Director
Version 3.1

User's Guide

Director
Version 3.1

User's Guide

Contents

Preface

The *Director User's Guide* provides the installation and start-up instructions for the Director product. It also describes the Director environment and the many tasks and services available to help you manage your network.

Who should read this guide.

This User's Guide is intended for webmasters and IT administrators in small-to medium-sized independent businesses, responsible for installing, configuring, and maintaining local area network (LAN) environments with hundreds of PCs and other network devices.

Readers should have a general knowledge of operating systems, network operations, and database functions.

What this guide contains

This User's Guide is organized into the following chapters:

- Chapter 1, "Introduction," on page 1, describes how Director works and introduces the various tasks available to the network administrator.
- Chapter 2, "Planning," on page 15, discusses considerations for network setup and management that should be addressed before installation and network administration tasks are performed with Director.
- Chapter 3, "Installation and configuration," on page 39, lists the prerequisites and restrictions that apply to Director and provides step-by-step instructions for installing and configuring Director component and agent software. The uninstallation procedure for Director is also provided.
- Chapter 4, "Upgrading Director and IBM Director Extensions," on page 73, details the procedures used for upgrading your previous versions of Director to Director 3.1.
- Chapter 5, "Using the Director Management Console," on page 79, describes the Director Management Console graphical user interface (GUI).
- Chapter 6, "Inventory Management," on page 91 through Chapter 31, "System Availability," on page 361, describes the various administrative tasks available in Director for managing the hardware and software in your network.
- Chapter 32, "Troubleshooting," on page 371, describes some problems you might encounter and possible ways to resolve them.
- Appendix A, "Resource monitor attributes," on page 379, contains information on the process used to establish trust relationships between the Director server and IBM Director agents when the network is brought up. This appendix describes the process and files used by IBM Director to implement agent-server security.

- Appendix B, “Creating the ODBC entry for the default database,” on page 395 through Appendix D, “Defining table property files,” on page 399, list steps to manually create the default Microsoft® Jet database and contains information on converting database support from the default Microsoft Jet database to any of the other supported databases and for converting between those databases.
- Appendix E, “Agent-server security,” on page 417, describes the process and files used by Director to implement agent-server security, and provides guidelines.
- Appendix F, “IBM Director Agent,” on page 421,through Appendix K, “Upward Integration Modules,” on page 505 provides you with basic information for installing and using IBM Director Agent on your system.

For related Director terminology, go to `http://www.networking.ibm.com/nsg/nsgmain.htm`. You can search for terms and download Portable Document Format (*.pdf) and PostScript (*.ps) glossary files from this Web page.

Chapter 1. Introduction

Director is a management product for the systems administrator in small- to medium-sized growth companies. Director solves the problem of managing today's Windows® and Intel®-centric PC and LAN networks and addresses real system management issues, enabling you to focus on managing your primary business.

The Director functions enable you to manage all aspects of the computing infrastructure, including software applications, network operating system (NOS) services, networks, and hardware.

Director provides integrated management across the network, Internet workgroup control and management, and highly automated, almost unattended operation.

The following terms are used in the guide:

Native Systems that have the Tivoli® management agent or IBM Director Agent installed and that communicate with the Director server.

SNMP The Simple Network Management Protocol. Defines a schema for representing network resources. SNMP devices are detected separately from Director native systems.

RMON Remote network monitor for SNMP devices. Further defines the SNMP schema and enables the collection of hundreds of additional network monitor statistics.

CIM Common Information Model. Defines a schema for representing network resources.

DMI Desktop Management Interface. Defines a schema for representing network resources.

Cluster A group of network resources whose ownership can be switched between managed systems.

How Director works

Director operates in a distributed environment. It consists of the following main components:

- Director Management Console

  The Director Management Console is the graphical user interface (GUI) from which administrative tasks are performed. It is your primary interface to the various Director tasks.

The Director Management Console GUI is fully Java™-based with all state information stored on the server. It runs as a locally installed Java application in a Java virtual machine (JVM) environment.

- Director Server

  The Director server is the heart of the Director product. The server engine, the management data and the management application logic reside in the Director server. It is a Java and native C++ application. Director provides basic functionality such as discovery of the network components, persistent store of inventory data, Structured Query Language (SQL) database support, presence checking, security and authentication, management console support, and support for each of the administrative tasks.

- Managed systems

  The Director server manages systems and devices in your network by communicating with the IBM Director Agent, Tivoli® management agent or with other agents installed on these *managed systems*. The agent provides all the code and interfaces necessary for the system to be managed by Director. Director recognizes two types of managed systems:

  — Native managed systems

  These are managed systems that have the IBM Director Agent installed, which acts as a passive, non-intrusive application. Users of these systems do not have access to a GUI, but users can communicate with Director by using a Web browser for certain network status information.

  — SNMP devices

  These are network devices, printers, or PCs that have an SNMP agent installed or embedded.

  Note: SNMP agents are not provided with Director.

Introducing IBM Director Extensions tools

IBM Director Extensions tools expand the flexibility and management capabilities of the Universal Manageability tool, Director. These systems-management tools are provided free-of-charge with Director. The following is a description of each IBM Director Extensions tool.

Management Processor Assistant

With Management Processor Assistant, 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 time-outs. 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 bottlenecks, and recommends ways to improve performance. It can generate a report, which enables effective planning of future capacity, such as microprocessor, disk, network or memory upgrades.

The Capacity Manager Report Generator tasks can be executed immediately or scheduled through the Director Task 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 can receive alerts to any event in a cluster through e-mail, electronic page, or starting another program, thereby reducing downtime.

Fuel Gauge Monitor

Fuel Gauge Monitor warns you about conditions that 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 know 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 critical.

Rack Manager

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

Software Rejuvenation allows you to schedule the restart of managed systems. The Software Rejuvenation interface allows you to drag-and-drop a system icon onto a calendar and set the time of the rejuvenation. You can also configure Predictive Rejuvenation, where resource utilization will be monitored and systems can be rejuvenated automatically before utilization becomes critical.

System Availability

System Availability is used to analyze the availability of a system or a group of systems. It can be used to provide statistics on the availability of systems. Additionally, it may be used as a means to graphically prove that Software Rejuvenation improves system availability. From the Director Management Console, drag-and-drop the System Availability icon onto a system or a group of systems. After the System Availability window opens, for a single system, it displays the system availability statistics for that system. For a group of systems, it displays combined averages of the group.

Accessing the IBM Director Extensions tools

Each task is started from the Director Management Console. Each IBM Director Extensions tool is accessed by dragging the appropriate icon from the Tasks pane of the Director Management console and dropping it onto a system that supports that IBM Director Extensions tool in the Group Contents panel. Or, you can right-click a system that supports the IBM Director Extensions tool in the Group Contents pane and select the appropriate 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 IBM Director Extensions tool service icon (or icons) and the placement of the icon within the Director console.

Cluster Systems Management

Icon	Location
	Cluster Systems Management is within the Cluster Tools parent description in the **Tasks** section of the Director Management console.

Management Processor Assistant

Icon	Location
	Management Processor Assistant is in the **Tasks** section of the Director Management console.

Capacity Manager

Icon	Location
	Capacity Manager is an IBM Director Extensions service description and not a task. The supported tasks are located within this heading in the **Tasks** section of the Director console.
	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.
	Monitor Activator is within the Capacity Manager parent description in the **Tasks** section of the Director Management console.
	Report Generator is within the Capacity Manager parent description but is an x Series Extensions service description and not a task. Click this description to expand the type of report selections.
	Report Viewer is within the Capacity Manager parent description in the **Tasks** section of the Director Management console.

Fuel Gauge Monitor

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

Rack Manager

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

ServeRAID Manager

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

Software Rejuvenation

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

System Availability

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

IBM Director Extensions tools and the Event Builder in Director

Some of the IBM Director 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 IBM Director Extensions events.

Descriptions of event filters are found in:

- Chapter 25, "Management Processor Assistant," on page 237.
- Chapter 26, "Capacity Manager," on page 277.
- Chapter 30, "Software Rejuvenation," on page 335.

What is new in this release

The following section describes the new functional details of this release of Director.

Integrated IBM Director Extensions and IBM Director Agent

In previous releases, IBM Director Extensions and IBM Director Agent were installed separately. IBM Director Extensions are an extension to Director with the same configuration as Director. In this release, Director 3.1 installs and upgrades IBM Director Extensions and IBM Director Agent (formerly known as Universal Manageability Services). If Director Server is installed, IBM Director Agent will also be installed with the server. (This is also true for all the components.) IBM Director Extensions are also installed in the same directory as Director.

OEM console and server installation

Director 3.1 can now be installed on any Intel-based system that supports SMBIOS 2.1 or above and meets all other hardware and software requirements. For more information, see "Hardware requirements" on page 39.

Linux installation

The following IBM Director Extensions tools can be installed using the Red Hat Package Manager (RPM): ServeRAID, Management Processor Assistant, Software Rejuvenation, Capacity Manager, and System Availability. For information about these installation scripts, see “Installing Director on a system running Linux” on page 54.

The Director CD contains theLinux\DirAgent-3.10-1.i386.rpm file. This file complies with the Red Hat Package Manager format. For information on how to install rpm packages, see the operating systems user guide. DirAgent-3.10-1.i386.rpm will install the IBM Director Agent, SNIA CIMON 1.0, and if desired, the agent portion of any server extension tool. The selection of server extension tools are ServeRAID, Management Processor Assistant, Rack Manager, SMBIOS instrumentation and SNMP access and trap forwarding.

Apache Web server

The internal web server of Director Agent has been replaced by the Apache Web server to improve performance and enhance security capability of Director Agent. The Apache Web server, a freeware application, developed by the Apache Software Foundation, can be used to install and test Director. For more information about the Apache Web server, see the Apache Software Foundation Web site at www.apache.org.

Package for the Web Console

A file named DirectorConsole.exe is provided in PackageForTheWeb format. It contains the Director console and the IBM Director Extensions for Director console extensions. You must copy DirectorConsole31.exe to a directory on an existing web server (for example, inetpub\wwwroot for IIS). This gives you the ability to point your web browser at the web address for the package and download it. It will automatically extract and install the Director management console and IBM Director Extensions tools, and then re-start the system. Afterward, you can use the console to point to a known Director server and manage its environment.

Event configuration

You can use the Health Configuration service to enable and disable event actions, set minimum and maximum threshold values for temperature and drive space and set the severity associated with each threshold. IBM Director Agent features six event actions and five categories of events with their supported severities

System Health

This new task provides an integration of status from the various subsystems into an overall system health. It provides an overall status and individual system health for servers on the Director console, Rack Manager, and IBM Director Agent Web-based facilities. The status presented on the Director console provides drill-down capability. The subsystem with a problem is identified. When that subsystem is selected, a list of events produced by that subsystem is presented in chronological order with individual severity in the right hand panel.

Mass Configuration

Mass Configuration is a three-tier architecture used to manage configurations across large groups of workstations and servers. This function is an extension to the IBM Director Agent. With Mass Configuration you can configure large numbers of workstations and servers from a single workstation. The Mass Configuration service also provides scheduling, conflict resolution, logging and a robust environment for the setting of CIM properties and the invocation of CIM methods. Mass Configuration can be used with the following Director tasks: Network Configuration, Alert on LAN™, and Asset ID™.

ServeRAID enhancements

ServeRAID has enhanced firmware, BIOS, device drivers, and utilities. It now supports a broader range of operating systems.

Capacity Manager enhancements

In the past, performance analysis was performed if the necessary monitors existed. This could be a time-consuming operation. A mechanism to disable performance analysis has been added.

A new mechanism called Group support has been added to Director 3.1. It enables you to define a set of systems as a group with Capacity Manager and treat that entity has a unit.

Supported communication protocols

Director relies upon a multi-protocol transport layer that enables the server engine to communicate with the Director Management Console and the managed systems.

Director uses TCP/IP to communicate between the Director Management Console and the Director server.

Director uses the following protocols to communicate between the Director server and its native managed systems and SNMP devices:

- Server and native managed system:
  - NetBIOS
  - IPX
  - TCP/IP
  - SNA
- Server and SNMP device:
  - IPX

— TCP/IP

Note: TCP/IP is the only protocol supported by Director servers and UNIX or Linux® agents systems.

Modem connections

For Director managed systems on Win32 platforms, the transport can detect when a modem connection becomes active. When it detects that a modem connection has been activated, the managed system will send a message to all known servers with its current IP address. The server can then update the address of that managed system for communications. This feature is only supported on Win32 systems.

Managing native systems

Director supports a comprehensive set of tasks for managed systems operating as full-function Director management agents. These agents communicate directly with the Director server, enabling the following tasks to be performed:

Asset ID™

Director collects extensive hardware information on all your IBM hardware equipped with the Asset ID chip. From serial numbers to lease information on your specific system hardware, Director displays this information as a agent-based task. See Chapter 16, “Asset ID,” on page 163 for more information.

Inventory management

Director collects information from discovered managed systems and stores the information in the inventory database. You can then view and analyze collected hardware and software inventory data and customize the display for your needs. See Chapter 6, “Inventory Management,” on page 91 for details.

Remote Control

Remote control provides faster and more accurate problem resolution by remotely controlling the desktop of a native managed system, sending keystrokes and mouse commands to the remote system, and displaying the remote system’s desktop on the Director Management Console. It is useful for training and educating new network administrators as well. See Chapter 7, “Remote Control,” on page 103 for details.

Resource Monitoring

Resource monitoring enables you to view statistics and usage of resources on the network. Information on attributes such as the central processing unit (CPU), disk, file, memory, and network protocols are collected and monitored. You can also set thresholds, record monitor data, generate graphs, and generate events when thresholds are exceeded. See Chapter 8, “Resource Monitoring,” on page 109 for details.

Event Management
Enables you to view a log of events that have occurred for a managed system or group of systems and to create event action plans to associate an event with a desired action, such as sending an e-mail, starting a program, or logging to a file. See Chapter 9, “Event management,” on page 117 for details.

Software Distribution
Enables you to collect IBM created file packages that can then be applied to one or more managed systems for immediate or scheduled delivery. See Chapter 10, “Software Distribution,” on page 127 for details.

File Transfer
Enables you to perform basic file transfer tasks on remote systems, such as manipulating files, updating device drivers, and replacing system files. Included is the feature that allows for “wildcard” filename transfers. See Chapter 11, “File Transfer,” on page 137 for details.

Process Management
Enables you to start, stop, and monitor applications and processes on remote native systems. You can have Director watch a particular process or application and generate an event if the application or process is started or terminated. See Chapter 21, “Process management,” on page 195 for details.

Task Scheduler
Enables you to schedule non-interactive tasks such as software distribution and inventory collection. You can schedule tasks on an hourly, daily, weekly, monthly, or yearly basis. Tasks can be triggered by changes in the state of managed systems or by the discovery of new hardware or software in the network. In addition, you can schedule tasks for individual managed systems or groups of managed systems. See Chapter 22, “Task Scheduler,” on page 201 for details.

Managing DMI-enabled and CIM-enabled native systems

Director can manage Win32 native systems configured for DMI support or CIM support. The following tasks can be applied to DMI- and CIM-enabled native systems:

- DMI Browser and CIM Browser
- Inventory
- Resource Monitors
- Event Management

Refer to Chapter 12, “DMI management,” on page 143 and to Chapter 16, “Asset ID,” on page 163 for information on tasks you can perform on DMI and CIM data.

Managing cluster-enabled native systems

Director can manage Windows NT®-native systems configured with Microsoft Clustering Service (MSCS). The following tasks can be applied to cluster-enabled native NT systems:

- Event Action Plans
- Cluster Browser
- Resource Monitors
- Event Management

See Chapter 20, “Cluster Management,” on page 191 for information on viewing cluster data.

Managing SNMP devices with Director

Director can also manage network devices, printers, and PCs that have SNMP agents installed or embedded. Tasks that can be performed on SNMP devices include:

- Event Action Plans
- Inventory
- Resource Monitors
- SNMP Browser

Basic monitor data can be collected from SNMP managed systems. Additional monitor data can be collected from SNMP managed systems that support the RMON MIB.

See Chapter 14, “SNMP Management,” on page 153 for details.

Multi-homed support has been added as well. A multi-homed device has two or more physical connections and requires multiple TCP/IP addresses, one corresponding to each of the device's network connections. Refer to Chapter 14, “SNMP Management,” on page 153 for more information.

Additional features in Director

Security

The Director server uses the security subsystem of the operating system for validating user IDs and passwords. Each Director Administrator has a unique login profile. This enables different users to log in to the Director Management Console. Refer to the section “Security” on page 26 for more information.

Database Management

Director supports the storage of hardware and software inventory data, and device information to the Microsoft Jet database(Access). For more advanced database needs, Director also supports Microsoft SQL Server, Oracle, IBM DB2 Server and MSDE.

Chapter 2. Planning

This chapter provides information that you should consider before you begin to install and configure IBM Director.

Director server

The Director server is where most of the Director processing occurs. Therefore, it requires more computer resources than the Director Management Console or the management agent software. Depending on your server, configuration, and the number of systems to be managed, you might need to dedicate an entire server in your network to act as the Director server.

Management server prerequisites

The following requirements are necessary for the installation:

- Sufficient free disk space must be available on the target system. See "Hardware requirements" on page 39.
- You must have administrator rights on the target system.
- Supported operating system. See "Supported platforms for Director 3.1" on page 40.

For more information, see Chapter 3, "Installation and configuration," on page 39 or refer to the product README file for the minimum hardware and software requirements for the Director server.

The TCP/IP networking transport and a network adapter that supports the TCP/IP is also required. The adapter must also support NetBIOS, IPX, or SNA, depending on the transport that is needed to communicate with the managed systems. See "Transport support" on page 23 for information on supported versions of transports.

Database support

Director supports the following databases and versions:

- Microsoft Access (Jet)
- Microsoft SQL Server 2000, Service Pack 1
- IBM DB2® 6.1, 7.1, 7.2
- Microsoft SQL Server 6.5 with Service Pack 5A and 7.0 with Service Pack 3
- Microsoft Data Engine (MDSE)
- Oracle Server 7.3.4 through 8.1.7 and 9i

Refer to the product README file for the supported operating systems for these databases.

You can use any of these for your database needs, depending on your systems management requirements. This database stores inventory data and any new tables created as part of a third party application extension to Director. Monitor and event data is stored in data files. To access the database, the Java Database Connectivity (JDBC) is used. In addition, Microsoft requires Open Database Connectivity (ODBC) APIs.

This driver change is required on the Director server. Before you run the command, make sure all DB2 services are stopped.

Planning to use the Jet database

Director comes with and uses by default the Microsoft Jet database.

The Jet database is a single database file and must be installed on the same system as the Director server. The Jet database has a maximum size of 1 GB.

Relocating the Jet database

The Jet database cannot be split. After the Director server is installed, it is possible to move the Jet database to another subdirectory in addition to *\database*, but this must be done manually with the server shut down at the time. You also must change the ODBC entry manually to make it point to the new file location. The name of the ODBC entry to be changed is the name that you selected when you installed the Director server (either the default or another name that you selected). Refer to the Windows NT online help for OBDC, or see your database administrator for assistance, if needed.

Planning to use the DB2 Universal database

Depending on the requirements of your environment, you may want to use the IBM DB2 Universal database instead of the default Jet database. The DB2 Server has additional storage capability and is more impervious to unwanted access attempts. Before installing the Director server software, your network must be configured to use the DB2 Server database. Install the DB2 Agent Application Enabler to access the DB2 Server. Make sure that the DB2 Java Enablement option is installed and that CLASSPATH points to the director that contains the DB2 Java.zip files. The following may require preliminary action:

- If you have a remote connection to the DB2 Server, do the following:
  - Set up a trusted connection or give proper login access to the Director server database user ID.
  - Create a node entry for the remote DB2 Server.
- Make sure you have sufficient licenses for the DB2 Server, as this is a separate product from Director and is not included in Director Licensing requirements.

Setting up trusted connections

Director may use trusted connections when logging in to the DB2 Server. The database administrator can set the database server security to support trusted connections. Refer to the *DB2 Administration Guide* for information on trusted agent scenarios.

Note: The JDBC 1.22 driver is the default JDBC driver for DB2 release 6.1, 7.1 and 7.2. Director now requires the JDBC 2.0 driver. To install the JDBC2.0 driver for Windows 32 bit operation systems, enter the usejdbc2 command from the sqllib\java 12 directory. This command performs the following tasks:

- Creates a sqllib\java11 directory for the 1.22 driver files.
- Backs up the JDBC 1.22 driver files into the sqllib\java 11 directory.
- Copies the JDBC driver files from the sqllib\java 12 directory into the appropriate directories.

DB2 Server login access for Windows

Your database administrator and your system administrator must configure security so that the Director Management Server Database User ID is able to login on the server that will be used for the DB2 database and has at least user-level login privileges for the DB2 Server. You may need to set up a trusted relationship between domains if the Director management server and the DB2 Server are on different domains. The Director user ID must be a domain account and must be authorized to login (see your NT system administrator or documentation for details).

Creating the DB2 Server database

Your database administrator may choose to create the database manually, or enable the database to be created automatically during Director server installation. Your database administrator should consider the following:

- The Director Management Server Database User ID must be given user access to the database server.
- To create the database automatically, the Director Management Server Database User ID must be given Create Database permission on theDB2 server database. If this level of authority is not desired, then the Administrator should create the database manually and either transfer ownership of the database to the Director Database User ID, or minimally give the user Create Table permission and User-level access to the database.
- When the database is created automatically, it will use the default values specified in the DB2 Administration Guide.

  An initial size of 100 MB is recommended for the database to hold data for 250 to 500 managed systems. More space might be required if you are managing more systems or if your software inventory data is extensive. If the DB2 database default size is not sufficient for your needs, then the database administrator can either modify the default values or create the database manually with the desired size. The size can be increased later, if necessary. Your database administrator should monitor this database and adjust its size as needed.

Whether the database is created manually or automatically, your database administrator should provide the name of the server where the database is located, and the name of the database itself. You are now ready to proceed with Director Management Server installation.

Note: DB2 has size restrictions on items such as user ID and table names. Refer to the *DB2 SQL Reference* guide for more information.

Planning to use the SQL Server

Depending on the requirements of your environment, you may want to use the Microsoft SQL Server instead of the default Jet database. SQL Server has additional storage capability and is more impervious to unwanted access attempts. Before installing the Director server software, your network must be configured to use the SQL Server database by:

- Setting up a trusted connection or giving proper login access to the Director server database user ID.
- Creating the SQL Server database manually or during Director server installation.

Note: Be sure you have sufficient licenses for Microsoft SQL Server, as this is a separate product from Director and is not included in Director licensing requirements.

Setting up trusted connections

Your database administrator must set the database server's security to support trusted connections. The recommended configuration is *mixed security*.

SQL Server login access for Director

Your database administrator and your NT system administrator must configure security so that the Director server user ID can:

- login to the NT Server that will be used for the SQL database.
- have at least user level login privileges for the SQL Server.

You may need to set up a trusted relationship between domains if the Director server and the SQL Server are on different domains. The Director user ID must be a domain account and must also be authorized to login (see your NT system administrator or documentation for details).

Creating the SQL Server database

Your database administrator may choose to create the database manually, or enable the database to be created automatically during Director server installation. Your database administrator should consider the following:

- The Director management server user ID must be given user access to the master database.
- To create the database automatically, the Director management server user ID must be given Create Database permission in the master database. If this level of authority is not desired, then the administrator should create the database manually and either transfer ownership of the database to the Director user ID, or give at least user-level access to the database, as well as Create Table permission.
- When the database is created automatically, the size of the database will default to the larger of:
  - The size of the model database
  - The default database size specified in the SQL Server configuration options (sp_configure).

    An initial size of 100 MB is recommended for the database to hold data for 250 to 500 managed systems. You might find you need more space if you are managing more systems or if your software inventory data is extensive. If the SQL Server default size is not sufficient for your needs, then the database administrator can either modify the default values or create the database manually with the desired size. The size can be increased later, if desired. Your database administrator should monitor this database and adjust its size as needed.
  - For SQL 6.5 only, when the database is created automatically, the database and the transaction log can be placed on a single device. You will be prompted to select the available device. If your database requirements call for further customizing, such as spanning the database across multiple devices, the database administrator should create the database manually, and configure it for multiple devices as desired.

Whether the database is created manually or automatically, your database administrator should tell you the name of the server where the database is located, and the name of the database itself. If the database will be created automatically during installation, your database administrator should also tell you the name of the devices to use for the database and the transaction log. You will use this information during the Director server installation.

You are now ready to proceed with Director server installation.

Planning to use the Oracle Server database

Depending on the requirements of your environment, you may want to use the Oracle Server Database instead of the default Jet Database. Oracle Server has additional storage capability and is more impervious to unwanted access attempts. Before installing the Director server software, your network must be configured to use the Oracle Server database. The following might require preliminary action:

- If you do not have a User ID, one is created during the database configuration process.
- The JDBC Thin agent-side driver is used for database connection. This is a JDBC Type 4 driver that uses Java to connect directly to Oracle. It emulates the Oracle SQL *Net, Net8, and TTC adapters using its own TCP/IP-based Java socket implementation. The JDBC Thin agent-side driver does not require Oracle agent software to be installed. However, it does require the server to be configured with a TCP/IP Listener.
- Make sure you have sufficient licenses for Oracle Server, as this is a separate product from Director and is not included in Director licensing requirements.

Oracle Server login access

If you do not have a User ID, one is created during the Database Configuration process. In addition, a role (TWG_ROLE) is created. The User ID is defaulted to use the table spaces that are created and TWG_ROLE for security.

Configuring the Oracle TCP/IP Listener

The Oracle TCP/IP Listener must be configured and started before you run the Database Configuration function.

Using unlimited rollback segments (Oracle Server version 7.3.4 only)

If you are running Oracle version 7.3.4, you must edit the initdirector.ora file in /opt/oracle/admin/director/pfile to allow the use of unlimited rollback segments (where director is the instance name). Add the following line: unlimited_rollback_segments = true

Log into Oracle and issue a shutdown and startup before attempting to run the Oracle Database Configuration dialog.

Setting the compatible parameter (Oracle Server version 7.3.4 only)

If you are running Oracle version 7.3.4, the COMPATIBLE parameter must be set to 7.3.0.0 or greater. To set this, edit the inidirector. ora file in /opt/oracle/admin/director/pfile

where director is the instance name. Uncomment the following line:

```
# compatible = "7.1.0.0"
```

and change it to:

```
compatible = "7.3.0.0"
```

Log in to Oracle and issue a shutdown and startup before attempting to run the Oracle Database Configuration dialog.

Creating the Oracle Server database

Your database administrator may choose to create the tablespaces manually, or allow the tablespaces to be created automatically during Director server installation. Your database administrator should consider the following:

- If you do not have a User ID, one is created during the Database Configuration process.
- The administrator ID allows the Database Configuration process to create the tablespaces and roles, as well as assign defaults for User ID and password. However, administrator information, such as its User ID and password, are not saved.
- When the tablespaces are created automatically, they will present default values.

An initial size of 100 MB is recommended for the database to hold data for 250-500 managed systems. More space may be required if you are managing more systems or if your software inventory data is extensive. If the Oracle tablespace defaults are not correct for your needs, then the database administrator can either modify the default values or create the tablespaces manually. If the tablespaces are created manually, they must be entered on the tablespace panel to be used. Your database administrator should monitor the tablespaces and adjust their size as needed.

Whether the tablespaces are created manually or automatically, your database administrator should provide the Oracle TCP/IP Listener Port, Host Name, and System Identifier (SID). You are now ready to proceed with Director Management Server installation.

Additional considerations

Depending on the devices you will be managing, one or more of the following may also apply to your network:

Novell NetWare Managed Systems

If you are managing systems running under Novell NetWare, the Internetwork Packet Exchange (IPX) networking transport must be installed and configured.

Systems using NetBIOS

If you are managing systems that use the NetBIOS networking transport, NetBIOS must be installed and configured.

SNMP devices

If you are using Director to manage Simple Network Management

Protocol (SNMP) devices, and you have not installed and configured the Windows NT SNMP service, you must seed SNMP with the IP address and subnet mask of an SNMP entity. For information on setting SNMP discovery parameters, see Chapter 14, “SNMP Management,” on page 153.

Web server

If you are using Director to manage a Web server, the Microsoft Peer Web Server, Microsoft Internet Information Server, or Netscape FastTrack or Enterprise Web Server must be installed. The server must have access to the file system of the Director server before the Director server software is installed.

News and Mail servers

If you plan to use Director to post event information to a news group, you need to install a Network News Transfer Protocol (NNTP) server. If you plan to send this information via e-mail, you need to install a Simple Mail Transfer Protocol (SMTP) server.

Message paging

If you plan to use Director to send event information to a user using the paging action, you will need modems installed and operational.

Wake On LAN

Director supports Wake On LAN, an advanced power management feature on many of today’s systems. If this feature is enabled during the Tivoli Management Agent® installation procedure, Director can send a “magic packet” to a managed system that is powered off. The packet is decoded by the system interface and the system is initialized, which usually causes the system to boot itself automatically into an operating system.

Wake-On-LAN support enables you to perform remote maintenance on a system, even when it has been turned off or powers itself off with its power management software. Wake-On-LAN is also used to control automatic server systems that are powered on for a specific function and then powered off by the power management software.

To use the Wake-On-LAN feature, a managed system must have a network card installed that supports it.

Discovering managed systems over bridges and routers

If you are using TCP/IP and are having problems discovering Director agents that reside across a bridge or router, make sure that all bridges and routers that you intend to do discoveries across do not block broadcast transmissions for port number 14247. Likewise, if you are using IPX, make sure that port 4490 (hex) for read and port 4491 (hex) for write are not blocked.

Director Management Console

The Director Management Console is installed when you install the Director server but it can also be installed independently almost anywhere in your network. You can operate multiple Management Consoles concurrently and a Management Console can coexist with other applications running on the same system.

Refer to the product README file for the minimum hardware and operating system requirements for the Director Management Console.

Director management agent

The Director management agent contains the executable files required to perform tasks on systems managed by the Director server.

To enable communication with the Director server, the managed system must have one of the following network transports installed. See "Transport support" for information on the following supported versions of transports:

- TCP/IP
- NetBIOS
- IPX
- SNA

Note: For the Unix server and agent, TCP/IP is the only protocol available for use.

Operating-system platform support

Please refer to "Supported platforms for Director 3.1" on page 40 for the operating systems supported by the three main components of Director.

Transport support

The Director server communicates with the Director Management Console using TCP/IP only. You can use TCP/IP, NetBIOS, SNA, or IPX to establish communication between the Director server and a managed system.

Supported transport software is *not* included as part of Director; the transport must already be installed. The following table lists support by protocol.

Note: For the Unix server and agent, TCP/IP is the only protocol available for use.

Protocol	Supported Versions
TCP/IP	All WinSock-compatible versions of TCP/IP supported by Windows 98, Windows NT Server 4.0, Windows NT Workstation 4.0, OS/2 Warp Server for eBusiness, NetWare 5.0,5.1, 6.0, Linux, Sco UnixWare
NetBIOS	Native NetBIOS versions supported by Windows 95, Windows 98, Windows NT Server 4.0, Windows NT Workstation 4.0, and Warp Server for eBusiness 4.5
IPX	IPX versions supported by NetWare 5.0, and 5.1, Windows 95, Windows 98, Windows NT Server 4.0, and Windows NT Workstation 4.0
SNA	**Windows NT:** Microsoft SNA 4.0 with Service Pack 1 Microsoft SNA 3.0 with Service Pack 2 IBM Communication Server 5.0 for Windows NT IBM Personal Communications (PCOMM) 4.2 or later for Windows NT **OS/2:** IBM Communications Server 5.0 for OS/2 IBM Personal Communications (PCOMM) 4.2 or later for OS/2

Discovery

Director discovery operates by sending out a discovery request from the server and then listening for responses from any Director agents. Agents listen for this request and then reply to the server that sent the request. Four distinct kinds of discovery can be used:

- Broadcast discovery

  Broadcast discovery sends out a general broadcast packet over the local area network. The destination address of this packet depends on the particular protocol used to communicate with the managed systems. For TCP/IP systems, for example, the destination address for the packet is 255.255.255.255. Thus, the server will discover any agents which can be reached by the broadcast packet.

  Broadcast discovery can also send out a broadcast packet to specific subnetworks by adding a discovery seed address. If you enter the IP address and subnet mask for a system in the subnet for which discovery is to be performed, Director will send a broadcast packet to that specific subnet. All agents on that subnet will be discovered.

- Multicast discovery

  Multicast discovery operates by sending a packet to the multicast address. Director uses 224.0.1.118 as the multicast address. Agent systems listen on this address and respond to the multicast from the server. Multicasts are defined with maximum Time to Live (TTL), and after the TTL expires, the packet is destroyed.

  Multicasts are useful for networks that filter broadcasts but, do not filter multicasts. Multicasting applies only to TCP/IP systems.

- Unicast discovery

  Unicast discovery sends a directed request to a specific address or range of addresses. This method can generate significant network traffic but is useful in networks where both broadcasts and multicasts are filtered.

  In many cases, Remote Access Servers (RAS) do not forward any broadcast traffic. To discover certain types of managed systems (for example, dial-up systems), it may be necessary to use unicast discovery. Unicast discovery is only available for TCP/IP systems.

- Broadcast relay agents

  Broadcast relay allows the server to discover TCP/IP and IPX agent systems when the systems are not directly reachable by broadcast packets due to network configuration. This situation can occur in networks where the server and agent are in separate subnets, and the network between them does not allow broadcast packets to pass from one subnet to the other. This option generates less network traffic than Unicast discovery and avoids many of the problems associated with filtered broadcasts.

  In broadcast relay, the server sends a special discovery request message to a particular agent, instructing the agent to perform a discovery on the local subnet using a general broadcast. When agents on that subnet receive the discovery request, they reply to the server that made the original request.

The server performs all types of discovery simultaneously. Enter as many broadcast, broadcast relay, or unicast addresses as needed to discover managed systems by selecting **Options →Discovery Preferences →System Discovery (IP)**. The Addressing Properties pane of this tab allows entry of the IP addresses and subnet masks and shows a list of existing discovery filters. The Properties pane shows the discovery settings for the local network. For detailed information on configuring system discovery preferences, refer to the online help.

Migration support

If you are upgrading your network to Director Version 3.1 and you intend to continue using previous version(s) of Director components (server, console, or agent), see Chapter 4, "Upgrading Director and IBM Director Extensions," on page 73 for the upgrade paths.

Security

To protect your network from unauthorized access, Director implements two levels of security: user-logon security and agent-server security. *User-logon security* is the user ID/password verification process supported by the operating system and used to validate users of the system. *Agent-server* security is an authentication process used to establish trust relationships between the Director server and Director agents when the network is brought up. This section describes user-logon security, which you need to establish immediately after installing the Director server and IBM Management Console. Agent-server security is described in Appendix E, "Agent-server security," on page 417.

Director user-logon security

Director provides multilevel console security that enables you to define and edit user IDs and specify access privileges for each user ID. Using the Console Security feature on the Director Management Console, you can:

- Add, edit, and delete user IDs
- Define general access privileges for each user ID
- Define group access and task access privileges for each user ID
- Manage authorization privileges of Windows NT users.

To set up user-logon security for your network, click **Options →Console Security** on the Director Management Console.

Director console authorization

Authorization to the console can be administered through user management facilities of the underlying operating system, or through the Console Security function of the management console. The Console Security function can manage console authorization for users that are defined to the operating system as well as users that are not defined to the operating system. For users who are defined to the operating system, the following procedures are used to control authorization to the console.

For console login with basic administrator authority:

- On systems running Windows NT, the user must be a member of the Administrators group or the TWGAdmins group.
- On systems running UNIX®, the user must be a member of group *root* or group *tdadm*.

For console login with superuser authority (authority to administer console users through the Console Security function):

- On systems running Windows NT, the user must be a member of the Administrators group or the TWGSuperAdmins group.
- On systems running UNIX, the user must be a member of group *root* or group *tdsupadm*.

Director accounts

Because the Director server runs on an operating system that already has account administration defined, it is necessary to recognize and support the accounts already defined for that system. These accounts are referred to as *native user accounts*. Native user accounts are recognized by Director but not administered by Director. This means that Director does not edit the user information for that account (such as changing the password or the user description) but you can modify the Director-specific information. To add or remove those specific accounts or to change the password, use the specific operating system user administration function.

Director also provides the capability to create accounts for which the Director server handles the administration. These accounts are called *non-native user accounts*. These accounts do not appear on the operating system user lists because they are defined only to Director. All administration of these accounts is done through the Director Console Security task.

Listing Director users

When you start the Director Console Security task, the window that opens shows a list of all users who are authorized to log in to the Director server. The main information for each user is presented here, including name, full name, description, and whether the user is currently logged in.

You can also look at a list of all unauthorized server users. These are users which have accounts on the native operating system server but have not been given authorization to access Director. To view these accounts, select **User → Show Unauthorized Server Users**. The task window creates a split window, showing the user information on top pane and the unauthorized user list in the bottom pane.

Note: All accounts on the server with Administrator authority are automatically given authorization to access Director.

Creating a new user

You can create a new Director user by performing one of the following methods:

- Select the **New User** option from the main menu or right-click the User Information table to bring up a context menu and click the **New User** option. Because Director needs the user ID and password information to create an account, the User Editor appears to allow you to enter this information.
- Right-click a user listed in the Unauthorized Server Users table and select the **Authorize User** option. This option creates an Director account for this user

using the server information and the current set of user defaults. Because Director already has the user ID and password information, the account is automatically created without presenting a dialog. To change any of the information from the defaults, just Edit the account after creation.

The accounts created are initialized with the default information that is defined by the User Defaults template. If you need to create a lot of accounts with the same types of authority or access, it is recommended that you first update the User Default settings with the authority or access you desire.

Director allows you to set up each individual user with specific information. This information is specified from within the User Editor window. This dialog is presented as a tabbed panel and contains four separate pages of information that you can modify.

The first page is the User Properties panel. It contains the general information about the user, including the user ID and password information. In order to create a new user, you must specify a unique user ID and provide a password. Optionally, you can specify the user's Full Name, Description, Mail Address, and Pager information. Check the **Superuser authority** check box if you wish to make a user a superuser, which grants the user full authority on Director (all privileges, access to all groups and tasks).

The next page is the Privileges page. Privileges govern the authority to perform specific kinds of activities on the system. The default privileges provided by the default user template grants all Director-supplied privileges except for modifying the cluster settings, database configuration, and the ability to perform user account administration. You can grant privileges to a user by dragging the privilege from the Available Privileges side of the list and dropping it into the Privileges Granted to User, or you can just select one or more privileges in the Available Privileges side of the list and press the **Add** button. To remove privileges from a user, just select the privileges you wish to remove from the Privileges Granted to User side of the list and press the **Remove** button.

The third page is the Group Access page. The settings here govern which groups a user can access (for example, which ones will appear on the Director Management Console in the Groups pane). The default settings provided by the default user template grants access to all groups. If you want to allow access to all groups but do not want the user to have the capability to create new groups, select the **Limit user to read-only access of groups** check box. If you wish to limit which groups the user can access, select the **Limit user access only to the groups listed** check box. This enables the panels below, showing all of the groups to pick from in the Available Groups section. Drag the groups that the user should be allowed to access over to the Groups User Can Access section, or select them in the Available Groups section and press the **Add** button. To remove access to certain groups from a user, select those groups in the Groups User Can Access section and press the **Remove** button. Note that when you limit the groups a user can access, the user is automatically prevented from creating his own groups.

The last page is the Task Access page. The settings here govern which tasks a user can access. The default settings provided by the default user template grant access to all tasks. If you want to limit a user to specific tasks, select the **Limit User Access Only to the Tasks Listed** check box. This enables the panels below, showing all of the tasks available in the Available Tasks section. As with the Group Access page, select items and press **Add** or **Remove** or perform drag and drop actions to set up the Tasks User Can Access section with the tasks you want to allow the user to perform.

Editing user accounts

To edit an existing Director user, right-click on the **User Information** table on the user you want to edit. This will bring up a context menu that contains the Edit option. Alternatively, you can select a user in the table and then select **User → Edit** from the main menu.

When the user editor is shown, you can modify the user attributes presented. These are described in “Creating a new user” on page 27. Note that native users with Administrator authority on the native server are automatically granted superuser authority. These accounts cannot be edited except to provide email and pager information.

Changing user defaults

Director provides a default template of attributes that is used to set up new user accounts when created. You can modify this template by clicking **User → User Defaults** on the main menu. This brings up a dialog that looks similar to the User Editor, allowing you to set up the default settings for the users that are created from this point on, until changed again. If you are setting up two types of users, first set the template for one type of user, create those users, and then modify it for the second type of user and create those users. Using the User Defaults editor will make your job easier if you are setting up a lot of accounts. For more information on each of the pages in the editor, see “Creating a new user” on page 27.

Changing user passwords

To change a user's password, edit the user account and type in the new password in both the password field and the confirm password field. You can only change the password for non-native accounts. If you try to change the password on a native account, you will see that the password fields are missing when you bring up the editor. To change the password on native accounts, use the user editor on the operating system.

Deleting user accounts

To delete an Director user, right-click the user in the User Information table and select the **Delete** option. If confirmation is turned on, you are prompted if you want to delete the user. If you answer yes, the account is deleted. If you delete a non-native account, it is removed. If you delete a native account that is not an Administrator account on the native system, the account becomes an unauthorized server account. You cannot delete a native user that has Administrator authority.

Planning for Director tasks

This section describes concepts, setup and usage considerations, and usage restrictions for Director tasks, such as Software Distribution, Remote Control, Event Management, and Inventory Management.

Software Distribution

This section describes methods of software distribution and the limitations that apply to various distribution scenarios, and operating systems.

Director supports the following to help optimize the use of network resources in distributing software distribution packages:

Redirected distributions
You can distribute packages using redirected distribution in two ways. If a package is from a UNC-based or FTP-based share, then you can copy the contents of a package from that share to the local managed system. If a package is placed on a UNC-based server share, then the package can be installed in the managed system directly from that share.

Streaming
You can stream (copy) packages directly from the server to the managed system.

Network resource allocation
You can limit the number of systems you distribute at once, as well as limit the network bandwidth that you use to distribute to those systems.

Redirected and streamed installations

This section describes the methods you can use to install software distribution packages through Director.

Distributing packages using redirection

Many of today's software packages are tens or hundreds of megabytes in size. Distributing software of this magnitude across a large network can cause bottlenecks in network data transmission. To help alleviate this problem, Director takes advantage of the standard file sharing feature by enabling you to set up a share (shared subdirectory) on a server in your network. A share is any location defined by a file distribution server. This product support UNC-based and FTP-based file distribution servers and does not require the installation of the Director server or Director management agent software. When the share is established, large software packages can be distributed by sending most of the package to the share. The managed system received only the minimum of installation code needed to access the share and install the software from the Director server.

This method, known as a redirected installation, greatly reduces the software distribution traffic in your network, and is the recommended method. This document does not describe how to set up server shares; refer to your server documentation for procedures on setting up a shared subdirectory on a server in your network. The share should allow full read/write access to the Director server and allow read access to all potential target systems. Refer to "Configuring Director to use file distribution servers" on page 67 for information on configuring the Director server to use file distribution servers.

Redirection limitation: If a redirected installation of a software distribution package is interrupted, for example, if the connection is lost, the installation must be started over.

Distributing packages using streaming

Streaming is the copying of a file package to a managed system. If no file distribution server shares are defined, streaming will occur.

If a server share is configured, Director attempts to use it. By default, if a managed system cannot access the share, the package is streamed directly to the managed system. However, you can override the default so that redirected distribution will fail. To do this, select the redirected distribution option. Do not stream distribution if redirected distribution fails. If you have multiple shares defined, Director tries to use each share before streaming the packages directly to the managed system. If a managed system can access the share and you have configured Director to always stream (copy) to the systems from the server share, the package is first sent to the share and then copied to the target systems set up to use that share.

In some cases, you might prefer to stream the entire software distribution package, either from the Director server or from a server share, to a managed system, for example:

- You might have an unreliable or slow network link.
- You might have a mobile dial-up managed system.

If a network connection is broken during a redirected installation, you must restart the installation. If a network connection is broken during a steamed installation, Director attempts to resume the connection from the point at which the transmission was interrupted. If the streaming operation can be resumed, retransmission time is saved. Refer to "Always streaming software to the managed system" on page 32 for more information on specifying streaming to Director.

Streaming limitation: Streaming requires that the directory on the target system have sufficient free storage to receive the entire package and use the temporary space required during installation. To ensure a successful streamed installation, allocate disk storage equivalent to twice the size of the software distribution package.

Memory and storage management for redirected installations

Software distribution treats file distribution server shares as a software package cache. A software package cache is a storage location, in this case a share, for software distribution packages. Once a package has been cached on a share, the cached package can be reused for future distributions, except in cases noted below. Use of a cached package can decrease the amount of time required to distribute a package through a redirected install. The amount of time saved varies, but generally, the larger the package the greater the savings.

Management of the cache is done entirely through the Director server. A software package is only cache on a share when the package is distributed, not when the package is created. If a software package is edited and saved, the cache entry is removed for any share where the managed system only evaluates shares if you choose to restrict distribution to only the shares in the list of the managed system, and one or more of those shares is in the list of the server. If you do not restrict the share preference list of the managed system, it can evaluate the shares in the list of the server that are not in its list. To restrict the list, do the following:

1. Under Distribution Preferences, define a subset for this managed system.
2. Set the Configuration option to **Restrict share selection to list.**

A managed system evaluates shares by trying to access them. If shares are accessible, the managed system identifies those shares to the Director server. From this list, the server chooses a share to act as a package cache and notifies the managed system which share is used for the distribution. The server share used to stream the package is evaluated the same way a share is evaluated for a redirected distribution.

Always streaming software to the managed system

To force the streaming of software distribution package for an individual managed system or group, you can select to always stream a package from the Director server. Refer to “Configuring distribution preferences for managed systems” on page 69 for instructions on accessing the appropriate option through the Management Console.

Specifying the transport for server shares

If a server to which the server share is set up is also configured as an FTP server, you can specify to use FTP when transferring packages from the Director server to the share.

Note: For OS/2, FTP is supported only for file transfer between the Director server and a server share. FTP *cannot* be used to distribute a software package from a server share to an OS/2 managed system.

An FTP server must be running on the file distribution server and a user ID and password that grants read and write access to the FTP server must be defined. Optionally, for OS/2 and Windows managed systems, the directory where the package is put can be shared and the targeted managed systems must have read access to the share. FTP is used to copy the package's contents to the remote file distribution server share. For OS/2 systems and optionally for Windows systems, the home directory for the FTP login should be the same directory as the file distribution server. (The home directory is not required for other supported platforms.) For example, if c:\stuff\swd_share is mapped to \\server\swd_share, then c:\stuff\swd_share should be the home directory for the FTP user ID login used on the FTP file distribution server configuration screen. Refer to "Configuring distribution preferences for managed systems" on page 69 for instructions on specifying the FTP protocol to Director.

Limitations on software distribution

This section lists the software distribution restrictions that you should review before you attempt software distribution in your network.

Limitations on software distribution to managed systems

The following restrictions apply to both streamed and redirected software distributions to managed systems:

- Director management agents for SCO UnixWare and NetWare do not support the software distribution task.
- To distribute a software package that uses InstallShield to a Windows NT® 4.0 managed system, the target system must have Service Pack 4 or greater installed.
- To distribute a software package to a FAT-based drive on an OS/2 managed system, all files within the package must have an 8.3 filename format.
- To distribute a software package over a WAN to a managed system on the other side of a firewall, TCP/IP session support must be disabled for that system. Disable session support by creating a tcpip.ini file in the

  ```
  \tivoliwg\bin
  ```

  directory of the agent system. This .ini file must contain the following line:

  ```
  session_support=0
  ```

Note: If more than one TCP/IP option is listed in the agent's Network driver Configuration panel, create a tcpip.ini file for each entry. The file-naming scheme should be tcpip.ini, tcpip2.ini, tcpip3.ini, and so on. After creating the appropriate files, restart the agent system or stop and restart the IBM Director agent.

Limitations on redirected installations

The following restrictions apply to using redirection:

- To distribute a software package from a file distribution server on Windows NT to a Windows 98 managed system that does not have a logon session (no one is logged on to the target system), you must first run TWGSHAREon the file distribution server. Refer to "Enabling UNC-based share access to Windows managed systems" on page 67 for instructions.

- To distribute a software package to an OS/2 managed system using redirection, the target system must have a logon session (a user ID must be logged on to the system).

Configuring security for UNC-based server shares

To access server shares, the Director management agent passes credentials (user ID and password) to the server where the share resides in order to gain security access to the share. The credentials used to access the share are determined by the security context (account) the agent running in. You must configure security on the server where the share resides to authorize Director management agents to access it with the credentials supplied. The credentials used by the Director management agent are determined as follows:

- On Windows NT, the Director management agent runs as a service that logs on to the account configured for the service. The default is the system account, which causes null credentials to be used to access server shares. You can change the account used by the service at installation time or by selecting the Services icon from the Windows NT Control Panel folder.
- On Windows 95 and Windows 98, the Director management agent runs under the security context of the user currently logged on to the system. When a user is logged on to the system, the user's credentials are used to access server shares. When no user is logged on, null credentials are used to access server shares.

When Director management agents use null credentials to access a server share, the server share must be configured to allow null credentials. The TWGSHARE command can be used to configure a share residing on Windows NT to allow null credentials. Refer to "Enabling UNC-based share access to Windows managed systems" on page 67 for information on TWGSHARE.

Note: NetWare servers, and OS/2 servers do not support access to shares using null credentials.

You can now specify a user ID and password to access server shares via Distribution Preferences. For more information on configuring distribution preferences for managed systems, see Chapter 3, "Installation and configuration," on page 39.

Limiting network resources for Software Distribution

You can control the dedication of network resources to software distributions by:

- Using redirection, where practical, to perform Software Distributions
- Limiting the number of concurrent redirected distributions
- Limiting the number of concurrent streamed distributions
- Limiting the bandwidths used to stream (copy) packages from the Director server to managed systems
- Limiting the bandwidth used to stream (copy) packages from the Director server to file distribution servers

- Limiting the bandwidth used to stream (copy) packages from file distribution servers to managed system

Redirected Software Distribution is designed to minimize the network bandwidth dedicated to a package installation. If the Director server puts a software package on a server share, managed systems can be configured to use that share. The number of managed systems installing the software package at one time does not exceed the limit defined for the maximum number of concurrent users. Therefore, other managed systems are queued and distributions occur as active managed systems finish. Refer to "Defining the maximum number of concurrent redirected distributions" on page 70 for instructions on setting the maximum number of concurrent distributions.

You can control the dedication of network resources to a Software Distribution streaming operation by limiting the number of concurrent streaming distribution and by limiting the amount of bandwidth that can be dedicated to a streamed package transfer. You can limit the streaming bandwidth for an individual managed system or group and for all streaming operations from the Director server. If you set a bandwidth limitation for all managed systems and for a specific managed system or group, the lowest bandwidth setting is used for streaming to the managed system.

Refer to "Defining the maximum number of concurrent streamed distributions" on page 70 for instructions on limiting concurrent streamed distributions. Refer to "Defining the maximum number of concurrent redirected distributions" on page 70 for instructions on limiting the bandwidth for all managed systems . Refer to "Configuring distribution preferences for managed systems" on page 69 for instructions on specifying the bandwidth for a managed system or group.

Remote control

This section lists the restrictions and conditions that apply to using remote control. Refer to Chapter 7, "Remote Control," on page 103 for information on using remote control.

- The remote control task can be performed only on native managed systems running under the following operating systems:
  - Windows NT 4.0
  - Windows 98
  - Windows ME
  - Windows 2000
  - Windows XP
  - OS/2 WARP® Server for eBusiness
- You can concurrently monitor or control two or more remote systems from a single Director Management Console.
- If multiple Director Management Consoles are connected through the same server to a remote system, only one console can send keyboard and mouse information to the remote managed system.

- Within the overall network, multiple Director Management Consoles can remotely control multiple managed systems concurrently; however, the overhead load generated can cause system response to degrade significantly.
- Only one Director server can communicate with a remote system through remote control. If more than one Director server attempts remote control communication, the communication is rejected and an error message is displayed on the Director Management Console from which the communication is initiated.
- Do not use remote control over a slow connection; when large amounts of data are transferred, they require greater network throughput than slow connections can accommodate.
- To reduce the amount of data transferred from a remote system, remote control reduces the display information of all images to 16 colors. As a result, the image displayed on the management console can differ from the image displayed on the remote system desktop.
- Remote control does not support full-screen graphic modes, including Win-OS/2 full screen graphics mode. You cannot use remote control foe such tasks as playing graphic-intensive games from a remote console.
- Certain keyboard restrictions apply; refer to "Sending keyboard information to a remote system" on page 105.
- To start a remote control session over a WAN on a managed system that is on the other side of a firewall, TCP/IP session support must be disabled for that system. Disable session support by creating a tcpip.ini in the \tivolig\bin directory of the agent system. This.ini file must contain the following line:

  ```
  SESSION_SUPPORT=3
  ```

  Note: If more than one TCP/IP option is listed in the agent Network Driver Configuration panel, create a tcpip.ini file for each entry. The file naming scheme should be tcpip.ini, tcpip2.ini, tcpip3.ini, and so on. After creating the appropriate files, reboot the agent system or stop and restart the IBM Director agent.
- If TME 10™ Remote Control or Director Remote Control has already been installed on a system, the Director server or management agent software can be installed on that system if, during installation of Director, the option to install remote control is disabled.
- If the NetWare IPX agent software has been installed on an OS/2 system, the Director management agent software can be installed on that system if, during installation of Director, the option to install remote control is disabled.
- Logging in to a remote system through remote control requires that the Require user authorization for screen access setting is disabled on the remote system. If this setting is modified on the remote system to allow remote control, Director must be stopped and restarted for the change to take effect.

Event management

This following sections describe requirements for enabling support for CIM and SNMP events.

CIM event support

The Director event server does *not* automatically detect and present CIM events for filtering. The *SDK* provides information on how to set up managed systems to map CIM events to Director events. When the mapping file is defined, Director can detect and present CIM events for filtering.

SNMP trap support

Director recognizes SNMP traps and generates a corresponding SNMP event if an SNMP trap is sent to the Director server. The Event Type field in the Event Filter Builder window is updated to include the SNMP filtering category if the Director server receives an SNMP trap. You can use this filtering category to create an event filter to respond to SNMP traps. To set up your network to use Director for SNMP trap recognition, configure the network devices that generate SNMP traps to specify the IP address of the Director server as a trap destination.

Following is an example of an SNMP trap event (cold start) entry in the Director event log. The Event Type value will extend as far as MIBs have been compiled. In this example, the text in brackets ([]) indicates the type of information that is included, it is not the actual data.

```
Event Details
Keywords               Values

Date                   16-Nov-1998
Time                   12:01:58 PM
Event Type             SNMP.iso.org.dod.internet.6.3.1.1.5.1
Event Text             Cold Start
System Name            [name of managed system for which the event was
generated]
Severity               Unknown
Category               Alert
Group Name



Sender Name          [IP address of the source from which the event was
sent]
1.3.6.1.6.3.1.1.4.3.0 [snmpTrapEnterprise.0]
```

Inventory management

Director collects inventory information from managed objects and stores it in database tables within the server database. The formats of these tables cannot be changed. With the addition of inventory collectors for the extensible sources CIM and DMI, and from static MIF files, some facility for allowing the end user to define custom tables became necessary.

Our approach to solving this problem uses property files that follow the Java property file format. These property files describe the contents of a custom database table. The property files, one per table, contain the table's name, names and types for each of the columns of the table, and other information.

For detailed information on defining these tables, see Appendix D, "Defining table property files," on page 399.

The inventory database tables are HTML files (one each for each table). The HMTL files now exist in the Help Index under the Inventory component. The online help contains a list of the inventory database tables and a description of the data they contain. Each table has a unique table name that is followed by one or more rows defining the name, type, and description of the data in each table.

Additional columns of provider information are listed, with an "X" in each cell signifying that inventory data can be obtained from the provider.

Some fields will be identified with the term ENUM. This signifies that data returned in these fields will consist of one of several specific text strings. For each data item identified with ENUM, a list of the valid text strings values is shown immediately after the table.

DB2 has the following limitations:

- Database CHAR columns are limited to 254 characters.
- Table names are limited to 17 characters.
- Field names are limited to 18 characters.
- All keys combined cannot exceed a 254-byte limit. (Therefore, the INSTALL_PATH column of the TWG_SOFTWARE table has been shortened to CHAR(154).)
- Because of these limitations, short names are used in databases where these limits apply (for example, DB2 Version 5). These short names have been added to the Table Name heading and the Field Name column of the following tables, where applicable. These short names are included in parentheses, following their standard named.

Chapter 3. Installation and configuration

Director is divided into the following components:

- Director server
- Director management console
- IBM Director Extensions
- IBM Director Agent

See Chapter 2, “Planning,” on page 15 for information about prerequisites for each component before you begin installation. When a prerequisite has not been installed or has been installed at the incorrect level, you may receive an error message informing you that the prerequisite is not present. You can continue the installation; however, the function dependent on that prerequisite may not work or can yield unpredictable results.

Director 3.1 requires administrator authority and the following hardware components and platforms.

Hardware requirements

- Intel® Pentium® class microprocessor, 266 MHz or faster.
- Supported SVGA driver (800x600 resolution).
- Director Console: with all IBM Director Extensions: 128 MB RAM, 160 MB disk space.
- Director Console, Server and Agent with all IBM Director Extensions: 256 MB RAM, 300 MB disk space.
- Director Agent with all IBM Director Extensions:128 MB RAM
- Director Agent with base components require 100 MB disk space. Additional disk space is required for the following Director and IBM Director Extensions tools:
  - Software Distribution: 100 MB
  - Capacity Manager: 251 KB
  - Management Processor Assistant: 1628 KB
  - ServeRAID Manager: 9765 KB
  - Cluster Manager: 683 KB
  - Rack Manager: 8 MB
  - System Availability: 1012 KB
  - Software Rejuvenation: 981 KB

- A network adapter that supports TCP/IP. The adapter must also support NetBIOS, IPX, or SNA, if the transport is needed to communicate with the managed systems.

Notes:

1. If insufficient disk space is available, a message box is displayed and the setup is stopped.

2. For Director installation, msvcrt.dll version 6.0 or later is required. This requirement primarily affects systems running Windows NT® and Windows 98 that have not had Microsoft software installed or updates to Internet Explorer. Microsoft provides a package with the necessary update to the msvcrt.dll. The package is named vcredist.exe. can be found on the Microsoft support site.

   You must restart your system after the update has been applied.

Supported platforms for Director 3.1

The following table list the operating-system platforms supported by the three main components of Director.

Note: This table applies only to version 3.1 levels of Director components. If you are upgrading from an earlier version of Director, refer to Chapter 4, "Upgrading Director and IBM Director Extensions," on page 73 to view the supported upgrade paths.

Operating System	Server	Console	Agent
Windows NT 4.0 Server	Yes	Yes	Yes
Windows NT 4.0 Enterprise	Yes	Yes	Yes
Windows NT 4.0 Workstation	No	Yes	Yes
Windows NT Terminal Server	No	Yes	Yes
Windows NT Metaframe	No	Yes	Yes
Windows 2000 Metaframe	No	Yes	Yes
Windows 98	No	Yes	Yes
Windows 98 SE	No	Yes	Yes
Windows Millennium Edition (Me)	No	No	Yes
Windows XP Professional	No	Yes	Yes
Windows 2000 Server	Yes	Yes	Yes

Operating System	Server	Console	Agent
Windows 2000 Professional	No	Yes	Yes
Windows 2000 Advanced Server	Yes	Yes	Yes
Windows 2000 DataCenter	No	No	Yes
Netware 5.0 (Service Pack 1 or later), 5.1, 6.0	No	No	Yes
Warp Server for ebusiness with Convience Pack 1 and 2	No	No	Yes
SCO UnixWare 7.1.1	No	No	Yes
Caldera OpenUnix 8.0	No	No	Yes
Red Hat Linux 6.2, 7.1	No	No	Yes
SuSe Linux 7.1, 7.2	No	No	Yes
Caldera Linux 2.3.1, 3.1	No	No	Yes
Turbo Linux 6.0.5, 6.5	No	No	Yes

Supported platforms for IBM Director Extensions

These IBM Director Extensions Agents are supported on the following operating-system platforms.

Management Processor Assistant

OS/2 Warp 4 Server for eBusiness, Convenience Pack 1 and 2, Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, NetWare 5.0 (Service Pack 1 or later), 5.1, 6.0, SCO UnixWare 7.1.1, Red Hat Linux 6.2, 7.1, SuSe Linux 7.2

Capacity Manager

Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, Windows Advanced Server through Service Pack 2, NetWare 5.0 (Service Pack 1 or later), 5.1, 6.0, Warp Server for eBusiness with Convenience Pack 1 and 2, Red Hat Linux 6.2, 7.1, SuSe Linux 7.1, 7.2, Caldera Linux 2.3.1, 3.1, Turbo Linux 6.0.5, 6.5

Rack Manager

OS/2 Server for eBusiness Convenience Pack 1 and 2, Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, Windows Advanced Server through Service Pack 2, NetWare 5.0 (Service Pack 1 or later), 5.1, 6.0, SCO UnixWare

7.1.1, Red Hat Linux 6.2, 7.1, SuSe Linux 7.1, 7.2, Caldera Linux 2.3.1, 3.1, Turbo Linux 6.0.5, 6.5, Caldera Open Unix 8.0

ServeRAID

Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, Windows Advanced Server through Service Pack 2, NetWare 5.0 (Service Pack 1 or later), 5.1, 6.0, SCO UnixWare 7.1.1, Caldera OpenUnix 8.0, Red Hat Linux 6.2, 7.1, SuSe Linux 7.1, Turbo Linux 6.5, Caldera Linux 3.1

Software Rejuvenation

Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, Windows Advanced Server through Service Pack 2, Red Hat Linux 6.2, 7.1, SuSe Linux 7.1, 7.2, Caldera Linux 2.3.1, 3.1, Turbo Linux 6.0.5, 6.5

System Availability

Windows 2000 DataCenter, Windows NT 4 Service Pack 4 or later, Windows 2000 through Service Pack 2, Windows Advanced Server through Service Pack 2, Red Hat Linux 6.2, 7.1, SuSe Linux 7.1, 7.2, Caldera Linux 2.3.1, 3.1, Turbo Linux 6.0.5, 6.5

Installing Director

To install the Director components (Console, Agent, and Server), use the following procedure:

1. Place the *Director CD* in the CD-ROM drive of the system to which you will be installing. If autorun is enabled, the Director Agent and Director Extensions window opens.

   Note: If your system is not configured for autorun:

   a. Click **Start** →**Run**.
   b. In the **Open** field, type

   ```
   e:\setup.exe
   ```

   where *e* is the location of the CD-ROM drive. (The Welcome window opens.)

2. Select **Install Director**. The Welcome window opens.

3. Click **Next** through the Welcome window and accept the License Agreement.

   The Select Components window opens.

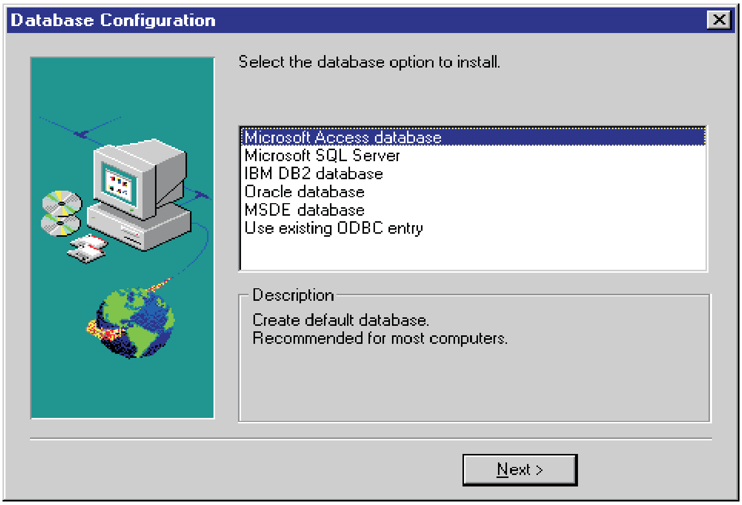

Select the database option to install.

- Microsoft Access database
- Microsoft SQL Server
- IBM DB2 database
- Oracle database
- MSDE database
- Use existing ODBC entry

Description

Create default database.
Recommended for most computers.

The four choices from the Select Components window are:

Server Select this choice to install the server, console, and agent.

Console Select this choice to install the remote user interface for the Director Servers.

Agent Select this choice to install the agent on the systems to be managed.

Workgroup/Enterprise Integration
Select this choice to install IBM Director Agent integration for management environments.

Installing the Server

Selecting **Server** from the Select Components window installs the Server, Agent, and Console. Use the following procedure:

1. Click the **Server** button from the Select Components window.

   The Director Agent Configuration window opens.

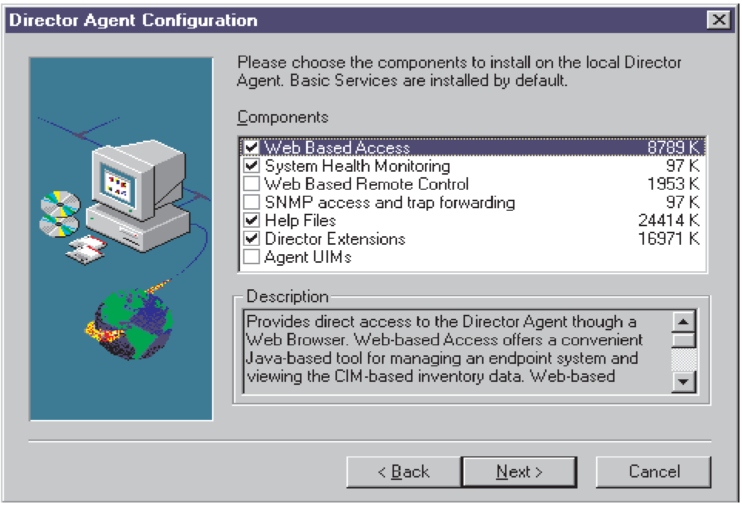

Director Agent Configuration

Please choose the components to install on the local Director Agent. Basic Services are installed by default.

Components

☑ Web Based Access	8789 K
☑ System Health Monitoring	97 K
☐ Web Based Remote Control	1953 K
☐ SNMP access and trap forwarding	97 K
☑ Help Files	24414 K
☑ Director Extensions	16971 K
☐ Agent UIMs	

Description

Provides direct access to the Director Agent though a Web Browser. Web-based Access offers a convenient Java-based tool for managing an endpoint system and viewing the CIM-based inventory data. Web-based

The following optional components are available.

Web Based Access
Web Based Access offers a convenient Java-based tool for managing an agent system and for viewing the CIM-based inventory data. If you install Web Based Access, a hypertext transport protocol (HTTP) daemon is installed and requires that a user name and password be entered during the installation. The user name and password are used to limit access to the HTTP daemon. With Web-based Access installed on the agent system, the agent can be managed from any remote computer with a supported Web browser. No software other than a Web browser is needed on the remote system.

System Health Monitoring
System Health Monitoring provides active monitoring of critical system functions, such as disk space available, system temperature, fan functionality, power supply voltage, and system cover removal. System Health Monitoring enables you to detect system problems early, before system failures occur. System administrators are notified of a system problem by a CIM event, SNMP trap (SNMP traps are available only if SNMP access and trap forwarding is also selected), or SMS Status Message (Microsoft SMS 2.0 only). Critical problems also result in a pop-up message appearing on the display of the agent and a status change in System Health GUI.

Web Based Remote Control
Web Based Remote Control enables a remote systems administrator using a Web browser or MMC console to take control of the agent system desktop, enhancing the administrator's ability to diagnose system problems and troubleshoot the system.

Note: You must install the Web Based Access component to install the Web Based Remote Control component.

SNMP access and trap forwarding

This feature enables CIM information to be accessed from a system using the Simple Network Management Protocol (SNMP). If System Health Monitoring is enabled, this option also enables System Health to forward CIM events as SNMP traps. This component requires that you have the SNMP service (provided with the operating system) installed on the endpoint. If the SNMP service is not installed, the system prompts you to insert the operating system installation media and install SNMP during the IBM Director Agent installation.

Help Files

This component installs online documentation. Do not select this option if you are concerned about disk space or do not need online documentation installed on every agent.

IBM Director Extensions

These tools are used to expand the flexibility and management capabilities of Director.

Agent UIMs

Select this feature to install LANDesk Common Base Agent and Tivoli Management Agent.

2. Select the check box beside any component that you want to install.
3. Click **Next**. If you selected Director Extensions as an install option, then the Director Extensions Install Options window opens.

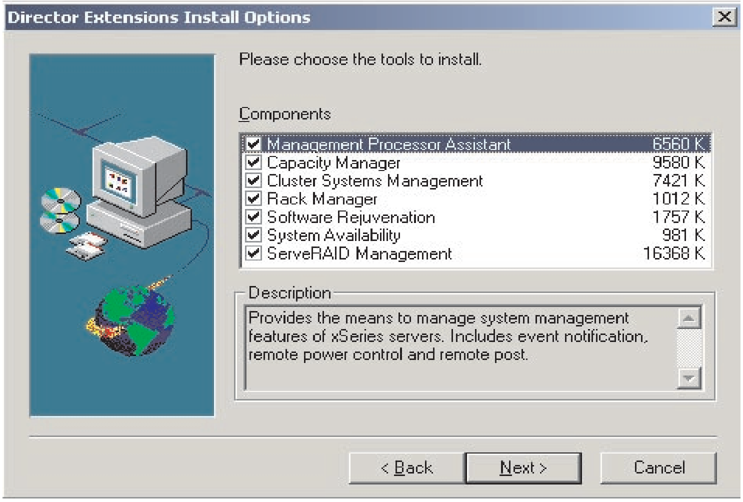

Director Extensions Install Options

Please choose the tools to install.

Components

☑ Management Processor Assistant	6560 K
☑ Capacity Manager	9580 K
☑ Cluster Systems Management	7421 K
☑ Rack Manager	1012 K
☑ Software Rejuvenation	1757 K
☑ System Availability	981 K
☑ ServeRAID Management	16368 K

Description

Provides the means to manage system management features of xSeries servers. Includes event notification, remote power control and remote post.

4. Select the check box beside any component that you want to install.
5. Click **Next**. If you selected Agent UIM as an install option, then the Agent UIM Install Options window opens.
6. Select the check box beside any component that you want to install.

7. Click **Next**.
8. Select an IP port number for the Apache Web server.
9. Click **Next**. The Add icons for IBM Director Agent window opens.
10. Select **Yes** or **No** to place IBM Director Agent icons on the start menu. The Choose Destination Location window opens for the location of the IBM Director files.
11. Click **Next** to accept the default directory (**C:\Program Files\Director\**), or click **Browse** to choose a different directory.

    The **Choose Destination Location** window opens again. This time a directory needs to be specified for a Software Distribution packages creation directory.
12. Click **Next** to accept the default directory (**C:\Program Files\Director\SwDistPk**), or click **Browse** to choose a different directory.

    After creating the Software Distribution packages directory, another **Choose Destination Location** window opens. This directory will be the location for Software Distribution packages that are installed on this system.
13. Click **Next** to accept the default directory (**C:\Program Files\Director\SwDistPkInst**), or click **Browse** to choose a different directory.
14. Select **Yes** or **No** to enable this agent for Remote Control. Selecting Yes will install additional software to enable Director to perform remote control operations on this system.
15. Click **Next**. The system begins installing the necessary files. The Windows Account information window opens. The Domain and User Name for the system you are using is displayed.

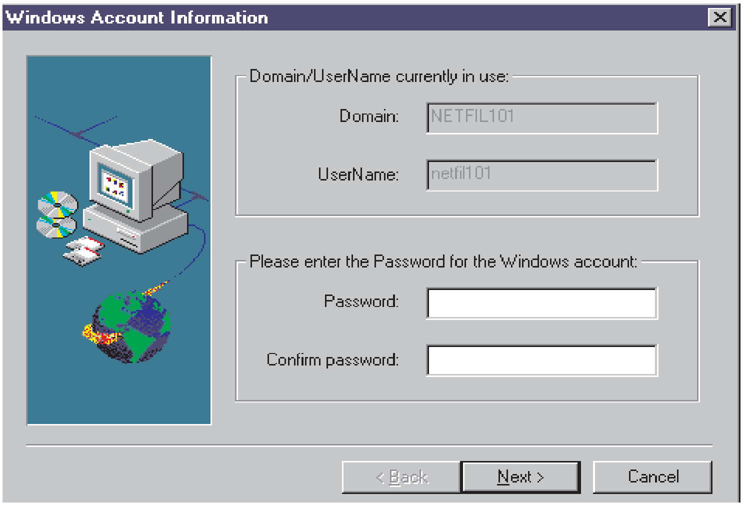
Windows Account Information

Domain/UserName currently in use:

Domain: NETFIL101

UserName: netfil101

Please enter the Password for the Windows account:

Password:

Confirm password:

16. Type your Windows account password in the **Password** field and type it again in the **Confirm Password** field.

17. Click **Next** to continue.

    The Director Database Configuration window opens.

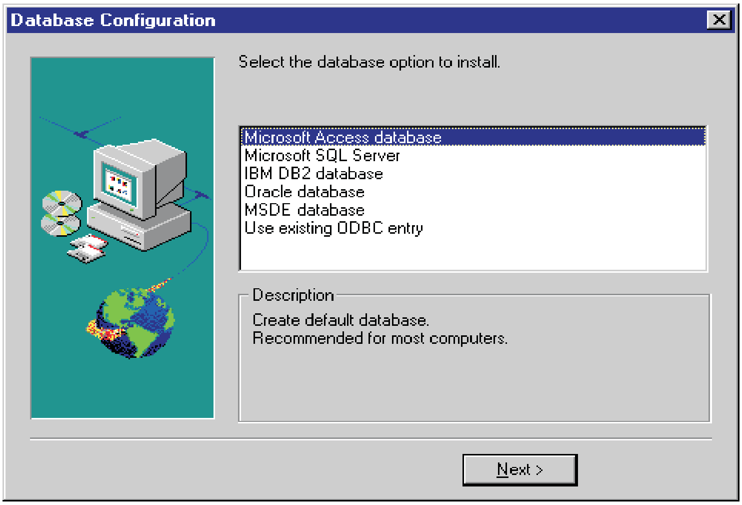

Database Configuration

Select the database option to install.

- Microsoft Access database
- Microsoft SQL Server
- IBM DB2 database
- Oracle database
- MSDE database
- Use existing ODBC entry

Description

Create default database.
Recommended for most computers.

18. Highlight the appropriate database and click **Next**. The options are:

    Microsoft Access
    : Creates a default database. This option is recommended for most users. (In order to use Microsoft Access with Windows NT, install the `mdac_typ.exe` located on the CD if Microsoft data are not already installed. The `mdac_typ` file is located in `director\win32\install\files\nfd\server\windows\files.` Launch this install and follow the directions.)

    MS SQL
    : Creates a link to the MS SQL Server database. A Microsoft SQL Server database must be installed and configured for the network.

    MSDE
    : Microsoft Database Engine must be installed and configured on your network.

    IBM DB2
    : Creates a link to the DB2 database. DB2 agent for server must be installed and configured in the network.

    Oracle
    : Oracle database must be installed and configured.

    Use Existing Database
    : Uses a database from a list of existing data sources.

Note: If you already have a database installed, the system allows you to use that database. The supported database configurations are:

- Microsoft Data Engine (MSDE), Service Pack 3
- Microsoft Access (Jet)
- Microsoft SQL Server 6.5 with Service Pack 5A and 7.0 with Service Pack 3
- Microsoft SQL 2000, Service Pack 1
- IBM DB2 database 6.1, 7.1, 7.2
- Oracle database 7.3.4 through 8.1.7, 9i

The Network Driver Configuration window opens.

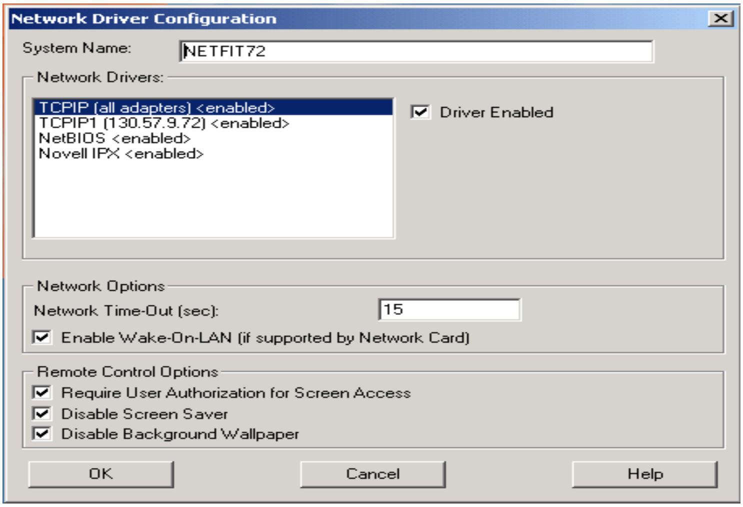

System Name: NETFIT72

Network Drivers:

TCPIP (all adapters) <enabled>
TCPIP1 (130.57.9.72) <enabled>
NetBIOS <enabled>
Novell IPX <enabled>

☑ Driver Enabled

Network Options

Network Time-Out (sec): 15

☑ Enable Wake-On-LAN (if supported by Network Card)

Remote Control Options

☑ Require User Authorization for Screen Access
☑ Disable Screen Saver
☑ Disable Background Wallpaper

The window defines the network transport options for a Director Server. The options are:

- **System Name** - The name of the IBM Director Server.
- **Network Drivers** - The box lists all network transport protocols defined in the system protocol list. They appear as either enabled or disabled. To enable a network transport for use with Director, click on the driver name and check the **Driver Enabled** box.
- **Network Time-out** (sec) - 15 seconds is the default time-out.
- **Enable Wake On-LAN** - Select this box if the network card supports Wake On-LAN.

- **Disable Screen Saver** - Select this box to disable the screen saver during a remote control session.
- **Disable Background Wallpaper** - Select this box to disable desktop wallpaper while system is being remote controlled.

19. Enable the appropriate network drivers by selecting the driver from the **Network Drivers** list and check the **Driver Enabled** checkbox.
20. You may change the Network Timeout if desired.
21. Select the **Enable Wake On-LAN** checkbox if the agent system has Wake On-LAN capability.
22. Select the **Require User Authorization for Screen Access** checkbox if you want to give agent users the authority to deny the system administrator remote control access to their systems. This option allows users to control who accesses their systems.
23. Select **OK** to continue. The file transfer and the Director Extensions tool installation process begins. When the installation process is complete, the Director Setup is Complete window opens.
24. Restart the computer now or Restart later. If you choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the installation program closes. However, you must restart and log in to the system to begin using Director.
25. Click **Finish**.

Installing the Console

Select the **Console** button from the Select Components window to install the Console files only. When the Director Extensions Install Options Window opens, complete the following steps:

1. Select the check box beside any Director Extensions component that you want to install.
2. Click **Next**.

   The Choose Destination Location window opens.
3. Click **Next** to accept the default directory (**C:\Program Files\Director**), or click **Browse** to choose a different directory.

   The IBM Director Extensions files are installed on the system.

   The Setup is Complete window opens.
4. Restart the computer now or Restart later. If you choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the installation program closes. However, you must restart and log into the system to begin using Director.
5. Click **Finish**.

Installing the Agent

Note: If you are installing Management Processor Assistant, you must first install the Management Processor Assistant driver. If you are installing ServeRAID, you must first install the RAID driver.

Select the Agent button from the Select Components window to install the Agent files only. Complete the following steps:

1. Click **Next**.

   The **Choose Destination Location** window opens.
2. Click **Next** to accept the default directory (**C:\Program Files\ UMS**) or click **Browse** to choose a different directory. The Director Agent Configuration window opens.
3. Select the checkbox beside any component that you want to install on the agent system. All of the components, except Director Support, are described under“Installing the Server” on page 43.

   Director support is an additional configuration option for the agent installation process only. Director is an advanced Intel processor-based workgroup hardware manager, with centralized agent and group management console and server services. Selecting this feature enables the system to be managed in a Director environment by installing a Director Agent on this system.
4. Click **Next**.
5. Select an IP port number for the Apache Web server. Click **Next**. (On systems running Windows 98 and Window Millennium, the User ID window opens. Enter the Administrator′s user ID and password for web access to IBM Director Agent.) The Add icons for Director Agent window opens.
6. Select **Yes** or **No** to add IBM Director Agent icons to your start menu. The Director Remote Control Question window opens.
7. Select **Yes** or **No** to install files for remote control. The file transfer and the IBM Director Extensions tool installation process begins. The Director setup is Complete window opens.
8. Restart the computer now or restart it later. If you choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the IBM Director Agent installation program closes. However, you must restart and log into the system to begin using Director.
9. Click **Finish**.

Installing Workgroup/Enterprise Integration

Select the **Workgroup Enterprise Integration** button from the Select Components window to install the IBM Director Agent integration for management environments. The Integration Selection window opens.

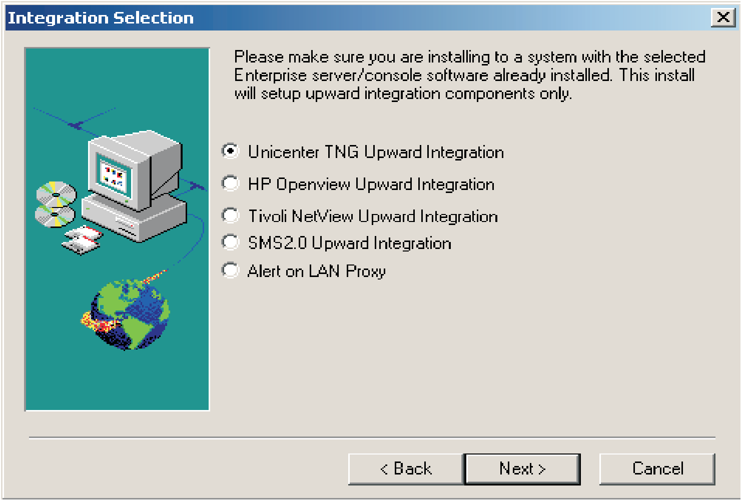

Please make sure you are installing to a system with the selected Enterprise server/console software already installed. This install will setup upward integration components only.

● Unicenter TNG Upward Integration

○ HP Openview Upward Integration

○ Tivoli NetView Upward Integration

○ SMS2.0 Upward Integration

○ Alert on LAN Proxy

The following integration modules are provided:

- Unicenter TNG Upward Integration
- HP Openview Upward Integration
- Tivoli Netview Upward Integration
- SMS 2.0 Upward Integration
- Alert on Lan™ Proxy

An Alert on Lan Proxy is provided in the Workgroup/Enterprise Integration section. The Alert on Lan Proxy application allows the user to set up a system to catch Alert on Lan alerts and send SMNP traps. You can install the proxy on any system on the network. Alert on Lan agents must be configured to send their alerts to the proxy, which can then forward those alerts to other management applications.

For information on the Workgroup/Enterprise Integration options, refer to Appendix K, “Upward Integration Modules,” on page 505.

Installing Director Agents on non-Windows platforms

Director, as a highly integrated workgroup hardware manager, allows the management of a heterogeneous environment through the use of the Director Agent.

The Agent is installed on a Microsoft Windows system as part of the Director installation process. For OS/2, Netware, Linux and SCO UNIX® systems that are managed by Director, use the *Director with IBM Director Agent* CD-ROM to install the Agent.

Installing Director Agent on a system running OS/2

To install the Director Agent software on a system running OS/2, use the following procedure:

1. Insert the *Director CD* into the CD-ROM drive of the system.
2. Change the directory to the *D:\Director\OS2* subdirectory, where *d* is the drive letter of the CD-ROM drive.
3. Run **setup.cmd** to start the installation utility.
4. The default location of the Director Agent files are displayed in the **Directory to Install** from field. Unless you have moved the files to another location, simply accept the default. Otherwise, enter the drive letter and file path location where the agent files are located. A subdirectory,`\SwPkInst` is automatically created in the selected directory, where software distribution packages will be placed on the system for installation.
5. The default target location for the agent files is `c:\TivoliWg\`. If you want to install the files in another location, replace the default drive and file path with the alternative location.
6. Select the **remote control agent** option if you want to enable the desktop of the local system to be taken over from a remote location.
7. Select **Install**. The files are installed in the specified directory. You can cancel the installation at any point by selecting **Cancel**.
8. You should then see the Network Driver Configuration dialog. Enter a name for the system in the System Name field. The Director administrator will use the system name to identify this system on the network.
9. Select one of the available network drivers that the managed system will use to communicate with the Director management server.

   When you select **NetBIOS**, a default network address is assigned. You can change this address, but ensure that the name you specify is 1 to 12 characters in length and unique on the network; otherwise, the managed system cannot start properly. Note that this address is case sensitive.
10. Select **Driver Enabled** to activate the network driver when the system starts. If the system has multiple network drivers available, you can select another driver at this point and repeat the steps for this dialog.
11. The Network Time-Out value specifies the number of seconds the Director management server attempts to establish communication with this system if it is not responding. You may not need to change the default setting.
12. The Require User Authorization for Screen Access option enables you to specify whether a remote user can access, and take over control of, the local system without the permission of the local user. If this option is enabled and a Director administrator attempts to use remote control to access the local system, a message window is displayed on the local system indicating that a remote user is attempting remote control access. You can then allow or disallow access.
13. When you have finished, click **OK** to save the settings.

14. The installation utility displays the changes that must be made to the config.sys and config.rps files. Select **Yes** if you want the installation utility to automatically include the configuration entries in these files. Select **No** to store the changes in config.new.

    Note: The changes must be included in config.sys and config.rps for the Director managed system to run correctly. If you select **No**, you must add the entries manually.

15. When Installation is complete, click **OK** to save the settings.

Unattended Installation of the Director Agents on OS/2

Director supports unattended installations on OS/2. You do not have to be present to provide responses to the various prompts during the installation process. Instead, a response file is automatically read, and the installation process proceeds normally.

The response files for the unattended installation for OS/2 are included in the OS/2 language subdirectory. For example, the sample English language response file, diragent.rsp, is located in the `e:\director\win32\install\files\NfD\Agent\OS2\en` subdirectory, where *e* is the drive letter of the CD-ROM drive.

Comments within the response files begin with a semicolon in the first column. All entries can be changed. The response file contains comments detailing the usage of each entry.

To launch an unattended installation of the OS/2 agent, do the following:

1. Copy and modify the sample response file (**diragent.rsp**).
2. Change the directory to the *e:*\director\win32\install\files\nfd\agent\OS2 subdirectory, where *e* is the drive letter of the CD-ROM drive.
3. Execute the unattended install with the following command:
   ```
   install.exe /R:filename
   ```
   (where *filename* is the fully-qualified response file)

Installing Director Agents on a system running NetWare

Notes:

1. Director Agent is supported only on NetWare 5.0, 5.1 and 6.0.

2. On a Windows 2000 or Windows NT 4.0 system that is logged into the NetWare server, you must map a drive to the SYS volume of the NetWare agent that you are installing.

To install the Director Agent software on a Novell NetWare system:

1. Insert the *Director* CD into the CD-ROM drive of the windows system.
2. Change the directory to the *e:*\Director\NetWare subdirectory, where *e* is the drive letter of the CD-ROM drive.

3. Run setup.bat.
4. Select the appropriate drive that is mapped to the sys volume of the targeted Novell server. The target location for the agent files is the `\tivoliwg` directory. The required files are copied and the following lines are added to the Autoexec.ncf file:

   ```
   ;********Director Agent*********
   Search add sys:tivoliwg
   load twgipc
   ;********Director Agent*********
   ```

5. The final window is displayed, listing several manual tasks that you must perform on the Netware server before using this application:
   - Type the following:

     ```
     Search add sys:tivoliwg
     ```

   - Configure the Director agent by typing: `load twgipccf`
   - Start the agent by typing: `load twgipc`

Installation is now complete. The agent automatically runs on the next Novell server restart.

Installing Director on a system running Linux

The `dirinstall` script which can be found in the director/linux folder, installs the agent code and all of the IBM Director Extensions by default. In order to use this script, you will need to mount the CD-ROM using the following command:

```
# mount -t iso9660 -o map=off /dev/cdrom /mnt/cdrom
```

The required RPM's to be installed are ITDAgent-3.10-3.i386.rpm, DirAgent-3.10.i386.rpm, and lincimom-1.0-1.i386.rpm. These RPM's install the agent portion of Director. For more information, please see "Installing IBM Director Extensions on a system running Linux" on page 62.

Installing Director on a system running SCO UNIX

The Director CD contains a file called director/sco/ITDAgent.pkg. This file complies with the UnixWare package format. The package resides in the SCO folder on the source media. For information on how to install the UnixWare package, see your operating-system user's guide. These packages must be run manually. Install the ITDAgent package first. ITDAgent.pkg installs the IBM Director Agent, and is required before installing the agent portion of any IBM Director Extensions tools. For more information, please see, "Installing IBM Director Extensions on a system running SCO UnixWare" on page 62.

To install the Director on a system running SCO UNIX, use the following procedure:

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

3. Type:

   ```
   pkgadd -d /CD-ROM_1/Director/sco/ITDAgent.pkg
   ```

   where CD-ROM_1 is the mount point of the device file created in step 2.
4. When the installation is complete, unmount the CD-ROM drive.
5. Remove the CD from the drive.

Note: The following patches must to be applied to SCO Unixware 7.1.1 before Director can be installed: ptf7045, ptf7410, ptf7441, ptf7602, ptf7603, ptf7631, freefont-2.0, udkrtfs. All of theses are available from the Caldera website.

Uninstalling Director on a system running Windows

To remove Director, use the following procedure:

1. Click **Start→ Settings→ Control Panel →Add/Remove Programs**. Select **Director** from the list of installed programs. Click the **Change/Remove** button. The message `"Are you sure you want to uninstall?"` is presented. Click **Yes** to uninstall, or click **No** to quit. The uninstall begins by stopping the server and removing components. If you have installed the Director server or agent with Director support, the system displays a second message to verify that you want to delete the configuration data and the database contents. The uninstallation program is automated and prompts you when the process is finished.
2. Restart the computer now or restart it later. If you click **Restart Now**, the system shuts down and restarts immediately.

Uninstalling Director on a system running OS/2

To uninstall Director components from a system running OS/2:

1. Change to the Director management agent directory. Typically this is c:\tivoliwg.
2. From a command prompt, type `bmuninst` and press **Enter**.
3. When prompted, click **Yes**, confirming that you wish to uninstall the agent.
4. When the uninstall completes, reboot your system.
5. After the system is restarted, remove the c:\tivoliwg directory and all of its contents. (Substitute the appropriate directory where you installed the Director management agent.)

Uninstalling Director on a system running NetWare

To uninstall Director components from a Novell NetWare system:

1. Unload Director by entering `unload twgipc.`
2. From a Windows NT 4.0 or Windows 2000 system that is logged into the NetWare server, map a drive to the sys volume and delete the tivoliwg directory.

3. Edit the autoexec.ncf file on the Novell server and remove the IBM Director section.

Uninstalling Director on a system running Linux

To uninstall Director on a system running Linux, use the `diruninst` script. This script may be found in the director\linux folder. It executes the uninstall command for the various installed RPM's. You may customize this script to uninstall one or more RPM's.

Uninstalling Director on a system running SCO UNIX

To uninstall Director on a system running SCO UNIX, run the following command:

```
pkgrm TivITDA
```

Installing IBM Director Extensions on non-Windows platforms

Installation of IBM Director Extensions on a system running NetWare, OS/2 or UnixWare is performed remotely on a system running Windows. Installation of Linux and SCO UnixWare is performed on a local system.

For NetWare, OS/2, Linux and SCO UnixWare installations, the tools are selectable and any combination of IBM Director Extensions tools can be installed. At least one IBM Director Extensions tool must be selected or an error message will be displayed. After you select the target directory, the installation program ensures that the target drive is the correct operating system.

The IBM Director Extensions tools for a system running NetWare are:

- Capacity Manager
- ServeRAID Manager
- Management Processor Assistant

See"Installing Director Agents on a system running NetWare" on page 53.

The IBM Director Extensions tools for a system running OS/2 are:

- Capacity Manager
- Management Processor Assistant

See "Installing IBM Director Extensions on a system running OS/2" on page 59.

The IBM Director Extensions tools for a system running Linux are:

- ServeRAID
- Software Rejuvenation

- Capacity Manager
- System Availability
- Management Processor Assistant

See“Installing IBM Director Extensions on a system running Linux” on page 62.

The IBM Director Extensions tools for a system running SCO UnixWare are:

- Management Processor Assistant
- ServeRAID

See “Installing IBM Director Extensions on a system running SCO UnixWare” on page 62.

To uninstall IBM Director Extensions on OS/2, NetWare, Linux, and SCO UnixWare, see the *Director 3.1 User's Guide*.

Installing IBM Director Extensions on a system running NetWare

To install the IBM Director Extensions tools on a system running NetWare, you must run IBM Director Extensions Installer remotely from a system running Windows NT. Windows 98 will not work for this purpose.

With NetWare, you must map a drive between the system running NetWare and the system running Windows NT 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 60.

To install IBM Director Extensions tools on a system running NetWare, use the following procedure:

1. From the NetWare console, type `unload twgipc` to shut down IBM Director.
2. From the Windows Start menu, click **Run**.
3. In the Run field, type**:**

   ```
   d:\umse\os2netw\xsesetup.exe
   ```

   where *d* and *umse* are the drive and the temporary directory in which the decompressed IBM Director Extensions tools files are located. The os2netw directory contains 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** button.

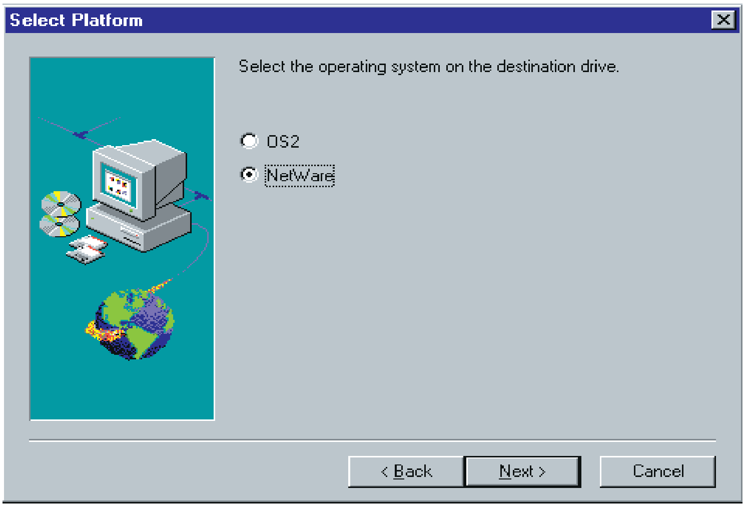

Select Platform

Select the operating system on the destination drive.

○ OS2

● NetWare

7. Click **Next**.

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

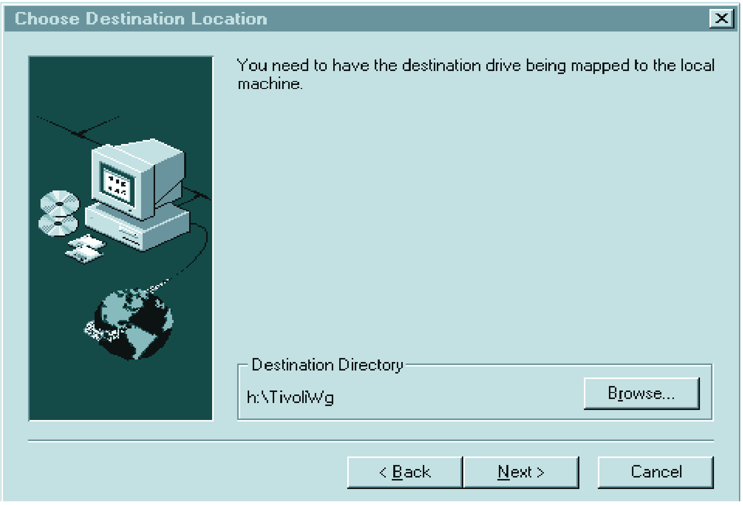

Choose Destination Location

You need to have the destination drive being mapped to the local machine.

Destination Directory

h:\TivoliWg

8. Click **Next** if you are using the default destination directory. Skip to step 11 on page 54.

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

The installation program displays the Choose Directory window.

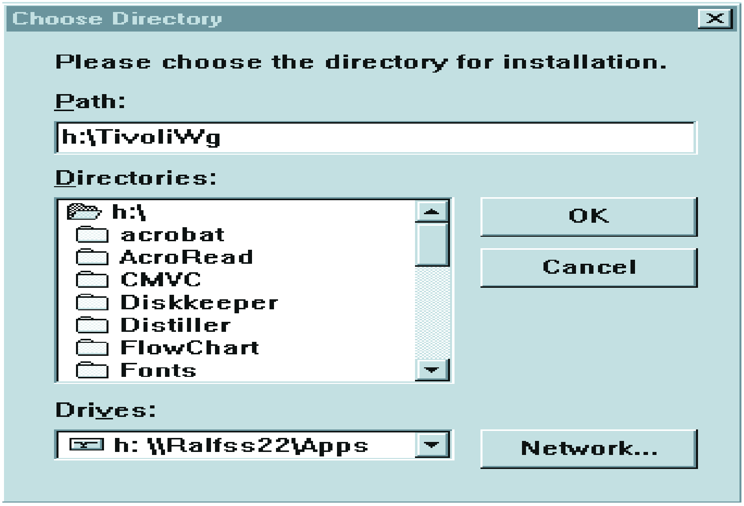
Choose Directory

Please choose the directory for installation.

Path:

h:\TivoliWg

Directories:

h:\
acrobat
AcroRead
CMVC
Diskkeeper
Distiller
FlowChart
Fonts

Drives:

h: \\Ralfss22\Apps

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 *d:*\TivoliWg\

    where *d* is the mapped drive, and TivoliWg is the default directory.

11. Click **OK**.

    IBM Director Extensions Installer installs the following tools:

    - Management Processor Assistant
    - Capacity Manager
    - ServeRAID Manager

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

Installing IBM Director Extensions on a system running OS/2

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

You must map the drive from the system running OS/2 to the system running 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[d|*][\\computername\sharename]
```

where *d* represents any available drive on the system running Windows NT, and the *computername* is the name of the OS/2 system and share name is the name that the shared directory on the OS/2 system is shared as.

Installation instructions for a system running OS/2

Perform the following steps to install the IBM Director Extensions tools on a system running OS/2:

1. On the target system, run the following command: twgipc shutdown. This will stop Director and enable the installation to proceed.
2. In the **Run** field, type:

   ```
   d:\umse\os2netw\xsesetup.exe
   ```

   where *d* and *umse* are the drive and the temporary directory where the decompressed IBM Director Extension tools files are located. The os2netw directory contains the installation files for an OS/2 or NetWare installation.
3. Click **OK**.

   Two Welcome windows open.
4. Click **Next** through both Welcome windows. The Select Uplifters window opens.
5. Select the **OS/2** button.
6. Click **Next**.

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

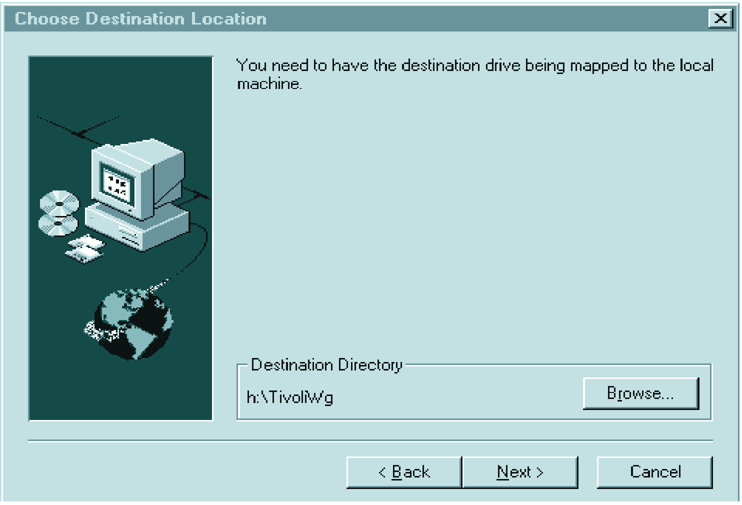

You need to have the destination drive being mapped to the local machine.

Destination Directory

h:\TivoliWg

7. Click **Next** if you are using the default destination directory. Skip to step 10 on page 57.

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

   The installation program displays the Choose Directory window.

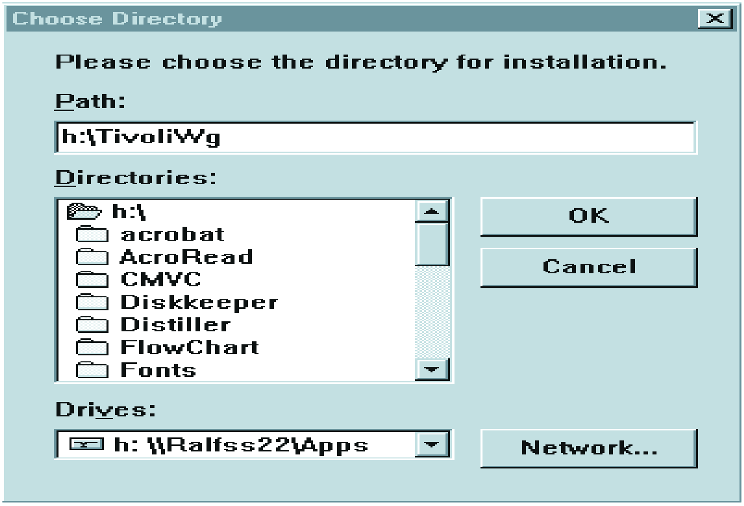

Please choose the directory for installation.

Path:

h:\TivoliWg

Directories:

h:\
acrobat
AcroRead
CMVC
Diskkeeper
Distiller
FlowChart
Fonts

Drives:

h: \\Ralfss22\Apps

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 60".

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

    IBM Director Extensions Installer installs the following Extensions tools:

    - Management Processor Assistant
    - Capacity Manager
11. When the installation is complete, from the console of the OS/2 system, type `twgipc` to restart Director.

Using command lines to install IBM Director Extensions

You install IBM Director Extensions on non-Windows systems using Red Hat Package Manager (rpms) for Linux and packages for SCO UnixWare. See "Installing IBM Director Extensions on a system running Linux" and "Installing IBM Director Extensions on a system running SCO UnixWare" for examples of the installation procedure.

Installing IBM Director Extensions on a system running Linux

To install IBM Director Extensions on a system running Linux, use the `dirinstall` script which can be found in the Director\Linux directory. The `dirinstall` script by default installs all of the IBM Director Extensions tools. You can modify the script to prevent installation of one or more IBM Director Extensions. For more information go to the comments in the dirinstall file in the Director\linux directory.

The available IBM Director Extensions tools for linux are:

- ServeRAID
- Management Processor Assistant
- Software Rejuvenaton
- Capacity Manager
- System Availability

Installing IBM Director Extensions on a system running SCO UnixWare

The available IBM Director Extensions tools for SCO UnixWare are:

- ServeRAID

- Management Processor Assistant

The Director UnixWare ASM Agent provides the Management Processor Assistant support for IBM Director on UnixWare 7.1.1. For installation, removal, and setup, use the following procedure:

1. Insert the CD into the CD-ROM drive.
2. Mount the CD-ROM drive.
3. Press **Enter** and type:

   ```
   pkgadd -d /CD-ROM_1/umse/unix/asmuwag.ds
   ```

   where *CD-ROM_1* is the mount point of the device file created in step 2.
4. Press **Enter**.
5. When the installation is complete, unmount the CD-ROM drive.

   ```
   umount /CD-ROM_1
   ```
6. Press **Enter** and remove the CD from the drive.

To install from the Web, use the following procedure:

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

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

   where *tmp* is the temporary 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 Director UnixWare RAID Agent provides the ServeRAID support for the Director on UnixWare 7.1. This package requires that the Director UnixWare Agent be installed. For installation, removal and setup, use the following procedure:

1. Insert the CD into the CD-ROM drive.
2. Mount the CD-ROM drive.
3. Press **Enter**, and then type:

   ```
   pkgadd -d /CD-ROM_1/umse/unix/RAIDUwAg.pkg
   ```

   where *CD-ROM_1* is the mount point of the device file created in step 2.
4. Press **Enter**.
5. When the installation is complete, unmount the CD-OM drive.
6. Press Enter. You can remove the CD from the drive.

The Director UnixWare RAID Agent can be downloaded. Use the following procedure:

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

   ```
   pkgadd -d /tmp/RAIDUwAg.pkg
   ```

   where *tmp* is the temporary 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
   ```

Uninstalling IBM Director Extensions tools

There are two uninstallation procedures for IBM Director Extensions tools. The procedure that you use will depend on whether you installed IBM Director Extensions tools locally or remotely.

Uninstalling IBM Director Extensions on a system running Windows

To uninstall IBM Director Extensions tools on a system running Windows, you must either uninstall Director or reinstall Director and uncheck the IBM Director Extensions tools which you wish to install.

Uninstalling IBM Director Extensions on a system running OS/2

To uninstall IBM Director Extensions from a system running OS/2, please refer to “Uninstalling Director on a system running OS/2” on page 55.

Uninstalling IBM Director Extensions on a system running NetWare

To uninstall IBM Director Extensions from a system running NetWare, please refer to “Uninstalling Director on a system running NetWare” on page 55.

Uninstalling IBM Director Extensions on a system running Linux

To uninstall IBM Director Extensions on a system running Linux, use the `diruninst` script. This script may be found in the director\linux folder. It executes the uninstall command for the various installed RPM’s. You may customize this script to uninstall one or more RPM’s. Please refer to the comments located in the diruninst file for information on modifying the script.

Uninstalling IBM Director Extensions on SCO UnixWare

To uninstall IBM Director Extensions on a system running SCO UNIX, run the following commands:

ServeRAID

```
pkgrm RAIDUwAg
```

Managment Processor Assistant

```
pkgrm asmuwag
```

After you uninstall the Director Agent Extensions, use the following command to uninstall the Director Agent:

```
pkgrm TivITDA
```

Configuring Director to use Oracle Server or DB2 Universal Databases

It is recommended that you use the graphical interface database installation process. However, you may use the command line to configure the DB2 or Oracle Director interface.

Configuring the DB2 Universal Database

To configure the DB2 database from the command line, use the following procedure:

1. In the IBM Director /data directory, edit or create the TWGServer.prop file. Add the following lines (where *test20* is the database name):

   ```
   twg.database.odbc.name=test20

   twg.database.jdbc.driver.name=com.ibm.db2.jdbc.app.DB2Driver

   twg.database.jdbc.subprotocol=db2

   twg.database.jdbc.user=bender
   ```

2. From the IBM Director /bin directory, issue the `dbpasswd` command to set your password:

   ```
   dbpassw -user <userid> -pwd <password> -confirmpwd <confirmpassword>
   ```

   A line (similar to the following) will be added to the TWGServer.prop file displaying an encrypted password:

   ```
   twg.database.jdbc.password=82A2697BA5E99212
   ```

Configuring the Oracle Server

To configure the Oracle Server from the command line, use the following procedure:

1. In the IBM Director /**data** directory, edit or create the TWGServer. prop file. Add the following lines (where *gotb-2* is the hostname, *1521* is the TCP/IP Listener port number, and orcl is the system identifier below):

   ```
   twg.database.odbc.name=thin:@goth-2:1521:orcl


   twg.database.jdbc.driver.name=oracle.jdbc.driver.OracleDriver


   twg.database.jdbc.subprotocol=oracle


   twg.database.jdbc.user=bender
   ```

2. From the IBM Director/bin directory, issue the `dbpasswd` command to set your password:

   ```
   dbpasswd -user <userid> -pwd <password> -confirmpwd
   <confirmpassword>
   ```

   A line (similar to the following) will be added to the TWGServer.prop file displaying an encrypted password:

   ```
   dbpasswd - user <userid> -pwd <password> -confirmpwd
   <confirmpassword>
   ```

   A line (similar to the following) will be added to the TWGServer.prop file displaying an encrypted password:

   ```
   twg.database.jdbc.password=82A2697BA5E99212
   ```

Defining server preferences for database properties

You can view and change various database information from the Database page on the Server Preferences window.

On the IBM Director Management Console, click **Option→ Server Preferences** and then from the Server Preferences window, click the **Database** tab. This page displays database name, vendor, version and current status information, as well as JDBC driver, version and subprotocol information. In addition, you can also change your password, where applicable, depending on the database you are using. Some databases do not require a password.

Configuring Director to use file distribution servers

If you have defined one or more servers to act as file distribution servers for software distribution, read the guidelines and restrictions described in Chapter 10, "Software Distribution," on page 127 before you attempt to use the file distribution servers (server shares).

Enabling UNC-based share access to the Director server

The User ID under which the Director server was installed must have read/write access to a share. Distribution defaults to streaming if the appropriate access is not established. Ensure that the file distribution server is a member of the same domain as the Director server, or has a trust relationship with that domain.

Enabling UNC-based share access to managed systems

All IBM Director managed systems must have read access to the server shares they intend to use.

Enabling UNC-based share access to Windows managed systems

If you intend to distribute software to system running Windows and you have not specified user IDs and passwords under Distribution Preferences to access your file distribution server shares, you must complete one additional step if your file distribution server is a Windows NT server.

The Director management agent runs under the System account on Window NT systems. When the Director management agent tries to access the file distribution server, it logs in with a set of null credentials. Microsoft restricts access to systems that try to read or write to a shared drive using the System account with null credentials. In order for Windows NT managed systems to access the file distribution server, the `TWGSHARE` utility must be run on the file distribution server.

In the BIN subdirectory where you installed the Director server you will find a program named TWGSHARE.EXE. Copy this program to your file distribution server. Run the utility on the file distribution server with the following parameters:

```
TWGSHARE-A



       SHARENAME
```

where *SHARENAME* is the name of the share you created on the file distribution server.

This utility alters a registry setting on the file distribution server to allow the share to be accessed by systems with null credentials. For more information on null credentials and the System account see Microsoft article Q122702 on the `http://support.microsoft.com` homepage. For a list of other parameters supported by TWGSHARE.EXE, just run the program with no parameters specified.

Defining server preferences

After your file distribution server has been configured, you need to configure the Director server to use it.

On the Director Management Console, select **Options →Server Preferences** and then from the Server Preferences window, select the **File Distribution Servers** tab. This tab displays a list of all configured file distribution servers.

Click **Add** to add a server to the list. The Add Share Name dialog is displayed.

In the Share Name field, enter the name of a shared server that can be accessed by the managed systems to which you wish to send software packages. Use Universal Naming Convention (UNC) format; for example, \\SRVR0001 as the name of the file distribution server and Sys45NT as the network name of the shared resource.

To specify an FTP file distribution server, use the following:

ftp:\\server_name

In this window you also specify:

- The maximum disk space allowed to be utilized by IBM Director on this server
- The maximum number of concurrent managed system connections
- A limit on the bandwidths when copying files from a file package on the IBM Director server to the identified share. You may want to limit the bandwidth when a dedicated connection, such as ISDN, is used for copying the files from the server to the share.
- User ID and password required to access the standard FTP server.

Refer to the online help for more information on these options. Click **OK** to continue. The Server Preferences window is displayed once more, this time with the File Distribution Servers tab containing the data you entered in the Add Share window.

If you have multiple file distribution servers, you can repeat this procedure to define each server share. When you are finished, click **OK** to save and close the Server Preferences window.

Configuring distribution preferences for managed systems

You can use Distribution Preferences to assign unique policies to both groups and individual managed systems. For example, if you have configured Distribution Preferences for a dynamic group, as managed systems become members of that group, the policy is assigned automatically. File distribution server shares configured in Distribution Preferences must already be defined in the Server Preferences.

By default, a managed system is set up to attempt to access all shares that have been defined to the IBM Director server. If you have set up file server shares for redirected installations or for streaming of software distribution packages and you want to:

- Restrict access to the shares for specific manages systems or groups
- Specify streaming (copying) only from the IBM Director server to specific managed systems or groups
- Specify streaming (copying) only from specific server shares to specific manages systems or groups
- Specify FTP server shares on all systems, except OS/2
- Specify user ID and password to access identified server shares (if anonymous FTP access is not supported)

then, in the Director Management Console, select the managed system, managed systems, or group for which you want to set up one or more of these distribution preferences and right-click to display the context menu.

Select **Distribution Preferences** from the context menu and the Set Managed System Distribution Preferences window appears.

Select **Always stream to Managed System(s)** if you want to copy packages directly from the Director server to the system for which you opened the window.

Select **Stream from File Distribution Server** if you want to copy packages from the server shares specified in the Shares field to the systems for which you opened the window.

Select **Restrict share selection to list** if you want to limit the shares that can be accessed by the selected systems to only the shares you specify in this window. If you do not select this option and the selected systems have access to other shares that are defined to Director for software distribution (through the **Server Preferences→ File Distribution Servers** menu option), then the other shares can be used for package distribution if the shares defined in this window are not available. In this case, UNC-based shares will be accessed via null credentials and FTP-based shares will be accessed anonymously.

Select **Enter streaming bandwidth (kbps) for managed systems** to limit the bandwidth when copying packages from file distribution servers to the managed system.

Note: This value is also used to determine the streaming rate between the IBM Director server and the managed system.

Other options are available to enable you to add, remove, and edit shared directory entries. Refer to the online help for details on these procedures.

Defining the maximum number of concurrent redirected distributions

Redirected software distribution is designed to minimize the usage of network bandwidth during a distribution. If a software package has been placed on a share by the Director server, then Director managed systems are assigned to use that share. The number of managed systems installing the software package at one time does not exceed the number of concurrent users defined under **Options →Server Preferences →File Distribution Servers**. The default limit is 10 concurrent managed systems per share. If the set value is reached, additional managed systems are queued and distributions occur as active distributions are completed.

To obtain higher distribution concurrency, individual manages systems should be configured to use other shares. Spreading the distribution load over multiple shares allows more managed systems to install the software concurrently. However, care must be taken so that the network is not overloaded from managed systems accessing shares that are located on the same physical part of the network.

Defining the maximum number of concurrent streamed distributions

You can set an integer representing the maximum number of managed systems that the Director server can stream software packages to concurrently. Use this number to help limit the amount of network traffic generated by streaming. To set a limit, from the Director Management Console select **Options →Server Preferences →Software Distribution**. The default limit is 3 concurrent managed systems.

Limiting the bandwidth for streamed distributions

You can specify the maximum number of kilobytes per second (kbps) that can be used for a streamed distribution. This value can be set for all streamed distributions from the Director server and for individual managed systems and groups. To set a value for all systems, select **Options→ Server Preferences→ Software Distribution**. To set a value for an individual managed system or group, right-click on the system or group and select Distribution Preferences. If both the Director server and a managed system bandwidth are set, the lower value is used. Refer to the online help for descriptions of the fields. See “Limiting network resources for Software Distribution” on page 34 for more information on limiting the bandwidth of a distribution.

Restricting Access Check

If you select **Restrict Server Access Check**, the Director server will verify access only for file distribution server shares configured for the systems targeted for distribution.

Specifying Do Not Stream Distribution if redirected distribution fails

If you select this option, if redirected distribution fails, the software distribution job will not attempt streaming to complete the job.

Defining the automatic time-out for remote control sessions

You can specify the inactivity time-out for remote control consoles. Console inactivity is defined as no mouse or keyboard input through the console. Any input restarts the timer, so this value only applies to consoles in active mode. Each agent has a separate timer set for its connection to any console. All timers restart if the time-out value is changed while the remote control server is running.

A value of 0 in this field deactivates all timers. When any of the timers expires, a message is sent to all consoles to notify them of the automatic time-out.

Changing the network transport

To change the network transport driver configuration used by the IBM Director server or agents, select **Start** → **Programs**→ **IBM Director** → **Network Driver Configuration**. The Network Driver Configuration window is displayed, enabling you to modify any of the options originally set during the initial installation. For non-Windows systems, you need to stop and restart the service or reboot the system for configuration changes to take effect.

To change the network transport driver configuration on OS/2 systems, open the **IBM Director Agent for OS2** and double-click on **Network Driver Configuration**. To activate the changes, the TWGIPC.EXE program must be shut down and reloaded, or the system must be rebooted.

To change the network transport driver configuration on systems running NetWare, access the NetWare server console either locally or by remote control. From the console, load TWGIPCCF from the NetWare server console or a remote NetWare console. After changing any of the desired values and saving, the user must then unload (if presently running) and then load TWGIPC from the NetWare server console.

Saving, restoring, and resetting program files in UNIX

Before uninstalling, back up the program files for the Director management agent, the management console, or the server. Later you can restore the program files from the backup files, if necessary. Or you can reset the system to reflect its state after initial installation and configuration.

Use the following commands (from the bin directory) to backup, restore, and reset the Director:

- twgsave

  This command saves the contents of the data directory and, on servers, it also saves the SwDistPk directory. The data files are placed in a directory at the same level as the tivoliwg directory. This directory is named tivoliwg.save*n*, where *n* is incremented by one each time this command is used. Use the optional -s parameter to prevent saving of software distribution packages (in the SwDistPk directory) on servers.

  This command runs automatically as part of the uninstall process. To prevent the uninstall from saving data, edit the Uninstall.properties file and change the SaveUserDataAtUninstall and SavePackagesAtUninstall variables.

- twgrestore

  This command copies the files saved by the twgsave command back into the data directory, or into the SwDistPk directory on servers. You must include the directory containing the saved data (tivoliwg.save*n*) as a parameter. This command operates by executing the twgreset command to erase any old files from the data directory, then restoring the saved data.

  Use the -t parameter if you do not want to restore the system identification data which is contained in files that include the system's name and access keys. If you erase these files, your system will no longer be known to Director servers.

- twgreset

  This command restores the system to its initially configured state. It deletes all files from the data directory except for the originally installed data files and the system identification files. Use the optional -i parameter to delete the system identification files; use the -d parameter to delete the tables in the database.

Chapter 4. Upgrading Director and IBM Director Extensions

This chapter contains detailed information for upgrading to Director 3.1 and IBM Director Extensions from a previous release. The installation program checks for previous versions and, depending on the type of installation, upgrades the necessary Director components.

There is an upgrade install available for the agent and server portion of Director 3.1. If a previous version (2.2) of IBM Director Agent is on the system, the install will upgrade, and a list of additional options will be displayed. You can then add any options you have not previously installed or just update the current configuration without adding additional functionality.

Upgrading to Director 3.1

The supported upgrade paths for Director versions are:

- Director 2.2.1 to 3.1
- Director 2.2 to 3.1
- Director 2.2 to 2.2.1 to Director 3.1

Note: A Director 3.1 server requires a Director 3.1 console. Also, 3.1 agents require a 3.1or later server.

Upgrading the Director Server

Follow the same steps for upgrading Director Server as you would to install it. Insert the *Director* CD into the CD-ROM drive of the Director system that you are upgrading. Follow step 1 through step 3 of the installation section, see “Installing the Server” on page 43.

To upgrade the Director server to version 3.1, use the following procedure:

1. Select **Server** from the Select Components window.

   Installation detects a previous installation. A Question window opens.
2. Click **Yes** from the Question window to proceed with the upgrade. The Current Install window opens.

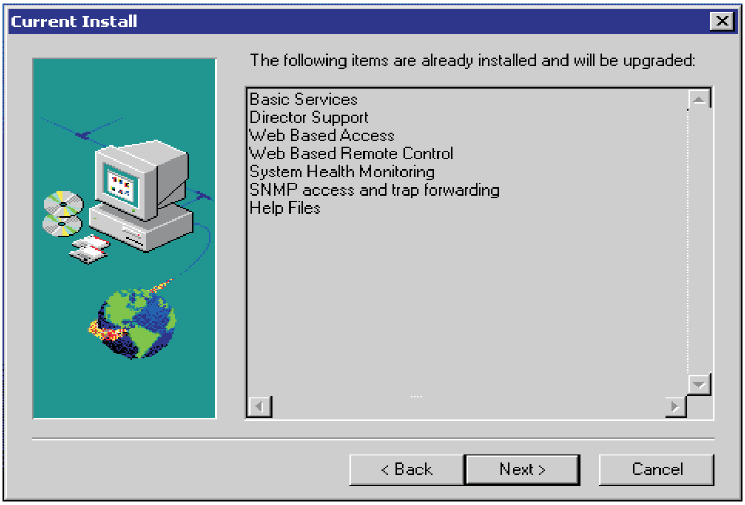

Current Install

The following items are already installed and will be upgraded:

Basic Services
Director Support
Web Based Access
Web Based Remote Control
System Health Monitoring
SNMP access and trap forwarding
Help Files

Items that were installed with the previous version are listed in this window.These items are upgraded with the new version.

3. Click **Next**. The Additional Components for Upgrade window opens.

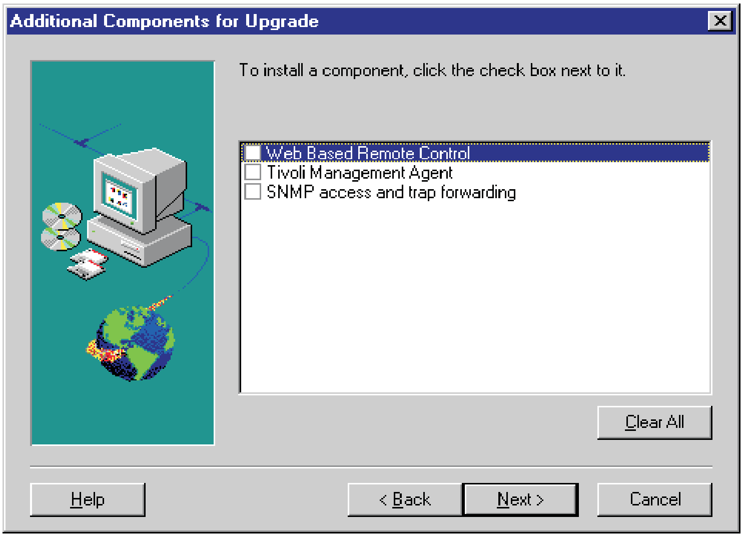

Additional Components for Upgrade

To install a component, click the check box next to it.

□ Web Based Remote Control
□ Tivoli Management Agent
□ SNMP access and trap forwarding

Features that are either new to Director or were not installed in the previous version are listed in this window. Select the check box beside each feature you want to add.

4. Click **Next**. The Director Extensions Install Options window opens.
5. Select the check box beside any component that you want to install.

6. Click **Next.** If you did not select any agent UIM components in the previous install, the Agent UIM Install Options window opens.
7. Select the check box beside any component that you want to install.
8. Click **Next**.
9. Select an IP port number for the Apache Web server.
10. Click **Next**.
11. Click **Yes** in the Add IBM Director Agent Icons window for the additional menu choices provided with the upgrade.
12. Click **Next** to accept the default directory (**C:\Program Files\Director\SwDistPk**), or click **Browse** to choose a different directory. The upgrade program detects the current Director database. Click **Yes** to use the current database with the upgrade.
13. Click **Next**. The system begins installing the necessary files. The Windows Account information window opens. The Domain and User Name for the system that you are using is displayed.
14. Type your Windows account password in the Password field and type it again in the Confirm Password field.
15. Click **Next**. A Question window opens.
16. Select **No** to install a new database. The Database Configuration window opens. Highlight the appropriate database. Select **Yes** to use the current database. The Network Configuration window opens.
17. Click **Next**.
18. Complete steps 20 through 23 of "Installing the Server" on page 43.
19. Select **OK** to continue. When the installation is complete, the Setup is Complete window opens.
20. Restart the computer now or Restart it later. If you choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the installation program closes. However, you must restart and log into the system to begin using IBM Director.

Upgrading the Director Console

After the Welcome and End User License Agreement windows open, select **Console** from the **Select Components** window.

The upgrade follows the previous installation path upgrading existing software, removing obsolete files and directories, and installing the new console components.

Restart the computer now or later. If you click **Restart Now**, the system shuts down and restarts immediately. If you click **Restart Later**, the IBM Director Agent installation program closes. However, you must restart and login to the system to begin using Director.

Upgrading the Director Agent

Note: Netware, OS/2, Linux and SCO do not support an upgrade. You must uninstall the previous version.

To upgrade the Director Agent for a selected system, do the following:

1. After the Welcome and End User License Agreement windows open, select **Agent** from the **Select Components** window. Installation detects the previous version of the Director Agent and the upgrade prompt appears.

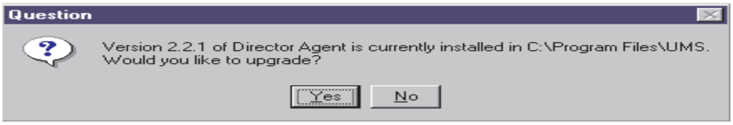

Version 2.2.1 of Director Agent is currently installed in C:\Program Files\UMS.
Would you like to upgrade?

2. Click **Yes** to begin. The installation program detects the agent components from the previous version of the installed IBM Director Agent. The Current Install window opens.

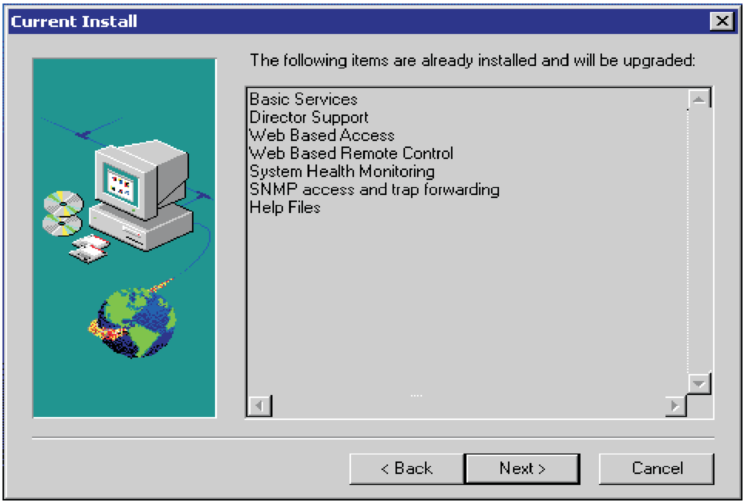

The following items are already installed and will be upgraded:

Basic Services
Director Support
Web Based Access
Web Based Remote Control
System Health Monitoring
SNMP access and trap forwarding
Help Files

The installation program upgrades the agent components listed in this window.

3. Click **Next** to continue. The Additional Components for Upgrade window opens.

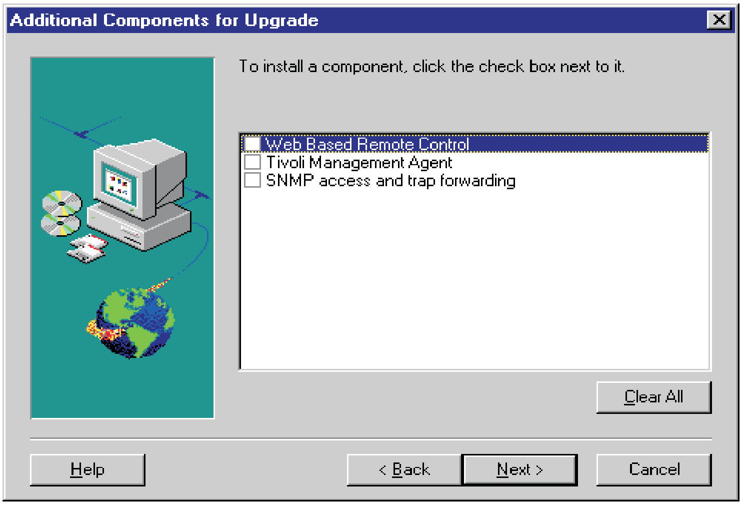
Additional Components for Upgrade

To install a component, click the check box next to it.

- □ Web Based Remote Control
- □ Tivoli Management Agent
- □ SNMP access and trap forwarding

Select additional components to add to the upgraded components of the IBM Director Agent. Click **Next**. The Director Extensions Install Options window opens.

4. Select Director Extensions components.
5. Click **Next**. If you did not select any agent UIM components in the previous install, the Agent UIM Install Options window opens.
6. Select the check box beside any component that you wish to install.
7. Click **Next**. Select an IP port number for the Apache Web server.
8. Click **Next**.
9. If you have not installed icons in the previous version, a prompt to add IBM Director Agent icons appears. Click **Yes**.I f you have added icons in the previous version, your icons will be updated.
10. The program upgrades the existing components and installs the selected new components. Restart the computer now or later. If you click **Restart Now**, the system shuts down and restarts immediately. If you click **Restart Later**, the IBM Director Agent upgrade program closes. However, you must restart and log into the system to begin using Director.

Upgrading IBM Director Extensions

An upgrade installation is available for the agent and console portion of IBM Director Extensions. If a previous version of Director Extensions is on the system, the installation will upgrade it to IBM Director Extension 3.1. All options previously installed will be upgraded, and a list of additional options will be displayed. You can add any options that have not been previously installed or just update the current configuration without adding any additional functionality.

Chapter 5. Using the Director Management Console

The Director Management Console is your interface into the Director environment. From here you can perform all of the administrative tasks as well as define how your various network elements are grouped together and managed.

This chapter describes the various parts of the Director Management Console. It also shows you examples of the tasks that you can perform. First, you need to become familiar with managed systems.

Managed systems

The operation of Director is built around the concept of *managed systems*. Managed systems can consist of various systems and devices. Each managed system has tasks and properties associated with it. Director recognizes two types of managed systems:

Native systems
Systems that have the IBM Director Agent or Tivoli management agent code installed

SNMP devices
Network devices, printers, or PCs that have SNMP agents installed or embedded

Director enables you to organize these managed systems into groups based on their specific attributes and properties. From the Director Management Console you can perform tasks on a single managed system or on a group of managed systems.

Starting Director Management Console

To start the Director Management Console, use the following procedure:

1. Select **Start→ Programs→Director → Management Console**. After the console starts, the Director login dialog appears.
2. Enter the name of your Director server, your user ID, and your password. The server name is the TCP/IP host name or the address of the Director server. The user ID and password must be an authorized user account on the Director management server. Your user account maintains your Director Management Console configuration preferences since the last time you were logged in, including status and security settings.

Your Director Management Console can communicate with only one Director server at a time. Multiple Director Management Consoles, however, can be open at the same time, each communicating with the same or a different Director server.

After your login information is accepted, the Director Management Console is opened and displayed in a window, similar to the following example:

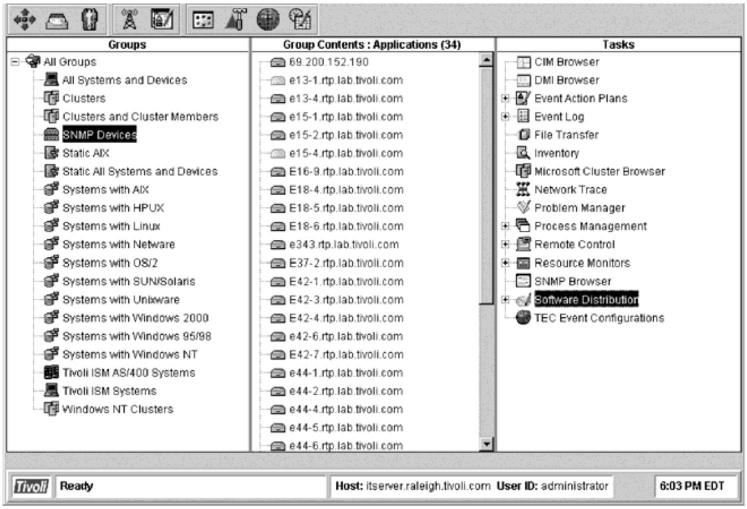

Ready

Host: itserver.raleigh.tivoli.com **User ID:** administrator

Navigating in Director

You can invoke many Director tasks and operations in several different ways. You perform some tasks by dragging and dropping icons or by selecting operations from pull-down menus. Your mouse buttons have different functions assigned to them.

Note: Throughout this User’s Guide and the online help, you might see references to clicking and right-clicking the mouse buttons to perform operations. This assumes that your mouse button configuration is set to right-handed, which uses the left button for normal select and drag functions, and the right button for context menus and special drag operations. If your mouse is configured for left-handed operation, you will have to transpose the meaning of clicking and right-clicking to the right and left buttons, respectively.

This section gives you an idea of how to navigate from one display to the next. You should try these techniques in each window while you use Director. Use the method that is most convenient for you.

Using drag and drop

Several windows displayed in Director consist of two or more panes. In most instances you can drag and drop task and target icons between these panes. However, you cannot perform drag and drop operations between two separate Director windows.

To execute a task on a single managed system on the Director Management Console, use the following procedure:

1. Drag the system icon from the Group Contents pane and drop it onto the desired task icon in the Tasks pane.
2. Drag the desired task icon from the Tasks pane and drop it onto the desired managed system icon in the Group Contents pane.

To apply the task to more than one system at a time:

1. Hold down the **Shift** key and click on a range of systems. This action highlights several systems.
2. Drag the task from the Tasks pane and drop it on any of the highlighted managed systems in the Group Contents pane. This action invokes the task on all the highlighted systems.

Similarly, you can hold down the **Ctrl** key and highlight individual managed systems, skipping over those you wish to omit from the selection. You then drag the desired task icon and drop it onto one of the highlighted system icons.

To invoke a task on all available managed systems in a group, use the following procedure:

1. Drag the group icon from the Groups pane and drop it onto the desired task icon in the Tasks pane.
2. Drag the desired task icon from the Tasks pane and drop it onto the desired group icon in the Groups pane.

You can use this drag-and-drop technique throughout Director. Examples include the following:

- The file transfer task (see Chapter 11, “File Transfer,” on page 137). This action enables you to drag files and subdirectories from one system to another.
- The event management task (see Chapter 9, “Event management,” on page 117). In this task, you can drag filter and action icons and drop them onto an event action plan icon to create event action plans.
- The software distribution task (see Chapter 10, “Software Distribution,” on page 127), where you can drag a software distribution package and drop it onto a managed system or group of managed systems. This action invokes the download and install of new software packages.

This type of task activation is referred to as *targeted* activation, because the tasks are being applied to specific managed systems or groups of systems.

See the online help for more details on performing specific operations for each task.

Using the double-click function of the mouse

You can double-click specific tasks such as Inventory. This action performs an *untargeted* activation of the Inventory task (see Chapter 6, “Inventory

Management," on page 91). The Inventory Query Browser displays the inventory on all discovered systems and devices.

Note: Untargeted activations are not applied to specific systems or groups of systems. Be careful using this technique—applying a task to all discovered systems and devices in a large network can be expensive and time-consuming.

You can also double-click icons displayed in tree structures that have branches containing additional icons representing subtasks or other associations. Double-clicking the icon will expand or collapse the tree structure, enabling you to manage the view in the pane. You can also just click on the plus (+) or minus (-) symbol next to the icon to expand or collapse the tree view.

Using context menus

You can right-click on some task icons or system icons and be presented with a pop-up context menu, enabling you to select one of several operations to be performed, depending on the context of where you are in the Director product.

Using the Add and Remove buttons

Some windows in Director contain add and remove buttons, such as the Inventory Query Builder window (see the figure on page 94).

- To add a selected item from the source pane to the target pane, click **Add**.
- To remove the selected item from the target pane, click **Remove**.

Managing columns of information

Many panes of information in Director are displayed in tabular format. You can tailor the view of this information by using one of the following techniques:

- Change the width of each column by dragging the edge of the column header left or right, enabling you to view the data more easily.
- Move whole columns at a time by dragging the center of a column header left or right. The entire column then moves with it. Adjacent columns shift automatically to fill the space.
- To perform the following operations, you can also right-click within some columns:
  - — Restore a hidden column: Position the mouse pointer over Show Columns in the context menu and click on the column you want to restore.
  - — Arrange data in a column: click **Arrange** and then click **Ascending** or **Descending**.

Monitoring the task in process

When you initiate a task or service, an animated IBM icon at the bottom left of the window indicates that Director is busy performing the designated activity.

Across the bottom of the window is a text-based status field which will inform you of the status as the task or service progresses.

Using keyboard arrow keys

You can move up and down a list of tree structure in a pane using the up and down arrow keys instead of using a mouse. When you want to expand a branch, press the right arrow key and the next level is displayed. Press the left arrow key to collapse the tree view again. When you reach the icon you want, select it by pressing Enter.

Saving files

In tasks where you generate data that you want to save in a file, select the **File** option from the menu bar at the top of the window, and then select **Save As** or **Export** to save the data to a new file. You will be prompted for a file name and may be asked if you want to save the file in your local file system or on the Director server.

If you are updating an existing file, select the **Save** option.

You can specify one of several formats to save to, such as comma-separated values (CSV), Hypertext Markup Language (HTML), or Extensible Markup Language (XML) format in the Inventory task.

Using the Director Management Console

The main portion of the Director Management Console contains the Groups, Group Contents, and Tasks panes.

Group Contents

The middle pane in the Director Management Console is the Group Contents pane. It displays the managed systems that are members of the group you selected from the Groups pane (see “Groups” on page 84).

You can use the drag and drop methods described earlier to perform tasks on managed systems, or select an option from the context menu of the system. Refer to the online help for detailed information on each available option.

An additional context menu is available for the Group Contents pane that enables you to identify new systems, perform searches for a particular system in the list, change the view in the pane, sort the order of systems displayed and group them by various common attributes (see “Associations” on page 84), or initiate a new discovery of systems in the network. See the online help for details.

The title bar in the Group Contents pane contains additional information. A number appearing in parentheses after the Group Contents title indicates the total number of managed systems in the selected group, that is, the group highlighted in the Groups pane on the left. Words in parentheses after the number indicate the current association that has been applied to the selected group.

Associations

You can define an association between sets of managed systems to group them in a more logical manner in the Group Contents pane.

1. Select **Associations** from the menu bar at the top of the Director Management Console window.
2. From the context menu that is displayed, select the association you want. This action organizes managed resources according to their role in the selected application or operating system.

For example, you might want to display all managed systems that exist in a Windows NT domain, or all managed systems that are identified as native Director agents, or some other system type. You might also want to see managed systems that have no particular association with other managed systems in that group.

To turn off the associations, select **Associations** → **None**. If no association is selected, the managed resources are listed alphabetically.

Groups

Groups consist of logical sets of managed systems. An example of a group might be one that contains only desktop PCs with 486 processors that have Windows 95 installed.

When you first log into the Director server with the Director Management Console, a minimum number of default groups is created. Included in this default list is the All Systems and Devices group, which contains everything in the network. Additionally, known groups will appear dynamically as systems of these types are discovered. Examples of some of these known groups are Systems with Linux, Systems with Windows 2000 and SNMP Devices.

You can create new dynamic groups if you are authorized. All changes that you make to these groups are global and are applied to all users.

Note: There is no implied hierarchy or relationship among groups of managed systems in the view. They are simply grouped logically for your convenience.

To select a group as the current group, click its group icon. Managed systems that are members of that group appear in the Group Contents pane. You can have an empty group, that is, a group icon that does not contain any managed systems meeting the criteria of the group.

You can only select one group at a time. To perform tasks simultaneously on multiple groups, create a new group and include all of the desired managed systems.

Dynamic and static groups

All default groups are considered to be *dynamic*. This term means that after the criteria is set, Director automatically updates the group when the set of systems known to the Director Server changes. Director adds and deletes managed systems when their attributes and properties change to match the group’s criteria.

While this operation covers most management needs, occasions arise when you need to add or remove systems or a group’s systems. These groups are then designated as *static*. This term means that the Director server does not automatically update the contents of the group.

You can copy a dynamic group into a new static group. Director does not automatically update this new static group. However, you can add and remove managed systems from the dynamic group.

Creating a dynamic group

You create a dynamic group by defining criteria that allow specific managed systems with specific attributes and properties to become members of that group.

To create a new dynamic group, use the following procedure:

1. From the Groups pane context menu, select **New Dynamic** (right-click in any empty space in the Groups pane).
   The Dynamic Group Editor window appears.
2. Expand the tree structure in the Available Criteria pane and select one or more criteria to define the group.

   You can drag the criteria and drop it anywhere in the Selected Criteria pane or use the Add button to add it to the list. You can then use the Boolean operators AND or OR to create a tree structure. Based on the structure you create, managed systems are added or removed from the group.

   Within the Selected Criteria pane you can move these criteria to redefine the logical association as desired.
3. To delete a highlighted criterion from the Selected Criteria pane, click **Remove**.

   You can further refine each selected criteria by specifying its logical value from its own context menu (right-click on an icon in the Selected Criteria pane), defining whether the selection criteria is equal to, not equal to, greater than, or less than, and so on.
4. Select **File→ Save As** to save the new dynamic group under a name you choose.

   Director dynamically populates the group with all managed systems that meet the specified criteria. When the Director Management Console refreshes itself, the new dynamic group appears in the Groups pane. You can immediately select it to see the managed systems that match your criteria listed in the Group Contents pane.

Creating a static group

You create a static group by selecting specific managed systems to become members of the group, regardless of their specific attributes or properties. Because static groups have no criteria on which to accept or reject members, the group consists of all the systems you add to it.

To create a static group from the Groups pane of the Director Management Console, follow one of these methods:

- Select **New Static** from the Groups pane context menu (right-click in any empty space in the Groups pane). This causes the Groups pane to split. The Static Group Editor then appears in the lower half of the Groups pane.
- Right-click an existing dynamic group and select **Copy as Static** from the context menu. Select **Edit** from the context menu of the newly created static group to bring up the Static Group Editor.
- Right-click an existing static group and select **Copy** from the context menu. Select **Edit** from the context menu of the newly created static group to bring up the Static Group Editor.

You can drag specific managed systems from the Group Contents pane and drop them into the Static Group Editor to add the system to the group. You can change to a different group in the Groups pane and continue to select managed systems from that group, mixing and matching systems as you need. Select **Save** to save the entire group. To close the Static Group Editor, press **Done**.

Group Category Editor

The Group Category Editor provides a means of organizing large numbers of groups by allowing you to create group categories. However, since group categories are by definition static, you cannot drag and drop a task onto a category for execution.

To create a user-defined category of groups from the Groups pane of the Director Management Console, select **New Group Category** from the Groups pane context menu (right-click in any empty space in the Groups pane). This causes the Groups pane to split. The Group Category Editor then appears in the lower half of the Groups pane. Drag and drop the groups you want to add to the new category and click **Save** to name the new category. The category and its group will be displayed as a subcategory.

For more information, see the online help.

Task Based Group Editor

The Task Based Group Editor enables you to create a new dynamic group based on the types of tasks for which the group of systems is enabled. For more details, see the online help.

Group export/import

You can also export groups for later import on another server, for example, or for archival or backup purposes. Only dynamic and task groups can be imported or exported. See online help for details on how to perform this operation.

Managing your groups

You can perform other operations on your dynamic and static groups in the Groups pane. Examples include searching for a particular group, changing the view of the icons, and sorting the groups by name and type.

Bringing up the context menu for a specific group enables you to perform a number of operations on that group, depending on your authority and the type of group you select. Refer to the online help for details.

Tasks

The Tasks pane lists all of the main tasks you are authorized to perform on managed systems. Each user ID has its own level of user authority as part of its configuration.

Several different kinds of tasks can be shown:

- One-to-one actions, such as file transfer, that can only operate on one system at a time.
- One-to-many actions, such as software distribution, that distribute software to many managed systems at once.
- Interactive actions, such as remote control.
- Non-interactive actions, such as software distribution, which could be a scheduled task.
- System actions that are built in or standard and that cannot be deleted.

You can drag and drop task icons onto groups or onto specific managed systems in the Groups and Group Contents panes, or you can drag and drop groups and managed systems onto the tasks you want to perform. You will usually be presented with another window, where you can enter parameters needed for the selected task.

The Tasks pane shows the top level of administrative tasks you can perform. Some tasks have lower level tasks that can be performed after the main task has been selected.

If you have event tasks defined for a group of managed systems, and a new system is added to the group, Director automatically adds the system to the event task.

Using the Scheduler feature of Director, you can define tasks to be performed immediately. You can also schedule tasks to be performed at a later date and time and repeat at a defined interval.

Right-click on an empty space in the pane to perform actions on the task icons. Examples include searching for a particular string, changing the view of the icons (large icons, small icons, list and tree views), and displaying the tasks in ascending or descending order.

Additional Director Management Console features

A menu bar appears across the top of the Director Management Console window. Just below that, a tool bar of icons provides access to console functions.

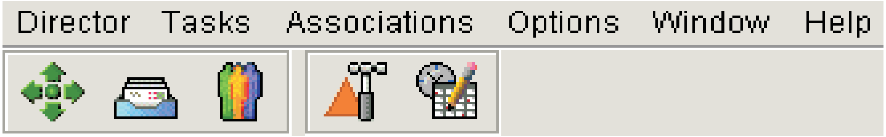

Using the menu bar

From the menu bar, you can perform various tasks. Examples include viewing inventory, performing console security and license administration, and setting user and server preferences. Refer to the online help for details.

Using the tool bar

These icons have the following tasks (listed from left to right across the tool bar):

- **Discover all systems**: Initiates a discovery of all Director and SNMP systems on the network. Inventory will be collected on newly discovered systems.
- **Message Browser**: Brings up the Message Browser window. This window displays messages that have been sent to this system, possibly as a result of an event action.
- **Console Security**: Brings up the Console Security interface. This interface enables you to manage new user accounts and authority for logging into the Director server.
- **Event Action Plan Builder**: Brings up the Event Action Plan Builder window. This window enables you to create event action plans. See Chapter 9, “Event management,” on page 117 for more details.
- **Scheduler**: Enables you to schedule any non-interactive task to occur at another time, such as software distribution. See Chapter 22, “Task Scheduler,” on page 201 for more information.

Using the status bar

At the bottom left corner of the window the IBM logo appears. This logo acts as a progress indicator to let you know the system is performing a task. You will notice a yellow ball move slowly back and forth across the logo as Director performs its tasks.

Across the bottom of the window is a status bar broken into three smaller information windows. These windows show the current console status, in this case `Ready`. This status means the console is idle and waiting for action. In addition, the window shows the server and user identifiers as well as the current time of day.

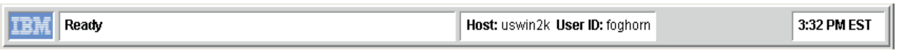

Ready

Host: uswin2k User ID: foghorn

3:32 PM EST

Using the ticker tape

One feature of the Director Management Console is the scrolling ticker tape area near the bottom of the window. You can monitor system attributes without having to view a separate Monitor Console.

A scrolling “ticker tape” line containing information about specific systems or conditions also appears near the bottom, in the space between the status bar and the three main panes. This ticker tape feature serves as a status indicator, providing real-time monitoring of critical resources. You can drag this information from an active monitor’s console to this part of the Director Management Console. See Chapter 8, “Resource Monitoring,” on page 109 for details.

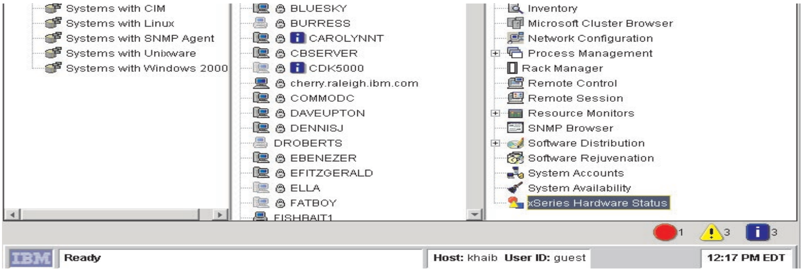

Ready

Host: khaib User ID: guest

12:17 PM EDT

To change the scrolling speed, left-click the scrolling ticker tape to slow it down. Click again to resume the normal speed.

When you right-click on the ticker tape area another context menu is displayed, enabling you to remove monitor attributes from the ticker tape one at a time, or all monitors at once. You can also bring up the Message Browser window to view messages generated by event action plans.

Chapter 6. Inventory Management

Inventory management enables you to quickly and easily display the hardware and software currently installed on your network. Its flexible queries can be used to search for specific CPU types, disk drives, word processors, applications, and installed memory in the Director inventory database. Reports can be saved to an HTML file, an XML file, or a file in Comma Separated Values (.CSV) format. The inventory function includes a dictionary file with many predefined software product profiles, called *product definitions*, that enable you to inventory and track key applications installed on your network systems.

Inventory is collected when a managed system is initially discovered, and during regular intervals. All of this data then becomes valid criteria for configuring a filter when creating a new group. You can set your own frequency of inventory collection, for example, daily or weekly.

You can also select a managed system and invoke an inventory update for it immediately.

Note: Refer to "Navigating in Director" on page 80 for tips on navigating your way through this task, or see the online help for detailed assistance.

Performing an inventory collection

Inventory is collected on all discovered managed systems in the network at system discovery and during regular intervals. You can also perform an inventory collection on a managed system and have it perform the collection immediately.

Note: CIM, DMI, and Static MIF data must be defined to the Director server *before* the Inventory task can collect and present this information. See "Setting up the server to inventory CIM and DMI information" on page 399 for information on setting up CIM, DMI, and Static MIF files.

To provide DMI data, managed systems must be running under Windows Me, Windows 98, Windows 2000, Windows XP, or Windows NT 4.0. They must have a DMI Service Provider (version 2.0 or later) installed. The DMI Service layer is installed with the CIM to DMI Mapper. You can download the DMI mapper from the following website:

`http://www.pc.ibm.com`

The Service Layer does not have to be present when the Director management agent is installed. The Service Layer can be added to a managed system after Director is installed. When the managed system is restarted, it is enabled for DMI operations.

Starting the Inventory Query Browser

Starting the Inventory task using normal drag and drop techniques displays the Inventory Query Browser window in targeted mode. You can also start the Inventory Query Browser in the untargeted mode (for all systems and devices) by double-clicking the Inventory icon. This window is divided into two panes:

Available Queries

Contains a Custom folder, which you can use to store customized queries that you or other users create, and a Standard folder of *default* queries that are defined for you during the installation. Selecting a query causes the corresponding inventory data to be displayed for the managed systems you have selected.

Note: The System User and System Location standard queries retrieve data from user-defined ASCII files. Refer to the Inventory online help for information on setting up the data files for these queries. To access this information quickly, select **standard queries** from the online help index.

Query Results

Displays the results of the queries you select. The query results include only data that is valid for the managed systems targeted.

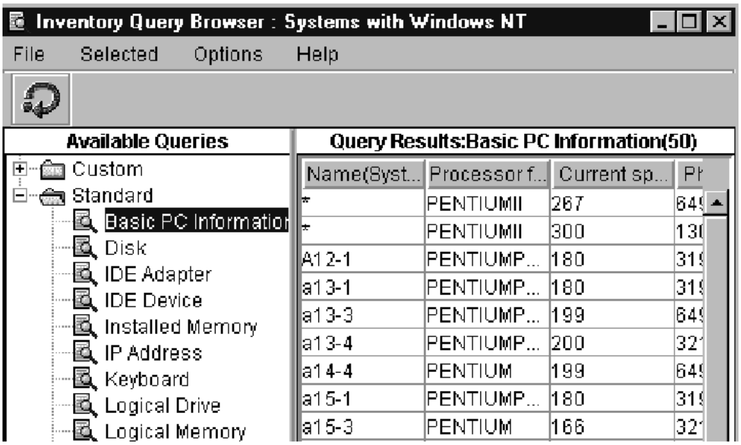

Query Results:Basic PC Information(50)

Name(Syst...	Processor f...	Current sp...	Ph
*	PENTIUMII	267	64
*	PENTIUMII	300	13
A12-1	PENTIUMP...	180	31
a13-1	PENTIUMP...	180	31
a13-3	PENTIUMP...	199	64
a13-4	PENTIUMP...	200	32
a14-4	PENTIUM	199	64
a15-1	PENTIUMP...	180	31
a15-3	PENTIUM	166	32

Additional Inventory Query Browser features

From the context menus of the Inventory Query Browser you can:

- **Copy** a standard query to create a new custom query, which you can edit.
- **Perform** the query as often as required.
- **Modify**, **Rename**, and **Delete** custom queries as desired.

Custom queries are created by selecting **Build Custom Query** (see “Building a customized query” for details).

Updating the list of available queries

Click on the **Refresh Queries** icon in the tool bar at the top of the Inventory Query Browser window to refresh the queries from the inventory database. This updates the view to show custom queries created by other authorized users. This is the same as the Refresh context menu option.

Managing your inventory query results

When an inventory query completes, the results are displayed in the Query Results pane. The results are shown in tabular columns in the order in which they were defined when the query was built or modified, or when the view was modified, whichever occurred last.

You can change the view of this data, re-order columns, hide and show columns, and change the size of the columns, using the techniques described in “Managing columns of information” on page 82. You can save your inventory query results using the standard techniques described in “Saving files” on page 83.

Using menu bar options

Many of the operations included in the menu bar selections File, Selected, Options, and Help have already been described, such as Perform Query, Refresh, Modify, Copy, Rename, Delete, and Build Custom Query.

You can also use the **Export** option to save inventory results in.CSV (spreadsheet), .HTM (HTML), or .XML document format, using the standard techniques described in “Saving files” on page 83. You can also select **Edit Software Dictionary** to add, edit, or remove entries in the software dictionary (see “Using the Inventory Software Dictionary editor” on page 94 for more information).

Building a customized query

There are many useful default queries defined in the Standard folder in the Inventory Query Browser window. If they do not quite match your needs, you can build your own custom inventory queries, using the Inventory Query Builder window.

Using the Inventory Query Builder

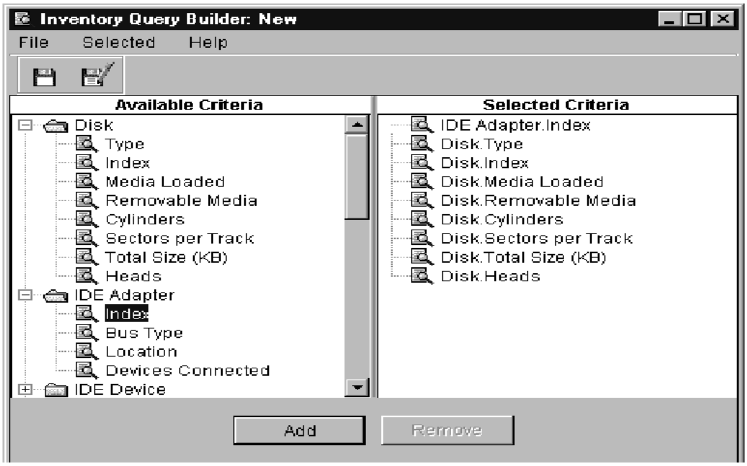

The Inventory Query Builder is divided into two main panes: Available Criteria and Selected Criteria. Drag the desired data items from the Available Criteria pane to the Selected Criteria pane, or click the **Add** and **Remove** buttons to create your query in the Selected Criteria pane (see “Using the Add and Remove buttons” on page 82). You can mix and match and order your query choices however you like. You can select entire folders or individual data items in each folder. You can have multiple tables open at once, and can move back and forth between them, selecting items to add to the query being built.

The managed systems that you initially select for inventory tasks have associated sets of *database tables*, which contain the relevant inventory data. See the online help entry at “Inventory Database Tables” for more information.

Using the Inventory Software Dictionary editor

You can use the Inventory Software Dictionary editor to track software on your PC-managed systems.

The inventory software dictionary editor enables you to associate the name of a software application with one or more specific files on a PC-managed system. You can also specify exact file sizes, last-modified dates, and so on, to refer to a specific level or software release.

Using this file matching technique, you can collect software inventory information on your systems and know exactly what applications are installed and what levels, so you can determine if upgrades are needed or if other maintenance actions should be performed.

The Inventory Software Dictionary Editor window consists of two panes: Available Entries and Entry Definition. The Available Entries pane contains a tree listing of all available software categories representing thousands of applications that may or may not reside on your PC-managed systems. Expand these category folders to show groups of applications, and then select the application of interest.

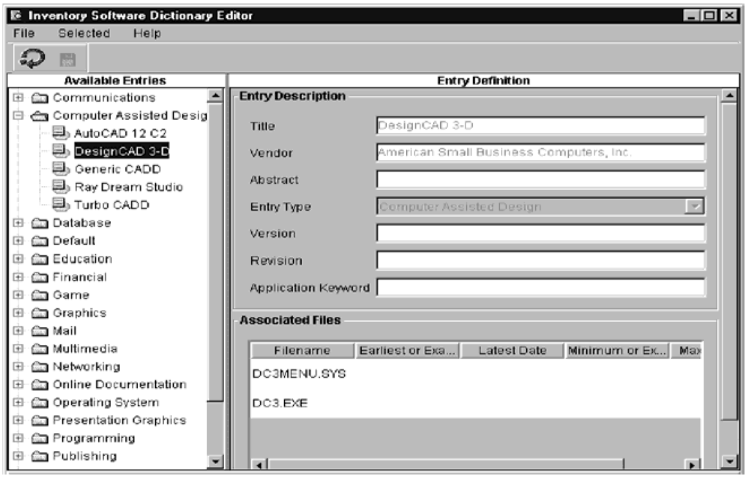

Title: DesignCAD 3-D

Vendor: American Small Business Computers, Inc.

Abstract:

Entry Type: Computer Assisted Design

Version:

Revision:

Application Keyword:

Filename	Earliest or Exa...	Latest Date	Minimum or Ex...	Max
DC3MENU.SYS				
DC3.EXE				

The pertinent information about that application is then retrieved and displayed in the Entry Definition pane, showing the name, vendor, and so on, entered for that application in the Entry Description fields. In the Associated Files area, a list of files which have been associated with this particular application and software level are also shown. Some files will also have dates of when they were last modified, or specific file sizes to further distinguish one specific software level from another.

When software inventory is performed on a system, these files are detected and compared to the information in the Inventory Software Dictionary. When an exact match is found, that particular application is listed in the displayed inventory for the system.

You can add new entries to this dictionary, specifying which category it should be grouped under. You can add associated files manually or from a file list, and specify exact dates and sizes to distinguish this entry from all others.

Managing Your Software Dictionary entries

Select the **Modify** operation from the context menus to change the information in the Entry Definition section.

You can modify entries in the Associated Files section, using the Edit and Remove buttons which appear when you highlight an entry. You can modify the associated file name, file date, and file size information of each associated file or delete the entire row. You can also re-order the file names or show, hide, and resize columns as desired, using the standard techniques.

Select the **Delete** operation to delete entries and associated files from the library.

Select the **Refresh** operation from the context menus or the **Refresh** icon in the tool bar to refresh the list of applications listed in the Available Entries pane. This is useful to see changes made by other authorized users.

Select **File** →**Close** to close the Inventory Software Dictionary window.

Performing batch operations on the Software Dictionary file

To maximize performance and conserve disk space, the Director software dictionary file is maintained in a binary format that cannot be edited. To add entries to the file in batch mode and to convert the dictionary entries into an editable format, Director provides the `TWGCLI` utility to perform the following software dictionary file operations:

- Export all entries to a Java properties file
- Import entries from a Java properties file
- Import relevant information from a Microsoft package definition file (PDF)
- Merge two software dictionary files

Requirements for using TWGCLI

The following requirements apply to using `TWGCLI`:

- Stop the Director server to release control of the software dictionary file before you use `TWGCLI`
- Run `TWGCLI` on the Director server. You cannot execute `TWGCLI` from the console

When a `TWGCLI` operation is finished, restart the Director server.

Exporting entries to a properties file

This function generates a Java properties file from a software dictionary file. You can export the entries to a properties file, use a text editor to add, delete or change the properties file entries, and then use the import function to convert the properties file back into a software dictionary file.

Command Syntax: `TWGCLISWDictionaryReader[`*`target`*`][-options]`

where *target* is the path and name of the properties file to be written. This file must have an extension of .properties. The default name is mastrsid.properties.

Each of the ***options*** must be preceded by a hyphen (**-**) or slash (**/**) character, and can be the following:

-h,-?,-help
Displays the syntax of the `TWGCLISWDictionaryReader` and associated options.

`-dict` *file*
Specifies the path and name of the software dictionary file to be read from. This file must have a file type of .sid. `c:\TivoliWg\Classes\com\tivoli\twg\inventory\defautl.sid` is the default file.

-sid *file* Same as **-dict *file***

-d *dir* Specifies the name of the directory of the properties file to which the converted dictionary entries are written. The default is `\TivoliWg\data\`. If *target* specifies an absolute path name, this option is ignored.

The following examples assume a Director installation directory of `c:\TivoliWg`:

```
TWGCLI SWDictionaryReader
```

Reads the default software inventory dictionary c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid and writes results to the default properties file `c:\TivoliWg\Data\mastrsid.properties.`

```
TWGCLI SWDictionaryReader -sid
D:\Data\Dictionaries\other.sid
```

Reads the specified software inventory dictionary `D:\Data\Dictionaries\other.sid` and writes results to the default properties file `c:\TivoliWg\Data\mastrsid.properties.`

```
TWGCLI SWDictionaryReader dict.properties
```

Reads the default software inventory dictionary c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid and writes results to the specified properties file (dict.properties) in the default output directory `c:\TivoliWg\data.`

```
TWGCLI SWDictionaryReader -d D:\Data
```

Reads the default software inventory dictionary `c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid` and writes results to the default properties file (mastrsid.properties) in the specified output directory `D:\Data.`

Importing entries from a properties file, Microsoft PDF, or Software Dictionary file

This function imports the contents of a text properties file, Microsoft Package Definition File (PDF), or software dictionary file and adds the imported entries to a target software dictionary file.

Command Syntax: `TWGCLISWDictionaryWriter source[-options]` where *source* is the path and name of the file from which the software dictionary entries are imported. This file must have an extension of .properties, .pdf, or .sid. This parameter is required.

Each of the *options* must be preceded by a hyphen (**-**) or slash (**/**) character. Options are not case sensitive. Which options are available depend on the type of file being imported:

Global Options:

-h, -?, -help
Displays the syntax of the `TWGCLISWDictionaryWriter` and associated options.

-dict *file* Specifies the path and name of the software dictionary file to be changed (read to). `c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid` is the default file. The target software dictionary file is backed up to a file with the name *target_N*, where *N* is a positive integer.

-sid *file* Same as **-dict** *file*

-d *dir* Specifies the name of the directory where the target software dictionary file is written.The default is `\TivoliWg\data\.` If *target* specifies an absolute path name, this option is ignored.

Properties File Options:

-n, -new Specifies to create a new software dictionary file using the source properties file. All existing entries in the software dictionary file are cleared.

PDF Options:

-cat *category*
Specifies the application category for the entries imported from this file. ***Category*** codes are:

Application Category	Value
CAD	19
Communications	2
Database	5
Default	0
Desktop Publishing	4
Education	13
Financial	9
Game	10
Graphics	12

Application Category	Value
Mail	6
Multimedia	11
Networking	1
Online Documentation	18
Operating System	14
Presentation Graphics	16
Programming Tools	15
Server	7
Spreadsheet	8
System Management	17
Word Processing	3

Examples:

```
TWGCLI SWDictionaryWriter word50.pdf -cat 3
```

Reads the specified Microsoft PDF file (word50.pdf) and writes results to the default software inventory dictionary c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid,(C:\TivoiiWg\, using application category 3, Word Processing.

```
TWGCLI SWDictionaryWriter new.properties -new
```

Reads the specified properties file (new.properties) and writes results to the default software inventory dictionary (C:\TivoliWg\Classes\com\tivoli\twg\inventory\default.sid), erasing the current contents of the file.

```
TWGCLI SWDictionaryWriter update.sid -dict
D:\Data\Dictionaries\Other.sid
```

Reads the specified software inventory dictionary (update.sid) and merges its entries with the specified software inventory dictionary D:Data\Dictionaries\Other.sid.

Modifying inventory collection preferences

You can configure how often inventory data is refreshed as well as the response time for attempted inventory refreshes, by selecting **Options→ Server Preferences**, and selecting the **Inventory collection preferences** tab.

In the Timeout Period settings field, enter the number of minutes to wait for an inventory refresh to be completed. If no response is received by this time limit the refresh is abandoned. The default value is 10 minutes.

In the Refresh Interval settings field, enter the interval of time desired between automatic refreshes of the inventory database. The default is set at 7 days.

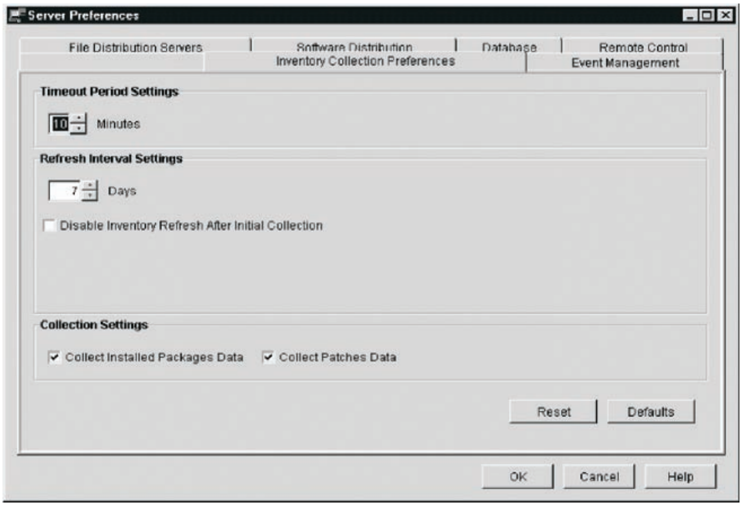

Timeout Period Settings

10 Minutes

Refresh Interval Settings

7 Days

☐ Disable Inventory Refresh After Initial Collection

Collection Settings

☑ Collect Installed Packages Data ☑ Collect Patches Data

You should select the **Disable Inventory Refresh After Initial Collection** check box if you do not want to automatically refresh the inventory database. If this is selected, only the initial inventory after the discovery of systems is performed. No other automated inventory update will occur.

Inventory collection consumes a significant amount of processor resources on the managed systems, so certain kinds of data are not collected by default. The Collection Settings boxes allow you to enable collection of these kinds of data. The kinds of data collected depend on the operating system of each managed system. Therefore the following options may not apply to specific managed systems:

- The **Collect Installed Packages Data** box enables collection of inventory data by querying operating system-specific APIs or system log files to determine which software packages have been installed on applicable managed systems.
- The **Collect Patches Data** box enables collection of inventory data about installed patches on applicable managed systems.

If you have changed any of these settings but wish to return to the values which were in effect at the last setting, press the **Reset** button. If you wish to return to the default values of 10 minutes and 7 days, press the **Defaults** button. To view the online help with this window, press the **Help** button.

When you are finished with these selections, press **OK** to save the settings or **Cancel** to quit without saving any changes. The Server Preferences window will then be closed.

Chapter 7. Remote Control

Remote Control enables you to manage a remote system by displaying the desktop of a remote managed system within a Director Management Console and by sending keyboard and mouse information to the remote managed system. You can also view a listing of all the consoles that have remote sessions with the managed system, and see the controlling state of each. See “Viewing a listing of current remote control sessions” on page 108 for more information.

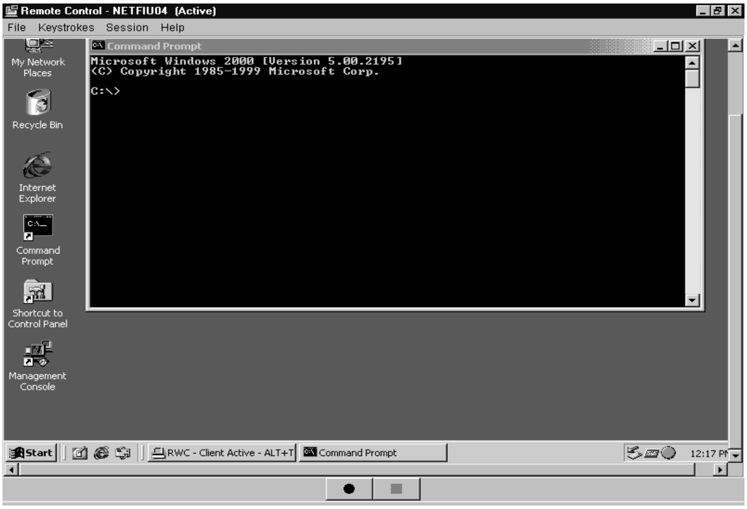

```
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

C:\>
```

Control states

Remote Control uses three control states to manage remote systems:

Active state
(Default) Remote control mode. A managed system in the active state can be controlled from a Director Management Console at a remote location. When a Director Management Console assumes control of a managed system in the active state, the screen image of the managed system is displayed within the Director Management Console and keyboard and mouse information originating from the console are passed through and executed on the remote system. Only one console can be in control of a specific remote system in the active state; all other attached consoles can only monitor the display of the system. The screen image is updated automatically at the active console when a change occurs on the display of the remote system.

Monitor state
View-only mode. A managed system in the monitor state is not under the control of a Director Management Console. Either the local user or

the active console has control of the managed system. If a change occurs on the display of the managed system, the screen image updates automatically on any console that has a remote control session in the monitor state with that managed system.

Suspend state

View-only mode without image refresh. A managed system in the suspend state does not update its screen image on any attached Director Management Consoles if the screen image changes. The user of the managed system has control of his or her desktop. When a managed system enters the suspend state, all attached consoles do not receive updates if the screen image of the managed system changes.

Overriding and changing control states

During initial configuration, all managed systems installed with the Director management agent are set to start up in an active state. Any remote Director Management Console can then assume initial control over any accessible managed system by establishing an active remote control session with that system.

Control states can be set by the Director Management Console and by the native managed system itself.

Requesting active control from a management console

If a console already has a remote control session with a managed system in the active state, you can request the controlling console to turn over control to your console. When you ask for control, the controlling console can refuse the request. If a time-out occurs before the request for control is processed, the default action is to transfer control to the requesting console and the original controlling console is put into the monitor state for the managed system.

Changing control states from the managed system

The user of the managed system can change or regain control at any time by pressing ALT+T on the managed system. When ALT+T is pressed a pop-up message appears which allows you to select control states for the managed system. There are four control states: active, monitor, suspend or terminate. When a managed-system user changes the control state of the system, the change affects all remote control sessions that are established with the managed system at the time the control state of the system is changed.

Control state scenarios

Scenario 1: Assume that a native managed system is in an active mode and multiple Director Management Consoles have remote control sessions with the system. In this scenario, only one console can be in a controlling active state with the managed system and all other consoles must be in either monitor or suspend state with the managed system. If the console in an active state changes to monitor state, the state of the managed system automatically changes to a monitor state. At this point, any attached console can assume control of the managed system by changing the session state to active.

Scenario 2: Assume that a managed system is in monitor state and multiple Director Management Consoles have remote control sessions with the managed system in either monitor or suspend states. The managed system can change its state to active, which would force the state of the first console that is notified into a controlling active state. All other attached consoles would remain in either monitor or suspend state.

Scenario 3: Assume that a managed system is in an active state, and multiple Director Management Consoles have remote control sessions with the managed system. If the user of the managed system changes the state of the system to suspend, all attached consoles automatically change to the suspend state. However, any of the attached consoles can change the state of a remote control session to either active or monitor state.

Remote control usage restrictions

There are several restrictions in using remote control. These are itemized in the section "Remote control" on page 35. Please refer to this section before attempting to perform remote control on your managed systems.

Remote access security

During the configuration of network drivers, the Remote User Authorization for Screen Access option can be enabled. This can be achieved during the process of installing the Director management agent or by bringing up the Network Driver Configuration window(**Start →Programs →IBM Director →Configuration**), or using the icon in OS/2. If you attempt remote control access to a managed system that has this option enabled, the user of the remote system can accept or reject the access attempt. If the user does not respond to the request within 15 seconds, your attempt is rejected.

Sending keyboard information to a remote system

When remote control is in an active state, nearly all key and key combinations are automatically passed through to the remote system. However, operating system requirements restrict the use of certain key combinations, for example, **Ctrl+Alt+Del**, which typically generates an interrupt that is intercepted and processed by the operating system of the local system.

To bypass certain key restrictions, select the desired key combination from the Keystrokes option in the menu bar at the top the window. The following selections are available:

- Alt+Esc
- Alt+Tab
- Ctrl+Esc
- Ctrl+Alt+Del

Numeric keys sent from the numeric keypad (typically on the right-hand side of the keyboard) are not differentiated from the numeric keys at the top of the keyboard.

During a remote control session, restricted keys such as the Tab key and the F1 through F12 function keys are displayed at the bottom of the screen for you to select as needed. You can click on one of these keys to perform the same function as when you press the key on the keyboard.

Remote control and inventory

Remote control is somewhat dependent on the inventory function of Director to provide information about the managed system. Be sure to run the inventory collection task against any systems on which you plan to perform remote control operations.

Type of operating system

If you sent a Ctrl+Alt+Del key sequence to a remote system running Windows 98, the remote system would lock up. An inventory of the managed system tells Director what type of operating system the managed system is running, and the Ctrl+Alt+Del capability will be enabled or disabled appropriately.

Code page for screen transfer

Taking inventory of the managed system tells Director which code page to use for proper screen transfer information from the managed system. Therefore, you should always perform an inventory on your remote managed systems before using remote control.

Restrictions on pointer and cursor support

Because the remote control service operates in the Java environment, pointer changes on the managed system may not be displayed on the controlling console. For example, the managed system may change the pointer to the up/down sizing arrows when it is over the border of a window, but the controlling console will continue to show the pointer in its normal state.

A console which has a session with a remote managed system in monitor mode will not see the cursor movement of the remote system, but will see screen changes as they occur on the desktop of the remote managed system.

Performing remote control tasks

For information on starting and stopping the remote control task and performing remote control operations, select **Help →Topics → Remote Control** from the Director Management Console. The tasks are also described briefly here:

- Starting and stopping a remote control session with a remote managed system
- Changing the control state of a remote control session
- Recording a remote control session
- Viewing a listing of current remote control sessions
- Changing the refresh rate for current remote control sessions

You can start a remote control session from the Director Management Console by using the drag and drop method between managed systems and the remote control icon the Tasks pane. When you first initiate a remote control session, the display window is placed in the active state. To change to another state, select the state from the Session menu. To view the list of current remote control sessions, select **Console List** from the Session menu. To end a session and close the remote control service, close the Remote Control window.

Starting a remote control session

You can start a remote control session from the Director Management Console by using the normal drag-and-drop methods between managed systems and the Remote Control icon in the Tasks pane, or from the context menu of the managed system. See “Navigating in Director” on page 80 for tips on navigating your way through this task, or see the online help for detailed assistance.

Stopping a remote control session

You can end a remote control session by:

- Closing the remote managed system window.
- Selecting **File** → **Close** from the top of the window.
- Using the Alt+T key combination directly on the managed system.

Changing the control state of a session

You can change the control state of the session by clicking **Session** at the top of the controlling console and then selecting a control state (Active, Monitor, or Suspend).

Recording a remote control session

You can record the screen output of a remote control session into a file for playback later. To begin saving the screen images, select **File** → **Start Session Logging**. Enter a name for the log file you are creating. The remote control session is then continuously recorded until you end the session log by selecting **File** → **Stop Session Logging**.

After you end the remote control session log, the log file appears as a subtask under the Remote Control icon on the Director Management Console. To replay a log file, double-click the selected log file icon.

If a usable data file is found, the remote control session recording is played back at normal speed. The playback utility can also pause or stop the recording. The Stop button resets playback to the beginning of the file.

Viewing a listing of current remote control sessions

You can view a list of all the Director Management Consoles which have remote sessions with the managed system and see which one is in control. Select **Session** → **Console List…** at the top of the Remote Control window, and the Remote Control Console List window will be displayed.

You can sort (in ascending or descending order) these entries for easier viewing by right-clicking anywhere in the window.

Changing the refresh rate for current remote control sessions

You can adjust the refresh rate for Director Management Consoles that have active remote control sessions. The refresh rate determines how often the screen image is captured and displayed to the controlling console. To change the refresh rate, select **Session** → **Refresh Rate** and choose from the options list:

- Fastest – screen refresh with no delay
- Fast – screen refresh every two seconds
- Medium – screen refresh every 10 seconds
- Slow – screen refresh every 30 seconds

You can change the refresh rate only for consoles in the Active state. If the console is in the Monitor state, you can see the current setting but cannot change it. If the monitor is in the Suspend state, the Refresh Rate menu entry is disabled.

Chapter 8. Resource Monitoring

The Resource Monitoring task enables you to view statistics on critical system resources, for example, CPU, disk, file, and memory usage.

When monitor data indicates a problem or potential problem with network resources, you can set thresholds and trigger events according to the requirements of your site. You can respond to resource monitor events by specifying event action plans. See Chapter 9, “Event management,” on page 117 for more information on event action plans.

Other monitors can be set up to monitor specific processes and system applications. See Chapter 21, “Process management,” on page 195 for details.

Understanding monitors

Director monitors use monitoring agents on the managed systems to enable the gathering of data at the Director server. These monitoring agents gather and forward sampled data to the Director server where it is stored for viewing. Gathered data is time-stamped and refreshed at regular intervals.

Monitoring categories are found in the Available Resources pane. Each monitoring category expands to show subcategorizes. Each subcategory expands to show *attributes*. For example, the File Monitors is a monitor category. ACROBAT.PDF is a subcategory and checksum is an attribute.

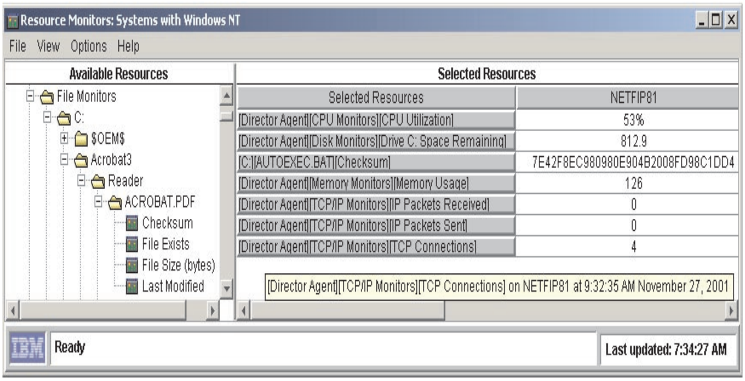

Selected Resources	NETFIP81
[Director Agent][CPU Monitors][CPU Utilization]	53%
[Director Agent][Disk Monitors][Drive C: Space Remaining]	812.9
[C:][AUTOEXEC.BAT][Checksum]	7E42F8EC980980E904B2008FD98C1DD4
[Director Agent][Memory Monitors][Memory Usage]	126
[Director Agent][TCP/IP Monitors][IP Packets Received]	0
[Director Agent][TCP/IP Monitors][IP Packets Sent]	0
[Director Agent][TCP/IP Monitors][TCP Connections]	4

[Director Agent][TCP/IP Monitors][TCP Connections] on NETFIP81 at 9:32:35 AM November 27, 2001

Ready

Last updated: 7:34:27 AM

Most attribute data is displayed in numerical format, for example, to indicate percentages or numbers of occurrences. Some attribute data is displayed in text format, for example, online or offline, to indicate the status of the system or application.

The Director server can monitor data from native managed systems and SNMP devices.

Monitoring data on native managed systems

You can monitor data for native managed systems running on remote systems using any of the supported operating systems.

The number of attributes you can monitor on native managed systems varies depending on the operating system that is running on the system. The following monitors are generally present on all native managed systems:

- File Monitors
- CPU Monitors
- Memory Monitors
- Disk Monitors

If the system is running Windows NT 4.0, the Director monitoring agent uses the Windows NT performance monitors to provide thousands of additional attributes.

Monitoring data on native managed systems configured with additional services

The Director monitoring agent will also interface with the APIs of the following services on native systems:

Desktop Management Interface (DMI)
The DMI service layer can be accessed to present corresponding attributes under DMI Monitors. It is installed with the CIM to DMI Mapper. Please refer to your DMI mapper utility for instructions on installing and configuring the DMI mapper for reporting data. To provide DMI data, managed systems must be running under Windows 98, Windows 2000, or Windows NT 4.0 and must have a DMI Service Provider (version 2.0 or later) installed.

Common Information Model (CIM)
CIM services can be accessed to present corresponding attributes under CIM Monitors. To provide CIM data, managed systems must be running under Windows 98, Windows 2000, or Windows NT 4.0 and must have Windows Management Interface (WMI) Core Services Version 1.1. installed.

Microsoft Clustering Service (MSCS)
The MSCS can be accessed to present corresponding attributes under Cluster Monitors. To provide cluster data, managed systems must be running under Windows 2000 or Windows NT 4.0 with Service Pack 5 or 6 and must have Microsoft Clustering Service installed.

Monitoring data on SNMP devices

To monitor data for an SNMP device, the remote system must be using either IP or Internet Packet Exchange (IPX) transport protocols to communicate with the Director server.

SNMP devices have a basic set of attributes available for monitoring. SNMP devices with the RMON Management Information Base (MIB) provide even more attributes for monitoring. See Chapter 14, “SNMP Management,” on page 153 for more details.

Monitoring data on Windows NT devices and services

You can also monitor the status of a Win32 device or device service by setting an individual threshold.

Starting Resource Monitors

The Resource Monitors task is started from the Director Management Console by using the standard drag-and-drop methods or by selecting **Resource Monitors** from the context menu of the managed system. (See “Navigating in Director” on page 80 for tips on navigating your way through this task, or see the online help for detailed assistance.)

The Resource Monitors task has two subtasks: All Available Recordings and All Available Thresholds. These subtasks provide you with a quick overview of the data recordings and thresholds that have been set, and enable you to perform operations, such as ending a recording or removing a threshold. Refer to the online help for information on performing operations through these subtasks.

You can also create additional views of specific monitored attributes. These views are also placed under the Resource Monitors icon in the Tasks pane as subtasks:

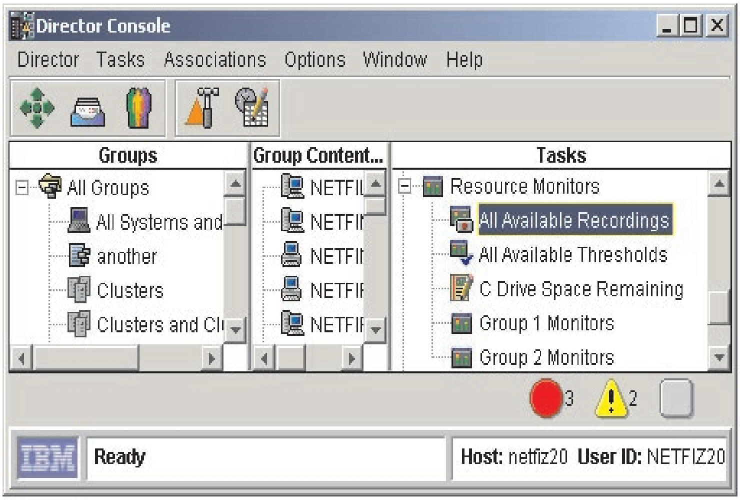

You can start a Resource Monitors subtask by dragging it to a managed system.

Using the Resource Monitor window

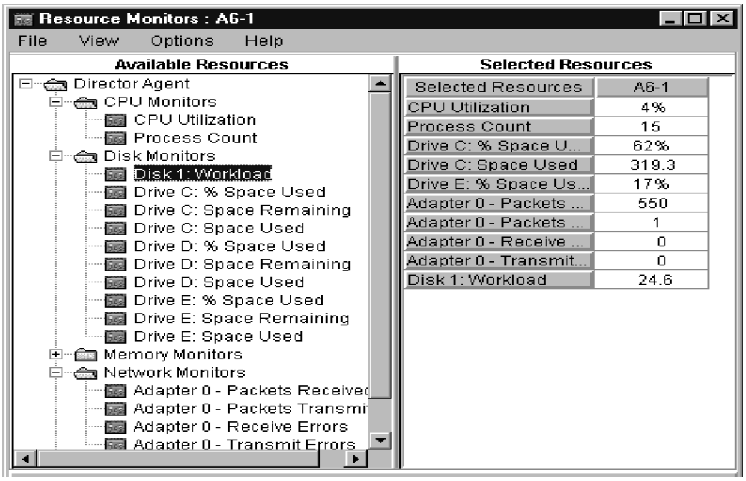

Selected Resources	A6-1
CPU Utilization	4%
Process Count	15
Drive C: % Space U...	62%
Drive C: Space Used	319.3
Drive E: % Space Us...	17%
Adapter 0 - Packets ...	550
Adapter 0 - Packets ...	1
Adapter 0 - Receive ...	0
Adapter 0 - Transmit...	0
Disk 1: Workload	24.6

The Resource Monitors window consists of two panes: Available Resources and Selected Resources.

Note: The attributes that are displayed include all those which are available on the targeted managed systems which can be accessed. If the accessibility of targeted managed systems changes, the available attributes may also change, and will be reflected in the Resource Monitors window when the attributes are refreshed.

The Selected Resources pane displays a table identifying the targeted system names across the top row and the corresponding attributes in the left-hand column.

Initiating a Resource Monitor

You can select attribute data from the Available Resources pane and scroll through the resulting monitor data displayed in the Selected Resources pane using the normal methods (see “Navigating in Director” on page 80).

Viewing monitor data on the ticker tape

You can monitor your managed systems from the Director Management Console using the ticker tape display feature (see “Using the ticker tape” on page 89 and the online help for details).

Setting monitor thresholds

If you assign a threshold for a given attribute, an event is generated when the threshold is met for the system to which the attribute applies.

For example, you could set a threshold on a file server to generate an event if there is less than 100 MB of free space on the disk drive. When the threshold is set, the free space on the server is monitored and when it goes below 100 MB, the event is generated. This event could then be sent to an alpha-numeric pager so you could be notified immediately. You can also create the same threshold on multiple systems. Refer to Chapter 9, “Event management,” on page 117 for more information on events and actions.

Most thresholds are numeric in value, expressed either as a discrete number or a percentage. You can also set a text string threshold, where a particular text string is monitored and an event is generated if the text changes from what is desired or expected. For example, if a critical system must always be up, you can set a threshold to trigger when the system goes offline.

You can set thresholds for a specific managed system. You can also create threshold plans, which are a collection of thresholds. A threshold plan can then be exported to a file that can be imported at a later time for use on other systems or for archival purposes. A threshold plan task allows you to drag and drop a threshold plan onto another system as well. See the online help for detailed assistance on creating, exporting, and importing threshold plans.

You can view individual thresholds set on selected resources, or you can view all thresholds as well as enable and disable individual thresholds. You can sort the order of the thresholds, highlight and delete any of the thresholds, refresh the view, adjust the column width and placement as desired, and modify the view by setting the level of attributes in a path to display in the Selected Resources pane. See the online help for details.

Setting numeric thresholds

When you set a numeric threshold for a single managed system, you are presented with the System Threshold window.

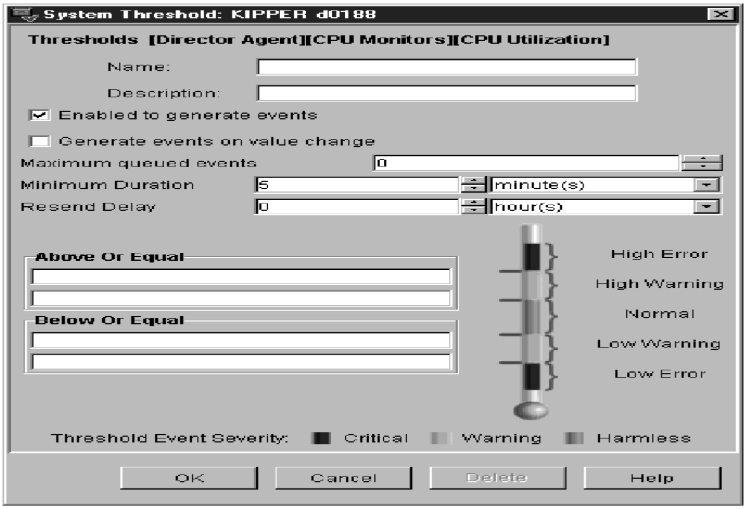
System Threshold: KIPPER d0188

Thresholds [Director Agent][CPU Monitors][CPU Utilization]

Name:

Description:

☑ Enabled to generate events

□ Generate events on value change

Maximum queued events: 0

Minimum Duration: 5 minute(s)

Resend Delay: 0 hour(s)

Refer to the online help for details on setting thresholds.

The event type generated is listed at the top of the System Threshold window. In the example shown above, the event type is set to:

```
[Director Agent][CPU Monitors][CPU Utilization]
```

Depending on which threshold value is exceeded, [High] or [Low] will be appended at the end of the event type, along with the particular severity of [Warning] or [Error].

Setting text string thresholds

When you set a string threshold for a single managed system, you are presented with the System Threshold window.

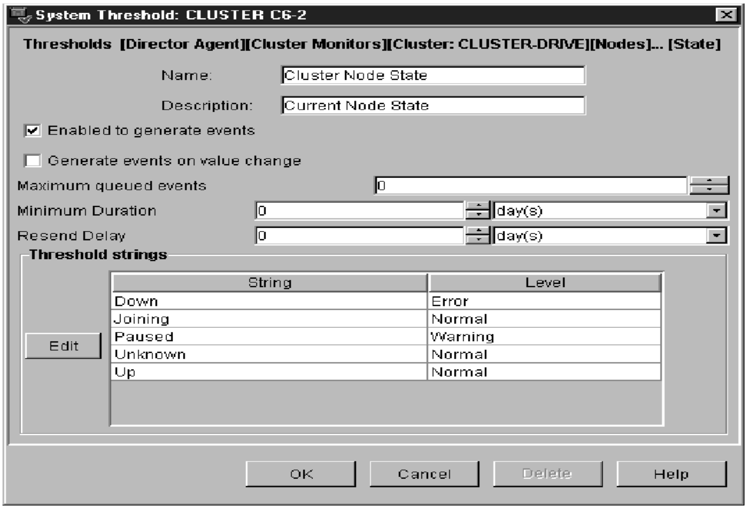

System Threshold: CLUSTER C6-2

Thresholds [Director Agent][Cluster Monitors][Cluster: CLUSTER-DRIVE][Nodes]... [State]

Name: Cluster Node State

Description: Current Node State

☑ Enabled to generate events

□ Generate events on value change

Maximum queued events: 0

Minimum Duration: 0 day(s)

Resend Delay: 0 day(s)

Threshold strings

String	Level
Down	Error
Joining	Normal
Paused	Warning
Unknown	Normal
Up	Normal

See the online help for details on setting string thresholds. Use the All Available Thresholds subtask to view the threshold setting.

Recording monitor data

Your selected monitor data is refreshed and displayed in the Selected Resources pane of the Resource Monitors window at regular intervals, but it only shows the most recent value since the last refresh.

You can set up a time period during which each refreshed monitor reading is recorded. Any time during or after the recording period, you can generate simple line graphs or export the data to a file in .CSV (Spreadsheet), .HTM (HTML), or .TXT (flat ASCII) format. Use the **All Available Recordings** subtask to view recordings.

Managing your monitored resources

After you have created a set of monitor attributes in the Selected Resources window, you can save them and apply them again later to other managed systems.

You can run multiple Monitor Consoles at the same time, by dragging systems to the Monitor Console icon, or conversely. Each time you do this, a new Monitor Console window is opened.

See the online help for details on other operations you can perform on your monitored resources in the Selected Resources pane.

Chapter 9. Event management

The Director event-management task enables you to identify and categorize network events, and automatically initiate actions in response to those events.

For example, you might have used the resource monitor task (see Chapter 8, “Resource Monitoring,” on page 109) to configure a threshold on your file server to generate an event when the remaining free space on the main data drive drops below 100 MB. Now, using event management, you can configure an event action plan that causes you to be automatically paged when the threshold is reached. As an administrator, you will know when the file server hard drive is approaching its capacity and can take corrective action before your users are impacted.

New terms in this chapter

The following terms are used in this chapter:

Event An event is a means of identifying a change of state of a process or a device on the network. For example, an event identifies when a workstation changes from an online state to an offline state in the network, or when a critical resource threshold, such as virtual memory utilization, is met. It is a notification that something has occurred.

Event filter
An event filter describes a set of characteristics (for example, severity and event type) which are used to differentiate a single event. Director provides predefined event filters and a utility that enables you to create custom filters.

Actions Actions define the steps to take in response to an event, for example, entering the event in the event log or executing a command. Director provides a set of predefined actions that you can customize for your network needs.

Event action plan
An event action plan binds an event filter to one or more actions. For example, an event action plan can be created to send a page to the network administrator’s pager if an event with a severity level of critical or fatal is received by a Director server. You can include as many event filter and action pairs as needed in a single event action plan.

Understanding event management

The Event Action Plan task enables you to:

- Create and apply new event action plans.

Using the Event Action Plan Builder, you can create new event action plans, event filters, and customize actions. Event filters and customized actions can then be logically associated to form event action plans. The resulting event action plans can then be applied to one or more managed systems or groups to perform actions in response to specific events.

- Manage event action plans.

  The **Associations→ Event Action Plans** selection in the Group Contents pane of the Management Console enables you to determine the systems to which an event action plan has been applied. You can also remove applied event action plans in the Group Contents pane. Event actions and event filters are edited and deleted in the Event Action Plan Builder window.

- Log and view event details.

  Events are recorded in the Event Log and you can view, sort, and delete these log entries as desired. You can also tailor the view to show only those events that occur on specific systems, or limit the view to only show predefined event action plans.

Creating an event action plan

You can create a new event action plan using the Event Action Plan Builder. You build the event action plan by associating event filters and customized actions to the event action plan.

The Event Action Plan Builder has three panes:

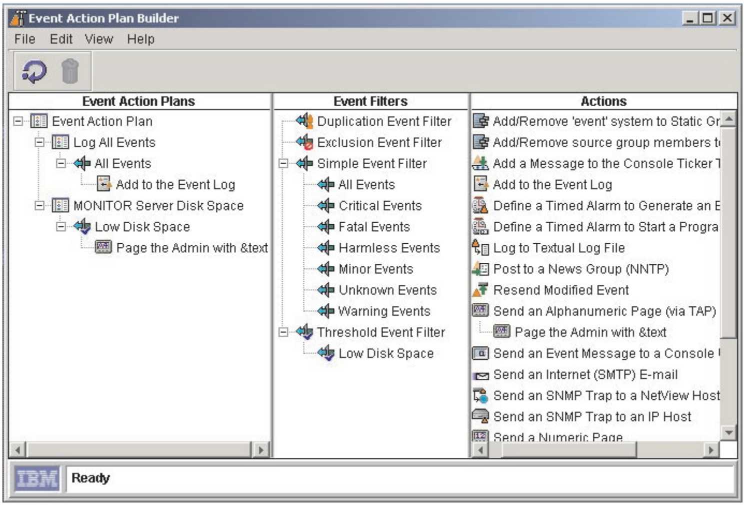
Event Action Plan Builder
File
Edit
View
Help
Event Action Plans
Event Action Plan
Log All Events
All Events
Add to the Event Log
MONITOR Server Disk Space
Low Disk Space
Page the Admin with &text
Event Filters
Duplication Event Filter
Exclusion Event Filter
Simple Event Filter
All Events
Critical Events
Fatal Events
Harmless Events
Minor Events
Unknown Events
Warning Events
Threshold Event Filter
Low Disk Space
Actions
Add/Remove 'event' system to Static Gr
Add/Remove source group members t
Add a Message to the Console Ticker T
Add to the Event Log
Define a Timed Alarm to Generate an E
Define a Timed Alarm to Start a Progra
Log to Textual Log File
Post to a News Group (NNTP)
Resend Modified Event
Send an Alphanumeric Page (via TAP)
Page the Admin with &text
Send an Event Message to a Console
Send an Internet (SMTP) E-mail
Send an SNMP Trap to a NetView Host
Send an SNMP Trap to an IP Host
Send a Numeric Page
IBM
Ready

Event Action Plans
: contains the Event Action Plan templates and all user-defined plans, with associated event filters and actions in a tree structure.

Event Filters
contains the list of predefined event filters and user-created event filters.

Actions contains the list of event action templates supplied by Director. You select a template and customize it to perform a specific action. When you save the new action, it is added as a subtask under the template used to define the action.

Note: On Unix, the Send a Numeric Page and Send an Alphanumeric Page (through TAP) event action templates do not function. These actions are disabled to avoid contention problems over the modem with other applications.

Building an event action plan is simply a matter of creating a new event action plan, dragging one or more event filters from the Event Filters pane and dropping them onto the desired event action plan icon in the Event Action Plans pane, and then dragging one or more customized actions from the Actions pane and dropping them onto the desired event filter associated with that event action plan. You can expand the tree structure under the event action plan icon and show all of the event filters associated with it. You can then do the same for the event filter icon and see the actions associated with that event filter. Note that the drag-and-drop function is one-directional; you can drag actions and filters to event action plans, but you cannot drag an event action plan icon over to an event filter or action.

Using predefined event filters

Predefined event filters are supplied by Director and listed in the Event Filters pane. They are designed to meet many of the basic monitoring requirements of your network environment; however, you can modify them to suit your particular needs as well.

See “Assigning an event filter to an event action plan” on page 120 to associate predefined event filters to an event action plan.

Creating an event filter

Use the Event Filter Builder window to create filters that meet the needs of your networking environment. Select the Event Action Plan Builder icon in the Management Console to display the Event Action Plan Builder window. To open the Event Filter Builder window, right-click in the Event Filters pane and select **New** from the context menu. Choose one or more event categories in the Event Filter Builder window, such as the time and day the event occurred, severity of the event, originator of the event, type of event, and extended attributes.

To create a targeted event filter for an event that has already occurred, open the event Log, right-click on the event, and click **Create**. Note that the Event Type category corresponding to the event you selected is already highlighted (selected).

See the online help for procedures on selecting event filtering criteria.

Assigning an event filter to an event action plan

You can associate an event filter to your event action plan using normal drag and drop and context menu selection techniques (see “Navigating in Director” on page 80 for tips on navigating your way through this task, or see the online help for detailed assistance).

When you add a filter to a plan, the filter icon appears under the event action plan icon in the Event Action Plans pane of the Event Action Plan Builder window.

Customizing an action

Each event filter you assign to an event action plan can have one or more actions associated with it. When an event occurs in the network that satisfies the filtering criteria, the action is performed.

Director provides predefined action templates that you can copy and customize for your particular needs. These templates are shown in the Actions pane of the Event Action Plan Builder window. See the online help for a description of each action template.

When you right-click or double-click an action template, a Customize Action window is displayed, enabling you to fill in the particular information to customize that action for your event action plan. Each action template has its own unique Customize Action window.

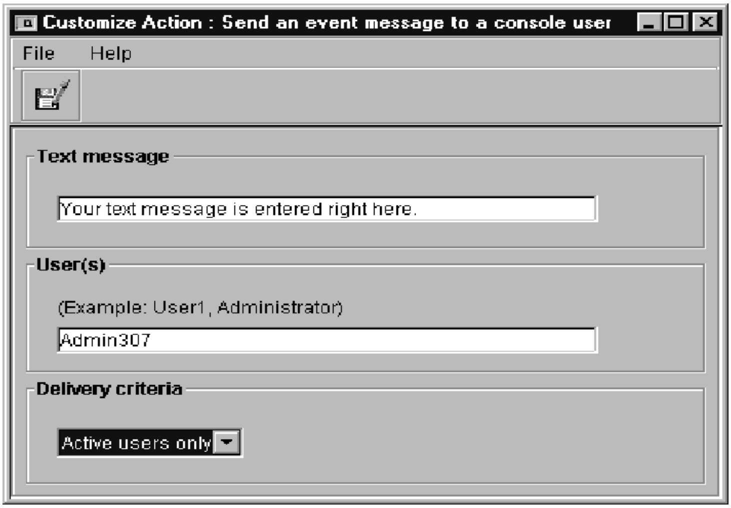

Text message

Your text message is entered right here.

User(s)

(Example: User1, Administrator)

Admin307

Delivery criteria

Active users only

When you save your customized action it appears under the action template in the Event Action Plan Builder window.

Testing an action

You can test the execution of many actions before having them initiated by an event. Right-click the new action and select **Test** from the context menu. Depending on the purpose of the action, you can use the Message Browser window or the Action History window to verify the results of the action. Some targeted actions, such as, Update the Status of the event System, cannot be tested because the input requirement to start the action cannot be met. The Test option is not included in the action list for these types of actions.

Assigning an action to an event filter

You can associate a customized action to an event filter in an event action plan using the normal drag and drop or selection from context menu techniques. See the online help for details.

The action icon appears under the event filter icon in the Event Action Plans pane of the Event Action Plan Builder window. You can continue to add event filters and customized actions to your event action plan as you like.

Saving an event action plan

When you finish building an event action plan, the plan is also added under the Event Action Plans icon in the Tasks pane of the Management Console.

Note that you still have not activated an event action plan or associated it with any managed systems. Refer to “Activating event action plans”for details on applying and activating an event action plan.

Activating event action plans

An event action plan is inactive until you apply it to managed systems. To apply a plan, drag and drop the plan from the Tasks pane in the Director Management Console to the appropriate managed systems.

Displaying applied event action plans

The Event Action Plans association in the Director Management Console enables you to see which event action plans have been applied. After you apply the plan to one or more systems, select the systems for which you want to view the applied plan, right-click in the Group Contents pane to display the context menu, and then select **Associations→ Event Action Plans**. The expansion icon is displayed beside each system in the Group Contents pane to which the plan has been applied.

Performing maintenance tasks

There are several maintenance tasks you can perform from the Event Action Plan Builder window, such as:

- Modifying and deleting event action plans, filters, and actions
- Archiving event action plans for backup
- Importing event action plans from archive
- Exporting event action plans to HTML and XML format for browsing and printing

The tool bar also has refresh and delete icons you can select.

In the three main panes you can perform typical operations such as rename, copy, edit, delete, find, expand, and collapse event action plans, event filters, and actions. You can build new event action plans and event filters, view the action history of an event action, and enable or disable the recording of action history.

See the online help for procedures on performing these operations.

Managing event action plans

In the Director Management Console, you can view which plans have been applied to systems in the network. The Event Action Plans association must be enabled to view applied plans. Right click in open space in the Group Contents plan and select **Associations→ Event Action Plans** from the context menu.

You can also perform the following operations to help manage event action plans:

- You can delete an event action plan that has been applied to a managed system.
- You can initiate a search for a particular system or event action plan.
- You can bring up the Event Action Plan Builder window and use **Expand All…** and **Collapse All…** to view the tree structures and see all the filters and actions associated with each event action plan.

Refer to the online help for more information on these operations.

Viewing event details in the event log

Using the event log, you can view details on all events or subsets of events that have been received and logged by the Director server. The event log is started from the Event Log icon in the Tasks pane of the Director Management Console.

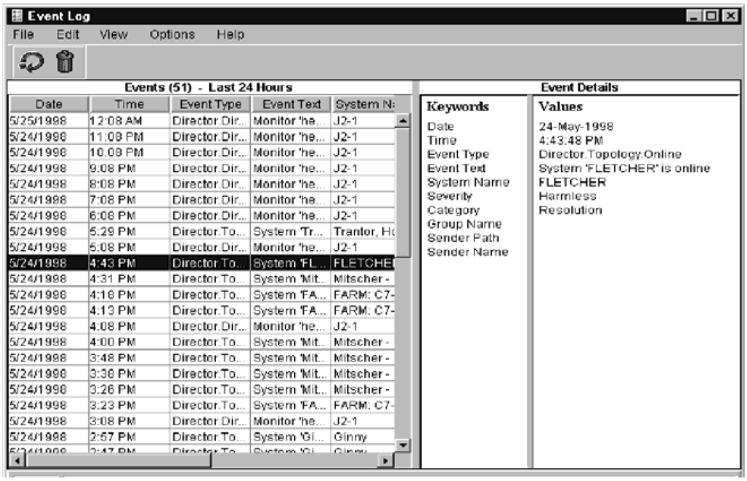

Events (51) - Last 24 Hours

Date	Time	Event Type	Event Text	System N:
5/25/1998	12:08 AM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	11:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	10:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	9:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	8:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	7:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	6:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	5:29 PM	Director.To...	System 'Tr...	Trantor, Hc
5/24/1998	5:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	4:43 PM	Director.To...	System 'FL...	FLETCHEI
5/24/1998	4:31 PM	Director.To...	System 'Mit...	Mitscher -
5/24/1998	4:18 PM	Director.To...	System 'FA...	FARM: C7-
5/24/1998	4:13 PM	Director.To...	System 'FA...	FARM: C7-
5/24/1998	4:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	4:00 PM	Director.To...	System 'Mit...	Mitscher -
5/24/1998	3:48 PM	Director.To...	System 'Mit...	Mitscher -
5/24/1998	3:38 PM	Director.To...	System 'Mit...	Mitscher -
5/24/1998	3:26 PM	Director.To...	System 'Mit...	Mitscher -
5/24/1998	3:23 PM	Director.To...	System 'FA...	FARM: C7-
5/24/1998	3:08 PM	Director.Dir...	Monitor 'he...	J2-1
5/24/1998	2:57 PM	Director.To...	System 'Gi...	Ginny
5/24/1998	2:47 PM	Director.To...	System 'Gi...	Ginny

Event Details

Keywords	Values
Date	24-May-1998
Time	4:43:48 PM
Event Type	Director.Topology.Online
Event Text	System 'FLETCHER' is online
System Name	FLETCHER
Severity	Harmless
Category	Resolution
Group Name	
Sender Path	
Sender Name	

Each entry in the event log is subdivided into fields containing the filter criteria associated with the event. See the online help for details on these fields.

Viewing all logged events

By default, the Add event to the event log action is coupled to the last 100 events received by the Director management server in the last 24 hours. 100 events and 24 hours are defaults you can change using **Options→ Set Log View Count** and **Set Time-Range**. The maximum number of entries that can be presented in the event log viewer is 20000; however, the log holds up to 100000 entries. When you start the Event Log without specifying a filter or managed system, all events are displayed.

Viewing events by filter characteristics

You can use the predefined filters or your user-defined filters to refine the events included in the log to only those that meet the filtering criteria. Double-click on the desired event filter icon under the Event Log icon.

Viewing events by system

To view a filtered list of events from a single managed system, drag the icon onto the desired event filter icon (or drag the filter icon onto the target system icon).

Using the Action History window

The Action History window enables you to view the history of event actions that have been initiated. To activate the action history, right-click on a customized action and select **Action History→ Enable** from the context menu. To view the history, right-click on the customized action after enabling the action history and select the **Action History →Show** option. The Action History window is displayed. It contains two main panes, Actions and Action Details. The Actions pane contains a table of every execution of the customized action which occurred during a given time range. Each row represents one execution of the customized action. The Action Details pane contains two sub-panes, Keywords and Values, which show the details of a selected occurrence of the action.

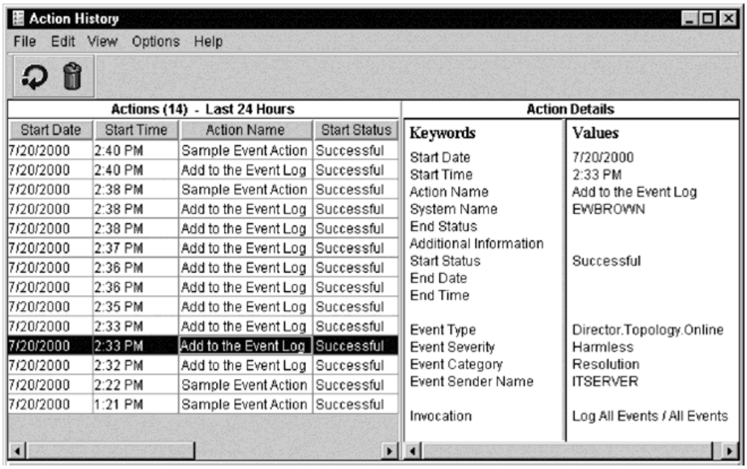

Actions (14) - Last 24 Hours

Start Date	Start Time	Action Name	Start Status
7/20/2000	2:40 PM	Sample Event Action	Successful
7/20/2000	2:40 PM	Add to the Event Log	Successful
7/20/2000	2:38 PM	Sample Event Action	Successful
7/20/2000	2:38 PM	Add to the Event Log	Successful
7/20/2000	2:38 PM	Add to the Event Log	Successful
7/20/2000	2:37 PM	Add to the Event Log	Successful
7/20/2000	2:36 PM	Add to the Event Log	Successful
7/20/2000	2:36 PM	Add to the Event Log	Successful
7/20/2000	2:35 PM	Add to the Event Log	Successful
7/20/2000	2:33 PM	Add to the Event Log	Successful
7/20/2000	2:33 PM	Add to the Event Log	Successful
7/20/2000	2:32 PM	Add to the Event Log	Successful
7/20/2000	2:22 PM	Sample Event Action	Successful
7/20/2000	1:21 PM	Sample Event Action	Successful

Action Details

Keywords	Values
Start Date	7/20/2000
Start Time	2:33 PM
Action Name	Add to the Event Log
System Name	EWBROWN
End Status	
Additional Information	
Start Status	Successful
End Date	
End Time	
Event Type	Director.Topology.Online
Event Severity	Harmless
Event Category	Resolution
Event Sender Name	ITSERVER
Invocation	Log All Events / All Events

You can perform the following operations:

- Select any row in the Actions pane and the details of that action are shown in the Action Details pane. See the online help for more information on these action details.
- Use the menu bar option **Set Time Range** to define the time range, in hours, for which you want actions displayed, and **Set History Count** to specify the maximum number of action entries to display.
- Using the menu bar, tool bar and context menu options, you can select one or more entries and delete them from the display, or refresh the view, perform a search for a particular entry, and sort the entries in ascending or descending order.
- You can hide and show columns, adjust sizes of columns and panes, and re-order columns using standard techniques described in “Navigating in Director” on page 80.

Generating your own events

The Director genevent utility enables you to generate events. By default, user-defined events are directed to the server or servers known to be managing the agent from which the event is sent. The genevent must be used from a command prompt on the Director server or a managed system; it is not available through the Director Management Console.

Use the following syntax to run genevent from a command prompt.

From the operating system command line, specify the following:

genevent/*required_parameters* /*optional_parameters*

You must specify the following *required_parameters*:

type:*type*
: where *type* is a dot-delimited string in the same format used to indicate event type, for example, Director.Topology.Online. Refer to the online help for details on keyword information and usage.

text:*text*
: where *text* is a descriptive string you supply to identify the cause of the event.

You can also specify the following *optional_parameters:*

sev:*severity*
: where *severity* indicates the urgency of this event. Specify one of the following:

- fatal
- critical
- minor
- warning
- harmless

If unspecified, *severity* defaults to unknown. These categories are described in the online help.

dest:@EventServer
: @EventServer (Default) designates that the event should be directed to the server or servers known to be managing this agent.

dest:*protocol*::*name*
: where *protocol* is the transport used between this managed system and the Director server to which this event will be sent and *name* is the name of the targeted Director management server used by the specified protocol, for example, NETBIOS::TWGSVR1. Valid values for *protocol* are: netbios, tcpip, and ipx.

The default destination is @EventServer. @EventServer designates that the event should be directed to the server or servers known to be managing this agent.

Chapter 10. Software Distribution

The Software Distribution tasks enables you to distribute IBM created file packages and install them on native systems on your network. For example, IBM Director Agent, which is included in the service pack on the CD, or downloaded from the web, is a IBM file package.

This chapter describes how to import and distribute IBM created file packages using the software distribution task. Before you attempt to distribute a package, make sure you read the guidelines for software distribution in "Planning for Director tasks" on page 30.

Importing a file package

To import a file package that has been exported, you must use the Director File Package wizard. When you import a file package using this wizard, you are prompted to specify the location of the package.

Distributing a file package

To perform a distribution of a software distribution file package, drag the file package icon from the Tasks pane of the Director Management Console and drop it on the desired system icon or group of systems. Only Director managed systems are valid targets for software distribution file packages. Refer to Chapter 32, "Troubleshooting," on page 371 for more help on distributing software distribution packages.

Scheduled distributions

When you initiate a software distribution task, you have the option to perform the task immediately or schedule it to occur at some later time. You can set up a software distribution to take place after business hours, for example, or when network traffic is lighter.

Refer to "Starting the Task Scheduler task" on page 201 for more information about scheduling software distribution packages.

Immediate distributions

When you perform an immediate software distribution, the Immediate Distributionwindow opens.

The status information in the top pane gives you a summary of the distribution status of the various systems you have targeted. The bottom pane lists the various systems under the different status levels.

If you require more detail concerning a distribution, you can display a log that contains additional information. Select **File→ View Log** to display the log. Using the selections on the menu bar, you can copy the log to a clipboard, refresh the log, request dynamic updating of the log, set the detail level of the log, and close the log. If you only need to view a log for a specific system, select the system and then select **Selected→ View System Log**, or double-click on the system icon.

Distributing file packages on systems running Linux

To perform a distribution of a software file package on systems running Linux, use the following procedure:

1. Import the Linux file package.
2. Drag-and-drop the file package onto the desired system or systems. This action will only transfer the file package, it will not install the package.
3. Create a Process Task. See Chapter 21, "Process management," on page 195.

   Note: The command to enter in your Process Task should be:
   Userid and Password with root permissions, and

   ```
   cd  /tmp;./dirinstall
   ```
4. Expand the Process Management task and double-click Process Tasks.
5. From the command line, enter the name of the script file as it resides on the Linux system. For example, /opt/tivoliwg/SwPkInst/upgradeITD. Do not click Log.
6. Save the task.
7. Drag the new task to the Linux system and Execute immediately.

Viewing package content information

The Package Summary window enables you to view the contents of a package, including the package files, the operating system platform for which the package was created, and whether the target system is to be rebooted after package installation. To access the window, in the Tasks pane of the Management Console, expand **Software Distribution** to view the list of software distribution packages. Right-click on a package, then select **Package Information** in the context menu.

Viewing distribution history

To view the distribution history for a selected software distribution package, in the Tasks pane of the Management Console, expand **Software Distribution** to view the list of software distribution packages. Right-click on a package, then select **Distribution History** in the context menu.

Hover help gives you the date/time stamp of the last distribution.

You can use Associations to view distribution history on a system-by-system basis. To view the last distribution status, double-click on the package under the Associations tree.

For more information, refer to the online help.

Renaming packages

To rename a software distribution package, in the Tasks pane of the Director Management Console, expand **Software Distribution** to view the list of software distribution packages. Right-click on a package, then select **Rename** in the context menu.

Viewing package audit activity

The Package Audit Log enables you to determine the status of software distribution package creation and distribution. Three levels of detail are provided to assist you in tracking and troubleshooting. You can also cut and paste entries into other files for printing. To access the log, in the Management Console, right-click on **Software Distribution** in the Tasks pane, then select **Package Audit Log** in the context menu. Refer to the online help for more information.

Deleting a file package

To delete a file package, right-click on the file package icon and then select **Delete** from the context menu.

If you receive a message indicating that the package is locked by another process, this usually means that it is being copied to a file distribution server. The package remains locked until the other process completes. It is possible for a package to remain locked when no process or user is using it. In these cases the package should become available again in approximately five to ten minutes.

Using File Distribution Servers Manager

File Distribution Servers Manager enables you to view details on file distribution servers and software packages. You can:

- View the file distribution server maintenance log
- Check access to file distribution servers
- Refresh packages from the server share
- Delete packages from the server share

To access the window, in the Tasks pane of the Director Management Console, right-click on **Software Distribution**, and then select **File Distribution Servers Manager** in the context menu. For more information, refer to the online help.

Editing the Package Configurator

Director installs a batch file named SDPConfig.bat to the Director\bin directory on servers. This batch file launches the Edit Software Distribution Package window. All imported IBM Director Win32 packages are listed in the Edit Software Distribution Package window. You are notified if no packages are imported.

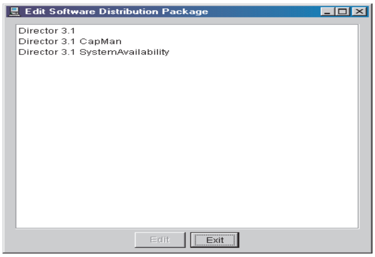
Edit Software Distribution Package

Director 3.1
Director 3.1 CapMan
Director 3.1 SystemAvailability

To edit the install options of a particular package, select the package from the list and click **Edit**. The Modify Install Parameters window opens with two tabbed panes. The current install options for the selected packages are preselected. However, you may modify the install options for the IBM Director 3.1 and for IBM Director Extensions 3.1 by selecting the options on the tabbed panes. The various install options available for both the IBM Director 3.1 and IBM Director Extensions 3.1 are described in the following tables. After editing the install options, click **SAVE** to save and exit the options panel or click Exit to exit without saving.

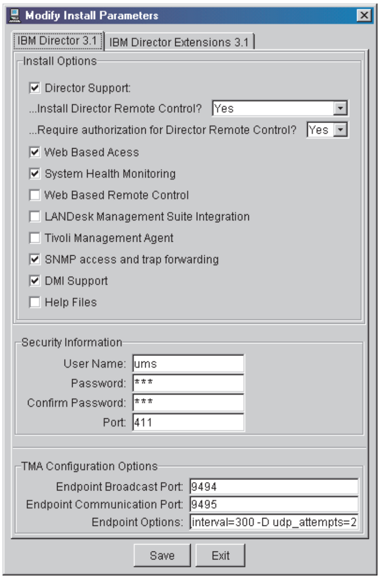
Modify Install Parameters

Install Options

☑ Director Support:

...Install Director Remote Control? Yes

...Require authorization for Director Remote Control? Yes

☑ Web Based Acess

☑ System Health Monitoring

□ Web Based Remote Control

□ LANDesk Management Suite Integration

□ Tivoli Management Agent

☑ SNMP access and trap forwarding

☑ DMI Support

□ Help Files

Security Information

User Name: ums

Password: ***

Confirm Password: ***

Port: 411

TMA Configuration Options

Endpoint Broadcast Port: 9494

Endpoint Communication Port: 9495

Endpoint Options: interval=300 -D udp_attempts=2

The following table lists the Modify Install Parameters for IBM Director.

Item	Description
Install Options	
Director Support	Director Support is an additional configuration option for the agent installation only. Director is an advanced Intel-processor-based workgroup hardware manager, with centralized agent and group management console and server services. Selecting this feature enables the agent system to be managed in a Director environment by installing IBM Director Agent on this system.
Install Director Remote Control	Remote Control enables you to manage a remote system by displaying the desktop of a remote managed system within a Director Management Console and by sending keyboard and mouse information to the remote managed system. You can also view a listing of all the consoles that have remote sessions with the managed system, and see the controlling state of each.
Require authorization for Director Remote Control	The Require User Authorization for Director Remote Control option enables you to specify whether a remote user can access, and take over control of, the local system without the permission of the local user. If this option is enabled and a Director administrator attempts to use remote control to access the local system, a message window is displayed on the local system indicating that a remote user is attempting remote control access. You can then allow or disallow access.

Item	Description
Web Based Access	Web Based Access offers a convenient Java-based tool for managing a agent system and for viewing the CIM-based inventory data. If you install Web Based Access, a hypertext transport protocol (HTTP) daemon is installed and requires that a user name and password be entered during the installation. The user name and password are used to limit access to the HTTP daemon. With Web-based Access installed on the agent system, the agent can be managed from any remote computer with a supported Web browser. No software other than a Web browser is needed on the remote system.
System Health Monitoring	System Health Monitoring provides active monitoring of critical system functions, such as disk space available, system temperature, fan functionality, power supply voltage, and system cover removal. System Health Monitoring enables you to detect system problems early, before system failures occur. System administrators are notified of a system problem by a CIM event, SNMP trap (SNMP traps are available only if SNMP access and trap forwarding is also selected), or SMS Status Message (Microsoft SMS 2.0 only). Critical problems also result in a pop-up message appearing on the display of the agent and a status change in System Health GUI.
Web Based Remote Control	Web Based Remote Control enables a remote systems administrator using a Web browser or MMC console to take control of the agent system desktop, enhancing the administrator’s ability to diagnose system problems and troubleshoot the system.
LANDesk Management Suite Integration	LANDesk Management Suite Integration installs the Intel Common Base Agent on the agent system. This enables the system administrator to use IBM Director Agent with LANDesk Management Suite.
Tivoli Management Agent	Tivoli Management Agent installs support on the agent system that enables it to be managed by the Tivoli Enterprise system-management platform.

Item	Description
SNMP access and trap forwarding	This feature enables CIM information to be accessed from a system using the Simple Network Management Protocol (SNMP). If System Health Monitoring is enabled, this option also enables System Health to forward CIM events as SNMP traps. This component requires that you have the SNMP service (provided with the operating system) installed on the endpoint. If the SNMP service is not installed, the system prompts you to insert the operating system installation media and install SNMP during the IBM Director Agent installation.
DMI Support	DMI Support is an vendor-neutral interface for collecting and manipulating network management information.
Help Files	This component installs online documentation. Do not select this option if you are concerned about disk space or do not need online documentation installed on every agent.
Security Information	
User Name	A unique character string that identifies the user (32 characters maximum).
Password	The user password (32 characters maximum, case sensitive). There are no restrictions on the characters that can be used in passwords.
Confirm Password	his field must contain the same character string as the **Password** field (32 characters maximum, case sensitive).
Port	
TMA Configuration Options	
Endpoint Broadcast Port	The key `szEdit1` specifies the port through which the gateway communicates. It is 9494 by default.
Endpoint Communication Port	The key szEdit2 specifies the port through which the endpoint communicates. It is 9495 by default.
Endpoint Options	See "Deploying endpoints with IBM Director Agent "on page 510.

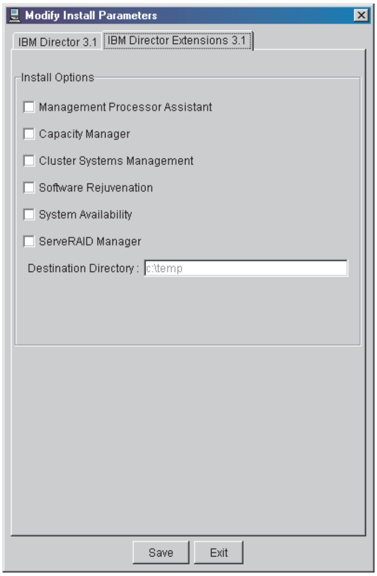

□ Management Processor Assistant

□ Capacity Manager

□ Cluster Systems Management

□ Software Rejuvenation

□ System Availability

□ ServeRAID Manager

Destination Directory : c:\temp

The following table list the Modify Parameters for IBM Director Extensions.

Item	Description
Management Processor Assistant	The Management Processor Assistant service is used 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.
Capacity Manager	Capacity Manager provides proactive management of hardware resources by gathering and presenting historic data or formulating trend analyses that identify and predict system performance bottlenecks.
Cluster Systems Management	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.

Item	Description
Software Rejuvenation	The Software Rejuvenation tool is used to reduce the number and impact of unplanned outages due to software aging. The result is an increase in the reliability of managed systems. This is achieved through scheduled software rejuvenations (restarts) on each system. You can implement software rejuvenation in either of two ways: manually or automatically.
System Availability	System Availability is used to analyze the availability of a system or a group of systems. It can be used to provide statistics on the availability of large sets of systems. Additionally it can be used as a means to graphically prove that Software Rejuvenation improves system availability.
System Health Monitoring	System Health Monitoring provides active monitoring of critical system functions, such as disk space available, system temperature, fan functionality, power supply voltage, and system cover removal. System Health Monitoring enables you to detect system problems early, before system failures occur. System administrators are notified of a system problem by a CIM event, SNMP trap (SNMP traps are available only if SNMP access and trap forwarding is also selected), or SMS Status Message (Microsoft SMS 2.0 only). Critical problems also result in a pop-up message appearing on the display of the agent and a status change in System Health GUI.
Web Based Remote Control	Web Based Remote Control enables a remote systems administrator using a Web browser or MMC console to take control of the agent system desktop, enhancing the administrator's ability to diagnose system problems and troubleshoot the system.
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 server.
Destination Directory	The default destination directory has been set to c:\temp. Currently, the user cannot modify this option.

Chapter 11. File Transfer

The File Transfer task enables you to transfer files from multiple locations, delete files, create directories, view file properties, edit the contents of a file, and synchronize files, directories, or drives.

You can transfer and receive individual files and directories between:

- The Director Management Console local system and the Director server system
- The Director Management Console local system and a native managed system
- The Director server system and a native managed system

File transfer between two managed systems is not supported directly. However, it is possible to receive a file from one managed system to a Director Management Console or Director Server, and then send that file to a different managed system.

File transfer is a one-to-one interactive task that provides a tool for troubleshooting and repairing a problem system. The purpose of file transfer is *not* to perform software distribution.Use it to send and receive small numbers of files to solve isolated problems in your network, or to help configure a particular system. You cannot schedule a file transfer to occur at a later time, because it is an interactive task.

Starting a File Transfer session

Open the File Transfer window from the Director Management Console by double-clicking the task or by using normal drag-and-drop techniques. Refer to "Navigating in Director" on page 80 for tips on navigating your way through this task, or see the online help for detailed assistance.

Director takes a few seconds to query the files on your local system and on the target system, and then displays the File Transfer console.

This window has a Source File System pane and a Target File System pane. The root directories for your local system or the Director server appear in a tree structure in the Source File System pane, and the root directories for the selected managed system or server appear in the tree structure in the Target File Systems pane.

Just under the File System title near the top of the Source pane there is a system pull-down menu where you can select between your local system and the Director server. If you started the file transfer by a drag-and-drop operation to a specific system, the system pull-down menu in the Target pane displays the file system of the managed system. If you opened the task without specifying a system, the Target pane displays the file system of the Director server.

The wild card feature

The file transfer task allows for multiple files to be transferred that, while may not have the same filename, has the same file extension (such as,.txt,.pdf,.dll) or, same filename and different extension.

The File Transfer console automatically enables the wildcard feature. In the Filename: field, the search opens with *.*. All files within a selected drive and expanded folder are revealed. Use the wildcard feature to transfer like files to the target system.

Selecting files to transfer

Select any of the drive icons in the File System pane on either side. The contents of that drive expand and appear in the pane, showing subdirectories and files. You can continue to expand and collapse additional subdirectories to go further down the tree structure.

You can transfer files or entire subdirectories using any of the following methods:

Drag-and-drop operation

1. Drag a file or subdirectory icon from one file system pane to the other file system pane.
2. Drop the icon on the destination subdirectory or drive.

Transfer file(s) to target

1. Select a file or subdirectory in the source pane (local system or Director server).
2. Select the drive or subdirectory in the target pane.
3. From the menu bar, select **Actions →Source → Transfer File(s) to target** to transfer a file or subdirectory from the local system or the Director management server to the target drive or subdirectory.

Transfer file(s) to source

1. Select a drive or subdirectory in the source pane (the remote system or server).
2. Select a file or subdirectory in the target pane.
3. From the menu bar, select **Actions → Target→Transfer File(s) to Source** to transfer the file or subdirectory from the target pane to the local system or the Director server.

You can select multiple files for transfer by pressing and holding the **Shift** key while clicking on the desired files with the mouse. As you select the last file in the group to be transferred, do not release the mouse button. Release only the Shift key, and while still holding the mouse button down, drag the cursor to the target File System pane.

Transferring files between managed systems

To transfer files from one managed system to another, you must first transfer the files from one managed system to your local system or the Director server, and then transfer the files from the local system or server to the desired target managed system.

After you transfer the file from the originating system to your local system or server, you will see the file or subdirectory refreshed to contain the transferred files. Now you can drag it or transfer it to the target managed system as usual.

Choosing a new target

To dynamically select a new target (agent) from within the File Transfer window, click **Other** beside the target drop-down list. The Choose Target window is displayed, listing all available systems that support file transfer. Select the system you want to transfer files to or from and click **OK**. The system is now selected for file transfers and is added to the target list. You can now transfer files to and from the selected system.

Note: Only six systems can be added to the drop-down list at one time. If you add more than six, the system added earliest is removed from the list.

Synchronizing files, directories, or drives

Synchronizing means making file contents, directory contents, or the contents of an entire drive identical across multiple managed systems. Synchronizing provides a simpler method for ensuring the consistency of files that reside on multiple systems.

Synchronizing involves only the target system and the source system. You can synchronize files, directories, and drives on as many systems as necessary, but you must synchronize them individually. You cannot synchronize multiple systems from a source system at the same time.

To synchronize files, directories, or drives, use the following procedure:

1. Select a source object as explained in “Starting a File Transfer session” on page 137.
2. Select a target object.

   Note: If you want to make the target directory identical to the source directory, select **Target → Synchronize from Source**. If you want to make the source directory identical to the target directory, select **Source → Synchronize from Target**.

3. You may receive a message stating that the selected directory names are different. Select **Yes** to continue.
4. You will receive a message stating that this action may delete some files and directories. Select **Yes** to continue.
5. The selected directories are now synchronized.

   Notes:

   a. When you synchronize a file, directory, or drive, its contents are deleted. Then the drive or directory from which you are synchronizing is copied to replace the original.

   b. Only similar objects (files, directories, or drives) can be synchronized. That is, a file can only be synchronized with another file, a directory with another directory, and so on.

Additional file transfer features

The file transfer task is not intended to be a full-function file manager, but you do have some limited capabilities, such as making new directories, deleting files and directories, renaming files, viewing file properties, and editing simple text files. Refer to the online help for details.

Precautions when using file transfer

There are a few precautions you should keep in mind when performing file transfers:

- You cannot use a file as the target of a transfer.
- If the network drives on the Director server or managed system are mapped using a different user name or password than the user name/password specified for the Director service during installation (that is, the user name/password of the Director support service), the network drives will be unavailable due to access limitations.
- The File Transfer task can only be applied to a single managed system at a time.
- You cannot transfer the entire contents of a drive by dragging the icon of the drive. You can only transfer files and directories using drag-and-drop operations.
- The contents of each subdirectory are discovered as the subdirectory is expanded when you click on it in the File Systems pane. The discovery process can be especially slow when using the Details view on a remote server or managed system.

- If you transfer a file which is the same name as an existing file on the destination system, the file is overwritten.
- If your file transfer session with the remote system is broken while performing a file transfer, you must re-establish the session and transfer the files again.
- If you select multiple files for a file transfer using a drag-and-drop action, be sure to hold the mouse button down as you select the files. Do not release the mouse button until you move the mouse to the destination. If you release the mouse button too soon, only the last file selected will be transferred.
- If you highlight multiple systems in the Group Contents pane of the Director Management Console, and then attempt to drag the File Transfer Console icon to one of the systems, an error message will be displayed.
- If you hold down the Shift key to highlight multiple systems in the Director Management Console and, while holding down the Shift key, right-click one of the managed systems highlighted, the File Transfer task will not appear in the list of available tasks in the context menu. File transfers can only be set up with a single managed system at a time.
- If the target managed system is a NetWare system and has DOS drives (A:\, B:\, C:\, and so on), these volumes are not displayed in the File System pane.

Chapter 12. DMI management

Director provides Desktop Management Interface (DMI) support for the browser, inventory, resource monitoring, and event management tasks. The DMI is an vendor-neutral interface for collecting and manipulating network management information. The Desktop Management Task Force, Inc. (DMTF) develops and maintains DMI specifications.

This chapter describes how to use the DMI Browser to isolate DMI components and view and change attribute values. For information on isolating DMI data for the inventory, resource monitoring, and event management tasks, see the following chapters:

- Chapter 6, “Inventory Management,” on page 91.
- Chapter 8, “Resource Monitoring,” on page 109.
- Chapter 9, “Event management,” on page 117.

The Director Management Console does not automatically display DMI-enabled systems as a separate group of systems. To create a dynamic group of DMI-enabled systems, follow the procedure described in “Creating a DMI dynamic group”. You can also use the Static Group Editor to create a group of one or more systems that are DMI-enabled in your network.

DMI requirements

To provide DMI data, managed systems must be running under Windows Me, Windows 98, Windows 2000, or Windows NT 4.0. They must have a DMI Service Provider (version 2.0 or later) installed. The DMI Service layer is installed with the CIM to DMI Mapper. The Service Layer can be added to a managed system after Director is installed. When the managed system is restarted, it is enabled for DMI operations.

Creating a DMI dynamic group

You can use the Task Based Group Editor to create new dynamic group filters based on combinations of tasks that apply to managed systems. This procedure assumes you want to create a filter that isolates systems that are DMI-enabled.

To create a dynamic group for DMI-enabled systems, use the following procedure:

1. Right-click the **Groups** pane of the Director Management Console to display the context menu.
2. Select **New Task Based** to display the Task Group Editor.

3. In the Available Resources pane, select **DMI Browser** and click **Add** to add the selection to the Selected Criteria pane. Selecting DMI Browser creates a filtering criteria for managed systems that are DMI-enabled.
4. Click **Save As** to save the new group with a name of your choosing.
5. In the dialog that appears, enter a descriptive name for the group, for example, "DMI-enabled systems."
6. Select **Close Group Editor** to save your group and exit the dialog.
7. Refresh the Director Management Console with a discovery operation, and the new group appears in the Groups pane.
8. Select your new group to see which managed systems match the DMI criteria. DMI-enabled systems, if discovered, are listed in the Group Contents pane.

Performing DMI Browser tasks

The DMI Browser enables you to perform the following tasks:

- View the DMI components and groups for a selected DMI-enabled system
- View attribute values for selected group classes
- Set values for individual attributes
- Define browser subtasks for specific group classes.

When you apply the DMI Browser to a managed system, the information is gathered directly from the target system and displayed. If you change an attribute value, Director issues a request to the Service Layer on the target system to update the specified attribute's value.

Starting the DMI Browser and viewing information

To start the DMI Browser and view information for a single managed system, use the following procedure:

1. Select the managed system for which you want to view information and drag it to DMI Browser in the Tasks pane. The DMI Browser window appears. The systems you selected are displayed as a tree (hierarchical) view in the DMI Components pane.
    - If a system is not configured for DMI, a message appears. It indicates that the target system does not support the task.
    - If the system is inaccessible, for example, if it is offline, the DMI Browser window is opened, but the DMI tree for the system cannot be expanded.
    - To open the browser for two or more systems, select the managed systems for which you want to view information and drag the **DMI Browser** from the Tasks pane to any system in the set of systems highlighted. The DMI Components pane displays the systems selected.
    - If one or more of the systems is not configured for DMI, a message indicates that at least one of the target systems does not support the task.

- If one or more of the systems is inaccessible, the DMI Browser window is opened, but one or more of the systems is shown as unavailable, and its DMI tree cannot be expanded.

2. Double-click a system to display the components of the system, and then click a component to display descriptive information in the right pane.
3. To view the group classes of a component, double-click the component name.
4. To view the attributes of a group class, click the group class name. A description of the group class appears in the upper-right pane labeled **Groups** and the associated attributes and methods are displayed in the lower right-hand pane.
5. To reverse the order of the properties, right an item and select **Sort→Descending**.
6. When you have finished viewing information, click **File → Close** to close the window.

Setting an attribute value for a DMI group

It is strongly recommended that you do not change an attribute value unless you are thoroughly familiar with the structure and manipulation of DMI data. Improperly setting a system value can cause unpredictable results on the target system.

To change an attribute value, use the following procedure:

1. Navigate to the attribute for which you want to change a value using the procedure described in “Starting the DMI Browser and viewing information” on page 144.
2. Right-click on the attribute row and select **Set Value** in the context menu. The Set Value dialog appears with the current value.
3. Enter the new value and select **OK** to enact the change. If you do not want Director to attempt to change the value, cancel the window.

   If Director is unable to change the value on the target system, a message indicates the failure.

Defining DMI Browser subtasks

A user-defined subtask is a fast path to a specific DMI group class. After it is defined, a browser subtask is applied directly to a managed system to view only information associated with the specified group class.

To define a browser task, use the following procedure:

1. In the Director Management Console, apply the DMI Browser task to a managed system to display the DMI Browser window.
2. Double-click the managed system to display the associated components.
3. Double-click a component to display the contained group classes.

4. Right-click the group class name to display the context menu, and click **Create task for group class**. A window appears and uses the name of the group class as the default name.
5. You can enter a new name or keep the default name. To keep the default name, click **OK**. The new task is entered as a subtask under DMI Browser in the Director Management Console.
6. Apply the browser subtask to a DMI-enabled managed system that has the same group class registered with its DMI service layer and view the associated data.

   Notes:

   a. If you create a subtask for a group class and then apply it to a system with two or more DMI components containing the same group class, separately tabbed panels are displayed for each component containing the group class. For example, if you create a subtask for the Component ID group class and then apply the subtask to a system with two or more DMI component IDs, separately tabbed pages are displayed for each component ID that is defined.

   b. The error message `The targeted system does not support this class` appears if a user-defined subtask for a group class is applied to a system that does not have registered components containing the group class.

Chapter 13. Configuring SNMP Agent

The SNMP Configuration tool is used to assign or reassign simple network management protocol (SNMP) trap destinations. This tool works in Windows 98, Windows NT, and Windows 2000 environments.

Director Agent forwards the SNMP traps only if the SNMP Agent is installed on both the Server and the Agent workstation. You can install and configure the SNMP Agent as part of the Director Agent installation process or afterwards by using the Windows 98 CD, Windows NT or Windows 2000 CD.   However, if you install the SNMP Agent for Windows 98 after Director Agent has been installed, you must use the SNMP Trap Configuration tool to add the trap destination to the registry.

You can mass configure large numbers of managed systems by building Profiles of the systems from a single location and using the mass configuration option.

Starting the configure SNMP agent task

From the Director Management Console, drag and drop the Configure SNMP Agent task icon on to a managed system in the Group contents pane. The SNMP window opens.

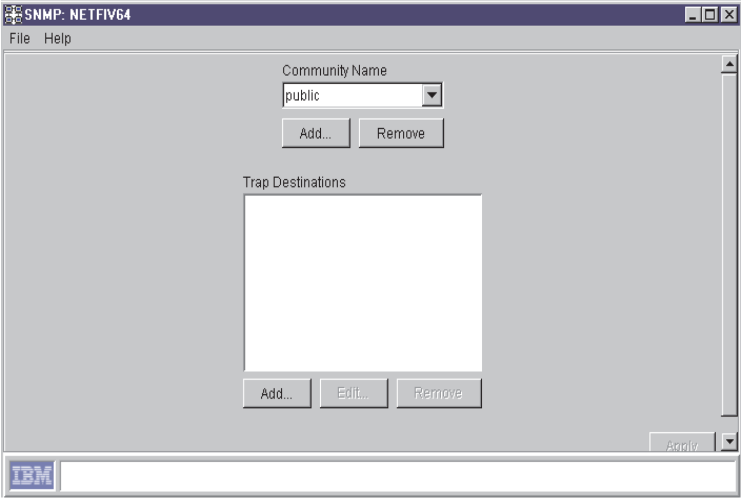

The following table describes the fields in the SNMP window.

Item	Description
Community Name	This drop-down list contains names assigned to one or more SNMP trap destinations. These names are used in the configuration. Multiple community names may be defined. A community name that has been configured may be selected from the list.
Add	Use this button to insert a new Community Name.
Remove	Use this button to remove an existing Community Name.
Trap Destination	The IP address or host name of the computer that will receive the SNMP traps. Multiple IP addresses and host names may be assigned.
Add	Use this button to insert a new SNMP Trap Destination.
Edit	Use this button to change an existing SNMP Trap Destination.
Remove	Use this button to remove a SNMP Trap Destination.
Apply	Use this button to activate any additions, edits, or removals you make. If you choose not to activate your information, exit this screen without pressing Apply.

For more information on adding, changing or removing a Community Name or a Trap Destination, please see online Help.

Using the Profile Builder

In the Director Management Console, right-click Configure SNMP Agent in the Task pane and select **Profile Builder**. The Profile window opens.

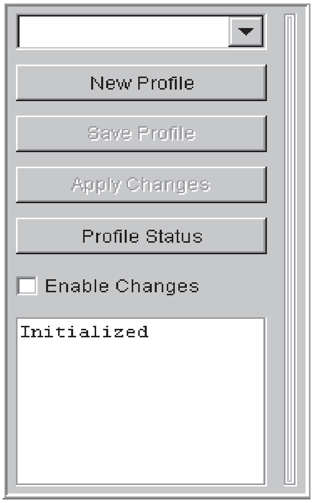

```
Initialized
```

From this window, you can create a profile.

Creating a profile

There are several way to configure a profile. You can create a new profile or modify an existing profile.

Creating a new profile

To create a new profile use the following procedure:

1. From the Profile window, click **New Profile**. The Input window opens.
2. Enter a New Profile name.
3. Click **OK**. The SNMP window opens.

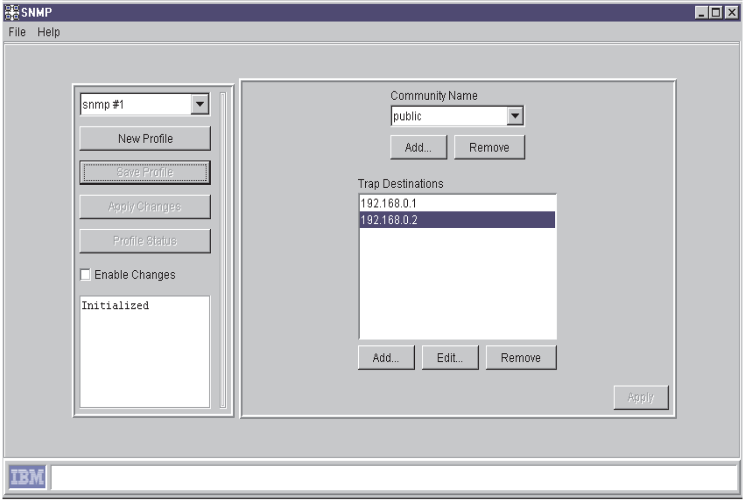

snmp #1

□ Enable Changes

Initialized

Community Name
public

Trap Destinations
192.168.0.1
192.168.0.2

4. Enter the SNMPdata.
5. Click **Save Profile**. The Save Profile window opens.
6. Click **Yes** to save the profile or click **No**.

Modifying an existing profile

From the Profile window, click the drop-down list to display all profiles. Select an existing profile and click **Profile Manager**. The Status window opens. Select **Profile**. Modify the existing profile by editing any data in the SNMP interface. . After the modified profile is created, click the **Save** button.

Using the Status window

The Status window displays a list of profiles, and groups assigned to those profiles. It also displays the status of a profile. To start the Status window, from the Director Management console, right-click any saved profile. A context menu will appear. Click **Profile Manager** to start the manager.

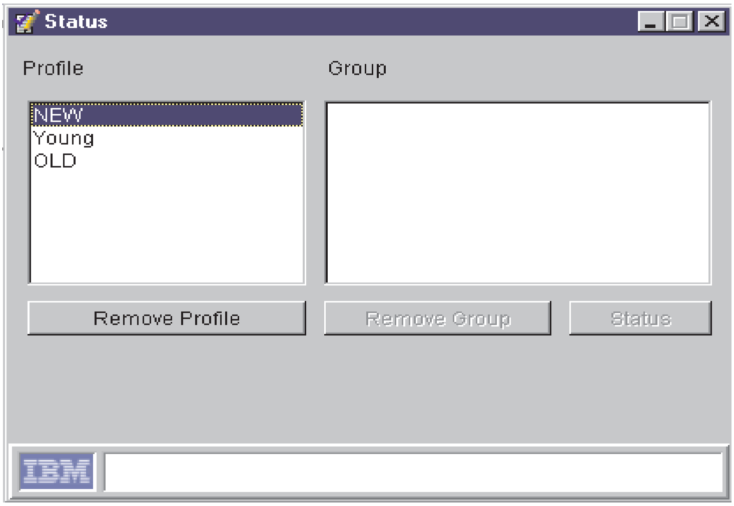

After a profile has been created, one or more Director groups might be associated with the profile. Profile Manager enables you to determine which groups are associated with any given profile.

When the Profile Manager is started all profiles are listed in the Status window. The first profile in the list is highlighted. All groups associated with a selected profile appear under Groups in the Status window. The Remove Group button and the Status button are disabled until a group is selected.

If a profile is not currently associated with any groups, you can select it and click **Remove Profile** to remove it. If you try to remove a profile associated with another group, a window will open stating that the operation is unsuccessful because a group is a member of another profile.

Viewing the status of a profile

To view the status of a group, select a profile and a group; click the **Profile Status** button.

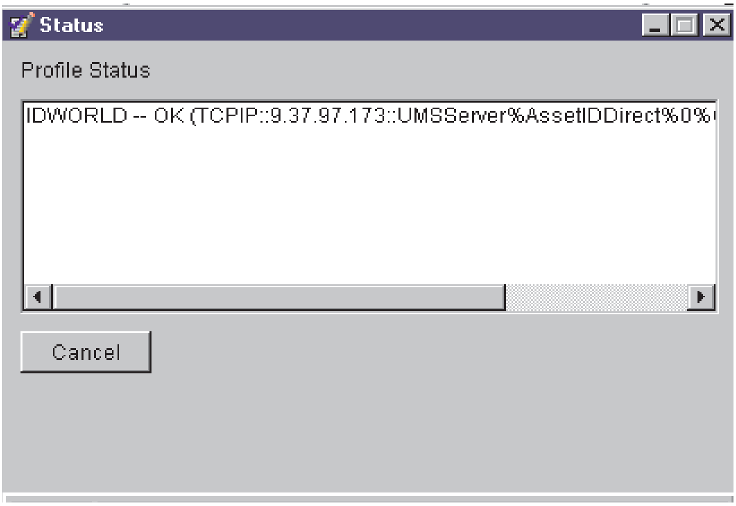

IDWORLD -- OK (TCPIP::9.37.97.173::UMSServer%AssetIDDirect%0%

The dialog box displays a list of systems within a group and the current status of the managed systems. A managed system will display one of the following status indicators:

- OK: The system has been configured correctly.
- Failed: Mass configuration was able to communicate with the system but unable to configure it correctly, or mass configuration has exhausted all attempts to communicate with the target system.
- Pending: The target system is currently offline and mass configuration will attempt to configure it at a later time, or mass configuration has scheduled the configuration attempt for a later time.

Using Mass Configuration

Before you can use the Mass Configuration task, a profile must be created using the Profile Builder. For more see “Using the Profile Builder” on page 148. To start the mass configuration task, double-click the Configure SNMP Agent icon to expand the Profiles. Drag-and-drop a Profile onto a managed system that you wish to mass configure.

Chapter 14. SNMP Management

Director includes Simple Network Management Protocol (SNMP) support that enables you to isolate SNMP devices for the event management, inventory, and resource monitor services and mass configure the settings. For information on using the Director Management Console to specify SNMP devices for these tasks, see the following chapters:

- Chapter 6, “Inventory Management,” on page 91.
- Chapter 8, “Resource Monitoring,” on page 109.
- Chapter 9, “Event management,” on page 117.

Director includes an SNMP browser that enables you to view detailed information on SNMP devices and managed groups. For example, if the performance of a network server, hub, router, or concentrator begins to degrade, you can use the SNMP browser to view the status of critical resources on selected systems configured for SNMP management.

Understanding SNMP management

SNMP functions require that information be structured using System Management Information (SMI), Version 1, format. Management Information Bases (MIBs) conforming to SMI Version 1 are used by manufacturers of SNMP manageable devices to specify the device attributes that can be accessed by end users. In addition, a MIB is used as a translation reference for the SNMP browser. Without MIBs, you cannot set attributes, such as text strings.

MIB Requirements for the SNMP browser

The SNMP Browser ships management information base (MIB) files associated with the MIB2 and RMON tables, as well as Microsoft LAN Manager; however, Director provides a MIB compiler that enables you to specify and compile MIBs that are not supplied by Director. Compiled MIBs enable the SNMP browser to more elegantly display the information associated with the MIB, and to set associated values on the SNMP device. Refer to the online help for details on the compilation procedure.

MIB requirements for Director services

Director recognizes MIBs in the System Management Information (SMI) Version 1 format. Director ships with a few MIBs necessary to recognize resource monitor devices and to aid the acquisition of certain inventory items. The MIBs shipped with Director compile the first time the Director management server starts. Additional MIBs may be compiled as needed from the Director Management Console.

Performing SNMP tasks

From the Director Management Console, you can:

- Specify SNMP discovery parameters to pinpoint devices and device groups in your network
- Specify community names for device access
- Compile new MIBs on the Director server
- Invoke the SNMP browser to view SNMP-formatted data

Understanding SNMP discovery

Director will discover SNMP devices in your network according to discovery parameters which you can specify. You can set SNMP discovery parameters to search for specific SNMP devices or groups of devices.

Note: SNMP devices must use either the IP or IPX network transport to be discovered. For example, SNMP devices that use NetBIOS as their sole network transport cannot be discovered and viewed through Director. Refer to "Installing the Server" on page 43 for details on configuring your SNMP device's network transport.

The process used to discover SNMP devices in your network uses lists of initial IP addresses, community names, and subnet masks.

The IP addresses should include the Domain Name Server of your network, the address of the system that acts as the router of your network, other addresses for network bridges (if they are configured for SNMP), and Windows NT Primary Domain Servers. These are locations in your network that contain information about the various systems and devices in your network, and will point to other addresses of additional SNMP devices for Director to discover.

SNMP devices and agents use *community names* to control their access. A community name can be any case-sensitive text string. By default, the community name of an SNMP device is set to public, indicating that access is not restricted. If specific SNMP devices in your network have unique community names to restrict access, you can specify the correct name to gain access to the device. Ideally, your list of community names should have the most publicly accessible names at the top of the list, down to the community names with the least public access. This allows Director to find the most desirable community name for your device.

Note: Be sure your community names are valid names that your device understands, otherwise Director will presume this is a non-SNMP address.

The subnet mask allows you to further refine the scope of the discovery process, limiting the search to certain subnets in the network. The default subnet mask is set to the subnet of each corresponding initial IP address.

Using your lists of IP addresses, community names, and subnet masks, a series of SNMP GET statements are performed against port 161 of the IP address to determine if the address is a valid SNMP device of some kind. If it is determined to be a valid SNMP device, another series of SNMP GET statements are sent to obtain information in the atTable, where additional IP addresses can be used to discover even more SNMP devices. The search continues until no new addresses are located.

Note: This discovery process only applies to SNMP devices using the IP network transport. Devices using IPX are simply discovered by Director, applying the community names as appropriate.

Setting SNMP discovery parameters

From the menu bar of the Director Management Console, select **Options** → **Discovery Preferences**. When the Discovery Preferences window is displayed, select the **SNMP Discovery** tab.

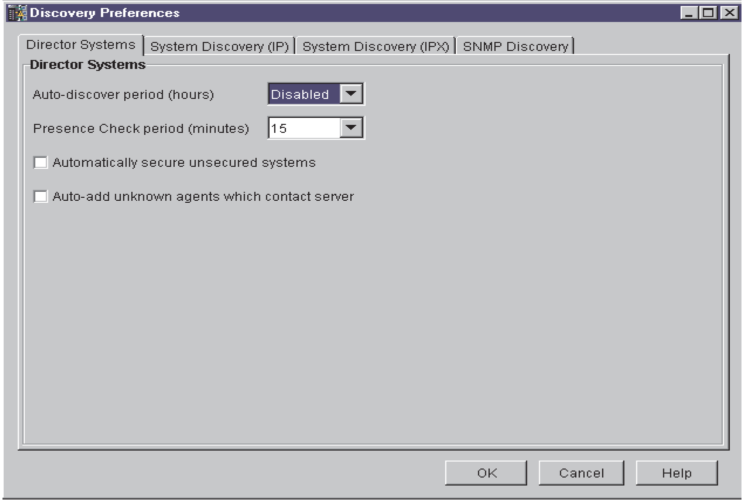

Auto-discover period (hours) Disabled

Presence Check period (minutes) 15

□ Automatically secure unsecured systems

□ Auto-add unknown agents which contact server

Use the Add, Replace, and Remove buttons under each pane to create your lists of IP addresses, corresponding subnet masks, and community names. Make sure the IP addresses use the standard dotted decimal numeric format, and that they lead to devices with SNMP agents on them. Ideally, they should go to the domain name server, or the router of the network, or the domain server.

Your subnet mask should be the same as what is used throughout the network. You can find this for your NT system by bringing up the context menu for the **Network Neighborhood** on your desktop. Choose **Properties**, then select the **Protocols** tab and double-click on **TCP/IP**. The subnet mask will be displayed. You can also specify 0.0.0.0, which is equivalent to using the device’s own subnet mask.

Note: For more information on network masks and how they work, refer to `http://www.freesoft.org/CIE/Topics/24.htm`, which contains details about subnetting and how subnet masks work (as documented in RFC 950).

Your community names should be ordered from most public access on top, to least public access on the bottom. Ensure that at least one community name gives access to the atTable of the router. See the online help for the procedure to set SNMP community names.

You can also set an Auto-discover period, in hours, and a **Presence Check period**, in minutes. These are disabled, by default. Refer to the online help for details.

Creating a new SNMP device

You can create a new SNMP device in your network and make it available for discovery by Director.

In the Group Contents pane of the Director Management Console, select **New→ SNMP Devices** from the context menu. The Add SNMP Devices window is displayed.

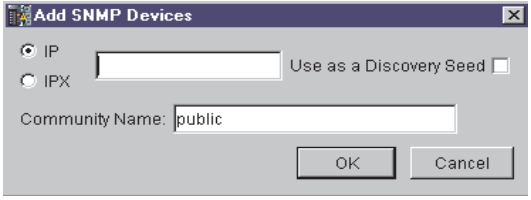

● IP

○ IPX

Use as a Discovery Seed □

Community Name: public

Select either the IP or IPX network transport, and then enter the network address. For IP, the dotted decimal address must be specified. Specify a community name for the device (be sure it is a valid name the router will recognize, and remember the case-sensitivity), and check if you want this device address used as a Discovery Seed, or an initial address for discovering additional SNMP devices.

Click **OK** to add the SNMP device to the Group Contents pane or click **Cancel** to quit.

Starting the SNMP Browser

You can view the attributes of SNMP and RMON devices using the SNMP Browser.

Start the SNMP Browser from the Director Management Console by normal drag and drop methods between the SNMP Browser icon in the Tasks pane and the desired managed systems or group icons. You can also select **SNMP Browser** from the context menu of the SNMP device or SNMP or RMON group.

Viewing SNMP information

The SNMP Browser is displayed and initially shows a tree view of the MIB structure for the SNMP or RMON devices you selected. You can expand the tree view for active systems and see their corresponding attributes. If a system is not active, its tree view cannot be expanded.

If no compiled MIBs are on the Director management server to format the information, or if the device returns information not found in a compiled MIB, the information is displayed in a dotted-decimal numerical format. If the information corresponds to a compiled MIB, the information is displayed in text format.

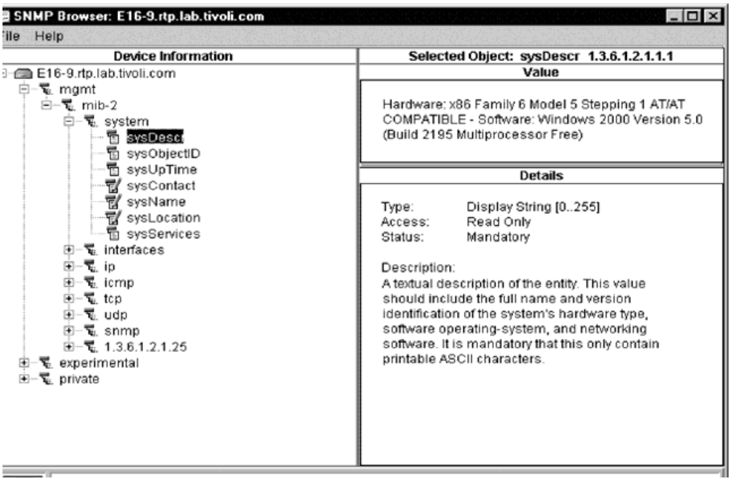

Selected Object: sysDescr 1.3.6.1.2.1.1.1

Value

Hardware: x86 Family 6 Model 5 Stepping 1 AT/AT COMPATIBLE - Software: Windows 2000 Version 5.0 (Build 2195 Multiprocessor Free)

Details

Type: Display String [0..255]
Access: Read Only
Status: Mandatory

Description:
A textual description of the entity. This value should include the full name and version identification of the system's hardware type, software operating-system, and networking software. It is mandatory that this only contain printable ASCII characters.

In the Device Information pane, information is displayed in a tree view. Device attributes are displayed in the Selected Object pane. You can expand the tree until a specific device and its corresponding attributes are displayed.

The Selected Object pane is now divided into two sections that contain details about the selected attribute from an SNMP device. The Value section (top) shows the value of the selected attribute, and the Details section (bottom) displays the characteristics of the selected attribute. This information includes, for example, the type and access status of the device attribute and a description of the device attribute.

If a “snap-in” is available for the selected attribute, then it appears on the right side of the SNMP browser in place of the Selected Object pane.

In the example shown in the preceding figure, the highlighted attribute, sysDescr, cannot be set to a value. It is a read-only attribute, and is listed as such in the lower section of the Selected Object pane. Other attributes, such as sysContact, sysName, and sysLocation, can be set to a value and are listed as read/write capable. Notice the two different icons for these attributes.

You can enter a value for those read/write attributes that have compiled MIBs by entering or changing the value in the box in the top portion of the Selected Object pane. After entering or changing values, click on the **Set** button to save the changes.

Multi-homed support

Discovery filters out certain types of transient TCP/IP addresses, like those associated with dial-up connections, on multi-homed devices.

A multi-homed device has two or more physical connections and requires multiple TCP/IP addresses, one corresponding to each network connection on the device.

To open a multi-homed device, right-click the device on the **Group Contents** pane and then click **Open...**. There will be more than one TCP/IP address listed for the device.

When viewing inventory on a multi-homed device, the IP address table will have multiple rows.

Chapter 15. CIM management

Director provides Common Information Model (CIM) support for the browser, inventory, resource monitoring, and event management tasks. The CIM is an implementation-neutral object-oriented schema for describing network management information. The Desktop Management Task Force, Inc. (DMTF) develops and maintains CIM specifications. For in-depth information on CIM, refer to `http://www.dmtf.org` on the Web.

This chapter describes how to use the CIM Browser to view and change property values and execute methods of specific class instances. For information on isolating CIM data for the inventory and resource monitoring, see the following chapters:

- Chapter 6, “Inventory Management,” on page 91.
- Chapter 8, “Resource Monitoring,” on page 109.

Unlike DMI events, CIM events are not automatically detected by Director. The Director Software Development Kit provides information on how to set up managed systems to map CIM events to Director events. When the mapping file is defined, Director can detect and present CIM events for filtering.

CIM requirements

To provide CIM data, managed systems must be running under Windows ME, Windows 98, Windows 2000, or Windows NT 4.0. They must have Windows Management Interface (WMI) Core Services Version 1.1 installed. WMI Core Services do not have to be present when the Director management agent is installed. You can add WMI to a managed system after Director is installed. When the managed system restarts, it is enabled for CIM operations.

Performing CIM Browser tasks

The CIM Browser enables you to perform the following actions:

- View the CIM structure for a selected CIM-enabled system.
- View property values for selected classes.
- Set values for individual properties.
- Execute the methods of selected class instances
- Define browser subtasks for specific CIM classes.

When you apply the CIM Browser to a managed system, the information is gathered directly from the target system and displayed. If you change a property value, Director attempts to update the value on the target system.

Starting the CIM Browser and viewing information

To start the CIM browser and view information for a single managed system, use the following procedure:

1. Select the managed system for which you want to view information and drag it to CIM Browser in the Tasks pane. The CIM Browser window appears. It uses the name of the system you selected in the CIM Classes pane.

   - If a system is not configured for CIM, a message appears indicating that the target system does not support the task.
   - If the system is inaccessible, for example, if it is offline, the CIM Browser window is opened but the system's CIM tree cannot be expanded.
   - If one or more of the systems is not configured for CIM, a message appears indicating that at least one of the target systems does not support the task.
   - If one or more of the systems is inaccessible, the CIM Browser window is opened but one or more of the systems is shown as grayed out and its CIM tree cannot be expanded.

   The CIM Classes pane displays the systems you selected.
2. To open the browser for two or more systems, select the managed systems for which you want to view information. Drag the CIM Browser from the Tasks pane to any system in the set of systems highlighted.
3. To turn the displaying of system classes on or off, right-click on a system and select **Display System Classes** from the context menu.

   A check mark indicates that displaying is set on. You can toggle on or off the displaying of CIM system classes. System classes are indicated by a double underscore that precedes the class name (__*classname*).
4. Double-click the system to display the CIM name spaces of the system. Double-click on a name space to display its classes.

   You can continue to expand each class by double-clicking until you reach the leaf classes.
5. To view an instance of a class, click on the class name.

   If an instance of the class is found, it appears in the upper right-hand pane labeled Instances: and the associated properties and methods appear under the Properties and Methods tabs in the lower right-hand pane. A class does not have to be a leaf class to have associated properties or methods.
6. To reverse the order of the properties or methods, right-click any line item and select **Sort → Descending**.
7. When you finish viewing information, select **File → Close**.

Setting a property value for a CIM class instance

It is strongly recommended that you do not change the value of a property unless you are thoroughly familiar with the structure and manipulation of CIM data. Improperly setting the value of a system can cause unpredictable results on the target system.

To change the value of a property, use the following procedure:

1. Navigate to the property for which you want to change a value using the procedure described above in "Starting the CIM Browser and viewing information" on page 160.
2. Right-click the value on the property row and click **Set Value** in the context menu. The Set Value window opens with the current value.
3. Enter the new value and click **OK** to enact the change. If you do not want Director to attempt to change the value, close the window or click **Cancel**.

   If Director cannot change the value on the target system, a message indicates the failure.

Executing a method for a CIM class instance

It is strongly recommended that you do not execute a method unless you are thoroughly familiar with the structure and manipulation of CIM data. Executing a method can cause the connection to the target system to be lost.

To execute a method for a CIM class, use the following procedure:

1. Using the procedure described in "Starting the CIM Browser and viewing information" on page 160, navigate to the class that has the method you want to execute. The associated methods appear on the Methods page in the lower-right pane.
2. Right-click a method and click **Execute** from the context menu. The Execute Method window opens.
3. If the method receives any input arguments, one or more Input fields appear. Enter the arguments in these fields.
4. Click **Execute** at the bottom of the Execute Method window to run the method. If you do not want to run the method, close the window. If Director is unable to run the method on the target system, a message indicates the failure.

Defining CIM Browser subtasks

You can define two types of browser subtasks:

- User-selected class that, when applied to a system, displays only the instances, properties, and methods associated with the specified class on the selected system
- User-selected method that, when applied to a system, executes the method on the selected system.

By creating browser subtasks, you can bypass navigating through the class tree to reach a specific class or method.

Defining a browser subtask for a CIM class

To define a browser subtask for a specific class, use the following procedure:

1. Navigate to the class for which you want to create a subtask using the procedure described in "Starting the CIM Browser and viewing information" on page 160.
2. Right-click anywhere on the class name and click **Create browser task for class**. A window opens with the name of the class entered as the default name.
3. You can enter a new name or keep the default name. To keep the default name, click **OK**. Enter the new subtask under **CIM Browser** in the Director Management Console window.
4. Apply the browser subtask to a CIM-enabled managed system that has the instances, properties, and methods associated with those in the subtask.

Defining a browser Subtask for a CIM class method

To define a browser subtask for a specific method, use the following procedure:

1. Use the procedure described in "Starting the CIM Browser and viewing information" on page 160 to navigate to the CIM class that has the method for which you want to create a subtask. The associated methods appear on the Methods page in the lower-right pane.
2. Right-click a method and select **Execute** from the context menu. The Execute Method window opens.
3. If the method receives any input arguments, one or more I**nput** fields appear. Enter the arguments in these fields.
4. Click **Save** at the bottom of the Execute Method window. A window opens with the name of the method entered as the default name.
5. You can enter a new name or keep the default name. To keep the default name, click **OK**. Enter the new subtask under **CIM Browser** in the Director Management Console window.
6. To run the method on a selected system, apply the browser subtask to a CIM-enabled managed system that supports the method you are attempting to run.

   Because method subtasks are not interactive, you can either run the task immediately or use the task scheduler to schedule the subtask to run at a specified time. Refer to "Starting the Task Scheduler task" on page 201 for information on task scheduling.

Chapter 16. Asset ID

Asset ID makes it possible to track lease, warranty, user, and system information, as well as serial numbers for major system components. You can use Asset ID to create personalized data fields for additional asset tracking. In addition, you can mass configure large numbers of managed systems by building a Profile of a system from a single location and using the mass configuration option.

You retrieve Asset ID information from the IBM Director Agent installed on any Director-managed system. The IBM Director Agent reads Asset ID data from systems that have the Enhanced Asset Information Area EEPROM. Systems without the EEPROM can store Asset ID settings in a file on the hard drive.

The Asset ID task can be used on systems running Windows and on systems running Linux. The interface GUI will change depending on the operating system that the target system is using.

Note: IBM Director Agent writes to and retrieves some Asset ID data from the Desktop Management Interface (DMI) on a Director-managed system that does not include the Enhanced Asset Information EEPROM.

Using the Asset ID Interface

To display the Asset ID interface, drag and drop the Asset ID task onto a managed system in the Group contents pane of the Director Management Console. Upon activation, the Asset ID window opens to display the data polled from an IBM Asset ID equipped system or another DMI-enabled system.

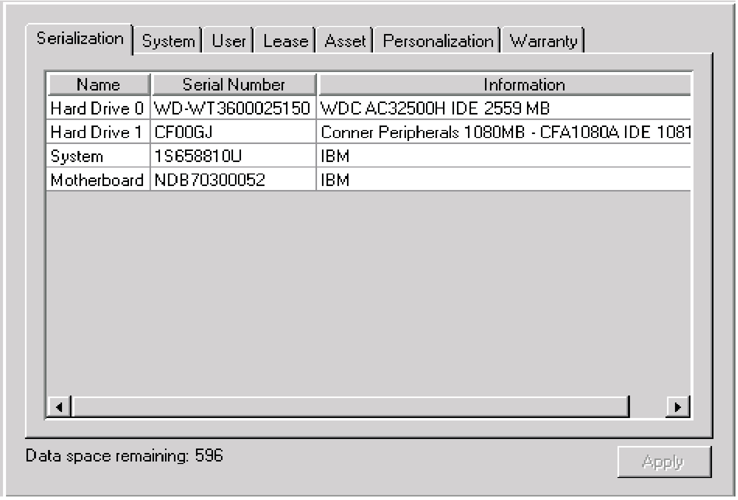

Name	Serial Number	Information
Hard Drive 0	WD-WT3600025150	WDC AC32500H IDE 2559 MB
Hard Drive 1	CF00GJ	Conner Peripherals 1080MB - CFA1080A IDE 1081
System	1S658810U	IBM
Motherboard	NDB70300052	IBM

Data space remaining: 596

The Asset ID interface contains the following tabbed interfaces:

Tab	Description
Serialization	Displays the serial numbers for the agent system hardware.
System	Displays the current agent system characteristics: system name, MAC address, user login name, operating system, GUID address, IBM LAN Agent Control Manager Profile.
User	Displays the user profile: user name, telephone number, work location, department, and professional position.
Lease	Displays the information on the lease agreement for the agent system hardware.
Asset	Displays the inventory factors that are related to the agent system.
Personalization	Displays the free-form window where you can add information on your systems, users, or computers.
Warranty	Displays the information on the warranty agreement for the agent system hardware.
Serialization Interface	
Name	The hardware component name.
Serial Number	The serial number for the hardware component.
Information	Descriptive information for the hardware component.
System Interface	
System Name	The NetBEUI name of the agent system (the computer name as it appears under **Network Properties**). NetBEUI is NetBIOS extended user interface, and NetBIOS is network basic input/output system.
MAC Address	The unique hexadecimal character string that identifies the network adapter in the agent system.
Login Name	The user ID that the system administrator assigned at installation.
Operating System	The operating system (for the management server or for the computer where IBM Director Agent resides).
System GUID	The agent system Global Unique Identifier (GUID). This is your BIOS unique ID number.
LCCM Profile	The profile name of the IBM LAN agent Control Manager (LCCM), if applicable.
User’s Interface	
Name	The user login name.
Phone	The user phone number.
Location	The user office location.

Tab	Description
Department	The user department name or number.
Position	The user job title.
Lease Interface	
Start Date (mm/dd/yy)	The date that the lease agreement began.
End Date (mm/dd/yy)	The date that the lease agreement ends. If a Lease End date is specified, a Warning alert will be generated when the Lease expires.
Term (months)	The number of months for which the agent system is leased.
Amount	The total price of the lease agreement.
Lessor	The name of the company that leased the agent system.
Personalization	
Purchase Date (mm/dd/yy)	The date the agent system was purchased
Last Inventoried (mm/dd/yy)	The date of the last inventory check.
Asset Number	A unique number that is assigned to the agent system for inventory purposes.
RF-ID	The radio-frequency identification (RF-ID) number that was encoded in the agent system by the manufacturer. Not all computers have RF-ID capabilities. This is a fixed field and cannot be changed.
Warranty Interface	
Duration (months)	The duration of the warranty agreement.
Cost	The total cost of the warranty.
End Date (mm/dd/yy)	The date that the warranty ends. If a Warranty End Date is specified, a Warning level alert will be generated when the Warranty expires.

At the bottom of the Asset ID window is the Data space remaining information line. This information is an indicator of the amount of remaining available data storage on the EEPROM, where nnn represents this storage as a number of characters that can be entered. After the limit has been reached, the Data space remaining line turns red. At this point, any further information that is entered is discarded.

Click **Apply** to write to the EEPROM any information added in the Asset ID window.

Click **File** → **Close** to close the Asset ID window.

Click **Help** → **Window Help** to open the online help window.

Using the Profile Builder

From the Task pane in the Director Management Console, right-click **Asset ID**. and select **Profile Builder** . The Profile window opens.

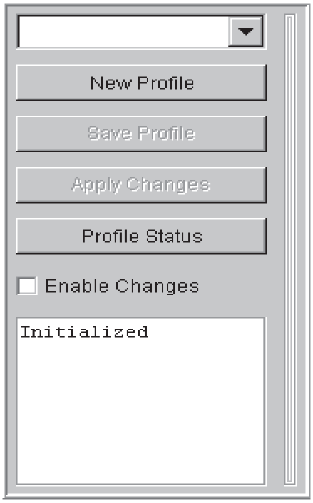

```
Initialized
```

Creating a profile

There are several way to configure a profile. You can create a new profile or modify an existing profile.

Creating a new profile

To create a new profile use the following procedure:

1. From the Profile window, click **New Profile**. The Input window opens.
2. Enter a New Profile name.
3. Click **OK**. The Asset ID window opens.

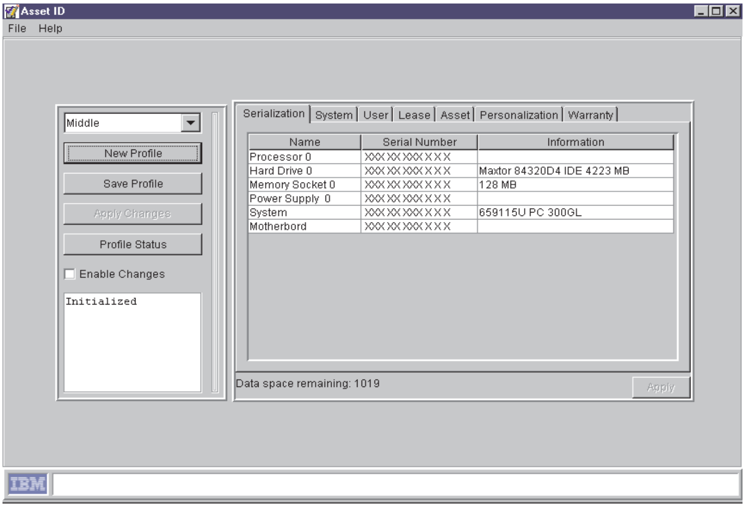

□ Enable Changes

Initialized

Name	Serial Number	Information
Processor 0	XXX XX XXX X X X	
Hard Drive 0	XXX XX XXX X X X	Maxtor 84320D4 IDE 4223 MB
Memory Socket 0	XXX XX XXX X X X	128 MB
Power Supply 0	XXX XX XXX X X X	
System	XXX XX XXX X X X	659115U PC 300GL
Motherbord	XXX XX XXX X X X	

Data space remaining: 1019

4. Enter Asset ID data. See “Using the Asset ID Interface” on page 163.
5. Click **Save Profile**. The Save Profile window opens.
6. Click **Yes** to save the profile or No.

Modifying an existing profile

From the Profile window, click the drop-down list to display all profiles. Select an existing profile and click **Profile Manager**. The Status window opens. Select **Profile**. Modify the existing profile by editing any data in the Asset ID interface. See “Using the Asset ID Interface” on page 163. After the modified profile is created, click the **Save** button.

Using the Status window

The Status window displays a list of profiles, and groups assigned to those profiles. It also displays the status of a profile. To start the Status window, from the Director Management console, right-click any saved profile. A context menu will appear. Click **Profile Manager** to start the manager.

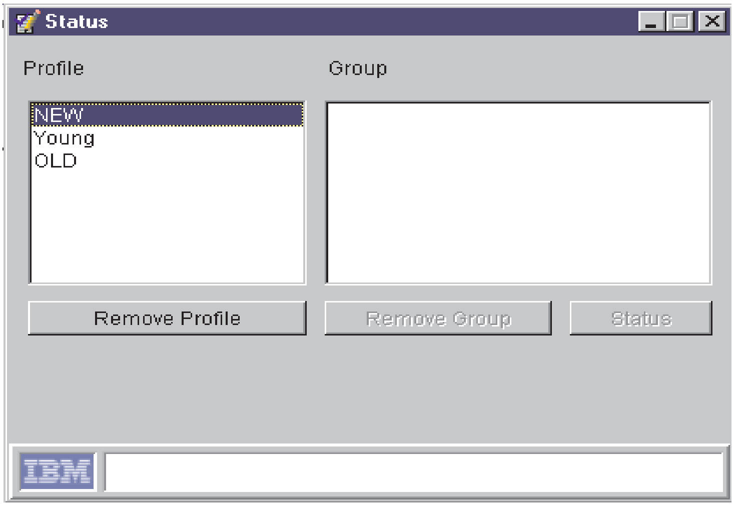

After a profile has been created, one or more Director groups might be associated with the profile. Profile Manager enables you to determine which groups are associated with any given profile.

When the Profile Manager is started all profiles are listed in the Status window. The first profile in the list is highlighted. All groups associated with a selected profile appear under Groups in the Status window. The Remove Group button and the Status button are disabled until a group is selected.

If a profile is not currently associated with any groups, you can select it and click **Remove Profile** to remove it. If you try to remove a profile associated with another group, a window will open stating that the operation is unsuccessful because a group is a member of another profile.

Viewing the status of a profile

To view the status of a group, select a profile and a group; click the **Profile Status** button.

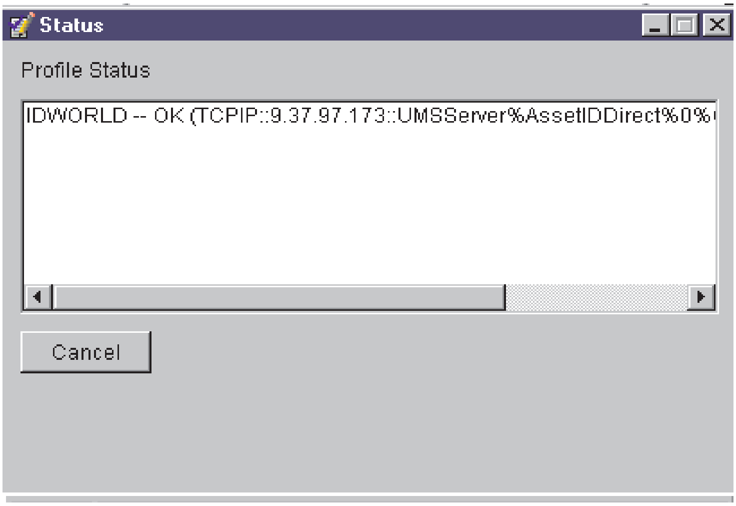

IDWORLD -- OK (TCPIP::9.37.97.173::UMSServer%AssetIDDirect%0%

The dialog box displays a list of systems within a group and the current status of the managed systems. A managed system will display one of the following status indicators:

- OK: The system has been configured correctly.
- Failed: Mass configuration was able to communicate with the system but unable to configure it correctly, or mass configuration has exhausted all attempts to communicate with the target system.
- Pending: The target system is currently offline and mass configuration will attempt to configure it at a later time, or mass configuration has scheduled the configuration attempt for a later time.

Using Mass Configuration

Before you can use the Mass Configuration task, a profile must be created using the Profile Builder. For more information see “Using the Profile Builder” on page 166. To start the mass configuration task, double-click the Network icon to expand the Profiles list. Drag-and-drop a Profile onto a managed system that you wish to mass configure.

Chapter 17. Alert on LAN

A user with administrative security-status can use the Alert on LAN task to set the options related to network system alerts. Alert on Lan makes it possible to track network information. The Alert on LAN options can be mass configured.

Using the Alert on Lan interface

To configure Alert on LAN, from the Director Management Console, drag-and drop-the Configure Alert on Lan icon onto a managed system in the Group Contents pane.

When you select **Alert on LAN** and apply it to an applicable system, the following screen is displayed.

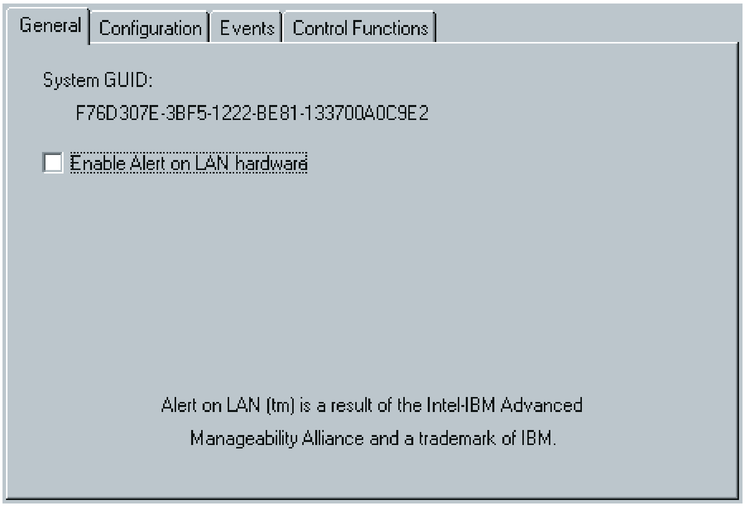

System GUID:
F76D307E-3BF5-1222-BE81-133700A0C9E2

□ Enable Alert on LAN hardware

Alert on LAN (tm) is a result of the Intel-IBM Advanced Manageability Alliance and a trademark of IBM.

The following items are available on the Alert on LAN screen.

Item	Description
General Tab	
System GUID	A Global Unique ID (GUID) is assigned to each system board for system-management purposes. The GUID is stored in the BIOS on the system board.
Enable Alert on LAN hardware	This option determines whether the system alerts are on or off. Select the check box to enable system alerts.

Item	Description
Configuration Tab	
Proxy server (IP address port)	The internet protocol address for the server you use to communicate with the agent systems. The IP address is assigned by the system administrator. (default port is 5500.)
Heartbeat timer period	The Alert on LAN proxy computer verifies that the agent system is running. This is the number of seconds between system checks. The default value is 32. The enabled heartbeat timer period values range from 43 to 5461 seconds and can be set in intervals of 43 seconds.
Watchdog Timer Period	If the watchdog timer indicates that a agent system has stopped, the watchdog timer automatically sends a message to the proxy computer. This is the period between polls for the watchdog timer (measured in seconds). The default value is 43. The watchdog timer period values range from 86 to 5461 seconds and can be set in intervals of 86 seconds.
Transmission attempts	The number of retries for transmission after the agent system stops. The default value is 30.
Event Polling Period	The polling period for software problems. The default value is 30.
Events Tab	
Cover Tamper	If the cover of the managed system has been opened or removed, an event message is generated.
LAN Leash Tamper	LAN Leash detects if a agent system is disconnected from the LAN, even when the computer is off. If a agent system is disconnected from the LAN, an event message is generated.
Temperature Out of Specification	If the microprocessor temperature is out of the specified range, an event message is generated.
Watchdog	If the operating system of the managed system is not functioning, or is in a suspended state, an event message is generated.
Voltage Out Specification	If there is a dramatic change in the voltage of the power supplied to any part of the agent system that is an event message is generated.

Item	Description
Auto-clear events	If this option is enabled, the agent system sends an alert each time the condition is present (multiple alerts). If this option is disabled, the system sends an alert for a condition only once (no reminder alerts).
Events Enabled	Selecting this option enables all events to be monitored. To select an individual event, select the particular event in the Enable row.
Clear All Events	Select this option and click **Apply**, to clear the events log.
Control Functions Tab	
Power Down	Receives this message as a system state report.
Power Up	Receives this message as a system state report.
Reboot	Receives this message as a system state report.
Presence Ping	Returns the message that the system is not on but is still connected to the network.

If you make any changes to your default user options for Alert on LAN, click **Apply** to save your options.

Using the Profile Builder

From the Task pane in the Director Management Console, right-click **Asset ID**. and select **Profile Builder.** The Profile window opens.

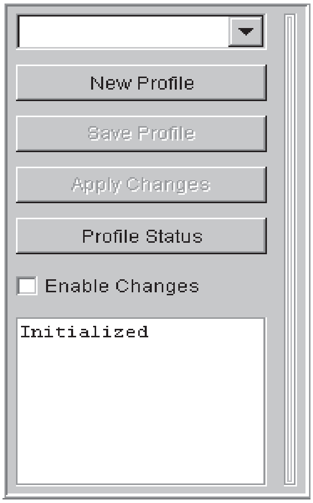

□ Enable Changes

```
Initialized
```

From this window, you can create a profile.

Creating a profile

There are several way to configure a profile. You can create a new profile or modify an existing profile.

Creating a new profile

To create a new profile use the following procedure:

1. From the Profile window, click **New Profile**. The Input window opens.
2. Enter a New Profile name.
3. Click **OK**. The Asset ID window opens.

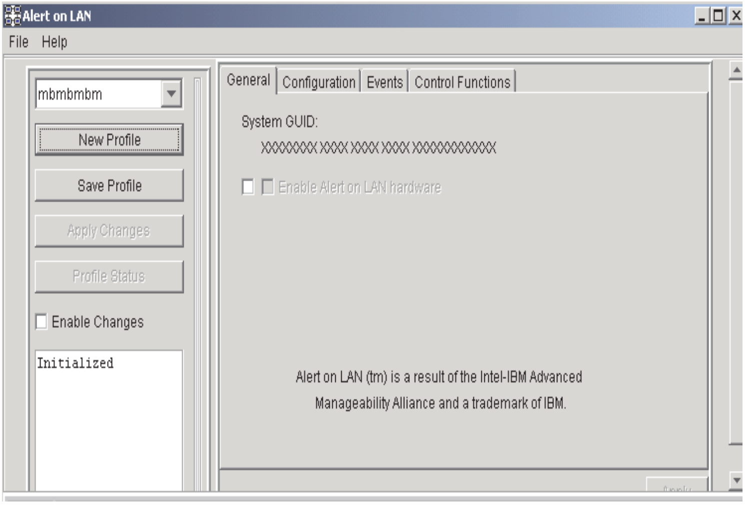

Initialized

System GUID:
XXXXXXXX XXXX XXXX XXXX XXXXXXXXXXXX

Alert on LAN (tm) is a result of the Intel-IBM Advanced Manageability Alliance and a trademark of IBM.

4. Enter the Alert on Lan data. See "Using the Alert on Lan interface" on page 171.
5. Click **Save Profile**. The Save Profile window opens.
6. Click **Yes** to save the profile or No.

Modifying an existing profile

From the Profile window, click the drop-down list to display all profiles. Select an existing profile and click **Profile Manager**. The Status window opens. Select **Profile**. Modify the existing profile by editing any data in the Asset ID interface. See "Using the Alert on Lan interface" on page 171. After the modified profile is created, click the **Save** button.

Using the Status window

The Status window displays a list of profiles, and groups assigned to those profiles. It also displays the status of a profile. To start the Status window, from the Director Management console, right-click any saved profile. A context menu will appear. Click **Profile Manager** to start the manager.

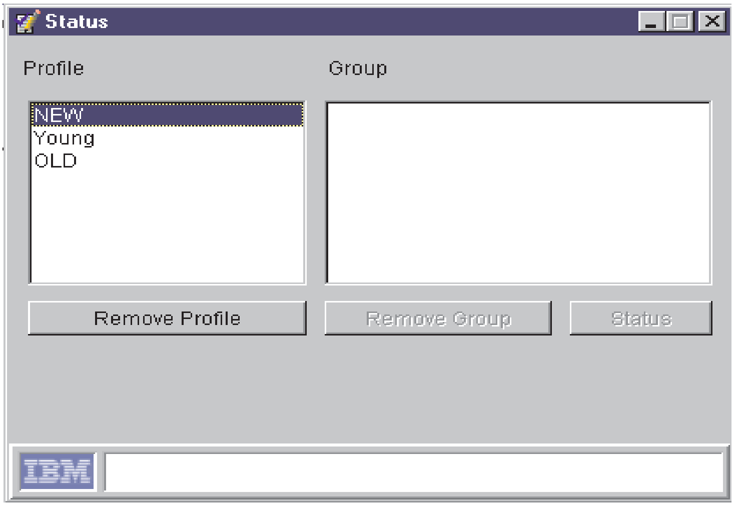

After a profile has been created, one or more Director groups might be associated with the profile. Profile Manager enables you to determine which groups are associated with any given profile.

When the Profile Manager is started all profiles are listed in the Status window. The first profile in the list is highlighted. All groups associated with a selected profile appear under Groups in the Status window. The Remove Group button and the Status button are disabled until a group is selected.

If a profile is not currently associated with any groups, you can select it and click **Remove Profile** to remove it. If you try to remove a profile associated with another group, a window will open stating that the operation is unsuccessful because a group is a member of another profile.

Viewing the status of a profile

To view the status of a group, select a profile and a group; click the **Profile Status** button.

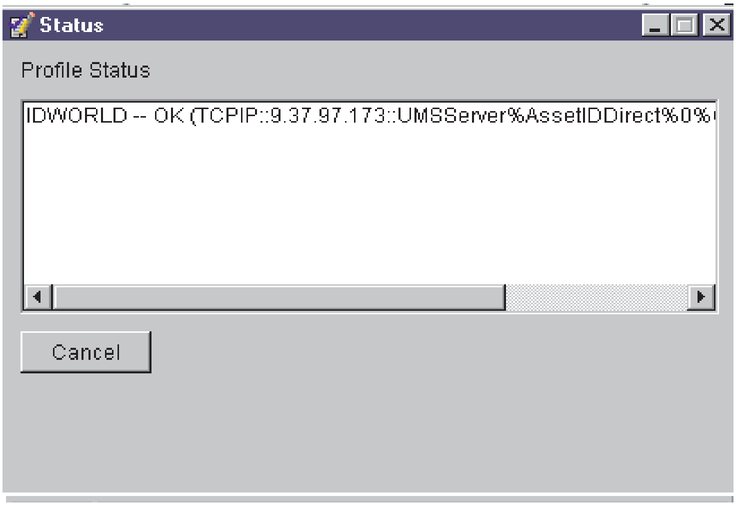

IDWORLD -- OK (TCPIP::9.37.97.173::UMSServer%AssetIDDirect%0%

The dialog box displays a list of systems within a group and the current status of the managed systems. A managed system will display one of the following status indicators:

- OK: The system has been configured correctly.
- Failed: Mass configuration was able to communicate with the system but unable to configure it correctly, or mass configuration has exhausted all attempts to communicate with the target system.
- Pending: The target system is currently offline and mass configuration will attempt to configure it at a later time, or mass configuration has scheduled the configuration attempt for a later time.

Using Mass Configuration

Before you can use the Mass Configuration task, a profile must be created using the Profile Builder. For more information, see “Using the Profile Builder” on page 173. To start the mass configuration task, double-click the Alert on Lan icon to expand the Profile list. Drag-and-drop a Profile onto a managed system that you wish to mass configure.

Chapter 18. Network Configuration

The Network configuration task provides the following information about the agent system: routing, distributed database system that is used to map domain names to IP address, Windows Internet Naming Service (WINS) server, the domain or workgroup for the agent and the modems that are installed on the agent system. These options can be mass configured.

Using the Network interface task

To start the Network configuration task from the Director console, drag-and-drop the Network icon onto a group of managed systems. When you select the Network Task and apply it to an applicable system, the following screen is displayed.

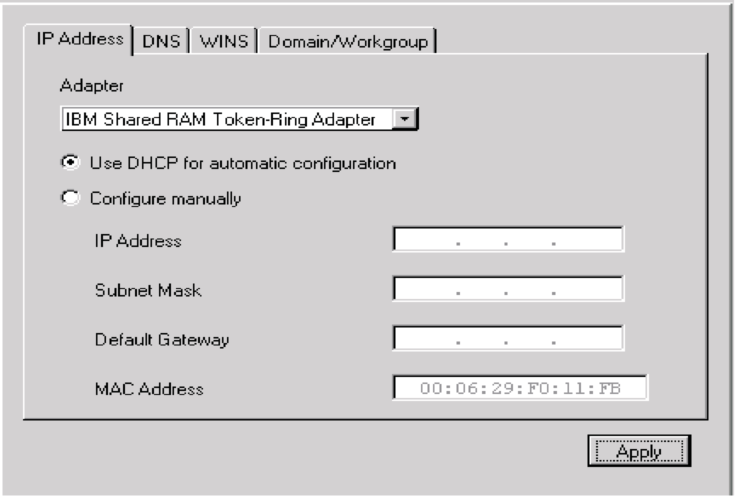

Adapter: IBM Shared RAM Token-Ring Adapter

● Use DHCP for automatic configuration

○ Configure manually

IP Address: . . .

Subnet Mask: . . .

Default Gateway: . . .

MAC Address: 00:06:29:F0:11:FB

Notes:

1. You cannot configure Network Options for Windows 98 and Windows ME agents. However, you can view the settings.

2. The Network configuration task can be used on systems running Linux. The task requires a NIC on the target system. It provides an interface for setting the parameters for each network adapter installed in the system.

3. The interface GUI will change depending on the operating system that the target system is using.

The following items are available on the Network interface.

Item	Description
IP Address Interface	
Network Adapter	Select the appropriate network adapter from the list.
Use DHCP for automatic configuration	Select this option to configure IP addresses automatically.
Configure manually	Select this option to configure IP addresses manually. When this option is selected, the remaining entry fields are enabled.
IP Address	The IP address of the agent system. If you do not use DHCP to obtain an IP address, you must type the values into the **IP Address** and **Subnet Mask** fields manually.
Subnet Mask	A bit mask that is used to identify which bits in an IP address correspond to the network address and which bits correspond to the subnet portions of the address. The address mask has ones in positions corresponding to the network and subnet numbers and zeros in the host-number positions.
Default Gateway	The IP address for the default gateway server that you are using to communicate with other networks.
MAC Address	The unique hexadecimal number that identifies the network adapter in the agent system. (Read-only)
DNS Interface	
Servers	The alphabetic identifier for your server with the network domain (IP address).This option is available on systems running Linux and Windows.
Suffixes	The text strings for the domain.
Apply	Changes are not saved until you click the **Apply** button.
Wins Interface	
WINS Servers	Registered NetBIOS name with the associated IP address. This option is available on systems running Windows only.
Primary	The address for the primary server. This option is available on systems running Linux only.

Item	Description
Secondary	The address for the secondary server. This option is available on systems running Windows only.
Apply	Click **Apply** to save changes.
Domain/Workgroup	
Computer Name	The name given to the agent system. This naming scheme enables identification of the computer. This options requires that the system be a member of a Windows NT or Windows 2000 domain or Linux.
Domain	The agent is a member of the domain. A domain requires that user's validate their account before they can log onto the network.
Workgroup	The agent is a member of a workgroup. The workgroup is a collection of agents and servers with no centralized logon validation.
Apply	Changes are not saved until the **Apply** button is clicked.
Modem Interface	
Modem	The drop-down list provides the name of the modem. Only installed modems are listed.
Com port	Lists the specific port that the modem is using.
Max Baud Rate	The maximum rate at which the modem operates.
Device Type	Describes the type of modem (internal or external).

Using the Profile Builder

From the Task pane in the Director Management Console, right-click **Network**. and select **Profile Builder**. The Profile window opens.

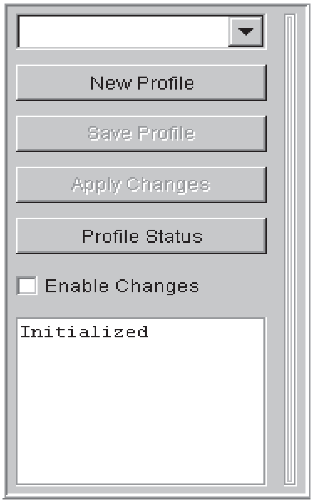

Initialized

Creating a profile

There are several way to configure a profile. You can create a new profile or modify an existing profile.

Creating a new profile

To create a new profile use the following procedure:

1. From the Profile window, click **New Profile**. The Input window opens.
2. Enter a New Profile name.
3. Click **OK**. The Network window opens.

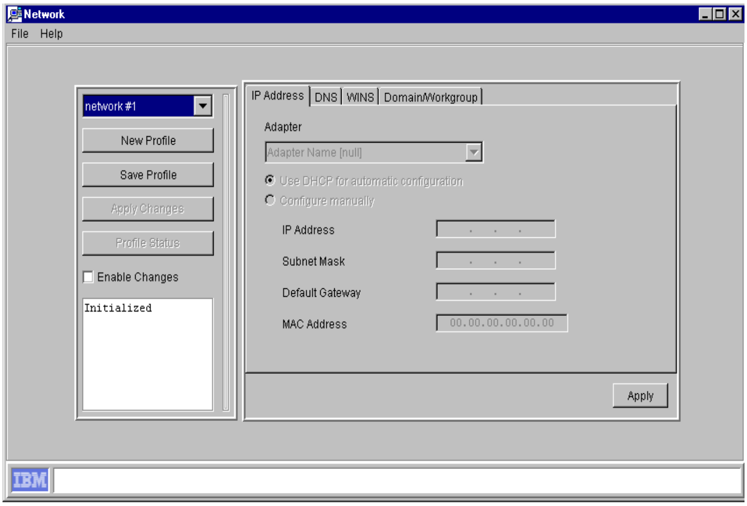

network #1

☐ Enable Changes

Initialized

Adapter

Adapter Name [null]

● Use DHCP for automatic configuration

○ Configure manually

IP Address . . .

Subnet Mask . . .

Default Gateway . . .

MAC Address 00.00.00.00.00.00

4. Enter Network data. See “Using the Network interface task” on page 179.
5. Click **Save Profile**. The Save Profile window opens.
6. Click **Yes** to save the profile or No.

Modifying an existing profile

From the Profile window, click the drop-down list to display all profiles. Select an existing profile and click **Profile Manager**. The Status window opens. Select **Profile**. Modify the existing profile by editing any data in the Network interface. See “Using the Network interface task” on page 179. After the modified profile is created, click the **Save** button.

Using the Status window

The Status window displays a list of profiles, and groups assigned to those profiles. It also displays the status of a profile. To start the Status window, from the Director Management console, right-click any saved profile. A context menu will appear. Click **Profile Manager** to start the manager.

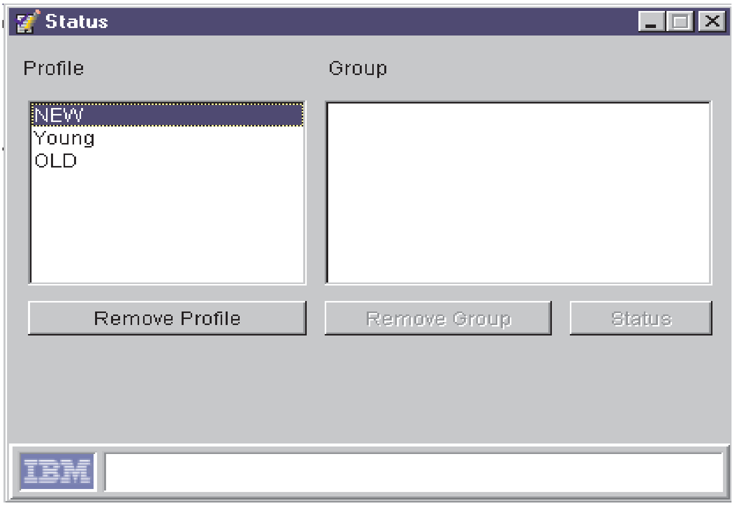

After a profile has been created, one or more Director groups might be associated with the profile. Profile Manager enables you to determine which groups are associated with any given profile.

When the Profile Manager is started all profiles are listed in the Status window. The first profile in the list is highlighted. All groups associated with a selected profile appear under Groups in the Status window. The Remove Group button and the Status button are disabled until a group is selected.

If a profile is not currently associated with any groups, you can select it and click **Remove Profile** to remove it. If you try to remove a profile associated with another group, a window will open stating that the operation is unsuccessful because a group is a member of another profile.

Viewing the status of a profile

To view the status of a group, select a profile and a group; click the **Profile Status** button.

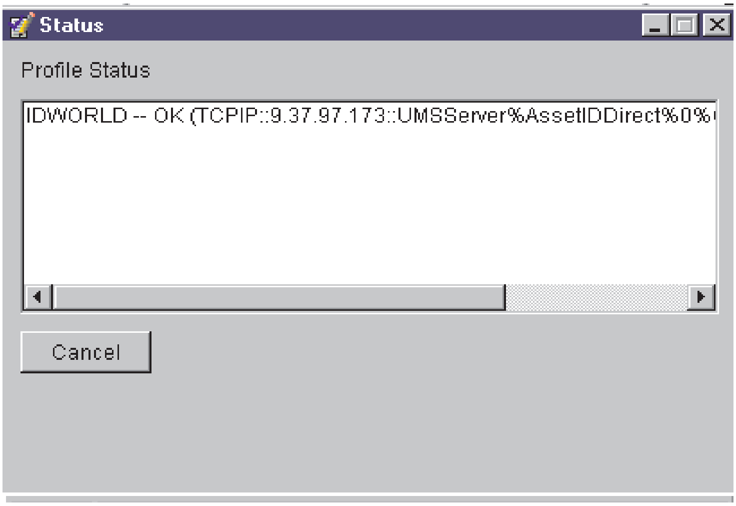

IDWORLD -- OK (TCPIP::9.37.97.173::UMSServer%AssetIDDirect%0%

The dialog box displays a list of systems within a group and the current status of the managed systems. A managed system will display one of the following status indicators:

- OK: The system has been configured correctly.
- Failed: Mass configuration was able to communicate with the system but unable to configure it correctly, or mass configuration has exhausted all attempts to communicate with the target system.
- Pending: The target system is currently offline and mass configuration will attempt to configure it at a later time, or mass configuration has scheduled the configuration attempt for a later time.

Using Mass Configuration

Before you can use the Mass Configuration task, a profile must be created using the Profile Builder. For more information see “Using the Profile Builder” on page 181. To start the mass configuration task, double-click the Network icon to expand the Profiles list. Drag-and-drop a Profile onto a managed system that you wish to mass configure.

Chapter 19. System Accounts

The System Accounts task provides remote administration of user security and group security. The System Accounts task can be used on systems running Windows and on systems running Linux. The interface GUI will change depending on the operating system that the target system is using.

Using the System Account interface

To start the System Accounts task, from the Director Management console, drag-and-drop the System Accounts icon onto a managed system in the Groups Contents pane.

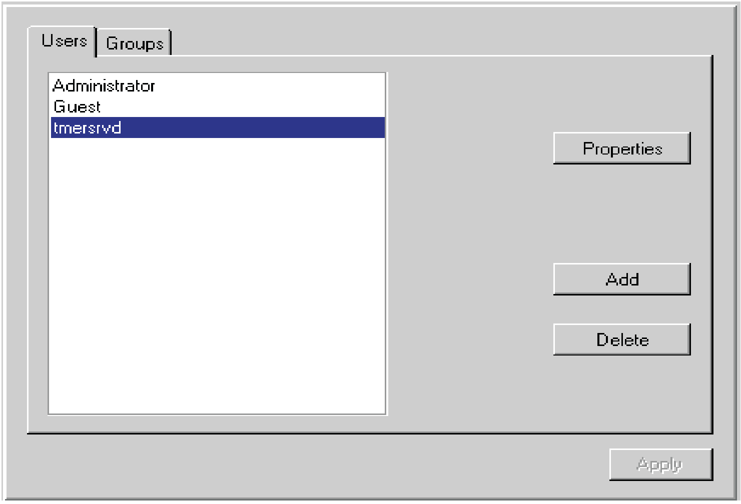

The following items are available on the System Account interface.

Item	Description
The Users interface provides information about the items that are described in the following table.	
Properties	Edit or view user properties
Add	Click the **Add** button to add a new user.
Delete	Click the **Delete** button, to delete a user.
The Groups interface enables the administrator to review and edit members within the group. Click the Groups page to display a list of all groups. The Properties and Delete buttons are initially disabled. They become enabled when a group is selected in the list.	

Item	Description
Groups	List of global groups.
Properties	Edit or review group properties.
Add	Click the **Add** button to add a new group.
Delete	Click the **Delete** button to delete a group.

Group Properties

The following interface is used to edit or review group properties.

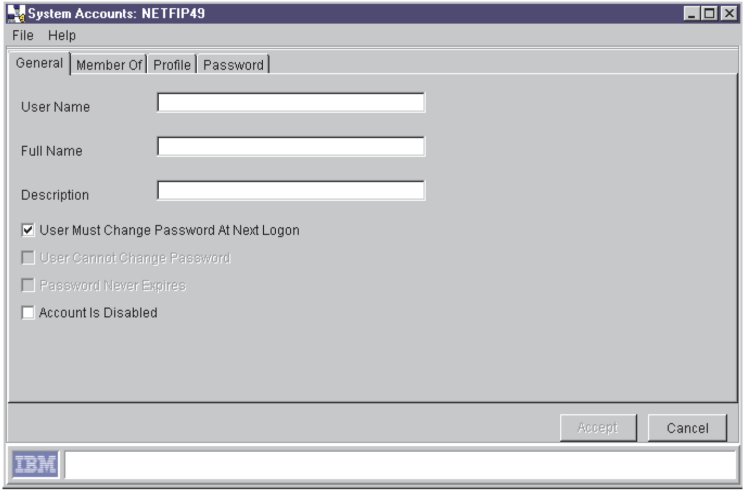

User Name

Full Name

Description

☑ User Must Change Password At Next Logon

□ User Cannot Change Password

□ Password Never Expires

□ Account Is Disabled

Item	Description
User’s Name	A unique character string that identifies the user (32 characters maximum).
Full Name	User’s complete name.
Description	Information about the user, such as title, department, or reason for granting access to the Director program on systems running Windows this is in text (32 characters maximum.).
User Must Change Password at Next Logon	Select this check box if you want to force the user to change the password the next time the user accesses the Director program.

Item	Description
User Cannot Change Password	Select this check box if you want to prevent the user from changing the password. If this check box is selected, only someone with Administrator privileges can change the password.
Password Never Expires	Select this check box if you do not require the password to be changed at scheduled intervals.
Account is Disabled	Select this check box if you want to temporarily disable a user's access to the Director program. As an Administrator, you cannot disable your own account. This ensures that at least one account with Administrator privileges remains active.
Accept	Click the Accept button to save changes.
Cancel	Select the **Cancel** button to cancel changes.
Member Of The Member Of interface displays a group membership list. Members are listed on the right pane, and non-member groups are listed in the left pane.	
Member groups	A list of users within the group.
Non-member Groups	A list of users who are not members of the group.
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.
Profile The Profile page on systems running Windows and the Group ID page on systems running Linux provide information about the items that are described in the following table.	
Path	The path to the user's profile folder. In Windows, type a network path such as \\*server name**profile folder name**user name*.
Logon Script	In Windows, a script assigned to a user account that runs each time the user logs on. In Linux, a shell script.
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.
Linux On systems running Linux, use the Linux page to enter password information.	
Date of last password change	
Must keep password	
Warn about password change	

Item	Description
Disable Account after Password Expires	
Use the Password page to enter a new password or change an existing password.	
New Password	The user’s new password (32 character maximum, case sensitive).
Confirm Password	This field must contain the same character string as the New Password field (32 character maximum, case sensitive).
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.

Chapter 20. Cluster Management

Director enables you to isolate clusters for viewing and for the resource monitoring and event management tasks. This chapter describes how to use the Cluster Browser task to view the members and member statuses of clusters. For information on isolating clusters for the resource monitoring and event management tasks, see the following chapters:

- Chapter 8. Resource Monitoring
- Chapter 9. Event management

Understanding cluster management

In Director, a *cluster* is a representation of a collection of network resources. Implementing clusters can enable you to determine the status of a logical collection of resources (*resource groups*) that you can distribute across nodes in a network or across network boundaries. For example, a Web server resource group might consist of individual resources, such as an IP address, physical disk containing the server files, and an application that defines how the server is started. One purpose of this resource group might be to ensure and enable redundancy of the Web server such that the resources could be transferred from one system to another if the Web server goes down.

Director supports only the clustering implementation of Microsoft Clustering Service (MSCS). For Windows NT systems configured with MSCS, a Director managed system interfaces with this service to obtain and present basic cluster data, including the name of the cluster, individual member nodes of the cluster, resource groups, and the resources defined for each group.

You can use the Resource Monitors task to define thresholds and use the Event Action Plans task to create event action plans for reported cluster and cluster resource statuses.

The Director Software Development Kit (SDK) provides additional programming information that can be used to extend the basic cluster support.

For more information on the Microsoft cluster implementation, visit the Web site at http://www.microsoft.com.

Cluster requirements

To provide cluster data:

- Cluster nodes must be running Windows NT Enterprise Edition (EE))4.0, Windows 2000 Advanced Server, or Windows 2000 DataCenterand must have Microsoft Clustering Service installed.

- Each node in a cluster should have the Director management agent installed.

Performing Cluster Browser tasks

The Cluster Browser task enables you to:

- Determine the structure, nodes, groups, networks, and resources associated with a cluster
- Determine the status of cluster resources
- View the associated properties of cluster resources
- Perform operations on cluster nodes, resources, and groups

Understanding cluster discovery

The Director Management Console displays clusters as both the Clusters and Windows NT Clusters groups. Because Director supports only the MSCS implementation of clusters, these groups contain the same cluster members. The Clusters group is intended as a placeholder for other cluster implementations. The Clusters and Cluster Members group contains the cluster name and the individual member nodes that contain resources defined for a cluster.

The Cluster Browser task displays data in real time, and is applied only to cluster names, not to individual cluster member nodes. When you apply the Cluster Browser to a cluster, the information is gathered from the associated member nodes and used to determine the status of the cluster and cluster resources. This status is based on the availability of the member nodes and resources assigned to the cluster.

Normal Online
One or more nodes are online and all resource groups are online and available.

Error Online
One or more nodes are online and one or more resource groups are unavailable.

Error Offline
All nodes in the cluster are online but one or more resources or resource groups are unavailable.

Normal Offline
No systems are online.

These statuses apply to the cluster objects in the Director Management Console, not in the Cluster Browser. The Cluster Browser does not display the status of a cluster as a whole. Instead, it displays the statuses of individual cluster resources, such as resource groups, nodes, networks, and network interfaces.

Starting the Cluster Browser and viewing information

To determine the individual member nodes of a cluster, in the Groups pane, click on **Clusters and Cluster Members**. All detected clusters and corresponding member nodes appear in the Group Contents pane. Follow the steps described below to view information on an individual cluster.

To start the Cluster Browser task and view the objects and object statuses of a single cluster:

1. In the Groups pane, select **Windows NT Clusters** to display all discovered clusters in the Group Contents pane.

   You can browse only cluster names. If a managed system is a member node of a cluster, the message “The targeted system does not support this task” appears.

   The Clusters group displays the same information as Windows NT Clusters. However, it has been included as a placeholder for types of clusters other than those detected through MSCS. To support cluster implementations other than those detected through MSCS, you need to programmatically extend Director by using the guidelines in the Director SDK.
2. In the Group Contents pane, select the cluster for which you want to view information and drag it to the **Cluster Browser** task in the Tasks pane.

   The Cluster Browser window appears with the cluster you selected in the Clusters pane. The cluster appears as the root of a tree structure.

   - To view a cluster’s status and description, double-click a cluster name.
   - To view information on the resources assigned to the cluster, expand the properties tree.
3. To reverse the order of the cluster names, right-click on any line item and select **Sort** → **Descending**.
4. When you finish viewing information, select **File** → **Close** to close the window.

Chapter 21. Process management

Director enables you to manage individual processes on remote systems. The process management task enables you to start, stop, and monitor applications and processes. You can set up a monitor on a particular process or application so when that process or application changes state, an event is generated.

The process management task is an interactive task that applies only to native managed systems. SNMP devices do not have the capability to be monitored and managed to this level of detail.

The process management task enables you to:

- View information about processes running on a system
- Execute commands on a selected system
- Create a non-interactive task which can be scheduled
- Close applications running on a selected system
- Create and save monitors for applications and services
- Initiate a monitor for specific applications and services
- Start, stop, pause, and continue system services on Windows NT, Windows 2000 and Windows XP systems

Starting the Process management window

You can start the main Process management window from the Director Management Console using drag and drop and the context menu techniques (refer to "Navigating in Director" on page 80).

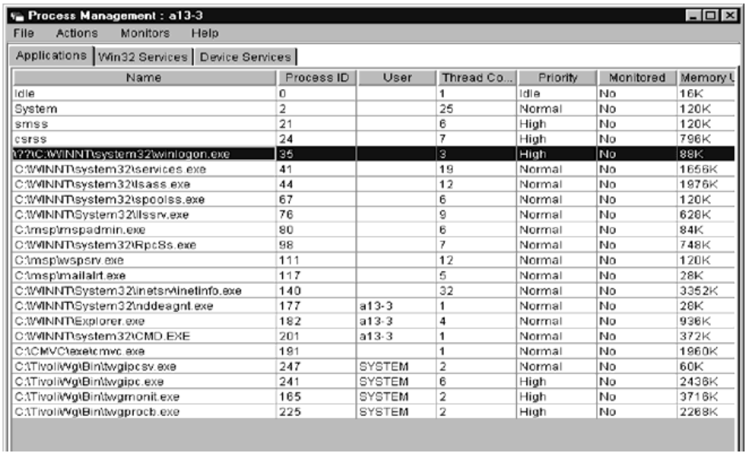

Name	Process ID	User	Thread Co...	Priority	Monitored	Memory U
Idle	0		1	Idle	No	16K
System	2		25	Normal	No	120K
smss	21		6	High	No	120K
csrss	24		7	High	No	796K
\??\C:\WINNT\system32\winlogon.exe	35		3	High	No	88K
C:\WINNT\system32\services.exe	41		19	Normal	No	1656K
C:\WINNT\system32\lsass.exe	44		12	Normal	No	1976K
C:\WINNT\system32\spoolss.exe	67		6	Normal	No	120K
C:\WINNT\System32\llssrv.exe	76		9	Normal	No	628K
C:\msp\mspadmin.exe	80		6	Normal	No	84K
C:\WINNT\system32\RpcSs.exe	98		7	Normal	No	748K
C:\msp\wspsrv.exe	111		12	Normal	No	120K
C:\msp\mailalrt.exe	117		5	Normal	No	28K
C:\WINNT\System32\inetsrv\inetinfo.exe	140		32	Normal	No	3352K
C:\WINNT\System32\nddeagnt.exe	177	a13-3	1	Normal	No	28K
C:\WINNT\Explorer.exe	182	a13-3	4	Normal	No	936K
C:\WINNT\system32\CMD.EXE	201	a13-3	1	Normal	No	372K
C:\CMVC\exe\cmvc.exe	191		1	Normal	No	1960K
C:\Tivoli\Wg\Bin\twgipcsv.exe	247	SYSTEM	2	Normal	No	60K
C:\Tivoli\Wg\Bin\twgipc.exe	241	SYSTEM	6	High	No	2436K
C:\Tivoli\Wg\Bin\twgmonit.exe	165	SYSTEM	2	High	No	3716K
C:\Tivoli\Wg\Bin\twgprocb.exe	225	SYSTEM	2	High	No	2268K

The Process management window of all operating systems contains an Applications tab. Windows NT, Windows 2000 and Windows XP contain two additional tabs: Win32 Services and Device Services.

Applications
Enables you to perform tasks on processes with which you can interact, such as program applications. Most process management tasks are performed on applications. You can add an application to the agent’s process monitors and configure its monitor to generate an event if the application stops or starts or fails to start. You can also close an application.

Win32 Services
(Windows NT, Windows 2000 and Windows XP only) Enables you to interact with Win32 services. You can start, stop, pause, and continue services, and you can also set monitors on services. See “Controlling Windows NT, Windows 2000, Windows XP system and device services” on page 200 for more information.

Device Services
(Windows NT, Windows 2000 and Windows XP only) Enables you to interact with Windows NT, Windows 2000 or Windows XP device services. Device services are the non-interactive programs that enable high-level applications to perform various functions. For example, the I/O drivers running on a system serve as support programs for application suites that perform word processing, database, and print functions. You can start and stop most driver services, as well as set monitors on device services. See “Controlling Windows NT, Windows 2000, Windows XP system and device services” on page 200 for more information.

Notes:

1. Not all services can be controlled in this manner.

2. You should exercise caution when starting or stopping Win32 and device services. Make sure you are familiar with the service and understand the impact of starting, stopping, pausing, and continuing these system applications.

Viewing application information

When you bring up the Process management window, the Applications tab is shown, with information about each application.

Every operating system uses a subset of the following:

Name Identifies the name of the application showing where the program resides on the system.

Process ID
Identifies the operating system's internal identification value for this process.

Command Line
Identifies the command that was used to launch this process.

Job Number
Identifies the 6-digit job number assigned to a job.

Parent Process ID
Identifies the operating system's internal identification value for the process or program that started this process.

User Identifies the logon ID of the user that started the process.

Type Describes the job type.

Session ID
Identifies the ID of the session under which the command is executing.

Description
Identifies the application with a short description.

Version Identifies the version number of the application.

Date Identifies the date of the application.

Thread Count
Identifies the number of program threads that this process is using.

Priority Identifies the relative importance of the process with regard to receiving attention from the processor.

Monitored
Identifies whether a process is being monitored. Note that this is not the same as the resource monitors that Chapter 8, “Resource Monitoring,” on page 109 discusses.

%CPU Identifies the percentage of total processor time used by an application.

Status Describes the 4-letter code representing the status of a job.

Memory Usage
Identifies the current memory usage, in KB, for the selected system.

Subsystem
Identifies the subsystem in which a job is running.

Viewing Windows services information

For Windows NT, Windows 2000, Windows XP, Win32 and device services, the following information is shown on the Win32 Services tab and the Device Services tab for each service:

Name
The name of the service.

Service Status
The current status of the service (stopped, paused, or running).

Executing commands on selected systems

You can use the Process management task to execute a command on a targeted managed system. You can do this right from the Process Management window. See the online help for details.

Creating non-interactive tasks to execute commands

You can use the process manager task to send individual commands to a selected system or group. You can send only one command at a time.

When the command executes, descriptive information is stored in the Scheduler task associated with the non-interactive task. This information might include the target system, command name, and completion status, and standard output and standard error information being executed.

See the online help for details on creating non-interactive tasks to execute commands.

Restricting anonymous command execution

By default, commands are executed on the target system as either administrator or root. There are provisions on Windows NT, Windows 2000 and Linux to disable this feature and always require a user ID and password to be specified.

Note: Keep in mind that this ability exists for current agents only. Earlier versions of IBM Director do not use this feature.

For systems running Windows NT, Windows 2000 and Windows XP

To enable or disable this feature for Windows NT, Windows 2000 or Windows XP modify the registry as follows:

1. From a command line, run **regedit**. The Registry Editor appears.

2. Navigate to the registry entry HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Director\CurrentVersion.
3. Double-click **RestricAnonCmdExe**.
4. In the Value data field, type one of the following values, depending on what you want to do:
   - To allow users to interact without an ID or password, type `0`.
   - To require users to use an ID and password, type 1.
5. Select **OK**. Your registry entry is saved.

For systems running Linux

To require users to enter ID and password on Linux systems, use the following procedure in a linux shell:

1. Change to the directory where the managed system is installed. By default, this is /opt/tivoliwg. To do this, type

   ```
   cd data
   ```

   then

   ```
   vi ProcMgr.properties
   ```

2. Change the line

   ```
   RestrictAnonCmdExec=false
   ```

   to

   ```
   RestrictAnonCmdExec=true
   ```

3. Save the file; changes take effect immediately.

Closing applications

You can use the process manager to close an application running on a target system. See the online help for details.

Note: Use this function with extreme caution; closing an application can cause a loss of data and halt the operating system. Also, note that not all applications can be closed this way.

Adding new process monitors

The Process management task enables you to create a process monitor that generates an event if a specified application starts, stops, or fails to start running within a specified amount of time after system startup or after the monitor is sent to an agent.

The Process Monitors window is used to create the process monitor. From the window, you can also edit and delete process monitor definitions. See the online help for all of the detailed procedures.

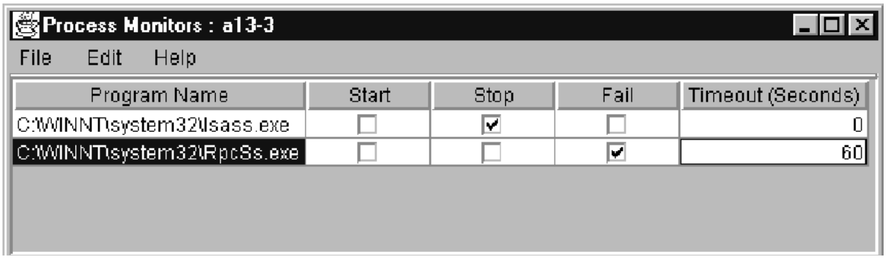

Program Name	Start	Stop	Fail	Timeout (Seconds)
C:\WINNT\system32\lsass.exe	☐	☑	☐	0
C:\WINNT\system32\RpcSs.exe	☐	☐	☑	60

When you activate a process event on a system, a monitor is started for the specified application. You might want to use this monitoring task to view statistics on the running application after you have applied an event-generation task. Refer to Chapter 8, “Resource Monitoring,” on page 109 for details on viewing your process monitor.

Controlling Windows NT, Windows 2000, Windows XP system and device services

You can use the Process Management task to start, stop, pause, and resume system services on Windows NT, Windows 2000 and Windows XP systems. See the online help for the procedures.

Removing process monitors

You can remove all process monitors defined for a given managed system by using the Remove Process Monitors subtask under the Process Management task in the Director Management Console. You can drag and drop this icon on the target managed system or group and the defined process monitors will be removed.

Adding service and device service monitors

You can monitor the status of services and device services. To do this, select the service or device you want to monitor, then right-click and select Add Threshold from the context menu. This will open the Resource Monitor Threshold dialog. You can then set alert levels for each possible state of the service or device. See Chapter 8, “Resource Monitoring,” on page 109 for more information on setting thresholds.

Chapter 22. Task Scheduler

The task scheduler feature of Director enables you to schedule sets of non-interactive tasks to occur at some time in the future. You can specify an exact date and time when you want the tasks to be started, and you can define tasks to repeat automatically at a given interval, such as "Every Saturday at 2:00 a.m.," "Every month on the 15th at midnight," and so on. You can also define a specific number of repeats, such as "Every Saturday at 2:00 a.m. for the next six weeks."

Only non-interactive tasks can be scheduled. A non-interactive task is a task that does not require interaction with the user. Most non-interactive tasks may be performed on multiple systems at once, such as software distribution and inventory. Still other non-interactive tasks are related to a single system or the Director server.

Interactive tasks require direct interaction with a user and cannot be scheduled. Examples of interactive tasks include remote control and file transfer.

Starting the Task Scheduler task

To schedule a task from the Director Management Console window, drag and drop between the task to be executed and a targeted managed system or group.

When you select a targeted non-interactive task to be performed, you need to specify whether you want the task to be performed immediately or if you want to schedule it to occur later:

1. To activate the job immediately, click **Execute Now**.
2. To set up a time and date for a job to activate, click **Schedule**.

   The New Scheduled Job dialog box prompts you for basic scheduling information:

   - Scheduled Job: Enter a title for the scheduled job. All scheduled jobs require a name.
   - Date: This is the date that you want the job to be executed. Click the calendar icon to the right of this field to display the calendar window.

     Use the arrows at the top and bottom to scroll the months and years and then click the desired date. The Date: field on the New Scheduled Job dialog is updated automatically.
   - Time: This is the time of day when you want to start the scheduled job. Enter the time in the field or use the pull-down menu to select a time in 15-minute increments.
3. Click **OK** to save your scheduled job.

4. Select **Advanced** to bring up a second New Scheduled Job window. This window enables you to customize your job by setting special job properties, such as generating events when the job completes, or specifying when the job will repeat.
5. You can also select **Cancel** to cancel the scheduled job creation or **Help** for online help information.

Customizing your scheduled job

The advanced New Scheduled Job window enables you to customize your scheduled job, enabling you to specify date and time, repeat intervals, the specific tasks to execute, the systems to which it is applied, and several other parameters.

Using the date and time page

This page enables you to:

- Specify a date and time for the scheduled job to be activated. If you already specified a date and time in the previous New Scheduled Job window (see "Starting the Task Scheduler task" on page 201) as part of the scheduler activation, those values are copied here. These fields operate identically to the ones described in the above reference.

  Note: Ensure that the Windows NT or the Windows 2000 server time matches the Director Management Console time; otherwise, the scheduled job will not propagate at the correct interval.
- You can also enable or disable the **Schedule the task to execute on a date and time** check box. If you leave this box unchecked, then a date and time is not assigned to the scheduled job. It is added to the job database with the other scheduled jobs, but is not activated automatically. You must activate it manually when you want to execute the job.
- Select the **Repeat** button to open the Repeat window, where you can create sophisticated schedules for re-executing your job.

Using the Repeat window

The Repeat window has several scheduling functions that, when combined, give you a powerful and flexible way to set up your repeating scheduled jobs.

The Repeats pane enables you to specify how often the job is repeated. Use the two drop-down lists to specify hourly, daily, weekly, monthly, or yearly intervals. Further define the repetition by specifying incremental hours, days, weeks, months, and so on. If you specify Custom in the first drop-down list, the Custom Dates pane becomes enabled. You can enter discrete dates to repeat the scheduled job, giving you complete flexibility.

The Duration pane enables you to enter a specific start and stop date and time. This action sets limits on how many times the job repeats or whether the job repeats forever. You can specify your own dates and times or use the pull-down calendar and clock panels to select the desired date and time by performing the following steps:

1. Specify a starting date and time and an ending date and time.
2. In the text box next to For, specify the interval in numbers of hours, days, weeks, months, or years.
3. If your scheduled job interval falls on a weekend, include special handling in the On weekends drop-down list. You can specify that the task be moved to Friday, Monday, or the nearest weekday; to not move it at all; or to delete the execution altogether if it falls on a weekend.
4. In the Your selection pane, the complete repeat interval is expressed in text so you can verify that it is what you intended.
5. When you finish, select **OK** to save your selection and return to the Date/Time page or select **Cancel** to close this window.

Using the Task page

The Task page enables you to select tasks from a list of all tasks that can be scheduled. Double-click a task to move it from the Available pane to the Selected Task pane. You can also highlight the desired task and then press the **Select** button.

You can select multiple tasks for a single job. Once you select the tasks and save the job, each task is processed in the order in which it appears on the Selected Tasks pane.

Using the Targets page

You can select the target systems or groups of systems from a list on the Targets page. The scheduled task is performed on each one of these systems, and the status of each is tracked during the execution of the scheduled job.

You have the options of using an entire group as the target for the scheduled job, or you can specify a list of managed systems as the targets. Find the two options below the pages:

- Use a group as the target

  This option enables you to select a group from a list of all groups. You can select only one group. If you select a second group, it replaces the first group.
- Specify a list of systems as targets

  This option enables you to select one or more systems from a list of all systems.

Using the Options page

Under the Options page, you can select **Special Execution Options** to address offline systems and systems that join the target group after the job has started executing. You can also limit the amount of status tracking and log information to only the most recent activity, and you can generate an event to occur upon the success or failure of a scheduled job or a particular system.

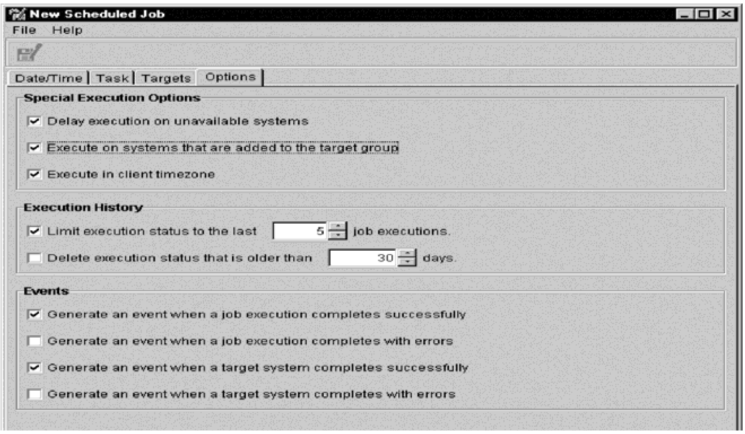

Special Execution Options

☑ Delay execution on unavailable systems

☑ Execute on systems that are added to the target group

☑ Execute in client timezone

Execution History

☑ Limit execution status to the last 5 job executions.

☐ Delete execution status that is older than 30 days.

Events

☑ Generate an event when a job execution completes successfully

☐ Generate an event when a job execution completes with errors

☑ Generate an event when a target system completes successfully

☐ Generate an event when a target system completes with errors

See the next section for more information on these special execution options.

Understanding the Special Execution options

It is important to understand the usefulness of the three Special Execution options:

- Delay execution on unavailable systems
- Execute on systems that are added to the target group
- Execute in agent time zone

Delay execution on unavailable systems

When You Do Not Check This Option: When this check box is not selected, only targeted systems that are online at the time of activation will have the task performed on them. Any targeted system that is offline at the time the task is activated is assigned a status of unavailable.

When all systems have been assigned a completed status or a failed status, the overall status for the execution of the job is changed to `complete` or `complete with errors`.

When You Check This Option: If this check box is selected, then when the scheduled task is activated, only targeted systems that are online at the time of activation will have the task performed on them.

Even after all online targeted systems are assigned a completed status or a failed status, the execution of the job will stay in the `in progress` state. It waits for the offline systems to come back online. When a system does come online, the task is activated on the systems that just came online.

When all of the targeted systems have been assigned a completed status or a failed status, the overall status for the execution of the job is changed to `complete` or `complete with errors`.

If this is a repeating job, and there are targeted systems that have still not been run (because they were offline) when the scheduled repeat time arrives, then the overall status of the job execution is changed to `incomplete`. A new execution of the job is activated.

Execute on systems that are added to the target group

When You Do Not Check This Option: When you do not select this option, the scheduled job is performed on all of the systems that are part of the target group at the time of activation. Any systems that join the group later do not have the scheduled job performed on them.

When this is a repeating scheduled job, any systems that have joined the target group since the last activation will then be included in the target group the next time the job is activated. Any systems that have left the target group since the last activation will not be included.

When You Check This Option: When you select this option, any new systems added to the target group are detected and the scheduled job is activated on the systems that have just been added. Checking this box will cause the execution of a one-time (non-repeating) job to stay active until you explicitly cancel it. Note that this option is selectable only if the target is a group of systems, not a list of specific systems, that you selected on the Targets page.

If this is a repeating scheduled job, the execution remains active and waits for new systems to be added, until the next repeat time is reached, and a new execution of the job is activated.

Execute in agent time zone

When You Do Not Check This Option: When you do not check this option, the scheduled job will execute on all selected targets when the server reaches the specified time and date.

When You Check This Option: When you check this option, tasks will execute according to the time zone in which the target system resides.

Notes:

1. You cannot create a job to repeat hourly and be executed in the time zone of the managed system.

2. One job activation record per 24 hours is created when the Execute in agent time zone option is selected. The job activation dynamically updates as managed-system agents move from pending to active when their time-zone window occurs.

3. If the first scheduled time zone start date occurs before the server date, the job cannot be created.

4. Job activations that are delayed because their target systems are in later time zones are classified as pending, much the same way jobs are classified until activated.

Saving your scheduled job

Saving your scheduled job is accomplished by either selecting **File → Save As** from the menu bar, or selecting the **Save As** icon from the tool bar. Specify a title for the scheduled job and then save it.

All scheduled jobs must have a title, but the titles do not have to be unique. For example, you may have two different jobs with the title of "test job."

Managing scheduled jobs

You can manage your scheduled jobs from the Director Management Console using the Scheduler task in the menu bar or the Scheduler icon in the tool bar. The Scheduler window appears with two pages, Calendar and Jobs.

You can use the Scheduler window menu bar to begin the scheduling of a new job. See "Customizing your scheduled job" on page 202 for details on using the New Scheduled Job window.

Using the Calendar pages

There are three Calendar pages. The Calendar pages shows when all jobs have been scheduled to execute, as well as status information for job executions. On the Month Calendar page, the current month appears in calendar format. Use the arrows at the top and bottom edges of the Calendar page to go to the desired month and year. The current week appears on the Week Calendar page, and the current day appears in calendar format on the Day Calendar page.

Note: The calendars are independent of each other. This means that changing the date on one calendar does not change the date on another calendar. Also, selecting a job on one calendar does not select it on other calendars.

You can begin the scheduling of a new job for a specific day by double-clicking the day in the calendar or selecting **New Job** from the day's context menu. See "Customizing your scheduled job" on page 202 for details on using the New Scheduled Job window.

Viewing job properties

To view the properties of a scheduled job, select the **Open Job Properties** menu bar option (or from the job's context menu).

The Scheduled Job window appears for the job, with four pages: Date/Time, Task, Targets, and Options. These pages have the same function as those in the New Scheduled Job window. See "Customizing your scheduled job" on page 202 for details.

The Scheduled Job window enables you to change the properties of the job and then save it as another scheduled job. IBM Director does not permit saving changes to an existing job; they must always be saved as a new job.

Viewing scheduled jobs

You can view the information about an execution of a scheduled job by selecting the **Open Execution History** menu bar option (or from the context menu of the execution history).

The Execution History window displays the overall status of the job. The top portion of the window displays a summary of the status for targeted systems. Targeted systems are also grouped together based on the status of each target for this execution and appear in the bottom portion of the window. For example, if five targeted systems completed the scheduled job successfully, then the top portion will have a count of 5 for Complete and the systems are listed together under Complete in the bottom portion.

In addition, a job can be executed again on selected groups and individual systems. To do this, select the system or group and then select **Execute Now** from the context menu. A selected job's execution history results can be exported to a CSV or HTML file as well. See the online help for more details.

Viewing execution history logs

You can view the entire log for an execution history by selecting **View → Log** from the menu bar or the context menu of the execution history.

You can also view only the log entries for an execution history that are related to a specific system by highlighting the system and selecting the **View System Log** option from the menu bar or the system's context menu.

When viewing either log, you can control the level of detail displayed by using the menu bar options. By default, the log displays the lowest level of detail.

See the online help for additional operations you can perform on your jobs.

Using the jobs page

Select the **Jobs** page to display a list of all scheduled jobs as well as status information for job executions. This information is displayed in a tree structure down the left side of the window.

Selecting a scheduled job causes information about the job to be displayed in the right side of the window. The information includes the number of executions that are active or complete, the next date that the job will execute, the tasks that the job will perform, and any options that have been specified for the job.

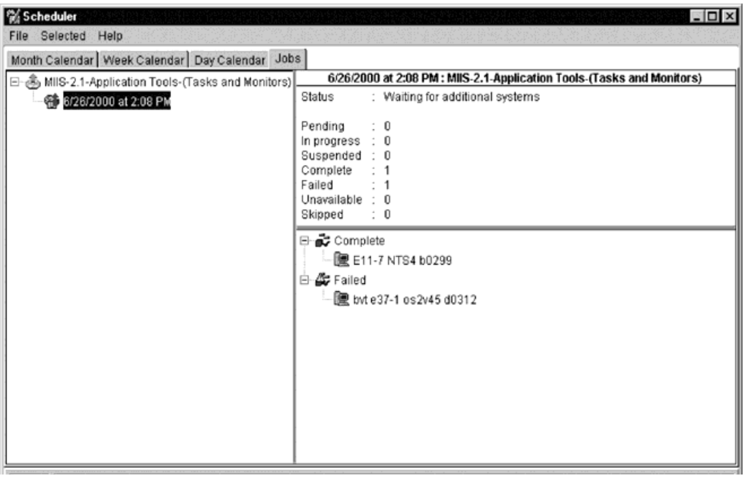

6/26/2000 at 2:08 PM : MIIS-2.1-Application Tools-(Tasks and Monitors)

```
Status      :  Waiting for additional systems

Pending     :  0
In progress :  0
Suspended   :  0
Complete    :  1
Failed      :  1
Unavailable :  0
Skipped     :  0
```

Selecting an execution of a scheduled job causes information about the job execution to be displayed in the right side of the window. This information is identical to the information displayed in the Execution History window. See “Viewing scheduled jobs” on page 207 for details.

See the online help for details on other menu bar and context menu options.

Viewing scheduled job information

The Execute Now button immediately starts a non-interactive task (see “Starting the Task Scheduler task” on page 201), and an Execution History window shows information about the execution. See “Viewing scheduled jobs” on page 207 for details.

The scheduler also maintains the execution history information from immediate executions. This information appears the same way as scheduled job execution history and it appears on the Calendar and Jobs pages for later reference.

Chapter 23. Hardware Status

The Hardware Status task makes it easy to identify systems that need attention. It provides an integration of status from Director Agent subsystems into an overall system health. This overall status is provided on the Director management console and the IBM Director Agent web-based GUI. The status icons categorizes hardware status into three groups: critical (red), warning (yellow) and informational (blue). Director Agent subsystems with a problem are identified and the appropriate subsystem tool can be launched to allow analysis.

Starting the Hardware Status task

To start the Hardware Status task, use one of the following methods:

- To view specific information about a managed system, double-click the status icon next to the managed system in the Group Content pane of the Director Management console.
- To view all managed systems within a particular status group, double-click the status icon at the bottom of the Director Management console.
- To view information about a specific managed system or information about all managed systems, right-click a status icon at the bottom of the Director Management console, then from the context menu, select Launch Hardware Status or select a managed system.

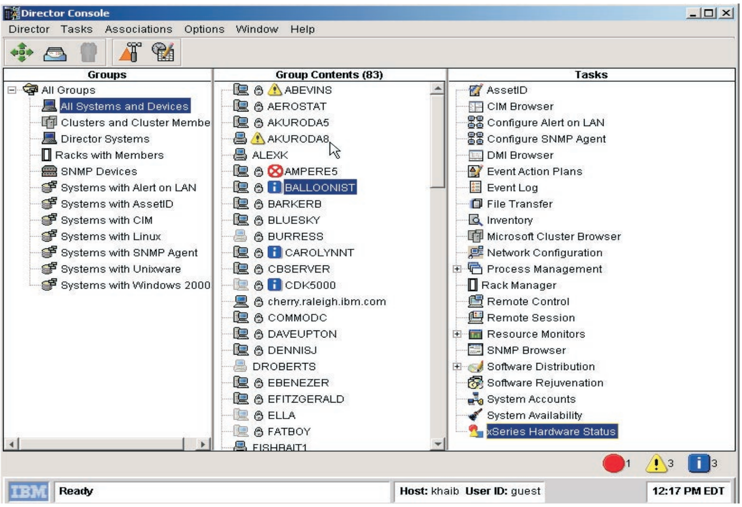

Note: The critical status icon blinks until disabled. To disable the icon, right-click and select **Disable Blinking**.

Viewing system-environment factors

Hardware Status automatically monitors Director Agent systems for changes in the following subsystem environment factors:

- Storage (Disk Space Low and SMART Drive)
- Network (LanLeash, Redundant Network Adapter)
- Environment (Fan, Power Supply, Temperature)
- Security (System Enclosure)

The Hardware Status window to displays these system-environment factors.

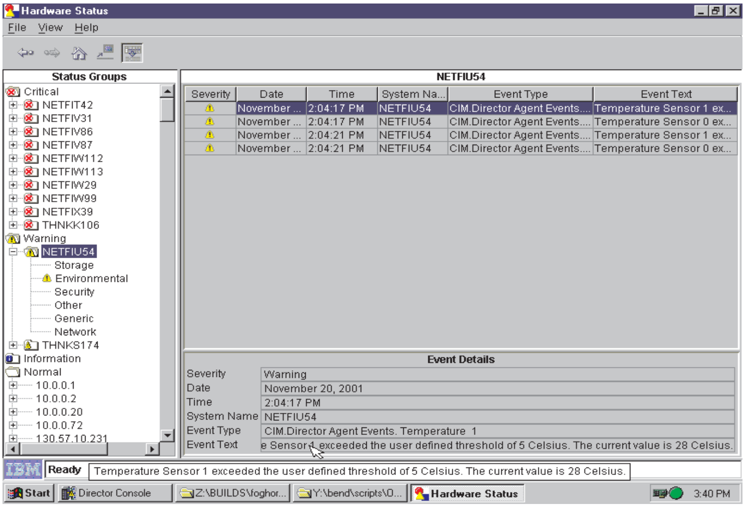

NETFIU54

Severity	Date	Time	System Na...	Event Type	Event Text
⚠	November ...	2:04:17 PM	NETFIU54	CIM.Director Agent Events....	Temperature Sensor 1 ex...
⚠	November ...	2:04:17 PM	NETFIU54	CIM.Director Agent Events....	Temperature Sensor 0 ex...
⚠	November ...	2:04:21 PM	NETFIU54	CIM.Director Agent Events....	Temperature Sensor 1 ex...
⚠	November ...	2:04:21 PM	NETFIU54	CIM.Director Agent Events....	Temperature Sensor 0 ex...

Event Details

Severity	Warning
Date	November 20, 2001
Time	2:04:17 PM
System Name	NETFIU54
Event Type	CIM.Director Agent Events. Temperature 1
Event Text	e Sensor 1 exceeded the user defined threshold of 5 Celsius. The current value is 28 Celsius.

Temperature Sensor 1 exceeded the user defined threshold of 5 Celsius. The current value is 28 Celsius.

The Hardware Status window is divided into three parts: Status Groups pane, Results Area and the Event Details section.

All managed systems are listed under the Status Groups pane. Status Groups appear in an expanded tree where managed systems are placed in one of four status severity categories (critical, warning, informational or normal). Highlight a managed system in the Status Group pane. The *Result Area* or right pane gives you detailed information about the condition, date, time, components, and events for the selected managed system. The title name, at the top of the Result Area, is the name of the selected managed system. For specific information, select a managed system in the Status group. Details about that system also appear in the Event Details area. To view the status of several managed systems simultaneously, use the Detach view. Select a managed system and click the **Detach** button in the menu bar, or click **View→ Detach Window.** Continue this process to detach as many Result areas as desired.

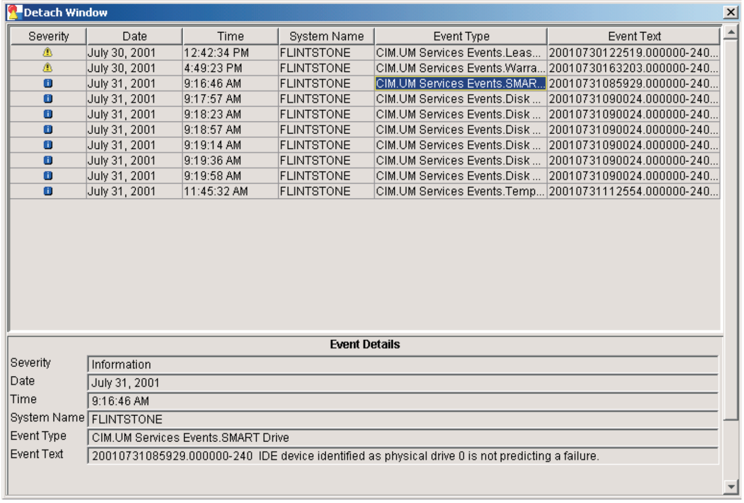

Detach Window

Severity	Date	Time	System Name	Event Type	Event Text
⚠	July 30, 2001	12:42:34 PM	FLINTSTONE	CIM.UM Services Events.Leas...	20010730122519.000000-240...
⚠	July 30, 2001	4:49:23 PM	FLINTSTONE	CIM.UM Services Events.Warra...	20010730163203.000000-240...
ℹ	July 31, 2001	9:16:46 AM	FLINTSTONE	CIM.UM Services Events.SMAR...	20010731085929.000000-240...
ℹ	July 31, 2001	9:17:57 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	9:18:23 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	9:18:57 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	9:19:14 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	9:19:36 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	9:19:58 AM	FLINTSTONE	CIM.UM Services Events.Disk ...	20010731090024.000000-240...
ℹ	July 31, 2001	11:45:32 AM	FLINTSTONE	CIM.UM Services Events.Temp...	20010731112554.000000-240...

Event Details

Severity	Information
Date	July 31, 2001
Time	9:16:46 AM
System Name	FLINTSTONE
Event Type	CIM.UM Services Events.SMART Drive
Event Text	20010731085929.000000-240 IDE device identified as physical drive 0 is not predicting a failure.

Note: All text fields support mouse over.

Working with status icons

To notify you of changes in environmental factors, Hardware Status icons are displayed beside managed systems. The severity of the situation is indicated and specific information describing the failure is available. The following three status states are reported: critical, informational and warning.

Icon	Name	Description
⊗	Critical icon	Indicates a significant problem that you should investigate.
ℹ	Informational/ harmless icon	A routine event designed to provide information.
⚠	Warning icon	A moderately significant event you should consider investigating.

Your managed systems status icons can be ignored or cleared. Right-click the status icon of a server in the Status Group pane. The context menu opens. Select **Ignore** Events and the server status icons disappear from the Status Group pane. However, the status icon of the server remains in the Result area. To make the icon reappear in the Status Group pane, right-click and click **Enable**. To clear the status icon in both the Status Group panel and in the Result area, click **Clear** and all events are cleared.

Chapter 24. Cluster Systems Management

The IBM Cluster tools are installed during the IBM Director Extensions installation and are seamlessly integrated into the Director Management 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 Director Management Console as a user with general access, group access, and task access privileges before you can use IBM Cluster Tools.

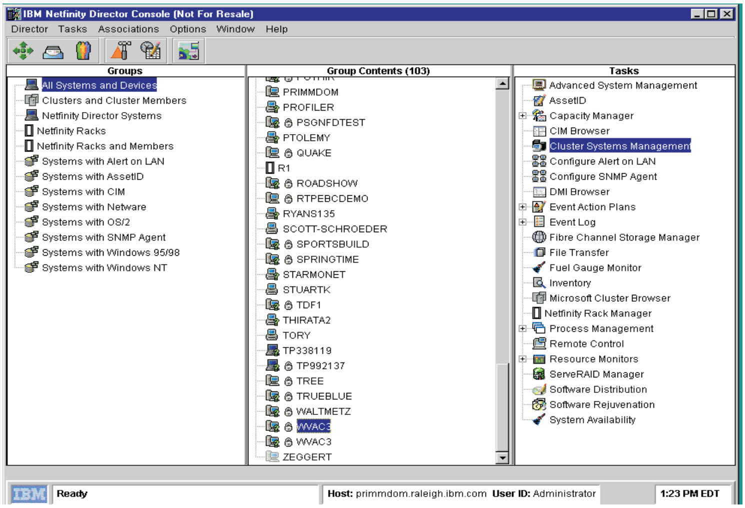

Cluster Systems Management is a graphical user interface (GUI) agent program that you can use to initiate cluster-related operations and manage cluster resources in a Microsoft Cluster Service (MSCS) based cluster. The program displays all cluster components, including nodes, groups, resources, networks, 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 agent/server applications continuous access to the storage device.)

Resource
A resource offers a service to agents in a agent/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 agent/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 agents 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 the Cluster Systems Management window

To start Cluster Systems Management from the Director Management console, 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 Director Management Console, and then click **Cluster Systems Management** from the menu.

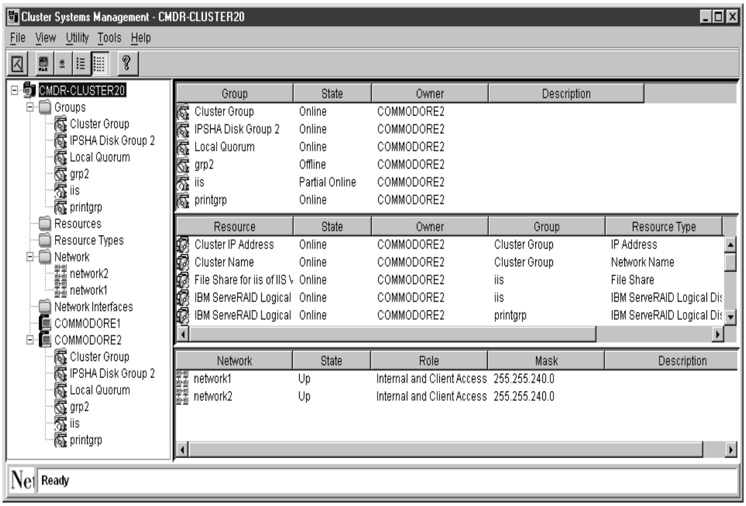

Group	State	Owner	Description
Cluster Group	Online	COMMODORE2	
IPSHA Disk Group 2	Online	COMMODORE2	
Local Quorum	Online	COMMODORE2	
grp2	Offline	COMMODORE2	
iis	Partial Online	COMMODORE2	
printgrp	Online	COMMODORE2	

Resource	State	Owner	Group	Resource Type
Cluster IP Address	Online	COMMODORE2	Cluster Group	IP Address
Cluster Name	Online	COMMODORE2	Cluster Group	Network Name
File Share for iis of IIS \	Online	COMMODORE2	iis	File Share
IBM ServeRAID Logical	Online	COMMODORE2	iis	IBM ServeRAID Logical Dis
IBM ServeRAID Logical	Online	COMMODORE2	printgrp	IBM ServeRAID Logical Dis

Network	State	Role	Mask	Description
network1	Up	Internal and Client Access	255.255.240.0	
network2	Up	Internal and Client Access	255.255.240.0	

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.

Command	Use this command to:
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.
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 has 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.
	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.

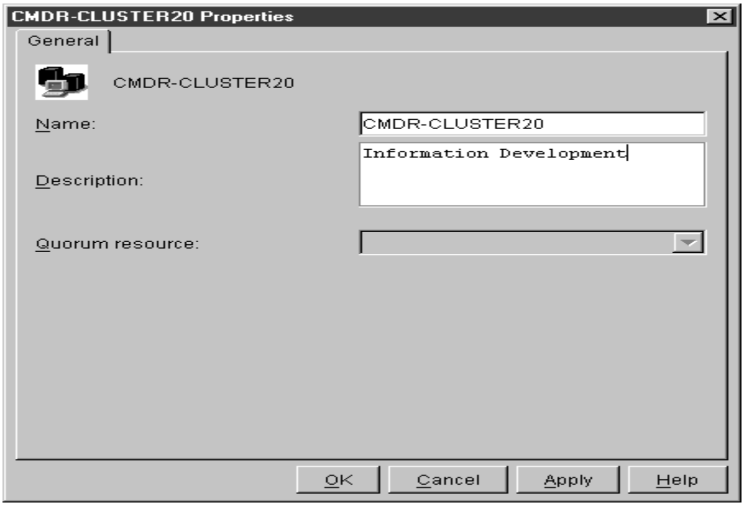
CMDR-CLUSTER20 Properties

CMDR-CLUSTER20

Name: CMDR-CLUSTER20

Description: Information Development

Quorum resource:

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 the General properties dialog box, 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**.

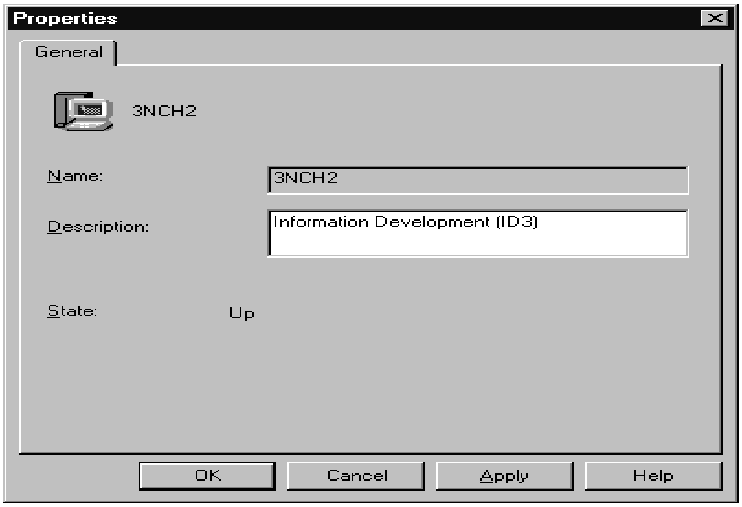

3NCH2

Name: 3NCH2

Description: Information Development (ID3)

State: Up

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, use the following procedure:

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

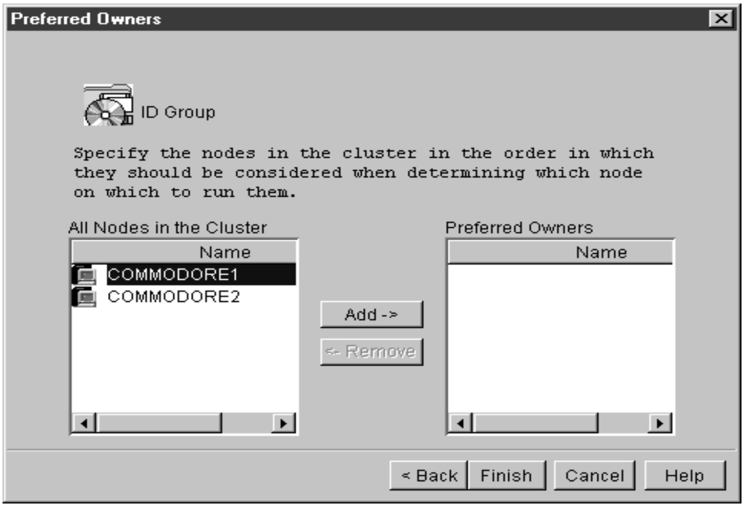

Preferred Owners

ID Group

Specify the nodes in the cluster in the order in which they should be considered when determining which node on which to run them.

All Nodes in the Cluster

Name
COMMODORE1
COMMODORE2

Preferred Owners

Name

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.

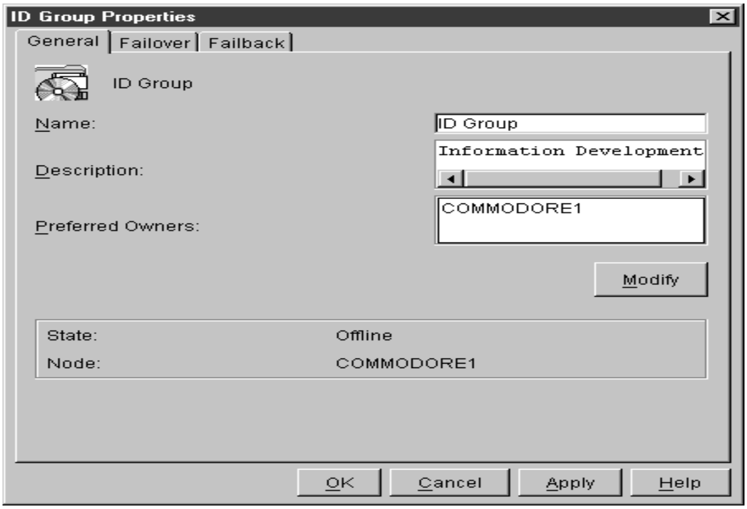

Name: ID Group

Description: Information Development

Preferred Owners: COMMODORE1

State: Offline

Node: COMMODORE1

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.

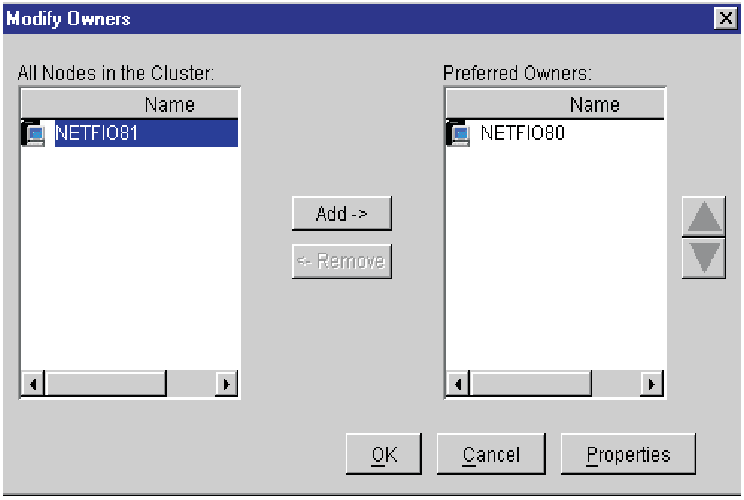

4. In the Modify Owners window, do the following:
   a. From the **All Nodes in the Cluster** drop-down list, select the name of the node that you want to be the preferred owner.
   b. Click the right arrow button to add the name to the **Preferred owners** list and to specify the preferred owner in 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.

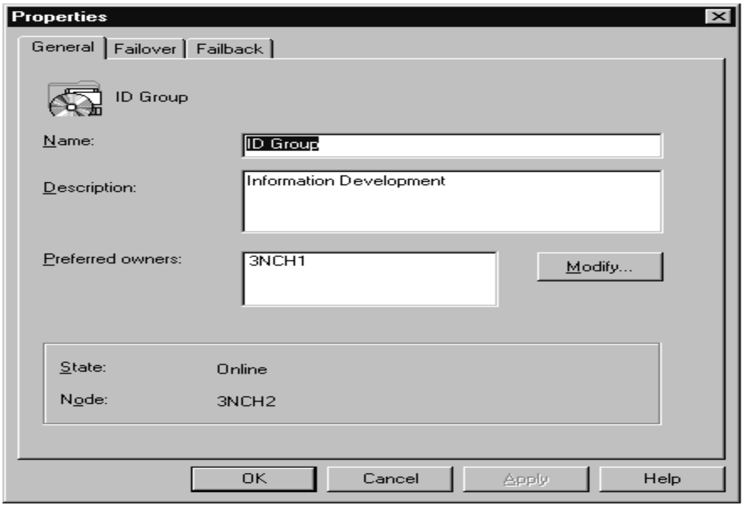

Name: ID Group

Description: Information Development

Preferred owners: 3NCH1

State: Online

Node: 3NCH2

3. Click the **Failover** tab.

   The **Failover** page window opens.

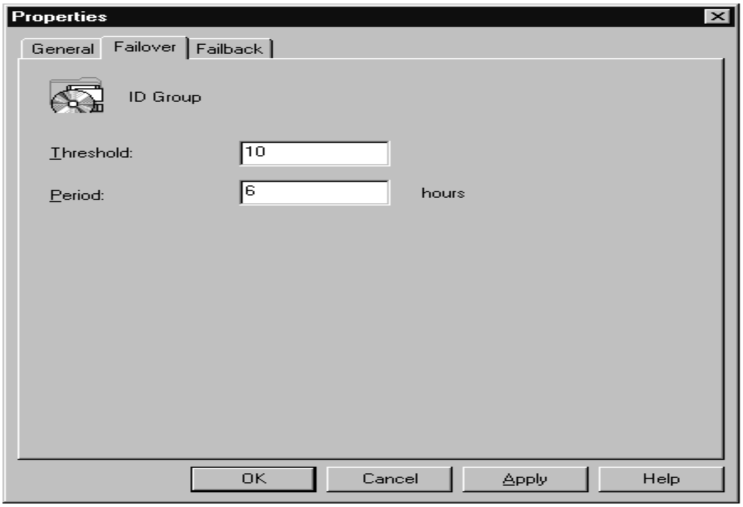

Threshold: 10

Period: 6 hours

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**.

Changing 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.

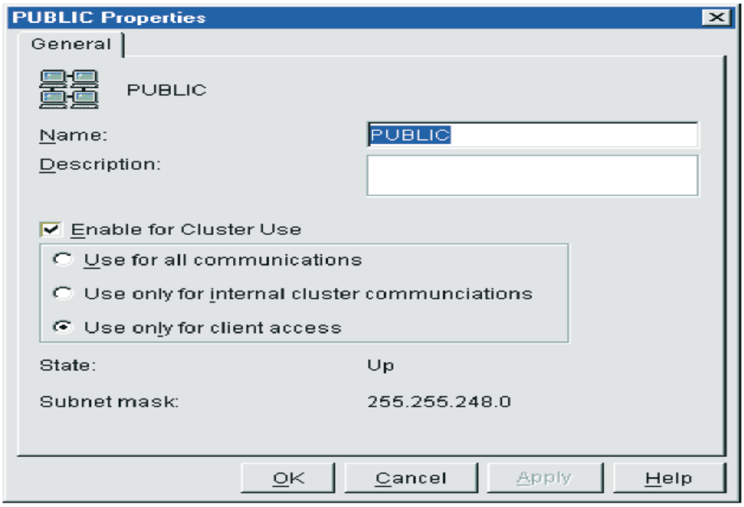

PUBLIC Properties

General

PUBLIC

Name: PUBLIC

Description:

☑ Enable for Cluster Use

○ Use for all communications

○ Use only for internal cluster communciations

● Use only for client access

State: Up

Subnet mask: 255.255.248.0

3. Select the **Enable for Cluster Use** check box.
4. Select one of the option buttons to specify cluster use in the network.
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.

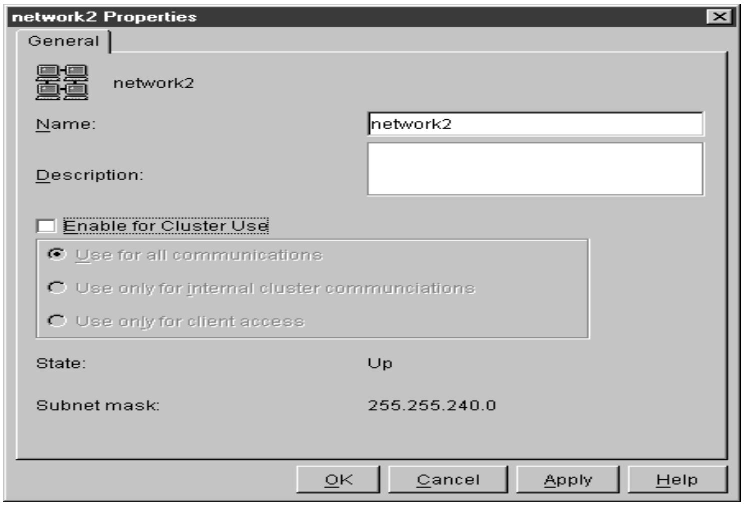

network2

Name: network2

Description:

☐ Enable for Cluster Use

● Use for all communications

○ Use only for internal cluster communciations

○ Use only for client access

State: Up

Subnet mask: 255.255.240.0

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 agents 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.

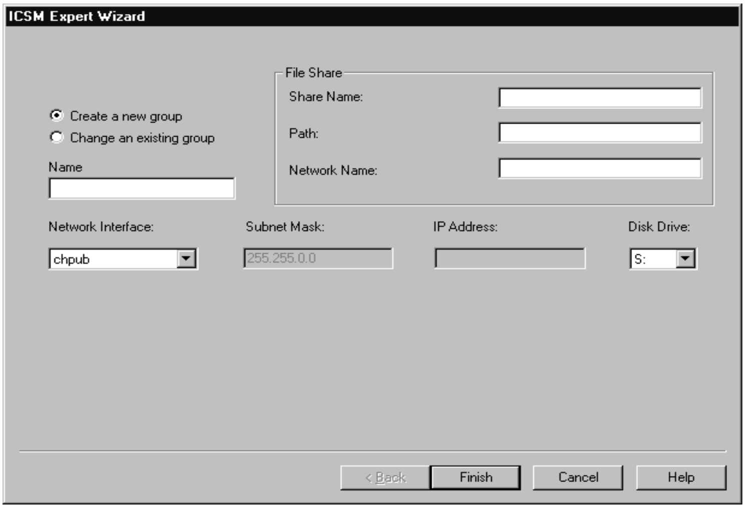

ICSM Expert Wizard

● Create a new group
○ Change an existing group

Network Interface: chpub

Subnet Mask: 255.255.0.0

Disk Drive: S:

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.

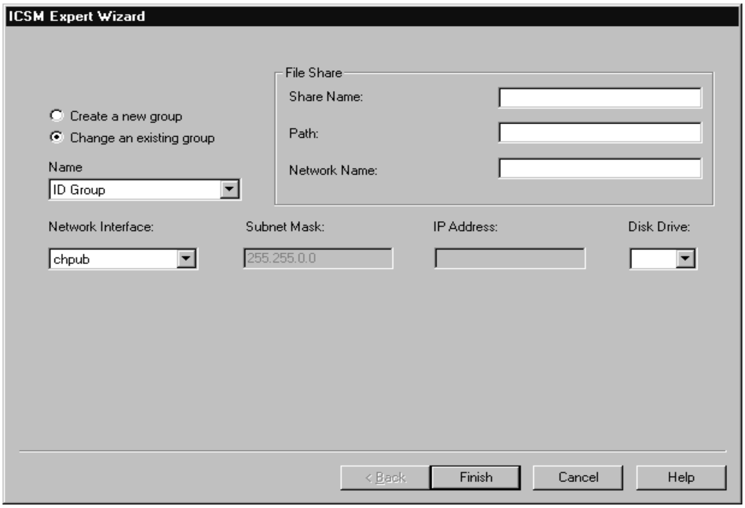
ICSM Expert Wizard

○ Create a new group
● Change an existing group

Name: ID Group

File Share
Share Name:
Path:
Network Name:

Network Interface: chpub
Subnet Mask: 255.255.0.0
IP Address:
Disk Drive:

2. Select **Change an existing group** option button.
3. Select the group name from the **Name** list.
4. Select the physical disk drive letter from the **Disk Drive** 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.

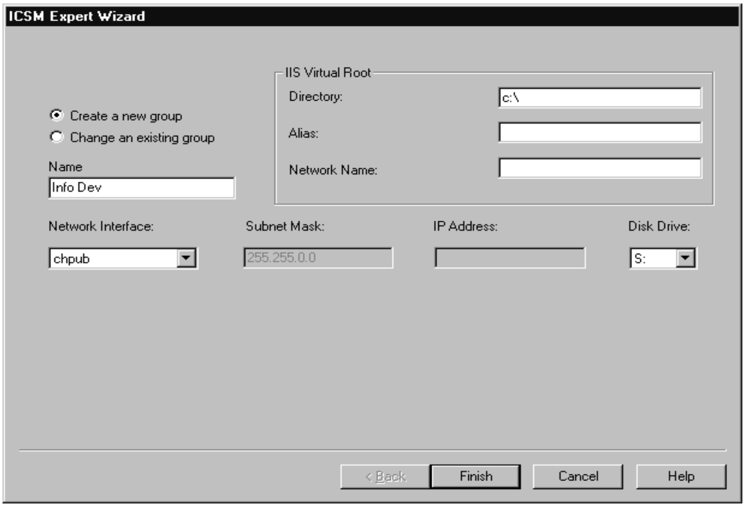

● Create a new group
○ Change an existing group

Name: Info Dev

Directory: c:\

Alias:

Network Name:

Network Interface: chpub

Subnet Mask: 255.255.0.0

IP Address:

Disk Drive: S:

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** list.
5. Select the physical disk drive letter from the **Disk Drive** 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.

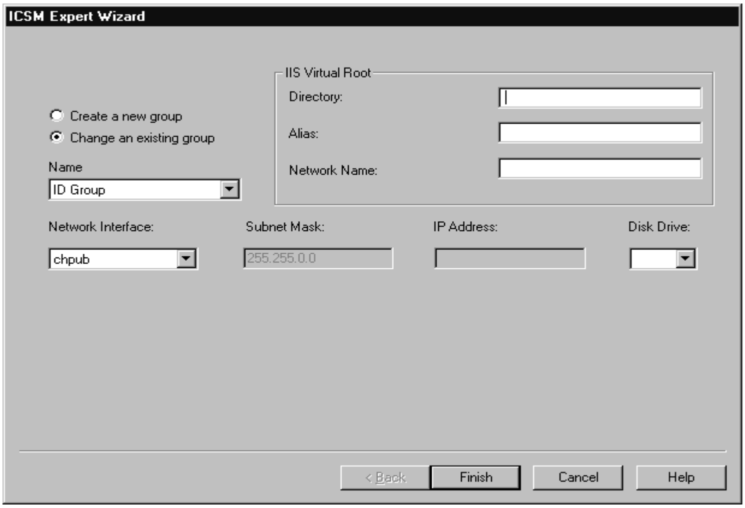

○ Create a new group

● Change an existing group

Name: ID Group

Network Interface: chpub

Subnet Mask: 255.255.0.0

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 list.
5. Select the physical disk drive letter from the **Disk Drive** 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.

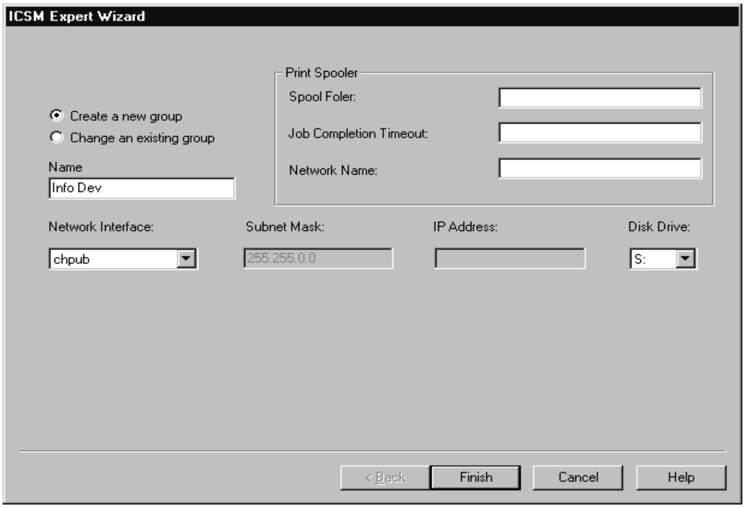

ICSM Expert Wizard

● Create a new group
○ Change an existing group

Name: Info Dev

Print Spooler
Spool Foler:
Job Completion Timeout:
Network Name:

Network Interface: chpub

Subnet Mask: 255.255.0.0

IP Address:

Disk Drive: S:

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 list.
5. Select the physical **Disk Drive** letter from the 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.

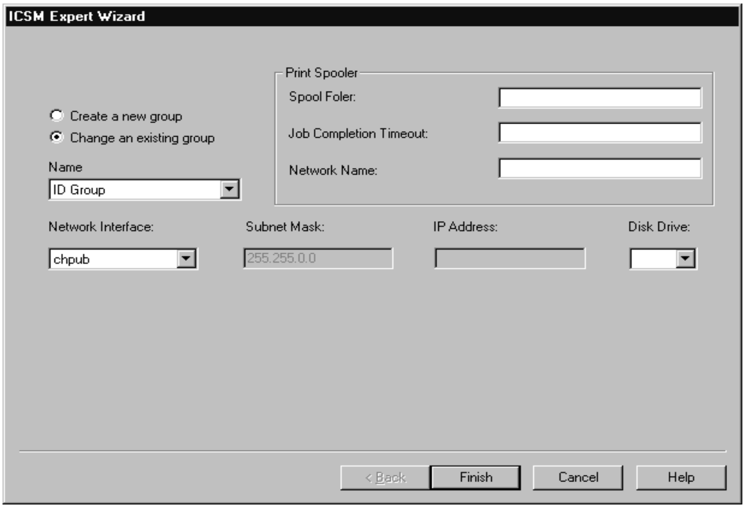
ICSM Expert Wizard

○ Create a new group
● Change an existing group

Name: ID Group

Print Spooler
Spool Foler:
Job Completion Timeout:
Network Name:

Network Interface: chpub
Subnet Mask: 255.255.0.0
IP Address:
Disk Drive:

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 list.
5. Select the physical **Disk Drive** from the 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.

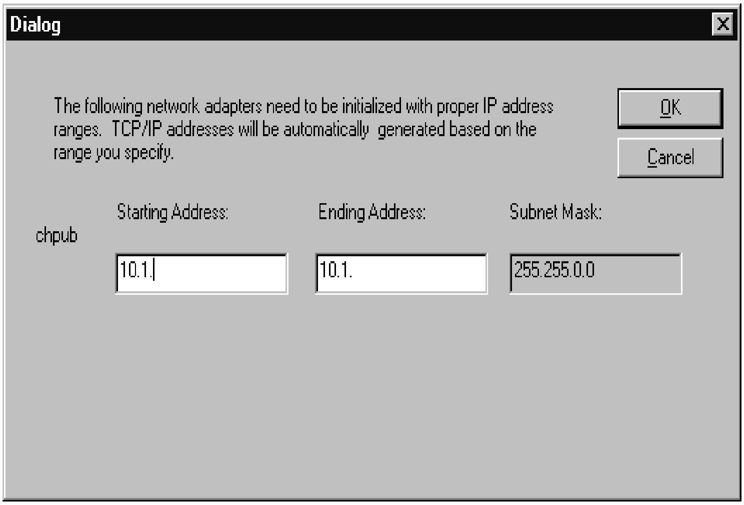

Dialog

The following network adapters need to be initialized with proper IP address ranges. TCP/IP addresses will be automatically generated based on the range you specify.

chpub

Starting Address: 10.1.

Ending Address: 10.1.

Subnet Mask: 255.255.0.0

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**.

Chapter 25. Management Processor Assistant

Use the Management Processor Assistant service for 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 Management Processor Assistant 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 Management Processor Assistant service can be configured to automatically forward an event in one of five ways:

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

With this service, you can manage the Management Processor hardware that is installed in your own system, or you can use Management Processor Assistant to connect with the management processor that is installed in a remote server. You can connect with remote Management Processor Assistant 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 Management Processor Assistant PCI adapter from your Director Console system.
- If your Advanced System Management PCI adapter or processor is connected to an Management Processor Assistant Interconnect network, or if the remote Management Processor Assistant PCI adapter or processor to which you have connected using a TCP/IP connection is connected to an Management Processor Assistant Interconnect network, you can use this connection to access and manage the Management Processor Assistant PCI adapter or processor of any other system that is connected to the Management Processor Assistant Interconnect network.
- If serial connection is established to the Advanced System Management PCI adapter or processor using modems or a null cable from the Director Console system.

In addition, with Management Processor Assistant, 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 Management Processor Assistant 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.**

Note: When managing a system on a RS485 network, you can use a Remote Supervisor Adapter, an Advanced System Management PCI Adapter or an Advance System Management Processor. However, when managing a system locally, you must use an Advance System Management PCI Adapter.

Starting the Management Processor Assistant task

There are two ways to initiate the Management Processor Assistant service: 1)through the Management Processor Assistant integrated Console or,2) through the Management Assistant Processor Assistant Web Based Management task.

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

Note: This starts Management Processor Assistant on this system and will enable you to configure and manage the Management Processor Assistant PCI adapter or Management Processor Assistant processor. To access, configure, and manage an Management Processor Assistant PCI adapter or processor in a remote system, you must first use Management Processor Assistant to establish a connection with the remote system.

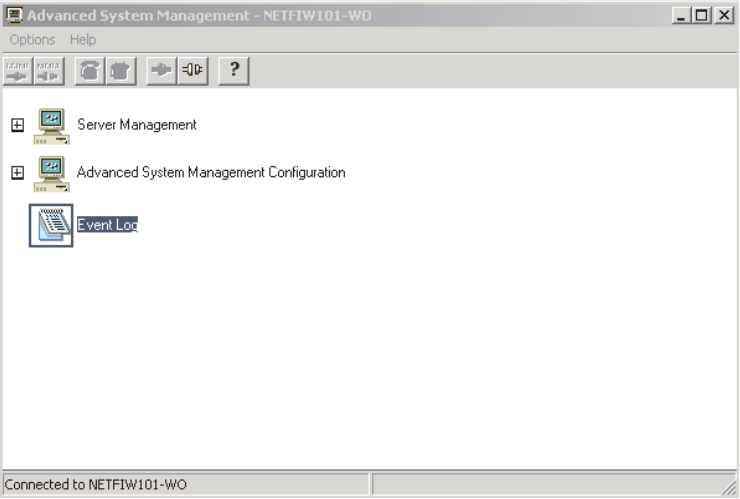
Advanced System Management - NETFIW101-WO
Options Help
?
Server Management
Advanced System Management Configuration
Event Log
Connected to NETFIW101-WO

Using the Management Processor Assistant window

Double-click any of the following selections available in the Management Processor Assistant window to access the function or the configuration information that you need.

Server Management

- **Operational Parameters** -double-click 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 Management Processor Assistant adapter. For more information, see “Operational parameters” on page 241.
  - Temperatures
  - Voltages
  - System Status
- **Configuration Information**
  - System Vital Product Data
  - System Card Vita Product Data
  - DIMM Product Data (on some systems)
- **System Power Control** - Double -click System Power Control to instruct the Management Processor Assistant adapter to power-off the system, restart the system, or power-on the system. For more information, see “System Power Control” on page 242.
- **Remote POST Console**- 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 Management Processor Assistant adapter. For more information, see “Remote POST console” on page 243.

Management Processor Assistant Configuration (Management Processor Assistant Processor Configuration)

- **System Management Processor**
- **Configuration Settings** - Double-click to expand the Configuration Settings tree and then double-click a component to configure the Management Processor Assistant adapter. 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, and Remote Alert Settings. For more information, see “Configuration settings” on page 245.
- **Network Protocols**(Network Settings, SNMP Settings, PPP Settings
- **Restart Management Processor** - Double-click this option to restart the Management Processor. Some changes in the configuration of the Management Processor require a restart before a change takes effect. The following configuration groups require a restart: Network Settings, SMNP Settings, PPP Configuration, and DNS Configuration.

Event Log - Double-click to view the contents of the Management Processor Assistant 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 271.

Management Processor Assistant Web-based management

The Management Processor Assistant Web-based task starts 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 started 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 prompted to supply new information.

To start the Management Processor Assistant service using the Management Processor Assistant Web Based Management task, use the following procedure:

1. Drag the Management Processor Assistant Web Based Management icon from the Tasks pane of the Director Console and drop it on a system that supports Management Processor Assistant in the Group Contents pane. The Management Processor Assistant Web Based Management window opens.
2. Type the IP Address or Host name, and click **OK**.

Updating microcode

Use this option to update the server POST/BIOS firmware and that of the Management Processor Assistant processor.

- To update the microcode on your Management Processor Assistant PCI adapter or processor, click **Options→ Update Microcode →System Management**. For more information, see "Updating PCI adapter or processor microcode" on page 273.
- To update the system POST/BIOS microcode on a system that includes an Management Processor Assistant PCI adapter or processor, click **Options →Update Microcode→ System POST/BIOS.** For more information, see "Updating system POST/BIOS microcode" on page 274.

Server Management

The following options enable you to manage your server.

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.

  Note: Monitored system components vary by Management Processor adapter or processor.
- Voltages

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

  VRM voltages are monitored but not displayed.
- System Status

  Includes system state, including operating system started, operating system running, POST started, POST stopped (error detected), 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) such as total system power, startup count, system state (starting POST, Operating System activity, fan status, microprocessors, system board, and hard disk drive backplane.

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 Management Processor Assistant subsystems:

- **System Vital Product Data**

  Provides information about the system such as 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, about all cards including microprocessor memory and power supplies. The information includes 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, microprocessor units (CPU), memory cards, power supplies, power backplane, front panel, I/O backplane, I/O adapter, DASD backplane, and system management subsystem.

- **Memory DIMM Information**

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

System Power Control

From the Management Processor Assistant console, click the + beside Configuration Settings or double-click **Server Management** to expand the Configuration Settings tree and view the Configuration Settings components. Double-click **System Power Control** to open the System Power Control window. Use the System Power Control window to enable the Management Processor Assistant PCI adapter or processor to turn off, restart, or turn on the system. 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.

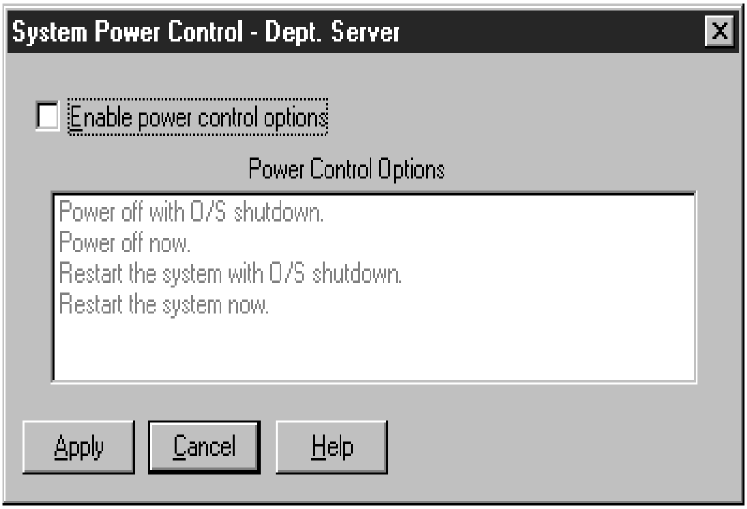
System Power Control - Dept. Server

□ Enable power control options

Power Control Options

Power off with O/S shutdown.
Power off now.
Restart the system with O/S shutdown.
Restart the system now.

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

If directly connected to the Advanced System Management PCI adapter through a TCP/IP, serial or RS-485 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

Note: You must be connected through a TCP/IP, or serial link.

You can use the Management Processor Assistant 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 Management Processor Assistant PCI adapter or processor.
2. Open the Remote POST window.
3. Restart the remote system by using the Management Processor Assistant 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 Management Processor Assistant PCI adapter or processor.

Management Processor Assistant Configuration

The Management Processor Assistant configuration option allows you to manage the service processor card.

System management processor information

The system management processor information option provides information about the Management Processor Assistant PCI adapter or processor, including Management Processor Assistant processor microcode build ID, revision number, file name, and date, device driver version number and Management Processor Assistant processor hardware revision number.

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

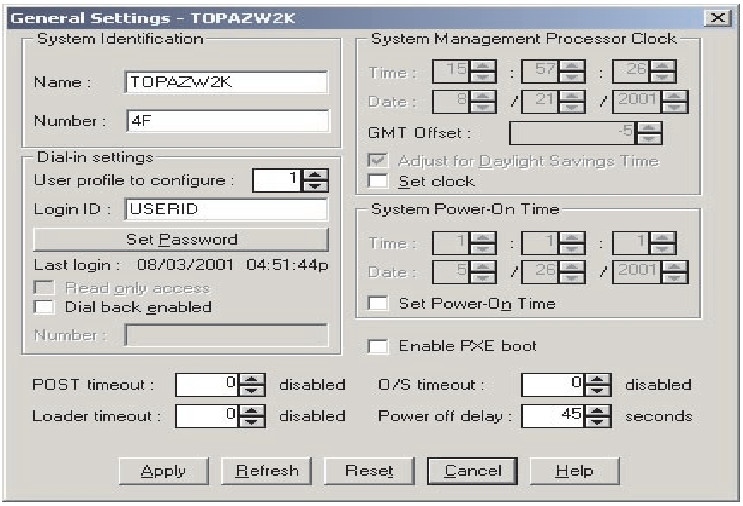

General Settings - TOPAZW2K

System Identification

Name : TOPAZW2K

Number : 4F

Dial-in settings

User profile to configure : 1

Login ID : USERID

Set Password

Last login : 08/03/2001 04:51:44p

□ Read only access

□ Dial back enabled

Number :

System Management Processor Clock

Time : 15 : 57 : 26

Date : 8 / 21 / 2001

GMT Offset : -5

☑ Adjust for Daylight Savings Time

□ Set clock

System Power-On Time

Time : 1 : 1 : 1

Date : 5 / 26 / 2001

□ Set Power-On Time

□ Enable PXE boot

POST timeout : 0 disabled

Loader timeout : 0 disabled

O/S timeout : 0 disabled

Power off delay : 45 seconds

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 time-out
- Loader time-out

- O/S time-out
- Power off delay

The System Identification group: The System Identification group box contains two fields to help you identify the system that contains the Management Processor Assistant 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 Management Processor Assistant 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.
Login ID	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.) **Note:** To remotely access the Management Processor Assistant PCI adapter specify a login ID.

Item	Description
Set Password	A password must be provided along with the login ID to allow a remote user to access the Management Processor Assistant 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.) **Note:** This password must be 5 to 8 characters in length and must contain a non-alphabetic character.
Last Login	This shows the date and time of the last successful login by a remote user.
Read only access	If the Read Only Access check box is selected, the user whose profile is selected will not be able to alter any of the Management Processor Assistant 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.
Dial Back Enable	If the check box is selected, the Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant PCI adapter or processor.

Note: The Management Processor Assistant 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 Management Processor Assistant 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 time-out*:* The POST time-out text box shows the number of seconds that the Management Processor Assistant PCI adapter or processor will wait for the system power-on self-test (POST) to complete before generating a POST time-out event. If POST takes longer than the configured amount of time to complete and the POST Time-out check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the Management Processor Assistant 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 Time-out is automatically disabled until the system is properly shut down and restarted.

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

To set the POST time-out value, use the drop-down list beside the POST Time-out text box to set the number of seconds that the IBM Management Processor Assistant PCI adapter or processor will wait for POST to complete. Then, click **Apply** to save this value. The maximum POST time-out value that you can set is 7650 seconds. Set this value to 0 to disable POST time-out detection.

Loader time-out*:* The Loader time-out field displays the number of seconds that the Management Processor Assistant PCI adapter or processor will wait for the system loading process to complete before generating a Loader time-out event. The Loader time-out 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 time-out 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 time-out 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 time-out check box and the system detects a Loader time-out.

To set the Loader time-out value, use the Loader time-out 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 time-out event. Then, click **Apply** to save this value. The maximum Loader time-out value that can be set is 7650 seconds. Set this value to 0 to disable Loader time-out detection.

O/S time-out*:* A periodic signal is sent from the Management Processor Assistant PCI adapter or processor to the operating system to confirm that the operating system is running properly. The operating-system time-out event occurs when the operating system does not respond to the signal within 6 seconds. The O/S Time-out text box shows the number of seconds that the Management Processor Assistant PCI adapter or processor will wait between operating system time-out checks. If the operating system fails to respond within 6 seconds, the Management Processor Assistant PCI adapter or processor will attempt to restart the system, and if the O/S Time out check box (in the Enabled Events group box of the Remote Event Settings window) is selected, the IBM Management Processor Assistant 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 time-out check box and the system detects an O/S time-out, the system will restart, but no event will be forwarded.

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

Power off delay*:* The Power off delay text box displays the number of seconds that the IBM Management Processor Assistant PCI adapter or processor will wait for the operating-system shutdown process to complete before turning off the system.

When the Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant 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 settings back to their default values, including configuration settings, dial-out settings, and advanced dial-out settings. **Note:** All previously configured Management Processor Assistant settings will be permanently lost.
Cancel	Click **Cancel** to close this window without saving any changes.
Apply	Click **Apply** to save changes.

Modem settings

From the Management Processor Assistant 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

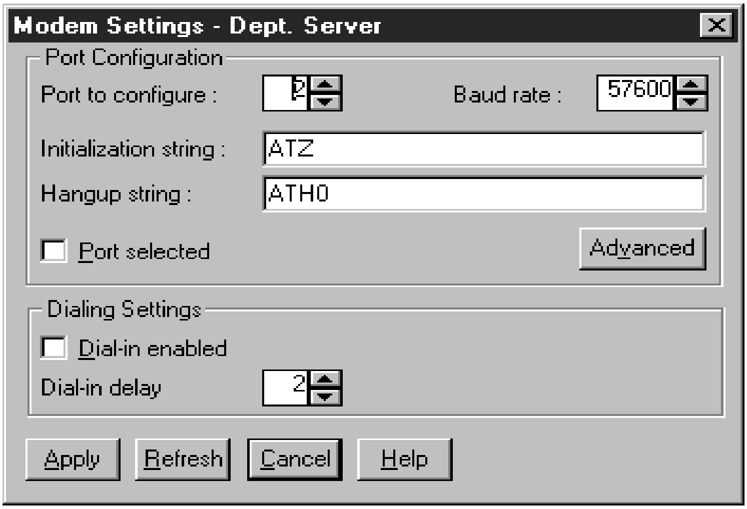

Modem Settings - Dept. Server

Port Configuration

Port to configure : 2

Baud rate : 57600

Initialization string : ATZ

Hangup string : ATH0

□ Port selected

Dialing Settings

□ Dial-in enabled

Dial-in delay 2

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 Management Processor Assistant event occurs. The Port Configuration group box contains the following items.

Item	Description
Port to configure	Use the scroll 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 Management Processor Assistant PCI adapter or processor. The port that you select to use affects the availability of the modem for use by either the Management Processor Assistant PCI adapter or processor or the operating system. The ports that are available to the system or to the Management Processor Assistant service vary depending on your hardware configuration.
Baud Rate	Use the scroll 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 (ATZ) 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 (ATH0) is provided. Do not change this string unless your dial-out functions are not working properly.

Item	Description
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 Management Processor Assistant PCI adapter or processor. Select this check box if you want to configure the Management Processor Assistant 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 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 is used before the number to be dialed. The default is ATDT.
Dial postfix string	Type the initialization string that STRING 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 is used to tell modem to answer the phone when it rings. The default is to answer after two rings or ATS0=2.
Autoanswer stop	Type the initialization string that STOP is used to tell the modem to stop answering the phone automatically when it rings. The default is ATS0=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.

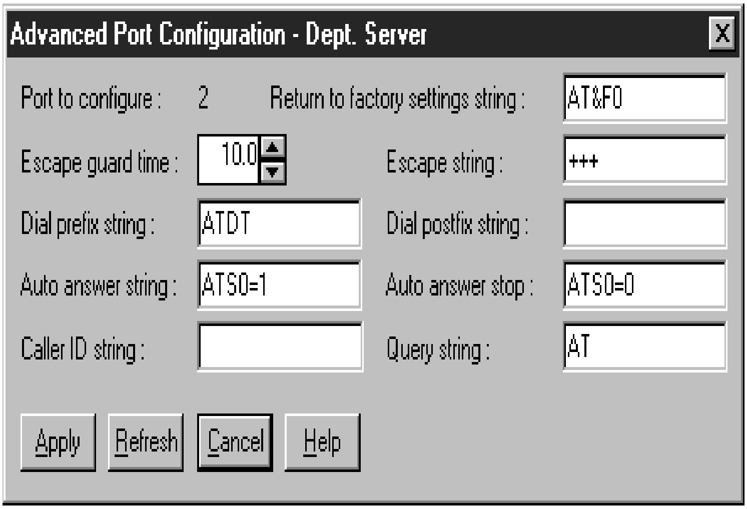

Advanced Port Configuration - Dept. Server

Port to configure : 2

Return to factory settings string : AT&F0

Escape guard time : 10.0

Escape string : +++

Dial prefix string : ATDT

Dial postfix string :

Auto answer string : ATS0=1

Auto answer stop : ATS0=0

Caller ID string :

Query string : AT

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

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

Physical Ports (as labeled)	A	B	C	MODEM	COM_AUX
Ports available to Processor Assistant	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:

1. The operating system recognizes shared ports when the system is running. The Management Processor Assistant 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 Management Processor Assistant PCI adapter or processor, when the system is started with DOS.

2. In a system with the Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant PCI adapter. If this check box is cleared, remote users will be unable to remotely access the Management Processor Assistant PCI adapter. Click **Apply** after selecting or clearing this check box to save the new setting.

Item	Description
Own port on startup	Select this check box to reserve a serial port for exclusive use by the IBM Management Processor Assistant 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. **Notes:** 1. 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 Management Processor Assistant 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. 2. Port C is dedicated for use by the Management Processor Assistant PCI adapter only. The operating system uses Port A. Port A is available to the Management Processor Assistant PCI adapter only when one of the following conditions exists: • The server is off. • The Management Processor Assistant PCI adapter needs a port to perform a critical enabled event (in this case, the IBM Management Processor Assistant 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).
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 Management Processor Assistant 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 Management Processor Assistant 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

Remote event settings

From the Management Processor Assistant 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 Management Processor Assistant adapter or processor event forwarding functions. If you configure a remote event entry, the Management Processor Assistant 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.

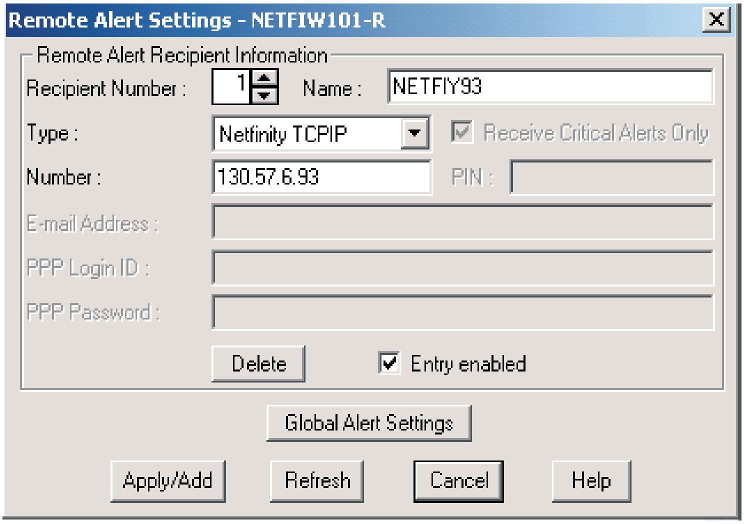
Remote Alert Settings - NETFIW101-R

Remote Alert Recipient Information

Recipient Number : 1

Name : NETFIY93

Type : Netfinity TCPIP

☑ Receive Critical Alerts Only

Number : 130.57.6.93

PIN :

E-mail Address :

PPP Login ID :

PPP Password :

☑ Entry enabled

Select Global Setting to open the following window:

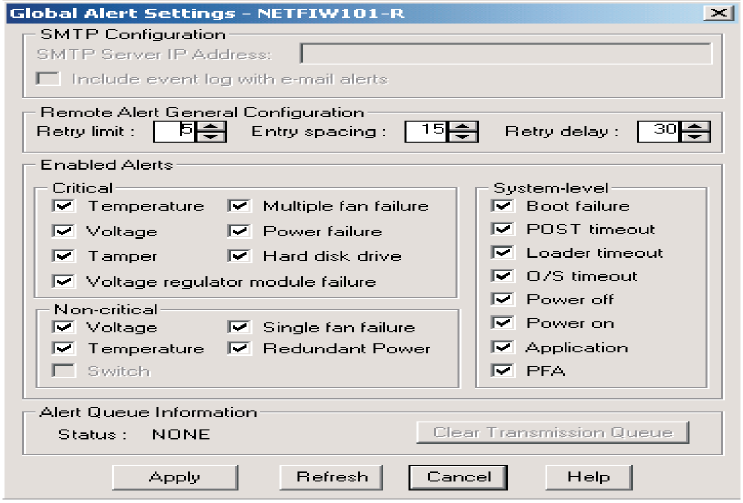
Global Alert Settings - NETFIW101-R

SMTP Configuration

SMTP Server IP Address:

□ Include event log with e-mail alerts

Remote Alert General Configuration

Retry limit : 5 Entry spacing : 15 Retry delay : 30

Enabled Alerts

Critical
- ☑ Temperature
- ☑ Multiple fan failure
- ☑ Voltage
- ☑ Power failure
- ☑ Tamper
- ☑ Hard disk drive
- ☑ Voltage regulator module failure

Non-critical
- ☑ Voltage
- ☑ Single fan failure
- ☑ Temperature
- ☑ Redundant Power
- □ Switch

System-level
- ☑ Boot failure
- ☑ POST timeout
- ☑ Loader timeout
- ☑ O/S timeout
- ☑ Power off
- ☑ Power on
- ☑ Application
- ☑ PFA

Alert Queue Information

Status : NONE

Before Management Processor Assistant can forward events to SNMP communities, you must configure the Management Processor Assistant SNMP settings. To configure SNMP settings, double-click the **SNMP Settings** component of the Network Protocols expandable tree. The SNMP Settings window opens.

You can configure the Management Processor Assistant PCI adapter or processor to forward events to multiple pagers or Director systems in response to any established 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 Management Processor Assistant PCI adapter or processor abandons the currently active event-forwarding operation and moves to the next one.

You can configure your Management Processor Assistant 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 Recipient 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 Management Processor Assistant PCI adapter or **Falcon**) 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 Management Management Processor Assistant 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), TCP/IP or SNMP trap or **IP** (for using a TCP/IP link to connect to a remote system; this is available only on systems with an Management Processor Assistant PCI adapter or **Falcon**).
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. Click **Apply/Add** to save changes.
7. Click Global settings.
8. Select dial-out events from the Enabled Events group box. If any of the selected events occur, the Management Processor Assistant 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.

9. 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 Management Processor Assistant will attempt to forward an event if an attempt fails, the amount of time that Management Processor Assistant will allow between event generation attempts, and the amount of time that Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant PCI adapter will attempt to contact each of these entries sequentially. Use the drop-down list to specify the number of seconds for the Management Processor Assistant 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.
Retry delay	Use the drop-down list to specify the number of seconds that Management Processor Assistant 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 Management Processor Assistant 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.

Management Processor Assistant 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 260 through 250.

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 Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant 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 Failures.
Power failure	The Management Processor Assistant PCI adapter or processor will forward an event if the power supply fails.	04	Power Supply Failure.
Hard disk drive	The Management Processor Assistant 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 Management Processor Assistant 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
Voltage	The Management Processor Assistant PCI adapter or processor will forward an event if any monitored voltage exceeds its threshold value.	13	Voltage

Item	Description	Numeric Code	Text String
Single fan failure	The Management Processor Assistant PCI adapter or processor will forward an event if one of the systems cooling fans fails.	11	Single Fan Failure
Redundant Power	The Management Processor Assistant 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 Management Processor Assistant 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.
POST time-out	The Management Processor Assistant PCI adapter or processor will forward an event if the time-out value (specified in the Configuration Settings window) is exceeded.	20	POST/BIOS Watchdog expired. System Restarted.
O/S time-out	The Management Processor Assistant PCI adapter or processor will forward an event if the O/S System time-out value (specified in the Configuration Settings window) is exceeded.	21	Operating system watchdog expired. System restarted.
Loader time-out	Management Processor Assistant PCI adapter or processor time-out value (specified in the Configuration Settings window) is exceeded.	26	Loader watchdog expired. System restarted.

Item	Description (if checked)	Numeric Code	Text String
Power off	The Management Processor Assistant PCI adapter or processor will turn OFF an event if the system is powered off.	23	System complex powered off.
Power on	The Management Processor Assistant PCI adapter or processor will forward an event if the system is turned on.	24	System complex powered on.
Application	The Management Processor Assistant PCI adapter or processor will forward an event if it receives an event.	22	Application logged event.
PFA	Management Processor Assistant PCI adapter will forward an event if it receives a Predictive Failure Analysis event from the system.	27	PFA

Network protocols

The Network Protocols option allows you to configure your Network and SNMP settings. From the Management Processor Assistant console, click the + beside Configuration Settings or double-click Management Processor Assistant Configuration to expand and view the protocols components

Network Settings

Double-click **Network Settings** to open the Network Settings window. Use the Network Settings window to specify network settings used by the Management Processor Assistant PCI adapter.

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

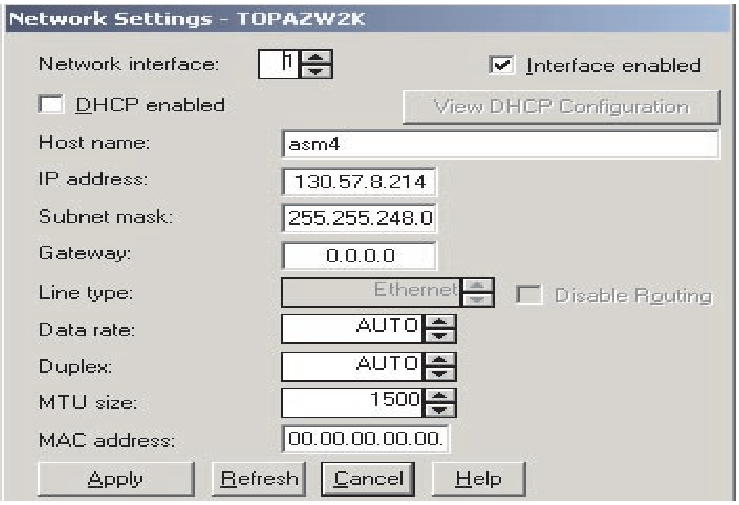

Network Settings - TOPAZW2K

Network interface: 1

☑ Interface enabled

□ DHCP enabled

Host name: asm4

IP address: 130.57.8.214

Subnet mask: 255.255.248.0

Gateway: 0.0.0.0

Line type: Ethernet

□ Disable Routing

Data rate: AUTO

Duplex: AUTO

MTU size: 1500

MAC address: 00.00.00.00.00.

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 Management Processor Assistant PCI adapter.
IP address	Type the IP address that will be used by the Management Processor Assistant PCI adapter.
Subnet mask	Type the subnet mask that will be used by the Management Processor Assistant PCI adapter.
Gateway	Type the TCP/IP address of the gateway that will be used by the Management Processor Assistant PCI adapter.
Line type	Use the drop-down list to select the line type that will be used by the Management Processor Assistant 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 Management Processor Assistant 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 Management Processor Assistant PCI adapter. Available selections are AUTO, FULL and HALF.

Item	Description
MTU size	Use the drop-down list to specify the maximum transmission unit (MTU) value that will be used by the Management Processor Assistant PCI adapter.
MAC address	Type the media access control (MAC) address of the network adapter that is being used by the IBM Management Processor Assistant PCI adapter.
DHCP	Use the list to configure DHCP settings.

Attention: If you have installed the IBM Management Processor Assistant Token Ring Connection, you must not enable or use the Ethernet port that is included on your Management Processor Assistant 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 Management Processor Assistant PCI adapter.

SNMP settings

From the Management Processor Assistant console, click Management Processor Assistant→Network Protocols→SNMP Settings. Double-click **SNMP Settings** to open the SNMP Settings window. Use the SNMP Settings window to specify SNMP settings for the IBM Management Processor Assistant PCI adapter. These settings must be configured in order for the Management Processor Assistant PCI adapter to forward events to SNMP managers on the network.

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

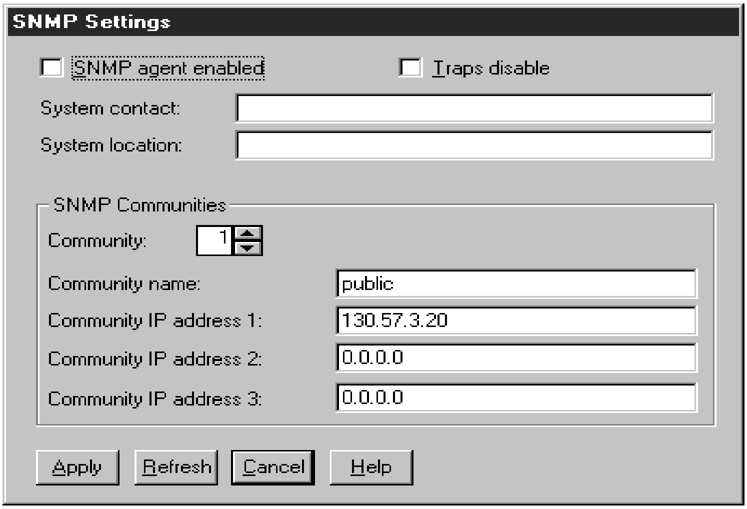

SNMP Settings

☐ SNMP agent enabled ☐ Traps disable

System contact:

System location:

SNMP Communities

Community: 1

Community name: public

Community IP address 1: 130.57.3.20

Community IP address 2: 0.0.0.0

Community IP address 3: 0.0.0.0

The SNMP Settings window contains the following items.

Item	Description
SNMP agent enabled	Select this check box to enable the Management Processor Assistant 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.
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 Management Processor Assistant PCI adapter will take effect.

PPP Settings

Some Management Processor Assistant processors implement a PPP server. Utilizing PPP agent software you can establish a TCP/IP connection to the Management Processor Assistant over a modem. This allows the use of TCP/IP based applications, such as configuring the Management Processor Assistant processor through the web interface it provides. The following fields should be configured if you wish to enable the PPP interface on the Management Processor Assistant processor.

Item	Description
PPP enabled	Select this field to enable or disable your Point-to-Point (PPP) Interface. The PPP Interface is enabled when this box is checked.
Local IP address	Select this field to define the IP address assigned to the Management Processor Assistant PPP Modem connection. It must be entered as a xxx.xxx.xxx.xxx formatted string.
Remote IP Address	Select this field to define the IP address to the PPP connection that dialed into the Management Processor Assistant. It must be entered as a xxx.xxx.xxx.xxx formatted string.
Subnet Mask	Select this field to configure the Subnet mask for this IP Interface. This bit mask is used in conjunction with the Local IP Address to determine a range of IP addresses that can be communicated with over this interface.
Authentication Protocol	Select this field to specify the type of authentication protocol that will be negotiated when a PPP connection is attempted. To set the PPP Authentication protocol, click the pull-down button and select one of the following: • PAP - This option uses a 2-way handshaking procedure to validate the identity of the originator of the connection. This is a weaker authentication protocol, but it is necessary if a plaintext password must be available to simulate a login at a remote host. • CHAP - This option uses a 3-way handshaking procedure to validate the identity of the originator of the connection upon connection or any time later. This is a stronger authentication protocol that protects against playback and \"trial and error"\ attacks. • CHAP then PAP - This option tries to authenticate using CHAP first. If the connecting server doesn't support CHAP, then PAP will be tried as a secondary authentication protocol. The default for this field is CHAP then PAP. Also, MD5-CHAP is the only supported algorithm.
Apply	When this button is selected, the configuration as currently displayed will be written to the Management Processor Assistant.

Item	Description
Refresh	Selecting this option will result in a query of the current configuration from the Management Processor Assistant by the software. Any changes made in the dialog since the last Apply will be lost.
Cancel	Select this button to close the PPP Configuration dialog. Any changed made since the last Apply will be lost.

Note: When the PPP interface is enabled, standard serial communications are disabled. If the PPP interface is enabled, it is not possible to use the Serial Connection option of the Management Processor Assistant tool to communicate with the Management Processor Assistant processor.

Restarting the Management Processor

Some changes in the configuration of the Management Processor require a restart of its firmware to take effect. Restarting the Processor will cause it to reset all of its communication interfaces, therefore the current session will be closed.The configuration groups that require a restart are:

- Network Settings
- SNMP Settings
- PPP Configuration
- DNS Configuration

Remote management

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

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

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

Notes:

1. TCP/IP linking over a network connection is available only when you are using Management Processor Assistant to directly access an Management Processor Assistant PCI adapter that is connected to your network.

2. Management Processor Assistant Interconnect connections are available only in the following situations:
   - You are using Management Processor Assistant to directly access the Management Processor Assistant PCI adapter that is installed in your own system (Netfinity 7000 M10 only).
   - You are using Management Processor Assistant to access an Management Processor Assistant adapter or processor that is connected to the same Management Processor Assistant Interconnect network that your Management Processor Assistant adapter or processor is connected to.
   - You have first established a TCP/IP or serial link with a remote Management Processor Assistant PCI adapter that is connected to other Management Processor Assistant adapters or processors on an Management Processor Assistant Interconnect network.

Configuring and establishing a TCP/IP connection

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

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

   The Establish TCP/IP Connection window opens.

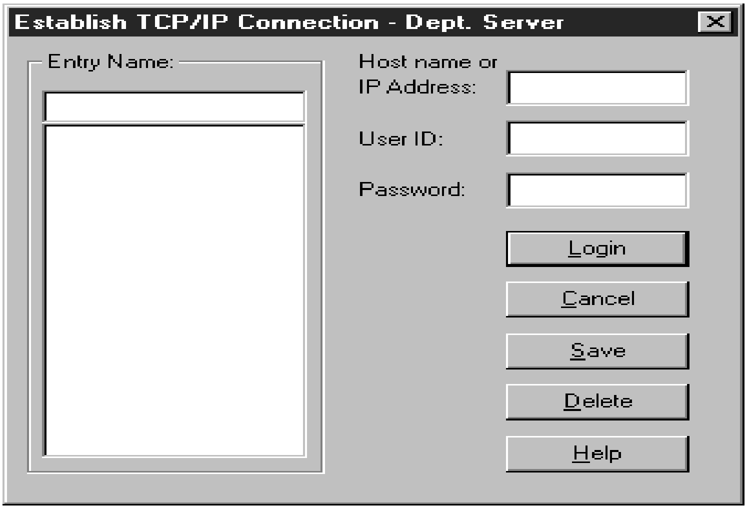

3. Select a TCP/IP connection entry from the **Entry Name** list, or create a new entry and then select the new entry. To create a new entry:
   a. In the **Entry Name** field, type a name for the entry.
   b. In the **Host Name** or **IP Address** field, type the TCP/IP address or host name that is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

      This must match a user ID and password combination that has been configured, using the Management Processor Assistant service, to allow access to the Management Processor Assistant 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 Management Processor Assistant PCI adapter. When the connection is established, use the Management Processor Assistant service to manage the remote Management Processor Assistant PCI adapter.

Establishing a Management Processor Assistant Interconnect connection

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

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

   The Establish Interconnect window opens.

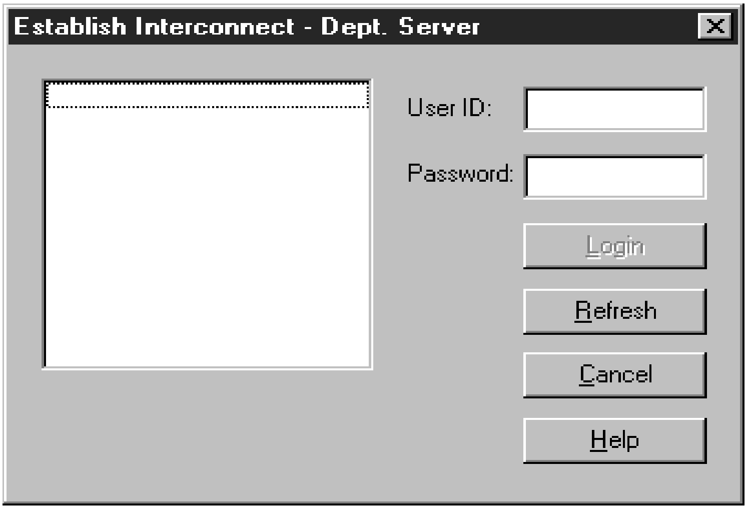

3. Select a system from the **Establish Interconnect** the list.
4. In the **User ID** and **Password** fields, type a user ID or password for logging on to the remote Management Processor Assistant adapter or processor.

   Type a user ID and password that will allow access to the remote Management Processor Assistant adapter or processor. This must match a user ID and password combination that has been configured, using the Management Processor Assistant service, to allow access to the Management Processor Assistant adapter or processor.
5. Click **Login** to establish a Management Processor Assistant Interconnect connection with the selected system. After the connection is established, use the Management Processor Assistant service to manage the Management Processor Assistant PCI adapter or processor in the remote system.

Event Log

From the Management Processor Assistant console, click **Event Log** to open the Event Log window. This window contains all entries that are currently stored in the Management Processor Assistant PCI adapter or processor Event Log. The Management Processor Assistant adapter or processor Event Log records information about all remote access attempts and dial-out events. System Health status icons appear beside managed systems that need attention. For more information see Chapter 23, “Hardware Status,” on page 209.

Notes:

1. If you are using the Management Processor Assistant service with an Management Processor Assistant 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.

2. If you are using the Management Processor Assistant 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.

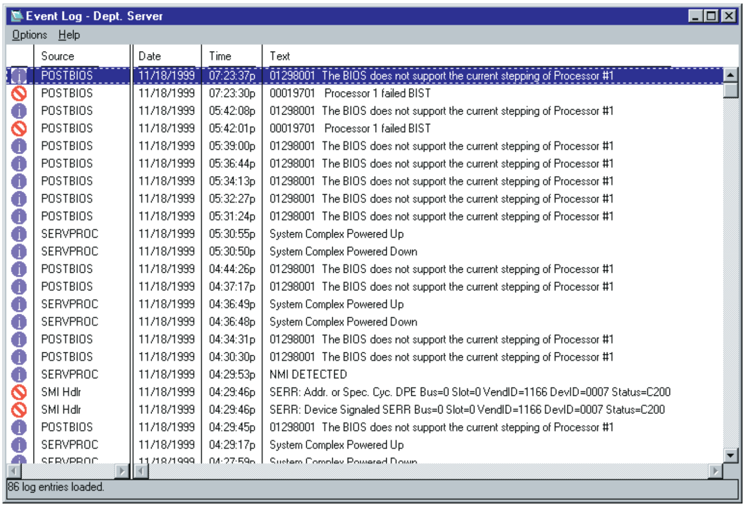

	Source	Date	Time	Text
ⓘ	POSTBIOS	11/18/1999	07:23:37p	01298001 The BIOS does not support the current stepping of Processor #1
⊘	POSTBIOS	11/18/1999	07:23:30p	00019701 Processor 1 failed BIST
ⓘ	POSTBIOS	11/18/1999	05:42:08p	01298001 The BIOS does not support the current stepping of Processor #1
⊘	POSTBIOS	11/18/1999	05:42:01p	00019701 Processor 1 failed BIST
ⓘ	POSTBIOS	11/18/1999	05:39:00p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	05:36:44p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	05:34:13p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	05:32:27p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	05:31:24p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	SERVPROC	11/18/1999	05:30:55p	System Complex Powered Up
ⓘ	SERVPROC	11/18/1999	05:30:50p	System Complex Powered Down
ⓘ	POSTBIOS	11/18/1999	04:44:26p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	04:37:17p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	SERVPROC	11/18/1999	04:36:49p	System Complex Powered Up
ⓘ	SERVPROC	11/18/1999	04:36:48p	System Complex Powered Down
ⓘ	POSTBIOS	11/18/1999	04:34:31p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	POSTBIOS	11/18/1999	04:30:30p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	SERVPROC	11/18/1999	04:29:53p	NMI DETECTED
⊘	SMI Hdlr	11/18/1999	04:29:46p	SERR: Addr. or Spec. Cyc. DPE Bus=0 Slot=0 VendID=1166 DevID=0007 Status=C200
⊘	SMI Hdlr	11/18/1999	04:29:46p	SERR: Device Signaled SERR Bus=0 Slot=0 VendID=1166 DevID=0007 Status=C200
ⓘ	POSTBIOS	11/18/1999	04:29:45p	01298001 The BIOS does not support the current stepping of Processor #1
ⓘ	SERVPROC	11/18/1999	04:29:17p	System Complex Powered Up
ⓘ	SERVPROC	11/18/1999	04:27:59p	System Complex Powered Down

86 log entries loaded.

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.

Option	Description
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.

Selecting an event source

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

- Driver

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

- TCP/IP

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

  Note: This selection is available only on systems that use a Management Processor Assistant PCI adapter.

- Serial
  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 Director events.

Updating PCI adapter or processor microcode

To update the Management Processor Assistant PCI adapter or processor microcode, use the following procedure:

1. From the Management Processor Assistant 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 Management Processor Assistant PCI adapter or processor microcode, the Management Processor Assistant service will apply the microcode update to the Management Processor Assistant PCI adapter or processor.

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

Note:

This option is not available with an Management Processor Assistant Interconnect.

Updating system POST/BIOS microcode

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

1. From the Management Processor Assistant 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 Management Processor Assistant service will apply the microcode update to the system that contains the Management Processor Assistant PCI adapter or processor to which you are connected.

Note: This option is only available with direct TCP/IP connection to the Management Processor Assistant card or direct serial connection.

Management Processor Assistant events in Director

The Management Processor Assistant tool adds the following events to the event builder log:

- Critical Enabled Events, see page 260.
- Non-critical Enabled Events, see page 261.

- System Enabled Events, see page 262.

Using the Management Processor Assistant PCI adapter as a network gateway

In systems that contain both an Management Processor Assistant processor and an Management Processor Assistant PCI adapter, the adapter acts as an Ethernet or token-ring network gateway (or as a shared modem resource). In this configuration, the Management Processor Assistant processor generates all events, time-outs, and other systems-management information. This data is relayed to the Management Processor Assistant PCI adapter using the Management Processor Assistant 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 Management Processor Assistant PCI adapter and the Management Processor Assistant processor, all systems-management settings (such as remote event settings and time-out settings) must be configured on the Management Processor Assistant processor. However, before using TCP/IP to communicate with your Management Processor Assistant adapter, you must first establish an Management Processor Assistant Interconnect connection with the adapter and configure the network settings.

To establish another Management Processor Assistant Interconnect connection to another Management Processor Assistant PCI adapter through a serial port, use the following procedure:

1. From the Management Processor Assistant Integrated Console, drag and drop the Management Processor Assistant icon onto a system. The Management Processor Assistant Console window opens.
2. Click **Options →Change Connection →Serial**.
3. In the **Entry Name** field, type the name of the connection.
4. Type the phone number of the remote Management Processor Assistant 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.

Chapter 26. Capacity Manager

Capacity Manager provides proactive management of hardware resources by gathering and presenting historic data or formulating trend analyses that identify and predict system performance bottlenecks.

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. Capacity Manager identifies potential bottlenecks in a network, enabling effective planning of future capacity needs, such as microprocessor, disk, network 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. However, you can collect data from any remote system running Agent Services for IBM Director for OS/2, Windows 2000, Windows NT, or NetWare and Linux.

Using Capacity Manager tasks

Capacity Manager contains the following features:

- 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 Windows 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.
  - System status is tabulated during report generation. 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 timeout to elapse, Capacity Manager will use its 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: Use it to select a report duration, a global sampling frequency, and the days and times for collecting data.
  - The Method of Generating a Report pane: Use it to choose between generating a report to the viewer, or to a file. To help you keep track of report files and to know which files can be merged, include in the file name, the date and time of report generation as well as the report definition name.
  - The Monitor Selection pane: Use it 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): Use it to set the length of time each system has to respond for inclusion in a report.
- The Report Viewer has the following options:
  - 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 save 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

Monitor Activator allows you to select which monitors are active on a managed system. By default, CMAgent collects data on a predefined set of default monitors. CMAgent runs on each managed system and continuously collects data from the monitors that are active on the system. This data is stored on the managed system in two .slt files. One slot file contains daily information by minute. The other slot file contains monthly information that is collected every five minutes. Monitor Activator detects all monitors that are available on selected systems and displays their status.

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:

1. Not all monitors are present on all systems. For example, drive E is present only on systems that have it installed.

2. 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. When an agent does not have active monitors, you will receive a "no monitors active" error message.

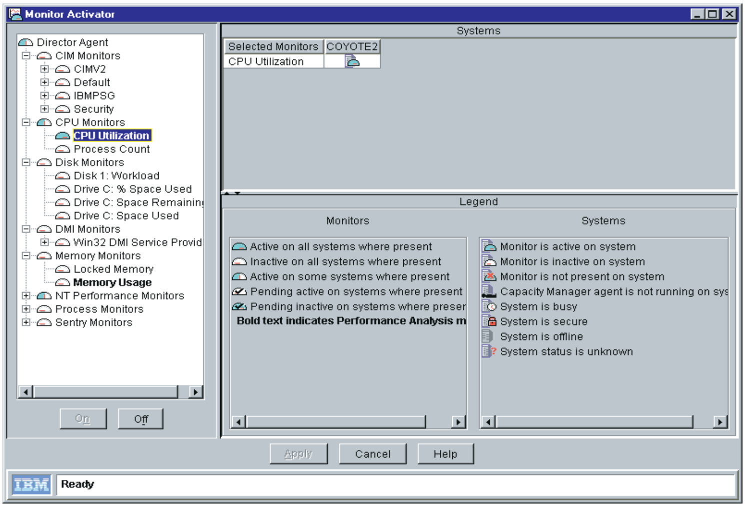

Selected Monitors	COYOTE2
CPU Utilization	

Monitors:

- Active on all systems where present
- Inactive on all systems where present
- Active on some systems where present
- Pending active on systems where present
- Pending inactive on systems where preser

Bold text indicates Performance Analysis m

Systems:

- Monitor is active on system
- Monitor is inactive on system
- Monitor is not present on system
- Capacity Manager agent is not running on sys
- System is busy
- System is secure
- System is offline
- System status is unknown

The Monitor Activator window 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.

The Monitors pane opens 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. It 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 provide only 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. Select 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.

Systems	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.

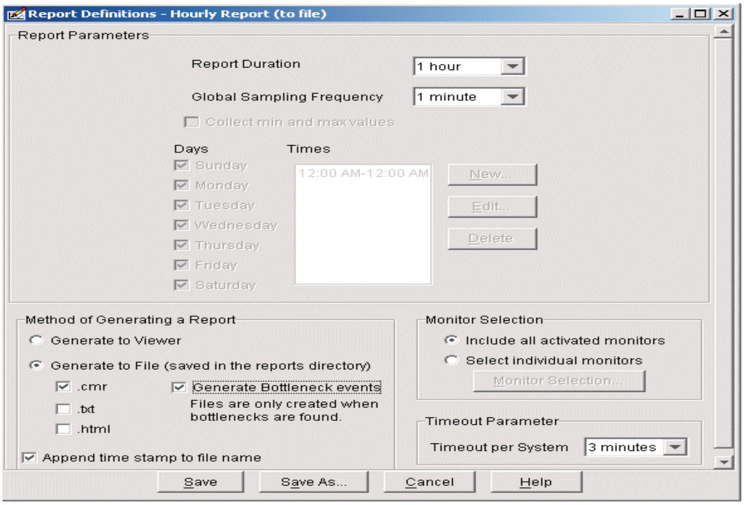

Report Parameters

Report Duration: 1 hour

Global Sampling Frequency: 1 minute

☐ Collect min and max values

Days: ☑ Sunday ☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday ☑ Friday ☑ Saturday

Times: 12:00 AM-12:00 AM

Method of Generating a Report

○ Generate to Viewer

● Generate to File (saved in the reports directory)

☑ .cmr

☐ .txt

☐ .html

☑ Generate Bottleneck events

Files are only created when bottlenecks are found.

☑ Append time stamp to file name

Monitor Selection

● Include all activated monitors

○ Select individual monitors

Timeout Parameter

Timeout per System: 3 minutes

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

Item	Description
Report Duration	Indicates the time span of a report, starting from the time of report generation and going back. 1 day, ending at the beginning of the current hour. 3 hours, includes data from the previous 3 hours including the current hour. 8 hours, includes data from the previous 8 hours including the current hour. 1 week, ending at the beginning of the current day. 1 month, ending at the beginning of the current day. 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.
Global Sampling Frequency	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).

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 a slower sampling frequency. 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.
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.

Item	Description
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 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 click **Delete**.

Generating a report

To generate a report, drag and drop the Report Definition icon onto the selected systems or clusters or groups that need to be analyzed. 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 window 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 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 Report 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 **CMView** from the command line. This is useful for running the Report Viewer outside of 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.

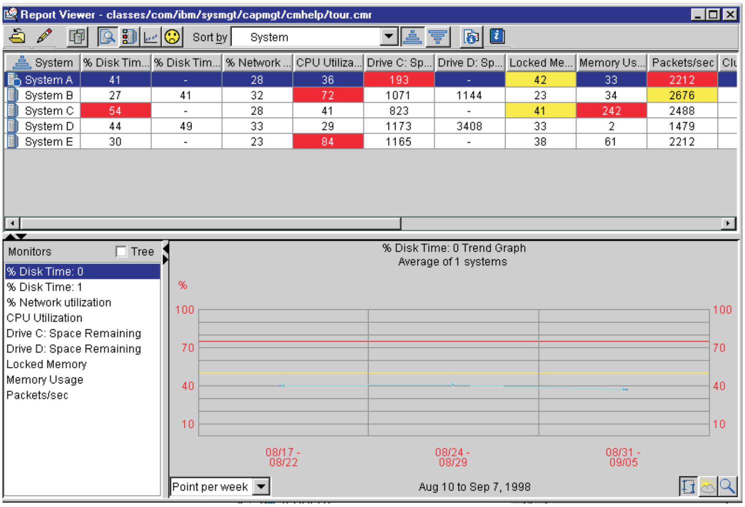
Report Viewer - classes/com/ibm/sysmgt/capmgt/cmhelp/tour.cmr

System	% Disk Tim...	% Disk Tim...	% Network ...	CPU Utiliza...	Drive C: Sp...	Drive D: Sp...	Locked Me...	Memory Us...	Packets/sec	Clu
System A	41	-	28	36	193	-	42	33	2212	
System B	27	41	32	72	1071	1144	23	34	2676	
System C	54	-	28	41	823	-	41	242	2488	
System D	44	49	33	29	1173	3408	33	2	1479	
System E	30	-	23	84	1165	-	38	61	2212	

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. Within the Report viewer window, systems can be viewed 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 cluster view, 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 view, you can change from Table view to either Icon or Hypertrophy 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 a red horizontal line 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. Within the Graph pane, you can use the following tools:

- 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 predicted data based on least-squares 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 300.

  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. Use these tabs to configure 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 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. The data lines for multiple-system line graphs or trend graphs are not bracketed. The maximum line represents the highest average values at each data point, and the minimum line represents the lowest. The Trend button activates the trend graph. By default, you can select up to nine systems to graph individually. A trend graph plots an average monitor value for all graphed systems at each time point. The marks on the vertical line depict the range of system monitor values at a given time. 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, to open a Settings notebook to configure the view and to Enable Performance Analysis.
- 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

For Windows 2000 and Windows NT, the algorithm uses the following monitors:

- Memory usage
- % disk time
- CPU utilization and, depending on the operating system you use, any of the following monitors to reflect LAN adapter performance:
- Packets/sec
- Total bytes/sec
- % network utilization

For Linux, the algorithm uses the following monitors:

- Memory Used Non-Cached (MBytes)
- Disk IO operations/second

- CPU Utilization and, any of the following monitors to reflect LAN adapter performance
- Bytes/second
- Packets/second

By default, Capacity Manager will activate all the required Performance Analysis (PA) monitors that are present on your systems, but only systems running Windows and Linux have all the required monitors; therefore Performance Analysis is not available for all operating systems. The packets/sec and CPU "X" utilization monitors are not required, but omitting them might cause Performance Analysis to miss some system problems. To disable Performance Analysis, from the menu, select **Edit** and deselect **Enable Performance Analysis** from the drop-down menu.

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:

1. Not all systems have the packets/sec monitor, but if this monitor is present on your system, it will be activated with your other Windows 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.

2. There is a CPU "X" Utilization monitor for each processor device on your system, but the monitor for CPU utilization monitors all processor-time devices on your system.

3. When an agent does not have active monitors, you will receive a "no monitors active" message.

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 c*onstrained* 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 (PA) function probes for bottlenecks in server hardware performance, diagnoses the problems, and suggests ways to improve performance. Bottlenecks are detected when one or more monitors exceed their threshold settings. You can adjust these threshold settings, but the default settings, particularly those that are important to the integrity of the Performance Analysis can be changed without consequences to the Performance Analysis report. To help you identify the PA monitors, their names are in bold on the Monitor Activator window, the Monitor Selection window in the Report Definition task, and on the Monitor tab for the Settings notebook in the Report Viewer.

Performance Analysis icons appear in the toolbar of the Report Viewer window. The Performance Analysis function buttons appear as one of six icons, each of which represent 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 has been analyzed and the aggregate set contains the required set of PA monitors, but no individual system has all of the monitors.
🙁	Your Performance Analysis report is ready, and you have system bottlenecks.
❌	Your Performance Analysis report could not be prepared; Enable Performance Analysis not selected under the Edit menu.

To see a Performance Analysis report of the system data, click **Edit→Enable Performance Analysis** from the menu. If this is not checked, no performance analysis is done and the performance analysis icon will be replaced with an x-ed out face.

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 you need 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.

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.

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.

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.

Group support

With Group Support, you can define a set of systems as a group within Capacity Manager and treat that entity as a unit. Capacity Manager will build on the cluster support that exists already and the unit group will be treated like a cluster. The cluster views will extend to support the user-defined group.

To define a Group, use the Director Management console. When a Group is dragged directly to a Capacity Manager report generator, the Director Management console group name is used as the name for the systems in that Group.

From the Director Management console each user-defined group is listed under All Groups. All Groups can be expanded to reveal rows for each system in the group when individual system management is desired. The cells from the group reflect the overall status of the group of systems and contain the average value. The color of a cell is red if one or more systems have a bottleneck or the average value exceeds the critical threshold. The color of the cell is yellow if one or more systems have latent bottlenecks or the average value exceeds the warning threshold.

When a group is selected all individual systems in that group will be graphed in the Report Viewer.

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 the bottleneck.

There are four types of single bottlenecks that are detected by Performance Analysis and each of the PA monitors detects one of these four bottleneck types:

- The CPU monitors detect a CPU bottleneck.
- The % disk-time monitor(Windows only) detects a disk bottleneck.
- The I/O operations/second monitor (Linux only) detects a disk bottleneck.
- The memory-usage monitor detects a memory bottleneck.
- The Used Non-Cached (MBytes) monitor (Linux only) detects a memory bottleneck.
- Total Bytes/sec, the Packets/sec and the % Network Utilization monitors (Windows only) detect a LAN adapter bottleneck.
- The Bytes/second and the Packets/second monitors (Linux only) 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 CPU - 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 functions, use the Report Viewer Tour.

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

1. Generate events when a bottleneck occurs. See “Generating an event”.
2. Create an event filter that filters out all events except bottlenecks. See “Creating an event filter” on page 299.
3. Define an action and an action plan that automatically occurs when the plan detects a bottleneck. See “Defining an action and action plan” on page 299.
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 online helps 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 the current 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 HMTL 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.

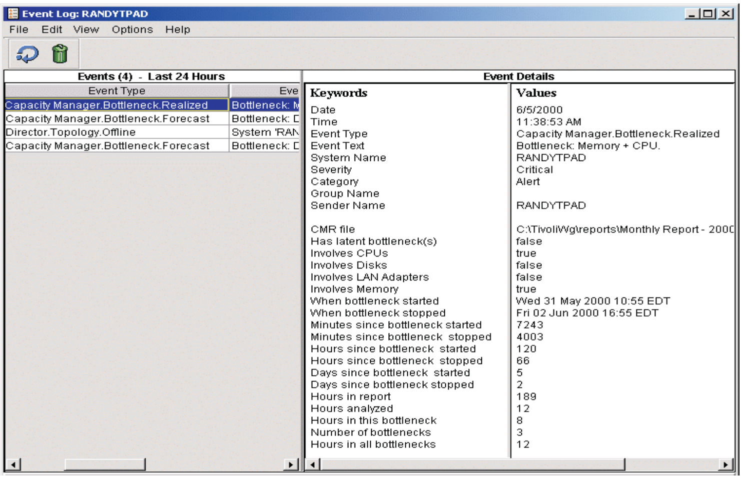

Events (4) - Last 24 Hours

Event Type	Eve
Capacity Manager.Bottleneck.Realized	Bottleneck: M
Capacity Manager.Bottleneck.Forecast	Bottleneck: D
Director.Topology.Offline	System 'RAN
Capacity Manager.Bottleneck.Forecast	Bottleneck: D

Event Details

Keywords	Values
Date	6/5/2000
Time	11:38:53 AM
Event Type	Capacity Manager.Bottleneck.Realized
Event Text	Bottleneck: Memory + CPU.
System Name	RANDYTPAD
Severity	Critical
Category	Alert
Group Name	
Sender Name	RANDYTPAD
CMR file	C:\TivoliWg\reports\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

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 a tree structure 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. That attribute is: 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 296.

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”, 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 displays the monitor for the current selection. To see another forecast, click its name in the monitor box.

Notes:

1. You cannot use the Zoom tool and the Forecast tool at the same time; they are mutually exclusive.

2. 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 the following 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 within 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

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.

Chapter 27. Fuel Gauge Monitor

Fuel Gauge Monitor warns you about conditions that could lead to preventable down time. These conditions involve the power subsystem and the load presented by that 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 form 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 managed system loads comes back down below these thresholds.

A second condition arises when a managed system that has multiple, plug-in power modules, experiences an increase in system load which takes power subsystem utilization above a limit know 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 an *over current situation* the server is at that time.

All Fuel Gauge events are forwarded to the Director event log viewer.

Starting the Fuel Gauge task

From the Director Management console, drag and drop the Fuel Gauge Monitor icon onto one or more systems. The Fuel Gauge window 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 window.
- 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.

Gathering information data

The following parameters are received from a selected system and can be displayed either graphically (Pie chart, Bar chart) or in a tabular format (Table view). The view displays data polled from the system at the time the task is started. Each tabbed view displays the same data.

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.
12v Utilization(U12v)	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.

Parameter	Description
5v Utilization(U5v)	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.
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.
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 Table view is the default view of the Fuel Gauge Monitor console. To refresh the data, click **Update**.

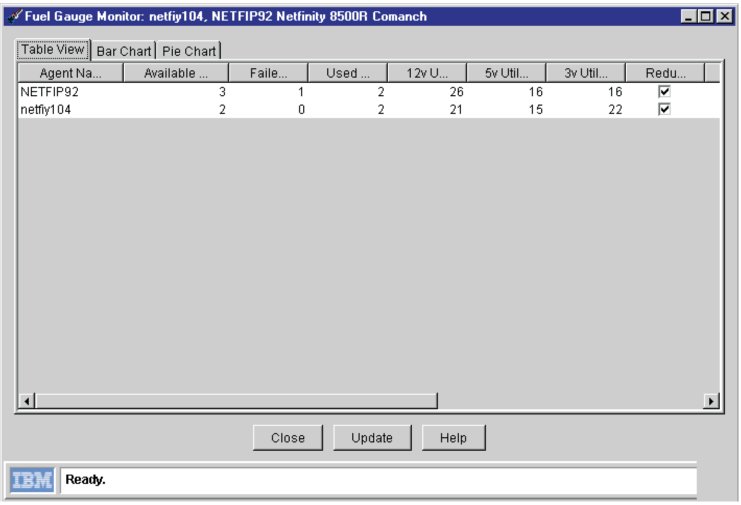
Fuel Gauge Monitor: netfiy104, NETFIP92 Netfinity 8500R Comanch
Table View
Bar Chart
Pie Chart
Close
Update
Help
IBM

Agent Na...	Available ...	Faile...	Used ...	12v U...	5v Util...	3v Util...	Redu...
NETFIP92	3	1	2	26	16	16	☑
netfiy104	2	0	2	21	15	22	☑

Ready.

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, click the bar chart tab and select a parameter from the drop down list. Each bar on the chart represents a system. A single system is displayed with a single bar. Multiple systems display as ratios with separate data.

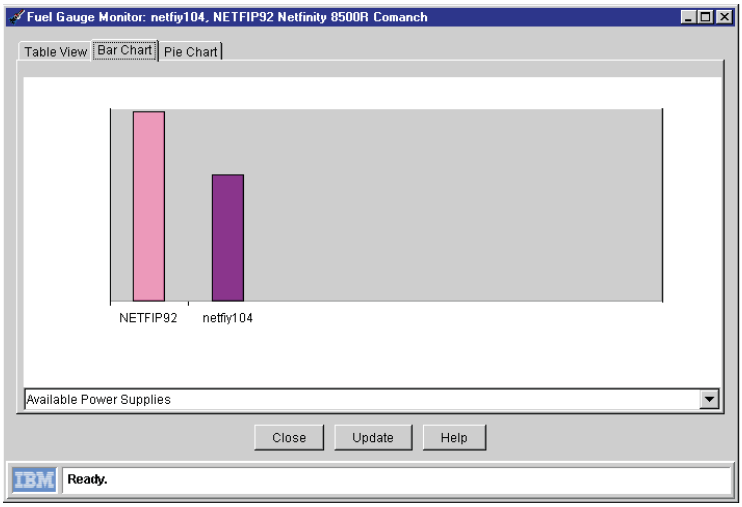
Fuel Gauge Monitor: netfiy104, NETFIP92 Netfinity 8500R Comanch
Table View
Bar Chart
Pie Chart
NETFIP92
netfiy104
Close
Update
Help
IBM

Available Power Supplies

Ready.

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, click the pie chart tab and 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.

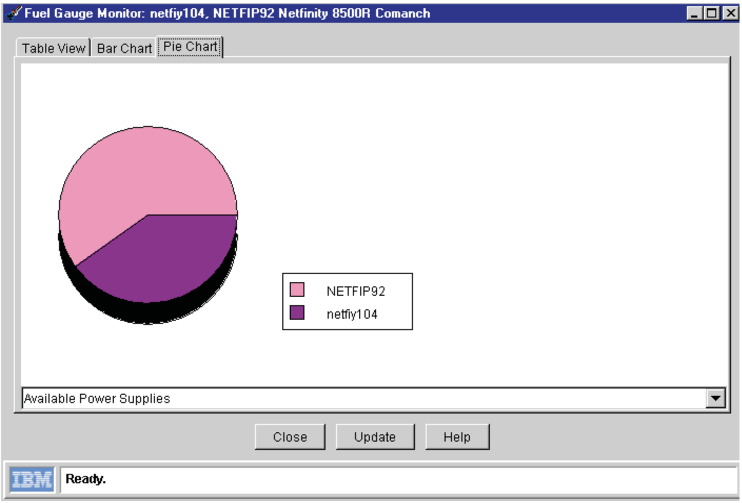

Chapter 28. Rack Manager

Rack Manager is a flexible, easy-to-install solution for consolidating and configuring 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.

Rack Manager has these advantages:

- 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.

Starting the Rack Manager task

To start the Rack Manager task, from the Director Management console, use one of the following techniques:

- From the Group pane of the Director Management console, select a group of managed systems. Drag and drop the group onto the Rack Manager task icon or right-click and select Rack Manager.
- From the Group Contents pane, select a managed system. Drag and drop the managed system on to the Rack Manager icon or right-click and select Rack Manager from the context menu.
- From the Group pane or the Group Contents pane, select multiple systems. Drag and drop the managed systems onto the Rack Manager task icon.

The Rack Manager Window opens. The selected managed systems are listed under the Floor division of the topology page in the Control pane.

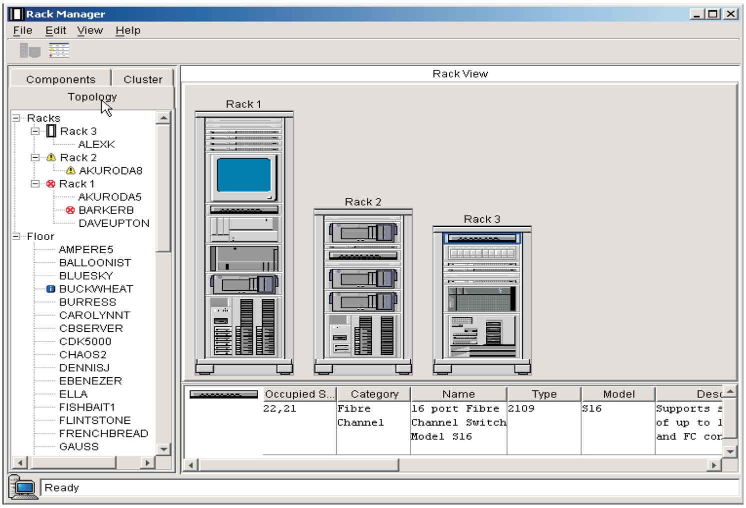

	Occupied S...	Category	Name	Type	Model	Desc
	22,21	Fibre Channel	16 port Fibre Channel Switch Model S16	2109	S16	Supports s of up to 1 and FC cor

The Rack Manager window consists of the following components:

- **Menu**- The menu contains File, Edit, View and Help. Using the menu, you can close a task, delete a rack, add racks, change views or obtain information from the online 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** - The Control Pane contains the Rack Topology page, Component page and the Cluster page. Use these pages to configure and manage the rack and rack components.
- **Properties Window** - The Properties Window appears below the Rack Manager Workspace 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.

Viewing the Control Pane

The Control Pane consists of the **Topology**, **Component** and **Cluster** pages. Each page displays an expandable tree. Use the Control Pane to select and build racks and rack components.You will perform most of your Rack Manager tasks by selecting a server or component from one of these expandable trees.

Topology

Click the Topology page to display an expandable tree of racks and servers. Racks are organized alphabetically. To display information about a server, expand or collapse the + beside the server. If the server is not a member of a rack, it can be found at the bottom of the expandable tree.

Click a rack or a server to select it. Clicking will highlight the rack or the server in the rack view and display property data in the Properties window. Right-clicking displays a menu for the rack or the server. Use the menu to associate or disassociate the server from a particular rack. 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 object) has not been created, 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 page.

Component

Click the **Component** page to display a list of components that can be used to configure racks.

Rack enclosures and components categories are listed in an expandable tree. To collapse or expand component categories, click the + or - to the left of the category icon. For a list of supported rack components, see “Supported rack components” on page 319.

When you select a component in the Component page, details about the component, including the name and technical specifications, are displayed in the Properties window. There are separate tables for each of the different component categories.

To add a component to a configuration, click the component and drag and drop the component onto the rack in the Rack view.

For each component, information is displayed so that you can determine which component to use. Measurements for components are shown in US units, and EIA units.

Cluster

Click the Cluster page to display a group of clusters. Components not included in a cluster are considered independent. It is important to note that although classed as independent, these components might in fact be in a rack.

Selecting a cluster highlights all the devices in the cluster that are currently on display in the Rack Manager workspace.

Using 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.

Creating and configuring a rack

To create a rack, from the Component page, use one of the following procedures:

- Select a rack and drag it to the Rack View.
- From the menu bar, click **File→ New Rack**.
- Place the cursor in the Rack View area of the Rack Manager workspace and right-click. From the context menu, click **New Rack**.

The 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 page, expand the last branch of the expandable tree. The servers represent managed objects discovered by Director. Drag servers from the last branch of the tree onto the rack in the Rack view. From the Component page, you can also expand a category of components available to be placed in a rack and drag the component onto the rack. Components can be associated to a particular predefined 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 warning message appears. 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 click **Associate Managed Object**.

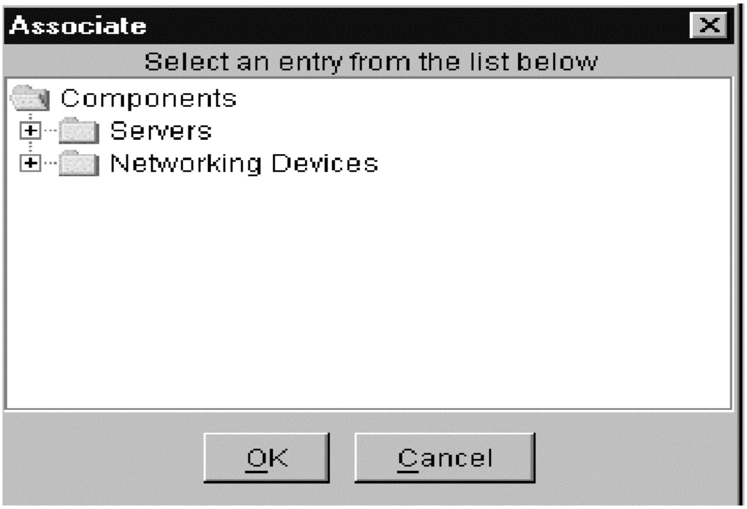

Removing a rack or a rack component

To remove a rack, right-click the rack component in the Rack view and click **Delete** from the menu. You can also left-click to select a rack and from the Edit menu, select **Delete**. 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.If there is a managed system associated with a managed object, it will appear on the floor

Associate component

When a component is not rack mountable, only the system type and model are listed. In addition, data collection might not have been performed successfully and an Inventory Data Not Available statement appears in the Properties window.

To provide more information about a particular component that includes additional properties, you can associate the component with a predefined IBM component from the items in the Components page.

To associate a component to an IBM predefined component, use the following procedure:

1. From the Floor division in the Topology page, select a component and right-click the component.
2. Click the **Associate**.
3. Select the component from the Floor division and dragging it to the Rack view, or select the appropriate component match from either the Servers or Networking Devices folders.
   The Associate context menu appears.

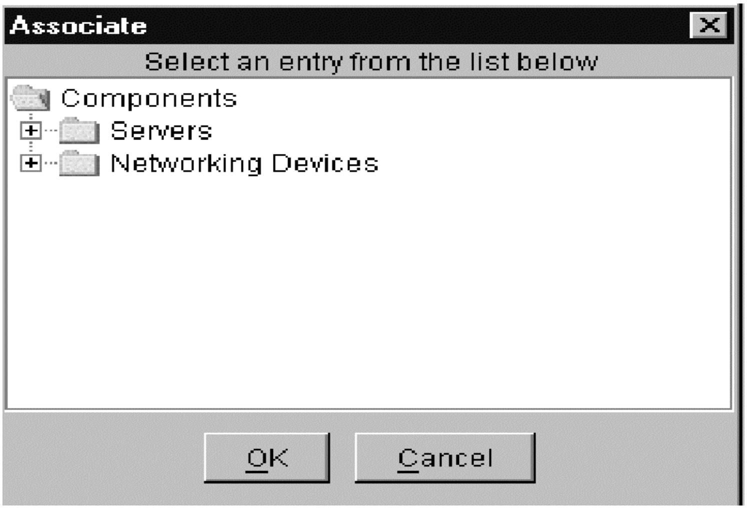

Disassociating or Canceling a component

You might want to cancel a component association or disassociate a component when:

- You have used an incorrect component association.
- Component inventory collection has been performed.
- The association is no longer valid.

To cancel a component, use the following procedure:

1. From the rack in the Rack view, right-click the component.
2. Click **Delete**. The component will appear in alphabetical order as a leaf below the floor division of the Topology page.

To disassociate a component, use the following procedure:

1. From the Topology page under the floor division, right-click the component.
2. Select **Disassociate**. (You can also right-click a component in the rack view and select **Disassociate** from the edit menu.)

The component information in the Properties window reverts to the information created at the time of discovery through Director.

Reassociate components

You can use this option when a managed system is on the Floor of the control pane and not in a Rack. Right-click and select **Reassociate** component.

Viewing Components

From the Director Management console, select a component and right-click. A context menu opens. Click **View Inventory**. The Inventory Query Browser opens. This window contains the available query. Select query from the list.

Viewing a rack

The Rack view displays the physical rack. By default, the Rack view is open when you initially start Rack Manager. When you select or drag and drop 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 units 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 (an error message will appear) and place it back onto the hierarchical tree or the component is return to the place it originated. The graphical representation of the component displays the coordinates of the units (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.The name of the removed unit will appear in the floor division of the hierarchical tree in the Topology page. To change to the Rack view, click the rack and monitor icon on the toolbar or from the system menu, select **View →Rack View**.

Right-click on a managed system in the Rack view and the following variable menu appears:

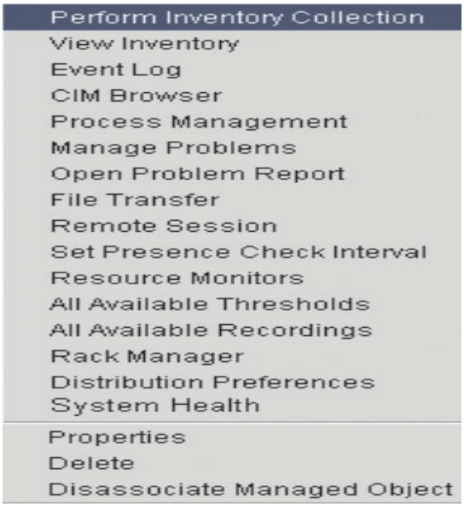

Use this menu to perform the following tasks:

- Perform Inventory Collection

- View Inventory
- CIM Browser
- Process Management
- Manage Problems
- Open Problem Report
- File Transfer
- Remote Session
- Set Presence Check Interval
- Resource Monitors
- All Available Thresholds
- All Available Recordings
- Rack Manager
- Distribution Preferences
- Hardware Status
- Properties
- Delete
- Disassociate Managed Object

Viewing the table

The Table view displays the data that is represented in a tabular format. The columns contain the rack, the rack position, device name, status, and component description. These positions can be changed by clicking a column heading and dragging it to the new position. To change the view, click the Table icon on the toolbar or select **View →Table view**.

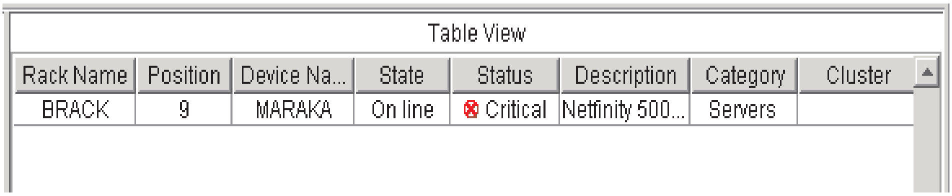

Table View

Rack Name	Position	Device Na...	State	Status	Description	Category	Cluster
BRACK	9	MARAKA	On line	⊗ Critical	Netfinity 500...	Servers	

Monitoring Hardware Status

Health Status events and their associated icons are available to Rack Manager. See “Viewing system-environment factors” on page 210 for more information on hardware status events. The events alert you to important changes in managed-object status so that you can take some steps to respond to those changes. For example, when a server generates a hardware status event, the Rack Manager task displays the associated icon alongside the representation of the managed-object in the control pane and in the table view.

To notify you of changes in environmental status, hardware status alert icons are displayed in the Status column in the table view in the Control Pane. The severity of the situation is indicated and specific information describing the failure is available. Three states are reported: warning, critical, and informational.

Any component, in the Topology page or Cluster page, with a warning, critical, or information health status will have the associated icon displayed alongside its entry. The parent server will also display that icon as will as the parent rack. When more than one status is identified, the icon for the most significant indicator will be assigned to the managed-object parent. Selecting a managed-object subcomponent in the expandable tree will have the same effect as selecting the managed-object parent in the server. That is, the server in the associated rack will be highlighted.

To obtain more information about hardware status, right-click on a server in the rack view. From the context menu, select **Hardware Status**. Rack Manager detects and displays the following hardware status icons in rack or table view:

Icon	Name	Description
⊗	Critical icon	Indicates a significant problem that you should investigate.
⚠	Warning icon	A moderately significant event you should consider investigating.
ℹ	Informational/ harmless icon	A routine event designed to provide some information.

For more information on hardware status, see Chapter 23, “Hardware Status,” on page 209.

Viewing Properties

Select a rack and right-click. A context menu opens. Click **Rack Properties**. Rack Properties windows are displayed in the Properties window. The Properties window can contain multiple rows, one for each managed server in the rack.

The Inventory Database (Properties window) details specifications about a selected rack or selected rack components. These specifications appear when

Director discovers the systems. A rack is not discovered by Director but must be created by the Director Administrator.

Selecting a component displays these descriptions in the 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 at 120 volts (amps)
- Current at 230 volts (amps)
- Weight
- Height
- Width
- Depth
- Power cords
- Component
- Maximum processors
- Maximum cache
- Memory

Unique to servers:

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 might not appear if the component is not a rack-mountable component or if the component does not have the IBM Director Agent installed.

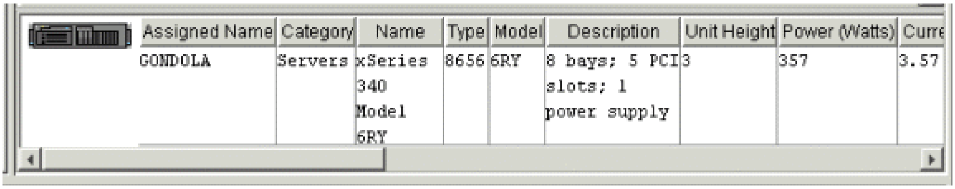

Assigned Name	Category	Name	Type	Model	Description	Unit Height	Power (Watts)	Curre
GONDOLA	Servers	xSeries 340 Model 6RY	8656	6RY	8 bays; 5 PCI slots; 1 power supply	3	357	3.57

Supported rack components

Rack Manager supports the following rack components:

- Rack enclosures
  IBM Netfinity Rack 42u cabinet Model 90
  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-inch Monocrome 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 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

Chapter 29. 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.

This chapter also provides information on obtaining ServeRAID updates, updating ServeRAID BIOS and firmware code, and configuring ServeRAID controllers.

Starting ServeRAID Manager

To start the ServeRAID Manager task, from the Director Management console, use one of the following techniques:

- Drag and drop the ServeRAID icon onto a managed system in the Group Contents pane.
- From the Group Contents pane, select a managed system. Drag and drop the managed system onto the ServeRAID icon.
- Right-click a managed system. From the context menu, select ServeRAID.

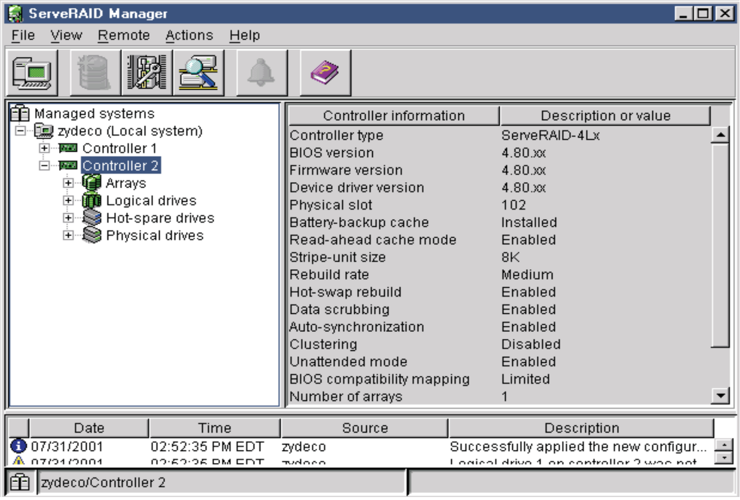

Controller information	Description or value
Controller type	ServeRAID-4Lx
BIOS version	4.80.xx
Firmware version	4.80.xx
Device driver version	4.80.xx
Physical slot	102
Battery-backup cache	Installed
Read-ahead cache mode	Enabled
Stripe-unit size	8K
Rebuild rate	Medium
Hot-swap rebuild	Enabled
Data scrubbing	Enabled
Auto-synchronization	Enabled
Clustering	Disabled
Unattended mode	Enabled
BIOS compatibility mapping	Limited
Number of arrays	1

Date	Time	Source	Description
07/31/2001	02:52:35 PM EDT	zydeco	Successfully applied the new configur...
07/31/2001	02:52:35 PM EDT	zydeco	Logical drive 1 on controller 2 was not

zydeco/Controller 2

The ServeRAID Manager window opens. The ServeRAID Manager window consists of the following components:

- Menu Bar
- Tool Bar
- Expandable Tree
- Main panel
- Event Viewer
- Status Bar

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 illustration on page 323 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
	Configure RAID
	Configure for clustering
	Scan for new or removed ready drives

Icon	Command
🔔	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.

Viewing the Event Viewer

The Event Viewer, located below the Main Panel, 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. To view the Configuration Event Detail window, double-click an event icon.

Viewing the Status Bar

The Status bar, located below the event viewer, 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 pane displays progress information.

Viewing ServeRAID controllers and subsystems

You can use ServeRAID Manager to view information about ServeRAID controllers and the ServeRAID subsystem (such as arrays, logical drives, hot-spare drives, and physical drives).

To view information, expand the ServeRAID Manager tree; then, click the relevant tree object. Detailed information about the selected device appears in the right pane.

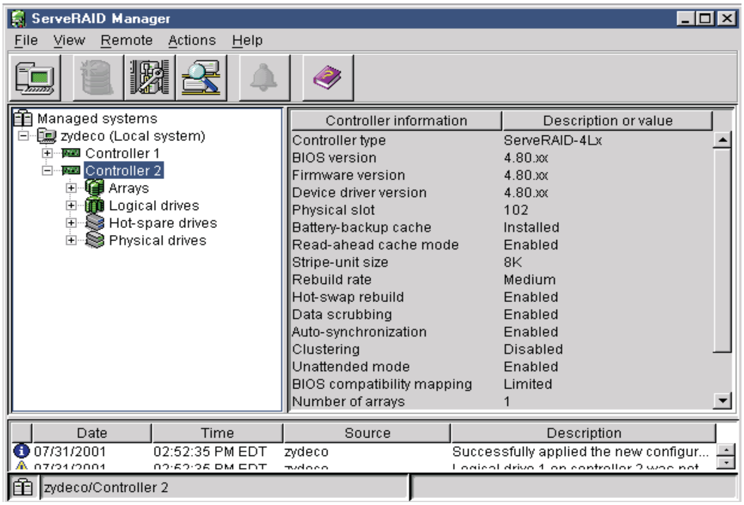

Controller information	Description or value
Controller type	ServeRAID-4Lx
BIOS version	4.80.xx
Firmware version	4.80.xx
Device driver version	4.80.xx
Physical slot	102
Battery-backup cache	Installed
Read-ahead cache mode	Enabled
Stripe-unit size	8K
Rebuild rate	Medium
Hot-swap rebuild	Enabled
Data scrubbing	Enabled
Auto-synchronization	Enabled
Clustering	Disabled
Unattended mode	Enabled
BIOS compatibility mapping	Limited
Number of arrays	1

	Date	Time	Source	Description
ℹ	07/31/2001	02:52:35 PM EDT	zydeco	Successfully applied the new configur...
⚠	07/31/2001	02:52:35 PM EDT	zydeco	Logical drive 1 on controller 2 was not

To display available actions for an item, click the item in the ServeRAID Manager tree and click **Actions.**

Using the Configuration Wizard

You can use the Configuration wizard to create up to eight arrays and up to eight logical drives for each ServeRAID controller. The Configuration wizard provides two configuration options: Express and Custom. Express configuration automatically configures your ServeRAID controller, and Custom configuration enables you to configure your controller manually. If you want to use RAID level-1E, RAID-level-5E, or RAID level-x0, you must use Custom configuration. For more information about RAID levels, refer to "Understanding RAID technology" in the IBM ServeRAID-4 Ultra160 SCSI Controller User's Reference on the IBM ServeRAID Support CD for more information

Using Express configuration

Express configuration automatically configures your ServeRAID controller. This feature does the following:

- Creates arrays by grouping together same-sized physical drives.
- Creates one logical drive per array.
- Assigns a RAID level based on the number of physical drives in an **array**:
  - An array with a single physical drive is assigned RAID level-0.
  - An array with two physical drives is assigned RAID level-1.
  - An array with three or more physical drives is assigned RAID level-5.
- Designates a hot-spare drive for the controller. If one or more arrays has four or more physical drives, the largest-sized drive from those arrays is designated the hot-spare drive.

To use Express configuration, do the following:

1. In the ServeRAID Manager tree, click the ServeRAID controller that you want to configure.
2. Click the configure RAID icon.
3. Select Express configuration.
4. Click Next. The Configuration summary window opens.
5. Review the information that is displayed in the Configuration summary window. To change the configuration, click **Modify arrays** or **Modify logical** drives.

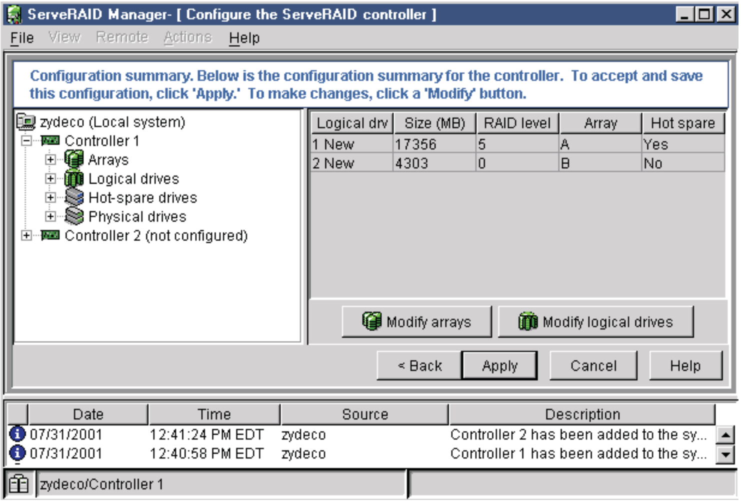

Configuration summary. Below is the configuration summary for the controller. To accept and save this configuration, click 'Apply.' To make changes, click a 'Modify' button.

zydeco (Local system)
Controller 1
Arrays
Logical drives
Hot-spare drives
Physical drives
Controller 2 (not configured)

Logical drv	Size (MB)	RAID level	Array	Hot spare
1 New	17356	5	A	Yes
2 New	4303	0	B	No

Date	Time	Source	Description
07/31/2001	12:41:24 PM EDT	zydeco	Controller 2 has been added to the sy...
07/31/2001	12:40:58 PM EDT	zydeco	Controller 1 has been added to the sy...

zydeco/Controller 1

6. Click **Apply**, then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the ServeRAID controller and on the physical drives.
7. If you have multiple controllers, do the following:
   a. From the toolbar, click the configure RAID icon.
   b. Repeat steps 3 on page 327 through step 6 on page 328 for each controller.

Using Custom configuration

Choose Custom configuration to configure your controller manually. To use Custom configuration, do the following:

1. In the ServeRAID Manager tree, click the ServeRAID controller that you want to configure.
2. Click the configure RAID icon.
3. Select Custom configuration.
4. Click Next. The Create arrays window opens. If you want to create spanned arrays, go to step 2 of “Creating Spanned arrays” on page 331.

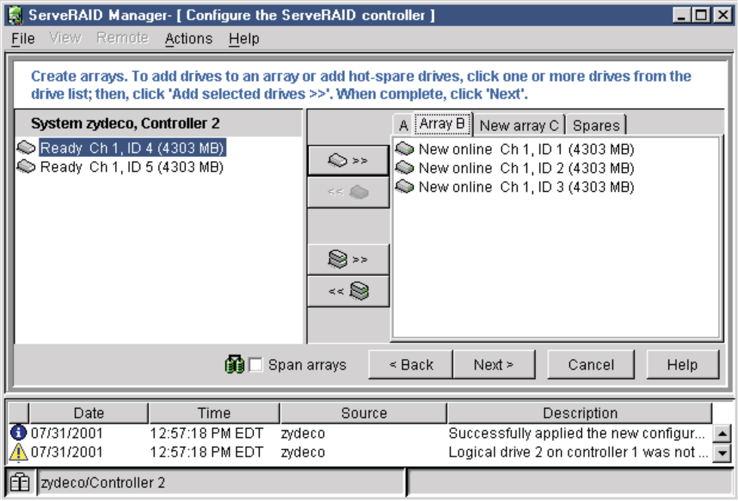
ServeRAID Manager- [Configure the ServeRAID controller]

Create arrays. To add drives to an array or add hot-spare drives, click one or more drives from the drive list; then, click 'Add selected drives >>'. When complete, click 'Next'.

System zydeco, Controller 2

Ready Ch 1, ID 4 (4303 MB)
Ready Ch 1, ID 5 (4303 MB)

New online Ch 1, ID 1 (4303 MB)
New online Ch 1, ID 2 (4303 MB)
New online Ch 1, ID 3 (4303 MB)

□ Span arrays

Date	Time	Source	Description
07/31/2001	12:57:18 PM EDT	zydeco	Successfully applied the new configur...
07/31/2001	12:57:18 PM EDT	zydeco	Logical drive 2 on controller 1 was not ...

zydeco/Controller 2

5. Click the appropriate tab in the right panel; then, from the list of ready drives, select the drives you want to move to the array.
6. Click the Add selected drives icon to add the drives to the array. You can click the Add all drives icon to move all ready drives to an array.
7. Repeat steps 4 and 5 for each additional array or hot-spare drive that you want to configure.
8. After you select the ready drives for your arrays and hot-spare drive, click Next. The Create logical drives window opens.

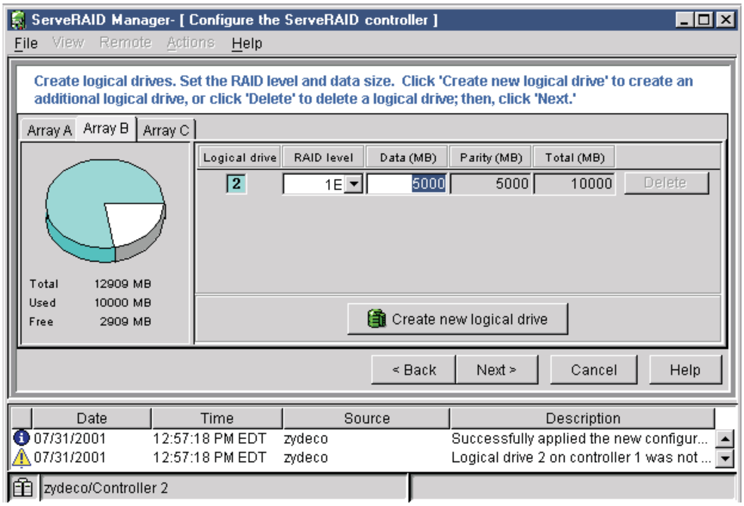

Create logical drives. Set the RAID level and data size. Click 'Create new logical drive' to create an additional logical drive, or click 'Delete' to delete a logical drive; then, click 'Next.'

Logical drive	RAID level	Data (MB)	Parity (MB)	Total (MB)
2	1E	5000	5000	10000

Total	12909 MB
Used	10000 MB
Free	2909 MB

```
Date         Time             Source   Description
07/31/2001   12:57:18 PM EDT  zydeco   Successfully applied the new configur...
07/31/2001   12:57:18 PM EDT  zydeco   Logical drive 2 on controller 1 was not ...
```

zydeco/Controller 2

9. Click the appropriate Array tab.
10. Select a RAID level from the drop-down list. (Refer to "Understanding RAID technology" in the IBM ServeRAID-4 Ultra160 SCSI Controller User's Reference on the IBM ServeRAID Support CD for more information.)

    Notes:

    a. RAID level-5E allows only one logical drive per array.

    b. If you are configuring a spanned array, you can set the RAID level only for the first logical drive you create.
11. If you do not want to use the maximum size for the logical drive, type the size in the Data (MB) field.

    Notes:

    a. You can define up to eight logical drives per controller. There are two exceptions:

       - If an array contains a logical drive assigned RAID level-5E
       - If you want to use the logical-drive migration feature

       In the above exceptions, one logical drive slot must be left free; therefore, you cannot define more than seven logical drives.

    b. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical drive is appropriate for you operating system. For more detailed information, see your operating system documentation.

c. Typically, the first logical drive defined on the first ServeRAID controller found by system BIOS during startup will be your startup (boot) drive.

12. If you have free space available and want to create additional logical drives. click **Create new logical drive**.

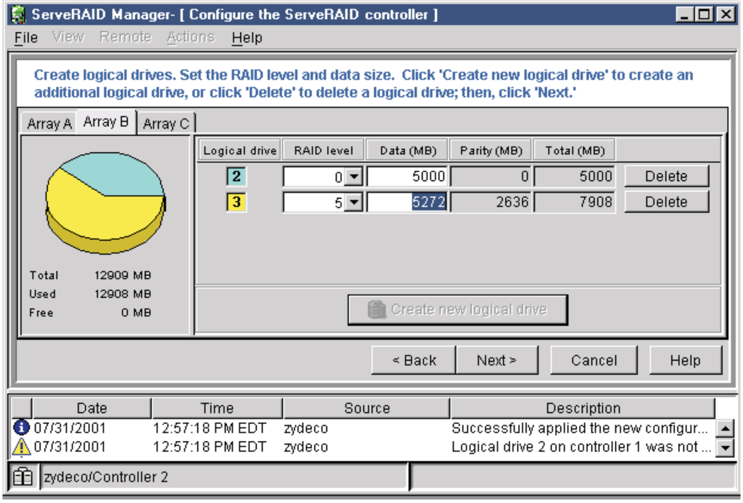

Create logical drives. Set the RAID level and data size. Click 'Create new logical drive' to create an additional logical drive, or click 'Delete' to delete a logical drive; then, click 'Next.'

Logical drive	RAID level	Data (MB)	Parity (MB)	Total (MB)	
2	0	5000	0	5000	Delete
3	5	5272	2636	7908	Delete

Total	12909 MB
Used	12908 MB
Free	0 MB

Date	Time	Source	Description
07/31/2001	12:57:18 PM EDT	zydeco	Successfully applied the new configur...
07/31/2001	12:57:18 PM EDT	zydeco	Logical drive 2 on controller 1 was not ...

zydeco/Controller 2

13. Repeat steps 9 on page 329 through step 11 for each logical drive that you want to define in this array.

14. Repeat steps 8 on page 329 through step 12 for each additional array that you want to configure.

15. Click **Next**. The Configuration summary window opens.

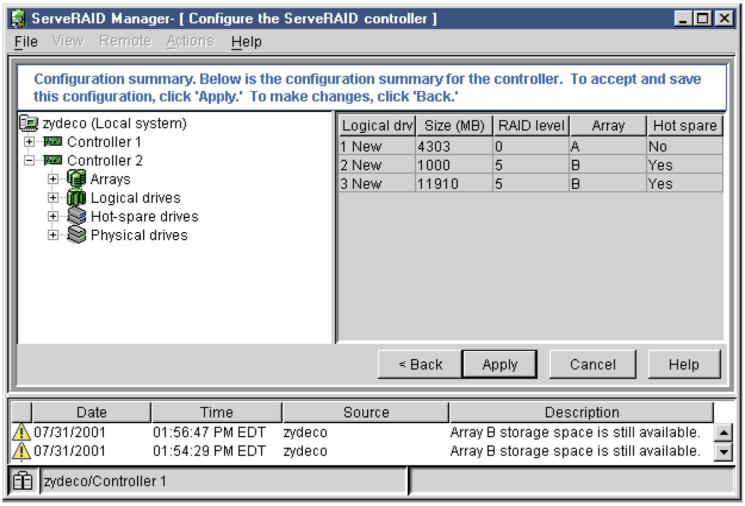

Configuration summary. Below is the configuration summary for the controller. To accept and save this configuration, click 'Apply.' To make changes, click 'Back.'

Logical drv	Size (MB)	RAID level	Array	Hot spare
1 New	4303	0	A	No
2 New	1000	5	B	Yes
3 New	11910	5	B	Yes

Date	Time	Source	Description
07/31/2001	01:56:47 PM EDT	zydeco	Array B storage space is still available.
07/31/2001	01:54:29 PM EDT	zydeco	Array B storage space is still available.

16. Review the information that is displayed in the Configuration summary window. To change the configuration, click **Back**.
17. Click Apply; then, click Yes when asked if you want to apply the new configuration. The configuration is saved in the ServeRAID controller and in the physical drives.
18. If you have multiple controllers, do the following:
    a. Click the ServeRAID controller that you want to configure.
    b. From the toolbar, click the Configure RAID icon.
    c. Repeat steps 2 , on page 328 through step 17, on page 331, for each controller.

Creating Spanned arrays

If you want to assign RAID level-x0 to an array, you must create a spanned array. For more information about spanned arrays, refer to "Understanding RAID technology" in the IBM ServeRAID-4 Ultra160SCSI Controller User’s Reference on the IBM ServeRAID Support CD for more information.

Note: Spanned arrays are supported only by IBM ServeRAID-4 Ultra160 SCSI controllers.

To create one or more spanned arrays, do the following:

1. If you have not completed step 1 through step 4 of “Using Custom configuration” on page 328, do so now.
2. Create identical arrays by doing the following:
    a. Click the **Array** tab in the right panel, then, from the list of ready drives, select the drives you want to move to the array.

b. Click the Add selected drives icon to add the drives to the array.

c. Repeat steps a and step b for each additional array that you want to configure.

   Note: To create a spanned array, the arrays to be spanned must be identical (that is, they must have the same number of physical drives).

d. Select the Span arrays check box, then, click Next. The Create spanned arrays window opens.

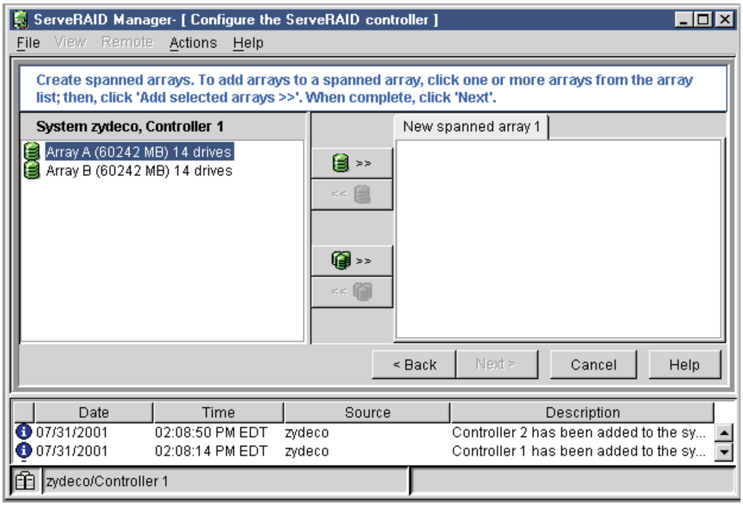

ServeRAID Manager- [Configure the ServeRAID controller]

Create spanned arrays. To add arrays to a spanned array, click one or more arrays from the array list; then, click 'Add selected arrays >>'. When complete, click 'Next'.

System zydeco, Controller 1

- Array A (60242 MB) 14 drives
- Array B (60242 MB) 14 drives

New spanned array 1

	Date	Time	Source	Description
ℹ	07/31/2001	02:08:50 PM EDT	zydeco	Controller 2 has been added to the sy...
ℹ	07/31/2001	02:08:14 PM EDT	zydeco	Controller 1 has been added to the sy...

zydeco/Controller 1

3. Create spanned arrays by doing the following:

   a. In the list of arrays, click the arrays that you want to add to your spanned array.

   b. Click the Add selected arrays icon to add the arrays to the spanned array. You can click the Add all arrays icon to move all arrays to the spanned array.

   c. To create additional spanned arrays, click the New spanned array tab in the right pane; then, repeat steps a and b.

4. Click Next; the Create Logical drives window opens. Continue with step 8 of “Using Custom configuration” on page 328.

Getting Assistance

For more information about ServeRAID Manager, consult the online help system. To start the help system, either click the **Information about this window** icon or select an item from the Help menu.

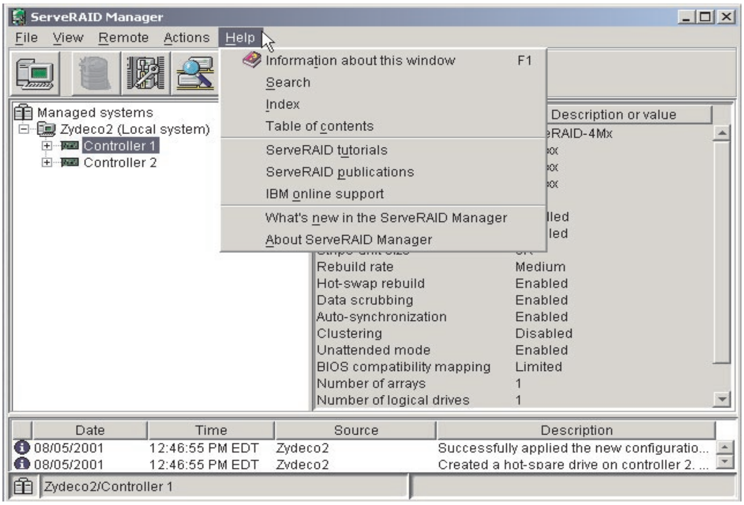

Rebuild rate	Medium
Hot-swap rebuild	Enabled
Data scrubbing	Enabled
Auto-synchronization	Enabled
Clustering	Disabled
Unattended mode	Enabled
BIOS compatibility mapping	Limited
Number of arrays	1
Number of logical drives	1

Date	Time	Source	Description
08/05/2001	12:46:55 PM EDT	Zydeco2	Successfully applied the new configuratio...
08/05/2001	12:46:55 PM EDT	Zydeco2	Created a hot-spare drive on controller 2. ...

Zydeco2/Controller 1

To learn more about the ServeRAID Manager tree objects and the actions that apply to them, use the Hints and Tips feature. Select a tree object and click **Actions →Hints and tips**. ServeRAID Assist will start, and information about the tree object will appear in the right pane of ServeRAID Manager.

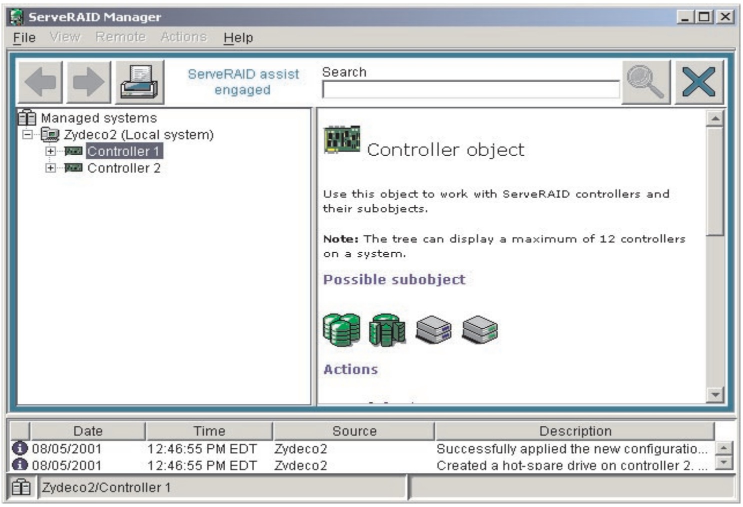

Controller object

Use this object to work with ServeRAID controllers and their subobjects.

Note: The tree can display a maximum of 12 controllers on a system.

Possible subobject

Actions

Date	Time	Source	Description
08/05/2001	12:46:55 PM EDT	Zydeco2	Successfully applied the new configuratio...
08/05/2001	12:46:55 PM EDT	Zydeco2	Created a hot-spare drive on controller 2. ...

Zydeco2/Controller 1

Chapter 30. Software Rejuvenation

The Software Rejuvenation tool is used to reduce the number and impact of unplanned outages due to software aging. The result is an increase in the reliability of managed systems. This is achieved through scheduled software rejuvenations (restarts) on each system. You can implement software rejuvenation in either of two ways: manually or automatically. For example, manually, you might elect to rejuvenate a server every Saturday night at 11 p.m. Automatic software rejuvenation is scheduled by prediction. For example, a rejuvenation is scheduled as needed automatically based on actual runtime data. In this case, resource utilization is monitored. If it is predicted that a resource will be exhausted, a rejuvenation is automatically scheduled before the time that the resource is predicted to be exhausted. Through the use of scheduling options, this automatic rejuvenation can be controlled or possibly avoided altogether through early notification of an administrator. The Software Rejuvenation program provides high availability by limiting the number and frequency of outages while ensuring that peak time availability is not compromised.

Starting the Software Rejuvenation task

To start the Software Rejuvenation program from the Director Management Console, use one of the following procedures:

- Drag-and-drop the Software Rejuvenation task icon onto a cluster or system (IBM Availability Extensions for MSCS or MSCS).
- Drag-and-drop a cluster or system onto the Software Rejuvenation task icon.
- Right-click a cluster name or a systems name in the Director Management console, and then click **Software Rejuvenation** from the menu.

The Software Rejuvenation window contains two panes. The left pane contains an expandable tree with server folders. The folders expand to display clusters and systems. The right pane contains a calendar with rejuvenation dates. The Software Rejuvenation window also contains the following components:

- Menu bar
- Calendar
- Server time

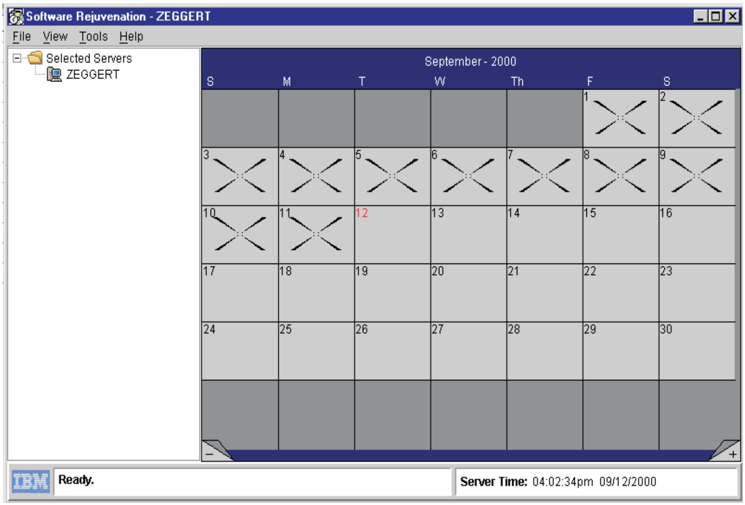

Ready.

Server Time: 04:02:34pm 09/12/2000

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 menus. The menu bar options are:

File Use this option to save a rejuvenation schedule or to close the rejuvenation calendar.

View Use this option to refresh the current month on the calendar, to view the previous month or to view the next month.

Tools

- Prediction: Use this option to configure prediction of resource exhaustion. This option is system-specific. You must first select a system.
- Trend Viewer: Use this option to view graphically actual and predicted resource utilization. This option is system-specific. You must first select a system.
- Service Rejuvenation: Use this option to manually schedule the rejuvenation of a Windows service or Linux DAEMON. This option is system-specific.
- Schedule Filter: Use this option to set Invalid Days and specify the way in which cases with Invalid Days schedule conflicts are to be handled.
- Rejuvenation Options: Use the Rejuvenation options to set parameters that control rejuvenation operations.

Help Use this option to view the Table of Contents and to learn more about Software Rejuvenation.

Calendar

Use the calendar to schedule and to edit system rejuvenations. The calendar also provides a monthly overview of scheduled system rejuvenations.

The calendar opens to the current month and year, with the current date highlighted. The lower corners of the calendar display a minus tab 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 system is scheduled for rejuvenation, a system icon is displayed on every day for which the rejuvenation is scheduled.

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.

Title bar

The title bar is located at the top of the Software Rejuvenation window. It displays the name of the program and the selected managed system.

Using Software Rejuvenation

Use the Software Rejuvenation program to selectively restart managed systems. In restarting a managed system, you rejuvenate or refresh software resources. This program not only restarts a system in a selected cluster, it also enables you to schedule restarts, which includes scheduling multiple systems on various dates.

The Software Rejuvenation program recognizes and can rejuvenate single MSCS clusters and a stand-alone system.

Scheduling systems for Software Rejuvenation

Using the Software Rejuvenation program, you can schedule rejuvenation for systems in several ways. The following sections describe these ways.

Note: The Software Rejuvenation program uses the Director Server time to validate and duplicate rejuvenation schedules and uses the server system clock to execute schedules.

The Software Rejuvenation task must be enabled to schedule a rejuvenation. In addition, for Prediction, the agent must be installed and configured on individual systems.

Scheduling a rejuvenation for a single system

To schedule rejuvenation for a single system, use the following procedure:

1. Select the system from the expandable tree in the Software Rejuvenation window.
2. Drag-and-drop the system onto the calendar date of the scheduled rejuvenation.

   The Repeat Schedule window opens.

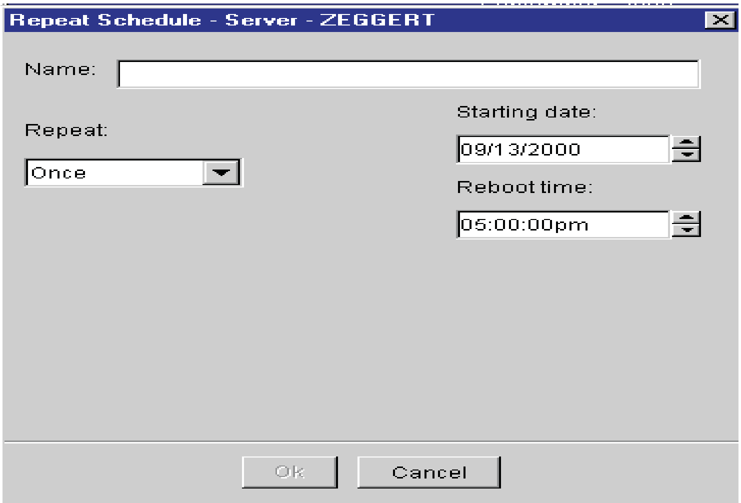
Repeat Schedule - Server - ZEGGERT

Name:

Repeat: Once

Starting date: 09/13/2000

Reboot time: 05:00:00pm

3. Type the name of the schedule. The name is user defined and is used to differentiate between multiple schedules for the same system.
4. From the **Repeat** list, select the preferred schedule type for the system.
5. In the **Starting Date** list, specify the date you want the rejuvenation to start.
6. From the **Reboot Time** list, select the rejuvenation start time.

   Note: The number of days set in the **Minimum Reboot Interval** field will override any repeat scheduling made through the calendar.
7. Click **OK**. A system 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 346.

Scheduling a single rejuvenation

To rejuvenate a system as a single instance, use the following procedure:

1. Select the system (icon) from the expandable tree in the Software Rejuvenation window.
2. Drag-and-drop the system onto the calendar date of the first scheduled rejuvenation.

   The Repeat Schedule window opens.

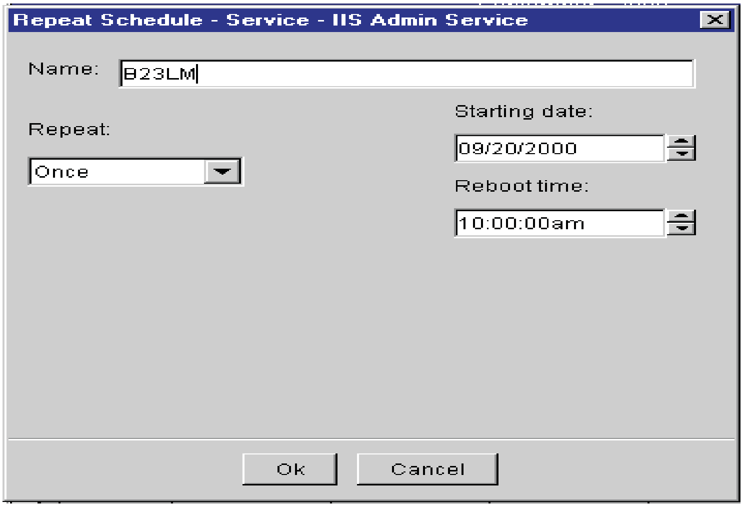

Name: B23LM

Repeat: Once

Starting date: 09/20/2000

Reboot time: 10:00:00am

3. Type the user-defined schedule name.
4. From the **Repeat** menu, select **Once** as the schedule type for the system.
5. In the **Starting date** list, select the date you want the rejuvenation to start.
6. In the **Reboot time** 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 system.
8. Press **Al+F** and **Alt+S** to accept changes.

Scheduling a daily rejuvenation

To rejuvenate a system daily, use the following procedure:

1. Select the system from the expandable tree in the Software Rejuvenation window.
2. Drag-and-drop the system onto the calendar date of the first scheduled rejuvenation.

   The Repeat Schedule window opens.

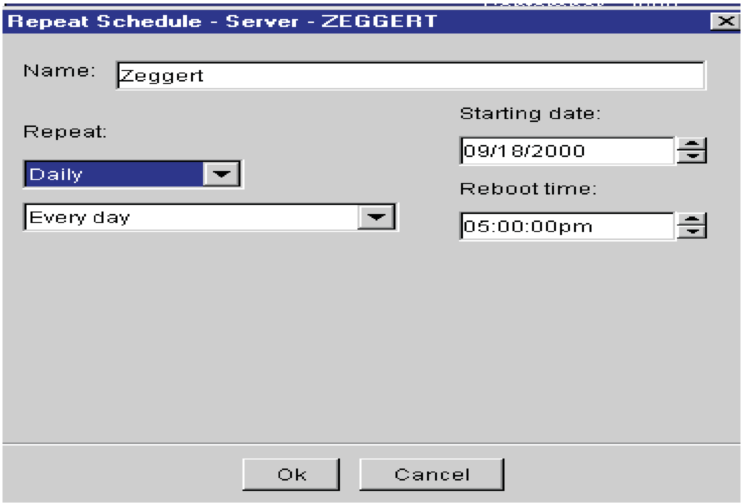

Repeat Schedule - Server - ZEGGERT

Name: Zeggert

Repeat: Daily

Every day

Starting date: 09/18/2000

Reboot time: 05:00:00pm

3. Type the user-defined schedule name.
4. From the **Repeat** list, select **Daily**.

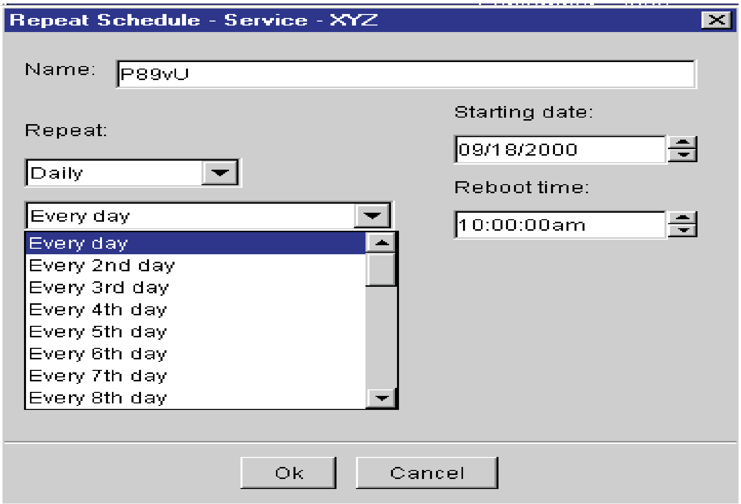

Repeat Schedule - Service - XYZ

Name: P89vU

Repeat: Daily

Every day

Every day
Every 2nd day
Every 3rd day
Every 4th day
Every 5th day
Every 6th day
Every 7th day
Every 8th day

Starting date: 09/18/2000

Reboot time: 10:00:00am

5. A list appears under **Daily**. From the list, select the rejuvenation frequency for the system.

   If you select **Every day**, the system is rejuvenated each day; while if you select **Every 2nd day**, the system 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** list, select the rejuvenation time. Time is specified in 15-minute intervals.
8. Click **OK**.
9. Press **Al+F** and **Alt+S** to accept the changes.

Scheduling a weekly rejuvenation

To rejuvenate a system as a weekly occurrence, use the following procedure:

1. Select the system from the expandable tree in the Software Rejuvenation window.
2. Drag-and-drop the system on to the calendar date of the first scheduled rejuvenation.

   The Repeat Schedule window opens.

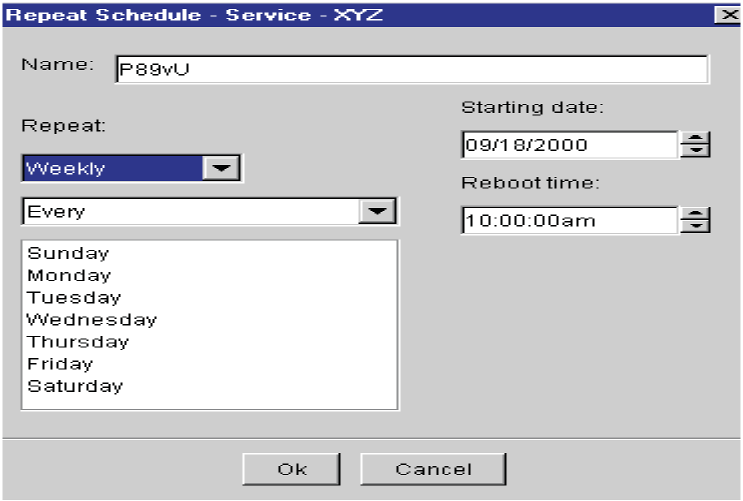

Repeat Schedule - Service - XYZ

Name: P89vU

Repeat: Weekly

Every

Sunday
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday

Starting date: 09/18/2000

Reboot time: 10:00:00am

3. Type the user-defined schedule name.
4. From the **Repeat** list, select **Weekly**.
5. A list appears under **Weekly**. From the list, 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 system each selected day of the week.
6. In the **Starting date** list, select the rejuvenation start date.
7. In the **Reboot time** list, select the rejuvenation time. The time is specified 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 system as on a specific date each month, use the following procedure:

1. Select a system from the expandable tree in the Software Rejuvenation window.
2. Drag-and-drop the system onto the calendar date of the first scheduled rejuvenation.

   The Repeat Schedule window opens.

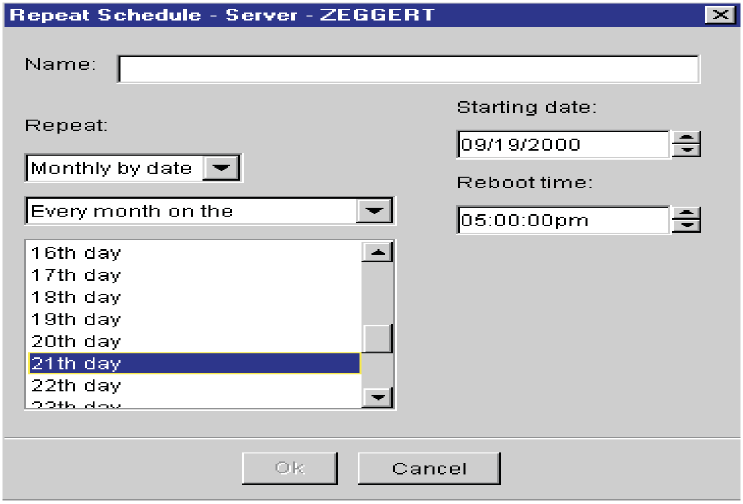

Repeat Schedule - Server - ZEGGERT

Name:

Repeat:

Monthly by date

Every month on the

21th day

Starting date: 09/19/2000

Reboot time: 05:00:00pm

3. From the **Repeat** list, select **Monthly by Date.**
4. A list appears under **Monthly by date**. From the list, select the day of the month for the system rejuvenation.
5. In the **Starting date** list, select a system rejuvenation date.
6. From the **Reboot time** list, select a rejuvenation time. The time is specified 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 system on specific days each month, use the following procedure:

1. Select the system from the expandable tree in the Software Rejuvenation window.

2. Drag-and-drop the system on to the calendar date of the first scheduled rejuvenation.

   The Repeat Schedule window opens.

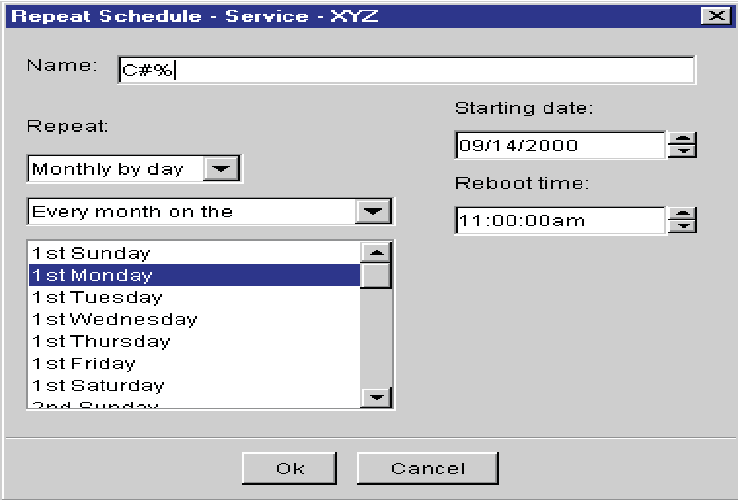

Repeat Schedule - Service - XYZ

Name: C#%

Repeat: Monthly by day

Every month on the

1st Sunday
1st Monday
1st Tuesday
1st Wednesday
1st Thursday
1st Friday
1st Saturday
2nd Sunday

Starting date: 09/14/2000

Reboot time: 11:00:00am

3. Type a user-defined schedule name.
4. From the **Repeat** list, select **Monthly by day**.

   New fields appear in the current window.
5. A new list appears under **Monthly by day**. Select the day of the month for system rejuvenation.
6. From the **Starting date** list, select the start date for the system rejuvenation.
7. From the **Reboot time** list, select the start time for the rejuvenation. The time is specified in 15-minute intervals.
8. Click **OK**.
9. The **Skipping Some Dates** window might appear. Click **OK**.
10. Press **Alt+F** and **Alt+S** to accept changes.

Editing scheduled rejuvenations

Use the Software Rejuvenation program to edit scheduled rejuvenations at the cluster level and the system level.

Editing a Scheduled Rejuvenation

To change the date, time, or frequency of scheduled system rejuvenations, use the following procedure:

1. From the Director Management console, drag-and-drop the system onto the Software Rejuvenation task icon.

   The Software Rejuvenation window opens.

   The calendar displays the current month and year. The current date is highlighted. A system icon is displayed on the calendar day for which a system is scheduled for rejuvenation. If multiple rejuvenations are scheduled on the same day, the icons are shown cascaded.

2. In the Calendar, right-click the system icon you want to edit.
   The Edit Schedule menu appears.

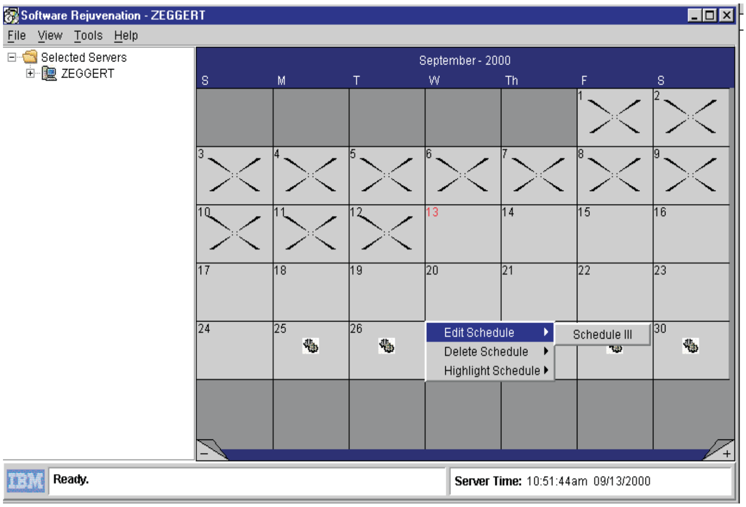

Ready.

Server Time: 10:51:44am 09/13/2000

3. Click **Edit schedule< Schedule<Schedule 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.

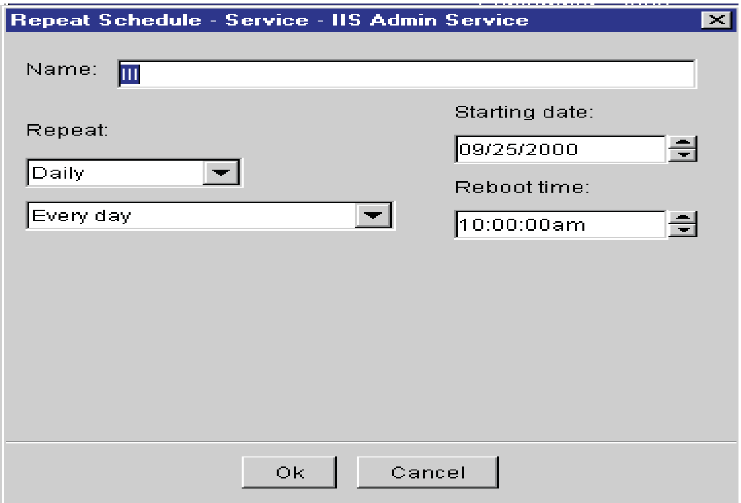

Repeat Schedule - Service - IIS Admin Service

Name:

Repeat:

Daily

Every day

Starting date:

09/25/2000

Reboot time:

10:00:00am

4. From the **Repeat** Schedule window, edit the rejuvenation schedule settings.
5. Click **OK**.
6. Press **Alt-F** and **Alt-S** to accept changes.

Removing a rejuvenation schedule

If a schedule is repeating and covers more than one date, removing the schedule from one date on the calendar removes it from all dates covered by that schedule. To remove a rejuvenation schedule, use the following procedure:

1. From the Director Management console, drag-and-drop the system onto the Software Rejuvenation task icon.

   The Software Rejuvenation window opens.

   The calendar displays the current month and year. The current date is highlighted. A system icon is displayed on every day that a system is scheduled for rejuvenation. If multiple systems are scheduled for rejuvenation on the same day, the system icons are shown cascaded on the specified date.
2. In the Calendar, right-click the system you want to remove. The **Edit** menu appears. **Select Delete →Schedule<ScheduleName>Verify Remove**. The Verify Remove window opens. If the schedule being deleted is a repeating schedule (that is, it is a schedule that applies to more that one day), removing the schedule from one day will cause that schedule to be removed from all days. If the system rejuvenation that you are removing from the calendar is scheduled to occur only once, this message is not displayed.

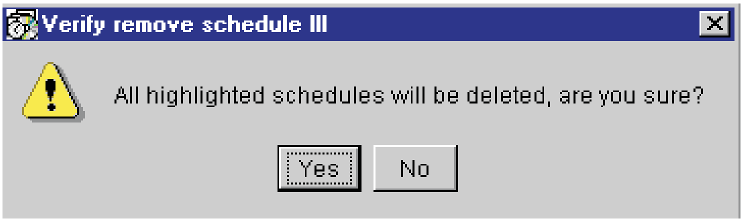

All highlighted schedules will be deleted, are you sure?

3. Click **Yes**.
4. 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 selected system.
Ctrl+V	Paste a copied system on a selected calendar date.
Ctrl+E	Access the Repeat Schedule window for the highlighted schedule.
Ctrl+D	Delete a selected scheduled.
Ctrl+H	Select (highlight) a schedule.

Rejuvenation options

Use the Rejuvenation Options window to set parameters that rejuvenate operations on a cluster level. To set parameters for cluster availability checks, to disable the rejuvenation logic, or to set a minimum reboot interval, do the following:

1. From the Software Rejuvenation window, click **Tools →Rejuvenation Options**.

   The Rejuvenation Options window opens.

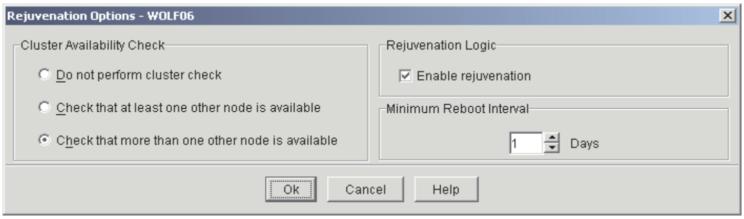

Rejuvenation Options - WOLF06

Cluster Availability Check

○ Do not perform cluster check

○ Check that at least one other node is available

● Check that more than one other node is available

Rejuvenation Logic

☑ Enable rejuvenation

Minimum Reboot Interval

1 Days

2. The Cluster Availability Check setting specifies the rules for rejuvenating a system that is a member of a cluster. Rejuvenation occurs if the systems in the cluster satisfy the selected criteria. To set the Cluster Availability Check option, use one of the following procedures:
   - Click **Do not perform cluster check** if you do not want the rejuvenation program to check the availability of other systems in the cluster.
   - Click **Check that at least one other node is available** if you want the rejuvenation program to check for at least one other available system in the cluster. If another cluster member is not available, then the rejuvenation will not occur.
   - Click **Check that more than one other node is available** if you want the rejuvenation program to check for more than one other available system in the cluster. If not, the rejuvenation will not occur. If the cluster has only two nodes, this option is equivalent to Check for One.
3. Select **Enable** or **Disable Rejuvenation**. This setting is maintained at the Director server and applies to all rejuvenations scheduled through that server.
4. In the **Minimum Reboot Interval** field specify the number of days that must elapse between rejuvenations.

   For example, if 3 days is selected, a system will not be rejuvenated again in a period of less than 3 days after the previous rejuvenation. Selecting 0 days allows a system 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.

Predictive rejuvenation

Resource exhaustion for a system can be predicted based on trends in resource utilization. When resource exhaustion is predicted, a notification is sent and a rejuvenation can be automatically scheduled. Before prediction can be started, it must be configured.

To configure a mode of operation, use the following procedure:

1. From the Software Rejuvenation window, select a cluster or single system, and click **Tools→ Prediction→ Configuration Wizard** to start the

configuration wizard. The Modify Configuration Forecasting Data window opens.

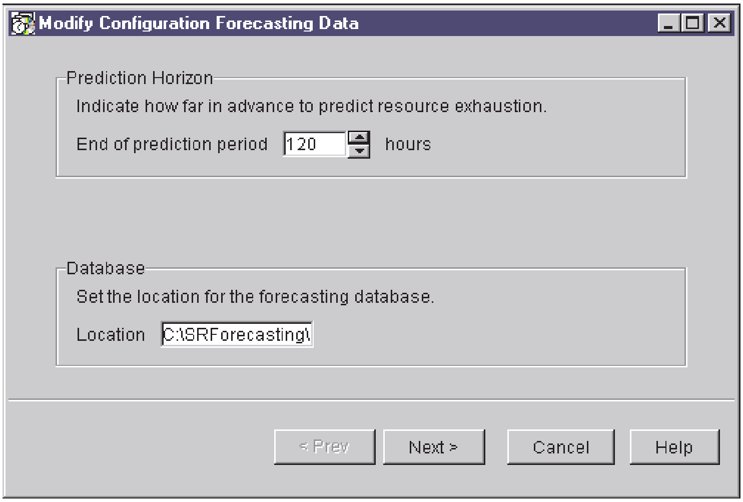

Modify Configuration Forecasting Data

Prediction Horizon

Indicate how far in advance to predict resource exhaustion.

End of prediction period 120 hours

Database

Set the location for the forecasting database.

Location C:\SRForecasting\

2. Enter the Prediction period and Database location.
3. Click **Next**. The Modify Configuration Notification and Scheduling window opens.

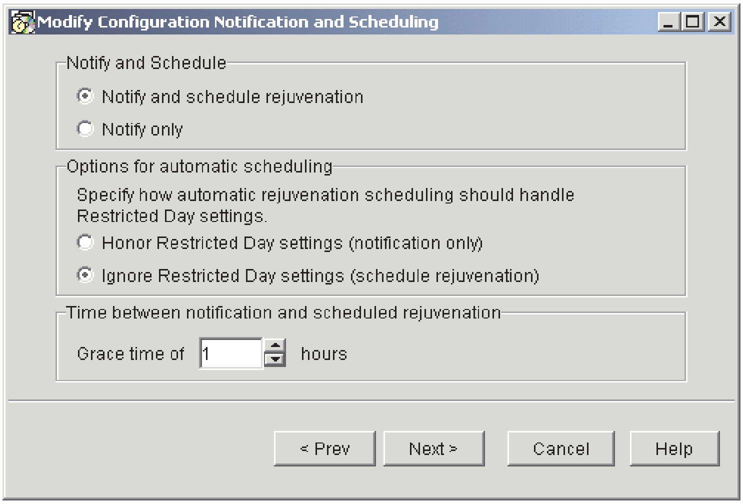

Modify Configuration Notification and Scheduling

Notify and Schedule

● Notify and schedule rejuvenation

○ Notify only

Options for automatic scheduling

Specify how automatic rejuvenation scheduling should handle Restricted Day settings.

○ Honor Restricted Day settings (notification only)

● Ignore Restricted Day settings (schedule rejuvenation)

Time between notification and scheduled rejuvenation

Grace time of 1 hours

4. Select the following information:
   - Notify and Schedule
   - Options for automatic scheduling

- Time between notification and scheduled rejuvenation

5. Click **Next**. The Modify Configuration Action Plan window opens.

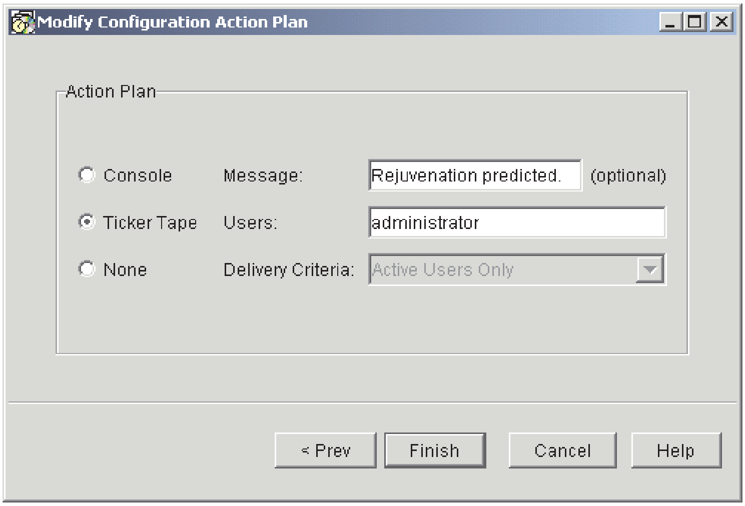

Action Plan

○ Console Message: Rejuvenation predicted. (optional)

● Ticker Tape Users: administrator

○ None Delivery Criteria: Active Users Only

6. Select **Console**, **Ticker Tape**, or **None**. If Console or Ticker Tape is selected, enter the names of users to be notified and optional message text.
7. Select the Delivery Criteria.
8. Click **Finish**.

Using the schedule filter

You can specify that rejuvenations should not occur on certain days of the week.

To use the Schedule Filter, use the following procedure:

1. From the Software Rejuvenation window, click **Tools**→ **Schedule Filter**. The Schedule Filter window opens.

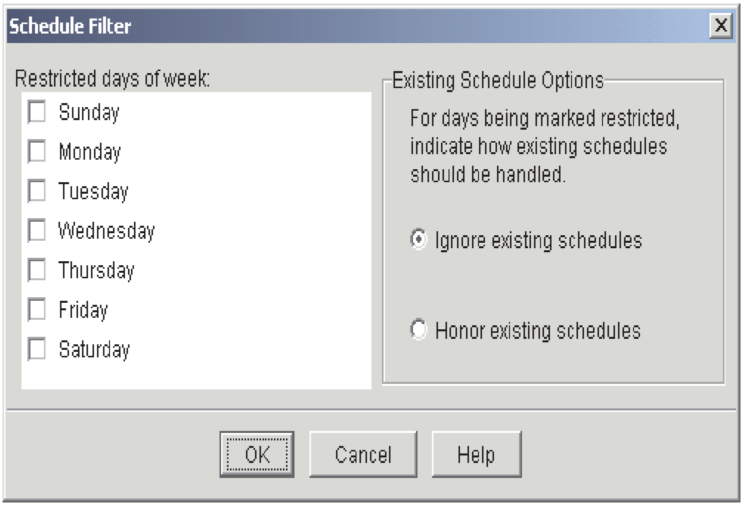
Schedule Filter

Restricted days of week:

- □ Sunday
- □ Monday
- □ Tuesday
- □ Wednesday
- □ Thursday
- □ Friday
- □ Saturday

Existing Schedule Options

For days being marked restricted, indicate how existing schedules should be handled.

● Ignore existing schedules

○ Honor existing schedules

2. Select the restricted days of the week for a rejuvenation. Under **Existing Schedule Options**, select **Ignore existing schedule** or **Honor existing schedule**.
3. Click **OK**.

Using the Trend Viewer

With the Trend Viewer you can graphically view the mathematics involved in the predictive analysis for a given system.

To view the trending information for a system that has been configured to run prediction, use the following procedure:

1. From the Software Rejuvenation window, select one or more systems.
2. Highlight the system and select **Tools** →**Trend Viewer**.
   The Trend Viewer window opens.

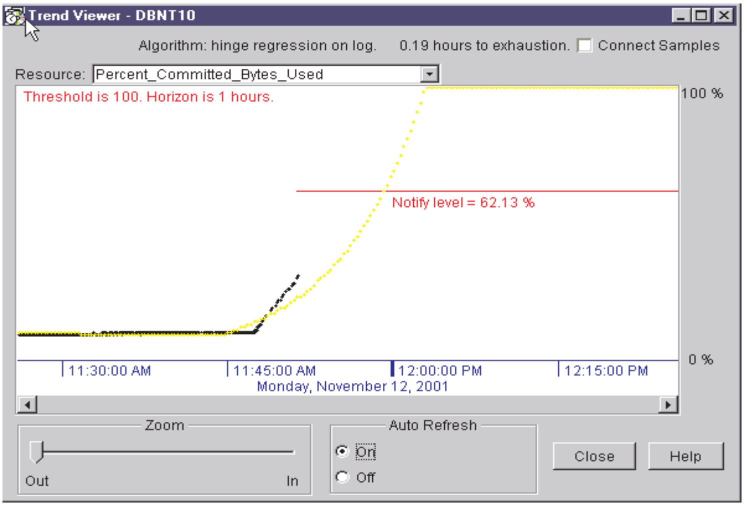

Algorithm: hinge regression on log. 0.19 hours to exhaustion. □ Connect Samples

Resource: Percent_Committed_Bytes_Used

Threshold is 100. Horizon is 1 hours.

Auto Refresh

● On

○ Off

3. Select **Auto Refresh On** to enable a continuous refresh of the displayed data or **select Auto Refresh Off** to freeze the display.

The Trend Viewer Window

The Trend Viewer Window consists of the following options:

- **System Name**: The name of the system whose data is being displayed.
- **Time Axis Labels**: The Trend Viewer calibrates the data and displays accurate axis labels. The horizontal x-axis is date and time. You select the granularity of the x-axis for displaying time. When the Zoom slider is moved farthest toward Out, the maximum time span is displayed. When the Zoom slider is moved farthest toward In, the window shows a more narrow view of the data.
- **Resource Axis Labels**: The vertical y-axis calibrates as a percentage of the resource used. 100% is at the top and 0% is at the bottom.
- **Auto Refresh Controls**: There are two buttons (**Auto. Refresh On** and **Auto. Refresh Off**).
- **Trend Area**: The collected sample points are drawn in black and the predicted data points are drawn in green. When a trend is found that looks like it will exhaust the resource within the horizon period, it will draw the prediction line in orange and draw a Notification line in red showing the percentage at which the resource will be considered exhausting. If and when the collected sample points reach this Notification line, then the prediction line will be drawn in red and two resource exhaustion events will be logged containing the hours to exhaustion that is displayed at the top of the window and the top consumers of the exhausting resource. Remember there is a Grace time that allows you to take action upon notification, prior to the system being restarted.

- **Resource**: With the Resource selection, you can decide which system resource you want the Trend Viewer to graph and trend on.
- **Connect samples**: Checking this box allows the collected and predicted data points to be connected with lines.

Using Service Rejuvenation

Service Rejuvenation enables you to manually schedule the rejuvenation of a Windows service or a Linux DAEMON. The procedure is similar to scheduling the manual rejuvenation of a system.

To schedule the rejuvenation of a Windows service or Linux DAEMON, use the following procedure:

1. Select **Tools**→ **Service Rejuvenation**.
   The Service Rejuvenation window opens.

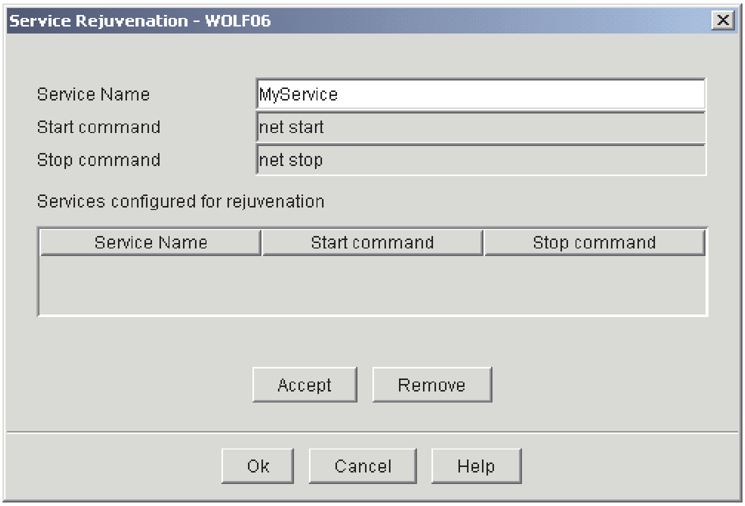

Service Rejuvenation - WOLF06

Service Name: MyService
Start command: net start
Stop command: net stop

Services configured for rejuvenation

Service Name	Start command	Stop command

2. In the **Service name** field, type the name of a Windows service or Linux DAEMON.
3. In the **Start Command** field, type the command used to start the DAEMON. (For Windows this will already be filled with Net Start and cannot be changed.)
4. In the Stop Command field, type the command use to stop the DAEMON. (In Windows, this will already be filed with **Net Stop** and cannot be changed.)
5. Click **Accept**. The service or DAEMON name, start command, stop command will appear in the list of services configured for rejuvenation. A maximum of five services or DAEMONs can be configured.
6. Click **OK** to complete configuration.

Removing a Service Rejuvenation

To remove a Service Rejuvenation, use the following procedure:

1. Click **Tools →Service Rejuvenation**.
2. From the **List of Configured Services** pane, select the service or DAEMON that you want to remove. Click **Remove**.
3. Click **OK**.

Closing the Software Rejuvenation program

To close the Software Rejuvenation program, click **File →Close**. If you did not save your scheduled event or had nothing to save, the "Verify application Close" window opens. Click **No** and go back and save your scheduled events using **File →Save**, and then **File →Close**.

Creating action plans

The ICSM and Software Rejuvenation tools add new rejuvenation event filters to the Director Management console. Using the Events Action Plan Builder, you can create filters for specific rejuvenation events and further qualify the events using the Extended Attributes option.

To create an event action, use the following procedure:

1. From the Director Management console, click **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 window opens.

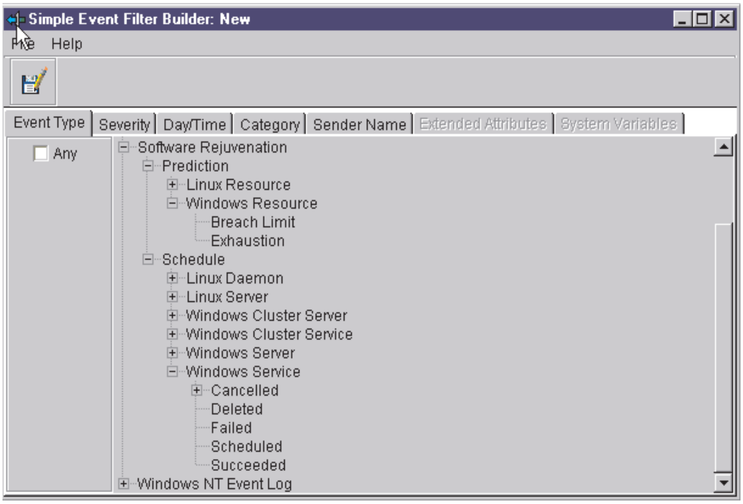

4. On the **Event Type** page, clear the **Any** check box.
5. Click **Software Rejuvenation** to expand the directory tree.
6. Under **Prediction** or under **Schedule**, select an action to create an event schedule.
7. Click **File Save As**.
8. In the **Event Filter** field, type a descriptive name for the event, and then click **OK**.

The new filter lists is shown in the **Event Filters** pane of the **Event Action Plan Builder** window.

Viewing Director Event Log entries

The following log entries are recorded by Software Rejuvenation. The entries below reflect entries originating from a Windows agent or Linux agent.

Software Rejuvenation.Schedule.Windows Cluster Server.Failed

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> failed with return code <x>."
Severity =Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Failed

"Software Rejuvenation <schedule name> for service <service name> on node <node name> failed with return code <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Server.Failed

"Software Rejuvenation <schedule name> for node <node name> failed with return code <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Linux Server.Failed

"Software Rejuvenation <schedule name> for node <node name> failed with return code <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Service.Failed

"Software Rejuvenation <schedule name> for service name> on node <node name> failed with return code <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Failed

"Software Rejuvenation <schedule name> for service <service name> on node <node name> failed with return code <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Succeeded

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> completed with return code <x>"
Severity=Harmeless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Succeeded

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> completed with return code <x>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Succeeded

"Software Rejuvenation <schedule name> for node <node name> completed with return code <x>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Server.Succeeded

"Software Rejuvenation <schedule name> for node <node name> completed with return code <x>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Succeeded

"Software Rejuvenation <schedule name> for node <node name> completed with return code <x>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Succeeded

"Software Rejuvenation <schedule name> for node <node name> completed with return code <x>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Scheduled

"Scheduled Software Rejuvenation <schedule name> of service <service name> on node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Scheduled

"Scheduled Software Rejuvenation <schedule name> for node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Server.Scheduled

"Scheduled Software Rejuvenation <scheduled name> for node <node name>."
Severity=Harmless Category=Alert

Sosftware Rejuvenation.Schedule.Windows Service.Scheduled

"Scheduled Software Rejuvenation <scheduled name> of service <service name> on <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Scheduled

"Scheduled Software Rejuvenation <schedule name> of service <service name> on node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server. Deleted

"Deleted Software Rejuvenation schedule <schedule name> for node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Deleted

"Deleted Software Rejuvenation schedule <schedule name> for service <service name> on node <node name>."

Software Rejuvenation.Schedule.Windows Server.Deleted

"Deleted Software Rejuvenation schedule <schedule name> for node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Server.Deleted

"Deleted Software Rejuvenation schedule <schedule name> for node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Deleted

"Deleted Software Rejuvenation schedule <schedule name> for service <service name> on node <node name>."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedulee.Linux Daemon.Deleted

"Deleted Software Rejvenation schedule <schedule name> for service <service name> on node <node name>."
Serverity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Cancelled. Node State

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because node state was <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Cancelled.Node State

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because node state was <x>."

Software Rejuvenation.Schedule.Windows Server.Cancelled.Node State

"Software Rejuvenation <schedule name> for node <node name> cancelled because node state was <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Linux Server.Cancelled.Node State

"Software Rejuvenation <schedule name> for node <node name> cancelled because node state was <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Service.Cancelled.Node State

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because node state was <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Cancelled.Node State

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because node state was <x>."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service. Cancelled.Peer State

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because peer node <node name> was in state <x>."

Software Rejuvenation.Schedule.Windows Server.Cancelled.Peer State

"Software Rejuvenation <schedule name> for node <node name> cancelled because peer node is unavailable."
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Cancelled. No Peers

(text.no.peers)
Severity=Warning Category=Alert

Software Rejuvenation.Schedule. Windows Cluster Service.Cancelled.No Peers

(text.no.peers)
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Server.Cancelled. No Peers

(text.no.peers2)
Severity=Warning Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Cancelled.Minimum Reboot Interval

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because it was rejuvenated more

recently than the minimum specified interval of <x> days."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Cancelled.Minimum Reboot Interval

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because it was rejuvenated more recently than the minimum specified interval of <x>days."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Cancelled.Minimum Reboot Interval

"Software Rejuvenation <schedule name> for node <node name> cancelled because it was rejuvenated more recently thant the minimum specified interval of <x> days."
Severity=Harmless Category=Alert

Software Rejuveation.Schedule.Linux Server.Cancelled.Minimum Reboot Interval

"Software Rejuvenation <scheduled name> for node <node name> cancelled because it was rejuvenated more recently than the minimum specified interval of <x> days."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Cancelled. Minimum Reboot Interval

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> cancelled because it was rejuvenated more recently than the minimum specified interval of <x>days."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Server.Cancelled.Missed

"Software Rejuvenation <scheduled name> for node <node name> in cluster <cluster name> at <date> was missed because target server was unavailable."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Cluster Service.Cancelled.Missed

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> at <date> was missed because target server was unavailable."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Cancelled.Missed

"Software Rejuvenation <schedule name> for node <node name> at <date> was missed because target server was unavailable."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Server.Cancelled.Missed

"Software Rejuvenatoin <schedule name> for node <node name> at <date> was missed because target server was unavailable."
Serverity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Cancelled.Missed

"Software Rejuvenation <schedule name> for node <node name> in cluster <cluster name> at <date> was missed because target server was unavailable."
Serverity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Cancelled.Restricted

"Software Rejuvenation <schedule name> for node <node name> cancelled because the day was restricted."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Server.Cancelled.Restricted

"Software Rejuvenation <schedule name> for node <node name> cancelled because the day was restricted."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Cancelled.Restricted

"Software Rejuvenation <schedule name> for service <service name> on node <node name> cancelled because the day was restricted."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Cancelled.Restricted

"Software Rejuvenation <schedule name> for service <service name> on node <node name> cancelled because the day was restricted."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Server.Cancelled.Disabled.

"Software Rejuvenation <scheduled name> for node <node name> cancelled because the Rejuvenation Logic is diabled."
Severity=Harmless Catergory=Alert

Software Rejuvenation.Schedule.Linux Server.Cancelled.Diabled

"Software Rejuvenation <schedule name> for node <node name> cancelled because the Rejuvenation Logic is disabled."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Windows Service.Cancelled.Disabled

"Software Rejuvenation <schedule name> for service <service name> on node <node name> cancelled because the Rejuvenation Logic is disabled."
Severity=Harmless Category=Alert

Software Rejuvenation.Schedule.Linux Daemon.Cancelled.Disabled

"Software Rejuvenation <schedule name> for service <service name> on node <node name> cancelled because the Rejuvenation Logic is disabled."
Severity=Harmless Category=Alert

Software Rejuvenation.Prediction.Windows Resource.Exhaustion

"<Resource type> will exhaust in <x> days <x> hours <x> hours <x> minutes and <x> seconds from <date time> (system date/time."
"Top <resource type> consumers [(process) pid size]:<process pid size> <...>
Severity=Critical Category=Alert

Software Rejuvenation.Prediction.Linux Resource.Exhaustion

"<Resource type> will exhaust in <x> days <x> hours <x> minutes and <x> seconds from <date time> (system date/time)."
"Top <resource type> consumers [(process) pid size]: <process pid size> <...>
Severity=Critical Category=Alert

Software Rejuvenation.Predication.Windows Server.Near Limit

"Resource <resource type> exceeded 80% of its threshold at <date time>."
Severity=Critical Category=Alert

Software Rejuvenation.Prediction.Linux Server.Near Limit

"Resource <resource type> exceeded 80% of its threshold at <date time>."
Severity=Critical Category=Alert

Software Rejuvenation.Predication.Windowsw Server.Breach Limit

"Resource <resource type> exceeded its threshold at <date time>."
Severity=Critical Category=Alert

Software Rejuvenation.Predication.Linux Server.Breach Limit

"Resource <resource type> exceededd its threshold at <date time>."
Severity=Critical Category=Alert

Chapter 31. System Availability

System Availability is used to analyze the availability of a system or a group of systems. It can be used to provide statistics on the availability of large sets of systems. Additionally it can be used as a means to graphically prove that Software Rejuvenation improves system availability.

Starting the System Availability task

To start the System Availability task, from the Director Management Console, drag-and-drop the System Availability icon onto a system or a group of systems. After the System Availability window opens, for a single system, it displays the system availability statistics for that system. For a group of systems, it displays the combined averages of the group.

The graphical interface in System Availability makes it easy for you to create, delete, and view availability graphs.

Viewing the menu

The menu bar contains the following options: File, View and Help.

The File menu contains the Set Time and Exit options. Set Time is used for defining a specific interval for the System Availability Analysis.

The View menu provides the following list of options: Detach View, Distribution of System Outages, Distribution of System Uptime, Availability Report, System Outages by Day of Week and System Outages by Hour of Day.

The tool bar contains a menu of system availability options for an assigned system or systems, a **Graphic** button, a **Report** button and a **Detach** button.

Using hover help

Hover help displays a detailed explanation of a section when you move the cursor over a section of a pie chart or bar chart.

Using System Availability

The System Availability window consists of two panes: a navigational area (left pane) and a work area (right pane). Use the navigational area to drill the reporting of a group down one or more systems. Click any node on the tree and use Ctl-click on other nodes to add them to the report. Selecting the root node, All Systems, will deselect all the leaf nodes and include every system in the report. Click on a system, under All Systems to display a graph or report of that system.

The work area is the largest area of the System Availability window. For each server or group of servers, you can display one of four possible graphs and a report:

- Distribution of System Outages (default)
- Distribution of System Uptime
- System Outages by Day of Week
- System Outages by Hour of Day

The System Availability report can be viewed in a graphical view or report view. Select a view from the menu bar and click **View** or from the list, select a report or graph.

Using the Detailed List of Records

The Detailed List of Records shows the entire System Availability operational record applied for a section of a pie chart or bar chart. There are two ways to view the Detailed List of Records:

- Double-click a section of a pie chart or bar chart
- Right-click a section of a pie chart or bar chart and click **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
- Event Type (Uptime, Planned or unplanned outage)

Detaching a 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. You can compare/contrast different system availability views by using the detach view function.

Click the **X** in the upper-right corner of the window to close the detached window. Closing the detached window will not close the System Availability task.

Creating a report

System Availability creates and maintains a database for storing the availability data. This provides availability data since the last time availability was performed on a given agent in the event an agent is down. Additionally, the database will provide the archival data if one or more of the agent event /message logs is cleared. The availability data is collected to the database on demand when you activate the System Availability task on one or more agents. The console receives all the data from the database and any new data from the agent since the last report. There is a possibility that availability data on Windows systems could be lost if the event log is erased prior to executing System Availability on that system. If an agent is down, the database will be able to provide availability information only since the last time the System Availability task was performed for that agent which is stored in the database.

Viewing the System Outages graph

The Distribution of System Outages graph represents the percentage of various duration times a system or systems were unavailable following a system outage. It indicates downtime of a system or systems.

The calculation is performed by measuring the duration time between stop and start events in the Windows NT system log or Linux messages file, specifically event (6006) stop time to event (6005) start time. The Time Marks measurements are shown in the following table:

Time Marks		
30 minutes	30*60	1800s
60 minutes	60*60	3600s
3 hours	3*60*60	10,800s
1 day	24*60*60	86,400s
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)

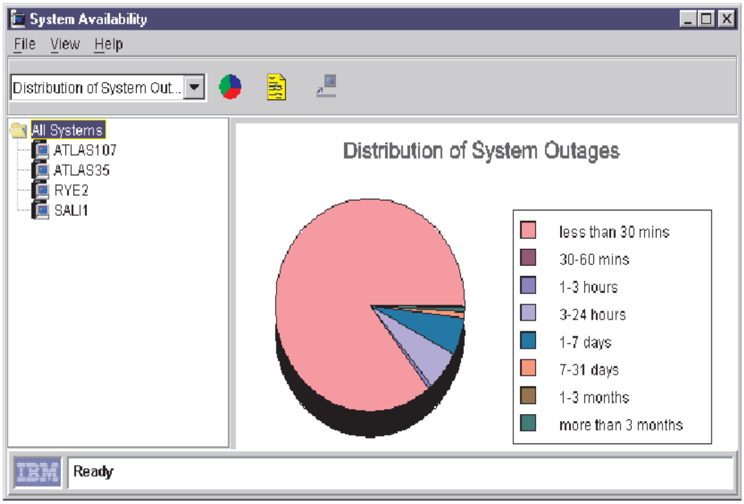

Differentiating planned and unplanned outages

There are two ways to differentiate Event types (planned and unplanned outages):

- Color: on bar charts you may see two colors used on each vertical bar (See Viewing the System Outages by Day of Week graph).
- Hover help

Viewing the System Uptime graph

The Distribution of System Uptime graph represents the analysis of system availability and group availability. It shows percents of various time duration that a system or systems were available.

The calculation is performed by measuring the duration time between start and stop events in the Windows NT system log or Linux messages file, specifically event(6005) start Event Log time to event (6006) stop Event Log time for Windows.

The System Uptime time marks and calculations are the same as the Distribution of System Outages.

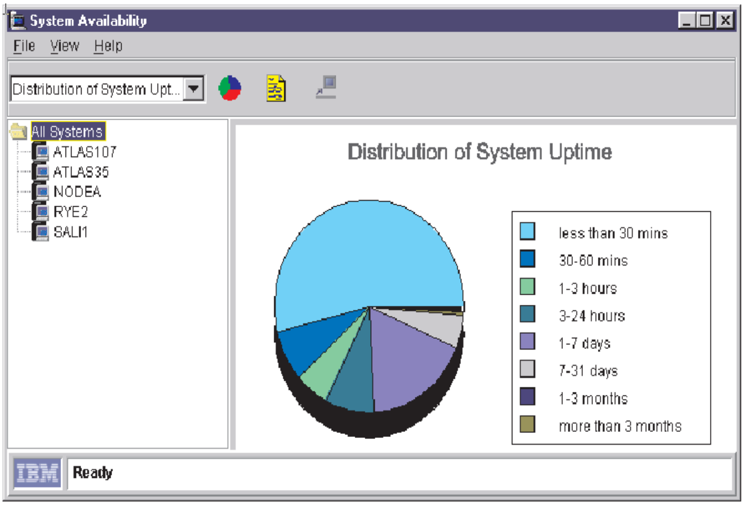

Viewing the System Availability Report View

The report shows some statistics and uses rows and columns to show a snapshot of system availability. It provides measurements for your systems or clusters.

The System/Cluster Availability measurements are:

- System(s) name: Name of the system from which statistics are gathered.
- Date report from: Starting date of the report.
- Date report to: Ending date of the report.
- 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 restarts: Count of startup events (6005).
- Mean time between outages: Total uptime/Total restarts.
- Average days analyzed per system: Sum of analyzed days for all systems divided by the number of systems.
- Number of systems restarted: The number of systems in this report that were restarted.
- Mean time between unplanned outages: Sum of Total uptime/total unplanned restarts(6008). An unplanned outage is an unexpected downtime where there was an unorderly shutdown of the system.

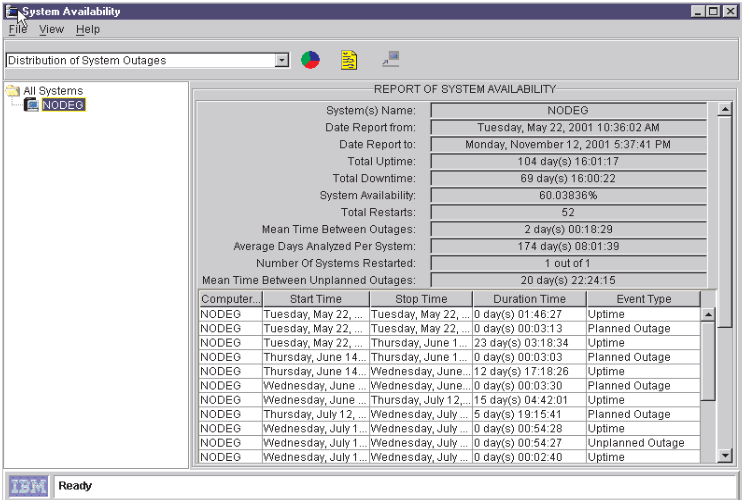

REPORT OF SYSTEM AVAILABILITY

System(s) Name: NODEG
Date Report from: Tuesday, May 22, 2001 10:36:02 AM
Date Report to: Monday, November 12, 2001 5:37:41 PM
Total Uptime: 104 day(s) 16:01:17
Total Downtime: 69 day(s) 16:00:22
System Availability: 60.03836%
Total Restarts: 52
Mean Time Between Outages: 2 day(s) 00:18:29
Average Days Analyzed Per System: 174 day(s) 08:01:39
Number Of Systems Restarted: 1 out of 1
Mean Time Between Unplanned Outages: 20 day(s) 22:24:15

Computer...	Start Time	Stop Time	Duration Time	Event Type
NODEG	Tuesday, May 22, ...	Tuesday, May 22, ...	0 day(s) 01:46:27	Uptime
NODEG	Tuesday, May 22, ...	Tuesday, May 22, ...	0 day(s) 00:03:13	Planned Outage
NODEG	Tuesday, May 22, ...	Thursday, June 1...	23 day(s) 03:18:34	Uptime
NODEG	Thursday, June 14...	Thursday, June 1...	0 day(s) 00:03:03	Planned Outage
NODEG	Thursday, June 14...	Wednesday, June...	12 day(s) 17:18:26	Uptime
NODEG	Wednesday, June ...	Wednesday, June...	0 day(s) 00:03:30	Planned Outage
NODEG	Wednesday, June ...	Thursday, July 12,...	15 day(s) 04:42:01	Uptime
NODEG	Thursday, July 12, ...	Wednesday, July ...	5 day(s) 19:15:41	Planned Outage
NODEG	Wednesday, July 1...	Wednesday, July ...	0 day(s) 00:54:28	Uptime
NODEG	Wednesday, July 1...	Wednesday, July ...	0 day(s) 00:54:27	Unplanned Outage
NODEG	Wednesday, July 1...	Wednesday, July ...	0 day(s) 00:02:40	Uptime

Viewing the System Outages by Hour of Day graph

The System Outages by Hour of Day graph shows the distribution of outages by time of day.

Moving the mouse over the bars in the graph invokes Hover Help and supplies detailed information about each outage: Total, Planned and Unplanned.

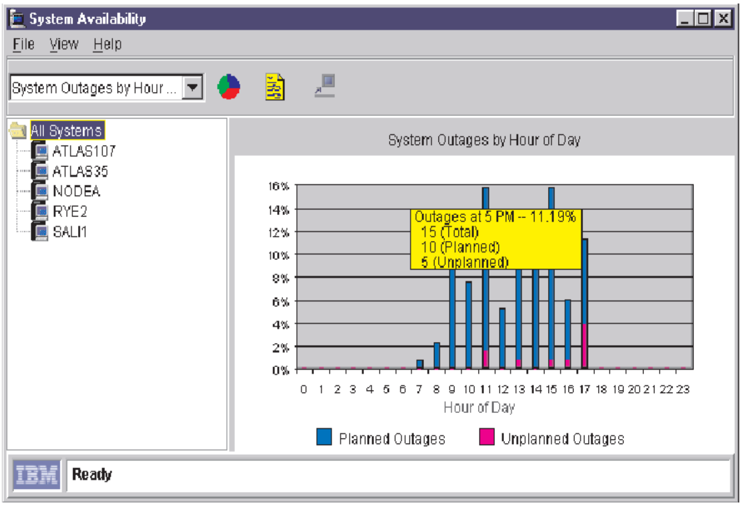

Viewing the System Outages by Day of Week graph

The System Outages by Day of Week graph shows the distribution of outages through a histogram for each day of the week.

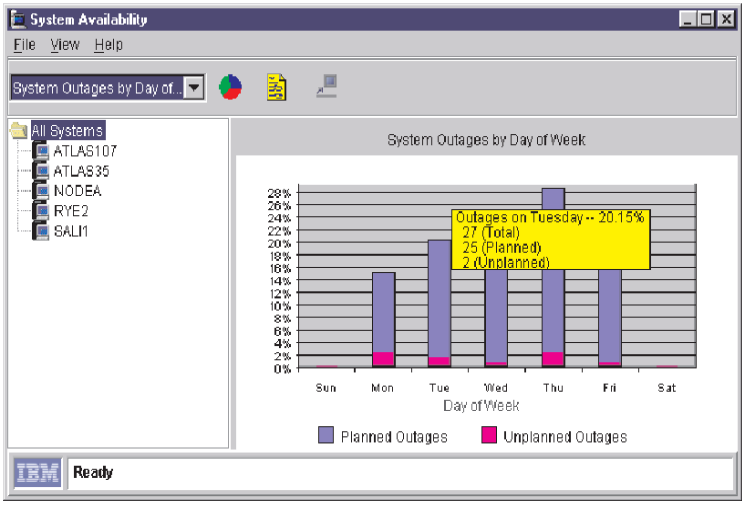

Outages on Tuesday -- 20.15%
27 (Total)
25 (Planned)
2 (Unplanned)

Displaying data within a defined interval

To display data within a defined time interval, customize the data displayed in the Customization of Graph Dates window. To select a specific interval, use the following procedure:

1. From the System Availability window, click **File→Set time**.

   The Customization of Graph Dates window opens.

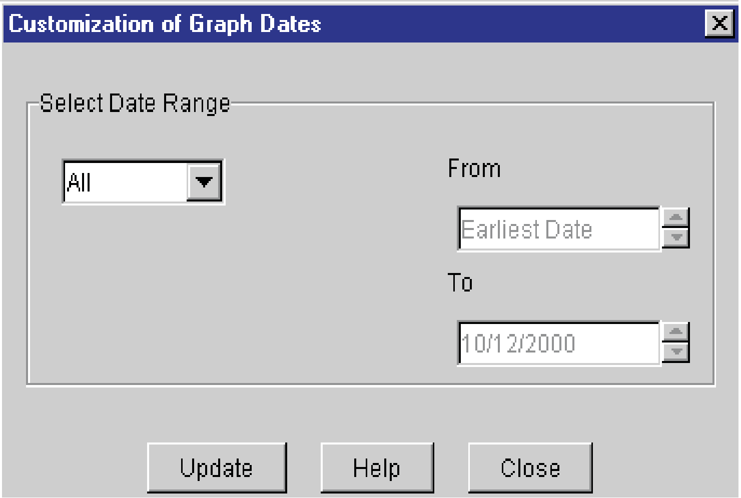

Select Date Range

All

From

Earliest Date

To

10/12/2000

2. From the Select Date 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 369.

   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.

Chapter 32. Troubleshooting

Although every effort has been made to provide you with a simple and easy to use interface, you have find problems when running Director.

Your Director reseller has training and experience in helping you solve your system management problems. This chapter includes some common situations you might encounter when you use Director.

Q: My Jet database is full. What can I do?

A: The Microsoft Jet database has a maximum limit of 1 GB. If your database is less than 1 GB, try to free up additional space, up to 1 GB. To do so, move some files from the drive where the \directory subdirectory was installed. You can also move the Jet database to another, larger drive that has at least 1 GB available space. See Appendix B for more information on moving the Jet database.

Q: My Jet database is at the 1 GB limit. How can I get more space?

A: You should switch database support to the more robust databases, such as DB2, Oracle or Microsoft SQL. For information on switching database support from Jet to SQL, see:

- Chapter 2, “Planning,” on page 15.
- Chapter 3, “Installation and configuration,” on page 39.

Q: Why are my component installations failing even though I have verified that sufficient space is available?

A: Director uses temporary disk space on the target system during installation. You must have sufficient space available for the temporary directory as well as the target installation directory. Use the following list to determine the amount of free space required in the temporary directory for installation and uninstallation of the various components. Note that TMA indicates the Tivoli Management Agent.

Component	Installation Space Required (in bytes)	Uninstallation Space Required (in bytes)
Novell Agent	1420331	N/A
Windows 98 Agent	4999233 with TMA; 3727506 without TMA	72192
Windows NT or Windows 2000 Agent	5073303 with TMA; 3727506 without TMA	72192
Windows 98 Console	1420331	72192
Windows NT or Windows 2000 Console	1420331	72192

Component	Installation Space Required (in bytes)	Uninstallation Space Required (in bytes)
Windows NT or Windows 2000 Server	3727506	2706431

Q: Why are some SNMP devices not being discovered?

A: Verify that the Director management server is running the SNMP service. If not, another system on the same subnet must be running an SNMP agent, and must be added as a seed device. In this case, the Director Management Server should be removed as a seed device.

Verify that the seed devices and devices to be discovered are running an SNMP agent.

Verify that the community names specified in the Director Discovery Preferences window allow Director to read the **mib-2.system** table of the devices to be discovered, and the **mib-2.at** table on seed devices.

Verify that the correct network masks have been configured for all systems that are to be discovered.

Verify that the correct addresses have been entered for the seed devices. The most effective seed devices are routers and domain name servers. To configure these devices, from the Director Management Console window select **Options→ Discovery Preferences**. SNMP discovery will not discover 100% of the systems. If a system has not communicated with other systems, it might not be discovered.

Q: When I open the SNMP browser for my device, it does not display the specific MIB that I requested. How can I get it to do so?

A: Verify that Director is using a community name that allows read access to the MIB you wish to view. Note that some SNMP devices enable you to hide certain MIBs behind certain community names.

Check that the device or agent implements the MIB in question.

Q: Why does Director not allow me to change an MIB value?

A: Check the following:

- Verify that Director is using a community name that allows write access to the MIB you wish to set.
- Verify that the MIB is writable. Director uses an icon shaped like a pencil to indicate the MIB is writable.
- Verify that you have compiled a MIB associated with the value you want to change.

Q: Director describes setting a particular MIB value to a hexadecimal/octal/binary value, but it will not accept my number. Why?

A: Director expects all values to be added in decimal. You must convert the number from hexadecimal/octal/binary to decimal.

Q: What protocols does Director use for sending and receiving SNMP traps?

A: This version of Director can send and receive traps over TCP/IP only.

Q: Why are some TCP/IP management agent systems not being discovered?

A: For systems to be discovered on subnets other than the one that the Director server resides, seed devices must be configured. Note that:

- You should use only one system for each subnet.
- The Director server must be able to ping the seed devices.
- The seed devices must be able to ping the Director server.

Configure these from the Director Management Console window. Select **Options →Discovery Preferences → Director System Discovery (IP)**.

In addition, discovery requires that any routers or bridges between the Director server and the target agent have port 14247 open. They also must allow IP broadcasts on that port.

Q: Why are some IPX management agent systems not being discovered?

A: For systems to be discovered on networks other than the one that the Director server resides, a network server that has access to the ROUTEs of the networks to be discovered must be the favored server of the Director server. Another method would be that seed devices are configured. Note that:

- You should use only one system for each network.
- The Director server must be able to respond to IPXPING requests from the favored NetWare server.
- The seed networks must be able to respond to IPXPING requests from the favored NetWare server.

Configure these from the Director Management Console. Select the **Options → Discovery Preferences →System Discovery (IPX)**.

Q: I am receiving incorrect inventory data back from my query. Why?

A: Verify that the hardware is returning the correct information.

Q: When I attempt a hardware inventory, a blue screen suddenly appears. Why?

A: If the Director server is running under Windows NT Service Pack 4, the symc810.sys device driver is probably causing the blue screen. Reinstate the original NT 4.0 symc810.sys device driver or obtain the latest symbios drivers from the Symbios Web site, www.symbios.com.

Q: When I start up the console I receive an error message: "IO error connecting to server." What can I do?

A: This usually occurs if you are attempting to bring up the console before the Director management server is completely up. Check the Director management server status to verify that it is ready.

Q: I receive errors when I try to log in to the server from the console.

A: Verify that the server name as well as your user ID and password are valid and that the server is up and running.

Q: Why do some of my managed systems appear "grayed out" on the Director Management Console?

A: Check the following:

- Verify that the system is powered on.
- Verify that the agent is running.
- Increase the Network Time-Out value on the Director management server system as well as the managed system (you must restart the system after making this change).

Q: Why is there a padlock on some of my managed system icons?

A: This denotes that the system is another Director management server. By default, you cannot manage other Director management servers. To enable other servers to manage your server, select **Unsecure System** from the context menu in the Group Contents pane of the Director Management Console window.

Q: Why are certain options not available on the context menu of my managed system?

A: Perhaps that managed system does not support the option, or inventory might not have been collected on that managed system yet.

Q: Why do some of my managed systems become unavailable on the console?

A: Perhaps the time-out value for Director to access the system needs to be increased. Modify your Network Time-Out value in the Network Driver Configuration window (select **Start Programs →Director → Network Configuration**).

Q: Why do I see a \~twgtemp subdirectory on my console system?

A: If a console machine fails while writing a locally created software distribution package to the server, there may be temporary files left on the console. These files are in the \~twgtemp subdirectory in the root of the drive on which you installed the Director Management Console. Delete this directory while the console is not running to reclaim lost disk space.

Q: Why is software distribution package creation failing on large packages?

A: Check the available disk space on your local (console) system. Packages are created locally before being written to the server, so if there is insufficient disk space on the local machine, package creation fails.

Q: Why do I have problems starting remote control sessions or distributing software packages when managed systems are on the other side of a firewall?

A: Remote control and software distribution both use session support to increase data flow. Session support within TCPIP causes data to flow through a different port than the one that Director normally uses for communications. Most firewalls do not allow the data to flow through this other port.

You can disable session support by creating an.ini file on the agent system. In the agent's \tivoliwg\bin directory, create a file named tcpip.ini that contains the following line:

```
SESSION_SUPPORT=0
```

If there is more than one TCP/IP option in the agent's Network Driver Configuration panel, you must create an .ini file for each entry. Name these files tcpip.ini, tcpip2.ini, tcpip3.ini, and so on. After creating the files, reboot the agent system or stop and restart the Director agent.

Q: Why does my system slow down when using resource monitors?

A: The system might slow down if many monitors are running. This will also occur if many systems are being monitored.

Q: Why does the performance drop when I run multiple Director Management Consoles?

A: You can run multiple instances of the Monitor Console. However, the overhead required to maintain multiple instances may degrade the response performance of the console, depending on the number of unique attributes and the number of systems being monitored. When multiple consoles are viewing the same attribute data, the performance degradation is minimized.

Q: While trying to use a share for software distribution of a particular software package, I received an error message of the form:

Managed System (system name) has detected that software package (Package Name) was not found on share (\\server\share).

What is wrong?

A: Software distribution packages are deleted from the Director Management Console. When a package is deleted and the package has been cached on a share, then Director also removes the package from the share.

The software packages are stored on the shares in a directory that is unique to the software package. This directory is maintained by the Director server and should not be modified by a user. If a software package directory is deleted through a means other than the Director Management Console, any managed system that attempts to use the share for that software package reports the error message you received.

To recover from this situation, the software package should be refreshed by using the File Distribution Servers Manager.

Q: Why are my software distributions not using the redirected drives?

A: There must be a trust relationship between the Primary Domain Controller and the server that is being used as the redirected drive.

Q: Why can I not use a server share to redistribute a software distribution package to an OS/2 managed system (it always defaults to a streamed installation)?

A:A user ID *must* be logged on to the target OS/2 machine to redistribute the package. When a user is not logged on to OS/2, the distribution defaults to streaming.

Q: I'm try to distribute a software package from Windows NT to OS/2 and it is failing, but software distribution to OS/2 is supported. Why is it failing?

A: The target OS/2 system might be using FAT-based drives. If so, the files within the software distribution package must be in 8.3 format to be installed on a FAT-based drive.

Q: The streaming of a software distribution package to an OS/2 managed system was suspended and resumed, but all of the package had to be retransmitted. Why?

A: If your OS/2 managed systems contain FAT-based drives and the DISKCACHE setting is enabled for Lazy Write, suspended distributions will not resume properly. To solve this problem, on the target systems, remove the Lazy Write (LW) parameter from the DISKCACHE statement in the config.sys file and restart the systems. This problem does not occur on HPFS-based drives.

Q: How can I change the software distribution package install location?

A: You must reinstall the Director agent, specifying a different drive and directory.

Q: Why is Director not starting up?

A: This is usually due to security issues or the database is not properly configured. If the administrator's password has changed, then you must change the password for the Director Support Program, in the Services section of the Control Panel.

You must have logged in with an administrator's ID when installing Director. If your ID is being validated by a domain, then it must be a domain administrator's ID. If you are using a local ID, then it must have administrator privileges.

If you have switched from your primary domain controller to your backup domain controller, you must create a local administrator's account on the BDC, to match the account that was used when Director was installed.

Q: Why are my redistributed installs not working properly?

A: If IBM AntiVirus is installed on the Director management server, redirected distributions will fail. You must uninstall the IBM AntiVirus, delete the packages that fail, and re-create the packages.

Q: Why do I get a Stack Fault dialog on a Windows NT 4.0 managed system after a distribution?

A: Installing NT 4.0 Service Pack 3 on the managed system should resolve the stack fault.

Q: When I create a dynamic group using the not equal to operator as part of the selected criteria, not all of the managed systems that do not possess that criterion are returned.

Why does this happen?

A: When you create a dynamic group by selecting certain criteria, each criterion only searches the rows in the table with which it is associated. For example, if you select a criterion of Inventory (PC) / SCSI Device/Device Type = TAPE, only those managed systems that appear in at least one row in the SCSI_DEVICE table that also have a value of TAPE in the DEVICE_TYPE column will be returned.

Likewise, if you select Inventory (PC) / SCSI Device/Device Type ^= TAPE as a criterion, only managed systems that appear in at least one row of the SCSI_DEVICE table, of which none of those rows have a value of TAPE in the DEVICE_TYPE column, will be returned. **It does not necessarily return all managed systems that do not have SCSI tape drives.** In other words, only managed systems that appear in a particular table and that meet the criteria for that table are returned.

Another example is a dynamic group created by specifying the following two criteria:

- Inventory/SCSI Device/Device Type ^= TAPE
- Inventory/ Operating System/Type = WINDOWS NT

Using these criteria, a Windows NT managed system with no SCSI devices would not be returned, because such a managed system would not appear in the SCSI_DEVICES table. However, if a Windows NT managed system had a SCSI hard drive but no SCSI tape drive, it would be returned, because such a system would appear in the SCSI_DEVICES table.

Q: I get an error when I try to run the Database Configuration process on Oracle. What might be the problem?

A: The Oracle TCP/IP Listener must be configured and started prior to running the Database Configuration dialog.

Q: I'm having trouble configuring Oracle 7.3.4. What should I do?

A:

- If you are running Oracle Version 7.3.4, you must edit the **initdirector.ora** file in **/opt/oracle/admin/director/pfile** to allow the use of unlimited rollback segments (where **director** is the instance name). Add the following line:

  ```
  unlimited_rollback_segments = true
  ```

  Log into Oracle and issue a shutdown and startup before attempting to run the Oracle Database Configuration dialog.

- If you are running Oracle Version 7.3.4, the COMPATIBLE parameter must be set to 7.3.0.0 or greater. To set this, edit the **initdirector.ora** file in **/opt/oracle/admin/director/pfile** (where **director** is the instance name). Uncomment the following line:

  ```
  # compatible = “7.1.0.0”
  ```

  and change it to:

  ```
  compatible = “7.3.0.0”
  ```

  Log into Oracle and issue a shutdown and startup before attempting to run the Oracle Database Configuration dialog.

Appendix A. Resource monitor attributes

This appendix contains a list of the attributes that can be monitored by the Director resource monitors task on managed systems that have the Tivoli management agent installed.

Monitor collection rates are every 30 seconds, unless otherwise noted.

Windows 98

CPU monitors

- CPU Utilization
- Process Counts

Disk monitors

Notes:

1. The disk drive monitors will repeat for each local non-removable logical drive found.
2. Monitor data collection rate is every 60 seconds.

- Disk 1 Workload
- Drive C: % Space Used
- Drive C: Space Remaining
- Drive C: Space Used

File monitors

File monitor attributes can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
2. Monitor data collection rate is every 60 seconds.

Directory

- Directory Exists
- Last Modified

File

- Checksum
- File Exists
- File Size
- Last Modified

Memory monitors

- Locked Memory
- Memory Usage

Registry monitors

Notes:

1. Each registry entry is an attribute.
2. Monitor data collection rate is every 60 seconds.

TCP/IP monitors

- Interface 0 - Broadcast Packets Received
- Interface 0 - Broadcast Packets Sent
- Interface 0 - Bytes Received
- Interface 0 - Bytes Sent
- Interface 0 - Unicast Packets Received
- Interface 0 - Unicast Packets Sent
- IP Packets Received
- IP Packets Received with Errors
- IP Packets Sent
- TCP Connections
- UDP Datagrams Received
- UDP Datagrams Sent

Process monitors

Notes:

1. The number of applications or executable files monitored by the process monitors is variable and configured by the Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.

2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start
- Times started
- Times stopped
- Total execution time
- Yesterday's execution time
- Yesterday's new executions

CIM monitors

Note: Monitor data collection rate is every 15 seconds.

- CIMV2
- Default
- Security
- WM1

Windows NT operating system

CPU monitors

- CPU Utilization
- CPU 'x' Utilization (on SMP machines)
- Process Counts

Device and service monitors

Note: Monitor data collection rate is every 15 seconds.

- State

Disk monitors

Notes:

1. The disk drive monitors will repeat for each local non-removable logical drive found.

2. Monitor data collection rate is every 60 seconds.

- Disk 1 Workload
- Drive C: % Space Used
- Drive C: Space Remaining
- Drive C: Space Used

DMI monitors

Notes:

1. DMI monitors are only on IBM systems

2. Monitor data collection rate is every 15 seconds.

File monitors

File monitor attributes can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.

2. Monitor data collection rate is every 60 seconds.

Directory

- Directory Exists
- Last Modified

File

- Checksum
- File Exists
- File Size
- Last Modified

Memory monitors

- Locked Memory
- Memory Usage

Windows NT Performance Monitors

Note: The number of Windows NT Performance Monitors can vary. These monitors are gathered directly from the Windows NT Performance Monitor (PerfMon) subsystem. These monitors change dynamically. On a typical Windows NT system over 3500 different attributes can be monitored under the Windows NT Performance Monitors.

Registry monitors

Notes:

1. Each registry entry is an attribute.
2. Monitor data collection rate is every 60 seconds.

TCP/IP monitors

- Interface 0 - Broadcast Packets Received
- Interface 0 - Broadcast Packets Sent
- Interface 0 - Bytes Received
- Interface 0 - Bytes Sent
- Interface 0 - Unicast Packets Received
- Interface 0 - Unicast Packets Sent
- IP Packets Received
- IP Packets Received with Errors
- IP Packets Sent
- TCP Connections
- UDP Datagrams Received
- UDP Datagrams Sent

Process monitors

Notes:

1. The number of applications or executable files monitored by the process monitors is variable and configured by the Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.

2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start
- Times started
- Times stopped
- Total execution time
- Yesterday’s execution time
- Yesterday’s new executions

Windows 2000 operating system

CIM monitors

Note: Monitor data collection rate is every 15 seconds.

CPU monitors

- CPU Utilization
- CPU ‘x’ Utilization (on SMP machines)
- Process Counts

Device and service monitors

Note: Monitor data collection rate is every 15 seconds.

- State

Disk monitors

Notes:

1. The disk drive monitors will repeat for each local non-removable logical drive found.

2. Monitor data collection rate is every 60 seconds.

- Disk 1 Workload
- Drive C: % Space Used
- Drive C: Space Remaining
- Drive C: Space Used

DMI monitors

Notes:

1. Only for IBM systems

2. Monitor data collection rate is every 15 seconds.

File monitors

File monitor attributes can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.

2. Monitor data collection rate is every 60 seconds.

Directory

- Directory Exists
- Last Modified

File

- Checksum
- File Exists
- File Size
- Last Modified

Memory monitors

- Locked Memory
- Memory Usage

Windows NT Performance Monitors

Note: The number of NT Performance Monitors can vary. These monitors are gathered directly from the Windows NT Performance Monitor (PerfMon) subsystem. These monitors change dynamically. On a typical Windows NT system over 3500 different attributes can be monitored under the Windows NT Performance Monitors.

Registry monitors

Notes:

1. Each registry entry is an attribute.
2. Monitor data collection rate is every 60 seconds.

TCP/IP monitors

- Interface 0 - Broadcast Packets Received
- Interface 0 - Broadcast Packets Sent
- Interface 0 - Bytes Received
- Interface 0 - Bytes Sent
- Interface 0 - Unicast Packets Received
- Interface 0 - Unicast Packets Sent
- IP Packets Received
- IP Packets Received with Errors
- IP Packets Sent
- TCP Connections
- UDP Datagrams Received
- UDP Datagrams Sent

Process monitors

Notes:

1. The number of applications or executables monitored by the process monitors is variable and configured by the Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.
2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start
- Times started
- Times stopped
- Total execution time

- Yesterday’s execution time
- Yesterday’s new executions

OS/2 operating system

APM monitors

Note: The APM Monitors are only supported on laptop systems with the correct vendor supplied drivers.

- Battery Remaining
- Percent

CPU monitors

- CPU ‘x’ Utilization
- CPU ‘x’ Utilization (on SMP machines)
- Process Count (1 minute refresh rate)
- Thread Count (1 minutes refresh rate)
- CPU Cache Hit Rate (Pentium processors only)
- Floating Point Operation Rate (Pentium processors only)
- Integer Instructions Rate (Pentium processors only)
- Interrupt Rate (Pentium processors only)
- Memory I/O Rate (Pentium processors only)
- Port I/O Rate (Pentium processors only)

Disk monitors

Notes:

1. The disk drive monitors will repeat for each local non-removable logical drive found.
2. Monitor data collection rate is every 60 seconds.

- Drive C: % Space Used
- Drive C: Space Remaining
- Drive C: Space Used

File monitors

File monitor attributes can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.

2. Monitor data collection rate is every 60 seconds.

Directory

- Directory Exists
- Last Modified

File

- Checksum
- File Exists
- File Size
- Last Modified

Memory monitors

- Locked Memory
- Memory Usage
- ECC Memory (if installed)

OS/2 server monitors

Note: Monitor data collection rate is every 30 seconds.

- Big Buf Shortage
- Bytes Received
- Bytes Sent
- Connections
- Logons
- Opens
- Print Jobs Queued
- Response Time
- Request Buf Shortage
- Sessions
- Shares

OS/2 swap file monitors

- Swap File Size

- Swap Space Remaining

Process monitors

Notes:

1. The number of applications or executables monitored by the Process Monitors is variable and configured by the Director administrator from the Process Manager console. Each one of the attributes under Process Monitors will be present for each executable being monitored.

2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start
- Times started
- Times stopped
- Total execution time
- Yesterday's execution time
- Yesterday's new executions

NetWare operating system

CPU Monitors

- CPU Utilization
- CPU 'x' Utilization (on SMP machines)
- Process Count (1 minute refresh rate)
- Thread Count (1 minute refresh rate)

Disk monitors

Notes:

1. The disk volume monitors will repeat for each volume detected on a NetWare Server.

2. Monitor data collection rate is every 60 seconds.

- Volume SYS: Space Remaining

- Volume SYS: Space Used

File monitors

File monitor attributes can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.

2. Monitor data collection rate is every 60 seconds.

Directory

- Directory Exists
- Last Modified

File

- Checksum
- File Exists
- File Size
- Last Modified

Memory monitors

- Cache Blocks in Use
- Percent of Cache in Use

Process monitors

Notes:

1. The number of applications or executables monitored by the Process Monitors is variable and configured by the Director administrator from the Process Manager console. Each one of the attributes under Process Monitors will be present for each executable being monitored.

2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start

- Times started
- Times stopped
- Total execution time
- Yesterday’s execution time
- Yesterday’s new executions

UNIX and Linux operating systems

CPU monitors

- CPU Utilization
- Process Count (1 minute refresh rate)

Disk monitors

Notes:

1. The list of file systems will appear first; the following attributes will appear under each file system.

2. Monitor data collection rate is every 60 seconds.

- Blocks Available
- Blocks Used
- Inodes Available
- Inodes Used
- Percentage Blocks Available
- Percentage Blocks Used
- Percentage Inodes Available
- Percentage Inodes Used
- Percentage Space Available
- Percentage Space Used
- Space Available (MB)
- Space Used (MB)

File system monitors

Notes:

1. The monitor attributes listed are useful Unix directories. If one of these directories does not exist on a given Unix system, then it will not appear as a monitor attribute.

2. Monitor data collection rate is every seconds.

- /
- /bin
- /dev
- /etc
- /home
- /lib
- /lost+found
- /sbin
- /tmp
- /usr
- /var

List of directory contents

- Directory Attributes
- Directory Exists
- Directory Owner
- Directory Size (bytes)
- Last Modified
- Object Type

The above elements can be files or directories. See the appropriate heading below for the corresponding list of monitors.

Notes:

1. If there are additional directories, additional sub-elements will be present.
2. It is possible that directories that contain a large number (more than several hundred) of sub-elements will take longer than 5 seconds to open.

File: •Checksum

- File Attributes
- File Exists
- File Owner
- File Size (Bytes)
- Last Modified
- Object Type

Directory: •Directory Attributes

- Directory Exists
- Directory Owner
- Directory Size (Bytes)
- Last Modified
- Object Type

Memory monitors

- Available (Bytes)
- Used (Bytes)

Process monitors

Notes:

1. The number of applications or executables monitored by the process monitors is variable and configured by the Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.

2. Monitor data collection rate is every 15 seconds.

- Current Active Processes
- Maximum running at once
- Maximum running yesterday
- New executions counted
- Times failed to start
- Times started
- Times stopped
- Total execution time
- Yesterday's execution time
- Yesterday's new executions

UNIX system monitors

Note: These monitors duplicate the CPU, Disk, and Memory monitors and their attributes detailed above. They are included to maintain backwards compatibility with a SCO UNIX agent previously released.

- CPU Monitors
- Disk Monitors
- Memory Monitors

Appendix B. Creating the ODBC entry for the default database

Use the following steps to manually create the default Microsoft Jet database:

1. Shut down the Director server and ensure that you are logged on with the Director user ID.
2. Go to the ODBC administrator by selecting **Start**→ **Settings**→**Control Panel** and then select the **ODBC** icon.
3. Select the **User DSN** tab.
4. Click the **Add** button.
5. Select the **Microsoft Access** driver (*.mdb) and then click **Finish**.
6. Enter **Director** as the Data Source name.
7. Click the **Create** button.
8. Enter **Director.mdb**, select the **Database** directory under the Director installation directory (for example, **c:\Tivoliwg\Database**), and click **OK**.
9. Click **OK** on the Access Setup window.
10. Click **OK** on the ODBC Data Source Administrator window.
11. Close the ODBC window.
12. Create file **TWGServer.Prop** in the **Data** directory under the Director installation directory (for example, **c:\Tivoliwg\Data**) with the following entry:**twg.database.odbc.name=Director**.
13. Restart the Director server and perform an inventory collection to fill the database.

Appendix C. Converting to other supported databases

This appendix contains information on converting database support from the default Microsoft Jet database to any of the other supported databases and for converting between those databases.

When you originally installed the Director server you should have specified that you wanted to use the default Microsoft Jet database that comes with IBM Director. Using this appendix, you can now convert to the other supported databases.

If you are currently using one of the other supported databases, you can also use this appendix to convert to another supported database (except Jet). If you want to convert back to the Jet database, see Appendix B, "Creating the ODBC entry for the default database," on page 395.

Note: This process only provides you the ability to use a different database. It does not transfer the contents of the database.

Preliminary steps

Refer to "Database support" on page 15, paying particular attention to the planning information for the database to which you are converting.

Using the database configuration window to convert to another database

To convert to another database, run the respective command below (corresponding to the database you want to convert to), to display the Database Configuration window. For more information on using the Database Configuration process, see Chapter 3, "Installation and configuration," on page 39, and the online help.

- **cfgmssql** - Microsoft SQL Server Database
- **cfgdb2** - IBM DB2 Universal Database
- **cfgoracle** - Oracle Server Database

Appendix D. Defining table property files

This appendix contains information on converting database support from the default Microsoft Jet database to any of the other supported databases and for converting between those databases.

When you originally installed the Director server you should have specified that you wanted to use the default Microsoft Jet database that ships with Director. Using this appendix, you can now convert to the other supported databases.

If you are currently using one of the other supported databases, you can also use this appendix to convert to another supported database (except Jet). If you want to convert back to the Jet database, see Appendix B, "Creating the ODBC entry for the default database," on page 395.

Note: This process only provides you the ability to use a different database. It does not transfer the contents of the database.

Preliminary steps

Refer to "Database support" on page 15 , paying particular attention to the planning information for the database to which you are converting.

Setting up the server to inventory CIM and DMI information

Director collects inventory information from managed systems and stores it in database tables in the server's database. The formats of these tables cannot be changed. With the addition of inventory collectors for CIM, DMI, and from static MIF files, a facility for allowing the end user to define custom tables is necessary.

The approach described here to solve this problem uses property files that follow the Java property file format. These property files describe the contents of a custom database table. The property files, one per table, contain the table's name, names and types for each of the columns of the table, and other information. For information on the syntax of the property file, see "Table property file format" on page 400.

Because the created tables may be viewed in any locale supported by Director, one may wish to have table names, column names, and some column values translated for different languages and locales. Files containing these translated strings can be supplied along with the table property files. These files will be read and their strings used where appropriate in the product. These files are explained in the section "NLS file format" on page 404".

In addition to table property files, you must supply files that specify associations between Director inventory collectors and the custom tables. These files will follow the Java property file standard as well. This file format is explained in the section “Inventory extension property file format” on page 406. Without these files, Director will not know how to map the data from the CIM, DMI, or MIF inventory collectors into the custom tables. The section “Static MIF Data Collection” explains how to set up a managed system to generate MIF files used by the collector.

The only user interface to the custom table facility is through the property files. Table and inventory extension property files are read when the Director server starts up. The server looks in two predetermined subdirectories in the server's directory for these files, loads all of the table property and inventory extension files it finds, and then creates or initializes tables defined by these files. Thus, if you need to make changes to the table or extension files, you must stop and restart the server before those changes take effect. There are important restrictions on how table property files can be changed, as well as special procedures the server follows with regard to new, removed, or changed table property files. These restrictions and procedures are explained in the section “Server initialization and table property fIles” on page 411.”

As the table property files are parsed by the server, the status of this parsing is written to text files in the same directory as the table property files. These status files explain what errors were encountered in parsing the file, if any. The error messages are designed to give as much information as possible, not requiring further explanation in this document. To help you with creating valid files, the section “Examples” gives some sample property files.

Here is a sample property file:

```
software = Director
hardware.type = Generic workstation
with 128MB RAM.
#video = VGA
```

It defines three properties: `software`, `hardware.type`, and `with` (`with` is defined unintentionally because the value for `hardware.type` takes up two lines, so `with` is read as a new property by the parser). The line `#video = VGA` is ignored because it is read as a comment. More examples are given in the “Examples” section.

Table property file format

A Table Property File can be created and edited using any ASCII editor. These files are placed in the User Tables subdirectory of the server's data directory; this path will usually be C:\Program Files\Director\Data\Tables. The syntax of a property file consists of one property name followed by its associated value on a single line, the two separated by an equal sign. Text on a new line implies a new property. Leading or trailing white space is ignored. Spaces within the property value are preserved. The first equal sign or space is assumed to be the separator between property name and value; any following equal signs or spaces in a property definition are added to the string for that property's value, except for white space which surrounds a separator. If a property is listed more than once in a file, each successive definition overwrites the previous one. Comment lines can be inserted in the file by starting the line with a hash character (#).

This format has a few subtleties that can cause unexpected side effects. If a property's value exceeds a line, the remainder of the value will be interpreted as one or more new property definitions. If a space is inserted into a property name, part of its name will be misinterpreted as its value. While the parser tries to catch errors, some errors can be interpreted as valid properties, and simple file editing mistakes could cause unexpected behavior.

Here is a sample property file:

```
software = Director
hardware.type = Generic workstation
with 128MB RAM.
#video = VGA
```

It defines three properties: `software`, `hardware.type`, and `with` (`with` is defined unintentionally because the value for `hardware.type` takes up two lines, so `with` is read as a new property by the parser). The line `#video = VGA` is ignored because it is read as a comment. More examples are given in the "Examples" section.

As a custom table property file is processed, its status is written to a text file with the same name as the table property file, but with a ".status" extension, in the same directory as the table property file. This status file contains a list of properties as the server parsed them (so you can check for formatting mistakes) as well as descriptions of errors that were encountered during the processing of the file.

The properties in the Table Property File are listed below. Property names must be entered with the same capitalization (all lowercase) as shown. Each of the values for tokens, realnames, and shortnames can contain only these characters:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvqxyz0123
456789-_).
```

This restriction is based on restrictions set by the supported DBMSs.

table.token: name of the table used internally in the Director server. Optional. If missing, table.token will be set to the property file's filename, without the extension or leading path.

table.realname: name of table as stored in the database. If NLS support is enabled, this name is also part of a key into an NLS resource file (see nls.X options below) that is used to obtain the user-readable version of this name from the resource file. Optional. If missing, table.realname will be set to table.token. This value must not be an SQL keyword for the database system used.

table.shortname: name of the table as stored in the database. If the database would truncate the real name in an undesirable way, one can specify the name for that database to use with this property. Optional. This value must not be an SQL keyword for the database system used.

table.displayname: name of table as displayed to the user Inventory Query Browser and the Dynamic Group Editor dialog boxes. If an NLS file is specified for the current locale (see nls.X properties below), and the table's name is defined in the NLS file, then that name is used instead; display names are used as a last resort. Optional. If missing, table.displayname will be set to table.realname.

table.filterprompt.alltrue: string displayed in the Dynamic Group Editor when adding a column from this table to a filter; this string appears for the "all true" option. Optional. If not specified, the default string will be used, which has already been translated to the locales Director supports. In English, this string is "All true (AND)."

table.filterprompt.anytrue: string displayed in the Dynamic Group Editor when adding a column from this table to a filter; this string appears for the "any true" option. Optional. If not specified, the default string will be used, which has already been translated to the locales Director supports. In English, this string is "Any true (OR)."

table.filterprompt.alltrueforsame: string displayed in the Dynamic Group Editor when adding a column from this table to a filter; this string appears for the "all true for same row" option. This option will only appear if there is more than one column designated as a key value (this includes the MANAGED_OBJ_ID column automatically added by the server, which is a key). Having more than one key in a table makes it possible to have more than one row for a managed system, so this prompt is shown to you in filter creation. Optional. If not specified, the default string will be used, which has already been translated to the locales Director supports. In English, this string is "All true for the same row."

table.filterprompt.eachtrueatleastone: string displayed in the Dynamic Group Editor when adding a column from this table to a filter; this string appears for the "each true for at least one row" option. This option will only appear if there is more than one column designated as a key value (this includes the MANAGED_OBJ_ID column automatically added by the server, which is a key). Having more than one key in a table makes it possible to have more than one row for a managed system, so this prompt is shown to you in filter creation. Optional. If not specified, the default string will be used, which has already been translated to the locales Director supports. In English, this string is "Each must be true for a least one row."

nls.X.locale: name of a locale for which a file of translated strings is provided. X is an integer index representing the Xth locale. The indices X may start at 0 and need not be sequential. The locale MUST follow this syntax: two-letter language code, OR two-letter language code, underscore, two-letter country code, OR two-letter language code, underscore, two-letter country code, underscore, variant code. Example: fr for French; en_us for US English Windows variant. Optional. If no NLS locales are specified, the displaynames of the table, columns, and values will be displayed.

nls.X.filename: path to a file that lists the literal strings that correspond to realnames in the property file (the table and column labels and values). This file has its own format, as described later. This file corresponds to the Xth locale as defined with the nls.X.locale property. The path should be relative to the directory in which the table property files are stored.

column.X.token: name of a column of data used internally in the Director server. X is an integer index representing the Xth column. The indices X may start at 0 and need not be sequential. Required.

Note: It is **not** necessary to define a column called MANAGED_OBJ_ID. This column is created automatically as the first column because it is required in every table.

column.X.realname: name of column X as stored in the database. If NLS support is enabled, this name is also part of a key into an NLS resource file (see nls.X options) that is used to obtain the user-readable version of this name from the resource file. Optional. If missing, column.X.realname will be set to column.X.token. This value must not be an SQL keyword for the database system used.

column.X.shortname: name of column X as stored in the database. If the database would truncate the real name in an undesirable way, one can specify the name for that database to use with this property. Optional, this value must not be an SQL keyword for the database system used.

column.X.displayname: name of column X as displayed to the user Inventory Query Browser and the Filter creation dialog boxes. If an NLS file is specified for the current locale (see nls.X properties), and this column's name is defined in the NLS file, then that name is used instead; displaynames are used as a last resort. Optional. If missing, column.X.displayname will be set to column.X.realname.

column.X.key: set to true or false. If the value is true, specifies that column X is a key. Optional.

column.X.type: type of data stored in column X. The type must be one of the following: SMALLINT, INTEGER, REAL, DOUBLE, CHAR, VARCHAR, DATE, DATETIME. If CHAR or VARCHAR is specified, there must also be a column.X.length property. This type MUST match the type of data returned by the CIM, DMI, or MIF collector that will be put into this column. Required.

column.X.metatype: meta type of data stored in column X. The meta type allows you to specify additional information about the data. The only currently-supported meta type is IPAddress for CHAR columns. This meta type defines the data stored in the CHAR column as a TCP/IP address. This additional information is necessary for sorting and filtering.

column.X.length: If column.X.type is CHAR, this property is required, as it specifies the fixed length of the character field. If column.X.type is VARCHAR, this property is also required, and specifies the maximum size of the variable-length character field.

column.X.value.Y.token: If column.X.type is CHAR or VARCHAR, you may supply strings that represent possible values of these columns. The indexes Y may start at 0 and need not be sequential. The reason a user would want to specify possible values is if you want to display strings to the end user other than the raw collected information. These strings are defined in column.X.value.Y.displayname properties. If a column.X.value.Y.token property is defined, a single corresponding column.X.value.Y.realname property MUST be defined.

column.X.value.Y.displayname: This is the string displayed to the end user when the value of column X is the string listed in column.X.value.Y.token. There must be one and only one displayname per token per column. If a displayname is not specified for a value token, then that token is displayed to the end user as-is. If a column contains a value that does not match to a token listed in the property file, then that value is displayed to the end user as-is.

NLS file format

For each locale specified in the table property file, there needs to be an associated NLS file created. The NLS files are used to build resource bundles as used in Java to provide locale support. Thus, these files follow a strict format, which is explained later. These resource bundles contain names and values like the table property files, where the names represent realnames of the table, its columns, and its column values, and the values associated with the names are the translated strings for those realnames. These strings are displayed to you in the Inventory Query Browser and the Dynamic Group Editor dialog boxes. The resource bundles are built in a hierarchy so that if a name is missing from a bundle, that bundle's parent bundle is searched for the name, and so on.

Generally, the bundle of a locale specified for simply a language, such as “pt” (for Portuguese), will be the parent of a bundle of a locale specified for a language and country, such as “pt_br” (for Brazilian Portuguese). That bundle will in turn be the parent of a bundle of a locale specified for language, country, and variant, such as “pt_br_WIN” (for Brazilian Portuguese, Windows variant).

When the server is started, it will automatically create an NLS file in the user tables directory with the table's file name (without leading path and extension) and the “.defbundle” extension. This file is used to build the default bundle. The values in the default bundle are created from the displayname properties defined in the table properties files. The server tries to make the default bundle the parent of all bundles of locales that only specify a language. For example, the default bundle will be made the parent of “pt” but not “br_pt,” which already has a parent “pt.” However, if a locale is missing, such as “pt,” and a more specific locale exists, such as “br_pt,” the default bundle is made the parent of the more specific locale.

After the NLS resource bundles are set up, the Director server looks through them to find strings to display in the Inventory Query Browser and the Dynamic Group Editor dialog boxes. It uses the search order defined by Java's NLS support: if a bundle is supplied that exactly matches the current locale, that bundle is used, and if a key is missing from that bundle, its parent bundles are followed until a match is made for that key. If no bundle exactly matches the locale, then the current locale is made more general (first the variant is removed if supplied, then the country, then the language) until it matches a bundle. So, for example, if NLS files are supplied for the locale “pt_br” but not “pt,” then if the program is run in the “pt” locale, the NLS default bundle, NOT the “pt_br” bundle, is used.

The NLS file format is strict but simple. In each of the examples below, the user creating the file must fill in his own values for the italicized pieces. The non-italicized pieces must be copied exactly.

To specify the table's display string:

```
TableName.TWGDbUserTable?tableTokenName = translated string for
table name
```

To specify a column's display string:

```
ColumnName.TWGDbUserTable?tableTokenName.columnTokenName =
translated column name
```

To specify a column value's display string:

```
ColumnName.TWGDbUserTable?tableTokenName.columnTokenName.columnVa
lueToken = translated value name
```

Note: If the string used for “columnValueToken” contains spaces, the spaces MUST be replaced with the string {0} (open bracket-zero-close bracket). For example, Default System BIOS becomes Default{0}System{0}BIOS. This substitution is necessary because of the way these files are parsed-a space on the left side of an equal sign signifies the end of the property name, and since columnValueToken is part of the property name, it cannot contain spaces. When the property name is processed by the server (after parsing), the {0} strings will be replaced by spaces. This space substitution is not done for any other property names.

To specify the filter prompt string for “All True:”

```
FilterTablePrompt.AllTrue.TWGDbUserTable?tableTokenName =
translated string
for “all true” for this table
```

To specify the filter prompt string for “Any True:”

```
FilterTablePrompt.AnyTrue.TWGDbUserTable?tableTokenName =
translated string
for “any true” for this table
```

To specify the filter prompt string for “All True For Same:”

```
FilterTablePrompt.AllTrueForSame.TWGDbUserTable?tableTokenName =
translated string
for “all true for same” for this table
```

To specify the filter prompt string for “Each True For At Least One:”

```
FilterTablePrompt.EachTrueForAtLeastOne.TWGDbUserTable?tableTok
enName = translated string
for “all true” for this table
```

The easiest way to create an NLS file is to start the server with the table property file in place in the UserTables directory. The default bundle file will be created as the server initializes. Stop the server, then copy the default bundle file for each locale for which support is needed. In this file, all of the correct keys have been created-just replace the values with the translated values for that locale. Note that the FilterTablePrompt keys are not created in the default bundle file because they have acceptable default values built into the server.

Inventory extension property file format

Once the server has loaded the table property files and has defined those tables, it must associate the data collected by inventory collectors with columns in the custom tables. These associations, called groups, are listed explicitly in inventory extension property files provided by you. The extension files are placed in the InvExtension subdirectory of the server’s data directory; this path will usually be C:\Program Files\Director\Data|Table. One group represents the association between one collector and one table; there can be more than one group per file, but all the properties for a group should be in the same file. An extension file can be one of three types: CIM, DMI, or MIF, with these extensions respectively:.CIMInvExt,.DMIInvExt, or.MIFInvExt. As property files, they can be created and edited with an ASCII text editor and follow a strict syntax. The DMI and MIF collectors extract attribute ID, type, and value data from groups and tables-other fields (name, description, etc.) are not currently supported.

The properties in the Inventory Extension File are listed below. Property names must be entered with the same capitalization as shown.

Group.X.ComponentName: (DMI and MIF only) name of a component in a DMI or MIF namespace from which the data is collected. X is an integer index representing the Xth group. The indices X MUST start at 1 and be sequential within each extension file. These indices do not remain in effect across different extension files; i.e., Group 1 in one file has nothing to do with Group 1 in another file. These indices are used strictly for parsing the files. Required if this extension file is for a DMI or MIF collector.

Group.X.NameSpace: (CIM only) CIM name space from which to retrieve the class name specified in the Group.X.ClassName property. Any slashes in this property must be forward slashes, for example, root/cimv2. Required if this extension file is for a CIM collector.

Group.X.ClassName:

- For CIM: name of a class in a CIM namespace from which the data is collected. This value should be the name of the "leaf" class. Names of any higher level classes should not need to be included. X is an integer index representing the Xth group. The indices X MUST start at 1 and be sequential within each extension file. These indices DO NOT remain in effect across different extension files; i.e., Group 1 in one file has nothing to do with Group 1 in another file. These indices are used strictly for parsing the files. Required if this extension file is for a CIM collector.
- For DMI or MIF: name of the class in DMI or MIF component specified in Group.X.ComponentName. Class names typically follow a Manufacturer|Component|Version format. Required if this extension file is for a DMI or MIF collector.

Group.X.DbTable: token name of the custom table in which to store the data. This name is defined by the table.token property in the table property file. Required.

Group.X.Attrib.Y.Property: name of a CIM property to collect from the class specified in the Group.X.ClassName property. Y is an integer index representing the Yth property for this group's list of attributes. The indices Y MUST start at 1 and be sequential within each list of attributes. This is required if this extension file is for a CIM collector.

Group.X.Attrib.Y.AttributeId: numeric ID of a DMI or MIF property to collect from the class specified in the Group.X.ClassName property. Y is an integer index representing the Yth property for this group's list of attributes. The indices Y MUST start at 1 and be sequential within each list of attributes. Required if this extension file is for a DMI or MIF collector.

Group.X.Attrib.Y.DbColumn: token name of a column in the custom table in which to store the property specified by Group.X.Attrib.Y.Property. Required.

Group.X.Attrib.Y.ScaleBy: scaling factor for numeric values that will be multiplied by the returned value. Optional. If missing, this value is 1 (no effect on the value).

Group.X.Attrib.Y.AdjustBy: scaling factor for numeric values that will be added to the returned value after the value is multiplied by the ScaleBy value. Optional. If missing, this value is 0 (no effect on the value).

All CIM properties collected will be stored (by default) in the database based on the mappings in the following table.

Cim Type	Default Database Type
EMPTY STRING	CHAR

Cim Type	Default Database Type
SINT8 UINT8 SINT16 UINT16 SINT32 UINT32 SINT64 UINT64 BOOLEAN	INT
REAL32	REAL
REAL64	DOUBLE
DATETIME	DATETIME
REFERENCE CHAR16 OBJECT	IGNORED

All DMI and MIF properties collected will, by default, attempt to be stored in the database based on the following mappings:

DMI or MIF Type	Default Database Type
OCTETSTRING DISPLAYSTRING	CHAR

DMI or MIF Type	Default Database Type
DATATYPE_0 COUNTER COUNTER64 GAUGE DATATYPE_4 INTEGER INTEGER64 DATATYPE_9 DATATYPE_10	INT
DATE	DATETIME

Static MIF data collection

The syntax of the extension files for DMI and static MIF is identical, except for the file names. However, collecting data from a MIF file requires some more preparation in the form of specifying how to generate the MIF file. Each agent from which MIF data will be collected will need an initialization file, called MIFGEN.INI, that specifies what program to run to refresh the static MIF data and from what MIF files to collect data. This method allows agents of many operating system types to run different programs to update the static MIF files. The MIFGEN.INI file resides in the same directory as the file DMIPARSE.DLL on the Windows agent (most likely C:\Program Files\Director\ Data\Tables). Be sure to verify that the MIF generation program can be executed successfully from a command line from the \bin directory. It may be necessary to provide an absolute path to the generation program.

The MIFGEN.INI file uses the standard Windows INI file format. There can be many sections in the INI file. Each section starts with a tag enclosed in square brackets and represents a different MIF file. The section contains three properties: `filename`, `command`, and `refresh`. Each property name is followed by an equal sign and the property's value, as in the other property files. The section ends where another section begins, or where the file ends. Section tags and property names are not case sensitive. The value of the tag must be unique within the set of tags in that file, and it is used as the filename if the filename property is missing from that section. If more than one section have equal section tags, only the settings from the first section will be applied. A line beginning with a semi-colon is considered a comment and is ignored by the INI file. The comment continues to the end of the line.

When the Director agent on a managed system is notified that MIF inventory is being collected, it will read the MIFGEN.INI file. For each section, it checks the refresh property. The refresh value can be ALWAYS or NEVER. If the value is ALWAYS, then the command specified by the command property is run and the MIF file specified by the filename property is generated. If the refresh value is NEVER, the command specified by the command property is run ONLY if the file specified by the filename property does NOT already exist-in other words, the file is generated once and never refreshed. If problems are encountered in generating a MIF file, verify that the target file can be created (for example, that no read-only file by the same name exists and that the filename is composed of legal filename characters).

For a section, if the value specified for the refresh property is not ALWAYS or NEVER, or the value isn't specified, then the default of ALWAYS is used. If the filename is not specified, then the section tag is used as the filename. If multiple sections define identical target MIF files (remember that names are NOT case sensitive), the settings from the first section defining that target MIF file will be applied. If the command specified by the command property fails, the previous version of the MIF file is used if it exists. If attempts to create the MIF file fail, and it doesn't exist, then MIF collection for the agent will fail for this MIF file, but collection from other MIF files on the same agent will not be affected.

A sample MIFGEN.INI file is included below. Notice that a command in a section does not need to run a MIF generator. You may create sections to move old MIF files, for example. In the example below, `genmif` is an imaginary MIF generator. You must supply your own name.

```
[DUPLICATE SECTION SAMPLE]
; Comments may be inserted in the middle of a section without
breaking the section
filename = bob.mif
refresh = NEVER
command = genmif bob.mif

[duplicate section sample]
command = This command does not get executed

[DUPLICATE ENTRY SAMPLE]
filename = joe.mif
refresh = always
filename = This entry is ignored; joe.mif is used as the filename
command = genmif joe.mif

[SAMPLE]
filename = frank.mif
refresh = never
```

```
command = cp mifs\default2.mif frank.mif

[MIFS\TESTTABLE2.MIF]
refresh = Never
command = genmif commandtest
```

If you encounter problems with the .MIFInvExt file, the following suggestions may help:

- Verify that the Group.xx.ComponentName and Group.xx.ClassName properties specified in the.MIFInvExt file match the component name and class name attributes from the MIF file exactly. Spacing and capitalization are significant.
- Verify that the Group.xx.DbTable property (specified in the .MIFInvExt file) matches the table.token property specified in the .TWGdbt file.
- Verify that the Group.xx.Attrib.yy.AttributeId properties (specified in the.MIFInvExt file) match the desired attribute IDs from the MIF file.
- Verify that the Group.xx.Attrib.yy.DbColumn properties for the desired MIF attributes (specified in the.MIFInvExt file) match the corresponding column.zz.token properties specified in the .TWGdbt file.
- Verify that the column.xx.type properties specified in the .TWGdbt file are appropriate to store values retrieved from the MIF file. The default MIF attribute-to-database type mappings are described in the section "Inventory extension property file format" on page 406.

Server initialization and table property fIles

When the Director server starts, it searches the UserTables subdirectory of the server's data directory (usually C:\Program Files\Director\Data\Tables) and loads all of the user table files, which have the extension.TWGDbt, that it finds. It is important to know that Director uses both a third-party DBMS to store data about managed systems, as well as its own persistent storage that contains information related to the server's functions; the table properties are stored in each and must be kept synchronized.

When the server goes through the table property files, if for a given file no matching table is found in the server's persistent storage, a new table is created in the database via the interface to the DBMS, and information about the table's properties is put into the server's persistent storage. If a matching table is found in the persistent store, it is initialized within the server. If a table is found in the server's persistent store but the table property file is missing, that table is removed from persistent storage **and** removed from the database. Therefore, you should be careful about removing table property files for tables you want to keep in the database. If a table property is not processed correctly due to syntactic errors, but enough of the file is correct so that the table's token name can be read, then that table will not be initialized in the server, but its contents in the database will remain intact. As the table is initialized, warnings and errors will be printed to the table status file (located in the table property file directory) as described in a previous section.

Remember that once the server is initialized, a custom table cannot be changed. To make changes to a table, you must stop the server, modify its property file, and restart the server. If the table property file has been changed since the last time the server was started, the table will be changed to reflect changes made to the property file. There are very important restrictions on changes that can be made to a table property file:

1. The following properties **cannot** be changed in a table property file once that table has been successfully initialized within the Director Server: Table token, realname, and shortname; and column tokens, realnames, key values, types, and lengths.
2. The following properties can be changed: any displaynames, any "nls." properties, any "table.filterprompt" properties, and any "column.X.value" properties.
3. Columns cannot be deleted.
4. The indices of the columns cannot be changed.
5. Columns may be added, but new columns must have a higher index than all existing columns.

If you want to make changes to table files that fall under any of the restrictions above, you must remove the old table and then recreate it with the changes. Any data in the table will be lost. The following procedure is recommended:

1. Stop the Director server.
2. Use a database management tool to remove the table from the database.
3. Make changes to the.TWGDbt file, as necessary.
4. Restart the server.

When the server starts up, it will re-create the table using the new property file.

If you cannot manipulate the database to remove the table, use this procedure:

1. Stop theDirector server.
2. Delete the property file for the table.
3. Start the server. The server will remove the table from the database for you when it does not find its property file.
4. Stop the server.
5. Replace the property file for the table with the new changes.
6. Restart the server.

Note: The server will not start up unless **all** database tables are successfully initialized, including custom user tables. Thus, errors in the user table property files can cause the server to not initialize, or cause the inventory or database components to stop (for example, if types in the table property file do not match those of the collected data).

There are no restrictions on how inventory extension property files can be changed, as long as they stay valid. You must be careful with the use of comments; if a group attribute property is “commented out,” the remaining attributes must have their indices changed so that the indices of the remaining attributes start at 1 and increase sequentially, or else all attributes after the “commented out” attribute will not be found.

Examples

These and other examples can be found in the /TivoliWg/Data/UserTables and /TivoliWg/Data/InvExtension directories. In these directories, the filenames end in “.sample.” To run these samples, the files must be renamed to remove the “.sample” file extension.

Example 1: CIM BIOS: Collect information from four fields in the Win32_BIOS class:

```
CIM_BIOS.TWGdbt:
table.token=CIM_BIOS
table.realname=CIM_BIOS
table.displayname=CIM BIOS Default

nls.0.locale=en
nls.0.filename=CIM_BIOS.en

column.1.token=BUILD_NUMBER
column.1.realname=BUILD_NUMBER
column.1.displayname=Build Number Default
column.1.type=CHAR
column.1.length=80

column.2.token=RELEASE_DATE
column.2.realname=RELEASE_DATE
column.2.displayname=Release Date Default
column.2.type=DATETIME

column.3.token=VERSION
column.3.realname=VERSION
column.3.displayname=Version Default
column.3.type=CHAR
column.3.length=80

column.4.token=DESCRIPTION
column.4.realname=DESCRIPTION
column.4.displayname=Description Default
column.4.type=CHAR
column.4.length=80

CIM_BIOS.en:
TableName.TWGDbUserTable?CIM_BIOS=CIM BIOS English
ColumnName.TWGDbUserTable?CIM_BIOS.BUILD_NUMBER=Build Number
English
```

```
ColumnName.TWGDbUserTable?CIM_BIOS.RELEASE_DATE=Release Date
English
ColumnName.TWGDbUserTable?CIM_BIOS.VERISON=Version English
ColumnName.TWGDbUserTable?CIM_BIOS.DESCRIPTION=Description
English

CIM_BIOS.CIMInvExt:
# This CIM ClassName exists under:
#
# root\CIMV2
#  CIM_ManagedSystemElement
#   CIM_LogicalElement
#    CIM_SoftwareElement
#     CIM_BIOSElment
#

Group.1.ClassName=Win32_BIOS
Group.1.NameSpace=root/cimv2
Group.1.DbTable=CIM_BIOS

Group.1.Attrib.1.Property=BuildNumber
Group.1.Attrib.1.DbColumn=BUILD_NUMBER

Group.1.Attrib.2.Property=ReleaseDate
Group.1.Attrib.2.DbColumn=RELEASE_DATE

Group.1.Attrib.3.Property=Version
Group.1.Attrib.3.DbColumn=VERSION

Group.1.Attrib.4.Property=Description
Group.1.Attrib.4.DbColumn=DESCRIPTION
```

Example 2: DMI Component ID: Collect information from five fields in the ComponentID class of the Win32 Service Layer component:

```
DMI_WIN32_COMP_ID.TWGdbt:
table.token=DMI_WIN32_COMP_ID
table.realname=DMI_WIN32_COMP_ID
table.displayname=DMI Component ID Default

nls.0.locale=en
nls.0.filename=DMI_WIN32_COMP_ID.en

column.1.token=MANUFACTURER
column.1.realname=MANUFACTURER
column.1.displayname=Manufacturer Default
column.1.type=CHAR
column.1.length=80

column.2.token=PRODUCT
column.2.realname=PRODUCT
column.2.displayname=Product Default
column.2.type=CHAR
column.2.length=80

column.3.token=VERSION
column.3.realname=VERSION
column.3.displayname=Version Default
column.3.type=CHAR
column.3.length=80

column.4.token=SERIAL_NUMBER
column.4.realname=SERIAL_NUMBER
column.4.displayname=Serial Number Default
column.4.type=CHAR
column.4.length=80

column.5.token=INSTALL_DATE
column.5.realname=INSTALL_DATE
column.5.displayname=Install Date Default
column.5.type=DATETIME

DMI_WIN32_COMP_ID.en:
TableName.TWGDbUserTable?DMI_WIN32_COMP_ID=DMI Name Table English

ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.MANUFACTURER=Manufac
turer English
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.PRODUCT=Product
English
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.VERSION=Version
English
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.SERIAL_NUMBER=Serial
Number English
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.INSTALL_DATE=Install
Date English

DMI_WIN32_COMP_ID.DMIInvExt:
Group.1.ComponentName=Win32 DMI Service Provider
Group.1.ClassName=DMTF|ComponentID|001
Group.1.DbTable=DMI_WIN32_COMP_ID

Group.1.Attrib.1.AttributeId=1
Group.1.Attrib.1.DbColumn=MANUFACTURER

Group.1.Attrib.2.AttributeId=2
Group.1.Attrib.2.DbColumn=PRODUCT

Group.1.Attrib.3.AttributeId=3
Group.1.Attrib.3.DbColumn=VERSION

Group.1.Attrib.4.AttributeId=4
Group.1.Attrib.4.DbColumn=SERIAL_NUMBER
Group.1.Attrib.5.AttributeId=5
Group.1.Attrib.5.DbColumn=INSTALL_DATE


Group.2.ComponentName=DMTF Developers
Group.2.GroupName=DMTF|DevNames|1.0
Group.2.DbTable=DMI_NAME_TABLE

Group.2.Attrib.1.AttributeId=1
Group.2.Attrib.1.DbColumn=INDEX

Group.2.Attrib.2.AttributeId=2
Group.2.Attrib.2.DbColumn=NAME

Group.2.Attrib.3.AttributeId=3
Group.2.Attrib.3.DbColumn=COMPANY

Group.2.Attrib.4.AttributeId=4
Group.2.Attrib.4.DbColumn=OP_SYS
```

Using the database configuration window to convert to another database

To convert to another database, run the respective command below (corresponding to the database you want to convert to), to display the Database Configuration window. For more information on using the Database Configuration process, see Chapter 3, “Installation and configuration,” on page 39, and the online help.

- **cfgmssql** - Microsoft SQL Server Database
- **cfgdb2** - IBM DB2 Universal Database
- **cfgoracle** - Oracle Server Database

Appendix E. Agent-server security

Agent-server security is an authentication process used to establish trust relationships between the Director server and Director agents when the network is brought up. This appendix describes the process and files used by Director to implement agent-server security, and provides guidelines for:

- Initializing managed systems securely
- Determining the origin of public or private keys
- Recovering lost public or private key files

How Director agent-server security is implemented

Director provides a means of security by which a managed system configured with Director management agent (agent) can authenticate an Director server (server) attempting to manage it. Authentication enables an agent to accept only management operations from servers authorized to manage it. Authentication protects agents and servers from access by unauthorized servers or “rogue” agent applications.

Agent-server security is different than the user-logon security used for controlling an administrator's access to an Director server, which controls the administrator's ability to issue requests to the Director server and agents through the Director Management Console.

Agent-server security is based on two core concepts: agent secure/unsecure state and public-private signature authentication. Agent secure/unsecure state refers to the willingness of the agent to accept *any* authorized Director server. If an agent is ***unsecure***, *any* Director server is allowed to manage the system. If the agent is ***secure***, only Director servers that pass authentication are allowed to manage the system.

Public-private authentication is the method used by an Director agent to authenticate an Director server once the agent is secure. Director authentication is based on the DSA digital signature scheme, a public-private key based algorithm that allows holders of a public key to verify the signature for a digital document which has been signed by a holder of the corresponding private key. In Director, when a server attempts to access an agent, the server “bids” the public keys corresponding to the private keys that it holds. An agent checks these keys, and if any are considered trusted by the agent, the agent replies with a challenge consisting of one of the trusted public keys and a random data block. The server then generates a digital signature of the random data block using the private key corresponding to the public key included in the challenge and sends the signature back to the agent. The agent then uses the public key to verify that the signature is a valid signature for the random block using the selected public key and grants access if the signature checks. If access is not granted, the server marks the system inaccessible (which is displayed as a small padlock icon next to the system icon on the Director Management Console).

The benefit of this scheme, versus a userid-password scheme, is that the public keys stored on the agents are usable only for verifying access, not for requesting access. Also, generating a private key corresponding to a given public key is cryptographically improbable, requiring on order of 2^128 or more operations to accomplish (i.e., theoretically, all the computers in the world working for billions of years or more). Also, the use of the random data block for signing makes replay attacks unusable.

The configuration information for agent-server security is stored in several files on both the server and agent machines. On Windows, Windows 9*x*, Windows NT, and Novell NetWare systems, the files are in the *x:*\tivoliwg\data directory. On OS/2, the files are in the *x:*\tivoliwg directory. The secure/unsecure state data is stored in the **secin.ini** file, which is generated if needed when the **twgipc.exe** first starts on a system. On Director servers, this file is initialized as secure by default, while on agents it is initialized as unsecure.

The public keys trusted by the agent (and the server, which is a super-set of the agent) are stored in files named dsa*xxxxx*.pub, where *xxxxx* is a unique identifier matching the name of the corresponding private key file (i.e., dsa23ef4.pub is the public key corresponding to the private key stored in dsa23ef4.pvt). The private keys held by a server are stored in files named dsa*xxxxx*.pvt. When an Director server is started, if no dsa*.pvt files are found, it randomly generates a matching set of public and private key files. The server then loads any dsa*.pvt files, and uses them for proving its identity. When any type of Director agent starts (including a server), it loads any dsa*.pub files it finds, and considers these keys trusted.

Note: The files are only loaded at the startup of **twgipc.exe**; adding or deleting files has no effect until the agent is restarted.

The contents of **secin.ini** are also loaded and used to control whether the agent is secured or unsecured.

When an Director server first communicates with an agent, including discovery and when the agent is first found to be online, it requests access. If access is granted (either due to the agent being unsecured or the server having a private key matching one of the public keys trusted by the agent), the server delivers copies of the public key corresponding to each of its private keys. This action assures that the agent will continue to trust the server if the agent is currently unsecure but is later secured. Next, if the **Automatically secure unsecure systems** option on the Director Systems tab of the Server Preferences window has been set, the agent is ordered to become secure. This order causes future servers which have private keys not currently trusted by the agent to be denied access, but allows any servers currently trusted to continue to access the agent (that., securing an agent does not revoke access by other trusted servers, only access by untrusted servers). Agents can also be secured or unsecured using the Secure System and Unsecure System context menu choices on the Director Management Console.

Installing Director agents in a secure state

The Director management console supports a "request access" function to initiate an access request from the Director server to Director management agents running in a secure state on Windows NT. This function is a context menu item that can be used as an alternative to copying *.PUB files from an Director server to a Director management agent in a secure state. Refer to the online help for more information.

To install Director agents in a secure state, use the following procedure:

1. Install and start any Director servers you want to use to administer the agents. Each server will create a set of dsa*.pub and dsa*.pvt files, as well as a **secin.ini** file set to secure. Get a copy of the dsa*.pub file from each server, as well as a **secin.ini** from one of the servers. Place these files onto a file server or similar location which will be accessible to the agent installation procedures.
2. After each agent is installed, but before the system is restarted, copy the dsa*.pub files and the **secin.ini** file into the appropriate directory (*x:*\Program Files\IBM Director Agent\Director\data for Windows Agents, *x:*\Program Files\Director\data for Windows Server, *x:*\tivoliwg for OS/2 and Novell). When started, the agent will be secure and only trust the desired servers.
3. If an agent has previously been started unsecurely, stop the agent (using **net stop twgipc** on WinNT, **twgipc shutdown** on Windows 9*x*, and OS/2, and **unload twgipc** on Novell), delete all dsa*.pub files, and copy the desired dsa*.pub and **secin.ini** files into the directory. When restarted (**net start twgipc** on WinNT, **start twgipc** on Windows 9*x*, **twgipc start** on OS/2, **load twgipc** on Novell NetWare), the agent will be secure and only trust the desired server(s). This procedure can be used in logon scripts or other automatic execution mechanisms. To add another trusted server to an existing secure environment, you can do any of the following:
   a. Setup the new server, and copy its dsa*.pvt file to one of the other trusted servers. Restart the other server. As the trusted server initializes, it begins delivering the dsa*.pub corresponding to the new server to all of its trusting agents, which causes them to trust the new server as well.
   b. Setup the new server and copy the dsa*.pvt file from an existing trusted server. This immediately allows the new server to authenticate itself to the other servers' trusting agents. The new server will also be trusted by the other server.
   c. Include the dsa*.pub generated by the new server in the initialization procedure described above. Once completed and restarted, the agents will trust the new server.

Determining the origin of a public or private key

The public and private key files are binary files, but they contain textual data which can be used to show their origin. If a dsa*.pub or dsa*.pvt file is printed using the **type** command at a command prompt, the first line of the data displayed will show:

- A 4-character header
- DSAp*string* for public key files
- DSAP*string* for private key file

Immediately after the 4-character header is a string corresponding to the computer name of the server which generated the key file (for example, `DSAPITDIRECTOR2` indicates a private key file generated by a computer named ITDIRECTOR2).

Recovering lost public and private keys files

It is *very important* to back up and protect the dsa*.pvt files. If lost, these files cannot be regenerated. (If they could be regenerated, they would not be secure.) If the private key file for a server is lost, you need to repeat one of the previously described procedures for initializing security or adding a new trusted server, using either another existing trusted dsa*.pvt key or using the new key generated by the server when it restarts without its private key file.

If a public key file is lost, it can be regenerated by having the server (which holds the corresponding private key) discover, add, or access any unsecured agent (the key file will be generated on the agent). The server does not require the dsa*.pub files corresponding to its own private key files because the private key files include all the information from the public key files and the server always trusts any agent holding a private key matching any of its public or private key files.

Appendix F. IBM Director Agent

Appendix H through Appendix J provides basic information for installing and using IBM Director Agent on your system. Appendix K provides information for installing and using IBM Director Agent Upward Integration Modules (UIMs) on systems running supported system-management applications.

Who should read this guide

This guide is intended for the individuals who are responsible for installing and using IBM Director Agent (on their systems or on remote agent systems in a network environment) and Upward Integration Modules (on supported system-management platforms). This guide assumes that you have extensive knowledge of server, mobile, and desktop system hardware; operating systems; Windows networking; desktop environments; system-management tasks; and the features and functionality that are provided by the supported system-management platforms.

What this guide contains

This guide contains the following appendices:

- Appendix G, "Preparing to install IBM Director Agent," on page 423 provides directions for starting IBM Director Agent on your IBM system.
- Appendix H, "About IBM Director Agent," on page 429, provides an overview of IBM Director Agent.
- Appendix I, "Installing IBM Director Agent," on page 431, guides you through the installation process for installing IBM Director Agent on supported operating systems.
- Appendix J, "Using IBM Director Agent," on page 437, provides an overview of the IBM Director Agent console and includes brief descriptions of IBM Director Agent Information and Task functions.
- Appendix K, "Upward Integration Modules," on page 505, provides information on installing and using UIMs with supported system-management applications.

Getting help

IBM Director Agent includes an online troubleshooting guide that contains solutions for many common installation and usage issues. This troubleshooting guide is installed when you install IBM Director Agent, and it can be accessed from the **Troubleshooting Guide** option in your **Start** menu.

Additional information and IBM Director Agent product updates, when available, can be obtained from the IBM Web site at:

http://www.pc.ibm.com

Appendix G. Preparing to install IBM Director Agent

This section contains instructions for installing IBM Director Agent on the agent systems. Depending on your system-management environment, you can choose to install all of IBM Director Agent, or you can choose to install only selected portions of IBM Director Agent.

The IBM Director Agent installation program also includes support for installing Workgroup/Enterprise Integration on agent systems running supported system-management platforms. Workgroup/Enterprise Integration adds an Upward Integration Module (UIM) to a supported system-management platform, such as CA Unicenter TNG Framework or Microsoft SMS. With UIMs, you can use your system-management software to manage agents running the IBM Director Agent agent software. Workgroup/Enterprise Integration installation procedures are explained in Appendix K, "Upward Integration Modules," on page 505.

Preparing to install IBM Director Agent

As noted in the previous chapter, IBM Director Agent includes a number of optional components that add value in a variety of system-management environments. Before you begin your installation, determine whether you will be using IBM Director Agent as a stand-alone agent-based system-management solution, or whether you will be using IBM Director Agent to gather data for a supported system-management platform.

Also, IBM Director Agent requires certain minimum hardware and software for installation. These requirements are listed on page 425.

Supported system-management environments

The IBM Director Agent components that you choose to install are largely determined by the system-management environment in which you are installing IBM Director Agent. Some components are of use to most system-management platforms, while others are of use only to specific platforms. The following sections offer example component selections based on the system-management platforms that are used to manage systems in the network.

IBM Director Agent console

If you will be using a Web browser or MMC to manage IBM Director Agent, agent systems install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- Web Based Remote Control

Note: Systems using a Web browser or MMC to access IBM Director Agent locally require 64 MB of RAM to function properly.

Tivoli Enterprise

If you use Tivoli Enterprise to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- Tivoli Management Agent
- SNMP Access and Trap Forwarding

Tivoli NetView

If you use Tivoli NetView® to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- Web Based Remote Control

Microsoft 2.0

If you use Microsoft SMS 2.0 to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring

Note: For SMS 2.0, the UIM must be installed on all site servers and all management consoles

CA Unicenter TNG Framework

If you use CA Unicenter TNG to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- SNMP Access and Trap Forwarding

LANDesk Management Suite

If you use LANDesk Management Suite to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- LANDesk Management Suite Integration (installs Common Base Agent)

HP Openview

If you use HP Openview to manage the agent systems in your network, install the following IBM Director Agent components on the agent systems:

- Web Based Access
- System Health Monitoring
- SNMP Access and Trap Forwarding

Installation requirements

Before you install IBM Director Agent, consider the following installation requirements:

- Hardware requirements
- Supported operating systems
- Supported browsers

Hardware requirements

The IBM Director Agent agent in a Microsoft® Windows® operating system requires the following hardware, memory, and disk space:

- An IBM server, IBM desktop computer, IBM IntelliStation computer, IBM ThinkPad mobile computer, or an OEM system that supports SMBIOS 2.1 or later.

  Note: Agent systems must support SMBIOS version 2.0 or later.
- An Intel Pentium 200 MHz or faster processor
- 75 MB of hard disk space on the agent systems
- A minimum of 32 MB random access memory (RAM), or the recommended minimum for the operating system

Supported operating systems

The following operating systems support IBM Director Agent:

- Windows 2000—Server or Advanced Server
- Windows 2000 Professional
- Windows NT® Server 4.0 (with Service Pack 4 or later)
- Windows NT Workstation 4 (with Service Pack 4 or later)
- Windows 98
- Windows Millennium Edition

- Windows XP Professional

Supported browsers

A World Wide Web browser is needed on the system from which you plan to manage remote IBM Director Agent systems and is required only if you are planning to install the IBM Director Agent Web Based Access or Web Based Remote Control options.

The following browsers support the IBM Director Agent console:

- Microsoft Internet Explorer 4.01 or later with 56-bit or 128-bit encryption.

  Notes:

  1. If you are using Internet Explorer 5.x or 6.x you must install the optional Java Virtual Machine (VM) support to access an agent system running IBM Director Agent.
  2. If you are using Internet Explorer and you reinstall Internet Explorer after installing IBM Director Agent, you must reapply the Microsoft VM update. The IBM Director Agent requires the Microsoft VM Build 3165 or later. Download the latest Microsoft VM from http://www.microsoft.com/java.
  3. The supported browser must have the file and URL associated with opening the browser.
- Microsoft Management Console (MMC) 1.1 or later.

  If you install IBM Director Agent before you install MMC, a Microsoft Management Console icon will not appear in the IBM Director Agent section of your **Start** menu.
- Netscape Navigator or Netscape Communicator 4.51 with 56-bit or 128-bit encryption.

Additional installation guidelines

Before you install IBM Director Agent, consider these additional restrictions, requirements, and installation options:

- **Default directory**
  By default, the IBM Director Agent installation program installs the IBM Director Agent program files in C:\Program Files\IBM\IBM Director **Agent**. If you do not want to install the program files in the default location, be prepared to provide an alternative installation drive and directory.
- **Microsoft Management Console (MMC) 1.1 or later**
  You can use IBM Director Agent after installing MMC on systems running Windows 98, Windows 2000, Windows Millennium Edition or Windows NT.
- **Running an unattended installation**
  IBM Director Agent supports unattended installation. If you want to use

your software distribution facility to install IBM Director Agent, you can set up an unattended installation for IBM Director Agent. IBM Director Agent can then be installed across your network from one central location. The format of the IBM Director Agent response file, named SETUP.ISS, is described in “Modifying the SETUP.ISS File Manually” in the *Director User's Guide* on the *Director with IBM Director Agent* CD.

- **Using Microsoft Internet Information Server**

  IBM Director Agent does not provide automatic integration with Microsoft Internet Information Server. The IIS configuration is a manual process. The configuration steps are as follows:

  1. Create an IBM Director Agent Web Site in IIS.

     From Internet Services Manager, select your host computer and select **create a New Web Site**. This starts the Web Site Creation wizard. Name the site *IBM Director Agent*. If you will be accessing IBM Director Agent only through IIS, assign port 411 to this site.   If you will be configuring the IBM Director Agent HTTPserv service to start automatically at startup, you might want to assign a different port to the IIS site. The primary IBM Director Agent provider should use port 411.
  2. Assign the home directory path for the site to point to the IBM Director Agent httpserv directory. If you choose the installations default, this is c:\Program Files\UMS \httpserv. Disallow anonymous access to the site by clearing the **Allow anonymous access to this Web site** check box.
  3. Assign this directory read permissions only.
  4. Enable Server-Side Includes for the IBM Director Agent site.

     Display the IBM Director Agent sites properties and click the **Home Directory** tab. Click the configuration button that is in the middle right of the bottom section of the window. Make a new entry to start ssinc.dll from the System32\inetsrv subdirectory of your System directory. You can specify that ssinc.dll will be started only on GET, POST actions.
  5. Enable **Read Access** and **Script Execution**.

     The Home Directory settings for the site must allow read access. Set **Execute Permissions** to **Scripts only**.
  6. Disable Anonymous Access.

     The IBM Director Agent windows will not be displayed properly if Anonymous Access is enabled for the root. If you did not turn off anonymous access when creating the site, do so now. Display the IBM Director Agent sites properties and click the **Directory Security** tab. In the **Anonymous access and authentication** control section, click the **Edit** tab. Clear the **Anonymous Access** check box. Depending on your security requirements, you might also want to enable Basic Authentication for access from Netscape.
  7. Enable CGI Script Execution for cgi-bin.

Expand the IBM Director Agent site tree and right-click the Directory to display the properties for the cgi-bin directory. In the middle of the window, remove read access. At the bottom of the page, change the **Execute Permissions** to **Scripts and Executables**.

8. Create an IBM Director Agent CGI account (optional).

   At this point, users will not be able to view any IBM Director Agent information under any of the following conditions:

   — The user's account is not a member of the Administrative group.

   — The browser is Internet Explorer.

   — The browser system is not using Windows 2000 or the user is authenticated through the NT Challenge Response mechanism (NTLM) or Kerberos (Negotiate), not Basic Authentication (BASE64).

   If you would like users to be able to view information, you must configure an account to access the WMI. Name the account Director AgentCGI and set the password to ibmdirectoragentcgiGuest.

Appendix H. About IBM Director Agent

IBM Director Agent is a lightweight agent that resides on managed systems. It provides a suite of graphical user interfaces that enhance the local or remote administration, monitoring, and maintenance of IBM systems, such as ThinkPad® computers, IntelliStation® computers, and @server xSeries servers.

With IBM Director Agent, an agent-system user or remote system administrator can use a Web browser or the Microsoft® Management Console (MMC) and the IBM Director Agent console support to inventory, monitor, and troubleshoot IBM systems on which IBM Director Agent is installed.

This “point-to-point” system-management approach enhances support and enables a system administrator to effectively maintain IBM systems without having to install additional system-management software on the administrator console.

IBM Director Agent also includes support for Upward Integration Modules (UIMs). A system administrator who uses any supported system-management platform (such as Tivoli® Enterprise, CA Unicenter TNG Framework, Microsoft Systems Management Server, Intel LANDesk Management Suite or HP OpenView) can use UIMs to integrate portions of IBM Director Agent into the administrator console. Because it is designed to use industry-standard information-gathering technologies and messaging protocols (such as Common Information Model, Desktop Management Interface, and simple network management protocol), IBM Director Agent adds value to any of these supported workgroup or enterprise system-management platforms.

Please refer to Appendix I, “Installing IBM Director Agent,” on page 431 for a description of the available components (services) that you can install on local IBM Director Agent systems.

Appendix I. Installing IBM Director Agent

This section guides you through the installation process for installing IBM Director Agent on supported operating systems.

To install IBM Director Agent, use the following procedure:.

1. Insert the *Director with IBM Director Extensions* CD in the CD-ROM drive. The Director with IBM Director Extensions window opens.
2. Click **Install Director**. The Welcome window opens.
3. Click **Next.** The License Agreement window opens. Click **Yes** to proceed. You must agree to the terms of the License Agreement to install IBM Director Agent. If you click **No**, the installation program will close.
4. Click **Next**. The Select Components window opens. The installation choices are Server, Console, Agent, and Workgroup/Enterprise Integration.

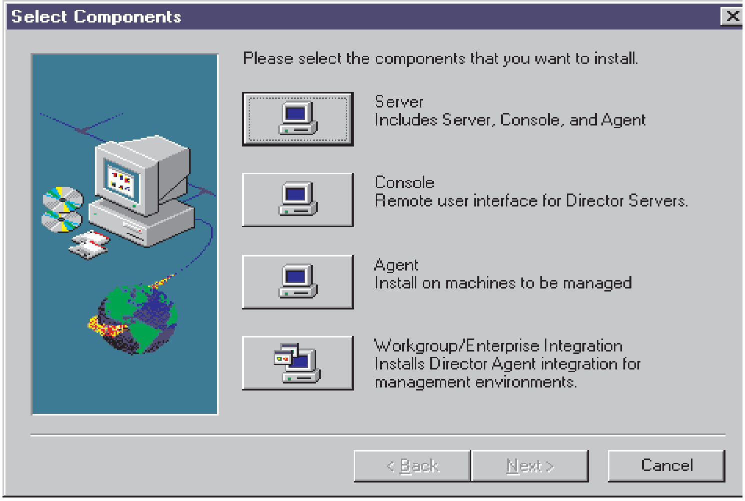

Select Components

Please select the components that you want to install.

Server
Includes Server, Console, and Agent

Console
Remote user interface for Director Servers.

Agent
Install on machines to be managed

Workgroup/Enterprise Integration
Installs Director Agent integration for management environments.

5. Click **Agent**.

   The IBM Director Agent Configuration window opens.

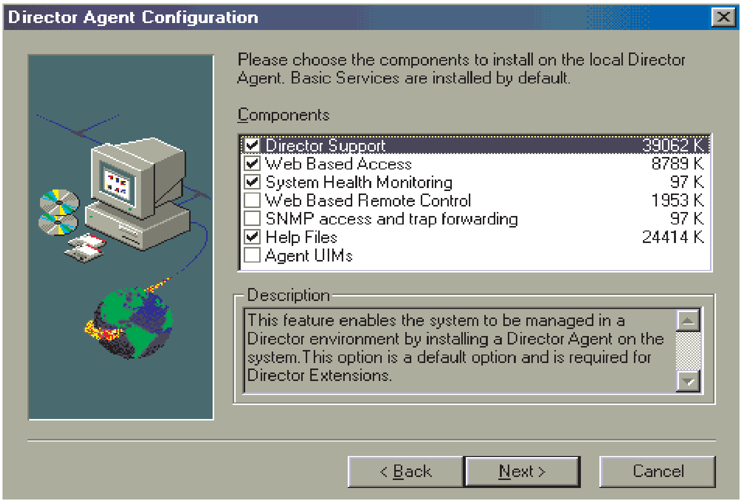

Please choose the components to install on the local Director Agent. Basic Services are installed by default.

Components

☑ Director Support	39062 K
☑ Web Based Access	8789 K
☑ System Health Monitoring	97 K
□ Web Based Remote Control	1953 K
□ SNMP access and trap forwarding	97 K
☑ Help Files	24414 K
□ Agent UIMs	

Description

This feature enables the system to be managed in a Director environment by installing a Director Agent on the system.This option is a default option and is required for Director Extensions.

6. Select the check box beside any of the components that you want to install on the agent system.

   The following optional components are available. The components that are selected by default are indicated as such.

   Director Support (default)
   Director Support is an additional configuration option for the agent installation only. Director is an advanced Intel-processor-based workgroup hardware manager, with centralized agent and group management console and server services. Selecting this feature enables the agent system to be managed in a Director environment by installing IBM Director Agent on this system.

   Web Based Access (default)
   Web Based Access offers a convenient Java-based tool for managing an agent system and for viewing the CIM-based inventory data. If you install Web Based Access, a hypertext transport protocol (HTTP) DAEMON is installed and requires that a user name and password be typed during the installation. The user name and password limit access to the HTTP DAEMON. With Web-Based Access that is installed on the agent system, the agent system can be managed from any remote computer with a supported Web browser. The Web browser is the only software that is needed on the remote system.

   System Health Monitoring (default)
   System Health Monitoring provides active monitoring of critical system functions, such as disk space available, SMART drive alerts, system temperature, fan functionality, power supply voltage, and system cover removal (dependent upon the hardware options of a selected managed system). You can use System Health Monitoring to detect system problems early, before system failures occur. System administrators are notified of a system problem by a CIM

event, and SNMP trap (SNMP traps are available only if **SNMP access and trap forwarding** is also selected), or an SMS status message (Microsoft SMS 2.0 only). Critical problems also cause a message to be displayed on the monitor of the agent system.

Web Based Remote Control
Web Based Remote Control enables a remote system administrator using a Web browser or MMC console to take control of the agent system desktop, enhancing the administrator's ability to diagnose system problems and troubleshoot the system.

Note: You must install the Web Based Access component to install the Web Based Remote Control component.

LANDesk™ Management Suite Integration
LANDesk Management Suite Integration installs the Intel Common Base Agent on the agent system. This enables the system administrator to use IBM Director Agent with LANDesk Management Suite.

Tivoli Management Agent
Tivoli Management Agent installs support on the agent system that enables it to be managed by the Tivoli Enterprise system-management platform.

SNMP Access and trap forwarding (default)
This feature enables CIM information to be accessed from systems that use the simple network management protocol (SNMP). If System Health Monitoring is enabled, this option also enables System Health to forward CIM events as SNMP traps. This component requires that you have the SNMP service (provided with the operating system) installed on the endpoint. If the SNMP service is not installed, the system prompts you to insert the operating system installation media and install SNMP during the IBM Director Agent installation.

Help Files (default)
Selecting this component installs online documentation. Do not select this option if you are concerned about disk space or do not need online documentation installed on every agent system.

7. Click **Next.** If you selected IBM Director Extensions as an installation option, then the IBM Director Extensions Install Option window opens.
8. Select the check box beside any component that you want to install.
9. Click **Next**. If you selected Agent UIM as an installation option, the Agent UIM Install Option window opens.
10. Select the check box beside any option that you want to install.
11. Click **Next**.
12. Select an IP port number for the Apache server.
13. Click **Next**. The "Add icons for IBM Director Agent" window opens.

14. Click **Yes** or **No** to place IBM Director Agent icons on the start menu. The installation program begins installing the necessary files. The Director Setup is complete window opens when the process is complete.
15. When the Director Remote Control window opens, click **Yes** or **No**. A second Director Remote Control window opens.
16. Click **Yes** or **No** to require user authorization for access.
17. Restart the computer now or restart later. If you click **Restart now**, the system shuts down and restarts immediately. If you choose to Restart later, the installation program closes, However, you must restart and log in to the system to begin using IBM Director Agent.
18. Click **Finish**.

Uninstalling IBM Director Agent

You can uninstall IBM Director Agent through the Add/Remove Programs feature in the Windows Control Panel.

To uninstall IBM Director Agent, use the following procedure:

1. Click **Start** → **Settings** →**Control Panel**. The Control Panel opens.
2. Click the **Add/Remove Programs** icon. The **Add/Remove Programs Properties** window opens.
3. From the **Install/Uninstall** list, click **Director**, and then click **Add/Remove**.
4. Select Director, and click the **Add/Remove** button.

The uninstallation process might take a while to complete.

Starting the IBM Director Agent Console

If the Web Based Access or Web Based Remote Control optional component is installed on a agent system, you can use a Web browser or MMC to access and manage the agent locally or remotely.

Note: IBM provides a Java class library (Swing/JFC with IBM Director Agent. Install this library on the Web browser before you access IBM Director Agent data. The first time you use a Web browser for IBM Director Agent, you will be asked to download the installation programs for this file. Click the HTML link to begin installation of the library. For Windows 98, you will have to restart your IBM Director Agent Web browser for this Java library to take effect.

Starting the IBM Director Agent Browser on a local system

During installation, if you clicked **Yes** to create icons on the Start menu, you can start IBM Director Agent locally from this menu.

Note: Systems using a Web browser or MMC to access IBM Director Agent locally require 64 MB of RAM to function properly.

To start IBM Director Agent, from the local system click **Start → Programs → Director Agent →Director Agent Browser**.

The system starts the default Web browser and opens it to the Web address

http://localhost:tcpip_port

where *tcpip_port* is the TCPIP port that you selected during installation. You must type your user ID and password in the IBM Director Agent User ID and Password window.

Starting IBM Director Agent browser remotely

You can start IBM Director Agent remotely using a supported Web browser. In the address field of your browser, type

http://systemname:tcpip_port

where *systemname* is the TCP/IP address or the Internet Protocol (IP) address of the agent and *tcpip_port* is the port number that is assigned for use by the IBM Director Agent console during IBM Director Agent installation. Port number 423 is the default, but if this port is being used by another application, IBM Director Agent could also have been configured to use port number 6411, 6500, 6600, or 6611.

Starting IBM Director Agent from Microsoft Management Console

If you installed Microsoft Management Console (MMC) 1.1 and then installed IBM Director Agent on the agent system, an icon is created on the Start menu from which you can start IBM Director Agent in the MMC.

To start IBM Director Agent, click **Start→Programs→Director Agent →Microsoft Management Console.**

The Microsoft Management Console is available as part of the Windows NT Option Pack 4 or from:

http://www.microsoft.com/MANAGEMENT/MMC

Starting IBM Director Agent from a UIM Management Console

If IBM Director Agent is integrated into Tivoli Enterprise 3.6, 3.62, 3.7.1, Tivoli NetView 5.1.1, SMS 2.0, CA Unicenter TNG 2.4 (Windows 98, Windows 2000, Windows XP, or Windows NT only), or HP OpenView, you can start IBM Director Agent directly from the management console. The management console starts either the default Web browser or the Microsoft Management Console (whichever is appropriate for the workgroup or enterprise environment). For more information, see Appendix K, “Upward Integration Modules,” on page 505.

Appendix J. Using IBM Director Agent

When you have connected to an agent system, the IBM Director Agent console opens in your Web browser or MMC. The console has two panes.

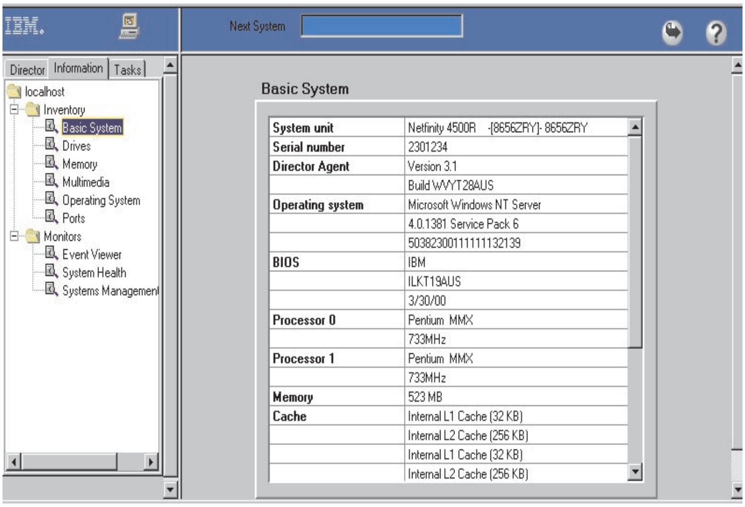

Basic System

System unit	Netfinity 4500R -[8656ZRY]- 8656ZRY
Serial number	2301234
Director Agent	Version 3.1
	Build WVYT28AUS
Operating system	Microsoft Windows NT Server
	4.0.1381 Service Pack 6
	50382300111111132139
BIOS	IBM
	ILKT19AUS
	3/30/00
Processor 0	Pentium MMX
	733MHz
Processor 1	Pentium MMX
	733MHz
Memory	523 MB
Cache	Internal L1 Cache (32 KB)
	Internal L2 Cache (256 KB)
	Internal L1 Cache (32 KB)
	Internal L2 Cache (256 KB)

The Services pane is on the left side of the IBM Director Agent console and contains three pages. Each page contains a list of IBM Director Agent services that are available on the agent system. The following pages are available from the Services pane:

- **Director** - This page contains an expandable view of Hardware Status. It only appears on systems that also have Director server installed.
- **Information —** This page contains an expandable tree view of the IBM Director Agent services that are used to gather hardware and software information from the agent system.
- **Tasks —** This page contains an expandable tree view of the IBM Director Agent services that are used to perform system-management and system-configuration tasks on the agent system.

The Display pane is on the right side of the IBM Director Agent console and is a dynamic view that displays the interfaces and data that are associated with the service that is selected from the **Director**, **Information,** or **Tasks** page.

The following field and command icons also appear in the Display pane of the IBM Director Agent console:

Next System field

You can use the same browser window or MMC to access multiple IBM Director Agent systems. In this field, type the TCP/IP address of another agent system running IBM Director Agent, and then press Enter to access another agent system without opening another Web browser. This field is not available on MMC.

Export icon ➡

With IBM Director Agent, you can create comma-separated-value (CSV) data files from the hardware and software data that is collected by many of the IBM Director Agent interfaces. These CSV files can be imported into many database programs so that you can create a centralized repository for data that is collected by IBM Director Agent. To create a CSV file, select a service from the Services pane. When IBM Director Agent has completed loading data, click the **Export** icon. A browser window opens with a **File Download** box. You must choose to open the file from its current location, save the file to a disk, or cancel the export. To save the data, click **Save** from the **File** menu.

Help icon ?

Online help is available for IBM Director Agent. To access online help for a service, select the service from the Services pane and, when it has finished loading, click the **Help** icon.

Viewing Director services

The **Director** page with the Hardware Status function is available when you browse a system with the Director server installed.

The Hardware Status function requires a login to the Director server before access is allowed. A login screen is presented, where you must type a network user name and password. This user name and password must have been previously established on the Director Server.

The Hardware Status task identifies systems that require attention. It combines the status of systems into an overall system health that is presented in Hardware Status. The status categorizes system health into four groups: critical, warning, information, or normal. When an event is recorded, an icon is activated for the appropriate severity, and the system is identified in a list under the respective icon.

To access additional information, click an icon to see a list of systems being monitored, or double-click a listed system to receive data specific to that system.

Hardware Status monitors systems for changes in the following environments:

- Generic
- Network
- Storage
- Environmental

- Security

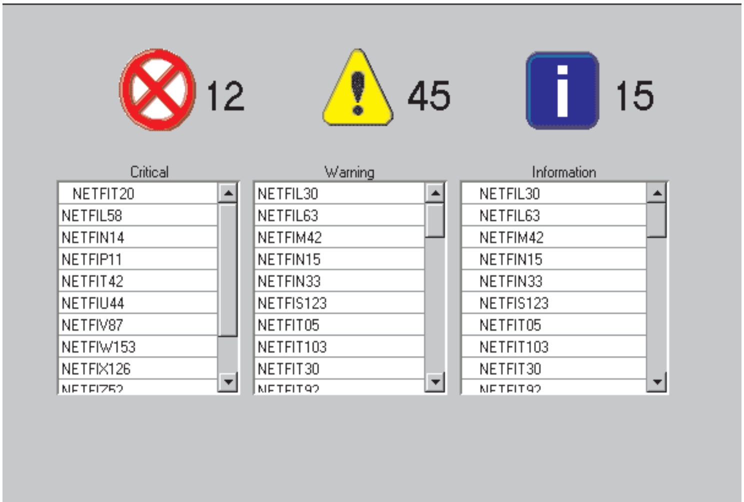

⊗ 12 ⚠ 45 ℹ 15

Critical	Warning	Information
NETFIT20	NETFIL30	NETFIL30
NETFIL58	NETFIL63	NETFIL63
NETFIN14	NETFIM42	NETFIM42
NETFIP11	NETFIN15	NETFIN15
NETFIT42	NETFIN33	NETFIN33
NETFIU44	NETFIS123	NETFIS123
NETFIV87	NETFIT05	NETFIT05
NETFIW153	NETFIT103	NETFIT103
NETFIX126	NETFIT30	NETFIT30
NETFIZ52	NETFIT92	NETFIT92

The Hardware Status interface provides information about the items that are described in the following table.

Item	Description
Icons	An overview of all systems is presented with these icons. When there are no events, the icon is outlined, but not displayed. When there are events, the number of events is listed to the right of the displayed icon.
List Box	The systems with events are listed in this area.
Operating system	The name, version number, and service pack level (if applicable) of the operating system that is running on the agent system.
⊗	The Critical icon is activated when a critical event has been reported. A critical event is defined as an event that requires immediate attention and action.
⚠	The Warning icon is activated when a warning event has been reported. A warning event is defined as an event that needs attention soon.
ℹ	The Information icon is activated when an information event has been reported. An information event is defined as an event that notifies control about a situation.

Viewing Information services

The services that are available from the Information page gather hardware information and software information from the agent system. This data is gathered directly from the agent system and represents the physical components of the system or the current, monitored state of the agent system as reported by monitoring hardware and software in the agent system. The data that is presented in the Information Service interface is static. An IBM Director Agent user cannot change or configure the data.

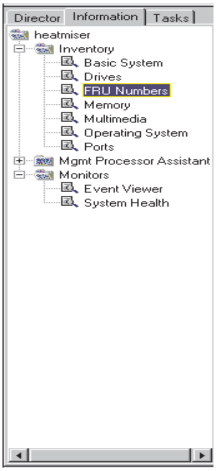

There are two categories in the Information page:

- Inventory
- Monitors

The sections that follow describe each of the services that are available from the Information page.

Inventory Services

Inventory Services gather information about the physical devices that make up the agent system (such as disk drives, multimedia adapters, video adapters, and memory) or the operating system of the agent system. The following Inventory Services are available:

- Basic system
- Drives
- FRU Numbers
- Memory
- Multimedia
- Operating system

- Ports

Descriptions of each of the Inventory services that you can select from the Information page follow.

Basic System

The Basic System service gathers and displays general information about the agent system hardware and operating system.

Note: Not all agent systems have all of the items that can be displayed in the Basic System interface. If an agent system does not have a particular item, the field that is associated with that item will not appear in the Basic System interface.

To start the Basic System service, click **Information→Inventory →Basic system** in the Services pane. The following interface opens in the Display pane.

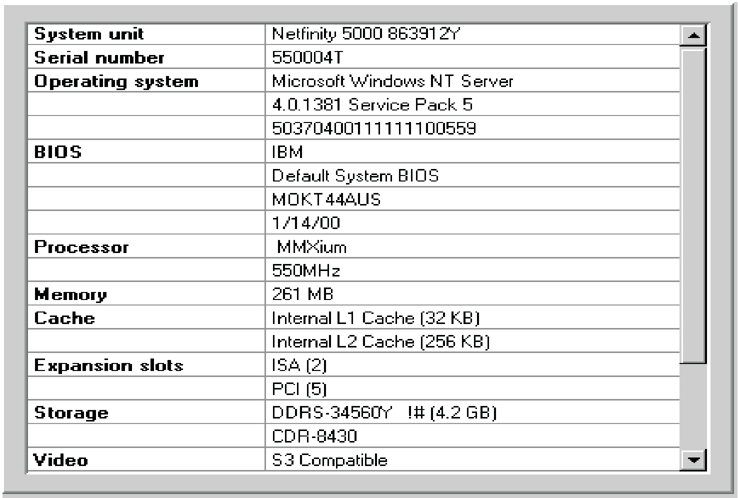

System unit	Netfinity 5000 863912Y
Serial number	550004T
Operating system	Microsoft Windows NT Server
	4.0.1381 Service Pack 5
	503704001111111100559
BIOS	IBM
	Default System BIOS
	MOKT44AUS
	1/14/00
Processor	MMXium
	550MHz
Memory	261 MB
Cache	Internal L1 Cache (32 KB)
	Internal L2 Cache (256 KB)
Expansion slots	ISA (2)
	PCI (5)
Storage	DDRS-34560Y !# (4.2 GB)
	CDR-8430
Video	S3 Compatible

The Basic System interface provides information about the items that are described in the following table.

Item	Description
System unit	The manufacturer and model of the agent system.
Serial number	The serial number of the agent system.
System GUID	The unique identification number for the system.
IBM Director Agent	The version and build number of IBM Director Agent installed on the system.

Item	Description
Operating system	The name, version number, and service pack level (if applicable) of the operating system that is running on the agent system.
BIOS	The version and completion date of the basic input/output system (BIOS) of the agent system.
Processor 0	The type and clock speed of the processor that is installed on the system board of the agent system.
Processor 1	If a second processor is installed, the type is displayed with the MHz Clock Speed.
Memory	The amount of random access memory (RAM) installed in the agent system, in megabytes (MB).
Cache	The amount of microprocessor memory cache that is available to the microprocessor of the agent system.
Expansion slots	The number and type (for example, PCMCIA or PCI) of expansion slots in the agent system.
Storage	The type and size, in megabytes (MB stands for 1 000 000 bytes), of storage devices that are installed in the agent system, such as hard disk drives, CD-ROM drives, or CD-RW (read/write) drives.
Video	The type of video adapter that is installed in the agent system.
Monitor	The monitor type of the system.
Audio	The name of the audio adapter of the system.
Communications	The network interface adapter or modem that is installed in the agent system, the media access control MAC address of the adapter, and the Internet protocol (IP) address of the agent system.
Keyboard	The type of keyboard that is attached to the computer.
Pointing device	The type of pointing device, such as a mouse, trackball, or TrackPoint® device that is attached to the agent system.
AC Line Status	Displays On-Line if the system is plugged into an ac outlet.
Docking Station	If the ThinkPad system is docked into an applicable docking station, *Docked* is displayed. If the ThinkPad is not docked into a docking station, *Not Docked* is displayed.
UPS	Indicates whether or not an uninterruptible power supply is attached to the system. It also indicates the method of attachment.

Drives

The Drives service gathers and displays information about the physical and logical disk drives that are installed in the agent system. To start the Drives service, click **Information→Inventory →Drives** in the Services pane. The following interface opens in the Display pane.

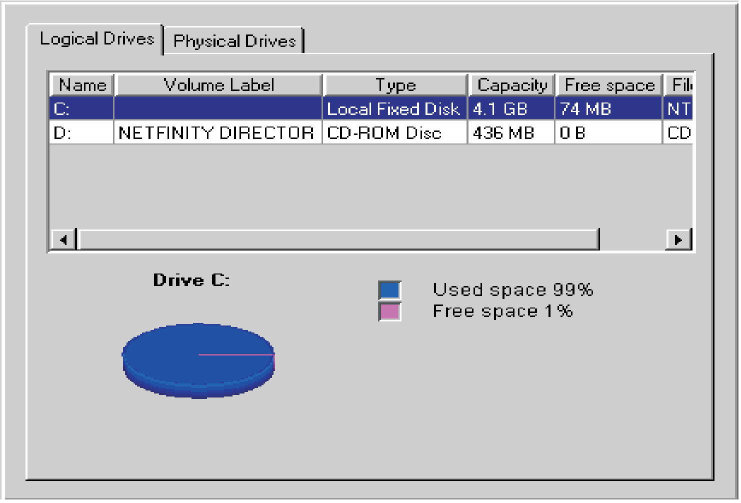

Name	Volume Label	Type	Capacity	Free space	Fil
C:		Local Fixed Disk	4.1 GB	74 MB	NT
D:	NETFINITY DIRECTOR	CD-ROM Disc	436 MB	0 B	CD

The Drives interface contains two pages:

Logical Drives
Click the **Logical Drives** page to display information about the logical drives that are configured on the agent system. The Logical Drives interface is displayed by default. For additional information about each entry on the Logical Drives interface, click any disk row. The interface displays a pie chart that shows used space and free space on the selected logical drive. Used space contains the applications and files that are on the disk, and free space is available for adding files or applications.

The Logical Drives interface provides information about the items that are described in the following table.

Item	Description
Name	The drive letter that is assigned to the logical drive or partition.
Volume Label	The name or label of the storage volume.
Type	The type of logical drive, such as a removable drive.

Item	Description
Capacity	The size of each logical drive, measured in megabytes (MB stands for 1 000 000 bytes) or gigabytes (GB stands for 1 000 000 000 bytes).
Free Space	The amount of disk space available on the logical drive.
File System	The file system in use by the partition.

Physical Drives

Click the **Physical Drives** page to display information about the physical drives installed in the agent system. The Physical Drives page displays the Physical Drives interface.

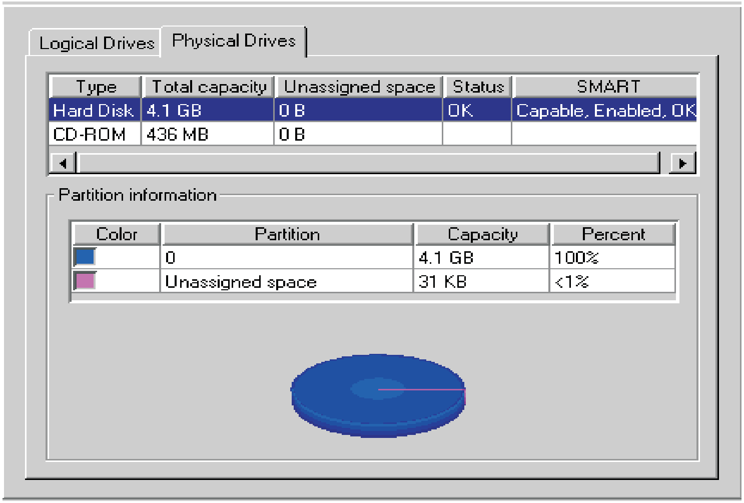

Type	Total capacity	Unassigned space	Status	SMART
Hard Disk	4.1 GB	0 B	OK	Capable, Enabled, OK
CD-ROM	436 MB	0 B		

Color	Partition	Capacity	Percent
🟦	0	4.1 GB	100%
🟪	Unassigned space	31 KB	<1%

The Physical Drives interface shows the type; capacity; unassigned space; and self-monitoring, analysis, and reporting technology (SMART) data of each physical drive that is installed in the agent system. To see if a physical hard disk has partitions, click any disk row. If the selected disk has partitions, information about the partitions is displayed in the **Partition information** section of the **Physical Drives** interface.

The partition information is displayed as a pie chart, showing the portion of the total physical disk that is used by each partition.

The Physical Drives interface provides information about the items described in the following table.

Item	Description
Type	The type of physical drive.
Total capacity	The total amount of data that can be stored by each physical drive, measured in Kilobytes (KB stands for 1 000 bytes), Megabytes (MB stands for 1 000 000), or Gigabytes (GB stands for 1 000 000 000).
Unassigned space	The amount of space on a hard disk that is not used by logical partitions.
Status	The condition of the drive as reported by SMART.
SMART	Health alerts, generated by a status monitor, for a physical drive if potential failure exists.
Partition Information	
Color	The color of the partition being displayed on the graphic.
Partition	The drive letter that is mapped to the partition.
Capacity	The total amount of data that can be stored by each partition, measured in MB or GB.
Percent	Partition percentage of the entire drive.

FRU Numbers

Click **Information** →**Inventory** →**FRU**, to view the FRU page which displays information about the Field Replaceable Unit information installed on the agent system. The FRU information is specific to the model type of the system. The list of supported systems includes the following: xSeries 200, xSeries 220, xSeries 232, xSeries 240, xSeries 250, xSeries 330, xSeries 340, xSeries 342, XSeries 350, xSeries 370, xSeries 300, xSeries 420.

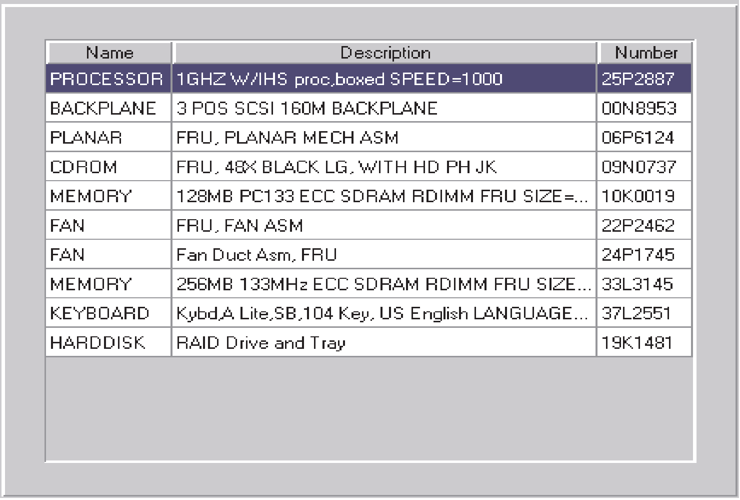

Name	Description	Number
PROCESSOR	1GHZ W/IHS proc,boxed SPEED=1000	25P2887
BACKPLANE	3 POS SCSI 160M BACKPLANE	00N8953
PLANAR	FRU, PLANAR MECH ASM	06P6124
CDROM	FRU, 48X BLACK LG, WITH HD PH JK	09N0737
MEMORY	128MB PC133 ECC SDRAM RDIMM FRU SIZE=...	10K0019
FAN	FRU, FAN ASM	22P2462
FAN	Fan Duct Asm, FRU	24P1745
MEMORY	256MB 133MHz ECC SDRAM RDIMM FRU SIZE...	33L3145
KEYBOARD	Kybd,A Lite,SB,104 Key, US English LANGUAGE...	37L2551
HARDDISK	RAID Drive and Tray	19K1481

The FRU Numbers service displays Field Replaceable Unit Information for the following system components:

- RAID Drives and Tapes
- CPUs
- Memory DIMMs
- Keyboard
- System Board
- CD-ROM Drive
- Floppy Disk Drive
- Service Processor
- Fans
- Backplanes

The FRU information for these servers can be obtained automatically through an anonymous ftp from the IBM Server Support site if the system has firewall access through the standard ftp port. In this case, a data file with the FRU information will be retrieved from the Support Site after the reboot of the Director Agent install. In addition, there is a command-line program named `getfru.exe` in the `%SystemRoot%\system32` directory that can be scripted so that the files can be retrieved from a site internal to the firewall automatically. The program's usage is :

```
getfru -s <ftp server name> -d <directory of fru files>
```

If no parameters are specfied, the default command line is:

```
getfru -s ftp.pc.ibm.com -d/pub/pccbbs/bp_server
```

Memory

The Memory service gathers information about the physical memory that is installed in the agent system and provides information about memory upgrade options that are available for the agent system. To start the Memory service, click **Information→ Inventory→Memory** in the Services pane. The following interface opens in the Display pane.

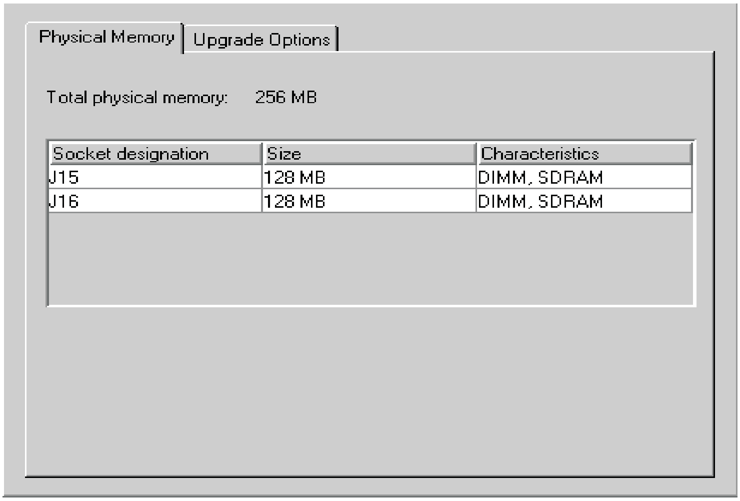

Total physical memory: 256 MB

Socket designation	Size	Characteristics
J15	128 MB	DIMM, SDRAM
J16	128 MB	DIMM, SDRAM

Note: On IBM xSeries 330 systems when the memory compression drivers are loaded, the “Memory is enabled” message will appear on the interface.

The Memory interface contains two pages:

Physical Memory
Click the **Physical Memory** page to display information about the physical memory that is installed in the agent system.

Upgrade Options
Click the **Upgrade Options** page to display information about memory upgrade options for the agent system.

The Physical Memory interface is displayed by default and provides information about the items that are described in the following table.

Item	Description
Socket designation	The type and number assignment of memory sockets that hold the memory modules.
Size	The size of the memory module currently installed in a given socket.
Characteristics	Details about the modules installed.

On servers that support memory compression technology, such as the series 330, a string indicating that compression is enabled will be displayed.

The Upgrade Options interface shows the current physical RAM that is installed in the computer and the maximum capacity of the agent system, which refers to the total RAM that can be installed in the computer. If you want to install additional memory in the agent system, select the amount of memory that you want to install to display additional information on proper memory configuration.

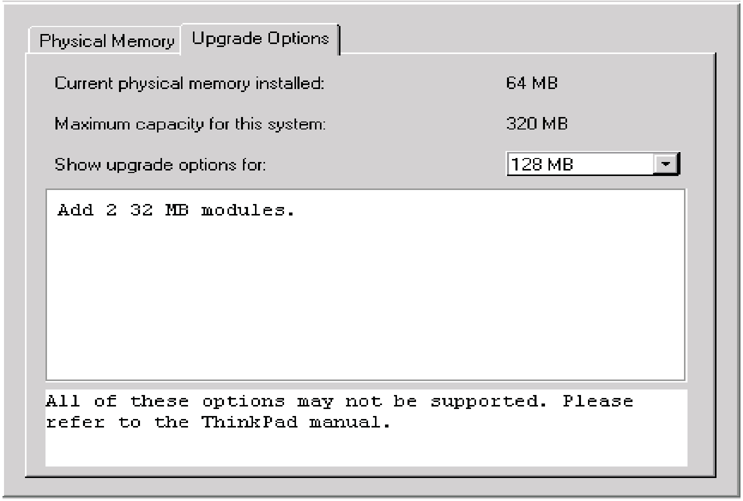

Current physical memory installed: 64 MB

Maximum capacity for this system: 320 MB

Show upgrade options for: 128 MB

```
Add 2 32 MB modules.
```

```
All of these options may not be supported. Please
refer to the ThinkPad manual.
```

The Upgrade Options interface provides information about the items that are described in the following table.

Item	Description
Current physical memory installed	The amount of total physical random access memory (RAM) that is installed on the system board.
Maximum capacity for this system	The number of the socket and the type of memory module the socket can hold. For example, DIMM-2 refers to a dual online memory module (DIMM) in socket 2, and SIMM-3 refers to a single online memory module (SIMM) in socket 3.
Show upgrade options for a total of	The size (in MB) of the memory module that is currently installed in each socket.

Multimedia

The Multimedia service gathers information about the multimedia adapter that is installed in the agent system. To start the Multimedia service, click **Information→Inventory→ Multimedia** in the Services pane. The following interface opens in the Display pane.

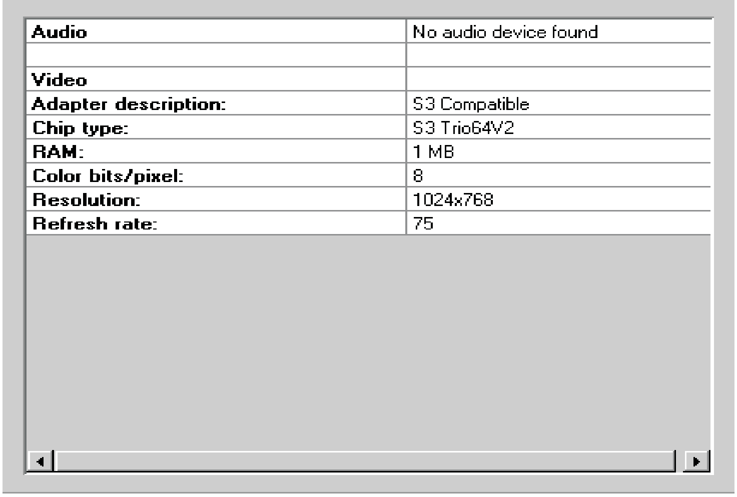

Audio	No audio device found
Video	
Adapter description:	S3 Compatible
Chip type:	S3 Trio64V2
RAM:	1 MB
Color bits/pixel:	8
Resolution:	1024x768
Refresh rate:	75

The Multimedia service has one interface that provides information about the system audio and video, described in the following table. If an audio or video adapter is not installed in the agent system or if information from the adapter is unavailable, the field that is associated with the missing data will not be shown in the interface.

Item	Description
Audio	The name of the audio adapter that is installed in the agent.
Video	The name of the video adapter that is installed in the agent.
Adapter description	The name of the video adapter that is installed on the computer.
Chip type	The type of video chip that is used by the video adapter.
DAC type	The type of digital-to-analog connector.
RAM	The amount of random access memory (RAM) available for use by the video subsystem.
Color bits/pixel	The number of color bits per picture element (pixel) that can be displayed by the video adapter.
Resolution	The picture element (pixel) resolution that is currently displayed by the video adapter (for example, 640 X 480 or 800 X 600).

Item	Description
Refresh rate	The frequency, in megahertz (Mhz), with which the monitor screen is cleared and redrawn.

Operating System

The Operating System service gathers information about the operating system that is installed and running on the agent system. To start the Operating System service, click **Information→ Inventory →Operating System** in the Services pane. The following interface opens in the Display pane.

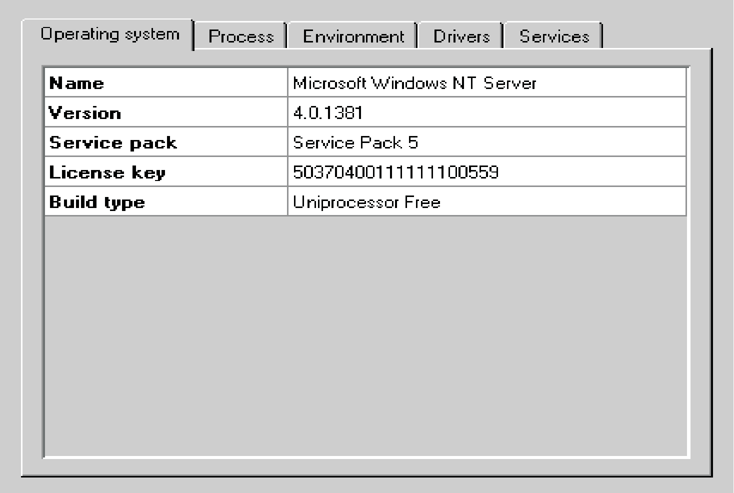

Name	Microsoft Windows NT Server
Version	4.0.1381
Service pack	Service Pack 5
License key	50370400111111100559
Build type	Uniprocessor Free

The Operating System interface contains the following interfaces:

Operating System
Click the **Operating System** page to display general information about the operating system, including name, version, and service pack level.

Process
Click the **Process** page to display information about the processes or tasks that are currently running on the agent system.

Environment
Click the **Environment** page to display information about the environment variables that are used by the agent operating system.

Drivers
Click the **Drivers** page to display information about the device drivers that are used by the agent system. (You can view the Drivers page only on agent systems running Windows NT, Windows 2000, or Windows XP.)

Services

Click the **Services** page to display information about the current state and start mode of services that are installed on the agent system. (You can view the Services page only on agent systems running Windows NT, Windows 2000, and Windows XP.

The Operating System interface is displayed by default and provides information about the items that are described in the following table.

Item	Description
Name	The name of the operating system.
Version	The version number of the operating system.
Service pack	The operating system service pack level installed on the system, if any.
License key	The license key number or code that was specified when the operating system was installed. Depending on the screen resolution, you might need to move the horizontal scroll bar to the right to view this item completely.
Build type	The operating system build type. Build type can refer to the processor configuration that the operating system is designed to run on (uniprocessor or multiprocessor), or to whether the build is a retail (or Free) version or debug (or Checked) version. Depending on the screen resolution, you might need to move the horizontal scroll bar to the right to view this item completely.

Click the **Process** page to display the Process interface.

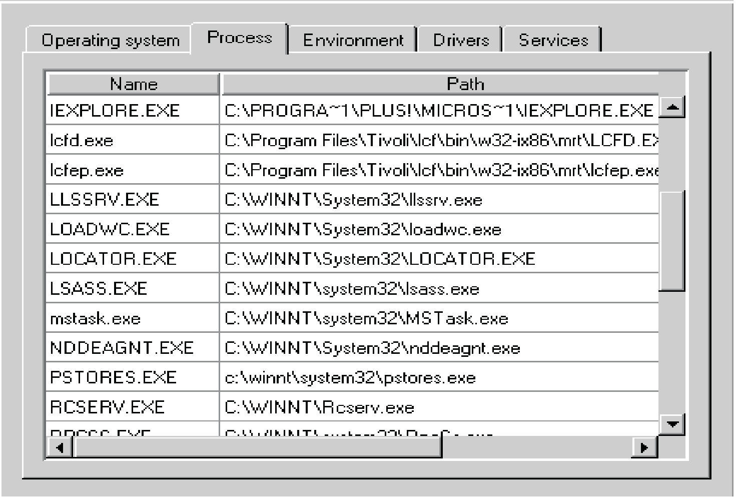

Name	Path
IEXPLORE.EXE	C:\PROGRA~1\PLUS!\MICROS~1\IEXPLORE.EXE
lcfd.exe	C:\Program Files\Tivoli\lcf\bin\w32-ix86\mrt\LCFD.E
lcfep.exe	C:\Program Files\Tivoli\lcf\bin\w32-ix86\mrt\lcfep.exe
LLSSRV.EXE	C:\WINNT\System32\llssrv.exe
LOADWC.EXE	C:\WINNT\System32\loadwc.exe
LOCATOR.EXE	C:\WINNT\System32\LOCATOR.EXE
LSASS.EXE	C:\WINNT\system32\lsass.exe
mstask.exe	C:\WINNT\system32\MSTask.exe
NDDEAGNT.EXE	C:\WINNT\System32\nddeagnt.exe
PSTORES.EXE	c:\winnt\system32\pstores.exe
RCSERV.EXE	C:\WINNT\Rcserv.exe

The Process interface provides information about the items that are described in the following table.

Item	Description
Name	The name of the executable process.
Path	The complete path to the executable file.
Kernel mode time	The amount of time that the computer processor spends in kernel mode because of this process. Depending on the screen resolution, you might need to move the horizontal scroll bar to the right to view this item completely. This item appears only for agent systems running Windows NT, Windows 2000 and Windows XP.
Process ID	The identifying number that is assigned to the process by the system according to startup sequence. Depending on the screen resolution, you might need to move the horizontal scroll bar to the right to view this item completely.

Click the **Environment** page to display the Environment interface.

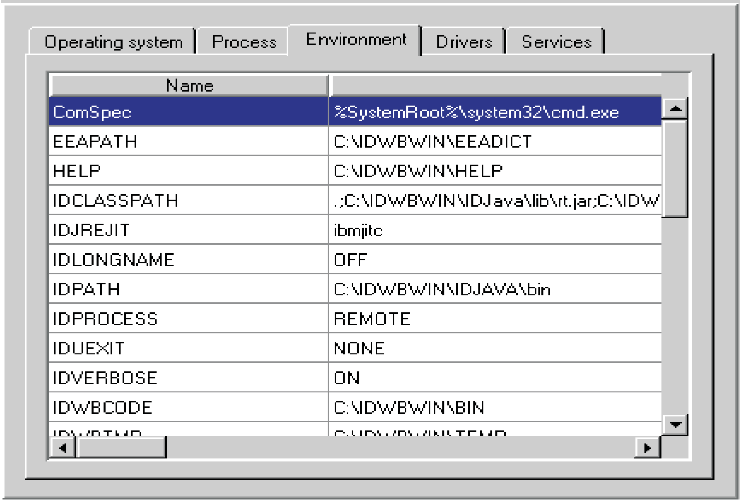

Name	
ComSpec	%SystemRoot%\system32\cmd.exe
EEAPATH	C:\IDWBWIN\EEADICT
HELP	C:\IDWBWIN\HELP
IDCLASSPATH	.;C:\IDWBWIN\IDJava\lib\rt.jar;C:\IDW
IDJREJIT	ibmjitc
IDLONGNAME	OFF
IDPATH	C:\IDWBWIN\IDJAVA\bin
IDPROCESS	REMOTE
IDUEXIT	NONE
IDVERBOSE	ON
IDWBCODE	C:\IDWBWIN\BIN
IDWBTMP	C:\IDWBWIN\TEMP

The Environment interface provides information about the item that is described in the following table.

Item	Description
Name	The name of the environment variables that are used by the agent system.
Value	The settings information for each environment.

Item	Description
Detailed Value	An in-depth display of the name which includes the path.

Click the **Drivers** page to display the Drivers interface. You must have administrator privileges to update the Driver Start mode and the Start/Stop options.

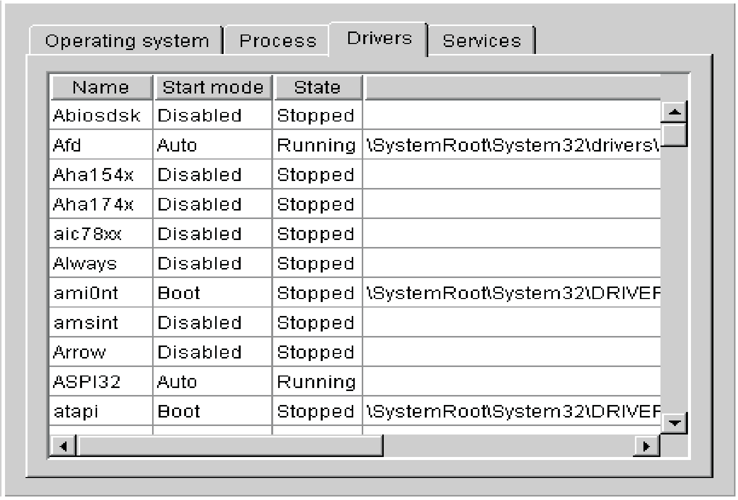

Name	Start mode	State	
Abiosdsk	Disabled	Stopped	
Afd	Auto	Running	\SystemRoot\System32\drivers\
Aha154x	Disabled	Stopped	
Aha174x	Disabled	Stopped	
aic78xx	Disabled	Stopped	
Always	Disabled	Stopped	
ami0nt	Boot	Stopped	\SystemRoot\System32\DRIVEF
amsint	Disabled	Stopped	
Arrow	Disabled	Stopped	
ASPI32	Auto	Running	
atapi	Boot	Stopped	\SystemRoot\System32\DRIVEF

The Drivers interface provides information about the items that are described in the following table.

Item	Description
Name	The name of each device driver in the operating system directory.
Start mode	The start mode that is assigned to each device driver. Depending on which mode is selected, a device driver is incorporated or not incorporated into the operating environment. *Disabled* means that the device driver is not added to the operating environment. *Auto* means that the device driver is automatically started when the operating system is started. *Boot* means that the device driver is initialized during the operating system startup (boot) sequence. *Manual* means the user must start the device driver. *System* means the system which occurs when the operating system starts the device driver.
Start	To start a device driver, highlight the device driver and click the **start** button.

Item	Description
Stop	To stop a driver, highlight the driver and click the **stop** button.
State	The current run state of each device driver (Running or Stopped). This item applies only to agent systems running Windows NT or Windows 2000. State values are not displayed on agent systems running Windows 98 or Windows ME.
Command line	The complete path to the device driver, such as C:\System Root\System32\adapti.sys**.** To view the complete command line, move the horizontal scroll bar to the right.

Click the **Services** page to display the Services interface.

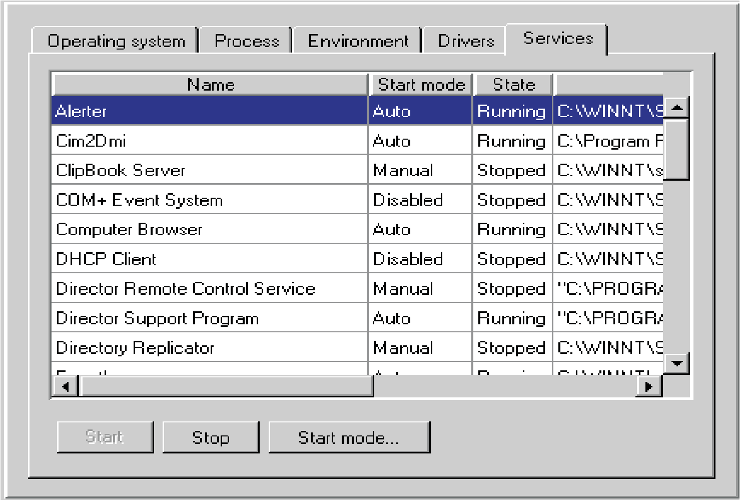

Name	Start mode	State	
Alerter	Auto	Running	C:\WINNT\S
Cim2Dmi	Auto	Running	C:\Program F
ClipBook Server	Manual	Stopped	C:\WINNT\s
COM+ Event System	Disabled	Stopped	C:\WINNT\S
Computer Browser	Auto	Running	C:\WINNT\S
DHCP Client	Disabled	Stopped	C:\WINNT\S
Director Remote Control Service	Manual	Stopped	"C:\PROGRA
Director Support Program	Auto	Running	"C:\PROGRA
Directory Replicator	Manual	Stopped	C:\WINNT\S

The Services interface provides information about the items that are described in the following table.

Item	Description
Name	The name of the service (for example, EventLog or Remote Control Service).
Start mode	The start mode of the service. For each service, the start mode can be Auto for automatic, Manual for manual start, or Disabled when the service is turned off or is unavailable.
State	The current run state of each service (Running or Stopped).
Start	To start a service, highlight the name and click **start**.
Stop	To stop a service, highlight the name and click **stop**.

Item	Description
Start Mode	Click to select service start mode. The modes are Auto, Manual, and Disabled.
Command line	The complete path to the device driver, such as C:\System Root\System32\adapti.sys. To view the complete Command line, move the horizontal scroll bar to the right.

Ports

The Ports service gathers information about the input and the output ports and connectors on the agent system. To start the Ports service, click **Information→ Inventory →Ports** in the Services pane. The following interface opens in the Display pane.

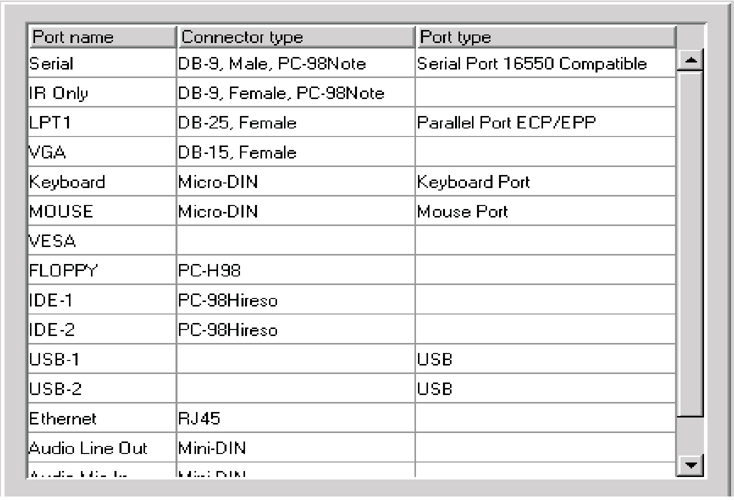

Port name	Connector type	Port type
Serial	DB-9, Male, PC-98Note	Serial Port 16550 Compatible
IR Only	DB-9, Female, PC-98Note	
LPT1	DB-25, Female	Parallel Port ECP/EPP
VGA	DB-15, Female	
Keyboard	Micro-DIN	Keyboard Port
MOUSE	Micro-DIN	Mouse Port
VESA		
FLOPPY	PC-H98	
IDE-1	PC-98Hireso	
IDE-2	PC-98Hireso	
USB-1		USB
USB-2		USB
Ethernet	RJ45	
Audio Line Out	Mini-DIN	
Audio Mic In	Mini-DIN	

The Ports service provides information about the items that are described in the following table.

Item	Description
Port name	The name of the input or output port (for example, LPT1, Keyboard, or Ethernet).
Connector type	The type of connector for each port (for example, DB-9 or DB-25 Female).
Port type	Type of port (for example, serial, parallel, or Universal Serial Bus).

Monitor Services

The Monitor Services uses system monitoring hardware and software that is included with the IBM Director Agent to gather data about the current operational state of the agent system, such as temperature, battery time remaining, and contents of the Windows NT, Windows 2000, or Windows XP Event Log on the agent system. The three Monitor Services are as follows:

- Event Viewer
- Battery
- System Health

Windows NT Event Log

Applications, device drives, operating systems, and IBM Director Agent record hardware events and software events in the Windows NT, Windows 2000, or Windows XP Event Log. The IBM Director Agent Event Viewer displays these events.

Event Viewer

The Event Viewer service shows the contents of the Windows NT, Windows 2000 or Windows XP Event Log. To start the Event Viewer service, click **Information→Monitors→Event Viewer** in the Services pane. The following interface opens in the Display pane.

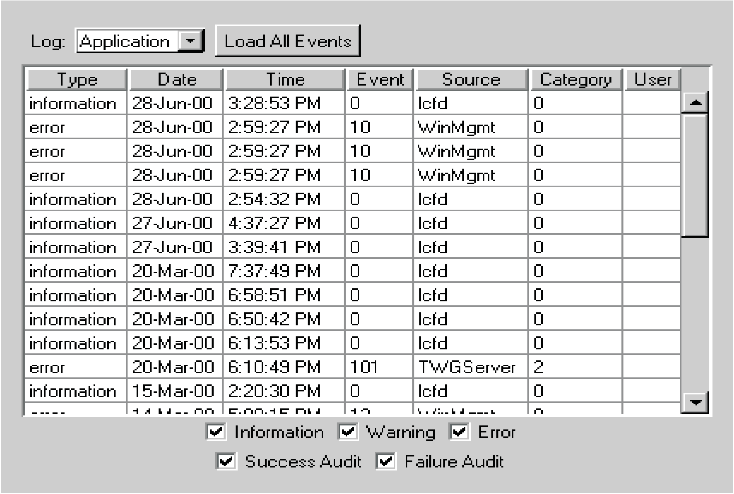

Log: Application

Type	Date	Time	Event	Source	Category	User
information	28-Jun-00	3:28:53 PM	0	lcfd	0	
error	28-Jun-00	2:59:27 PM	10	WinMgmt	0	
error	28-Jun-00	2:59:27 PM	10	WinMgmt	0	
error	28-Jun-00	2:59:27 PM	10	WinMgmt	0	
information	28-Jun-00	2:54:32 PM	0	lcfd	0	
information	27-Jun-00	4:37:27 PM	0	lcfd	0	
information	27-Jun-00	3:39:41 PM	0	lcfd	0	
information	20-Mar-00	7:37:49 PM	0	lcfd	0	
information	20-Mar-00	6:58:51 PM	0	lcfd	0	
information	20-Mar-00	6:50:42 PM	0	lcfd	0	
information	20-Mar-00	6:13:53 PM	0	lcfd	0	
error	20-Mar-00	6:10:49 PM	101	TWGServer	2	
information	15-Mar-00	2:20:30 PM	0	lcfd	0	
error	14-Mar-00	5:09:15 PM	13	WinMgmt	0	

☑ Information ☑ Warning ☑ Error

☑ Success Audit ☑ Failure Audit

The Event Viewer interface provides information about the items that are described in the following table.

Item	Description
Type	The log category (information, warning, error, success audit, or failure audit).
Date	The date when the event took place, in mm/dd/yy format.
Time	The time when the event occurred, in hh:mm:ss format for AM or PM.
Event	The identification number that is automatically given to an event, with related events getting the same number. For example, Service Control Manager has 7001 and 7002 event numbers, regardless of the time of the event.
Source	The program, application, system, or security problem that led to the event (for example, WinMgmt, DCOM, Simple Network Management Protocol (SNMP), AOLAgent, or IBM Director Agent).
Category	A number that identifies the category that the event falls into. This is used to organize the events. To view this field, move the horizontal scroll bar to the right.
User	The ID of the user. To view this field, move the horizontal scroll bar to the right.

The Event Log can contain a large number of entries. With the Event Viewer service, you can filter the contents of the Event Log before viewing the entries. Before Event Viewer loads the contents of the Event Log, you must select a log category. These broad categories help limit the number of Event Log entries that will be loaded into the Event Viewer. From the **Log** menu, select an Event Log category that corresponds to the Event Log entries you want to view, or click **Load All Events** to display all log entries.

Note: The Event Log can contain thousands of individual entries. Clicking **Load All Events** can result in significant delays while the contents are loaded into the Event Viewer.

The available selections are as follows:

Application
Displays the 30 most recent log entries that result from software issues or application issues, faults, and problems.

System
Displays the 30 most recent log entries that result from system issues or hardware issues, faults, and problems.

Security
Displays the 30 most recent log entries that result from security

problems, such as invalid user ID or password entries and other attempted security violations.

Use the check boxes at the bottom of the **Event Viewer** window to filter the contents of the Event Viewer by event type. The event type provides a general description of the severity of the event. The following event-type check boxes are available:

Information
Displays rows of informational entries that are related to the Application, System, or Security Event Log category that you selected.

Warning
Displays rows of warning entries that indicate a severe problem to resolve for an application, system, or security problem.

Error
Displays logs that result from security issues, such as password or user ID failures or other access problems, or attempted security violations. It also displays log errors for application and system.

Success Audit
Displays event information on successful events.

Failure Audit
Displays event information on unsuccessful events.

Only entries that correspond to a selected check box will be displayed in the Event Viewer. For example, if you want to view only entries that are the results of system errors, select the **Error** check box and leave the other selections cleared. If you select an event-type check box and no information is displayed, it means that there are no Event Log entries that correspond to the selected event type.

You can use Event Viewer to display additional information about any entry that appears in the window. To display additional information about any entry, click the entry to highlight it, and then double-click the entry. A window opens, containing additional information about the event.

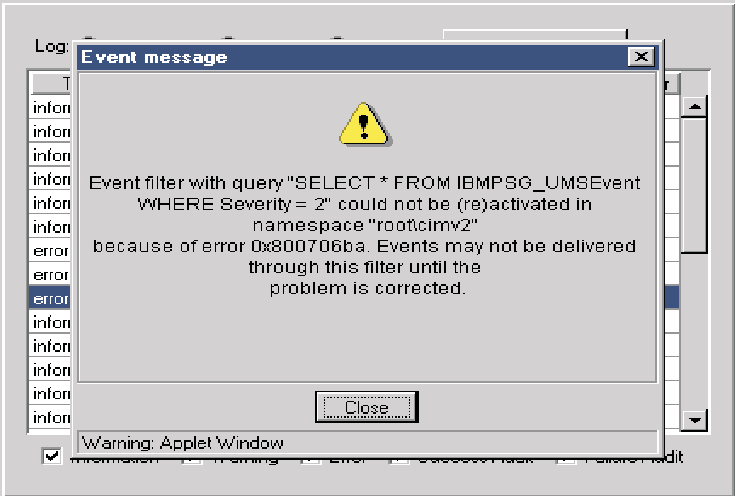

Event filter with query "SELECT * FROM IBMPSG_UMSEvent WHERE Severity = 2" could not be (re)activated in namespace "root\cimv2"
because of error 0x800706ba. Events may not be delivered through this filter until the problem is corrected.

Battery

Note: This option is not available on systems running Windows NT.

The Battery service gathers and displays information about ThinkPad battery power source. To start the Battery service, click **Information→Monitors → Battery** in the Services pane. The following interface opens in the Display pane.

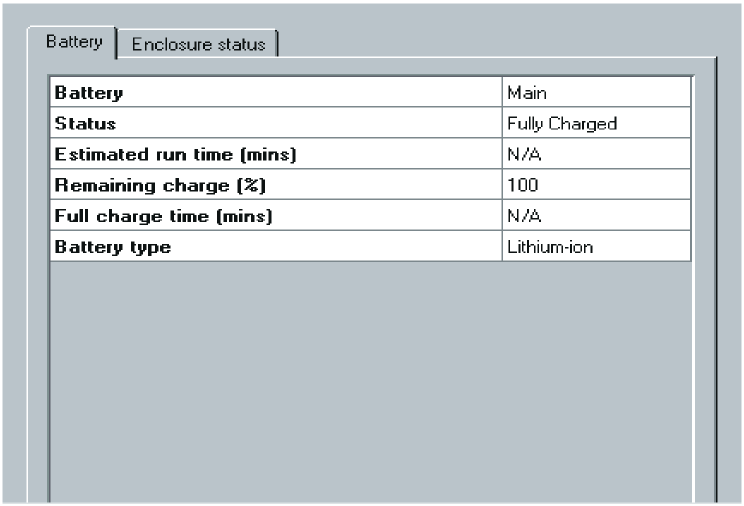

Battery	Main
Status	Fully Charged
Estimated run time (mins)	N/A
Remaining charge (%)	100
Full charge time (mins)	N/A
Battery type	Lithium-ion

The Battery interface provides information about the items that are described in the following table.

Item	Description
Battery	The battery that is being used by the ThinkPad computer (Main or Backup).
Name	The name of this battery.
Manufacturer	The manufacturer of the battery.
Status	The charge status of the battery (Fully Charged, Partial, High, Low, Critical, Charging, Charging High, Charging Low, Charging Critical, Unknown).
Remaining charge (%)	The approximate percentage of battery charge remaining, running from 100% to 0%. If the ThinkPad computer is plugged into an electrical outlet using an adapter, the Remaining Charge will continue to show the percentage of battery life that remained at the time the system was plugged in.
Battery Type	The type of battery in use. For instance, ThinkPad systems use lithium-ion batteries.
Full Charge Capacity	The number of units of run time left on the battery. However, if the ThinkPad system is plugged into an AC outlet through an adapter, the Estimated Run Time will show as N/A.
Design Capacity	This displays the number of minutes of run time specified in the design specification.
Low Battery Alert	The number of remaining units of battery life after a low life alert is issued.
Warn Battery Alert	The number of minutes of run time when an alert is issued.
Units	The units of measurement used in this chart, such as mWH (milliwatthour).

System Health

IBM Director Agent automatically monitors the agent systems for changes in a variety of system-environment factors, including temperature and voltage. Each monitored value has a System Health normal range. If the monitored value stays within normal range, the assumption is that the System Health is normal. However, if any of these monitored values falls outside of acceptable System Health parameters, IBM Director Agent can automatically generate five forms of output to alert the system administrator of this state change. The following alert output can be generated by IBM Director Agent:

- System Health window in IBM Director Agent
- Alert messages
- Alert messages that are sent as simple network management protocol (SNMP) traps
- Alert messages that are sent as System Management Server (SMS) status messages
- Common Information Model (CIM) events
- Alert messages that are sent as TEC events
- Alert messages that are sent as Director server events

You can use the System Health service to check the status of all health monitors that are supported by the agent system. To start the System Health service, click **Information→ Monitors→System Health** in the Services pane. The following interface opens in the Display pane.

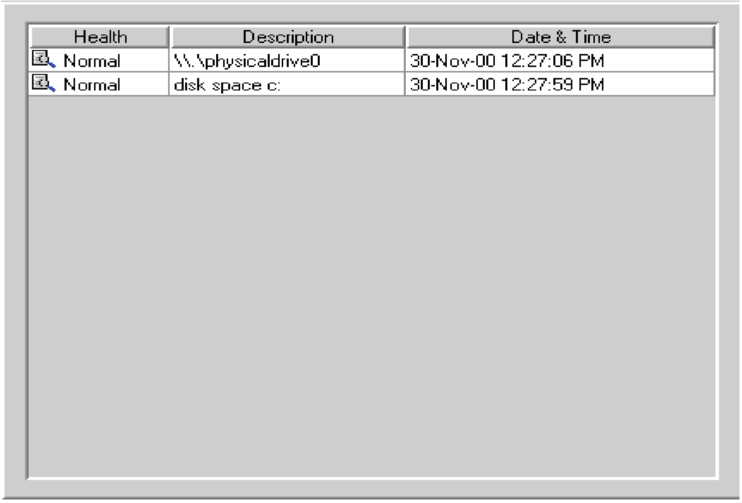

Health	Description	Date & Time
Normal	\\.\physicaldrive0	30-Nov-00 12:27:06 PM
Normal	disk space c:	30-Nov-00 12:27:59 PM

The System Health interface provides information about the items that are described in the following table.

Item	Description
Health	The current state of the monitored device (Normal, Warning, or Critical).
Description	A description of the monitored device.

Item	Description
Date & Time	The date and time stamp applied to the health event. The format is dd/mm/yy hh:mm:ss (AM or PM).

Health reports are gathered from a variety of system devices. One of these devices is the LM Sensor, which performs environmental monitoring. The health reports that are available on an agent system are dependent on the availability of components that contribute to health reports. The following are some examples of potential System Health event messages and the circumstances that cause them:

Chassis Intrusion

If the system chassis has been opened, a Critical System Health event is generated, regardless of the reason.

Fan Failure

If the system cooling fan fails, a Critical System Health event is generated. This might be the only prediction of a temperature-related event.

Memory PFA

Available on select servers. Indicates a predictive failure from a memory DIMM.

Processor PFA

Available on select servers. Indicates a predictive failure from a CPU.

LAN Leash

LAN Leash detects if an agent system is disconnected from the LAN, even when the computer is off. A Critical System Health event is generated if a agent system is disconnected from the LAN.

Low Disk Space

If free disk space is low, a Warning or Critical System Health event is generated.

Processor Removed

If the microprocessor is removed from the agent system, a Warning System Health event is generated.

Temperature Out of Specification

If the microprocessor temperature is out of the specified range, a Warning System Health event is generated.

Voltage Out of Specification

If there is a dramatic change in the voltage of the power that is supplied to any part of the agent system, a Warning or Critical System Health event is generated.

Predicted Failure Alert (PFA)
SMART-drive enabled systems generate events if operational thresholds on the hard disk drive are exceeded.

Redundant NIC
If the redundant network interface card is configured for automatic failover, a Warning or Critical System Health event is generated. The Redundant NIC feature is supported only on Windows NT and Windows 2000.

Management Processor Assistant

On systems with a Management Processor installed, there are ten informational parameters associated with Management Processor Assistant task:

- Event Log
- Fan Speeds
- Power/Restart Activity
- Server Time Outs
- Temperature
- Voltages
- VPD Components
- VPD Machine
- VPD Management Product
- VPD POST/BIOS

Event Log

The Event Log interface displays entries currently stored in the system management event log which are associated with the Management Processor Assistant adapter. To access the Event Log interface, click **Information→Mgmt Processor Assistant→Event Log.**

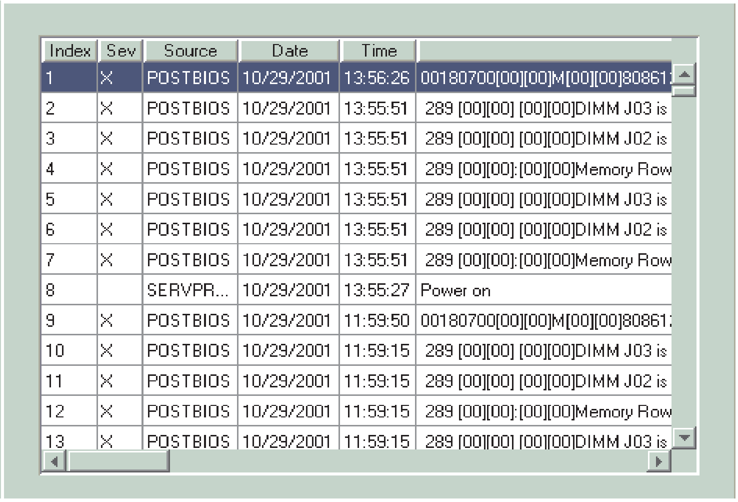

Index	Sev	Source	Date	Time	
1	X	POSTBIOS	10/29/2001	13:56:26	00180700[00][00]M[00][00]80861
2	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00] [00][00]DIMM J03 is
3	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00] [00][00]DIMM J02 is
4	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00]:[00][00]Memory Row
5	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00] [00][00]DIMM J03 is
6	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00] [00][00]DIMM J02 is
7	X	POSTBIOS	10/29/2001	13:55:51	289 [00][00]:[00][00]Memory Row
8		SERVPR...	10/29/2001	13:55:27	Power on
9	X	POSTBIOS	10/29/2001	11:59:50	00180700[00][00]M[00][00]80861
10	X	POSTBIOS	10/29/2001	11:59:15	289 [00][00] [00][00]DIMM J03 is
11	X	POSTBIOS	10/29/2001	11:59:15	289 [00][00] [00][00]DIMM J02 is
12	X	POSTBIOS	10/29/2001	11:59:15	289 [00][00]:[00][00]Memory Row
13	X	POSTBIOS	10/29/2001	11:59:15	289 [00][00] [00][00]DIMM J03 is

The Event Log interface provides information about the items that are described in the following table

Item	Description
Index	Entries are displayed in chronological order.
Sev	Event severities are informational, indicated by a blank; warning, indicated by an !; and Error, indicated by an x.
Source	Displays the specific firmware.
Date	Date of the event.
Time	Time of the event.
Text	A brief description of the event.
Note: All events are informational unless noted as an error or warning event.	

Fan Speeds

The Fan Speed interface displays the current speed of system fans as a percentage of the maximum.To access the Fan Speeds interface, click **Information→ Mgmt Processor Assistant→ Fan Speeds.**

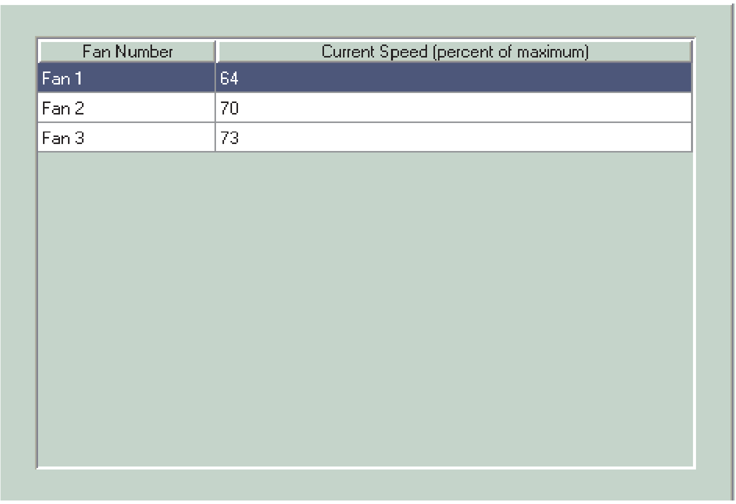

Fan Number	Current Speed (percent of maximum)
Fan 1	64
Fan 2	70
Fan 3	73

The Fan Speeds interface provides information about the items that are described in the following table

Item	Description
Fan Number	The number of fans.
Current Speed (% of maximum)	The current speed of the fan expressed as a percentage of the maximum speed of the fan.

Power/Restart Activity

The Power/Restart Activity interface displays information about any access attempt. To access Power/Restart Activity interface, click **Information→Mgmt Processor Assistant→Power/Restart Activit**y. The following interface opens.

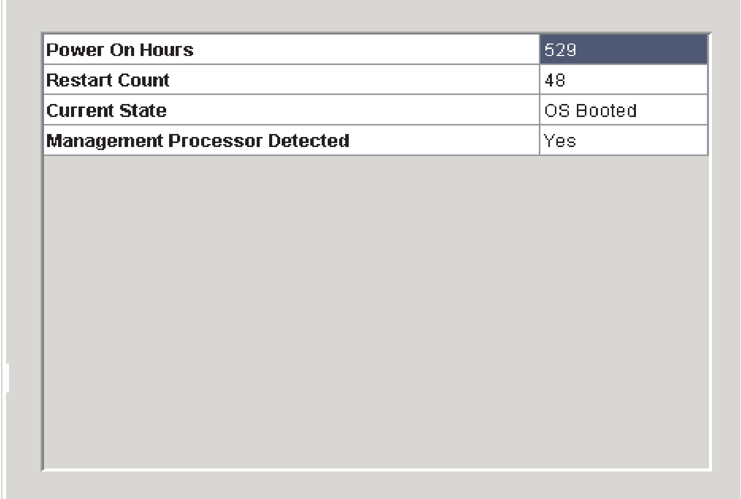

Power On Hours	529
Restart Count	48
Current State	OS Booted
Management Processor Detected	Yes

The Power/Restart Activity interface provides information about the items that are described in the following table

Item	Description
Power on Hours	The total number of hours this server has been powered on.
Restart Count	The number of times the system has restarted. This counter is reset to 0 each time the Management Processor Assistant subsystem is cleared to factory defaults.
Current State	This shows the state of the system when this web page was generated. Possible states include: System power off/State unknown In Post Stopped in POST (Error detected) Booted Flash or System partition Booting OS or in OS (could be in the Operating System if the Operating System or application does not report the new system state.) In OS CPUs held in reset System power on/Before POST
Management Processor Detected	Indicated Yes or No to the Management Processor Assistant being detected.

Server Timeouts

The Server Timeouts interface displays the current temperature readings for various hardware components. All temperature readings are in degrees Celsius. Some of the links are shown as hyperlinks that display the present temperature threshold values for the corresponding component. To access Server Timeouts interface, click **Information→Mgmt Processor Assistant→Server Timeouts**. The following interface opens.

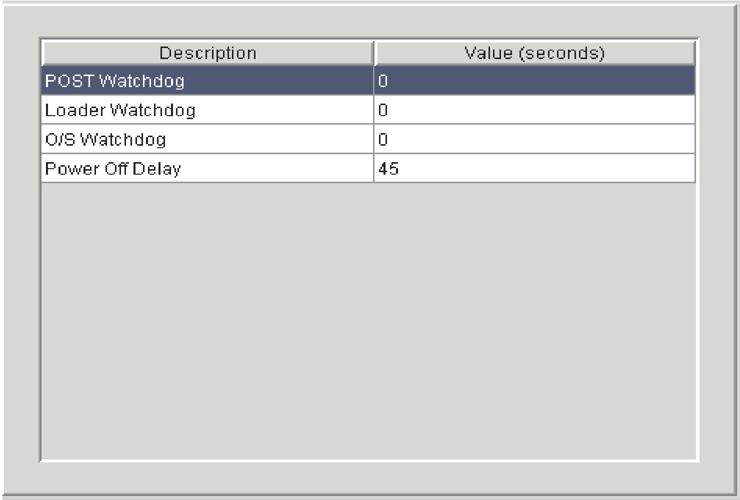

Description	Value (seconds)
POST Watchdog	0
Loader Watchdog	0
O/S Watchdog	0
Power Off Delay	45

The Server Timeouts interface provides information about the items that are described in the following table

Item	Description
POST Watchdog	This field specifies the number of minutes that the Management Processor Assistant subsystem will wait for this system to complete Power-on Self Test (POST). If this system fails to complete POST within this time, the Management Processor Assistant subsystem generates a POST timeout alert and automatically restarts the system one time. Once the system is restarted, the POST watchdog is automatically disabled until the operating system is shutdown and the server is power cycled.
Loader Watchdog	Use this field to specify the number of minutes that the Management Processor Assistant subsystem will wait between the completion of POST and the end of loading of the operating system. If this interval is exceeded, the Management Processor Assistant subsystem will generate a Loader Timeout alert.

Item	Description
O/S Watchdog	This file specifies how often, in minutes, the Management Processor Assistant subsystem will check to confirm that the operating system is running properly. If the operating system fails to respond within 6 seconds to one of these checks, the Management Processor Assistant subsystem will generate an O/S Timeout alert and automatically restart the system one time. After the system is restarted, the O/S watchdog is automatically disabled until the operating system is shutdown and the server is power cycled.
Power Off Delay	Use this field to specify the number of minutes that the Management Processor Assistant subsystem will wait for the operating system to shutdown before powering off the system.

Temperature

The Temperature interface displays the current temperature readings for various hardware components. All temperature readings are in degrees Celsius. To access Temperature interface, click **Information→Mgmt Processor Assistant→ Temperature**. The following interface opens.

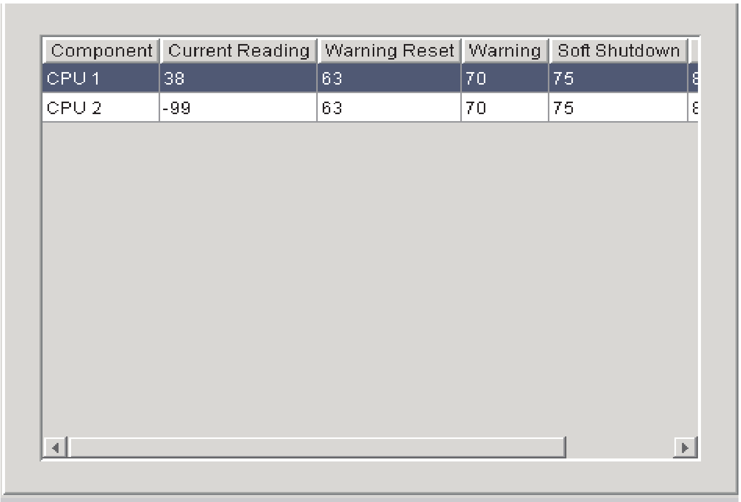

Component	Current Reading	Warning Reset	Warning	Soft Shutdown
CPU 1	38	63	70	75
CPU 2	-99	63	70	75

The Temperature interface provides information about the items that are described in the following table

Item	Description
Component	The part of the system being monitored.
Current Reading	The current temperature of the component.
Warning Reset	If the temperature was above the Warning threshold and then dropped below this value, any active temperature events are cleared.
Warning	If the temperature reaches this value, a warning event is generated.
Soft Shutdown	If the temperature reaches this value, a critical event is generated and the server is powered off after the operating system is shut down.
Hard Shutdown	If the temperature reaches this value, a critical event is generated and the server is powered off immediately.

Voltages

The voltage interface displays the current voltage readings for the system board and VRMs. Each voltage threshold is defined as a Low, High value pair. The voltage thresholds are defined in the following table. To access Voltages interface, click **Information→Mgmt Processor Assistant→Voltages**. The following interface opens.

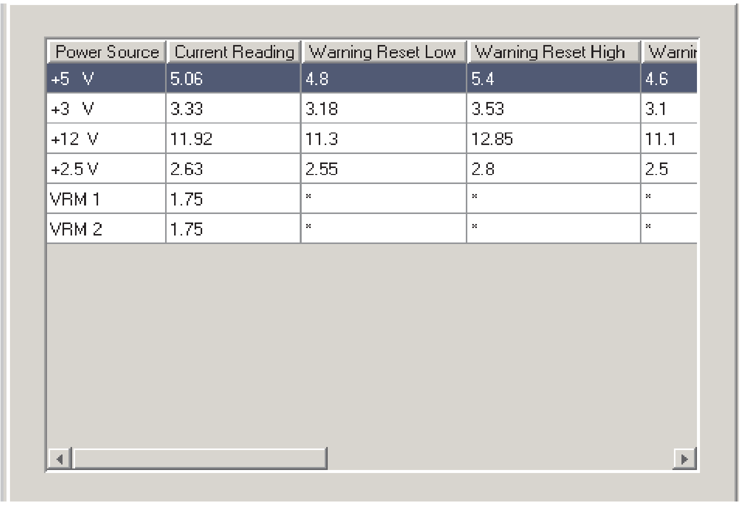

Power Source	Current Reading	Warning Reset Low	Warning Reset High	Warnir
+5 V	5.06	4.8	5.4	4.6
+3 V	3.33	3.18	3.53	3.1
+12 V	11.92	11.3	12.85	11.1
+2.5 V	2.63	2.55	2.8	2.5
VRM 1	1.75	*	*	*
VRM 2	1.75	*	*	*

The Voltages interface provides information about the items that are described in the following table

Item	Description
Power Source	The power source being monitored.
Current Reading	The current voltage of the source being monitored.
Warning Reset Low	If the voltage reading was outside the warning threshold range and then changed a value within this range.
Warning Reset High	If the voltage reading was outside the warning threshold range, and then changed to a value within this range, any active voltage events are cleared.
Warning High	If the voltage rises above the High value, a warning event is generated.
Warning Low	If the voltage drops below the low value, a warning event is generated.
Soft Shutdown Low	If the voltage drops below this value, a critical event is generated and the server is powered off after the operating system is shut down.
Soft Shutdown High	If the voltage rises above this value, a critical event is generated and the server is powered off after the operating system is shut down.
Hard Shutdown Low	If the voltage drops below this value, a critical event is generated and the server is powered off immediately.
Hard Shutdown High	If the voltage rises above this value, a critical event is generated and the server is powered off immediately.
Warning Reset	If the voltage reading was outside the warning threshold range, and then changed to a value within this range, any active voltage events are cleared.

Vital Product Data (VPD) Components

The VPD Components interface displays information about the Management Processor Assistant card components.

Item	Description
Description	Description of the component.

Item	Description
Build ID	The build ID for the firmware.
FRU No.	The field replaceable unit part number of the component.
Model	The model number of the component.
Manufacturer	The manufacturer of the component.
Other	Other information regarding the component.

Vital Product Data (VPD) Machine

The VPD Machine interface displays general information about the Management Processor Assistant card.

Item	Description
Machine Type	The four-digit machine type of machine type of the system.
Machine Model	The model of the system.
Serial Number	The system serial number.
UUID	The Universal Identification number of the system.

Vital Product Data (VPD) Management Product

The VPD Management Product interface displays information about the Firmware type of Management Processor Assistant card.

Item	Description
Description	A description of the firmware.
Build ID	The build ID for the firmware.
Revision Number	The revision number of the firmware.
File Name	The file name for firmware
Release Note	The release date for the firmware.

Vital Product Data (VPD)POST/BIOS

The VPD Product Data interface displays information about the Firmware BIOS.

Item	Description
Version	The version of BIOS.
Build Level	The build level of this BIOS.
Build Date	The build date of this BIOS.

Viewing Tasks services

The services that are available on the **Tasks** page help the system administrator manage the agent systems. Users with less than system-administrator authority can view the available screens, but only system administrators can change or update system configurations and use the available tools.

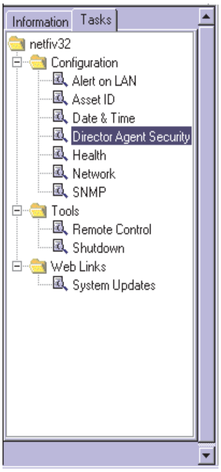

IBM Director Agent displays only the tasks that are associated with the components that are installed on an agent system. For example, if the Web Based Remote Control feature is not installed on an agent system, the task Remote Control (under **Tools**) is not displayed for that system. Requirements and optional installations are noted under each task heading. Certain security levels are required so that users can view or edit selected features in the IBM Director Agent program. See “IBM Director Agent Security” on page 484 for additional information.

There are three main categories in Tasks services:

- Configuration
- Tools
- Web Links

The sections that follow describe each of the services that are available on the **Tasks** page.

Configuration

There are seven setup options that are associated with the Configuration task.

- Alert on LAN
- Asset ID
- Date and Time
- Director Agent Security (Windows 98 and Windows ME only)
- Health
- Network
- System Accounts

(Windows NT, Windows 200 and Windows XP only)

Alert on LAN

A user with administrative security status can use the Alert on LAN task to set the options that are related to network system alerts.

To start the Alert on LAN service, click **Tasks→Configuration→Alert on LAN**. The following interface opens in the Display pane.

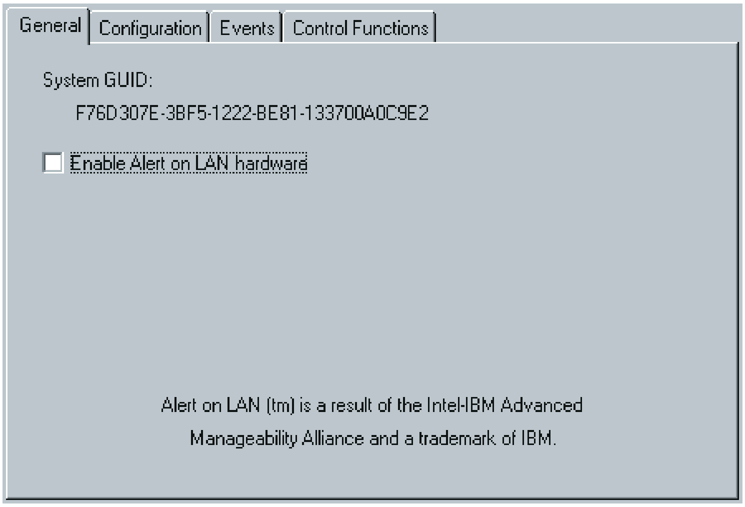

System GUID:
F76D307E-3BF5-1222-BE81-133700A0C9E2

□ Enable Alert on LAN hardware

Alert on LAN (tm) is a result of the Intel-IBM Advanced Manageability Alliance and a trademark of IBM.

The Alert on LAN interface provides information about the items that are described in the following table. Each section in the table denotes a tabbed window within the Alert on LAN task.

Item	Description
General page	
System GUID	A Globally Unique ID (GUID) is assigned to each system board for system-management purposes. The GUID is stored in the BIOS on the system board.
Enable Alert on LAN hardware	This option determines whether the system alerts are on or off. Select the check box to enable system alerts.
Configuration page	
Proxy server (IP address port)	The Internet protocol address for the server you use to communicate with the agent systems. The IP address is assigned by the system administrator. (Default port is 5500.)
Heartbeat timer period	The Alert on LAN proxy computer verifies that the agent system is running. This is the number of seconds between system checks. The default value is 32. The enabled heartbeat timer period values range from 43 to 5461 seconds and can be set in intervals of 43 seconds.
Watchdog Timer Period	If the watchdog timer indicates that a agent system has stopped, the watchdog timer automatically sends a message to the proxy computer. This is the period between polls for the watchdog timer (measured in seconds). The default value is 43. The watchdog timer period values range from 86 to 5461 seconds and can be set in intervals of 86 seconds.
Transmission attempts	The number of retries for transmission after the agent system stops. The default value is 3.
Event Polling Period	The polling period for software problems. The default value is 30.
Events page	
Cover Tamper	If the cover of the managed system has been opened or removed, an event message is generated.
LAN Leash Tamper	LAN Leash detects if an agent system is disconnected from the LAN, even when the computer is off. If an agent system is disconnected from the LAN, an event message is generated.

Item	Description
Temperature Out of Specification	If the microprocessor temperature is out of the specified range, an event message is generated.
Watchdog	If the operating system of the managed system is not functioning or is in a suspended state, an event message is generated.
Voltage Out Specification	If there is a dramatic change in the voltage of the power supplied to any part of the agent system an event message is generated.
Processor 0	Click to enable the notification of a watchdog.
Processor 1	Click to enable the notification of a processor missing sensor report.
Auto-clear events	If this option is enabled, the agent system sends an alert each time the condition is present (multiple alerts). If this option is disabled, the system sends an alert for a condition only once (no reminder alerts).
Clear All Events	Select this option and click **Apply**, to clear the events log.
Control Functions page	
Power Down	Click to power down the system. You will receive this message as a system state report.
Power Up	Click to restart the system. You will receive this message as a system state report.
Reboot	Click to restart your system. You will receive this message as a system state report.
Presence Ping	Returns the message that the system is not on but is still connected to the network.

If you make changes to any of the Alert on LAN default user options, click **Apply** to save the changes and return to the IBM Director Agent main window.

Asset ID

The Asset ID service contains the hardware information for the agent system. To start the Asset ID service, click **Tasks→Configuration→Asset ID**. The Asset ID interface contains the following interfaces:

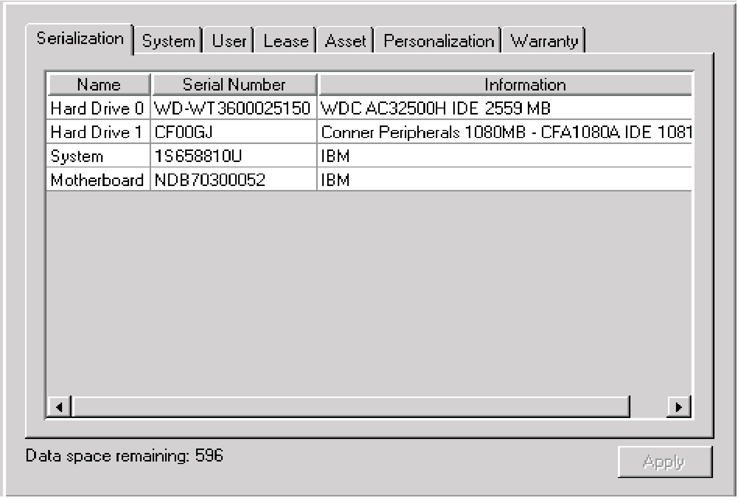

Name	Serial Number	Information
Hard Drive 0	WD-WT3600025150	WDC AC32500H IDE 2559 MB
Hard Drive 1	CF00GJ	Conner Peripherals 1080MB - CFA1080A IDE 1081
System	1S658810U	IBM
Motherboard	NDB70300052	IBM

Data space remaining: 596

Serialization
Click the **Serialization** page to display serial numbers for the agent system hardware.

System
Click the **System** page to display the current agent system characteristics: system name, MAC address, user login name, operating system, GUID address, IBM LAN Client Control Manager (LCCM) Profile.

User Click the **User** page to display the user profile: user name, telephone number, work location, department, and professional position.

Lease Click the **Lease** page to display the information on the lease agreement for the agent system hardware.

Asset Click the **Asset** page to display the inventory factors that are related to the agent system.

Personalization
Click the **Personalization** page to display the free-form window where you can add information on your systems, users, or computers.

Warranty
Click the **Warranty** page to display the information on the warranty agreement for the agent system hardware.

When you click **Asset ID** from the IBM Director Agent task list, the Serialization window opens. To access any of the other Asset ID windows, click the appropriate page.

Serialization
Click the **Serialization** page to display the Serialization interface. The

Serialization interface displays the serial numbers for the various components in the agent system.

The Serialization interface provides information about the items that are described in the following table.

Item	Description
Name	The hardware component name.
Serial Number	The serial number for the hardware component.
Information	Descriptive information for the hardware component.

System

Click the **System** page to display the System interface. The System interface displays information about the agent system.

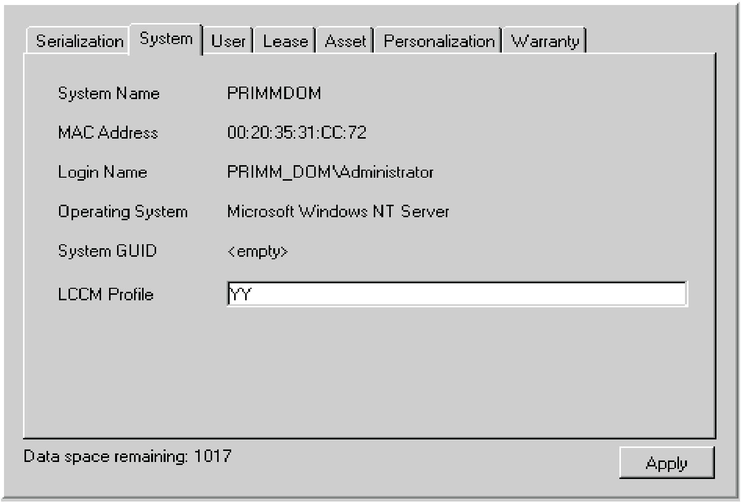

System Name PRIMMDOM

MAC Address 00:20:35:31:CC:72

Login Name PRIMM_DOM\Administrator

Operating System Microsoft Windows NT Server

System GUID <empty>

LCCM Profile YY

Data space remaining: 1017

The System interface provides information about the items that are described in the following table.

Item	Description
System Name	The NetBEUI name of the agent system (the computer name as it appears under **Network Properties**). NetBEUI is NetBIOS extended user interface, and NetBIOS is network basic input/output system.

Item	Description
MAC Address	The unique hexadecimal character string that identifies the network adapter in the agent system.
Login Name	The user ID that the system administrator assigned at installation.
Operating System	The operating system (for the management server or for the computer where IBM Director Agent resides).
System GUID	The agent system Global Unique Identifier (GUID). This is your BIOS unique ID number.
LCCM Profile	The profile name of the IBM LAN Client Control Manager (LCCM), if applicable.

User

Click the **User** page to display the User interface. The User interface displays information about the logged-in user.

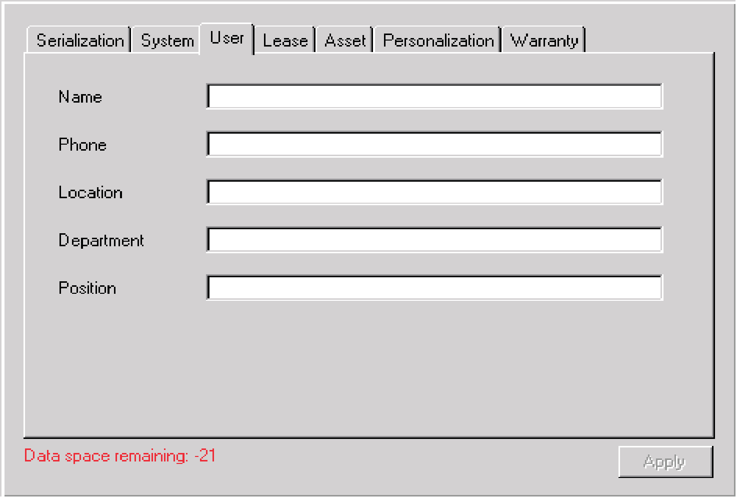

Name

Phone

Location

Department

Position

Data space remaining: -21

The User interface provides information about the items that are described in the following table.

Item	Description
Name	The user login name.
Phone	The user phone number.
Location	The user office location.

Item	Description
Department	The user department name or number.
Position	The user job title.

Lease

Click the **Lease** page to display the Lease interface. The Lease interface displays lease information for the agent system.

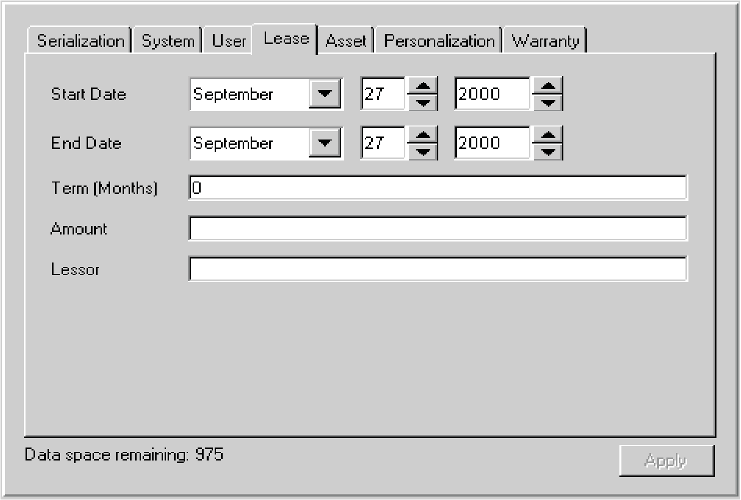

Start Date September 27 2000

End Date September 27 2000

Term (Months) 0

Amount

Lessor

Data space remaining: 975

The Lease interface provides information about the items that are described in the following table.

Item	Description
Start Date (mm/dd/yy)	The date that the lease agreement began.
End Date (mm/dd/yy)	The date that the lease agreement ends. If a Lease End Date is specified, a Warning alert will be generated when the lease expires.
Term (months)	The number of months for which the agent system is leased. In this field, you may enter an integer between 0 and 255.
Amount	The total price of the lease agreement. In this field, you may enter 20 characters or fewer.
Lessor	The name of the company that leased the agent system. In this field, you may enter 64 characters or fewer.

Asset

Click the **Asset** page to display the Asset interface. The Asset interface displays inventory information about the agent system.

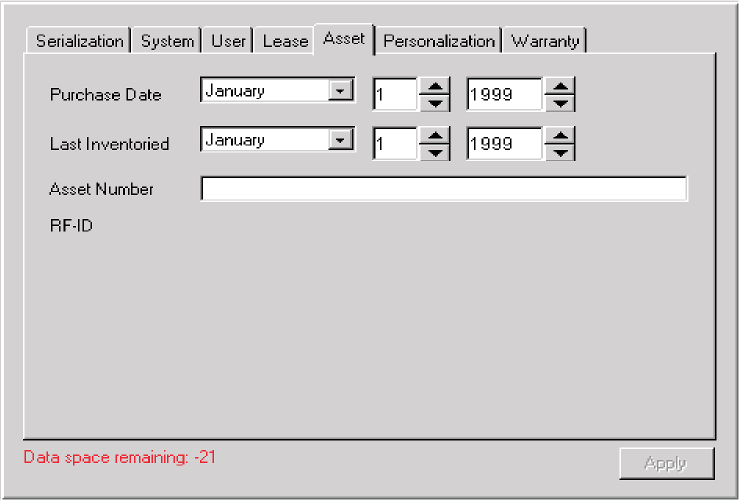

Purchase Date January 1 1999

Last Inventoried January 1 1999

Asset Number

RF-ID

Data space remaining: -21

The Asset interface provides information about the items that are described in the following table.

Item	Description
Purchase Date (mm/dd/yy)	The date the agent system was purchased.
Last Inventoried (mm/dd/yy)	The date of the last inventory check.
Asset Number	A unique number that is assigned to the agent system for inventory purposes. In this field, you may enter 64 characters or less.
RF-ID	The radio-frequency identification (RF-ID) number that was encoded in the agent system by the manufacturer. Not all computers have RF-ID capabilities. This is a fixed field and cannot be changed.

Personalization

Click the **Personalization** page to display the Personalization interface. The Personalization interface is a free-form window where you can type information about your users, system, or computer. There is a 64-character maximum for each of these fields.

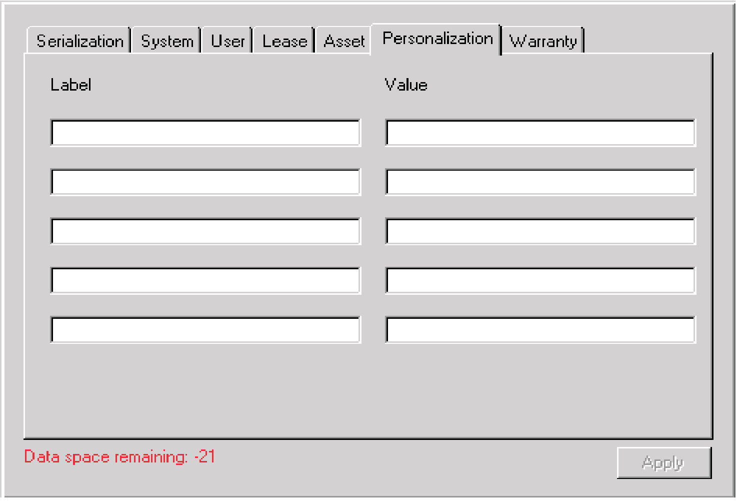

Label

Value

Data space remaining: -21

Warranty

Click the **Warranty** page to display the Warranty interface. The Warranty interface displays information about the warranty on the agent system.

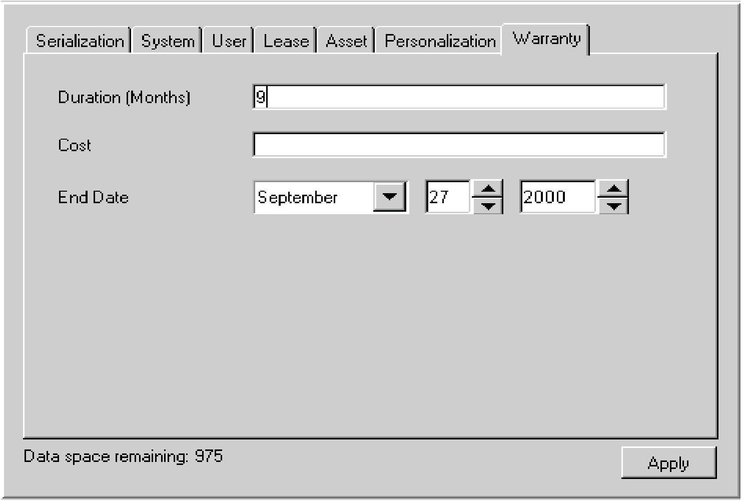

Duration (Months): 9

Cost:

End Date: September 27 2000

Data space remaining: 975

The Warranty interface provides information about the fields that are described in the following table.

Item	Description
Duration (months)	The duration of the warranty agreement. In this field, you may enter an integer between 0 and 255.
Cost	The total cost of the warranty. There is a 20 character maximum for this field.
End Date (mm/dd/yy)	The date that the warranty ends. If a Warranty End date is specified, a Warning level alert will be generated when the Warranty expires.

Date and Time

Use the Date and Time service to set the date and time that are displayed on the agent system. For the date, you have separate fields for month, day, and year. For the time, you have a field for the local time.

To start the Date and Time service, click **Tasks→Configuration→ Date and Time** in the Services pane. The following interface opens in the Display pane.

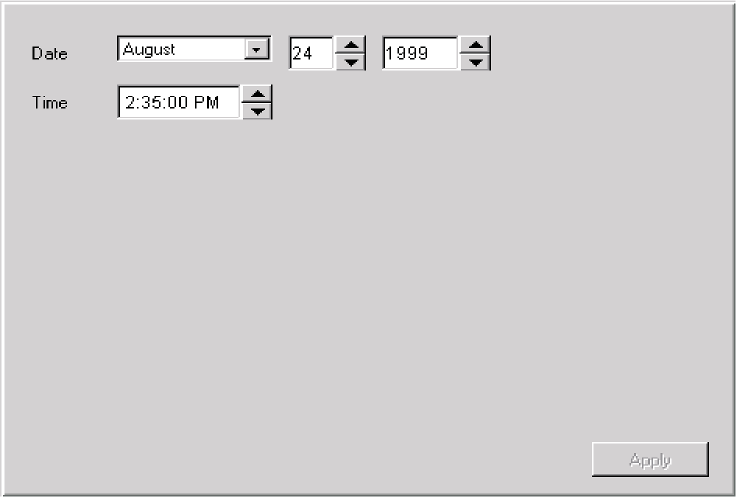

Date August 24 1999

Time 2:35:00 PM

Health Configuration

Use the Health Configuration service to enable and disable event actions, set minimum and maximum threshold values for temperature and drive space, and set the severities associated with each threshold. The Health Configuration tree is divided into two types: thresholds and bindings. Thresholds allows you to set boundaries within a manufactured range. Events will be triggered when these boundaries are reached.

To view the Health Configuration page, click **Tasks →Configuration →Health Configuration**. The Health Configuration page is divided into two sections. The left pane contains selectable items in a tree layout, and the right pane contains descriptive text or health configuration controls for the item selected on the left.

The Health Configuration tree is divided into two types: thresholds and bindings. The following list displays Thresholds and Binding settings.

Item	Description
Thresholds allows you to set boundaries within a manufactured range. Events will be triggered when these boundaries are reached. Thresholds are set for the following:	
Temperature Sensor	This is used to set thresholds for all temperature sensors in the system. The thresholds must be below manufacturer's limits. User-defined thresholds will not overwrite the manufacturer-defined thresholds, but will instead generate an additional alert when the user-defined threshold is exceeded. Thresholds can be defined for both warning and critical severities.
Disk Drives	This is used to set thresholds for all logical drives in the system. The thresholds indicate the amount of disk space remaining in terms of percentage remaining or as an absolute value in MB. Thresholds can be defined for both Warning and Critical severities.
Bindings allow you to enable or disable the reporting of events that are reported to the following:	
Event Log	Record the event in Windows NT, Windows 2000, or Windows XP. Event application logs.
Director	Generation of a corresponding Director server event.
Local Pop-up	Local popup dialog that displays a description of the event.
Microsoft SMS	Generation of a corresponding SMS status message.
TME Event Console	Generation of a corresponding TME 10 TEC event.
Director Agent Health	Record of the event's description in the System Health service.

The severities are selected by the user for a particular parameter corresponding to the impact the parameter has on the system. There are three supported severities: normal, warning, and critical.

IBM Director Agent Security

The IBM Director Agent Security displays user names, provides the ability to add and remove user names, and sets the security level and password options for each user name. (This page can be view only on systems running Windows 98 and Windows Millennium Edition.) Director Agent uses accounts configured through the operating system on Windows NT, Windows ME and Windows XP.

Security Levels
The security level that is assigned to a user affects that user's ability to access the IBM Director Agent program and the ability to view, manipulate, and access selected features in the program. The following security levels are associated with IBM Director Agent:

Disabled
A disabled user is not permitted to log in (usually a temporary state). The disabled state is not technically a security level but is included here because users with Administrator privileges can use it to override other security levels and temporarily prevent access to the IBM Director Agent program.

Browse, User, and Power User
Users with Browse, User, or Power User privileges can access much of the IBM Director Agent program. They cannot change settings or save typed information. In this release of the IBM Director Agent program, there are no practical differences between the Browse, User, and Power User privileges.

Administrator
Users with Administrator privileges have full control over the IBM Director Agent program. In addition to having full read and write access to the IBM Director Agent program, they can add new users, assign and change passwords, and assign security levels.

Click **Tasks**→ **Configuration** →**User Services Security** to display the following interface.

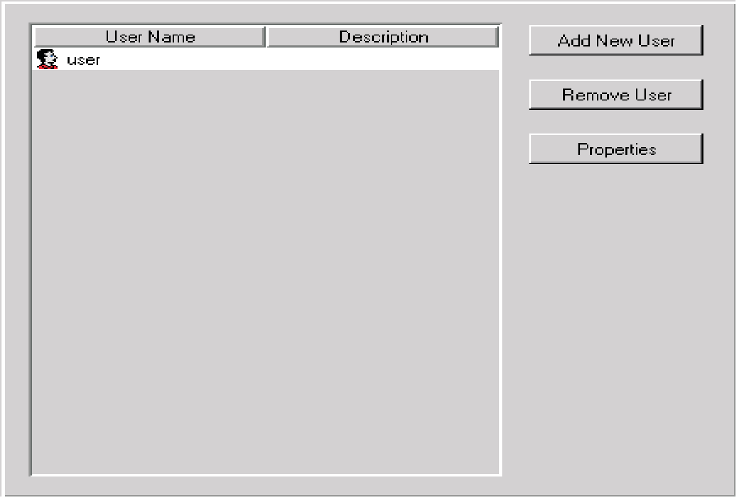

The IBM Director Agent Security interface provides information about the items that are described in the following table.

Item	Description
Add New User	To add a new user, highlight user name and click **Add New User.**
Remove User	To remove a user, highlight the user name and click **Remove User**.
Properties	To view or edit user properties, highlight the user name and click **Properties**

New User
Use the Add New User interface to add new IBM Director Agent users with the appropriate security levels and password options.

To display the Add New User interface, from the IBM Director Agent Security interface, click **Add New User**.

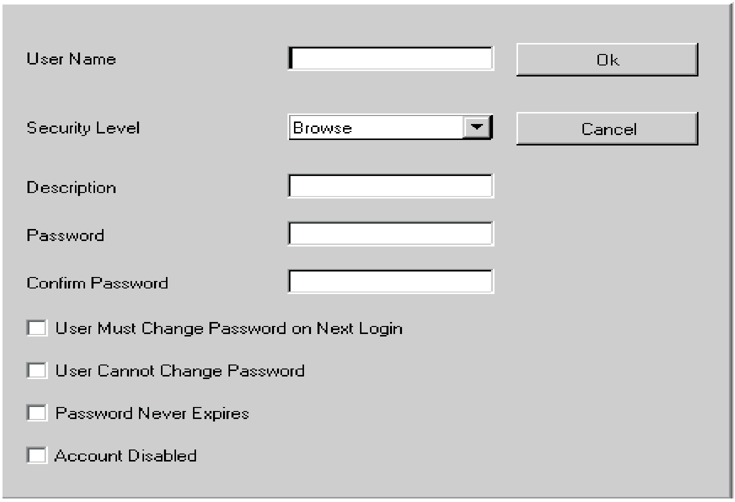

User Name

Security Level: Browse

Description

Password

Confirm Password

□ User Must Change Password on Next Login

□ User Cannot Change Password

□ Password Never Expires

□ Account Disabled

The Add New User interface contains the following items.

Item	Description
User Name	A unique character string that identifies the user (32 characters maximum).
Security Level	The level of system access authority the user is granted. From the list, select the security level that is appropriate for the user that you are adding.
Description	Information about the user, such as title, department, or reason for granting access to the IBM Director Agent program (32 characters maximum).
Password	The user password (32 characters maximum, case sensitive). There are no restrictions on the characters that can be used in passwords.
Confirm Password	This field must contain the same character string as the **Password** field (32 characters maximum, case sensitive).
User Must Change Password on Next Logon	Select this check box if you want to force the user to change the password the next time the user accesses the IBM Director Agent program.
User Cannot Change Password	Select this check box if you want to prevent the user from changing the password. If this check box is selected, only someone with Administrator privileges can change the password.

Item	Description
Password Never Expires	Select this check box if you do not require the password to be changed at regularly scheduled intervals.
Account Disabled	Select this check box if you want to temporarily disable this user’s access to the IBM Director Agent program. As a system administrator, you cannot disable your own account. This ensures that at least one account with Administrator privileges remains active.

The following table lists the user security levels.

Item	Description
User	Limited read/write
Power User	Limited read/write
Administrator	Read/write, lock/unlock, assign security levels, add new users and passwords

Note:

- Disabled users cannot access the IBM Director Agent program. Administrators can use the disabled state to override other security levels and temporarily prevent access to the IBM Director Agent program.
- Users with Browse, User, or Power User privilege have similar privileges in this release of the product. Though these users can access much of the IBM Director Agent program, they cannot change settings to save typed information.
- Users with Administrator privileges have the highest degree of control over the IBM Director Agent program. They can add new users, assign and change passwords, lock and unlock selected features, and assign security levels.

User

You can use the User Services Security interface to remove a user from IBM Director Agent. To remove a user, use the following procedure:

1. From the User Security interface, highlight the user that you want to remove.
2. Click **Remove User**. The following message is displayed:

   ```
   Are you sure you want to remove user?
   ```

3. Click **Yes**. The user is removed.

Viewing User Properties

You can use the IBM Director Agent Security interface to review or edit user properties such as description, security level, and password options. To view or edit user properties, use the following procedure:

1. From the **IBM Director Agent Security** interface, highlight the user that you want to view or edit.
2. Click **Properties**. The **User Properties** interface opens.

   You can view or edit the properties that are listed in this interface.

Network

The Network service provides information about your network. The following pages are available under **Network**:

IP Address	Provides routing information for your network.
DNS	Provides information on the distributed database system that is used to map domain names to IP addresses.
WINS	Provides information about the WINS server.
Domain/Workgroup	Provides information about the domain or workgroup for the agent.
Modem	Provides a list of modems installed on the agent system.

When you click **Network** from the IBM Director Agent Tasks page, the **IP Address** interface opens. Click the DNS page or the Modem page to view those windows.

IP Address

The IP Address interface provides routing information for your network.

Click **Tasks→Configuration→Network** to display the IP Address interface.

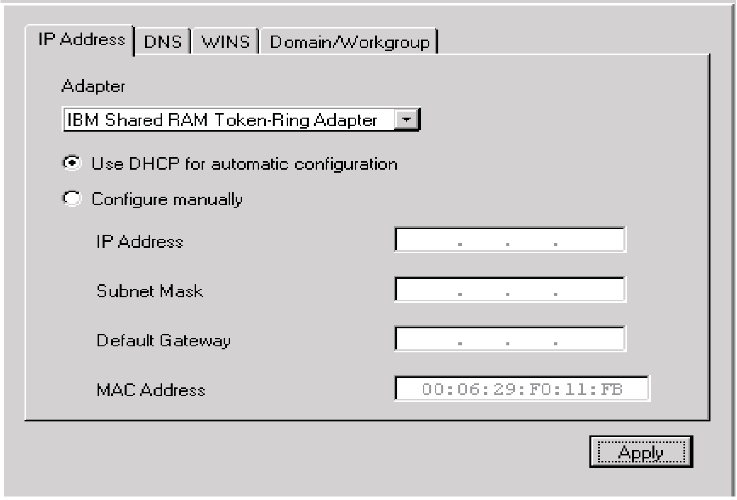

Adapter

IBM Shared RAM Token-Ring Adapter

● Use DHCP for automatic configuration

○ Configure manually

IP Address . . .

Subnet Mask . . .

Default Gateway . . .

MAC Address 00:06:29:F0:11:FB

The IP Address interface provides information about the items that are described in the following table.

Item	Description
Network Adapter	Select the appropriate network adapter from the list.
Use DHCP for automatic configuration	Select this option to configure IP addresses automatically.
Configure manually	Select this option to configure IP addresses manually. When this option is selected, the remaining entry fields are enabled.
IP Address	The IP address of the agent system. If you do not use DHCP to obtain an IP address, you must type the values into the **IP Address** and **Subnet Mask** fields manually.
Subnet Mask	A bit mask that is used to identify which bits in an IP address correspond to the network address and which bits correspond to the subnet portions of the address. The address mask has ones in positions corresponding to the network and subnet numbers and zeros in the host-number positions.
Default Gateway	The IP address for the default gateway server that you are using to communicate with other networks.

Item	Description
MAC Address	The unique hexadecimal number that identifies the network adapter in the agent system.

DNS

Domain Name System (DNS) is the distributed database system that is used to map domain names to IP addresses.

From the IBM Director Agent task list, click **Tasks** →**Configuration**→**Network** → **DNS** page to display the interface.

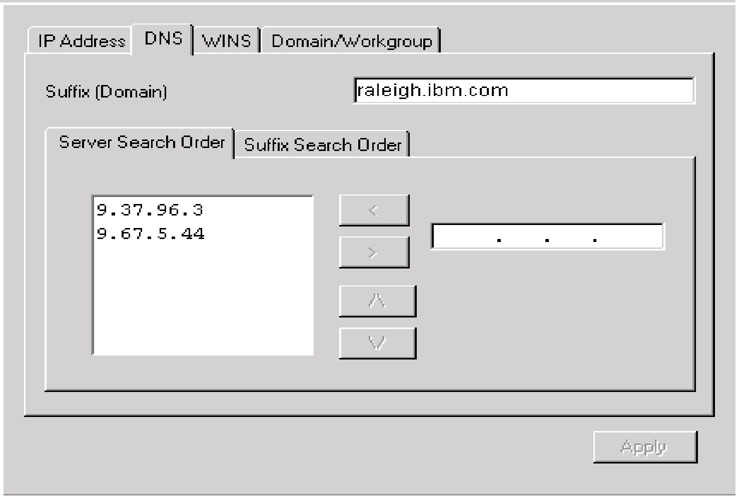

Suffix (Domain) raleigh.ibm.com

9.37.96.3
9.67.5.44

The DNS interface displays the following items.

Item	Description
Servers	The alphabetic identifier for your server with the network domain (IP address).
Suffixes	The text strings for the domain.
Apply	Changes are not saved until you click the **Apply** button.

Note: The up-arrow button is enabled when an item is selected in the list and is not the topmost entry. The down-arrow button is enabled when an item is selected in the list box and is not the bottom entry. Clicking the up-arrow or down arrow button moves the selected entry up or down one position in the list.

The **Remove** (>) button is enabled only when an item is selected in the list box. Removing an item from the list box removes the item from the list box and puts it into the text field.

WINS

The **Windows Internet Naming Service (WINS)** interface provides information about NetBIOS names and their corresponding IP addresses.

From the IBM Director Agent Task list, click **Tasks→Configuration→ Network→WINS** page to display the interface.

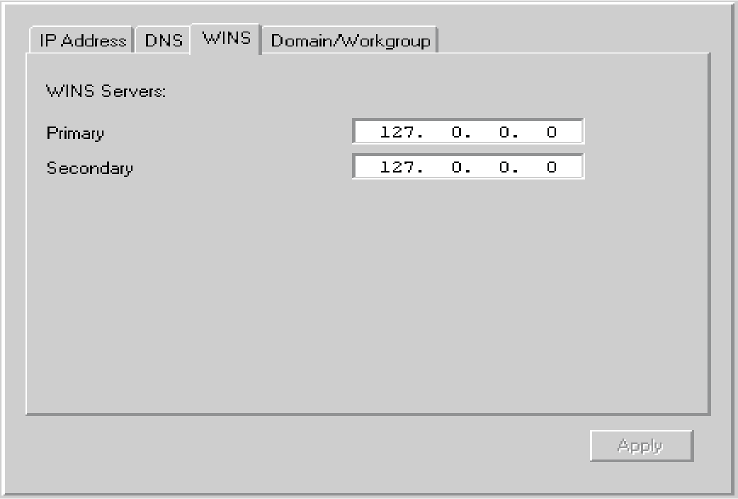

WINS Servers:

Primary 127. 0. 0. 0

Secondary 127. 0. 0. 0

The WINS interface provides information about the items that are described in the following table.

Item	Description
WINS Servers	Registered NetBIOS name with the associated IP address.
Primary	The address for the primary server.
Secondary	The address for the secondary server.
Apply	Click **Apply** to save changes.

Domain/Workgroup

The Domain/Workgroup interface lists the agent system with its associated domain or workgroup.

From the IBM Director Agent task list, click **Tasks→Configuration→Network→ Domain/Workgroup** page to display the interface.

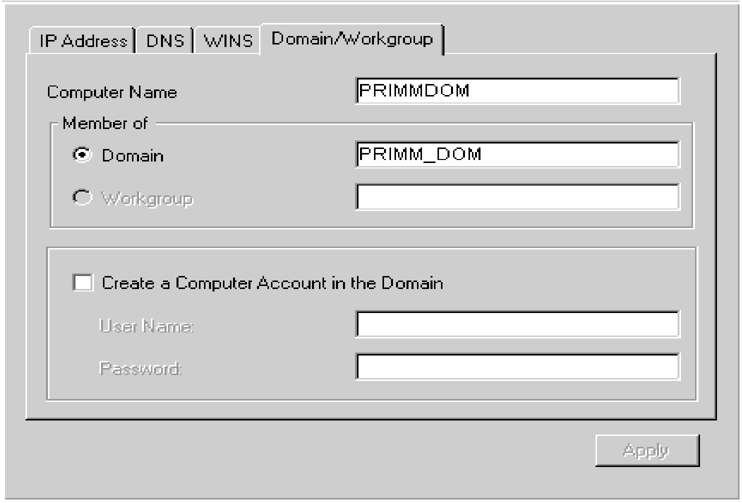

Computer Name: PRIMMDOM

Member of

● Domain: PRIMM_DOM

○ Workgroup:

□ Create a Computer Account in the Domain

User Name:

Password:

The Domain/Workgroup interface provides information about the items that are described in the following table.

Item	Description
Computer Name	The name given to the agent system. This naming scheme enables identification of the computer.
Domain	The agent is a member of the domain. A domain requires that users validate their accounts before they can log on to the network.
Workgroup	The agent is a member of a workgroup. The workgroup is a collection of agents and servers with no centralized logon validation.
Apply	Changes are not saved until the **Apply** button is clicked.

Modem

The Modem interface provides information about the modem type, speed, connector and device. From the IBM Director Agent task list, click **Tasks→Configuration→Network→Modem** page to display the interface.

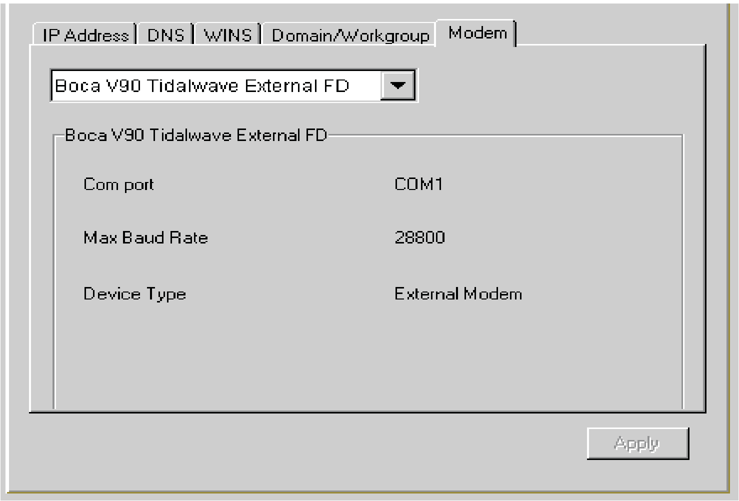

Boca V90 Tidalwave External FD

Boca V90 Tidalwave External FD

Com port COM1

Max Baud Rate 28800

Device Type External Modem

The modem interface provides information about the items that are described in the following table.

Item	Description
Modem	The drop-down list provides the name of the modem. Only installed modems are listed.
Com port	Lists the specific port that the modem is using.
Max Baud Rate	The maximum rate at which the modem operates.
Device Type	Describes the type of modem (internal or external).

SNMP

Note: The SNMP task appears on the task list only if the SNMP service is installed on the operating system.

The SNMP task provides the ability to work with community strings that are used in network communication and to set trap destination addresses.

Click **Tasks→Configuration→ Network→SNMP** to display the SNMP interface.

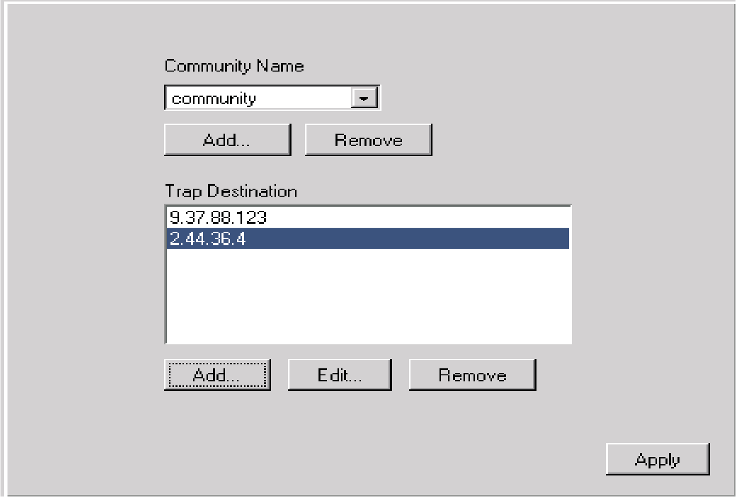
Community Name

community

Trap Destination

9.37.88.123
2.44.36.4

The **SNMP** screen displays the following the items.

Item	Description
Community name	A unique character string that identifies the community. The community name enables your network-management system (NMS) to verify that a server is authorized to take a specific action. If the server community name matches the community name that is assigned to the requested information or action, the NMS provides the information or action to the server. You can add or remove a unique community name.
Trap Destination	A list of network management system IP addresses to which the server can send alerts. You can add, remove, or edit a trap destination. To modify a trap destination, select an IP address and click **Edit**.

System Accounts

The System Accounts task provides remote administration of user security and group security within a Windows operating system. (This task can be used only on systems running Windows NT, Windows 2000, or Windows XP). To start the System Account service, click **Tasks →Configuration→ System Accounts** in the Configuration pane.

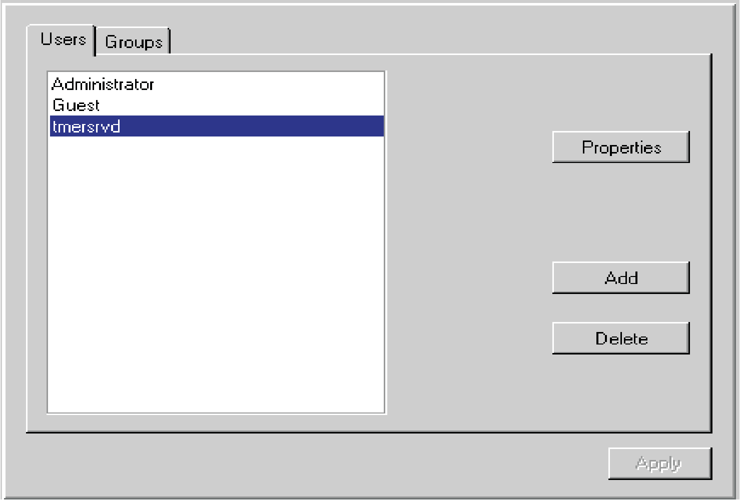

Users

The Users interface provides information about the items that are described in the following table.

Item	Description
Properties	Edit or view user properties
Add	Click the **Add** button to add a new user.
Delete	Click the **Delete** button to delete a user.

Groups

The Groups interface enables the administrator to review and edit members within the group. Click the **Groups** page to display a list of all groups.

Note: The **Properties** and **Delete** buttons are initially disabled. They become enabled when a group is selected in the list.

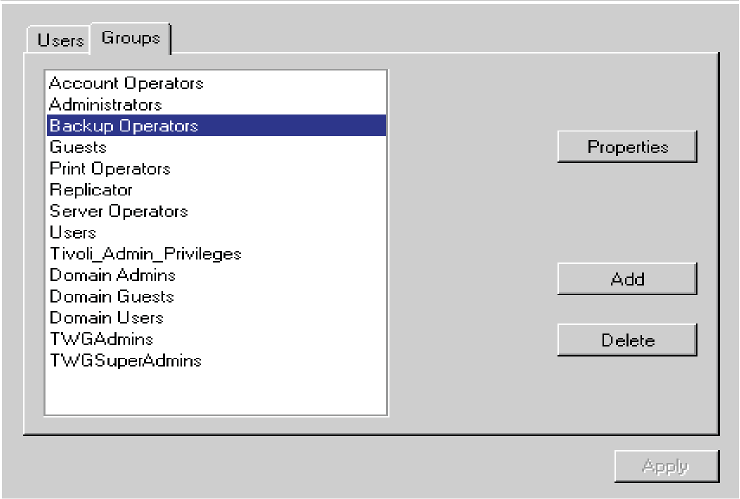

Account Operators
Administrators
Backup Operators
Guests
Print Operators
Replicator
Server Operators
Users
Tivoli_Admin_Privileges
Domain Admins
Domain Guests
Domain Users
TWGAdmins
TWGSuperAdmins

The **Groups** interface provides information about the items that are described in the following table.

Item	Description
Groups	List of global groups.
Properties	Edit or view group properties.
Add	Click the **Add** button to add a new group.
Delete	Click the **Delete** button to delete a group.

General

The General interface is used to give IBM Director Agent users the appropriate security levels and password options. Click the **Add** button on the Users or Groups interface to display the General page.

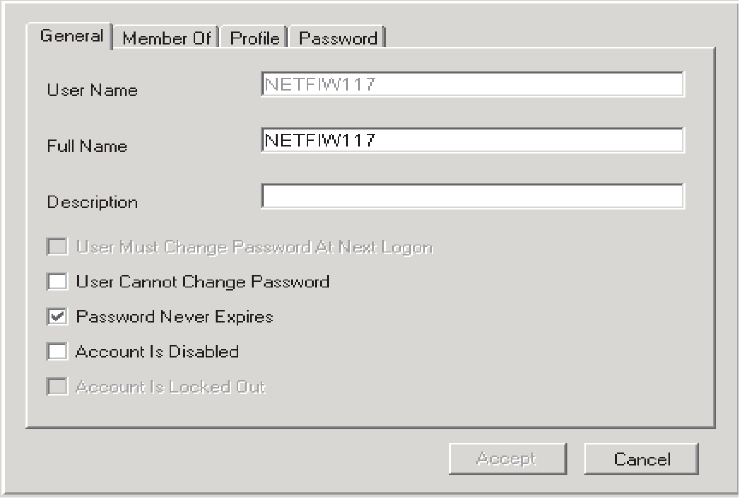

User Name: NETFIW117

Full Name: NETFIW117

Description:

☐ User Must Change Password At Next Logon

☐ User Cannot Change Password

☑ Password Never Expires

☐ Account Is Disabled

☐ Account Is Locked Out

Item	Description
User's Name	A unique character string that identifies the user (32 characters maximum).
Full Name	User's complete name.
Description	Information about the user, such as title, department, or reason for granting access to the IBM Director Agent program (32 characters maximum.
User Must Change Password at Next Logon	Select this check box if you want to force the user to change the password the next time the user accesses the IBM Director Agent program.
User Cannot Change Password	Select this check box if you want to prevent the user from changing the password. If this check box is selected, only someone with Administrator privileges can change the password.
Password Never Expires	Select this check box if you do not require the password to be changed at scheduled intervals.
Account is Disabled	Select this check box if you want to temporarily disable a user's access to the Director Agent program. As an Administrator, you cannot disable your own account. This ensures that at least one account with Administrator privileges remains active.

Item	Description
Account is Locked Out	Select this check box if you want to disable a user's access to Director.
Accept	Click the Accept button to save changes.
Cancel	Select the **Cancel** button to cancel changes.

Member Of

The Member Of interface displays a group membership list. Members are listed on the right pane, and non-member groups are listed in the left pane. To access this interface, from the IBM Director Agent task list, click **Tasks→Configuration→System Accounts→Add →Member Of**.

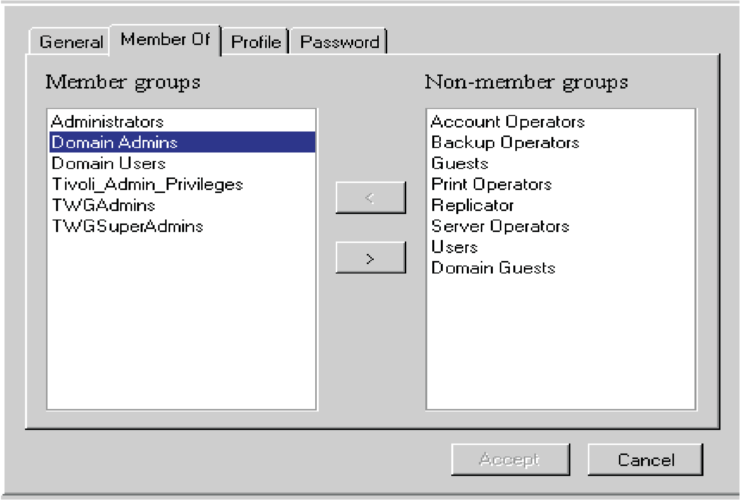

Member groups

Administrators
Domain Admins
Domain Users
Tivoli_Admin_Privileges
TWGAdmins
TWGSuperAdmins

Non-member groups

Account Operators
Backup Operators
Guests
Print Operators
Replicator
Server Operators
Users
Domain Guests

Clicking the < or > buttons will move group names from the **Member groups** list to/or from the **Non-member** groups list.

The Member Of interface provides information about the items that are described in the following table.

Item	Description
Member groups	A list of users within the group.
Non-member Groups	A list of users who are not members of the group.
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.

Profile

Use the Profile interface to configure user profiles. To access the Profile interface, from the IBM Director Agent task list, click **Tasks→Configuration→ System Accounts→Add →Profiles.**

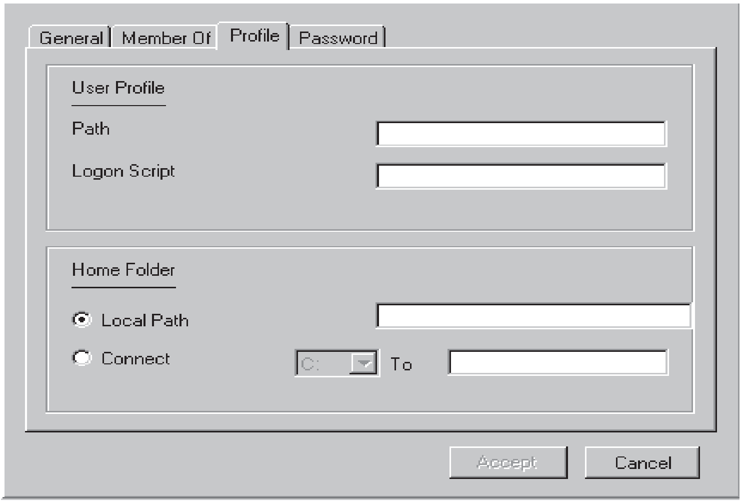

The Profile interface provides information about the items that are described in the following table.

Item	Description
Path	The network path to the user's profile folder. Type a network path in the form\\*server name**profile folder name**user name*.
Logon Script	A script assigned to a user account that runs each time the user logs on.
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.

Password

Use the Password interface to enter a new password or change an existing password. To access the Password interface, from the IBM Director Agent task list, click **Tasks→Configuration→System Accounts→ Add→Password**.

Item	Description
New Password	The user's new password (32 character maximum, case sensitive).
Confirm Password	This field must contain the same character string as the New Password field (32 character maximum, case sensitive).
Accept	Click the **Accept** button to save changes.
Cancel	Click the **Cancel** button to cancel changes.

Tools

On the IBM Director Agent task list there are two items under **Tools**:

- **Remote Control —**Provides a way to control one computer to from another computer.
- **Shutdown —** Provides three shutdown options.

Note: Administrator level of security is required to use Tools functions.

Remote Control

Notes:

1. The Remote Control service appears on the task list only if the Remote Control option is installed on the agent system. If you did not select **Web Based Remote Control** during the IBM Director Agent installation, the Remote Control option is not displayed.

2. Remote Control is not supported when you use a Web browser or MMC to manage the agent systems

   You can use the Remote Control interface to set configuration options for Remote Control sessions.

Click **Tasks** →**Tools** →**Remote Control** to display the Remote Control interface.

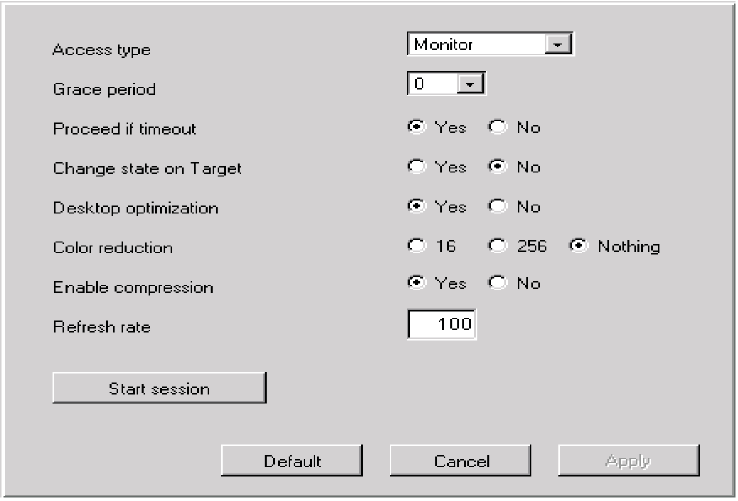

Access type: Monitor

Grace period: 0

Proceed if timeout: ● Yes ○ No

Change state on Target: ○ Yes ● No

Desktop optimization: ● Yes ○ No

Color reduction: ○ 16 ○ 256 ● Nothing

Enable compression: ● Yes ○ No

Refresh rate: 100

The Remote Control interface displays the following items.

Item	Description
Access type	The access type (Monitor or Active) determines whether you will monitor the agent system or actively control it.
Grace period	Number of seconds before the management server takes active control of the agent system. This is the number of seconds between the warning and the actual active control. If you are having trouble with the management server timing out, increase the grace period.
Proceed if timeout	This timeout option is associated with the grace period. If you click **Yes**, the management server automatically takes control of the agent system if the grace period times out before you get a response from the agent system.
Change state on Target	If you click **Yes**, the agent system can take back control from the remote computer.
Desktop optimization	If you click **Yes**, you can disable the background wallpaper of the agent system for faster transmission between computers.

Item	Description
Color reduction	For faster transmission between computers, you can compress the display on the agent system. This number (16, 256, or Nothing) represents the number of pixels to which the display will be compressed.
Enable compression	If you click **Yes**, the system compresses data for faster transmission between computers, but the user must wait for the compressed data to be decompress after the transfer. If you are having difficulty transferring the data, enable compression.
Refresh rate	A numeric value that represents the delay time between the controlling computer and the agent system. The default is 100 milliseconds.

Shutdown

The Shutdown service provides three options for shutting down your system:

- **Shutdown and Power Off —** Shut down and turn off the computer.

  Note: Shutdown and Power Off are available only on systems that support and have enabled Advanced Power Management.

- **Restart —** Shut down and restart your computer without turning it off.
- **Log Off —** Log off your operating system without shutting down the computer.

Click **Tasks →Tools→Shutdown** from the IBM Director Agent task list to display the following interface.

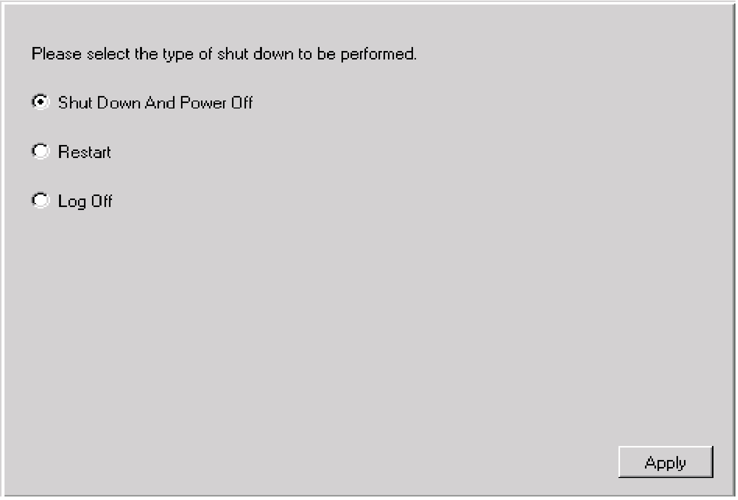

Please select the type of shut down to be performed.

● Shut Down And Power Off

○ Restart

○ Log Off

Web Links

The Web Link gives you immediate access to the latest drivers and news about your system.

System Updates

Use the **System Updates** option to connect to an IBM Web site that provides updated information for your specific system. This option works only if you have the ability to connect to the Internet. Click **Tasks** →**Tools** → **System Updates** from the IBM Director Agent task list. The System Updates interface opens.

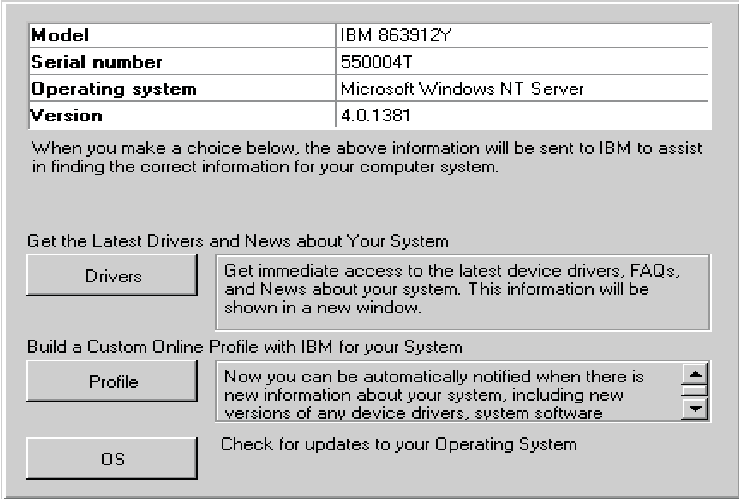

Model	IBM 863912Y
Serial number	550004T
Operating system	Microsoft Windows NT Server
Version	4.0.1381

When you make a choice below, the above information will be sent to IBM to assist in finding the correct information for your computer system.

Get the Latest Drivers and News about Your System

Get immediate access to the latest device drivers, FAQs, and News about your system. This information will be shown in a new window.

Build a Custom Online Profile with IBM for your System

Now you can be automatically notified when there is new information about your system, including new versions of any device drivers, system software

Check for updates to your Operating System

The following items are available from the System Updates window.

Item	Description
Table of machine information	The agent system model number, serial number, operating system, and version number.
Get the latest drivers and news about your system	Immediate access to the latest device drivers, technical information, and news about the agent system.
Build a custom online profile with IBM for your system	Be notified automatically when there is new information about the agent system.

Appendix K. Upward Integration Modules

This chapter provides information on installing and using Upward Integration Modules (UIMs) on supported system-management platforms.

UIMs enable workgroup-and enterprise-level system-management products to interpret and display data that is provided by agent systems running IBM Director Agent. The modules provide enhancements on the management server that enable the system administrator to start IBM Director Agent from within the system-management platform, collect IBM Director Agent inventory data, and view IBM Director Agent alerts. UIMs are provided for the following system-management platforms:

- Tivoli Enterprise™, including Tivoli Framework 3.6.2, 3.7.1 Tivoli Software Distribution 3.6, 3.6.2 and Tivoli Enterprise Console 3.6, 3.6.2, 3.7, Tivoli Inventory 4.0
- Tivoli NetView 5.1.1 and 6.0 for Windows NT
- CA Unicenter TNG for WIN32 Version 2.4, AIM IT Version 3.0, Asset Management Option 3.0, SHIP IT Version 2.0, Software Distribution Option 2.0
- Intel LANDesk Management Suite 6.4
- Microsoft SMS 2.0
- HP OpenView Network Node Manager

You can use the IBM Director Agent installation program to install the Alert on LAN proxy agent on your system-management platform. The Alert on LAN proxy is not an IBM Director Agent UIM, but this proxy agent must be installed on your system-management platform to receive Alert on LAN messages from Alert on LAN-enabled agent systems.

Installing Upward Integration Modules

The method that is used to install an Upward Integration Module depends on the system-management platform for which you are installing support.

- If you are installing the Tivoli Enterprise Plus Module, see “Installing the Tivoli Enterprise Plus Module” on page 506.
- If you are installing Intel LANDesk Management suite integration, see “Intel LANDesk Management Suite Integration” on page 520.
- If you are installing an Upward Integration Module for any other supported system-management platform or are installing the Alert on LAN proxy agent, use the IBM Director Agent installation program to install the UIM on the system-management platform. Copy the IBM Director Agent installation files to a directory on the system-management system and then go to the

same platform-specific section of this chapter for additional installation instructions and usage information.

Systems Management Platform	Additional Information
Tivoli NetView	“Tivoli NetView 6.0 Integration” on page 511
CA Unicenter TNG	“CA Unicenter TNG Integration” on page 518
Microsoft SMS	“Microsoft SMS integration” on page 521
HP OpenView Network Node Manager	"HP OpenView Integration" on page 529

- If you are installing the Intel Alert on LAN proxy on your system-management platform, see “Installing Intel Alert on LAN Proxy” on page 528.
- If you are installing IBM Director Agent Support on the HP OpenView Server, see page 529

Tivoli Enterprise Plus Module Integration

Tivoli Enterprise is part of Tivoli Management Environment (TME)10. The Tivoli Enterprise UIM adds a module that enables a system administrator using Tivoli Enterprise to manage agent systems that have IBM Director Agent installed. For example, the system administrator can shut down, restart, and wake up any selected agent system that has IBM Director Agent installed.

Installing the Tivoli Enterprise Plus Module

Follow these steps to install the Tivoli Enterprise Plus Module:

1. Copy the Tivoli Enterprise Plus Module to a temporary directory on the system that is running Tivoli Enterprise. You can download the Tivoli Enterprise Plus Module from

   http://www.pc.ibm.com/ww/alliances/lifecycle/ums/download.html
2. From the **Select Product** list, highlight **IBM Director Agent Plus-Tivoli** and click **Submit**.
3. Use a file decompression program that supports the TAR file compression format to extract the contents of the file to a temporary directory, for this example, IBM Director Agent Plus for Tivoli.
4. Use the Tivoli Desktop to install the Tivoli Enterprise Plus Module.
   a. From the Tivoli Desktop menu, click **Install** → **Install Product**.
   b. Select your host and directory. Select the temporary directory **IBM Director Agent Plus for Tivoli**, which contains the Tivoli Enterprise Plus Module files.

c. Click **Set Media → Close**.

d. Install the Plus Module Support link binaries first, and then install the IBM Director Agent Plus module for Tivoli. You must install the Plus Module on the Tivoli Management Region (TMR) and on any other managed nodes from which the Plus Module will be used.

   Notes:

   1) Because of a limitation in Tivoli Enterprise 3.6, the only administrator roles that can install the Plus Module are root (for systems running UNIX) and administrators (for systems running Windows NT).

   2) When installing the Tivoli Enterprise Plus Module, the administrator must use a fully qualified Tivoli login name (in *name@domain* format).

Enabling Additional Functions

After you install the Tivoli Enterprise Plus Module, use the following information to enable additional functionality.

- Enabling Wake on LAN®support
  To use the Wake on LAN component, you must have a Java Virtual Machine installed on the computer on which the Tivoli Enterprise Plus Module is installed. Also, you must install the Inventory module and collect inventory from the agent endpoints, before you attempt a wake-up.
- Enabling Software Distribution support

  To enable Software Distribution support, install the Software Distribution Gateway from the Software Distribution CD onto a managed node before installing the IBM Director Agent Plus Module so that Tivoli endpoints can be targeted. Additionally, you must install the Software Distribution product on any managed node where the Tivoli Plus Module is installed.

  1. Before you can distribute IBM Director Agent software, you need a source computer and a source directory that contains the IBM Director Agent installation files. This is your staging location for distributing software.You also need a destination drive and location (for example, C:\temp) on the target systems where the installation files will be copied. This destination drive and directory must exist on all agent systems before you distribute the software.

     Notes:

     a. The source path of the IBM Director Agent FilePack profile designates the root directory as the beginning of the installation. If you have changed or added to the installation file path /Win32/Install/en you must edit the correct path in the file UMS_fp_after.bat.

For example, if you mount the *Director CD* on a system running UNIX with a mount point of /cdrom, you add the change directory command of

```
CD CDROM
```

before the other change directory commands. The edited section of the batch file will look like this:

```
REM Next line is Unix CDROM mount

CD CDROM

CD Win32

CD Install

CD en
```

b. This destination differs from the final destination for the IBM Director Agent software (for example, C:\program files\ibm\ UMS, which must be configured with the UIMSETUP.ISS file. For more information on editing theUIMSETUP.ISS file, see “Modifying the UIMSETUP.ISS File Manually” on page 523.

2. Before doing any software distribution, run the **Prepare for IBM Director Agent Install** service to configure an IBM Director Agent file package. Take the following steps:

   a. Double-click the **Prepare for IBM Director Agent Install** icon.

   b. In the **Source Host** field, type the machine name of the system where the installation files are located.

   c. In the **Source Path** field, type the directory path where the installation files are located.

   d. In the **Destination Path**, type the directory path of the managed system where the files are to be distributed.

   e. Click **Set and Close**.

f. Right-click the **Install IBM Director Agent** icon. In the menu, click **Distribute**.

3. You can distribute IBM Director Agent to any supported Windows operating system (Windows 98, or Windows NT 4.0 or later). A log file (named umsinst.log) records the results of the software distribution. The file is located in %DBDIR%\..\tmp.

- Installing the Inventory Gateway product

  Install the Inventory Gateway product from the Tivoli Enterprise Inventory CD onto a managed node before you install the Plus Module. The Inventory Gateway product must be installed on a managed node where the Plus Module is installed.

- Enabling distributed monitors

  Distributed monitors are not supported on Windows 98 endpoints. To enable distributed monitors on systems running Windows NT, install the distributed monitor package on a managed node where the Plus Module is installed.

  IBM Director Agent monitors the IBM Director Agent HTTP DAEMON and the SNMP subagent processes.

- Activating Tivoli Enterprise Console (TEC) integration
  To activate TEC integration, run the Setup TEC Event Server for the IBM Director Agent task on the TEC server.

  — TEC events from IBM Director Agent Monitors

  The IBM Director Agent SNMPCheck and HTTPCheck monitors send TEC events when the IBM Director Agent SNMP subagent or HTTP DAEMON processes stop. The Events register, with the corresponding TEC indicator, changes its temperature icon depending on the severity of the event. Additionally, the events appear in the TEC console as members of the UM_Services_Plus Event Group.

  Note: The default configuration for each monitor is for critical responses only.

  To change the default settings in the Edit Monitor window, you first must select the critical response level. Otherwise, you will create a new configuration rather than change the existing one.

  — Automated actions in response to events sent by IBM Director Agent Monitors

  When the httpserv.exe and snmp.exe processes stop, the TEC server responds by restarting them automatically.

- Starting IBM Director Agent
  IBM Director Agent provides a Web browser-based console that you can use on any system that supports Netscape 4.5 or later, Internet Explorer 4.01 or later, and Java 1.1.7b or later. This includes UNIX-based Tier 1 nodes. However, because the Tivoli application does not import system-wide environment variables on UNIX platforms, such as CLASSPATH, that are needed by IBM Director Agent, you must add the CLASSPATH environment variable to the Tivoli setup_env.sh or setup_env.csh scripts to enable IBM

Director Agent startup support. Also, under AIX, the MOZILLA_HOME variable that Netscape uses must be included in the setup_env.sh or setup_env.csh scripts. After you set these variables, run setup_env.sh to enable the startup support for IBM Director Agent in the Plus Module.

- Adding Plus Module icons

  The Plus Module installation automatically places startup icons in the collection window for every managed node in the TMR that has the Plus Module installed. The launch functionality for a specific managed node works only when you double-click the startup icon for that managed node (for example, Launch IBM Director Agent@*hostname*.)

- Managing large numbers of managed nodes

  For TMRs with a large number of managed nodes running the Plus Module, the number of startup icons in the Plus Module collection window can become excessive. You can remove some of these icons from the collection window by clicking an icon to highlight it and clicking **Remove** from the **Edit** menu. Alternatively, you can create a separate collection just for the startup icons by clicking **Create** → **Collection**, and then dragging and dropping the individual icons into the new collection.

  Note: In this case, the icons still must be removed from the original view.

- Deploying endpoints with IBM Director Agent

  Deploying endpoints with a preconfigured gateway helps ensure their successful initial login. You can specify the gateway for a particular endpoint in the silent installation script tmasetup.iss that is included in the IBM Director Agent installation package. Open the file and scroll to

  ```
  [SdShowDlgEdit3-0]
  szEdit1=9494
  szEdit2=9495
  szEdit3=-d1
  Result=1
  ```

  The key

  `szEdit1`

  specifies the port through which the gateway communicates. It is 9494 by default.

  The key

  `szEdit2`

  specifies the port through which the endpoint communicates, also 9495 by default.

  The key

  `szEdit3`

  specifies any command lines to pass to the installation action program. Use this key with the following flags to specify a gateway:

```
szEdit3=-d1 -g gateway host+ gateway listening port
```

Using the Tivoli Enterprise Plus Module

You can use Tivoli Enterprise Plus Module to perform the following additional system-management tasks from the system where IBM Director Agent is installed:

- Configure IBM Director Agent for all platforms
- Conduct IBM Director Agent queries
- Obtain IBM Director Agent inventory on an agent computer
- Use the IBM Director Agent indicators for monitors

You can also select a remote system and perform any of the following tasks remotely:

- Restart (shut down and restart) an IBM Director Agent system
- Shut down an IBM Director Agent system
- Wake up an IBM Director Agent system

Installation instructions for the Plus module can be found in the file UM_Services_ Plus. pdf. The file is located in the installation media in the Director/UM Services Plus For Tivoli folder.

Tivoli NetView 6.0 Integration

Tivoli NetView can be used to manage agents running IBM Director Agent. To enable this functionality, you must use the IBM Director Agent installation program to install the IBM Director Agent NetView Upward Integration Module on the system that is running NetView Server.

Installing the Tivoli NetView Upward Integration Module

To install the Tivoli NetView UIM, use the following procedure:

1. Establish an appropriate remote access policy. The NetView system administrator must have access privileges on every remote agent system so that IBM Director Agent inventory data can be obtained for each agent system. The easiest way to do this is to create a Windows NT domain in which every agent is a member. The system administrator should be a member of the Domain Admins group. This enables the system administrator to access all the computers in the domain (and any trusted domains) without requiring further authentication.
2. Install the IBM Director Agent on the agent systems. You must install the Web Based Access component on the IBM Director Agent system to classify agent systems as IBM Director Agent-capable. It is important to do this before installing the NetView UIM on the NetView server. For more

information on installing IBM Director Agent on systems, see Appendix G, “Preparing to install IBM Director Agent,” on page 423.

3. Use **dcomcnfg.exe** to enable DCOM connections on agent systems running Windows 98.

   You can download **dcomcnfg.exe** from

   http://www.microsoft.com/com/tech/DCOM.asp

   Ensure that **CONNECT** level authentication and **IMPERSONATE** level impersonation are selected.

4. Stop the NetView server. Click **Programs**→ **NetView**→ **Administrator**→**Stop Server**.

5. Start the IBM Director Agent installation program on the NetView server, and insert the Director CD. The Director window opens.

6. Click **Install Director**.

7. Click **Next**.

   The License Agreement window opens. Click **Accept** to proceed.

   Note: You must agree to the terms of the License Agreement to install IBM Director Agent. If you click **Decline**, the installation program will close.

   The Select Components window opens.

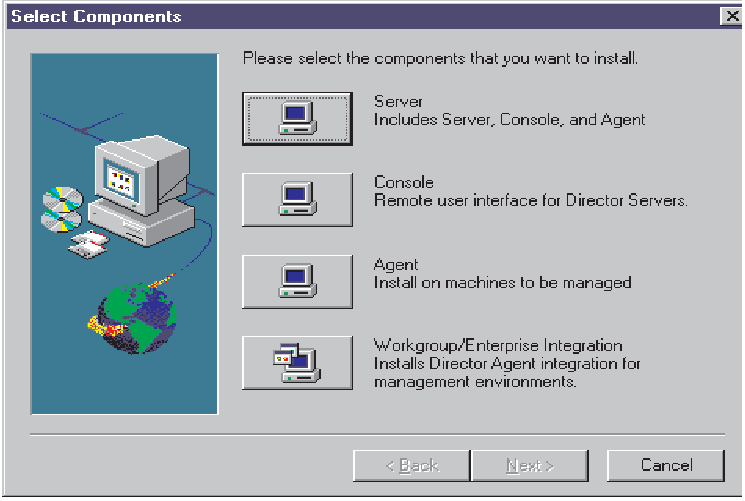

Select Components

Please select the components that you want to install.

Server
Includes Server, Console, and Agent

Console
Remote user interface for Director Servers.

Agent
Install on machines to be managed

Workgroup/Enterprise Integration
Installs Director Agent integration for management environments.

8. From the Setup Options window, click the **Workgroup/Enterprise Integration-- Installs IBM Director Agent integration for management environments** button.

   The Integration Selection window opens.

9. Click **Tivoli NetView Upward Integration**, and then click **Next**.

   The installation program adds the IBM Director Agent SmartSet, loads IBM Director Agent MIB files, adds trap filters for IBM Director Agent SNMP traps, and adds menu items for starting and inventory collection into the **NetView Tools** menu.

10. When the installation is finished, restart the server. The integration setup program configures **nvsniffer.exe** to run after the system is restarted and has populated the IBM Director Agent SmartSet.

Starting IBM Director Agent on NetView agents

NetView agents with the IBM Director Agent Web Based Access component installed can be targets for starting IBM Director Agent. To start IBM Director Agent on NetView agents, use the following procedure:

1. Open the IBM Director Agent SmartSet and then select an agent system.
2. From the **Tools** menu, click **IBM Director Agent Browser.**

   This selection will be disabled and unavailable if the agent system you selected does not have the Web Based Access component installed.

IBM Director Agent starts on the selected agent, using the management system default Web browser. A valid IBM Director Agent user ID and password are required to use IBM Director Agent to manage a remote agent system.

Using Tivoli NetView 5.1.1 and 6.0 to Obtain Inventory Data

NetView agents with the UMSCIM Object capability can be targets for collecting IBM Director Agent WBEM-based inventory. Inventory data from IBM Director Agent can be collected in one of the following ways:

- To create a new NetView sniffer configuration file, use the following procedure:
  1. Open the **nvsniffer.conf** file.
  2. Remove the # from the following line:

```
isumscim|||\usr\ou\bin\umscimtest.exe\usr\ou\bin\
umscimtest.exe\
```

3. Open the **nvsniffer_ums.conf** file
4. Remove the # from the following line:

   ```
   isumscim|||\usr\ou\bin\umscimtest.exe\usr\ou\bin\
   umscimtest.exe\
   ```

- Use the Windows NT Schedule service to run nvsniffer.exe.

  By default, NetView configures the Schedule service to run nvsniffer.exe daily at 1 a.m. However, because IBM Director Agent inventory collection relies on a remote DCOM connection to the agent system to access its WBEM repository, you must reconfigure the Schedule service to log in as a user with remote access privileges on the agent system.

  To reconfigure the NT Schedule service, use the following procedure:

  1. Click **Control Panel → Services**.
  2. Select the Schedule service.
  3. Click **Startup** and configure the service to log on as a user to join the Domain Admins group.
- From the NetView Console
  1. Open the NetView console.
  2. Open the SmartSets submap.
  3. Open the IBM Director Agent SmartSet.
  4. Select a system in the submap for which to collect inventory information.
  5. Open the **Tools** menu and select the **IBM Director Agent Inventory** item.

     A menu displays the different categories of WBEM-based inventory information that is provided by IBM Director Agent on the agent system.
  6. Click **Demand Poll** to have the data collected from each agent in the NetView database.
  7. To see the database-logged IBM Director Agent inventory data from a particular system, from the command prompt, type:

     ```
     ovobjprint -s <hostname>
     ```

     Note: If the agent system that you selected does not have a UMSCIM Object capability (there is no remote access policy to the computer from the NetView server), the IBM Director Agent Inventory item will be disabled.

Viewing IBM Director Agent SNMP Data from NetView

During the installation of IBM Director Agent NetView integration support, trap filters for IBM Director Agent SNMP traps are added to the NetView **trapd.conf** file. Thus, when an SNMP trap is sent from an IBM Director Agent system, it can be viewed in the NetView Event Browser. Only traps of critical severity are filtered, and the default action is to change the icon color of the source computer from green to red. Traps indicate an abnormal environment condition on the IBM Director Agent, such as chassis intrusion, a removed processor, or temperature out of range. To configure the NetView console to display advanced menu items, use the following procedure:

1. Click **Options** → **Advanced** to configure the NetView console to display advanced menu items and query this data.
2. Shut down and reopen the NetView console.
3. Open the IBM Director Agent SmartSet and select a system from which to view SNMP data.
4. To start the NetView SNMP browser, open the **Tools** menu and click **MIB** → **Browser**. Ensure that the selected system node name is displayed in the **Node Name** or **Address** field, and click **Get Values**.

   Note: To view specific Alert on LAN SNMP traps, you must use the SNMP V2 browser. Click **MIB** → **SNMP V2** → **Browser**.
5. The NetView SNMP collection DAEMON will contact the IBM Director Agent SNMP subagent on the agent system and query it for the data published in the IBM Director Agent MIB files.

Because SNMP support is an optional component of IBM Director Agent, and not a required component, not all systems in the IBM Director Agent SmartSet will have the IBM Director Agent SNMP subagent installed. Those that have the subagent installed will have UMSSnmp Object capability. The "is UMSSnmp" test is not enabled by default in nvsniffer.conf, but it can be enabled by opening the nvsniffer.conf and removing the comment symbol from the line that begins with "is UMSSnmp".

Forwarding SNMP Trap Information

IBM Director Agent forwards SNMP trap alerts to the NetView administrator for critical IBM system environmental conditions, low disk space, a failing hard disk drive, and a system being removed from a LAN. During the installation of the IBM Director Agent upward integration support for NetView, these traps are added to the NetView **trapd.conf** file with their Trap Properties configured using the **addtrap.exe** utility.

IBM Director Agent forwards the following SNMP traps to the NetView server workstation:

- iBMPSG_TemperatureEvent
- iBMPSG_VoltageEvent
- iBMPSG_ChassisEvent

- iBMPSG_FanEvent
- iBMPSG_StorageEvent
- iBMPSG_SMARTEvent
- iBMPSG_LANLeashEvent

The following list describes the default properties that are configured for IBM Director Agent SNMP traps:

Enterprise:
ibm

Trap type:
Specific <Last field of NOTIFICATION-TYPE OID>

Trap name:
<Label of NOTIFICATION-TYPE>

Display the trap category as:
Status Events

With severity:
Critical

From this source:
Load MIB

Object status for specific traps:
Critical/Down

Event description:
<Event> condition critical

Run this command when the trap is received:
“”

Run as: Hidden Application

The MIB file describing the traps is named umsevent.mib and is installed in the %NV_DRIVE%\USR\OV\SNMP_MIBS directory. Agent systems must have SNMP support installed and the UMSSnmp Object Property before they can forward IBM Director Agent SNMP traps.

Additionally, the installation program adds placeholders for the following traps, which will be implemented by IBM Director Agent in the future:

- iBMPSG_ProcessorEvent
- iBMPSG_AssetEvent
- iBMPSG_POSTEvent
- iBMPSG_ConfigChangeEvent
- iBMPSG_LeaseExpiration
- iBMPSG_WarrantyExpiration

Note: SNMP on agent must be set to forward traps to the server.

Alert on LAN 2.0 Traps

The Alert on LAN 2.0 networking hardware that is present on certain IBM systems, such as the IBM PC300® PL, also has the ability to send alerts when it detects abnormal environmental conditions or system tampering. These alerts are sent to the AOL 2 Proxy tool that can be installed from the IBM Director Agent install program by clicking **Workgroup/Enterprise Integration** and then **AOL Proxy**. (For more information, see "Installing Intel Alert on LAN Proxy" on page 528.)

After the tool is installed, the administrator can configure AOL 2 agents to forward their alerts to the system with the AOL 2 Proxy tool. After the alerts reach the AOL 2 Proxy system, they are converted to SNMP traps and can be forwarded to the NetView Event Browser.

The traps that are forwarded by AOL 2 Proxy are defined in the **aolntrap.mib** and **aolnpet.mib** files. The files are loaded into NetView SNMP MIB loader when the IBM Director Agent upward integration support for NetView is installed. Both MIB files describe the same traps except in different formats. Therefore, it is recommended that one of the MIB files should be unloaded. Specifically, **aolntrap.mib** defines traps whose packets are formatted in the traditional way, and **aolnpet.mib** defines traps whose packets are formatted in a way that complies with the newer Intel pET standard.

The traps that are defined in both MIB files are added to **trapd.conf**. To view the traps, use the following procedure:

1. From the NetView console, click **Trap Settings**.
2. For aolntrap.mib traps, select **Enterprise Intel, ID 1.3.6.1.4.1.343**.

   For aolnpet.mib traps, select **Enterprise Intel, ID 1.3.6.1.4.1.3183**.
3. Click **Event Details**.

Director Traps

During the installation of the IBM Director Agent upward integration support for NetView, a trap filter for IBM Director is added to **trapd.conf**. This filter enables administrators to view SNMP traps that are forwarded from the IBM Director Management Server and that were received from a IBM Director agent. IBM Director supports only one trap type.

To view the details of the trap, including the description, severity, and origin, use the following procedure:

1. From the NetView console, click **Trap Settings**.
2. Select **Enterprise IBM, ID 1.3.6.1.4.1.2.6.146**.
3. Click **Event Details**.

MIB Browsing

By default, the IBM Director Agent upward integration support for NetView installation program loads the IBM Director Agent MIBs using the **loadmib.exe** utility. These MIBs comply with the SMIv1 standard, and therefore the SNMP MIB Browser must be used when browsing IBM Director Agent systems. In addition, target systems must be installed with the IBM Director Agent SNMP Support installation option. See Appendix I, "Installing IBM Director Agent," on page 431 for more information.

To uninstall the IBM Director Agent MIBs, click **Start→Programs→NetView→IBM Director Agent**. Run the uninstallation program. NetView 5.1.1-6.0 must unload SNMP version 2 MIBs.

CA Unicenter TNG Integration

You can use CA Unicenter to manage agent systems that have IBM Director Agent. The IBM Director Agent integration for Unicenter TNG supports the generation of custom MIF files for the AimIT or Asset Management Option repositories. You can also discover IBM agent systems running IBM Director Agent, start IBM Director Agent, and create an IBM Director Agent software distribution package for installing on agent systems. Additional information is available in the README included with the installation.

Configuring CA Unicenter TNG Framework

To receive SNMP traps from IBM Director Agent, you first must configure the SNMP trap server to receive IBM Director Agent alerts (critical only) from IBM computers that have IBM Director Agent installed. To activate the SNMP trap service, perform the following steps:

1. Click **Start** → **Programs** → **Unicenter TNG Enterprise Management** . Then click **Enterprise Managers**. A window with a **Windows NT** computer icon opens.
2. Double-click the **Windows NT** icon.
3. Double-click the **Configuration** icon. The **Settings** window opens.
4. Double-click the **Settings** icon. The **Settings** window, with tabs, a table, and a list of items, opens.
5. Click the **Component Activation Flag** at the bottom, click the **Client Preferences** tab on the right, and then scroll to the row with **SNMP Trap Server Activated** in the **Description** field.
6. In the **Settings** column, the value should be **YES**. If it is not, select the **Setting** check box. Click the **YES** option.
7. Click the **Server Preferences** tab, and scroll to the row with **SNMP Trap Server Activated** in the **Description** column.
8. In the **Settings** column, the value should be **YES**. If it is not, select the **Setting** check box. Click the **YES** option.

9. Set the Windows NT SNMP Trap service to **Manual** by using the following procedure:
   a. Click **Start** → **Settings** → **Control Panel.**
   b. Double-click **Services**.
   c. Double-click **SNMP Trap Service**.
   d. Set the startup type to **Manual**.
10. Start all CA Unicenter TNG Enterprise Management services. Open a command prompt window, and type:

    ```
    UNICNTRL START ALL
    ```
11. Close any Unicenter TNG applications that are running.
12. If you have not restarted the computer since the CA Unicenter TNG was installed, do so now.

Installing the CA Unicenter TNG Upward Integration Module

To install the CA Unicenter Upward Integration Module, use the following procedure:

1. Start the IBM Director Agent installation program on the CA Unicenter server. The Director window opens.
2. Click **Install Director**.

   The **Welcome** window opens.
3. Click **Next**.

   The **License Agreement** window opens. Click **Accept** to proceed.

   Note: You must agree to the terms of the license agreement to install IBM Director Agent. If you click **Decline**, the installation program will close.

   The Select Component window opens.
4. From the **Setup Options** window, click the **Workgroup/Enterprise Integration — Installs IBM Director Agent integration for management environments** button. The Integration Selection window opens.
5. Click **CA Unicenter TNG Upward Integration,** and then click **Next**. The installation program adds the IBM Director Agent UIM for CA Unicenter to the Unicenter server. A command window opens.
6. A notification is displayed and informs you that the installation is complete. Click **OK** to close this notification window.
7. You must restart the system to activate the CA Insinuator TNG UIM.
8. The `restarts`
9. command should be run from a command prompt to activate the integration.

Uninstalling the CA Unicenter TNG Upward Integration Module

To remove the IBM Director Agent UIM from the Unicenter TNG server, click **Start →Programs→ IBM Integration with Unicenter TNG→ IBM Integration with Unicenter TNG →Uninstall IBM Integration with Unicenter TNG**.

Intel LANDesk Management Suite Integration

Unlike other IBM Director Agent Upward Integration Modules, LANDesk Management Suite integration requires that an additional component be installed on each IBM Director Agent system. From the Director Agent Configuration window, click **Agent UIMs**. Click Next and select **LanDesk Management Suite.**

Important: Do *not* select Web Based Remote Control when installing IBM Director Agent on systems that you will manage using LANDesk Management Suite. LANDesk Management Suite includes a remote control service that is not compatible with the Web Based Remote Control service that is included with IBM Director Agent.

For more information on supported system-mangement, see 505 and Appendix I, "Installing IBM Director Agent," on page 431.

Note: You can install IBM Director Agent on systems with the LANDesk integration option enabled, even if you have not yet installed LANDesk Management Suite on your network.

You do not need to install additional software to your LANDesk Management Suite administration system to manage IBM Director Agent systems that have the LANDesk Management Suite component installed. To integrate IBM Director Agent systems into your LANDesk Management Suite environment, configure the batch file ldinv.bat to run periodically on each IBM Director Agent system. The ldinv.bat file generates custom MIF files that can be used by the LANDesk Management Suite inventory functions. Be sure to run **ldinv.bat** before **ldiscn32.exe** inventory collection. One method that you can use to accomplish this task is described in the following procedure:

1. From a login script, run **ldinv.bat**, and then run **ldiscn32.exe**.
2. Create a login script that connects the system to the \LDLOGON share of the LDMS server.
3. Copy the contents of the **%UMS_HOME%\inventory\ldinv.bat** file into the script.
4. Remove the comment symbols from the command line for **ldiscn32.exe** that is included in the batch file and configure it with the name of the inventory server and its network address and any other desired settings. With this line active, whenever a user logs in, the login script generates an MIF file, outputs it to **c:\dmi\dos\mifs** (by default), and trigger an inventory scan that will update the LDMS inventory database.

5. Use the LANDesk scheduler to run **ldinv.bat** on each agent at a predefined time or run **ldinv.bat** from the **Startup** folder of each agent system. Ensure that **ldiscn32.exe** runs from **ldinv.bat** and that no other copies of **ldiscn32.exe** are run from the **Startup** folder

Microsoft SMS integration

Microsoft SMS can be used to manage agent systems running IBM Director Agent. To enable it, use the IBM Director Agent installation program to install the IBM Director Agent Microsoft SMS Upward Integration Module on the Microsoft SMS server.

Note: After you have installed the Microsoft SMS 2.0 Server or Console, you can use the SMS Software Distribution function to distribute IBM Director Agent to your SMS 2.0 agent systems. A special installation program designed to facilitate this process, named **umsw32un.exe**, can be downloaded from the Web at

http://www.pc.ibm.com/ww/software/applications/ums

Installing the Microsoft SMS Upward Integration Module

During installation of the UIM on the SMS 2.0 server, the SMS console is configured with the queries, collections, and tools specific to IBM Director Agent. Also, the installation provides a Microsoft Management Console (MMC) snap-in module that adds a context to the agent systems. The IBM Director Agent specific menu items appear only on systems that have IBM Director Agent installed.

1. Start the IBM Director Agent installation program on the Microsoft SMS Console or Server.
2. Insert the Director with IBM Director Agent CD into the CD-ROM drive. The Director with IBM Director Extensions window opens.
3. Click **Next**.

   The **License Agreement** window opens. Click **Accept** to proceed.

   Note: You must agree to the terms of the License Agreement to install IBM Director Agent. If you click **Decline**, the installation program will close.

   The Select Components window opens.
4. From the Select Components window, click **Workgroup/Enterprise Integration - Installs IBM Director Agent integration for management environments.**

   The Integration Selection window opens.
5. Click **Microsoft SMS Upward Integration** and then click **Next**.

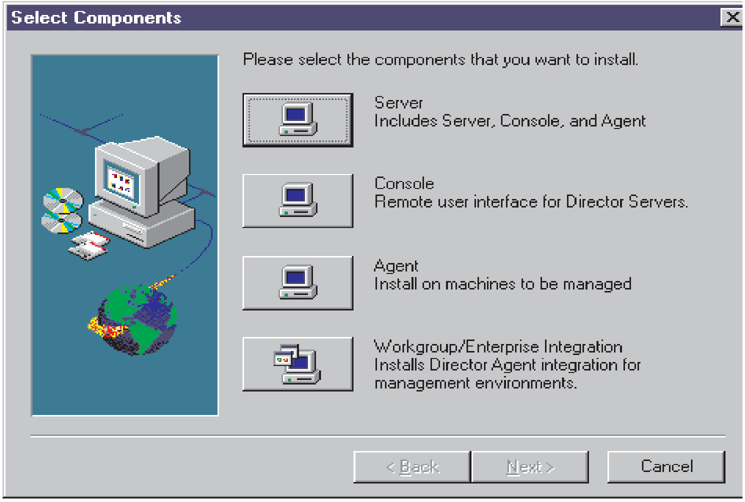

Select Components

Please select the components that you want to install.

Server
Includes Server, Console, and Agent

Console
Remote user interface for Director Servers.

Agent
Install on machines to be managed

Workgroup/Enterprise Integration
Installs Director Agent integration for management environments.

6. Select the version of Microsoft SMS for which you are installing support, and then click **Next**. The installation program adds the IBM Director Agent UIM to your Microsoft SMS system.

Customizing the SMS Installation

Use the following procedure to modify the IBM Director Agent installation routine and to create a new IBM Director Agent installation executable file. This procedure requires a compatible compression utility program, such as WinZip, to decompress the original file and to create a new executable file with preset command-line instructions.

Note: To customize the IBM Director Agent installation file (**UMSW32UN.EXE**) properly, the compression utility must meet the following basic requirements:

- Has the ability to extract files into a temporary directory
- Enables the user to edit and delete compressed files
- Contains a self-extracting process that creates a user-defined installation executable file, with the ability to set up specific startup commands

To create a new IBM Director Agent installation file, use the following procedure:

1. Use Windows Explorer to locate the file **UMSW32UN.EXE**.
2. Using the compatible compression utility program, extract the component files of **UMSW32UN.EXE** into a temporary directory.
3. Edit the extracted file, **UIMSETUP.ISS**, to select the installation components to install for your particular configuration.

See “Modifying the UIMSETUP.ISS File Manually” on page 523 for more information.

4. To reduce the size of the new package, delete **Tivoli.z** or **Netfin.z** from the temporary directory if they are not required for downloads.
5. Select all the extracted files, including **UIMSETUP.ISS**, and create a new compressed file with the file name **UMSW32UN**.
6. Open your self-extracting executable file program using the compressed **UMSW32UN** file.
7. Label the new file **UMSW32UN.EXE**.
8. In the **Command To Issue** field (after the self-extraction operation is completed), type

   ```
   en/um_setup.exe REBOOT-S-SMS
   ```

   Note: To have the agent system restart after installation type

   ```
   en/um_setup.exe REBOOT-S-SMS
   ```
9. In the **Wait For** field, type

   ```
   _ISDEL
   ```
10. Click **OK** to save the new file.

A customized IBM Director Agent installation executable file is created.

Modifying the UIMSETUP.ISS File Manually

This section describes the contents of the IBM Director Agent response file, UIMSETUP.ISS, which can be used for the silent installation of IBM Director Agent on agent systems. You can use this response file for software distribution from the Upward Integration Module environments and other silent installation scenarios.

Note: You must install the UIM before using this process.

The response file is a text file that includes a number of variables that specify installation selections, such as which components are installed or the drive and directory to which the program files will be copied, that would ordinarily be selected during an attended installation. Some portions of the response file must not be changed by the user; making changes to these sections will cause the installation program to fail. All portions of the UIMSETUP.ISS file, including sections that the user should not change, are described in the following pages.

This first four sections of the UIMSETUP.ISS response file provide information about the installation process to the installation program. These entries must not be changed by the user:

```
[InstallShield Silent]
Version=v3.00.000
File=Response File

[Application]
Name=UMS
Version=3.1
Company=IBM

[DlgOrder]
Dlg1=SdAskOptions-0
Dlg2=AskDestPath-0
Dlg3=AskSecurInfo-0
Count=3

[SdOptionsButtons-0]
Result=103

[SdOptionsButtons-1]

Result = 101
```

The next section of the UIMSETUP.ISS response file *can* be customized by the user. This section determines the IBM Director Agent components that will be installed on the agent system.

To change these options, change the value of `Component-count` to the total number of components that you want to install, and list the components in sequential order, starting at `Component-0`.

```
UMS detected on system. Would you like to upgrade? Yes or No.

1=Yes

0=No

[UpgradeYesNo]

Result=1
```

The following is the list of components that you can install. If your UIMSETUP.ISS file contains all of these examples as shown, all selectable IBM Director Agent components will be installed on the agent system. Include only the components that you want to install. *X* is the total number of components to be installed.

```
Component-x=0 Director Support

Component-x=1 Web Based Access

Component-x=2 System Health & Monitoring

Component-x=3 Web Based Remote Control
```

```
Component-x=4 LANDesk (TM) Management Suite Integration

Component-x=5 Tivoli Management Agent

Component-x=6 SNMP access and trap forwarding

Component-x=7 Help Files
```

Add components below if desired. Recommended defaults are already set below:

```
[SdAskOptions-0}


Component-type=string


Component-count=4


Component-0=0


Component-1=1


Component-2=2


Component-3=6



Component-4=7


Result=1
```

The next section of the response file displays the path where you want to install IBM Director Agent. The default path is shown. You can change the installation path if necessary.

```
[AskDestPath-0]
szPath="C:\Program Files\IBM\UMS"
Result=1
```

The next section of the response file configures the security information. In this section of the file, type your user ID and password; both items are case-sensitive.

Type your password again in the `svConfirm` line. You can also use this section to specify the TCPIP port number that will be used by IBM Director Agent.

```
[AskSecurInfo-0]
svUser=ums

svPassword=ums

svConfirm=ums

svPort=411



Result=1
```

The default port number is 411. You can change this default port if necessary. Other valid port numbers are 6411, 6500, 6600, and 6611.

The last two sections of the SETUP.ISS response file provide information about the installation process to the installation program. You must not change the following entries.

```
Do you want icons on the start menu?

[icons]

Result=0

If you installed Director support, would you like Director Remote
Control?

[NFDRemote]

Result=1

Do you want to require authorization for Director remote Control?

[NfDreqAuth]

Result=1

AutoReboot machine when install is finished?

[AutoReboot]

Result=0

Terminal services installation question.

Setup could not detect if system is in install mode. Are you sure the
system is in install mode?

Continue with install.
```

```
            1=YES

            0=NO

[TerminalServices}

Result=0

Use IIS as web server for UMS?

Setup detected IIS installation. Do you want to use IIS as web server
for UMS?

           1=YES

           0=NO

[Use IIS]

Result=o
```

Using Microsoft SMS to View Agent System Inventory

You can use Microsoft SMS2.0 to view IBM Director Agent inventory data.

Microsoft SMS 2.0

The Microsoft SMS 2.0 UIM enables the SMS server to gather inventory data directly from the CIM agent on agent systems running Windows 98 or Windows NT.

The Microsoft SMS 2.0 UIM extends the SMS 2.0 **Collections** tree in the SMS 2.0 console to include IBM Director Agent agents. The SMS 2.0 UIM also extends the **Queries** tree, so that it can retrieve IBM Director Agent-specific inventory data. The **Tools** tree is extended as well, so that you can start the IBM Director Agent console on a agent system.

SMS 2.0 does not support SNMP trap listening. However, the Microsoft SMS 2.0 UIM translates CIM notifications that are generated by IBM Director Agent into SMS 2.0 status messages.

To view the IBM Director Agent inventory data from the SMS 2.0 Console, click the **Collections** tree, and then use following procedure:

1. Click **All Systems with IBM Director Agent**
2. Right-click a agent system in the window on the right side.
3. Click **All Tasks** → **Start Resource Explorer**.
4. Click the **Hardware** node. The IBM Director Agent inventory data is under **IBM Director Agent**.

Installing Intel Alert on LAN Proxy

To install the Intel Alert on LAN Proxy on your system-management system, use the following procedure:

1. Start the IBM Director Agent installation program on the system-management system that will receive Alert on LAN messages.
2. Insert the Director with IBM Director Extensions CD in the CD-ROM drive. The Director with IBM Director Extensions window opens.
3. Click **Install Director**.
4. Click **Next**.

   The **License Agreement** window opens. Click **Accept** to proceed.

   Note: You must agree to the terms of the License Agreement to install IBM Director Agent. If you click **Decline**, the installation program closes.

   The Select Components window opens.

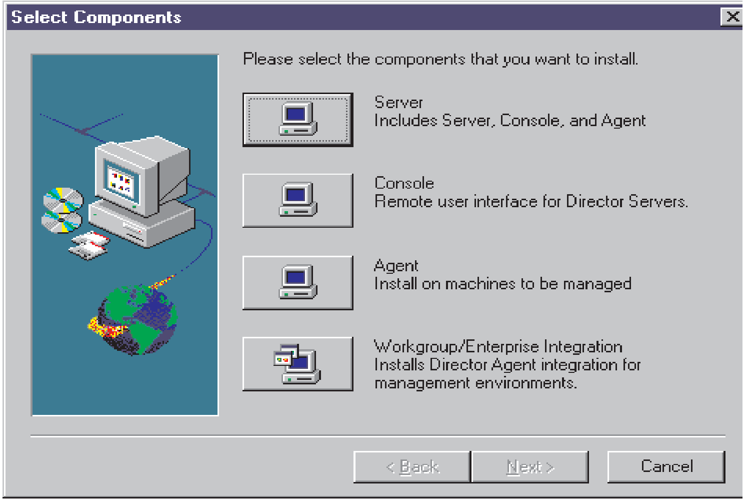

Select Components

Please select the components that you want to install.

Server
Includes Server, Console, and Agent

Console
Remote user interface for Director Servers.

Agent
Install on machines to be managed

Workgroup/Enterprise Integration
Installs Director Agent integration for management environments.

5. From the Select Components window, click the **Workgroup/Enterprise Integration — Installs IBM Director Agent integration for management environments** button.

   The Integration Selection window opens.
6. Click **Intel Alert on LAN Proxy** and then click **Next**.

7. Select a TCPIP port for use by the Alert on LAN Proxy. Then, click **Next** to finish the installation process.

HP OpenView Integration Module

IBM Director Agent provides the following integration with HP OpenView Network Node Manager:

- Visibility of IBM Director Agent Inventory data from the OpenView Console
- Availability of a current IBM Director Agent submap
- Accessibility of IBM Director Agent from the OpenView Console

Installing IBM Director Agent Support on the OpenView Server

From the **Enterprise** option, the installation program automatically installs the files necessary for IBM Director Agent installation for OpenView.

1. Insert the Director with IBM Director Agent CD into the CD-ROM drive. The Director with IBM Director Extensions window opens.
2. Click **Install Director**. The Welcome window opens.
3. Click **Add New Program**→ **IBM Director Agent**.
4. Click **Next**. The License Agreement window opens. Click **Accept** to proceed. You must agree to the terms of the License Agreement to install IBM Director Agent. If you click **Decline**, the installation program will close.

   The Select Components window opens.

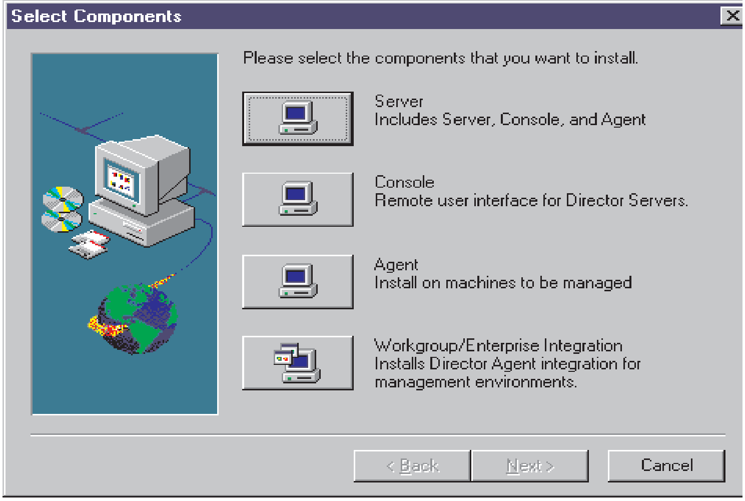
Select Components
Please select the components that you want to install.
Server
Includes Server, Console, and Agent
Console
Remote user interface for Director Servers.
Agent
Install on machines to be managed
Workgroup/Enterprise Integration
Installs Director Agent integration for management environments.
< Back
Next >
Cancel

5. From the Select Components window, click the **Workgroup/Enterprise Integration** to install IBM Director Agent integration for management environments. The Integration Selection window opens.

6. Click **HP Openview Integration** and then click **Next**. The installation program installs the files necessary for IBM Director Agent for HP OpenView.

The following events are associated with IBM Director Agent installation to Openview:

1. Install **\openview\snmp_mibs\*.mib** (SNMP MIB files for IBM Director Agent systems).
2. Add lines to **\openview\alerts\umstraps.conf** (configures IBM Director Agent events in to HP OpenView)
3. Install **\openview\registration\c\ums.reg** (adds map files).
4. Install **\openview\registration\c\umsinv.reg.umsinv62.reg** (adds tools menu options)
5. Install\openview\registration\c\umsdb.reg(adds database support for IBM Director Agent)
6. Install **\openview\registratoin\c\umstrap.reg** (loads IBM Director Agent events into the event configuration window)
7. Install **\openview\bin\*** (adds map support for IBM Director Agent)
8. Install **\openview\bitmaps\c\universal_service\*** (bitmaps for IBM Director Agent icons in OpenView Console)
9. Install **\openview\fields\c\universal_service** (fields for IBM Director Agent in OpenView Console

Accessing IBM Director Agent from the OpenView Console

You can access IBM Director Agent from the OpenView Server by using either Microsoft Internet Explorer (Version 4.01 or later) or Netscape (Version 4.51 or later).

1. Select an agent node in the OpenView Console.
2. Click **Tools →Director Agent →Director Agent Browser**.
3. Click **IBM Director Agent Browser**.

Viewing IBM Director Agent Inventory Data from the OpenView Console

To view IBM Director Agent Inventory Data from the HP OpenView Console:

1. Select a agent node in the OpenView Console.
2. Click **Tools→ Director Agent→ Director Agent Inventory.**
3. Click an option under **IBM Director Agent Inventory**.

Adding IBM Director Agent to the IBM Director Agent Submap

When the **hpovums.exe** program detects agent systems that have IBM Director Agent installed.

- IBM Director Agent system become members of the IBM Director Agent Submap.
- The **OpenView Console Tools** menu is expanded to include items enabling access IBM Director Agent inventory options.

Populating the IBM Director Agent Submap

The IBM Director Agent upward integration module for OpenView defines three attributes that **hpovums.exe** checks before creating IBM Director Agent: isUniversalService, isUmServicesCim, and LaunchHttpUmServices.

isUniversalService

An SNMP Get() operation is performed against the Enterprise OID 1.3.6.1.4.1.2.6.159 on a remote system. This OID is owned by IBM Director Agent, and a valid return value will add the IBM Director Agent Snmp Object Capability to the system properties and add the system to the Director Agent Submap.

isUMServicesCim

A WMI ConnectServer() call is performed against a remote system **winmgt.exe** process in the **root\cim\2 namespace**. If the user who is invoking the connection has remote access privileges on the system, the IBM Director Agent HTTP port is queried, the UMServicesCim Object Capability is added to the system properties, and the system is added to the Director Agent Submap. The isUMServicesCim test requires WMI on the agent system (installed automatically by IBM Director Agent) and remote access privileges for the OpenView administrator. Note that a system must have the UMServicesCim capability to retrieve IBM Director Agent-specific inventory information, because Windows Management instrumentation supplies this data.

Launch HttpIBM Director Agent

After the agent systems with IBM Director Agent are found, **hpovums.exe** checks whether each agent system supports starting of Web-based IBM Director Agent by sending an SNMP request on the HTTP port. If a valid HTTP port is returned, this attribute is enabled; if not, this attribute is disabled.

Based on this attribute, **Tools** menu option **Director Agent Browser**, which is used to start HTTP-based IBM Director Agent, is enabled or disabled.

Forwarding IBM Director Agent Events

IBM Director Agent forwards the following SNMP traps to the OpenView server workstation:

- iBMPSG_TemperatureEvent
- iBMPSG_VoltageEvent
- iBMPSG_ChassisEvent

- iBMPSG_FanEvent
- iBMPSG_StorageEvent
- iBMPSG_SMARTEVENT and
- iBMPSG_LANLeashEvent
- iBMPSGRedundantNetworkAdapterEvent
- iBMPSGRedundantNetworkAdapterSwitchoverEvent
- iBMPSGRedundantNetworkAdapterSwitchbackEvent

These traps alert the OpenView administrator to critical environmental conditions in IBM systems, low disk space, a failing hard disk drive, and a system being removed from a LAN. During the installation of the IBM Director Agent upward integration support for OpenView, these traps are added to OpenView **trapd.conf** file with their trap properties configured using the **addtrap.exe** utility. In addition, the installation program adds placeholders for the following traps, which will be implemented by future versions of IBM Director Agent:

- iBMPSG_ProcessorEvent
- iBMPSG-AssetEvent
- iBMPSG_POSTEvent
- iBMPSG_ConfigChangeEvent
- iBMPSG_LeaseExpiration
- iBMPSG_WarrantyExpiration

The following list describes the default properties that are configured for IBM Director Agent SNMP traps:

- Enterprise: ibm
- Trap-Type: Specific <Last field of NOTIFICATION - TYPE OID>
- Trap Name: <Label of NoTIFICATION - Type>
- Display the Trap Category as: Status Events
- With Severity: Critical
- From this Source: Load MIB
- Object Status for Specific Traps: Critical/Down
- Event Description: <Event> condition critical
- Run this command when the trap is received: " "
- Run as: Hidden Application

The MIB file describing the traps is named **umsevent.mib** and installed in the **%OV_DRIVE%\openview\snmp_mibs** directory. Agent systems must have SNMP support installed before they can forward IBM Director Agent SNMP traps.

Alert on Lan 2 Traps

The Alert on LAN 2 networking hardware which is present on certain IBM systems, such as the PC300 PL, also has the ability to send alerts when it detects abnormal environment conditions or system tampering. These alerts are sent to the AOL 2 Proxy tool that can be installed from the IBM Director Agent installation program by selecting **Workgroup/Enterprise Integration** and **AOL Proxy**. After the tool is installed, the administrator can configure AOL 2 agent systems to forward their alerts to the system with the AOL 2 Proxy tool. After the alerts reach the AOL 2 Proxy system, they are converted to SNMP traps and can be forwarded to OpenView Event Browser.

The traps that are forwarded by AOL 2 Proxy are defined in the **aoltrap.mib** and **aolpet.mib** files, which are loaded into OpenView SNMP MIB loader when the IBM Director Agent upward integration support for OpenView is installed. Because the MIB files describe the same traps, just in different formats, one of the MIB files should be unloaded. Specifically, **aolntrap.mib** defines traps whose packets are formatted in the traditional way, and **aoimpet.mib** defines traps whose packets are formatted in a way that complies with the newer pET standard from Intel.

The traps that are defined in both MIB files are added to **trapd.conf** and can be viewed by opening the Trap Settings window and scrolling to **Enterprise intel, ID 1.3.6.1.4.1.343** of the **aointrap.mib** traps and **ID 1.3.6.1.4.1.3183** for the **aoinpet.mib** traps.

Director Traps

During the installation of the IBM Director Agent upward integration support for OpenView, a trap filter for the Director product is added to **trap.conf**. This filter enables administrators to view SNMP traps forwarded from IBM Director management server on behalf of a IBM Director agent. IBM Director supports only one trap type, so to view the details of the trap, including the description, severity, and origin, select the trap and open its Event Details. The trap settings for the Director_Trap type can be viewed by opening the **Trap Settings** window and scrolling to **Enterprise ibm, ID 1.3.6.1.4.1.2.6.146**.

MIB Browsing

By default, the IBM Director Agent upward integration support for the HP OpenView Network Node Managed install program loads Management Information Bases (MIBs) using the **loadmib.exe** utility. These MIBs comply with the SMlv1 standard, and therefore, the SNMP MIB Browser must be used when browsing IBM Director Agent systems. In addition, target systems must be installed with the IBM Director Agent SNMP Support installation option. See Chapter 4 of the *IBM Director Agent User Guide* for more information.

Uninstalling OpenView Integration Support

To Uninstall OpenView integration support, from the HP Openview console click **Start→Programs→HP OpenView→UMServices Integration** and run Uninstall IBM Director Agent. You must unload SNMPVS MIBs.

Director Agent Database Support for NetView and OpenVIew

Director Agent supports Saving Inventory to Database. The database that the Director Agent supports is Microsoft SQL Server, version 7.0. Before you create a MS SQL Server database, review SQL Server Authentication levels with your Database Administrator. Use the following procedure to Save Inventory to Database:

From the main menu,

1. click **Tools →Director Agent →Get Director Agent Nodes**.
   This option sends all Director Agents to a file called ums_nodes. The file is created under the BIN directory of HP OpenView and NetView. When the option is successful, you will receive the `All Director Agents are collected and written to a ums_nodes file` message. You can edit this file manually to add or delete unwanted Director Agents. However, the file format must be maintained.
2. Select **Tools →Director Agent →Configure ODBC DataSource**.

   The Configure DataSource for Database Access window opens.
3. From the window, configure the following settings:

   ODBC Data Source Name:
   Enter the ODBC data source name.

   Server Name:
   Enter the host name of the server where the database resides.

   Database Name:
   Enter the database name. If one does not exist, it is created.

   User Name:
   Enter the valid user name for trusted connection.

   Password:
   Enter the password for the user name.
4. Click **Add**. If the configuration was successful, a data source and an inventory table is created. If not successful, a data source is created with no inventory tables. The status is displayed in the status field of the dialog box.
5. To delete the Data Source, enter all of the data in the dialog box and click **Delete**. To clear all of the data, click **Clear**. To close the dialog box, click **Close**.
6. From the main menu, select **Tools →Director Agent →Save Inventory to Database**.
   The Inventory Status window opens. The status of **Inventory Saving to Database** is shown. Your selected nodes and the inventory tables are shown. The dialog box contains the following fields:

Name or Address:

Director Agent Name or Address.

Inventory Tables:

All inventory tables used for saving inventory.

Messages:

Status of each Director Agent.

7. Click **Close**, after you receive the `Inventory is saved to Database` message.

Appendix L. Notices

This publication was developed for products and services offered in the U.S.A.

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

Edition notice

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

- Alert on LAN
- Chipkill
- EtherJet
- e-business logo
- HelpCenter
- HelpWare
- IBM
- Light Path Diagnostics
- NetBAY
- NetView
- OS/2 WARP
- Predictive Failure Analysis
- ServeRAID
- ServerGuide
- ServerProven
- TechConnect
- Tivoli
- Tivoli Enterprise
- Update Connector
- Wake on LAN
- xSeries

Lotus and Domino are trademarks of Lotus Development Corporation and/or IBM Corporation in the United States, other countries, or both.

Intel, Celeron, MMX, LANDesk, Pentium, Pentium II Xeon, Pentium III Xeon, and Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Linux is a registered trademark of Linus Torvalds.

Other company, product, or service names may be the trademarks or service marks of others.

Important notes

Processor speeds indicate the internal clock speed of the microprocessor; other factors also affect application performance.

CD-ROM drive speeds list the variable read rate. Actual speeds vary and are often less than the maximum possible.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1000000 bytes, and GB stands for approximately 1000000000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1000000 bytes, and GB stands for 1000000000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Maximum memory may require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven®, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

Unless otherwise stated, IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Index

File Number:

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.