



UM Server Extensions
User's Guide



UM Server Extensions User's Guide

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 MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication 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. This publication was developed for products and services offered in the United States of America and the United Kingdom. It is possible that this publication may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative. No part of this publication may be reproduced or distributed in any form or by any means without prior permission in writing from the International Business Machines Corporation.

©Copyright International Business Machines Corporation 1999, 2000. All rights reserved.

Note to U.S. Government Users—Documentation related to restricted rights—Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corporation.

Trademarks

Alert on LAN, IBM, IntelliStation, Netfinity, Netfinity Manager, Netfinity Director, ThinkPad, and Wake on LAN are trademarks of IBM Corporation.

Microsoft, Windows, Windows NT, WIN32, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

Java and all Java-based trademarks or logos are trademarks of Sun Microsystems, Inc.

Intel, Pentium, and LANDesk are trademarks of Intel Corporation in the United States and other countries.

Other company, product, and service names mentioned in this document may be trademarks or servicemarks of others.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. 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. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter 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
N. Castle Drive
Armonk, NY 10504-1785
U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (1) 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, Department 80D, P.O. Box 12195, 3039 Cornwallis, Research Triangle Park, NC 27709, U.S.A. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Contents

Chapter 1—Introducing UM Server Extensions Tools 15

Advanced Systems Management	15
Capacity Manager	16
Cluster Systems Management.....	16
Fibre Channel Storage Manager	16
Fuel Gauge Monitor	17
Rack Manager	17
ServeRAID Manager	17
Software Rejuvenation.....	18
System Availability.....	18
Accessing the UM Server Extensions Tools	18
Cluster Systems Management	19
Advanced Systems Management	19
Capacity Manager	19
Fibre Channel Storage Manager.....	20
Fuel Gauge Monitor	20
Rack Manager	20
ServeRAID Manager.....	20
Software Rejuvenation	21
System Availability	21
UM Server Extensions Tools and the Event Builder in Director	21

Chapter 2—Installing UM Server Extensions Tools. 23

Before You Begin	23
Supported Systems by Component	24
Hardware Requirements.....	25
Installing UM Server Extensions on Windows.....	25
Installing the UM Server Extensions Tools on OS/2.....	27
Mapping the Drive for a Remote Installation.....	28

Installation Instructions for a System Running OS/2	28
Installing the UM Server Extensions Tools on a System Running NetWare	30
Installation Instructions for a System Running NetWare	31
Deploying IBM UM Server Extensions 2.2 Clients through Tivoli IT Director..	34
Uninstalling the UM Server Extensions Tools	35
Local System Running Windows NT	35
Remote System	35
SCO UnixWare	36
Redhat Linux 6.x	38

Chapter 3—Cluster Systems Management 41

Using Cluster Systems Management	42
Starting Cluster Systems Management.....	43
Viewing the Menu Bar	44
File Menu	44
View Menu	45
Utility Menu	46
Help Menu	46
Viewing the Toolbar	46
Viewing the Status Bar	47
Managing Clusters	47
Renaming a Cluster	48
Changing the Description of a Cluster	48
Managing Nodes in a Cluster	49
Starting a Node.....	49
Stopping a Node.....	49
Pausing a Node.....	49
Resuming a Node	50
Adding Node Descriptions.....	50
Managing Resource Groups in a Cluster	51
Creating a New Group	51
Deleting a Group.....	53

Renaming a Group	53
Bringing a Group Online.....	53
Taking a Group Offline	54
Changing Group Description Properties	54
Changing Group Preferred Owners	55
Setting Group Failover Policy	56
Setting Group Failback Policy	57
Moving a Group to Another Node	58
Managing Networks and Network Interfaces	58
Changing a Network and Network Interface Description	59
Enabling a Network for Use in a Cluster	59
Disabling a Network for Use in a Cluster	60
Using the Cluster Expert Wizard	61
File-share Resource Groups	62
Creating a New File-share Resource Group	62
Changing a File-share Resource Group.....	63
Internet Information Server Resource Group.....	64
Creating an IIS Resource Group.....	64
Changing an IIS Resource Group.....	65
Print-spooler Resource Groups	66
Creating a Print-spooler Resource Group.....	66
Changing an Existing Print-Spooler Resource Group.....	67
Resetting IP Address Ranges	68
Closing Cluster Systems Management	69

Chapter 4—Advanced System Management 71

Using the Advanced System Management PCI Adapter as a Network Gateway.	72
Starting Advanced System Management.....	74
ASM Web Based Management.....	76
Remote Management	77
Configuring and Establishing a TCP/IP Connection.....	78
Establishing an ASM Interconnect Connection	80

Selecting an Event Source	81
Operational Parameters.....	82
Configuration Information.....	82
Configuration Settings	83
General Settings	84
The System Identification Group	85
The Dial-in Settings Group	85
The System Management Processor Clock Group.....	88
POST Timeout.....	88
Loader Timeout	89
O/S Timeout	89
Power Off Delay	90
Other Configuration Settings Functions.....	91
Remote Event Settings	91
Remote Event Entry Information Group	93
Remote Event Strategy Group.....	94
Enabled Events Group.....	95
Modem Settings	100
The Port Configuration Group Box	101
Dialing Settings	105
Initialization String Guidelines	107
Network Settings.....	108
SNMP Settings	110
Event Log.....	112
System Power Control	114
Remote POST Console	115
Updating PCI Adapter or Processor Microde	116
Updating System POST/BIOS Microcode	117
Chapter 5—Capacity Manager	119
Components of Capacity Manager	119
Capturing Data.....	122

Starting the Monitor Activator	122
Using the Monitor Activator	123
Activating or deactivating monitors	124
Cluster Monitors	125
Monitoring for Clusters	126
Creating a Status Table in the Systems Pane	126
Viewing Data	127
Generating a Report	131
Generating the Report to a File	131
Generating the Report to the Viewer.....	132
Settings Notebook.....	135
Line Graph.....	136
Trend Graph.....	136
The Toolbar	137
Analyzing Data	137
Performance Analysis Monitor Requirements	137
Performance Analysis	140
Performance Analysis Reports	141
Saving and Printing a Performance Analysis	142
Bottlenecks.....	142
Bottleneck Event Extended Attributes	144
Using Bottleneck Events	145
Generating an Event	145
Creating an Event Filter.....	147
Defining an Action and Action Plan.....	148
Forecasting Data	148
About the Calculations	149
Viewing a Performance Forecast for a Selected System.....	149
About the Forecast Display	150
About Warning Messages on the Forecast Graph.....	150

Chapter 6—Fibre Channel Storage Manager 153

Enterprise Management Window	154
Subsystem Management Window	156
Using the Subsystem Management	157
Renaming Storage Subsystems	158
Configuring Storage Subsystems	159
Partially Managed Storage Subsystem Condition.....	160
Correcting Partially Managed Condition	163
Adding a Storage Subsystem	164
Using Automatic Discovery	167
Storage Subsystem Configuration.....	169
Recovering from Storage Subsystem Problems.....	170
Problem Notification	170
Recovering from Storage Subsystem Problems	172
Recovery Guru.....	172
Specifying Storage Subsystem Cache Settings.....	174
Using the Script Editor	176
Script Usage Guidelines.....	177
Adding Comments to a Script.....	177
Show Statement.....	178
Verifying Script Syntax	178
Script Editor File and Edit Options.....	178
Interpreting the Script Execution Results	181
Executing Scripts	183
Configuration File Errors.....	183
Recovering from a Configuration File Corruption Error	184
Defining Security.....	185
Setting or Changing a Storage Subsystem Password.....	186
Entering a Storage Subsystem Password.....	187
Defining an Array	189
Defragmenting an Array	194
Adding Free Capacity to an Array	194

Redistributing Arrays	197
Logical Drive	197
Creating a Logical Drive	198
Creating a Logical Drive from Unconfigured Capacity	198
Creating a Logical Drive from Free Capacity	204
Failing a Drive.....	204
Reconstructing a Drive.....	206
Initializing a Drive	206
Channel Protection	207
Media Scan.....	207
Changing the Storage Subsystem Media Scan Settings.....	210
Hot Spare Drive	211
Assigning a Hot-Spare Drive	211
Deassigning a Hot Spare Drive	212
Downloading Firmware or NVSRAM.....	212
Controller Mode.....	215
Placing a Controller Offline or Online.....	216
Storage Partitions	218
Configuring Storage Partitions.....	218
Reconfiguring Storage Partitions Topology.....	219
Default Host Group	219
Host Group	220
Host	220
Host Port.....	220
Defining Logical Drive-to-LUN Mappings	221
Reconfiguring your Logical-to-LUN Mappings	222
Changing a Logical Drive-to-LUN Mapping.....	222
Deleting a Host Port, Host, Host Group, or Logical Drive Mapping.....	223
Moving a Host Port	227
Replacing a Host Port.....	228
Renaming a Host Group, Host, or Host Port.....	230

Showing Undefined Host Ports.....	230
Chapter 7—Fuel Gauge Monitor	231
Requirements	232
Fuel Gauge Console.....	232
Information Data.....	233
Table View.....	236
Bar Chart.....	237
Pie Chart	238
Chapter 8—Rack Manager.	241
Rack Manager Console.....	242
Control Pane	243
Topology Tab	243
Component Catalog.....	244
Rack Manager Workspace	245
Table View	246
Rack View.....	246
Component Properties Window	246
Creating and Configuring a Rack	248
Adding Components to a Rack	248
Removing a Rack or a Rack Component	249
Associate Component	249
Disassociate Component.....	251
Supported Rack Components	252
Chapter 9—ServeRAID Manager.	259
Using the ServeRAID Manager Program Interface.....	259
Viewing the Menu Bar.....	259
Viewing the Toolbar	260
Viewing the Expandable Tree.....	260
Viewing the Main Panel.....	261

Viewing the Event Viewer	261
Viewing the Status Bar.....	261
Using the ServeRAID Manager Program	262
Using the Configuration Mode.....	262
Learning about Express Configuration.....	262
Using Express Configuration.....	264
Learning about Custom Configuration	265
Using Custom Configuration.....	266
Using the Information Mode	269
Modifying the ServeRAID Controller Configuration.....	270
Modifying a Configuration.....	270
Changing RAID Levels.....	271
Increasing Free Space.....	272
Increasing Logical Drive Space	274
Setting User Preferences	275
Initialization Settings.....	275
Alarm Settings.....	275
Silencing the Repeating Alarm	276
Understanding Drive States	276
Physical Drive State Descriptions	276
Logical Drive State Descriptions	277
ServeRAID Events in Director	278

Chapter 10—Software Rejuvenation 281

The Software Rejuvenation Window	281
Title Bar.....	282
Calendar	282
Server Time	283
Managing Software Rejuvenations	283
Scheduling Nodes for Software Rejuvenation.....	283
Rejuvenation Options	291
Editing Scheduled Rejuvenations	292

Removing a Node from the Rejuvenation Schedule	295
Using Keyboard Commands	296
Predictive Rejuvenation	297
Schedule Filter	305
Trend Viewer.....	306
Closing the IBM Software Rejuvenation Program.....	306
Creating Action Plans	306
Chapter 11—System Availability	309
The System Availability Interface	309
The Menu Bar	310
The Information window	310
Hover Help.....	311
Detailed List of Records	311
Distribution of System Outages	311
Distribution of System Uptime	313
Availability Report.....	314
System Outages by Hour of Day.....	315
System Outages by Day of Week	317
Displaying Data within a Defined Interval	318
Customizing Availability Dates	319
Detach View	319
Index	321

1

Introducing UM Server Extensions Tools

UM Server Extensions tools expand the flexibility and management capabilities of the **Universal Manageability** tool, IBM Director. These systems-management tools are provided free-of-charge with IBM Director, and if you choose to upgrade to Tivoli IT Director, the Extensions tools remain intact.

The following is a description of each UM Server Extensions tool.

Advanced Systems Management

With Advanced Systems Management, you can monitor critical subsystems, restart logs, and troubleshoot Servers, even when the targeted system is not turned on.

You can configure this tool to send alerts to changes in events such as POST, loader, and operating system timeouts. If any of these events occur, this tool automatically forwards an alert in one of the following ways:

- To another system through a modem
- To a numeric or alphanumeric pager
- To a Manager system using a TCP/IP network connection, or to an SNMP Manager/Serial

Capacity Manager

Capacity Manager is a resource management planning tool that tracks resource utilization, identifies potential network bottlenecks, and recommends ways to improve performance. It can generate a report, which enables effective planning of future capacity, such as microprocessor, disk, or memory upgrades.

The Capacity Manage Report Generator tasks can be scheduled through the Director Scheduler.

Cluster Systems Management

Cluster Systems Management is the interface for IBM Availability Extensions for Microsoft Cluster Service (MSCS). You can manage all cluster-related operations and manage cluster resources in an MSCS-based Availability Extensions cluster from one interface. With this tool you can provide cluster operations for a single cluster, or you can group components onto a node by using the drag-and-drop method. A system administrator receives alerts to any event in a cluster through E-mail, electronic page, or starting another program, thereby reducing downtime.

Fibre Channel Storage Manager

Fibre Channel Storage Manager is a management tool that reduces the time that is needed to configure, administer, and monitor locally or remotely installed Fibre Channel storage subsystems. Some storage subsystem management activities include:

- Configuring available storage subsystem capacity into logical drives to maximize data availability, optimize application performance, and use storage resources
- Recovering from storage subsystem problems to maximize data availability
- Granting access to partitions to the application hosts in the enterprise
- Tuning for optimal application performance

Fuel Gauge Monitor

Fuel Gauge Monitor warns you about conditions which could lead to preventable down time. These conditions involve the power subsystem and the load presented by the system. Some of these conditions occur when the system load increases to the point where power subsystem specifications are being violated. Others occur when a server that has multiple, pluggable power modules experiences an increase in system load which takes power subsystem utilization above a limit known as the Loss of redundancy threshold.

Fuel Gauge Monitor provides a feedback mechanism which allows you to review the operation of the power subsystem at any point in time to determine how far from a 'loss of redundancy' or 'over current' situation the server is at that time. In addition Fuel Gauge Monitor has the following alerts: informational, warning and error.

Rack Manager

IBM Rack Manager provides a flexible, easy-to-install solution for consolidating IBM servers, storage devices, and other standard 19-inch rack equipment.

With Rack Manager, you can group your equipment together, enabling you to manage your system resources and to monitor your system functions more efficiently. Centralizing your equipment in integrated rack suites helps to reduce your "real estate" and support costs.

ServeRAID Manager

ServeRAID Manager is a management tool that reduces the time that is needed to configure, administer, and monitor ServeRAID controllers that are installed locally or remotely on servers. The following are features of the ServeRAID Manager:

- Hot spares are reduced, reserving space on existing drives for rebuilds.
- Data scrubbing and auto synchronization of the parity drive can start automatically.

-
- Migration from previous RAID levels is possible to increase free space and logical drive space.

Software Rejuvenation

With Software Rejuvenation, you can have precise control over the scheduling of a node restart for a specific cluster. The Software Rejuvenation interface resembles the IT Director Tasks Scheduling interface and offers drag-and-drop capability. You can schedule node restarts by dragging a node icon onto the calendar and specifying schedule parameters.

System Availability


System Availability provides graphical proof that Software Rejuvenation improves your system availability. This tool can be used, not only for the Software Rejuvenation users, but for any user with the responsibility for determining system availability for large sets of systems. This tool can be used on groups of nodes.

Accessing the UM Server Extensions Tools


Each task is started from the Director console. Each Extensions tool is accessed by dragging the appropriate icon from the Tasks pane of the Director console and dropped onto a system that supports that Extensions tool in the Group Contents panel. Or, you can right-click a system that supports the Extensions tool in the Group Contents pane and select the appropriate Extensions tool from the system context menu.

Some tasks contain multiple features. The task that has a + beside the corresponding icon denotes a collapsed selection tree. Click the + to expand the multiple features of that task. The following sections describe each Extensions tool service icon (or icons) and the placement of the icon within the Director console.






Cluster Systems Management

Icon	Location
	<p>Cluster Systems Management is within the IBM Cluster Tools parent description in the Tasks section of the Director console.</p>


Advanced Systems Management

Icon	Location
	<p>Advanced Systems Management is in the Tasks section of the Director console.</p>


Capacity Manager

Icon	Location
	<p>Capacity Manager is an Extensions service description and not a task. The supported tasks are located within this heading in the Tasks section of the Director console.</p>
	<p>Using Capacity Manager is not a system-enabled task but a set of online tutorials for using Capacity Manager. Click this description to expand the tutorial choices. Click one of the tutorial tasks to open the Capacity Manager online help.</p>
	<p>Monitor Activator is within the Capacity Manager parent description in the Tasks section of the Director console.</p>
	<p>Report Generator is within the Capacity Manager parent description but is an Extensions service description and not a task. Click this description to expand the type of report selections.</p>
	<p>Report Viewer is within the Capacity Manager parent description in the Tasks section of the Director console.</p>


Fibre Channel Storage Manager

Icon	Location
	The Fibre Channel Storage Manager is in the Tasks section of the Director console.


Fuel Gauge Monitor

Icon	Location
	The Fuel Gauge Monitor icon is displayed in the icon menu of the Director console.


Rack Manager

Icon	Location
	The Rack Manager icon is in the Tasks section of the Director console.


ServeRAID Manager

Icon	Location
	ServeRAID Manager is in the Tasks section of the Director console.

Software Rejuvenation

Icon	Location
	Software Rejuvenation is in the Tasks section of the Director console.

System Availability

Icon	Location
	System Availability is in the Tasks section of the Director console.

UM Server Extensions Tools and the Event Builder in Director

Some of the UM Server Extensions tools add event filters to the Event Builder of Director. These filters provide additional systems-management capabilities. An event such as a battery failure in a ServeRAID controller is now channeled through the Event Action Plan in Director. The Event Action Plan creates a single monitoring interface for UM Server Extensions events.

Descriptions of event filters are found in:

- "Advanced Systems Management events" on page 86.
- "ServeRAID Events in Director" on page 280

2

Installing UM Server Extensions Tools

This section contains instructions for installing UM Server Extensions tools on your client, server, or console. The Extensions tools work with IBM Director or IT Director to enhance your systems-management capabilities.

Before You Begin

The files that the UM Server Extensions Installer program installs are determined by the configuration of the resident IBM Director or IT Director program. Depending on whether you are using your system as a server, console, client, or any combination of the three, Extensions Installer detects your current configuration and automatically installs the appropriate files.

Although the installation program automatically detects your configuration, you must install the UM Server Extensions tools on your console system, server system, and each of your client systems separately.

Before you install UM Server Extensions, consider the following:

- Supported systems by component
- Hardware requirements

Supported Systems by Component

UM Server Extensions Server, which installs all components (Server, Console, and Client), is supported on the following operating systems:

- Microsoft Windows® 2000—Server or Advanced Server
- Windows NT® Server 4.0 (with Service Pack 4 or later)

The UM Server Extensions Console component supports the same operating systems as UM Server Extensions Server, including:

- Microsoft Windows® 2000—Server or Advanced Server
- Windows NT® Server 4.0 (with Service Pack 4 or later)
- Windows 2000 Professional
- Windows NT Workstation 4 (with Service Pack 4 or later)
- Windows 98
- Windows 95 (with OEM Service Release 2 (OSR2) or later)
- Novelle NetWare 4.x or later

Note: ServeRAID Manager console does not run in a NetWare 4.x environment. The Java console runs only in a NetWare 5.x environment. For more information on ServeRAID Manager, see “Chapter 9. ServeRAID Manager,” on page 259.

The following operating systems support UM Server Extensions Client components:

- Windows 2000—Server or Advanced Server
- Windows 2000 Professional
- Windows NT Server 4.0 (with Service Pack 4 or later)
- Windows NT Workstation 4.0 (with Service Pack 4 or later)
- IBM OS/2® Warp 4.0
- OS/2 Warp 5.0
- NetWare 4.x or later
- ASM and RAID for SCO and Linux

Hardware Requirements

UM Server Extensions tools require that you have IBM Director or IT Director installed and operational on your system. In addition, for server systems or console systems, the required components are:

- Intel® Pentium® processor, 200 MHz or faster.
- 64 MB of random access memory (RAM).
- 150 MB of virtual storage.
- 75 MB of free disk space.
- A network adapter that supports TCP/IP. The adapter must also support NetBIOS, IPX, or SNA. This support depends on the transport that is needed to communicate with the managed systems.

The UM Server Extension Client requires at least:

- An IBM Personal Computer or server with an Intel Pentium processor.
- 16 MB of RAM.
- A network adapter that supports TCP/IP. The adapter must support also NetBIOS, IPX, or SNA. This support depends on the transport that is needed to communicate with the managed systems.

The steps for each installation depend on the operating system to which you are installing. UM Server Extensions tools support console systems and server systems that run the Windows operating system only. For steps on installing on a Windows system, refer to “Installing UM Server Extensions Tools on a System Running Windows” on page 25.

UM Server Extensions tools support client systems that run Windows NT, OS/2, and NetWare. If you need to install an OS/2 or NetWare client, see steps under “Installing the UM Server Extensions Tools on OS/2” on page 27, or “Installing the UM Server Extensions Tools on a System Running NetWare” on page 30.

Installing UM Server Extensions on Windows

This section details the attended installation of UM Server Extensions tools in a Microsoft Windows operating-system environment. Follow

the supported operating-system guidelines for installing the various UM Server extensions components.

Perform the following steps to install the UM Server Extension tools on a system running windows:

1. From the Windows Start menu, click **Run**.
2. In the Run field, type:

drive letter:*directory*\Win\setup.exe

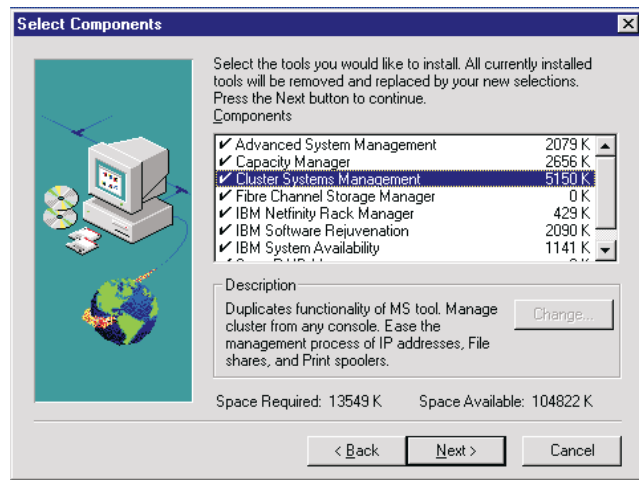
where *drive letter* and *directory* are the drive and the temporary directory where the decompressed UM Server Extensions tools files are located. The **Win** directory indicates the installation files for a Windows installation.

3. Click **OK**.

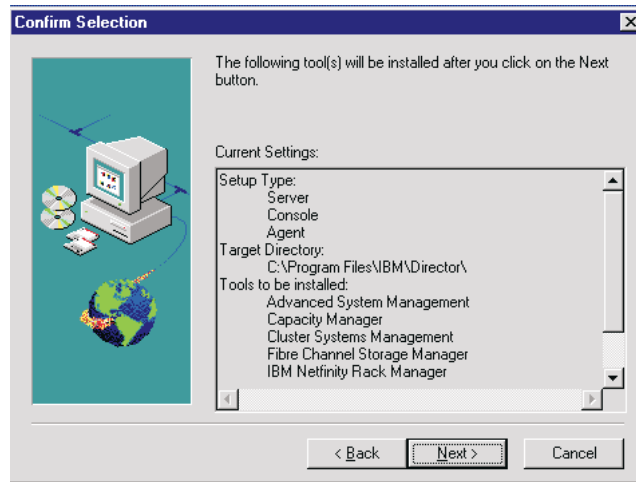
The installation program displays two Welcome windows.

4. Click **Next** through both windows.

The Select Components window opens.



5. Select the tools that you would like to install. Each selection has a description field.
6. Click **Next**. The Confirm Selection window opens.



7. Click **Next**. The UM Server Extensions Installer installs the selected UM Server Extensions tools.
8. If IBM Director is currently running on your system, the Question window opens. Click **Yes** to continue the installation.
9. When the installation is complete, you are prompted to restart your system.

Installing the UM Server Extensions Tools on OS/2

To install the UM Server Extensions tools on an OS/2 platform, you must run UM Server Extensions Installer remotely from a system running Windows NT. Systems running Windows 95 or Windows 98 do not work for this purpose.

You must map the drive from the OS/2 system to the Windows NT system. The best method of mapping the drive is to use the Net Use command in a DOS window. The installation program displays

windows that enable you to map the drive. Using the window method does not work in all remote installation situations.

Mapping the Drive for a Remote Installation

To map a drive from a DOS window, type the Net Use command as follows:

```
Net Use [driveletter|*] [\\computername\sharename]
```

where *driveletter* represents any available drive on the Windows NT system, and *computername\sharename* are the system name and share name of the OS/2 system.

Installation Instructions for a System Running OS/2

Perform the following steps to install the UM Server Extensions tools on an OS/2 system:

1. On the target system, run the following command: twqipc shutdown. This will stop Director and allow the installation to proceed.

2. In the Run field, type:

```
drive_letter:\directory\os2netw\setup.exe
```

where *drive_letter* and *directory* are the drive and the temporary directory where the decompressed UM Server Extension tools files are located. The **os2netw** directory indicates the installation files for an OS/2 or NetWare installation.

3. Click **OK**.

The installation program displays two Welcome windows.

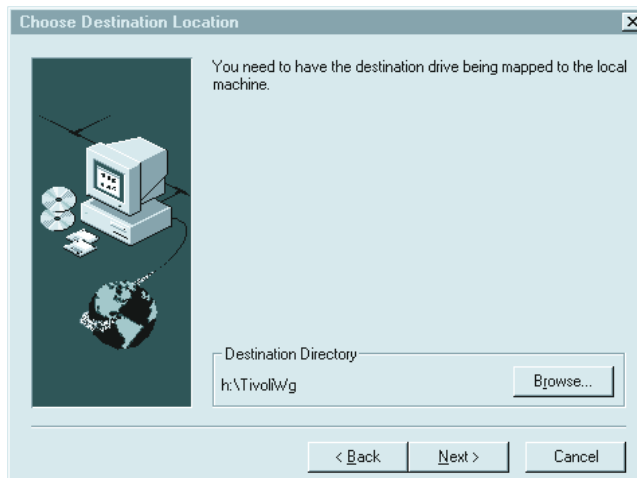
4. Click **Next** through both Welcome windows.

The system displays the Select Platform window.

5. Select the **OS/2** radio button.

6. Click **Next**.

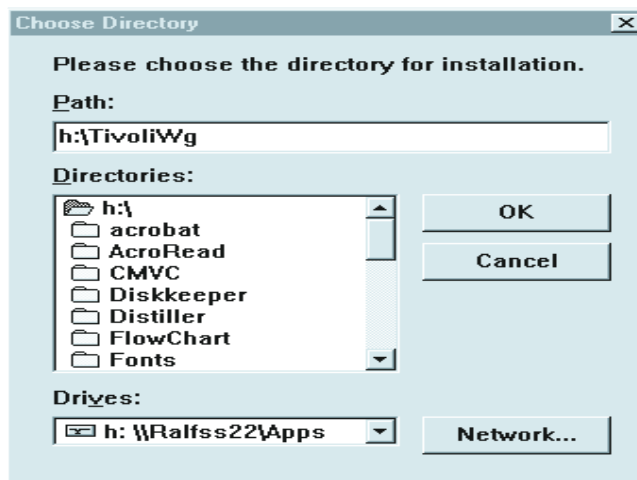
The Choose Destination Location window opens, reminding you that you must first map a drive for the OS/2 installation.



7. Click **Next** if you are using the default destination directory. Skip to step 10 on page 30.

Click **Browse** if you want to change the destination directory.

The installation program displays the Choose Directory window.



-
8. Click the down arrow beside the Drives field, and select the drive that you mapped for the OS/2 installation.

You can click the **Network** button to find an available drive to map; however, for this installation it is best to map the drive before you start the Extensions Installer program.

Refer to “Mapping the Drive for a Remote Installation” on page 28.

9. From the Directories list, select the directory where IBM Director is installed. The default is *x:\TivoliWg* where *x* is the mapped drive and *TivoliWg* is the default directory.
10. Click **OK**.

UM Server Extensions Installer installs the following Extensions tools:

- Advanced Systems Management
- Capacity Manager

11. When the installation is complete, from the console of the OS/2 system, type `twgipc` to restart IBM Director.

Installing the UM Server Extensions Tools on a System Running NetWare

To install the UM Server Extensions tools on a NetWare system, you must run UM Server Extensions Installer remotely from a Windows NT system. Windows 95 or Windows 98 systems will not work.

With NetWare, you must map a drive between the NetWare system and the Windows NT system before you start the Extensions Installer program.

The best method of mapping the drive is to use the Net Use command in a DOS window. The installation program displays windows that are intended to enable you to map the drive, but using the window method will not work for you in all remote installation situations.

For instructions on using the Net Use command to map a drive, refer to “Mapping the Drive for a Remote Installation” on page 28.

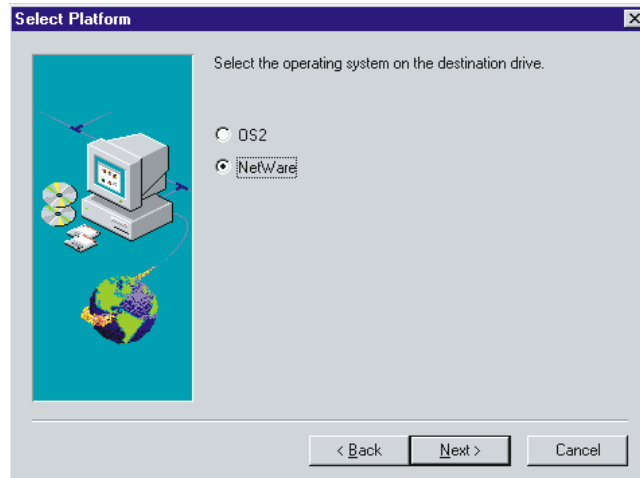
Installation Instructions for a System Running NetWare

Perform the following steps to install UM Server Extensions tools on a NetWare system:

1. From the NetWare console, type `unload SecMod01.sec` to shut down IBM Director.
2. From the Windows Start menu, click **Run**.
3. In the Run field, type:
`drive_letter:\directory\os2netw\setup.exe`

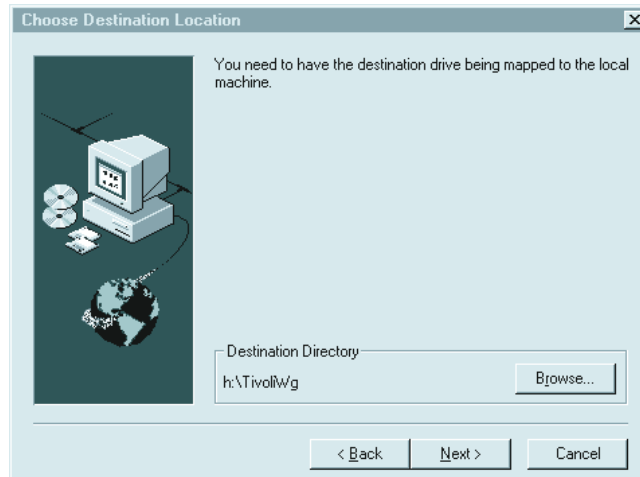
where *drive_letter* and *directory* are the drive and the temporary directory in which the decompressed UM Server Extensions tools files are located. The `os2netw` directory indicates the installation files for an OS/2 or NetWare installation.

4. Click **OK**.
The installation program displays two Welcome windows.
5. Click **Next** through both Welcome windows.
The system displays the Select Platform window.
6. Select the **NetWare** radio button.



7. Click **Next**.

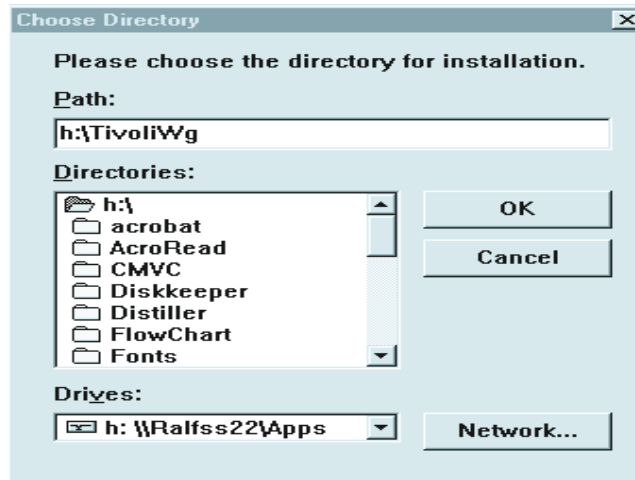
The Choose Destination Location window opens, reminding you that you must first map a drive for the NetWare installation.



8. Click **Next** if you are using the default destination directory. Skip down to step 11 on page 33.

Click **Browse** if you want to change the default destination directory.

The installation program displays the Choose Directory window.



9. Click the down arrow beside the Drives field, and select the drive that you mapped for the NetWare installation.

You cannot use the Network button to find an available drive to map, because with a NetWare installation, you must map the drive before you start the Extensions Installer program.

10. From the Directories list, select the directory where IBM Director is installed. The default is *x:\TivoliWg* where *x* is the mapped drive, and *TivoliWg* is the default directory.
11. Click **OK**.

UM Server Extensions Installer installs the following Extensions tools:

- Advanced Systems Management
- Capacity Manager
- ServerRAID Manager

-
12. When the installation is complete, from the NetWare system console, type `twgipc` to restart IBM Director.

Deploying IBM UM Server Extensions 2.2 Clients through Tivoli IT Director

Using Tivoli IT Director, you can distribute the UM Server Extensions Client code to managed systems through the Software Distribution task. When you install UM Server Extensions Server, you also install the distributable client file as well. You will use this file to build a software distribution package that you can deploy to your selected nodes.

To distribute the IBM UM Server Extensions 2.2, use the following procedure:

1. From the Tivoli IT Director Management Server, open the IT Director Console.
2. In the Tasks pane, right-click Software Distribution.
3. Expand the selection Wizards, and double-click IT Director File Package.

The IT Director File Package Builder window opens.

4. Select **Get files from the IT Director Server**, and then click **Browse**.

5. Find the file IBMUMSE.BFP, located in the directory

drive_letter:\tivoliwg\IBMSWDist\

where *drive_letter* is the drive in which the UM Server Extensions distribution package is located.

6. Click **Next**.
7. Click **Finish**.

You now have an IBM UM Server Extensions 2.2 icon under the Software Distribution task. Deploy the client by dragging this icon onto the target node. For more information on Tivoli IT Director Software Distribution, see your *Tivoli's IT Director Users Guide*.

Uninstalling the UM Server Extensions Tools

The procedure for uninstalling the UM Server Extensions tools depends on whether you installed the Extensions tools locally or remotely.

Local System Running Windows NT

To remove the UM Server Extensions tools from a local Windows NT system, use the following procedure:

1. Click **Start** → **Programs** → **UM Server Extensions** → **Uninstall IBM UM Server Extensions**.

The uninstallation window displays the message:

Are you sure you want to remove the application and its components?

2. Click **Yes**.

When the system is finished removing the files, it displays the message:

You need to reboot the machine to complete the uninstallation.

3. Click **OK**.

You are prompted to restart the system.

Remote System

Regardless of which operating system is running on the remote system, **Windows NT**, **OS/2**, or **NetWare**, the uninstallation procedure is the same. To remove the UM Server Extensions tools from a remote system, perform the following steps:

1. From a DOS window on the Windows NT system, map a drive for the remote target system. For steps on mapping the drive, refer to “Mapping the Drive for a Remote Installation” on page 28.
2. Change the directory to the target directory of your Director program.
3. Type `uninstex nddirectory`

where *nddirectory* is the target directory of your Director program (including drive and full path).

Note: If you do not type the correct target directory; the system displays a message that the uninstallation routine cannot find the appropriate log file.

The **Remove Programs from your Computer** window opens, confirming the deleted files.

4. Click **OK**.
5. The system displays the message:
You need to reboot your machine to complete the installation.
6. Click **OK**.
7. You must manually shut down and restart the system.

SCO UnixWare

The IBM Director UnixWare ASM Agent provides the Advanced System Management support for the IBM Director on UnixWare 7 and 7.x. For installation, removal and setup, use the following procedure:

1. Insert the CD into the CD-ROM drive.
2. Type

```
"mount -r -F cdfs/dev/cdrom/cdromdevicefile/mnt"
```

where *cdromdevicefile* is the specific device file for the CD-ROM block device. You can look in the */dev/cdrom* directory to determine what this is on your system. An example of what you might see in the directory is: *c0b0t610*.

3. Press Enter and type:

```
"pkgadd -d /mnt/unix/asmuwag.ds"
```

where *mnt* is the mount point of the device file created in step 2.

4. Press Enter.
5. When the installation is complete, type:

```
'umount /mnt'
```

-
6. Press Enter.
You can remove the CD-ROM from the system.

To install from the Web use the following procedure:

1. Download the IBM Director UnixWare ASM Agent from the following location:
http://www.pc.ibm.com/ww/eserver/xseries/systems_management/nfdir/serverext.html
2. Select the `asmuwag.ds` file and download it to a temporary directory, such as `/tmp`
3. To install this package, type the following:

```
"pkgadd -d /tmp/asmuwag.ds"
```

where `tmp` is the directory specified in step 2
Note: To install or remove the ASM package, you must have 'root' privileges. To remove this application, type:

```
"pkgrm asmuwag"
```

The IBM Director UnixWare RAID Agent provides the ServeRAID support for the IBM Director on UnixWare 7.x. This package requires that the IBM Director UnixWare Agent be installed. For installation, removal and setup, use the following procedure:

1. Insert the CD into the CD-ROM drive.
2. Type:

```
"mount -r -F cdfs/dev/cdrom/cdromdevicefile/mnt"
```

where `cdromdevicefile` is the specific device file for the CD-ROM block device. You can look in the `/dev/cdrom` directory to determine what file this is on your system. An example of what you might see in the directory is: `c0bot610`.
3. Press Enter, then type:

```
"pkgadd -d /mnt/unix/RAIDUwAg.pkg"
```

where `mnt` is the mount point of the device file created in step 2.
4. Press Enter.
5. When the installation is complete, type:

```
"umount /mnt"
```

-
6. Press Enter.

You can remove the CD-ROM from the system.

The IBM Director Unixware RAID Agent can be downloaded from the WEB. Use the following procedure:

1. Download from the following location:
http://www.pc.ibm.com/ww/eserver/xseries/systems_management/nfdir/serverext.html
2. Select the RAIDUwAg.pkg file and download it to a temporary directory, such as /tmp.
3. To install this package, type the following:

```
"pkgadd -d /tmp/RAIDUwaAg.pkg"
```

where tmp is the directory specified in step 2.

Note: To install or remove the ServeRAID Manager package, you must have 'root' privileges. To remove this application, type:

```
"pkgrm RAIDUwAg"
```

Redhat Linux 6.x

The IBM Director Linux ASM Agent provides the Advanced System Management Support for the IBM Director on RedHat Linux 6.x. Use the following installation procedure:

1. Insert the CD into the CD-ROM drive.
2. Type:

```
"mount /mnt/cdrom"
```

where cdromdevicefile is the specific device file for the CD-ROM block device. You can look in the /dev/cdrom directory to determine what this is on your system. An example of what you might see in the directory is: c0bot60.

3. Press Enter, then type:

```
"rpm -i /mnt/cdrom/unix/asmlxag-1.0.i386.rpm"
```


-
- where mnt is the mount point of the device file created in step 2.
4. Press Enter.
 5. When the installation is complete, type:

```
"umount /mnt/cdrom"
```
 6. Press Enter.
You can remove the CD-ROM from the system.

To install IBM Director Linux ASM Agent from the WEB, use the following procedure:

1. Download IBM Director UnixWare ASM Agent from the following location:
http://www.pc.ibm.com/ww/solutions/enterprise/sysmgmt/products/nfdir/ser_ext.html
2. Select the asmlxag.rpm file and download it to a temporary directory, such as /tmp.
3. To install this package, type the following

```
"rpm -i /tmp/asmlxag-1.0.i386.rpm"
```

where tmp is the directory specified in step 2.

Note: To install or remove the ASM package, you must have 'root' privileges. To remove this application, type:

```
"rpm -e asmlxag"
```

The IBM Director Linux RAID Agent provides ServeRAID Manager Support for the IBM Director on RedHat Linux 6.x. Use the following installation procedure:

1. Insert the CD into the CD-ROM drive.
2. Type:

```
"mount /mnt/cdrom"
```

where cdromdevicefile is the specific device file for the CD-ROM block device. You can look in the /dev/cdrom directory to determine what this is on your system. An example of what you might see in the directory is: c0bot60.

-
3. Press Enter, then type:

```
"rpm -i /mnt/cdrom/unix/raidlxAgm.rpm"
```

where mnt is the mount point of the device file created in step 2.
 4. Press Enter.
 5. When the installation is complete, type:

```
"umount /mnt/cdrom"
```
 6. Press Enter.
You can remove the CD-ROM from the system.

To install IBM Director Linux ServeRAID Agent from the WEB, use the following procedure:

1. Download IBM Director UnixWare ASM Agent from the following location:
http://www.pc.ibm.com/ww/eserver/xseries/systems_management/nfdir/serverext.html
2. Select the RAIDLxAg.rpm file and download it to a temporary directory, such as /tmp.
3. To install this package, type the following

```
"rpm -i /raidlxag.rpm"
```

where tmp is the directory specified in step 2.

Note: To install or remove the ASM package, you must have 'root' privileges. To remove this application, type:

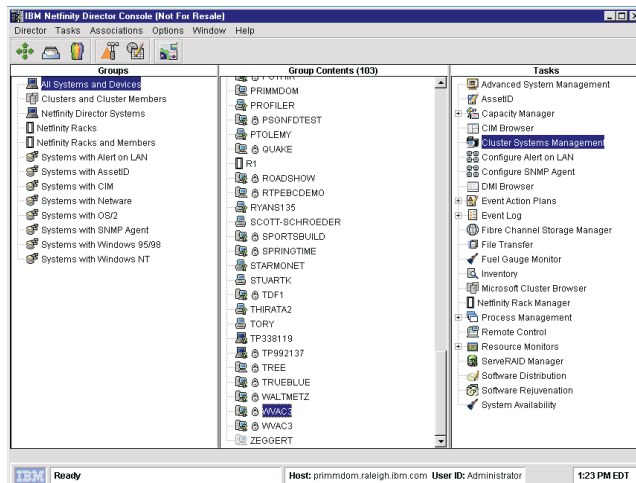
```
"rpm -e raidlxag"
```

3

Cluster Systems Management

The IBM Cluster Tools are installed during the UM Server Extensions installation and are seamlessly integrated into the IBM Director Console. Cluster Systems Management (ICSM) is the main component of the IBM Cluster Tools. This server program is used to administer high-availability cluster environments, for example IBM Availability for MSCS or MSCS clusters, and to increase reliability of cluster nodes.

You must be logged on to the IBM Director console as a user with general access, group access, and task access privileges before you can use IBM Cluster Tools.



Using Cluster Systems Management

Cluster Systems Management is a graphical user interface (GUI) client program that you can use to initiate cluster-related operations and manage cluster resources in a Microsoft Cluster Service (MSCS) based cluster. The program displays all cluster components, including nodes, groups, resources, networks, and network interfaces. You can provide cluster operations for a single cluster, or you can group components onto a node. The following describes the cluster entities.

Node

A node in the clustering environment represents a supported IBM server. Nodes can own resource groups.

Resource group

A resource group in the clustering environment is a collection of resources that are grouped together in a single system. State, Move, and Initiate Failure are some resource-group functions. When applied, these functions affect all resources in a group. (For example, if a node fails, another node in the cluster takes ownership of the group, providing client/server applications continuous access to the storage device.)

Resource

A resource offers a service to clients in a client/server application. Cluster Systems Management uses the resource types (for example, physical disks, IP addresses, and network names) that exist in the clustering environment to perform specific high-availability functions. For example, if a node fails, another node in the cluster takes ownership of the physical disk resources, providing client/server applications continuous access to the storage device.

Network and network interfaces

In the clustering environment, networks and network interfaces define an internal cluster communication between nodes and how clients access nodes in a cluster.

Names

Each cluster, node, resource, or any other component of Cluster Systems Management that is defined by the user contains some basic limitations. Each component name (such as cluster name,

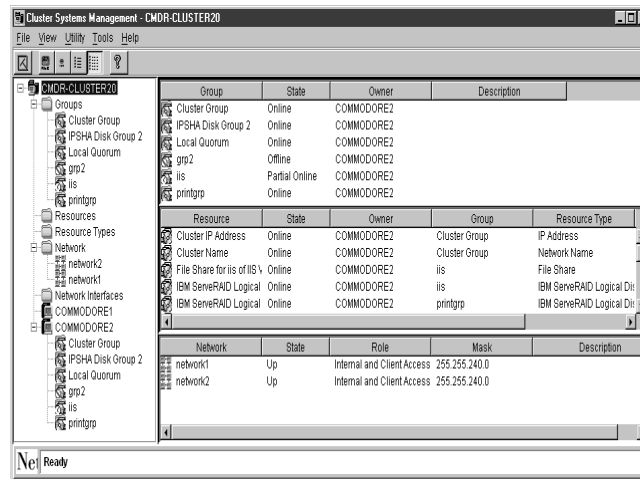
or group name) must be no more than 255 characters long, including spaces. If a component name includes a dash (-) or a hyphen(-), use the underscore (_) to denote these symbols. Failure to follow these naming conventions will cause the cluster systems management to display an erroneous subset of cluster information.

Starting Cluster Systems Management

To start Cluster Systems Management from the IBM Director main window, you must do one of the following:

- Drag-and-drop the Cluster Systems Management task icon onto a cluster (IBM Availability Extensions for MSCS or MSCS cluster).
- Drag-and-drop a cluster onto the Cluster Systems Management task icon.
- Right-click a cluster name in the IBM Director Console, and then click **Cluster Systems Management** from the menu.

The Cluster Systems Management window opens.



The Cluster Systems Management window includes the following components:

- Title bar
- Menu bar
- Toolbar
- Status bar

The ICSM window is divided into four panes and is populated with the cluster data. The left pane displays the Cluster Systems Management cluster tree structure (the cluster name at the highest level, followed by the groups, resources, resource types, nodes, networks, and network interfaces). The panes on the right display the Group, Resource, and Network views.

Viewing the Menu Bar

The menu bar is a set of menu names that are located directly below the title bar. The menu bar contains the following options.

File Menu

The **File** menu provides options that you can use to perform basic cluster tasks. The options of in this menu can vary, depending on the selected cluster entity.

Command	Use this command to:
New	Create a new cluster group or cluster resource.
Properties	Change the properties of a cluster entity.
Rename	Rename a cluster entity.
Bring Online	Bring a cluster group or cluster resource online.
Take Offline	Take a cluster group or cluster resource offline.
Initiate Failure	Initiate a resource failure.
Move Group	Move a cluster group or cluster resource to another location.
Change Group	Change the group for the selected resources.
View Resource Types	Inspect the various cluster resources types.

Command	Use this command to:
Pause Node	Pause the operations of the server (node) in a cluster.
Resume Node	Resume the operations of the server (node) in a cluster.
Start Cluster Service	Start the ICSM or MSCS service.
Stop Cluster Service	Stop the ICSM or MSCS service.
Delete	Permanently remove a resource or group from a cluster.
Exit	Close the ICSM program.

Note: These options can also be accessed from a menu by right-clicking an entity.

View Menu

You can use the View menu to change the appearance of items that are displayed in the Cluster Systems Management main window.

This menu contains the following commands.

Command	Use this command to:
Toolbar	Show or hide the toolbar in the main window.
Large Icons	Display the cluster entity in the main window.
Small Icons	Display the cluster entities as small icons in the main window.
List	List the cluster entities in the main window.
Details	List and display details (for example, owner and description) about the cluster entities in the main window.
Refresh	Refresh the main window.

Utility Menu

The **Utility** menu contains an option that you can use to manage and administer the ICSM clusters.

This menu contains the following option.

Command	Use this command to:
Cluster Expert Wizard	Create file-share, IIS, and print spooler resource groups. (IIS is for MSCS clusters only.)




Help Menu




The **Help** menu provides online information about Cluster Systems Management.

Viewing the Toolbar

The Toolbar is a set of buttons that are located directly below the menu bar. These buttons serve as shortcuts for many frequently used commands. When you first view the Cluster Systems Management window, some commands are disabled and are enabled only after you access certain menu commands.

The toolbar contains the following options.

Icon	Command	Use this command to:
	Refresh	Refresh the main window.
	View Large Icons	Display the cluster entity.
	View Small Icons	Display the cluster entities as small icon in the main window.

Icon	Command	Use this command to:
	View List	List the cluster entities in the main window.
	View Details	List and display details (for example, state owner and description) about the cluster entities in the main window.
	About	Invoke online information about ICSM.

Viewing the Status Bar

The Status bar at the bottom of the window displays a message line that provides information about a selected menu command.

Managing Clusters

A cluster name represents the top component in the Cluster Systems Management tree structure and owns all of the entities, such as resources, resource groups, nodes, networks, and network interfaces. You can create, delete, or move cluster entities, as well as change the properties of cluster entities.

Renaming a Cluster

To change the name of a cluster, use the following procedure:

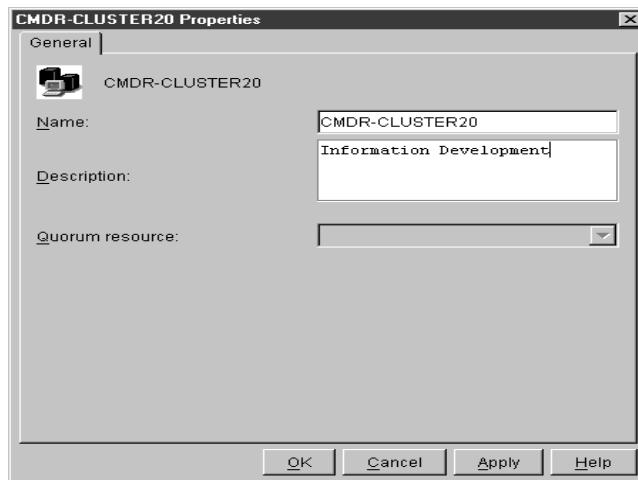
1. In the upper-left pane of the Cluster Systems Management window, right-click the cluster name that you want to change.
2. Click **Rename**.
3. Type the new name.
4. Press **Enter**.

Changing the Description of a Cluster

To change the description of a cluster, use the following procedure:

1. In the left pane of the Cluster Systems Management window, click the name of the cluster whose properties you want to change.
2. Click **File** → **Properties**.

The Properties window opens for the selected cluster.



3. Type a description of the cluster in the Description field.
4. Click **Apply**.
5. Click **OK**.

Managing Nodes in a Cluster

A node in the MSCS environment represents a supported IBM server. Nodes own resource groups, and a resource group can be owned by only one node at time. When a node starts, the cluster service starts automatically.

The following sections describe the various operations that you can apply to a node in a cluster.

Starting a Node

To start a node, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node that you want to start.
2. Click **Start Cluster Service**.
3. Click **Refresh**.

Stopping a Node

To stop a node, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node that you want to stop.
2. Click **Stop Cluster Service**.
3. Click **Refresh**.

Pausing a Node

To pause a node, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node that you want to pause.
2. Click **Pause Node**.

Resuming a Node

To resume a node, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node that you want to resume.
2. Click **Resume Node**.
3. Click **Refresh**.

Adding Node Descriptions

To add comments or a description of a node to its General properties information, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node whose description you want to change.
2. Click **Properties**.



3. Type a description of the node in the Description field.
4. Click **Apply**.
5. Click **OK**.

Managing Resource Groups in a Cluster

A resource group is a collection of resources. You can change the state of resources by changing the state of the resource group it is in.

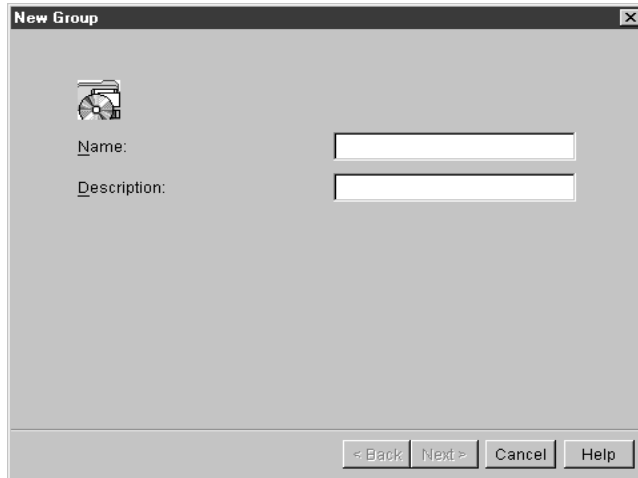
You can manage groups in a cluster by:

- Creating a new group
- Deleting a group
- Renaming a group
- Bringing a group online
- Taking a group offline
- Changing group description properties
- Changing group preferred owners
- Setting group failover policy
- Setting group failback policy
- Moving a group to another node

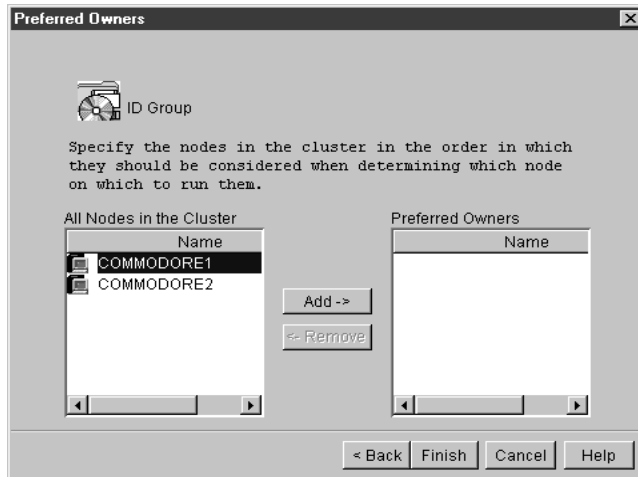
Creating a New Group

After defining the resources that you want to group together, you can create groups. To create a group in a cluster, do the following:

1. Click **File** → **New** → **Group**.
The New Group window opens.



2. Type the name and the description of the group. Supported group names must be less than 255 characters in length.
3. Click **Next**.



4. In the Preferred Owners window, use the following procedure:

-
- a. From the All Nodes in the Cluster menu, click the name of the node that you want to be the preferred owner.
 - b. Click **Add**.
 - c. Repeat steps a and b for each node that you want to be considered for use in the event of failure.
- Note:** To remove a node from the Preferred Owner list, click the name, and then click **Remove**.
5. Click **Finish**.

Deleting a Group

You can delete a group that is online only if it does not contain resources. To delete a cluster group, use the following procedure:

1. From the Cluster Systems Management window, click a group name.
2. Click **File** → **Delete**.
3. Click **Yes** to confirm the deletion.

Renaming a Group

You can assign a different name to each group in a cluster. To rename a group in a cluster, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the node that you want to change.
2. Click **Rename**.
3. Type the new name.
4. Press **Enter**.

Bringing a Group Online

To bring a group online in a cluster, use the following procedure:

1. From the Cluster Systems Management window, right-click a group name.
2. Click **Bring Online**.

Taking a Group Offline

To take a group offline in a cluster, use the following procedure:

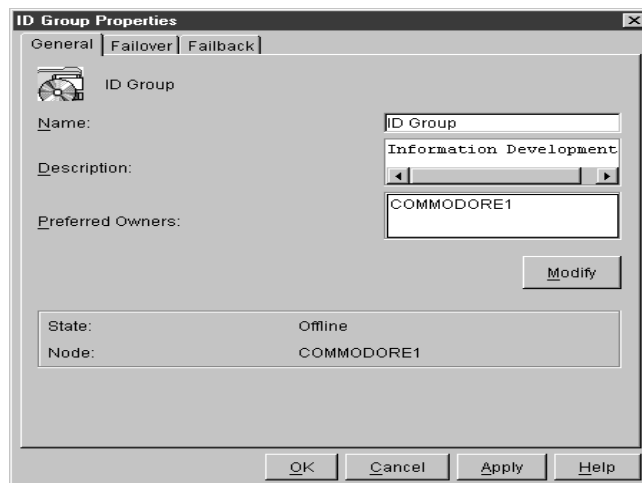
1. From the Cluster Systems Management window, right-click a group name.
2. Click **Take Offline**.

Changing Group Description Properties

To change the description of a cluster group, use the following procedure:

1. From the Cluster Systems Management window, select the group name.
2. Click **File** → **Properties**.

The ID Group Properties window opens.



3. In the Description field, type a description of the group.
4. Do one of the following:
 - Click the **General** tab, if you want to change the name, description, or preferred owners of the group.

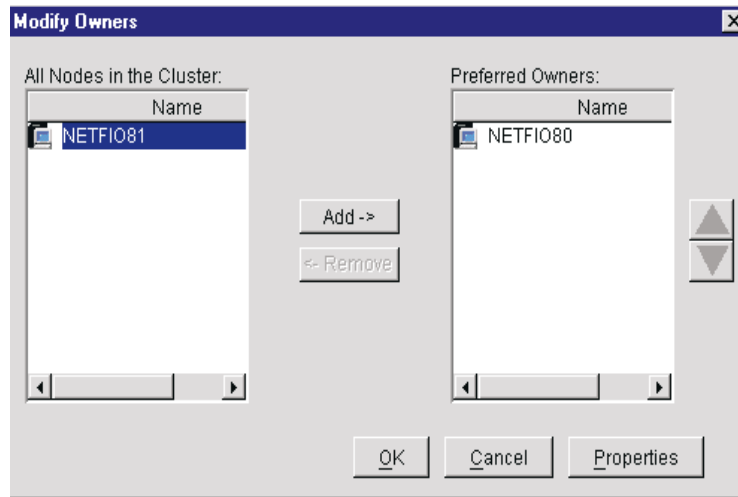
-
- Click the **Failover** tab, if you want to change the group failover Threshold or Period information.
 - Click the **Failback** tab, if you want to prevent, allow, or schedule an immediate or predefined time for failback.
5. Click **Apply**.
 6. Click **OK**.

Changing Group Preferred Owners

To add, remove, or change group preferred owners, use the following procedure:

1. Click a group.
2. Click **File** → **Properties**.
3. Click **Modify**.

The Modify Owners window opens.



4. In the **Modify Owners** window, do the following:
 - a. From the **All Nodes in the Cluster** drop-down list, select the name of the node that you want to be the preferred owner.

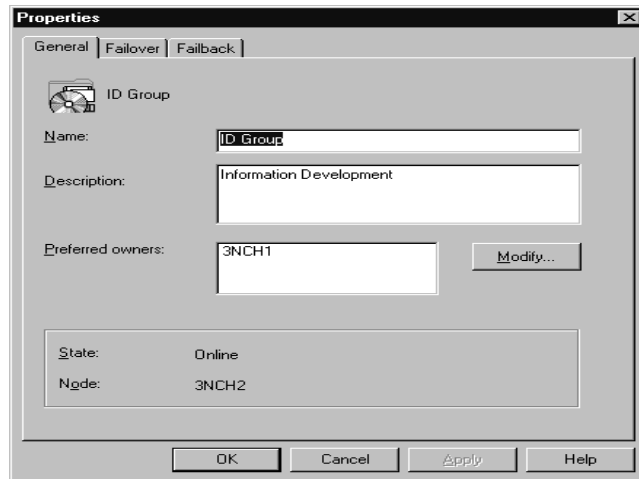
-
- b. Click the right arrow button to add the name to the **Preferred owners** list and to specify the preferred owner in the event of failure.
 - c. Repeat steps a and b for each node that you want to be considered for use in the event of failure.
5. Click **OK**.
 6. Click **Apply**.
 7. Click **OK**.

Setting Group Failover Policy

The failover policy for Cluster Systems Management is to rotate to the next node that is listed in the **Preferred Owners** list. To set the failover policy for a group in a cluster, use the following procedure:

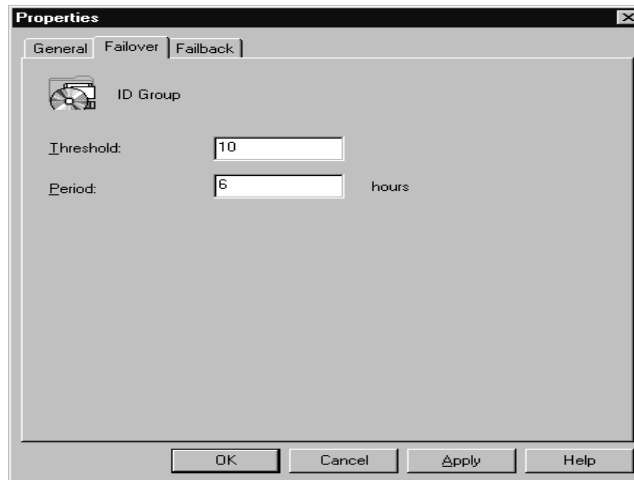
1. Click a group name.
2. Click **File** → **Properties**.

The **Properties** window opens.



3. Click the **Failover** tab.

The **Failover** page window opens.



4. Do the following:
 - Set the Threshold field to the maximum number of times the group is to failover.
 - Set the Period field to the maximum number of hours before you take the group offline.

For example, if the group failover threshold is 10 and the period is 6, the cluster software will be taken offline after the tenth attempt that occurs within six hours.
5. Click **Apply**.
6. Click **OK**.

Setting Group Failback Policy

When a node comes back online, the group will failback to that node only if it is the first node in the Preferred Owners list. To enable or disable the group failback policy, use the following procedure:

1. Click the group name.
2. Click **File** → **Properties**.
3. In the Properties window, click the **Failback** tab.

-
4. Do one of the following:
 - To initiate the Failback, click **Prevent Failback**.
 - To schedule a Failback, click **Allow Failback** and click either **Immediately** or **Failback Between**
Note: The Failback Between beginning and ending values must be from 0 through 23. If the beginning value is greater than the ending value, the failback will occur the following day.
 5. Click **Apply**.
 6. Click **OK**.

Moving a Group to Another Node

To move a group to a different node, use the following procedure:

1. In the upper-left pane of the Cluster Systems Management window, right-click the name of the group that you want to move.
2. Click **Move Group**.
3. Click the name of the node to which the group is to be moved.

Note: You can also use the drag-and-drop method to accomplish this task.

Managing Networks and Network Interfaces

You can manage a network by:

- Changing the network and network interface description (MSCS clusters only)
- Enabling the network for use in a cluster
- Disabling the network for use in a cluster

Changing a Network and Network Interface Description

To change the network description for MSCS clusters, use the following procedure:

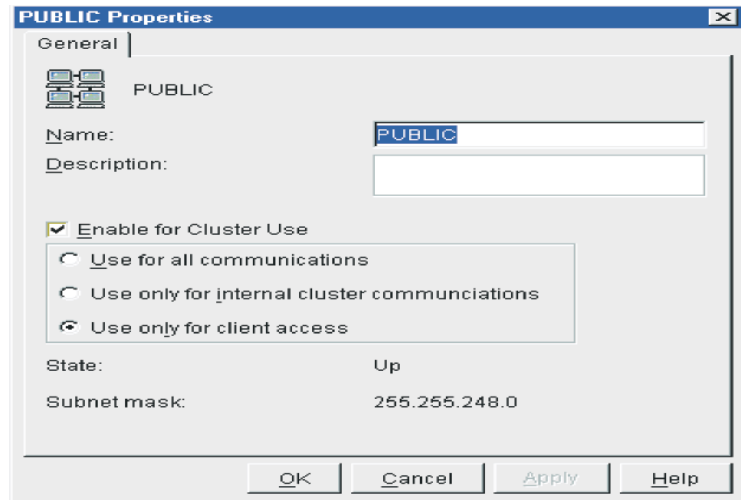
1. Click the network name.
2. Click **File** → **Properties**.
3. Type the new network name.
4. Type the new description.
5. Click **OK**.
6. Click **Apply**.
7. Click **OK**.
8. Click **Finish**.

Enabling a Network for Use in a Cluster

To enable a network for use in a cluster, use the following procedure:

1. Click the name of the network that you want to enable.
2. Click **File** → **Properties**.

The Public Properties window opens.



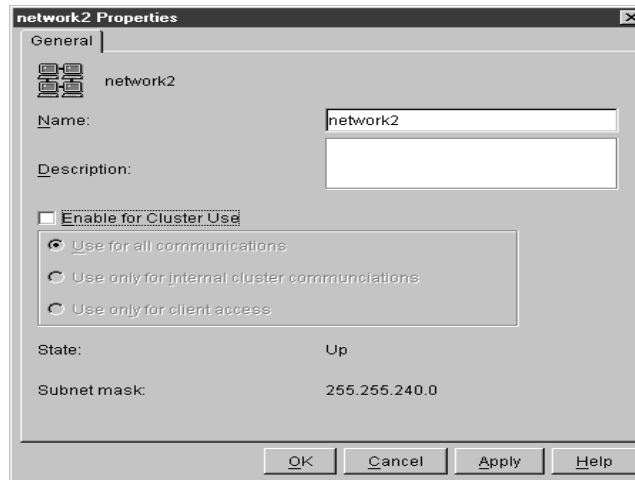
3. Select the Enable for cluster use check box.
4. Select one of the option buttons to specify how you want to use the network in the cluster.
5. Click **Apply**.
6. Click **OK**.

Disabling a Network for Use in a Cluster

To enable a network for use in a cluster, use the following procedure:

1. Click the name of the network that you want to disable.
2. Click **File** → **Properties**.

The Network window opens.



3. Clear the **Enable for Cluster use** check box.
4. Select one of the option buttons to specify the way in which you want to use the network in the cluster.
5. Click **Apply**.
6. Click **OK**.

Using the Cluster Expert Wizard

In the MSCS environment, you must define high-availability resource groups. MSCS must have information about which resources make up the resource groups and what their dependencies on each other are.

Every resource group must contain a virtual IP address for the clients to use when communicating with resource groups. MSCS makes virtual IP addresses and other resources highly available, to reduce the likelihood of failure. For example, if a server fails, another node will take ownership of the IP address and resources of that resource group.

For cluster groups to failover correctly, cluster resource groups must have the correct resources and dependencies.

The **Cluster Expert** wizard is a tool that you can use to support file-share and print-spooler resource groups. You can use it to create and

define new resource groups in existing resource groups. This is especially useful when you have a limited number of physical disks that need to serve multiple purposes in your environment. For example, to store data for multiple file shares, you can use a single physical disk.

During startup of the server, ICSM prompts you for a range of virtual IP addresses. For the server, there is a sequential range of IP addresses.

Note: Do not specify an IP address range, that includes an address that is currently assigned. For example, if the address 9.9.9.10 is the address that is currently assigned, use a range of 9.9.9.11-9.9.9.100.

The Cluster Expert wizard adds deleted IP addresses to the list of available IP addresses.

File-share Resource Groups

File-share resource groups share the directory on one of the shared disks in the configuration. This file-share resource group is highly available, so that if one node fails, another node takes ownership of the failed node resources.

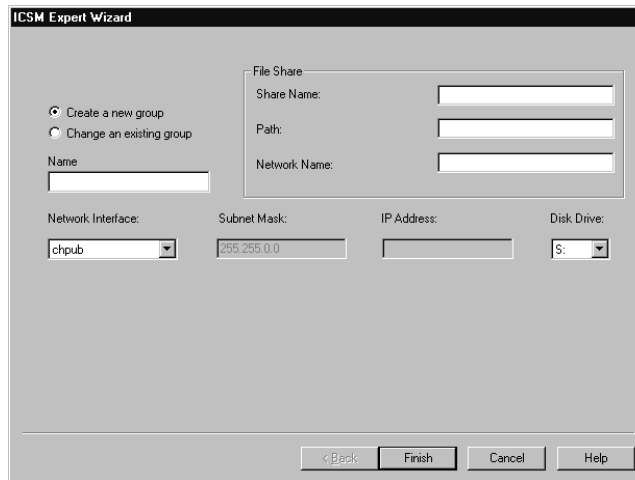
You can create a new file-share resource group or change an existing file-share resource group.

Creating a New File-share Resource Group

Before creating a new file-share resource, make certain that a physical disk drive is available, and create a directory on the drive. To create a new file-share resource group, use the following procedure:

1. Click **Utility** → **Cluster Expert Wizard** → **File Share**.

The ICSM Expert Wizard window opens.



2. Select **Create a new group**.
3. Type the group name in the **Name** field.
4. Type the share name, path, and network name in the **File Share** group box.
5. Select the network interface from the **Network Interface** drop-down list.
6. Select the physical disk drive letter from the **Disk Drive** drop-down list.
7. Click **Finish**.

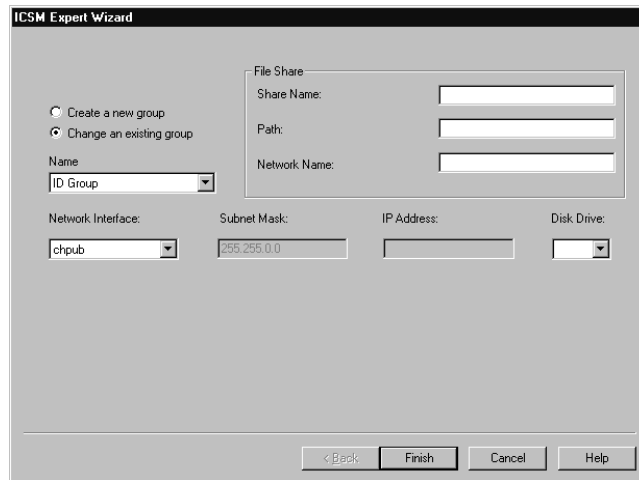
The Cluster Systems Management window displays the new group names.

Changing a File-share Resource Group

To change a file-share resource group, use the following procedure:

1. Click **Utility** → **Cluster Expert Wizard** → **File Share**.

The ICSM Expert Wizard window opens.



2. Select **Change an existing group** option button.
3. Select the group name from the **Name** drop-down list.
4. Select the physical disk drive letter from the **Disk Drive** drop-down list.
5. Click **Finish**.

Internet Information Server Resource Group

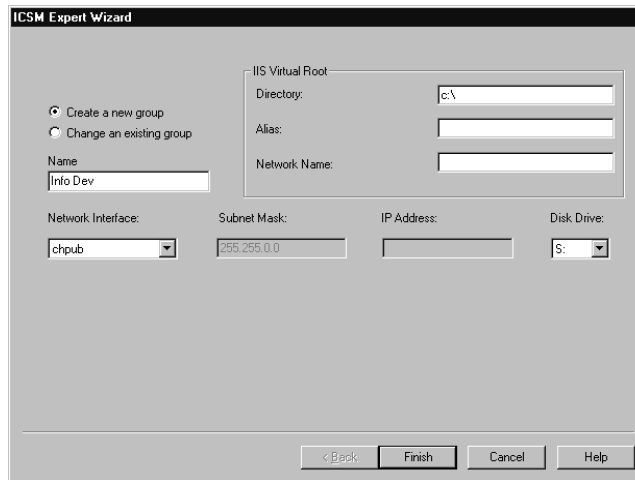
An Internet Information Server (IIS) resource group provides high-availability for the World Wide Web server, FTP, and Gopher components of the Microsoft Internet Information Server. IIS functionality is available only to MSCS clusters.

Creating an IIS Resource Group

To create an Internet Information Server resource group in an MSCS cluster, use the following procedure:

1. Click **Utility** → **Cluster Expert Wizard** → **IIS**.

The ICSM Expert Wizard window opens.



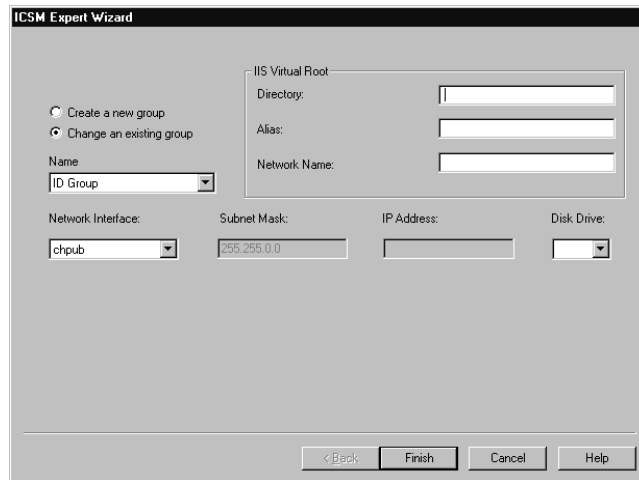
2. Select **Create a new group**.
3. For the **IIS virtual root**, type the **Directory**, **Alias**, and **Network Name**.
4. Select the network interface for the IIS, from the **Network Interface** drop-down list.
5. Select the physical disk drive letter from the **Disk Drive** drop-down list.
6. Click **Finish**.

Changing an IIS Resource Group

To change an Internet Information Server resource group in an MSCS cluster, use the following procedure:

1. Click **Utility** → **Cluster Expert Wizard** → **IIS**.

The ICSM Expert Wizard window opens.



2. Select **Change an existing group**.
3. Type the directory, alias, and network name for the IIS virtual root in the **IIS Virtual Root** group box.
4. Select the **Network Interface** for the IIS, from the drop-down list.
5. Select the physical disk drive letter from the **Disk Drive** drop-down list.
6. Click **Finish**.

Print-spooler Resource Groups

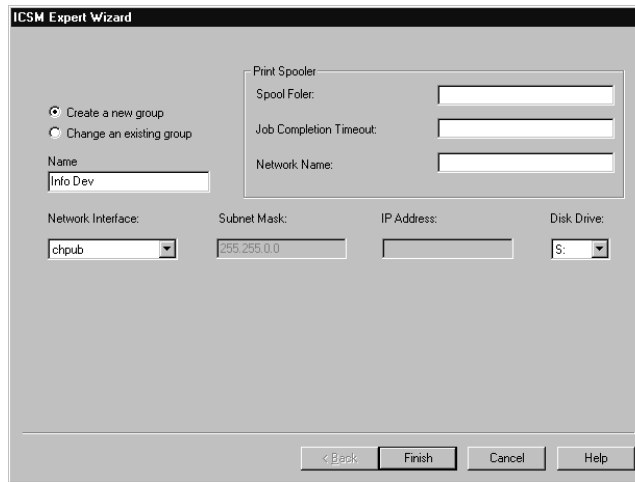
When a server functions as a print spooler, the server must specify where the print spooler stores its data. In a single-server environment where the server functions as the print spooler, the server must specify a spool directory for data storage.

Creating a Print-spooler Resource Group

To create a print-spooler resource group, do the following:

1. Click **Utility** → **Cluster Expert Wizard** → **Print Spooler**.

The ICSM Expert Wizard window opens.



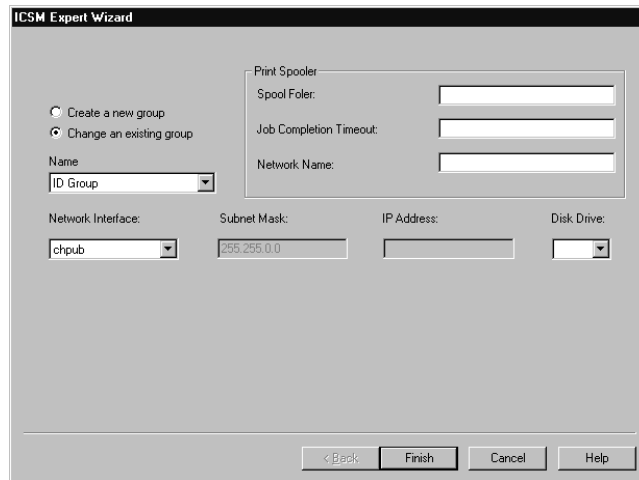
2. Select **Create a new group** option button.
3. Type the spool folder, job completion time-out, and network name information for the print spooler in the **Print Spooler** group box.
4. Select the **Network Interface** for the print spooler, from the drop-down list.
5. Select the physical **Disk Drive** letter from the drop-down list.
6. Click **Finish**.

Changing an Existing Print-Spooler Resource Group

To change an existing print-spooler resource group, use the following procedure:

1. Click **Utility** → **Cluster Expert Wizard** → **Print Spooler**.

The ICSM Expert Wizard window opens.

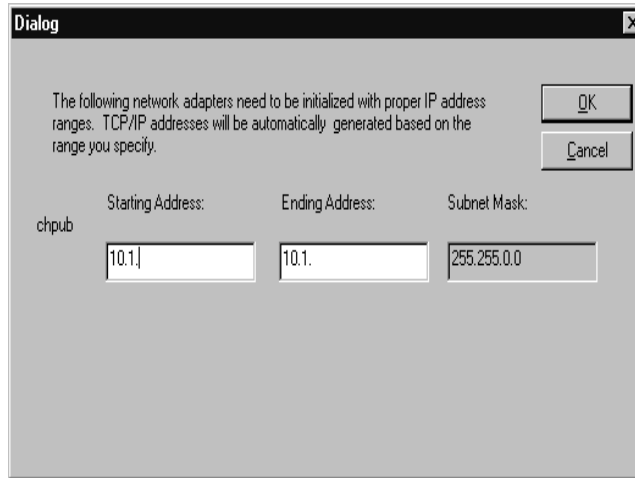


2. Click the **Change an existing group** option button.
3. Type the spool folder, job completion time-out, and network name for the print spooler.
4. Select the **Network Interface** for the print spooler from the drop-down list.
5. Select the physical **Disk Drive** from the drop-down list.
6. Click **Finish**.

Resetting IP Address Ranges

Network adapters require proper IP address ranges to be able to initialize. To reset the virtual IP address range for the Cluster Expert wizard, use the following procedure:

1. Click **Tools** → **Reset Expert Wizard IP Address Range**.
The Dialog window opens.



2. Type the preferred IP address range in the **Starting Address** and the **Ending Address** fields.
3. Click **OK**.

Closing Cluster Systems Management

To close the Cluster Systems Management program, from the Cluster Systems Management window click **File** → **Exit**.

4

Advanced System Management

Use the Advanced System Management (ASM) service for IBM Director to change the configuration, modem, network, and automatic dial-out settings of your IBM Advanced System Management PCI adapter or IBM Advanced System Management processor.

With the Advanced System Management service, you can configure system-management events (such as POST, loader, and operating system time-outs or critical temperature, voltage, and tampering events). If any of these events occurs, the Advanced System Management service can be configured to automatically forward an event in one of five ways:

- To a standard numeric pager
- To an alphanumeric pager
- To an IBM Director system using a TCP/IP network connection (available only when using Advanced System Management with an Advanced System Management PCI adapter)
- To a simple network management protocol (SNMP) based system management system in SNMP format (available only when using Advanced System Management with an Advanced System Management PCI adapter)
- To an IBM Director system with an attached modem

With this service, you can manage the Advanced System Management hardware that is installed in your own system, or you can use Advanced System Management to connect with the Advanced System Management PCI adapter or processor that is installed in a remote

server. You can connect with remote Advanced System Management hardware in one of three ways:

- If the remote system has an Advanced System Management PCI adapter installed and the adapter is connected to a TCP/IP network, you can open a TCP/IP link with the IBM Advanced System Management PCI adapter from your IBM Director Console system.
- If your Advanced System Management PCI adapter or processor is connected to an Advanced System Management Interconnect network, or if the remote Advanced System Management PCI adapter or processor to which you have connected using a TCP/IP connection is connected to an ASM Interconnect network, you can use this connection to access and manage the Advanced System Management PCI adapter or processor of any other system that is connected to the ASM Interconnect network.
- If serial connection is established to the Advanced System Management PCI adapter or processor using modems or a null cable from the IBM Director Console system.

In addition, with Advanced System Management, you can remotely monitor, record, and replay all text that is generated during power-on self-test (POST) on a remote system that includes an Advanced System Management PCI adapter or processor. While monitoring a remote system during POST, you can type commands from your keyboard that will then be relayed to the remote system.

Using the Advanced System Management PCI Adapter as a Network Gateway

In systems that contain both an Advanced System Management processor and an Advanced System Management PCI adapter, the adapter acts as an Ethernet or token-ring network gateway (or as a shared modem resource). In this configuration, the Advanced System Management processor generates all events, time-outs, and other systems-management information. This data is relayed to the Advanced System Management PCI adapter using the ASM Interconnect connection between the processor and the adapter. The adapter then forwards this information to other systems on the Ethernet or token-ring

network (or uses its modem to forward this data using a serial connection).

When configuring systems that have both the Advanced System Management PCI adapter and the Advanced System Management processor, all systems-management settings (such as remote event settings and time-out settings) must be configured on the Advanced System Management processor. However, before using TCP/IP to communicate with your Advanced System Management PCI adapter, you must first establish an ASM Interconnect connection with the adapter and configure the network settings.

To establish another ASM Interconnect connection to another Advanced System Management PCI adapter through a serial port, use the following procedure:

1. From the ASM Integrated Console, drag and drop the ASM icon onto a system.
The Advanced System Management Console window opens.
2. Click **Options** → **ChangeConnection** → **Serial**.
3. In the **Entry Name** text box, type the name of the connection.
4. Type the phone number of the remote Advanced System Management processor or PCI adapter in the text box.
5. Click **Dial Number** to connect.
6. At the prompt, type the **User ID** and **Password** in their respective text boxes.
7. Click **Cancel** to exit from the Establish Serial Connection window.
8. Click **Configure Entry** to open the **Configure Serial Connection Entry** window. From this window, configure the local dialing and modem settings.
9. Click **Save Entry** to save the serial connection.
10. Click **Delete Entry** to remove the highlighted entry in the list of available connections.

The Number text box at the bottom of the window displays the phone number to be dialed. The Serial text box at the bottom of the window

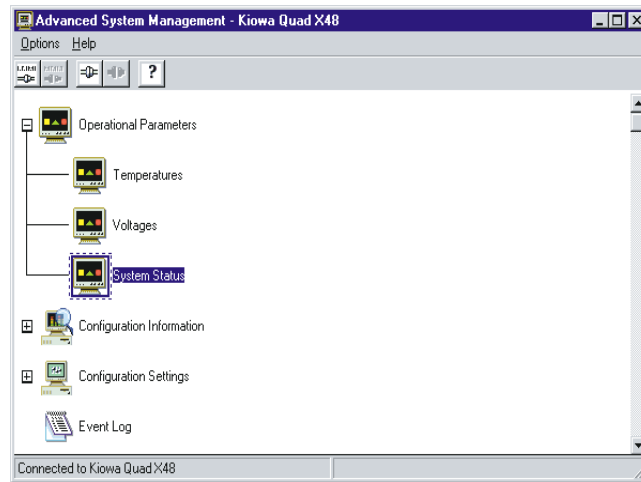
displays the type of serial connection device that will be used to make the connection.

Starting Advanced System Management

There are two ways to initiate the Advanced System Management service: through the ASM integrated Console or through the ASM Web Based Management task.

To start the Advanced System Management service from the ASM integrated Console, drag the Advanced System Management icon from the Tasks pane of the IBM Director Console and drop it on a system that supports Advanced System Management in the Group Contents pane or right-click a system that supports Advanced System Management in the Group Contents pane and click **Advanced System Management** from the menu. Then double-click any of the selections available in the **Advanced System Management** window to access the function or the configuration information that you need. To start the Advanced System Management task in a disconnected mode, right-click **Advanced System Management** and click **Open**. Use this method to establish a direct serial or TCP/IP connection to a remote Advanced System Management PCI adapter or processor.

Note: This starts Advanced System Management on this system and will enable you to configure and manage the Advanced System Management PCI adapter or Advanced System Management processor that is installed in your own system only. To access, configure, and manage an Advanced System Management PCI adapter or processor in a remote system, you must first use Advanced System Management to establish a connection with the remote system.



- Double-click **Operational Parameters** to expand the Operational Parameters tree, and then double-click a component to view the current values or status of many system components that are monitored by the Advanced System Management PCI adapter. For more information, see “Operational Parameters” on page 82.
- Double-click **Configuration Information** to expand the Configuration Information tree, and then double-click a component to view detailed information about the Advanced System Management PCI adapter or processor, including random-access memory (RAM) microcode, read-only memory (ROM) microcode, and device-driver information. For more information, see “Configuration Information” on page 82.
- Double-click **Configuration Settings** to expand the Configuration Settings tree, and then double-click a component to configure the Advanced System Management PCI adapter or processor. These features include General Settings (such as system identification data, dial-in security settings, the time and date reported by the system management processor clock, time-out and delay values), Modem Settings, Network Settings (Advanced System Management PCI adapter only), SNMP Settings (Advanced System Management PCI adapter only), and Remote Event

Settings. For more information, see “Configuration Settings” on page 83.

- Double-click **Event Log** to view the contents of the Advanced System Management PCI adapter Event Log or processor. Information about all remote access attempts and dial-out events that have occurred is recorded in the Event Log. For more information, see “Event Log” on page 112.
- Double-click **System Power Control** to instruct the Advanced System Management PCI adapter or processor to power-off the system, restart the system, or power-on the system. For more information, see “System Power Control” on page 114.
- Double-click **Remote POST Console** to remotely monitor, record, and replay all textual output that is generated during POST on a remote system that has an Advanced System Management PCI adapter. For more information, see “Remote POST Console” on page 115.
- To update the microcode on your Advanced System Management PCI adapter or processor, click **Options** → **Update Microcode** → **System Management**. For more information, see “Updating PCI Adapter or Processor Microcode” on page 116.
- To update the system POST/BIOS microcode on a system that includes an Advanced System Management PCI adapter or processor, click **Options** → **Update Microcode** → **System POST/BIOS**. For more information, see “Updating System POST/BIOS Microcode” on page 117.

ASM Web Based Management

The Advanced System Management Web based task launches a web browser against the Local Area Network (LAN) interface of the Advanced Systems Management Service Processor. It queries the service processor in-band through the IBM Director Agent for the IP address of its active interface. If the query is successful, the web browser is launched using this information. However, if the query is not successful (for instance, the target system is powered off), information from the management server is used. You are presented with this

information for confirmation, or You are allowed to supply new information.

To start the Advanced System Management service using the ASM Web Based Management task, use the following procedure:

1. Drag the ASM Web Based Management icon from the Tasks pane of the IBM Director and drop it on a system that supports Advanced System Management in the Group Contents pane. The ASM Web Based Management window opens.
2. Enter the IP Address and Hostname, click **OK**.
3. The browser is launched.

Remote Management

If you want to use the TCP/IP, serial or ASM Interconnect network connection from your system access and manage the Advanced System Management adapter or processor on a remote system, you can use Advanced System Management to:

- Establish a TCP/IP, serial, or ASM Interconnect link with the Advanced System Management PCI adapter that is installed in the remote server
- Establish an ASM Interconnect link with the Advanced System Management processor that is installed in the remote server

You can also establish a TCP/IP connection with a remote Advanced System Management PCI adapter and then “pass through” that Advanced System Management PCI adapter and remotely access and manage any Advanced System Management adapter or processor that is connected to the remote IBM Advanced System Management PCI adapter using an ASM Interconnect network.

Notes:

- TCP/IP linking over a network connection is available only when you are using Advanced System Management to directly access an Advanced System Management PCI adapter that is connected to your network.
- ASM Interconnect connections are available only in the following situations:

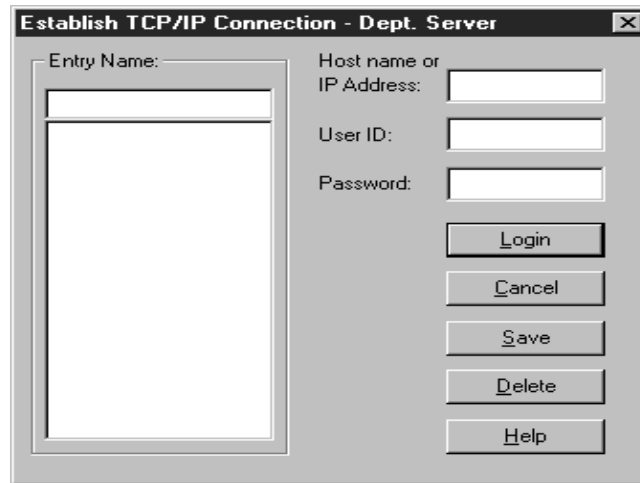
-
- You are using Advanced System Management to directly access the Advanced System Management PCI adapter that is installed in your own system (Netfinity 7000 M10 only).
 - You are using Advanced System Management to access an Advanced System Management adapter or processor that is connected to the same ASM Interconnect network that your Advanced System Management adapter or processor is connected to.
 - You have first established a TCP/IP link with a remote Advanced System Management PCI adapter that is connected to other Advanced System Management adapters or processors on an ASM Interconnect network.

Configuring and Establishing a TCP/IP Connection

To configure and establish a TCP/IP connection with the Advanced System Management PCI adapter in a remote server, use the following procedure:

1. From the ASM Integrated Console, drag and drop the **Advanced System Management** icon onto the selected system.
The Advanced System Management window opens.
2. Click the **TCP/IP Connection** icon or click **Options**→ **Change Connection**→**TCP/IP**.

The Establish TCP/IP Connection window opens.



3. Select a TCP/IP connection entry from the **Entry Name** selection list, or create a new entry and then select the new entry. To create a new entry:
 - a. In the Entry Name text box, type a name for the entry.
 - b. In the Host Name or IP Address text box, type the TCP/IP address or host name that is used by the remote Advanced System Management PCI adapter.
 - c. In the User ID and Password text boxes, type a user ID and password combination that will enable you to access the remote Advanced System Management PCI adapter or processor.

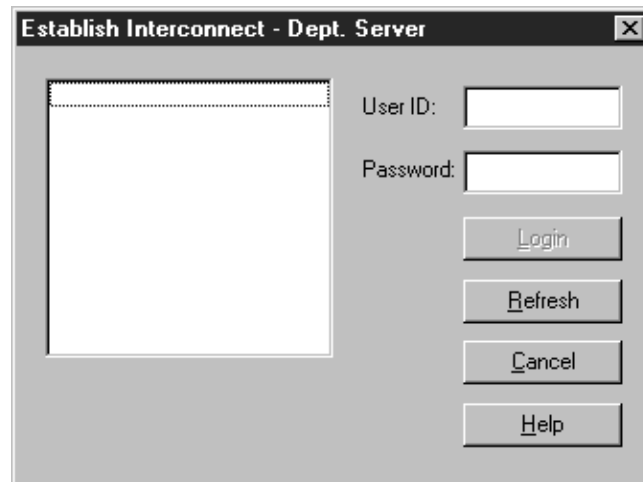
This must match a user ID and password combination that has been configured, using the Advanced System Management service, to allow access to the Advanced System Management adapter or processor.
 - d. Click **Save** to add this entry to the Entry Name selection list.
4. Click **Login** to establish the TCP/IP connection with the remote Advanced System Management PCI adapter. When the connection is established, use the Advanced System Management service to manage the remote Advanced System Management PCI adapter.

Establishing an ASM Interconnect Connection

Unlike the TCP/IP connection, ASM Interconnect connections require no additional configuration before you attempt to connect with other Advanced System Management adapters or processors on the ASM Interconnect network. To establish an ASM Interconnect connection, use the following procedure:

1. From the ASM Integrated Console, drag and drop the Advanced System Management icon onto the selected system.
The Advanced System Management window opens.
2. Double-click the **ASM Interconnect Connection** icon in the **Advanced System Management** menu or click **Options**→**Change Connection** →**Interconnect**.

The Establish Interconnect window opens.



3. Select a system from the Establish Interconnect the list.
4. In the User ID and Password text boxes, type a user ID or password for logging on to the remote Advanced System Management adapter or processor.

Type a user ID and password that will allow access to the remote Advanced System Management adapter or processor. This must match a user ID and password combination that has been

configured, using the Advanced System Management service, to allow access to the Advanced System Management adapter or processor.

5. Click **Login** to establish an ASM Interconnect connection with the selected system. After the connection is established, use the Advanced System Management service to manage the Advanced System Management PCI adapter or processor in the remote system.

Selecting an Event Source

Use the selections that are available from the Select Event Sources menu (in the **Options** menu) to select the sources of Advanced System Management events that will be received and managed by the Advanced System Management service. These are the three available options, one for each method by which the Advanced System Management PCI adapter or processor can report events:

- Driver (via driver)

This enables IBM Director to receive Advanced System Management events that are generated by the Advanced System Management device drivers. The IBM SP Driver (via Agent) option is always enabled.

- TCP/IP (via server)

Click **TCP/IP (via server)** to enable IBM Director to receive Advanced System Management events that are generated by the Advanced System Management PCI adapter and are forwarded using the adapter network connection.

Note: This selection is available only on systems that use an Advanced System Management PCI adapter.

- Serial (via Server)

This option enables you to select a configured serial line to use for receiving serial (dial-up) events using a modem.

If you do not enable an event source, events that are generated by this source will not be received and handled as IBM Director events.

Operational Parameters

Click the + beside Operational Parameters or double-click **Operational Parameters** to expand the Operational Parameters tree and view the Operational Parameters components. Use the Operational Parameters components to view the current status of system components, including:

- **Temperatures**

Includes current temperatures and threshold levels for system components such as far-end adapter, center adapter, microprocessors, system board, and hard disk drive backplane.

Notes:

- Monitored system components vary by Advanced System Management adapter or processor.
- VRM voltages are monitored but not displayed.

- **Voltages**

Includes +5.9, +3.3, +12.0, and -12.0 volt power- supply voltages.

- **System Status**

Includes system state, including operating system started, operating system running, POST started, POST stopped (error detected), and system turned off/state unknown, system power status (on or off), and power- on hours (the total number of hours that the system has been turned on, a cumulative count of all running hours, not a count of hours since the last system restart).

Configuration Information

Click the + beside Configuration Information or double-click **Configuration Information** to expand the Configuration Information tree and view the Configuration Information components. To view Configuration Information, double-click a Configuration Information component. Configuration Information is available for four Advanced System Management subsystems:

- **System Management Processor Information**

Provides information about the Advanced System Management PCI adapter or processor, including Advanced System

Management processor microcode build ID, revision number, file name, and date; device driver version number; and Advanced System Management processor hardware revision number.

- **System Vital Product Data**

Provides information about the system that contains the Advanced System Management PCI adapter or processor, including build ID, unique number, system board identifier, machine type and model power-controller microcode revision level, and front-panel microcode revision level.

- **System Card Vital Product Data**

Provides information, such as field replaceable unit (FRU) number, unique number, manufacturing ID, and slot number, about a variety of individual components that are installed in the remote system, including processor card, central processing units (CPU), memory cards, power supplies, power backplane, front panel, I/O backplane, I/O card, DASD backplane, and system management subsystem.

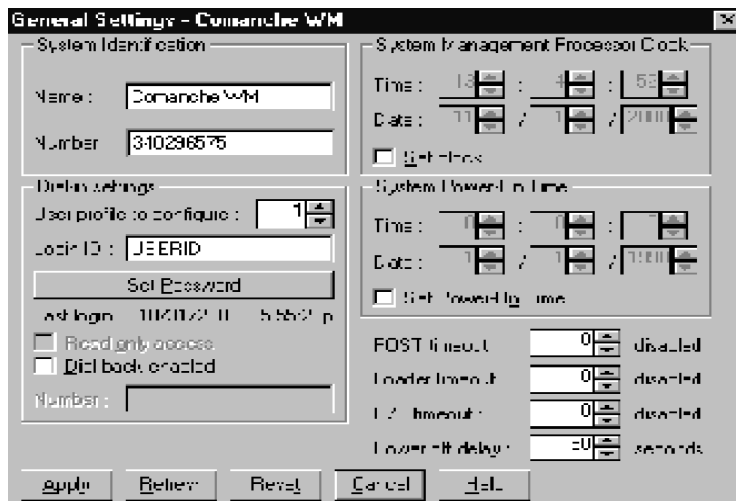
- **Memory DIMM Information**

Provides information on dual in-line memory modules that are installed in the system, such as memory type, size, and speed (frequency).

Configuration Settings

Click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Use the Configuration Settings components to configure:

- General settings
- Remote event settings
- Modem settings
- Network settings
- SNMP settings



Double-click one of the Configuration Settings components to view or change the configuration of the selected component. For more information, see the Configuration Settings component-specific sections that follow.

General Settings

Double click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click **General Settings** to open the Configuration Settings window. The Configuration Settings window contains the following groups or fields:

- System Identification
- Dial-in settings
- System Management Processor Clock
- POST timeout
- Loader timeout
- O/S timeout
- Power off delay

The System Identification Group

The System Identification group box contains two fields to help you identify the system that contains the Advanced System Management PCI adapter or processor.

Field	Description
Name	Use this field to provide a name for the system, the name of the system user, or the name of a contact. This information is included with forwarded events and with messages that are sent to alphanumeric pagers to help you identify the system that generated the event
Number	Use this field to identify the system with a specific serial or identification number, to record the phone number that is used to dial into the system, or to provide the phone number of a contact. This information is included with forwarded events and with messages that are sent to numeric pagers to help you identify the system that generated the event.

To change the System Identification information, use the following procedure:

1. In the Name or Number text box, type the system information that you want to record.
2. Click **Apply** to save this information.

The Dial-in Settings Group

Use the fields in the Dial-In settings group box to enable or disable dial-in support and to enable users to dial in and access the Advanced System Management PCI Adapter. The Dial-In setting group box contains the following items.

Item	Description
User Profile to Configure	Use the buttons to select the user profile that you want to configure. This service supports up to 12 separate profiles.

Item	Description
Login ID	<p>In this field, type the login ID that will be used by the remote user. Up to 12 login IDs can be configured. (This field is case sensitive.)</p> <p>Note: To remotely access the Advanced System Management PCI adapter specify a login ID.</p>
Set Password	<p>A password must be provided along with the login ID to allow a remote user to access the Advanced System Management PCI adapter or processor. After providing a login ID, click Set Password to open the Set password window. (The fields in the Set Password window are case sensitive.)</p> <p>Note: This password must be 5 to 8 characters in length and must contain a non-alphabetic character.</p>
Last Login	<p>This shows the date and time of the last successful login by a remote user.</p>
Read only access	<p>If the Read Only Access check box is selected, the user whose profile is selected will not be able to alter any of the Advanced System Management PCI adapter or processor settings when access is granted. The user will, however, be able to see all currently configured settings and values except passwords.</p>

Item	Description
Dial Back Enable	If the Dial Back Enabled check box is selected, the Advanced System Management PCI adapter will automatically terminate the connection as soon as the user whose profile is selected logs in, and will then use the telephone number in the Number text box to dial out and attempt to connect with a remote system.

To create a new login ID for a remote user, use the following procedure:

1. In the Login ID text box, type the ID that the remote user will use. This ID can be up to eight characters.
2. Click **Set Password** to open the Set Password window. Remote users must provide a password along with a login ID to access the Advanced System Management PCI adapter.
3. From the Set Password window:
 - a. In the Enter Password text box, type a password.
Note: This password must be 5 to 8 characters in length and must contain at least one non-alphabetic character.
 - b. In the Re-enter Password text box, type the same password that you typed in the Enter Password text box.
 - c. Click **OK** to save this password and close the Set Password window.
4. Click **Apply** to save the new user ID.

To delete the currently selected login ID, use the following procedure:

1. Use the buttons beside the User ID to Configure text box to select a previously configured user profile.
2. Click the Login ID text box.
3. Using the **Backspace** or **Delete** key, delete the currently displayed login ID.
4. Click **Apply** to remove the user ID.

The System Management Processor Clock Group

Use the selections that are available in the System Management Processor Clock group box to set the time and date that is reported by your Advanced System Management PCI adapter or processor.

Note: The Advanced System Management processor clock is separate from and independent of the system clock. Changes that are made to this setting will have no effect on the system clock.

To change the time or date, use the following procedure:

1. Verify that the **Set Clock** check box is selected. You must select this check box to enable the Advanced System Management service to change the currently stored time and date values.
2. Use the drop-down list beside each field to set the time or date.
 - The Time text boxes represent, when viewed from left to right, hours, minutes, and seconds.
 - The Date text boxes represent, when viewed from left to right, month, date, and year.
3. Click **Apply** to save the new time and date.

POST Timeout

The POST Timeout text box shows the number of seconds that the Advanced System Management PCI adapter or processor will wait for the system power-on self-test (POST) to complete before generating a POST timeout event. If POST takes longer than the configured amount of time to complete and the POST Timeout check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the IBM Advanced System Management PCI adapter or processor will automatically restart the system one time and will attempt to forward an event to all enabled remote event entries. After the system restarts, POST Timeout is automatically disabled until the system is properly shut down and restarted.

Note: If a POST timeout occurs and you have not selected this check box, the system will restart, but no event will be forwarded.

To set the POST timeout value, use the drop-down list beside the POST Timeout text box to set the number of seconds that the IBM Advanced

System Management PCI adapter or processor will wait for POST to complete. Then, click **Apply** to save this value. The maximum POST timeout value that you can set is 7650 seconds. Set this value to 0 to disable POST timeout detection.

Loader Timeout

The Loader timeout field displays the number of seconds that the Advanced System Management PCI adapter or processor will wait for the system loading process to complete before generating a Loader timeout event. The Loader timeout measures the amount of time that passes between the completion of POST and the end of operating system startup. If this takes longer than the configured amount box the Loader timeout check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the Advanced Systems Management PCI adapter will automatically restart the system one time and will attempt to forward an event to all enabled remote event entries. After the system restarts, Loader timeout is automatically disabled until the system is properly shut down and restarted.

Note: The system will restart, but no event will be forwarded, if you do not select the Loader timeout check box and the system detects a Loader timeout.

To set the Loader timeout value, use the Loader timeout drop-down list to set the number of seconds that the Advanced Systems Management PCI adapter or processor will wait between POST completion and operating system startup before generating a timeout event. Then, click **Apply** to save this value. The maximum Loader timeout value that can be set is 7650 seconds. Set this value to 0 to disable Loader time-out detection.

O/S Timeout

A periodic signal is sent from the Advanced System Management PCI adapter or processor to the operating system to confirm that the operating system is running properly. The operating system Timeout event occurs when the operating system does not respond to the signal within 6 seconds. The O/S Timeout text box shows the number of seconds that the Advanced System Management PCI adapter or processor will wait between operating system time-out checks. If the

operating system fails to respond within 6 seconds, the Advanced System Management PCI adapter or processor will attempt to restart the system, and if the O/S Timeout check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the IBM Advanced System Management PCI adapter will automatically restart the system one time and will attempt to forward an event to all enabled remote event entries.

Note: If you do not select the timeout check box and the system detects an O/S timeout, the system will restart, but no event will be forwarded.

To set the timeout value, use the O/S time-out drop-down list to set the number of seconds that the Advanced System Management PCI adapter will wait between checks. Then, click **Apply** to save this value. The maximum O/S timeout value that can be set is 255 seconds. Set this value to 0 to disable O/S timeout detection.

Power Off Delay

The “Power Off Delay” text box displays the number of seconds that the IBM Advanced System Management PCI adapter or processor will wait for the operating-system shutdown process to complete before turning off the system.

When the Advanced System Management PCI adapter or processor initiates a shutdown procedure and the **Power Off** check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the Advanced System Management PCI adapter will automatically attempt to forward an event to all enabled remote event entries. This occurs after the system is turned off and the Power Off Delay time has passed.

To set the power-off delay value, use the Power off delay drop-down list to set the number of seconds that the Advanced System Management PCI adapter or processor will wait for the operating system shutdown to complete before turning off the system. Then, click **Apply** to save this value. The maximum power- off delay value that can be set is 9999 seconds. Set this value to 0 to disable the power-off delay.

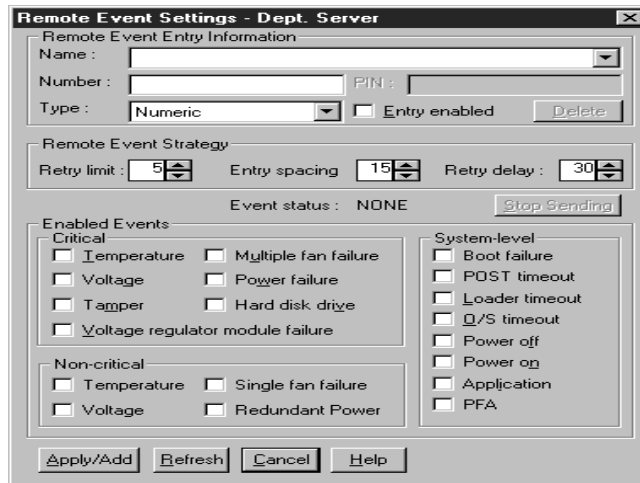
Other Configuration Settings Functions

The Configurations Settings window also includes these three buttons:

Button	Description
Refresh	Click Refresh to update all data that is shown on the Configuration Settings window, including date, time, and last login.
Reset	Click Reset to set all Advanced System Management settings back to their default values, including configuration settings, dial-out settings, and advanced dial-out settings. Note: All previously configured Advanced System Management settings will be permanently lost.
Cancel	Click Cancel to close this window without saving any changes.

Remote Event Settings

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click **Remote Event Settings** to open the Remote Event Settings window. Use the Remote Event Settings window to configure the Advanced System Management adapter or processor event forwarding functions. If you configure a remote event entry, the Advanced System Management adapter or processor will attempt to forward an event to a remote IBM system through a network connection, a numeric pager, an alphanumeric pager, or an SNMP community when any of the events that are selected from the Enabled Events group box occur. This event will contain information about the nature of the event that occurred, the time and date at which the event occurred, and the name of the system that generated the event.



Before Advanced System Management can forward events to SNMP communities, you must configure the Advanced System Management SNMP settings. To configure SNMP settings, double-click the **SNMP Settings** component of the Configuration Settings tree. The SNMP Settings window opens.

You can configure the Advanced System Management PCI adapter or processor to forward events to multiple pagers or IBM Director systems in response to individual critical and non-critical dial-out events. Therefore, the Event Status text will read **SENDING** as soon as the first event forwarding operation begins, change to **NONE** when the event forwarding operation is completed, change to **SENDING** again when the second event forwarding operation begins, change to **NONE** when the second event forwarding operation is completed, and so on. If you click **Stop Sending**, the Advanced System Management PCI adapter or processor abandons the currently active event-forwarding operation and moves to the next one.

You can configure your Advanced System Management PCI adapter with 12 or fewer separate remote event entries.

Remote Event Entry Information Group

To edit or create a remote event entry, use the following procedure:

1. In the Name text box, type the name of the person or system to which the event will be forwarded. The information in the Name text box is strictly for your use in identifying the remote event entry. If you are editing a previously configured remote event entry, select the entry that you want to edit from the Name drop-down list.
2. In the Number text box, type a telephone number (if you are forwarding the event to a pager) or an IP address (if you are forwarding the event to a system using the network adapter; this feature is supported only with an Advanced System Management PCI adapter) that will be used to forward an event.

Note: Depending on your paging service, you might need to increase the amount of time that this event action waits after dialing the telephone number before it transmits the numeric data. To increase the amount of time that will pass before the numeric data is transmitted, add one or more commas (,) to the end of the telephone number. Each comma will cause the modem to wait two seconds before transmitting the numeric data.

3. In the PIN text box, type the personal identification number that is required by your alphanumeric pager provider. This field will be active only if you select Alpha-numeric in the Type drop-down list.
4. From the Type drop-down list, select the type of connection the Advanced System Management PCI adapter or processor will attempt to make to forward the event notification. You can select **Numeric** (for standard pagers), **Alpha-numeric** (for alphanumeric pagers), or **IP** (for using a TCP/IP link to connect to a remote system; this is available only on systems with an Advanced System Management PCI adapter).
5. Select the **Entry Enabled** check box to activate this remote event entry. If the **Entry Enabled** check box is not selected, no events will be forwarded to this entry.

6. Select dial-out events from the Enabled Events group box. If any of the selected events occur, the Advanced System Management PCI adapter or processor will use the telephone number or IP address that is specified in the Number text box to forward an event describing the event using the method selected in the **Type** list.
7. Click **Apply/Add** to save these settings.

To remove a previously configured remote event entry, select the name of the entry from the Name drop-down list and then click **Delete**.

Remote Event Strategy Group

Use the selections in the Remove Event Strategy group box to specify the number of times that Advanced System Management will attempt to forward an event if an attempt fails, the amount of time that Advanced System Management will allow between event generation attempts, and the amount of time that Advanced System Management will wait between successive event-forwarding operations. The Remote Event Strategy group box contains the following items.

Item	Description
Retry limit	Use the drop-down list to select the number of additional times that Advanced System Management will attempt to forward an event. The dial-out retry limit applies only to attempts to forward the event information to an alphanumeric pager. If you are forwarding the event information to a numeric pager, only one attempt will be made to forward this information. The maximum value for this text box is 8.
Entry spacing	If you have configured more than one remote event entry to forward events, the Advanced System Management PCI adapter will attempt to contact each of these entries sequentially. Use the drop-down list to specify the number of seconds for the Advanced System Management PCI adapter or processor to wait between dial-out attempts for separate remote event entries. The minimum value for this field is 15 seconds, and the maximum value is 120 seconds.

Item	Description
Retry delay	Use the drop-down list to specify the number of seconds that Advanced System Management will wait before retrying a dial-out attempt. The minimum value for this field is 30 seconds, and the maximum value is 240 seconds.

Enabled Events Group

Use the selections box in the Enabled Events group box to specify which events will result in all currently configured entries being contacted by the Advanced System Management PCI adapter or processor. Any selected items will, if detected, result in an event describing the event being forwarded, using the method that is selected in the Type text box, to the recipient specified by the Name text box in the Remote Event Entry window.

Advanced System Management events that are forwarded to a pager will include information about the event that triggered the event. If the event is forwarded to a numeric (or standard) pager, the pager will include a code number that corresponds to the triggering event. If the event is forwarded to an alphanumeric pager, the page will include both a code number and a text string that describe the triggering event. For more information on the numeric codes and text strings that are transmitted to pagers, see the tables on pages 96 through 100.

All numeric codes and text strings are included in forwarded Manager events, regardless of whether they are forwarded using a serial or TCP/IP link. All information is also included in forwarded SNMP events.

The Enabled Events group box is divided into the Critical, Non-critical, and System groups.

The Critical Enabled Events group box contains the following items.

Item	Description (if checked)	Numeric Code	Text String
Temperature	The Advanced System Management PCI adapter or processor will forward an event and then automatically initiate a system shutdown if any monitored temperatures exceed their threshold values.	00	System Shutdown Due to Temperature
Voltage	The Advanced System Management PCI adapter or processor will forward an event if the voltages of any monitored power sources fall outside their specified operational ranges.	01	System Shutdown Due to Voltage
Tamper	The Advanced System Management PCI adapter or processor will forward an event if the voltages of any monitored power sources fall outside their specified operational ranges.	02	System Tamper Event
Voltage regulator module failure	The Advanced System Management PCI adapter or processor will forward an event and then automatically initiate a system shutdown if the voltage regulator module (VRM) fails.	06	VRM Failure

Item	Description (if checked)	Numeric Code	Text String
Multiple fan failure	The Advanced System Management PCI adapter or processor will forward an event if two (or more) of the cooling fans fail and will automatically initiate a system shutdown.	03	Multiple System Fan Failures
Power failure	The Advanced System Management PCI adapter or processor will forward an event if the power supply fails.	04	Power Supply Failure
Hard disk drive	The Advanced System Management PCI adapter or processor will forward an event if one or more of the hard disk drives in the system fail.	05	DASD Fault

The Non-critical Enabled Events group box contains the following items.

Item	Description	Numeric Code	Text String
Temperature	The Advanced System Management PCI adapter or processor will forward an event if any monitored temperature exceeds its threshold value. However, unlike the Critical Temperature event, this event will not initiate a system shutdown automatically.	12	Non-Critical Temperature Threshold Exceeded

Item	Description	Numeric Code	Text String
Voltage	The Advanced System Management PCI adapter or processor will forward an event if any monitored voltage exceeds its threshold value.	13	Voltage
Single fan failure	The Advanced System Management PCI adapter or processor will forward an event if one of the systems cooling fans fails.	11	Single Fan Failure
Redundant Power	The Advanced System Management PCI adapter or processor will forward an event if the redundant power system fails.	10	Power Redundancy has been compromised, please check the system management processor error log for more information.

The System Enabled Events group contains the following items.

Item	Description (if checked)	Numeric Code	Text String
Boot Failure	The Advanced System Management PCI adapter or processor will forward an event if the system fails to start.	25	Requires POST interaction. POST detected error which prevents the system from booting.

Item	Description (if checked)	Numeric Code	Text String
POST time-out	The Advanced System Management PCI adapter or processor will forward an event if the timeout value (specified in the Configuration Settings window) is exceeded.	20	POST/BIOS Watchdog expired. System restarted.
O/S time-out	The Advanced System Management PCI adapter or processor will forward an event if the O/S System timeout value (specified in the Configuration Settings window) is exceeded.	21	Operating system watchdog expired. System restarted.
Loader time-out	Advanced System Management PCI adapter or processor timeout value (specified in the Configuration Settings window) is exceeded.	26	Loader watchdog expired. System restarted.
Power off	The Advanced System Management PCI adapter or processor will turn OFF an event if the system is powered off.	23	System complex powered off.
Power on	The Advanced System Management PCI adapter or processor will forward an event if the system is turned on.	24	System complex powered on.
Application	The Advanced System Management PCI adapter or processor will forward an event if it receives an event.	22	Application logged event.

Item	Description (if checked)	Numeric Code	Text String
PFA	Advanced System Management PCI adapter will forward an event if it receives a Predictive Failure Analysis event from the system.	27	PFA

Modem Settings

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click **Modem Settings** to open the Modem Settings window. Use the Modem Settings window to specify modem and dialing settings. The Modem Settings window contains the following groups or fields:

- Port Configuration
- Dialing Settings

Modem Settings - Dept. Server

Port Configuration

Port to configure : 2 Baud rate : 57600

Initialization string : ATZ

Hangup string : ATH0

Port selected Advanced

Dialing Settings

Dial-in enabled

Dial-in delay : 2

Apply Refresh Cancel Help

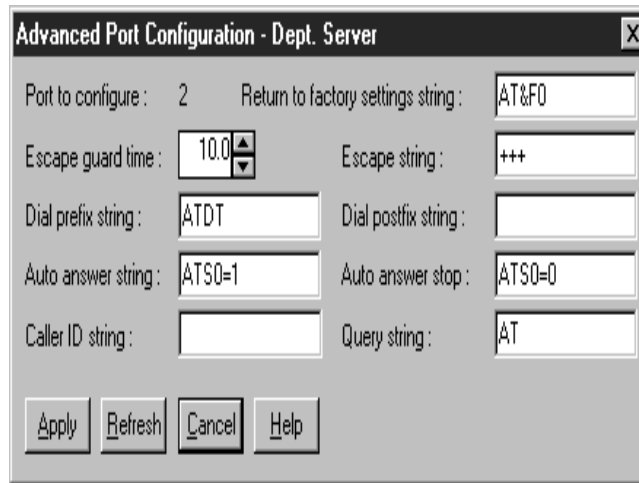
The Port Configuration Group Box

Use the Port Configuration group box to specify and configure the modem or port that will be used to forward an event when an Advanced System Management event occurs. The Port Configuration group box contains the following items.

Item	Description
Port to configure	Use the drop-down list to select the port that your modem is configured to use. This drop-down list will show only values that are available for use by your Advanced System Management PCI adapter or processor. The port that you select to use affects the availability of the modem for use by either the Advanced System Management PCI adapter or processor or the operating system. The ports that are available to the system or to the Advanced System Management service vary depending on your hardware configuration.
Baud Rate	Use the drop-down list to specify the baud rate for the serial port.
Initialization string	Type the initialization string that will be used for the specified modem. A default string (ATEO) is provided. Do not change this string unless your dial-out functions are not working properly.
Hangup string	Type the initialization string that will be used to instruct the modem to disconnect. A default string (ATHO) is provided. Do not change this string unless your dial-out functions are not working properly.
Port selected	This check box indicates whether the port number that is currently displayed in the Port to Configure list is the port that is currently designated for use by the Advanced System Management PCI adapter or processor. Select this check box if you want to configure the Advanced System Management PCI adapter or processor to use the currently displayed port number.
Advanced	Click this button to open the Advanced Port Configuration window.

The **Advanced Port Configuration** window contains the following items.

Item	Description
Return to factory settings string	Type the initialization string that returns the modem to its factory settings when the modem is initialized. The default is AT&F0.
Escape guard time	In this field, type the length of time TIME ³ before and after the escape string is issued to the modem. This value is measured in 10- millisecond intervals. The default value is 1 second.
Escape string	Type the initialization string that returns the modem to command mode when it is currently communicating with another modem (connected). The default is +++.
Dial prefix string	Type the initialization string that STRING ³ is used before the number to be dialed. The default is ATDT.
Dial postfix string	Type the initialization string that STRING ³ is used after the number is dialed to tell the modem to stop dialing. The default is the Carriage Return character or ^M.
Autoanswer string	Type the initialization string that STRING ³ is used to tell modem to answer the phone when it rings. The default is to answer after two rings or ATSO=2.
Autoanswer stop	Type the initialization string that STOP ³ is used to tell the modem to stop answering the phone automatically when it rings. The default is ATSO=0.
Caller ID string	Type the initialization string that will be used to get Caller ID information from the modem.
Query string	Type the initialization string that is used to find out if the modem is attached. The default is AT.



The port that you select to use affects the availability of the modem for use by either the Advanced System Management PCI adapter or processor or the operating system. The ports that are available to the system or to the Advanced System Management service vary depending on your hardware configuration.

- If the system has an Advanced System Management processor only, use this table to determine what ports are available.

Physical Ports (as labeled)	A	B	C
Ports available to Advanced System Management	Port 1 shared	N/A	Port 2 dedicated
Ports available to operating system	COM 1 shared	COM 2	N/A

- If the system has an Advanced System Management PCI adapter only, use this table to determine what ports are available.

Physical Ports (as labeled)	A	B	MODEM	COM_AUX
Ports available to Advanced System Management	N/A	N/A	Port 1 shared	Port 2 dedicated
Ports available to operating system	COM 1	COM 2	COM 3 shared	N/A

- If the system has an Advanced System Management PCI adapter and an Advanced System Management processor, use this table to determine what ports are available.

Physical Ports (as labeled)	A	B	C	MODEM	COM_AUX
Ports available to Advanced System Management	Port 1 shared	N/S	Port 2 dedicated	N/A	N/A
Ports available to operating system	COM 1 Shared	COM 2	N/A	N/A	N/A

Notes:

- The operating system recognizes shared ports when the system is running. The Advanced System Management PCI adapter or processor recognizes shared ports when the system starts up or turns off. The system recognizes the shared port, but the shared port is not recognized by the Advanced System Management PCI adapter or processor, when the system is started with DOS.
- In a system with the Advanced System Management PCI adapter only, the device driver must be running for the operating system to recognize COM3.

Dialing Settings

Use the Dialing Settings group box to specify settings that are related to the modem and to configure the modem that is used to forward an event when an Advanced System Management event occurs. The Dialing Settings group box contains the following items.

Item	Description
Dial-in enabled	Select this check box to enable remote users to dial into and access the Advanced System Management PCI adapter. If this check box is cleared, remote users will be unable to remotely access the Advanced System Management PCI adapter. Click Apply after selecting or clearing this check box to save the new setting.

Item	Description
Own port on startup	<p>Select this check box to reserve a serial port for exclusive use by the IBM Advanced System Management PCI adapter or processor. Selecting this check box will reserve one of the integrated communications ports in the adapter. Click Apply after selecting or clearing this check box to save the new setting.</p> <p>Notes:</p> <ul style="list-style-type: none"> • Select this check box if you are configuring your system for dial-in access. If this check box is not selected, you will be unable to dial into this system unless the adapter has reclaimed the port for a dial-out. If you want to configure the Advanced System Management PCI adapter or processor to always be dial-in enabled, regardless of whether the system is on, you must select this check box. When this check box is selected, you cannot configure the specified port for use by your system. • Port C is dedicated for use by the Advanced System Management PCI adapter only. The operating system uses Port A. Port A is available to the Advanced System Management PCI adapter only when one of the following conditions exists: <ul style="list-style-type: none"> ■ The server is off. ■ The Advanced System Management PCI adapter needs a port to perform a critical enabled event (in this case, the IBM Advanced System Management PCI adapter seizes control of the port from the operating system, dials out, and then turns off the server to avoid damage to your hardware).

Item	Description
Dial-in delay	The Dial-In Delay Minutes text box shows the number of minutes that must pass after an incorrect user ID or password has been used in six successive dial-in attempts before valid dial-in access will be permitted. After the sixth successive login failure, dial-in access is disabled for the number of minutes that you specify, the Advanced System Management PCI adapter or processor adds an entry in the Event Log noting that dial-in access was suspended because of six successive login failures, and the Advanced System Management PCI adapter or processor attempts to forward an event if the Tamper Enabled Events check box has been selected. The minimum value for this field is 4 minutes, and the maximum value is 240 minutes.

Initialization String Guidelines

If you need to provide a new initialization string, refer to the user's guide that comes with your modem. Your initialization string must contain commands that configure your modem as follows:

- Command echoing OFF
- Online character echoing OFF
- Result codes ENABLED
- Verbal result codes ENABLED
- All codes and connect messages with BUSY and DT detection
- Protocol identifiers added - LAPM/MNP/NONE V42bis/MNP5
- Normal CD operations
- DTR ON-OFF hang-up, disable AA and return to command mode
- CTS hardware flow control
- RTS control of receive data to computer
- Queued and nondestructive break, no escape state

The abbreviations in these commands have the following meanings:

AA Auto Answer

-
- CD** Carrier Detect
 - CTS** Clear to Send
 - DT** Data Transfer
 - DTR** Data Terminal Ready
 - RTS** Ready to Send

Network Settings

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click **Network Settings** to open the Network Settings window. Use the Network Settings window to specify network settings on the Advanced System Management PCI adapter.

Note: This window is available only when you are using the Advanced System Management service to manage a system that has a Advanced System Management PCI adapter or if you have used Advanced System Management to establish a TCP/IP, serial, or ASM Interconnect connection with a remote Advanced System Management PCI adapter.

The screenshot shows a window titled "Network Settings - Comanche WM". It contains the following fields and controls:

- Network interface: 1 (with a refresh icon) and a checked "Interface enabled" checkbox.
- Host name: (empty text field)
- IP address: 130.57.8.203
- Subnet mask: 255.255.248.0
- Gateway: 130.57.15.254
- Line type: Ethernet (dropdown menu) and an unchecked "Disable Routing" checkbox.
- Data rate: AUTO (dropdown menu)
- Duplex: AUTO (dropdown menu)
- MTU size: 1514 (dropdown menu)
- MAC address: 00:00:00:00:00:00 (text field)
- Buttons at the bottom: Apply/Restart, Refresh, Cancel, and Help.

The Network Settings window contains the following items.

Item	Description
Network interface	Use the drop-down list to select a network interface to configure. When you have selected a network interface, select the Interface Enabled check box.
Host name	Type the TCP/IP host name that will be used by the Advanced System Management PCI adapter.
IP address	Type the IP address that will be used by the Advanced System Management PCI adapter.
Subnet mask	Type the subnet mask that will be used by the Advanced System Management PCI adapter.
Gateway	Type the TCP/IP address of the gateway that will be used by the Advanced System Management PCI adapter.
Line type	Use the drop-down list to select the line type that will be used by the Advanced System Management PCI adapter. Available selections are Ethernet, PPP, and Token Ring. Select the Disable Routing check box if necessary.
Data rate	Use the drop-down list to select the data rate that will be used by the Advanced System Management PCI adapter. Available selections are AUTO, 4M, 16M, 10M, and 100M.
Duplex	Use the drop-down list to select the duplex method that will be used by the Advanced System Management PCI adapter. Available selections are AUTO, FULL and HALF.
MTU size	Use the drop-down list to specify the maximum transmission unit (MTU) value that will be used by the Advanced System Management PCI adapter.
MAC address	Type the media access control (MAC) address of the network adapter that is being used by the IBM Advanced System Management PCI adapter.

Attention: If you have installed the IBM Advanced System Management Token Ring Connection, you must not enable or use the Ethernet port that is included on your Advanced System Management PCI adapter. Enabling the Ethernet port while the IBM Turbo 16/4 Token Ring PCMCIA card is installed on your adapter will cause your system to become unstable. To enable, configure, or use the Ethernet port, you must first remove the IBM Turbo 16/4 Token Ring PCMCIA card from your Advanced System Management PCI adapter.

SNMP Settings

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click on **SNMP Settings** to open the SNMP Settings window. Use the SNMP Settings window to specify SNMP settings for the IBM Advanced System Management PCI adapter. These settings must be configured correctly for the Advanced System Management PCI adapter to forward events to SNMP managers on the network.

Note: This window is available only when you are using the Advanced System Management service to manage a system that has an Advanced System Management PCI adapter or if you have used Advanced System Management to establish a TCP/IP, serial, or ASM Interconnect connection with a Advanced System Management PCI adapter. This window is not available on systems that do not have an Advanced System Management PCI adapter.

The SNMP Settings window contains the following items.

Item	Description
SNMP agent enabled	Select this check box to enable the Advanced System Management PCI adapter to forward events to SNMP managers on your network.
Traps disable	Select this check box to prevent SNMP traps from being sent.
System contact	Type the name of the SNMP system contact in the text box.
System location	Type information regarding the location of your system in the text box.
Community	Use the drop-down list to select and define up to three SNMP communities. Note: SNMP events are sent only to the currently selected SNMP communities.
Community name	Type the name of the selected SNMP community in the text box.

Item	Description
Community IP address 1, 2 and 3	Type the IP addresses for the selected SNMP communities in the text boxes.

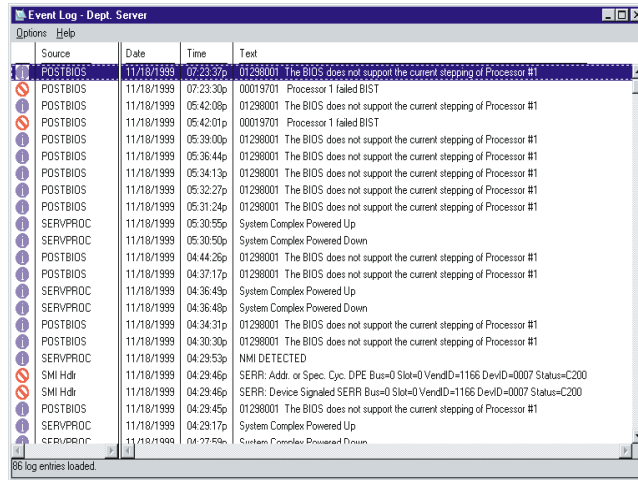
After making any changes to these settings, click **Apply** to save the changes. Then, close this window and click **Restart** in the Network Settings window. You must restart the adapter before changes to network settings on a Advanced System Management PCI adapter will take effect.

Event Log

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Click **Event Log** to open the Event Log window. This window contains all entries that are currently stored in the Advanced System Management PCI adapter or processor Event Log. The Advanced System Management adapter or processor Event Log records information about all remote access attempts and dial-out events.

Notes:

- If you are using the Advanced System Management service with an Advanced System Management PCI adapter installed in an IBM server the Event Log might contain entries that begin with the text `12 C Message`. These messages are normal. Services use these messages in the event of system problems.
- If you are using the Advanced System Management service with a Netfinity 8500 M10, 7000 M10, 7000 M20, 5600, 5500, 5500 M10, 5500 M20, or 5000 server, the Event Log will also include any POST error messages.



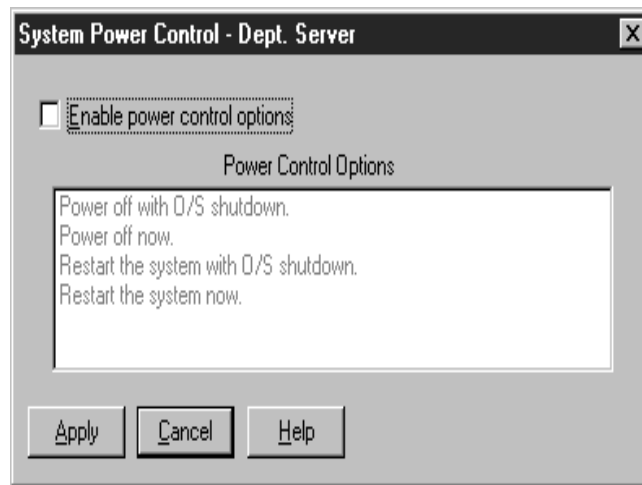
The following functions are available from the Options menu in the Event Log window:

Option	Description
Load	Refreshes the contents of the Event Log window.
Print to file	Saves the contents of the Event Log window to a text file.
Print to printer	Sends the contents of the Event Log window to a printer that is attached to your system.
Clear log	Erases all entries that are currently stored in the Event Log (including any entries that are not currently visible in the Event Log window).

Note: After you use **Clear Log** to erase the entries in the Event Log, they are permanently erased and cannot be retrieved.

System Power Control

From the Advanced System Management console, click the + beside Configuration Settings or double-click **Configuration Settings** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click on **System Power Control** to open the System Power Control window. Use the System Power Control window to instruct the Advanced System Management PCI adapter or processor to turn off the system, restart the system, or turn on the server. To initiate a power control options, select the **Enable Power Control Options** check box. If the check box is not selected, the Power Control Options text box will not be available.



The following System Power Control functions are available at all times.

Function	Description
Power off with O/S shutdown	Performs an operating-system shutdown before removing power from the system.
Power off now	Immediately removes power from the system.

Function	Description
Restart the system with O/S shutdown	Performs an operating system shutdown, removes power from the system, and then restores power to the system.
Restart the system now	Immediately removes power from the system, and then restores power to the system.

If directly connected to the IBM Advanced System Management PCI adapter through a TCP/IP link, the Power On Now selection will be available. This function turns on the server and enables the microprocessor to perform POST, loading, and operating-system startup procedures.

To initiate a system Power Control function, use the following procedure:

1. Select the **Enable Power Control Options** check box.
2. From the **Power Control Options** list, select the power-control option that you want to activate.
3. Click **Apply**.

Remote POST Console

You can use the Advanced System Management Remote POST Console function to remotely monitor, record, and replay all textual output that is generated during POST. To monitor and record the POST data on a remote system, use the following procedure:

1. Connect to the remote Advanced System Management PCI adapter or processor.
2. Open the Remote POST window.
3. Restart the remote system by using the Advanced System Management System Power Control functions.

The Remote POST Console displays and records all POST data as the remote system completes POST. While you are monitoring POST on a remote system, all local keystrokes are relayed automatically to the

remote system, enabling you to interact with the POST process on the remote system.

To review data after POST is completed, disconnect from the remote system and use the Replay functions.

Use the selections in the Replay menu to replay the text output that was captured during the last Remote POST operation. All text that was displayed by the remote system during POST will be displayed as it appeared on the remote system.

- To begin playing the recorded POST data, or to resume playing the recorded POST data after stopping playback, click **Replay Last POST**.
- To stop playback of the recorded POST data, click **Stop**.
- To resume viewing the recorded POST data from the beginning, click **Restart**.
- Click **Fast**, **Medium**, or **Slow** to specify the speed at which the recorded POST data is displayed in the Remote POST window.

Note: You can replay Remote POST data when not connected to a remote system Advanced System Management PCI adapter or processor.

Updating PCI Adapter or Processor Microcode

Note: Updating the Advanced System Management PCI adapter microcode, resets the default user name and password. If you had previously changed them, you will need to change them again.

To update the Advanced System Management PCI adapter or processor microcode, use the following procedure:

1. From the Advanced System Management console, click **Options** → **Update Microcode** → **System Management**.
An Insert Diskette window opens.
2. Insert the **System Management microcode** update diskette into the diskette drive.

3. Click **OK** to continue.
Warning notices will appear, asking that you verify if you want to continue. Click **OK** to continue or **Cancel** to stop the microcode update process.

When you have verified that you want to proceed with updating the Advanced System Management PCI adapter or processor microcode, the Advanced System Management service will apply the microcode update to the Advanced System Management PCI adapter or processor.

During this process, some of the monitoring functions of the Advanced System Management PCI adapter or processor (such as the environmental monitors) will be disabled. After you have updated the microcode, all system monitoring will resume.

Updating System POST/BIOS Microcode

To use Advanced System Management and your Advanced System Management PCI adapter or processor to update the system POST/BIOS microcode, use the following procedure:

1. From the Advanced System Management console click **Options** → **Update Microcode** → **System POST/BIOS**.
An Insert Diskette window opens.
2. Insert the **System POST/BIOS update** diskette into the diskette drive.
3. Click **OK** to continue.
Warning notices will appear, asking that you verify that you want to continue. Click **OK** to continue or **Cancel** to stop the microcode update process.

When you have verified that you want to proceed with updating the system POST/BIOS microcode, the Advanced System Management service will apply the microcode update to the system that contains the Advanced System Management PCI adapter or processor to which you are connected.

5

Capacity Manager

Capacity Manager is an easy-to-use resource-management and planning tool for network managers, and administrators. It enables remote performance monitoring of every server on the network. IBM Capacity Manager identifies potential bottlenecks in a network, enabling effective planning of future capacity needs, such as microprocessor, disk, or memory upgrades, thus preventing network slowdowns and downtime. With Capacity Manager you can plan for future hardware upgrades.

Capacity Manager includes extensive online help, including an online tour. The tour is an interactive help that guides you through the Capacity Manager functions, making it especially simple to learn and understand this service. To begin a tour, click **Report Viewer tour** from the Using Capacity Manager task on the IBM Director Console.

Note: The Capacity Manager interface is available for use only on systems running Windows NT. However, you can collect data from any remote system running Client Services for IBM Manager for OS/2, Windows 2000, Windows NT, or NetWare.

Components of Capacity Manager

Capacity Manager contains several features within the IBM Director Extension tools package:

- You can manage your systems from a server with Capacity Manager installed or access Capacity Manger functionality remotely.

-
- Capacity Manager can generate events if a bottleneck is detected. Each hour, if a new bottleneck begins to occur on any system in your network, Capacity Manager can take any action that you specify. For instance, it can notify you of the bottleneck via e-mail or pager.
 - The Monitor Activator task provides a single console where you can manage your systems. Actions within the task include:
 - All NT PerfMon monitors.
 - Monitor Activator informative icons that you can use to learn whether a monitor is active, inactive, or not present on a given system, whether Capacity Manager is running, or if a system is busy, secure, offline, or of an unknown status.
 - Automatic activation by default of the Performance Analysis monitors that are present on your systems. You can activate additional monitors and deactivate them at any later time without having to edit an initialization (.ini) file or restart your systems.
 - The Report Generator generates a report directly to the viewer for immediate viewing or generate a report to a file for later viewing.
 - Reports that are generated to the viewer are created very quickly. These reports are not saved to the disk. You view these reports before deciding to save them. This will keep your disk from being cluttered with every generated report.
 - Generation performance is faster because of the improved error detection. For example, when generating a report to the viewer, you will first see a status window that tabulates the status of each system during the generation. Without waiting for the time-out to elapse, Capacity Manager will use its improved diagnostics to report why any system does not respond, so that you can more quickly remedy the problem.
 - The Report Definition interface provides usability and functionality. The interface includes:
 - The Report Parameters pane, which you can use to select a report duration, a global sampling frequency, and the days and times for collecting data.

-
- The Method of Generating a Report pane which you can use to choose between generating a report to the viewer or to a file. To help you more easily keep track of report files and know which can be merged, you can include in the file name the date and time of report generation in addition to the report definition name.
 - The Monitor Selection pane, which you can use instead of including all activated monitors on your report. In the Monitor Selection pane, you can individually activate or deactivate monitors and select their sampling frequencies.
 - The Timeout parameter at the bottom of the Report Definition interface which you can use to set how long each system has to respond for inclusion in a report.
 - The Report Viewer has the following options and performance enhancements:
 - Forecasting has a wavelet transform technique. It transforms the observed monitor data before the linear regression computations. The result is a 95% prediction interval for the forecasting graphs.
 - You can choose between saving a report or a file in graphical image format (.gif) as either a remote file on the IBM Director server or as a local file on your workstation. You manage both remote files and local files through the console. You can also access remote files from other workstations. There is also an option to manage local files with the command-line tools CMView and CMReport.
 - The sort option in the Table view enables you to sort by clicking a column header. For example, you can click the column header CPU Utilization to see your systems sorted by their CPU-utilization values. The sort order is dependent on the sort-order button that is clicked on the toolbar.
 - Capacity Manager supports clusters. When you use Monitor Activator and Report Generator on a cluster, the nodes of the cluster are included. When you view a report that includes clusters, a new cluster mode is available for analyzing the clusters as if they were individual systems. The Table view aggregates the monitor

data from the cluster nodes, providing a single-system image of the cluster. The table can also expand the clusters to show the data for the individual nodes at the same time. The Icons view and Hypergraph view also show clusters as single-system images.

Capturing Data

After you have a list of discovered systems, you can use the Monitor Activator task to learn the status of the monitors on your selected systems and select which monitors you want to be active or inactive. Capacity Manager will collect data on all the monitors that are both activated and present on your selected systems. The Performance Analysis monitors that are present on your system are activated by default when you install Capacity Manager.

When you initiate the Monitor Activator task, Capacity Manager opens a Monitor Activator window that lists your selected systems and all the monitors for those systems. In the left pane, you can select which monitors to activate or deactivate on your selected systems. By selecting one or more monitors in the left pane, you create in the upper-right pane a table of each selected monitor status on each system that is listed. In the lower-right pane you can see the legend for all the icons that are used in the Monitor Activator task.

Notes:

- Not all monitors are present on all systems. For example, drive E is present only on systems that have it installed.
- If you add or remove disk drives or local area network (LAN) adapters, be sure to rerun the **Monitor Activator** task. Not rerunning the Monitor Activator task after changes in the drives and LAN adapters could compromise the accuracy of the Performance Analysis function.

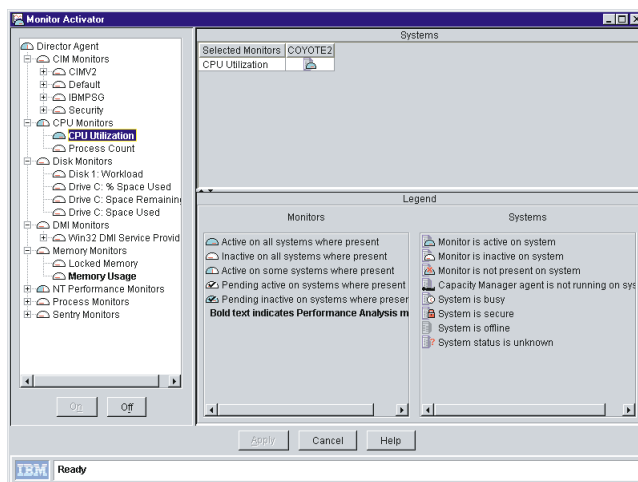
Starting the Monitor Activator

To start the Monitor Activator, drag and drop the Monitor Activator task onto a system or a group of systems. The Monitor Activator task includes all of the systems that you selected. If you drag and drop the Monitor Activator on a cluster, the nodes of the cluster will appear in the

Monitor Activator window. The cluster itself will not appear because clusters do not have Capacity Management agents. Cluster nodes have agents.




Using the Monitor Activator




The Monitor Activator has two main functions: to reveal the status of monitors on your selected systems, and to enable you to change the status of these monitors. These two functions are driven by actions that you take in Monitors pane. These actions are reflected in the Monitors pane or in the Systems pane.



When you first see the Monitors pane, it is in a collapsed tree structure. To see the individual monitors, expand the nodes of the tree. For example, the CPU Utilization monitor is at the end of the CPU Monitors branch node or group.

There is an icon to indicate the status of each monitor and each group. The meaning of the icon depends on whether it is beside a group or beside a monitor.

Group icons	
	All monitors in the group are active.
	Some monitors in the group are active.
	No monitors in the group are active.

Monitor icons:	
	This monitor is active on your selected systems where it is present.
	This monitor is inactive on your selected systems where it is present.
	This monitor is active on some of the systems where it is present. You will see this status if at least one system has this monitor in an inactive state.

Activating or deactivating monitors

To include a monitor, select it on the list of monitors, then click the **ON** button at the bottom of the pane. You can select more than one monitor at a time by pressing the **Control** key and clicking on multiple monitors. However, you cannot select multiple monitors by clicking on a group name. The **On** and **Off** buttons become available once you select one or more monitors on the list of monitors. You can double-click a monitor to toggle between on and off without clicking either the on or off



buttons. Once you include a monitor, the settings options become available.

To change the status of a monitor, click its name and click **on** to activate it or **off** button to deactivate it. To select multiple monitors at a time, **Ctrl**-click additional monitors.

As a safety feature, Capacity Manager will not allow you to select a group of monitors by clicking the group name. Select each monitor individually. After you have completed your changes, click **Apply** to apply the changes and close the window.

Note: Although you can go through and deactivate all your monitors on this panel and apply this change, the next time you open the monitor activator pane, you will see that a monitor still remains activated. You cannot delete all the monitors from the slot file.

If you decide to change the status of a monitor, the icon of the monitor will change to indicate its pending state.

Pending States	
	Your selected monitor will be activated on systems where it is present.
	Your selected monitor will be deactivated on systems where it is present.

Cluster Monitors






The Report Viewer Performance Analysis function probes for bottlenecks in server hardware performance, diagnoses the problem, and suggests ways to improve performance. To create this report on your selected system performance, the Performance Analysis (PA) function must have specific monitors activated in the Monitor Activator. These PA monitors are shown in bold text in the Monitors pane.




Monitoring for Clusters

The cluster monitors appear at the top level of the Monitor Activator tree. Unlike the other monitors, the cluster monitors provide textual, not numerical data. When you activate the cluster monitors, they display the ? character in the Report Viewer table. Because these monitors only provide textual data and cannot be graphed, you should not activate these monitors.

Creating a Status Table in the Systems Pane

To see the status of one or more monitors on each of your selected systems, build a status table in the Systems pane by selecting monitors in the Monitors pane. This table lists all your selected monitors on the left side and all your selected systems horizontally across the top.

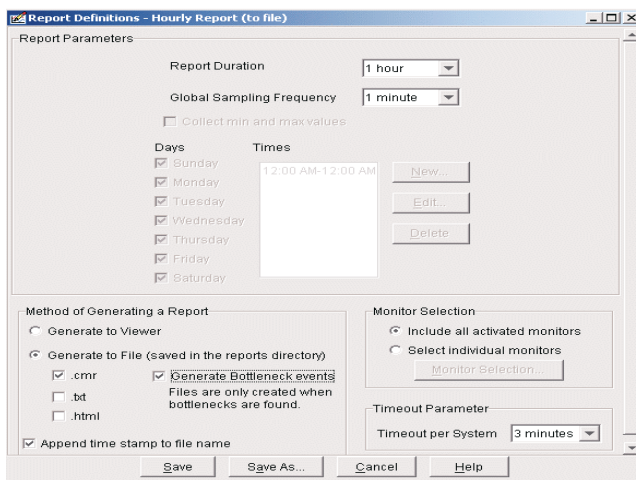
System pane icons	
	The monitor is present and active on this system.
	The monitor is present but not active on this system.
	The monitor is not present on this system.
	The agent is not running on this system. Either the Capacity Manager agent is not installed on this system or there is an error.
	This system is busy; the Capacity Manager agent is processing another request. Try again later.

	The system is secured.
	The system is offline.
	The system status is unknown because the Capacity Manager agent is not responding.

Viewing Data

Capacity Manager enables you to generate a file either to the Report Viewer for immediate viewing or to a file for later viewing. You must create a report definition before you can view the report.

To create a report definition, double-click **New Report Definition** under the Report Generator task. The Report Definition window opens.



The following table lists the parameters in the Report Definition window that are required to generate a report.

Item	Description
Report Duration	<p>Indicates the time span of a report, starting from the time of report generation and going back.</p> <p>1 day, ending at the beginning of the current hour.</p> <p>3 hours, includes data from the previous 3 hours including the current hour.</p> <p>8 hours, includes data from the previous 8 hours including the current hour.</p> <p>1 week, ending at the beginning of the current day.</p> <p>1 month, ending at the beginning of the current day.</p> <p>The monthly report can contain from 28 to 31 days, depending on the number of days in the previous month. For example, if today is the 24th and you select a duration of 30, a report will be generated using data that was collected from the 24th of the previous month through the 23rd of the current month.</p>
Global Sampling Frequency	<p>Indicates the sampling frequency of a report. Sampling frequency represents the quantity of data points that will be gathered per monitor and how often the data points will be collected (for example, once a minute or once an hour). Before you can input a new sampling frequency, you must first select the Override global settings option.</p>

Item	Description
Collect min and max values	Indicates whether CMAgent should include the minimum or maximum data points for the duration of the report period, or only the average. These data points include peak and trough values. Before you can elect to collect minimum or maximum data, you must first select the Override global settings option. An advantage to collecting the minimum and maximum data points is that you can collect data less frequently, reduce the size of the report, and still have informative system performance data.
Days	Indicates which days of the week to collect data.
Times	Indicates which hours of the day to collect data.
Method of Generating a Report	
Generate to Viewer	Indicates whether the report will be generated to view.
Generate to File (saved in the reports directory)	Indicates if the report will be generated to file.
.cmr	Indicates whether the output file should be written in CMR format.
.txt	Indicates whether the output file should be written in text format.
.html	Indicates whether the output file should be written in HTML format.
Generate Bottleneck events	Indicates whether a bottleneck event is generated when Performance Analysis finds a bottleneck.
Append time stamp to file name	Indicates whether the time stamp should be included in the name of the report.

Item	Description
Monitor Selection	
Include all activated monitors	All monitors that are activated in the Monitor Activator task will have the global sampling frequency that you set in the Report Parameters pane. This is the default.
Select individual monitors	Select the monitors that you want to be active among those monitors that you activated in the Monitor Activator task. Select the sampling frequency for specific monitors and the monitors from which you will collect minimum and maximum data.
Timeout Parameter	The time value represents the length of time Capacity Manager will wait for each system to respond when you generate a report.

Note: The following parameters affect the size of the generated report: Report Duration, Global Sampling Frequency, Min and Max values, Days, Times, and Monitor Selection. If you want the report to load faster, see the online help topic, “Improving Performance of the Report Viewer.”

In the Report Definition window, you can choose the time period for collecting data, the amount of data to collect, and the days and times for collecting data. You can also determine what monitor data will appear in the report. While you are in the report definition, you can also decide whether to generate the report to the viewer or to a file and to generate bottleneck events.

If you choose to generate the report to a file, the file is saved on the IBM Director server for later viewing. If you choose to generate the report to the viewer, this generated file will be kept in memory only for the time that you are viewing the report. You can also choose to generate bottleneck events.

Notice that the standard report definitions for weekly and hourly reports, are set to be generated to file and to viewer by default. You can change

these settings either by right-clicking the report definition name and selecting the alternative option, or you can double-click the report definition name and change the option on the Report Definition window.

If you are creating a new report definition with the New Report Definition task listed under Report Generator, click **to Viewer** or **to File** on the Report Definition window. Save your New Report Definition. These reports are labeled with your selected option. Your new report appears as a file under the Report Generator task.

To delete an existing report, right click the Reports icon and select **delete**.

Generating a Report

To generate a report, drag and drop the Report Definition icon onto the selected systems or clusters. When the Report Definition icon is dropped on a cluster, each of the cluster nodes is included in the report, and the cluster membership of each node is reported. However, if the report definition is dropped on a system that is a node in a cluster but not on the cluster itself, the node is treated as an individual system and its cluster membership is not recognized by Capacity Manager. If you drop the **Report Definitions** icon on both a cluster and on one of the nodes of the cluster, then the nodes are treated as part of the cluster.

Generating the Report to a File

When you initiate the generation step, a window opens, which will ask if you want to either create a scheduled job for this task or execute it immediately.

If you click **Schedule**, another dialog box opens, which will ask you for a job name, and the time and date for generating the report.

If you click **Execute Now**, a status box for the generation of your report is displayed. Any report that is saved to a file is saved automatically in the /reports subdirectory of the IBM Director installation directory.

You can create a report in more than one format, such as Capacity Manager Report (.cmr), text (.txt), or HTML (.html). Use the Capacity Manager Report Viewer to view .cmr and .txt files. A .cmr file loads more quickly than a .txt file, but most spreadsheet applications can use

.txt files. Some spreadsheet applications can use HTML files. The Report Viewer cannot use HTML files, but you can view an HTML file in a Web browser and then print it.

Generating the Report to the Viewer

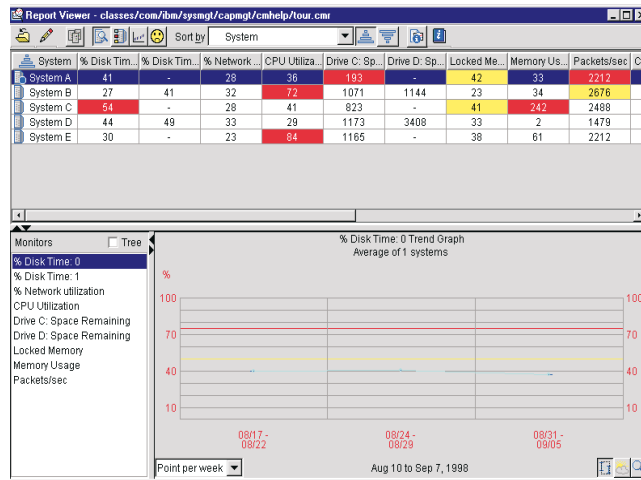
You can use the Report Viewer to see the performance of your systems, clusters or groups, to detect performance bottlenecks, and to see predictions of future performance. You can open, merge, and save reports and save files remotely (to the server) or locally (to the console system). You can also save graphs, look at report information, and open another Report Viewer window. There are two modes, system mode and cluster mode, for viewing performance information.

There are three ways to start the Report Viewer:

- Drag and drop the report definition task to be generated to a viewer that has been set to viewer onto a system, cluster or a group. This will generate a new report into the Report Viewer without saving it to disk.
- Double-click the Report Viewer task to view existing reports.
- Start the **Report Viewer** from the command line. This is useful for running the Report Viewer outside of IBM Director.

Click **Finish Now** to immediately terminate the report generation and open the Report Viewer to display the system data that was collected at that point.

The Report Viewer window has three panes (System, Monitor, Graph) and a toolbar.



The System pane occupies the upper portion of the viewer. In the System pane, you can see system information or cluster information in one of four views:

- The Table view is the most detailed, with a tabular listing of systems, monitors and parameters. Table cells for monitors are highlighted in red if the monitor value is above the critical threshold that you defined in the Settings window. A table cell is highlighted in yellow if the monitor value is above the warning threshold that you defined in the Settings window. System parameter cells are not highlighted.
- The Icons view enables you to see all systems on one panel.
- The HyperGraph view graphically displays the Table view cell values for a selected monitor or a system parameter for all systems in the report. Each system is represented by an icon on the graph.
- The Performance Analysis view displays the Performance Analysis (PA) report.

A table cell for a monitor used in Performance Analysis is highlighted in red if there is a bottleneck.

You can view systems in two modes, system mode or cluster mode. The system mode enables you to view systems individually regardless of

whether they belong to a cluster. To find out if a system belongs to a cluster, look for a cluster name in the Cluster Name column of the Table view. If the field is blank, the system does not belong to any cluster. The cluster mode enables you to view systems that are grouped together and belong to a particular cluster. In the Table view, scan the table for the Cluster Name column heading. This column indicates what cluster belongs to a system. If the field is blank, the system does not belong to any cluster. In the cluster mode, clusters are listed in the system column of the table. By clicking on the plus icon (in the field to the left of the cluster) you can expand the cluster and see the systems that are a part of the cluster. If you click on a cluster or a system within a cluster, all systems within the cluster appear on the graph. When you are in cluster mode, you can change from Table view to either Icon or Hypergraph view.

The Monitor pane, in the lower-left portion of the viewer, lists system monitors alphabetically, either in a flat list or in a tree structure. You can select a monitor from the list.

The Graph pane is in the lower-right portion of the viewer. If you select System mode, you will see either a line graph or a trend graph of the performance of your systems over the duration of the report. Line graphs and trend graphs have red horizontal lines at the critical threshold and a yellow horizontal line at the warning threshold level. The horizontal axis represents time, and the vertical axis represents the data values. If you select a cluster, its nodes are graphed as if a group of individual systems were selected together.

- The Resolution tool enables you to adjust the density of the points in the graph. It uses an average of the raw data points to present the requested number of points for a given period of time. To activate the tool, select an option in the Point Per drop-down list at the bottom left of the graph pane.
- The Trend graph button enables you to start the trend graph directly from the Report Viewer window. Use this button for instantaneous control over the trend display on the Report Viewer window.
- The Zoom tool enables you to expand a selected portion of the time line of the graph. When you activate the Zoom tool, you can zoom

out or scroll forward and backward to expand different portions of the original time span.

- The Forecast tool enables you to display wavelets as well as least-squares or linear regression calculations of future system performance. The time span of the projection is equivalent to that of the originally collected data. For example, if the report is for one month, the forecast will be for one month. A single-system forecast includes prediction intervals, which appear as dotted lines that bracket the forecast line. Prediction intervals indicate the reliability of the forecast. There is a 95% chance that the actual value will fall within the prediction intervals in the graph. The forecast line itself is a linear regression calculation. See “Viewing a Performance Forecast for a Selected System” on page 149.

If enough data for a valid forecast is not available, a warning message appears on the graph, stating that there is not enough data for a valid forecast. A valid forecast requires 24 days of data with the monitor being active at least half of the time. If the report includes only certain hours of the day, this does not affect the validity of the forecast. For example, if a report covers a month with only weekdays and only the hours from 9:00 a.m. to 5:00 p.m., and the system was on most of the time during the month, there will be enough data for a forecast. For Report Definition to create a report that can support valid forecasts, it must have a duration of a month.

When you select the forecast tool, the graph doubles the time that is displayed, with the actual data compressed to the left and the forecasted data on the right. The forecast line is a dashed version of the actual data line. Forecasting works for both line graphs and trend graphs. If there is insufficient data to make a forecast, a message will appear at the top of the graph.

Settings Notebook

The Settings notebook consists of three tabbed pages for configuring the appearance of the graph in the Graph pane, the appearance of the viewer, and threshold setting for each monitor. The configuration settings are saved in the internal storage of the IBM Director. You cannot edit this file.

For the graph, the configurable values are the number of systems to graph individually before a trend graph is triggered, the number of horizontal and vertical lines in the graph grid, whether the legend shows, and whether minimum and maximum lines show.

For monitors, the values are critical and warning threshold levels, whether data is graphed as absolute values or as a percentage of the maximum values if appropriate, and whether the graph should lower its maximum value close to the maximum data value. The Monitor page also displays information about each monitor.

Line Graph

A different-colored data line and symbol represent each system on a line graph. The symbol, such as a circle, triangle, or square, is the same color as the data line and is located at each data point on the line.

Trend Graph

A trend graph has a data line. If you select the min/max option in the Settings notebook, dotted minimum and maximum lines will bracket the data line for a graph of a single system, but not multiple-system line graphs or trend graphs. The maximum line represents the highest average values at each data point, and the minimum line represents the lowest. The Trend button can activate the trend graph as well as selecting too many systems. By default, you can select up to nine systems to graph individually; the tenth system activates the trend graph. A trend graph plots an average monitor value for all graphed systems at each time point. Marks on a vertical line that depicts the range of system monitor values at a given time point represent individual system values. Clusters of marks represent concentrations of system monitor values at a data point. By looking at the distribution of the individual system values around the average point, you can get useful information for system load balancing.

Note: When you close the generated report in the Report Viewer, Capacity Manager will not ask you if you want to save it. If you do not save the file before closing it, the file will be lost. You can save your report as a .cmr file or .txt file, or you can export it to HTML format. You also have the choice of saving the

report as a remote file on the IBM Director server, or as a local file on your workstation.

The Toolbar

The toolbar provides the following controls:

- A file menu to open, merge, and save reports, to export reports and graphs in formats that are viewable in a Web browser, to view current report information, to start a new report viewer, and to exit. The file options to open, merge, save, save as, and to export can be addressed to either a local or remote location. Local files reside on the console; remote files reside on the server.
- An Edit menu to select all systems at once or to open a Settings notebook to configure the viewer.
- A Latching toggle button that alternates the viewer window between System mode and Cluster mode.
- Four adjacent Latching toggle buttons to select one of the four possible views to display in the System pane.
- A sorting control consisting of a drop-down list of all monitors and system parameters and two latching toggle buttons for selecting either an ascending or descending sort direction.
- A Help button.

Analyzing Data

Capacity Manager probes for bottlenecks in server hardware performance, diagnoses the problem, and suggests ways to improve performance through the Performance Analysis function. The function indicates a bottleneck when one or more monitors exceed their thresholds settings for a significant amount of the report period. Performance Analysis will also look at performance trends and predict when a bottleneck might occur in the future or what latent bottlenecks will be revealed when current bottlenecks are resolved.

Performance Analysis Monitor Requirements

The algorithm uses the following monitors:

-
- Memory usage
 - % disk time
 - CPU utilization
 - % processor time
 - Packets/sec
 - % network utilization
 - Total bytes/sec

By default, Capacity Manager will activate all the required Performance Analysis (PA) monitors that are present on your systems, but only systems running Windows 2000 and Windows NT have all the required monitors; therefore Performance Analysis is not available for other operating systems. The packets/sec and % processor time monitors are not required, but omitting them might cause Performance Analysis to miss some system problems.

To help you identify the PA monitors, their names are in bold on the Monitor Activator task window, on the Monitor Selection window for the Report Definition task, and on the Monitor page in the Settings notebook in the Report Viewer. Each of the PA monitors has a critical threshold and a warning threshold; both thresholds are important to the generation of a Performance Analysis report. Any change to the threshold values for any of the PA monitors could adversely affect the results of Performance Analysis. Whenever you work with your monitor selections, remember to have all the PA monitors set to the same sampling frequency and active at the same time.

Notes:






- Not all systems have the packets/sec monitor, but if this monitor is present on your system, it will be activated with your other NT Performance monitors. There are other packets/sec monitors located under the Redirector, but these will not give the appropriate data for the Performance Analysis report.
- There is a % processor-time monitor for each processor device on your system, but the monitor for CPU utilization monitors all processor-time devices on your system. When a system has only one processor-time device, data values from the single %

processor-time monitor will be very similar to the CPU-utilization monitor data values. Any differences in these data values are attributable to the sampling times of the two monitors not being precisely in phase.

The following terminology is used in Performance Analysis Reports:

- A *device* is a system component such as memory, a processor, or a LAN adapter.
- A *constrained* or *overused* device occurs when one or more of the monitors is in the critical threshold for a significant length of the time.
- A *bottleneck* occurs on a system when one or more devices are constrained.
- A *realized bottleneck* is a bottleneck that is currently happening.
- A *latent bottleneck* is a bottleneck that might occur after you correct the realized bottleneck.

Performance Analysis

The Performance Analysis button appears as one of four icons, each of which represents a different meaning.	
	Your Performance Analysis report is ready and has no bottleneck recommendations, but the Details section of the report might discuss some bottlenecks or latent bottlenecks.
	Your Performance Analysis report is ready and will be displayed in a moment.
	Your Performance Analysis report is not available; click the Performance Analysis button to learn why.
	Your Performance Analysis report is still being prepared.
	Your Performance Analysis report is ready, and you have system bottlenecks.

To see a Performance Analysis report of the system data, click **Performance Analysis** from the menu.

To examine the report more closely, go to the top of the report. In the Recommendations section, click **Go to details** to see a detailed report of bottlenecks in that system.

Events (4) - Last 24 Hours		Event Details	
Event Type	Eye	Keywords	Values
Capacity Manager.Bottleneck.Realized	Bottleneck.C	Date	6/5/2000
Capacity Manager.Bottleneck.Forecast	Bottleneck.C	Time	11:38:53 AM
Director.Topology.Offline	System.RAN	Event Type	Capacity Manager.Bottleneck.Realized
Capacity Manager.Bottleneck.Forecast	Bottleneck.C	Event Text	Bottleneck: Memory + CPU.
		System Name	RANDYTPAD
		Severity	Critical
		Category	Alert
		Group Name	RANDYTPAD
		Sender Name	RANDYTPAD
		CMR file	C:\Tivoli\gtreports\Monthly Report - 2000
		Has latent bottleneck(s)	false
		Involves CPUs	true
		Involves Disks	false
		Involves LAN Adapters	false
		Involves Memory	true
		When bottleneck started	Wed 31 May 2000 10:55 EDT
		When bottleneck stopped	Fri 02 Jun 2000 16:55 EDT
		Minutes since bottleneck started	7243
		Minutes since bottleneck stopped	4003
		Hours since bottleneck started	120
		Hours since bottleneck stopped	66
		Days since bottleneck started	5
		Days since bottleneck stopped	2
		Hours in report	189
		Hours analyzed	12
		Hours in this bottleneck	8
		Number of bottlenecks	3
		Hours in all bottlenecks	12

Performance Analysis Reports

The Performance Analysis report consists of two main sections: Recommendations and Details. The Details section shows everything that was found, and the Recommendations section shows only the subset of details on which the user needs to act. The Details section includes links that enable you to see a graph of the performance of the monitor in question.

The systems with the most severe bottlenecks are first on the report list. A bottleneck that is reported in the Details section will appear in the Recommendations section if it meets one of the following criteria:

- It occurred on the last day of the report.
- It occurred more than 25% of the time, and it occurred more than any other bottleneck for that particular system.
- It probably will occur in the future. Performance Analysis must have enough data to make a reliable forecast.

Saving and Printing a Performance Analysis

To save a report summary as a local HTML file, click **File**→**Export Report to local HTML**. In the Save as local HTML window, select a preferred directory, type a new file name, and click **Save**.

A report saved as an HTML file will contain the following sections:

- A Table of Contents that contains links to the other sections: Report Table, the Report Information, and the Performance Analysis Recommendations and Details.
- A Report Table that presents the same monitor and the system data that is also available in the Report Viewer in the Table view.
- The Report Information, which includes the file name, the analysis start and end dates, days of the week and hours of coverage, the name of the report definition, and a list of any systems that were requested but not included in the report.
- The Performance Analysis Recommendations, which include recommendations for remedying the most serious bottlenecks.
- The Performance Analysis Details, which include information on the frequency and duration of both active and latent bottlenecks, and remedies for the bottlenecks. When you correct a current bottleneck, it becomes a latent bottleneck.

To print a Performance Analysis report, save it first as either a local HTML file or as a remote (on your server) HTML file, and then print it from your Web browser. A printed version of the report omits the links to the graphs but includes the monitor and system parameter information from the Table view.

Bottlenecks

Bottlenecks occur on a system when one or more devices become constrained. The monitors for the devices will detect the constraint, and Performance Analysis indicates a bottleneck.

There are four types of single bottlenecks that are detected by Performance Analysis: CPU, disk, memory, and LAN adapter. Each of the PA monitors detects one of these four bottleneck types:

-
- The CPU utilization and % processor-time monitors detect a CPU bottleneck.
 - The % disk-time monitor detects a disk bottleneck.
 - The memory-usage monitor detects a memory bottleneck.
 - The packets/sec and the % network utilization monitors detect a LAN adapter bottleneck.

These are not the only types of bottlenecks that can occur. There can be bottleneck combinations. For example, a disk bottleneck occurs if either the disk 1 or disk 2 device is constrained, and a memory bottleneck occurs if a memory-usage device is constrained. But if disk 1 becomes constrained along with memory usage, the memory bottleneck becomes a disk + memory bottleneck, and the recommendation for addressing a bottleneck with two devices might not be the same as addressing two bottlenecks individually.

In the example of a disk + memory bottleneck, the algorithm recognizes that insufficient memory can lead to disk thrashing, so the recommendation is to add memory and leave the disk drive unchanged. Devices often interact in this way, so each combination of device types (memory, disk, CPU, LAN) constitutes a separate bottleneck with its own recommendation.

Often, when one bottleneck occurs, others are not evident while the first bottleneck slows the system. The bottleneck that is occurring is a realized bottleneck. A latent bottleneck is a bottleneck that is not evident while the system slows down. The Performance Analysis will report a device as a latent bottleneck if it exceeds the warning threshold setting at least 50% of the time that another device is constrained. It is also possible for a device to contribute to a latent bottleneck for part of the time and to a realized bottleneck for part of the time.

The Performance Analysis algorithm scans for bottlenecks on each system. If no bottlenecks are found for a system, Performance Analysis will use a forecasting algorithm to look simultaneously at all the system monitors to predict what bottlenecks might occur and report the first bottleneck that it foresees. The forecast period is the same length as the report period. For example, a report period of one month can generate a forecast one month into the future.

Bottleneck Event Extended Attributes

Generated bottleneck events have extended attributes. These attributes appear in the Event Log. They also appear in the Simple Event Filter Builder, where you can use them in building more elaborate filters. The following extended attributes appear in the Event Log:

- CMR file, TXT file, HTML file - The names of the files that were saved when this event was generated. To understand an event, load the report file and read the Recommendations section of Performance Analysis. Remember that the event reflects the first recommendation only.
- Involves Memory, Involves Disk, Involves LAN Adapters, Involves CPUs - Each of these is true if the bottleneck involves the given type of device. For example, Involves Memory would be true for a memory bottleneck or for a memory + disk bottleneck.
- Cluster Node - This is true if the bottlenecked system is a node in a cluster.
- When bottleneck first started, When bottleneck last stopped - The time stamps for the beginning and end of the bottleneck, as reported in Performance Analysis. The bottleneck might stop and start again any number of times between these two time stamps.
- Minutes since bottleneck first started, Minutes since bottleneck last stopped, Hours since bottleneck first started, Hours since bottleneck last stopped, Days since bottleneck first started, Days since bottleneck last stopped - The number of minutes, hours, days since the bottleneck started or stopped. Each of these numbers is independent of the others. For example, if the number of days since the bottleneck started is 2, the number of hours would be greater than 48, and the number of minutes would be greater than 2880.
- Hours in this bottleneck - The number of hours that the bottleneck was active, as reported by Performance Analysis. If the bottleneck was active only part of the time between the start and end times, this number excludes the times it was not active. Therefore, this number might be less than the end time minus the start time

Using Bottleneck Events

The Capacity Manager program identifies bottlenecks whenever you generate a report. The program uses the Performance Analysis algorithm to determine when and where bottlenecks occur. For a demonstration of Performance Analysis and other Report Viewer, see the Report Viewer Tour.

To receive automatic notification when a bottleneck occurs, set an event to specify how you are to be notified. There is a four-step process for configuring Capacity Manager to use bottleneck events:

1. Generate events when a bottleneck occurs.
2. Create an event filter that filters out all events except bottlenecks.
3. Define an action and an action plan that automatically occurs when the plan detects a bottleneck.
4. View the bottleneck data.

You can use the bottleneck data to respond to situations that slow your network performance and to try to avoid future bottlenecks.

If you follow these steps, Capacity Manager notifies you every hour when a new bottleneck starts on any system. This is the most effective configuration for detecting bottleneck events. If you want to set up more elaborate configurations, see the IBM Director documentation for bottleneck events.

Generating an Event

You can generate an event whenever Performance Analysis recommends correcting a bottleneck. If you configure the program to generate an event for a bottleneck, you should also create an action plan to respond to the event.

To configure the program to check for bottlenecks every hour on selected systems and to generate a report if a bottleneck occurs, do the following:

1. Drag **Hourly Bottleneck Events** (under the Report Generator task) onto a group of systems or onto one or more selected systems. To generate a report on more than one system, use **[ctrl]+click** to select the systems, and then drag the report viewer onto one of the

selected systems

Note: Clusters do not generate bottleneck events; systems generate bottlenecks. By definition, a cluster is a group of systems that are unable to collectively generate events. Therefore, when you generate a bottleneck event report, specify systems, not clusters.

2. Click **Schedule**. The New Scheduled Job window opens.
3. In the New Scheduled Job group box, type a job name, and select a date and a time for the job to run. The default date is today's date, and the default time is 12:00 p.m. Select a date and time in the future; otherwise the schedule job will not run. The date and time that you select in the New Scheduled Job group box indicates the first time that the job runs.
4. Click **Advanced** to schedule the job to run at regular intervals. A four-tab pane opens where you can make your selections.
5. On the Date/Time page, click **Repeat** to schedule your report generation as a repeating event. In the Repeat pane, change Once to Hourly, and click **OK**.
6. Click **File**→**Save as** to save your selections. For example, a job name that you can use is *Generate bottleneck events*. Click **Help** to access IBM Director help for scheduling a new job.
7. Close the New Scheduled Job window.

Every hour, Performance Analysis checks the specified systems for bottlenecks. Each time Performance Analysis recommends correcting a bottleneck on any of the systems, two things happen:

- Performance Analysis generates a report. The report is stored on the IBM Director server in the \reports subdirectory. You can view it in the Report Viewer if it is either a .cmr file or a .txt file. If it is an HTML file, you can use your Web browser to view it.
- Systems with bottlenecks generate events. These events are associated with the first recommendation that is made in the Recommendations section of Performance Analysis. The IBM Director Event Log registers these events.

Creating an Event Filter

You can filter out all events except bottleneck events. The goal is to respond to new bottleneck events when they initially occur. The program should not respond to the same bottleneck every hour. To filter out all events except bottleneck events, use the following procedure:

1. From the IBM Director Console, click **Tasks** → **Event Action Plan Builder**.
2. Click **File** → **New** → **Simple Event Filter**.
3. Click the **Event Type** tab. In the left pane, clear the **Any** check box. In the right pane is a tree structure that lists applications. Under the Capacity Manager task, open the Bottleneck tree, and click **Recommendation**.
4. Click the **Extended Attributes** tab; clear the **Any** check box.
5. In the **Keywords** drop-down list click **Hours since the bottleneck first started**. In the **Operator** drop-down list, select **Equal To**. In the Values text box, type **2**.
6. Click **Add**.
7. Click **File Save As** to save your selections. For example, a filter name that you can use is *Bottleneck Events*.
8. Close the Event Action Plan Builder window.

The next step is to set up an action/action plan with this event filter. When this event filter isolates a new bottleneck, the action will run and notify you that a bottleneck occurred.

If a bottleneck repeatedly occurs every hour, this event filter will perform the action the first time only. However, the bottleneck event will still appear in the event log every hour, and the report file will still be saved every hour.

This filter uses one extended attribute: Hours since the bottleneck first started. If you want to create more event filters that use other extended attributes see Bottleneck Event Extended Attributes on page 144.

Defining an Action and Action Plan

To define an action that automatically occurs when a bottleneck is detected and to correlate the action to the bottleneck event filter see “Creating an event filter that filters out all events except bottlenecks” on page 147, use the following procedure:

1. From the IBM Director Console, click **Tasks**→ **Event Action Plan Builder**.
2. In the Actions pane (on the right side of the window), double-click any action to customize that action. For example, you can double-click **Add a Message to the Console Ticker Tape**.
3. Complete all selections. For the ticker-tape example, type a message and the user.
4. Click **File**→**Save As** to save your selections. For example, an action name that you might use is **Bottleneck Action**.
5. Click **File**→ **New** →**Event Action Plan**.
6. Type the name for the event action plan and click **OK**. For example, a name that you might use is *Respond to Bottleneck Events*.
7. Drag the event filter for bottleneck events from the center pane onto the event action plan that you named in the previous step.
8. Drag the **Bottleneck Action** from the Actions pane (on the right side of the window) onto the event filter that is in the Event Action Plan pane (on the left side of the window).
9. Click **File** →**Close** to exit from the Event Action Plan Builder.

Forecasting Data

The Forecast function enables you to see a prediction of the future performance of your selected systems.

Capacity Manager uses forecasting in the following components:

- In a Performance Analysis report. If there are no realized bottlenecks, Capacity Manager uses forecasting to predict, with a

level of confidence, if and when it foresees a monitor performance bottleneck.

- In your system monitor performance graph. On a graph of a selected monitor for one or more systems, you can click the **Forecast** button to see a forecast of the performance on your selected systems. The graph depicts both the observed data and the forecast.

About the Calculations

To create a forecast, Capacity Manager applies a wavelet transform to the monitor data prior to performing a least squares linear regression. With this transformed data, it computes a forecast line with a 95% prediction interval. The forecast duration is equal to the duration of the observed data. For the forecast to be valid, Capacity Manager needs a minimum of 24 days of previously collected data where the system monitors have been running at least 50% of the time.

Viewing a Performance Forecast for a Selected System

To see the performance forecast for your selected systems, click **Forecast** in the lower-right corner of the screen. The Capacity Manager forecast is for whatever monitor you currently have selected. To see a forecast for another Monitors pane, click its name in the monitor box.

Notes:

- You cannot use the Zoom tool and the Forecast tool at the same time; they are mutually exclusive.
- The forecast data is more meaningful for systems that are individually graphed rather than shown in a trend graph. To change from a trend graph to a graph of individual systems, either set your trend graph threshold to a higher number, or select fewer systems to graph at one time. For more information, see *Changing the number of systems graphed as individual systems* in the online help.

About the Forecast Display

The Forecast graph has these features:

- The Forecast line is a dashed line with an arrow at the end. This line describes possible future data values that are consistent with the prediction that an actual future data value will fall with equal probability above or below the forecast line. The forecast interval is a multiple of your data-collection period. The default prediction period is the same length as the data-collection period. For example, if you have a month of collected data, the forecast will be for a month into the future.
- The Prediction interval is represented by the dotted lines above and below the forecast line. The prediction interval represents the range of data values that are located above and below the forecast line and are consistent with the prediction that an actual future data value will fall within the interval with a probability of 95%. The width of the interval depends on the variability of the observed monitor data: the greater the variability, the wider the prediction interval. The prediction interval is displayed when you request a forecast of a single system. Graphs of multiple-system forecasts do not show prediction intervals.
- The vertical bar at the beginning of the forecast data depicts the range.
- The gap between the actual collected data and the beginning of the predicted data serves as a separator between these two data sets.
- If you do not know how to interpret a wide prediction interval for a forecast, select the **Point per** check box to call up a finer resolution of your data. Your data points might have a broad variance that is hidden by the averaging that occurs when data is displayed at a coarser resolution

About Warning Messages on the Forecast Graph

Capacity Manager will display one of two warnings if your forecast is not valid:

-
- The data collection period is too short for a valid forecast. To generate a valid forecast, you need at least 24 days of data.
 - System X does not have enough data for forecasting, or multiple systems do not have enough data for forecasting. The selected monitor must have been on at least 50% of the time during the data-collection period.

Note: Do not use these warnings to make decisions about your systems.

6

Fibre Channel Storage Manager

Fibre Channel Storage Manager is installed as part of UM Server Extensions and is accessed as a task through the IBM Director Console. Storage management commands are sent from the console over your network to the storage subsystem. The connection between the console and the storage subsystem controller is called the *network management connection*.

A storage subsystem is a storage entity that consists of a collection of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as arrays and logical drives). A storage subsystem can span multiple physical enclosures. The Fibre Channel Storage Manager task provides management capabilities only if the storage subsystem has been correctly installed and the appropriate network configuration tasks have been performed. Refer to the *IBM FASiT500 RAID Controller Enclosure Unit Installation Guide* and the *IBM FASiT Storage Manager for Windows NT Installation and Support Guide* for more details.

Fibre Channel Storage Manager program includes management features, such as:

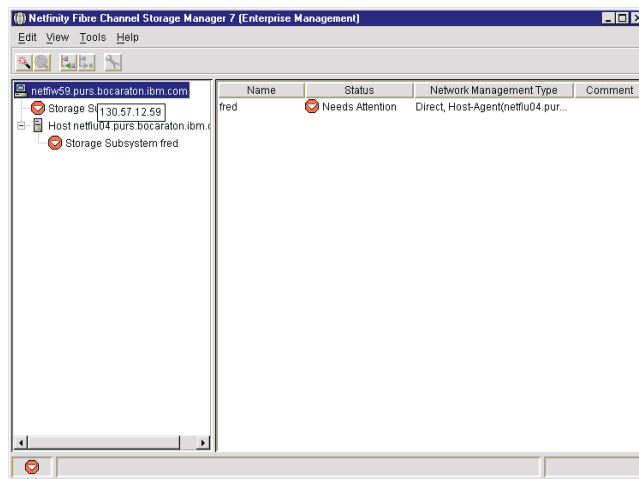
- Configuring available storage subsystem capacity into logical drives to maximize data availability, optimize application performance, and utilize storage resources
- Granting access to partitions to the application hosts in the enterprise

- Setting up a management domain, which is a collection of storage subsystems to manage
- Monitoring storage subsystems in the management domain for problems or conditions that require attention
- Configuring destinations to receive alert messages for critical problems concerning one or more storage subsystems in the management domain
- Recovering from storage subsystem problems to maximize data availability
- Tuning for optimal application performance

The Fibre Channel Storage Manager program has two main windows: the **Enterprise Management** window and the **Subsystem Management** window.

Enterprise Management Window

The Enterprise Management window is the first window that opens when you start the storage management task. After you add storage subsystems, the Enterprise Management window communicates with each storage subsystem (and controllers) for remote discovery and coarse-level management.



The **Enterprise Management window** provides the following primary storage-management functions:

- Discovers hosts and storage subsystems on your local subnetwork to add to your management domain
- Provides the ability to manually add hosts and storage subsystems to your management domain and also remove them
- Enables you to add storage subsystem comments to the device table
- Provides the ability to sort rows in the device table according to different criteria
- Stores your Enterprise Management window view preferences and management domain data in local configuration files. The next time you open the Enterprise Management window from the management station, data from your configuration files is used to display your custom management domain and other view preferences
- Monitors the condition of storage subsystems and reports a high-level status using appropriate icons
- Loads the appropriate Subsystem Management window for a selected storage subsystem to enable detailed configuration and management operations

Along with a menu bar and a tool bar, the Enterprise Management window consists of two viewing panes, the device tree and the device table. The Device Tree in the left pane of the window displays a picture of the management domain. The Device Tree shows discovered devices, added devices, and device status conditions. The Device Table in the right pane displays more detailed information for each storage subsystem.

Note: If you right-click a node in the device tree, a pop-up menu containing the appropriate menu options for that node is displayed.

There are two ways to populate a management domain:

- Click the **Automatic Discovery** icon to automatically discover directly managed and host-agent managed storage subsystems on the local subnetwork and add them to the management domain. Fibre Channel Storage discovers host-agent managed storage subsystems by first discovering the hosts that provide network management connections to the storage subsystems. Then the host and associated storage subsystems are displayed in the device tree.
- Click the **Add Device** icon. For a directly managed storage subsystem, type a host name or IP address for each controller in the storage subsystem. For a host-agent managed storage subsystem, type a name or IP address for the host that is attached to the storage subsystem.

Subsystem Management Window

The Subsystem Management window provides the options that you need to manage an individual storage subsystem in the management domain. The Subsystem Management window is accessed from the Enterprise Management window and provides the following primary storage subsystem management functions:

- Provides storage subsystem options such as renaming a storage subsystem, changing a password, or enabling a background media scan
- Provides the ability to configure logical drives from your storage subsystem capacity, define hosts and host groups, and grant host or host group access to sets of logical drives, called *storage partitions*
- Monitors the health of storage subsystem components and reports a detailed status using appropriate icons
- Provides you with the appropriate recovery procedures for a failed logical or physical component
- Presents a view of the storage subsystem Event Log
- Presents profile information about physical components such as controllers and drives

- Provides controller management options such as changing ownership of logical drives, setting a controller in active or passive mode, or placing a controller online or offline
- Provides drive management options such as assignment of hot spares, and drive initialization
- Provides monitoring of the storage subsystem performance
- Provides recovery from failures through the Recovery Guru

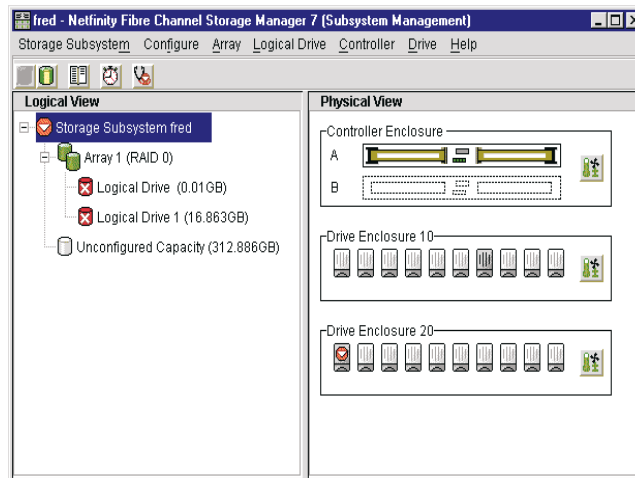
Using the Subsystem Management

Use the Subsystem Management window to manage (configure, monitor, recover, and tune) the logical and physical components that make up a storage subsystem. The Subsystem Management window is specific to an individual storage subsystem.

To open a Subsystem Management window for a selected storage subsystem, use the following procedure:

1. In the Enterprise Management window, select a storage subsystem.
2. Select **Tools** → **Manage Device**.

The Subsystem Management window opens.



Note: When you open Subsystem Management window, you can manage only the selected storage subsystem. However, you can open multiple Subsystem Management windows to manage other storage subsystems.

The Storage Subsystem window consists of two views, a Logical view and a Physical view. The Logical view in the left pane displays a picture of the storage subsystem. The storage subsystem is organized into logical entities called arrays and logical drives. The Physical view in the right pane displays a picture of the physical devices in the storage subsystem, such as controllers, drives, and other components. When you select a logical drive or other entity in the Logical view, you will see the associated components in the Physical view.


When you open the Subsystem Management window for the first time, the storage subsystems are unnamed. You must rename each storage subsystem from *<unnamed>* to the name that you want.

Renaming Storage Subsystems

When renaming a storage subsystem remember that Fibre Channel Storage Manager does not check for duplicate names. Check the Enterprise Management window to verify that the name that you have chosen is not already in use by another storage subsystem. There is a 30-character limit. All leading and trailing spaces will be deleted from the name. Use a unique, meaningful name that is easy to understand and remember. Avoid arbitrary names or names that would quickly lose their meaning in the future. When you have renamed the storage subsystem, the prefix *storage subsystem* is automatically added to the name shown in the Logical view and in the Enterprise Management window. For example, if you renamed the storage subsystem *Engineering*, it will display *storage subsystem Engineering*. You do not need to add the prefix *storage subsystem* to the name.

To rename a storage subsystem, use the following procedure:

1. From the Subsystem Management window, select **Storage Subsystem** → **Rename**.
2. The Rename Storage Subsystem window opens. The window displays the current name of the storage subsystem. In some cases, this is the default name that is assigned by Fibre Channel Storage.



3. Type a new name.
4. Click **OK**.
The storage subsystem is renamed. The name that is shown in the Logical view in the Enterprise Management window will have the prefix *storage subsystem* followed by the name.
5. Prepare a label with the new storage subsystem name, and place it on the storage subsystem.

Configuring Storage Subsystems

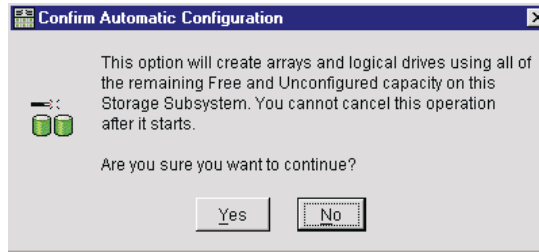
You can automatically configure your storage subsystems to use all of the remaining unconfigured storage subsystem capacity and any free capacity in existing arrays. The result will be a controller-defined set of logical drives, arrays, and hot-spare drives. The resulting configuration is well suited to the needs of many applications.

Do not use this option if:

- You want to optimize the configuration for a certain application
- You want logical drives with different RAID levels, cache settings, or other parameters
- You have specific hot-spare requirements or you want to keep free space available for future use

To automatically configure your storage subsystem, use the following procedure:

1. From the Subsystem Management window, click **Configure**→**Automatic Configuration**.
The Confirm Automatic Configuration window opens.



2. Click **Yes**.
A controller-defined configuration is applied to the storage subsystem.

Note:

- From the initial set of unassigned drives, one is configured as a hot spare.
- From the remaining set of unassigned drives, the maximum number of 5+1 RAID logical drives is configured.
- The Logical view and the Physical view are updated to reflect the new configuration.
- If one to five drives remain, they are configured as follows:

Number of remaining drives	Resulting configuration
From 3 to 5	An N + 1 RAID 5 logical drive is created.
2	A 1 + 1 RAID 1 logical drive is created.
1	A hot spare is created.

Partially Managed Storage Subsystem Condition

A partially managed storage subsystem condition occurs when only one controller was defined or could be reached when the storage subsystem

was added or discovered. You are provided several indications of this condition:

- When you close the **Add Device** window.
- When you attempt to manage a storage subsystem (that is, when you open its Subsystem Management window).
- When you select the **View Partially Managed Devices** drop-down list.
- When you activate the tooltip on the storage subsystem by placing your cursor over the storage subsystem in the device tree.

If a storage subsystem is partially managed, it will display "partially managed" in the tooltip. (This tooltip is available only for directly managed storage subsystems and is unavailable for host-agent managed storage subsystems.)

This condition has different causes, depending on how the storage subsystem is managed over the network:

- If the storage subsystem is managed through each Ethernet connection on the controllers (Directly Managed), either you have not provided a host name or IP address of one of the controllers, or there was a connection problem.
- If a storage subsystem is managed through a host network connection (host-agent managed), there was a problem detecting one of the controllers in the attached storage subsystem.

To correct a partially managed condition for a directly managed storage subsystem, use the following procedure:

1. From the Enterprise Management window, click **View** → **Partially Managed Devices**. The Partially Managed Devices window opens.
2. Make a note of the controllers that are listed (both known and unknown). Use this information to identify the unknown controller. The controller pairs are listed. The known controller is listed first, followed by the unknown controller. The name or IP address of the known controller, its associated storage subsystem,

and a description of the known controller location, such as the controller in slot A or slot B, are displayed.

3. From the Partially Managed Devices window, click **Add More**. The Partially Managed Devices window closes. The Add Device window opens.
4. From the Add Device window, specify the name or IP Address of the controller that was listed in the Partially Managed Devices window with an *unknown* designation, then select **Add**. One of the following connecting messages is displayed while Fibre Channel Storage Manager attempts to contact the specified name or IP address:

If	Then
The Add Device window opens again	You have resolved the partially managed condition for the particular controller, go to step 6.
A Connection Problem dialog displays as follows: The specified device was not accessible. Do you want to add it as an Unidentified Device?	Click No . There is either a problem with the physical connection or a faulty controller. Go to step 5.

5. Check to make sure of the following conditions:
 - The Ethernet cable into the controller is not visibly damaged and is securely connected.
 - The appropriate network configuration tasks have been performed (for example, IP address assigned to the controller, and so on). Refer to the *Software Installation Guide* for details.
 - The controller is network-accessible. One way to do this is to use the **ping** command to verify that the controller can be reached.

If	Then
There was a cable or network accessibility problem	Correct the problem and then repeat step 4.
There was not a cable or network accessibility problem	Turn the power to the controller enclosure off and on. Repeat step 4. If you continue to receive the Connection Problem dialog, you might have a faulty controller.

- Repeat step 4 for each controller that is associated with a directly managed storage subsystem that was listed with an **Unknown** designation.
- When you have added all names or IP addresses, click **Close** from the **Add Device** window. One of the following conditions will appear:

If	Then
The Partially Managed Device dialog is displayed again	Either you have not resolved all of the directly managed storage subsystems or the remaining controllers listed are associated with host-agent managed storage subsystems.
The Partially Managed Devices dialog does not display again	You have resolved a partially managed situation.

Correcting Partially Managed Condition

To correct a partially managed condition for a host-agent managed storage subsystem, use the following procedure:

- Select **View** → **Partially Managed Devices**.
The Partially Managed Devices window opens.
- Make a note of the controllers that are listed (both known and unknown). Use this information to identify the unknown

controller.

The display shows controller pairs with the known controller listed first, followed by the unknown controller. The display shows the name or IP address of the host, the associated storage subsystem, and a description of the known controller location such as the controller in slot A or the controller in slot B.

3. From the Partially Managed Devices window, click **Exit**. The Partially Managed Devices window is closed.
4. Verify the connections and rescan the controller using the following procedure:
 - a. Check the external cable between the host adapter and the controller to make sure that it is securely connected.
 - b. Restart the host agent software on the host through which you are having problems detecting the controller. Refer to the *Software Installation Guide* for information about restarting the host agent.
 - c. Highlight the Host in the device tree of the Enterprise Management window and then click **Tools** → **Rescan** to attempt to detect the controller in the storage subsystem.
 - d. To determine if the problem has been corrected, click **View** → **Partially Managed Devices** in the Enterprise Management window.

Note: There are two likely reasons that the host-agent software could not detect the controller. There can be a connection problem or the controller was installed or replaced since the last time the host-agent software was started.

5. Repeat step 4 for each controller that is associated with a host-agent managed storage subsystem that was listed with an *unknown* designation.

Adding a Storage Subsystem

Use this option to manually add a directly managed or host-agent managed storage subsystem to the management domain. It can take

several minutes for the Fibre Channel Storage Manager to connect to the specified device. Before using this option, verify that the appropriate network configuration tasks have been performed. Refer to the *Software Installation Guide* for details. To add a directly managed storage subsystem, make sure that you provide the host name or IP address of each controller in the storage subsystem. Note that the window allows you to add only one address at a time. To add a host-agent managed storage subsystem, add the host through which the storage subsystem is attached to the network. The host will be displayed in the device tree along with any attached storage subsystems.

To add a storage subsystem, use the following procedure:

1. From the Enterprise Management window, open the **Add Device** window in one of the following ways:
 - Click the **Add Host/Device** icon.
 - Click **Edit** → **Add Device** from the menu.
 - Right-click the node in the Enterprise Management window, and click **Add Device**.

The Add Device window opens.

2. If the storage subsystem is a directly managed storage subsystem, type a host name or IP address for one of the controllers in the storage subsystem. If the storage subsystem is a host-agent managed storage subsystem, type a name or IP address for the host through which the storage subsystem is attached to the network.
3. Click **Add**.
The Connecting message window opens while Fibre Channel Storage Manager attempts to contact the specified device.
4. The following table list device connection status.

Status	Description
The device was successfully accessed.	The device is added to the device Tree and the device table. Also, the Connecting message window closes and the Add Device window is displayed again. You can either add another device or click Close .
The device was <i>not</i> successfully accessed, because of an unresolved host name.	The Connection Error message window is displayed. Click OK to return to the Add Device window. You can either type a valid host name or select Close .
The device was not successfully accessed, because of an unsuccessful connection to the name or IP address of a controller in the Storage Subsystem or an unsuccessful connection to a Host.	<p>The Connection Problem message window is displayed.</p> <ul style="list-style-type: none"> ■ Click Yes to add the device to the device tree as an unidentified Node and close the window. ■ Click No to close the window without adding the device. <p>After you click yes or no, the Add Device window opens. You can either add another device or select Close.</p>

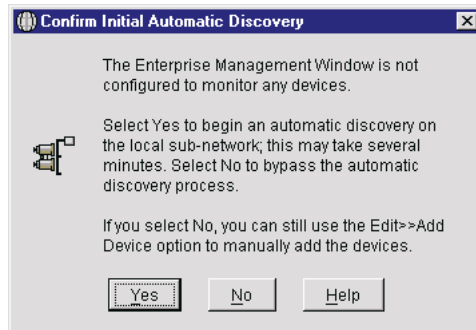
Note: If a Partially Managed Devices window opens after you have clicked **Close** in the Add Device window, one or more of the devices that you have added is only partially manageable (that is, only one controller in a storage subsystem has been defined). This window provides you with information to assist you in determining what devices still need to be defined.

Using Automatic Discovery

Use the **Automatic Discovery** icon to automatically discover directly managed and host-agent managed storage subsystems on the local subnetwork and add them to the management domain.

Fibre Channel Storage Manager discovers host-agent managed storage subsystems by discovering the hosts that provide network management connections to the storage subsystems. Then, the host is displayed in the device tree along with its associated storage subsystems.

Whenever you open the Enterprise Management window from a particular management station and the configuration files containing management domain data are not present or are empty, an Automatic Discovery window opens. This window gives you an opportunity to automatically add device (hosts and storage subsystems) from your local subnetwork to your management domain.



Note: This window also opens if the primary configuration file is unusable and the backup configuration file is empty. It is also displayed as part of the recovery sequence when the configuration files have been damaged and have been reset as empty files. Your selection will have one of the following results:

Selection	Result
Yes	An automatic discovery process begins. For more information, see "Discovery Storage Subsystems" in the online help.
No	To add devices to your management domain, use the Add Device option from the Enterprise Management window. For more information, see "Adding a Storage Subsystem" in the online help.

The automatic discovery process sends out a broadcast message across the local subnetwork and any device responding to the message is included in the device tree view.

To start an automatic discovery, use the following procedure:

1. Do one of the following to display the **Confirm Automatic Discovery** window:
 - Click the **Automatic Discovery** icon.
 - Select **Tools**→ **Automatic Discovery**.
 - Right-click a node from the Enterprise Management window, and click **Automatic Discovery**.
2. Click **OK**.

Fibre Channel Storage Manager begins the automatic discovery process. The Progress dialog displays the message *Performing automatic discovery of devices*, and the progress bar begins to display relative progress, shown as a percentage of the estimated time it will take to complete. During the Automatic Discovery process, the software updates the Progress dialog for each new host and storage subsystem that is discovered. At the end of the automatic discovery process, the progress bar is cleared to show no activity, and the Progress window displays the following messages:

- Automatic discovery completed. N new devices.
(where N is the number of devices found)

-
- Automatic discovery completed. No new devices were found.

Any discovered devices are displayed in the device tree and in the device table.

Storage Subsystem Configuration

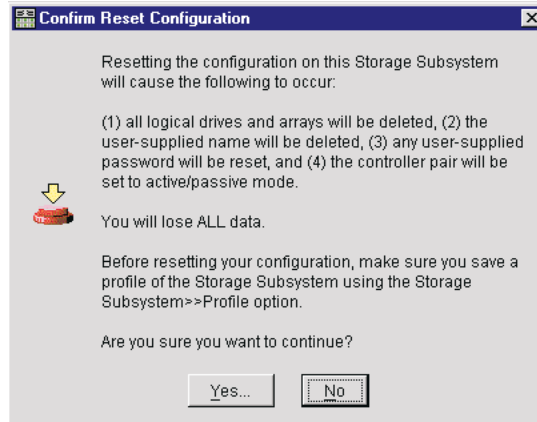
If you reset a storage subsystem configuration from a factory logical drive configuration before you manually or automatically configure the storage subsystem, you can reclaim disk space. Resetting a storage subsystem configuration will also correct major configuration errors that cannot be corrected using any other method.

All data currently stored on the storage subsystem will be lost. Do not attempt this operation unless all data on the storage subsystem has been backed up to tape or to another storage subsystem. Resetting the configuration on the storage subsystem will also cause the following to occur:

- All logical drives and arrays will be deleted.
- The user-supplied name will be deleted.
- If you have password-protected your storage subsystem for destructive operations, that password protection will be lost.
- The controller pair will be set to Active/Passive mode.

To reset a storage subsystem configuration, use the following procedure:

1. From the Subsystem Management window, click **Configure**→ **Reset Configuration**.
The Confirm Reset Configuration window opens.



2. Select **Yes**.
A second **Confirm Reset Configuration** is displayed.
3. If you are certain that you want to delete your entire storage subsystem configuration, click **Yes**.
The current configuration is deleted, destroying any existing data on the storage subsystem. A controller-defined configuration is applied (which may include a default logical drive). All other drives are unassigned. The Logical view and the Physical view are updated to reflect the new configuration.

Recovering from Storage Subsystem Problems

Use the Recovery Guru to help determine the cause of a problem with a storage subsystem and to obtain instructions on how to correct the problem. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components. See "Recovery Guru" on page 172.

Problem Notification

Indicators of a storage subsystem problem include:

- A Needs Attention storage subsystem icon displayed in:
 - The Overall Health Status field, device tree, or device table of the Enterprise Management window.

- The Subsystem Management window Logical view.

- The Recovery Guru Optimal icon in the Subsystem Management window changes to a flashing Recovery Guru Needs Attention toolbar button.
- Non-optimal component icons are displayed in the Subsystem Management window, in the Logical view and the Physical view.
- Critical SNMP or e-mail error messages are received.
- Hardware fault lights are received.

Recovering from Storage Subsystem Problems

To recover from storage subsystem problems, use the following procedure:

1. From the Subsystem Management window, do one of the following:
 - Click the **Recovery Guru** icon.
 - Click **Storage Subsystem** → **Recovery Guru** menu.
 - Right-click the storage subsystem in the Logical view and click **Recovery Guru** from the menu.
2. Select the first problem shown in the Summary view, then follow the recovery steps to correct it. Repeat the procedure for each listed problem. When all problems have been corrected, the storage subsystem icon will change from *Needs Attention* to *Optimal*. There are certain problems where a *Fixing* icon will be displayed while an operation such as reconstruction is in progress.
3. Click **Recheck** to verify the success of the completed recovery procedure.

Recovery Guru

The Recovery Guru window is divided into three views: Summary, Details and Recovery Steps.

The Summary view displays a list of storage subsystem problems. Multiple problems within a storage subsystem can be related, and the order in which the problems are corrected can affect the outcome. Select and correct the problems in the order in which they are listed in the Summary group box.

Note: Multiple power-supply canister failures are grouped and listed as one problem in the Summary group box. Multiple fan canister failures are also listed as one problem.

The Details view displays information about the selected problem in the Summary group box.

Note: You can use standard keyboard cut, copy, and paste operations on text in this area

The problem types are:

Problem type	Information Shown
All	Storage subsystem name
Controller	<ul style="list-style-type: none"> ■ Slot identification (where applicable, the appropriate controller slot (A or B) is listed). ■ Channel ID, for failed drive SCSI channels. ■ Serial number of controller.
Logical drive	<ul style="list-style-type: none"> ■ Logical drive group sequence number. ■ Drive enclosure identification. ■ Affected drive slots (if applicable). ■ Logical drive names. ■ RAID level. ■ Status.
Drive	<ul style="list-style-type: none"> ■ Drive enclosure identification. ■ Affected drive slots, listing slot numbers.
Hardware components	<ul style="list-style-type: none"> ■ Component type. ■ Drive enclosure identification

The Recovery Steps view lists the appropriate steps to follow for the selected problem in the Summary group box.

There are four headings or statements that might appear in a recovery procedure:

Heading	Description
What Caused the Problem	Summarizes the likely causes of the problems
Caution	Lists instructions or information that can prevent data loss, service interruption, or component problems
Important Notes	Lists important information about recovery process that should be read before proceeding with the recovery steps
Recovery Steps	Lists appropriate steps for a selected problem

Saving Recovery Information to a Text File

To save recovery information to a text file, use the following procedure:

1. Click **Save As**.
The **Save** window opens.
2. Select the subdirectory location in which to save the file.
3. Type the name of the new file in the appropriate text box, and click **OK**.
The **Save** window is closed, and the file is saved in HTML format. The file contains all of the information in the Details and Recovery Procedure group boxes.

Specifying Storage Subsystem Cache Settings

You can specify when unwritten cache data should be written to disk (flushed), and when a cache flush should stop, and you can specify cache block size.

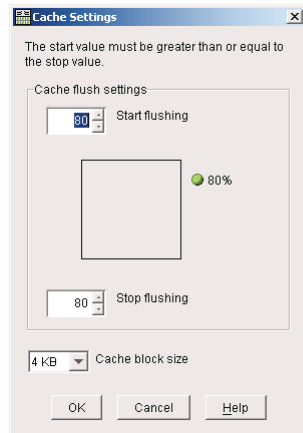
When the cache holds the specified start percentage of unwritten data, a flush is triggered. When the cache is flushed down to the specified stop percentage, the flush is stopped. For example, you can specify that the

controller start flushing the cache when the cache reaches 80% full or stop flushing the cache when the cache reaches 16% full. Both the Start flushing and Stop flushing parameters are percentages of the total cache memory.

To change cache settings, use the following procedure:

1. From the Subsystem Management window, click **Storage Subsystem** → **Change Cache Settings**.

The Change Cache Settings window opens.



2. In the **Start Flushing** drop-down list, select the percentage of unwritten data in the cache that will trigger a cache flush. A percentage that is too low increases the chance that data that is needed for a host read will not be in the cache. It also increases the number of disk writes necessary to maintain the cache level, increasing system overhead and decreasing performance.
3. In the **Stop Flushing** drop-down list, select the percentage of unwritten data that is held in the cache that will stop a cache flush in progress. The lower the percentage that you specify, the greater the chance that data for a host read would require a disk read than being present in the cache.
4. Select a cache block size that is appropriate for your applications from the Cache block size drop-down list. Available choices include 4 KB (a good choice for file system or database

applications use) and 16 KB (a good choice for applications that generate sequential I/O, such as multimedia).

5. Click **OK**.
The storage subsystem controller cache settings are updated.

Using the Script Editor

Many storage management commands that are available through the Subsystem Management window can also be sent to the storage subsystem using statements in scripts. Use the Script Editor to create or edit a script file, save a script file to the management station local disk, or load a script file from disk. The Script Editor has an underlying function that verifies statement syntax, interprets the statements, converts statements to the appropriate protocol-compliant commands, and passes the commands to the storage subsystem where they are run by the storage subsystem controller. For a list of command statements and how to use them, see *Script Editor Command Reference* in the online help.

To open the Script Editor, select the desired storage subsystem in the Enterprise Management window from either the device tree or device table. Then, click **Tools**→**Execute Script**.

When the Script Editor window opens, there are two views in the window:

- **Script view:** Provides an area for typing and editing script commands.
- **Output view:** Displays verification or execution results.

A horizontal splitter bar divides the window between Script view and the Output view. You can use the splitter bar to resize the views.

```
Flute1_Upper - Script Editor - newscrip.scr
File Edit Tools Help
/*The following command assigns hot spare drives.*/
set drives [1,2 1,3] hotspare=true;

Performing syntax check...
Encountered "(" at line 2, column 12.
Was expecting:
 "[" ...
Performing syntax check...
Syntax check complete.
```

Script Usage Guidelines

The following are script usage guidelines:

- All statements must end with a (;)
- Each base command and its associated primary and secondary parameters must be separated with a space.
- The Script Editor is not case sensitive.
- Each statement must be on a separate line.
- Comments that are added to your scripts make it easier for yourself and future users to understand the purpose of the command statements.

Adding Comments to a Script

The Script Editor supports the following comment formats:

- Text contained after two forward slashes // until an **Enter** character is reached. For example:

```
//The following command assigns hot spare drives"
set drives [1,2 1,3] hotspare=true;
```

You must end a comment beginning with // with an end-of-line character, inserted by pressing the **Enter** key. If the Script Editor does not find an end-of-line character in the script after processing a comment, an error message is displayed, and the script is terminated. This error commonly occurs when a comment is placed at the end of a script and you have forgotten to press **Enter**.

- Text is contained between the characters /* and */.

```
/* The following command assigns hot spare drives.*/  
set drives [1,2 1,3]hotspare=true;
```

A comment must begin with /* and end with */. If the Script Editor does not find both a beginning and ending comment notation, an error message is displayed, and the script execution is terminated.

Show Statement

Use the Show statement to embed comments in your script that will display in the Output view when the script runs. If you include Show *setting controller mode* statement in your script the setting controller mode is displayed in the Output view when the script is run. For more information about the Show statement, see *Script Editor Command Reference* in the online help.

Verifying Script Syntax

The Script Editor engine parses the statements, one line at a time, in your script file and verifies that they have the correct syntax. Any syntax errors are displayed in the Output view, reporting the line number of the error and description of the error. If the Script Editor encounters a syntax error, no further syntax verification is performed on the script. Correct the syntax error and rerun the Verify Syntax command to validate the error correction and check the remainder of the statements in the script.

To verify the script command syntax, from the Script Editor window, click **Tools** → **Verify Syntax**.

Script Editor File and Edit Options

The script editing options are:

Options	How to Perform
Open a new script	Click File → New Script . The Script view is cleared for new script input and the Output view is cleared.
Edit an existing script	<ol style="list-style-type: none"> 1. Click File → Load Script. The Load Script file selection window opens. 2. Click the script file you want to edit from your file system, and then click OK. The Load Script dialog closes and the selected script file is loaded into the Script view.
Save the script in the Script View without prompting for a new file name.	Click File → Save Script . If you exit from Script Editor without saving a modified script, you are prompted to save your script before exiting.

Options	How to Perform
<p>Save the script in the Script view with a new file name.</p>	<ol style="list-style-type: none"><li data-bbox="836 204 1225 413">1. Click File→ Save Script As. The Save Script As window opens. Any subdirectory folders and file names with a .scr extension is listed. To show all files, click All Files from the Files of type list box.<li data-bbox="836 430 1225 873">2. Select a subdirectory folder, name the script file, and click Save. The Save Script As window closes and the script file is saved in the selected directory. Note: You can specify any file extension, but the Save Script As dialog by default shows files with a .scr extension. The Script Editor does not automatically append the .scr to the end of the file.

Options	How to Perform
Save the script execution results to a local file	<ol style="list-style-type: none"> 1. Click File → Save Output As. The Save Output As window opens, listing any subdirectory folders and showing any filenames with a .txt extension. To show all files, click All Files from the Files of type list box. 2. Select a subdirectory folder, if desired, name the output file, and click Save. The Save Output As dialog closes, and the output file is saved in the selected directory. <p>Note: You can use any file extension, but the Save Output As dialog by default shows files with a .txt extension. The Script Editor does not automatically append the .txt extension.</p>
Cut selected text from the script in the Script view	Click Edit → Cut .
Copy selected text from the Script view or Output view to system memory	Click Edit → Copy .
Paste information from system memory into the script	Click Edit → Paste .
Clear the text in the Script view	Click Edit → Clear Script .
Clear the text in the Output view	Click Edit → Clear Output
Select all of the text in the Script view	Click Edit → Select All

Interpreting the Script Execution Results

During script execution, messages are displayed in the Output view beginning with

Executing script...

After a successful script execution, you see the message

```
Script execution complete.
```

If there is an error during the parse phase, an error message is displayed in the Output view giving the line and column number and a description of the syntax error.

If you type the statement in a script:

```
set controller[a] mod = passive;
```

in a script, the resulting syntax error

```
Encountered "mod" at line 2, column 19
was expecting one of...
    "mode"....
    "availability"...
    "NVSRAMbyte"...
```

is displayed in the Output view.

If there is an error during execution, a message is displayed in the output area stating that the command failed and reporting a description of the error.

If you type the statement

```
set logical drive [three] userLabel = "OneOne";
```

in a script, be sent to the Storage Subsystem in an Unresponsive state. The resulting error is displayed:

```
Unable to change the logical drive user label using the Set
Logical Drive command at line 1
error - 1 Could not communicate with the Storage Subsystem
to complete this request.
```

Certain execution errors, including the inability to communicate with the storage subsystem, always cause script execution to halt. In these cases, execution stops even if you have used the ON Error Continue statement. For more information, see *Recovering from an Unresponsive Controller or Storage Subsystem Condition*, in the online help.

Executing Scripts

The Script Editor parses the command statements in the script, interprets and converts the statements to the appropriate commands, and sends the commands to the storage subsystem controller. If a syntax error is encountered, the execution stops and an error message is displayed. Correct the error, and then use the **Verify Syntax** or **Verify** and **Execute** options to validate the error correction.

To run the currently loaded script, click **Tools** → **Verify** and **Execute** or **Tools** → **Execute Only** (Both options display an error message if a syntax error is encountered, but

Tools → **Verify** and **Execute** checks the script one line at a time for correct syntax before the script is executed, and it provides syntax error feedback in the Output view.

If an execution error occurs, the script might or might not continue to execute depending on the included **On Error** script statement.

- The On Error Stop statement stops the script if an execution error is encountered (this is the default).
- The On Error Continue statement allows the script to continue even after an execution error is encountered.

Configuration File Errors

There are three possible configuration file errors:

Configuration File Startup Error

Either configuration file (**emwdata.bin** and **emwback.bin**) is discovered to have a read or write permissions problem when you are starting a management session. Follow the instructions on the window to recover from the error.

Note: When you click **OK** on this window, the Enterprise Management window closes. You *must* correct this error before you can successfully open a management session.

Configuration File Write Error

Either configuration file (**emwdata.bin** and **emwback.bin**) is discovered to have a read or write permissions problem during a management session, or the write fails because the local disk is full. Follow the instructions on the dialog to recover from the error.

Note: Do not click **Cancel** if you want to attempt to correct the problem. Correct the problem and then click **Retry**.

Configuration File Corruption Error

Both configuration files (**emwdata.bin** and **emwback.bin**) are damaged and cannot be recovered, or one of the configuration files is damaged and the other one is missing.

Note: Make sure that you record the location of the directory path that is listed on the error dialog. You might need to use this directory path in the Recovery from a Configuration File Corruption Error procedure.

Recovering from a Configuration File Corruption Error

When a configuration file corruption error occurs, the only action you can initially take is to click **OK** because the current configuration files cannot be recovered. Clicking **OK** will replace the configuration files and reset them to empty files. This results in the Enterprise Management window displaying no devices, reflecting an empty management domain. After you click **OK**, you are given an opportunity to perform an automatic discovery to add devices back to the management domain and rebuild your files.

You have the following choices:

Use alternate configuration files

Use either a backup copy of your configuration files or configuration files from another management station:

1. Click **No** at the **Automatic Discovery** window.
2. Exit from the Enterprise Management window.
3. Copy the backup files from the alternative directory to the directory that is specified in the error dialog.
4. Restart the Enterprise Management window.
5. Verify the displayed devices and, if necessary, use the **Add Device** option to add any devices that are missing.

Note: Use backup copies of your own configuration files. If these are not available, you can use another system administrator's files that have been generated from either the same or a different management station.

Rebuild your configuration files

If you do not have backup copies of the configuration files, you will have to completely rebuild your management domain using one of the two listed alternatives:

- Click **Yes** if you want to perform an automatic discovery to repopulate the management domain.
- Click **No** if you want to manually add the devices back to the management domain using the **Add Device** option.

Defining Security

To maximize storage subsystem security use the password option. From the Subsystem Management window, select **Storage Subsystem** → **Change Password**.

Setting or Changing a Storage Subsystem Password

Use the Change Password option to establish a new or change an existing one.

Unless you specify a password using this option, all options are available within Fibre Channel Storage Manager. If you specify a password, any option that is destructive will be password protected. The controller firmware determines which options are destructive.

Implementing destructive commands on a storage subsystem can cause serious damage, including data loss. Use a long password with a minimum of 15 alphanumeric characters to increase security.

Notes:

- If no password has previously been set, no current password is required to establish a new password.
- The password is stored on the storage subsystem. Therefore, if desired, you will need to provide a password for each storage subsystem.
- You cannot change a storage subsystem password unless the current password is supplied first. The Storage Subsystem password must be cleared through a specific controller shell command.
- The maximum password length is 30 characters.
- Passwords are case sensitive. Remember your use of uppercase and lowercase letters when you change a password.
- Trailing white spaces are not removed from passwords. Be careful not to include trailing spaces in the new password, because they can be difficult to type accurately in the future.
- If you no longer want to have the destructive operations password protected, enter the current password, then leave the **New password** and **Confirm password** text boxes blank.

To set or change a Storage Subsystem password, use the following procedure:

1. From the Subsystem Management Window, click the **Storage Subsystem** → **Change Password**.
The Change Password window opens.
Note: If you are setting the password for the first time, you do not need to type the current password. Skip step 2 and go to step 3.
2. Type the current password in the **Current password** text box. Press the **Tab** key to move to the **New password** text box.
3. Type the new password in the **New password** text box. Press the **Tab** key on your keyboard to move to the **Confirm password** text box.
4. Type the new password in the **Confirm password** text box must match the password in the **New password** text box.
5. Click **OK** or press **Enter** to change the Storage Subsystem password.
The **Password Successfully Changed** confirmation window is displayed.
6. Click **OK** to close the **Password Successfully Changed** window.
Note: While you are in the current management session, you will not be asked for a password when you attempt a destructive operation because you have just set the password from that management station. However, the next time you start the Subsystem Management window, you will be prompted for a password the first time you attempt a destructive operation.

Entering a Storage Subsystem Password

The Enter Current Password window is displayed when:

- A storage subsystem password has been specified (as shown in Setting or Changing a Storage Subsystem Password on page 186).
- You are attempting to use a password protected option.

The options that are password protected are determined by the controller and are generally those that cause destructive Storage Subsystem operations.

Notes:

- Only asterisks are displayed when you type a password.
- Passwords are case sensitive.
- Trailing white spaces are removed from passwords when they are set; be careful to include white spaces if they were included as part of the password.
- There is a 30-character maximum password length.
- You will be asked for the password only once during a single management session. However, if a user who is managing the same storage subsystem from another management station changes the password while your session is in progress, you will be prompted for a password the next time you attempt a destructive operation.
- The first time a destructive operation is attempted during the management session, the Enter Current Password window is displayed containing the initial instructions. If an invalid password is typed, the Enter Current Password window is displayed again with instructions indicating that an invalid password was supplied. The Enter Current Password window will be repeatedly displayed until either the correct password is supplied or you click **Cancel**.
- There is no limit to the number of times a password can be tried.

To enter a storage subsystem password, use the following procedure:

1. In the Enter Current Password dialog, type the storage subsystem password.
2. Select **OK** to enter the password.
The Enter Current Password window is closed, and the attempted operation completes.

Note: If you do not know the password, click **Cancel**. An error message is displayed indicating that the operation cannot

be completed because a password has not been typed.
Click **OK** to close the window.

Using Passwords in the Script Editor

If destructive commands are contained in a script, a command line password must be entered in the script using the **SET** command. See "Using the Script Editor" on page 176.

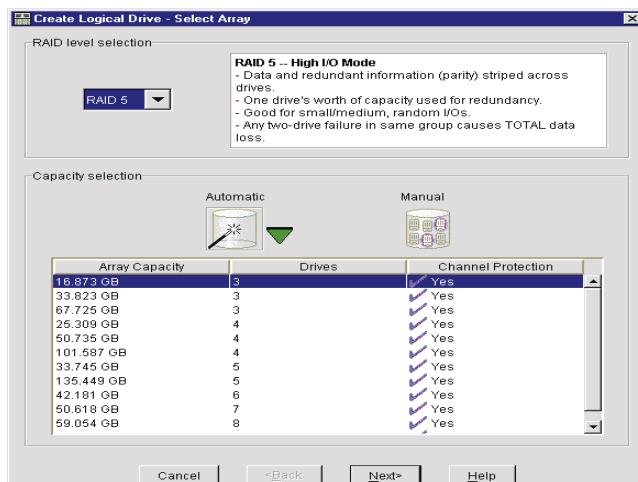
Defining an Array

The first step in creating an individual logical drive from unconfigured capacity is to define an array. When you define the array, you specify two parameters: RAID level and capacity.

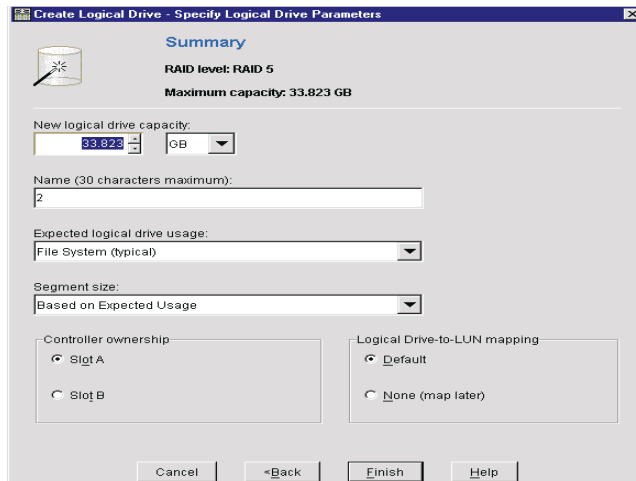
This window is available only when you are creating logical drives from the unconfigured capacity node. If you are creating a logical drive on an existing array using a free capacity node, this dialog is bypassed because the array (and its associated drives and RAID level) have already been specified. For a free capacity node, the only window that is displayed is the Specify Logical Drive Parameters window.

To create an array and specify Logical drive parameters, use the following procedure:

1. Right-click an unconfigured capacity in the Logical view and click **Create Array**
The Create Logical Drive - Select Array window opens.



2. In the **RAID level selection** drop-down list, specify the RAID level that will meet your data storage and protection requirements. RAID level choices include RAID 0, RAID 1, RAID 3, and RAID 5.
3. In the **Capacity selection** group box, select a method for defining the available drives to be used in the array. The methods are automatic and manual.
4. Select the available arrays capacity, drives and channel protection. See "Channel Protection" on page 207.
5. Click **Next**.
The Specify Logical Drive Parameters window opens.



6. Specify the capacity units and logical drive capacity. Use the drop-down list to indicate capacity units (in MG or GB) and use the list box to indicate the specific numeric capacity. You cannot specify a capacity greater than the value shown in the Summary area at the top of the window.
 - You can type up to three decimal places for the capacity.
 - If you type a capacity larger than the maximum, the value will default to the maximum capacity when you move to another area on the dialog or when you select **Finish**.
 - If you type a capacity that is less than the maximum, make sure that you leave enough remaining capacity to enable you to create one or more additional logical drives that meet your application needs.
7. In the **Name** text box, specify the logical drive name. Choose the unique, meaningful name that helps you identify the logical drive. Generally, a meaningful name describes the data that will be stored on the logical drive. There is a 30-character maximum. If you choose the name of a logical drive that duplicates the name of another logical drive in the storage subsystem, an error message is displayed and you are prompted to choose another name. Any leading and trailing spaces in the name are deleted.

-
8. In the **Expected logical drive usage** drop-down list specify the expected usage of the logical drive based on your application needs.
The expected logical drive usage that you specify is used by the controller firmware to indicate an appropriate default logical drive usage, including file system (the default), database, and multimedia.
 9. In the **Segment size** drop-down list, select an appropriate segment size for your application. You can select the default (based on expected usage), 8K, 16K, 32K, 64K, 128K, or 256K. (You can change the logical drive segment size later using the **Change Segment Size** option. For more information, see *Changing Segment Size* in the online help.)
 10. Under **Controller ownership**, specify controller ownership (a preferred controller). This option is available only if you are using dual active controllers. Choices include either the controller in slot A or the controller in slot B.
Review the following table.

If you are creating this logical drive from:	Then
An unconfigured capacity node	The storage subsystem controller firmware automatically selects a preferred owner. You can change the preferred owner by selecting the other controller radio button.
A free capacity node	The preferred owner is set automatically by determining the owner of the other logical drives in the array. You cannot change controller ownership because the same controller <i>must</i> own all logical drives in the same array.

11. Under **Logical Drive-to-LUN mapping**, specify whether the logical drive should have a default mapping or be left unmapped.

For more information, see *Introducing Storage Partitions* in the online help.

Default Mapping

The logical drive will be available to all host groups or hosts that do not have specific logical drive-to-LUN mappings. (All hosts or host groups in the default host group can access the logical drive.)

You can change the default mapping at a later time using the **Storage Partitions Definitions** window.

No Mapping (map later)

The logical drive cannot be accessed until you make a specific logical drive-to-LUN mapping using the Storage Partitions Definitions Window.

This option is not available if you have not enabled the Storage Partitions premium feature.

12. Click **Finish**.
If the logical drive is created successfully, the **Logical Drive Successful** window opens. If there are any errors, an error message is displayed.
13. Click **OK**.
A logical drive (or array/logical drive) is created with the specified RAID level and parameters. You are returned to the Storage Subsystem Management window. The Logical view is updated to reflect the addition of the new array, logical drive, or both and the subtraction of capacity from either the free capacity node or the unconfigured capacity node.
14. Perform any operating system modifications that are necessary on the application host so that the applications can use the logical drive. See the *Installation and Support Guide* specific to your operating system.

Defragmenting an Array

You can consolidate all free capacity on a selected array so that additional logical drives can be created from the maximum amount of free capacity. A fragmented array can result from the deletion of logical drives or from not using the entire available free capacity in a free capacity node during logical drive creation.

For example, you have two logical drives, named *New York* and *Dallas*. Each logical drive is 2 GB. You would like to create a new logical drive, named *Chicago* of 2 GB in the same array. There is 2 GB of space available in the array, but it is split between two Free Capacity nodes that are not contiguous and cannot be organized into a single logical drive.

To create a new logical drive *Chicago* of 2 GB, you must first defragment the array so that the space in the two free capacity nodes is contiguous.

To use the Defragment an array option, all logical drives in the array must be online and optimal, and there can be no logical drive modification operations (such as Change of Segment Size) in progress. You cannot cancel the operation after it begins. To defragment an array, use the following procedure:

1. From the Subsystem Management window, select an array in the Logical view, and click **Array Defragment**. The Confirm Defragment window is displayed.
2. Click **Yes**.
All free data blocks on the free capacity nodes are consolidated into one free capacity node containing a large set of contiguous data blocks. The Logical view is updated to display the larger, single free capacity node. You can select this free capacity node for logical drive creation.

Adding Free Capacity to an Array

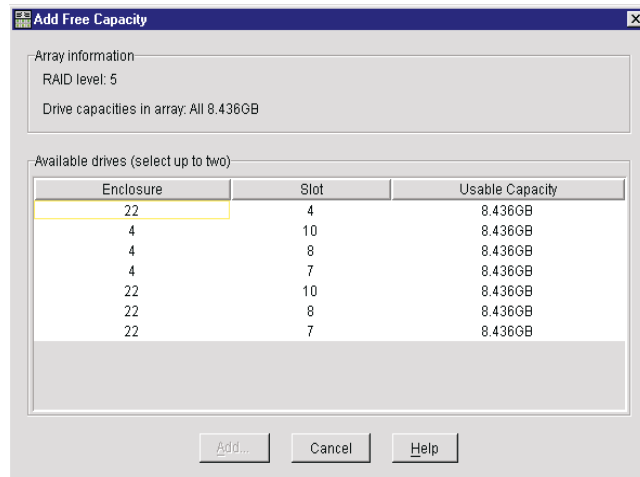
You can expand the capacity of a selected array by adding unassigned drives. The result of the expansion causes additional free capacity to be included in the array. You can use this free capacity to create additional logical drives.

Notes:

- You cannot cancel this operation after it begins.
- Your data remains available during this operation.
- The array must be in an *Optimal* state before you can perform this operation.
- This option will work only if there are unassigned drives in the storage subsystem.
- Existing logical drives in the array do not increase in size as a result of this operation. Existing logical drive capacity is redistributed over the larger number of drives.
- There is a controller firmware-determined limit to the number of drives that you can select for a single **Add Free Capacity** operation. Selecting more drives than the allowed number causes an error message to be displayed. However, after you have completed the first **Add Free Capacity** operation, you can repeat the option to add more drives until the desired capacity is reached.
- In a RAID 1 logical drive, you must add a minimum of two drives. (You must add an even number of drives, half of which are used for redundancy).
- Whenever possible, select drives that have a capacity equal to the current drive capacities in the array.
- If you must add unassigned drives with a smaller capacity, be aware that the usable capacity of each drive currently in the array will be reduced so that the drive capacity is the same across the array. If there is data on the current drives in the array that could be lost when the usable capacity is reduced, an error message is displayed.
- If you must add unassigned drives with a larger capacity, be aware that the usable capacity of the unassigned drives that you add will be reduced so that they match the current drive capacities in the array.

To add free capacity to an array, select an array in the Logical view and use the following procedure:

1. From the Storage Subsystem window, select **Array**→ **Add Free Capacity**.
The Add Free Capacity window opens. The window displays the RAID level and the existing drive capacities of the selected array. A table is also displayed that contains information about the unassigned drives that are available for selection. For Fibre Channel drives, the ID is the actual ID rather than the preferred ID.



Notes:

- If you are adding SCSI drives, it is best to select drives that are on different channels so that the impact of a channel failure is minimized.
 - Whenever possible, select drives that have a capacity equal to the current drive capacities in the array.
2. Under **Available drives**, select up to the maximum allowable number of drives in the list, using **Ctrl+ Click** or **Shift+ Click** for drives that are next to each other on the list.
 3. Select **Add**.
The Add Free Capacity Confirmation window opens.
 4. Click **Yes** to add the capacity to the array.
The Add Free Capacity operation begins.

For detailed information on Logical view changes and Physical view changes after adding free capacity to an array see online help.

Redistributing Arrays

By redistributing arrays, you can move arrays (and their associated logical drives) back to their preferred owners. Typically, multiple-path device drivers move arrays from their preferred controller owner when there is a failure in the connection between the host and the storage subsystem. This option is not available if all arrays are currently owned by their preferred controllers, or if no logical drives exist on the storage subsystem.

If you use this option without a multiple-path device driver installed on the hosts, stop I/O to the logical drives while this operation is in progress to prevent application errors. You might also need to make operating system modifications to recognize the new I/O path to the logical drives.

To redistribute arrays, use the following procedure:

1. From the Subsystem Management window, click **Storage Subsystem** → **Redistribute Array**.
The Redistribute Arrays confirmation window opens.
2. Click **Yes**.
All storage subsystem arrays are redistributed back to their preferred owners. Selecting an array in the Logical view displays an associated controller in the Physical view that is both the current owner and the preferred owner.
Note: Under certain application host operating system environments, a reconfiguration of the host device driver might be an additional requirement.

Logical Drive

A logical drive is the logical entity that is created to store your data. You can create a single logical drive from either unconfigured capacity or free capacity. With larger capacity-drives and the ability to share logical drives across multiple hosts, you might want to create more than one logical drive per array to better organize your data. Click the **Create Logical Drive** icon for each new logical drive that you want to create.

Note: For certain operating systems, you can create only a single logical drive on the array and it must use all of the available capacity. Fibre Channel Storage Manager enforces these restrictions when they apply to your operating system. There can be host operating-system-specific limits on the number of logical drives that the host can access. Keep this in mind when creating logical drives for use by a particular host. After creating one or more logical drives, run the host-based **SM7devices** utility to correlate logical drive names with host device names. For information on using this utility, see the *Installation and Support Guide*.

Creating a Logical Drive

To create a logical drive, select either a Unconfigured Capacity node or Free Capacity node from the Logical view in the Storage Subsystem menu and use one of the following procedures:

- Click the **Create Logical Drive** icon.
- Click **Configure** → **Create Logical Drive**.
- Right-click the **Unconfigured Capacity** or the **Free Capacity** and click **Create Logical Drive**.

Creating a Logical Drive from Unconfigured Capacity

To create an individual logical drive from unconfigured capacity you must first define an array and then specify the logical drive parameter. See the procedure in "Defining an Array" on page 189.

Deleting a Logical Drive

You can delete a logical drive that was created with the wrong parameters or capacity or no longer meets application storage needs. This will result in an increase of free capacity in the array or additional unconfigured capacity. Deleting a logical drive causes the loss of all data that is contained on the logical drive. Back up the data and stop all I/O before performing this operation. In addition, if a file system is mounted on the logical drive, unmount it before attempting this operation.

Unless the logical drive is in a Failed state, you cannot delete a logical drive that is initializing or reconstructing or that is part of an array that

contains a drive that is undergoing a copyback operation or undergoing a modification operation such as a change of segment size.

To delete the a Logical drive, use the following procedure:

1. Select the logical drive in the Logical view and click **Logical Drive →Delete**.
The **Confirm Delete** window opens.
2. Click **Yes**.
The logical drive is deleted and any data that was on it is destroyed.
The following table list additional considerations.

If the deleted logical drive was	Then
The only logical drive in the array	<ul style="list-style-type: none">■ The drives associated with the deleted logical drive return to an unassigned state in the Physical view.■ The array is removed from the Logical view.■ The raw capacity of the newly unassigned drives is added to the unconfigured capacity node in the Logical view.
One of the multiple logical drives in the array	The size of the free-capacity node is increased if the deleted logical drive was adjacent to a free-capacity node in the array. Otherwise, a new free-capacity node is created.

3. If you see system error messages relating to this deleted logical drive, reconfigure or restart you host system to permanently remove any system information about this logical drive.

Changing Logical Drive Parameters

Click **Create Logical Drive**→**Specify Logical Drive Parameters** window to specify the logical drive name, usage, desired capacity, controller ownership, and storage partition mapping preference.

Create Logical Drive - Specify Logical Drive Parameters

Summary

RAID level: RAID 5
Maximum capacity: 33.823 GB

New logical drive capacity:
33.823 GB

Name (30 characters maximum):
z

Expected logical drive usage:
File System (typical)

Segment size:
Based on Expected Usage

Controller ownership:
 Slot A
 Slot B

Logical Drive-to-LUN mapping:
 Default
 None (map later)

Cancel <Back Finish Help

If you select a free capacity node when selecting the **Create Logical Drive** option and there are no valid logical drive candidates, an error message is displayed and you are returned to the Storage Subsystem Management window.

To specify logical drive parameters, use the following procedure:

1. In the **New logical drive capacity** drop-down list, specify the capacity units and logical drive capacity.
Use the drop down list to indicate capacity units (in MG or GB) and use the list box to indicate the specific numeric capacity. The maximum value is shown in the Maximum capacity field.
 - You can type up to three decimal places for the capacity.
 - If you type a capacity larger than the maximum, the value will default to the maximum capacity when you move to another area on the dialog or when you click **Finish**.
 - If you type a capacity that is less than the maximum, make sure that you leave enough remaining capacity to enable you

to create one or more additional logical drives that meet your application needs.

2. In the **Name** text box, specify the logical drive name. Choose a unique, meaningful name that helps you identify the logical drive. Generally, a meaningful name describes the data that will be stored on the logical drive. There is a 30-character maximum. If you choose the name of a logical drive that duplicates the name of another logical drive in the storage subsystem, an error message is displayed and you are prompted to choose another name. Any leading and trailing spaces in the name are deleted.
3. From the **Expected logical drive usage** drop-down list, specify the expected usage of the logical drive based on your application needs.

The expected logical drive usage that you specify is used by the controller firmware to indicate an appropriate default logical drive usage, including file system (the default), database, and multimedia.
4. In the **Segment size** drop-down list, select an appropriate segment size for your application. You can select the default (based on expected usage), 8, 16, 32, 64, 128, or 256K. (You can change the logical drive segment size later using the Change Segment Size option. For more information, see *Changing Segment Size* in the online help.)
5. In the **Controller ownership** group box, specify controller ownership (a preferred controller). This option is available only if you are using dual active controllers. Choices include either the controller in slot A or the controller in slot B.

Review the following table.

If you are creating this logical drive from:	Then
An unconfigured capacity node	The storage subsystem controller firmware automatically selects a preferred owner. You can change the preferred owner by selecting the other controller radio button.
A free-capacity node	The preferred owner is set automatically by determining the owner of the other logical drives in the array. You cannot change controller ownership because the same controller must own all logical drives in the same array.

6. In the **Logical Drive-to-LUN mapping** group box, specify whether the logical drive should have a default mapping or be left unmapped. For more information, see *Introducing Storage Partitions* in the online help.

Default Mapping

The logical drive will be available to all host groups or hosts that do not have specific logical drive-to-LUN mappings. (All hosts or host groups in the default host group can access the logical drive.)

You can change the default mapping at a later time using the Storage Partitions Definitions window.

No Mapping (map later)

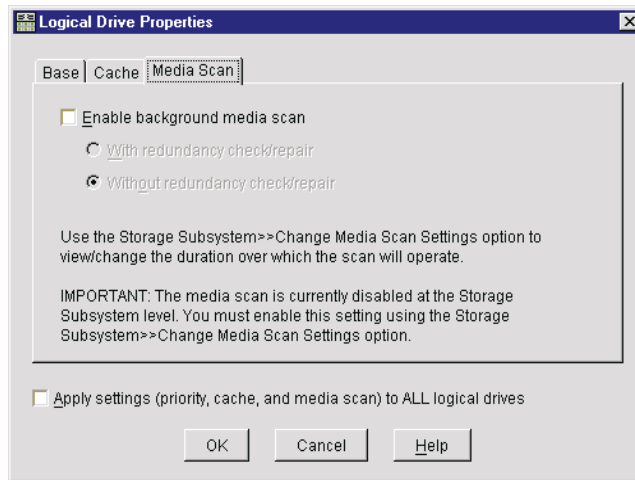
The logical drive cannot be accessed until you make a specific logical drive-to-LUN mapping using the **Storage Partitions Definitions** window.

This option is not available if you have not enabled the Storage Partitions premium.

-
7. Click **Finish**.
If the logical drive is created successfully, the Logical Drive Successful window opens. If there are any errors, an error message is displayed.
 8. Select **OK**.
A logical drive (or array/logical drive) is created with the specified RAID level and parameters. You are returned to the Storage Subsystem Management window. The Logical view is updated to reflect the addition of the new array, logical drive, or both and the subtraction of capacity from either the free-capacity node or the unconfigured-Capacity node.
To change other logical drive parameters, see *Specifying Logical Drive Properties* in the online help.
 9. Perform any operating system modifications that are necessary on the application host so that the applications can use the logical drive. See the *Installation and Support Guide* specific to your operating system.

Applying Individual Logical Drive Properties to All Logical Drives

To apply all settings of an individual logical drive (modification priority, cache, and media scan) to all other logical drives in a storage subsystem, right-click a logical drive and click **Properties**. The Logical Drive Properties window opens. Select the **Apply Settings to All Logical Drives** check box. Click **OK**. Use this check box only if you want all of the logical drives to share the same settings. This option overwrites any current settings on the other logical drives in the storage subsystem.



Creating a Logical Drive from Free Capacity

Because the free capacity is on an existing array only, the overall capacity (that is, the number of drives) and RAID level have already been specified. Therefore, the **Select Array** window is bypassed when you create a logical drive from free capacity. When you select a free-capacity node, the only step that is required is to specify the parameters for the individual logical drive that you plan to create on the array. Use the **Specify Logical Drive Parameters** window to specify those parameters. For more information, see "Changing Logical Drive Parameters", page 200.

Failing a Drive

Note: This operation could result in data loss or the loss of data redundancy. Do not use this operation unless advised to do so by the Recovery Guru. This operation will fail when there is a logical drive modification operation in progress

Data loss will occur under the following conditions:

- Failure of a single drive in a RAID 0 logical drive

- Failure of a single drive in a degraded RAID 1 logical drive, if the failed drive is the mirror drive of another failed drive
- Failure of a single drive in a degraded RAID 3 or 5 logical drive
- Failure of two or more drives in an Optimal RAID 3 or 5 logical drive

Loss of data redundancy will occur under the following conditions:

- Failure of a single drive in a RAID 1 logical drive
- Failure of a single drive in an Optimal RAID 3 or 5 logical drive

To fail a drive or drives, use the following procedure:

1. Select one or more drives in the Physical view.
 2. Click the **Drive** → **Fail**.
The Confirm Fail Drive window opens.
 3. Click **Yes**.
The drives fail and the window closes. The drive in the Physical view is shown in the Needs Attention state in the selected mode.
- Note:** If you are going to remove the drive, make sure at least 30 seconds have elapsed before you remove the drive.

If the drive was assigned to an array, the Logical view changes are shown in the following table.

If	Then
The logical drive was Optimal before the drive was failed	The logical drive changes to a Degraded state in the Logical view.
The logical drive was Degraded before the drive was failed.	The logical drive changes to a failed state in the Logical view.

Reconstructing a Drive

Start manual reconstruction of a drive only when instructed to do so by **Recovery Guru**.

This option is only available when a drive is assigned to a RAID 1, 3, or 5 array with either a Failed state or a Replaced state, and reconstruction did not automatically start.

To reconstruct a drive manually, use the following procedure:

1. Select one or more drives in the Physical view.
2. Click **Drive**→**Reconstruct**.
The Confirm Reconstruct window opens.
3. Click **Yes**.
The Confirm Reconstruct dialog is closed and the drive starts to reconstruct. The drive icon in the Physical view changes to Optimal. The associated logical drive icons in the Logical view change to Operation in Progress during the reconstruction, then to Optimal when the reconstruction is finished.

To view the reconstruction progress, select a logical drive in the Logical view that is associated with this drive, then click **Logical** →**Drive Properties**.

Initializing a Drive

You must initialize a drive when you have moved a subset of drives that were previously part of a multiple-drive array from one storage subsystem to another. Because you did not move the entire set of drives, the logical drive/array information on the drives is incomplete.

Therefore, to erase all logical drive/array information on these drives and make them available again for new logical drive creation, you must initialize these drives. This will erase the logical drive/array information and return the selected drives to an unassigned state, adding new or additional unconfigured capacity to the storage subsystem. You can use this capacity to create additional logical drives.

To initialize selected drives, use the following procedure:

1. From the Physical view, select one or more drives that you want to initialize.

2. Click **Drive**→ **Initialize** menu.
The **Initialize Drive** window opens.
3. Click **Yes**.
The drive initialization process begins, and the drive in the **Physical** view returns to an unassigned state.

The **Logical View** is updated as follows:

If	Then
An unconfigured-capacity node did not exist prior to the drive initialization operation	A new unconfigured capacity node is added to the Logical view, containing the raw capacity of the newly initialized drives.
An unconfigured-capacity node existed prior to the drive initialization operation	The raw capacity of the newly initialized drives is added to the existing unconfigured-capacity node.

Channel Protection

When creating logical drives from unconfigured capacity, the storage management software displays either **Yes** or **No** in the **Channel Protection** column for each array candidate. See "Defining an Array" on page 189.

In a SCSI environment, channel protection depends on the RAID level of the array and how many drives are present on any single drive channel. For example, a RAID 5 array would not have channel protection if more than one was present on a single drive channel.

In a Fibre Channel environment, channel protection is usually present for any array candidate, because when the storage subsystem is properly cabled there are two redundant Fibre Channel arbitrated loops for each drive.

Media Scan

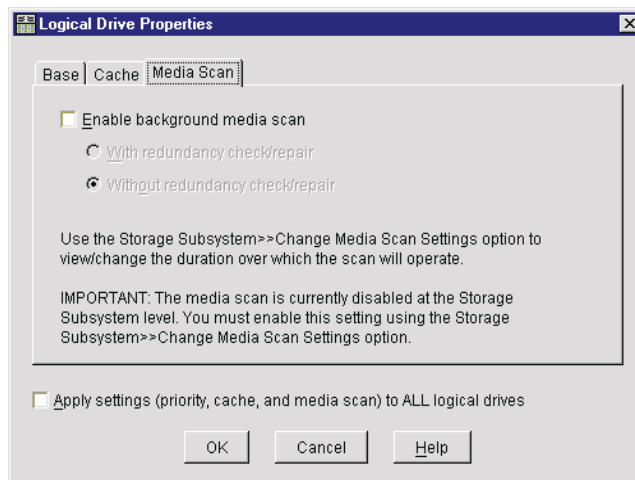
A Media Scan is a background process that runs on all logical drives in the storage subsystem for which it has been enabled, providing error detection on the drive media. The advantage of enabling the Media

Scan process is that the process can find media errors before they disrupt read and write operations. The Media Scan process scans all logical drive data to verify that it can be accessed, and optionally scans the logical drive redundancy data as well.

You can set the duration over which the Media Scan runs. Right-click a logical drive, and click **Properties**. The Logical Drive Properties window opens. Click the **Media Scan** tab.

To enable a background Media Scan for a logical drive, select the Enable background scan check box.

To enable a redundancy check during a Media Scan, select the **With redundancy check** option.



During the redundancy check, all data blocks in a logical drive are scanned, and:

- In a RAID 3 or 5 logical drive, redundancy is checked and repaired.
- In a RAID 1 logical drive, the data is compared between the mirrored drives and data inconsistencies are corrected.

Note: RAID 0 logical drives have no data redundancy.

A media scan discovers the following errors and reports them to the Event Log:

Error	Description
Unrecovered media error	<p>The drive could not read the requested data on its first attempt or on any subsequent retries.</p> <p>This occurs on Logical drives with redundancy protection. Data is reconstructed, rewritten to the drive, and verified. The error is reported to the Event log.</p>
Recovered media error	<p>The drive could not read the requested data on its first attempt, but succeeded on a subsequent attempt.</p> <p>The data is rewritten to the drive and verified. The error is reported on the Event log.</p>
Redundancy mismatches	<p>Redundancy errors are found and a media error is forced on the block stripe so that it is found when the drive is rechecked. If redundancy is repaired, this forced media error is removed.</p> <p>The first 10 redundancy mismatches that are found on a logical drive are reported to the event log.</p> <p>The Media Scan checks for redundancy only if the optional redundancy check box is enabled.</p>
Unfixable error	<p>The data could not be read, and parity or redundancy information could not be used to regenerate it. For example, redundancy information cannot be used to reconstruct data on a degraded logical drive.</p> <p>Errors are reported to the Event log.</p>

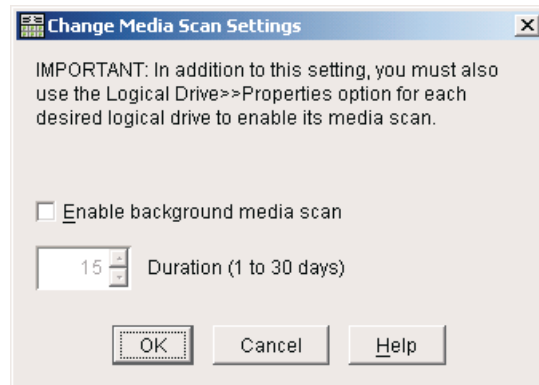
Changing the Storage Subsystem Media Scan Settings

You can enable a storage subsystem Media Scan and specify the duration period (in days) that you want the Media Scan to run. You can set the duration within minimum values that are supplied by the storage subsystem. When enabled, the Media Scan runs on all logical drives in the storage subsystem that are:

- Optimal
- Have no modification operations in progress
- Have the Media Scan parameter enabled in the **Logical Drive Properties**

To change the Media Scan Settings window, use the following procedure:

1. From the Subsystem Management window, click **Storage Subsystem** → **Change Media Scan Settings**. The Change Media Settings window opens.



2. Select the **Enable background media scan** check box.
3. Select a number in the drop-down list to specify the duration (in days) of the Media Scan.

4. Click **OK**.

The Media Scan settings for the storage subsystem is changed.

The Media Scan duration specifies the number of days over which the Media Scan should run on the eligible logical drives. For efficiency, the controllers in the storage subsystem complete the media scan at a faster rate during times when I/O is not present or is very light. Although the Media Scan might be completed more quickly than you have specified, Fibre Channel Storage Manager will not restart a complete media scan until the specified number of days have elapsed.

Hot Spare Drive

You can assign a selected unassigned drive as a hot-spare drive. A hot-spare drive contains no data and acts as a standby in case any drive fails in a RAID level 1, 2, or 5 array on the storage subsystem. The hot spare drive adds another level of redundancy to the storage subsystem. You can assign Optimal unassigned drives as hot-spare drives. Although the hot-spare drive is available to any RAID level 1,3, or 5 array in the storage subsystem, it is not specifically associated with any single array. When you assign a drive to be a hot spare, make sure that there is at least one other drive on the storage subsystem that is not assigned as a hot spare and that has a capacity equal to or smaller than the drive you assign as a hot spare. Select a drive with a capacity equal to or larger than the configured capacity of the drive that you want to cover with the hot spare. For example, if you have an 18 GB drive with configured capacity of 8 GB, you could use a 9 GB or larger drive as a hot spare. In general, however, you should not assign a drive to be a hot spare unless its capacity is equal to or greater than the capacity of the largest drive on the storage subsystem.

The maximum number of SCSI hot-spare drives that you can create is equal to the number of SCSI drive channels that are supported by the controller. The maximum number of Fibre Channel hot-spare drives that you can create is four.

Assigning a Hot-Spare Drive

To assign a hot-spare drive, select one or more drives in the Physical view from the available unassigned drives; then, click

Drive →Assign Hot Spare.

The drive changes from unassigned to Standby Hot-Spare. The capacity of the drive is subtracted from the unconfigured-capacity node in the Logical view.

Deassigning a Hot Spare Drive

To deassign a hot spare drive, select one or more Standby Hot-Spare drives in the Physical view. Click **Drive→ Deassign Hot.** The drive changes from Standby Hot-Spare to unassigned. The capacity of the drive is added to the unconfigured-capacity node in the Logical view.

Downloading Firmware or NVSRAM

Run this option from the Storage Subsystem Management window. This will cause a package file transfer to a controller. After receiving the package, the controller flash memory is updated with the new package and it is restarted. If you are downloading a firmware package, the Fibre Channel Storage Manager periodically polls the controller to see if the firmware version has changed to the new version. Next, the first controller in the download sequence passes the target package to the second controller. The flash memory of the second controller is updated with the new package, and it is restarted. If you are downloading a firmware package, the storage-management software polls the second controller periodically until the firmware version has changed to the new version.

Notes:


- If you do not have a multiple-path device driver installed on your application host operating system, stop I/O to the storage subsystem prior to the download to prevent application errors.
- The operating status of the storage subsystem controllers is checked when you perform a download operation, and an error message window is displayed if any controllers are not optimal. You are given the opportunity to stop or continue the download. Be sure to correct the non-optimal condition before continuing. However, if you choose to proceed with the download, the code

on the non-optimal controller will be synchronized after the non-optimal condition has been corrected.

- You might experience an erroneous failure message, stating that the download could not be complete, if a network problem occurs while Fibre Channel Storage Manager is polling a controller in the download sequence to verify that the firmware version has changed. Although the download might have been completed, the storage management software was unable to verify the operation success. Turn the power to the controller enclosure off and then on. Wait until the storage subsystem has finished initializing (this might take several minutes). Click **Storage Subsystem** → **Profile** in the Subsystem Management window to obtain the current firmware version. If the listed firmware does not match the version that you tried to download, retry the download.

To download a firmware or NVSRAM file, use the following procedure:

1. From the Subsystem Management window, click **Storage Subsystem** → **Download** → **Firmware** or **Storage Subsystem** → **Download** → **NVSRAM**.



2. Click **Download** → **Firmware** or **Download** → **NVSRAM** .
The Firmware Download or NVSRAM Download window opens.

-
3. The **Current configuration** group box lists the current firmware and NVSRAM versions in use by the storage subsystem controllers. The **File Selection** group box enables you to select a local or mapped directory and displays downloaded package files that are compatible with the current storage subsystem configuration that is stored in that directory. You can select a file to download from the displayed list.
 4. In the **File Selection** group box, change the home directory to the directory in which the file to be downloaded is located. Do this by clicking the directory folder to open it, or typing the folder name in the **File name** text box. Then, select the file.
 5. Select the package file to download from the **File Selection** group box doing by one of the following:
 - Double-click on the file.
 - Single-click the file, then click **OK**
 - Type the name of the file in the **File name** text box, then click **OK**

Notes:

- Only downloadable packages that are compatible with the current storage subsystem configuration are displayed.
 - When you select a file in the **File Selection** group box of the window, applicable attributes (if any) of the file are displayed in the **File Information** group box.
6. Click **OK**.

The Firmware or NVSRAM Download window is closed. The Confirm Download window opens.

If the selected file was not valid or compatible with the current storage subsystem configuration, the **File Selection Error** window opens. Click **OK** on this window to close it, and then close a compatible firmware or NVSRAM file.
 7. In the **Confirm Download** dialog, click **Yes**.

The Confirm Download window is closed. The download begins, and the Download Progress window opens.

8. Monitor the progress and completion status of the download to the controllers.

First, the file is transferred (downloaded) to the storage subsystem, and then each controller is updated with this new file, one at a time. If the update succeeds to the first controller, the second controller then receives the update. The status of the file transfer and the update to each participating controller is displayed in the dialog.

Note: The progress and status of optimal controllers participating in the download is displayed. Other controllers in the storage subsystem are not represented in the window.

Controller Mode

You can place a controller in an Active or Passive mode. The active mode is available only for a selected controller that is currently passive. Conversely, the passive mode is available only for a selected controller that is currently active.

Application errors will occur if you do not have a multiple-path device driver installed on the application host and you change a controller from Active or Passive mode while the application is using the associated logical drives.

Additionally, if you change an Active controller to Passive mode and the other controller of the pair is already in Passive mode, an error is displayed, and the operation will not be complete. At least one controller in the storage subsystem must be in Active mode.

To change a controller to Active mode, select the controller in Physical view and use the following procedure:

1. Click **Controller** → **Change Mode** → **Active**.
The Confirm Change to Active Mode window opens.
2. Click **Yes**.
The controller changes to **Active** mode. The controller icon in the **Physical View** changes to Active. (Arrays and associated logical drives are not moved to the new controller automatically).
3. Move one or more arrays and their associated logical drives to the newly active controller.

-
4. For more options, see the following table:

If	Then
The newly active controller previously owned arrays and their associated logical drives (before a previous switch to Passive mode).	Click Storage Subsystem → Redistribute Arrays to move the arrays and logical drives originally owned by this controller back to their controller.
The newly active controller never previously owned arrays and their associated logical drives.	Click Array → Change Ownership to move one or more arrays to this controller, or create new arrays and logical drives and specify this controller as the owner.

To change a controller to mode, select the controller in the Physical view and use the following procedure:

1. Click **Controller**→ **Change Mode**→ **Passive**.
The Confirm Change to Passive Mode window opens.
2. Select **Yes**.
The controller changes to Passive mode. The controller icon in the Physical view changes to passive mode. Any arrays and their associated logical drives that were owned by the controller change ownership to the other controller, and all I/O to these logical drives is automatically rerouted by the multiple-path device driver on the application host.

Note: If you do not have a multiple-path driver installed on the application host, make the appropriate host operating system modifications so that the applications can recognize the new I/O path

Placing a Controller Offline or Online

To place a controller offline, use the following procedure:

1. Select the controller in the Physical view and click **Controller**→ **Place Offline**.
The Confirm Place Offline window opens.

-
2. Click **Yes**.
The Place Offline operation begins. The controller icon changes to **Needs Attention, Passive**.
 3. If you do not have a multiple-path device driver installed on the application host, make the appropriate host operating system modifications so that the applications can recognize the new I/O path.

Notes:

- Use this option only when instructed to do so by **Recovery Guru**.
- If you do not use write cache mirroring, data in the cache of the controller that you placed offline will be lost.
- If you take a controller offline and you have controller failover protection through a host multiple-path device driver, the other controller in the pair takes over. Arrays and their associated logical drives that were assigned to the offline controller are automatically reassigned to the remaining good controller. If you do not have a multiple-path device driver installed on the application host and you take a controller offline while the application is using the associated logical drives, application errors will occur.

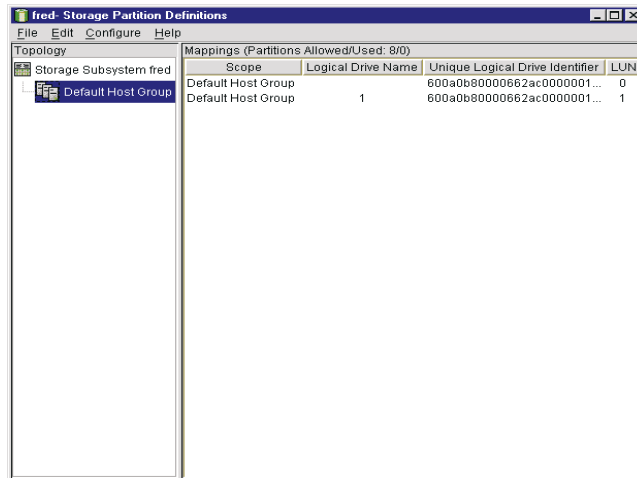
To place a controller online, use the following procedure:

1. Select the offline controller in the Physical view, and click **Controller → Place Online**.
The Place Online option begins. The controller icons changes to Optimal, Passive.
2. Click **Controller → Change Mode** to change the controller to Active mode.
The controller icon changes to Optimal, Active.
3. Click **Storage Subsystem → Redistribute Arrays** to move the formerly owned logical drives back to this controller owner.

Storage Partitions

A storage partition is a logical entity consisting of one or more storage subsystem logical drives that can be shared among hosts that are part of a host group or accessed by a single host. A storage partition is created when you define a collection of hosts (a host group) or a single host and then define a logical drive-to-logical unit number (LUN) mapping. This enables you to define what host group or host will have access to a particular logical drive in your storage subsystem.

Storage partitions are a premium feature of Fibre Channel Storage Manager and therefore must be enabled either by you or your storage vendor supplier. There is a maximum number of partitions that can be created on a storage subsystem depending on the premium feature that has been enabled. Benefits of partitioning include ease of management, amortization of costs, scalability, and flexibility.



Configuring Storage Partitions

To configure storage partitions, use the following procedure:

1. Create logical drives on your storage subsystem. See "Creating a Logical Drive" on page 198.

2. Define the host groups and hosts (and their associated host ports) for which you want to define specific logical drive-to-LUN mappings.
3. Grant access to one or more logical drives to a defined host group or host. Do this by defining a specific logical drive-to-LUN mapping using the following procedure:
 - a. Select a defined host group or host in the Topology view of the Storage Partition window.
 - b. Click **Define New Mappings**.
 - c. Select a logical drive that you want the host group to be able to access
 - d. Select the logical unit number (LUN) that you want the host group or host to use to access the logical drive.

Each host group or host is granted a unique view of partitioned storage. A defined host group or host can either see:

 - Logical drives with default logical drive-to-LUN mappings. In this case, the host group or host will be part of the default host group.
 - Logical drives to which it has been granted access through a specific logical drive-to-LUN mapping.

Reconfiguring Storage Partitions Topology

A storage partitions topology is reconfigurable. You can:

- Move a host port
- Replace a host port
- Move a host from one host group into another host group
- Delete a host group, host, or host port
- Rename a host group, host, or host port

Default Host Group

A default host group is a standard node in the Topology view that designates all host groups, host, and host ports that do not have any

specific logical drive-to-LUN mappings, and that share access to any logical drive that were automatically assigned default LUNs by the controller firmware during logical drive creation.

Host Group

A host group is an optional topological element that you define if you want to designate a collection of hosts that will share access to the same logical drives. For more information, see *Defining a Host Group* in the online help.

Host

A host is a computer that is attached to the storage subsystem and accesses various logical drives on the storage subsystem through its host ports. You can define a host as a separate entity or as part of a host group. You can define specific logical drive-to-LUN mappings to an individual host as well as have the host be part of a host group that shares access to one or more logical drives.

Host Port

A host port is a physical connection on the host adapter in a host. When the host adapter has only one physical connection (host port), the terms *host port* and *host adapter* are synonymous.

The host ports are automatically discovered by Fibre Channel Storage Manager. A host port is the actual physical connection that enables a host to gain access to the logical drives in the storage subsystem. Therefore, if you want to define specific logical drive-to-LUN mappings for a particular host and create partitions, you must define its associated host ports.

Initially, all discovered host ports belong to the default host group and have access to any logical drives that were automatically assigned default LUNs by the controller firmware during logical drive creation.

Any host port in the default host group can automatically access the following logical drives:

-
- Logical drives that when created were given default logical drive-to-LUN mappings when they are created.
 - Any logical drives that were created using previous versions of this storage management software. These logical drives are automatically given default logical drive-to-LUN mappings. You can change these mappings using the Change Mapping option.

Defining Logical Drive-to-LUN Mappings

After you have defined your storage partitions topology, you can select a defined host group or host in the Topology view and click the **Define New Mapping** to assign a specific logical drive and logical unit number (logical drive-to-LUN mapping). This designates that only the selected host group or host has access to the particular logical drive through the assigned LUN.

Fibre Channel Storage Manager manages the logical drive-to-LUN mappings. Host port requests for data that are issued for specific LUNs are routed to the appropriate logical drive by the controller firmware.

Only one logical drive-to-LUN mapping is allowed per logical drive. To grant logical drive access to more than one host, you must first group the hosts together in a logical host group. Then, you can grant logical drive access to that host group thereby enabling the hosts in the host group to share access to the logical drive.

Because each host has its own LUN address space, you can use the same LUN in more than one logical drive-to-LUN mapping, if that LUN is available for use by each host that is participating in the mapping.

You can define new logical drive-to-LUN mappings only for logical drives that currently have no mappings (that is, they do not have default mappings or specific mappings). If any of the host groups or hosts have specific logical drive-to-LUN mappings defined, they will no longer appear under the default host group; they will move under the root, storage subsystem, node. Also, they will no longer be able to access logical drives with default logical drive-to-LUN mappings. A host group or host can access either logical drives with default mappings or specific mappings, but not both.

Logical drive-to-LUN mappings are shared between controllers in the storage subsystem.

Reconfiguring your Logical-to-LUN Mappings

A logical drive-to-LUN mapping is reconfigurable. You can:

- Change a logical drive-to-LUN mapping.
- Delete a logical drive-to-LUN mapping.

Changing a Logical Drive-to-LUN Mapping

To change either the assigned logical unit number (LUN) for a logical drive or the associated host group or host for the logical drive, use the following procedure:

1. From the Storage Partition Definitions window, select a single logical drive from the Mappings view. Click **Configure** → **Logical Drive-to-LUN Mappings** → **Change Mapping**. The Change Mapping window opens. The name and unique identifier of the selected logical drive are listed on the window.
2. Change the mapping as follows:
 - **LUN**
Select a single logical unit number (0 through 31) from the LUN drop-down list. By default, the current LUN that is associated with the selected logical drive is displayed. The full range of logical unit numbers is displayed because the same LUN can potentially participate in multiple mappings, because the host has its own LUN address space. However, a logical drive can be mapped only to a single LUN and single host group or host.
 - **Host group or host**
Select a single name from the drop-down list. The drop-down list displays an alphabetical list of all defined host groups or host names in the form Host Group <Host Group Names> or Host <Host Names>. By default, the current host group or host that is associated with the selected logical drive is displayed

-
3. Click **OK** to save the logical drive-to-LUN mapping and close the window.

The logical drive-to-LUN mapping is checked for validity and saved. The Mappings view is updated to reflect the new mapping. The Topology view is also updated to reflect any movement of host groups or hosts from under the default host group if the logical drive-to-LUN mapping was changed to a host group or host that did not previously have any specific logical drive-to-LUN mappings.

- Note:** If you change the assignments for a host group or host that has only one logical drive-to-LUN mapping, the host group or host will move back to the default host group and inherit any default mappings.

If you want to change the last mapping of a host group that also has an associated host with specific mappings, you must delete or move the host mappings first.

Deleting a Host Port, Host, Host Group, or Logical Drive Mapping

You can delete a host port, host, or host group from the Topology view. You can also delete a logical drive-to-LUN mapping. However, you cannot delete the root (storage subsystem) node or the Default Host Group node. If you are deleting the last logical drive-to-LUN mapping for a host group and the host group also has hosts with specific logical drive-to-LUN mappings, make sure that you delete or move those mappings before deleting the last logical drive-to-LUN mapping for the host group.

To delete a node or logical drive mapping, use the following procedure:

1. From the Storage Partition Definitions window, select a host group, host, or host port from the Topology view, or a single logical drive from the Mappings view; then do one of the following:
 - Click **Edit** → **Delete**.
 - Right-click a host group, host, or host port and click **Delete** from the menu.

-
- Press the **Delete** key.
The Confirm Delete window opens.
 - 2. Click **Yes** to delete the node or the logical drive mapping and close the Confirm Delete window.
 - 3. If you selected a host group, host or host port under the Default Host Group node see the following table.

If you selected a host group, host, or host port under the Default Host Group node	Then
Host Group	<p>The host group and its associated hosts and host ports are removed from the Topology view.</p> <p>Any default logical drive-to-LUN mappings that are associated with the default host group are unaffected and remain in the Mappings view.</p> <p>Any host ports that were associated with the host group become undefined. For more information on how to view the undefined host ports, see <i>Showing Undefined Host Ports</i> in the online help.</p>
Host	<p>The host and its associated host ports are removed from the Topology view. Any default logical drive-to-LUN mappings that are associated with the default host group are unaffected and remain in the Mappings view. Also, if the host is part of a host group, the host group is unaffected.</p> <p>Any host ports that were associated with the host become undefined. For information on how to view the undefined host ports, see <i>Showing Undefined Host Ports</i> in the online help.</p>

If you selected a host group, host, or host port under the Default Host Group node	Then
Host port associated with a host group or a host	<p>The host port is removed from the Topology view and is no longer associated with a host group or host. However, any host group or host that is associated with the host port is unaffected.</p> <p>Any default logical drive-to-LUN mapping associated with the default host group are unaffected.</p> <p>For information on how to view the undefined host ports, see <i>Showing Undefined Host Ports</i> in the online help.</p>

If you selected a host group, host or host port that has specific logical drive-to-LUN mappings	Then
Host Group	<p>The host group and its associated hosts and host ports are removed from the Topology view.</p> <p>Any specific logical drive-to-LUN mappings are deleted from the Mappings view and the associated logical drives are available for a new mapping using the Define New Mapping option.</p> <p>Any host ports that were associated with the host group become undefined.</p> <p>For information on how to view the undefined host ports, see <i>Showing Undefined Host Ports</i> in the online help.</p>

If you selected a host group, host or host port that has specific logical drive-to-LUN mappings	Then
Host	<p>The host and its associated host ports are removed from the Topology view. Any specific logical drive-to-LUN mappings are deleted from the Mappings view and the associated logical drives are available for a new mapping using the Define New Mapping option.</p> <p>If the host is part of a host group that has its own specific logical drive-to-LUN mappings, the host group is unaffected. However, if the host is part of a host group that does not have any other mappings, the host group and any other associated hosts or host ports will be moved back to the default host group and inherit any default mappings.</p> <p>Any host ports that were associated with the host become undefined. For information on how to view the undefined host ports, see <i>Showing Undefined Host Ports</i> in the online help.</p>
Host Port	<p>The host port is removed from the Topology view and is no longer associated with a host group or host. However, any host group or host that is associated with the host port is unaffected.</p> <p>Any specific logical drive-to-LUN mappings that are associated with the host group or host are unaffected. However, if this was the last host port, these logical drive-to-LUN mappings will be inaccessible until you redefine associated hosts and host ports.</p>

If you selected a logical drive-to-LUN mapping in the Mappings view	Then
The only logical drive-to-LUN mapping or the last logical drive-to-LUN mapping that is associated with a host group or a host	The logical drive-to-LUN mapping is removed from the Mappings view. The associated logical drive is available for a new mapping using the Define New Mapping option. The host group and/or hosts and host ports that are associated with the mapping are moved back under the default host group in the Topology view and inherit any default mappings. If this mapping is associated with a host group and one or more of its hosts has specific mappings, the deletion will not be allowed until you move or delete the specific mappings of the hosts.
One of many logical drives that are associated with a host group or a host	The logical drive-to-LUN mapping is removed from the Mappings view and is no longer associated with its respective host group or host. The associated logical drive is available for a new mapping using the Define New Mapping option. The host group or host is unaffected.

Moving a Host Port

You can move a host adapter (host port) from one host to another host. If you are replacing a failed host adapter with a new host adapter, do not use this option. Use the Replace Host Port option instead. Use the following procedure when moving a host port:

1. Physically remove the host adapter from the host.
2. Use the Move Host Port option to move the host adapter (host port) to the desired host in the Topology view.
3. Physically insert the host adapter into the new host.

-
4. Restart both hosts so that the changes can be recognized by their operating systems.

To move a host port, use the following procedure:

1. From the Storage Partition Definitions window, select a host port in the Topology view; then click **Configure** → **Topology** → **Move Host Port**.
The Move Host Port window opens. The Host Port name and Host port identifier (worldwide name) are listed on the window.
2. Select the desired host from the **Host name** drop-down list. The list displays an alphabetical list of all defined hosts.
3. Click **OK**.
The changed host port-to-host definition is saved, and the window closes. The host port is removed from the previous host and is now displayed under the new host. The host port inherits the logical drive-to-LUN mappings.
If the host port was the only port that was listed under the previous host, the previous host is still displayed in the Topology view and retains its current logical drive-to-LUN mappings. However, these logical drive-to-LUN mappings will be inaccessible until you redefined at least one associated host port for this host.

Replacing a Host Port

It is extremely important that you use one of the following sequences when replacing the failed host adapter to ensure that the host will retain its current logical drive-to-LUN mappings. If you do *not* have any default logical drive-to-LUN mappings in the default host group, use procedure 1. If you *do* have default logical drive-to-LUN mappings in the default host group, use sequence 2.

Replacement procedure 1:

1. Physically replace the failed host adapter with the new host adapter and restart the host. Use the host start-up utility for the host adapter to obtain the host port identifier (worldwide name) of the new host adapter.

-
2. Use the **Replace Host Port** option to indicate the new host adapter (host port) that replaced your failed host adapter.

Replacement procedure 2:

1. Use the **Replace Host Port** option to indicate the new host adapter (host port) that will replace your failed host adapter. You must obtain the host port identifier (worldwide name) of the new host adapter before physically inserting it in the intended host.
2. Physically replace the failed host adapter with the new host adapter.
3. Restart the host so that the new host adapter is identified by the host operating system.

To replace a host port, use the following procedure:

1. From the Storage Partition Definitions window, select the failed host port from the Topology view; then click **Configure** → **Topology** → **Replace Host Port**.
The Replace Host Port window opens. The current host port identifier (worldwide name) is listed on the window.
2. Define a replacement host port using the **New Identifier** drop-down list. The drop-down list displays all undefined host ports (that is, all host ports that have not been associated with a host) in ascending identifier order. The first undefined host port is displayed by default. You can also type an identifier. The maximum length of the host port identifier cannot exceed 16 hexadecimal characters.
3. Type a host port name in the **Host port name** text box. By default, the text box displays the host port name that is associated with the current identifier. Keep the same name or use a similar name. The host port name is a required field. If the name of the host port duplicates the name of another host port in the storage subsystem, you will receive an error message and a prompt to type a new host port name.
4. Click **OK**.
The changed host port definition is saved and the window closes.

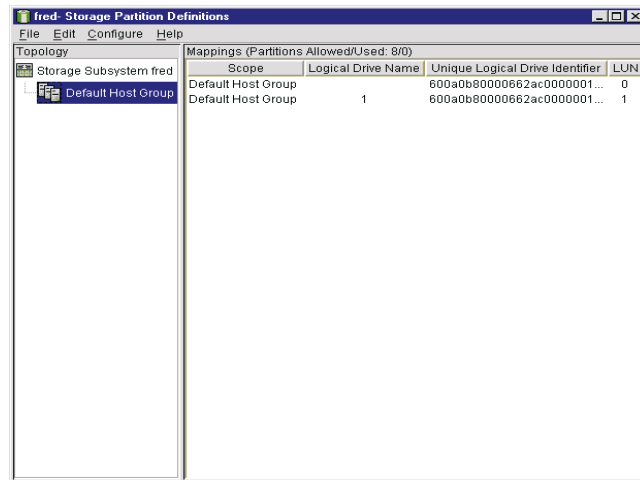
Renaming a Host Group, Host, or Host Port

You can rename a host group, host, or host port. The operation will fail if you choose a name that duplicates the name of another node of the same type. For example, if you are renaming a host that duplicates the name of another host in the Topology view, an error message is displayed and you are prompted to choose another name. The maximum length of the name cannot exceed 30 characters. Any leading and trailing spaces in the name are deleted.


Showing Undefined Host Ports

To view a list of discovered host ports that have not been currently defined (associated with any hosts), use the **Define New Host Port** option.

1. From the Storage Partition Definitions window, click **Default Host Group**.



2. Right-click **Default Host Group** and select the **Show Undefined Host Ports** from the pop-up menu. The Undefined Host Ports window opens.
3. Click **Close** to close the window.



Fuel Gauge Monitor

Fuel Gauge Monitor warns you about conditions which could lead to preventable down time. These conditions involve the power subsystem and the load presented by the system.

One condition occurs when the system load increases to the point where power subsystem specifications are being violated. For example, when too much current is drawn from the power subsystem and utilization is over 100%. This increase may result because of hardware configuration changes, backup devices or hardware failures. Operation above 100% utilization may result in the failure of the power subsystem and may remove the system from service. To avoid this situation, Fuel Gauge Monitor alerts you when maximum power utilization of the system is approaching, when the power subsystem is operating beyond rated specification, and when the system loads comes back down below these thresholds.

A second condition arises when a server that has multiple, pluggable power modules experiences an increase in system load which takes power subsystem utilization above a limit known as the *loss of redundancy* threshold. Below this threshold, power utilization is low and one entire power module is unused. The unused power module is, in essence, a *spare fuel tank* that could be pressed into service in the event of the failure of another power module. However, above this threshold, all available power modules are needed to meet the demands of the server and the failure of any power module places the power subsystem into an over current mode of operation, which may take the system down. Fuel Gauge Monitor alerts you when the system has

entered a state of non-redundant operation and notifies you when a state of redundant operation is restored.

In addition, Fuel Gauge Monitor enables you to review the operation of the power subsystem at any point in time to determine how far from a loss of redundancy or over current situation the server is at that time.

All Fuel Gauge events are forwarded to the IBM Director. To view an event use the IBM Director event log viewer.

Requirements

Fuel Gauge Monitor is supported on IBM systems that contains the following:

- BIOS with Fuel Gauge architecture
- Device Driver support for the Service Processor
- UM Server Extensions installed

Fuel Gauge Console

From the Director console, drag and drop the Fuel Gauge Monitor icon onto one or more systems. The Fuel Gauge Console opens. Power subsystem information is displayed graphically. There are three views available:

- Table View - The selected parameter is arranged in rows and columns.
- Bar Chart - The selected parameter is arranged in a tubular graph.
- Pie Chart - The selected parameter is arranged in a circular graph.

The console also has the following task buttons:

- Close - Use the Close button to exit the Fuel Gauge Monitor console.
- Update - Use the Update button to refresh the data polled from the selected system.
- Help - Use the Help button to view online help.
- Status Bar - Use the Status Bar to view the status of a selected system. If a system cannot be reached or does not support the Fuel Gauge monitor task, the error message appears here.

Information Data

Data received from the selected system is displayed either graphically (Pie chart, Bar chart) or in a tabular format (Table view) by choosing one of the view tabs. The display reports data polled from the system at the time the task is started. Each tabbed View displays the same data. The display received from the selected system falls into these categories.

Note: In the following table, *n* refers to a number.

Parameter	Description
Available Power (N1)	N of power supplies. N will always be 1 or greater.
Failed Power Supplies	N of failed power supplies. For systems that are designed so that any single power module can provide all standby power required by the system. No special Fuel Gauge reporting or alerting is required other than detecting a failure of a power supply in this area. When more than one power supply is available, they share the standby power demand, however any single supply could take over completely in the event of failure.
Used Power Supplies	N of power supplies in use.

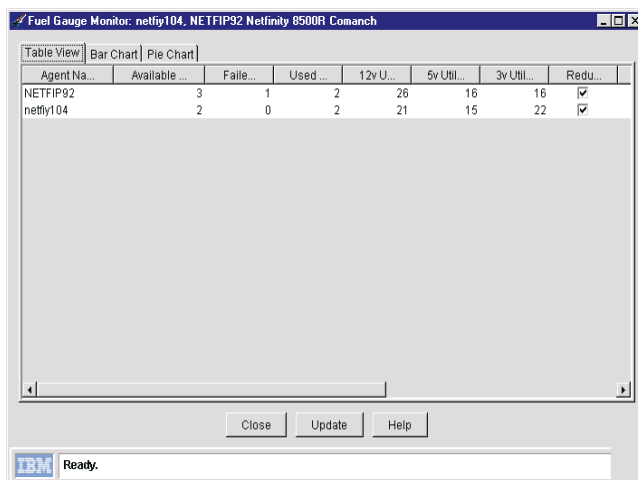
Parameter	Description
12v Utilization(U12v)	<p>12v utilization as a percentage (nnn%). Represents the percentage of 'full load' placed upon the 12VDC output of the power subsystem. Values presented to the service processor (by the power subsystem controller) range from 0 to 100% or more, with any value over 100% representing a condition where the power modules are being operated beyond specifications. The service processor constrains the values to the range of 0 to 100% before presenting them through any other software interface or user interface. U12v is based upon the number of operational power units and therefore will change drastically whenever a power unit fails or becomes available.</p>
5v Utilization(U5v)	<p>5v utilization as a percentage (nnn%). Represents the percentage of 'full load' placed upon 5VDC output of the power subsystem. Values presented to the service processor by the power subsystem range from 0 to 100% or more, with any value over 100% representing a condition where the power modules are being operated beyond specifications. The service processor constrains the value to the range of 0 to 100% before presenting them through any other software interface. U5v is based upon the number of operational power units and therefore will change drastically whenever a power unit fails or becomes available.</p>

Parameter	Description
3v Utilization(U3v)	3v utilization as a percentage (nnn%). Represents the percentage of 'full load' placed upon the 3VDC output of the power subsystem. Values presented to the service processor by the power subsystem range from 0 to 100% or more, with any value of 100% representing a condition where the power modules are being operated beyond specifications. The service processor constrains the values to the range of 0 to 100% before presenting it through any other software interface. The value of U3v is based upon the number of operational power units and therefore will change drastically whenever a power unit fails or becomes available.
Redundancy State(Sr)	This data refers to the enabled/disabled status of redundant power subsystem capabilities. The check box is marked in the Table view.
Low Fuel Threshold (TL)	Threshold is automatically computed by the service processor. The data refers to the percentage of power subsystem output until Low Threshold alert is reached.
Power Supply Capacity	The total amount of power in nnn watts. This is the sum of all power subsystems stated output.

Parameter	Description
Minimum Power Supply Required	The amount of wattage required for an operational system. The integer nnn watts indicates the number of power modules or power supplies required to meet the demand placed on the power subsystem by the current load without causing 'over current' operation of any power unit. The value of this attribute will never exceed the number of operational power units in the system even when the current load is causing 'over current' operation. In the case of an overload, this attribute will report that all available power units are needed and the overload is reported separately. This value is determined by the Service Processor by polling power subsystems instrumentation.
Maximum Available Power	The amount of wattage available to the system as reported by the service processor.
Low Fuel Status	This status is disabled by default. It indicates that power is below the threshold of potential failure. The check box is unmarked in the Table view.
Over Current Status	Power is over the threshold of long term recommendation.

Table View

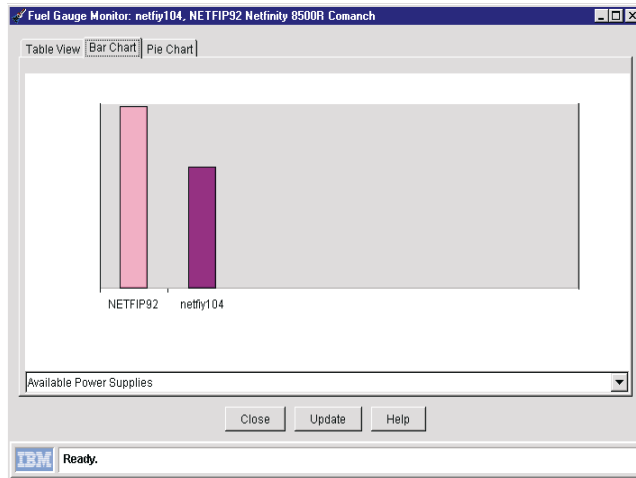
The Table view displays the data that is represented in a tabular format. The first column contains the rack name followed by the rack position, device name, status, and components description. The Table view is the default view of the Fuel Gauge Monitor console. To refresh the data, click **Update**.



Bar Chart

The Bar chart displays data in a vertical bar for each selectable parameter. If multiple systems are selected for the Fuel Gauge Monitor task, each system is represented by a separate bar.

To view information, select the parameter from the drop down list. Each bar represents a system. A single system is displayed with a single bar. Multiple systems display as ratios with separate data.

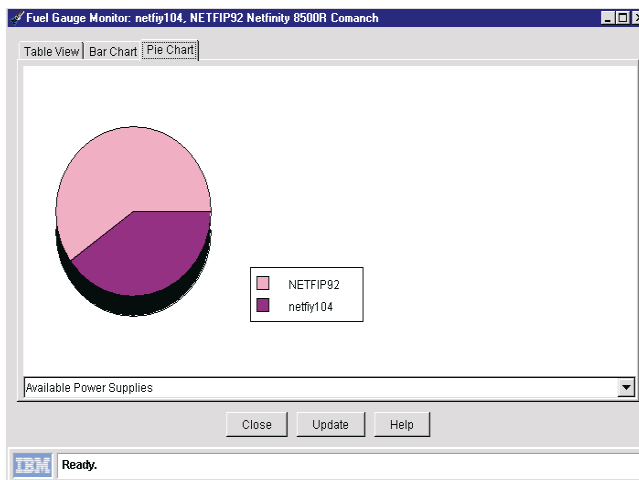



Pie Chart

The Pie chart displays data in a pie chart for each selectable parameter. If multiple systems are selected for the Fuel Gauge Monitor task, each system is represented as a separate slice of the pie. Each system is represented by a color reference labeled in the legend.

To view the data, select a parameter from the drop down list. Each section of the Pie chart represents a system. Use the Hover Help to view the parameters for any section within the Pie chart. A single system is displayed as a full circle.

Note: If the data is returned as a value of zero, the slice color is black or white.





Rack Manager

IBM Director Rack Manager is a flexible, easy-to-install solution for consolidating IBM servers, storage devices, and other standard 19-inch rack equipment.

With Rack Manager, you can group your equipment, enabling you to manage your system resources and to monitor your system functions more efficiently. Centralizing your equipment in integrated rack suites helps to reduce your "real estate" and support costs.

The program provides a realistic picture of nodes and devices within a physical rack for status monitoring and management of racks and components. The devices in the racks have been fully instrumented and are integrated into the IBM Director.

Highlights

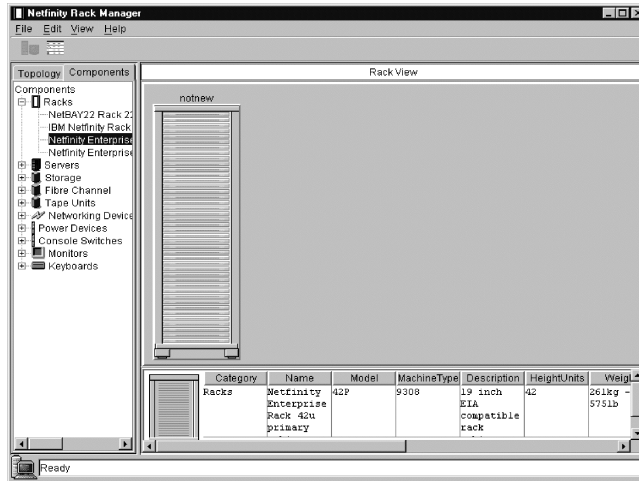
- Efficient, convenient space management. Reduce costs and clutter by consolidating currently installed equipment and new IBM servers. Attractive, well-designed rack enclosures maximize your use of space.
- Effective resource sharing. By installing a console selector switch, you can share a single monitor, keyboard, and mouse among multiple systems.
- Flexible design. You can choose from a variety of rack hardware components to configure your rack for your application.

Rack Manager Console

The Rack Manager console consists of the following components:

- **Menu**- The Menu contains File, Edit and View. Using the menu, you may close a task, delete/add racks, change views or obtain information from the online help.
 - File
 - Exit
 - New Rack
 - Find Rack
 - Edit
 - Delete
 - Associate
 - Disassociate
 - Add
 - View
 - Rack View
 - Table View
 - Help
- **Toolbar** - The Toolbar contains a Rack view icon and a Table view icon. Select an icon to change the Rack Manager workspace to either a tabular view or a graphical representation of a selected rack configuration.
- **Control Pane or Left Pane** - The Control Pane provides a rack topology and component tabs, which are used to configure and manage the rack and it's components.
- **Resource Information Window** - The Components Resource Window contains the properties for each rack and it's components. It appears below the results pane and provides property data on a selected rack or its components.
- **Rack Manager Workspace** - The workspace either identifies rack information as a listing (in the Table view) or as a graphical

representation (in the Rack view) of a rack with managed components. The Resource Information Window is also included in the workspace.



Control Pane

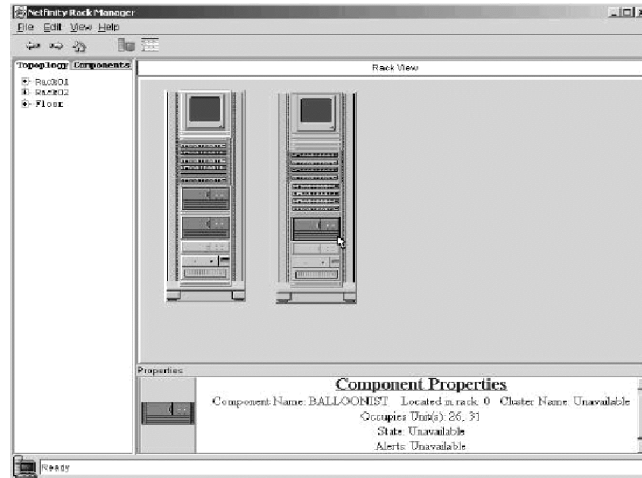
The Control Pane or left Pane consists of a component tree with a rack topology. Use the Control Pane to select and build racks and rack components.

Topology Tab

The Topology tab displays an Explorer tree view of racks and servers. A server appears as a leaf below a rack branch. Racks are organized alphabetically. To display a server, expand or collapse the Explorer tree view. If the server is not a member of a rack, then it is visible as a leaf below the floor branch.

Click on a rack or a server to select it. A left-click will highlight the rack or the server in the rack view and display property data in the Resource Information Window. A right-click displays a pop-up menu for the rack or the server. Use the menu to associate or disassociate the server from

a particular rack. In addition, servers that appear on the floor can be dragged and dropped onto a rack in the Graphic view.




Each rack device that is associated with a managed object is monitored for any changes in status. Rack components that are not associated with a managed object will not be monitored and will always appear in the normal state.

If the rack (or rack managed objects) has not been created, then the Topology list view will place all the servers and devices associated with the targeted managed object or group on the floor. Therefore, the servers appear as a leaf below the Floor division of the Topology tab.

Component Catalog

The Component Catalog provides a list of all components that you can use to configure racks in an Explorer view.

The Explorer view displays a tree with possible rack enclosures and components for configuring the rack view. Each branch describes a category of components belonging to a rack. You can expand the category to view the rack components. For a list of supported rack components, see "Supported Rack Components" on page 252.



In addition to components, the Component Catalog contains component categories, which also contain lists of components. To collapse or expand component categories, click on the plus box or minus box to the left of the category icon.

When you select a component in the Component Catalog by clicking on it, information about the component, including the name and technical specifications, is displayed in the Properties Pane window. There are separate tables for each of the different component categories.

To add a component to a configuration, click on the component and drag and drop the component on to the rack in the Rack View.

For each component, component information is displayed so that you can determine which component to use. Measurements for components are shown in US, metric units, and EIA units.

Rack Manager Workspace

The Rack Manager Workspace displays information about the physical rack. You can view a graphical or table representation of the rack information by selecting the Table view or the Rack view. This dual view of the physical rack will assist you in monitoring and problem determination.

Table View

The Table view displays the data that is represented in a tabular format. The first column contains the rack name followed by the rack position, device name, status, and components description. You can change the view by clicking the Table icon on the toolbar or from the system menu, select **View** → **Table view**.

Rack View

The Rack view displays the physical rack. By default, the Rack view is open when you initially start the Rack Manager. When you select or drop and drag a component, a graphical representation of that component appears in the Rack view. As you move the cursor containing the component into the graphical representation of the rack, a green shadow highlights the physical size of the component as unit's inside the rack. The green shadow appears only if there is legitimate space to support the placement of the rack component. Trying to place a component inside the rack where a green shadow does not appear will cause the Rack Manager to reject the component and place it back into the Floor division or where you selected the component. Also the graphical representation of the component displays the units coordinates (from highest units to the lowest) as you drag inside the rack. If you remove a component from the rack, the graphical representation of the component will no longer appear in the Rack view and instead appear under the Floor division in the Topology tab. To change to the Rack view, click the rack and monitor icon on the toolbar or from the system menu, select **View** → **Rack View**.

Component Properties Window

The Component Properties Window details specifications about a selected rack or selected rack components. These specifications appear when IBM Director discovers the systems. A rack is not discovered by IBM Director but must be created by the Director Administrator.

Selecting a component displays these descriptions in the Component Properties window:

Note: If the component properties being viewed are in a rack there will be one additional property

- Occupied Slots (non-Managed Object) or,
- Assigned Name (Managed Object)

Common to all categories:

- Category
- Name (IBM designation)
- Type
- Model
- Unit Height
- Power (Watts)
- Current @ 120 Volts (amps)
- Current @ 230 Volts (amps)
- Weight
- Height
- Width
- Depth
- Power Cords
- Component

Unique to Servers:

- Maximum Processors
- Maximum Cache
- Memory

In categories other than Servers:

- Viewable Image (Monitors)
- Maximum Resolution (Monitors)
- Number of Keys (Keyboards)
- EIA Capacity (Racks)

-
- Weight Capacity (Racks)
 - Ports (Console Switches)

Some of these values may not appear if the component is not a rack mountable component or if the component does not have the UM Services Agent installed.

Creating and Configuring a Rack

When displaying the Rack View initially, the cluster components will appear on the Floor division of the Topology tab. To create a rack, from the Component tab, select a rack and drag it to the Rack View or from the menu bar, select **File**→**New Rack** or place the cursor in the Rack View area of the Rack Manager workspace and right-click for a pop-up menu. Select New Rack.

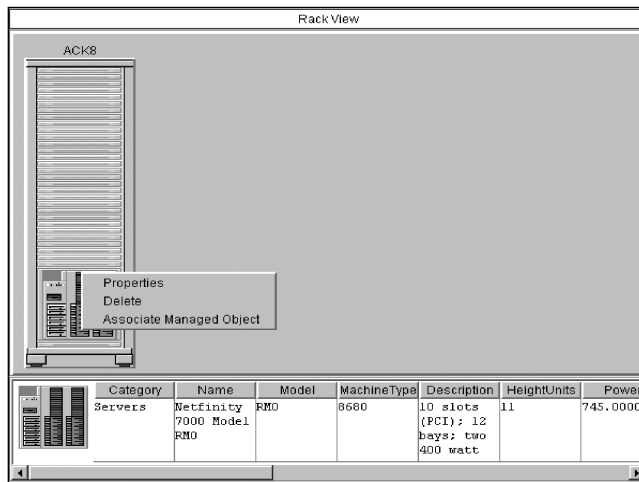
The Rack Properties window opens. Enter the rack name, location, and type. This creates a rack, and a graphical representation appears in the Rack View.



Adding Components to a Rack

From the Topology tab, expand the Floor division. The servers represent managed objects discovered by IBM Director. Drag servers from the Floor division onto the rack in the Rack View. From the Component tab, you can expand a category of components available to be placed in a rack and drag the component onto the rack. Also, a component can be associated to a particular pre-defined IBM component type. However, when you associate a server with a native managed object that does not match the model type in the rack, a pop-up window is displayed as a warning. The server model type is the only validation that is done. To

associate the component to a managed object, right click the component in the rack and select **Associate Managed Object**. When you associate a server with a native managed object that does not match the model type in the rack, a pop-up window is displayed as a warning. The server model type is the only validation that is done.



Removing a Rack or a Rack Component

To remove a Rack, right click the rack component in the Rack View and select **Delete** from the pop up menu. This will remove the rack from the Rack View and all associations to the rack. In addition, the rack-managed object will no longer appear on the IBM Director console.

To remove a Rack component, right click the component and select **Delete** from the pop-up menu. This deletes the rack component and any association with a managed object. If there is a server associated with a managed object, it will appear on the floor.

Associate Component

When a component is not rack mountable, only the machine type and model are listed. Also, data collection may not have been performed successfully and a **Inventory Data Not Available** statement appears in the Component Properties window.

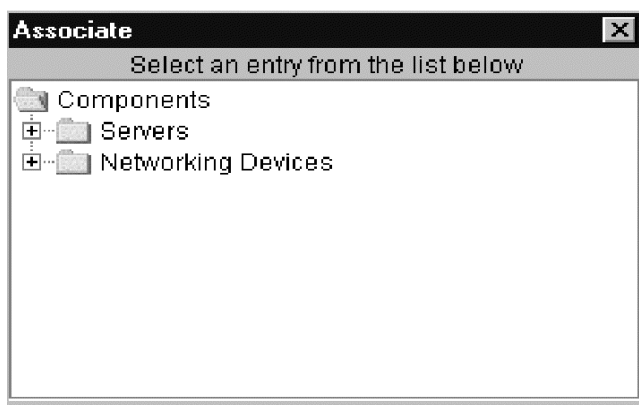
To provide more information about a particular component that includes additional properties, you can associate the component with a pre-defined IBM component from the items in the Components tab.

To associate a component to an IBM pre-defined component, use the following procedure:

1. From the Floor division in the Topology tab, right click the component.
2. Select the **Associate**.
3. Select the appropriate component match from either the Servers or Networking Devices folders.

You can also associate a component by selecting the component from the Floor division and dragging it to the Rack View. The Associate pop-up menu appears. Also you can change the association of a managed component by first disassociating the component and then associating with a different selection.





Note: If you associate a component with a managed object that does not match the model type in the rack, a warning dialog box appears. The server model type is the only validation. You have the option to ignore the warning.

Disassociate Component

You may want to cancel a component association (disassociate):

- You have used an incorrect component association
- Component inventory collection has been performed
- The association is no longer valid

If the component is in a rack, delete the component from the rack before you begin the following procedure:

1. From the rack in the Rack view, right click the component.
2. Select **Delete**.
The component will appear in alphabetical order as a leaf below the floor division of the Topology tab.

To disassociate a component, use the following procedure:

1. From the Topology tab under the floor division, right click the component.
2. Select **Disassociate**.

The component information in the Resource Information window reverts to the information created at the time of discovery through IBM Director.

Supported Rack Components

Rack Manager supports the following Rack components:

- Rack Enclosures
 - IBM Netfinity Rack 42u cabinet Model 900
 - Netfinity Enterprise Rack 42u expansion cabinet Model 42X
 - Netfinity Enterprise Rack 42u primary cabinet Model 42P
 - NetBay22 Rack 22u cabinet Model 200
- Storage Devices
 - EXP10 Storage Expansion Unit Model 1Rx
 - EXP15 Storage Expansion Unit Model 2RU
 - Netfinity Exp 200 with optional power Model 1RU
 - 3518 Enterprise Expansion Enclosure Model 001
 - 3519 SCSI Storage Unit Model R01
 - 7133 SSA Storage Unit Model 020
- Fibre Channel Devices
 - Fibre Channel Hub Model 1RU
 - Fibre Channel RAID Unit with Failsafe RAID Model 1RU
- Tape Unit Devices
 - 3447 DLT Tape Library Model 106
 - 3449 8 mm Tape Library Model 356
 - 3502 Tape Autoloader Model R14
 - NetMEDIA Storage Expansion EL Model 001
 - Magstar MP 3570 Tape Library Model B21
 - Magstar MP 3570 Tape Library Model B22
 - Magstar MP 3570 Tape Library Model C21
 - Magstar MP 3570 Tape Library Model C22
- Network Devices
 - 8230 Token-Ring Controlled Access Unit Model 04X
 - 2210 Nways Multiprotocol Router - Model 12x
 - 2210 Nways Multiprotocol Router - Model x4x

8270 Nways LAN Switch with redundant power supply Model 800
8271 Nways Ethernet Switch Model 712
8271 Nways Ethernet Switch Model E12
8271 Nways Ethernet Switch Model E24
8271 Nways Ethernet Switch Model F12
8271 Nways Ethernet Switch Model F24
8285 Nways ATM Workgroup Switch Base Model 00B
8237 Ethernet Stackable Hub Base 10BaseT Model 00x
3529 Netfinity SP Switch Model 1RY
8285 Nways ATM Workgroup Switch Expansion Model 00E

- Power Devices
 - High Voltage PDU Model 450
 - Low Voltage PDU Model 666
 - APC Smart-1400 UPS Model 674
 - APC Smart-3000 UPS Model 676
 - APC Smart-UPS 5000 RMB Model 861
- Console Switches
 - 4 port Console Switch Model 542
 - 8 port Console Switch Model 445
- Monitors
 - 9 in Monochrome Display Model E01
 - G42 Color Monitor Model xxx
 - G51 Color Monitor Model xxx
 - G52 Color Monitor Model xxx
 - G54 Color Monitor Model 4AN
 - G72 Color Monitor Model xxx
 - G74 Color Monitor Model 4AN
 - IBM T55a Flat panel monitor Model AG1
 - P50 Color Monitor Model xxx
 - P70 Color Monitor Model xxx
- Keyboards
 - Space saver Keyboard Model 644
 - Standard Keyboard Model 861
- Servers
 - Netfinity 4000R
 - Netfinity 4000R Model 11Y

Netfinity 4000R Model 21Y
Netfinity 4000R Model 22Y
Netfinity 4000R Model 31Y
Netfinity 4000R Model 41Y
Netfinity 4000R Model 42Y
Netfinity 4000R Model 43Y
Netfinity 4000R Model 44Y
Netfinity 4000R Model 51Y
Netfinity 4000R Model 61Y
Netfinity 4000R Model 62Y
Netfinity 4000R Model 63Y
Netfinity 4000R Model 64Y
Netfinity 4000R Model 1RY
Netfinity 4500R Model 2RY
Netfinity 4500R Model 3RY
Netfinity 4500R Model 4RY
Netfinity 4500R Model 5RY
Netfinity 5000
Netfinity 5000 Model 1SY
Netfinity 5000 Model 2SY
Netfinity 5000 Model 3RY
Netfinity 5000 Model 4RY
Netfinity 5000 Model 5RY
Netfinity 5000 Model 6RY
Netfinity 5000 Model 7RY
Netfinity 5000 Model 8RY
Netfinity 5100 Model 1RY
Netfinity 5100 Model 2RY
Netfinity 5100 Model 3RY
Netfinity 5100 Model 4RY
Netfinity 5100 Model 5RY
Netfinity 5500
Netfinity 5500 M10
Netfinity 5500 M10 Model 1RY
Netfinity 5500 M10 Model 2RY
Netfinity 5500 M10 Model 3RY
Netfinity 5500 M10 Model 4RY
Netfinity 5500 M10 Model 5RY

Netfinity 5500 M20
Netfinity 5500 M20 Model 3RY
Netfinity 5500 M20 Model 4RY
Netfinity 5500 M20 Model 5RY
Netfinity 5500 M20 Model 6RY
Netfinity 5500 Model 1RU
Netfinity 5500 Model 4RU
Netfinity 5500 Model 4SU
Netfinity 5500 Model 5RU
Netfinity 5500 Model 5SU
Netfinity 5500 Model 6RU
Netfinity 5500 Model 6SU
Netfinity 5500 Model 7SU
Netfinity 5600
Netfinity 5600 Model 1RY
Netfinity 5600 Model 2RY
Netfinity 5600 Model 3RY
Netfinity 5600 Model 4RY
Netfinity 5600 Model 5RY
Netfinity 5600 Model 6RY
Netfinity 5600 Model 7RY
Netfinity 6000R Model 1RY
Netfinity 6000R Model 2RY
Netfinity 7000 M10
Netfinity 7000 M10 Model 11Y
Netfinity 7000 M10 Model 1RU
Netfinity 7000 M10 Model 1SY
Netfinity 7000 M10 Model 21Y
Netfinity 7000 M10 Model 2RU
Netfinity 7000 M10 Model 2SY
Netfinity 7000 M10 Model 3RU
Netfinity 7000 M10 Model 3SY
Netfinity 7000 M10 Model 4RU
Netfinity 7000 M10 Model 5RU
Netfinity 7000 M10 Model 6RY
Netfinity 7000 M10 Model 7RY
Netfinity 7000 M10 Model 8RY
Netfinity 7000 Model RH0

Netfinity 7000 Model RM0
Netfinity 7100 Model 1RY
Netfinity 7100 Model 2RY
Netfinity 7100 Model 3RY
Netfinity 7100 Model 4RY
Netfinity 7600 Model 1RY
Netfinity 7600 Model 2RY
Netfinity 7600 Model 3RY
Netfinity 8500R
Netfinity 8500R Model 4RY
Netfinity 8500R Model 5RY
Netfinity 8500R Model 6RY
Netfinity 8500R Model 7RY
Netfinity 8500R Model 8RY
PC Server 325 Model 1RY
PC Server 325 Model 2RY
PC Server 325 Model RB0
PC Server 325 Model RS0
PC Server 325 Model xxx
PC Server 330 Model xxx
xSeries 200 Model 10X
xSeries 200 Model 11X
xSeries 200 Model 12X
xSeries 200 Model 13X
xSeries 200 Model 40X
xSeries 200 Model 41X
xSeries 200 Model 42X
xSeries 200 Model 50X
xSeries 200 Model 51X
xSeries 200 Model 52X
xSeries 220 Model 21X
xSeries 220 Model 22X
xSeries 220 Model 2AX
xSeries 220 Model 31X
xSeries 220 Model 32X
xSeries 220 Model 3AX
xSeries 220 Model 41X
xSeries 220 Model 42X

xSeries 220 Model 4AX
xSeries 230 Model 6RY
xSeries 240 Model 8RY
xSeries 340 Model 6RY

9

ServeRAID Manager

This chapter provides the information needed to start and use the ServeRAID Manager program. You can use ServeRAID Manager to easily configure and monitor your ServeRAID controllers.

The information in this chapter is a high level explanation of the ServeRAID Manager program and its capabilities. For instructions on specific processes using the ServeRAID Manager program, refer to the ServeRAID Manager online help.

Note: ServeRAID Manager works with systems using Microsoft Windows NT, Windows 95, Windows 98.

Using the ServeRAID Manager Program Interface

The graphical interface in ServeRAID Manager makes it easy for you to create, delete, change, view and monitor your ServeRAID configuration.

Before you begin, review the following illustration to become familiar with the layout of the ServeRAID Manager program windows.






Viewing the Menu Bar

The menu bar is a set of menu names that are located directly below the title bar. It provides commands from drop-down menus. The menu bar options include; File, View, Remote, Actions, and Help.

Viewing the Toolbar

The toolbar is a set of buttons that are located directly below the menu bar. These buttons serve as shortcuts for many frequently used commands. When you first view the ServeRAID Manager window, some commands are disabled and are enabled only after you access certain menu commands.

The toolbar includes the following commands.

Icon	Command
	Create arrays
	Configure for clustering
	Scan for new or removed ready drives
	Silence repeating alarm
	Help

Viewing the Expandable Tree

The ServeRAID Manager interface provides an expandable tree view of your ServeRAID subsystem.

You will perform most of your ServeRAID configuration and maintenance tasks by first selecting the ServeRAID controller, array, logical drive, hot-spare drive, or physical drive objects from this Main Tree.

Viewing the Main Panel

The ServeRAID Manager interface provides specific device information or configuration instructions. When ServeRAID Manager is in the Information mode and you select an object from the Main Tree, detailed Information about the object appears in this panel. When ServeRAID Manager is in the Configuration mode, the instructions needed to configure your ServeRAID subsystem appear in this panel.

Viewing the Event Viewer

The ServeRAID Manager interface provides advisory and progressive-status information and messages during the ServeRAID configuration process and while monitoring systems with ServeRAID controllers. Each message appears with a host name from where the event originated, a time stamp, a date stamp, and an icon that classifies the severity of the event. The event icons are:

- Information: An "i" inside a blue circle
- Warning: A "!" inside a yellow triangle
- Fatal: An "x" inside a red circle

Warning messages identify potential data-loss situations, and Fatal messages inform you when a failure has occurred. All Fatal messages will launch an audible alarm.

Viewing the Status Bar

The Status bar provides three types of information in a resizable pane. The panes contain the following information:

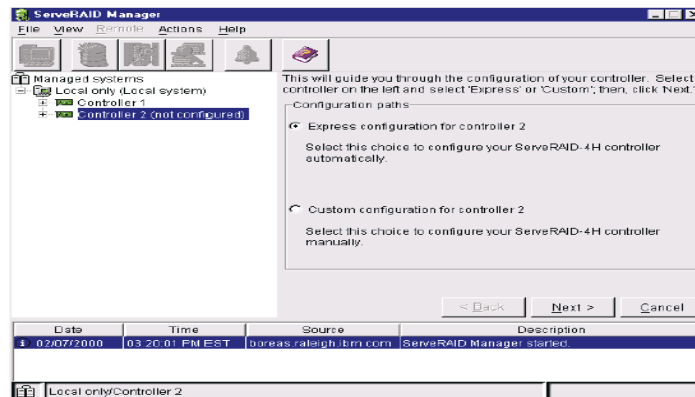
- The left pane displays the managed system status, which is either **No problems detected on any system** or **Problems detected on one or more systems**.
- The center pane displays the current tree path.
- The right panel displays a progress meter with the label of the currently selected system and the process that is taking place at the current time.

Using the ServeRAID Manager Program

The ServeRaid Manager program runs in two modes: Configuration mode and information mode. To start the ServeRAID Manager, drag the task to the system. ServeRAID Manager starts in Information mode.

In Information mode, you can view and modify the existing device and configuration information for your ServeRaid subsystem. When this mode is active, you can use the functions available from the menu and tool bars to customize settings for your ServeRaid controllers.

In Configuration mode, you can create disk arrays, create logical drives, and define your hot-spare drives. When this mode is active, a set of functions is available from the menu and tool bars.



Using the Configuration Mode

You can use the Configuration mode to create up to eight independent disk arrays for each ServeRAID controller. The Configuration mode provides two configuration paths: Express configuration and Custom configuration.

Learning about Express Configuration

Express configuration provides a quick and easy path for you to automatically configure your ServeRAID controller. This choice creates the most efficient ServeRAID configuration based on the number and capacity of the Ready drives available in your system. If four or more

Ready drives of the same capacity are available, this choice also will define a *hot-spare* drive for the ServeRAID controller. A hot-spare drive is a physical drive that is defined for automatic use when a similar drive fails.

The Express configuration choice groups all Ready drives of the same capacity into one or more disk arrays and defines one logical drive for each array. This choice defines the size of the logical drive, based on the amount of free space available, and it assigns the highest RAID level possible, based on the number of physical drives available.

For example, if your server contains one 1024 MB Ready drive, two 2150 MB Ready drives, and four 4300 MB Ready drives, Express Configuration will create three arrays and one hot-spare drive such as the following:

Array A: The total capacity of this array is 1024 MB (1 x 1024 MB), and it contains one 1024 MB RAID level-0 logical drive.

Array B: The total capacity of this array is 4300 MB (2 x 2150 MB), and it contains one 2150 MB RAID level-1 logical drive.

Array C: The total capacity of this array is 12900 MB (3 x 4300 MB), and it contains one 8600 MB RAID level-5 logical drive.

Hot Spare: Express Configuration defines one of the four 4300 MB drives as a hot-spare drive.

Notes:

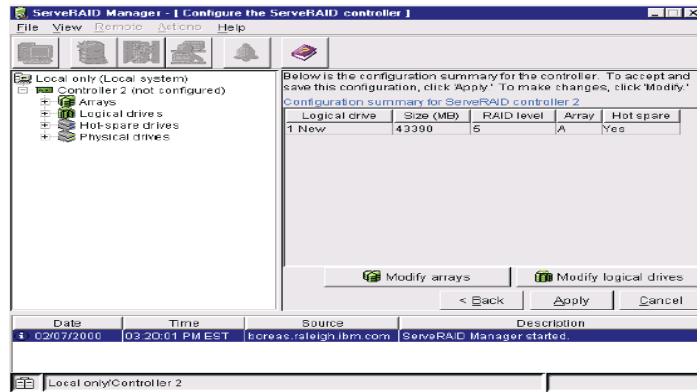
- When there are four or more Ready drives of the same capacity, Express configuration groups defines one drive as a hot spare and the remaining drives into one array (as in Array C).
- A hot-spare drive must be of equal or greater capacity than the drive that it is intended to replace. In this configuration, the 4300 MB drive can replace any failed drives in Array B or Array C. The drive in Array A cannot be replaced because it is RAID level-0.

Using Express Configuration

To use the Express configuration path, do the following:

Note: If the Express configuration mode/Custom configuration mode radio buttons do not appear on the right screen, click the tool bar.

1. Click the ServeRAID controller that you want to configure in the Main Tree.
2. If it is not already selected, click the **Express configuration** radio button.
3. Click **Next**.
The Configuration Summary window opens.
4. Review the information that appears on the **Configuration summary** screen.



Notes:

- Some operating systems have size limitations for logical drives. Before you save the configuration, you might want to verify that the size of the logical drive is appropriate for your operating system. See your operating-system documentation for more detailed information.
- To change the configuration, click **Modify arrays** or **Modify logical drives**.

-
5. Click **Apply** to accept and save the configuration. This saves the configuration in the ServeRAID controller and in the physical drives.

Learning about Custom Configuration

The Custom configuration path enables you to manually configure your ServeRAID subsystem. Using this path, you can select the drives that you want to include in each array, define the number and size of the logical drives for each array, and select the drives that you want to use as hot spares.

Before you select the Custom configuration path, consider the following:

- Each ServeRAID controller supports a maximum of eight arrays.
- Each ServeRAID controller supports a maximum of eight logical drives.

When you create an array, you group physical drives into one storage area. You can define this storage area as a *single* logical drive, or you can subdivide it into several logical drives. Each logical drive appears to the operating system as a single physical drive.

If you have only one array, you can define it as a single logical drive, or you can divide it into several logical drives. Typically, the first logical drive defined on the first ServeRAID controller found by the basic input/output system (BIOS) during startup will be your startup (boot) drive.

If you have two or more arrays, each array can be one logical drive, or you can divide each array into multiple logical drives, as long as the total number of logical drives for all of the arrays does not exceed eight.

Note: Independent of the RAID logical drives, most operating systems allow you to partition the logical drives further.

- The optimal way to create arrays is to use physical drives that have the same capacity.

Physical drive capacities influence the way you create arrays. Drives in an array can be of different capacities (1 GB, or 2 GB, for

example), but RAID controllers treat them as if they all have the capacity of the *smallest* disk drive.

For example, if you group three 1 GB drives and one 2 GB drive into an array, the total capacity of the array is 1 GB times 4, or 4 GB, not the 5 GB physically available. Similarly, if you group three 2 GB drives and one 1 GB drive into an array, the total capacity of that array is 4 GB, not the 7 GB physically available.

- A hot-spare drive is a disk drive that is defined for automatic use in the event of a drive failure. The hot-spare drive must be of equal or greater capacity than the drive that it is intended to replace. If a physical drive fails and it is part of a RAID level-1, RAID level-1E, RAID level-5, or RAID level-5E logical drive, the ServeRAID controller automatically starts to rebuild the data on the hot-spare drive.
- If you are using the ServeRAID-3L:
 - Eight physical drives are supported in an array if the stripe-unit size is set to 32 KB or 64 KB.
 - Sixteen physical drives are supported in an array if the stripe-unit size is set to 8 KB or 16 KB.

If you are using the ServeRAID-3H or ServeRAID-3HB, sixteen physical drives are supported in an array if the stripe-unit size is set to 32 KB or 64 KB.

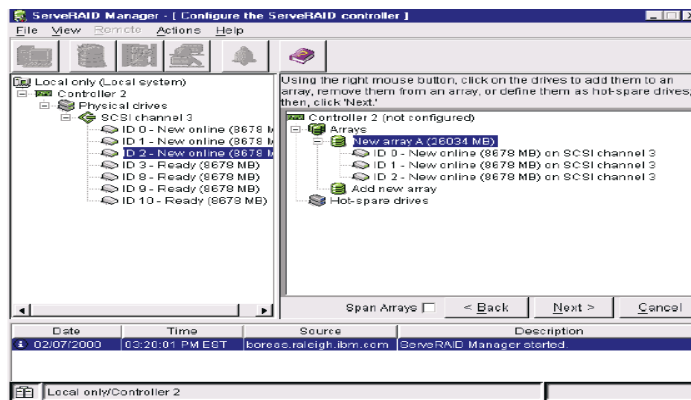
Using Custom Configuration

To use the Custom configuration path, do the following:

Note: If the **Express configuration mode/Custom configuration mode** radio buttons do not appear on the right screen, click on the tool bar.

1. Click the ServeRAID controller that you want to configure in the Main Tree.
2. Click the **Custom configuration** radio button.
3. Click **Next**. The screen similar to the following opens:
4. Do one of the following:

- Using the right mouse button, click on the drive or SCSI channel icons in the Main Tree to select the drives that you want to add to your arrays, delete from your arrays, or define as hot-spare drives; then, select a choice from the pop-up list.
 - Drag the drive or SCSI channel icon from the Main Tree and drop it on the Array or Hot-spare drive icon in the Main Panel on the right. If you change your mind, you can drag the icons back to the Main Tree to remove them from the configuration.
5. After you select the physical drives for your arrays and define your hot-spare drives, click **Next**.

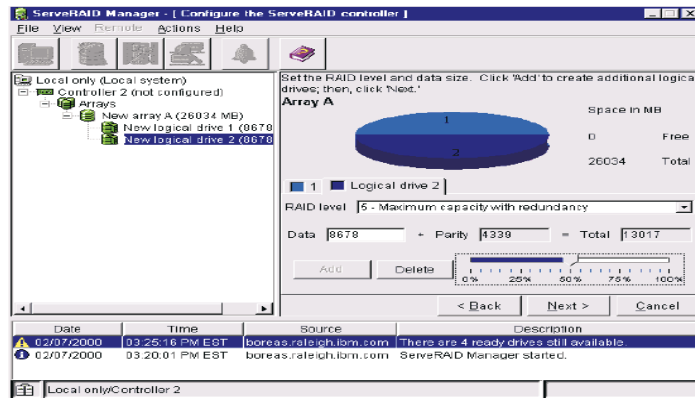


If you change your mind, you can remove a specific drive from a newly defined array or delete an entire newly defined array. To do this, click **Back**, use the right mouse button to click on the specific drive or Array icon in the Main Panel on the right, and then select **Remove from new array** or **Delete new array**.

6. Select a RAID level for the logical drive from the RAID menu in the Main Panel.
7. If you do not want to use the maximum size for the logical drive, move the slide bar in the Main Panel from right to left to allocate data and parity space for the logical drive, or you can type in the size in the data field.

Notes:

- You can define from one to eight logical drives.
 - Some operating systems have size limitations for logical drives. Before you save the configuration, you might want to verify that the size of the logical drive is appropriate for your operating system. See your operating-system documentation for more detailed information.
 - Typically, the first logical drive defined on the first ServeRAID controller found by system BIOS during startup will be your startup (boot) drive.
8. If free space is available and you want to define another logical drive, click the **Add logical drive** tab in the Main Pane.



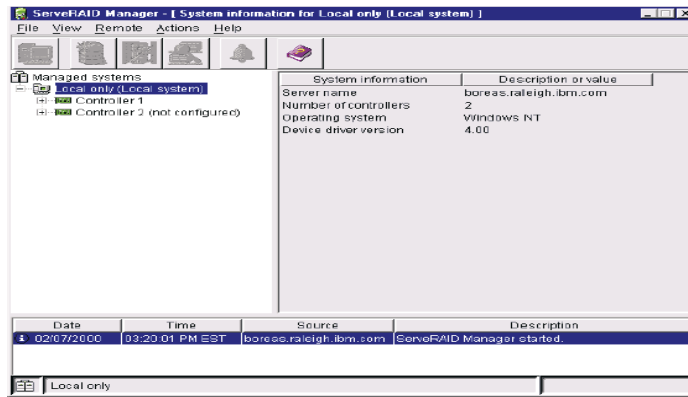
9. Repeat steps 6, 7, and 8 for each logical drive that you want to define in this array; then, continue with step 10.
10. If you are defining multiple arrays, click on the next New array tree node; then, repeat steps 6, 7, 8, and 9. When you are finished, go to step 11.
11. Click **Next**.
12. Review the information that appears on the **Configuration summary** dialog.

Note: To change the configuration, click **Modify arrays** or **Modify logical drives**.

13. Click **Apply** to accept and save the configuration. This saves the configuration in the ServeRAID controller and on the physical drives.
14. If you have multiple controllers, repeat steps 1 to 13 for each controller.

Using the Information Mode

You can use the Information mode to view the configuration information of ServeRAID controllers, arrays, logical drives, hot-spare drives, and physical drives that make up your ServeRAID subsystem. You also can use this mode to change some of the ServeRAID controller settings. Click an item in the Main Tree and select the **Actions** menu to display available actions on that item.



To view the current settings for the device and to view configuration information, do the following:

1. Click the plus (+) box next to the object in the Main Tree to expand that portion of the tree.
2. Click the icon for the server, ServeRAID controller, array, logical drive, hot-spare drive, or physical drive to view its current settings.

Detailed information about the selected device will appear in the Main Pane on the right.

Modifying the ServeRAID Controller Configuration

You can use the ServeRAID Manager program to configure and modify ServeRAID controllers. This section provides information about the following ServeRAID actions:

- Change RAID levels
- Increase free space
- Increase logical drive space

Modifying a Configuration

You can modify a configuration to change RAID levels, increase free space, or increase logical drive space. When you use these actions, you dynamically change the current logical drive structure.

To modify a configuration, you must have at least two logical drives available: one free logical drive and one source logical drive that is in the okay state. During the procedure, the ServeRAID controller changes the state of the free logical drive to system; then, temporarily uses the system drive to perform the migration. When the migration procedure completes, the ServeRAID controller changes the state of the system drive back to free.

While modifying a Logical Drive, consider the following:

- The Main Tree displays only physical drives that are at least the size of the smallest physical drive in the selected array. When adding to an existing array, you cannot add physical drives that are smaller than the smallest drive currently in the array.
- When adding to an existing array, do not add physical drives that are larger than the physical drives currently in the array. If you add physical drives that are larger than the physical drives currently in the array, you will not be able to use all the space on the new physical drives.
- Existing physical drives in the selected array are labeled Online. Newly added physical drives are labeled New online. You cannot remove online drives from an existing array.

-
- If you want to remove a new online drive from the array, drag it back to the Main Tree.

Changing RAID Levels

You can change the RAID levels of currently defined logical drives. To use the Change RAID level, all of the logical drives within the array must be the same RAID level.

The ServeRAID controllers support changing RAID levels as follows:

- Change two or more drives from RAID level-0 to RAID level-5 by adding one physical drive.
- Change two drives from RAID level-1 to RAID level-5 by adding one physical drive.
- Change from RAID level-5 to RAID level-0 by removing one physical drive.
- Change from RAID level-5E to RAID level-5.

To change the RAID level of an existing array, use the following procedure:

1. Click an array from the Main Tree (that is, the left tree).
2. Select **Logical drive migration** from the Actions menu.
3. Select from the available Change RAID levels on the menu.

If you select one of the following, another tree appears on the right. Go to step 4 on page 272.

- Change RAID level from RAID 0 to RAID 5
- Change RAID level from RAID 1 to RAID 5

Note: If the array already contains the maximum number of physical drives (based on the firmware level and the stripe-unit size), the two previous Change RAID levels actions are not available.

If you select one of the following, go to step 6 on page 272.

- Change RAID level from RAID 5 to RAID 0
- Change RAID level from RAID 5E to RAID 5

-
4. Add to the selected array by dragging one ready drive from the Main Tree to the array in the right tree.
 5. After you add the physical drives to your arrays, click **Next** to review your new logical drive configuration.

Note: If you change your mind and want to modify your configuration, you can click **Back** to return to this array tree window.

6. Review the configuration in the right pane. It describes how the configuration will affect the logical drives and free space before you click **Apply**.

Note: You can expand the tree for another view of the configuration.

7. Click **Apply** to approve the new configuration. The ServeRAID Manager program automatically changes the RAID levels of the logical drives in the array.

If you do not want to approve the logical drive configuration, click **Cancel** and return to the system management information window.

Note: Because changing RAID levels is a lengthy process, no other actions on the controller can be performed until the migration is complete. The ServeRAID Manager program displays a progress indicator in the status bar. You can use the logical drive during the change RAID levels process. You can also remove the system power. When you restore the power, the process restarts where it left off.

8. Click **Yes** when the confirm pop-up window opens.

Increasing Free Space

Using this feature you can add one, two, or three physical drives to an existing array so you can create another logical drive in the array.

When you create an array, logical drives are striped across all physical drives in that array. To create more free space in an existing array, define additional physical drives to add to the array. Then, the ServeRAID Manager program migrates the logical drives such that the data is spread across the existing and new physical drive

To increase free space, do the following:

1. Click an array from the Main Tree (that is, the left tree).
2. Select **Logical drive migration** from the Actions menu.
3. Select **Increase free space** from the menu.
4. Add to the selected array by dragging ready drives from the Main Tree to the array in the right tree.

Note: You can add up to three physical drives.

5. After you add the physical drives to your arrays, click **Next** to review your new logical drive configuration.
6. Review the configuration in the right panel. It describes how the configuration will affect the logical drives and free space when you click **Apply**. If you want to modify the configuration, click **Back** to return to the array tree window.

Note: You can expand the tree for another view of the configuration.

7. Click **Apply** to approve the new configuration. The ServeRAID Manager program automatically increases the free space.

If you do not want to approve the logical drive configuration, click **Cancel** and return to the system management information window.

8. Click **Yes** when the confirm pop-up window opens.

Note: Increasing free space is a lengthy process. You cannot perform any other actions on the affected controller until the process is complete. The ServeRAID Manager displays a progress indicator in the status bar. You can use the logical drive during the increase free space process. You can also remove the system power. When you restore the power, the increase free space process restarts where it left off.

9. When the increase free space process completes, you can create new logical drives in this array.

Increasing Logical Drive Space

Using this feature you can add one, two, or three physical drives to an existing array so you can expand the size of the logical drives in the array. When increasing logical drive space, the ServeRAID Manager program migrates the logical drives such that the logical drives gain additional space, much like adding paper to a notebook.

When you create an array, logical drives are striped across all of the physical drives in that array. To increase the size of the logical drives in an existing array, define additional physical drives to add to the array. Then, the ServeRAID Manager program migrates the logical drives such that the data is spread across the existing and new physical drives.

To increase logical drive space, do the following:

1. Click an array from the Main Tree (that is, the left tree).
2. Select **Logical drive migration** from the Actions menu.
3. Select **Increase logical drive space** from the menu.
4. Add to the selected array by dragging ready drives from the Main Tree to the array in the right tree.

Note: You can add up to three physical drives.

5. After you add the physical drives to your arrays, click **Next** to review your new logical drive configuration.
6. Review the configuration in the right pane. It describes how the configuration will affect the logical drives and free space when you click **Apply**. If you want to modify the configuration, click **Back** to return to the array tree window.

Note: You can expand the tree for another view of the configuration.

7. Click **Apply** to approve the new configuration. The ServeRAID Manager program automatically increases the logical drive space. If you do not want to approve the logical drive configuration, click **Cancel** and return to the system management information window.
8. Click **Yes** when the confirm pop-up window opens.

Note: Increasing logical drive space is a lengthy process. You cannot perform any other actions on the affected controller

until the process is complete. The ServeRAID Manager displays a progress indicator in the status bar. You can use the logical drive during the increase logical drive space process. You can also remove the system power. When you restore the power, the increase logical drive space process restarts where it left off.

Setting User Preferences

Select **User preferences** from the File menu. A window opens that you can use to specify the following settings:

- Initialization settings
- Alarm settings

Initialization Settings

Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive. The default setting is checked. When this mode is enabled, the ServeRAID Manager program automatically initializes each new logical drive and prevents access to the manual initialization function. To initialize settings, do the following:

1. Click the **Initialization settings** tab.
2. Click the check box to enable or disable automatic initialization of logical drives.
3. Click **OK**.

Note: If you uncheck (that is, disable) automatic initialization, the Manager will not initialize each logical drive (by writing zeros to the first 1024 sectors of the logical drive) when it is created. You must still initialize all newly defined logical drives before storing data.

Alarm Settings

All Warning events and all Error events cause an audible alarm to sound every five minutes, notifying you of the event. To adjust the alarm interval in the Alarm settings window, do the following:

-
1. Click the **Alarm settings** tab.
 2. Click the check box to enable or disable the repeating alarm.
Note: If you disable the alarm, you will not hear an audible alarm when you receive problem events.
 3. If you have enabled the alarm, you can adjust the time interval (in seconds) you want between each alarm. The default is 300 seconds (that is, 5 minutes).
 4. If you have enabled the alarm, you can adjust the length of time (in beeps) you want the alarm to continue. The default is 3 beeps.
 5. Click **OK**.
 6. Restart the Manager for these settings to take effect.

Silencing the Repeating Alarm

Silence the repeating alarm turns off the alarm for the current failure without disabling the feature.

- To silence the alarm, click on the tool bar.

Understanding Drive States

This section provides descriptions of the physical and logical drive states.

Physical Drive State Descriptions

The following table provides descriptions of the valid physical drive states.

DRIVE STATE	MEANING
Defunct	A physical drive in the Online, Hot-Spare, or Rebuild state has become defunct. It does not respond to commands, which means that the ServeRAID controller cannot communicate properly with the drive.

DRIVE STATE	MEANING
Hot Spare	A Hot-Spare drive is a physical drive that is defined for automatic use when a similar drive fails.
Online	The drive is Online. It is functioning properly and is part of an array.
Rebuilding	The drive is being rebuilt.
Ready	The ServeRAID controller recognizes a Ready drive as being available for definition.
Standby Hot Spare	A Standby Hot Spare is a Hot-Spare drive that the Spare ServeRAID controller has spun down. If an Online drive becomes Defunct and no suitable Hot-Spare drive is available, a Standby Hot-Spare drive of the appropriate size automatically spins up, and enters the Rebuild state.

Logical Drive State Descriptions

The following table provides descriptions of the valid logical drive states.

DRIVE STATE	MEANING
Blocked	<p>During a rebuild operation, the ServeRAID controller sets the state of any RAID level-0 logical drives associated with a failed array to the Blocked state; then, it reconstructs the data that was stored in RAID level-1, RAID level-1E, Raid level-5 and RAID level-5E logical drives.</p> <p>After the Rebuild operation completes, you can unblock the RAID level-0 logical drives, and access them once again. However, the logical drive might contain damaged data. You must either re-create, reinstall, or restore the data from the most recent backup disk or tape to the RAID level-0 logical drive.</p>
Critical Migrating	A logical drive in the critical state that is undergoing a logical drive migration (LDM).

DRIVE STATE	MEANING
Critical System	The ServeRAID controller uses this reserved state during a logical drive migration (LDM) when the logical drive is in the critical state.
Critical	A RAID level-1, RAID level-1E, RAID level-5, or RAID level-5E logical drive that contains a defunct physical drive is in the critical state. A critical logical drive is accessible, despite a physical drive failure.
Migrating	The logical drive is undergoing a logical drive migration; that is, a change in RAID levels, a change in logical drive size, an increase in free space, or a RAID level-5E compression or decompression.
Offline	<p>The logical drive is offline and not accessible. This state occurs when one of the following is true.</p> <ul style="list-style-type: none"> ■ One or more physical drives in a RAID level-0 logical drive is defunct. ■ Two or more physical drives in a RAID level-1, RAID level-1E, or level-5 logical drive are defunct. ■ Three or more drives in a RAID level-5E logical drive are defunct.
Okay	The logical drive is okay. It is in a good, functional state.
System	The ServeRAID controller uses this reserved state during logical drive migration (LDM).

ServeRAID Events in Director

The ServeRAID Manager tool adds the following events to the event builder log:

- Battery-Backup events: Failed and Need Replacement.
- Configuration Change events: Adapter, Array, Array (BIOS Mode, Failover), Import, Logical Drive, Merge Migration, Read Ahead Cache, Rebuild Rate, Removed, Replaced, Reset To, Scan, SCSI, Set To (Strip Size, Successful), Unattended Mode, Write Back, Write Cache, Write Through.

-
- Enclosure
 - Logical Drive
 - Physical Drive
 - State events: Added, Failed, Failover, Not Found, Replaced, Working.

10

Software Rejuvenation

The IBM Software Rejuvenation program is used to increase the reliability of "stand-alone" servers and nodes in a cluster. This is achieved through scheduled software rejuvenations (restarts) on each node. These rejuvenations can be planned for a specific time or time period (such as *every other month on the 1st Sunday* or *once, today at 5:15 p.m.*). In addition, there is the option to plan "no" rejuvenations during a specified period. With predictive rejuvenation, the program can predict when a resource requires a rejuvenation. This provides high availability by limiting the number and frequency of outages while ensuring that peak time availability is not compromised.

To start the Software Rejuvenation program from the IBM Director main window, you must do one of the following:

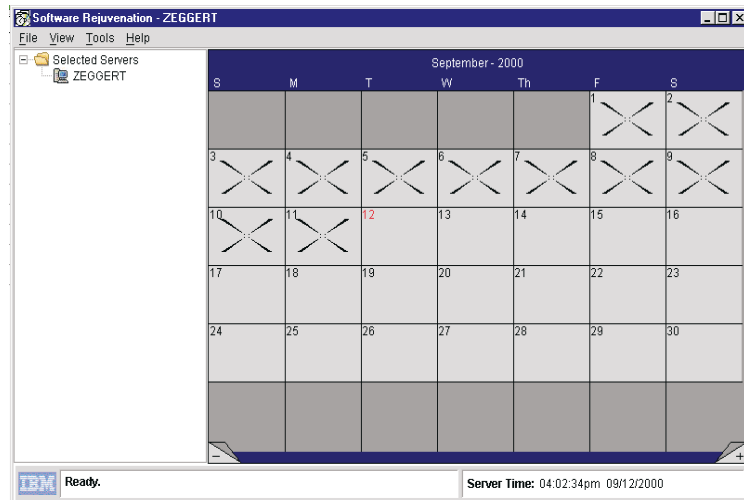
- Drag-and-drop the Software Rejuvenation task icon on to a cluster or group of nodes (IBM Availability Extensions for MSCS or MSCS).
- Drag-and-drop a cluster or one or more nodes onto the Software Rejuvenation task icon.
- Right-click a cluster name or one or more nodes names in the IBM Director console, then click Software Rejuvenation from the menu.

The Software Rejuvenation Window

The Software Rejuvenation window contains two panels. The left panel contains an Explorer tree view with server folders. These folders

expand to display clusters and nodes. The right panel contains a calendar with rejuvenation dates. The Software Rejuvenation window also contains the following components:

- Title bar
- Calendar
- Server Time



Title Bar

The title bar is located at the top of the window and displays the name of the program and the name of the selected cluster or nodes.

Calendar

Use the calendar to schedule and to edit node rejuvenations. The calendar also provides a monthly overview of scheduled node rejuvenations.

The calendar opens to the current month and year, with the current date highlighted. The lower corners of the calendar displays a minus and a plus tab. Click the minus tab (-) to revert the calendar to the previous month. Click the plus tab (+) to advance the calendar one month.

When a node is scheduled for rejuvenation, a node icon displays on every day for which the node is scheduled. If multiple nodes are

scheduled for rejuvenation on the same day, the node icons are shown cascaded on the specified date.

Server Time

The time and date is listed at the bottom of the Software Rejuvenation window. This is the current time and date of the server.

Note: The server time and the console time may be different.

Managing Software Rejuvenations

You can use the IBM Software Rejuvenation program to selectively restart nodes in a cluster or one or more stand-alone servers or nodes. In restarting a node, you rejuvenate or refresh software resources. This program not only restarts a node in a selected cluster, but also allows you to schedule restarts, which includes scheduling multiple nodes on various dates.

The IBM Software Rejuvenation program recognizes and can rejuvenate single MSCS clusters and one or more stand-alone nodes.

Scheduling Nodes for Software Rejuvenation

Using the IBM Software Rejuvenation program you can schedule nodes for rejuvenation in several modes. The following sections describe these ways.

Note: The Software Rejuvenation program uses the IBM Director console time to validate and duplicate rejuvenation schedules, and uses the server system clock to execute schedules.

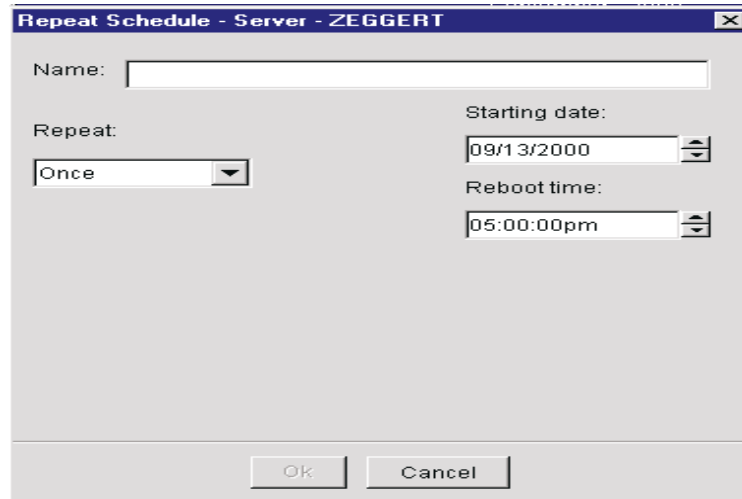
The Software rejuvenation must be enabled in order to schedule a rejuvenation. In addition, for Prediction, the agent must be installed and configured on individual nodes and the nodes of a cluster.

Scheduling a Rejuvenation for a Single Node

To schedule rejuvenation for a single node, use the following procedure:

-
1. Select the node from the Explorer view in the Software Rejuvenation window.
 2. Drag-and-drop the node onto the calendar date of the scheduled rejuvenation.

The Repeat Schedule window opens.



The screenshot shows a dialog box titled "Repeat Schedule - Server - ZEGGERT". It has a "Name:" text input field. Below it is a "Repeat:" dropdown menu currently showing "Once". To the right, there are two more dropdown menus: "Starting date:" showing "09/13/2000" and "Reboot time:" showing "05:00:00pm". At the bottom of the dialog are "Ok" and "Cancel" buttons.

3. Enter the Name of the Schedule.
The name is user defined.
4. From the Repeat drop-down list, select the preferred schedule type for the node.
5. In the Starting Date drop-down list, specify the date you want the rejuvenation to start.
6. From the Reboot Time drop-down list, select the rejuvenation start time.
Note: The number of days set in the Min. Reboot Interval field will override any repeat scheduling made through the calendar.
7. Click **OK**.
The node icon appears on the calendar.
8. Press **Alt-F** and **Alt-S**(File and Save) to accept the changes.

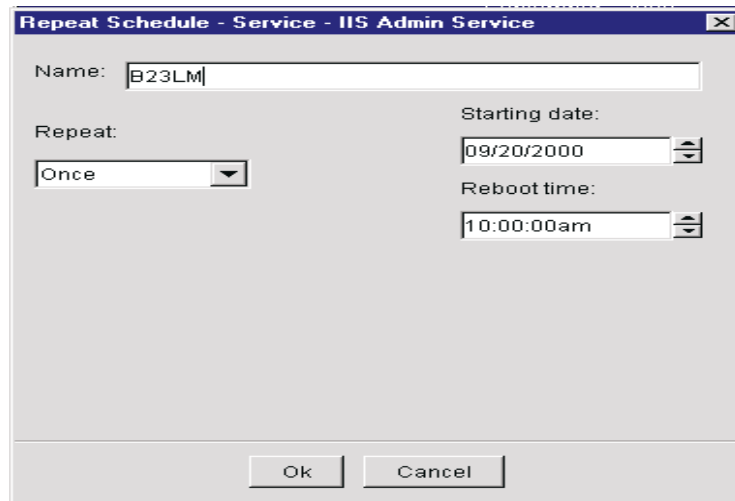
For more information on setting software rejuvenation parameters, see "Rejuvenation Options" on page 291.

Scheduling a Single Rejuvenation

To rejuvenate a node as a single instance, use the following procedure:

1. Select the node (icon) from the Explorer view in the Software Rejuvenation window.
2. Drag-and-drop the node onto the calendar date of the first scheduled rejuvenation.

The Repeat Schedule window opens.



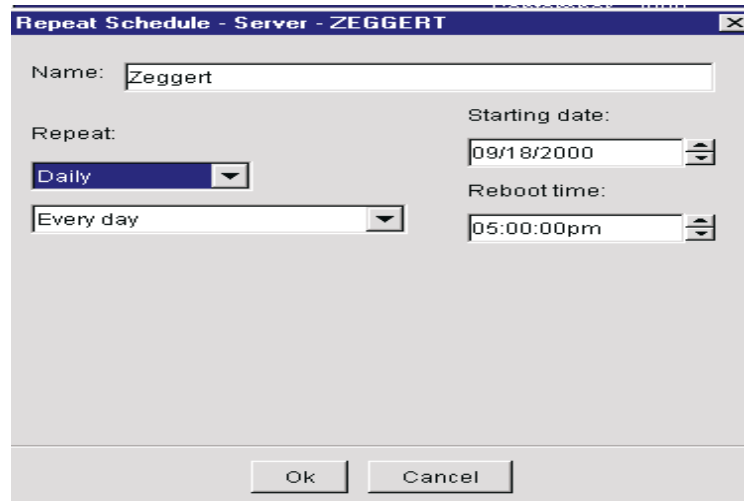
3. Enter the User-defined name.
4. From the Repeat drop-down, select **Once** as the schedule type for the node.
5. In the Starting date drop-down list, select the date you want the rejuvenation to start.
6. In the Reboot time drop-down list, specify the time you want the rejuvenation to start. Time values are in 15-minute intervals.
7. Click **OK** to schedule the rejuvenation for this node.
8. Press **Alt-F** and **Alt-S** to accept changes.

Scheduling a Daily Rejuvenation

To rejuvenate a node daily, use the following procedure:

1. Select the node from the Explorer view in the Software Rejuvenation window.
2. Drag-and drop the node on to the calendar date of the first scheduled rejuvenation.

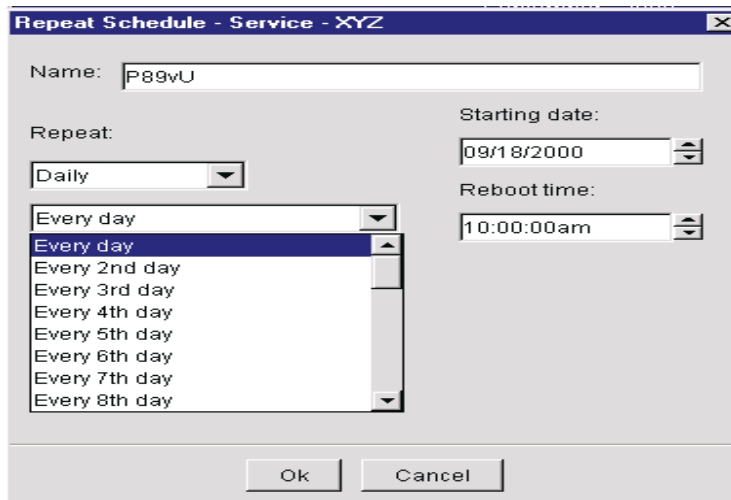
The Repeat Schedule window opens.



The screenshot shows a dialog box titled "Repeat Schedule - Server - ZEGGERT". It contains the following fields and controls:

- Name:** A text input field containing "Zeggert".
- Repeat:** A dropdown menu with "Daily" selected.
- Starting date:** A date input field containing "09/18/2000".
- Reboot time:** A time input field containing "05:00:00pm".
- Frequency:** A dropdown menu with "Every day" selected.
- Buttons:** "Ok" and "Cancel" buttons at the bottom.

3. Enter the User-defined name.
4. From the Repeat drop-down list, select **Daily**.



5. A drop-down list appears under **Daily**. From the drop-down list, select the rejuvenation frequency for the node.
If you select **Every day**, the node is rejuvenated each day; while if you select **Every 2nd day**, the node is rejuvenated every other day, and so on.
6. In the **Starting date** drop-down list, specify the date you want the rejuvenation to start.
7. From the **Reboot time** drop-down list, select the rejuvenation time. Time is measured in increments of 15-minute intervals.
8. Click **OK**.
9. Press **Alt-F** and **Alt-S** to accept the changes.

Scheduling a Weekly Rejuvenation

To rejuvenate a node as a weekly occurrence, use the following procedure:

1. Select the node from the Explorer view in the Software Rejuvenation window.
2. Drag-and-drop the node on to the calendar date of the first scheduled rejuvenation.

The Repeat Schedule window opens.

The screenshot shows a dialog box titled "Repeat Schedule - Service - XYZ". It has a "Name" field containing "P89vU". Below it is a "Repeat" dropdown menu set to "Weekly". Underneath "Weekly" is another dropdown menu set to "Every", with a list of days (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday) visible below it. To the right, there is a "Starting date" field with "09/18/2000" and a "Reboot time" field with "10:00:00am". At the bottom are "Ok" and "Cancel" buttons.

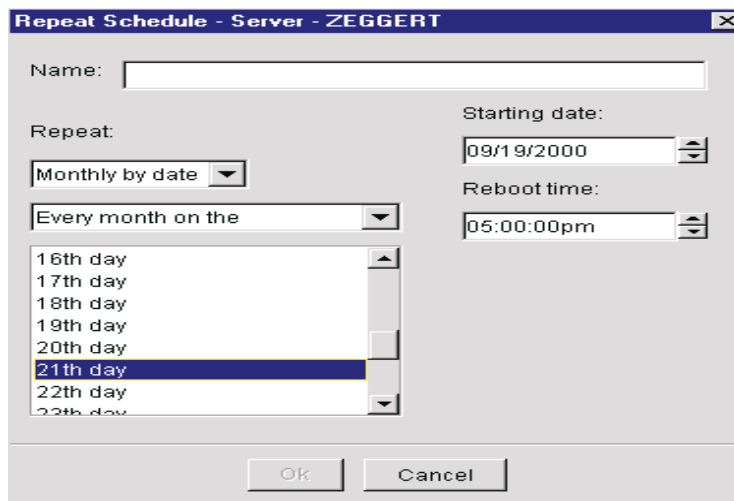
3. Enter a User-defined name.
4. From the **Repeat** drop-down list, select **Weekly**.
5. A drop-down list appears under **Weekly**.
From the drop-down, select a day of the week for rejuvenation.
You may select two or more days of the week for the rejuvenation
Selecting **Every** will rejuvenate the node each selected day of the week.
6. In the **Starting date** drop-down list, select the rejuvenation start date.
7. In the **Reboot** time drop-down list, select the rejuvenation time.
Time is displayed in 15-minute intervals.
8. Click **OK**.
9. Press **Alt-F** and **Alt-S** to accept the changes.

Scheduling a Monthly Rejuvenation by Date

To rejuvenate a node as on a specific date each month, use the following procedure:

1. Select the node from the Explorer view in the Software Rejuvenation window.
2. Drag-and-drop the node on to the calendar date of the first scheduled rejuvenation.

The Repeat Schedule window opens.



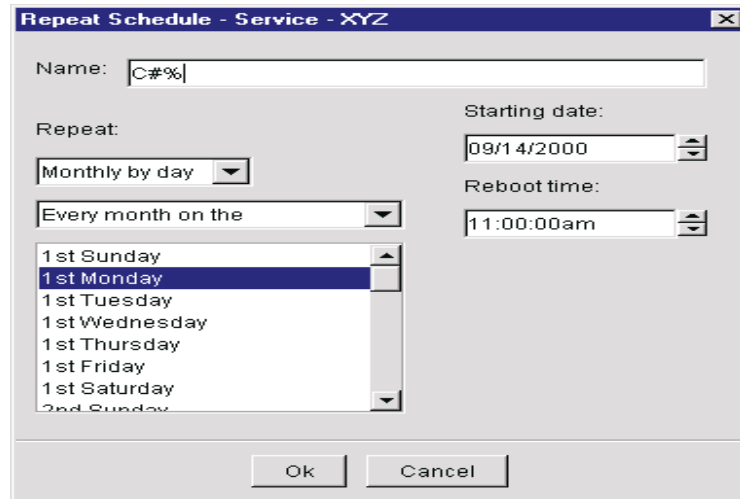
3. From the **Repeat** drop-down list, select **Monthly by date**.
4. A drop-down list appears under **Monthly by date**. From the drop-down list, select the day of the month for the node rejuvenation.
5. In the **Starting date** drop-down list, select a node rejuvenation date.
6. From the **Reboot time** drop-down list, select a rejuvenation time. Time is displayed in 15-minute intervals.
7. Click **OK**.
8. Press **Alt-F** and **Alt-S** to accept changes.

Scheduling a Monthly Rejuvenation by Day

To rejuvenate a node on specific days each month, use the following procedure:

-
1. Select the node from the Explorer view in the Software Rejuvenation window.
 2. Drag-and-drop the node on to the calendar date of the first scheduled rejuvenation.

The Repeat Schedule window opens.



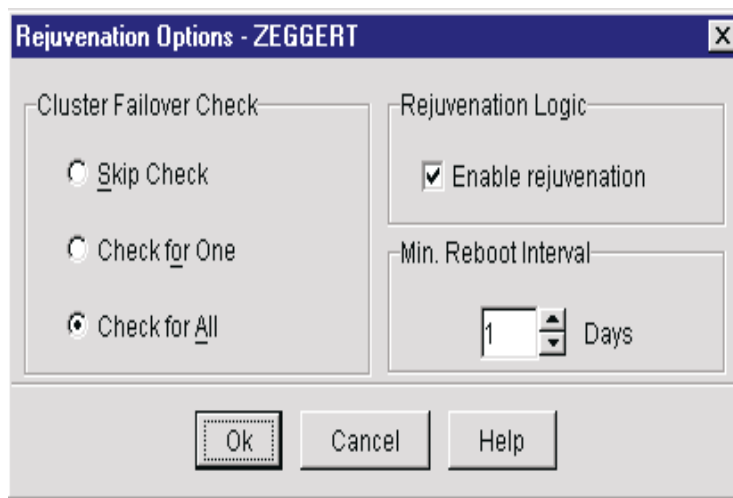
3. Enter a User-defined Name.
4. From the **Repeat** drop-down list, select **Monthly by day**.
New fields appear in the current window.
5. A new drop-down list appears under **Monthly by day**. Select the day of the month for node rejuvenation.
6. From the **Starting date** drop-down list, select the start date for the node rejuvenation.
7. From the **Reboot time** drop-down list, select the start time for the rejuvenation. Time is displayed in 15-minute intervals.
8. Click **OK**.
9. The **Skipping Some Dates** pop-up box may appear. Click **OK**.
10. Press **Alt-F** and **Alt-S** to accept changes.

Rejuvenation Options

Use the Rejuvenation Options window to set parameters for rejuvenation operations on a cluster level. To set parameters for a cluster failover check, cluster rejuvenation logic (errors), and a minimum restart interval, do the following:

1. From the Software Rejuvenation window, select **Tools → Rejuvenation Options**.

The Rejuvenation Options window opens.



2. The **Cluster Failover Check** option defines the failover rules for a node in the selected cluster. Rejuvenation occurs if the selected value matches the state of the cluster. To set the **Cluster Failover Check** option, use one of the following procedures:
 - Click **Skip Check** if you do not want the rejuvenation program to check for other nodes in the cluster. Selecting this value allows a rejuvenation to always occur.
 - Click **Check for One** if you want the rejuvenation program to check for another node in the cluster. If the program cannot find another node in the selected cluster, the selected node will not be rejuvenated.

-
- Click **Check for All** if you want the rejuvenation program to check the availability of all nodes in the cluster. If a single node in the cluster does not respond to the cluster, the selected node will not be rejuvenated.
3. To set the rejuvenation logic for the cluster, select the **Enable Rejuvenation** check box to enable error checking on the cluster.
 4. In the **Min. Reboot Interval** drop-down list, select the number of days that must elapse before a node in the selected cluster can be rejuvenated.

For example, if 3 days is selected, a node in the cluster will not be rejuvenated again in a period of less than 3 days after the previous rejuvenation. Selecting 0 days allows a node to be rejuvenated at any time without checking for the date of a previous rejuvenation.
 5. Click **OK**.
 6. Press **Alt-F** and **Alt-Save** to accept changes.

Editing Scheduled Rejuvenations

Use the IBM Cluster Software Rejuvenation program to edit scheduled node rejuvenations at the cluster level and the node level.

Editing a Scheduled Rejuvenation at the Node Level

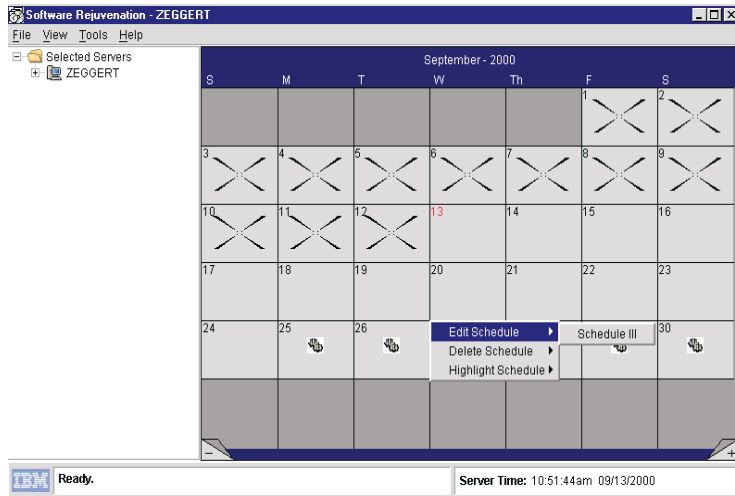
To change the date, time, or frequency of scheduled node rejuvenations, use the following procedure:

1. From the IBM Director Console, drag-and-drop the Cluster (the owner of the nodes you want to reschedule) onto the IBM Software Rejuvenation task icon.

The **Software Rejuvenation** window opens.

The calendar displays the current month and year. The current date is highlighted. A node icon is displayed on calendar day for which a node is scheduled for rejuvenation. If multiple nodes are scheduled for rejuvenation on the same day, the node icons are shown cascaded on the specified date.

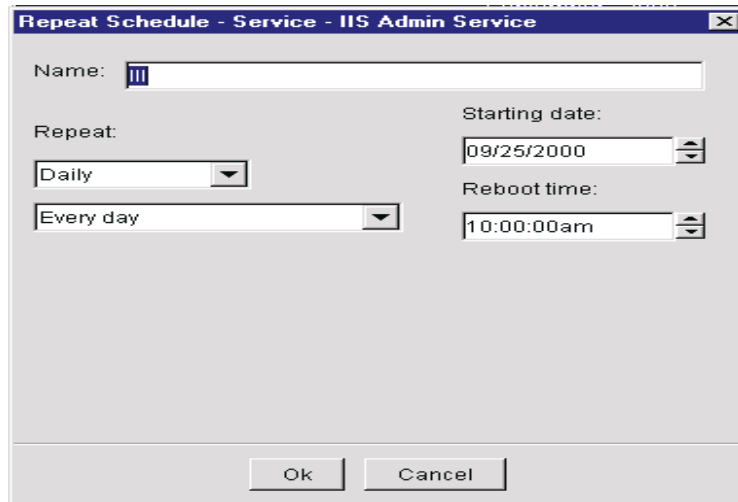
2. In the Calendar, right-click the node you want to edit.
The Edit Schedule menu appears.



3. Click **Edit schedule** → **Schedule<node name>**.

Note: Click the plus sign to advance the calendar a month, and click the minus sign to revert the calendar to the previous month.

The Repeat Schedule window opens.



4. From the **Repeat** Schedule window, edit the rejuvenation schedule settings.
5. Click **OK**.
6. Press **Alt-F** and **Alt-S** to accept changes.

Editing a Rejuvenation at the Cluster Level

To change the failover check, cluster rejuvenation logic (error), and minimum reboot interval for rejuvenation operations for a cluster, do the following:

1. From the IBM Director Console, drag-and-drop the cluster you want to edit onto the IBM Software Rejuvenation task icon.

The **Software Rejuvenation** window opens. The calendar displays the current month and year. The current date is highlighted.

2. From the Software Rejuvenation window, click **Tools**→ **Rejuvenation Option**.

The Rejuvenation Options window opens.

3. To set the Cluster Failover option, use the following procedure:

- Click **Skip Check** if you do not want the rejuvenation program to check for other nodes in the cluster. Selecting this value allows a rejuvenation to always occur.
- Click **Check for One** if you want the rejuvenation program to check for another node in the cluster. If the program cannot find another node in the selected cluster, the selected node will not be rejuvenated.
- Click **Check for All** if you want the rejuvenation program to check the availability for all nodes in the cluster. If a single node in the cluster does not respond to the cluster, the selected node will not be rejuvenated

Note:

The Cluster Failover Check option defines the failover rules for a node in the selected cluster. Rejuvenation occurs if the selected value matches the state of the cluster.

4. To set the rejuvenation logic for the cluster, select the **Enable Rejuvenation** check box to enable error checking on the cluster.
5. From the Min. Reboot Interval drop-down list, set the number of days for a node rejuvenation.

For example, if 3 days is selected, a node in this cluster will not be rejuvenated again in a period of less than 3 days after the previous rejuvenation. Selecting 0 days will allow a node to be rejuvenated at any time without checking for the date of a previous rejuvenation.

6. Click **OK**.
7. Press **Alt+F** and **Alt+S** to accept changes.

Removing a Node from the Rejuvenation Schedule


If a node is scheduled for rejuvenation on more than one date, removing the node from one date on the calendar removes it from all other dates and times on the calendar. To remove a node from the rejuvenation schedule, use the following procedure:

-
1. From the IBM Director Console, drag-and-drop the cluster (the owner of the node you want to remove) onto the IBM Software Rejuvenation task icon.

The Software Rejuvenation window opens.

The calendar displays the current month and year. The current date is highlighted. A node icon is displayed on every day that a node is scheduled for rejuvenation. If multiple nodes are scheduled for rejuvenation on the same day, the node icons are shown cascaded on the specified date.

2. In the Calendar, right-click the node you want to remove. The **Edit** menu appears.
3. Select **Delete** → **Schedule<Node Name>**. The **Verify Remove** box appears (If the node rejuvenation that you are removing from the calendar is scheduled to occur only once, you will not receive this message window).



4. Select **Yes**.
5. Press **Alt+F** and **Alt+S** to accept changes.

Using Keyboard Commands

The following table describes the keyboard commands that you can use to perform basic tasks.

Keyboard command	Use this shortcut to:
TAB / SHIFT-TAB	To move a selected date forward or backwards through the calendar interface.
CTRL+C	Copy a node.

Keyboard command	Use this shortcut to:
CTRL+V	Paste a copied node on a selected calendar date.
CTRL+E	Access the Repeat Schedule window and to edit the schedule.
CTRL+D	Delete a scheduled rejuvenation on a selected date.
CTRL+H	Select (highlight) a schedule.

Predictive Rejuvenation

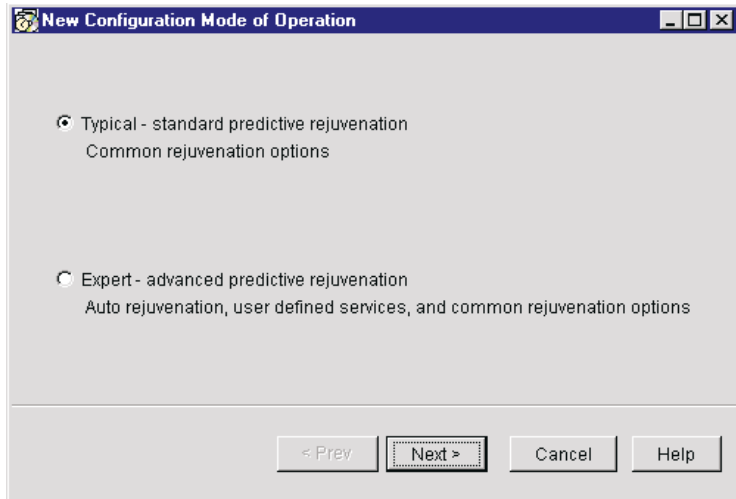
Predictive Rejuvenation defines the way in which PFA (Predictive Failure Analysis) logic is processed. It can be used with one or more nodes. In order to use predictive rejuvenation, you must first configure the nodes. From the Software Rejuvenation window, select **Tools**→**Prediction**→**Configuration Wizard** to start the configuration wizard.

Using the Predictive Rejuvenation Wizard

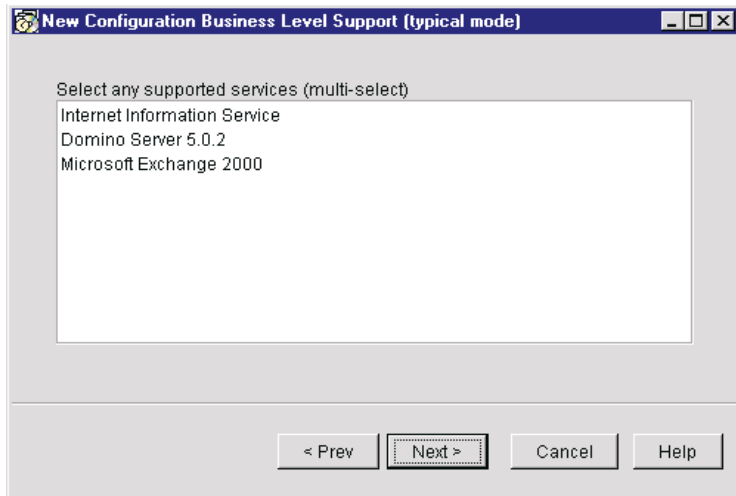
There are two Predictive Rejuvenation modes of operation: **Typical** and **Expert**. **Typical** predictive mode handles the common rejuvenation options. The **Expert** predictive mode handles the advanced predictive rejuvenation options such as, auto rejuvenation, user defined application, as well as common rejuvenation options.

To configure a **Typical Mode of Operation**, use the following procedure:

1. From the Software Rejuvenation window, select **Tools**→**Prediction**→**Configuration Wizard** to start the configuration wizard.
The **Configure Mode of Operation** window opens.



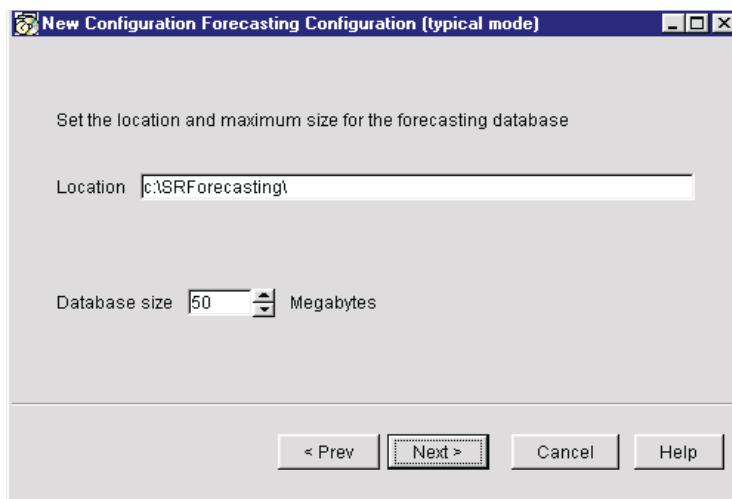
2. From the Configuration - Mode of Operation window, select **Typical**.
3. Click **Next**
The Configure Business Level Support window opens.



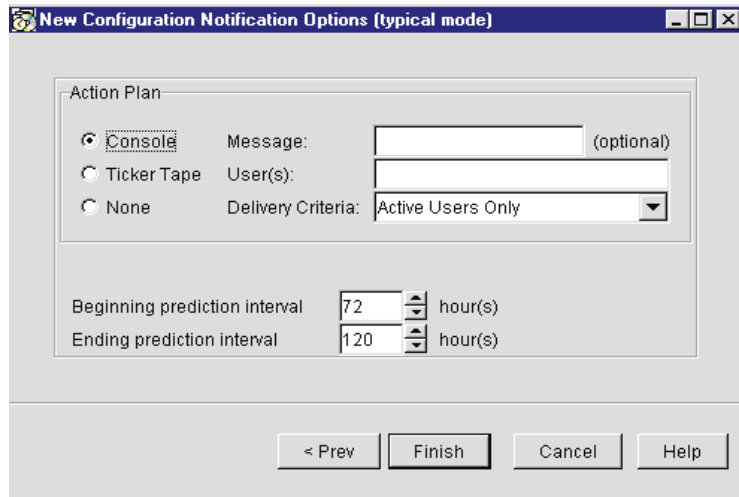
4. Configure Business Level Support.
Supported NT services will be listed. Select a service. You can

select multiple services. If you select IIS (Internet Information Serviceable predictions algorithms specific to IIS will be performed on this system. Supported Services are IIS v4/5, Domino v5.02, and Exchange 2000. If a Business-level service is not selected, no events or notification will be generated for that service.

5. Click **Next**.
The Forecasting window opens.



6. To configure Forecasting set the directory location and the maximum size for the database.
7. Click **Next**.
The **Notification Option** window opens.

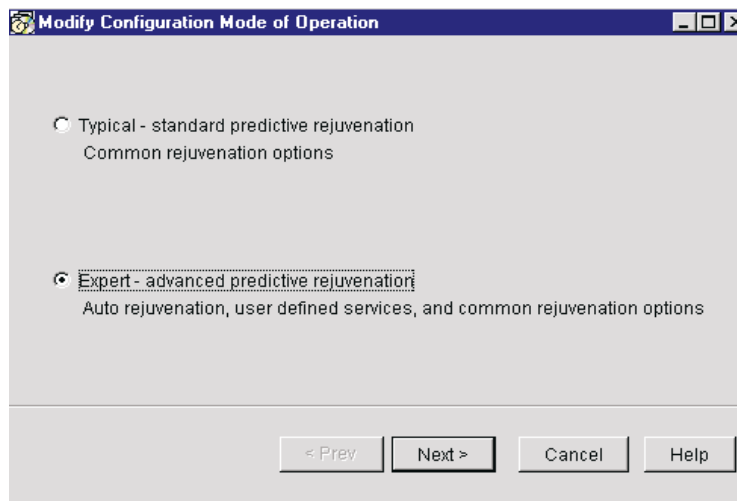


8. To configure Notification Options:
 - a. Under **Action Plan**, select **Console**, **Ticker Tape**, or **None**.
 - b. In the **Message** text box, enter a message (optional).
 - c. In the **User(s)** text box, enter user name.
 - d. From the **Delivery Criteria** drop-down list, select a criteria.
 - e. From the **Beginning** prediction interval drop-down list select the beginning prediction interval. From the **Ending** prediction interval drop-down list, select the ending prediction interval. This range represents the earliest and latest times (in the future) in which the forecasting logic will determine an exhaustion will occur.
The earliest time (in hours) is the earliest time that the forecasting logic will predict an exhaustion. The Latest time (in hours) is the latest time that the forecasting logic will predict an exhaustion. This range is configured so that the earliest time is always prior to the latest time. The smallest interval between earliest and latest is one day (24 hours).
9. **Click Finish.**

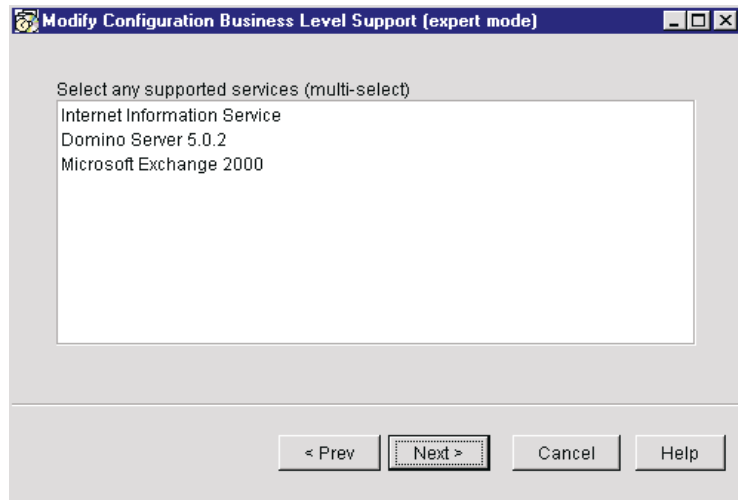
To configure an Expert Mode of Operation, use the following procedure:

1. From the Software Rejuvenation window, select **Tools**→**Prediction**→**Configuration Wizard** to start the configuration wizard.

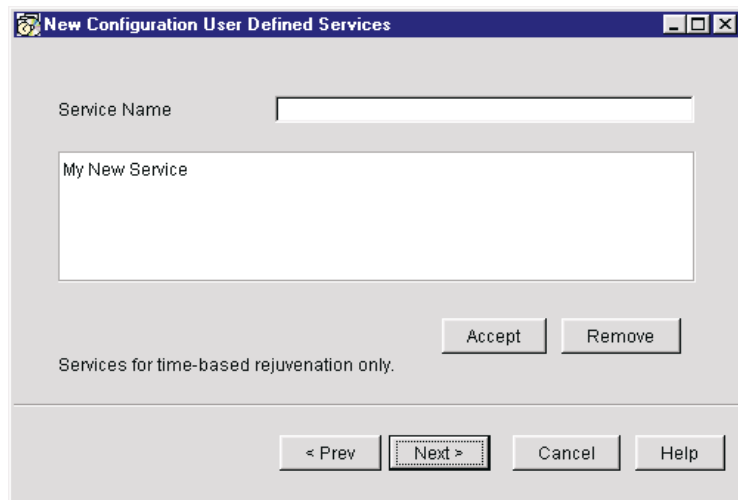
The Configuration Mode of Operation window opens.



2. From the Configuration - Mode of Operation dialog box, select **Expert**.
3. Click **Next**.
4. Configure the **Business Level** support. Supported NT services will be listed. Select a service. You may select multiple services. If you select IIS (Internet Information Serviceable) prediction algorithms specific to IIS will be performed on this system. Supported Services are IIS v4/5, Domino v5.02, and Exchange 2000. If a Business-level service is not selected, no events or notification will be generated for that service.



5. Click **Next**.
The User Defined Services window opens.

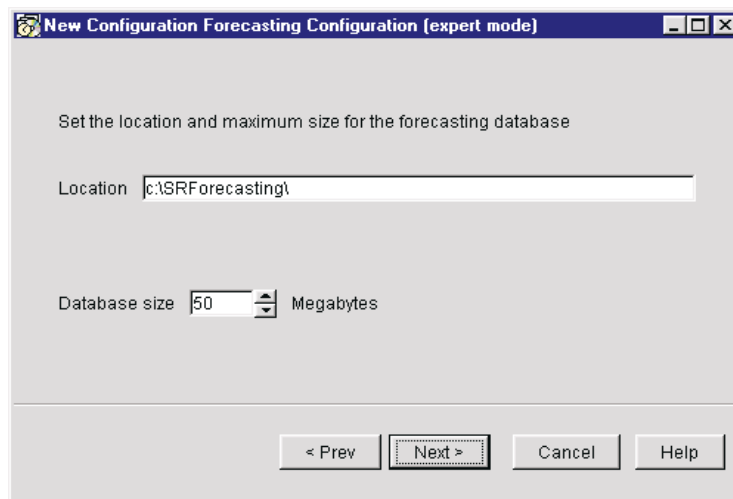


6. In the **Service Name** text box, enter a user-defined service.
Click **Accept**.
In **Expert mode**, in addition to listing supported NT services, you

can add additional services. If the service is found and the user has permission to stop or start the service.

7. Click **Next**.

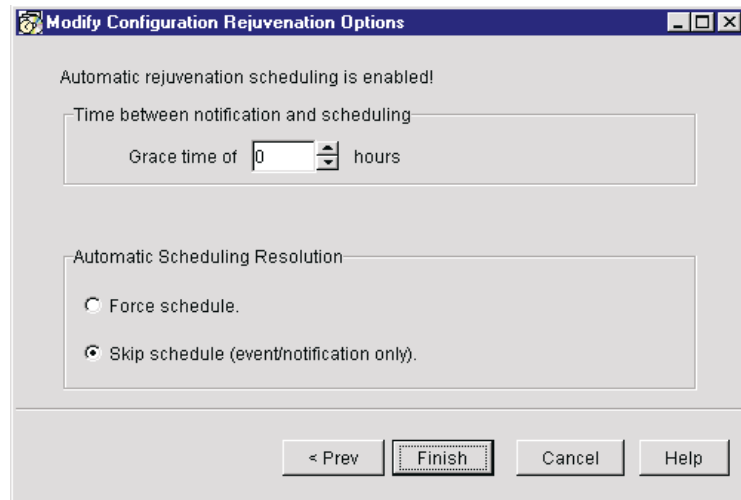
The Forecast Configuration window opens.



8. Configure Forecasting by setting the directory location and the maximum database size.
9. Click **Next**.
The Notification Options window opens.
10. To configure **Notification Options**, use the following procedure:
 - a. Under **Action Plan**, select **Console**, **Ticker Tape**, or **None**.
 - b. In the **Message** text box, enter a message (optional).
 - c. In the **User(s)** text box, enter user name.
 - d. From the **Delivery Criteria** drop-down list, select a criteria.
 - e. From the **Beginning** prediction interval drop-down list select the beginning prediction interval. From the **Ending** prediction interval drop-down list, select the ending prediction interval. This range represents the earliest and latest times (in the future) in which the forecasting logic will determine an exhaustion will occur.

The earliest time (in hours) is the earliest time that the forecasting logic will predict an exhaustion. The Latest time (in hours) is the latest time that the forecasting logic will predict an exhaustion. This range is configured so that the earliest time is always prior to the latest time. The smallest interval between earliest and latest is one day (24 hours).

11. Click **Next**.
The **Rejuvenation Options** window opens.



12. To configure Rejuvenation Options use the following procedure:
 - a. Select **System Level** or **Business level**.
A System level exhaustion means that the system will be restarted. A Business level exhaustion means that, if it's a service it will be stopped and started.
 - b. Enter the **Time between notification and scheduling** from the drop-down list.
 - c. Under **Automatic Schedule Resolution**, select **Force schedule** or **Skip schedule**.
You can choose to perform an automatic rejuvenation base on

the application predications and your "hands-off" configuration in the main Software Rejuvenation application.

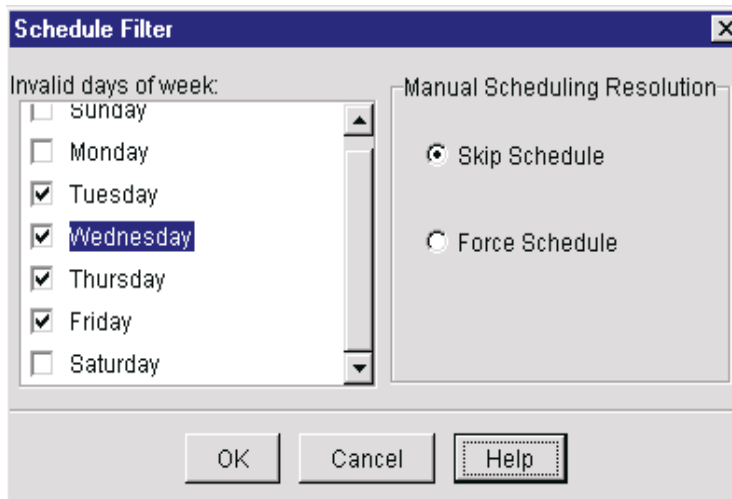
13. Click **Finish**.

Schedule Filter

You can schedule specific days for *no* rejuvenations. High availability of your resources is based on usage-rates. Therefore, most usage-rates are during a set of peak times. For example, a business may determine that peak usage times are during Wednesday's from 2:00pm through 7:00pm. You would schedule *no* rejuvenations during peak usage times. This mode of configuration is based on the concept that certain periods are considered *hands-off*.

To use the Schedule Filter, use the following procedure:

1. From the Software Rejuvenation window, select **Tools**→**Schedule Filter**
The Schedule Filer dialog box appears.



2. Select the invalid days of the week for a rejuvenation. Under Manual Scheduling Resolution, select **Skip Schedule** or **Force Schedule**.
3. Click **OK**.

Trend Viewer

To view the predictive information for a system use the Trend Viewer. First configure your system for prediction. See "Predictive Configuration" page 297.

To view the trend viewer information, use the following procedure:

1. From the Software Rejuvenation window, select one or more nodes.
2. Highlight the node and right click or select **Tools** → **Trend Viewer**.
3. Press **polling on** to enable a continuous refresh of the data gathering or press **polling off** to disable a continuous refresh of the data gathering. When a different parameter is selected, by default, polling is off is shown.

The Trend Viewer has the following areas: View Area, Parameter and Connect Samples.

- View Area: The charts x axis times the charts y axis represents the threshold. At the bottom of the chart, there is a line which represents the significant points in time. It also shows the exhaustion worst-case, expected-case, and the best-case.
- Parameter: With the Parameter selection, you can decide the parameters for the Trend Viewer information.
- Connect samples: With this option, you can *connect-the-dots* of sample points.

Closing the IBM Software Rejuvenation Program

To close the Software Rejuvenation program, select **File** → **Close**. If you did not save your file, the Verify application dialog box appears. Click **Yes** to close the file or click **No**, if you would like to save your file.

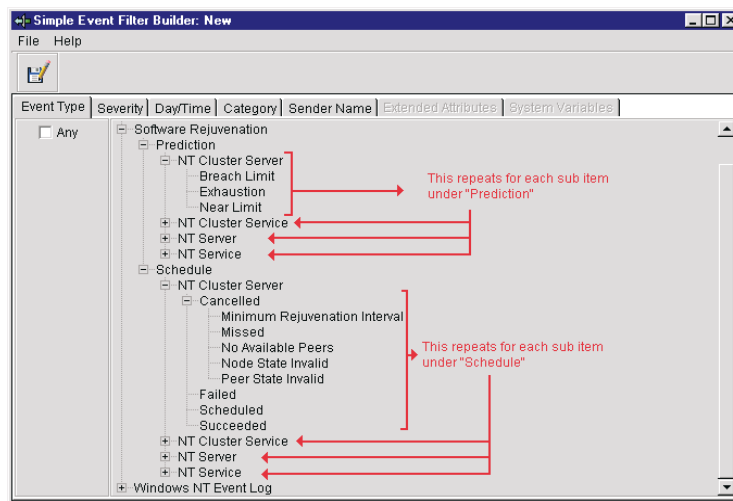
Creating Action Plans

The ICSM and Software Rejuvenation tools add new cluster event filters to the IBM Director console. Using the Events Action Plan Builder you

can create filters for specific cluster events, and further qualify the events using the Extended Attributes option.

To create an event action, do the following:

1. From the IBM Director console, select **Tasks** → **Event Action Plan Builder**.
2. Right-click in the **Event Filter** pane to display a menu.
3. Click **New** → **Simple Event Filter**.
The Simple Event Filter dialog box appears.



4. On the **Event Type** tab, clear the **Any** checkbox.
5. Click **Software Rejuvenation** to expand the directory tree.
6. Under **Prediction** or under **Schedule**, select an action to create an event schedule.
7. Click File **Save As**.
8. In the **Event Filter** box, type in a descriptive name for the event, and then click **OK**.

The new filter lists in the **Event Filters** pane of the **Event Action Plan Builder** window.

11

System Availability

System Availability provides graphical proof that Software Rejuvenation improves your system availability. This tool can be used, not only for the Software Rejuvenation users, but for any user with the responsibility of determining system availability for large sets of systems. This tool can be used on groups of nodes.

The System Availability Interface

Applying the System Availability task to a system with the UM Server Extension agent will open the System Availability interface. The interface consists of these elements:

- Menu bar
- An information window consisting of five tabs:
 - Distribution of System Outages
 - Distribution of System Uptime
 - Availability Report
 - System Outages by Day of Week
 - System Outages by Hour of Day
- A Status bar

The interface accepts multiple systems onto the same System Availability task.

The Menu Bar

The Menu bar consists of these selections:

- File
 - Set Time
 - Exit
- View
 - Detach View
 - Distribution of System Outages
 - Distribution of System Uptime
 - Availability Report
 - System Outages by Day of Week
 - System Outages by Hour of Day
- Help
 - Contents
 - About the Program

The Information window

The information window displays graphically the system availability of an assigned system or systems. By selecting a tab, you can view the following system availability selections:

- Distribution of System Outages
- Distribution of System Uptime
- Availability Report
- System Outages by Day of Week
- System Outages by Hour of Day

The Information window displays the combined averages of multiple systems when applied to the System Availability task.

Hover Help

When you move your mouse over a section of a pie chart or bar chart, hover help displays a detailed explanation of the hovered section.

Detailed List of Records

The information window graphically displays the system availability of an assigned system or systems. On a section of a pie chart or bar chart, there are two ways to display the Detailed List of Records, double click or right click and select **Detailed List of Records**.

The Detailed List of Records shows the entire System Availability operational record for that system or systems selected for the Availability Report. The list window is a reporting tool and cannot be edited. You can change the sorting structure of the report by clicking on one of the list topics.

The reporting topics are:

- Computer Name
- Start time
- Stop time
- Duration time
- Planned or unplanned outage

Note: The Planned outage list contains either a checked or unchecked box. A checked box means that the downtime was planned by a normal shutdown procedure. An unchecked box means that the system was shutdown unexpectedly.

Distribution of System Outages

The System Outages by Duration graph represents the percentage of various duration times a system/system is unavailable following a system outages. It indicates downtime of a system or systems.

The calculation is performed by measuring the duration time between stop and start events in the NT system log, specifically event (6006) stop time to event (6005) start time. The Time Marks measurements are:

Time Marks		
30 minutes	30*60	1800s
60 minutes	60*60	3600s
3 hours	3*60*60	10,800s
1 day	24*60*60	84,400s
7 days	7*24*60*60	604,800s
1 month	31*24*60*60	2,678,400s(assuming 31 days/ 1month)
3 months	92*24*60*60	7,948,800s(assuming 92 days/ 3 months)

In the calculations, T equals time, and Ti equals a percent: The value calculations are:

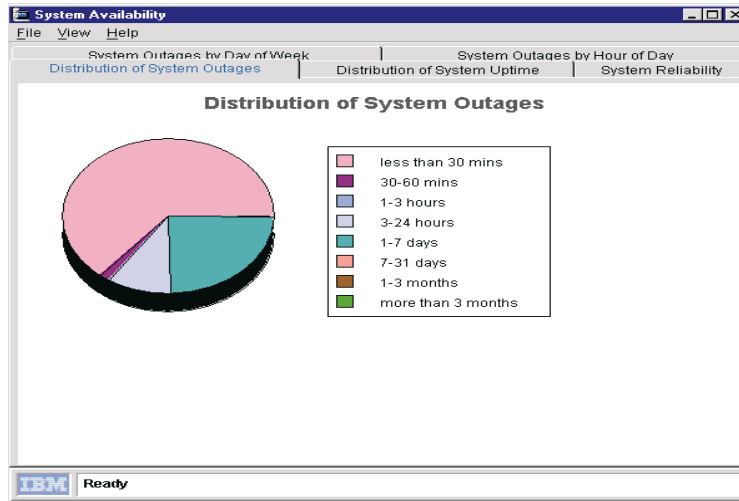
- T1: Total outages of system and cluster last < 30 min. in that period of time.
- T2: Total outages of system and cluster last >30 min. < 1 hour in that period.
- T3: Total outages of system and cluster last > 1hour < 3hours in that period.
- T4: Total outages of system and cluster last >3 hour < 1 day in that period.
- T5: Total outages of system and cluster last >1 day < 7 days in that period.
- T6: Total outages of system and cluster last >7 days < 1month. (1 month equal 31days).
- T7: Total outages of system and cluster > last 1 month < 3 months (92days).

- T8: Total outages of system and cluster last > 3 months.

$$\text{Total T} = T1+T2+T3+T4+T6+T7+T8$$

$$\text{Percent Ti} = Ti/\text{Total T} * 100$$

$$Ti = T1 \text{ or } T2 \text{ or...}T8$$

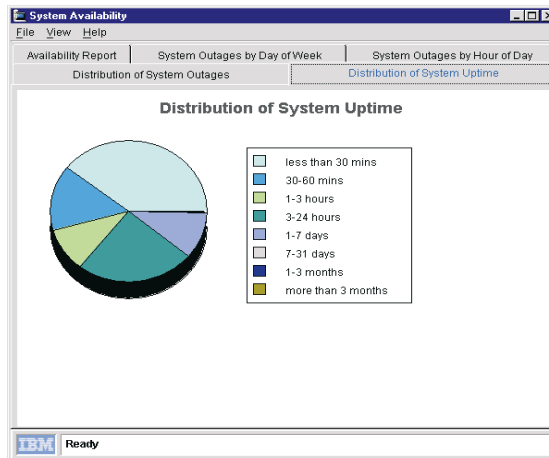


Distribution of System Uptime

The System Uptime graph represents the analysis of system availability and cluster availability. It shows percents of various time durations that a system and cluster are available.

The calculation is performed by measuring the duration time between start and stop events in the NT system log, specifically event(6005) start time to event (6006) stop time.

The System Uptime time marks and calculations are the same as the Distribution of System Downtime.



Availability Report

The Availability graph is a snapshot of system reliability. It provides measurements for system or systems.

The Servers Information measurements are:

- System name
- Date report from/to
- Total uptime: Sum of the times between startups and shutdowns (6005→6006)
- Total downtime: Sum of the times between shutdowns and startups (6006→ 6005)
- System availability: Total uptime (total uptime + total downtime).
- Total reboots: Count of startup events (6005).
- Mean time between reboots: Total uptime/Total reboots.
- Average days analyzed per system: Sum of analyzed days for all systems divided by the number of systems.
- Percentage of system rebooted: Sum of systems reboots/ Total number of system * 100.

-
- Mean Time between System crashes: Sum of analyzed days for all systems divided by the total number of unplanned reboots (6008).

System Outages by Hour of Day

The System Outages by Hour of Day graph shows the distribution of outages by time of day. The calculations are:

- T0: Total down time (in seconds) in period 12:00 a.m. through 1:00 a.m.
- T1: Total down time (in seconds) in period 1:00 a.m. through 2:a.m.
- T2: Total down time (in seconds) in period 2:00 am through 3:00 am.
- T3: Total down time (in seconds) in period 3:00 a.m. through 4:00 a.m.
- T4: Total down time (in seconds) in period 4:00 a.m. through 5:00 a.m.
- T5: Total down time (in seconds) in period 5:00 a.m through 6:00 a.m.
- T6: Total down time (in seconds) in period 6:00 a.m. through 7:00 a.m.
- T7: Total down time (in seconds) in period 7:00 a.m. through 8:00 a.m.
- T8: Total down time (in seconds) in period 8:00 a.m. through 9:00 a.m.
- T9: Total down time (in seconds) in period 9:00 a.m. through 10:00 a.m.
- T10: Total down time (in seconds) in period 10:00 a.m. through 11:00 a.m.
- T11: Total down time (in seconds) in period 11:00 a.m. through 12:00 p.m.
- T12: Total down time (in seconds) in period 12:00 p.m. through 1:00 p.m.

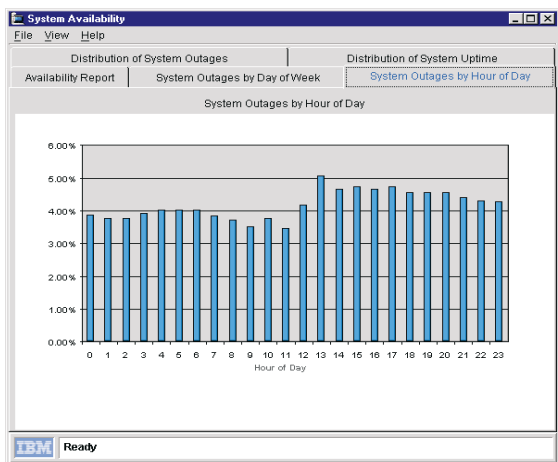
-
- T13: Total down time (in seconds) in period 1:00 p.m. through 2:00 p.m.
 - T14: Total down time (in seconds) in period 2:00 p.m. through 3:00 p.m.
 - T15: Total down time (in seconds) in period 3:00 p.m. through 4:00 p.m.
 - T16: Total down time (in seconds) in period 4:00 p.m. through 5:00 p.m.
 - T17: Total down time (in seconds) in period 5:00 p.m. through 6:00 p.m.
 - T18: Total down time (in seconds) in period 6:00 p.m. through 7:00 p.m.
 - T19: Total down time (in seconds) in period 7:00 p.m. through 8:00 p.m.
 - T20: Total down time (in seconds) in period 8:00 p.m. through 9:00 p.m.
 - T21: Total down time (in seconds) in period 9:00 p.m. through 10:00 p.m.
 - T22: Total down time (in seconds) in period 10:00 p.m. through 11:00 p.m.
 - T23: Total down time (in s) in period 11:00 p.m. through 12:00 a.m.

Total T = Total down time in the period (T0+T1+...T23)

Percent Ti = $T_i / \text{Total T} * 100$

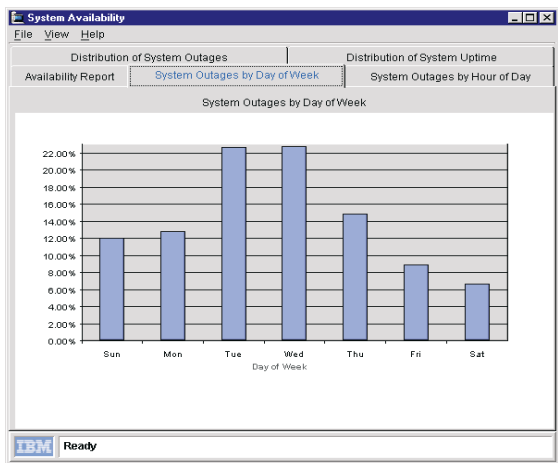
Ti = T1, T2....T23

Note: One hour is a period from xxh:00m:00s to xxh:59m:59s.



System Outages by Day of Week

The System Outages by Day of Week graph shows the distribution of outages by each day of a single week.

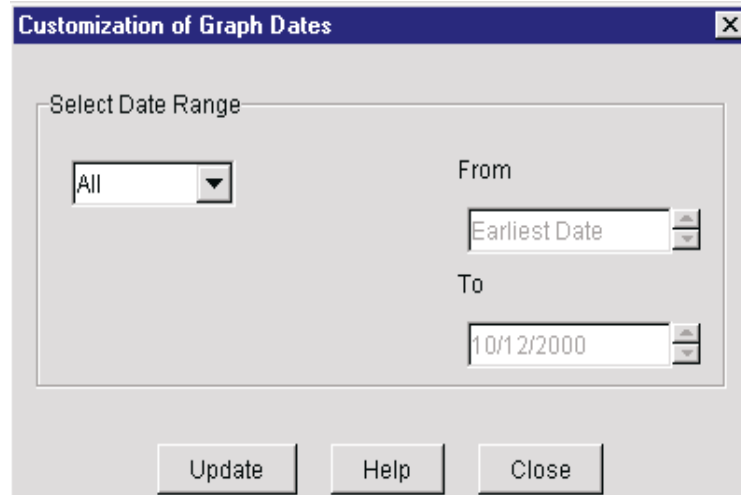


Displaying Data within a Defined Interval

To display data within a defined interval, customize the data displayed in the Information window. To select a specific interval, use the following procedure:

1. From the System Availability window, select **File**→**Set time**.

The Customization of Graph Dates dialog box opens.



2. From the Select Data Range drop-down list, select one of these periods:
 - All
Select this option to display system availability data from the earliest day to the present day. All is the default.
 - 1 week
Select this period to display system availability data from the present day to one week previous.
 - 1 month
Select this period to display system availability data from the present day to one month previous.
 - 3 months

Select this period to display system availability data from the present day to three months previous.

- 1 year

Select this period to display system availability data from the present day to one year previous.

- Customize

Select customize to display system availability data from a user defined period. See "Customizing Availability Dates" on page 319.

In the text boxes, enter a **From** date and a **To** date.

3. Click **Update** to update the Information window immediately.
4. Click **Close** to close the customization window.

Clicking **Close** will not automatically update the information window.

Customizing Availability Dates

To set a range of dates that are not listed as one of the options, select **Customize** from the drop-down list. In the **From** and **To** text boxes, click the up and down arrow to change the start date and end date.

Detach View

To detach the current tabbed selection of the Information window, from the System Availability window, select **View**→**Detach View**. The detached information window moves around the desktop independent of the System Availability task or the Director console.

Click the **X** in the upper right hand corner of the window to close the detached window. Closing the detached window will not close the System Availability task.

Index

A

- accessing the UM Server Extensions 18
 - accessing Advanced Systems Management 19, 20, 21
 - accessing Capacity Manager 19
 - accessing IBM Cluster Tools 19
 - accessing IBM SP Switch Administrator 20
 - accessing ServeRAID Manager 20
- Advanced System Management
 - configuration information 82
 - configuration settings 83
 - connecting remote hardware 72
 - event log 112
 - general settings 84
 - operational parameters 82
 - remote management 77
 - Remote Post Console 115
 - selecting an event source 81
 - starting 74
 - System Power Control 114
 - updating 116
 - updating system POST/BIOS Microcode 117
 - using the PCI adapter as a network gateway 72
- analyzing data with Capacity Management 137
- ASM Interconnect 77
 - establishing a connection 80

B

- bottlenecks identified with Capacity Management 142

C

- Capacity Management
 - analyzing data 137
 - bottlenecks 142
 - performance analysis 140
 - performance analysis monitors 137
 - performance analysis report 142
 - recommendations 141
 - capturing data 122
 - activating monitors 124
 - creating a status table 126
 - deactivating monitors 124
 - Monitor Activator 122, 123
 - monitoring for performance analysis 125
 - components of 119
 - forecasting 148, 149
 - the forecast display 150
 - viewing for selected systems 149
 - warning messages 150
 - viewing data 127
- capturing data for Capacity Management 122
- changing
 - cluster descriptions 48
 - IIS resource groups 65
 - network and network interface description 59
 - print-spooler resource groups 67
- Cluster Expert Wizard
 - changing resource groups
 - file-share 63
 - IIS 65
 - print-spooler 67
 - creating resource groups
 - file-share 62
 - IIS 64
 - print-spooler 66
 - resetting IP addresses 68
 - using 61– 69
- clusters

- changing description of 48
- disabling a network for use in 60
- enabling a network for use in 59
- managing 47– 48
- managing nodes in 49
- managing resource groups in 51– 58
- MSCS 42, 283
- renaming 48
- components of Capacity Manager
 - console 119
 - monitor activator 120
 - report definition 120
 - report generation 120
 - report viewer 121
- configuring ServeRAID Manager 262, 264, 265, 266
- connecting remote systems management hardware 72
- console for Capacity Manager 119
- creating
 - a file-share resource group 62
 - IIS resource groups 64
 - new resource groups 51
 - print-spooler resource groups 66

D

- data recommendations with Capacity Management 141
- drive states for ServeRAID Manager 276
 - logical 277
 - physical 276

E

- editing
 - scheduled rejuvenations 292
- Event Builder and UM Server Extensions 21
- event log for ASM 112
- event source, selecting for ASM 81

- event viewer for ServeRAID Manager 261

F

- file-share
 - changing a resource group 63
 - creating a new resource group 62
- forecast display, the 150
- forecast warning messages 150
- forecasting with Capacity Management 148
 - understanding calculations 149

G

- general settings for ASM
 - Configuration Settings 91
 - Dial-in Settings group 85
 - Loader Time-out 89
 - O/S Time-out 89
 - Post Time-out 88
 - Power Off Delay 90
 - System Identification group 85
 - System Management Processor Clock group 88
- group. See resource group

H

- hardware requirements, Netfinity Director Console 25
- hardware requirements, Netfinity Director Management Server 25

I

- IBM Cluster Systems Management 41
 - closing 42
 - GUI components of 43– 47

-
- IBM Netfinity Cluster Tools
 - components
 - using IBM Cluster Software Rejuvenation
 - using IBM Cluster Systems Management 42–??
 - IBM Netfinity Director Console 41
 - IBM Software Rejuvenation
 - closing 306
 - GUI components 281–??
 - IIS
 - changing resource groups 65
 - creating resource groups 64
 - increasing free space for ServeRAID Manager 272
 - increasing logical drive space for ServeRAID Manager 274
 - installing UM Server Extensions 23
 - on a NetWare system 30
 - on an OS/2 system 27
 - Internet Information Server 64
 - IP address
 - setting ranges 68
 - M**
 - main panel for ServeRAID Manager 261
 - managing
 - networks and network interfaces 58– 61
 - nodes in clusters 49– 50
 - resource groups in a cluster 51– 58
 - software rejuvenations 283, 306
 - mapping a drive 28
 - menu bar for ServeRAID Manager 259
 - Microsoft Cluster Service 42
 - modem settings for ASM 100
 - Initializing String Guidelines 107
 - The Dialing Settings group 105
 - The Port Configuration group 101
 - Monitor Activator for Capacity Management 122, 123
 - monitor activator for Capacity Manager 120
 - monitors
 - activating and deactivating 124
 - performance analysis 125
 - monitors, performance analysis 137
 - MSCS 42
 - MSCS. See also Microsoft Cluster Service 283
 - N**
 - NetWare system
 - installing UM Server Extensions on 30
 - mapping a drive to 28
 - network settings 108
 - networks and network interfaces
 - changing description of 59
 - disabling for use in a cluster 60
 - enabling for use in a cluster 59
 - managing 58– 61
 - nodes
 - entering descriptions of 50
 - managing in cluster 49– 50
 - pausing a 49
 - resuming 50
 - starting a 49
 - stopping a 49
 - O**
 - OS/2 system
 - installing UM Server Extensions on 27
 - mapping a drive to 28
 - P**
 - PCI adapter as a network gateway 72
 - performance analysis (Capacity Management) 140
-

- monitors 125, 137
- reports 142
- preferences (user) for ServeRAID Manager 275
- print-spooler
 - changing resource groups 67
 - creating resource groups 66

R

- RAID levels (changing) 271
- remote event settings for ASM 91
 - Enabled Events group 95
 - Remote Event Entry Information group 93
 - Remote Event Strategy group 94
- remote management
 - configuring and establishing a TCP/IP connection 78
 - rules for ASM Interconnect 77
- remote post console for ASM 115
- reports for Capacity Management
 - report definition 120
 - report generation 120
 - report viewer 121
- resource group
 - bringing online 53
 - changing description properties of 54
 - changing file-share 63
 - changing preferred owners of 55
 - creating a new 51
 - creating file-share 62
 - deleting 53
 - file-share 62
 - moving to another node 58
 - print-spooler 66
 - renaming a 53
 - setting failback policy for 57
 - setting failover policy for 56
 - taking offline 54

S

- scheduling
 - software rejuvenations 283
 - daily 286
 - for a single node 283
 - monthly by date 288
 - monthly by day 289
 - single 285
 - weekly 287
- ServeRAID Manager
 - configuring 262
 - custom configuration 265, 266
 - express configuration 262, 264
 - modifying the configuration 270
 - using the information mode 269
 - drive states 276
 - logical 277
 - physical 276
 - Events in Director 278
 - modifying 270
 - changing RAID levels 271
 - increasing free space 272
 - increasing logical drive space 274
 - the configuration 270
 - setting user preferences 275
- ServeRaid Manager
 - using the interface 259
 - the menu bar 259
 - the tool bar 260
 - viewing the event viewer 261
 - viewing the expandable tree 260
 - viewing the main panel 261
 - viewing the status bar 261
 - using the program 262
- SNMP settings 110
- software rejuvenations
 - command shortcuts for 296
 - editing 292
 - at cluster level 294
 - at node level 292
 - managing 283

removing a node from schedule 295
scheduling 283, 290
starting
 a node 49
 ICSM 43
status bar for ServeRAID Manager 261
status table for Capacity Management 126
system power control for ASM 114

W
warning messages for performance
 forecasting 150

T

tool bar for ServeRAID Manager 260
tree structure for ServeRAID Manager 260

U

UM Server Extensions
 Advanced Systems Management
 introduction 15
 Capacity Management introduction 16
 IBM Cluster Tools introduction
 Clusters Systems Management 16
 Software Rejuvenation 18
 installing 23
 introduction to 15
 ServeRAID Manager introduction 17
UM Services Client, supported systems 24
uninstalling UM Server Extensions
 from a local Windows NT system 35
 from a remote system 35
user preferences for ServeRAID Manager
 275

V

viewing data with Capacity Management
 127
viewing performance forecasts for selected
 systems 149



Part Number: UM021205

UM021205