



IBM Informix Storage Manager Administrator's Guide



IBM Informix Storage Manager Administrator's Guide

Note!

Before using this information and the product it supports, read the information in “Notices” on page D-1.

First Edition (December 2004)

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1996, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	vii
About This Manual	vii
Types of Users	vii
Software Dependencies	viii
Assumptions About Your Locale	viii
Demonstration Database	ix
New Features in ISM and Organizational Changes	ix
Documentation Conventions	x
Typographical Conventions	x
Feature, Product, and Platform	xi
Syntax Diagrams	xii
Example Code Conventions	xv
Additional Documentation	xvi
Installation Guides	xvi
Online Notes	xvi
Informix Error Messages	xviii
Manuals	xix
Online Help	xix
Accessibility	xix
IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set	xix
Compliance with Industry Standards	xxii
IBM Welcomes Your Comments	xxiii
Chapter 1. Informix Storage Manager	1-1
What Is ISM?	1-1
ISM Server Tasks	1-2
Y2K Compliance	1-3
How Do ON-Bar and ISM Fit Together?	1-3
What Happens During a Backup Operation?	1-3
What Happens During a Restore Operation?	1-5
Configuring ISM	1-6
Setting Up ISM on UNIX	1-8
Setting Up ISM on Windows	1-9
Setting Up ISM Volume Pool Names	1-10
Setting ISM Environment Variables and ONCONFIG Parameters	1-10
Updating the onbar Script	1-11
Uninstalling ISM	1-11
Using ISM with DNS Disabled	1-13
Migrating ISM 1.0 to ISM 2.2	1-14
Reverting from ISM 2.2 to ISM 1.0	1-17
Using ISM	1-17
Using ISM Administrative Privileges	1-18
Starting and Stopping ISM	1-18
Reviewing ISM Logs	1-18

Chapter 2. Configuring ISM Server Properties	2-1
Setting the Number of Parallel Data Streams.	2-1
Setting Target Sessions for a New Device	2-2
Setting the Retention Period for All Data	2-3
Adding, Changing, or Deleting Administrative Users	2-4
Adding the First Administrative User to the ISM Server	2-4
Choosing Acceptable ISM User and Group Names	2-5
Deleting an Administrative User	2-6
Connecting to a Different ISM Server	2-6
Using ISM Configuration Parameters to Change the Volume Pool Name	2-6
Using ISM Environment Variables	2-8
Using ON-Bar Environment Variables with ISM.	2-9
Chapter 3. Managing Storage Devices	3-1
Choosing and Configuring Storage Devices	3-2
Adding and Deleting Storage Devices	3-3
Enabling and Disabling a Storage Device	3-4
Changing Storage-Device Properties	3-5
Designating a Storage Device as Read-only	3-5
Using Tape Drives and Optical Disk Drives as Storage Devices	3-5
Using Filesystems as Storage Devices	3-6
Using Storage Devices for Logical-Log Backups.	3-7
Reserving a Storage Device for Cloning	3-7
Automating Media Management for Tape and Optical Devices.	3-9
Enabling Automatic Media Management	3-9
Selecting a Storage Volume	3-9
Relabeling Volumes Automatically.	3-10
Using Unrecognizable Volumes.	3-10
Chapter 4. Managing Storage Volumes	4-1
Managing the Data Life Cycle	4-1
How ISM Uses the Retention Period	4-2
How ISM Uses Volume Pools	4-5
Managing Storage Volumes	4-8
Labeling a Volume	4-8
Mounting a Volume	4-9
Choosing Volumes for Mounting.	4-9
Unmounting Volumes	4-10
Finding a Volume Name	4-10
Relabeling a Volume	4-10
Overriding Automatic Media Management	4-12
Disabling a Storage Volume	4-12
Protecting a Storage Volume from Automatic Recycling.	4-13
Deleting Volumes from the ISM Catalog	4-13
Cloning Storage Volumes and Save Sets	4-14
Tracking Cloned Data	4-15
Cloning Storage Volumes	4-15
Cloning Save Sets	4-16
Chapter 5. Recovering Save Sets with ISM	5-1

Restoring a Save Set with ON-Bar	5-1
Locating Volumes That Contain Partial Save Sets	5-1
Re-Creating the ISM Catalog Before You Use ON-Bar	5-1
Recovering Data from a Cloned Volume	5-2
Recovering Data After a Disaster	5-3
If You Have the ISM Server Bootstrap Printout	5-3
If You Do Not Have the ISM Server Bootstrap Printout	5-4
Creating and Backing Up the ISM Bootstrap	5-4
Transferring Data with the Imported Restore.	5-5
Step 1: Backing Up the Source Database Server to ISM	5-5
Step 2: Porting ISM to Your Target Database Server	5-8
Step 3: Restoring Data from ISM to the Target Database	5-13
Appendix A. ISM Command-Line Utilities	A-1
Appendix B. Accessibility	B-1
Glossary	C-1
Notices	D-1
Index	X-1

Introduction

About This Manual	vii
Types of Users	vii
Software Dependencies	viii
Assumptions About Your Locale	viii
Demonstration Database	ix
New Features in ISM and Organizational Changes	ix
Documentation Conventions	x
Typographical Conventions	x
Feature, Product, and Platform	xi
Syntax Diagrams	xii
How to Read a Command-Line Syntax Diagram	xiii
Keywords and Punctuation	xiv
Identifiers and Names	xv
Example Code Conventions	xv
Additional Documentation	xvi
Installation Guides	xvi
Online Notes	xvi
Locating Online Notes	xvii
Online Notes Filenames	xviii
Informix Error Messages	xviii
Manuals	xix
Online Manuals	xix
Printed Manuals	xix
Online Help	xix
Accessibility	xix
IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set	xix
Compliance with Industry Standards	xxii
IBM Welcomes Your Comments	xxiii

In This Introduction

This introduction provides an overview of the information in this manual and describes the conventions it uses.

About This Manual

This manual describes IBM Informix Storage Manager (ISM), which manages storage devices and media for your Informix database server. Use ISM with ON-Bar to back up and restore database server data.

Types of Users

This manual is written for the following users:

- Database administrators

- Database server administrators
- Backup operators

This manual is written with the assumption that you have the following background:

- Familiarity with the ON-Bar backup and restore utility
For more information about ON-Bar, refer to the *IBM Informix: Backup and Restore Guide*
- A working knowledge of your computer, your operating system, and the utilities that your operating system provides
- Some experience working with relational databases or exposure to database concepts

If you have limited experience with relational databases, SQL, or your operating system, refer to the *IBM Informix: Getting Started Guide* for your database server for a list of supplementary titles.

Software Dependencies

ISM supports the following database servers:

- IBM Informix Dynamic Server, Version 7.3x
- IBM Informix Dynamic Server with Advanced Decision Support and Extended Parallel Options, Version 8.2x
- IBM Informix Extended Parallel Server, Version 8.3x
- IBM Informix Extended Parallel Server, Version 8.4
- IBM Informix Extended Parallel Server, Version 8.5
- IBM Informix Dynamic Server, Version 9.2x
- IBM Informix Dynamic Server, Version 9.3
- IBM Informix Dynamic Server with J/Foundation, Version 9.3
- IBM Informix Dynamic Server, Version 9.40
- IBM Informix Dynamic Server, Version 10.0

Assumptions About Your Locale

IBM Informix products can support many languages, cultures, and code sets. All the information related to character set, collation, and representation of numeric data, currency, date, and time is brought together in a single environment, called a Global Language Support (GLS) locale.

The examples in this manual are written with the assumption that you are using the default locale, **en_us.8859-1**. This locale supports U.S. English format conventions for date, time, and currency. In addition, this locale supports the ISO 8859-1 code set, which includes the ASCII code set plus many 8-bit characters such as é, è, and ñ.

If you plan to use nondefault characters in your data or your SQL identifiers, or if you want to conform to the nondefault collation rules of character data, you need to specify the appropriate nondefault locale.

For instructions on how to specify a nondefault locale, additional syntax, and other considerations related to GLS locales, see the *IBM Informix: GLS User's Guide*.

Demonstration Database

The DB–Access utility, provided with the database server products, includes one or more of the following demonstration databases:

- The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in IBM Informix manuals are based on the **stores_demo** database.

Extended Parallel Server

- The **sales_demo** database illustrates a dimensional schema for data-warehousing applications. For conceptual information about dimensional data modeling, see the *IBM Informix: Database Design and Implementation Guide*.

End of Extended Parallel Server

Dynamic Server

- The **superstores_demo** database illustrates an object-relational schema. The **superstores_demo** database contains examples of extended data types, type and table inheritance, and user-defined routines.

End of Dynamic Server

For information about how to create and populate the demonstration databases, see the *IBM Informix: DB–Access User's Guide*. For descriptions of the databases and their contents, see the *IBM Informix: Guide to SQL Reference*.

The scripts that you use to install the demonstration databases reside in the **\$INFORMIXDIR/bin** directory on UNIX and in the **%INFORMIXDIR%\bin** directory on Windows.

New Features in ISM and Organizational Changes

The following list provides information about the new features in ISM, Version 2.2, and other changes to this manual. To go to the desired page, click the chapter number. For a description of all new features for your database server, see your *IBM Informix: Getting Started Guide*.

- The ISM setup procedures included in Chapter 1
- The procedure for migrating from ISM, Version 1.0, to ISM, Version 2.2, and reverting from ISM, Version 2.2, to ISM, Version 1.0, in Chapter 1
- The procedure for performing an imported restore in Chapter 5
- The **ism_chk.pl** command for collecting information about ISM, ON-Bar, and database server processes in Appendix A
- ISM support on 32-bit and 64-bit operating systems and on Windows 2000 (see Chapter 1)

Important: The ISM Administrator GUI is no longer supported. However, you can continue to use ISM commands to manage ISM. For the list of ISM commands, see Appendix A.

Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other volumes in the documentation set.

The following conventions are discussed:

- Typographical conventions
- Other conventions
- Syntax diagrams
- Command-line conventions
- Example code conventions

Typographical Conventions

This manual uses the following conventions to introduce new terms, illustrate screen displays, describe command syntax, and so forth.

Convention	Meaning
KEYWORD	All primary elements in a programming language statement (keywords) appear in uppercase letters in a serif font.
<i>italics</i> <i>italics</i> <i>italics</i>	Within text, new terms and emphasized words appear in italics. Within syntax and code examples, variable values that you are to specify appear in italics.
boldface <i>boldface</i>	Names of program entities (such as classes, events, and tables), environment variables, file and pathnames, and interface elements (such as icons, menu items, and buttons) appear in boldface.
monospace <i>monospace</i>	Information that the product displays and information that you enter appear in a monospace typeface.

Convention	Meaning
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.
>	This symbol indicates a menu item. For example, "Choose Tools > Options " means choose the Options item from the Tools menu.

Tip: When you are instructed to "enter" characters or to "execute" a command, immediately press RETURN after the entry. When you are instructed to "type" the text or to "press" other keys, no RETURN is required.

Feature, Product, and Platform

Feature, product, and platform markup identifies paragraphs that contain feature-specific, product-specific, or platform-specific information. Some examples of this markup follow:

Dynamic Server

Identifies information that is specific to IBM Informix Dynamic Server

End of Dynamic Server

Extended Parallel Server

Identifies information that is specific to IBM Informix Extended Parallel Server

End of Extended Parallel Server

UNIX Only

Identifies information that is specific to UNIX platforms

End of UNIX Only

Windows Only

Identifies information that is specific to the Windows environment

End of Windows Only

This markup can apply to one or more paragraphs within a section. When an entire section applies to a particular product or platform, this is noted as part of the heading text, for example:

Table Sorting (Linux Only)

Syntax Diagrams







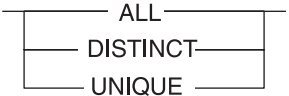
This guide uses syntax diagrams built with the following components to describe the syntax for statements and all commands other than system-level commands.

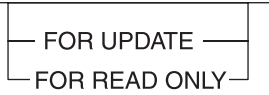
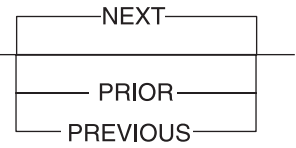
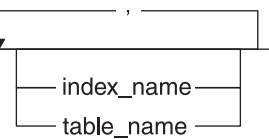

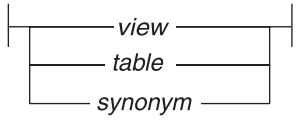
Note: Starting in 2004, syntax diagrams have been reformatted to conform to the IBM standard.

Syntax diagrams depicting SQL and command-line statements have changed in the following ways:

- The symbols at the beginning and end of statements are now double arrows instead of a vertical line at the end.
- The symbols at the beginning and end of syntax segment diagrams are now vertical lines instead of arrows.
- How many times a loop can be repeated is now explained in a diagram footnote instead of a number in a gate symbol.
- Syntax statements that are longer than one line now continue on the next line instead of looping down with a continuous line.
- Product or condition-specific paths are now explained in diagram footnotes instead of icons.

The following table describes syntax diagram components.

Component represented in PDF	Component represented in HTML	Meaning
	>>-----	Statement begins.
	----->	Statement continues on next line.
	>-----	Statement continues from previous line.
	-----<<	Statement ends.
	-----SELECT-----	Required item.
	--+-----+--- '-----LOCAL-----'	Optional item.
	--+-----+--- +--DISTINCT-----+ '-----UNIQUE-----'	Required item with choice. One and only one item must be present.

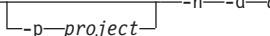
Component represented in PDF	Component represented in HTML	Meaning
	<pre>---+-----+--- +--FOR UPDATE-----+ '--FOR READ ONLY--'</pre>	Optional items with choice are shown below the main line, one of which you might specify.
	<pre>.---NEXT----- -----+-----+--- +---PRIOR-----+ '---PREVIOUS-----'</pre>	The values below the main line are optional, one of which you might specify. If you do not specify an item, the value above the line will be used as the default.
	<pre>.-----,----- v----- ----- -----+-----+--- +---index_name---+ '---table_name---'</pre>	Optional items. Several items are allowed; a comma must precede each repetition.
	<pre>>>- Table Reference -><</pre>	Reference to a syntax segment.
<p>Table Reference</p> 	<pre> --+-----view-----+-- +-----table-----+ '-----synonym-----'</pre>	Syntax segment.

How to Read a Command-Line Syntax Diagram

The following command-line syntax diagram uses some of the elements listed in the table in the previous section.


Creating a No-Conversion Job

```

>>-onpladm create job-job--n-d-device--D-database-

```

```

--t-table-
 (1)
 Setting the Run Mode


```

Notes:

- 1 See page 17-4

The second line in this diagram has a segment named “Setting the Run Mode,” which according to the diagram footnote, is on page 17-4. This segment is shown in the following segment diagram (the diagram uses segment start and end components).

Setting the Run Mode:



To construct a command correctly, start at the top left with the command. Follow the diagram to the right, including the elements that you want. The elements in the diagram are case sensitive.

The Creating a No-Conversion Job diagram illustrates the following steps:

1. Type **onpladm create job** and then the name of the job.
2. Optionally, type **-p** and then the name of the project.
3. Type the following required elements:
 - **-n**
 - **-d** and the name of the device
 - **-D** and the name of the database
 - **-t** and the name of the table
4. Optionally, you can choose one or more of the following elements and repeat them an arbitrary number of times:
 - **-S** and the server name
 - **-T** and the target server name
 - The run mode. To set the run mode, follow the Setting the Run Mode segment diagram to type **-f**, optionally type **d**, **p**, or **a**, and then optionally type **l** or **u**.
5. Follow the diagram to the terminator.

Your diagram is complete.

Keywords and Punctuation

Keywords are words reserved for statements and all commands except system-level commands. When a keyword appears in a syntax diagram, it is

shown in uppercase letters. When you use a keyword in a command, you can write it in uppercase or lowercase letters, but you must spell the keyword exactly as it appears in the syntax diagram.

You must also use any punctuation in your statements and commands exactly as shown in the syntax diagrams.

Identifiers and Names

Variables serve as placeholders for identifiers and names in the syntax diagrams and examples. You can replace a variable with an arbitrary name, identifier, or literal, depending on the context. Variables are also used to represent complex syntax elements that are expanded in additional syntax diagrams. When a variable appears in a syntax diagram, an example, or text, it is shown in *lowercase italic*.

The following syntax diagram uses variables to illustrate the general form of a simple SELECT statement.

►—SELECT—*column_name*—FROM—*table_name*—►

When you write a SELECT statement of this form, you replace the variables *column_name* and *table_name* with the name of a specific column and table.

Example Code Conventions

Examples of SQL code occur throughout this manual. Except as noted, the code is not specific to any single IBM Informix application development tool.

If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...

DELETE FROM customer
    WHERE customer_num = 121
...

COMMIT WORK
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using DB–Access, you must delimit multiple statements with semicolons. If you are using an SQL API, you must use EXEC SQL at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement.

Tip: Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the manual for your product.

Additional Documentation

For additional information, refer to the following types of documentation:

- Installation guides
- Online notes
- Informix error messages
- Manuals
- Online help

Installation Guides

Installation guides are located in the **/doc** directory of the product CD or in the **/doc** directory of the product's compressed file if you downloaded it from the IBM Web site. Alternatively, you can obtain installation guides from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Online Notes

The following sections describe the online files that supplement the information in this manual. Please examine these files before you begin using your IBM Informix product. They contain vital information about application and performance issues.

Online File	Description	Format
TOC Notes	The TOC (Table of Contents) notes file provides a comprehensive directory of hyperlinks to the release notes, the fixed and known defects file, and all the documentation notes files for individual manual titles.	HTML
Documentation Notes	The documentation notes file for each manual contains important information and corrections that supplement the information in the manual or information that was modified since publication.	HTML, text
Release Notes	The release notes file describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. For some products, this file also contains information about any known problems and their workarounds.	HTML, text
Machine Notes	(Non-Windows platforms only) The machine notes file describes any platform-specific actions that you must take to configure and use IBM Informix products on your computer.	text
Fixed and Known Defects File	This text file lists issues that have been identified with the current version. It also lists customer-reported defects that have been fixed in both the current version and in previous versions.	text

Locating Online Notes

Online notes are available from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>. Additionally you can locate these files before or after installation as described below.

Before Installation

All online notes are located in the `/doc` directory of the product CD. The easiest way to access the documentation notes, the release notes, and the fixed and known defects file is through the hyperlinks from the TOC notes file.

The machine notes file and the fixed and known defects file are only provided in text format.

After Installation

On UNIX platforms in the default locale, the documentation notes, release notes, and machine notes files appear under the **\$INFORMIXDIR/release/en_us/0333** directory.

Dynamic Server

On Windows the documentation and release notes files appear in the **Informix** folder. To display this folder, choose **Start > Programs > IBM Informix Dynamic Server version > Documentation Notes** or **Release Notes** from the taskbar.

Machine notes do not apply to Windows platforms.

End of Dynamic Server

Online Notes Filenames

Online notes have the following file formats:

Online File	File Format	Examples
TOC Notes	<i>prod_os_tocnotes_version.html</i>	ids_win_tocnotes_10.0.html
Documentation Notes	<i>prod_bookname_docnotes_version.html/txt</i>	ids_hpl_docnotes_10.0.html
Release Notes	<i>prod_os_relnotes_version.html/txt</i>	ids_unix_relnotes_10.0.txt
Machine Notes	<i>prod_machine_notes_version.txt</i>	ids_machine_notes_10.0.txt
Fixed and Known Defects File	<i>prod_defects_version.txt</i>	ids_defects_10.0.txt client_defects_2.90.txt
	<i>ids_win_fixed_and_known_defects_version.txt</i>	ids_win_fixed_and_known_defects_10.0.txt

Informix Error Messages

This file is a comprehensive index of error messages and their corrective actions for the Informix products and version numbers.

On UNIX platforms, use the **finderr** command to read the error messages and their corrective actions.

Dynamic Server

On Windows, use the Informix Error Messages utility to read error messages and their corrective actions. To display this utility, choose **Start > Programs > IBM Informix Dynamic Server version > Informix Error Messages** from the taskbar.

End of Dynamic Server

You can also access these files from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Manuals

Online Manuals

A CD that contains your manuals in electronic format is provided with your IBM Informix products. You can install the documentation or access it directly from the CD. For information about how to install, read, and print online manuals, see the installation insert that accompanies your CD. You can also obtain the same online manuals from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Printed Manuals

To order hardcopy manuals, contact your sales representative or visit the IBM Publications Center Web site at <http://www.ibm.com/software/howtobuy/data.html>.

Online Help

IBM Informix online help, provided with each graphical user interface (GUI), displays information about those interfaces and the functions that they perform. Use the help facilities that each GUI provides to display the online help.

Accessibility

IBM is committed to making our documentation accessible to persons with disabilities. Our books are available in HTML format so that they can be accessed with assistive technology such as screen reader software. The syntax diagrams in our manuals are available in dotted decimal format, which is an accessible format that is available only if you are using a screen reader. For more information about the dotted decimal format, see the Accessibility appendix.

IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set

The following tables list the manuals that are part of the IBM Informix Dynamic Server, Version 10.0 and the CSDK Version 2.90, documentation set. PDF and HTML versions of these manuals are available at <http://www.ibm.com/software/data/informix/pubs/library/>. You can order hardcopy versions of these manuals from the IBM Publications Center at <http://www.ibm.com/software/howtobuy/data.html>.

Table 1. Database Server Manuals

Manual	Subject
Administrator's Guide	Understanding, configuring, and administering your database server.
Administrator's Reference	Reference material for Informix Dynamic Server, such as the syntax of database server utilities onmode and onstat , and descriptions of configuration parameters, the sysmasters tables, and logical-log records.
Backup and Restore Guide	The concepts and methods you need to understand when you use the ON-Bar and ontape utilities to back up and restore data.
DB-Access User's Guide	Using the DB-Access utility to access, modify, and retrieve data from Informix databases.
DataBlade API Function Reference	The DataBlade API functions and the subset of ESQL/C functions that the DataBlade API supports. You can use the DataBlade API to develop client LIBMI applications and C user-defined routines that access data in Informix databases.
DataBlade API Programmer's Guide	The DataBlade API, which is the C-language application-programming interface provided with Dynamic Server. You use the DataBlade API to develop client and server applications that access data stored in Informix databases.
Database Design and Implementation Guide	Designing, implementing, and managing your Informix databases.
Enterprise Replication Guide	How to design, implement, and manage an Enterprise Replication system to replicate data between multiple database servers.
Error Messages file	Causes and solutions for numbered error messages you might receive when you work with IBM Informix products.
Getting Started Guide	Describes the products bundled with IBM Informix Dynamic Server and interoperability with other IBM products. Summarizes important features of Dynamic Server and the new features for each version.
Guide to SQL: Reference	Information about Informix databases, data types, system catalog tables, environment variables, and the stores_demo demonstration database.
Guide to SQL: Syntax	Detailed descriptions of the syntax for all Informix SQL and SPL statements.
Guide to SQL: Tutorial	A tutorial on SQL, as implemented by Informix products, that describes the basic ideas and terms that are used when you work with a relational database.
High-Performance Loader User's Guide	Accessing and using the High-Performance Loader (HPL), to load and unload large quantities of data to and from Informix databases.
Installation Guide for Microsoft Windows	Instructions for installing IBM Informix Dynamic Server on Windows.
Installation Guide for UNIX and Linux	Instructions for installing IBM Informix Dynamic Server on UNIX and Linux.

Table 1. Database Server Manuals (continued)

Manual	Subject
J/Foundation Developer's Guide	Writing user-defined routines (UDRs) in the Java programming language for Informix Dynamic Server with J/Foundation.
Large Object Locator DataBlade Module User's Guide	Using the Large Object Locator, a foundation DataBlade module that can be used by other modules that create or store large-object data. The Large Object Locator enables you to create a single consistent interface to large objects and extends the concept of large objects to include data stored outside the database.
Migration Guide	Conversion to and reversion from the latest versions of Informix database servers. Migration between different Informix database servers.
Optical Subsystem Guide	The Optical Subsystem, a utility that supports the storage of BYTE and TEXT data on optical disk.
Performance Guide	Configuring and operating IBM Informix Dynamic Server to achieve optimum performance.
R-Tree Index User's Guide	Creating R-tree indexes on appropriate data types, creating new operator classes that use the R-tree access method, and managing databases that use the R-tree secondary access method.
SNMP Subagent Guide	The IBM Informix subagent that allows a Simple Network Management Protocol (SNMP) network manager to monitor the status of Informix servers.
Storage Manager Administrator's Guide	Informix Storage Manager (ISM), which manages storage devices and media for your Informix database server.
Trusted Facility Guide	The secure-auditing capabilities of Dynamic Server, including the creation and maintenance of audit logs.
User-Defined Routines and Data Types Developer's Guide	How to define new data types and enable user-defined routines (UDRs) to extend IBM Informix Dynamic Server.
Virtual-Index Interface Programmer's Guide	Creating a secondary access method (index) with the Virtual-Index Interface (VII) to extend the built-in indexing schemes of IBM Informix Dynamic Server. Typically used with a DataBlade module.
Virtual-Table Interface Programmer's Guide	Creating a primary access method with the Virtual-Table Interface (VTI) so that users have a single SQL interface to Informix tables and to data that does not conform to the storage scheme of Informix Dynamic Server.

Table 2. Client/Connectivity Manuals

Manual	Subject
Client Products Installation Guide	Installing IBM Informix Client Software Developer's Kit (Client SDK) and IBM Informix Connect on computers that use UNIX, Linux, and Windows.
Embedded SQLJ User's Guide	Using IBM Informix Embedded SQLJ to embed SQL statements in Java programs.

Table 2. Client/Connectivity Manuals (continued)

Manual	Subject
ESQL/C Programmer's Manual	The IBM Informix implementation of embedded SQL for C.
GLS User's Guide	The Global Language Support (GLS) feature, which allows IBM Informix APIs and database servers to handle different languages, cultural conventions, and code sets.
JDBC Driver Programmer's Guide	Installing and using Informix JDBC Driver to connect to an Informix database from within a Java application or applet.
.NET Provider Reference Guide	Using Informix .NET Provider to enable .NET client applications to access and manipulate data in Informix databases.
ODBC Driver Programmer's Manual	Using the Informix ODBC Driver API to access an Informix database and interact with the Informix database server.
OLE DB Provider Programmer's Guide	Installing and configuring Informix OLE DB Provider to enable client applications, such as ActiveX Data Object (ADO) applications and Web pages, to access data on an Informix server.
Object Interface for C++ Programmer's Guide	The architecture of the C++ object interface and a complete class reference.

Table 3. DataBlade Developer's Kit Manuals

Manual	Subject
DataBlade Developer's Kit User's Guide	Developing and packaging DataBlade modules using BladeSmith and BladePack.
DataBlade Module Development Overview	Basic orientation for developing DataBlade modules. Includes an example illustrating the development of a DataBlade module.
DataBlade Module Installation and Registration Guide	Installing DataBlade modules and using BladeManager to manage DataBlade modules in Informix databases.

Compliance with Industry Standards

The American National Standards Institute (ANSI) and the International Organization of Standardization (ISO) have jointly established a set of industry standards for the Structured Query Language (SQL). IBM Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of IBM Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL Common Applications Environment (CAE) standards.

IBM Welcomes Your Comments

We want to know about any corrections or clarifications that you would find useful in our manuals, which will help us improve future versions. Include the following information:

- The name and version of the manual that you are using
- Section and page number
- Your suggestions about the manual

Send your comments to us at the following email address:

`docinf@us.ibm.com`

This email address is reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact IBM Technical Support.

We appreciate your suggestions.

Chapter 1. Informix Storage Manager

What Is ISM?	1-1
ISM Server Tasks	1-2
Y2K Compliance	1-3
How Do ON-Bar and ISM Fit Together?	1-3
What Happens During a Backup Operation?.	1-3
What Happens During a Restore Operation?.	1-5
Configuring ISM	1-6
Setting Up ISM on UNIX	1-8
Setting Up ISM on Windows	1-9
Setting Up ISM Volume Pool Names	1-10
Setting ISM Environment Variables and ONCONFIG Parameters	1-10
Updating the onbar Script	1-11
Uninstalling ISM.	1-11
Using ISM with DNS Disabled	1-13
Migrating ISM 1.0 to ISM 2.2	1-14
Reverting from ISM 2.2 to ISM 1.0.	1-17
Using ISM.	1-17
Using ISM Administrative Privileges	1-18
Starting and Stopping ISM	1-18
Reviewing ISM Logs	1-18

In This Chapter

This chapter provides an overview of how IBM Informix Storage Manager (ISM) works with the ON-Bar utility, which you can use to perform backups and restores. For more information, see the *IBM Informix: Backup and Restore Guide*.

ISM provides data-storage management services for your Informix database server. ISM receives backup and restore requests from ON-Bar and directs your data to and from storage volumes that are mounted on storage devices. ISM tracks backed-up data through a data life cycle that the database administrator or system administrator determines and can automatically manage your storage devices and storage volumes.

What Is ISM?

ISM includes the following components:

- The ISM server for exchanging backup and recovery information between storage devices and ON-Bar

- The ISM catalog, which maintains up-to-date records of the backup operations that have been performed and the media on which the backed-up data is stored

The ISM server resides on the same computer as ON-Bar and the Informix database server; your storage devices must be attached to this computer as well. The ISM server runs on UNIX or Windows. ISM supports both 32-bit and 64-bit operating systems.

ISM Server Tasks

Figure 1-1 shows the interaction between ON-Bar, ISM server, and ISM commands. ISM must be running before you can use the ISM commands.




Figure 1-1. ON-Bar and ISM Interaction

The ISM server performs the following tasks:

- Receives backup and restore requests from ON-Bar
- Manages the data life cycle of your data
- Maintains the ISM catalog

The ISM catalog allows you to store related save sets in volumes and to track related volumes by means of volume pools. A volume is a specific backup medium, such as a disk or reel of magnetic tape. A volume pool is a named set of volumes.

- Creates and saves information that is needed to restore the state of the ISM server after a disk failure.

The ISM command-line utilities allow you to perform the following tasks:

- Add, change, and delete administrative users
- Connect to and administer multiple ISM servers on your network

- Configure up to four storage devices for each ISM server
- Configure up to four active data streams
- Label and mount storage volumes on your storage devices
- Determine how long the ISM server should track backed-up data
- Manage the storage volumes
- Create exact copies of data

For a complete description, see Appendix A, “ISM Command-Line Utilities,” on page A-1.

Y2K Compliance

ISM is Y2K compliant. ISM supports dates in the year 2000 and beyond. All internal dates are stored in an internal format that allows representation of dates from January 1, 1970 through December 31, 2037. ISM correctly interprets the year 2000 as a leap year. When a year is entered as a two-digit specification, ISM interprets it as follows:

- A year entered in the range of 70 - 99 is interpreted as 1970 - 1999
- A year entered in the range of 00 - 37 is interpreted as 2000 - 2037
- A year entered in the range of 38 - 69 is invalid

How Do ON-Bar and ISM Fit Together?

ON-Bar and ISM communicate through an X/Open Backup Services Application (XBSA) Programming Interface, which enables ISM to provide media-management services for the database server. Figure 1-1 on page 1-2 illustrates how ON-Bar and ISM interact on a network to back up the database server.

You can initiate ON-Bar backup and restore operations with the **onbar** command line, the **onbar** script, or the ON-Bar option in IBM Informix Server Administrator (ISA). For information on how to run ON-Bar commands from ISA, see the ISA online help.

What Happens During a Backup Operation?

The database server administrator or operator uses ON-Bar to request a backup operation. When the ISM server receives a backup request from ON-Bar, it creates one or more save sets on the storage media as the data is backed up. The ISM server writes the backup data to storage volumes mounted on storage devices that are attached to the ISM server. ISM uses volume pools to sort specific data onto preselected storage volumes in order to help manage your data. For example, by default ISM writes data from dbspaces to a collection of volumes in the ISMData volume pool and the logical-log files to volumes in the ISMLogs volume pool.

During a backup or restore operation, ISM requests that you mount a labeled storage volume from the appropriate volume pool on one of the storage devices attached to the ISM server. You can label and mount a storage volume when you receive the request from the ISM server, or you can prelabel several volumes and mount them as you need them.

To further automate the labeling and mounting process, you can designate that a specific storage device should handle all storage media that is loaded in it exclusively as ISM media. This arrangement means that you can load a blank tape in a storage device, and the ISM server labels and mounts it when it needs a new storage volume for an ON-Bar backup. For more information, see Chapter 4, “Managing Storage Volumes,” on page 4-1.

Figure 1-2 on page 1-5 illustrates how data and requests move between ON-Bar, the ISM server, and the database server during a backup operation. Backup requests are processed in the following sequence:

1. The user uses ON-Bar to start a backup operation.
2. ON-Bar determines what data needs to be backed up.
3. ON-Bar connects to the ISM server.
4. ON-Bar retrieves the data from the appropriate storage spaces or logical logs on the Informix database server, builds a list of backup objects, and passes it to the ISM server.
5. The ISM server creates a save set from the data and writes that save set to the proper storage volumes.
6. The ISM server enters a record of the save sets and the storage volumes in the ISM catalog.

The ISM server uses the ISM catalog to track the save sets and the volumes to which they are backed up.

7. After the backup operation finishes, ON-Bar disconnects from ISM.
8. ON-Bar creates a save set called the ISM server bootstrap, which contains configuration information about the ISM server for disaster recovery purposes.

The ISM server writes the ISM server bootstrap to a volume from the ISMData volume pool. You can send the ISM server bootstrap information to the printer that you designated during installation.

9. ON-Bar writes information about the backup to the database server message log and ON-Bar activity log. ISM writes information about the backup to the `ISM\2.20\logs\messages` log.

Use the `ism_watch` command to track ISM server activity.

Warning: Because ISM server bootstrap information is vital for disaster recovery, keep the most recent bootstrap printout in a safe place.




Figure 1-2. Data Movement During an ON-Bar Backup

What Happens During a Restore Operation?

The database server administrator or operator uses ON-Bar to request a restore operation. When the ISM server receives a restore request from ON-Bar, ISM searches the ISM catalog to identify the required storage volumes. When ISM recovers data, ISM requests a specific storage volume by name. To fulfill this request, mount the appropriate storage volume on one of the storage devices that is attached to the ISM server.

Figure 1-3 illustrates how data and requests move between ON-Bar, the ISM server, and the database server through XBSA during a restore operation. Restore requests are processed in the following sequence:

1. The user uses ON-Bar to start a restore operation.
2. ON-Bar determines what data needs to be restored.
3. ON-Bar connects to the ISM server.
4. The ISM server searches the ISM catalog for specific restore objects that ON-Bar requests. ISM maps each restore object to its offset in a save set.
5. ON-Bar retrieves the data from the ISM server and passes it to the Informix database server, object by object.
6. The database server writes the data to the proper storage spaces or logical logs.
7. After the restore operation finishes, ON-Bar disconnects from ISM.




Figure 1-3. Data Movement During an ON-Bar Restore

Configuring ISM

This section contains ISM configuration information. Before you can use ISM with ON-Bar to back up and restore data, you must configure it.

The ISM server files are installed but not started when you install the database server.

To configure ISM:

1. Copy the `sm_versions.std` file to `sm_versions` and verify that the latest ISM information is in the file.
2. Verify that the `BAR_BSA_PATH` configuration parameter points to the correct XBSA shared library for ISM. For details, see the *IBM Informix: Backup and Restore Guide*.
3. Set ISM environment variables.
To set up ISM, see “Setting Up ISM on UNIX” on page 1-8 or “Setting Up ISM on Windows” on page 1-9.
4. Add administrative users. Ensure that your administrative users have privileges on both the computer that runs the ISM server software (either UNIX or Windows) and the computer that runs the ISM commands (Windows).
For details, see “Adding the First Administrative User to the ISM Server” on page 2-4.
5. Determine whether you want to back up data to disk (also called *file-type devices*), tape, or optical drives.
If you want to back up data to file-type devices, see “Using Filesystems as Storage Devices” on page 3-6.
If you want to back up data to tape or optical drives, see “Adding and Deleting Storage Devices” on page 3-3.

6. Configure a storage device for storage-space backups and another one for logical-log backups. The devices must be attached to the computer where ISM is installed.

For details, see “Using Storage Devices for Logical-Log Backups” on page 3-7.

7. Label your storage volumes.
 - ISMData is the default volume pool name for storage spaces. ISMLogs is the default volume pool name for logical-log files.
 - Use the ISMDiskData and ISMDiskLogs volume pools if you are using file-type storage devices.

For more information, see “Using Volume Pools and Volume Labels” on page 4-6. For information on how to change which volume pool is used for which data, see “Using ISM Configuration Parameters to Change the Volume Pool Name” on page 2-6.

8. Mount the storage volumes on the storage devices.
9. Establish a retention period for the backed-up data that the ISM server tracks. The default value is six months.

For details, see “Setting the Retention Period for All Data” on page 2-3.

10. Determine if you want the ISM server to automatically manage the storage volumes loaded in a storage device.

See “Automating Media Management for Tape and Optical Devices” on page 3-9.

11. Designate a safe place to keep the ISM server bootstrap printouts. You need the information in the bootstrap printout for disaster recovery purposes.

For details, see “If You Have the ISM Server Bootstrap Printout” on page 5-3.

After you configure the ISM server and storage devices and label volumes for your database and logical-log backups, you are ready to initiate a backup operation with ON-Bar.

Warning: Install only one storage manager per computer to prevent possible conflicts with the XBSA shared library. Do not run ISM and Legato NetWorker on the same computer because they write files to the `/nsr` directory.

Extended Parallel Server

Extended Parallel Server allows multiple ISM instances but only one instance per coserver node.

End of Extended Parallel Server

Tip: If the database server is installed on UNIX, the ISM utilities are stored in the **bin** directory with the other executables. If the database server is installed on Windows, the ISM utilities are stored in the ISM directory.

Setting Up ISM on UNIX

To start ISM on UNIX for the first time, issue the **ism_startup -init** command. The following C shell script is an example of a simple setup for making ON-Bar storage-space and logical-log backups to disk files:

```
#!/bin/csh
# setup for ISM storage manager

echo "ISM setup begins"

# the user must edit these two lines to select
# appropriate disk locations for storage manager use.
setenv SM_DISKDEV1 /backups/dbspaces
setenv SM_DISKDEV2 /backups/logfiles

# start ISM server
echo "initialize ISM server, must be done as root"
ism_startup -init
```

Important: To make **informix** the ISM user and thereby ensure **informix** has access to conduct ISM operations during daily routines, you must run the following command after any **ism_startup -init** or ISM installation:

```
ism_add -admin informix@server

# create some devices and media
# the ISM user, defined in ism_add -admin, can do these tasks
echo "create and mount ISM devices and pools"
ism_add -device $SM_DISKDEV1 -type file
ism_add -device $SM_DISKDEV2 -type file
ism_op -label $SM_DISKDEV1 -pool ISMDiskData -volume ISMData
ism_op -label $SM_DISKDEV2 -pool ISMDiskLogs -volume ISMLogs
ism_op -mount $SM_DISKDEV1
ism_op -mount $SM_DISKDEV2
echo "end of ISM setup"
```

To configure ISM to start up automatically:

1. Add **\$INFORMIXDIR/bin/ism_startup** to one of the **rc** scripts in **/etc/rc2.d** (on Solaris) to start the ISM server when UNIX boots. Make sure the **PRINTER** environment variable is set correctly.
2. Add **\$INFORMIXDIR/bin/ism_shutdown -q** to one of the **rc** scripts in **/etc/rc0.d** to shut down the ISM server when UNIX shuts down.

After you configure ISM, you can set the **LOG_BACKUP_MODE** configuration parameter to **CONT** or **MANUAL** to back up the logical logs. Until ISM is configured and available to back up the logical logs, **LOG_BACKUP_MODE** should be set to **NONE**.

To stop and start ISM manually:

- After you shut down the database server, to shut down ISM, use the **ism_shutdown** command.
- Before you restart the database server, issue the **ism_startup** command to start ISM.

Setting Up ISM on Windows

The database server installation program creates the following ISM services that start automatically on system startup:

- ISM Local Execution
- ISM Portmapper
- ISM Server

On Windows, the database server installer automatically initializes ISM. The Service Control Manager automatically starts the ISM server when the computer reboots and stops ISM when the computer shuts down.

You can use a batch file to set up ISM operations with ON-Bar on Windows. This batch file configures two file-type devices to back up logical logs and storage spaces. This configuration allows you to back up logical logs and storage spaces to disk. You must edit the locations for the **SM_DISKDEV1** and **SM_DISKDEV2** environment variables as needed.

Important: Install and enable the Netware IPX and SPX networking protocols on the computer where the Informix database server is installed only if you encounter network flooding.

You must be a member of the **Informix-Admin** group to perform these operations. To open a command window for running ISM commands, select **Start > Programs > Informix > ol_servername**.

```
REM setup for IBM Informix Storage
Manager
echo "ISM setup begins"
```

```
REM this step should be done only once
ism_startup -init
```

Important: To make **informix** the ISM user and thereby ensure **informix** has access to conduct ISM operations during daily routines, you must run the following command after any **ism_startup -init** or ISM installation:

```
ism_add-admin informix@server

REM user must edit these two lines to select appropriate disk
REM locations for storage manager use.
set SM_DISKDEV1=\backups\dbspaces
set SM_DISKDEV2=\backups\logfiles
```

```

REM we need to create some devices and media
echo "create and mount ISM devices and pools"
ism_add -device %SM_DISKDEV1% -type file
ism_add -device %SM_DISKDEV2% -type file
ism_op -label %SM_DISKDEV1% -pool ISMDiskData -volume ISMData
ism_op -label %SM_DISKDEV2% -pool ISMDiskLogs -volume ISMLogs
ism_op -mount %SM_DISKDEV1%
ism_op -mount %SM_DISKDEV2%
echo "end of ISM setup"

```

Setting Up ISM Volume Pool Names

If you use ISM, you can specify the volume pool names for storage spaces and logical logs in the `ISM_DATA_POOL` and `ISM_LOG_POOL` configuration parameters in the `ONCONFIG` file. If you do not set these parameters, `ISMData` is the default volume pool name for storage spaces and `ISMLogs` is the default volume pool name for logical logs.

If you are using `ISMDiskData` and `ISMDiskLogs` for the volume pool names, update the `ISM_DATA_POOL` and `ISM_LOG_POOL` entries in the `ONCONFIG` file with those names.

To back up storage spaces and logical logs to the same storage device, set the `ISM_DATA_POOL` and `ISM_LOG_POOL` parameters in the `ONCONFIG` file to the same volume pool name. However, for best restore performance, we recommend that you back up storage spaces and logical logs to different volume pools on different storage devices.

Setting ISM Environment Variables and ONCONFIG Parameters

When you use ISM, you must set certain environment variables that affect the way in which the ISM server handles requests. For information, see “Using ON-Bar Environment Variables with ISM” on page 2-9.

Dynamic Server

You set these environment variables in the `onbar` script or in your environment.

End of Dynamic Server

Extended Parallel Server

You can set these environment variables in your environment if you enter `onbar -w` at the command prompt, or before you start the database server set them in `start_worker.sh` or `start_worker.bat`.

End of Extended Parallel Server

Updating the onbar Script

If you configure ISM and ON-Bar to use a storage pool other than ISMData, you need to edit the **onbar** or **onbar.bat** script to direct **ism_catalog** to the storage pool that you are using. For example, if you use the ISMDiskData pool, make the following changes in the **ism_catalog** line of the **onbar** script:

UNIX Only

On UNIX, the command in the **onbar** script is:

```
\${INFORMIXDIR}/bin/ism_catalog -create_bootstrap -pool  
ISMData
```

End of UNIX Only

Windows Only

On Windows, the command in the **onbar.bat** script is:

```
%ISMDIR%\bin\ism_catalog -create_bootstrap -pool ISMDiskData
```

End of Windows Only

Warning: If the **onbar** script selects a nonexistent storage pool, ON-Bar seems to hang and a message appears in the ISM log that an operation is waiting for the ISMData pool to be mounted.

For more information on the **onbar** script, see the *IBM Informix: Backup and Restore Guide*.

Uninstalling ISM

The following two sections provide procedures for uninstalling ISM on UNIX and Windows. You must uninstall your currently running ISM in these situations:

- Before you reinstall ISM on UNIX or in Windows
- Before you install Legato NetWorker or another storage manager

You must not install ISM on a computer in which a NetWorker server or a NetWorker client is already installed. If you attempt such an installation, it fails with the following message:

```
The ISM setup script detected that another storage manager is  
installed in /nsr, which conflicts with ISM. If you want to setup  
ISM, you need to de-install the other storage manager and then  
run $INFORMIXDIR/bin/ism_startup -init.
```

For ISM to be installed on a computer that has a NetWorker server or client on it, the NetWorker server or client has to be completely uninstalled first.

To uninstall your currently running ISM on UNIX:

1. Log in as user **informix**.
2. Change to the ISM executables directory:

```
cd $INFORMIXDIR/bin
```
3. Become user **root**.
4. Shut down the ISM daemons:

```
./ism_shutdown -q
```
5. Verify that the daemons are shut down:

```
ps -ef | grep nsr
```
6. Move the ISM executables to a temporary directory:

```
mkdir ./ism.tmp  
mv ism* ism.tmp  
mv nsr* ism.tmp  
mv mm* ism.tmp  
mv save* ism.tmp  
mv scanner ism.tmp  
mv usam ism.tmp  
mv recover ism.tmp
```
7. Move the ISM catalogs to a temporary directory:

```
cd $INFORMIXDIR  
mv ism ism.bak
```
8. Remove the symbolic link to the ISM catalogs:

```
rm /nsr
```
9. Remove the call to **ism_catalog** from the **\$INFORMIXDIR/onbar** script.

Now you are ready to install your new storage manager on UNIX.

Tip: The **ism_shutdown** command shuts down the ISM daemons **nsrd**, **nsrexecd**, **nsrmmdbd**, and **nsrindexd**. It does not uninstall ISM, that is, it does not remove the ISM executables in **\$INFORMIXDIR/bin**, the symbolic link **/nsr**, or the various ISM files in **\$INFORMIXDIR/ism**.

To uninstall your currently running ISM on Windows:

1. Shut down ISM. On a command line enter:

```
ism_shutdown -deinstall
```
2. Shut down the Informix database server:

```
onmode -yuk
```
3. Check **Settings > Control Panel > Services** to see if the database server is down. If it is not down, wait until it is.

4. Change the user environment variable **PATH** so that **%ISMDIR%\bin** is no longer in the environment variable **PATH**. You might find:
`PATH=c:\WINDOWS\system32;d:\ISM\1.00\bin;c:\msdev;`
Change it to:
`PATH=c:\WINDOWS\system32;c:\msdev;`
5. Rename the directory that contains the ISM executables:
`rename %ISMDIR% ism.bak`
6. Delete the user environment variable **ISMDIR**.
7. Start the Windows registry. On the command line enter:
`regedt32`
8. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > Informix > ISM**.
With the cursor on **ISM**, click to select and then choose **Edit > Delete** from the registry menu.

Warning: Make sure that ISM is selected. You must have the correct entry. Be careful when you edit the registry because mistakes can corrupt Windows.
9. Go to **HKEY_LOCAL_MACHINE > SOFTWARE > XBSA**.
With the cursor on **XBSA**, click to select and then choose **Edit > Delete** from the registry menu.

Using ISM with DNS Disabled

Windows Only

The domain name service (DNS) translates the hostnames into internet protocol (IP) addresses. This means that when your computer is not on the network, DNS is disabled and you need to provide the IP address and hostname in the hosts file.

End of Windows Only

Dynamic Server

For nonnetwork ISM installations with Dynamic Server, create a Windows hosts file with the default local loopback entry as follows:

```
127.0.0.1    localhost
```

The hosts file is located in **C:\WINNT\system32\drivers\etc\hosts**. If this entry is not present in the hosts file, you might get the following error when you use ISM:

```
nsmmd #1: unauthorized READ_LABEL call to mmd 1: Access violation -  
[] unknown host: 127.0.0.1
```

End of Dynamic Server

Migrating ISM 1.0 to ISM 2.2

You can either upgrade ISM 1.0 to ISM 2.2 alone or upgrade ISM with the database server version.

Migration reinstalls the ISM binaries but maintains the ISM data (the catalogs and tape volumes that contain the save sets).

The following section explains how to migrate ISM 1.0 to ISM 2.2. ISM 2.2 includes changes to the format of data in the ISM catalogs and volumes. Begin the following procedure with ISM 1.0 running on your old database server version.

Important: Do not use ISM 1.0 storage media for future backups after you migrate from ISM 1.0 to ISM 2.2.

To migrate ISM:

1. Complete a full backup of your system with one of the following commands:

```
onbar -b -w  
onbar -b -L 0
```
2. Create a bootstrap of your ISM 1.0 server with the following command:

```
ism_catalog -create_bootstrap -pool ISMData
```

The bootstrap is a copy of the files and directories in **\$INFORMIXDIR/ism/mm**, **index**, and **res** (UNIX) or **%ISMDIR%\mm**, **index**, and **res** (Windows). These directories are backed up into a single save set, called the bootstrap.
3. Shut down the ISM 1.0 server.

UNIX Only

On UNIX:

```
ism_shutdown
```

End of UNIX Only

Windows Only

On Windows:


```
ism_shutdown -deinstall
```

End of Windows Only

4. Remove the ISM 1.0 catalogs.

Warning: Do not remove the **res** directory.

UNIX Only

On UNIX:

```
rm -rf $INFORMIXDIR/ism/index  
rm -rf $INFORMIXDIR/ism/mm
```

End of UNIX Only

Windows Only

On Windows:

```
del %ISMDIR%/index  
del %ISMDIR%/mm
```

End of Windows Only

If you have file-type devices configured in ISM, you cannot move, copy, or rename the directories that contain those devices.

5. Remove ISM 1.0. (This step is optional.)
Follow the instructions on how to remove ISM 1.0 in “Uninstalling ISM” on page 1-11. Use **regedt32** to check the registry keys.
6. On Windows, you must rename or remove the ISM 1.0 **bin** directory because the ISM 2.2 installer installs the ISM files in a different directory. Then move the ISM 2.2 **bin** directory to the ISM 1.0 **bin** directory location.
7. Install the new ISM 2.2 files, either separately or with the new database server version.

Warning: You must install the new ISM 2.2 files in the same directory as the ISM 1.0 files.

8. If you are upgrading ISM on Windows, follow these steps to ensure that ISM is properly configured.
 - a. The installer might have made a Windows command window script for the Informix environment. The filename of this script is *servername.CMD*.

- b. Edit this file to be sure that **ISMDIR** and **PATH** are correct for the location of the ISM 2.2 directory.
- c. Change the `%INFORMIXDIR%\bin\onbar.bat` file for any user-customized references to the ISM directory.
- d. If necessary, edit the `%INFORMIXDIR%\bin\SETISM.BAT` file to ensure that it refers to the ISM 2.2 directory.
- e. Check your database server configuration file (usually `%INFORMIXDIR%\etc\ONCONFIG.servername`). Be sure that the `BAR_BSALIB_PATH` parameter points to the `libbsa.dll` in the `bin` subdirectory of the ISM 2.2 directory.
- f. Check Windows system environment variable settings that affect the **PATH** environment variable or that set the **ISMDIR** environment variable.
- g. Copy the `sm_versions.std` file to create a new `sm_versions` file.
- h. If you changed the configuration files, you might need to reboot Windows.

Warning: An error message that an entry point cannot be found in `libnsr.dll` means part of the Windows configuration still references the old ISM installation.

9. Start the ISM server with the following command. Do not initialize the server.
`ism_startup`
10. Place the tape that contains the bootstrap in a device and mount it, if it is not already mounted.
11. Create an index for your host with the following command:
`nsrck -c`
12. Locate the bootstrap on the tape and note the save-set ID: **ism_catalog -find_bootstrap device_name**.
13. Recover your bootstrap with the following command:
`ism_catalog -recover`

Important: Do not replace the `res` directory with the `res.R` directory. Wait for the preceding command to issue a message that the index was fully recovered.

14. Remove all the defined devices with the following command. You must unmount each device individually.
`ism_op -unmount device_name`
15. Segregate all ISM 1.0 volumes. Make backups of file-type devices.

Important: For future use, you must store the tape with the bootstrap that you created in step 2. Without the bootstrap, you cannot revert to ISM 1.0 if you need to.

16. Label new volumes.

ISM 2.2 must not write to any ISM 1.0 volumes because they would be unreadable by ISM 1.0 if you chose to revert.

```
ism_op -label device_name -volume volume_name -pool  
ISMData
```

17. Mount the new volumes with the following command for each device:

```
ism_op -mount device_name
```

18. Create a new bootstrap to back up the converted indexes.

```
ism_catalog -create_bootstrap -pool ISMData
```

19. Upgrade the database server to the new version if necessary, and then start the database server.

20. Immediately perform a level-0 backup.

```
onbar -b -L 0
```

Reverting from ISM 2.2 to ISM 1.0

When you revert the database server to an earlier version, do *not* revert to ISM 1.0. All versions of the database server, up through Version 9.3, support ISM 2.2. Also, ISM 2.2 is year-2000 compliant, but ISM 1.0 is not. Versions of the database server without ON-Bar are not compatible with ISM.

- Because the database server installers install ISM, be sure to preserve the ISM 2.2 directory by renaming it.
- To install the earlier database server version, follow the instructions in the *IBM Informix: Installation Guide*.
- To revert to the earlier database server version, follow the instructions in the *IBM Informix: Migration Guide*.
- To restore the ISM directory, copy the new ISM files to the directory that you renamed earlier and rename the directory to its original name.

If you need to revert the database server and perform a point-in-time restore of the earlier database server version, ISM 2.2 might have the original backups in its catalog (if you followed the procedure in this manual for upgrading ISM). If the backups are no longer in the ISM catalog, recover the catalogs from the backup media after you revert the database server.

Using ISM

To have access to and be able to run ISM, you must correctly configure certain privileges and system requirements.

Using ISM Administrative Privileges

By default, ISM grants administrative user privileges to **root** or **informix** for ISM servers on UNIX and to **Administrator** for ISM servers on Windows. Only users included in the ISM server administrator list can make configuration changes to the ISM server. For more details, see “Adding the First Administrative User to the ISM Server” on page 2-4.

Important: If the ISM server software is running on UNIX, you must have **root** or user **informix** privileges to recover data. For more information, see Chapter 5, “Recovering Save Sets with ISM,” on page 5-1.

Starting and Stopping ISM

To start ISM, issue the following command:

```
ism_startup
```

To stop ISM, issue the following command:

```
ism_shutdown
```

Reviewing ISM Logs

The ISM server maintains two log files in which it records its activities. The first log records the requests received and operations that the ISM server performs. The second log records the XBSA messages that the ISM server sends or receives.

To review either of the log files directly, make a copy and open the copy with a text editor. The log files reside in the following locations:

UNIX Only

- On UNIX, the activity log is located in a file named **\$INFORMIXDIR/ism/logs/daemon.log**. The XBSA messages log is located in **\$INFORMIXDIR/ism/applogs/xbsa.messages**.

If you have different Informix database servers installed in different directories on the same computer, you can find the location of the activity log for the active ISM server with the following command:

```
ls -ls /nsr
```

End of UNIX Only

Windows Only

- On Windows, the activity log is located in a file named **%ISMDIR%\logs\daemon.log**. The XBSA messages log is located in

`c:\nsr\applogs\xbsa.messages.`

End of Windows Only

For more information on the activity log and XBSA messages log, see “Using ISM Environment Variables” on page 2-8.

Chapter 2. Configuring ISM Server Properties

Setting the Number of Parallel Data Streams.	2-1
Setting Target Sessions for a New Device	2-2
Setting the Retention Period for All Data	2-3
Adding, Changing, or Deleting Administrative Users	2-4
Adding the First Administrative User to the ISM Server	2-4
Choosing Acceptable ISM User and Group Names.	2-5
Deleting an Administrative User.	2-6
Connecting to a Different ISM Server	2-6
Using ISM Configuration Parameters to Change the Volume Pool Name	2-6
Using ISM Environment Variables	2-8
Using ON-Bar Environment Variables with ISM.	2-9

In This Chapter

This chapter describes how to customize the ISM server properties to suit your environment.

The ISM server connects your Informix database server to your storage devices and enables you to monitor the progress of your backup and restore operations, manage backed-up data, and control your storage media and storage devices.

Use the **ism_config** command to change the ISM server properties. For more information, see “ism_config” on page A-8.

Setting the Number of Parallel Data Streams

ISM supports up to four parallel data streams. You can adjust the number of data streams that can be active at one time. When you set this value, it allows you to specify the number of save sets that are backed up to the ISM server simultaneously. Backing up save sets in parallel allows your storage devices to operate at their most efficient levels because they receive a steady stream of data, which is written onto your storage volumes. For example, suppose that you have four storage devices and set the number of parallel data streams to four. During a backup to all four devices, each device receives one data stream.

To set or adjust the number of data streams that can be backed up in parallel:

1. To display the ISM server properties, use the following command:

```
ism_show -config
```

The **parallelism** field in the **ism_show -config** output indicates the number of parallel data streams.

2. To change the number of parallel data streams, use the following command:

```
ism_config -streams #datastreams
```

To back up only one data stream at a time, set the value to 1.

```
ism_config -streams 1
```

In the following situations, adjust the number of data streams that you back up in parallel:

- The ISM server is taking too long to back up all the data streams.
Increase the number of data streams.
- The ISM server is so busy backing up data streams that it is not available for any other tasks.
Decrease the number of data streams.
- The ISM server runs out of swap space or memory.
Decrease the number of data streams.

Setting Target Sessions for a New Device

When you define and mount multiple devices for the ISMData pool, all storage-space backups still go to the first device. You can send storage-space backups or logical-log backups to multiple devices at the same time.

To use multiple devices:

1. Start ISM in the usual way (**ism_startup**).
2. Define and label all devices (up to 4).
3. Shut down ISM in the usual way (**ism_shutdown**).
4. Enter the following command at the prompt:

```
nsradmin
```
5. At the **nsradmin** prompt, enter the following command:

```
print NSR device
```
6. To update the target sessions variable for your configured devices, enter the following command. You need to enter this command only once to apply it to all configured devices.

```
update target sessions: 1
```
7. The following information is displayed for each device:

```
update resource id: resource_id  
target sessions: 1;  
Update?  
Enter y.
```


8. Restart ISM in the usual way (**ism_startup**).

Setting the Retention Period for All Data

You can specify how long the ISM server tracks the data on your storage volumes before it designates that a volume is eligible for reuse.

The ISM server maintains an entry for each backed-up save set in the ISM catalog for the length of time that you set as the retention period. You can use ON-Bar to recover a save set at any time during the retention period for a save set.

When the retention period for a save set expires, the ISM server no longer maintains the save-set entry in the ISM catalog.

To obtain information about the expiration date and retention period:

1. To display the expiration date for a save set, issue the **ism_show -volumes** command (see the **expires** field).
2. To display the retention period for the ISM server, issue the **ism_show -config** command (see the **retention policy** field).

To recover an expired save set, use the **ism_catalog -recreate_from** command at any point until the storage volume that contains the save set is relabeled. For details on how to recover expired save sets, see “Re-Creating the ISM Catalog Before You Use ON-Bar” on page 5-1.

To set the retention period for the ISM server:

1. To set the length of time before the data is marked Expired, issue the following command:

```
ism_config -retention #days
```
2. Specify the retention period in days. The default is six months.

Important: Because changing the retention period for the ISM server applies the new retention period to all the data that the ISM server tracks, some data might become eligible for automatic recycling either sooner or later than you expected. Be careful about changing the retention period for the ISM server because you could potentially overwrite recent backups.

For more information on how ISM manages the life cycle of data, see “Managing the Data Life Cycle” on page 4-1. For details on how to issue ISM commands at the command line, see Appendix A.

Adding, Changing, or Deleting Administrative Users

To make configuration changes to the ISM server, label and mount storage volumes, and initiate cloning operations, you must have administrative user privileges on the ISM server.

Adding the First Administrative User to the ISM Server

To administer your ISM server, you need administrative user privileges for that ISM server based on the system on which you run your administrative programs. Administrative user privileges can be granted to individual user accounts, UNIX groups, or Windows groups. They are typically restricted to an individual system.

When you installed ISM, the local **root@server_name** for UNIX ISM servers or the local Windows Administrators group, **Administrators@server_name**, for Windows, ISM servers was added as an administrative user by default. To add additional administrative users, you must log in as either user **root** or **informix** on UNIX or as **Administrators@server_name** on Windows.

For each administrator entry in the list of administrators, the ISM server compares the hostname of the administrator with the hostname of the account. If the hostnames match, the ISM server compares the administrator name with both the user and group name of the account. If the administrator name and the user name or group name match, the ISM server grants administrative privileges to that administrator.

For example, if your ISM server is running on a computer called **nova**, but you use ISM commands on a computer called **jupiter**, you must log on to **nova** and add the user name or group name that you use on **jupiter** to the **nova** list of ISM administrative users. The following numbered steps illustrate:

- How to add a user or user group on **jupiter** as an administrative user to the ISM server on **nova**
- How to add a user or user group on **nova** (other than **root**, **informix**, or **Administrator**) as an administrative user to the ISM server on **nova**

To add a user or user group:

1. Log on to **nova** as user **root**, **informix**, or **Administrator**.
2. To add a user or group on **jupiter** as an administrative user for the ISM server, run one of the following commands at the command line:
 - To add a user named **karen** on **jupiter**, run the following command:

```
ism_add -admin karen@jupiter
```

You might need to qualify the hostname as follows:
`jupiter.informix.com`
 - To add a group on **jupiter**, run the following command:

```
ism_add -admin group_name@jupiter
```

This user or group can now use any of the ISM command-line utilities on the computer **jupiter** to administer the ISM server on **nova**.

3. To add other users or groups on nova (the computer that is running the ISM server) as administrative users for the ISM server, run one of the following commands at the command line:

- To add a user on **nova**, run the following command:

```
ism_add -admin user_name@nova
```

- To add a group on **nova**, run the following command:

```
ism_add -admin group_name@nova
```

This user or group can now use the ISM command-line utilities on **nova** to administer the ISM server on **nova**.

If you are not on the administrative users list and try to make configuration changes to the ISM server, you receive the following message:

```
user user_name needs to be on administrator's list
```

Choosing Acceptable ISM User and Group Names

If you need to restrict the administrative users list so that it does not include **root**, **informix**, or **Administrators**, make sure that the administrative users list includes access to the user who started the ISM server. Typically this user is the system account.

The following list shows acceptable ISM user and group names.

Acceptable Usage	Example
<i>local_group</i> @ <i>hostname</i>	root@jupiter; Administrators@jupiter
<i>global_group</i> @ <i>hostname</i>	Domain_Admins@jupiter
<i>account_name</i> @ <i>hostname</i>	jsmith@jupiter@informix.com
<i>wildcard</i> @ <i>hostname</i>	*@jupiter
<i>account_name</i> @ <i>wildcard</i>	jsmith@*

Important: Be careful about using wildcards to grant users or groups administrative user privileges. If you use a wildcard before the @ sign, you grant ISM administrative user privileges to any user logged onto that system. If you use a wildcard after the @ sign, you grant ISM administrative user privileges to any group or account that is named before the @ sign, regardless of location. It is recommended that you avoid using wildcards when you grant administrative user privileges.

ISM does not recognize the following names:

@Domain name
@unresolved TCP/IP name

ISM displays the following error message:

Cannot establish a connection with server '*hostname*':
Server *hostname*: Unknown host.

Deleting an Administrative User

Every user or group in the Administrator list has permission to modify the ISM server configuration.

To delete an administrative user from the ISM server:

1. Log on as user **root**, **informix**, or **Administrator**.
2. Specify the user name or group and hostname in the following command:

```
ism_rm -admin name@hostname
```

If you remove yourself as an administrative user, you can no longer administer the ISM server.

Connecting to a Different ISM Server

If you have more than one ISM server on your network, you can open a connection to a different ISM server from your current session. The following commands allow you to specify an ISM server on a different computer:

- **ism_add**
- **ism_catalog**
- **ism_clone**
- **ism_config**
- **ism_op**
- **ism_rm**
- **ism_show**
- **ism_watch**

Using ISM Configuration Parameters to Change the Volume Pool Name

The following parameters, when listed in the ONCONFIG configuration file for the database server, affect how the ISM server handles backup and restore requests.

ONCONFIG Parameter	Effect on ISM Server
ISM_DATA_POOL	If present in the ONCONFIG file, this parameter specifies the volume pool to be used for backing up dbspaces and other storage spaces. The value for this parameter

can be any volume pool that ISM recognizes. If this parameter is not present, ISM uses the ISMData volume pool.

ISM_LOG_POOL

If present in the ONCONFIG file, this parameter specifies the volume pool to be used for backing up logical logs. The value for this parameter can be any volume pool that ISM recognizes. If this parameter is not present, ISM uses the ISMLogs volume pool.

The following volume pool names are recommended:

- ISMData for storing storage spaces on nonfile-type devices
- ISMDiskData for storing storage spaces on file-type devices
- ISMLogs for storing logical logs on nonfile-type devices
- ISMDiskLogs for storing logical logs on file-type devices

You can direct storage-space and logical-log backups to the same device. In the following example, ISMDiskData is the volume pool name for both storage-space and logical-log backups. However, you can specify any volume pool name in the ONCONFIG file as follows:

```
ISM_DATA_POOL ISMDiskData
ISM_LOG_POOL ISMDiskData
```

If you change the name of either ISMData or ISMLogs, you must also complete the following steps:

- Update ISM_DATA_POOL and ISM_LOG_POOL in the ONCONFIG file with the new names.
- Change the create-bootstrap command in the **onbar** script (`$INFORMIXDIR/bin/onbar` or `onbar.bat`).

For example, change `ism_catalog -create_bootstrap -pool ISMData` to `ism_catalog -create_bootstrap -pool new_pool_name`.

If you choose not to use ISM, remove the `create_bootstrap` command from the `onbar` script or `onbar.bat` script.

————— UNIX Only —————

On Solaris, if you are a member of the **bargroup**, it should be your primary group.

————— End of UNIX Only —————

For information on the **onbar** script, see the *IBM Informix: Backup and Restore Guide*.

Using ISM Environment Variables

The following environment variables control the size and number of log files that ISM maintains.

Environment Variable	Description
ISM_MAXLOGSIZE	If this variable is not set in the environment of the ISM server, its value is 0, which is interpreted as no limit. That is, the system resource is the limit. If this variable is set, the value set is the maximum log size. When the size of the log reaches this ISM_MAXLOGSIZE value, the log is truncated and it starts over with the next log message.
ISM_MAXLOGVERS	If this variable is set in the environment of the ISM server, the variable specifies the maximum number of activity-log files to be preserved by the ISM server. If this environment variable is not set, the default number of files to be preserved is four. If the variable is set to a null value, the ISM server preserves no activity log files.
ISM_DEBUG_FILE	If this variable is set in the environment of the ISM server, the variable specifies where the XBSA messages should be written. If you do not set this variable, the XBSA messages log is in \$INFORMIXDIR/ism/applogs/xbsa.messages on UNIX or c:\nsr\applogs\xbsa.messages on Windows.

When ISM starts up, it checks the length of the activity log. On UNIX, this log resides in the **\$INFORMIXDIR/ism/logs/daemon.log** file. On Windows, the file is **%ISMDIR%\logs\daemon.log**. If the log file exceeds the threshold that the **ISM_MAXLOGSIZE** environment variable specifies, the ISM server renames it to **daemon.001**. If a file named **daemon.001** is already present, the ISM server renames that file to **daemon.002**, and so on, for all other existing log files up to the maximum number that **ISM_MAXLOGVERS** specifies.

The ISM server does not monitor the size of the XBSA messages log, which resides in a file named **\$INFORMIXDIR/ism/applogs/xbsa.messages** on UNIX. This log resides in a file named **c:\nsr\applogs\xbsa.messages** on Windows. Use the **ISM_DEBUG_FILE** environment variable to change the location of the XBSA messages log.

To manage the growth of this file, set the **ISM_DEBUG_LEVEL** environment variable to 1 (the default) or 0, thereby minimizing or eliminating updates. If you use a higher level of detail, or you want complete assurance that the XBSA messages log can never fill up the disk, you can use the **cron** utility on UNIX or the **Winat** facility on Windows to save a periodic snapshot of the file and then truncate it.

For more information about ISM logs, refer to “Reviewing ISM Logs” on page 1-18. For information about **ISM_DEBUG_LEVEL**, see “Using ON-Bar Environment Variables with ISM” below.

Using ON-Bar Environment Variables with ISM

The following environment variables, when set in the ON-Bar environment, determine whether ISM uses compression or encryption when backing up data.

Dynamic Server

You can set these environment variables in the **onbar** script file. For example:
`ISM_COMPRESSION=TRUE; export ISM_COMPRESSION`

End of Dynamic Server

Extended Parallel Server

You can set these environment variables in the **start_worker** script file. Insert the lines that contain the environment variables before the line that starts the **onbar_w** utility.

End of Extended Parallel Server

Environment Variable in Effect When ON-Bar Issues a Request	Effect on ISM Server Processing for That Request
ISM_COMPRESSION	If this variable is set to TRUE in the environment of the onbar process making a request, the ISM server uses a data-compression algorithm to store or retrieve the data specified in that request. If it is set to FALSE or is not present, the ISM server does not use compression.
ISM_ENCRYPTION	If this variable is set to TRUE or XOR in the environment of the onbar process making a request, the ISM server uses encryption to store or retrieve the data specified in that request. If it is set to NONE or is not present, the ISM server does not use encryption.

Environment Variable in Effect When ON-Bar Issues a Request	Effect on ISM Server Processing for That Request
ISM_DEBUG_LEVEL	<p>If this variable is set in the ON-Bar environment, it controls the level of reporting detail recorded in the XBSA messages log. The XBSA shared library writes to this log. You can specify a value between 0 and 9. If this environment variable is not set, has a null value, or has a value outside of this range, the detail level equals 1 by default. A detail level of 0 suppresses all XBSA debugging records. A level of 1 reports only XBSA failures.</p>

Chapter 3. Managing Storage Devices

Choosing and Configuring Storage Devices	3-2
Adding and Deleting Storage Devices	3-3
Enabling and Disabling a Storage Device	3-4
Changing Storage-Device Properties	3-5
Designating a Storage Device as Read-only	3-5
Using Tape Drives and Optical Disk Drives as Storage Devices	3-5
Using Filesystems as Storage Devices	3-6
Using Storage Devices for Logical-Log Backups.	3-7
Reserving a Storage Device for Cloning	3-7
Automating Media Management for Tape and Optical Devices.	3-9
Enabling Automatic Media Management	3-9
Selecting a Storage Volume	3-9
Relabeling Volumes Automatically.	3-10
Using Unrecognizable Volumes	3-10

In This Chapter

This chapter discusses the commands that you use to configure storage devices to use ISM.

Choosing and Configuring Storage Devices

Table 3-1 shows the types of storage devices that ISM supports.

Table 3-1. Types of Storage Devices

Device Type	Media Type
Half-inch magnetic tape drives	himt
Quarter-inch cartridge tape drives	qic
Four-millimeter (DAT) tape drives	4mm
Eight-millimeter tape drives	8mm
Eight-millimeter five-gigabyte tape drives	8mm 5GB
3480 tape drives	3480
3570 tape drives	3570
4890 tape drives	4890
9490 Timberline tape drives	9490
Digital linear tape drives	dlt
Filesystem	File
Optical disk drives	Optical
VHS tape drives	VHS

Important: ISM does not support autochangers or high-speed devices. Autochangers are storage devices such as jukeboxes, carousels, libraries, near-line storage, datawheels, silos, and autoloaders.

Keep the following issues in mind when you configure your storage devices:

- ISM supports up to four storage devices.
- It is recommended that you reserve a storage device for your logical-log backups.
- You cannot use automatic rewind tape devices.

The combination of tape drives, optical disk drives, and filesystems that you use depends on your environment needs. When you use a filesystem (file-type) as a storage device, some limitations are imposed on your configuration. For more information, see “Using Filesystems as Storage Devices” on page 3-6.

Adding and Deleting Storage Devices

After you install ISM, you can add up to four storage devices. When you first install the ISM software, you must add the devices that you want to use as storage devices.

To add a new tape drive or optical disk drive to the ISM server:

1. Create a new device as follows:

```
ism_add -device device_path -type type
```

- a. Specify the complete pathname of the storage device that is connected to the ISM server.

UNIX Only

For ISM servers on UNIX computers, the format is as follows:

```
/dev/rmt/0mbn
```

End of UNIX Only

Windows Only

For ISM servers on Windows, the format is as follows:

```
\\.\Tape0
```

End of Windows Only

- b. Specify the device type.

The **Media Types** column in Table 3-1 on page 3-2 shows the valid device types for tape and optical drives.

2. To verify that you created the new device, issue the command:

```
ism_show -devices
```

By default, the new storage device is enabled for use in backup and restore operations.

For more information, see “Using Tape Drives and Optical Disk Drives as Storage Devices” on page 3-5.

Important: A no-rewind backup device is required when you use tape devices.

To add a filesystem as a storage device to the ISM server:

1. On the ISM server, create a new directory or folder where you want to store the data.

Be sure to create it on media with free space that is larger than the size of your backup. It is recommended that you create the file-type device on a different drive from the one that contains your operating system so that you do not run out of disk space.

2. Create a new device as follows:

```
ism_add -device device_path -type file
```

- a. Specify the pathname of the directory or folder that you just created.
- b. Specify a device type of **file**.

3. To verify that you created the new device, issue the command:

```
ism_show -devices
```

By default, the new storage device is enabled for use in backup and restore operations.

For more information, see “Using Filesystems as Storage Devices” on page 3-6.

Important: You cannot include file-type volumes used with a file-type device in a pool that includes other types of volumes. ISM does not allow volumes of other types to be recycled or added to a pool of file-type volumes. For details on how ISM uses volume pools, see “How ISM Uses Volume Pools” on page 4-5.

To delete a storage device:

1. Remove the storage device as follows:

```
ism_rm -device device_path
```

2. To verify that the device was deleted, issue the command:

```
ism_show -devices
```

You cannot remove a storage device that has a storage volume mounted on it.

Enabling and Disabling a Storage Device

Before the ISM server can use a storage device for backup or restore operations, you must enable it. To determine if a storage device is enabled or disabled, use the **ism_show -devices** command or the **nsradmin -c** command. You can use the **nsradmin -c** command in full-screen mode to enable or disable a storage device as shown in the following steps.

Using the nsradmin command to enable or disable a storage device:

1. Log in as user **root**.
2. Enter the following command:

```
nsradmin -c
```
3. Choose the command **Select**.
4. Under type, choose **NSR device**.

5. Choose the command `Next` until you locate the storage device you want to modify.
6. Choose the command `Edit`, and then press `RETURN`.
7. Under `enabled`, select `Yes` or `No`.
8. Press `ESC` to leave the menu.
9. Press `RETURN` to save your changes.

You cannot enable or disable a storage device that has an ISM storage volume mounted on it.

Changing Storage-Device Properties

If a storage device already has a storage volume mounted on it, you cannot change any of the device properties.

Use the `ism_config` command to change the storage device properties. For details, see “`ism_config`” on page A-8.

Designating a Storage Device as Read-only

If you want to designate a storage device associated with a volume as read-only, follow these steps:

Designating a storage device as read-only:

1. Unmount the storage volume on the device that you want to designate as read-only:

```
ism_op -umount device_path
```
2. To designate the storage device as read-only, use the command:

```
ism_config -volume name -only_restore yes
```

You can use this volume only for restores, not for backups.
3. To redesignate the storage device as writable, use the command:

```
ism_config -volume name -only_restore no
```
4. Remount the storage volume on the device.

You can now use this storage volume for backups.

Using Tape Drives and Optical Disk Drives as Storage Devices

Using tape drives or optical disk drives as storage devices provides a long-term, inexpensive way to store your data. You can store tapes and optical disks offsite or in a protected vault. They provide virtually unlimited storage capacity because you can always purchase new media. Tape and optical disk drives access data sequentially, which slows down backup and recovery.

If you are using tape drives or optical disk drives as storage devices for backups of your storage spaces, it is recommended that you reserve one

device for your continuous logical-log file storage volumes. The other devices are available to receive data from storage-space backups.

Important: When the currently mounted tape becomes full, ISM requires operator intervention to change tapes. The request to change tapes appears in the ISM log. You use the `ism_watch` command to monitor ISM for tape-change requests. To automate tape changes during unattended backups, you can use NetWorker or other storage managers that support advanced devices.

Warning: You must use no-rewind tape devices because ISM writes a file mark on the storage volume at the end of each backup and then appends data onto the storage volume based on the position of the file mark. (A backup operation might contain several save sets.) If you use a rewind tape device, your backups will appear to have completed successfully but your tapes will contain only the last save set. The current save set overwrites the previous save set and data is lost.

Using Filesystems as Storage Devices

Backup and restore speed are important considerations when you plan your backup strategy. File-type storage devices such as filesystems or the hard drive back up data more quickly than some tape devices. This feature is especially important if you have a limited time for backup. You can also access your data more quickly during restore operations because filesystems allow random access to your data.

The maximum size of a file-type device on many systems is 2 gigabytes.

Disadvantages of file-type storage devices include expense (disk drives are more expensive than tapes) and the inability to store the data on the file-type storage device offsite.

You must mount and label storage volumes for file-type devices, even though there is no physical volume to mount. When ISM writes data to this file-type volume, the volume is treated as one or more open files. If you enter a value for the file-type volume, ISM writes data to the file-type volume until that value is reached. For specific instructions on how to label storage volumes, see “Labeling a Volume” on page 4-8.

Although you can specify up to four file-type devices, we recommend a maximum of two file-type devices so that you can reserve two tapes for copying the data. The copying process is called *cloning*. Once a file-type device is full, you must clone the data to tape before you can overwrite the file-type device with new data. Alternately, you could clone your backups to

tape after your backups complete. You must recover the data to the cloned copy on tape. For more information, see “Reserving a Storage Device for Cloning” on page 3-7 and “Cloning Storage Volumes” on page 4-15.

Important: You cannot include file-type devices in a volume pool that includes other types of storage volumes. ISM does not allow volumes of other types to be recycled or added to a pool of file-type volumes. You cannot clone bootstrap records. Bootstrap records are always backed up to the ISMData volume pool, even when all other data is backed up to ISMDiskData or ISMDiskLogs.

Using Storage Devices for Logical-Log Backups

During a backup, the ISM server automatically routes storage spaces to volumes in the ISMData volume pool and logical-log files to volumes in the ISMLogs volume pool. (For more information, see “Using ISM Configuration Parameters to Change the Volume Pool Name” on page 2-6.) You can configure ON-Bar to back up each logical-log file as soon as it becomes full. Thus the timing of logical-log file backups is not always predictable.

When you keep a storage volume from the ISMLogs volume pool mounted on one of the storage devices, you ensure that a storage device is always available to accept logical-log data that the ISM server routes. It is recommended that you keep an ISMLogs labeled storage volume mounted on one of the storage devices at all times. For more information on volume pools, see “How ISM Uses Volume Pools” on page 4-5.

Reserving a Storage Device for Cloning

Cloning data to tape ensures that you can successfully recover your data after the file-type volumes are relabeled. This procedure requires a second tape device, in addition to the one tape device used for logical-log file backups, as Figure 3-1 shows. For more information, see “Using Filesystems as Storage Devices” on page 3-6.




Figure 3-1. ISM Storage-Device Allocation with One File-Type Device

If you are using two file-type devices, reserve two tape devices so you can clone your file-type volumes to tape, as Figure 3-2 shows.




Figure 3-2. ISM Storage-Device Allocation with Two File-Type Devices

To recover data when the file-type volume mounted on the file-type device fills:

1. Make sure that the storage device is not active.
2. Protect the storage volume from backup operations. If the storage volume is not mounted, change it to Backup Disabled:
 - a. Unmount it and note the volume name:


```
ism_op -unmount device_path
```
 - b. Change the backup status of the specified storage volume to Backup Disabled:


```
ism_config -volume name -only_restore yes
```


- c. Mount the storage volume once again:
`ism_op -mount device_path`
3. Label and mount a tape volume in the ISMData Clone pool:
 - a. Create a new device as follows:
`ism_add -device device_path -type tape`
 - b. Label the device:
`ism_op -label device_path -pool pool -volume name`
 - c. Mount the volume:
`ism_op -mount device_path`
4. Clone the file-type volume (or save set) to tape:
`ism_clone -volume name -pool pool`

Depending on the type of media that you use, you might need more than one clone volume to hold all the data from the file-type device.
5. Relabel the file-type volume so that it can be used for new backups.

When you relabel the file-type volume, you can no longer access the original data. If you need to recover data, ISM directs you to mount the cloned volume.

Automating Media Management for Tape and Optical Devices

To simplify the backup procedure and avoid the task of prelabeling and mounting a storage volume before backup, use the **Automatic Media Management** option.

If the ISM server detects a writable storage volume from the appropriate volume pool loaded in a storage device, the ISM server automatically mounts the volume and uses it for the backup operation. If the ISM server detects a volume without an ISM label, it considers the volume to be blank and automatically labels, mounts, and overwrites the volume with new data. ISM labels the storage volumes with the next sequential label that is associated with the appropriate pool of volumes.

Enabling Automatic Media Management

To enable Automatic Media Management from the command line, issue the `ism_config -volume volume_name -recycle auto` command. This feature does not apply to file-type storage devices.

Selecting a Storage Volume

When you enable the **Automatic Media Management** option, ISM uses the following hierarchy to choose a storage volume to write data to:

- A writable storage volume from the appropriate volume pool
- A storage volume from the appropriate pool that is eligible for recycling
- An unlabeled storage volume

Relabeling Volumes Automatically

If ISM detects a storage volume from the appropriate pool with the Recycle status of Eligible, ISM mounts, relabels, and overwrites the storage volume with new data. However, the new label uses the same name and sequence number as the original label. Therefore, newer data might be on a storage volume with a lower label number. ISM labels recycled storage volumes in this way to provide you with an easy way to determine the age of your storage volumes. Lower-numbered storage volumes typically are older than higher-numbered storage volumes, even though the data on the older storage volume might be more recent.

For example, suppose that over time you add ten storage volumes, ISMData.0001 through ISMData.0010, to the ISMData volume pool. When the retention periods for all the save sets on ISMData.0001 expire, ISMData.0001 becomes eligible for automatic recycling. During a backup operation, the ISM server searches for a writable storage volume from the ISMData volume pool. If ISMData.0001 is the only volume from the ISMData pool loaded in a storage device, or if the ISM server needs more than one volume to complete the backup operation and no other writable volumes from the ISMData pool are detected, the ISM server automatically relabels and mounts ISMData.0001 and continues the backup operation. ISM relabels ISMData.0001 with the same name and sequence number. However, the original data on ISMData.0001 can no longer be recovered.

Although the data on ISMData.0001 might be more recent than the data on ISMData.0010, you can tell at a glance that the physical media on ISMData.0001 is older than ISMData.0010. Because storage media can wear out over time, this labeling scheme enables you to monitor the age of your media.

ISM gives each physical volume a default expiration date of two years after it is labeled. If you relabel a volume, it gets a new two-year expiration date. The expiration date of a storage volume refers to the physical media, not the data on the volume. When a volume reaches its expiration date, ISM automatically changes the Backup status of the volume to Disabled. This status means that no new data can be written to the volume. However, the data on the volume can still be recovered until the volume is relabeled.

To determine the expiration date of your physical media, execute the **ism_show -volumes** command at the command line. If you relabel volumes within the same volume pool, however, you can easily determine the relative age of a volume by the volume sequence number.

Using Unrecognizable Volumes

ISM does not recognize storage volumes without a label.

ISM does not recognize the following volume labels and considers these storage volumes as unlabeled:

- Storage volumes that are labeled with something other than an ISM label
- Storage volumes with ISM labels that were written in a density different from that of the device in which they are currently loaded

For example, if you performed ISM backups on a new 8-millimeter tape drive, an older 8-millimeter tape drive might not be able to read the volumes from the newer device because the device uses a different density to write the data.

If ISM does not recognize a storage volume, relabel the storage volume. For details, see “Relabeling a Volume” on page 4-10.

Warning: Be careful about sharing non-ISM storage volumes or ISM volumes written in different densities among different backup devices if you use the **Automatic Media Management** option. Otherwise, you might potentially relabel and overwrite valuable data.

Adjusting the number of active data streams on the ISM server can also affect how ISM writes data to the storage volumes. For more information, see “Setting the Number of Parallel Data Streams” on page 2-1.

Chapter 4. Managing Storage Volumes

Managing the Data Life Cycle	4-1
How ISM Uses the Retention Period	4-2
What Is the Save-Set Retention Status?.	4-2
What Is a Recycling Status of Eligible?.	4-3
What Is a Recycling Status of Manual?.	4-4
How ISM Uses Volume Pools	4-5
Writing to Volumes of Different Media Types	4-6
Using Volume Pools and Volume Labels	4-6
Using Expiration Dates for Volume Labels	4-7
Managing Storage Volumes	4-8
Labeling a Volume	4-8
Mounting a Volume	4-9
Choosing Volumes for Mounting.	4-9
Unmounting Volumes	4-10
Finding a Volume Name	4-10
Relabeling a Volume	4-10
Relabeling a Volume Eligible for Recycling	4-11
Relabeling a Volume Before It Is Eligible for Recycling	4-11
Relabeling a File-Type Volume	4-11
Recycling a Volume or Adding More Volumes to a Pool	4-12
Overriding Automatic Media Management	4-12
Disabling a Storage Volume	4-12
Protecting a Storage Volume from Automatic Recycling.	4-13
Deleting Volumes from the ISM Catalog	4-13
Cloning Storage Volumes and Save Sets	4-14
Tracking Cloned Data	4-15
Cloning Storage Volumes	4-15
Cloning Save Sets	4-16

In This Chapter

This chapter discusses how ISM manages the life cycle of your data and describes how to manage your storage volumes.

Managing the Data Life Cycle

ISM manages your data in the following two ways:

- Uses a data-management policy called a retention period to control the life cycle of your data in the ISM catalog
- Uses volume pools and volume labels to sort and track data in the ISM catalog

How ISM Uses the Retention Period

The retention status of a save set and the recycling status of a storage volume indicate how ISM manages your data. A save set can consist of storage-space backups or logical-log backups. To see the status of a save set, use the **ism_show -volume *name*** command. For more information, see “ism_show” on page A-11.

The ISM server tracks your data based on the retention period that an administrative user sets. For details, see “Setting the Retention Period for All Data” on page 2-3. Every time that a backup is completed, the ISM server creates entries for each backed-up save set in the ISM catalog. The ISM catalog allows ON-Bar to reconstruct the data to look exactly as it did at a specific time.

What Is the Save-Set Retention Status?

Retention, expiration, recycling, and labeling are volume-level operations.

The retention policy applies to both volumes and save sets. ISM retains the volume until all the save sets on that volume expire. A volume expires when the newest save set stored on the volume expires. If a volume is empty, it expires when its expiration date is reached.

For example, if you set the retention period to one month, an empty volume expires in one month. Suppose you save a save set to this volume each week until it holds six save sets. This volume expires after 10 weeks, one month after the last save set is saved during the sixth week.

ON-Bar can restore a save set any time during the retention period. To determine if ON-Bar can restore a specific save set, use **ism_show -volume *name*** to check the save-set status. If the retention status of the save set is blank, the retention period for the save set has not yet expired, and ON-Bar can restore it. For information on using ON-Bar, refer to the *IBM Informix: Backup and Restore Guide*.

If the retention status of a save set is marked Expired, ON-Bar can no longer restore the save set. You need to re-create the save-set entry before you can restore it with ON-Bar. To re-create the save-set entry in the ISM catalog, execute **ism_catalog -recreate_from** from the command line. For details on how to recover expired save sets, see “Re-Creating the ISM Catalog Before You Use ON-Bar” on page 5-1 and “ism_catalog” on page A-2.

When a volume expires, ON-Bar can no longer restore any of the save sets. To recover an expired volume, execute **ism_catalog -recover_from** the command line. After you recover the volume, you can restore the save sets on this volume with ON-Bar.

The status of a save-set entry does not necessarily change to Expired on the same day that its retention period is exceeded. ISM changes the retention status for a save set when *all* the save sets that are dependent on it also exceed their retention periods. Save sets that contain level-1 and level-2 backups are dependent on the save set for the level-0 backup. This extra time ensures that an administrator can reconstruct a save set to a specific time in the retention period.

In Figure 4-1, the retention period is set for one week on the ISM server. The retention period for the save set for the level-0 backup that was performed on September 1 expires on September 8. The level-1 or level-2 backups that were performed during the rest of the week depend on this level-0 backup, so the entry for the level-0 backup does not expire until the entries for the level-1 and level-2 backups have also expired. During this time, you can use ON-Bar to restore the save sets.

On September 15, all the save sets from the backup cycle for the first week change status to Expired.

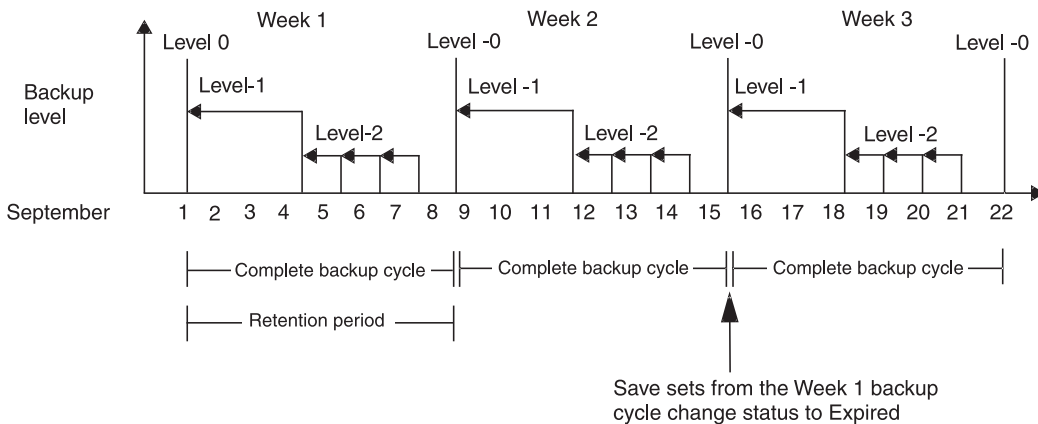


Figure 4-1. One-Week Retention Period

What Is a Recycling Status of Eligible?

When the retention status of *every* save set on a volume has changed to Expired, ISM changes the recycling status of the volume to Eligible. Because a volume might contain many save sets, the recycling status of a volume might not change to Eligible for some time. All entries for Expired save sets remain listed in the **ism_show -volume** command until the volume is relabeled (either automatically or manually) or deleted from the ISM catalog. All the data on the volume remains available for recovery with the **ism_catalog -recreate_from** command until the volume is relabeled.

Therefore, the change in the recycling status of a volume to Eligible is a reminder to the administrator that the volume can be relabeled if conditions are right. If the administrator mounts the volume on a storage device with the **Automatic Media Management** option enabled, relabeling can occur automatically, and the existing data on the volume is no longer accessible to ISM. At this point, you can no longer recover the data on the storage volume.

When an administrator manually deletes a volume, the ISM server removes the entry for the volume from the ISM catalog. This action effectively releases the volume from the control of ISM. However, the data on a deleted volume is still available for recovery with the **ism_catalog -recreate_from** command. This command retrieves the information that is needed to re-create entries in the ISM catalog. Once entries are re-created in the ISM catalog, an administrator can use ON-Bar to recover the data. For details on recovering volumes that have been deleted from the ISM catalog, see “Re-Creating the ISM Catalog Before You Use ON-Bar” on page 5-1.

When a volume is automatically relabeled for reuse within the same pool, the volume name remains unchanged on the volume label. At this point, the volume is ready for new data. All original data is inaccessible and cannot be recovered. For more information on how ISM can automatically recycle storage volumes, see “Relabeling Volumes Automatically” on page 3-10.

What Is a Recycling Status of Manual?

Occasionally, you might want to keep data on a storage volume longer than the retention period that is specified for the ISM server. Changing the recycling status of a storage volume to Manual allows you to protect the volume from becoming eligible for automatic recycling. As long as the storage volume has a recycling status of Manual, ISM does not automatically recycle this volume, nor does it set the recycling status to Eligible. ISM continues to track the volume in the ISM catalog, and you can recover the data on the volume with the **ism_catalog -recreate_from** command. The storage volume is protected from automatic recycling until you remove the Manual recycling status or relabel the volume manually.

You can also use the Manual recycling status to create your own volume recycling policy. For example, suppose that you have a limited number of storage volumes and want to reuse them in a specific order, regardless of the age of the data on the volumes. You can apply the Manual recycling status to each volume and then relabel each volume as needed. The ISM server tracks these volumes in the ISM catalog until you relabel or delete them from the ISM catalog. For storage devices that have the **Automatic Media Management** option enabled, the volumes are protected from automatic recycling.

How ISM Uses Volume Pools

ISM uses pools of volumes to sort and organize your storage space and logical-log backups. A volume pool is a collection of storage volumes. Any tape and optical disk storage volume can belong to the same volume pool. When you initiate an ON-Bar backup, ISM directs the data to a specific volume pool rather than to a specific volume. If ISM detects that a volume from the correct pool is mounted on a storage device, ISM writes data to that volume, regardless of media type.

Every storage volume belongs to a specific pool of volumes. ISM uses two types of volume pools, Backup and Backup Clone. Data cannot be mixed between pools of differing type. This situation enables ISM to keep original and cloned data separate. ISM ships with the following pools for use in backup and cloning operations:

- **ISMData pool.** Storage volumes in this pool accept backups of your storage spaces.
- **ISMLogs pool.** Storage volumes in this pool accept backups of your logical-log files.
- **ISMDiskData pool.** Storage volumes in this pool accept backups of your storage spaces when you use a file-type storage device. Volumes in this pool are of media type *file* and cannot be mixed with volumes of other media type.
- **ISMDiskLogs pool.** Storage volumes in this pool accept backups of your logical-log files when you use a file-type storage device. Volumes in this pool are of media type *file* and cannot be mixed with volumes of other media type.
- **ISMData Clone pool.** Storage volumes in this pool accept clones of your storage-space backups.
- **ISMLogs Clone pool.** Storage volumes in this pool accept clones of your logical-log backups.

We recommend that you use the ISMData and ISMData Clone pools for storage spaces and the ISMLogs and ISMLogs Clone pools for logical logs. If you use file-type storage devices, use the ISMDiskData and ISMDiskLogs pools for the file-type volumes.

ISM also ships with several other preconfigured pools that are included for forward compatibility, as follows:

- Default
- Default Clone
- Full
- NonFull
- Offsite

Writing to Volumes of Different Media Types

Figure 4-2 illustrates how ISM writes data to volumes of different media type that both belong to the same pool.




Figure 4-2. Volume Pools That Contain Varying Media Types

You cannot include storage volumes for file-type devices in a volume pool that includes other types of volumes. ISM does not allow volumes of other types to be recycled or added to a pool of file-type volumes. If you use file-type devices, it is recommended that you use the ISMDiskData and ISMDiskLogs pools with file-type volumes. You can then clone the data from your ISMDiskData backups to tape volumes.

For details on the recommended storage-device configuration when you use file-type devices, see “Using Filesystems as Storage Devices” on page 3-6. For details on the cloning feature, see “Cloning Storage Volumes” on page 4-15. For information on how to choose which volume pool is used for backups, see the “Using ISM Environment Variables” on page 2-8.

Using Volume Pools and Volume Labels

The volume label determines the volume pool to which a storage volume belongs. ISM maintains a record of each volume in the ISM catalog. ISM uses the ISM catalog to identify volumes that it needs for backing up and recovering data.

Each volume pool has a matching label format associated with it. ISM automatically uses the label format associated with the pool that you choose. These label formats provide a method for consistently naming and labeling volumes.

ISM automatically labels a volume with the next sequential label available for that volume pool. The label names are recorded internally on the media, creating an electronic volume label.

Important: Attach an adhesive label to each volume that matches the internal volume label so that you can easily identify your storage volumes.

The preconfigured ISMData and ISMLogs label formats automatically use the ISM server hostname, the appropriate pool name, and a numbering scheme that labels the volumes from 0001 to 9999, as the following table shows.

Label	Use
nova.ISMData.0024	This volume is from the ISMData pool and contains storage spaces from a host named nova .
nova.ISMLogs.0012	This volume is from the ISMLogs pool and contains logical logs from a host named nova .

The following table lists the pools that you should use for your storage spaces and logical logs and the format of the resulting volume label.

Data for Backup	Pool Name	Numbering Sequence	Volume Label Format
Storage spaces	ISMData	0001-9999	<i>hostname.ISMData.0001</i>
	ISMDiskData	0001-9999	<i>hostname.ISMDiskData.0001</i>
	ISMData Clone	0001-9999	<i>hostname.ISMDataClone.0001</i>
Logical logs	ISMLogs	0001-9999	<i>hostname.ISMLogs.0001</i>
	ISMDiskLogs	0001-9999	<i>hostname.ISMDiskLogs.0001</i>
	ISMLogs Clone	0001-9999	<i>hostname.ISMLogsClone.0001</i>

Use the **ism_op -label** *device* command to label a volume. After a volume is labeled and mounted, it is available for backups. You can label several volumes at once without dismissing the dialog box, if you prefer. For details, see “ism_op” on page A-9.

Using Expiration Dates for Volume Labels

Whenever you label or relabel a storage volume, ISM automatically assigns a two-year expiration date to the volume. This expiration date applies to the physical media and is different from the retention period that you assign to your data. When a volume reaches its two-year expiration date, ISM marks the Backup status for the volume as Disabled. To minimize the failure rate of older media, ISM cannot write new data to the volume. The ISM server still

tracks the data on the volume and applies the retention period to the save sets on the volume. Until the volume is relabeled, you can recover the data. When you relabel a storage volume, it receives a new two-year expiration date.

Managing Storage Volumes

ISM requires that you label and mount a volume before you can use it to back up data. Labeling a volume associates the volume with a specific volume pool and provides ISM with a unique name for tracking and recognizing the media. ISM uses the volume labels to maintain a record of each volume in the ISM catalog. ISM uses the ISM catalog to identify volumes that it needs to back up and recover data.

If you have the **Automatic Media Management** option enabled for a storage device, you can have ISM automatically label and mount volumes on that storage device. For more information, see “Automating Media Management for Tape and Optical Devices” on page 3-9. You can also prelabel a volume without mounting it.

Tip: You can put several save sets on one volume if they are intended for the same pool and space exists on the volume.

Labeling a Volume

When you label a volume, you associate the volume with a specific volume pool, which means that the volume is eligible to receive data when ISM needs a volume from that pool.

To label a blank volume:

1. Load the volume in a storage device.
2. To label a blank volume, use the following command:

```
ism_op -label device -pool pool -volume name
```

 - a. Specify the pathname of the device.
 - b. Use the ISMData and ISMData Clone pools for storage spaces and the ISMLogs and ISMLogs Clone pools for logical logs.
If you use a file-type device, use the ISMDiskData and ISMDiskLogs pools to create file-type storage volumes.
 - c. Specify the volume name.
 - d. Optionally, use the **-set_capacity** option to set the estimated capacity for the volume:

For *tape volumes*, if you do not set the capacity explicitly, the ISM server uses a default setting based on the media type. ISM writes data to the tape volume, however, until the volume is full, regardless of the estimated capacity that is set.

For *file-type* and *optical disk volumes*, if you do not set a value, ISM writes data to the file-type device or optical device until the device is full. If you enter a value, however, ISM writes data to the file-type or optical disk device until that value is reached.

Mounting a Volume

Before you can back up or recover data from a volume, you must mount the volume on a storage device that the ISM server recognizes.

If you use a file-type storage device, you must still mount a labeled volume on the file-type device before ISM can write backups to the file-type device. The labeling and mounting processes work the same way for file-type volumes as for tape volumes and optical disks.

To mount a volume in a specific device:

1. Load the storage volume in the device.
2. Label the storage volume (see “Labeling a Volume” on page 4-8).
3. To mount the storage volume, use the following command:

```
ism_op -mount device
```

The Mount status of the storage volume changes to Mounted, and the volume name and pool name appear in their respective columns.

If you try to initiate an ON-Bar backup when a volume is not mounted, ISM requests a writable volume with the following message:

```
media waiting: backup to pool 'ISMDData' waiting for 1 writable  
backup tape or disk
```

You receive one of three messages to request that you mount a volume, relabel a volume, or label a new volume.

If you are recovering data, ISM requests the volume that you need to mount in the device by its volume name, as the following example shows:

```
media waiting: recover waiting for 8mm 5GB ISMDData.0034
```

Important: To perform an unattended backup with a stand-alone device, you must premount the volume if the storage device does not have the **Automatic Media Management** option enabled. The storage device must also be enabled.

Choosing Volumes for Mounting

When ISM requires a volume for backup, it uses the following hierarchy to select volumes that are not currently in use from the appropriate volume pool:

1. Already mounted storage volume available for backups

2. Already mounted storage volume that is eligible for recycling on a device with the **Automatic Media Management** option enabled
3. Unmounted, unlabeled storage volume in a device with the **Automatic Media Management** option enabled
4. Unmounted storage volume that is available for backups
5. Unmounted storage volume that is eligible for recycling in a device with the **Automatic Media Management** option enabled

Unmounting Volumes

You can only unmount a storage volume when the storage device is not active.

To unmount a volume from a specific storage device:

1. To unmount the storage volume, use the following command:

```
ism_op -unmount device
```

The Mount status of the storage volume changes to Unmounted.
2. Unload the volume from the backup device.

Finding a Volume Name

ISM can detect whether a selected storage device has a volume already loaded and whether the volume already has an ISM label. If the physical label on the volume is missing or illegible, perform the following steps to determine its name.

To detect a storage volume:

1. Load the volume in a storage device.
2. To detect the storage volume, use the following command:

```
ism_op -detect device
```

ISM displays the volume name and pool name.

Relabeling a Volume

ISM uses volume labels to track your data. Relabeling a volume means that you can no longer recover the original data on the volume. To relabel a volume, use the **ism_op -label** command. You can relabel a volume only if it is unmounted. If you delete a volume from the ISM catalog, it is recommended that you do not use the deleted volume name for a new volume unless you are sure that you will never need to recover the data on the original volume.

You can relabel a volume whenever you want to. But to ensure that your data is protected, we recommend that you relabel a volume only under the following conditions:

- When a volume is eligible for recycling
- Before a volume is eligible for recycling

- When the volume is a file-type volume
- When you add more volumes to a pool

Relabeling a Volume Eligible for Recycling

When the retention periods for all the save sets on a volume expire and their retention status is marked Expired, the recycling status of the storage volume changes to Eligible. At this point, you know that the retention period set for the ISM server for all the data on the volume has expired. If you use the **Automatic Media Management** option, ISM automatically relabels the volume and overwrites it with new data. However, the new label uses the same name and sequence number as the original label. Therefore, newer data might be on a volume with a lower label number. ISM labels recycled volumes in this way so that you can easily determine the age of your volumes. Each time that a volume is relabeled, it receives a new two-year expiration date. Lower-numbered volumes typically are older than higher-numbered volumes, even though the data on the older volume might be more recent.

Eligible volumes are recycled within the same pool. For example, an eligible volume from the ISMData pool remains in the ISMData pool when it is relabeled. If you want to recycle a volume to a different pool, you must manually label the volume and change the volume pool selection. For more information, see “Recycling a Volume or Adding More Volumes to a Pool” on page 4-12.

Relabeling a Volume Before It Is Eligible for Recycling

You can manually relabel a storage volume at any time. If you relabel a storage volume before it is eligible for recycling, the retention period for some of the save sets on the volume have not yet expired, and you risk losing valuable data. Once a storage volume is relabeled, you can no longer recover the original data.

Warning: If you want to relabel a storage volume before its recycling status changes to Eligible, be sure that you do not need any of the save sets on the volume for a restore operation. Relabeling a storage volume overwrites the data on the volume, and you cannot recover the original data.

Relabeling a File-Type Volume

If you use a file-type device, you need to relabel the file-type volume mounted on the file-type device when the volume fills with data. ISM tracks file-type volumes in the same way that it tracks other volumes. If you relabel the file-type volume, you can no longer recover the original data. To ensure that the data is protected, we recommend that you clone the data on your file-type volume to tape *before* you relabel the volume. ISM tracks cloned

volumes in the same way that it tracks source volumes. When you clone your data to tape before you relabel a file-type volume, you ensure that you can recover the data.

After the file-type volume is relabeled, it appears to the ISM server that a new file-type volume is available to be mounted.

For more information on cloning volumes, see “Cloning Storage Volumes” on page 4-15.

Recycling a Volume or Adding More Volumes to a Pool

Both recycling eligible volumes and adding more volumes to a pool have advantages. When you recycle, you reuse the same volumes and do not add new volumes to the pool. However, the media might wear out over time and have a higher failure rate. ISM uses the original volume sequence number to relabel volumes within the same volume pool. This number enables you to easily track which volumes are older. When you relabel a volume to a different volume pool, the volume receives the next available sequence number for the new pool. Because each volume receives a new two-year expiration date when it is relabeled, you cannot determine how old a specific volume is when it is relabeled to a different volume pool.

On the other hand, if your site requires you to maintain the backups in storage for a certain period of time, you might have to add more media to the volume pool instead of recycling it. For example, if your volumes fill up in 90 days but your retention period is set for 180 days, you need to add new media to the pool until you can recycle the backup volumes that contain expired or old backups.

Overriding Automatic Media Management

ISM allows you the flexibility to change how ISM manages a storage volume, if necessary.

Disabling a Storage Volume

Typically, a storage volume can be used for both backup and restore operations. Occasionally, you might decide that a particular storage volume should not receive any more new data. In this case, you can change the Backup status to Disabled, which prevents ISM from using the storage volume during a backup operation. ISM uses this storage volume only for restore operations until you change the volume status.

If a storage volume fills with data during a backup operation, the ISM server automatically changes the Backup status to Disabled. This action ensures that the ISM server does not try to use the volume for subsequent backup operations. Whenever a storage volume gets full, ISM sets the Backup status

to Disabled until that volume is recycled. For the procedure, see “Enabling and Disabling a Storage Device” on page 3-4.

Important: The data on the volume with a Backup status of Disabled is still subject to the retention period that is set for the ISM server. The recycling status of the volume changes to Eligible after all the save sets on the volume have expired, and the volume is then eligible for automatic recycling.

Protecting a Storage Volume from Automatic Recycling

Occasionally, you might decide that you want to keep the data on a particular volume longer than the retention period for the ISM server specifies or because you created a different volume recycling plan. To override the retention period for the data on a volume, change the recycling status of a volume to Manual. The volume is not a candidate for automatic recycling until you *manually* relabel the volume or use the `ism_config -recycle manual` command. ISM continues to track the storage volume in its ISM catalog, and you can always recover the data on the volume with the `ism_catalog -recreate_from` command.

Deleting Volumes from the ISM Catalog

Removing a volume means that ISM no longer tracks the volume in the ISM catalog and effectively releases the volume from the control of ISM. Remove a volume from the ISM catalog only if you will never need the data on the volume again or if the volume has been physically damaged and is unusable.

You rarely need to remove a volume from the ISM catalog. ISM catalog entries take up a small amount of disk space. Leave the volume in the ISM catalog so that you do not accidentally label another volume with the same name.

If a volume is lost or destroyed, remove its entry from the ISM catalog. Otherwise, ISM might request the volume to recover data.

To delete a storage volume from the ISM catalog:

1. Remove the storage device as follows:
`ism_rm -volume name`
2. To verify that the storage volume was deleted, issue the command:
`ism_show -volumes`
The ISM server no longer tracks the storage volume and its associated save sets.

Important: If you accidentally delete a volume from the ISM catalog, you can still recover the data with the `ism_catalog -recreate_from` command as long as you have not overwritten the volume with new data.

Cloning Storage Volumes and Save Sets

Cloning volumes or save sets is beneficial for the following purposes:

- Sending data off-site
- Replacing old volumes
- Testing the integrity of your data

If you use a filesystem as a storage device, clone your data to tape when the file-type volume fills with data. This action ensures that your data is available for recovery if your file-type volume fills quickly.

Create clones of your media to store identical copies of your data off-site, or share or transfer important data to other company facilities easily.

Most media has a limited lifespan. Consequently, routinely replace your volumes. The expiration date of a storage volume indicates how old the physical media is, not the data on the volume. A volume has a default expiration date of two years after it is labeled. If you relabel a volume, it receives a new two-year expiration date. When a volume reaches its expiration date, ISM automatically changes the Backup status of the volume to Disabled. This status means that no new data can be written to the volume. However, the data on the volume can still be recovered until the volume is relabeled. If you have old volumes that contain important data, clone them before they expire to ensure the safety of the data. If you clone a volume to newer media, you can prevent media failures during recover operations. You can determine the expiration date of your physical media with the **ism_show -volumes** command at the command line.

When you clone a volume, you can determine the integrity of backed-up data. ISM displays an error message if it has difficulty reading the data from the original volume during the cloning process.

You can clone entire volumes from the command line. Because save sets can span multiple volumes, ISM might request more than one source volume to complete the cloning operation. Furthermore, because media varies in size and capacity, you might need more destination volumes to hold the cloned data than the source volumes require. Conversely, depending on the type of media that you use, you might be able to fit multiple source volumes on one destination volume.

Important: You cannot perform disaster recovery operations from cloned volumes. For instructions on how to recover the ISM catalog after a disaster, see “Recovering Data After a Disaster” on page 5-3.

Tracking Cloned Data

Cloned data on a volume is indistinguishable from the original, although the different capacities of storage media can mean that more or fewer volumes are required. Data-tracking information about each cloned save set is added to the existing entries that describe the source save sets. If administrators go outside of the ISM application and duplicate their volumes with a copy utility that the operating system or hardware provides, ISM has no knowledge of these volumes and cannot recover data stored on such volumes.

Both original and cloned volumes share the same ISM catalog entries, so the cloned data is subject to the same retention period as the original data. ISM automatically manages both original and cloned data based on the retention periods for the volumes. If you clone data a week before the retention period for the source data expires, the retention status for the cloned data changes to Expired at the same time that the original data does. If you want to keep data on a cloned volume beyond the retention period of the original data, you can change the recycling status of the cloned volume to Manual. ISM can keep and track the data on the cloned volume indefinitely.

To distinguish between cloned and original volumes, you must assign the cloned volume to a clone pool and include *clone* in the volume name. ISM provides preconfigured ISMData Clone and ISMLogs Clone pools for you to use when you clone your storage spaces and logical logs.

Cloning Storage Volumes

Cloning with ISM provides an additional layer of protection because it allows you to make identical copies of the data on your volumes as you need them.

The original volume is the *source* volume, and the volume that receives the cloned data is the *destination* volume. The destination volume must be a different volume from the source volume. Cloned data belongs to a clone pool. Use the ISMData Clone pool and the ISMLogs Clone pool for your cloned volumes.

The requirements for cloning storage volumes with ISM follow:

- Two backup devices
- Source and destination volumes

To clone a file-type storage volume before it is full:

1. If the volume is already mounted, you must unmount it to change the Backup status. This action ensures that no new data is written to the file-type volume during the cloning operation.

If the volume is already unmounted, change the Backup status for the volume to Disabled.

2. To clone a volume, use the following command:

```
ism_clone -volume name
```

- a. Specify the name of the volume to clone.
 - b. Specify the name of the clone pool (Default Clone, ISMData Clone, or ISMLog Clone).
3. Mount the selected storage volume and cloned volume.

Cloning Save Sets

In addition to cloning all the save sets on a storage volume, you can clone individual save sets. ISM tracks these individual cloned save sets in the same way that it tracks the originals. The cloned save sets are subject to the same retention period as the original save sets and must be written to a volume that belongs to a clone pool. When the original save sets expire, so do the cloned save sets.

To clone a save set:

1. Mount a storage volume from the appropriate clone pool on a storage device.
2. To display the save set IDs for a volume, use the following command:

```
ism_show -volume name
```
3. To clone a save set, use the following command:

```
ism_clone -saveset ssid
```

 - a. Specify one or more save set IDs.
 - b. Specify the name of the clone pool.
4. Click the **Clone** button in the toolbar or choose **Save Set > Clone**.
5. Mount the selected storage volume and cloned volume.

Chapter 5. Recovering Save Sets with ISM

Restoring a Save Set with ON-Bar	5-1
Locating Volumes That Contain Partial Save Sets	5-1
Re-Creating the ISM Catalog Before You Use ON-Bar	5-1
Recovering Data from a Cloned Volume	5-2
Recovering Data After a Disaster	5-3
If You Have the ISM Server Bootstrap Printout	5-3
If You Do Not Have the ISM Server Bootstrap Printout	5-4
Creating and Backing Up the ISM Bootstrap	5-4
Transferring Data with the Imported Restore.	5-5
Step 1: Backing Up the Source Database Server to ISM	5-5
Step 2: Porting ISM to Your Target Database Server	5-8
Step 3: Restoring Data from ISM to the Target Database	5-13

In This Chapter

This chapter describes how to recover data at various points during the data life cycle.

Restoring a Save Set with ON-Bar

As long as the retention status for a save set has not changed to Expired to indicate that the retention period for the save set has passed, you can use ON-Bar to initiate a restore operation for that save set. For instructions on how to initiate a restore operation with ON-Bar, refer to the *IBM Informix: Backup and Restore Guide*.

Locating Volumes That Contain Partial Save Sets

Depending on the size of your storage spaces, a save set might require more than one storage volume during a backup. The **Partial** column in the **ism_show -volume** command indicates if a save set spans more than one storage volume. Typically, save sets do not span more than two volumes. If you need to recover a save set, you might want to preload all the volumes before you begin the restore operation.

Re-Creating the ISM Catalog Before You Use ON-Bar

The ISM server tracks the backed-up data in the ISM catalog. You can re-create entries in the ISM catalog for expired save sets and for volumes that have been deleted from the ISM catalog from the command line.

After you re-create entries in the ISM catalog, you can use ON-Bar to recover the data. For details on how to use the ISM command-line utilities, see Appendix A, “ISM Command-Line Utilities.”

Before you use ON-Bar to recover data, use the **ism_catalog -recreate from** command to re-create the entries in the ISM catalog in the following situations:

- Recovering expired save sets
After a save set has expired, you need to re-create its entry in the ISM catalog before you can restore the data with ON-Bar.
- Recovering volumes that are eligible for recycling
When the Recycling status for a storage volume has changed to Eligible, the retention periods for all the save sets on the storage volume have expired. You need to re-create the entries in the ISM catalog before you can restore the data with ON-Bar.
- Recovering volumes that have been deleted from the ISM catalog
If you remove a storage volume from the ISM catalog, but the data on the volume is still intact, you can re-create the entries in the ISM catalog and then use ON-Bar to recover the data on the volume.

If the ISM server is on UNIX, you must have **root** privileges to execute the **ism_catalog** command. If the ISM server is on Windows, you must be a member of the **Informix-Admin** group.

To recover entries in the ISM catalog:

1. Load the volume in a storage device.
2. At the command line, initiate the following command with the full pathname of the storage device on which you loaded the storage volume:

```
ism_catalog -recreate_from device
```
3. Use ON-Bar to restore the data.

When you recover a storage volume in this way, ISM automatically marks the Backup status for the volume as *Disabled* to ensure that no new data is written to the storage volume. To protect the volume from automatic recycling, change the recycling status of the volume to Manual. ISM does not automatically recycle a volume marked Manual.

Recovering Data from a Cloned Volume

Cloned volumes provide added security because they allow you to recover your data in case the original volumes become damaged or destroyed.

During a restore operation, ISM is configured to request the original source volume. However, if you know that the original volume is unavailable (perhaps it is stored off-site), you can force ISM to request a cloned volume.

To force ISM to request a cloned volume:

1. Use the **ism_rm -volume** command to delete the original storage volume. Alternatively, change the Restore status for each save set on the original volume to Disabled:

```
ism_config -volume name -disable_restore ssid
```

To obtain a list of the **ssid** numbers, issue the **ism_show -volume *name*** command.

2. Use **ism_catalog** to recover unexpired save sets.
If the save sets on the cloned volume have expired, follow the instructions in “Re-Creating the ISM Catalog Before You Use ON-Bar” on page 5-1 to re-create the ISM catalog.
3. Use ON-Bar to restore the data.

When you delete the original volume from the ISM catalog or change the Restore status for each save set on the original volume to Disabled, ISM knows that the original volume is unavailable for the restore operation and requests the cloned volume instead.

Recovering Data After a Disaster

Hardware malfunctions rarely occur at convenient times, but if you have a disaster recovery plan in place and maintain regular backups of your database server, you are well prepared to recover critical data in a timely manner.

Before you can restore your data with ON-Bar, you must re-create the ISM catalog in the following situations:

- The ISM catalog is corrupted
- You reinstalled the operating system and the Informix database server software, which includes the ISM software

If You Have the ISM Server Bootstrap Printout

The following steps overwrite all existing information for this ISM server and recover the ISM catalog.

To recover the ISM catalog:

1. Locate the most recent bootstrap printout for the ISM server.
The bootstrap printout lists which volume contains the ISM server bootstrap save set.

2. Mount the volume that contains the most recent bootstrap save set on a storage device and issue the following command:

```
ism_catalog -recover
```

ISM re-creates the ISM catalog for this ISM server.
3. After you re-create the ISM catalog, use ON-Bar to restore your data.

If You Do Not Have the ISM Server Bootstrap Printout

If you do not have the most recent ISM server bootstrap printout, you must locate the storage volume that contains the most recent bootstrap save set before you can recover the ISM catalog.

To recover the ISM catalog:

1. Mount the storage volume that contains the most recent backup from the ISMData volume pool.
This volume contains the most recent ISM server bootstrap save set.
If you do not know which volume contains the most recent backup from the ISMData volume pool, you must scan each volume for the most recent bootstrap save set. At the command line, enter the following command with the full pathname of the storage device on which you loaded the storage volume:

```
ism_catalog -find_bootstrap device
```

ISM displays the **ssid** numbers for all bootstrap save sets on the volume. You can use bootstraps to restore the ISM catalog only from volumes of type Backup. You cannot use bootstraps from volumes of type Backup Clone in the restore process.
2. After you locate the volume with the most recent bootstrap save set, issue the following command:

```
ism_catalog -recover
```

ISM re-creates the ISM catalog for this ISM server.
3. After you re-create the ISM catalog, use ON-Bar to restore your data.

Important: You cannot perform disaster recovery from a cloned bootstrap save set.

Creating and Backing Up the ISM Bootstrap

To create the ISM bootstrap, use one of the following methods:

- **onbar** script

If you use the **onbar** script to perform backups, it automatically calls the **ism_catalog -create_bootstrap** command, which creates and backs up the ISM bootstrap. For more information, see “Updating the onbar Script” on page 1-11.

- **ism_catalog -create_bootstrap** command

For more information, see “ism_catalog” on page A-2.

Important: We recommend that you back up the ISM bootstrap after each level-0 backup.

Transferring Data with the Imported Restore

With the imported restore feature, you can transfer all the data from one instance of the database server to the same instance on a foreign host. For example, you can back up data to ISM on one computer and restore it to ISM on a different computer. You can perform imported restores using either whole system (serial) or storage-space (parallel) backups. The parallel backup must include all storage spaces and logical logs. You must also use compatible versions of XBSA libraries and ISM for both operations.

Important: You cannot use a backup from one database server version to restore on a different version. You can, however, change the database server name in an imported restore.

The imported restore is a subset of the following three processes:

- Disaster recovery
- Database server upgrade
- Initialization of High-Availability Data Replication (HDR)

As its name implies, you can also use the imported restore solely to complete an imported restore.

An imported restore consists of these three steps:

1. Backing up the source database server to ISM
2. Porting ISM to a target computer that is a different physical host
3. Restoring the data from ISM to the target database server with the same instance on the target computer

These procedures are discussed in detail in the following sections.

Important: The imported restore procedure is the same for both UNIX and Windows. However, when you set the environment variables, you must use the appropriate UNIX or Windows syntax.

Step 1: Backing Up the Source Database Server to ISM

Important: If the **hostname** command outputs a fully-qualified hostname, use fully qualified hostnames in the imported restore procedure. An example of a qualified hostname is **mycomputer@informix.com**. If the command outputs the short hostname (**mycomputer**, for example), use the short form.

To back up the source database server to ISM:

1. Execute the **hostname** command on your source and target computers.
2. Bring the database server to quiescent mode with the following command:
onmode -s
3. Set the following ISM environment variables:

UNIX Only

On UNIX:

```
ISM_SERVER = source_hostname
ISM_DATA_VOLUME_POOL = ISMData
ISM_LOG_VOLUME_POOL = ISMLogs
export ISM_SERVER
export ISM_DATA_VOLUME_POOL
export ISM_LOG_VOLUME_POOL
```

End of UNIX Only

Windows Only

On Windows:

```
set ISM_SERVER = source_hostname
set ISM_DATA_VOLUME_POOL = ISMData
set ISM_LOG_VOLUME_POOL = ISMLogs
```

End of Windows Only

4. Perform a full backup of the source database server:
onbar -b -L 0
5. Back up copies of the following critical files located in the **\$INFORMIXDIR/etc** (UNIX) or **%INFORMIXDIR%\etc** (Windows) directory. When you perform an imported restore to another computer system, you will need to restore these files.

Dynamic Server

On Dynamic Server, copy the following files:

```
oncfg_source_dbservername.servernum
ixbar.servernum
# for UNIX users only
$ONCONFIG
sqlhosts
# for Windows users only
%ONCONFIG%
```

End of Dynamic Server

Windows Only

On Windows, use **regedit** to copy the **sqlhosts** information from the source computer to the target computer. Use the following registry entry:
HKEY_LOCAL_MACHINE/SOFTWARE/Informix/SQLHOSTS/...

End of Windows Only

Extended Parallel Server

On Extended Parallel Server, copy the following files:

```
onconfig.sourcecomputer  
Bixbar.hostname.servernum  
# for all coserver numbers, copy the oncfg files  
oncfg_source_dbservername.servernumber.coservernumber  
sqlhosts  
xcfg_sourcecomputer.servernumber
```

End of Extended Parallel Server

6. Complete the following steps to make your security portable to the new computer:

UNIX Only

On UNIX:

```
ism_add -admin root@target_hostname  
ism_add -admin informix@target_hostname
```

End of UNIX Only

Windows Only

On Windows:

```
ism_add -admin Administrators@target_hostname  
ism_add -admin informix@target_hostname
```

End of Windows Only

7. Save the bootstrap of your source ISM server:

UNIX Only

On UNIX:

```
$INFORMIXDIR/bin/ism_catalog -create_bootstrap \  
-pool $ISM_DATA_VOLUME_POOL
```

End of UNIX Only

Windows Only

On Windows:

```
cd %INFORMIXDIR%\bin  
ism_catalog -create_bootstrap -pool %ISM_DATA_VOLUME_POOL%
```

End of Windows Only

Step 2: Porting ISM to Your Target Database Server

You need to port and replicate ISM on the target database server.

To port ISM to your target database server:

1. Log in as user **root** (UNIX) or **Administrator** (Windows) on the target database server.
2. Shut down the Informix database server and ISM.

Dynamic Server

On Dynamic Server:

```
onmode -ky  
ism_shutdown
```

End of Dynamic Server

Extended Parallel Server

On Extended Parallel Server:

```
xctl onmode -ky  
ism_shutdown
```

End of Extended Parallel Server

3. Unset the **ISM_CLIENT** environment variable as follows:

```
unset ISM_CLIENT
```

4. Edit the **ism_startup** script or **ism_startup.bat** script:

UNIX Only

On UNIX:

- a. Locate the following line in the `$INFORMIXDIR/bin/ism_startup` script and type **hostname** verbatim:

```
hostname=$(hostname)
```

- b. Add the following line immediately after that line:

```
hostname=source_hostname
```

End of UNIX Only

Windows Only

On Windows:

- a. Locate the following line in the `%INFORMIXDIR%\bin\ism_startup.bat` script:

```
Set nextLabel = ask_init
```

- b. Add the following line immediately before that line:

```
set hostname = source_hostname
```

End of Windows Only

5. Reinitialize ISM:

```
ism_startup -init
```

6. Using a text editor, create a temporary file and enter the following information:

UNIX Only

On UNIX:

```
create type: NSR client; name: source_hostname; remote  
access: root@target_hostname, informix@target_hostname
```

End of UNIX Only

Windows Only

On Windows:

```
create type: NSR client; name: source_hostname;  
remote access: Administrator@target_hostname,  
informix@target_hostname
```

|----- End of Windows Only -----|

7. Run the following command against the temporary file:

```
nsradmin -s target_hostname -i temporary_file
```

Important: After you run this command, you might receive an error message similar to the following message. Ignore this message and complete the procedure as outlined.

```
create failed: save set of All for client source_hostname in group  
Default excludes any clients with the same name from the group.
```

To replicate ISM with the same resources that you have on the source host, you must recover your bootstrap:

1. If your source **INFORMIXDIR** does not match your target **INFORMIXDIR**, you need to create a symbolic link to recover the bootstrap from the source computer.

For example, if **INFORMIXDIR** on your source computer is **/usr2/informix** and **INFORMIXDIR** on your target computer is **/usr/local/informix**, create the **/usr2** directory on the target computer and symbolic link as follows:

|----- UNIX Only -----|

On UNIX:

```
mkdir /usr2  
ln -s /usr/local/informix /usr2/informix
```

|----- End of UNIX Only -----|

|----- Windows Only -----|

On Windows:

Create a shortcut from your source **INFORMIXDIR** to your target **INFORMIXDIR**. For information on how to create a shortcut, see your operating-system documentation.

|----- End of Windows Only -----|

2. Add the device that contains the volume pool that contains the bootstrap file:

```
ism_add -device device_name -type device_type
```
3. Add the remaining devices for the other volume pools.
4. Mount the storage volume that contains the most recent backup from the ISMData volume pool.

5. Mount the remaining storage volumes.
6. Scan the volume for the most recent bootstrap save set:

```
ism_catalog -find_bootstrap device_name
```

ISM displays the **ssid** numbers for all bootstrap save sets on the volume. Record the most recent **ssid** number because you will need to use this number again in step 8.
7. Recover the most recent bootstrap save set:

```
ism_catalog -recover
```
8. Enter the most recent bootstrap save set number you recorded earlier in step 6.
Wait until the following message appears. Ignore the warning and proceed to the next step:

```
nsrindexasm: Didn't find earlier versions to recover!  
WARNING: The online index for hostname was NOT fully  
recovered...
```
9. Shut down ISM:

```
ism_shutdown
```
10. Recover the **res.R** file from the source database server.
The **res.R** file contains the resource configuration information for the source database server.

UNIX Only

On UNIX:

```
cd $INFORMIXDIR/ism  
rm -rf res  
mv res.R res
```

End of UNIX Only

Windows Only

On Windows:

```
cd %ISMDIR%  
del res  
rmdir res  
move res.R res
```

End of Windows Only

Important: Perform step 11 for disaster recovery only if you did not issue the **ism_add -admin** command on the source computer because it was unavailable.

11. Create the following temporary file and then run the **nsradmin** command against that file.

UNIX Only

On UNIX:

```
print type:NSR;
```

```
update administrator:informix@source_host,  
root@source_host,informix@target_host,  
root@target_host
```

```
nsradmin -f $INFORMIXDIR/ism/res/nsr.res -i temporary_file
```

End of UNIX Only

Windows Only

On Windows:

```
print type:NSR;
```

```
update administrator:informix@source_host,  
Administrator@source_host,informix@target_host,  
Administrator@target_host
```

```
nsradmin -f %ISMDIR%\res\nsr.res -i temporary_file
```

End of Windows Only

12. Start ISM:

```
ism_startup
```
13. Remove the device entries from the source computer that do not exist on the target computer:

```
ism_show -devices  
ism_rm -device device_name
```
14. Create any additional media device entries that do not appear in ISM to match those that exist on your target host:

```
ism_show -devices  
ism_add -device device_name -type device_type
```
15. Grant permission for local users to access the indexes for the source client.
Using a text editor, create a temporary file and enter the following information:

Important: You must enter the command exactly as it appears, including the space between the two lines.

|----- UNIX Only -----|

On UNIX:

```
print type: NSR client; name: source_hostname  
  
update remote access: informix@target_hostname,  
root@target_hostname
```

|----- End of UNIX Only -----|

|----- Windows Only -----|

On Windows:

```
print type: NSR client; name: source_hostname  
  
update remote access: informix@target_hostname,  
Administrator@target_hostname
```

|----- End of Windows Only -----|

16. Run the following command against the temporary file:

```
nsradmin -s target_hostname -i temporary_file
```

Step 3: Restoring Data from ISM to the Target Database

Important: Set up the target database server with exactly the same disk layout and device allocations as the source database from which you are transferring the data. Make sure that you assign the same database server number that was used on the source computer.

To restore the data from ISM to the target database server:

1. Log in as user **informix**.
2. If the Informix database server is online, shut it down.

|----- Dynamic Server -----|

On Dynamic Server:

```
onmode -ky
```

End of Dynamic Server

Extended Parallel Server

On Extended Parallel Server:

```
xctl onmode -ky
```

End of Extended Parallel Server

3. Configure the following ISM environment variables for this session. For example:

UNIX Only

On UNIX:

```
ISM_CLIENT = source_hostname
ISM_SERVER = target_hostname
ISM_DATA_VOLUME_POOL = same_value_on_source_host
ISM_LOG_VOLUME_POOL = same_value_on_source_host
export ISM_CLIENT ISM_SERVER
export ISM_DATA_VOLUME_POOL
export ISM_LOG_VOLUME_POOL
```

End of UNIX Only

Windows Only

On Windows:

```
set ISM_CLIENT = source_hostname
set ISM_SERVER = target_hostname
set ISM_DATA_VOLUME_POOL = same_value_on_source_host
set ISM_LOG_VOLUME_POOL = same_value_on_source_host
```

End of Windows Only

4. Recover your media database for each of your data volume pools (ISMData, ISMLogs):

```
ism_catalog -recreate_from device_name
```
5. Synchronize the media index database:

```
nsrim -v
```
6. To mark your media as appendable for each of your volume pools (ISMData, ISMLogs), complete the following steps:

```
ism_show -devices
ism_op -unmount device_name
ism_show -volumes
ism_config -volume volume_name -only_restore no
ism_op -mount device_name
```

Dynamic Server

7. On Dynamic Server, restore copies of the following critical files located in the **\$INFORMIXDIR/etc** (UNIX) or **%INFORMIXDIR%\etc** (Windows) directory:

```
oncfg_source_dbservername.servernum
ixbar.servernum
# for UNIX users only
$ONCONFIG
sqlhosts
# for Windows users only
%ONCONFIG%
```

Windows Only

Use **regedit** to copy the **sqlhosts** information from the source computer to the target computer. Use the following registry entry:

```
HKEY_LOCAL_MACHINE/SOFTWARE/Informix/SQLHOSTS/...
```

Rename the **oncfg*** file and replace the host database server name with the target database server name. For example:

```
oncfg_target_dbservername.servernum
```

End of Windows Only

End of Dynamic Server

Extended Parallel Server

8. On Extended Parallel Server, restore copies of the following critical files located in the **\$INFORMIXDIR/etc** directory from each coserver on the target computer:

```
onconfig.sourcecomputer
Bixbar.hostname.servernum
oncfg_source_dbservername.servernumber.coservernumber
sqlhosts
xcfg_sourcecomputer.servernumber
```

For the following files, replace the host database server name with the target database server name:

```
oncfg_source_dbservername.servernumber.coservernumber
xcfg_hostname.servernumber
Bixbar.hostname.servernum
```

Edit the `xcfg_hostname.servernumber` file to change the source database server name to the target database server name.

End of Extended Parallel Server

9. Update the `sqlhosts` file or registry entry and include the proper shared memory settings for the target database server.

UNIX Only

In the `sqlhosts` file, change the database server name in the first and fourth entries on each line, if necessary. Change the hostname in the third entry to the target hostname.

End of UNIX Only

10. Update the ONCONFIG file and replace the source database server name with the target database server name. For example:

```
DBSERVERNAME target_dbservername
```

Extended Parallel Server

For Extended Parallel Server only, change the hostnames in the NODE lines. Delete the SADDR, LADDR, and HADDR lines that follow each NODE line in the ONCONFIG file.

End of Extended Parallel Server

11. Create new dbspaces if you have not already done so.
Make sure they reside at the same path locations as they do on the source database server. User **informix** and group **informix** also should own these dbspaces.
12. Perform the imported restore on the target database server:

```
onbar -r
```
13. Reconfigure your ISM settings for daily backup and restore operations.
Do not change the settings of the `ISM_CLIENT` and `ISM_SERVER` environment variables.

Important: If the imported restore fails with an error about the source hostname not having access, grant permission for local users to access the indexes for the source client. For details, see steps 15 and 16 on page 5-12.

Appendix A. ISM Command-Line Utilities

The ISM server supports several command-line utilities for managing administrative users, the ISM catalog, storage devices, and media volumes. If an option contains multiple letters, you can abbreviate it to the shortest unique string. For example, you can specify the **ism_catalog -recreate_from** option as **-recr**.

The table below lists all available commands and gives a brief explanation of their function. The following sections contain a more detailed explanation of the commands.

Command	Used to
ism_add	Add administrative users and storage devices.
ism_catalog	Re-create entries in, or recover, the ISM catalog, create a new bootstrap, or find the ISM server bootstrap.
ism_chk.pl	Collect information about ISM, ON-Bar, and the database server.
ism_clone	Initiate a clone operation on a storage volume or save set.
ism_config	Configure the ISM server properties and change storage-volume parameters.
ism_op	Label, mount, and unmount storage volumes.
ism_rm	Remove an administrative user or storage device from the ISM server or remove a storage volume from the ISM catalog.
ism_show	View information about the ISM administrators, media volumes, and storage devices.
ism_shutdown	Shut down the ISM server.
ism_startup	Start an ISM server.
ism_watch	Display a terminal-based screen utility to monitor the ISM server.

ism_add

The **ism_add** command enables you to add administrative users and storage devices to the ISM server.

```
► ism_add [-server hostname] [-admin name@ host] [-device name --type type] ►
```

Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-admin <i>name@host</i>	Adds an administrative user to the ISM server	This user can make configuration changes to the ISM server.
-device <i>name</i>	Adds a storage device to the ISM server	ISM supports up to four active devices.
-type <i>type</i>	Designates the type of storage device	You must choose a device type that ISM supports. For more information, see “Adding and Deleting Storage Devices” on page 3-3.

ism_catalog

The **ism_catalog** command enables you to re-create entries in the ISM catalog for a volume that has been removed from the ISM catalog, find the ISM server bootstrap, create a new ISM bootstrap, and recover the entire ISM catalog.

Important: To run the **ism_catalog** command, you must have **root** privileges.

```
► ism_catalog [-server hostname] ►  
  
► [-recreate_from device] [-create_bootstrap] [-pool pool] [-find_bootstrap device] [-recover] ►
```

Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.

Element	Purpose	Key Considerations
-recreate_from <i>device</i>	Re-creates the ISM catalog directly from the storage media	Use to recover expired save sets or storage volumes that are eligible for recycling. To add the volume to a different volume pool from the one to which it originally belonged, you can use the -pool option.
-create_bootstrap	Creates a new bootstrap for this ISM server	This command writes a new bootstrap for the ISM server.
-pool <i>pool</i>	Designates a volume pool to which to write recovered data or a new bootstrap	None.
-find_bootstrap <i>device</i>	Scans the media looking for the start of the ISM server bootstrap save set After the media has been scanned, prints the save-set ID (<i>ssid</i>) and tape file location of the most recent bootstrap save set.	You need the bootstrap save set for disaster recovery. You use the bootstrap information to recover the ISM catalog directly from storage media with the ism_catalog -recover command. You then use Volume Inventory or the ism_show -volumes command to determine which volumes you want ON-Bar to restore.
-recover	Recovers the ISM catalog from storage media when the catalog has been lost or damaged	This command overwrites the existing ISM catalog. Before you execute this command, you need the bootstrap save-set ID that you obtained from the most recent bootstrap printout or from the ism_catalog -find_bootstrap command.


ism_chk.pl

The **ism_chk.pl** command collects information about the current state of ISM, ON-Bar processes, and the database server from various log files and utility programs. The **ism_chk.pl** report is useful when you are investigating backup or restore problems. For more information, see “Contents of the *ism_chk.pl* Report” on page A-6.

Although the **ism_chk.pl** command displays the report on the screen, you can redirect the report to a file. The **ism_chk.pl** command is a **perl** program that runs on both UNIX and Windows.

To display the usage message, enter the following:

```
perl ism_chk.pl --
```



Element	Purpose	Key Considerations
<code>-s starttime</code>	Allows you to supply a date and time range for scanning the logs Enter the date and time to begin scanning the logs.	If omitted, the default start date and time is "1980-01-01 00:00:00". Use double quotes only if you include the time in the range.
<code>-e endtime</code>	Allows you to supply a date and time range for scanning the logs Enter the date and time to stop scanning the logs.	If omitted, the default end date and time is "2999-12-31 23:59:59". Use double quotes only if you include the time in the range.
<code>-l</code>	Prints separate reports for the database server, ON-Bar, and ISM logs.	If omitted, the database server message, ON-Bar activity log, and ISM message logs are combined into a single report, ordered by the time stamp.
<code>-m maxlines</code>	Specifies the maximum number of lines for the log report	Choose this feature when the logs are large or you need only certain parts of the report. The default is 999999 lines.

Element	Purpose	Key Considerations
-w	<p>Prints a report with a table of contents formatted in HTML for display with a web browser</p> <p>If you print the basic report with the -w option, the following table of contents is displayed with links to various sections of the report:</p> <ul style="list-style-type: none"> • Environment variables • Log file list • Informix database server log • ON-Bar activity log • ON-Bar debug log • ISM daemon log • ISM message log • ISM savegrp log • ISM XBSA log • Combined logs 	<p>Set on to use HTML format. The default is HTML off.</p> <p>If you print the report with the -w, -x, and -c options, the following table of contents displays:</p> <ul style="list-style-type: none"> • Environment variables • Log file list • Informix database server log • ON-Bar activity log • ON-Bar debug log • ISM daemon log • ISM message log • ISM savegrp log • ISM XBSA log • Combined logs • ISM save set report • Bootstrap verification report • sysutils generated bootstrap • onstat output • Network status report • ISM resource report • Directory of executables
-c	Includes information on the ISM catalog, volume, and ixbar in the report	The default is off.
-x	Includes information on the network status and onstat output in the report	The default is off.
-v	Displays progress messages while the report is being constructed and sends the report to the screen	Set on to use verbose mode. The default is quiet mode that sends the report to the screen.

Specifying the Start and End Times for the Log Report

If you select a narrow time interval for the **ism_chk.pl** report, you can more easily find the information that you need in the various logs. For example, to produce a report on the logs for a time period from 8:00 A.M., June 1, 1999 to 9:00 P.M., June 2, 1999, specify:

```
ism_chk.pl -s "1999-06-01 08:00:00" -e "1999-06-02 21:00:00"
```

A shorthand way to specify the start and end times is:

```
ism_chk.pl -s "1999-06-01 08" -e "1999-06-02 21"
```

For example, to produce a report on the logs for all day June 3, 1999, specify:
`ism_chk.pl -s 1999-06-03 -e 1999-06-04`

Contents of the `ism_chk.pl` Report

Table A-1 describes the contents of each section of the `ism_chk.pl` report. To specify the date and time range for the log information, use the `-s` and `-e` options.

Table A-1. Contents of the `ism_chk.pl` Report

Report Section	Description
Environment variables	Lists the important environment variables that affect oninit , ISM, and ON-Bar.
Log file list	Lists the full path names of the various logs that contain information about ISM and ON-Bar operations.
Database server log ON-Bar activity log ON-Bar debug log ISM daemon log ISM message log ISM savegrp log ISM XBSA log	Lists the text from these logs.
ISM save set report	Lists all the save sets currently in the ISM catalogs. Use the <code>-c</code> option to print this report.
Bootstrap verification report	Verifies that the dbspaces and logical logs that are in the ON-Bar bootstrap file are also in the ISM catalog. Use the <code>-c</code> option to print this report.
sysutils generated bootstrap	Re-creates the ON-Bar bootstrap file that is generated from the ON-Bar catalogs. Use the <code>-c</code> option to print this report.
onstat -a output	Displays the output of onstat -a that shows the state of the database server. Use the <code>-x</code> option to print this report. For more information, see the onstat section of the utilities chapter in the <i>IBM Informix: Administrator's Guide</i> .
network status report	Displays the output of the netstat utility that reports on network load and resources. Use the <code>-x</code> option to print this report.
ISM resource report	Lists the ISM resource file that contains definitions for storage devices and volumes, and the retention policies for stored objects. Use the <code>-x</code> option to print this report.

Table A-1. Contents of the *ism_chk.pl* Report (continued)

Report Section	Description
directory of executables	Lists the contents of the Informix executables directory. Use this output to verify the software version. Use the <code>-x</code> option to print this report.

Sample *ism_chk.pl* Verbose Output

The following sample output is from the *ism_chk.pl* command with the `-v` verbose flag set. The `-m` option limits the report to the last 10 lines of the logs. The report is written to an ASCII file, *my_ism_report*.

```
$ perl ism_chk.pl -v -m 10 > my_ism_report
```


```
ism_chk: checking ISM/Onbar/Online logs
-s 1980-01-01 00:00:00
-e 2999-12-31 23:59:59
-m 10 lines will be displayed
-v report output selected
```

```
checking ISM resources...
checking ONCONFIG ONCONFIG.ol_mazama28...
checking Online log...
checking Bar Activity Log...
checking Bar Debug Log...
checking ISM daemon log...
checking Message log...
checking Savegrp log...
checking ISM catalog for ssid's...
checking Onbar catalog for ISM ssid's...
checking Onstat...
checking Network status...
```

```
ism_chk: all done now.
```

ism_clone

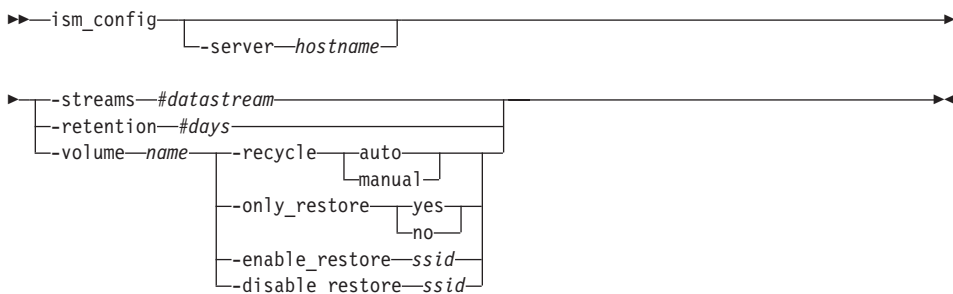
The *ism_clone* command allows you to initiate a clone operation on a storage volume or specific save set from the command line.



Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-volume <i>name</i>	Clones one or more storage volumes	You need two storage devices to clone a volume, one for the source volume and one for the clone volume.
-saveaset <i>ssid</i>	Clones one or more storage-specific save sets	When you clone save sets, you need the save-set ID (<i>ssid</i>). To obtain a save-set ID, execute the ism_show -volume name command.
-pool <i>pool</i>	Designates a volume pool	Clone volumes must belong to a clone pool.

ism_config

The **ism_config** command allows you to configure the ISM server properties, such as the number of active data streams and the retention period for the ISM server. You can also change storage-volume parameters such as overriding the retention period of the data on the volume, and enabling or disabling a save set for use in a restore operation.




Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-streams <i>#datastreams</i>	Changes the number of simultaneous data streams for this ISM server	You can configure up to four active data streams for your ISM server.

Element	Purpose	Key Considerations
-retention <i>#days</i>	Changes the retention period for all data that this ISM server tracks	Resetting the retention period applies the new retention period to all the data that this ISM server tracks. Changing the retention period could make some data eligible for automatic recycling sooner or later than you expected.
-volume <i>name</i>	Changes the parameters for a storage volume or save set	The -recycle and -backup parameters are valid only for storage volumes. The -enable_restore and -disable_restore parameters apply only to save sets and must be applied to save-set IDs.
-recycle <i>auto manual</i>	Changes the recycle status for a storage volume	The -recycle parameter is valid only for storage volumes. A volume with a recycle status of <i>auto</i> is eligible for automatic recycling. A volume with a recycle status of <i>manual</i> is protected from automatic recycling.
-only_restore <i>yes no</i>	Specify <i>yes</i> to change the mode of a volume to nonappendable. You can use this volume only for restores, not for backups. Specify <i>no</i> to change the volume to appendable. You can use this volume for backups.	The -only_restore parameter is valid only for storage volumes. It applies to all save sets in the volume. You can write save sets to an appendable volume but not to a nonappendable volume.
-enable_restore <i>ssid</i>	Enables the restore status for a save set	The -enable_restore parameter is valid only for save sets. To obtain a save-set ID, run the ism_show -volume name command.
-disable_restore <i>ssid</i>	Disables the restore status for a save set	The -disable_restore parameter is valid only for save sets. To obtain a save-set ID, run the ism_show -volume name command.

ism_op

The **ism_op** command enables you to perform operations such as labeling, mounting, and unmounting storage volumes.



Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-label <i>device</i>	Labels or relabels a storage volume	Storage media must be loaded in a storage device. When you relabel a volume, you can no longer recover the original data.
-volume <i>name</i>	Designates a volume name	When you relabel a volume within the same volume pool, ISM uses the same label and sequence number unless you designate a new volume name.
-pool <i>pool</i>	Designates the pool to which the storage volume should belong	When you relabel a volume, it still belongs to the same volume pool, unless you specify a different pool with the -pool option.
-force	Forces a relabel operation on a volume that has a recycle status of <i>manual</i> or contains save sets whose retention periods have not yet expired	Once a volume is relabeled, the original data can no longer be recovered.
-detect <i>device</i>	Displays the volume name and pool name for a volume that is loaded in a storage device	The volume cannot be mounted on the storage device.
-mount <i>device</i>	Mounts a storage volume on the specified storage device	Storage media must be labeled.
-unmount <i>device</i>	Unmounts a storage volume from the specified storage device	None.

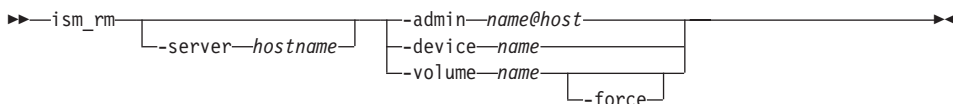
Note: The **-volume** and the **-pool** options are required when specifying a device that is being used for the first time.

For example:

```
ism_op -label /ismtape -volume testing1 -pool default
```

ism_rm


The **ism_rm** command enables you to remove an administrative user or storage device from the ISM server. You can also remove a storage volume from the ISM catalog.



Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-admin <i>name@host</i>	Removes an administrative user from the ISM server	If you remove yourself as an administrative user, you can no longer administer this ISM server.
-device <i>name</i>	Removes a storage device from the ISM server	You cannot remove a storage device that is currently in use.
-volume <i>name</i>	Removes a storage volume from the ISM catalog	ISM no longer tracks the data on this storage volume. You can still recover data on the volume with the ism_catalog -recreate_from command.
-force	Forces a storage volume to be removed from the ISM catalog that is marked Manual recycle or a volume that contains save sets whose retention periods have not yet expired	ISM no longer tracks the data on this storage volume.

ism_show


The **ism_show** command enables you to view information about the ISM administrators, media volumes, and storage devices.



Element	Purpose	Key Considerations
-server <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.
-admins	Shows all administrators for this ISM server	None.
-config	Shows all server parameters for this ISM server	None.
-devices	Displays all storage devices that this ISM server recognizes and the volumes that are currently mounted on the devices	None.
-volumes	Displays the name of each volume in the ISM catalog, the volume pool to which each volume belongs, whether the volume can be used only in restore operations, the number of bytes written to each volume, the percent of space used (or the word <i>full</i> to indicate that the volume is filled to capacity), and the volume expiration date	<p>Volume recycling status is reported in the first column of the report:</p> <ul style="list-style-type: none"> • E indicates that the volume is eligible for recycling. • M indicates that a volume has a Manual recycle status. • X indicates that a volume is eligible for recycling and has a Manual recycle status. <p>If a volume cannot be used in a backup operation, an -R flag appears in the fourth column of the display.</p>
-volume <i>name</i>	Displays the save-set IDs for complete save sets, the save-set name, whether a save set has been marked suspect, the total bytes in the save set, and the save-set creation date for all save sets on a volume	If a save set has been marked suspect, an -S flag appears in the third column of the display.

ism_shutdown

The **ism_shutdown** command shuts down an ISM server.



Notes:

- 1 Windows Only

Element	Purpose	Key Considerations
-q	Quietly shuts down the ISM server without asking for confirmation	In the absence of the -q option, ism_shutdown prompts for confirmation before it shuts down the ISM server.
-deinstall	Removes Windows services also	

Usage

You can run this command only from the computer on which the ISM server is running.

UNIX Only

On UNIX, you must log in as user **root** to use this command. Include this command in the UNIX shut-down script to exit the ISM server automatically whenever the operating system shuts down.

End of UNIX Only


Windows Only

On Windows, the Service Control Manager automatically starts the ISM server when the computer reboots and stops it when the computer shuts down. You must be the **Administrator** to run this command.

End of Windows Only

ism_startup

The **ism_startup** command starts an ISM server.



Element	Purpose	Key Considerations
-init	Initializes the ISM server	Use this option only once after installation. If you use this option thereafter, it produces a harmless error message and takes no other action. The -init option resets the retention period of the save sets on the volume and other values.

Usage

You can run this command only from the computer on which the ISM server is installed.

UNIX Only

On UNIX, you must log in as user **root** to use this command. Include this command in one of the UNIX startup scripts to start the ISM server automatically whenever the operating system is rebooted.

End of UNIX Only

Windows Only

On Windows, the Service Control Manager automatically starts the ISM server when the computer reboots and stops it when the computer shuts down. You must be the administrator to run this command.

End of Windows Only

ism_watch

The **ism_watch** command brings up a terminal-based screen utility for monitoring the activity of the ISM server.

```

▶▶ ism_watch ▶▶
    -p polltime
    -s hostname
  
```

Element	Purpose	Key Considerations
-p <i>polltime</i>	Specifies the interval in seconds between updates of displayed activities that the ISM server performs	None.
-s <i>hostname</i>	Specifies the name of the computer on which the ISM server is running	If you omit this option, this command looks first for an ISM server on the local computer and then for the first ISM server to respond over the network.

Figure A-1 illustrates a typical screen that the **ism_watch** command displays.

```

Server: dbservername.pdx.informix.com
Version: IBM Informix Storage
Manager 2.2 09/16 Network Edition/1

Saves: 0 session(s) Recovers: 0 session(s)
Device      type      volume
/dev/dsk1   file     Pebblebeach.ISMData.0003 mounted file disk Pebblebe
/dev/rmt/0mbn 8mm     Pebblebeach.ISMData.0003 mounted 8mm tape Default.002

Sessions:

Messages:
Mon 18:04:56 /dev/dsk1 label without mount operation in progress
Mon 18:04:58 /dev/dsk1 labeled file disk Pebblebeach.ISMData.0003
Mon 18:05:03 /dev/dsk1 mount operation in progress
Mon 18:05:04 /dev/dsk1 mounted file disk Pebblebeach.ISMData.0003
Mon 18:05:18 /dev/rmt/0mbn mount operation in progress
Mon 18:06:00 /dev/rmt/0mbn mounted 8mm tape Default.002
Pending:

```

Figure A-1. Monitoring the ISM Server

NSR Commands

For information on using **nsr** commands with ISM, see your Legato NetWorker documentation.

Appendix B. Accessibility

The syntax diagrams in the HTML version of this manual are available in dotted decimal syntax format, which is an accessible format that is available only if you are using a screen reader.

Dotted Decimal Syntax Diagrams

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), the elements can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read punctuation. All syntax elements that have the same dotted decimal number (for example, all syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, the word or symbol is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is read as 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol that provides information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements

must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this identifies a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to a separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? Specifies an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element (for example, 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! Specifies a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.
- * Specifies a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data-area, you know that you can include more than one data area or

you can include none. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.

+ Specifies a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times. For example, if you hear the line 6.1+ data-area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. As for the * symbol, you can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Glossary

This glossary contains terms and definitions that are found in this manual. Most of the terms are specific to IBM Informix Storage Manager.

8mm. Represents 8-millimeter magnetic tape.

administrative user. A user on the ISM Administrators list. This user can make changes to the ISM server properties.

backup volume. Backup media, such as magnetic tape or optical disk.

bootstrap. Information needed to recover the ISM catalog after a disk failure.

clone. The process by which ISM makes an exact copy of saved data (save sets). ISM can clone individual save sets or the entire contents of a storage volume.

clone volume. A duplicated volume. ISM tracks cloned volumes in the same way that it tracks the original volumes. Save sets of different types cannot be mixed on one volume; for example, backup versus clone-type volumes.

command-line interface. The prompt where you enter commands.

device. The storage device (tape drive, disk drive, optical disk drive, or file type) connected to the ISM server. The device is used to back up and recover data.

file-type device. A file system on the ISM server that is used as a storage device.

file-type volume. A storage volume that is mounted on a file-type device.

HDR. See High-Availability Data Replication.

High-Availability Data Replication. Copies data to a database server on another computer.

You can also use ON-Bar with HDR to restore data on another computer.

hmt. Half-inch magnetic tape.

Imported restore. Allows for the transfer of data from one instance of a database server to the same instance on a foreign host. In addition to completing the imported restore process, the imported restore is also a subset of these three processes: disaster recovery, database server upgrade, and initialization of HDR.

IBM Informix Server Administrator. A browser-based tool that provides web-based system administration for Informix database servers. For information, see the ISA online help and onscreen instructions.

IBM Informix Storage Manager. Receives backup and restore requests from ON-Bar and directs your data to and from storage volumes that are mounted on storage devices.

ISA. See IBM Informix Server Administrator.

ISM. See IBM Informix Storage Manager.

ISM command-line utilities. Commands that you enter at the command-line interface to administer the ISM server.

ISM server. The UNIX or Windows computer that runs the ISM server software.

ixbar.servernum. An emergency boot file that Dynamic Server uses. On Extended Parallel Server, the emergency boot files consist of a backup boot file, a restore boot file, and a merge boot file.

Legato NetWorker®. A network storage management solution that protects and helps to manage data across an entire network of computers. It manages your data and creates a database of saved data, to make it easy to locate data for recovery. As the network and amount of

data grow, Legato NetWorker provides the capacity and performance to handle the load.

media. Magnetic tape, hard disk, or optical disks used to back up database logs and storage spaces. *See also* backup volume.

ON-Bar bootstrap. Emergency boot files that contain information needed to perform a cold restore. These files are updated after every backup. They replace the **sysutils** tables during a cold restore so ON-Bar can request the correct backup object from ISM. *See also* ixbar.servernum.

qic. Quarter-inch cartridge tape.

recycle. To relabel a volume whose data retention period has expired.

recycle status. A volume status that indicates whether the retention periods for all the save sets on a volume have expired.

retention period. ISM server setting that determines how long save-set entries are retained in the ISM catalog and are thus recoverable by ON-Bar.

retention status. A save-set status that indicates whether the retention period for a save set has expired.

save. Saves files, including directories or entire filesystems, to the ISM server.

savefs. Saves a filesystem (using **save**) to an ISM server. Mount points are not crossed, and symbolic links are not followed.

savegrp. Starts a group of ISM clients saving their filesystems. The **savegrp** command runs a group of ISM clients through the process of saving their filesystems (using **save**). Another way to start a save group is to execute the **ism_catalog** command or the **onbar** script that includes the **ism_catalog** command.

save set. A collection of data that is produced by a single data stream and backed up to storage media. The save set is treated as a single entity on the storage media.

save-set ID . An internal identification number that ISM assigns to a save set. Use the save-set ID (ssid) for disaster recovery.

ssid. *See* save-set ID.

stand-alone device. A storage device that contains a single drive for backing up data. Stand-alone devices cannot store or automatically load backup volumes.

storage space. An area of disk space where data that the database server manages is stored. Dbspaces, blobspaces, and sbspaces are types of storage spaces.

XBSA. *See* X/Open Backup Services Application Programming Interface.

X/Open Backup Services Application Programming Interface. Provides a programming interface and functions that manage backup and restore operations. XBSA connects ISM with the database server.

volume. Backup media, such as magnetic tape, disk-drive partition, or optical disk. *See also* file-type volume.

volume name. The name that you assign to a backup volume when it is labeled.

volume pool. A feature that allows you to sort backup storage spaces to selected volumes. A volume pool contains a collection of backup volumes to which specific storage spaces have been backed up.

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. (enter the year or years).
All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

AIX; DB2; DB2 Universal Database; Distributed Relational Database Architecture; NUMA-Q; OS/2, OS/390, and OS/400; IBM Informix®; C-ISAM®; Foundation.2000™; IBM Informix® 4GL; IBM Informix® DataBlade® Module; Client SDK™; Cloudscape™; Cloudsync™; IBM Informix® Connect; IBM Informix® Driver for JDBC; Dynamic Connect™; IBM Informix® Dynamic Scalable Architecture™ (DSA); IBM Informix® Dynamic Server™; IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager); IBM Informix® Extended Parallel Server™; i.Financial Services™; J/Foundation™; MaxConnect™; Object Translator™; Red Brick™; IBM Informix® SE; IBM Informix® SQL; InformiXML™; RedBack®; SystemBuilder™; U2™; UniData®; UniVerse®; wintegrate® are trademarks or registered trademarks of International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names used in this publication may be trademarks or service marks of others.

Index

A

- Accessibility xix
 - dotted decimal format of syntax diagrams B-1
 - syntax diagrams, reading in a screen reader B-1
- Activity log, ON-Bar 1-4, A-4, A-6
- Adding administrative users 2-4
- Adding storage devices
 - file-type 3-3
 - tape or optical drives 3-3
- Administrative user
 - adding the first 2-4
 - deleting 2-6
 - ism_add 2-4, A-2
 - ism_rm A-11
 - privileges 1-18
- Automatic Media Management
 - described 3-9
 - labeling volumes 4-9
 - relabeling 3-10
 - volume-selection criteria 3-9

B

- Backup
 - See also* Logical-log backup.
 - bootstrap 5-4
 - imported restore 5-5
 - ON-Bar and ISM interaction 1-2
- Boldface type x
- Boot file.
 - See* Bootstrap.
- Bootstrap
 - backing up 5-4
 - defined 1-4
 - finding save set 5-4
 - ism_catalog A-2
 - recovering A-2
 - saving 1-7

C

- Catalog, ISM.
 - See* ISM catalog.
- Changing
 - storage device properties 3-5
 - volume pool name 2-7
- Choosing user and group names 2-5
- Cloned data
 - recovering 5-2
 - tracking 4-15

Cloning

- file-type volumes to tape 3-8, 4-15
 - how retention period applies 4-15
 - ism_clone A-7
 - pools 4-5
 - reserving a device 3-7
 - save sets 4-16
 - storage volumes 4-15
 - understanding feature 4-14
- Code, sample, conventions for xv
- Cold restore.
 - See* Restoring.
- Collecting ISM information A-3
- Command-line conventions
 - how to read xiii
 - sample diagram xiii
- Command-line utilities
 - ism_add A-2
 - ism_catalog A-2
 - ism_chk.pl A-3
 - ism_clone A-7
 - ism_config A-8
 - ism_op A-9
 - ism_rm 5-3, A-11
 - ism_show A-11
 - ism_shutdown A-12
 - ism_startup A-13
 - ism_watch A-14
 - using 1-2
- Compliance
 - with industry standards xxii
- Configuration parameters
 - ISM_DATA_POOL 1-10, 2-6
 - ISM_LOG_POOL 1-10, 2-7
- Configuring
 - devices 3-2
 - ISM on UNIX 1-8
 - ISM on Windows 1-9
 - ISM server properties 2-1
- Connecting to new ISM server 2-6
- Contact information xxiii
- Conventions
 - command-line xiii
 - documentation x
 - sample-code xv
 - syntax diagrams xii
 - syntax notation xii
 - typographical x
- Creating the bootstrap 5-4

Criteria for labeling blank volumes 3-11

D

daemon.log 2-8

Data life cycle

Automatic Media Management 3-9

manual status 4-4

recycle status 4-3

retention period 2-3, 4-2

save-set retention status 4-2

Data recovery.

See Restoring.

Data streams, setting 2-1

Database server upgrade 1-14, 5-5

Database server upgrade, imported restore 5-5

Default locale viii

Deleting

administrative users 2-6

storage devices 3-4

Dependencies, software viii

Detecting volume name 4-10

Devices

adding 3-3

allocating 3-7

Automatic Media Management 3-9

changing properties 3-5

cloning 3-7

configuring 1-7, 3-2

deleting 3-4

designating as read-only 3-5

enabling and disabling 3-4

file-type, optical, tape

adding 3-3

using 3-6

reserving for

cloning 3-7

logical logs 3-7

setting target sessions 2-2

types supported 3-2

Disabilities, visual

reading syntax diagrams B-1

Disabling storage devices 3-4

Disaster recovery

bootstrap A-2

imported restore 5-5

re-creating ISM catalog 5-3, A-2

Documentation conventions x

Documentation Notes xvii

Documentation set of all manuals xix

Documentation, types of xvi

machine notes xvii

online manuals xix

printed manuals xix

Dotted decimal format of syntax diagrams B-1

E

Eligible recycling status 4-3

en_us.8859-1 locale viii

Enabling storage devices 3-4

Environment variables x

ISM_COMPRESSION 2-9

ISM_DATA_POOL 1-10

ISM_DEBUG_FILE 2-8

ISM_DEBUG_LEVEL 2-10

ISM_ENCRYPTION 2-9

ISM_LOG_POOL 1-10

ISM_MAXLOGSIZE 2-8

ISM_MAXLOGVERS 2-8

ON-Bar 2-9

Error messages xviii

Expiration of media 3-10, 4-11, 4-14

Expired save sets

defined 4-2

recovering 5-2

F

File-type devices

adding 3-3

cloning to tape 3-8

configuring

UNIX 1-8

Windows 1-9

pool restrictions 4-6

using 3-6

File-type volumes

cloning to tape 3-8, 4-15

labeling 4-8

mounting 4-9

pool restrictions 4-5, 4-6

relabeling 4-11

unmounting 4-10

using 3-6

Filesystem.

See File-type devices.

Finding a volume name 4-10

Fixed and Known Defects File xvii

G

Global Language Support (GLS) viii

Group names, choosing 2-5

H

Help xix

High-Availability Data Replication (HDR), imported

restore 5-5

Hosts file, Windows 1-13

I

Imported restore 5-5, 5-16

Industry standards, compliance with xxii

Informix Dynamic Server documentation set xix
Informix Storage Manager.

See ISM.

Installation 1-6

Installation Guides xvi

Internal volume label 4-7

ISM

adding administrative users 2-4

changing server properties 2-3

command-line utilities A-1

connecting to a different server 2-6

deleting administrative users 2-6

installation 1-6

message logs 2-8, A-4, A-6

nsr commands A-15

overview 1-1

reports A-3

setup

UNIX 1-8

Windows 1-9

uninstallation 1-11

volume pool names 1-10

ISM catalog

deleting volume 4-13

re-creating entries 5-1

restoring catalog 1-5

ISM Local Execution service 1-9

ISM Portmapper service 1-9

ISM server bootstrap.

See ISMData pool, bootstrap save set.

ISM Server service 1-9

ism_add

-admin 1-8, 2-4

-device 1-8, 1-10

syntax A-2

ism_catalog

-create_bootstrap 2-7, 5-4

-find_bootstrap 5-4

-recover 4-2, 5-4

-recreate_from 2-3, 4-2, 5-2

syntax A-2

ism_chk.pl A-3

ism_clone A-7

ISM_COMPRESSION environment variable 2-9

ism_config

changing server properties 2-1

syntax A-8

ISM_DATA_POOL configuration parameter 1-10, 2-6

ISM_DEBUG_FILE environment variable 2-8

ISM_DEBUG_LEVEL environment variable 2-10

ISM_ENCRYPTION environment variable 2-9

ISM_LOG_POOL configuration parameter 1-10, 2-7

ISM_MAXLOGSIZE environment variable 2-8

ISM_MAXLOGVERS environment variable 2-8

ism_op A-9

ism_rm 5-3, A-11

ism_show

-volume 5-3

-volumes 3-10, 4-14

syntax A-11

ism_shutdown

syntax A-12

UNIX 1-12

Windows 1-12

ism_startup 1-6, A-13

ism_watch 1-4, A-14

ISMData pool

bootstrap save set 1-4

configuration parameter 2-6

renaming 2-7

setting environment variables 1-10

setting up 1-10

ISMDiskData pool 1-7

ISMDiskLogs pool 1-7

ISMLogs pool

configuration parameter 2-7

renaming 2-7

setting environment variables 1-10

setting up 1-10

ISO 8859-1 code set viii

K

Keywords

in syntax diagrams xiv

L

Labeling

Automatic Media Management 4-9

automating 1-4

criteria for blank 3-11

formats 4-6

from command line A-9

recyclable volumes 4-11

storage volumes 4-8

unrecognized labels 3-11

libnsr.dll 1-16

Life cycle of data 4-2

Locale viii

Locating partial save sets 5-1

Log files

daemon.log and environment variables 2-8

reserving device 3-7

viewing 1-18

xbsa.messages, preventing disk overflow 2-8

Logical logs.

See Log files.

Logical restore.

See Restoring.

Logical-log backup

allocating storage devices 1-7

Logical-log backup (*continued*)
operation 1-4
storage devices 3-7

M

Machine notes xvii
Managing media automatically 3-9
Manual recycling status 4-4
Media types, list of supported 3-2
Message file
 See also Activity log, ON-Bar.
 daemon.log 2-8, A-5
 xbsa.messages 2-8, A-5
Migrating to ISM 2.2 1-14
Mounting volumes
 from command line A-9
 from Devices window 4-9
 how ISM chooses 4-9

N

No-rewind tape devices 3-6
nsr commands A-15
nsr directory 1-11, 1-18
nsradmin command 2-2, 3-4, 5-10, 5-12
nsrck command 1-16
nsrd daemon 1-12
nsrexecd daemon 1-12
nsrim command 5-14
nsrindexasm 5-11
nsrindexd daemon 1-12
nsrmmd command 1-13
nsrmmdbd daemon 1-12

O

ON-Bar
 activity log 1-4, A-4, A-6
 backup operation 1-3
 connection to ISM 1-3
 environment variables 2-9
 re-creating the ISM catalog 5-1
 recovering with 5-1
 restore operation 1-5
 setting up ISM 1-6
onbar script
 backup and restore data 1-3
 re-creating bootstrap 5-3
ONCONFIG file parameters.
 See Configuration parameters.
Online help xix
Online manuals xix
Online notes xvi, xvii
Optical disk drives
 adding 3-3
 automating media management 3-9
 using 3-5

P

Partial save sets 5-1
Permissions, adding administrative users 1-18, 2-4
Physical restore.
 See Restoring.

Pools

 file-type devices, restrictions 4-6
 for cloning 4-5
 how ISM uses 4-5
 ISM volume name 1-10
 label format 4-6
 relabeling volumes 4-11
Premounting backup volumes 4-9

Printed manuals xix

Properties

 device 3-5
 ISM server 2-1

R

Re-creating

 ISM catalog 5-1
 save sets 4-2

Read-only designating 3-5

Recovering

 bootstrap A-2
 cloned data 5-2
 expired save sets 5-2
 expired volumes 4-2
 from a disaster 5-3
 locating partial save sets 5-1
 volumes
 deleted from ISM catalog 5-2
 eligible for recycling 5-2
 when to use ON-Bar 5-1

Recycle status

 eligible 4-3
 manual 4-4

Relabeling

 eligible volumes 4-10
 file-type volumes 4-11
 manually 4-11

Release Notes xvii

Renaming ISMData or ISMLogs 2-7

Restoring

 imported 5-5, 5-16
 ON-Bar operation 1-5
 save sets 5-1

Retention period

 clones 4-15
 save sets and volumes 4-2
 setting 2-3

Reverting from ISM 2.2 1-17

S

- sales_demo database ix
- Sample-code conventions xv
- Save sets
 - bootstrap 1-4, 5-4
 - cloning 4-16
 - locating partial 5-1
 - recovering expired 5-2
 - retention status 4-2
- Screen reader
 - reading syntax diagrams B-1
- Server, ISM
 - administrative users 2-4, 2-6
 - connecting to new 2-6
 - data streams 2-1
 - properties 2-1
 - retention period 2-3
- Services, ISM 1-9
- Setup procedures 1-8
- Sharing storage volumes 3-11
- Simultaneous data streams, setting 2-1
- Software dependencies viii
- SQL code xv
- Starting ISM automatically
 - UNIX 1-8
 - Windows 1-9
- Storage devices.
 - See* Devices.
- Storage volume
 - Automatic Media Management 3-9
 - cloning 4-15
 - criteria for labeling blank 3-11
 - defined 1-2
 - deleting 4-13, A-11
 - detecting name 4-10
 - estimated capacity 4-8
 - expiration 3-10, 4-11, 4-14
 - file-type 3-6
 - internal labels 4-7
 - labeling 4-6, 4-8
 - labeling blank volumes 3-10
 - locating partial save sets 5-1
 - managing from command line A-9
 - mounting 4-9
 - premounting 4-9
 - recovering deleted 5-2
 - recovering recyclable 5-2
 - recycle status 4-3, 4-4
 - relabeling
 - automatically 3-10
 - file-type volumes 4-11
 - when to do 4-10
 - sharing 3-11
 - supported media types 3-2
 - unlabeled 3-11

- Storage volume (*continued*)
 - unmounting 4-10
- stores_demo database ix
- superstores_demo database ix
- Syntax diagrams
 - conventions for xii
 - keywords in xiv
 - reading in a screen reader B-1
 - variables in xv
- Syntax segment xiv
- System requirements
 - database viii
 - software viii

T

- Tape drives
 - adding 3-3
 - automating media management 3-9
 - no-rewind 3-6
 - using 3-5
- Target sessions, multiple devices 2-2
- TOC Notes xvii
- Tracking cloned data 4-15
- Types of supported devices 3-2
- Typographical conventions x

U

- Uninstalling ISM 1-11
- UNIX operating system
 - configuring ISM 1-8
- Unlabeled volumes 3-11
- Unmounting volumes 4-10
- update target sessions 2-2
- Upgrading ISM 1-14
- Users, administrative
 - adding the first 2-4
 - deleting 2-6
 - ism_add 2-4, A-2
 - ism_rm A-11

V

- Variables, in syntax diagrams xv
- Visual disabilities
 - reading syntax diagrams B-1
- Volume pools
 - default names 1-10
 - defined 1-2
- Volume.
 - See* Storage volume.

W

- Windows Administrator group, permissions 2-4
- Windows, configuring ISM 1-9
- WORM.
 - See* Optical disk drives.

X

xbsa.messages log 2-8



Printed in USA

G251-2299-00

