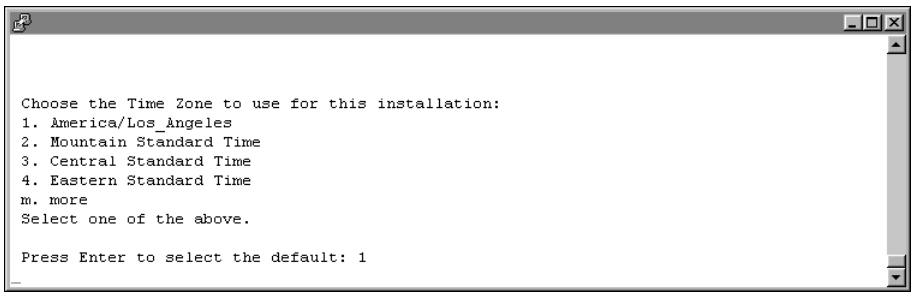

Figure 3-28 Specifying a time zone

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.

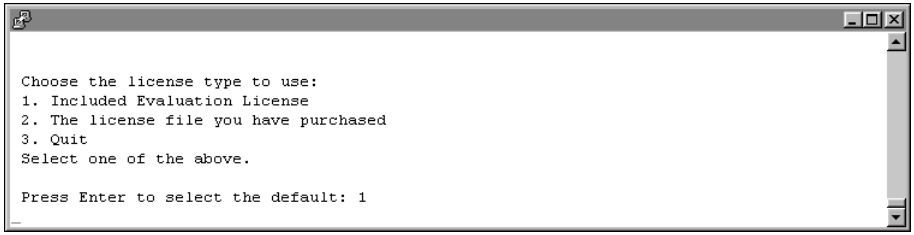

Figure 3-29 Specifying license type

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.

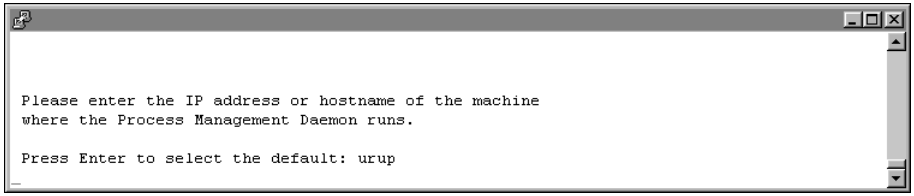

Figure 3-30 Specifying the hostname that Management Console uses to contact the PMD

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.

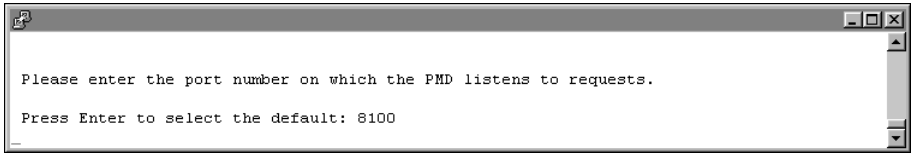

Figure 3-31 Specifying the port number on which the PMD listens

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.

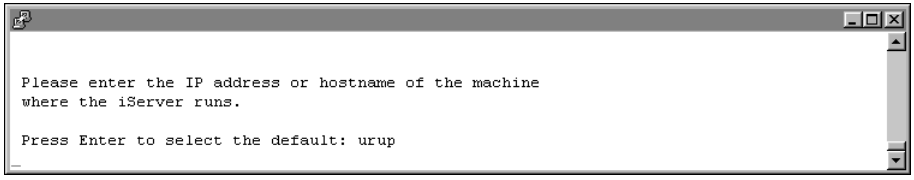

Figure 3-32 Specifying the machine on which the iServer runs

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

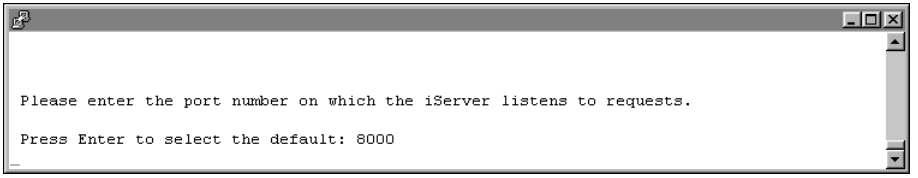

Figure 3-33 Specifying the port number on which the iServer listens

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.

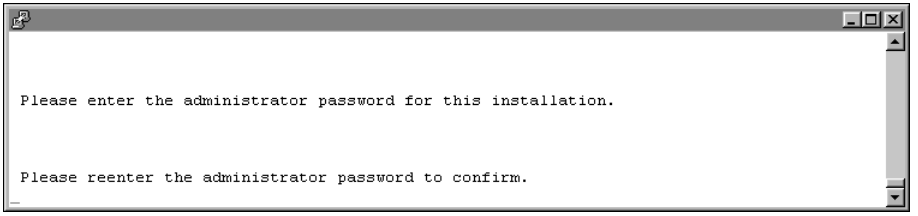

Figure 3-35 Re-entering the iServer administrator password

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.

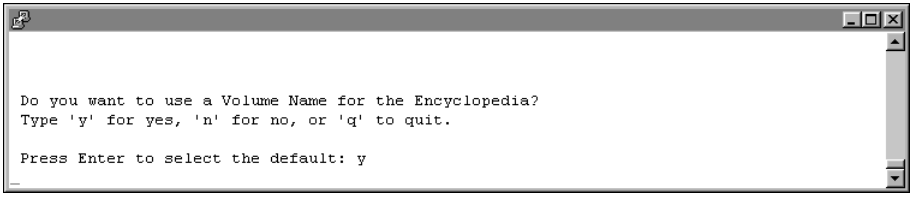

Figure 3-36 Specifying whether to use a volume name

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.

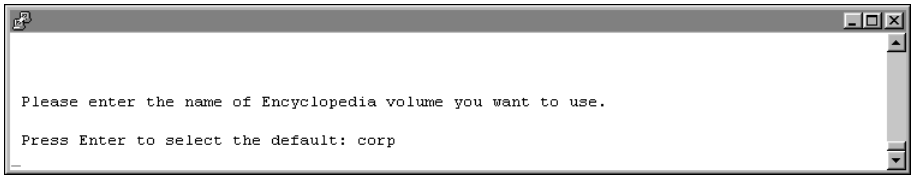

Figure 3-37 Specifying the Encyclopedia volume name

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

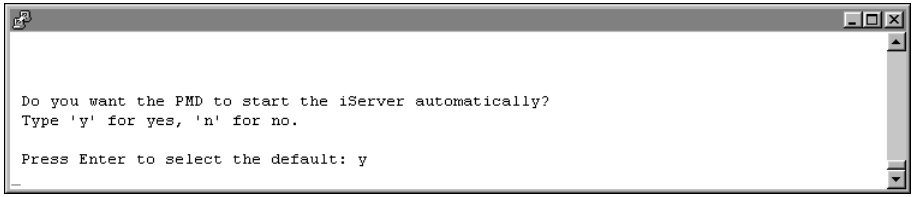

Figure 3-38 Specifying whether to start iServer automatically

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.

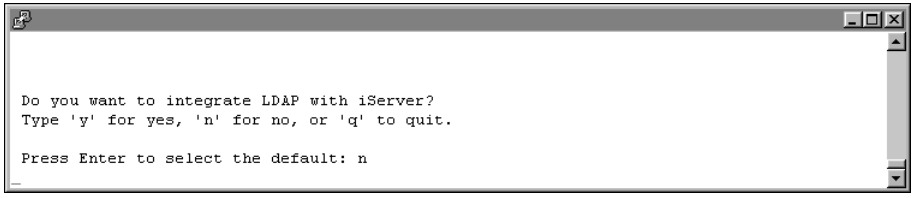

Figure 3-39 Specifying whether to integrate LDAP with iServer

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.

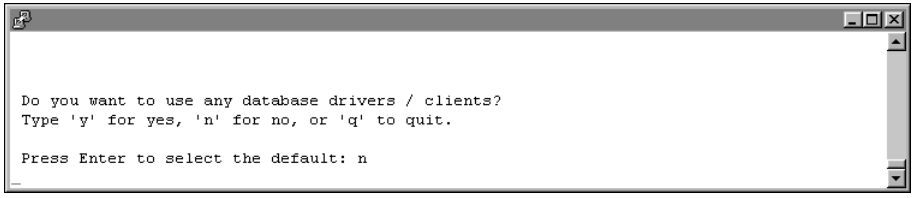

Figure 3-40 Specifying whether to use database drivers/clients

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.

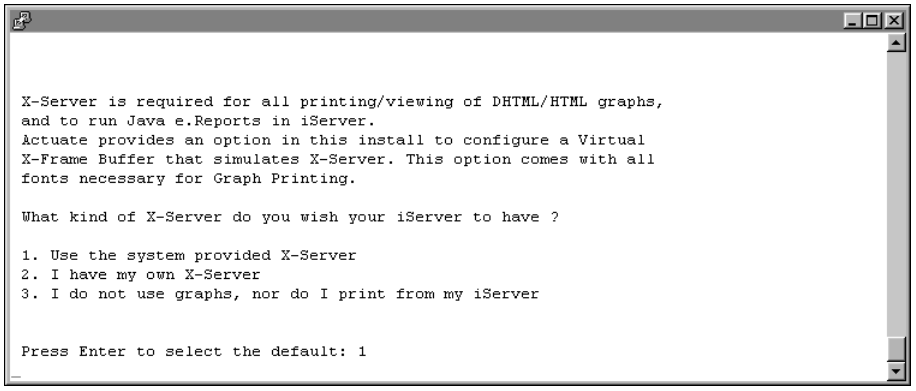

Figure 3-41 Specifying what kind of X-Server to use, if any

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.

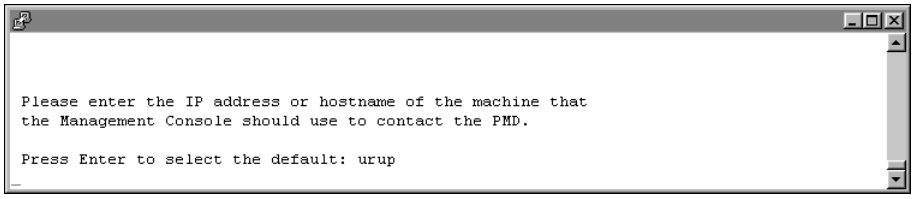

Figure 3-42 Specifying the machine host name that Management Console uses to contact the PMD

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.

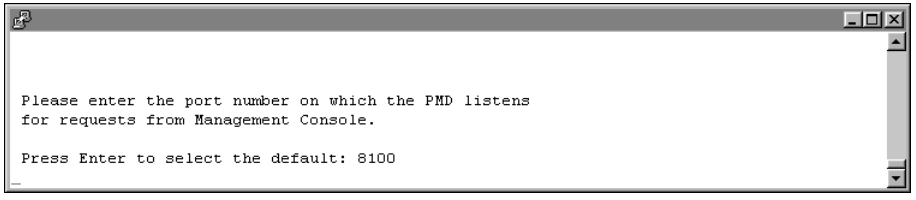

Figure 3-43 Specifying the port number for the PMD to listen for requests from Management Console

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.

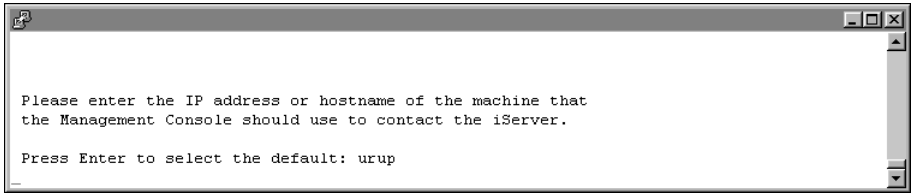

Figure 3-44 Specifying the name Management Console uses to contact iServer

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.

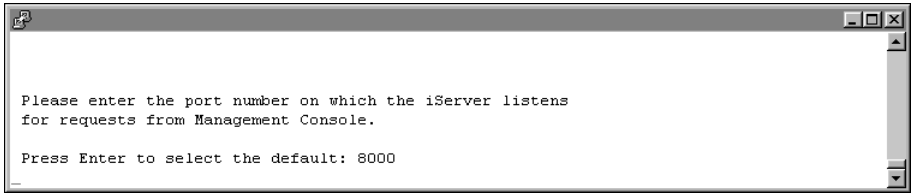

Figure 3-45 Specifying the port number for listening for requests from Management Console

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.

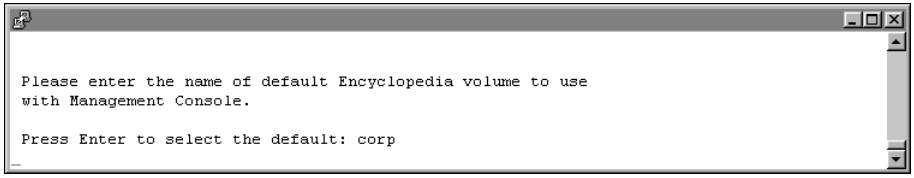

Figure 3-46 Specifying the name of the default Encyclopedia volume

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.

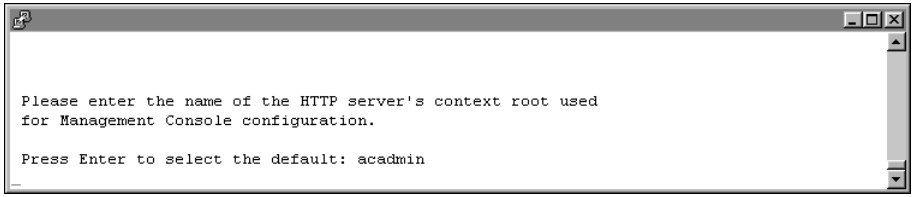

Figure 3-47 Specifying the name of the HTTP server context root

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.

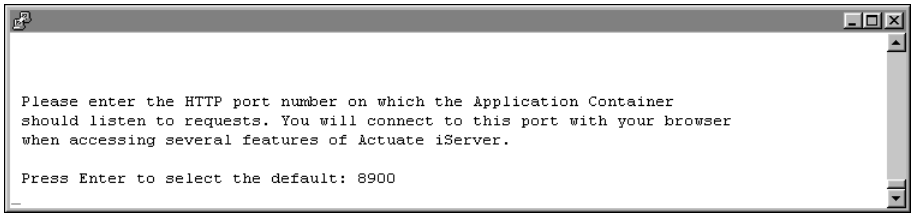

Figure 3-48 Specifying the application container listening port number

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.

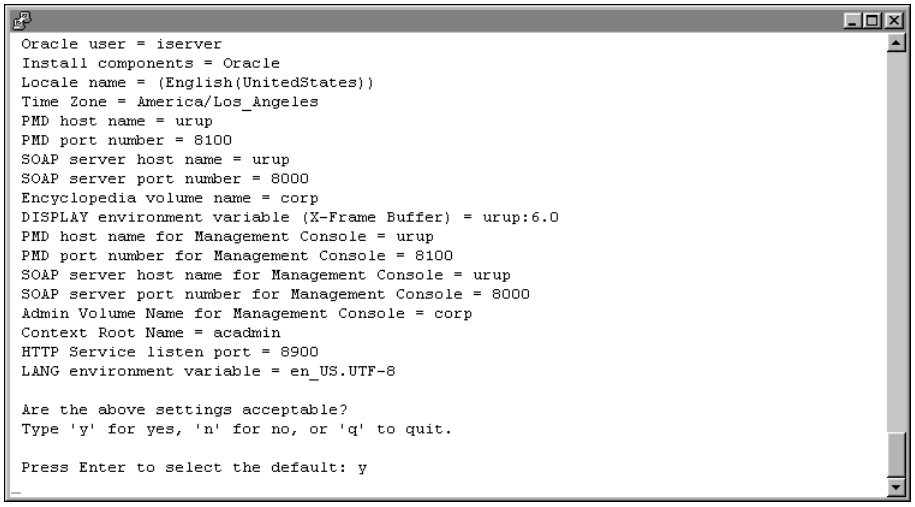

Figure 3-49 Reviewing your settings

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

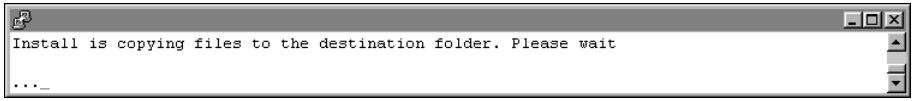

Figure 3-50 Copying iServer files to your destination folder

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.

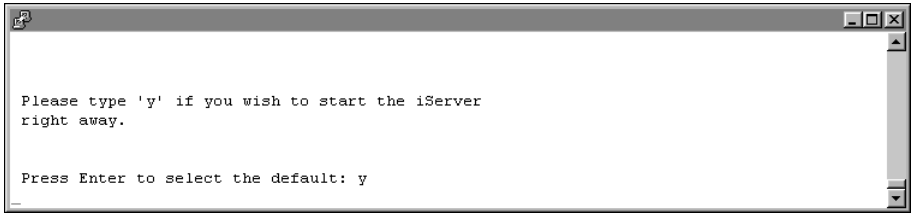

Figure 3-51 Specifying whether to start iServer

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.

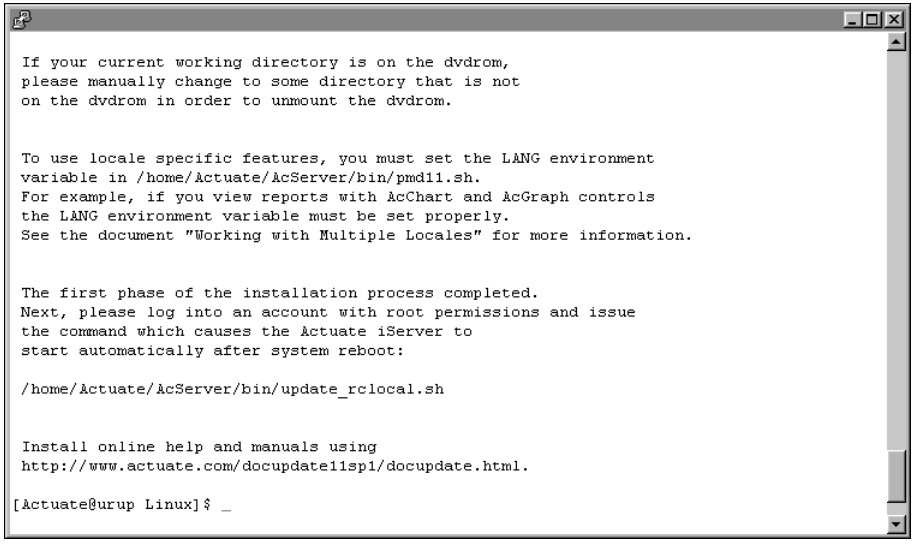

Figure 3-52 Specifying information about localization, logging in, and installing online help

Chapter

4

Upgrading BIRT iServer

This chapter discusses the following topics:

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

Preparing to upgrade BIRT iServer

When upgrading to BIRT iServer Release 11, the administrator must choose to use the out-of-the-box (OOTB) PostgreSQL database or another data store instance to store Encyclopedia volume metadata. The administrator upgrades iServer and the Encyclopedia volume database using one of the following options:

- Automatic in-place upgrade

  Upgrades an earlier iServer system in place, such as Release 10 Service Pack 1, automatically migrating one or more existing Encyclopedia volumes during the installation process.

- Manual in-place upgrade

  Upgrades an earlier iServer system from an earlier major release, such as Release 10 Service Pack 1, without automatically migrating any Encyclopedia volumes. During installation, the administrator chooses to migrate the volumes manually. After installation, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the volume.

  When updating an iServer system from an earlier Release 11 installation, such Release 11 Service Pack 1 to Service Pack 2, the administrator uses the Encyclopedia Data Store Upgrader utility to upgrade the volume.

- Manual side-by-side upgrade

  Upgrades a new Release 11 iServer system in the environment side-by-side with an earlier iServer System, such as Release 10 Service Pack 1. During installation, the administrator chooses whether to install the default volume with sample designs and documents. After installation, the administrator uses the Squirrel Data Exporter and Encyclopedia Data Store Administrator utilities to migrate the volumes manually from the older system to the Release 11 iServer.

Creating 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.

Before installing iServer, create a user account with 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 user account. Creating a special user account for running Actuate iServer System is the recommended practice. However, you can install the software from an existing account.

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

Backing up iServer system and Encyclopedia volume metadata

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

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

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

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

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

If you are a purchasing customer, you can download iServer from an Actuate ftp site or install from a product DVD. If you are evaluating BIRT iServer, you can download iServer from BIRT Exchange at the following location:

`http://www.birt-exchange.com`

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

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

Performing an automatic or manual in-place upgrade

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

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

The automatic option also performs the following operations during installation:

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

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

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

When updating an iServer system from an earlier Release 11 installation, such as Release 11 Service Pack 1 to Service Pack 2, the administrator uses the Encyclopedia Data Store Upgrader utility to upgrade the volume. This upgrade modifies the previous Release 11 to include any changes to the system since the major 11.0 release and earlier service packs. The iServer installation program creates a script to use for setting properties and running the utility from the

command line. The Data Store Upgrader must be run for each Encyclopedia data store schema.

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

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

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

1. Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
   - encyc directories from all nodes
   - acserverconfig.xml in the /etc directory
   - acpmdconfig.xml in the /etc directory
   - RSSE code and associated files if you use the Open Security option
2. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required files.
3. 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.
4. The license agreement appears, as shown in Figure 4-1.

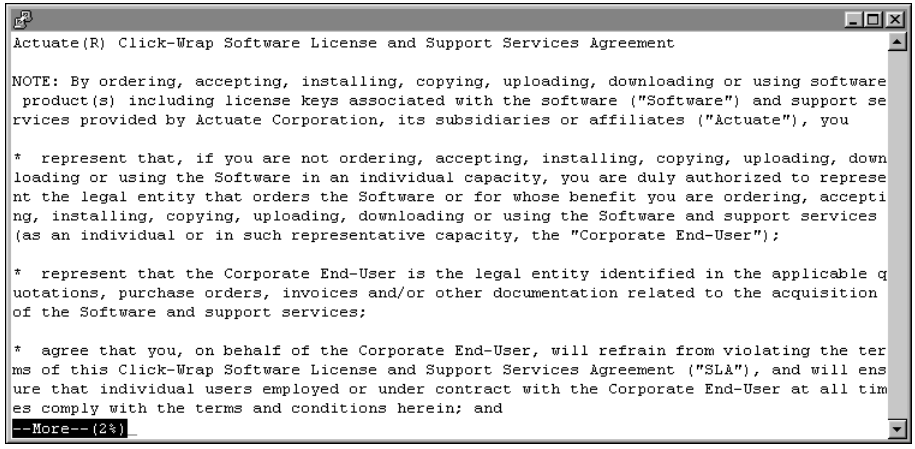

Figure 4-1 Reviewing the license agreement

5 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

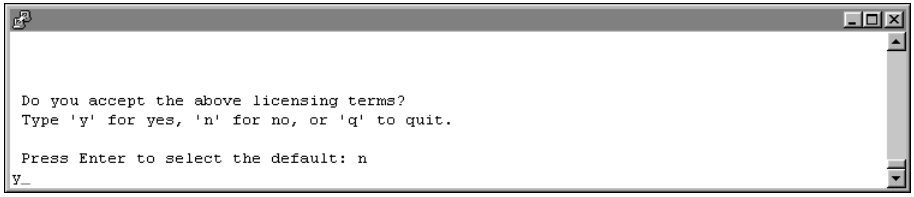

Figure 4-2 Accepting the licensing terms

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

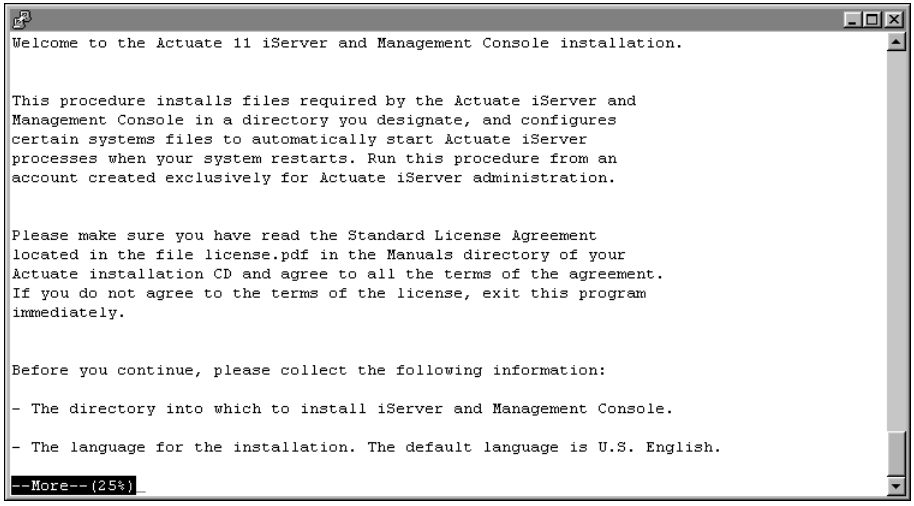

Figure 4-3 Reviewing the introductory information

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

8 Press Enter to select the default directory as the location of the Actuate 10 installation. Alternatively, type the appropriate path and press Enter, as shown in Figure 4-5. The name of the installation directory does not change when you upgrade. For example, the installation directory in Figure 4-5 is still /home/Actuate/AcServer after upgrading from Actuate 10 to Actuate 11.

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.

Unlike a new Actuate BIRT iServer 11 installation where all run-time data is in AC_DATA_HOME, the following files and folders remain in their original

directories under AC_SERVER_HOME rather than AC_DATA_HOME after the upgrade:

- iServer configuration
- Default encyclopedia folder
- iServer log files
- iServer temporary folder

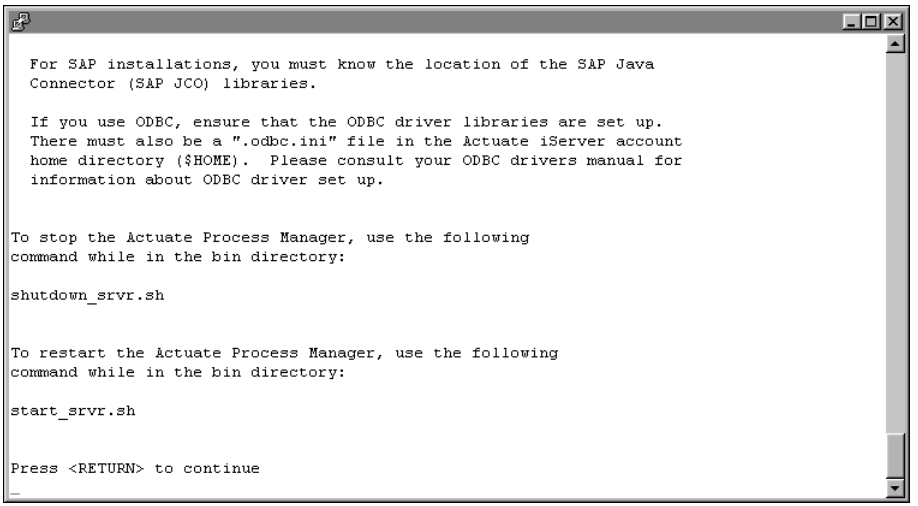

Figure 4-4 Finishing the review of introductory information

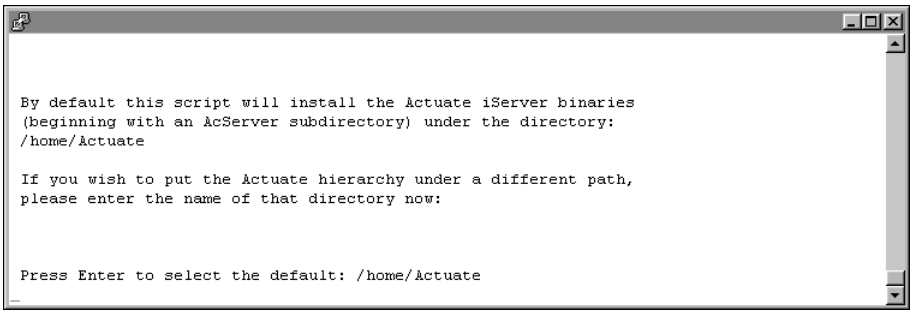

Figure 4-5 Specifying the installation directory

9 The installer detects the previous iServer version, and asks whether to perform an upgrade in place that overwrites the existing installation. Press Enter to accept the default option of upgrading the older release to iServer Release 11, as shown in Figure 4-6. Alternatively, type n for no, or q to quit.

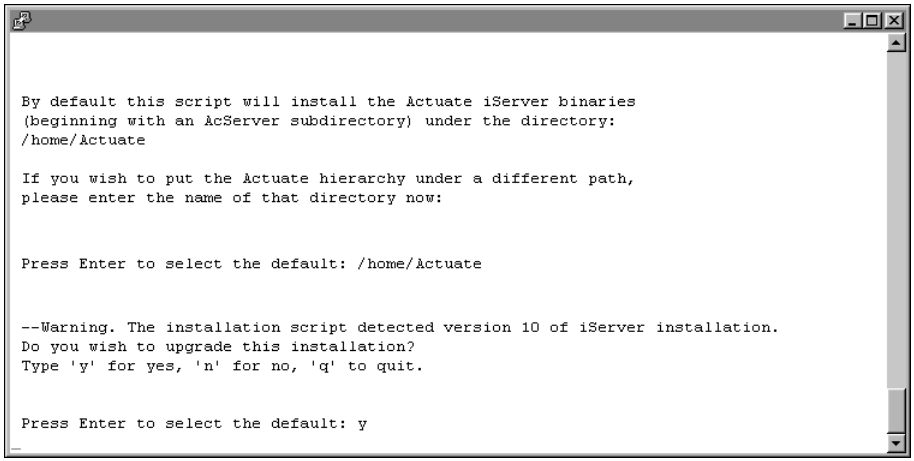

Figure 4-6 Choosing to upgrade iServer

10 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.

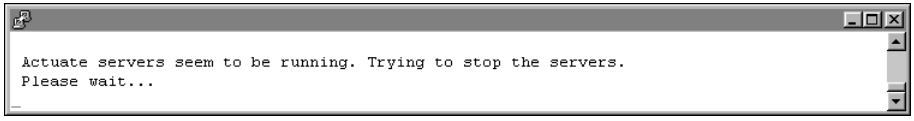

Figure 4-7 Shutting down the servers

11 Press Enter to accept the default installation directory, AC_SERVER_HOME/data, to install iServer data, as shown in Figure 4-8. Alternatively, choose a different directory for iServer data.

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

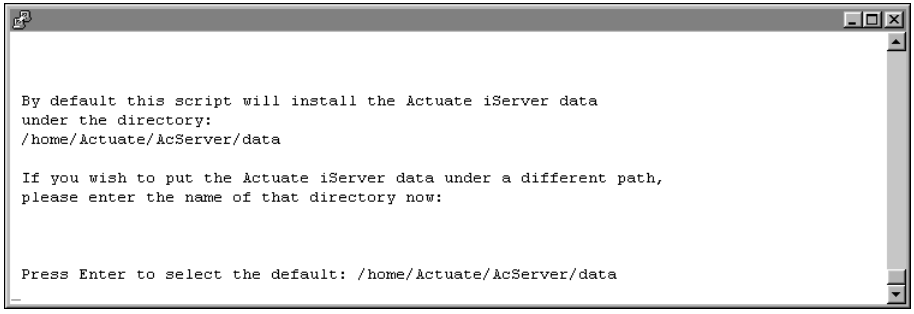

Figure 4-8 Specifying the data installation directory

12 Press Enter to accept the default option of creating the directory for AC_DATA_HOME, as shown in Figure 4-9. Alternatively, type n for no, or q to quit, and press Enter.

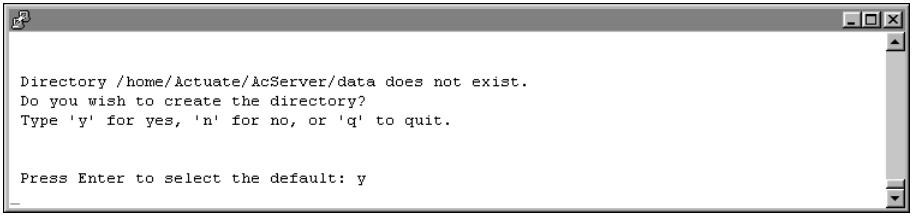

Figure 4-9 Creating the AC_DATA_HOME directory

13 The installer copies the prerequisite files to the destination folder, as shown in Figure 4-10.

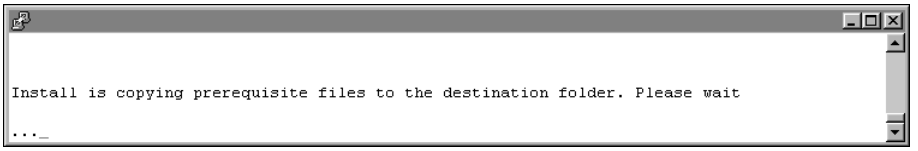

Figure 4-10 Copying prerequisite files

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.

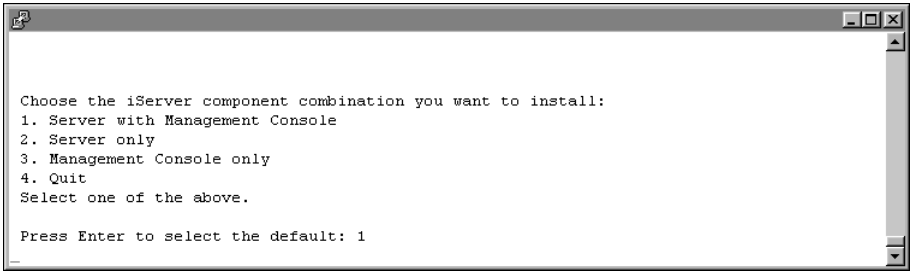

Figure 4-11 Choosing the iServer components to install

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

16 Type a name to use for the BIRT iServer System name, as shown in Figure 4-13. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.

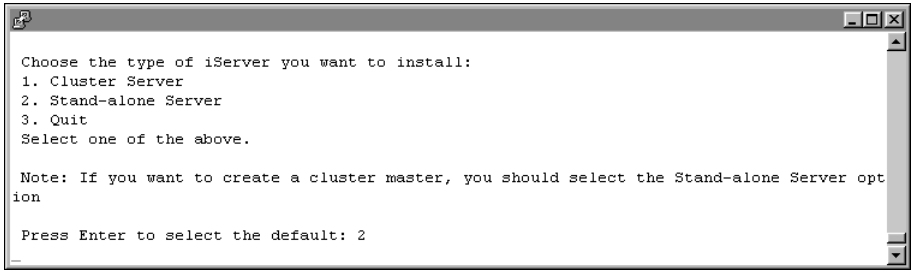

Figure 4-12 Choosing the iServer installation type

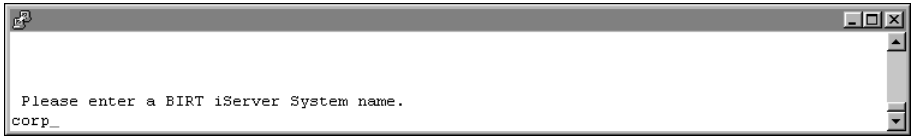

Figure 4-13 Specifying the BIRT iServer System name

17 Press Enter to choose the default embedded PostgreSQL database to store the Encyclopedia volume metadata, as shown in Figure 4-14.

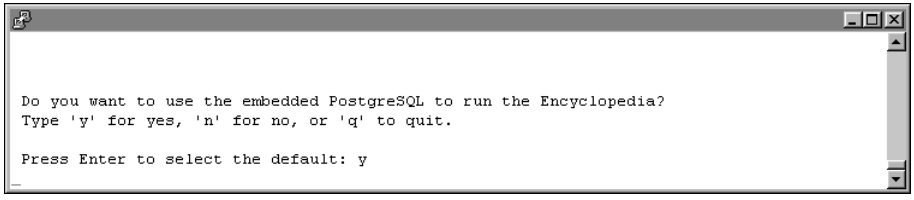

Figure 4-14 Choosing the embedded PostgreSQL

18 Press Enter to choose the default PostgreSQL superuser name, postgres, as shown in Figure 4-15. Alternatively type a different PostgreSQL superuser name.

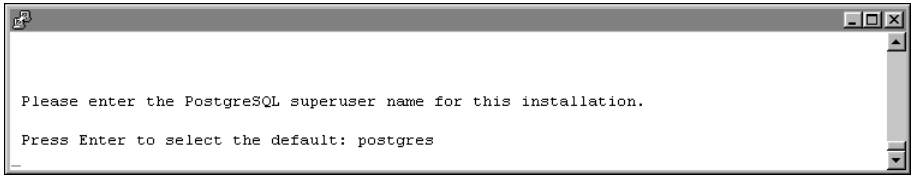

Figure 4-15 Choosing the PostgreSQL superuser name

19 Type a PostgreSQL superuser password that conforms to the password security policy requirements within your organization, then press Enter, as shown in Figure 4-16.

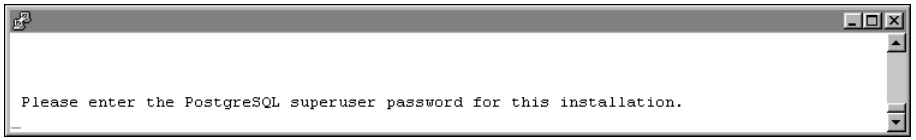

Figure 4-16 Typing the PostgreSQL superuser password

20 Re-enter the password for PostgreSQL superuser, as shown in Figure 4-17, and press Enter.

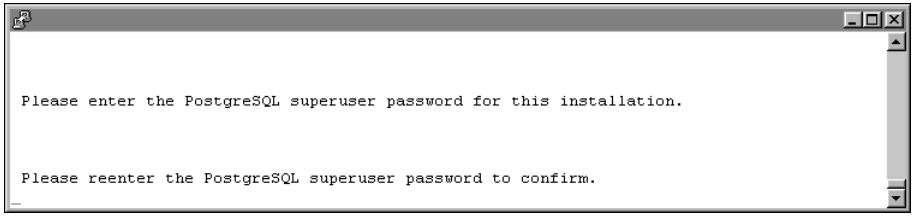

Figure 4-17 Re-entering the PostgreSQL superuser password

21 Press Enter to accept the default port number, 8432, at which the PostgreSQL database server listens, as shown in Figure 4-18. Alternatively, type a different port number.

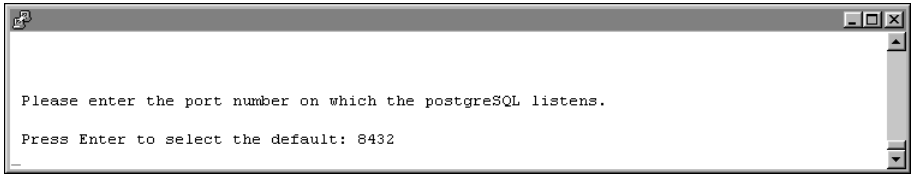

Figure 4-18 Entering the port number which PostgreSQL uses

22 Type y for yes to share the database schema for all encyclopedia volumes, as shown in Figure 4-19. Alternatively, accept the default value of n for no. If you choose no, you must manually set up a schema for each Encyclopedia volume using Configuration Console, then manually migrate the volume using the migration tools.

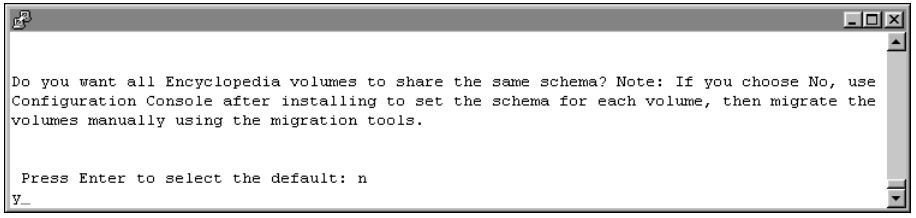

Figure 4-19 Choosing whether to share a schema for all volumes

23 Specify the iServer administrator password, as shown in Figure 4-20.

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

Figure 4-20 Specifying the iServer administrator password

24 Re-enter the password of the iServer administrator, as shown in Figure 4-21.

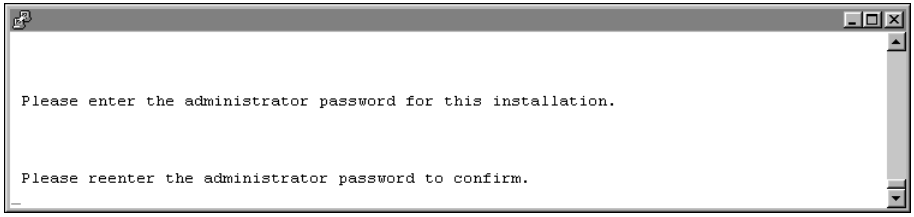

Figure 4-21 Re-entering the iServer administrator password

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

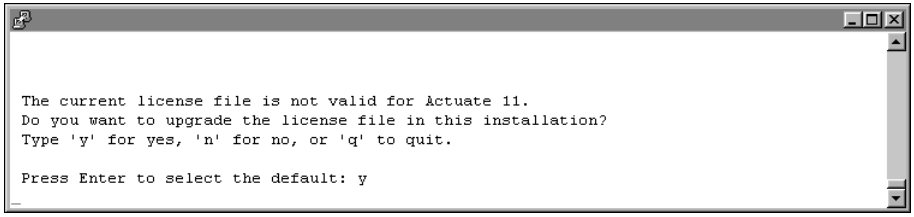

Figure 4-22 Accepting to upgrade the iServer license

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

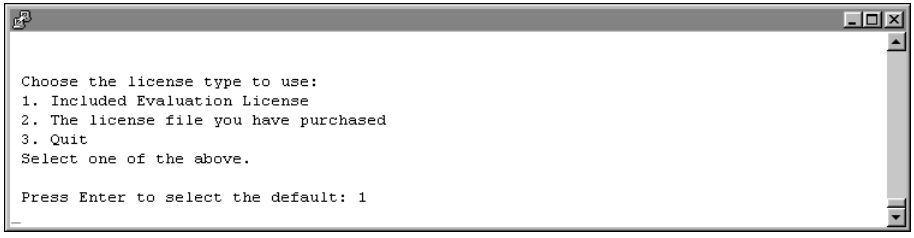

Figure 4-23 Specifying license type

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

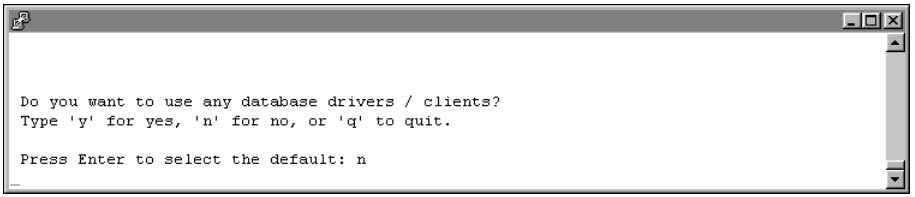

Figure 4-24 Specifying whether to use database drivers/clients

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

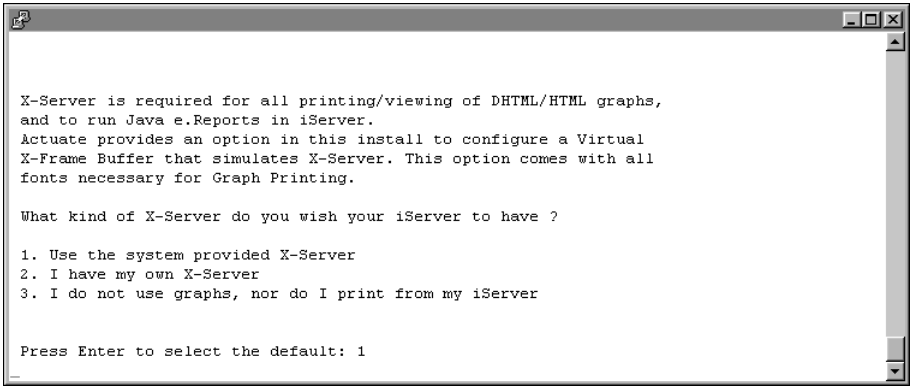

Figure 4-25 Specifying what kind of X-Server to use

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

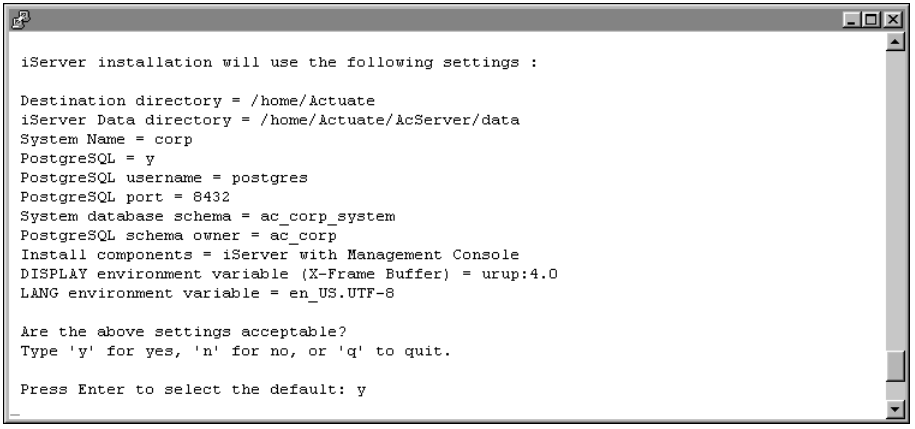

Figure 4-26 Reviewing your settings

30 If you chose yes when asked if you want all volumes to share the same schema, the following prompt appears. Press Enter to choose the default option of running the upgrade utility to migrate the encyclopedia automatically, as shown in Figure 4-27. Alternatively, choose n for no, to migrate the encyclopedia manually after the installation completes.

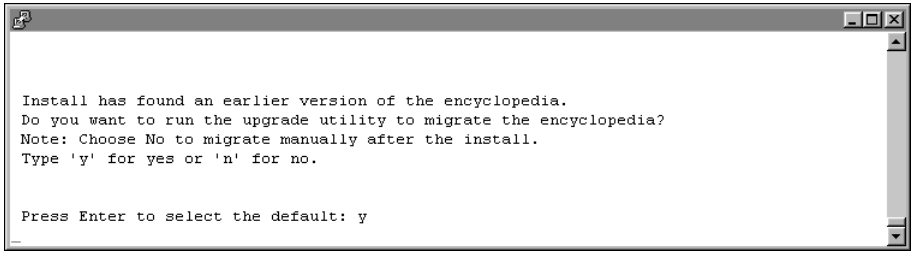

Figure 4-27 Choosing to migrate the encyclopedia automatically

31 The installation program installs iServer, displaying an indicator showing the progress of the installation, as shown in Figure 4-28.

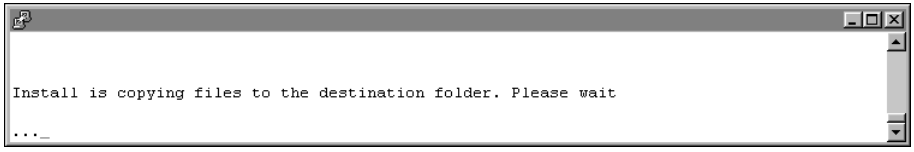

Figure 4-28 Copying iServer files to your destination folder

32 At the end of the installation, the program asks whether to start iServer. Accept the default, y for yes, to start the Process Management Daemon (PMD), as shown in Figure 4-29.

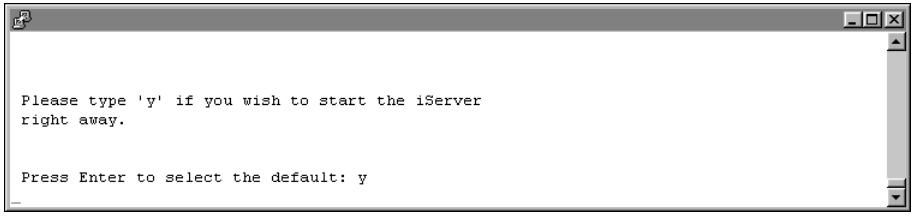

Figure 4-29 Specifying whether to start iServer

33 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 4-30.

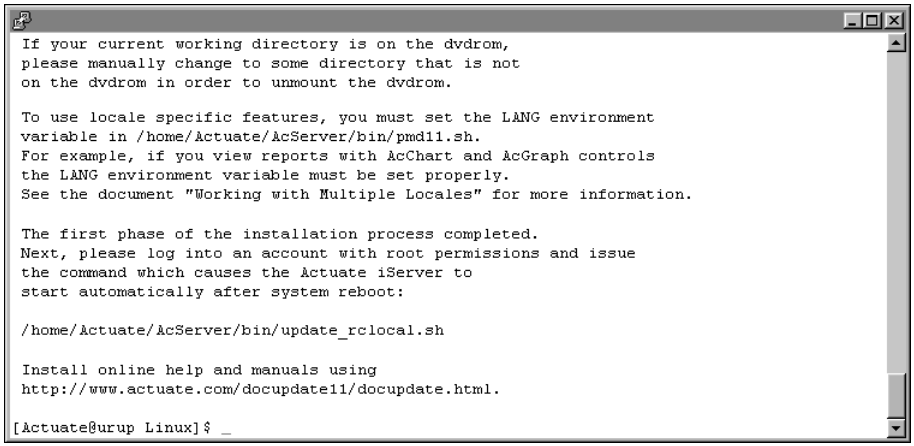

Figure 4-30 Specifying information about localization, logging in, and installing online help

34 If you performed an automatic in-place install, log in to Management Console. In Files and Folders, the data from the previous release appears. You performed an automatic in-place install if you chose yes to share the database schema for all encyclopedia volumes, and yes to run the upgrade utility to migrate the encyclopedia automatically.

About the /etc files saved by the upgrade process

Figure 4-31 shows the file system after upgrading to iServer Release 11. The date-stamped /etc directory contains the files saved during the upgrade from the older release to Release 11, which includes the following files:

- acserverconfig.xml in the /etc directory
- acpmdconfig.xml in the /etc directory

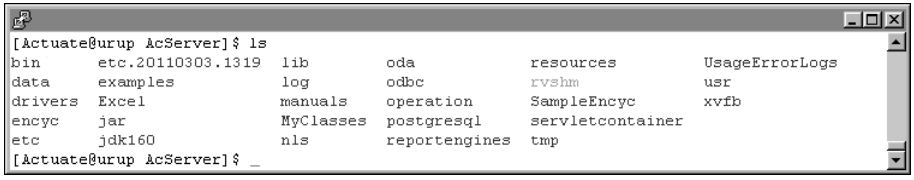

Figure 4-31 File system after upgrading to iServer Release 11

Performing a manual in-place upgrade

After installing BIRT iServer, if you choose to perform a manual upgrade of an Encyclopedia volume, you must perform additional operations to complete the installation. When upgrading an iServer system from an earlier major release in place, such as Release 10 Service Pack 1, the administrator uses the Squirrel Data

Exporter and Encyclopedia Data Store Administrator utilities to migrate the volume. When updating from an earlier Release 11 installation, such Release 11 Service Pack 1 to Service Pack 2, the administrator uses the Encyclopedia Data Store Upgrader utility to upgrade the volume.

The following sections describe how to perform manual upgrades to these in-place installations. These types of upgrades are typically performed by an OEM vendor, for example, who embeds iServer in products, or a customer with special requirements who cannot do an automatic upgrade.

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

How to perform a manual, in-place upgrade of an Encyclopedia volume from an earlier major release

1. Shut down iServer by performing the following tasks:
   1. In a web browser type:

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

      Log into Configuration Console as Administrator.
   2. Choose Advanced view.
   3. From the side menu, choose System. In System—Status, choose Stop to shut down iServer.

      Make sure the earlier iServer is offline before performing the next operation, as shown in Figure 4-32.

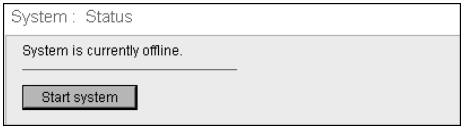

Figure 4-32 iServer is offline

2. Run the Squirrel Data Exporter by performing the following tasks:
   1. Create a properties file named SquirrelDataExporter.properties. You pass this file to the Squirrel Data Exporter utility when you execute it. Use the following example to create SquirrelDataExporter.properties:

```
AC_SERVER_HOME = /home/Actuate/AcServer
AC_DATA_HOME = /home/Actuate/AcServer/data
SQUIRREL_DATA_HOME = /home/Actuate/AcServer/encyc
SQUIRREL_EXPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
SQUIRREL_EXPORT_FORMAT = PostgreSQL
NEW_SCHEMA_NAME =  ac_corp
NEW_VOLUME_NAME =  corp
```

See Table 4-3 for descriptions of the required Squirrel Data Exporter properties later in this chapter. See Table 4-4 for descriptions of the optional Squirrel Data Exporter properties. See "Specifying Squirrel Data Exporter properties" and "Specifying Encyclopedia Data Store Administrator properties" for additional notes on property files. 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 Create a script to run the following commands. Adjust the location of AC_SERVER_HOME if the location is different in your environment:

```
#!/bin/sh
export JAVA_HOME=/home/Actuate/AcServer/jdk160
export AC_SERVER_HOME=/home/Actuate/AcServer
export LD_LIBRARY_PATH=$AC_SERVER_HOME/lib
export PATH=$JAVA_HOME/bin:$LD_LIBRARY_PATH:$PATH
export CLASSPATH=
   $AC_SERVER_HOME/jar/AcCommon.jar:
   $AC_SERVER_HOME/jar/JDBCDrivers.jar:
   $AC_SERVER_HOME/jar/com.actuate.common-server.jar:
   $AC_SERVER_HOME/jar/ServerAdminTools.jar
export SQUIRREL_DATA_EXPORTER=
   com.actuate.iserver.encyclopedia.datastore.admin
   .SquirrelDataExporter
java $SQUIRREL_DATA_EXPORTER SquirrelDataExporter.properties
```

3 Run the Squirrel Data Exporter script from the command prompt.

3 In the Advanced view of Configuration Console, on System—Status, choose Start system to restart iServer.

4 In this step, you create a new schema that you assign to the existing volume. In the Advanced view of Configuration Console, from the side menu, choose Volumes. On Volumes, point to the icon next to Default ActuatePostgreSQL MetadataDatabase and choose Add volume schema, as shown in Figure 4-33.

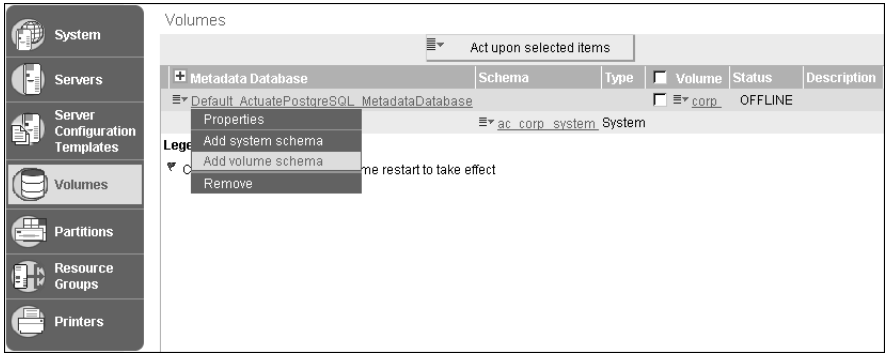

Figure 4-33 Choosing to add a new schema

On Volumes—New Schema, perform the following tasks:

1. In Schema name, type a name for the schema. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.
2. In Database schema name, type a name for the Database schema. The name must be less than 30 characters. Observe the same naming restrictions for this schema as the volume schema name.
3. In Database schema password, type a new password.
4. In Database schema password confirm, re-type the new password.
5. In Database superuser, type the database superuser name. For the PostgreSQL RDBMS that installs with iServer by default, the PostgreSQL superuser name is postgres.
6. In Database superuser password, type the superuser password. This password is the same password that the installer provides when installing iServer, as shown in Figure 4-34.

Choose OK.

Volumes—New Schema appears as shown in Figure 4-34.

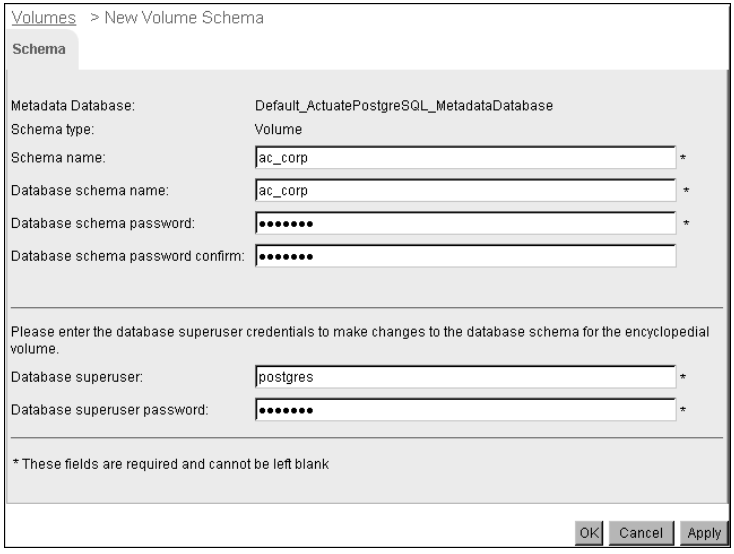

Figure 4-34 Creating a new schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-35.

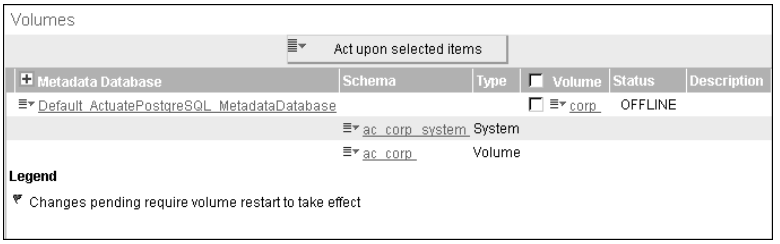

Figure 4-35 Viewing the new schema

5 Point to the icon next to the volume name and choose Properties, as shown in Figure 4-36.

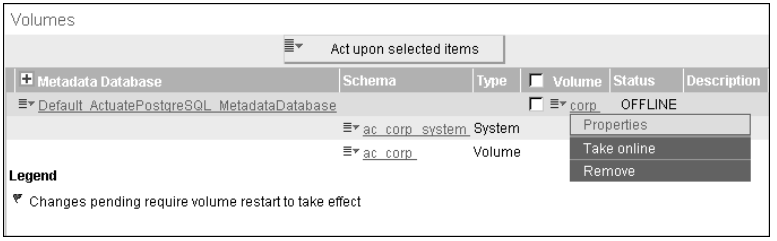

Figure 4-36 Choosing volume properties

6 On Properties—General, in Metadata database and schema, select the Database schema name from the list of available schemas, as shown in Figure 4-37. Choose OK.

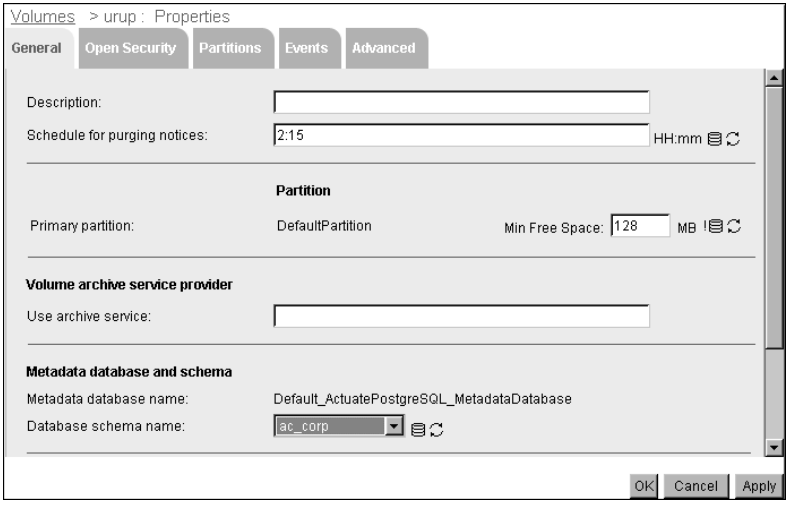

Figure 4-37 Choosing the database schema name

On Volumes, the Database schema name you chose for the volume appears on the same line as the volume, as shown in Figure 4-38.

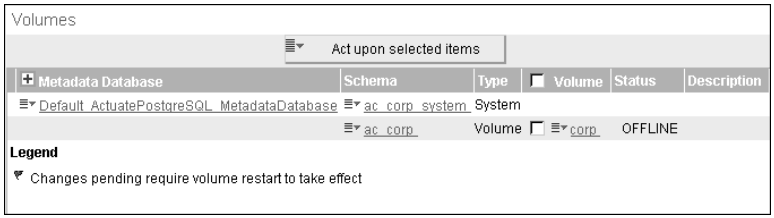

Figure 4-38 Viewing Volumes after binding the schema to the volume

7 Run the Encyclopedia Data Store Administrator utility by performing the following operations:

1 To create the Encyclopedia volume schema, perform the following tasks:

1 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
AC_DATA_HOME = /home/Actuate/AcServer/data
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
SCHEMA_NAME =  ac_corp
SCHEMA_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
```

See Table 4-5 for descriptions of the required Encyclopedia Data Store Administrator properties later in this chapter. See Table 4-6 for descriptions of the optional Encyclopedia Data Store Administrator properties. See "Specifying Squirrel Data Exporter properties" and "Specifying Encyclopedia Data Store Administrator properties" for additional notes on property files.

2 Create a script to run the following commands. Adjust the location of AC_SERVER_HOME if the location is different in your environment:

```
#!/bin/sh
export JAVA_HOME=
   /home/Actuate/AcServer/jdk160
export AC_SERVER_HOME=/home/Actuate/AcServer
export LD_LIBRARY_PATH=$AC_SERVER_HOME/lib
export PATH=$JAVA_HOME/bin:$LD_LIBRARY_PATH:$PATH
export CLASSPATH=
   $AC_SERVER_HOME/jar/AcCommon.jar:
   $AC_SERVER_HOME/jar/JDBCDrivers.jar:
```

```
   $AC_SERVER_HOME/jar/com.actuate.common-server.jar:
   $AC_SERVER_HOME/jar/ServerAdminTools.jar
export ENCYCLOPEDIA_DATASTORE_ADMINISTRATOR=
   com.actuate.iserver.encyclopedia.datastore.admin
   .EncyclopediaDataStoreAdministrator
java $ENCYCLOPEDIA_DATASTORE_ADMINISTRATOR
   VolumeImport.properties
```

3 Run the Encyclopedia Data Store Administrator script from the command prompt.

8 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-39.

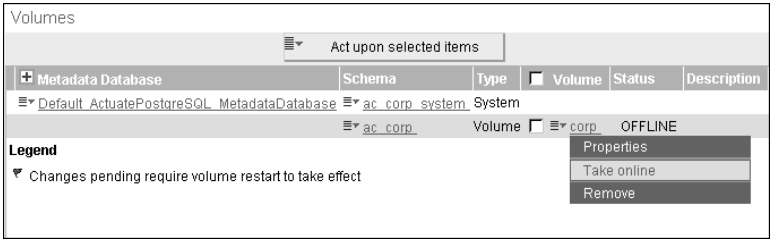

Figure 4-39 Viewing the new volume

The new volume comes online, as shown in Figure 4-40.

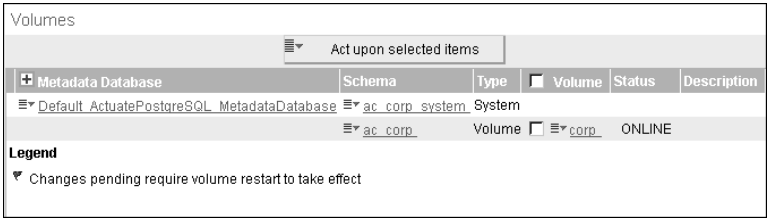

Figure 4-40 Viewing the new volume after it comes online

9 In web browser type:

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

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

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

When upgrading BIRT iServer from an earlier Release 11 version, you install the new iServer version in the same directory as the earlier version. After completing the install procedure, you run a utility that updates the volume schema, enabling the new iServer version to work with your existing Encyclopedia volume or volumes. Then, you bring the volume or volumes online.

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

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.

1. Although the install program saves these files during an upgrade, Actuate recommends that you make a backup copy of the following files before installing:
   - encyc directories from all nodes
   - acserverconfig.xml in the /etc directory
   - acpmdconfig.xml in the /etc directory
   - RSSE code and associated files if you use the Open Security option
2. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required files.
3. 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.
4. The license agreement appears, as shown in Figure 4-41.

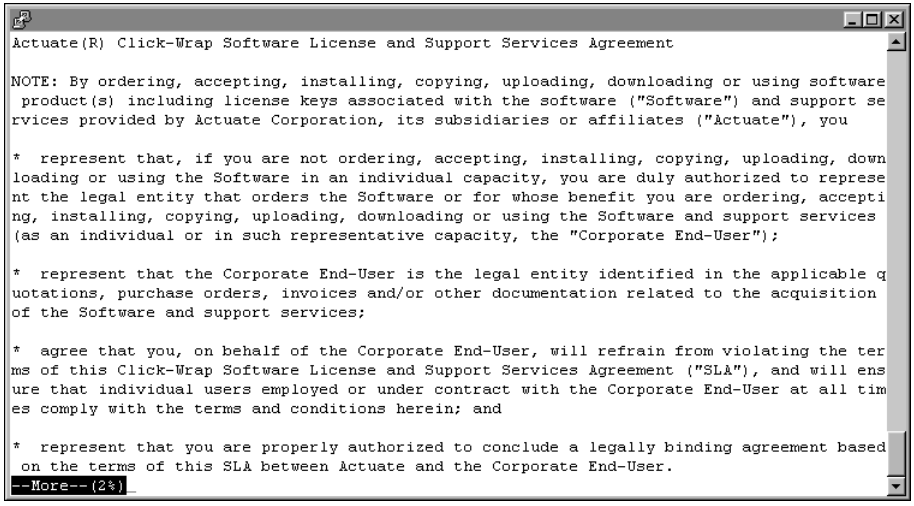

Figure 4-41 Reviewing the license agreement

5. 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 4-42

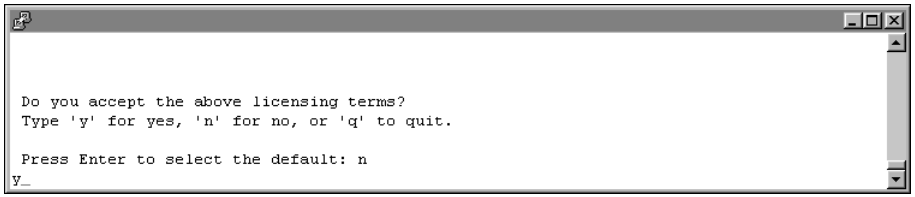

Figure 4-42 Accepting the licensing terms

6 The introduction to the installation appears, as shown in Figure 4-43.

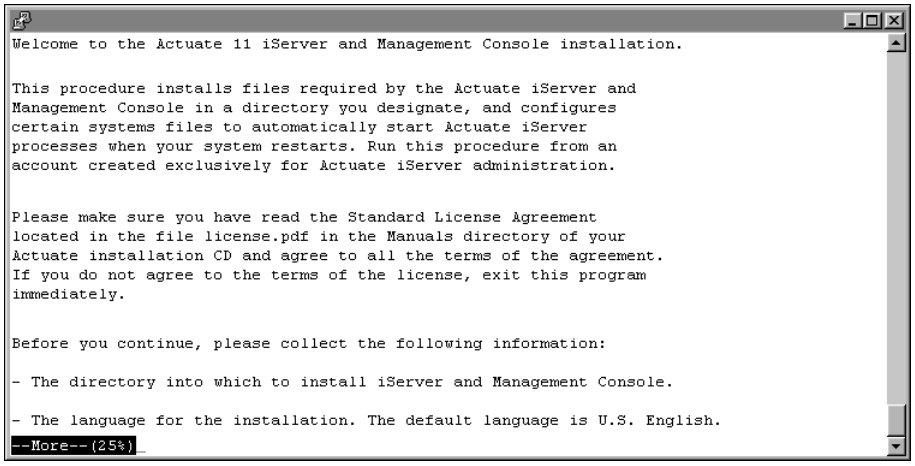

Figure 4-43 Reviewing the introductory information

7 Press Enter after reviewing the introductory information, as shown in Figure 4-44.

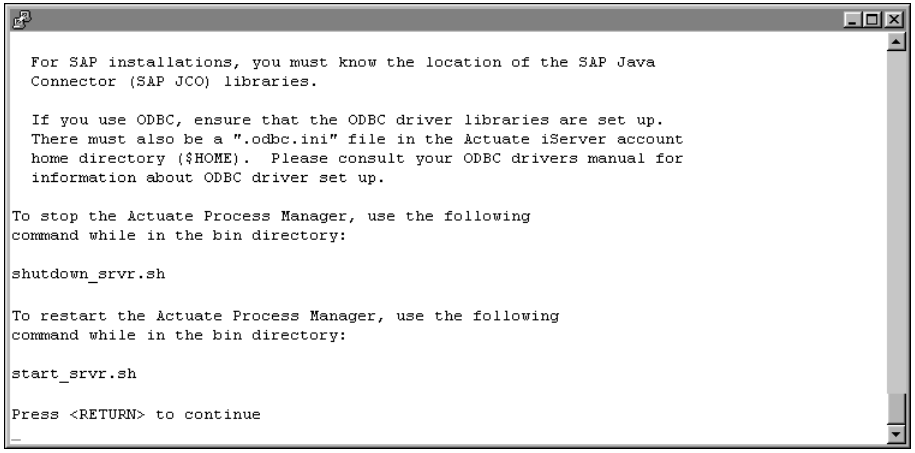

Figure 4-44 Finishing the review of introductory information

8 Press Enter to select the default directory as the location of the Actuate 11 installation. Alternatively, type the appropriate path and press Enter, as shown in Figure 4-5. The name of the installation directory does not change when you upgrade. For example, the installation directory in Figure 4-45 is still /home/Actuate/AcServer after upgrading from Actuate 11SP1 to Actuate 11 SP2.

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.

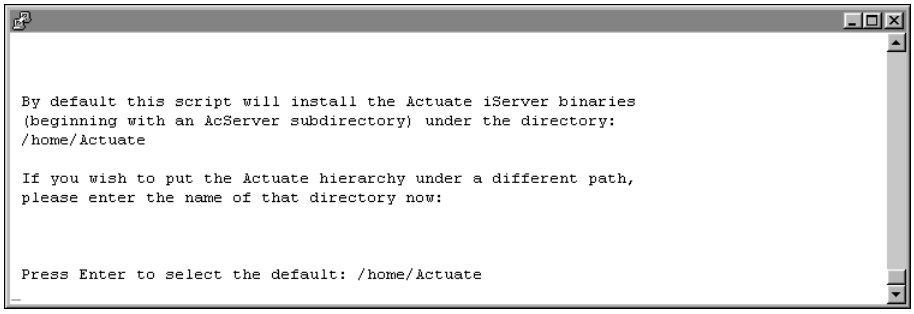

Figure 4-45 Specifying the installation directory

9 The installer detects the previous iServer version, and asks whether to perform an upgrade in place that overwrites the existing installation. Press Enter to accept the default option of overwriting the contents of the installation directory, as shown in Figure 4-46. Alternatively, type n for no, or q to quit.

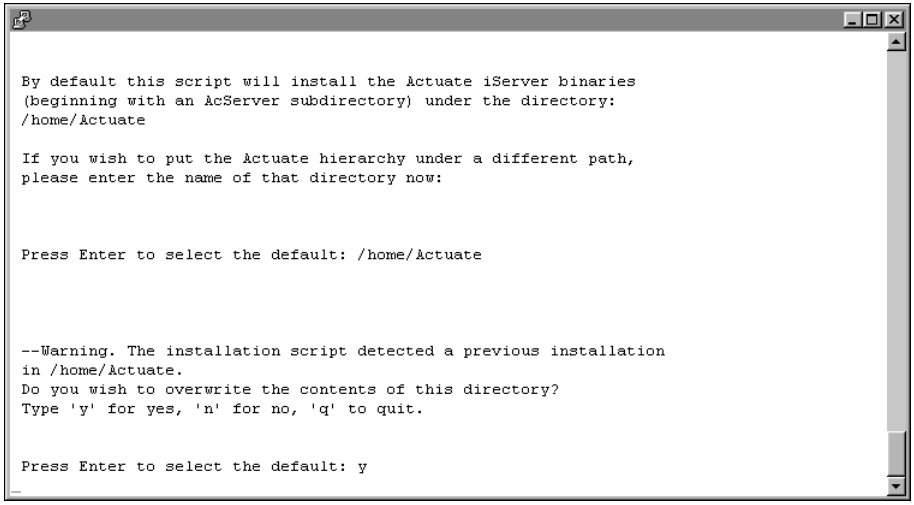

Figure 4-46 Choosing to upgrade iServer

10 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 4-47. This process might take a few minutes.

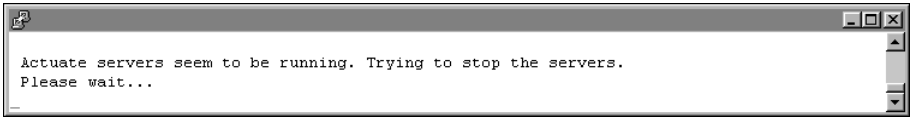

Figure 4-47 Shutting down the servers

11 Press Enter to accept the default iServer component combination, as shown in Figure 4-48. Alternatively, choose a different iServer component combination and press Enter.

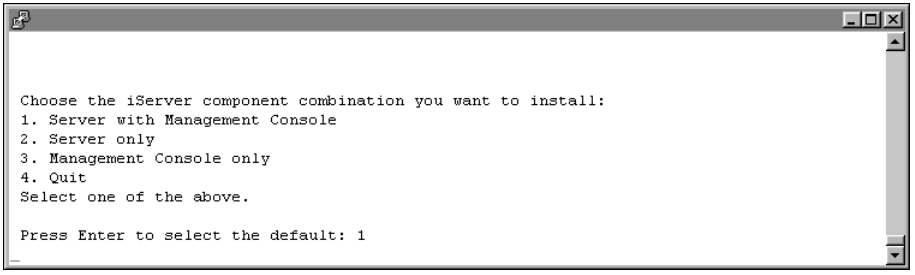

Figure 4-48 Choosing the iServer components to install

12 Press enter to accept the default stand-alone Server installation, as shown in Figure 4-49.

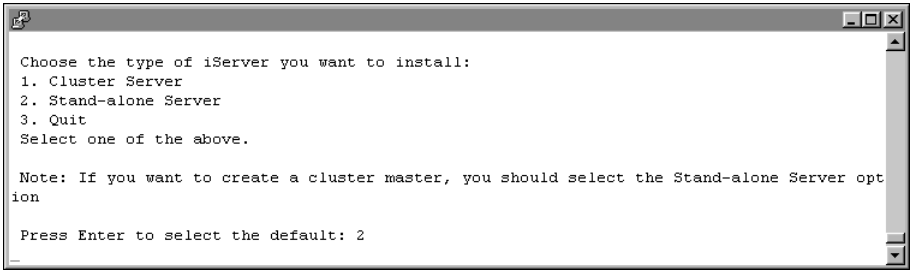

Figure 4-49 Choosing the iServer installation type

13 Press Enter to select the default locale, which is English, as shown in Figure 4-50. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.

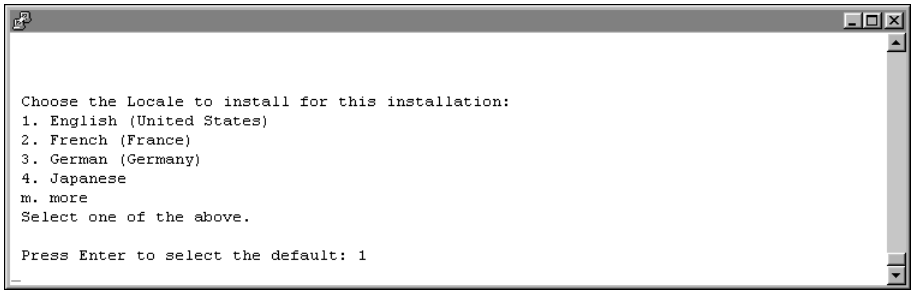

Figure 4-50 Specifying a locale

14 To evaluate the product using the included evaluation software license press Enter, as shown in Figure 4-51. Alternatively, type 2, then type the path to a purchased license file.

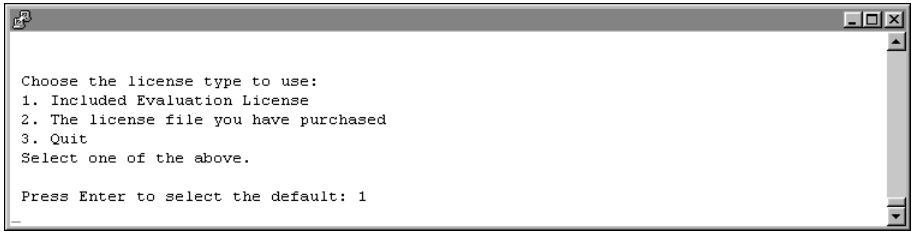

Figure 4-51 Specifying license type

15 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 4-52. Alternatively, type a different IP address or hostname.

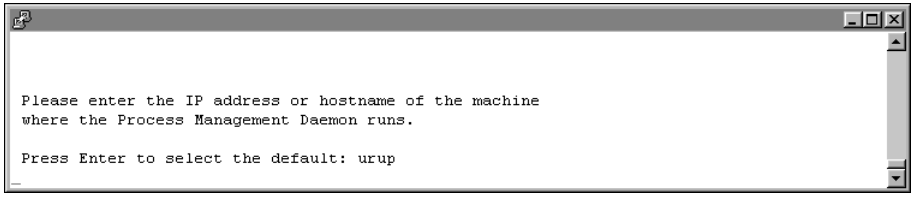

Figure 4-52 Specifying the hostname that Management Console uses to contact the PMD

16 Press Enter to accept the default port number, 8100, where Process Management Daemon (PMD) listens for requests, as shown in Figure 4-53. Alternatively, type a different port number and press Enter.

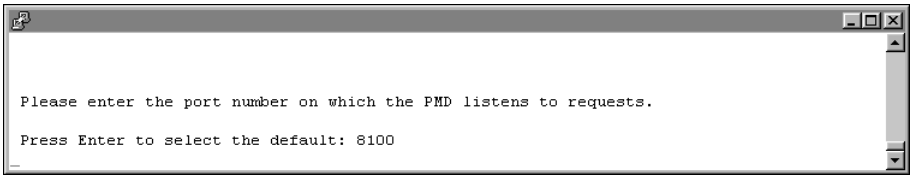

Figure 4-53 Specifying the port number on which the PMD listens

17 Press Enter to accept the default hostname, the name of the machine on which iServer runs, as shown in Figure 4-54. Alternatively, type a different hostname or IP address, then press Enter.

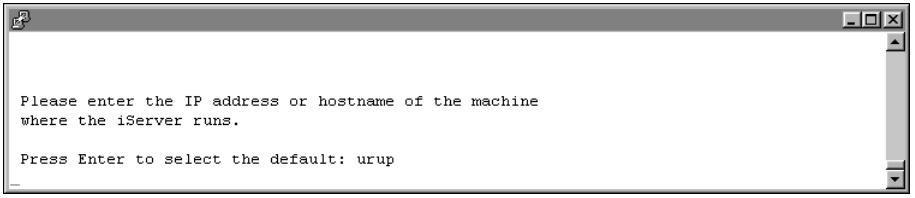

Figure 4-54 Specifying the machine on which the iServer runs

18 Press Enter to accept the default port number where iServer listens to requests, as shown in Figure 4-55. Alternatively, type a different port number and press Enter.

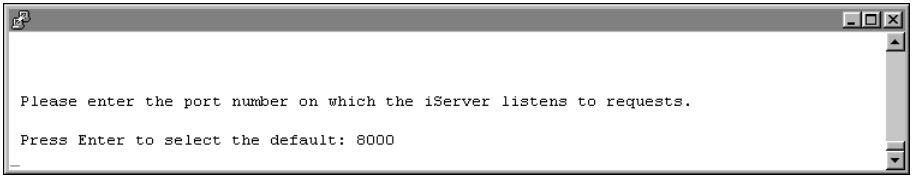

Figure 4-55 Specifying the port number on which the iServer listens

19 Specify the iServer administrator password, as shown in Figure 4-56.

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

Figure 4-56 Specifying the iServer administrator password

20 Re-enter the password of the iServer administrator, as shown in Figure 4-57.

You use this password to log in to Configuration Console.

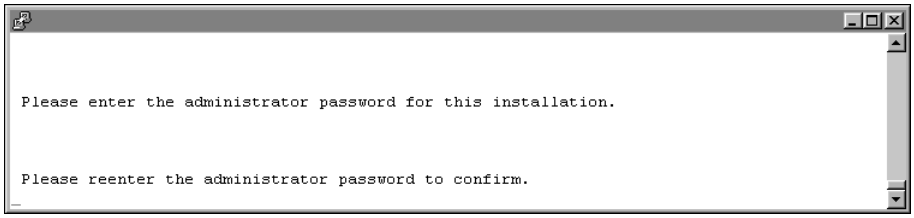

Figure 4-57 Re-entering the iServer administrator password

21 Press Enter to accept the default option to use a volume name for the Encyclopedia, as shown in Figure 4-58. Alternatively, type n for no to not use a volume name for the Encyclopedia, or q to quit the installation.

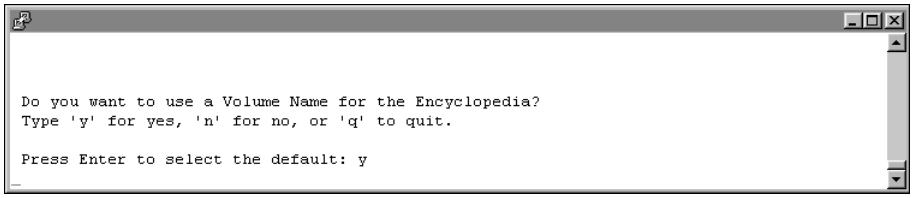

Figure 4-58 Choosing not to use a volume name

22 Press Enter to accept the name of the default Encyclopedia volume to use with Management Console, as shown in Figure 4-59. Alternatively, type a different name for the Encyclopedia volume.

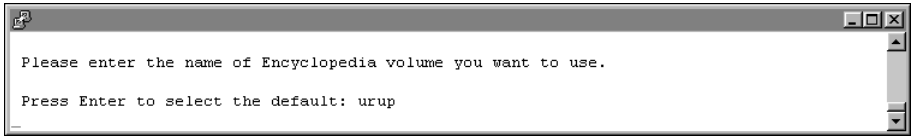

Figure 4-59 Specifying the name of the default Encyclopedia volume

23 Press Enter to accept the default option to start iServer automatically, as shown in Figure 4-60. Alternatively, type n for no.

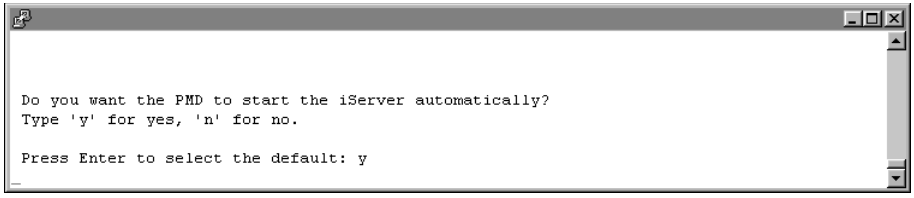

Figure 4-60 Specifying whether to start iServer automatically

24 Press Enter to accept the default option of not using any database drivers/clients, as shown in Figure 4-61. Alternatively, type y for yes, specify the database drivers/clients you want to use, and press Enter.

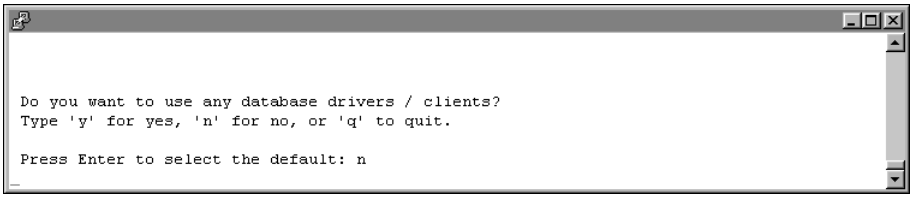

Figure 4-61 Specifying whether to use database drivers/clients

25 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-62.

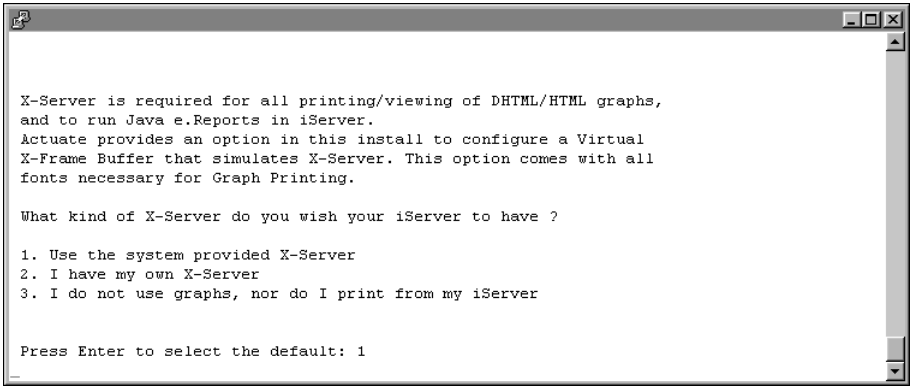

Figure 4-62 Specifying what kind of X-Server to use, if any

26 Review the settings, as shown in Figure 4-63, 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.

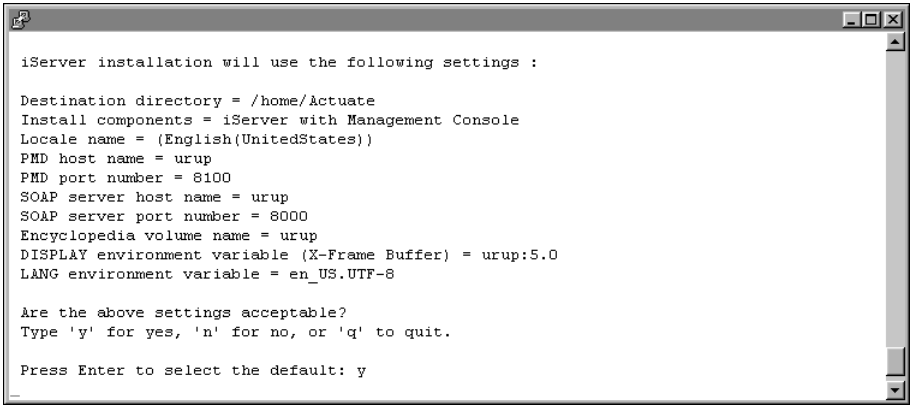

Figure 4-63 Reviewing the installation settings

27 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-64.

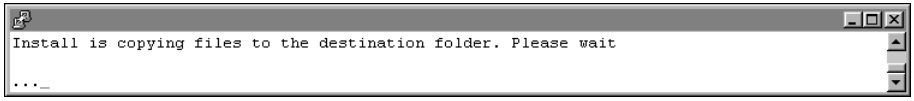

Figure 4-64 Viewing iServer installation progress

28 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 4-65.

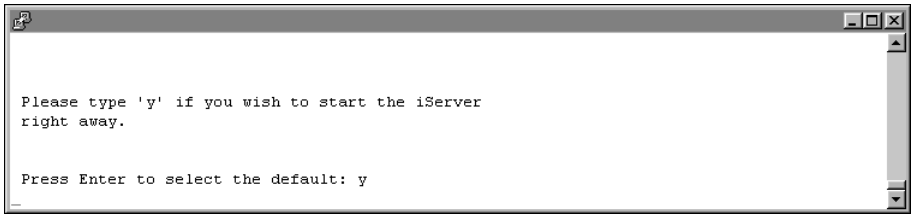

Figure 4-65 Specifying whether to start iServer

29 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 4-66.

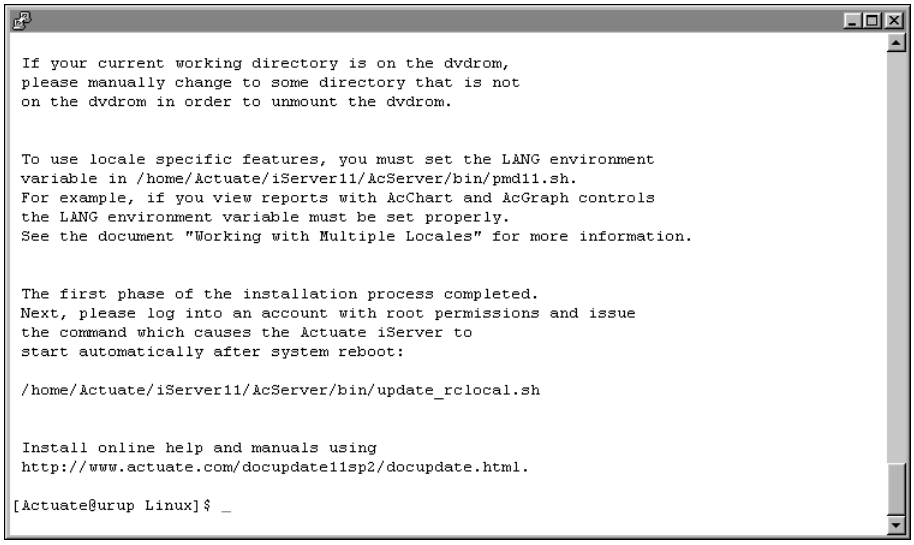

Figure 4-66 Viewing information about localization, logging in, and installing online help

How to perform an in-place upgrade of an Encyclopedia volume schema from an earlier Release 11 installation

1 In a web browser, type:

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

2 Log into 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 schema and choose Properties, as shown in Figure 4-67.

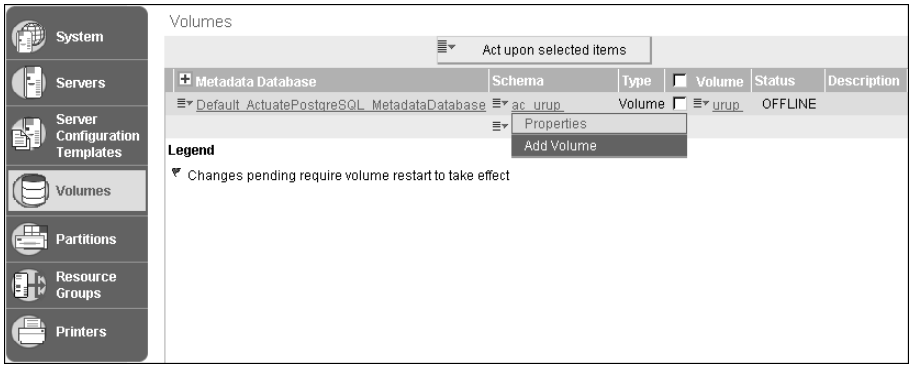

Figure 4-67 Choosing to view volume schema properties

3 On Properties, make note of the value that appears for Schema name. You specify the schema name in a later step. As shown in Figure 4-68, Schema name is ac_urup.

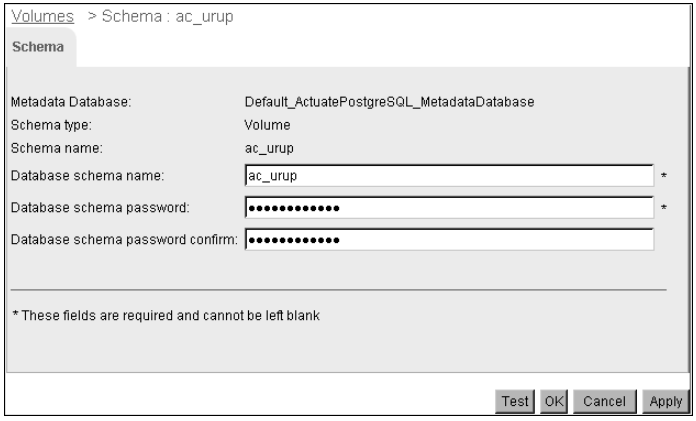

Figure 4-68 Making note of Schema name

From the side menu, choose System.

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

1 Navigate to AC_SERVER_HOME/bin. The default location of AC_SERVER_HOME is /home/Actuate/AcServer. Open the upgrade_encyclopedia_data_store.sh file in an editor. You can use this

script to execute the Encyclopedia Data Store Upgrader utility. The upgrade_encyclopedia_data_store.sh file sets the following properties and executes the utility after accepting the schema name as argument:

```
#!/bin/sh

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

fi

# Values stamped by the installer
AC_SERVER_HOME=/home/Actuate/AcServer
AC_DATA_HOME=/home/Actuate/AcServer/data
export JAVA_HOME=/home/Actuate/AcServer/jdk160

# 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
echo "PROMPT_FOR_PASSWORDS = true"  >> $PROPERTY_FILE

# Upgrade data store
unset JAVA_TOOL_OPTIONS
unset _JAVA_OPTIONS
export PATH=$JAVA_HOME/bin:$PATH
export CLASSPATH=$AC_SERVER_HOME/jar/AcCommon.jar\
:$AC_SERVER_HOME/jar/JDBCDrivers.jar\
:$AC_SERVER_HOME/jar/com.actuate.common-server.jar\
:$AC_SERVER_HOME/jar/ServerAdminTools.jar
ENCYCLOPEDIA_DATASTORE_UPGRADER=\
com.actuate.iserver.encyclopedia.datastore.admin.Encyclopedia
   DataStoreUpgrader

java $ENCYCLOPEDIA_DATASTORE_UPGRADER "$PROPERTY_FILE"
```

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

See Table 4-1 for descriptions of the required Encyclopedia Data Store Upgrader properties later in this chapter. See Table 4-2 for descriptions of the optional Encyclopedia Data Store Upgrader properties.

2 At the command line, enter the following command:

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

where <schema name> is the schema name you made note of in step 3.

3 Type the password for the schema, as shown in Figure 4-69.

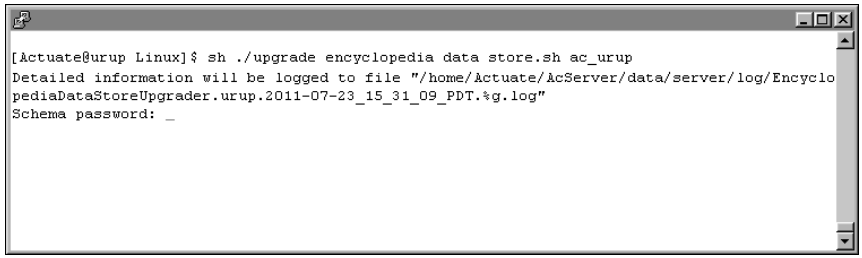

Figure 4-69 Typing the schema password

The utility runs, as shown in Figure 4-70.

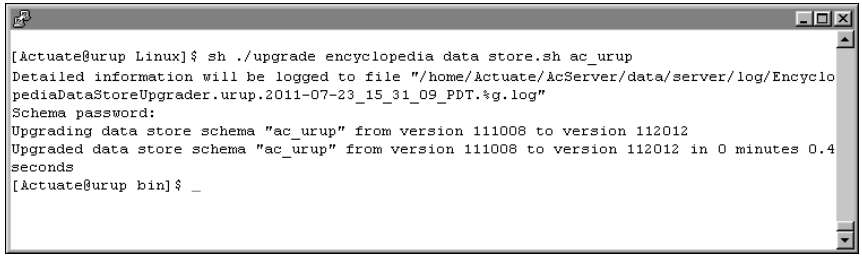

Figure 4-70 Running the utility

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

Point to the icon next to the volume name and choose Take online, as shown in Figure 4-71.

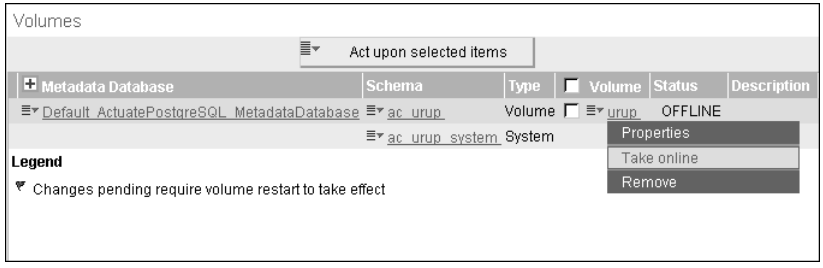

Figure 4-71 Taking the volume online

The volume comes online, as shown in Figure 4-72.

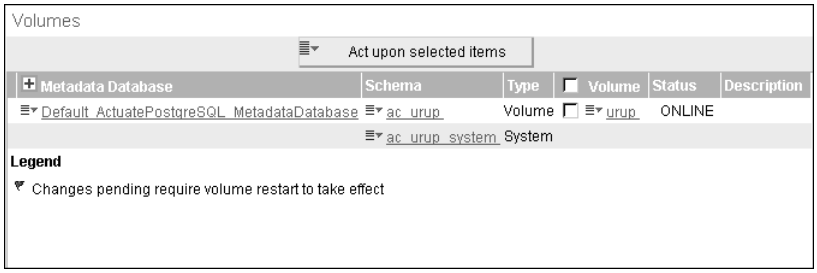

Figure 4-72 Viewing the status of the volume after it comes online

6 In a web browser, type:

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

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

Specifying Encyclopedia Data Store Upgrader properties

Table 4-1 describes the required Encyclopedia Data Store Upgrader properties to specify in the upgrade_encyclopedia_data_store.bat or other properties file for a manual volume upgrade operation.

Table 4-1 Required Encyclopedia Data Store Upgrader properties

Parameter	Description
AC_DATA_HOME	Points to the location of the iServer data, which you specify during the BIRT iServer Release 11 installation, as shown in Figure 4-8.
AC_SERVER_HOME	Points to the location of the iServer binaries, which you specify during the BIRT iServer Release 11 installation, as shown in Figure 4-8.
APPLICATION_USER	User ID used to connect to the database for normal operations.
DATABASE_TYPE	Type of relational database system that contains the data store. Actuate Release 11 currently supports PostgreSQL and Oracle.
LOG_FOLDER	Absolute path to the log folder.
SCHEMA_FILE_NAME	Base name of the file without the file extension that contains the meta-schema definition.

Table 4-1 Required Encyclopedia Data Store Upgrader properties

Parameter	Description
SCHEMA_NAME	Name of the target schema which the Encyclopedia Data Store Upgrader updates. 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 4-2 describes the optional Encyclopedia Data Store Upgrader properties to specify in the upgrade_encyclopedia_data_store.bat or other properties file.

Table 4-2 Optional Encyclopedia Data Store Upgrader properties

Parameter	Description	Default Value	Supported Databases
CONSOLE_LOG _LEVEL	Minimum logging level for messages sent to the console (stderr). Supported values are INFO, CONFIG, FINE, FINER, and FINEST.	INFO	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
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

(continues)

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

Parameter	Description	Default Value	Supported Databases
GENERATE_SCRIPTS	Set to true to generate scripts to perform operations instead of performing the operations manually.	False	All
FILE_LOG_LEVEL	The minimum logging level for messages sent to the log file. This parameter only works for AcLogger, and is a standard JUL Level name. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts for performing the 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
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

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

Parameter	Description	Default Value	Supported Databases
PROMPT_FOR _PASSWORDS	Indicates whether to prompt the user to type in a password instead of specifying it using a property value. Required if SCHEMA_PASSWORD is not present. Password prompting works only when running the tool from the command line.	False	All
SCHEMA_PASSWORD	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema.	{SCHEMA _NAME}	All

Performing a manual side-by-side upgrade

In a manual side-by-side upgrade, the administrator installs a new Release 11 iServer in a path separate from the earlier release. The installation program performs the following operations using a PostgreSQL relational database management system (RDBMS) as the example:

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

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

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

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

You can leave the iServer service for the earlier iServer release running when you perform the side-by-side upgrade procedure. The Release 11 install program detects the presence of the earlier iServer version and does not try to use any port number the earlier iServer release uses.

1. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required files.
2. To install the server files, execute the isinstall script:

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

   The script displays a number of prompts. Respond to the prompts as described in the following procedure.
3. The license agreement appears, as shown in Figure 4-73.

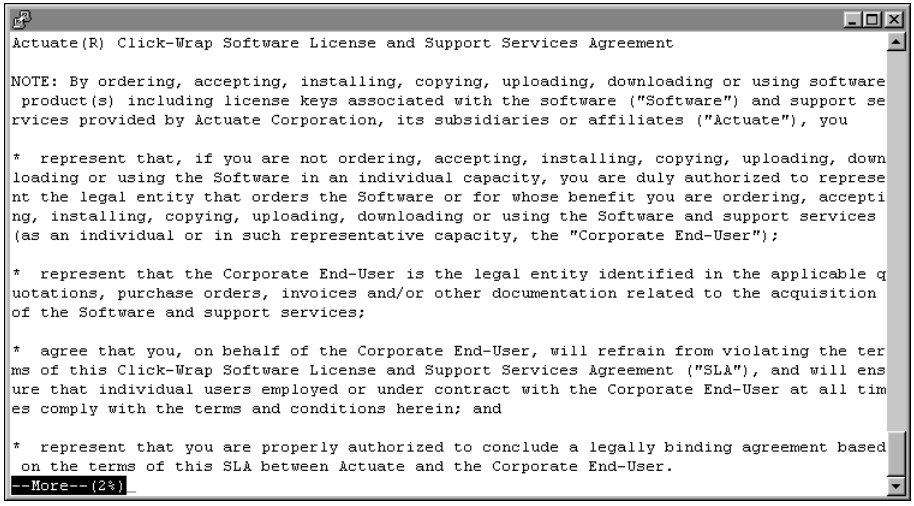

Figure 4-73 Reviewing the license agreement

4. Read the license agreement, then press Enter to continue the installation. At the prompt, type y for yes if you accept the licensing terms, as shown in Figure 4-74.

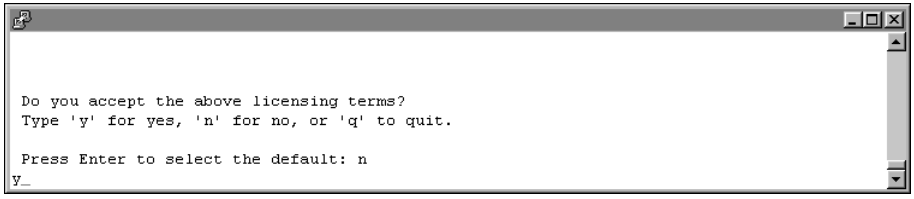

Figure 4-74 Accepting the licensing terms

5 The introduction to the installation appears, as shown in Figure 4-75.

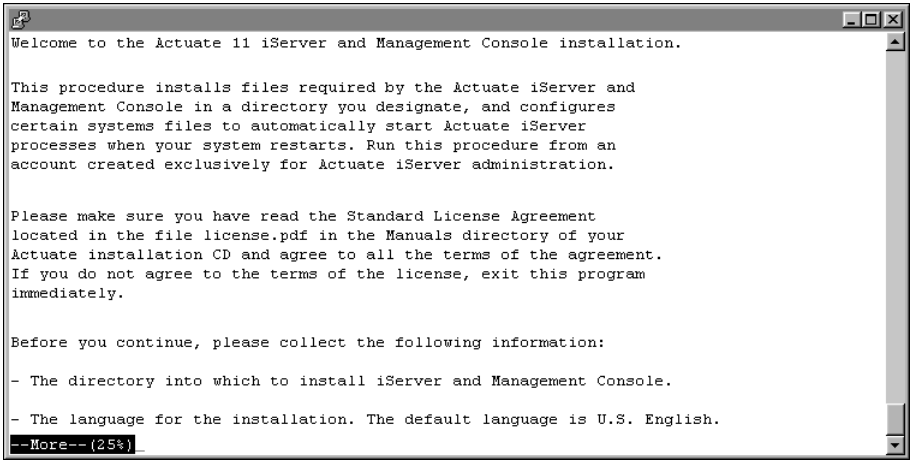

Figure 4-75 Reviewing the introductory information

6 Press Enter after reviewing the introductory information, as shown in Figure 4-76.

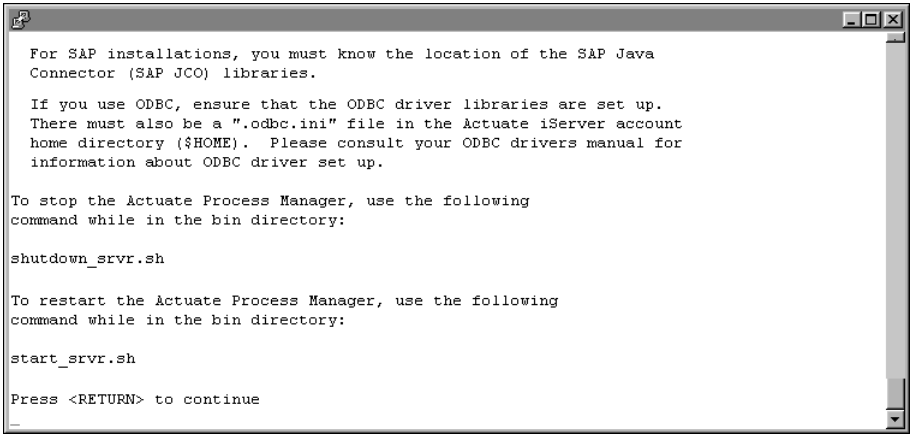

Figure 4-76 Finishing the review of introductory information

7 Type the name of a new directory in /home/Actuate, where the install program installs BIRT iServer Release 11. Specify both the path and the directory, and press enter, as shown in Figure 4-77.

iServer uses this location to resolve paths to all the binaries that it launches. The default path for this location is $HOME/AcServer, which is referred to in the iServer documentation by the environment variable AC_SERVER_HOME. For example, using the new directory name shown in Figure 4-77, AC_SERVER_HOME refers to /home/Actuate/iServer11/AcServer.

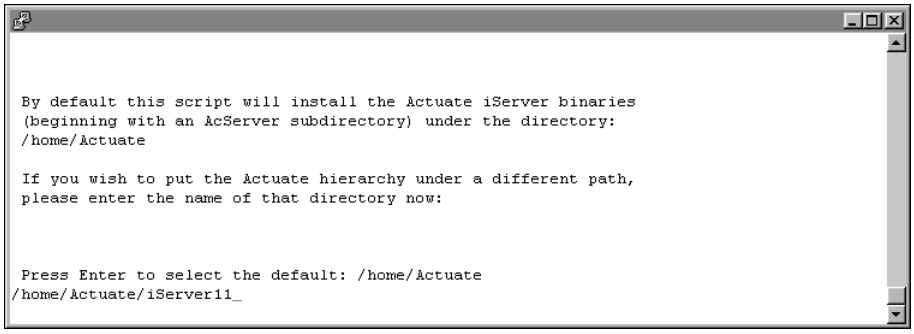

Figure 4-77 Specifying the installation directory

8 Press Enter to accept the default option of creating the folder you specified in the previous step, as shown in Figure 4-78. Alternatively, type n for no, or q to quit, and press Enter.

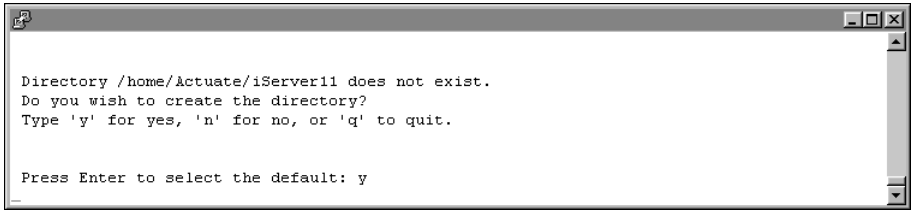

Figure 4-78 Creating the new installation directory

9 Press Enter to accept the default installation directory, AC_SERVER_HOME/data, to install iServer data, as shown in Figure 4-79. Alternatively, choose a different directory for iServer data.

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

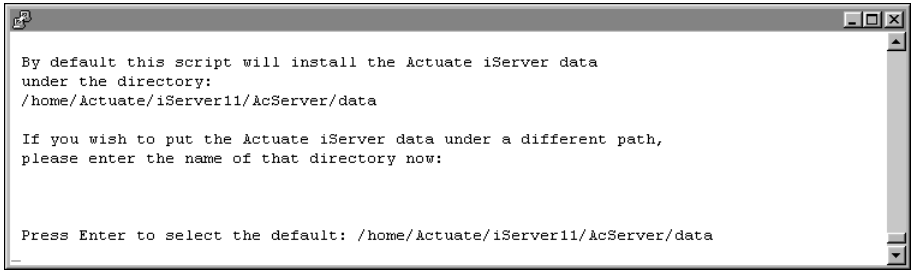

Figure 4-79 Specifying the data installation directory

10 The installer asks whether you want to create the directory for AC_DATA_HOME. Press Enter to accept the default option to create the

folder, as shown in Figure 4-80. Alternatively, type n for no, or q to quit, and press Enter.

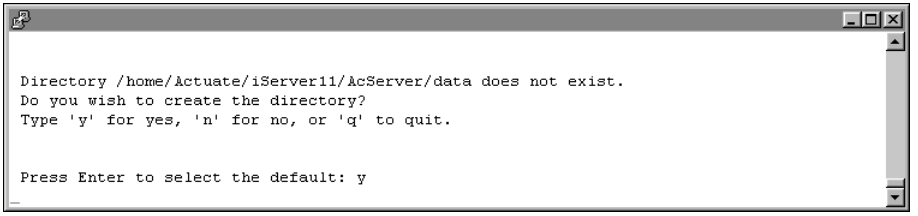

Figure 4-80 Creating the AC_DATA_HOME directory

11 The installer copies the prerequisite files to the destination folder, as shown in Figure 4-81.

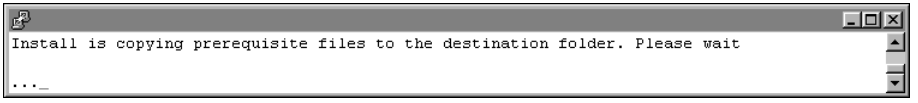

Figure 4-81 Copying prerequisite files

12 Press Enter to accept the default iServer component combination, as shown in Figure 4-82. Alternatively, choose a different iServer component combination and press Enter.

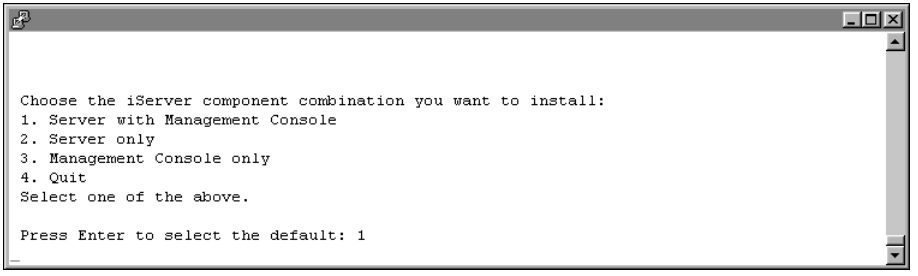

Figure 4-82 Choosing the iServer components to install

13 Press Enter to accept the default stand-alone Server installation, as shown in Figure 4-83.

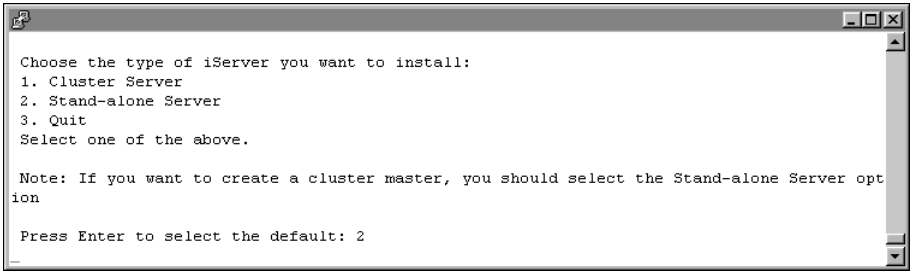

Figure 4-83 Choosing the iServer installation type

14 Type a name to use for the BIRT iServer System name, as shown in Figure 4-84. iServer assigns this name to the default Encyclopedia volume. Additionally, iServer inserts this name into the names iServer creates for the Encyclopedia volume schema and the iServer system schema.

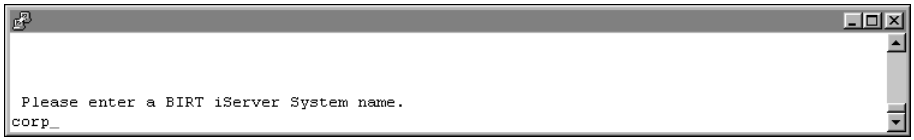

Figure 4-84 Specifying the BIRT iServer System name

15 Press Enter to choose the default embedded PostgreSQL database to store the Encyclopedia volume metadata, as shown in Figure 4-85.

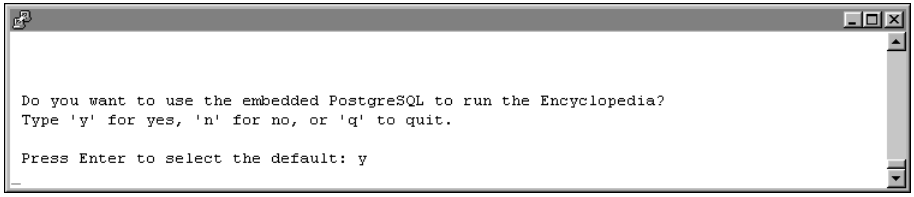

Figure 4-85 Choosing the embedded PostgreSQL

16 Press Enter to choose the default PostgreSQL superuser name, postgres, as shown in Figure 4-86. Alternatively, type a different PostgreSQL superuser name.

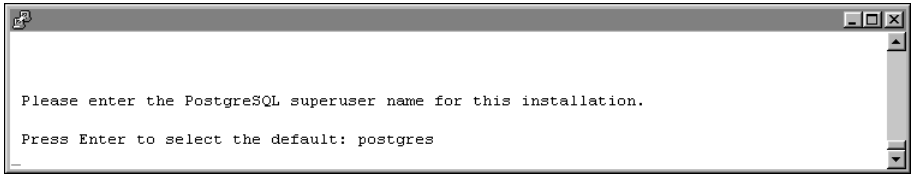

Figure 4-86 Choosing the PostgreSQL superuser name

17 Type a PostgreSQL superuser password that conforms to the password security policy requirements within your organization, then press Enter, as shown in Figure 4-87.

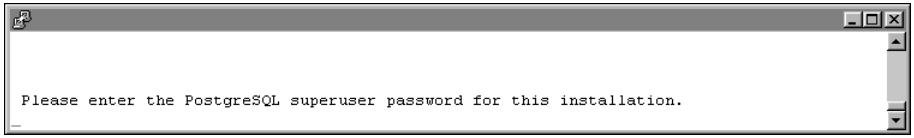

Figure 4-87 Typing the PostgreSQL superuser password

18 Re-enter the password for PostgreSQL superuser, as shown in Figure 4-88, and press Enter.

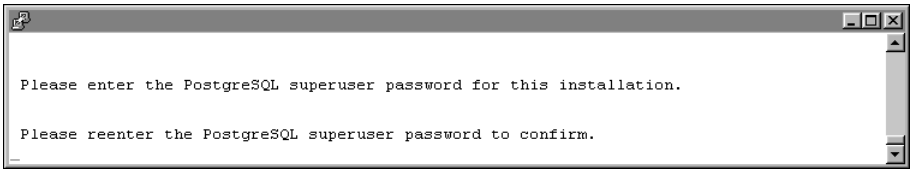

Figure 4-88 Re-entering the PostgreSQL superuser password

19 Press Enter to accept the default port on which the PostgreSQL database server listens for requests, as shown in Figure 4-89. Alternatively, type a different port number.

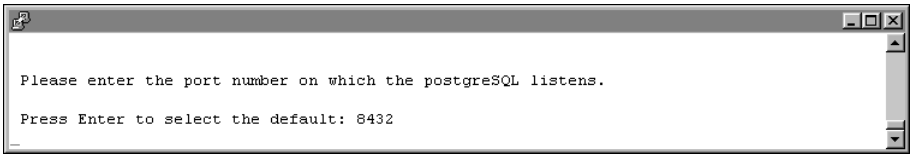

Figure 4-89 Entering the port number which PostgreSQL uses

20 Press Enter to select the default locale, which is English, as shown in Figure 4-90. Alternatively, select a different locale. If you do not see the locale for your region, type m for more and press Enter.

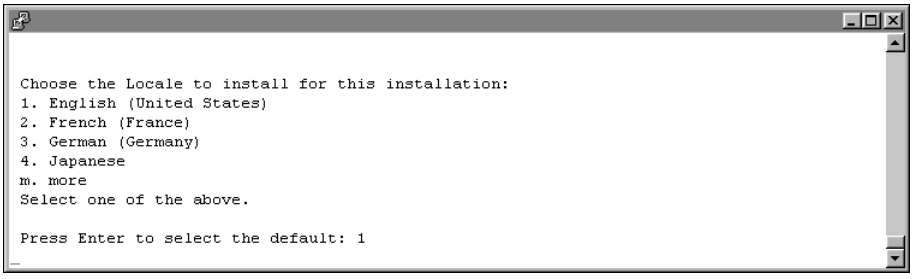

Figure 4-90 Specifying a locale

21 Press Enter to select the default time zone, which is America/Los_Angeles, as shown in Figure 4-91. Alternatively, select another time zone from the numbered list.

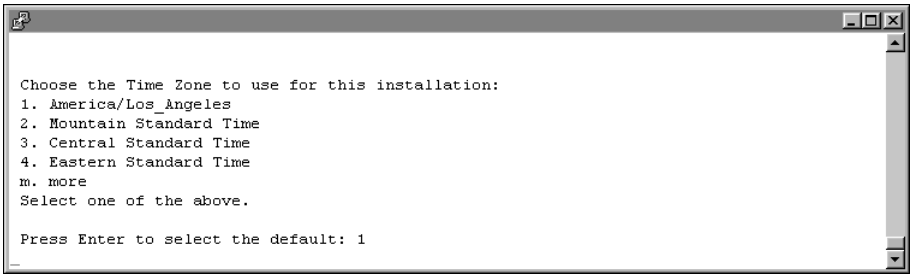

Figure 4-91 Specifying a time zone

22 To evaluate the product using the included evaluation software license press Enter, as shown in Figure 4-92. Alternatively, type 2, then type the path to a purchased license file.

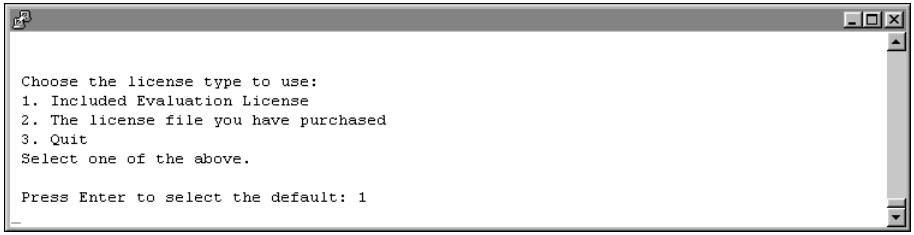

Figure 4-92 Specifying license type

23 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 4-93. Alternatively, type a different IP address.

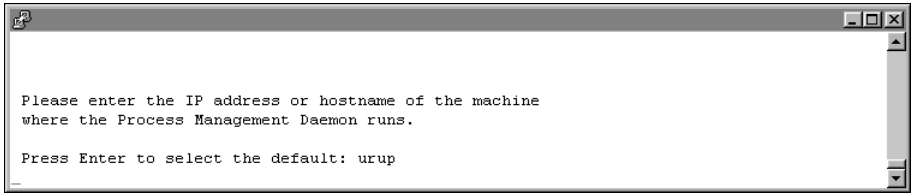

Figure 4-93 Specifying the hostname that Management Console uses to contact the PMD

24 Press Enter to accept the default port number, 8100, where Process Management Daemon (PMD) listens for requests, as shown in Figure 4-94. Alternatively, type a different port number and press Enter.

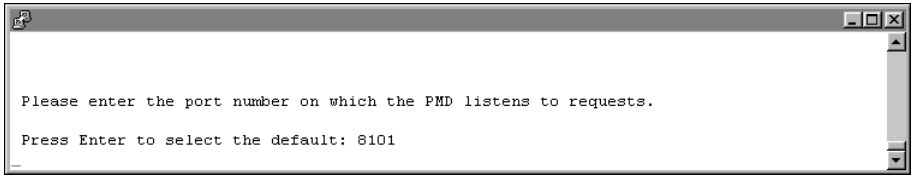

Figure 4-94 Specifying the port number on which the PMD listens

25 Press Enter to accept the default hostname, the name of the machine on which iServer runs, as shown in Figure 4-95. Alternatively, type a different hostname or IP address, then press Enter.

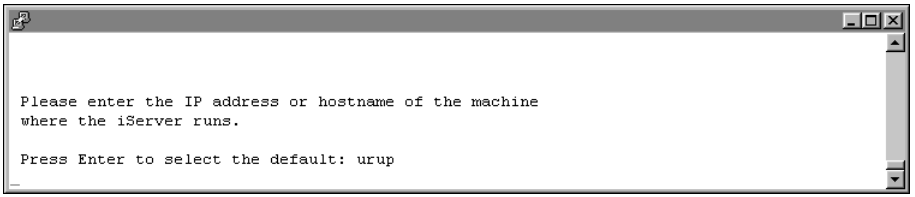

Figure 4-95 Specifying the machine on which the iServer runs

26 Press Enter to accept the default port number where iServer listens to requests, as shown in Figure 4-96. Alternatively, type a different port number and press Enter.

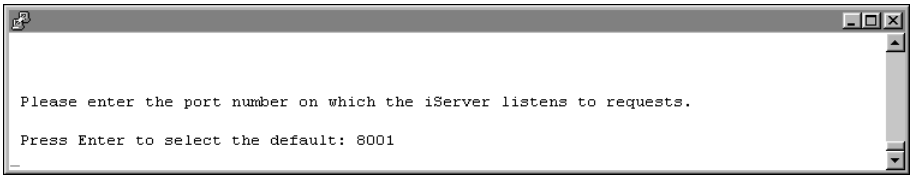

Figure 4-96 Specifying the port number on which the iServer listens

27 Specify the iServer administrator password, as shown in Figure 4-97.

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

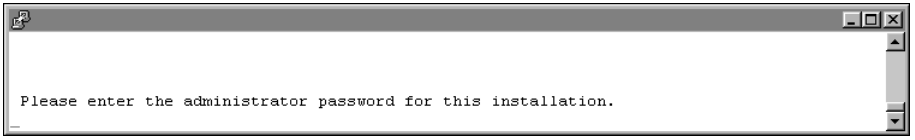

Figure 4-97 Specifying the iServer administrator password

28 Re-enter the password of the iServer administrator, as shown in Figure 4-98.

You use this password to log in to Configuration Console.

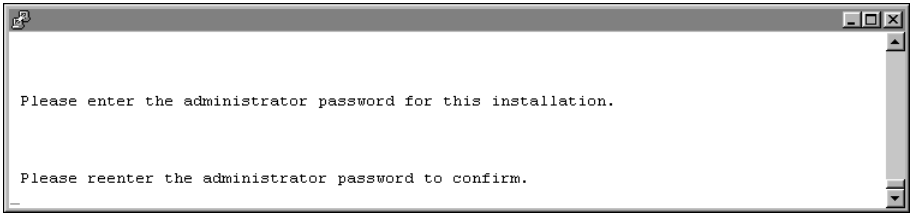

Figure 4-98 Re-entering the iServer administrator password

29 Choose n for no, to prevent iServer from using a volume name, as shown in Figure 4-99. In a side-by-side upgrade, you migrate your volumes from your earlier iServer release.

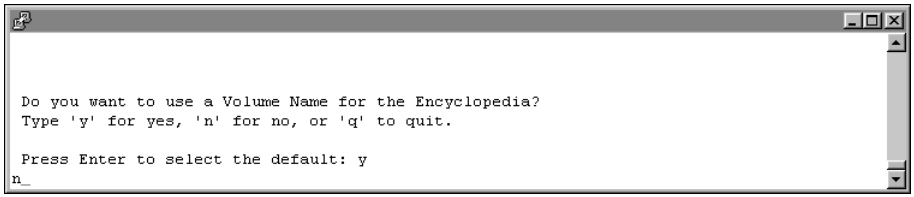

Figure 4-99 Choosing not to use a volume name

30 Press Enter to accept the default option to start iServer automatically, as shown in Figure 4-100. Alternatively, type n for no.

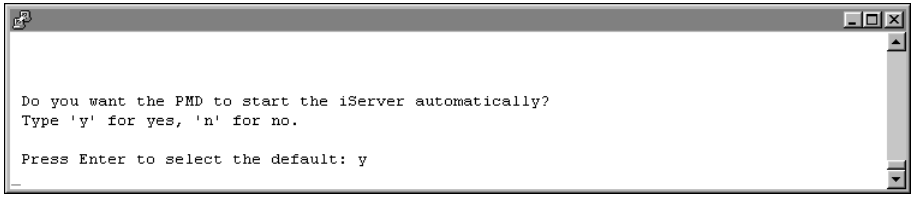

Figure 4-100 Specifying whether to start iServer automatically

31 Press Enter to accept the default option to not integrate LDAP with iServer, as shown in Figure 4-101. Alternatively, you can edit the setting.

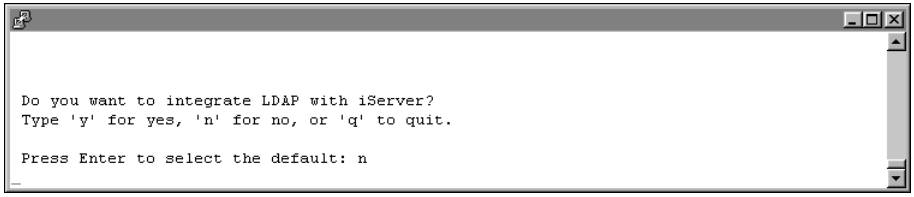

Figure 4-101 Specifying whether to integrate LDAP with iServer

32 Press Enter to accept the default option to not use any database drivers/clients, as shown in Figure 4-102. Alternatively, type y for yes, and specify the database drivers/clients you want to use.

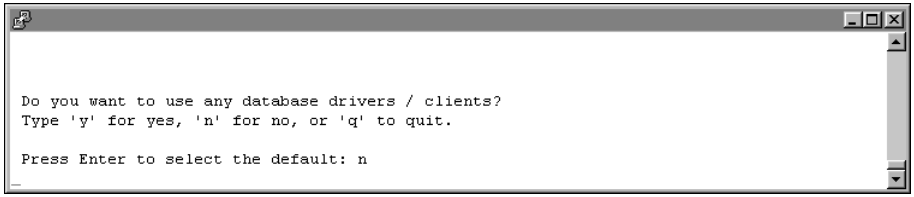

Figure 4-102 Specifying whether to use database drivers/clients

33 Specify what kind of X-Server you want to use, if any. To accept the default, press Enter, as shown in Figure 4-103.

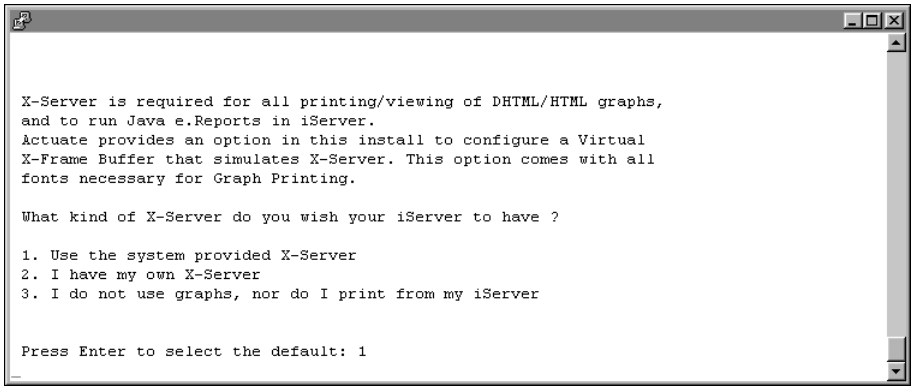

Figure 4-103 Specifying what kind of X-Server to use, if any

34 Press Enter to accept the hostname of the machine that Management Console uses to contact the Process Management Daemon (PMD), as shown in Figure 4-104. Alternatively, type a different IP address.

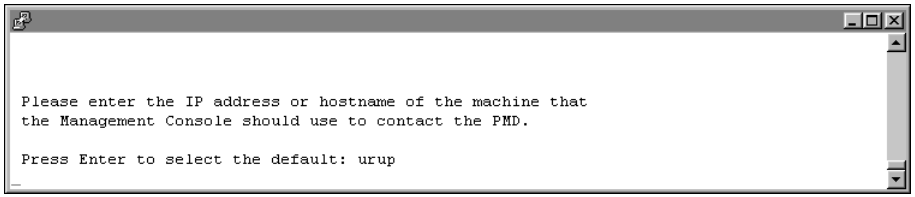

Figure 4-104 Specifying the hostname that Management Console uses to contact the PMD

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 4-105. Alternatively, type a different port number.

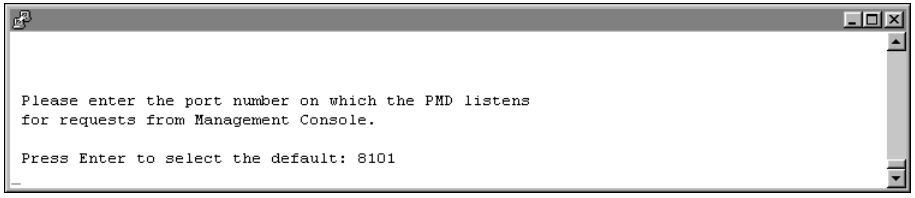

Figure 4-105 Specifying the port number for the PMD to listen for requests from Management Console

36 Press Enter to accept the hostname or enter the IP address of the machine that Management Console uses to contact iServer, as shown in Figure 4-106. Alternatively, type a different IP address.

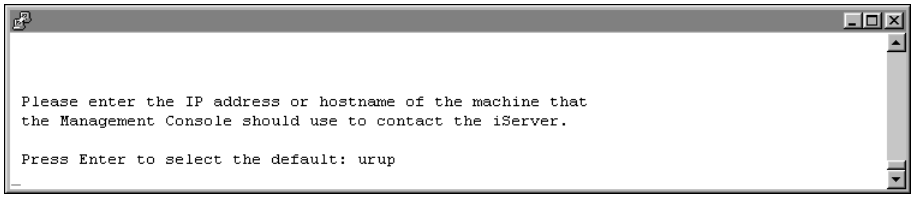

Figure 4-106 Specifying the hostname that Management Console uses to contact iServer

37 Press Enter to accept the default port number, 8000, on which iServer listens for requests from Management Console, as shown in Figure 4-107. Alternatively, type a different port number.

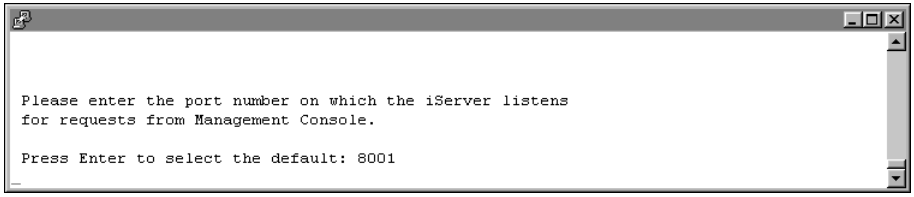

Figure 4-107 Specifying the port number on which iServer listens for requests from Management Console

38 Press Enter to accept the name of the default Encyclopedia volume to use with Management Console, as shown in Figure 4-108. Alternatively, type a different name for the Encyclopedia volume.

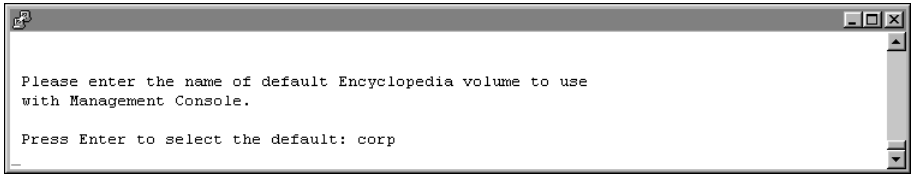

Figure 4-108 Specifying the name of the default Encyclopedia volume

39 Press Enter to accept the default name, acadmin, for the HTTP server context root for Management Console configuration, as shown in Figure 4-109. Alternatively, type a different name.

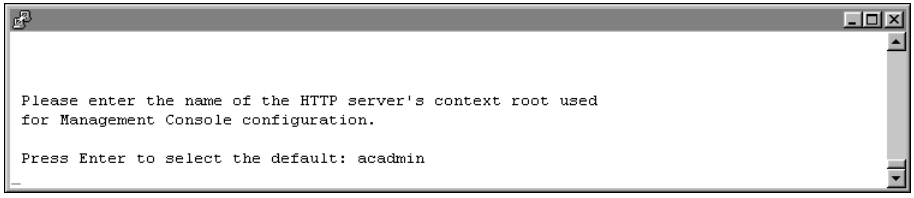

Figure 4-109 Specifying the name of the HTTP server context root

40 Press Enter to accept the default HTTP port number, 8900, on which the application container listens to requests, as shown in Figure 4-110. Alternatively, choose a different port.

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

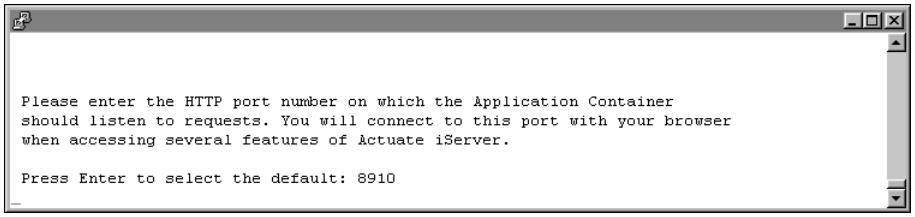

Figure 4-110 Specifying the application container listening port number

41 Review the settings, as shown in Figure 4-111, 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.

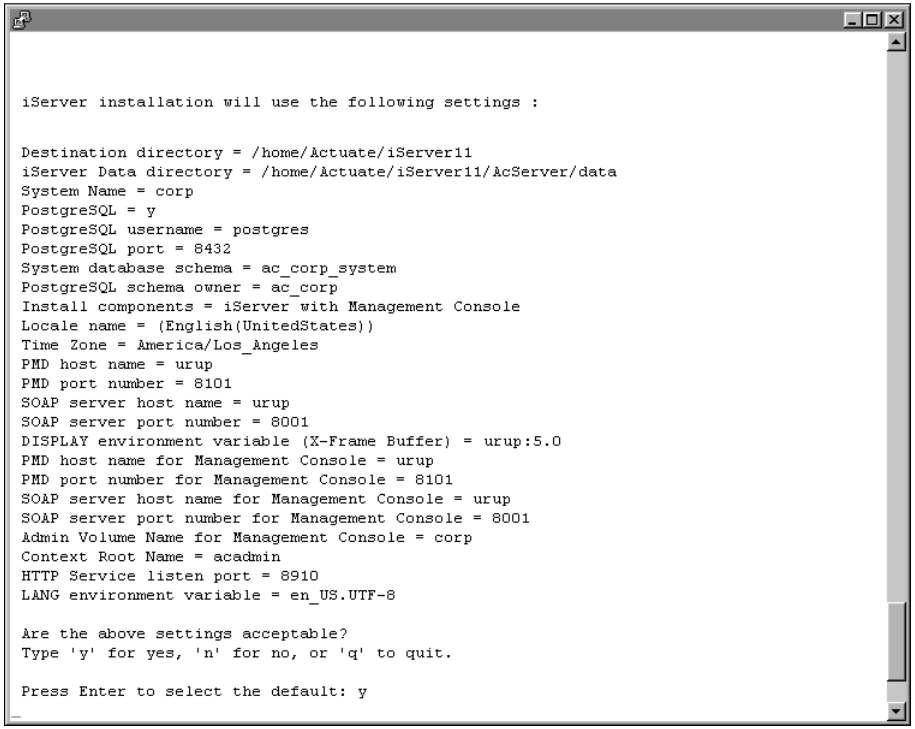

Figure 4-111 Reviewing the installation settings

42 The installation program installs iServer, displaying an indicator that shows the progress of the installation, as shown in Figure 4-112.

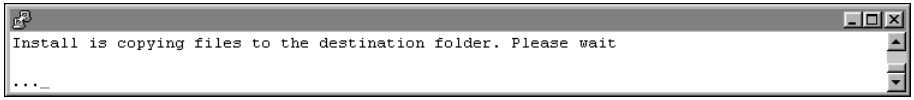

Figure 4-112 Viewing iServer installation progress

43 At the end of the installation, the program asks if you want to start iServer. Accept the default, y for yes, to start iServer, as shown in Figure 4-113.

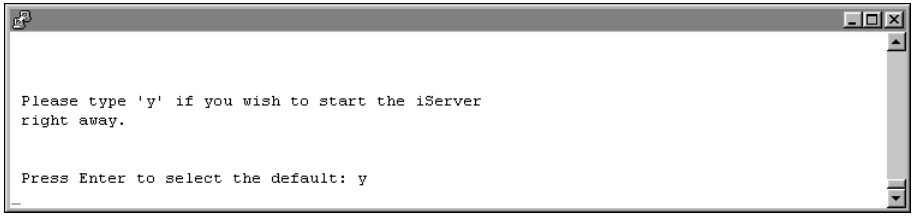

Figure 4-113 Specifying whether to start iServer

44 When the installation program finishes, it provides additional information about localization, logging in using an account with root permissions to start iServer, and installing online help and manuals, as shown in Figure 4-114.

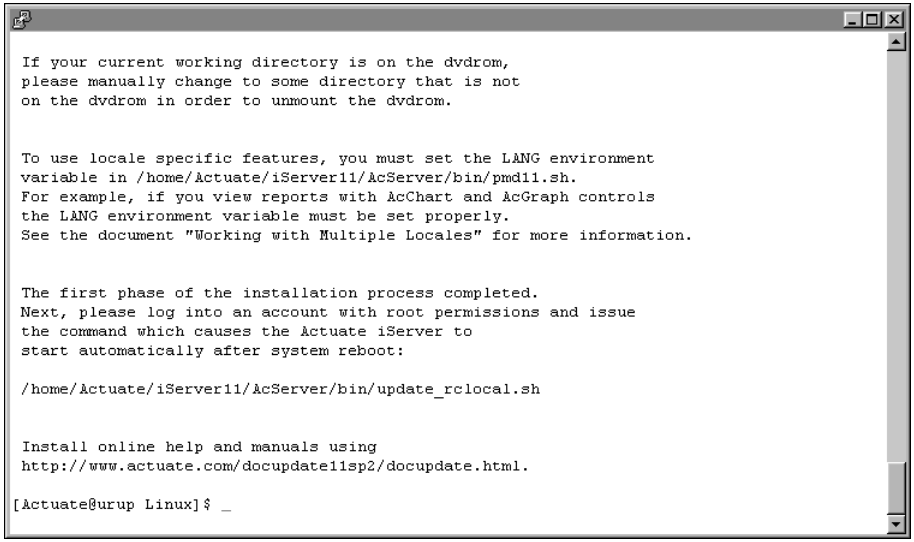

Figure 4-114 Viewing information about localization, logging in, and installing online help

Performing a manual side-by-side migration

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

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

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

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

How to perform a manual export of an Encyclopedia volume in a side-by-side installation

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

1. In a web browser, type:

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

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

3. If you have not migrated a volume from an earlier release to BIRT iServer Release 11 yet, you must first delete the default partition that the installation program creates. If you have already performed this step, skip to step 4. Otherwise, perform the following tasks:

   1. Choose Advanced view.

   2. From the side menu, choose Partitions. On Partitions, point to the arrow next to DefaultPartition and choose Delete, as shown in Figure 4-115.

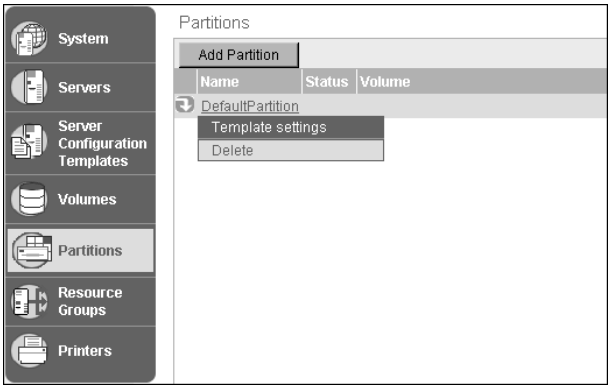

Figure 4-115 Deleting DefaultPartition

4 Create a partition on BIRT iServer Release 11 for the volume that you are going to migrate from the previous iServer release. To create a new partition, perform the following tasks:

1. In the operating system, create a folder for a new partition in AC_DATA_HOME, which points to the location of the iServer data specified during the BIRT iServer Release 11 installation.
2. Log in to Configuration Console.
3. Choose Advanced view.
4. From the side menu, choose Partitions, then choose Add Partition.
5. In Partition name, type the name of the earlier iServer release partition that specifies the path to the volume that you want to migrate. For example, if the name of the Encyclopedia volume partition on the earlier iServer release is ac_corp_partition, type that name in Partition name, as shown in Figure 4-116.

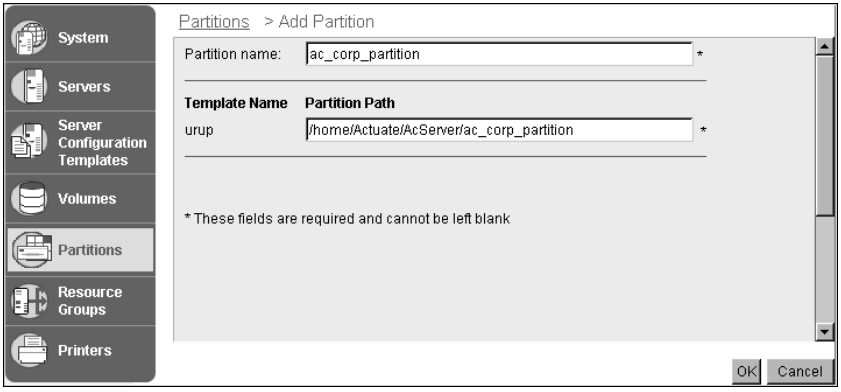

Figure 4-116 Creating a partition for a migrated volume

6 In Partition Path, type the path to the iServer Release 11 Encyclopedia volume folder, as shown in Figure 4-116. This path does not need to match the partition path to the same volume on the earlier iServer release. Choose OK.

5 Copy the file, filetype, and status folders from the old Encyclopedia volume folder to the iServer Release 11 Encyclopedia volume folder. For example, if you want to migrate the default volume from release 10SP1, the location of these folders is AC_SERVER_HOME/encyc. The default AC_SERVER_HOME path in release 10SP1 is /home/Actuate/AcServer.

6 In this step, create a new schema to use to create a new volume. In the Advanced view of Configuration Console, from the side menu, choose Volumes.

On Volumes, point to the icon next to Default ActuatePostgreSQL MetadataDatabase and choose Add volume schema, as shown in Figure 4-117.

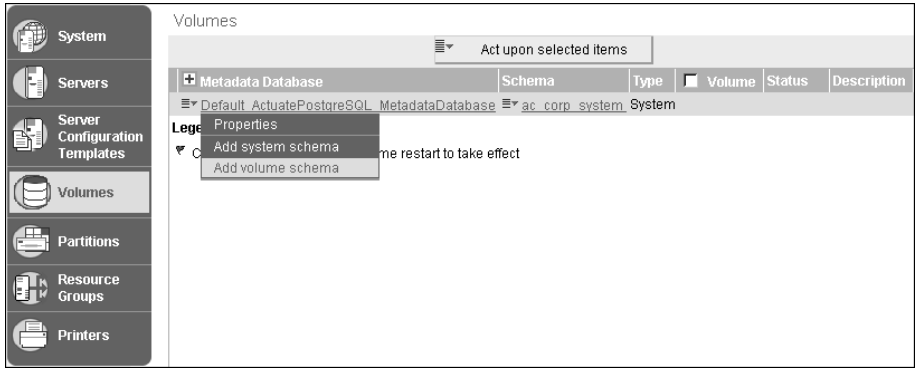

Figure 4-117 Choosing to add a new schema

On Volumes—New Schema, perform the following tasks:

1 In Schema name, type a name for the schema. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.

2 In Database schema name, type a name for the Database schema. The name must be less than 30 characters. Observe the same naming restrictions for this schema as the volume schema name.

3 In Database schema password, type a new password.

4 In Database schema password confirm, re-type the new password.

5 In Database superuser, type the database superuser name. For the PostgreSQL RDBMS that installs with iServer by default, the PostgreSQL superuser name is postgres.

6 In Database superuser password, type the superuser password. This password is the same password that the installer provides when installing iServer. Choose OK.

Volumes—New Schema appears as shown in Figure 4-118.

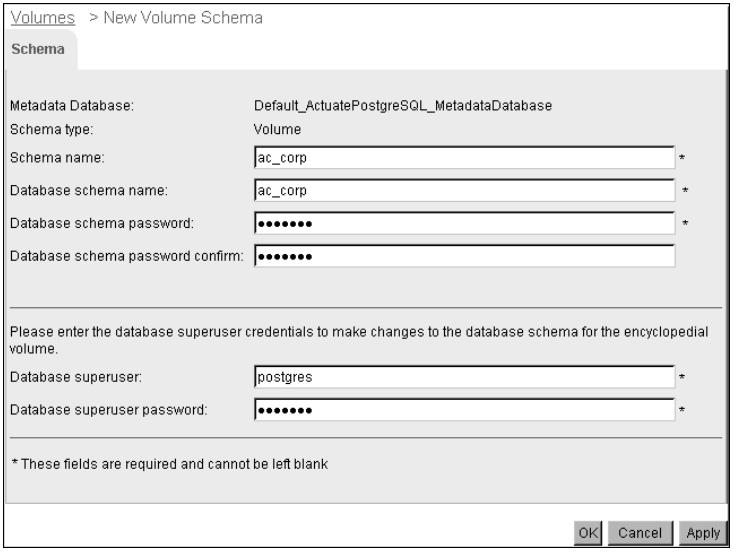

Figure 4-118 Creating a new schema

The new schema appears in the list of schemas on Volumes, as shown in Figure 4-119.

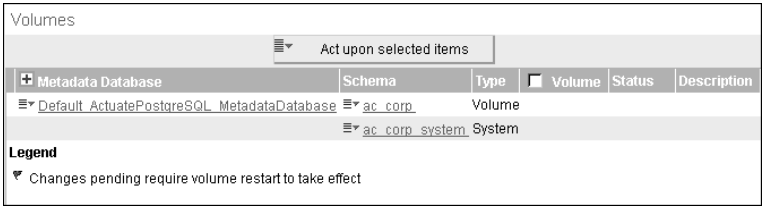

Figure 4-119 Viewing the new schema

7 Shut down the earlier iServer release by performing the following tasks:

1 In a web browser type:

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

2 Log into 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 4-120.

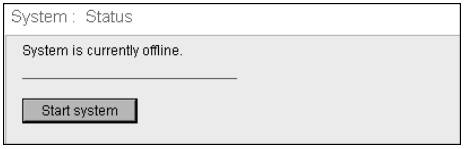

Figure 4-120 iServer is offline

8 Run the Squirrel Data Exporter by performing the following tasks:

1. Create a properties file named SquirrelDataExporter.properties. You pass this file to the Squirrel Data Exporter utility when you execute it. Use the following example to create SquirrelDataExporter.properties:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
AC_DATA_HOME = /home/Actuate/iServer11/AcServer/data
SQUIRREL_DATA_HOME = /home/Actuate/AcServer/ac_corp_partition
SQUIRREL_EXPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
SQUIRREL_EXPORT_FORMAT = PostgreSQL
NEW_SCHEMA_NAME =  ac_corp
NEW_VOLUME_NAME = corp
```

   See Table 4-3 for descriptions of the required Squirrel Data Exporter properties. See Table 4-4 for descriptions of the optional Squirrel Data Exporter properties. See "Specifying Squirrel Data Exporter properties" and "Specifying Encyclopedia Data Store Administrator properties" for additional notes on property files.

2. Create a script to run the following commands. Adjust the location of AC_SERVER_HOME if the location is different in your environment:

```
export JAVA_HOME=/home/Actuate/AcServer/jdk160
export AC_SERVER_HOME=/home/Actuate/iServer11/AcServer
export LD_LIBRARY_PATH=$AC_SERVER_HOME/lib
export PATH=$JAVA_HOME/bin:$LD_LIBRARY_PATH:$PATH
export CLASSPATH=
   $AC_SERVER_HOME/jar/AcCommon.jar:
   $AC_SERVER_HOME/jar/JDBCDrivers.jar:
   $AC_SERVER_HOME/jar/com.actuate.common-server.jar:
   $AC_SERVER_HOME/jar/ServerAdminTools.jar
export SQUIRREL_DATA_EXPORTER=
   com.actuate.iserver.encyclopedia.datastore.admin
   .SquirrelDataExporter
java $SQUIRREL_DATA_EXPORTER SquirrelDataExporter.properties
```

3. Open a command prompt and run the script.

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

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

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

1 Create a properties file named VolumeImport.properties. You pass this file to the Encyclopedia Data Store Administrator utility when you execute it. Use the following example to create VolumeImport.properties:

```
AC_SERVER_HOME = /home/Actuate/iServer11/AcServer
AC_DATA_HOME = /home/Actuate/iServer11/AcServer/data
DATABASE_TYPE = PostgreSQL
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
SCHEMA_NAME =  ac_corp
SCHEMA_PASSWORD = <your schema password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = /home/Actuate/SquirrelData/ac_corp
```

See Table 4-5 for descriptions of the required Encyclopedia Data Store Administrator properties. See Table 4-6 for descriptions of the optional Encyclopedia Data Store Administrator properties. See "Specifying Squirrel Data Exporter properties" and "Specifying Encyclopedia Data Store Administrator properties" for additional notes on property files.

2 Create a script to run the following commands. Adjust the location of AC_SERVER_HOME if the location is different in your environment:

```
export JAVA_HOME=/home/Actuate/AcServer/jdk160
export AC_SERVER_HOME=/home/Actuate/iServer11/AcServer
export LD_LIBRARY_PATH=$AC_SERVER_HOME/lib
export PATH=$JAVA_HOME/bin:$LD_LIBRARY_PATH:$PATH
export CLASSPATH=
   $AC_SERVER_HOME/jar/AcCommon.jar:
   $AC_SERVER_HOME/jar/JDBCDrivers.jar:
   $AC_SERVER_HOME/jar/com.actuate.common-server.jar:
   $AC_SERVER_HOME/jar/ServerAdminTools.jar
export ENCYCLOPEDIA_DATASTORE_ADMINISTRATOR=
   com.actuate.iserver.encyclopedia.datastore.admin
   .EncyclopediaDataStoreAdministrator
java $ENCYCLOPEDIA_DATASTORE_ADMINISTRATOR
   VolumeImport.properties
```

3 Run the Encyclopedia Data Store Administrator script from the command prompt.

2 Create a new volume for the migrated volume by performing the following tasks:

1 Log in to Configuration Console and choose Advanced view.

2 From the side menu, choose Volumes.

3 On Volumes, point to the icon next to a schema and choose Add Volume, as shown in Figure 4-121.

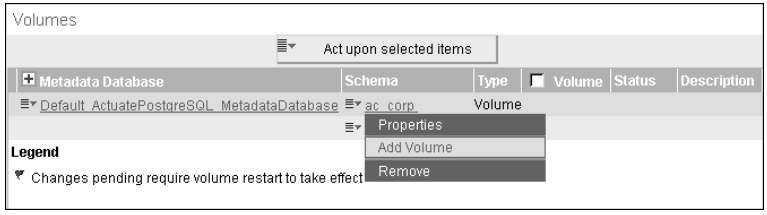

Figure 4-121 Choosing to add a volume

4 On New Volume—General, perform the following tasks:

1 In Volume name, type the name you specified for NEW_VOLUME_NAME in SquirrelDataExporter.properties.

2 In Primary partition, specify the partition you created previously, as shown in Figure 4-122.

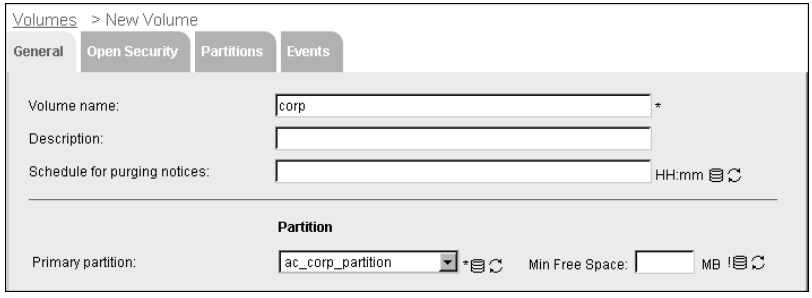

Figure 4-122 Specifying volume and partition name

Choose Partitions.

5 On New Volume—Partitions, perform the following tasks:

1 In Available partitions, select the partition that you specified in the previous step, then move it to Selected by choosing the right arrow.

2 In Selected partitions, select the partition. Choose Start, as shown in Figure 4-123.

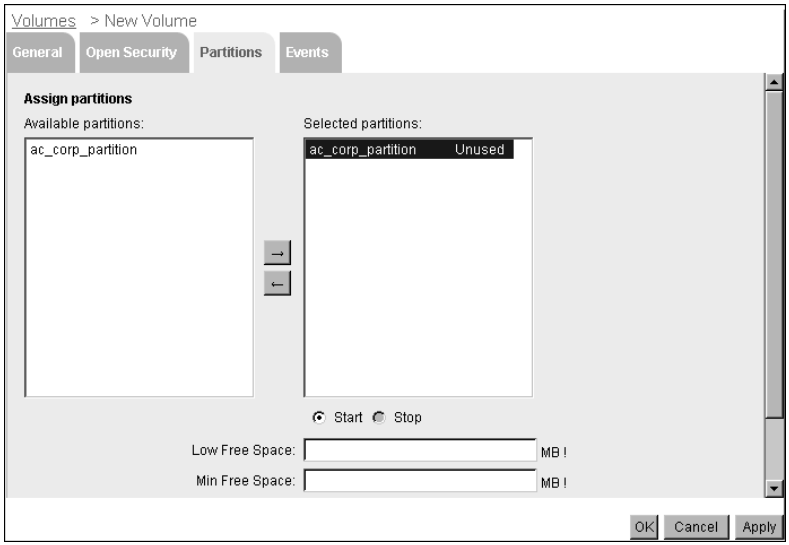

Figure 4-123 Assigning a partition

Choose OK.

3 On Volumes, point to the icon next to the new volume name and choose Take online, as shown in Figure 4-124.

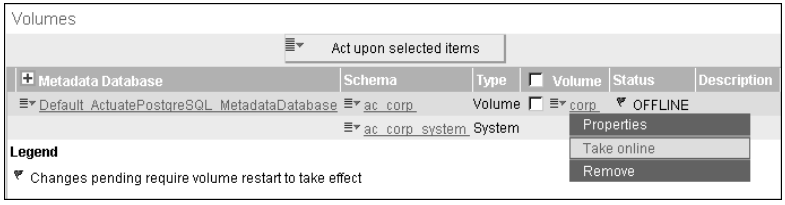

Figure 4-124 Viewing the new volume

The new volume comes online, as shown in Figure 4-125.

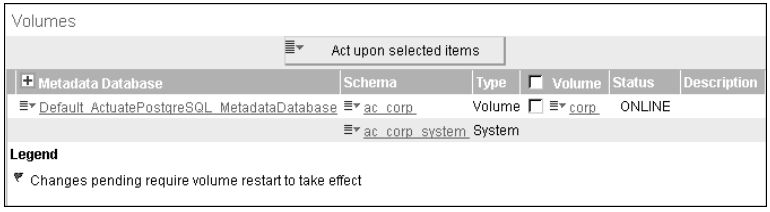

Figure 4-125 Viewing the new volume after it comes online

4 Log on to Management Console. In Files and Folders, the data from your previous release appears.

5 If you are satisfied that the contents of the migrated volume are correct and complete, delete the folder specified in SQUIRREL_EXPORT_FOLDER property in the SquirrelDataExporter.properties file.

Specifying Squirrel Data Exporter properties

Table 4-3 describes the required Squirrel Data Exporter properties used to configure the environment for a manual export operation.

Table 4-3 Required Squirrel Data Exporter properties

Parameter	Description
AC_SERVER_HOME	Points to the location of the iServer binaries, which you specify during the BIRT iServer Release 11 installation.
AC_DATA_HOME	Points to the location of the iServer data, which you specify during the BIRT iServer Release 11 installation.
SQUIRREL_DATA_HOME	Absolute path to the folder on your earlier iServer release containing the Encyclopedia volume you want to migrate.
SQUIRREL_EXPORT_FOLDER	Absolute path to the folder that Squirrel Data Exporter creates, containing the volume data from SQUIRREL_DATA_HOME, which Encyclopedia Data Store Administrator loads into the new volume on BIRT iServer Release 11.
SQUIRREL_EXPORT_FORMAT	Format of exported data. Specify Oracle or PostgreSQL.
NEW_SCHEMA_NAME	The name of the target schema into which the Squirrel Data Exporter loads data. Restrict the schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.
NEW_VOLUME_NAME	The name of the migrated volume on BIRT iServer Release 11. This name does not need to match the name this volume had on your earlier iServer release.

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

Table 4-4 Optional Squirrel Data Exporter properties

Property	Value	Default Value
FILE_LOG_LEVEL	The minimum logging level for messages sent to the log file. This parameter only works for AcLogger. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Data Store Administrator tools do not generate large logs.	5
LOG_FILE_SIZE	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
LOG_FOLDER	Absolute path to the log folder.	{AC_DATA_HOME/ server/log.

Specifying Encyclopedia Data Store Administrator properties

Table 4-5 describes the required Encyclopedia Data Store Administrator properties used to configure the environment for a manual migration operation.

Table 4-5 Required Encyclopedia Data Store Administrator properties

Parameter	Description
DATABASE_TYPE	Type of supported RDBMS that contains the data store: PostgreSQL or Oracle
LOG_FOLDER	Absolute path to the log folder
SCHEMA_FILE_NAME	Base name of the file without the extension that contains the meta-schema definition
SCHEMA_PASSWORD	Schema owner password

Table 4-5 Required Encyclopedia Data Store Administrator properties

Parameter	Description
SCRIPT_HOME	Absolute path to the root of the folder hierarchy that contains the scripts and the meta-schema definition

Table 4-6 describes optional Encyclopedia Data Store Administrator properties. The properties used depend on the type of operation performed and the installation environment.

Table 4-6 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
CONSOLE_LOG _LEVEL	The minimum logging level for messages sent to the console (stderr). Supported values are INFO, CONFIG, FINE, FINER, and FINEST.	INFO	

(continues)

Table 4-6 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default Value	Supported Databases
CREATE_SCHEMA	Set to true to create a new schema, false to use an existing schema. Specifies whether to create a new Encyclopedia schema. PostgreSQL only, not supported for Oracle. Creating a schema automatically creates the schema owner and application user if necessary.	False	PostgreSQL
DATA_EXPORT _FOLDER	Absolute path to the folder to which data is exported. This parameter is required if EXPORT_DATA is true. Folder is on the database machine.		All
DATA_EXPORT _FORMAT	Format of exported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE _TYPE}	All
DATA_IMPORT _FOLDER	Absolute path to the folder from which data is imported. This parameter is required if IMPORT_DATA is true. For PostgreSQL, the data must be on the database server.		PostgreSQL
DATA_IMPORT _FORMAT	Format of imported data. Specify 'Oracle' or 'PostgreSQL'.	{DATABASE _TYPE}	PostgreSQL
DATA_ INITIALIZATION _SCRIPT	Base name of the file without the extension that contains the data initialization script. This parameter is required if INITIALIZE_DATA is true.		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

Table 4-6 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default Value	Supported Databases
DATABASE_PORT	The port that the database server uses. This value is required for PostgreSQL. Required for Oracle if not using TNS.		All
DELETE_ALL_DATA	Set to true to delete all data from the data store. Does not delete data from other tables in the schema.	False	All
DELETE_DATA	Set to true to delete filtered data from the data store. This parameter only works when a filter value is available.	False	All
DEPOPULATE_SCHEMA	Set to true to delete data store elements such as tables, views, stored functions, and procedures from the schema. Does not remove any other objects in the schema.	False	All
DROP_SCHEMA	Set to true to delete the schema.	False	PostgreSQL
EXPORT_ALL_DATA	Set to true to export all data from the data store.	False	All
EXPORT_DATA	Set to true to export filtered data from the data store. This parameter only works when a filter value is available.	False	All
FILE_LOG_LEVEL	The minimum logging level for messages sent to the log file. This parameter only works for AcLogger, and is a standard JUL Level name. Supported values are CONFIG, FINE, FINER, and FINEST.	CONFIG	All
GENERATE_SCRIPTS	Set to true to generate scripts for performing the operations instead of performing the operations directly.	False	All

(continues)

Table 4-6 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default Value	Supported Databases
IMPORT_DATA	Set to true to import data into the data store.	False	All
INITIALIZE_DATA	Set to true to initialize the data in the data store, using the data initialization script.	False	All
LOG_FILE_COUNT	Maximum number of log files to create before starting to overwrite previous log files. Must be a valid integer. Data Store Administrator tools do not generate large logs.	5	All
LOG_FILE_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
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.	{SCHEMA _NAME}	All
NEW_SCHEMA _PASSWWORD	Password of the database superuser. This parameter is required if {NEW_SCHEMA_ NAME} is present.	{NEW _SCHEMA _PASSWORD}	All
NEW _VOLUME_NAME	Name of the new volume to be created, or the name of the target schema for data export. Do not use a hyphen in a volume name.	{NEW _VOLUME _NAME}	All
ORACLE_TNS _NAMES_FILE	Absolute path to the Oracle TNS names file. This path can be used instead of {DATABASE_HOST} and {DATABASE_PORT} to generate a JDBC URL.		Oracle

Table 4-6 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 ancillary structures in database.	{POPULATE _SCHEMA}	All
SCHEMA_NAME	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema. Restrict schema name to alphanumeric and underscore characters with an initial alphabetic character in the pattern [a-z][a-z 0-9]*. Do not use a hyphen.	{NEW _SCHEMA _NAME}	All
SCHEMA_PASSWORD	Name of the existing schema. Must be a legal SQL identifier. This parameter is required when performing operations on an existing schema.	{SCHEMA _NAME}	All

(continues)

Table 4-6 Optional Encyclopedia Data Store Administrator properties (continued)

Parameter	Description	Default Value	Supported Databases
SUPERUSER	User ID of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL Typically postgres in a PostgreSQL database.		PostgreSQL
SUPERUSER _PASSWORD	Password of the database superuser. This parameter is required if CREATE_SCHEMA is true or when importing data into PostgreSQL		PostgreSQL
SYSTEM_DATABASE _NAME	Name of the system database.		PostgreSQL
TABLESPACE_NAME	Name of the application tablespace.		PostgreSQL
TABLESPACE _LOCATION	Absolute path to the application tablespace folder.		PostgreSQL
TIME_ZONE	Local time zone for installation.		All
VOLUME_NAME	Name of the volume or target schema for data import.	{VOLUME _NAME}	All

Specifying System Data Store Administrator properties

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

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

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

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

- Import or export data.

The following script provides an example of how to set properties and run the System Data Store Administrator utility in the Linux environment:

```
#!/bin/sh
export JAVA_HOME=/home/Actuate/AcServer/jdk160
export JAVA=$JAVA_HOME/bin/java
export PATH=$JAVA_HOME/bin:$PATH
export AC_SERVER_HOME=/home/Actuate/AcServer
export CLASSPATH=$AC_SERVER_HOME/jar/AcCommon.jar\
:$AC_SERVER_HOME/jar/JDBCDrivers.jar\
:$AC_SERVER_HOME/jar/com.actuate.common-server.jar\
:$AC_SERVER_HOME/jar/ServerAdminTools.jar
export SYSTEM_DATASTORE_ADMINISTRATOR=\
   com.actuate.iserver.encyclopedia.datastore.admin
      .SystemDataStoreAdministrator
$JAVA -cp $CLASSPATH $SYSTEM_DATASTORE_ADMINISTRATOR
   systemdatastore.properties
```

In this example, the backslash is a line-continuation character.

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

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

Performing operations using System Data Store Administrator utility

The System Data Store Administrator utility supports a 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.

Creating and Populating a System Schema

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

Configure the properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
AC_DATA_HOME = /home/Actuate/AcServer/data
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>
NEW_SCHEMA_PASSWORD = <provide a password>
```

Populating a System Schema

Perform this operation on a system schema that already exists. Importing data into PostgreSQL requires specifying schema owner credentials and the iserver application user name.

Configure the properties as shown in the following example:

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

Performing operations using Encyclopedia Data Store Administrator utility

The Encyclopedia Data Store Administrator utility supports a wide range of import and export operations in the BIRT iServer environment. The following sections describe a mix of operational scenarios that an administrator can perform using this utility.

Importing One or More Volumes into a New Schema

This operation is only supported for PostgreSQL. NEW_SCHEMA_NAME does not need to match the schema name from which the Squirrel Data Exporter utility exported the data. Do not change the volume names, since these names are defined in the exported data.

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

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
AC_DATA_HOME = /home/Actuate/AcServer/data
DATABASE_TYPE = PostgreSQL
SYSTEM_DATABASE_NAME = postgres
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
DATABASE_HOST = localhost
DATABASE_PORT = 8432
CREATE_SCHEMA = true
NEW_SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
APPLICATION_USER_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {SQUIRREL_EXPORT_FOLDER}
```

Importing One or More Volumes into a Populated Schema

This operation is only supported for PostgreSQL. Typically, you perform this operation to import an additional volume into a shared schema that already contains one or more volumes.

SCHEMA_NAME does not need to match the schema name from which the Squirrel Data Exporter utility exported the data. Do not change the volume names, because these names are defined in the exported data. Importing data into PostgreSQL requires superuser privileges.

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
AC_DATA_HOME = /home/Actuate/AcServer/data
DATABASE_TYPE = PostgreSQL
SUPERUSER = postgres
SUPERUSER_PASSWORD = <your superuser password>
DATABASE_NAME = iserver
```

(continues)

```
DATABASE_HOST = localhost
DATABASE_PORT = 8432
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
IMPORT_DATA = true
DATA_IMPORT_FOLDER = {SQUIRREL_EXPORT_FOLDER}
```

Exporting All Volumes from a Schema

Squirrel Data Exporter generates bulk load script files in the same folder as the exported data. You can omit NEW_SCHEMA_NAME if it is the same as SCHEMA_NAME.

Configure these properties as shown in the following example:

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

Exporting A Single Volume from a Schema

Squirrel Data Exporter generates bulk load script files in the same folder as the exported data. 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
AC_DATA_HOME = /home/Actuate/AcServer/data
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
AC_DATA_HOME = /home/Actuate/AcServer/data
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
AC_DATA_HOME = /home/Actuate/AcServer/data
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
AC_DATA_HOME = /home/Actuate/AcServer/data
DATABASE_TYPE = Oracle
DATABASE_NAME = xe
DATABASE_HOST = localhost
DATABASE_PORT = 1521
SCHEMA_NAME = <provide a name>
SCHEMA_PASSWORD = <provide a password>
APPLICATION_USER = iserver
POPULATE_SCHEMA = true
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

Creating a New Volume in a Populated Schema

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

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

Configure these properties as shown in the following example:

```
AC_SERVER_HOME = /home/Actuate/AcServer
AC_DATA_HOME = /home/Actuate/AcServer/data
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
AC_DATA_HOME = /home/Actuate/AcServer/data
DATABASE_TYPE = PostgreSQL
SYSTEM_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>
INITIALIZE_DATA = true
NEW_VOLUME_NAME = <provide a name>
TIME_ZONE = America/Los_Angeles
```

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

When configuring Encyclopedia Data Store Upgrader, Squirrel Data Exporter, and Encyclopedia Data Store Administrator properties files, observe the following rules:

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

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

Chapter

5

Installing a BIRT iServer cluster

This chapter discusses the following topics:

- Preparing to install an iServer cluster
- Installing a BIRT iServer cluster node
- Finding the BIRT iServer home directory
- About the Java Development Kit

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 account with the right privileges

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.

About 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.

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 uses:

- 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 Bermuda:

```
# setenv DISPLAY bermuda: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 are BIRT iServer installation options in these environments. To use Xvfb in an HP-UX, you must install Xvfb before you install BIRT iServer.

About HP-UX 11i installation

Minimum hardware requirements for HP-UX 11i are a PA-RISC 2.0 processor and 1024 MB of RAM. For more information about HP-UX system requirements, see the Supported Products and Obsolescence Policy on the Actuate web site at the following URL:

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

Working with large files on HP-UX-11i

Actuate works with large files, greater than two gigabytes, on all platforms. You must reconfigure the HP-UX 11i file system as a large file system for this feature to take effect. The disk where the file resides must be a local disk. Contact HP Customer Support for information about how to reconfigure the file system.

Installing Xvfb

Actuate does not provide an Xvfb library for HP-UX 11i. If you choose to use an X frame buffer configured by the BIRT iServer installation instead of using an existing X server, your machine must have the Xvfb binary installed.

How to configure X frame buffer

To configure X frame buffer after the Actuate BIRT iServer System installation, perform the following steps:

1 Download and install the Xvfb patch from HP.

2 Create a symbolic link to the X frame buffer under AC_SERVER_HOME/xvfb/bin.

3 Start your Actuate BIRT iServer System.

Installing a BIRT iServer cluster node

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 Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required 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.

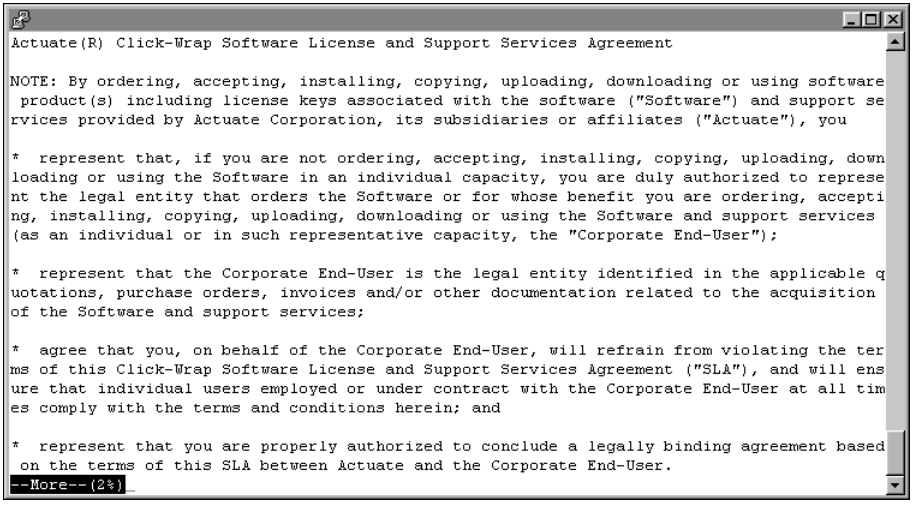

Figure 5-1 Reviewing the license agreement

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.

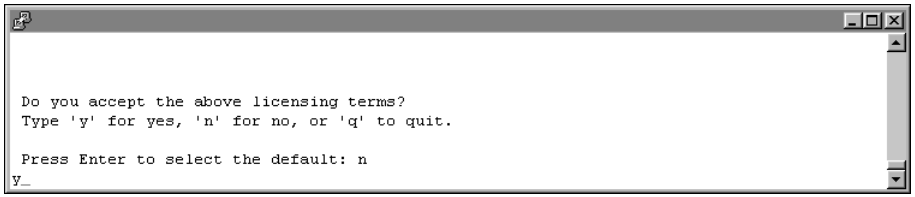

Figure 5-2 Accepting the licensing terms

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

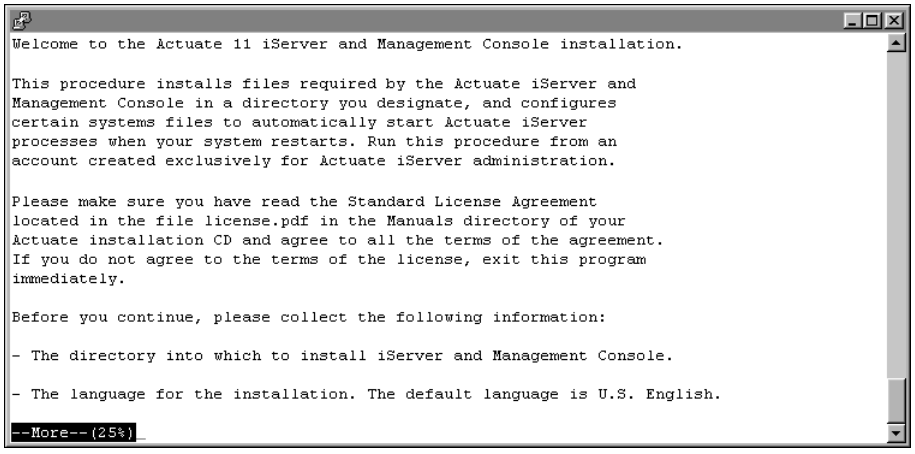

Figure 5-3 Reviewing the introductory information

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

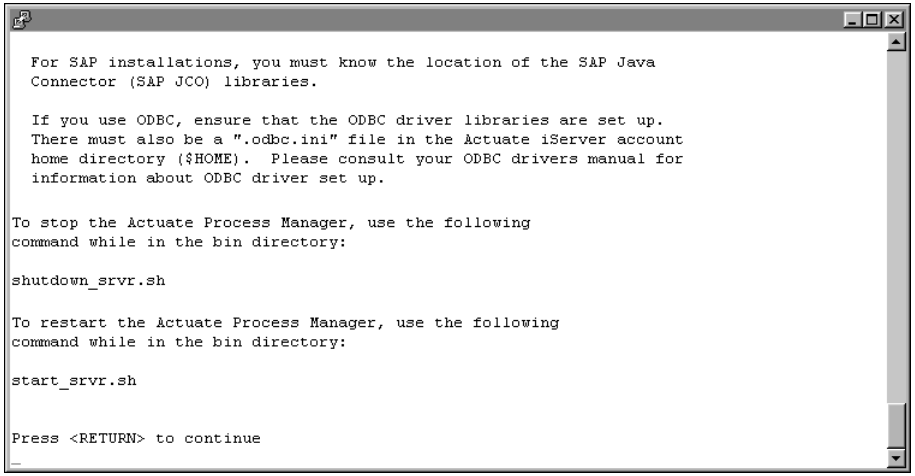

Figure 5-4 Finishing the review of introductory information

7 Press Enter to accept the default location for installation, 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.

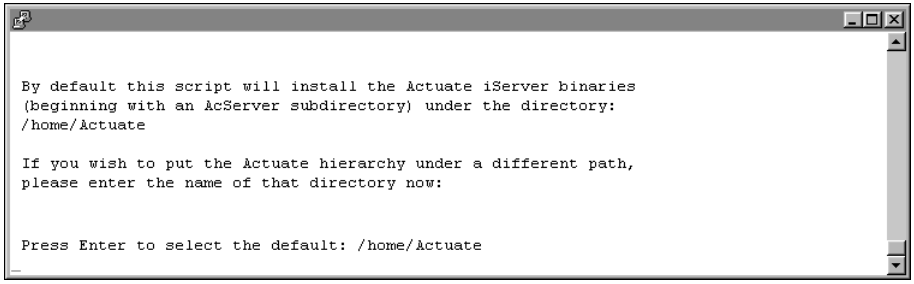

Figure 5-5 Specifying the installation directory

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.

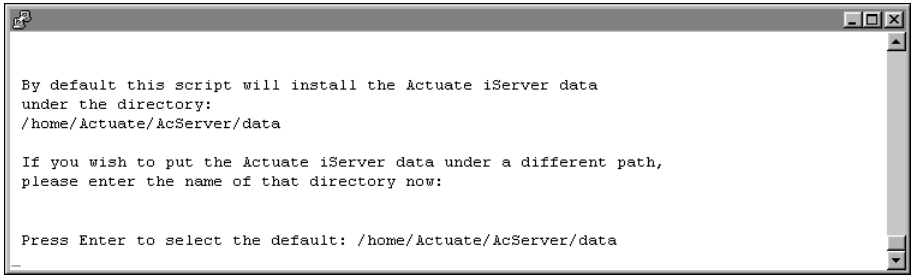

Figure 5-6 Specifying the data installation directory

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.

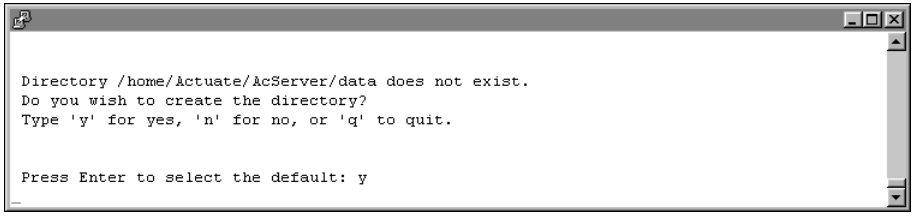

Figure 5-7 Creating the AC_DATA_HOME directory

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

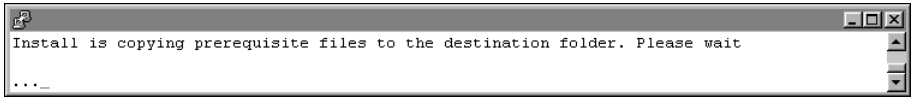

Figure 5-8 Copying prerequisite files

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.

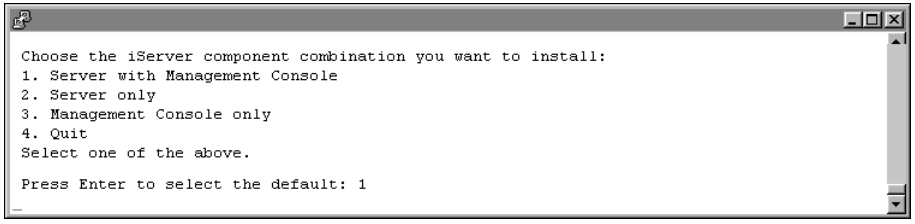

Figure 5-9 Choosing the components to install

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.

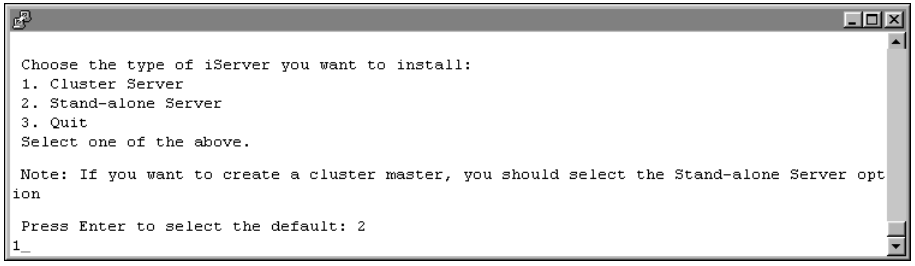

Figure 5-10 Specifying the type of iServer to install

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, acserverlicence.xml, and acconfigowner.lock.

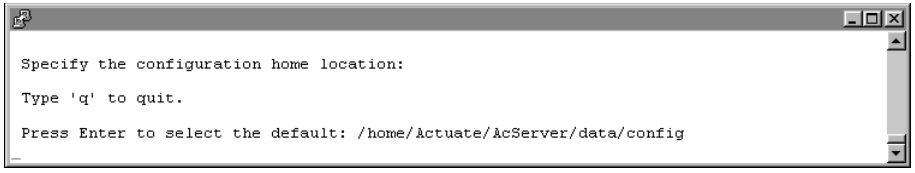

Figure 5-11 Specifying the configuration home location

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.

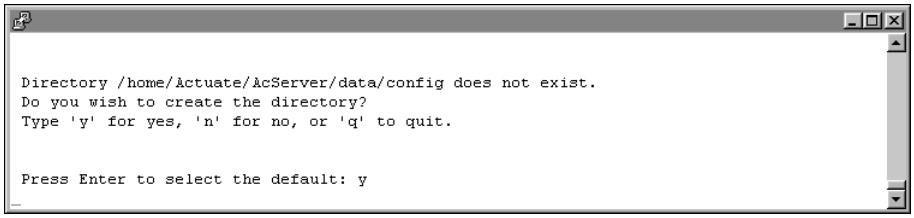

Figure 5-12 Creating the configuration home location

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.

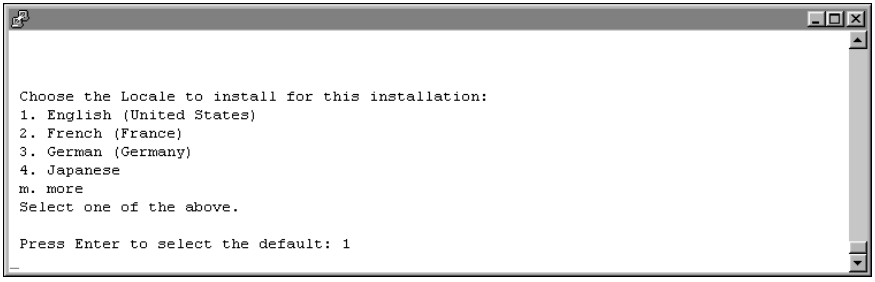

Figure 5-13 Specifying a locale

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.

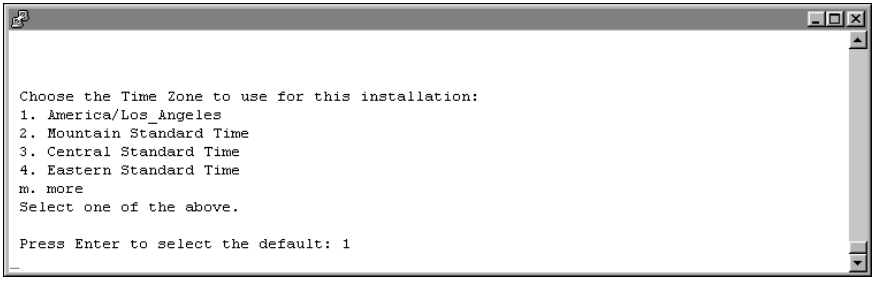

Figure 5-14 Specifying a time zone

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.

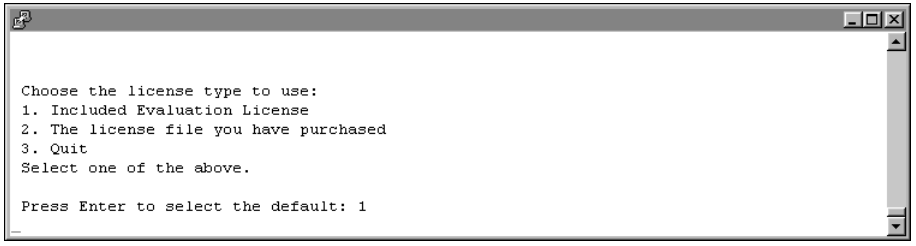

Figure 5-15 Specifying license type

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.

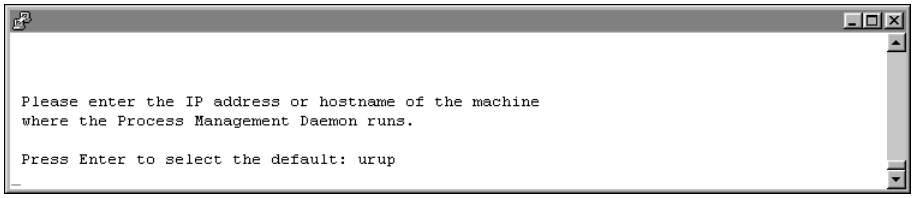

Figure 5-16 Specifying the machine on which the PMD runs

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.

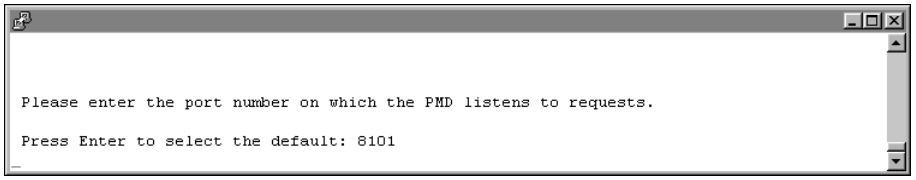

Figure 5-17 Specifying the port number on which the PMD listens

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.

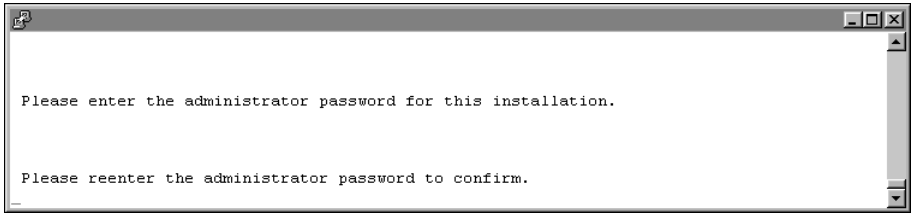

Figure 5-19 Re-entering the administrator password

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.

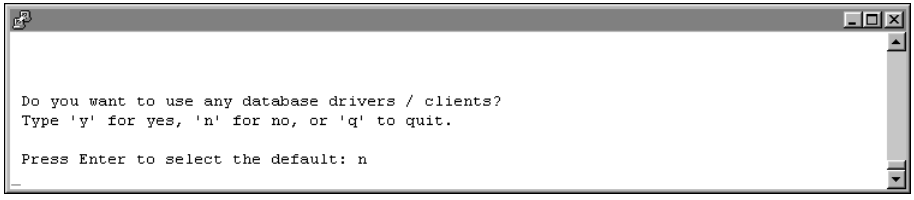

Figure 5-20 Specifying whether to use database drivers/clients

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.

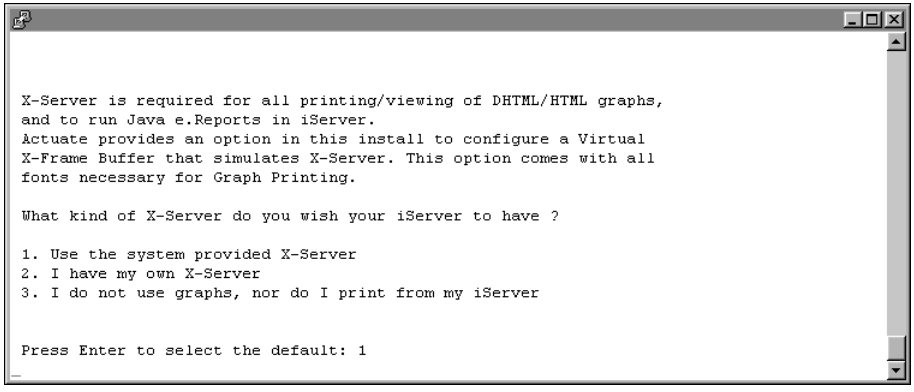

Figure 5-21 Specifying what kind of X-Server to use, if any

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.

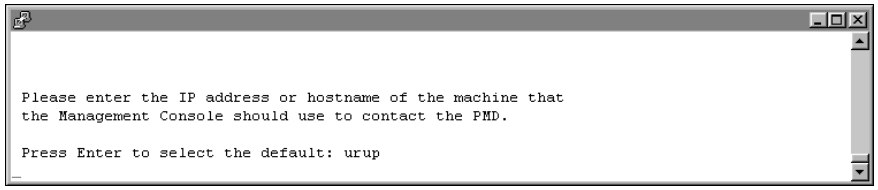

Figure 5-22 Specifying the machine host name that Management Console uses to contact the PMD

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.

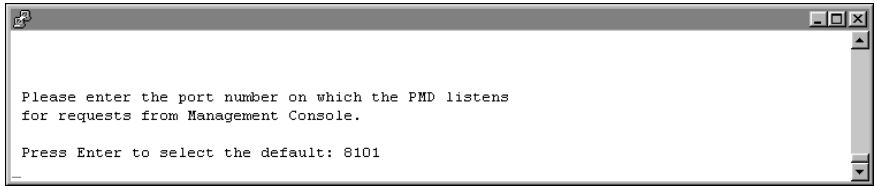

Figure 5-23 Specifying the port number for the PMD to listen for requests from Management Console

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.

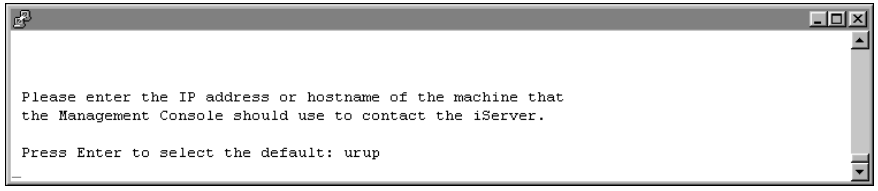

Figure 5-24 Specifying the name Management Console uses to contact iServer

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

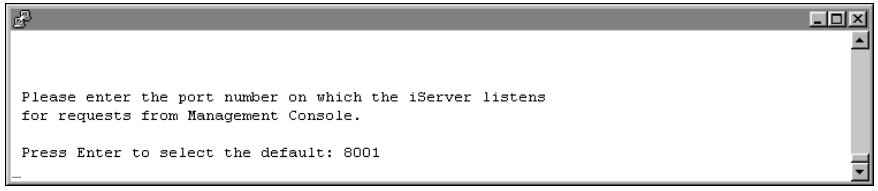

Figure 5-25 Specifying the port number for listening for requests from Management Console

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.

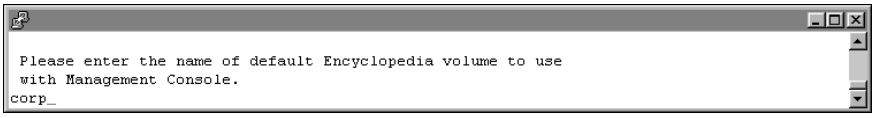

Figure 5-26 Specifying the name of the default Encyclopedia volume

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.

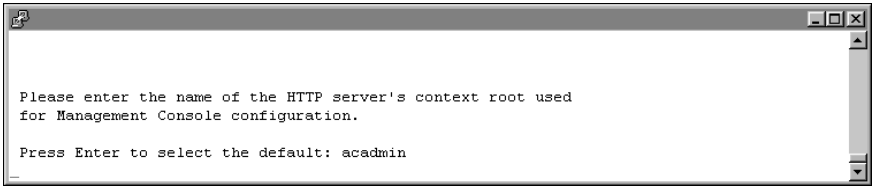

Figure 5-27 Specifying the name of the HTTP server context root

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.

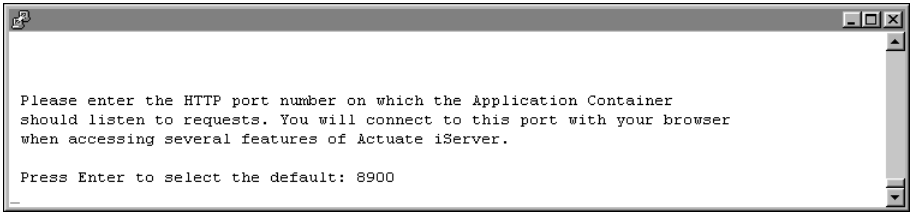

Figure 5-28 Specifying the application container listening port number

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.

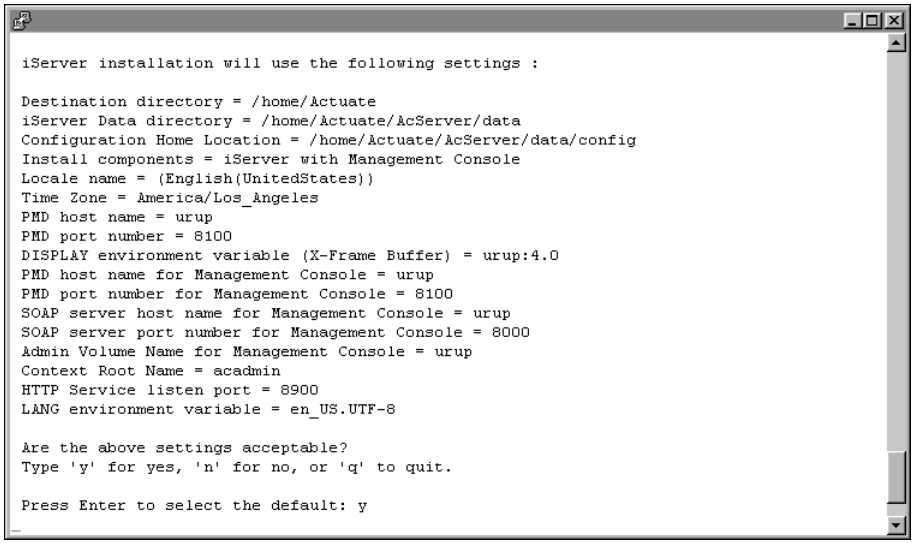

Figure 5-29 Reviewing settings for a Server with Management Console install

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

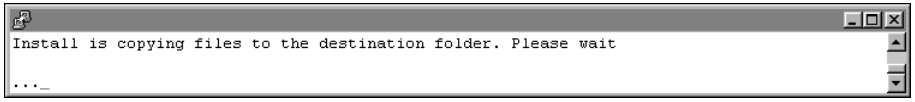

Figure 5-30 Copying iServer files to your destination folder

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.

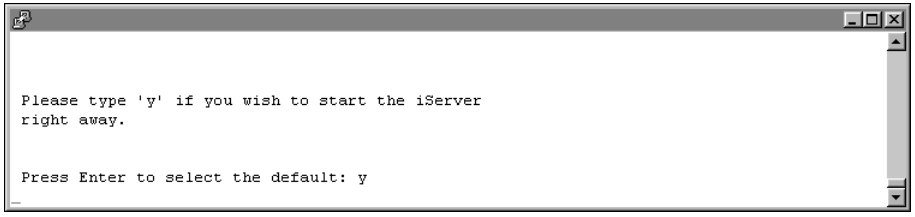

Figure 5-31 Specifying whether to start iServer

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.

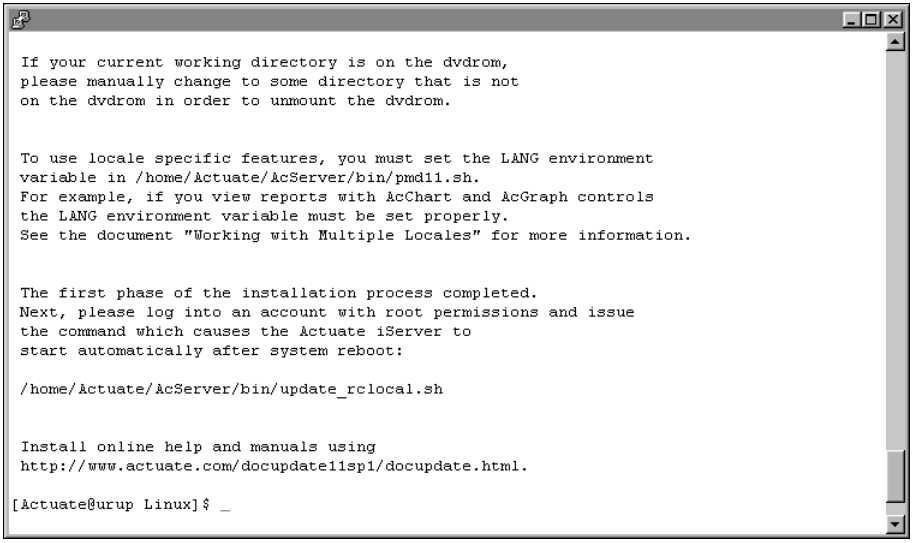

Figure 5-32 Viewing information about localization, logging in, and installing online help

Adding a node to a cluster

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

1. Shut down the node.
2. Open acpmdconfig.xml, by default located in AC_SERVER_HOME/etc.
3. Modify <ConfigHomeDirectory> to point to the shared configuration home directory.
4. Modify <AC_TEMPLATE_NAME> to use the server template from the available server templates in acserverconfig.xml.
5. Save acpmdconfig.xml.
6. Start the new node. The new cluster node automatically contacts the acserverconfig.xml in the shared configuration home directory to access the specified template and join the cluster.
7. To test data sharing between the cluster nodes, add a shared partition and create an Encyclopedia volume.

Finding the BIRT iServer home directory

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

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

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

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

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

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

About the Java Development Kit

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

```
AC_SERVER_HOME/jdk160
```

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

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

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

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

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

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 installing from an FTP download

If you download an Actuate product from the Actuate FTP site, keep all the files together in their original relative locations when you extract them. The installation program verifies that all necessary files are present before beginning the installation process.

If any files are missing, the installation program exits. Files can be missing if you download the DVD image, extract the files, move only some of the files to a new location, and attempt to install from that location.

About performing a full installation

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

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

To access online documentation, such as the online help and PDF files of the product manuals, you must install the documentation files, which are on the documentation DVD that ships with your Actuate Software.

Installing Information Console on Linux and UNIX

This section describes how to install Information Console for Linux and UNIX. Before you begin the installation process, ensure that you have Actuate administrator, system administrator, and web administrator privileges.If you download an Actuate product for Linux and UNIX from the Actuate FTP site, keep all the files together in their original relative locations after you extract them.

Information Console installation requires the following information:

- Installation directory
  $HOME, the account's home directory, is the default installation directory.
- Port used by the Apache Tomcat Information Console service

If you are using a firewall, ensure that the firewall passes the port number you select. The default port is 8900.

- Value of the SPINLOOPTIME environment variable
  If you use AIX in a multiple-CPU environment, Information Console installation sets value of the SPINLOOPTIME environment variable to 2000. For more information about SPINLOOPTIME, see your AIX documentation.
- Encyclopedia volume name that Information Console accesses
  The default is the current machine.

You can install Information Console in following ways:

- Use an installation script.
  The installation script configures Information Console, creates shortcuts, and extracts and installs all necessary files. Use this option for automated configuration.
- Deploy a WAR file to an Application Server.
  Deploying directly requires that you configure Information Console for your application server. Use this option if your application server supports configuration of an application from a WAR file.

Using the script to install

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

How to install using a script

1. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required 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.

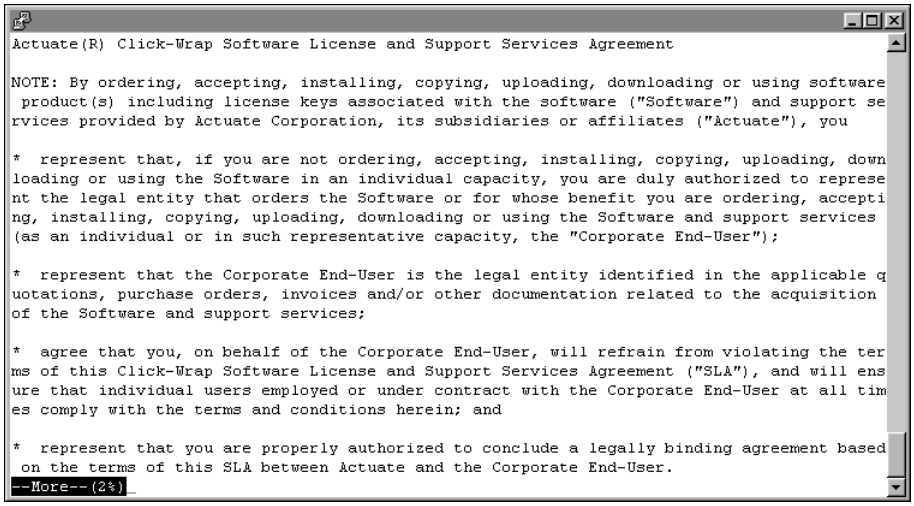

Figure 6-1 The license agreement

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.

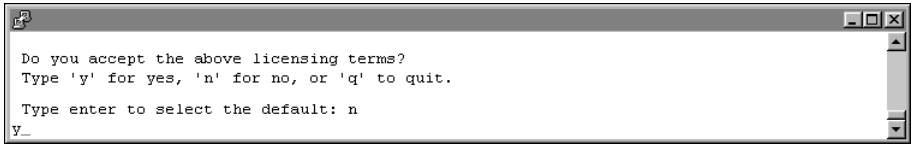

Figure 6-2 Specifying whether you accept the license agreement

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

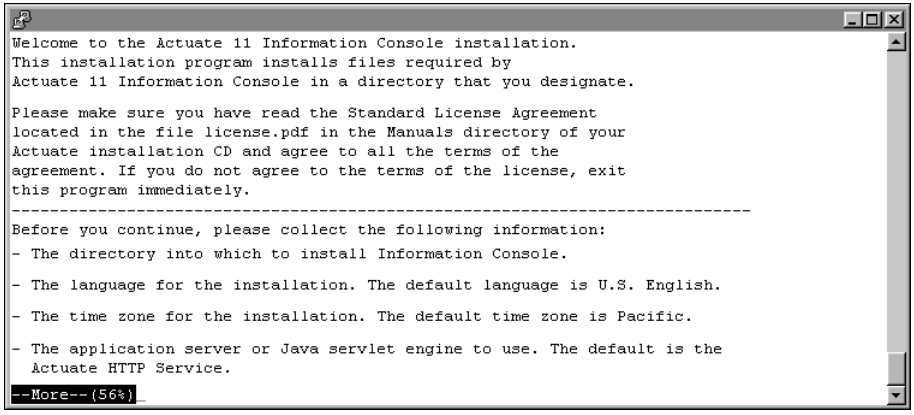

Figure 6-3 Viewing the introduction

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.

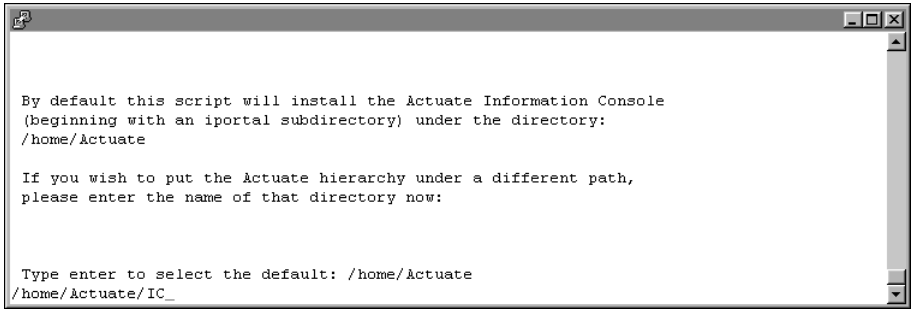

Figure 6-4 Specifying the Information Console install directory

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.

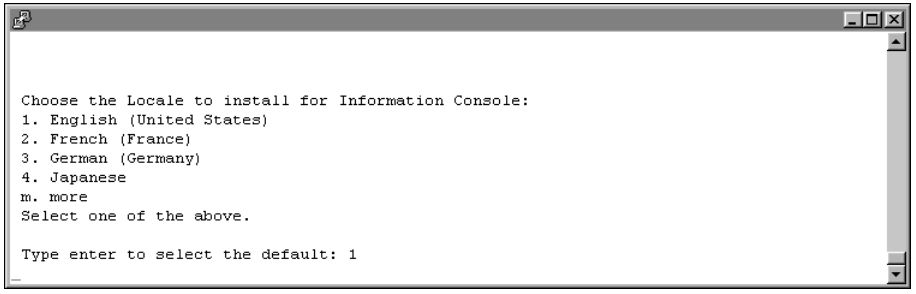

Figure 6-5 Specifying a locale

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.

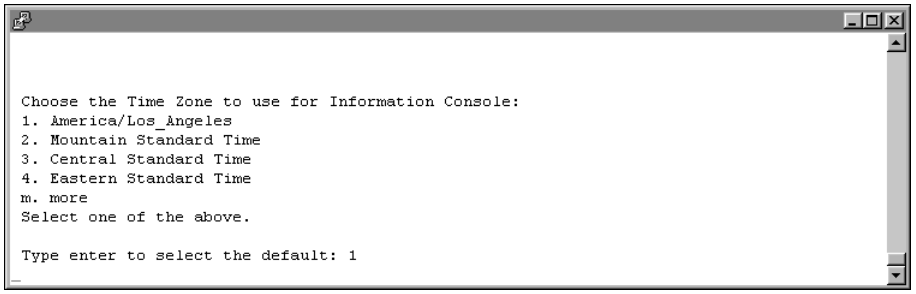

Figure 6-6 Specifying a time zone

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.

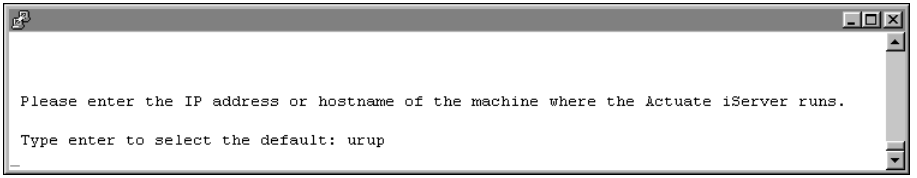

Figure 6-8 Specifying the machine on which the iServer runs

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

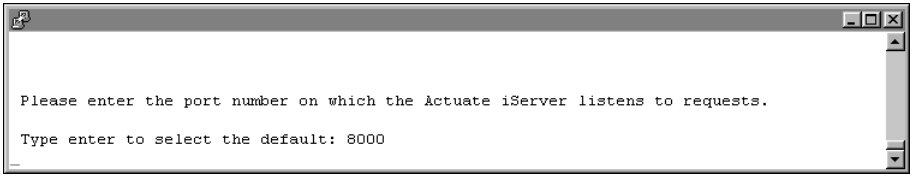

Figure 6-9 Specifying the port number on which iServer listens

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

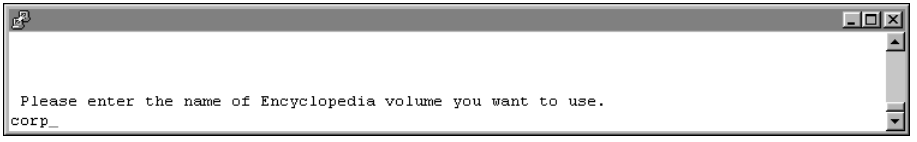

Figure 6-10 Specifying the Encyclopedia volume name

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.

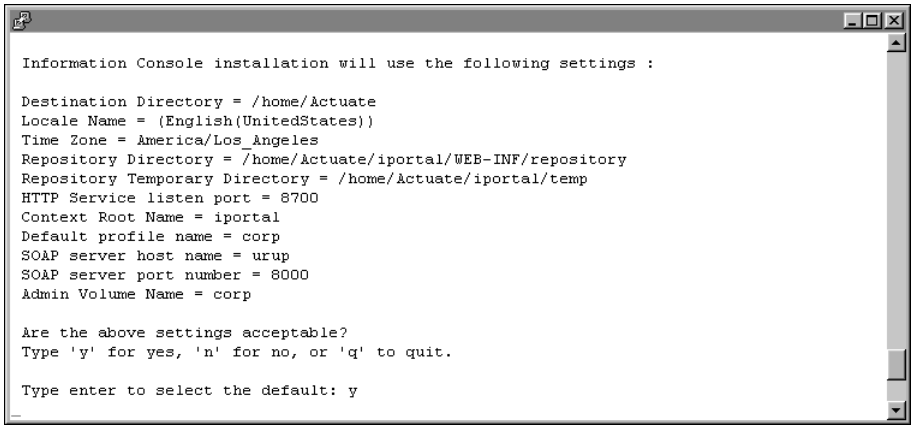

Figure 6-11 Reviewing your settings

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

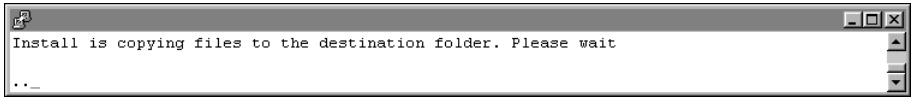

Figure 6-12 The install program copies files to your destination folder

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.

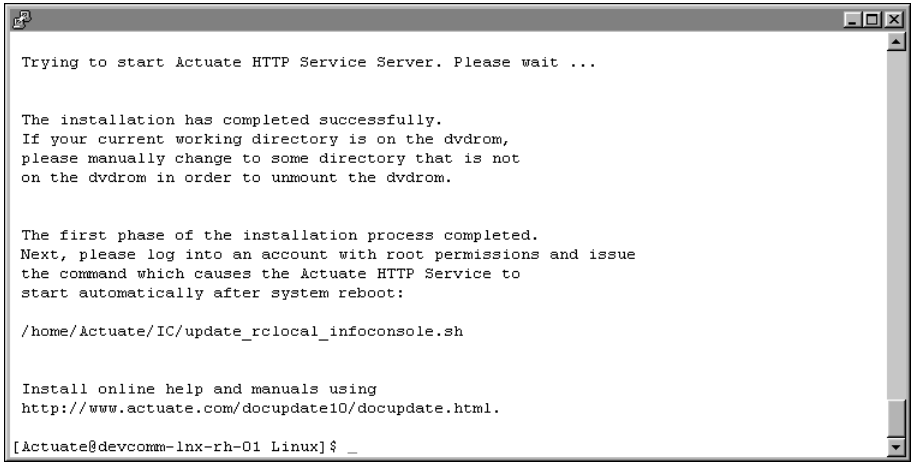

Figure 6-13 Typing the command to start the service

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 on the Actuate Support site at the following URL:

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

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

General deployment tasks

You must configure the Information Console WAR 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."
- 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.
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. Insert the iServer System installation DVD.
2. Create a temporary directory, such as /home/Actuate/ic_temp.

   If you use an existing directory, ensure that this directory is empty.
3. Copy Information Console WAR file to the temporary directory, as shown in the following example:

   ```
   cp WL_TOMCAT_ActuateInformationConsole.war /home/Actuate/
      ic_temp
   ```
4. Decompress the WAR file, as shown in the following example:

   ```
   jar -xf WL_TOMCAT_ActuateInformationConsole.war
   ```

   The Information Console files appear in the temporary directory.
5. 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 3, 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.
6. Save and close web.xml.
7. Type the following command:

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

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

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

About clusters of servers

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

Avoiding cache conflicts after installing

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

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

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

Testing the installation

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

How to test the installation

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

   You can use a URL similar to the following example:

   ```
   http://Actuate1:8900/iportal/getfolderitems.do
      ?repositoryType=Enterprise&volume=volume1
      &serverurl=http://iServer1:8000
   ```

   where

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

   The Information Console login page appears.

   On the Information Console login page:

   - For enterprise mode Information Console, in Volume, select an Encyclopedia Volume.

- In User name, type your login name.
- In Password, type your password.
- In Language, select a locale.
- In Time zone, select a time zone.

4 Choose Log In.

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. If you download an Actuate product for UNIX or Linux from the Actuate FTP site, keep all the files together in their original relative locations after you extract them.

How to install

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

1. Insert the installation DVD, mount a DVD on your system and set your working directory to the mounted DVD directory, or download the required 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.

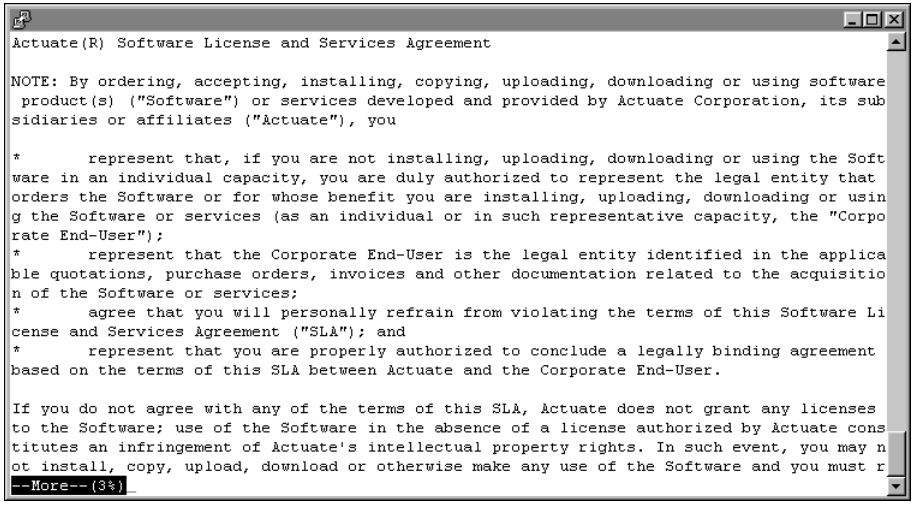

Figure 7-1 Reviewing the license agreement

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.

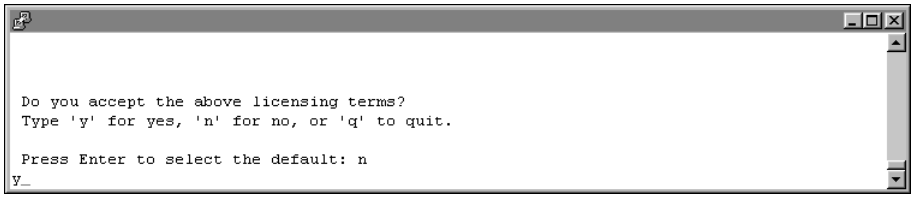

Figure 7-2 Accepting the license agreement

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

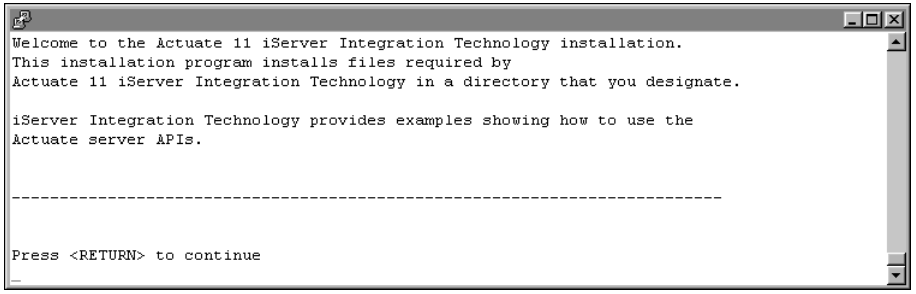

Figure 7-3 Reviewing introductory information

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.

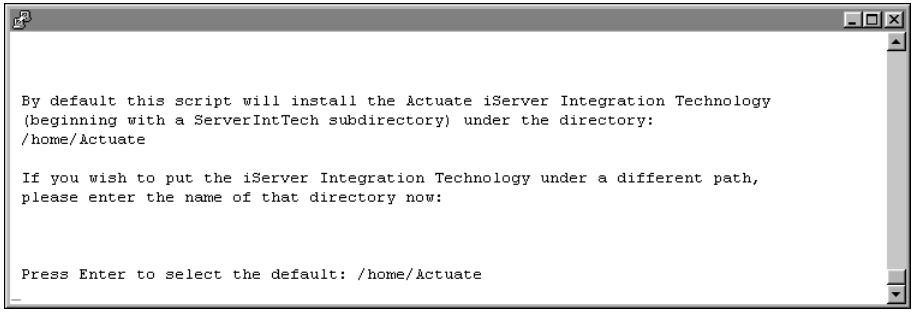

Figure 7-4 Specifying the installation directory

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.

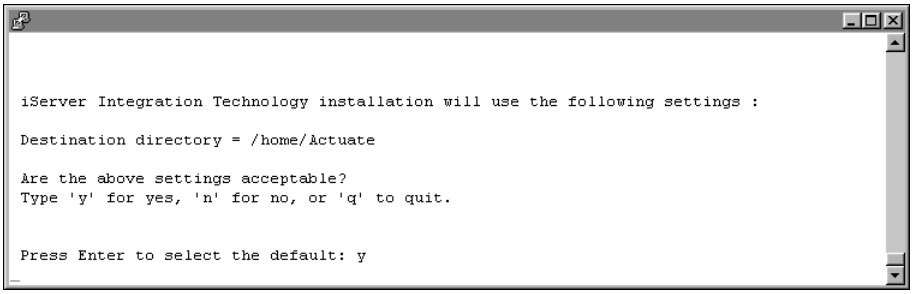

Figure 7-5 Reviewing settings before copying files

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

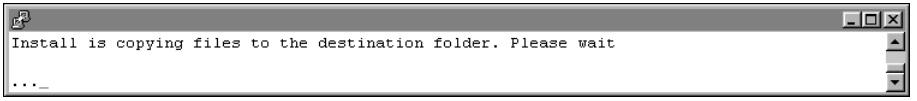

Figure 7-6 Copying files to your destination folder

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

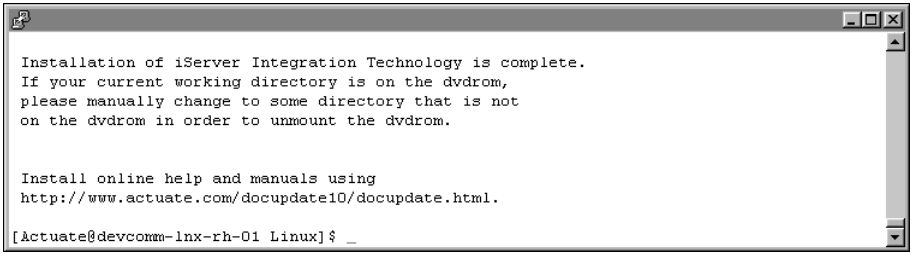

Figure 7-7 Viewing information about changing the working directory and installing online help

Installing the localization and documentation files

The information in the printed manuals is available as Adobe Acrobat PDF files and as a context-sensitive help system for Actuate products. After installing Actuate products, install the localization and documentation files from the Documentation DVD, which ships with Actuate software.

Both localization and documentation resource file updates can become available between releases. The Actuate Localization and Online Documentation Update tool provides replacements and additional files for PDF documentation, context-

sensitive help, and localization of installed Actuate products. The tool is available from the following Actuate web site location:

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

Updates to documentation in PDF form are available at the following Actuate web site locations:

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

```
http://www.actuate.com/docupdate11sp1/docupdate.html
```

If you do not see an update tool for your release on the Support site, no updates exist for the release.

How to install the localization and documentation files

To install the iServer localization and documentation files, perform the following 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.

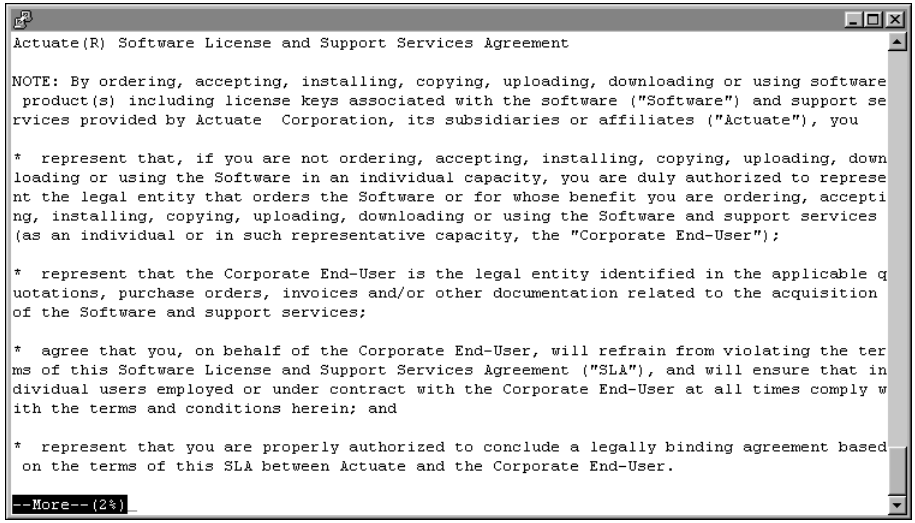

Figure 7-8 Reviewing the license agreement

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.

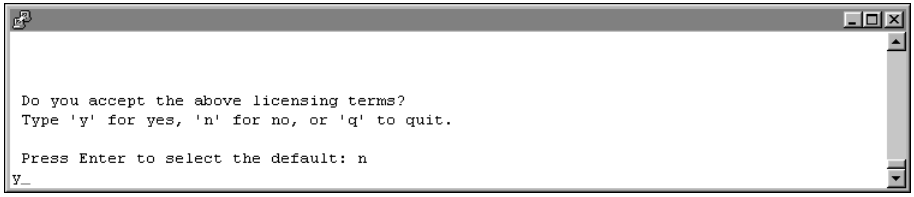

Figure 7-9 Accepting the license agreement

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.

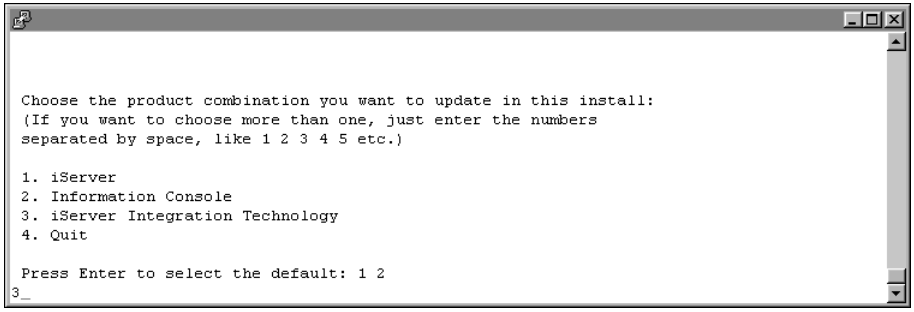

Figure 7-10 Selecting a product combination

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.

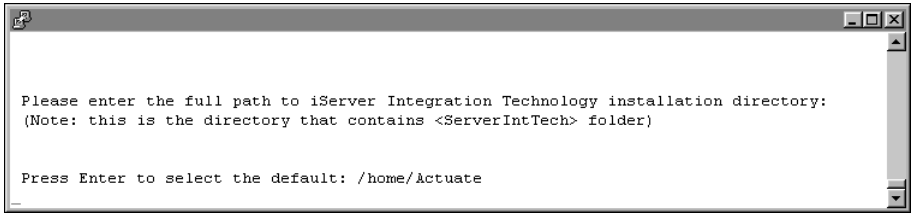

Figure 7-11 Specifying a product directory

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.

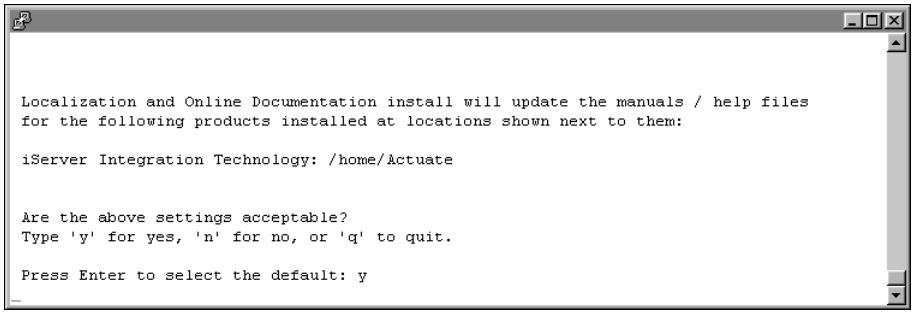

Figure 7-12 Reviewing settings before copying files

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

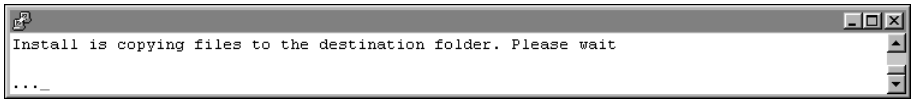

Figure 7-13 Copying files to your destination folder

9 When the installation program finishes, it provides additional information about changing the working directory, as shown in Figure 7-14.

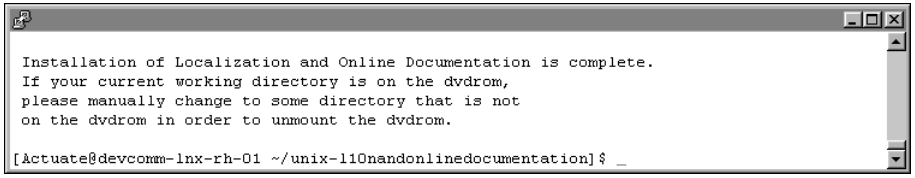

Figure 7-14 Viewing information about changing the working directory

Part Three

Licensing

Chapter

8

Licensing BIRT iServer

This chapter discusses the following topics:

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

Working with licensing

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

The following list describes the types of licenses:

- Named User License

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

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

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

- Platform License

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

- Work Unit (WU) License

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

Understanding the licensing options

Table 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 Analytics	Allows building a cube and displaying a cube report for the purpose of multidimensional analysis. The Analytic Option enables BIRT iServer to extract data from a database and build a compressed OLAP file. When you analyze the cube, you can aggregate or categorize data, summarize data, and create graphs based on data. You can save and share views of the analysis you perform in the Encyclopedia volume.	10, 11
Actuate Query	Supports retrieving information using an information object.	10, 11
BIRT	Allows a user to publish and run a BIRT design using BIRT iServer. This option is a requirement for BIRT Page Level Security Option.	10, 11
BIRT 360	Allows a user to create, execute, and view dashboard files.	11
BIRT Data Analyzer	Allows a user to create, view, and modify cubeview files.	11
BIRT Interactive Viewer	Allows a user who has the BIRT Option to use BIRT Interactive Viewer to view and interact with a BIRT document.	10, 11
BIRT Page Level Security	Controls access to structured content available on the web. This option works for reports created using BIRT Designer Professional and requires the BIRT Option. Access privileges are based on user name or security role.	10, 11
BIRT SmartSheet Security	Controls access to structured content available on the web. This option works for reports created using BIRT Spreadsheet Designer and requires BIRT Spreadsheet Option. Access privileges are based on user name or security role.	10, 11
BIRT Spreadsheet	Allows a user to deploy and run a spreadsheet built using BIRT Spreadsheet Designer. This tool enables customers to save reports as richly formatted Excel spreadsheets and manage them in an Encyclopedia volume.	10, 11
BIRT Studio	Allows a user to create a BIRT design and to run it in BIRT iServer. BIRT Studio Option supports access to an information object on BIRT iServer System.	10, 11

(continues)

Table 8-1 BIRT iServer System license options (continued)

Option	Description	Supported releases
Data Integration	Supports using an information object to combine data from two or more disparate sources. This option is required where data joining is not supported through separately licensed third-party databases or other separately licensed technology. BIRT iServer System Enterprise Information Integration (EII) services provide a standardized way to access data from different data sources.	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 Page Level Security	Controls access to structured content available on the web. This option works for reports created using Actuate e.Report Designer Professional and requires the e.Report (Actuate Basic Report) Option. Access privileges are based on user name or security role.	10, 11
Information Object Caching	Provides the ability to cache data from an information object in a third-party database. This option requires separate licensing of a third party database server for data storage. This option is not available for an information object based on Actuate Basic technology.	10, 11
Multi-Tenant	Allows a BIRT iServer System user to access more than one Encyclopedia volume. This option is available with an Unlimited User CPU License.	11

To determine the license options installed on iServer, log in to Configuration Console, and choose Show License. The license options appear, as shown in Figure 8-1.

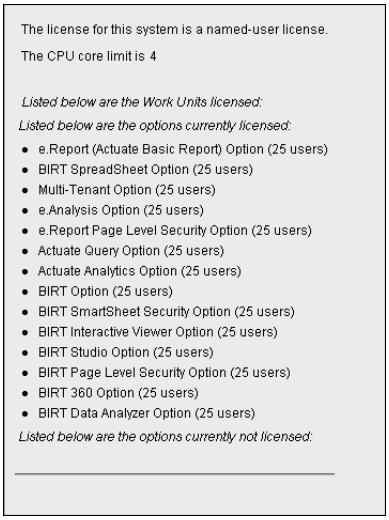

Figure 8-1 iServer License options

About a license file

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

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

This license key file specifies the available iServer license options and node-key information for the cluster nodes. This license file must be in a shared location, specified by the ConfigHomeDirectory variable in the acpmdconfig.xml file of each node, and accessible to all nodes in the cluster.

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

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

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

About CPU binding

BIRT iServer System supports CPU-binding on a machine with an appropriate CPU-based license. CPU-binding restricts a process or processes to run on a subset of CPUs. If you bind the BIRT iServer System to a subset of CPUs, only those CPUs count toward the total number of licensed CPUs. The CPU limit in the license file applies to all CPUs for all machines in the cluster.

The following sections provide a detailed description of how the license key process works in Actuate Releases 10 and 11.

Understanding node-key licensing

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

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

```
Actuate_iServer_key_xxxxx.xml
```

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

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

Obtaining the installation license key file

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

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

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

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

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

A license key file can be an expiring file that is valid until a specific date. If your license key file is an expiring file, you are reminded that the file expires on a certain date when you log in to the Configuration or Management Consoles. Reminders also appear in the system log file. To arrange for a permanent license key file, or if you have a problem with an expiring file, please contact Actuate Licensing at licensing@actuate.com.

You can easily modify your decisions about Actuate license options by contacting Actuate Licensing. If you are an Actuate international customer, please be aware that the an e-mail message sent to Actuate goes to Actuate headquarters, and we route your request to a team in the appropriate country.

Collecting machine information for a node-key license

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

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

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

The machine id required for the node key file can also be
   generated by using the acmachineid utility that can be found in
   the AC_SERVER_HOME/bin folder.

Press Enter to continue.
```

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

Make a note of the machine ID in the installation prompt and send it to Actuate Licensing. Actuate Licensing processes your request and sends two new license keys, the BIRT iServer System and node-key license files.

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

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

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

Installing the license key

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

The following message about expiration of the initial license key always appears on the login screen regardless of the node-key license status:

```
Reminder
   Your BIRT iServer license expires in [the number of days] days,
   on [the specified date]. When the current license expires, the
   iServer will require a new license to restart. Please contact
   Actuate to purchase a new license.
```

The following message about how to obtain the second set of license keys from Actuate Licensing appears until you install the new license keys issued by Actuate Licensing:

```
Reminder
   One or more iServers in your BIRT iServer System are in
   violation of the node locked BIRT iServer license. After the
   grace period expires the iServers that violate the node locked
   BIRT iServer license cannot be restarted. Please contact
   Actuate Licensing (licensing@actuate.com or
   http://www.actuate.com/licensing), or your representative, and
   request a new license file for the iServer nodes that are in
   violation. Please restart the iServers on the nodes after
   updating the node key file.
```

You have 45 days to apply for and install the node-key license after you install BIRT iServer System.

How to install the license key file

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

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

An Actuate license key is an XML file. Actuate Licensing sends this XML file to you with an appended .txt file extension because transmitting a file with an .xml extension can cause problems in an e-mail system. You must remove the .txt extension from the file name before installing the license key file in the BIRT iServer System. Make sure that the file name contains the unique five-digit number generated by Actuate Licensing.

2. Copy the Actuate_iServer_key_XXXXX.xml license file to the location of the temporary license file.
3. Log in to Configuration Console. For example, type http://localhost:8900 /acadmin/config/ in Address of a browser, and use the system configuration password you specified during installation.
4. Choose Update License.
5. Navigate to the location of the new license key file and select the file. Choose OK.
6. Restart any node where the node-key configuration changed.

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

Understanding CPU binding

CPU binding can work only on a machine with an appropriate CPU-based license. CPU binding restricts a process or processes to run on a subset of CPUs. If you bind the BIRT iServer System to a subset of CPUs, only those CPUs count toward the total number of licensed CPUs. Depending on the operating system and specific system command, you can restrict other processes from running on the processor to which you bind a process.

You can bind BIRT iServer processes to a specific set of processors on a machine that runs a Windows, Sun Solaris, or HP-UX 11i operating system. The default configuration does not bind BIRT iServer to a set of processors. In the default configuration, all processors on 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 into Configuration Console, and choosing System—Properties—Advanced—Cluster Operation—Administrative.

   For example, the following e-mail address sends e-mail to a user named admin at a company for which the domain is mycompany:

   ```
   admin@mycompany.com
   ```

3. Restart BIRT iServer System. Restarting applies the changes after you set or change the e-mail address.

Part Four

Backing Up

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

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 in disk partitions, and the acserverconfig.xml file containing iServer configuration information

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. It is necessary to back up all Encyclopedia volume metadata and data to ensure the recoverability of the volume in the event of failure.

The third-party database that contains Actuate Encyclopedia metadata is a critical component of Actuate iServer System. An Actuate system administrator must take all necessary precautions to ensure that this database is properly backed up and available to safeguard Encyclopedia volume metadata. Please consult Actuate Support at the time of installation if you have any questions about the backup, recovery, or failover procedures necessary to protect against the possibility of catastrophic failure.

Managing the backup and recovery of an Encyclopedia volume database 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 iServer disk partitions containing file data
- A copy of the acserverconfig.xml file containing iServer configuration information

Back up the Encyclopedia volume metadata in the RDBMS at the same time that you back up the disk partition data files. A carefully coordinated backup ensures that a one-to-one correspondence exists between each entry in the volume metadata database and the data files.

The Encyclopedia volume metadata backup on the RDBMS must be done before the backup of the data on the disk partitions. Files that are partially created when the metadata backup begins are either not yet registered in the database or are marked incomplete in the database. The metadata database does not retain a record of incomplete files.

When contacting Actuate Support to troubleshoot problems, it is best to provide a snapshot of the Encyclopedia volume configuration, including the following items and information:

- A database backup of the Encyclopedia volume schema containing the metadata
- The name of the Encyclopedia volume schema and user that iServer uses to connect to the RDBMS
- A copy of the acserverconfig.xml file containing iServer configuration information
- A copy of the iServer logs

Using RDBMS and file system backup utilities

The administrator must perform the Encyclopedia volume metadata backup using the tools provided or supported by the RDBMS. Copying the physical files of a database at the operating system level while an RDBMS is running does not create a valid backup.

Most RDBMS backup tools can be scripted and run while iServer is using the database. PostgreSQL and Oracle also provide graphical administration tools in addition to command-line tools. For more information on using these RDBMS tools, see "Backing up and restoring an Encyclopedia volume that uses a PostgreSQL database" and "Backing up and restoring an Encyclopedia volume that uses an Oracle database," later in this chapter.

How to perform an Encyclopedia volume backup

To back up an Encyclopedia volume, perform the following tasks:

1. Make sure that the autoarchive file purging process is not running.
2. Make an online backup of the volume schema using the tools provided by the RDBMS.
3. Backup the volume data files using the tools available in the operating system environment.

Avoiding conflict with the autoarchive file purging process

A metadata backup is consistent with a data backup only if the file purging process run 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.

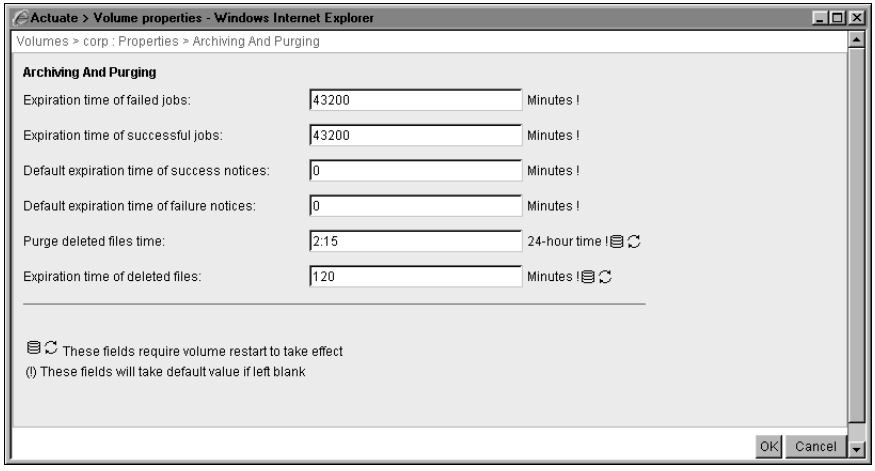

Figure 9-1 Configuring file purging properties

Choose OK.

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

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

PostgreSQL provides the 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.

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 backup each 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
```

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
```

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
```

Note that a backup of a PostgreSQL database is not portable across all operating systems.

The following sections describe how to backup and restore an Encyclopedia volume that uses the OOTB PostgreSQL database to store the metadata. This demonstration serves as a detailed reference example. Other RDBMS environments, such as a pre-existing PostgreSQL or Oracle RDBMS require similar procedures.

For more information on how to administer, back up, and restore a PostgreSQL RDBMS, see the vendor documentation. For more information on how to administer, back up, and restore an Oracle RDBMS, see "Backing up and restoring an Encyclopedia volume that uses an Oracle database," later in this chapter, and the vendor documentation.

Backing up and restoring an Encyclopedia volume

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 pg_dump PostgreSQL command-line utility
- Backs up Encyclopedia volume data and configuration files using operating system copy commands

To restore an Encyclopedia volume in the OOTB PostgreSQL RDBMS environment, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the pg_restore PostgreSQL command-line utility
- Restores Encyclopedia volume data and configuration files using operating system copy commands

The following sections show how to perform these backup and restore operations.

Backing up an Encyclopedia volume using pgAdmin

To back up an Encyclopedia volume, perform the following tasks:

- Create a folder to contain the metadata and volume data backup files.
- Backup Encyclopedia volume metadata using the PostgreSQL pg_dump utility.
- Stop Actuate 11 BIRT iServer service.
- Back up the acserverconfig.xml file and volume data folders to the backup folder
- Start Actuate 11 BIRT iServer service.

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`

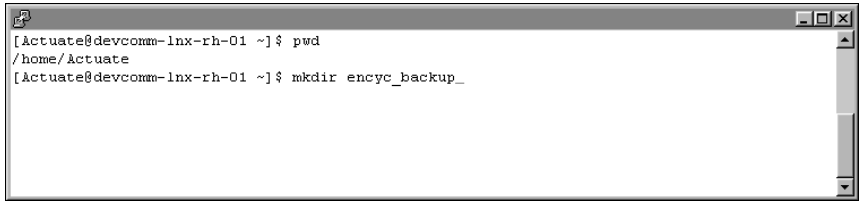

Figure 9-2 Creating a backup folder

Backup 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
```

3 The command line appears as shown in Figure 9-3.

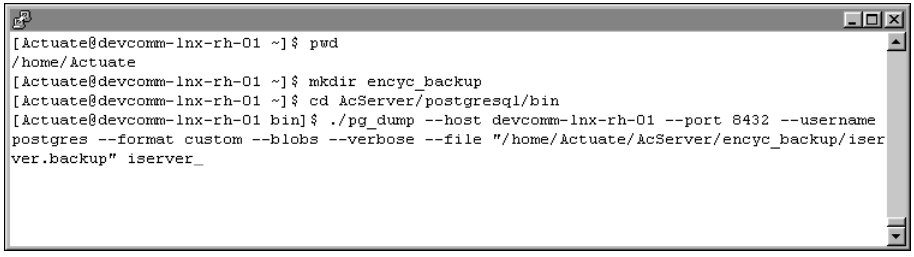

Figure 9-3 Entering the command to execute pg_dump

4 Type the postgres superuser password. The installer specifies this password during the iServer installation procedure.

pg_dump executes, writing status messages to the command prompt.

After backing up the Encyclopedia volume metadata, stop the Actuate 11 BIRT iServer service by performing the following tasks before backing up the volume data folders.

How to stop the Actuate 11 BIRT iServer service

1 Navigate to the following location:

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

2 Execute the following command, as shown in Figure 9-4:

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

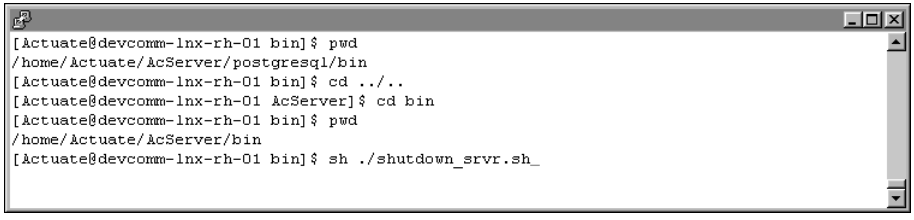

Figure 9-4 Stopping the BIRT iServer service

Confirm that you want to shutdown iServer by pressing enter. The command completes, and the iServer service stops, as shown in Figure 9-5.

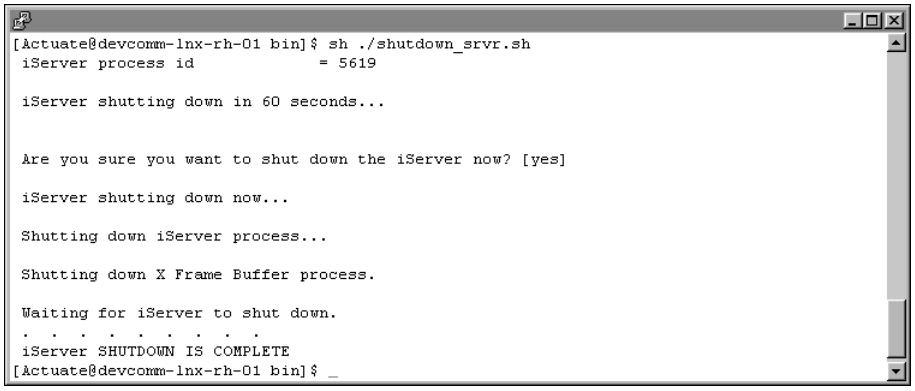

Figure 9-5 Confirming the iServer service shutdown

Back up the volume data folders 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 directory.

   Copy acserverconfig.xml to the following backup location, as shown in Figure 9-6:

   `/home/Actuate/encyc_backup`

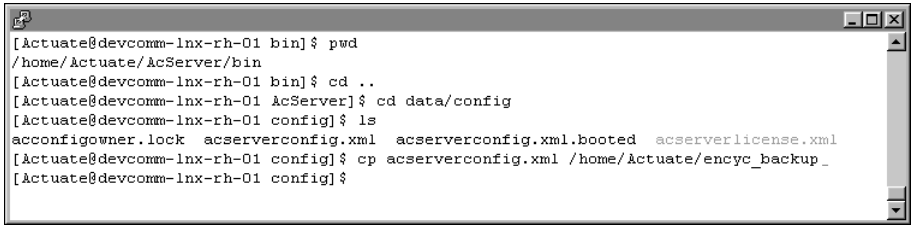

Figure 9-6 Copying acserverconfig.xml to the backup location

3. Navigate to AC_SERVER-HOME/encyc.

   Copy the file, fileType, status, and tempRov folders to the following backup location, as shown in Figure 9-7:

   `/home/Actuate/encyc_backup`

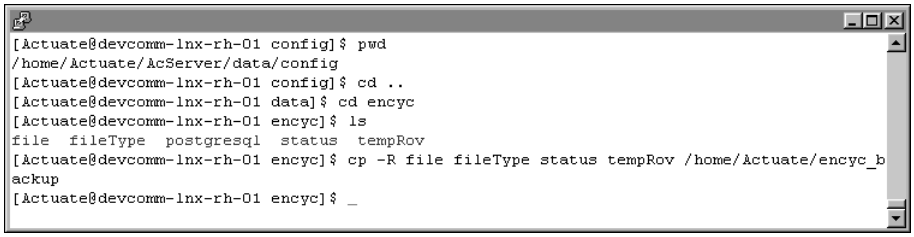

Figure 9-7 Copying the volume data folders to the backup location

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 9-8.

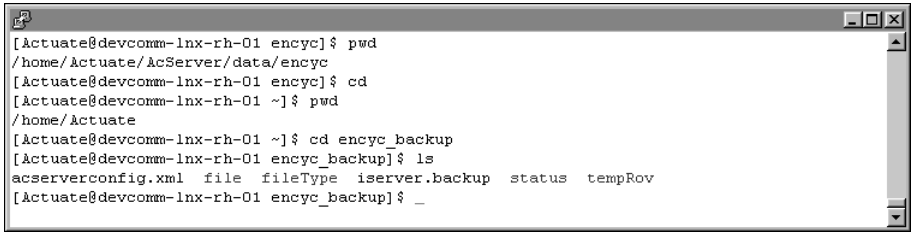

Figure 9-8 Viewing the contents of the backup folder

Next, start the Actuate 11 BIRT iServer service by performing the following tasks.

How to start the Actuate 11 BIRT iServer service

1. Navigate to the following location:

   `/home/Actuate/AcServer/bin`

2. Execute the following command, as shown in Figure 9-9. Press enter to return to the command prompt:

   `sh ./start_srvr.sh`

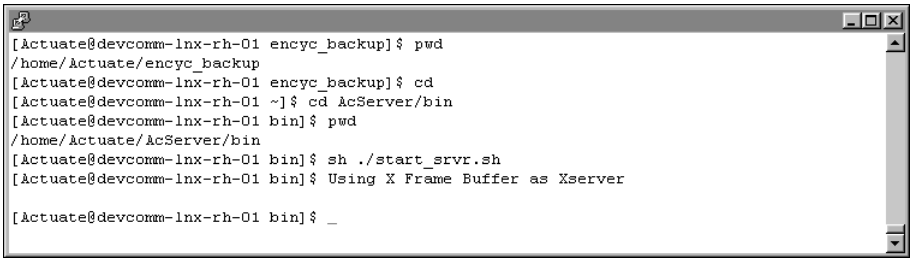

Figure 9-9 Starting the BIRT iServer service

Restoring an Encyclopedia volume

To restore a backed-up Encyclopedia volume, perform the following tasks:

- Stop Actuate 11 BIRT iServer service.
- 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.
- Restart Actuate 11 BIRT iServer service.

How to stop the Actuate 11 BIRT iServer service

1. Navigate to the following location:

   `/home/Actuate/AcServer/bin`

2. Execute the following command, as shown in Figure 9-10:

   `sh ./shutdown_srvr.sh`

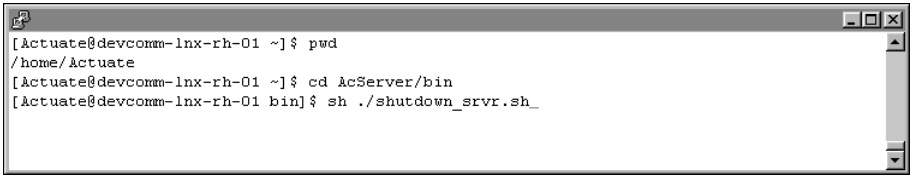

Figure 9-10 Stopping the BIRT iServer service

Confirm that you want to shutdown iServer by pressing enter. The command completes, and the iServer service stops, as shown in Figure 9-11.

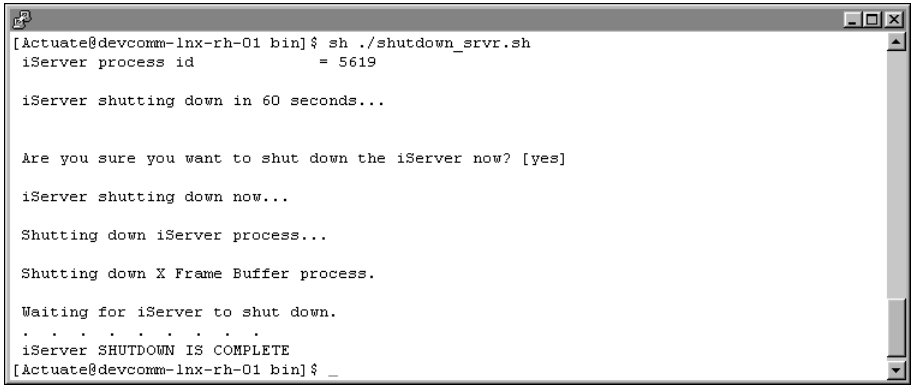

Figure 9-11 Confirming the iServer service shutdown

How to restore the backed up volume data folders

1 Navigate to AC_DATA_HOME/config.

Delete acserverconfig.xml, as shown in Figure 9-12.

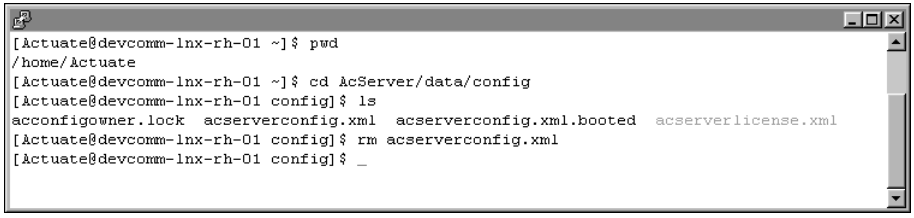

Figure 9-12 Deleting acserverconfig.xml

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 9-13.

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. Do not select the postgresql folder.

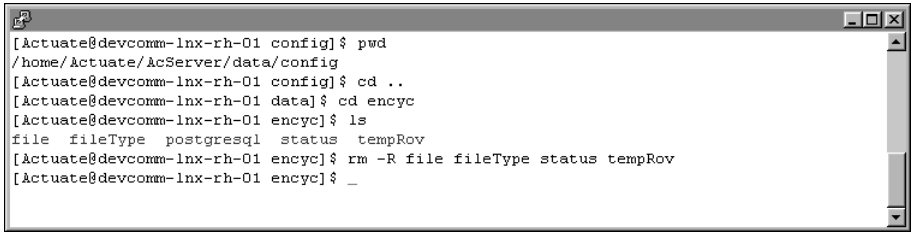

Figure 9-13 Deleting the file, filetype, status, and tempRov folders

3 Navigate to the following location:

```
/home/Actuate/encyc_backup
```

Copy acserverconfig.xml to AC_DATA_HOME/config, as shown in Figure 9-14.

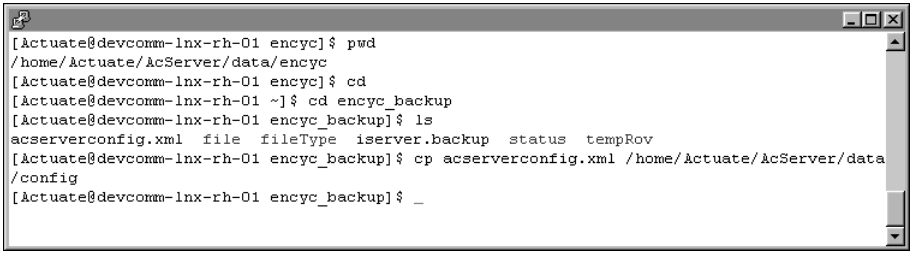

Figure 9-14 Copying acserverconfig.xml to AC_DATA_HOME/config

4 Copy the file, fileType, status, and tempRov folders to AC_DATA_HOME/encyc, as shown in Figure 9-15.

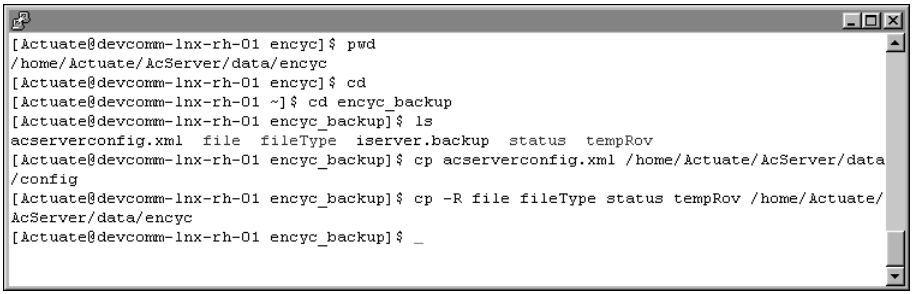

Figure 9-15 Copying the volume data folders to AC_DATA_HOME/encyc

The contents of AC_DATA_HOME/config appear as shown in Figure 9-16.

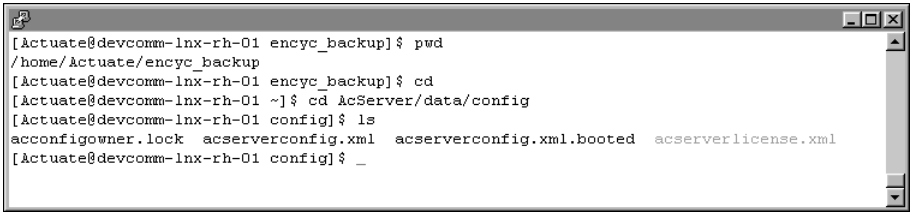

Figure 9-16 Viewing the contents of AC_DATA_HOME/config

The contents of AC_DATA_HOME/encyc appear as shown in Figure 9-17.

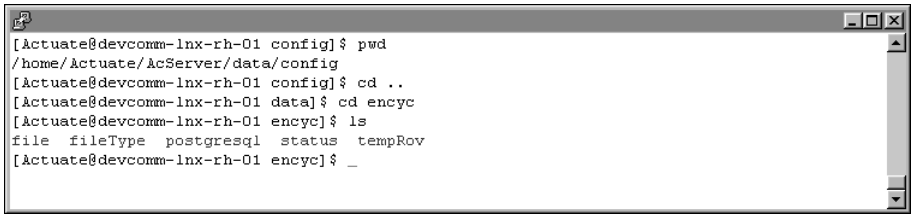

Figure 9-17 Viewing the contents of AC_DATA_HOME/encyc

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 9-18.

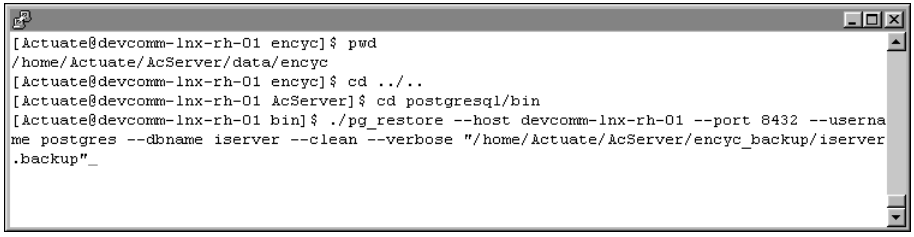

Figure 9-18 Entering the command to execute pg_restore

Next, start the Actuate 11 BIRT iServer service by performing the following tasks.

How to start the Actuate 11 BIRT iServer service

1. Navigate to the following location:

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

2. Execute the following command, as shown in Figure 9-19. Press Enter to return to the command prompt:

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

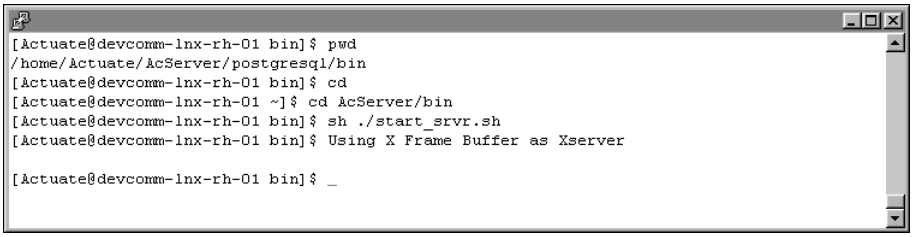

Figure 9-19 Starting the BIRT iServer service

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.

Backing up and restoring an Encyclopedia volume

To back up an Encyclopedia volume stored in an Oracle RDBMS, the administrator performs the following operations:

- Backs up Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) command-line utility
- Backs up Encyclopedia volume data and configuration file using operating system copy commands

To restore an Encyclopedia volume in an Oracle RDBMS, the administrator performs the following operations:

- Restores Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line utility
- Restores Encyclopedia volume data and configuration file using operating system copy commands

The following sections provide more information on how to perform these backup and restore operations.

Backing up an Encyclopedia volume using Oracle Data Pump Export (expdp)

Oracle Data Pump utilities require the database administrator to map a database dump directory to a physical directory in the file system. The Oracle RDBMS

writes to and reads from this directory when performing export and import operations.

Create a directory to contain the metadata and volume data backup files outside the iServer data installation environment. To provide protection against single-point media failure, it is best to store the backup files on a partition that is physically separate from the Encyclopedia volume data location.

To back up an Encyclopedia volume metadata using expdp, perform the following tasks:

- Create a directory to contain the metadata and volume data backup files using sqlplus.
- Backup Encyclopedia volume metadata using the Oracle pg_dump utility.
- Stop Actuate 11 BIRT iServer service.
- Back up the acserverconfig.xml file and volume data directories to the backup folder.
- Start the Actuate 11 BIRT iServer service.

How to create a back up of the Encyclopedia volume metadata

1. In Linux, open a command window.
2. Using sqlplus, an Oracle SQL editing tool, 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. Backup the Encyclopedia volume metadata using the Oracle Data Pump Export (expdp) utility, as shown in the following example:

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

   The expdp example uses the following arguments:

   - system/password@db_host/dbname.actuate.com

Specifies the system administrator and password, the host name of the machine where the Oracle server is running, and the full database domain name

- SCHEMAS
  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, stop the Actuate 11 BIRT iServer service before backing up the volume data folders by performing the following tasks.

How to stop the Actuate 11 BIRT iServer service

1. Navigate to the following location:

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

2. Execute the following command:

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

   Confirm that you want to shutdown iServer by pressing enter. The command completes, and the iServer service stops.

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. You 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.

   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.

Next, start the Actuate 11 BIRT iServer service by performing the following tasks.

How to start the Actuate 11 BIRT iServer service

1. Navigate to the following location:

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

2. Execute the following command:

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

   Press Enter to return to the command prompt.

Restoring an Encyclopedia volume using Oracle Data Pump Import (impdp)

To restore a backed-up Encyclopedia volume, perform the following tasks:

- Stop Actuate 11 BIRT iServer service.
- 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.
- Restart Actuate 11 BIRT iServer service.

How to restore the backed up volume data directories

1. In Linux, open a command window.
2. Stop Actuate 11 BIRT iServer Service.

   Wait for the service to stop before continuing to the next step. Leave the Services window open.

3. Navigate to AC_DATA_HOME/config directory and delete the acserverconfig.xml file.
4. Navigate to AC_DATA_HOME/encyc directory and delete the file, and 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.

5. Navigate to the following backup directory location:

   ```
   /home/Actuate/encyc_backup
   ```

6 From the backup directory location, perform the following tasks:

   1 Copy acserverconfig.xml to AC_DATA_HOME/config.

   2 Copy the file, filetype, status, and tempROV directories to AC_DATA_HOME/encyc.

Restore the Encyclopedia volume metadata using the Oracle impdp utility by performing the following tasks.

How to restore a back up of the Encyclopedia volume metadata

1 Restore the Encyclopedia volume metadata using the Oracle Data Pump Import (impdp) command-line 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.

2 Restart the Actuate 11 BIRT iServer service.

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
```

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