

IBM Director User's Guide

Version 2.2

January 2001

Contents

Preface	xix
1.Introduction	1
How IBM Director Works	2
What Is New in This Release.....	3
Expanded Managed System Support.....	3
Expanded Database Support.....	3
Expanded Transport Support.....	3
Expanded Associations Function	4
Expanded Support for Groups.....	4
Event Management Enhancements	5
Inventory Enhancements	5
Resource Monitoring Enhancements	5
Process Management Enhancements	6
Remote Control Enhancements.....	6
File Transfer Enhancements.....	6
Task Scheduler Enhancements.....	7
Removed Multiplatform Manager (MPM) Support.....	7
Supported Communication Protocols	8
Modem Connections	8
Managing Native Systems	8
Managing DMI-Enabled and CIM-Enabled Native Systems	10
Managing Cluster-Enabled Native Systems	11
Managing SNMP Devices with IBM Director.....	11
Additional Features in IBM Director.....	12
2.Planning	13
IBM Director Server	13
Management Server Prerequisites	13
Database Support.....	13

Planning to Use the Jet Database.....	14
Relocating the Jet Database.....	14
Planning to Use the DB2 Universal Database.....	15
Setting up Trusted Connections.....	15
DB2 Server Login Access for Windows NT.....	15
Creating the DB2 Server Database.....	16
Planning to Use the SQL Server.....	17
Setting up Trusted Connections.....	17
SQL Server Login Access for IBM Director.....	17
Creating the SQL Server Database.....	17
Planning to Use the Oracle Server Database.....	19
Oracle Server Login Access.....	19
Configuring the Oracle TCP/IP Listener.....	19
Using Unlimited Rollback Segments (Oracle Server Version 7.3.4 Only)	20
Setting the Compatible Parameter (Oracle Server Version 7.3.4 Only)	20
Creating the Oracle Server Database.....	20
Additional Considerations.....	21
IBM Director Management Console.....	23
Tivoli Management Agent.....	23
Operating System Platform Support.....	24
Transport Support.....	24
Discovery.....	26
License Management.....	27
Migration Support.....	28
Security.....	29
IBM Director User-Logon Security.....	29
IBM Director Console Authorization.....	30
IBM Director Accounts.....	30
Listing IBM Director Users.....	31
Creating a New User.....	31

Editing User Accounts	33
Changing User Defaults	33
Changing User Passwords	34
Deleting User Accounts	34
Investment Protection and Integration	34
Java Classes	34
Planning for IBM Director Tasks	35
Software Distribution	35
Redirected and Streamed Installations	35
Distributing Packages Using Redirection	36
Distributing Packages Using Streaming	36
Memory and Storage Management for Redirected Installations	37
Determining Which Share Is Used on Redirected Installs	38
Always Streaming Software to the Managed System.....	39
Specifying the Transport for Server Shares	39
Limitations on Software Distribution	40
Configuring Security for UNC-based Server Shares	41
Limiting Network Resources for Software Distribution	42
Remote Control	43
Event Management.....	45
CIM Event Support.....	46
SNMP Trap Support	46
Inventory Management	47
3.Installation and Configuration	49
Before You Begin	49
Supported Systems by Component	50
Hardware Requirements	51
Workgroup/Enterprise Integration	51
Installing IBM Director 2.2 Instructions.....	52
Installing IBM Director Server	53
Installing the IBM Director Console.....	64

Installing the UM Services Client for IBM Director	65
Supported Operating Systems	65
Hardware Requirements.....	66
Installing the Windows UM Services Client.....	66
Using the UM Services Client Attended Installation for OS/2.....	71
Using the UM Services Client Unattended Installation for OS/2	73
Installing the UM Services Client for NetWare.....	74
Installing UM Services Client for SCO UnixWare.....	75
Installing UM Services Client for Linux Red Hat	76
Reserving Agent Disk Space During Software Distribution (Windows and Unix)	76
Installing Oracle Server or DB2 Universal Databases Using the Command Line (Unix)	78
Installing the DB2 Universal Database	78
Installing the Oracle Server.....	78
Defining Server Preferences for Database Properties	79
Configuring IBM Director to Use File Distribution Servers	79
Enabling UNC-based Share Access to the IBM Director Server.....	80
Enabling UNC-based Share Access to Managed Systems.....	80
Enabling UNC-based Share Access to Windows Managed Systems.	80
Defining Server Preferences	81
Configuring Distribution Preferences for Managed Systems	82
Defining the Maximum Number of Concurrent Redirected Distributions .	83
Defining the Maximum Number of Concurrent Streamed Distributions ..	84
Limiting the Bandwidth for Streamed Distributions.....	84
Restricting Access Check.....	85
Specifying Do Not Stream Distribution if Redirected Distribution Fails...	85
Defining the Automatic Time-out for Remote Control Sessions.....	85
Changing the Network Transport	85
Saving, Restoring, and Resetting Program Files (Unix only)	86
Uninstalling IBM Director.....	87
Uninstalling IBM Director Components on OS/2	88

Uninstalling IBM Director Components on NetWare	88
4.Upgrading to IBM Director 2.2.	89
Upgrading the IBM Director Server	89
Upgrading the IBM Director Console.....	91
Upgrading the IBM Director Client.....	92
5.Using the Management Console.	95
Managed Systems	95
Starting the Management Console	96
Getting Around in IBM Director	97
Using Drag and Drop	98
Using Your Mouse’s Double-Click Function	99
Using Context Menus.....	100
Using Add and Remove Buttons	100
Managing Columns of Information.....	100
Monitoring the Task in Process.....	101
Using Keyboard Arrow Keys.....	101
Saving Files	101
Using the Management Console	101
Group Contents	101
Associations.....	102
Groups	103
Dynamic and Static Groups.....	103
Creating a Dynamic Group.....	104
Creating a Static Group	105
Group Category Editor	105
Task Based Group Editor	106
Group Export/Import	106
Managing Your Groups	106
Tasks.....	106
Additional Management Console Features	108

Using the Menu Bar	108
Using the Toolbar	108
Using the Status Bar	109
Using the Ticker Tape	109

6.Inventory Management 111

Performing an Inventory Collection	111
Using the Inventory Query Browser	112
Additional Inventory Query Browser Features	113
Updating the List of Available Queries	113
Managing Your Inventory Query Results	113
Using Menu Bar Options	114
Building a Customized Query	114
Using the Inventory Query Builder	115
Using the Inventory Software Dictionary Editor	115
Managing Your Software Dictionary Entries	117
Performing Batch Operations on the Software Dictionary File	117
Requirements for Using TWGCLI	118
Exporting Entries to a Properties File	118
Command Syntax	118
Importing Entries from a Properties File, Microsoft PDF, or Software Dictionary File	120
Command Syntax	120
Modifying Inventory Collection Preferences	122

7.Remote Control 125

Control States	126
Overriding and Changing Control States	127
Requesting Active Control from a Management Console	127
Changing Control States from the Managed System	127
Control State Scenarios	127
Scenario 1	127

Scenario 2	128
Scenario 3	128
Remote Control Usage Restrictions	128
Remote Access Security	128
Sending Keyboard Information to a Remote System	129
Remote Control and Inventory	129
Type of Operating System.....	129
Code Page for Screen Transfer.....	130
Restrictions on Pointer and Cursor Support.....	130
Performing Remote Control Tasks	130
Starting a Remote Control Session.....	131
Stopping a Remote Control Session.....	131
Changing the Control State of a Session.....	131
Recording a Remote Control Session.....	131
Viewing a Listing of Current Remote Control Sessions	132
Changing the Refresh Rate for Current Remote Control Sessions	132

8.Resource Monitoring 133

Understanding Monitors	133
Monitoring Data on Native Managed Systems	134
Monitoring Data on Native Managed Systems Configured with Additional Services 135	
Monitoring Data on SNMP Devices	135
Monitoring Data on Windows NT Devices and Services	136
Starting Resource Monitors	136
Using the Monitor Console.....	137
Initiating a Resource Monitor	137
Viewing Monitor Data on the Ticker Tape	138
Setting Monitor Thresholds.....	138
Setting Numeric Thresholds	139
Setting Text String Thresholds	140
Recording Monitor Data.....	140

Managing Your Monitored Resources	141
9.Event Management	143
New Terms in This Chapter.....	143
Understanding Event Management.....	144
Creating an Event Action Plan	145
Using Predefined Event Filters	147
Creating an Event Filter	147
Assigning an Event Filter to an Event Action Plan	148
Customizing an Action.....	148
Testing an Action	149
Assigning an Action to an Event Filter	149
Saving an Event Action Plan	150
Activating Event Action Plans	150
Displaying Applied Event Action Plans	150
Performing Maintenance Tasks	150
Managing Event Action Plans	151
Viewing Event Details in the Event Log	151
Viewing All Logged Events.....	152
Viewing Events by Filter Characteristics.....	153
Viewing Events by System	153
Using the Action History Window	153
Generating Your Own Events.....	154
10.Software Distribution	157
Importing a File Package	157
Distributing a File Package.....	157
Scheduled Distributions	158
Immediate Distributions.....	158
Viewing Package Content Information	159
Viewing Distribution History	159
Renaming Packages	159

Viewing Package Audit Activity	160
Deleting a File Package	160
Using File Distribution Servers Manager	160

11.File Transfer 161

Using the File Transfer Task.....	161
Starting a File Transfer Session	162
The Wild Card Feature	162
Selecting Files for Transfer	162
Transferring Files between Managed Systems.....	163
Choosing a New Target.....	164
Synchronizing Files, Directories, or Drives	164
Additional File Transfer Features	165
Precautions when Using File Transfer	165

12.SNMP Management 167

Understanding SNMP Management	167
MIB Requirements for the SNMP Browser	168
MIB Requirements for IBM Director Services	168
Performing SNMP Tasks	168
Understanding SNMP Discovery	168
Setting SNMP Discovery Parameters	170
Creating a New SNMP Device.....	171
Using the SNMP Browser	172
Starting the SNMP Browser	172
Viewing SNMP Information	173
Multi-homed Support.....	174

13.DMI Management 175

DMI Requirements.....	176
Creating a DMI Dynamic Group	176
Performing DMI Browser Tasks.....	177

Starting the DMI Browser and Viewing Information	177
Setting an Attribute Value for a DMI Group	178
Defining DMI Browser Subtasks	179
14.CIM Management.	181
CIM Requirements.....	182
Creating a CIM Dynamic Group	182
Performing CIM Browser Tasks.....	183
Starting the CIM Browser and Viewing Information	183
Setting a Property Value for a CIM Class Instance	184
Executing a Method for a CIM Class Instance	185
Defining CIM Browser Subtasks.....	185
Defining a Browser Subtask for a CIM Class.....	186
Defining a Browser Subtask for a CIM Class Method	186
15.Asset ID	189
The Asset ID Interface.....	190
Serialization	191
System.....	191
User.....	193
Lease	194
Asset	195
Personalization.....	196
Warranty	196
16.Alert on LAN	199
17.Cluster Management.	203
Understanding Cluster Management	203
Cluster Requirements	204
Performing Cluster Browser Tasks.....	204
Understanding Cluster Discovery	204

Starting the Cluster Browser and Viewing Information.....	205
---	-----

18.Process Management 207

Starting the Process Management Task	208
Viewing Application Information.....	209
Viewing Windows NT Services Information.....	211
Executing Commands on Selected Systems.....	211
Creating Non-Interactive Tasks to Execute Commands	211
Restricting Anonymous Command Execution.....	211
For Windows NT systems	212
For Unix systems	212
Closing Applications.....	213
Adding New Process Monitors	213
Controlling NT System and Device Services	213
Removing Process Monitors	214
Adding Service and Device Service Monitors.....	214

19.Task Scheduler. 215

Scheduling Tasks	215
Customizing Your Scheduled Job.....	216
Using the Date/Time Page.....	216
Using the Repeat Window	217
Using the Task Page	218
Using the Targets Page.....	218
Using the Options Page	219
Understanding the Special Execution Options.....	219
Delay execution on unavailable systems	220
When You Do Not Check This Option.....	220
When You Check This Option	220
Execute on systems that are added to the target group.....	221
When You Do Not Check This Option.....	221
When You Check This Option	221

Execute in client time zone.....	221
When You Do Not Check This Option	221
When You Check This Option	221
Saving Your Scheduled Job.....	222
Managing Scheduled Jobs	222
Using the Calendar Pages	222
Viewing Job Properties	223
Viewing Scheduled Jobs	223
Viewing Execution History Logs	224
Using the Jobs Page	224
Viewing Scheduled Job Information	225
20.Troubleshooting	227
A. Resource Monitor Attributes	239
Windows NT Operating System.....	239
CPU Monitors	239
Device and Service Monitors	239
Disk Monitors	239
DMI Monitors	240
File Monitors.....	240
Directory	240
File.....	240
Memory Monitors	241
NT Performance Monitors	241
Registry Monitors	241
TCP/IP Monitors	241
Process Monitors.....	242
Windows 2000 Operating System	242
CIM Monitors	242
CPU Monitors	242
Device and Service Monitors	243

Disk Monitors.....	243
DMI Monitors	243
File Monitors	243
Directory	243
File	244
Memory Monitors	244
NT Performance Monitors	244
Registry Monitors.....	244
TCP/IP Monitors	244
Process Monitors	245
Sentry Monitors.....	245
Windows 95 Operating System	246
CPU Monitors	246
Disk Monitors.....	246
File Monitors	246
Directory	246
File	246
Memory Monitors	247
Performance Statistics	247
File System	247
IPX/SPX compatible protocol	247
Kernel	248
Memory Manager	248
Microsoft Client for NetWare.....	248
Microsoft Network Client.....	249
Process Monitors	249
Registry Monitors.....	250
Sentry Monitors.....	250
OS/2 Operating System.....	250
APM Monitors.....	250
CPU Monitors	250

Disk Monitors	251
File Monitors.....	251
Directory	251
File.....	251
Memory Monitors	252
OS/2 Server Monitors	252
OS/2 Swapfile Monitors.....	252
Process Monitors.....	252
Sentry Monitors.....	253
NetWare Operating System	253
CPU Monitors	253
Disk Monitors	253
File Monitors.....	254
Directory	254
File.....	254
Memory Monitors	254
Process Monitors.....	254
Unix and Linux Operating Systems.....	255
CPU Monitors	255
Disk Monitors	255
File System Monitors	256
List of Directory Contents.....	256
File.....	257
Directory.....	257
Memory Monitors	258
Process Monitors.....	258
Sentry Monitors.....	258
Unix System Monitors	258

B. Creating the ODBC Entry for the Default Database 261

C. Converting to Other Supported Databases. . . . 263

Preliminary Steps	263
Using the Database Configuration Window to Convert to Another Database ..	264
D. Defining Table Property Files	265
Setting up the Server to Inventory CIM and DMI Information	265
Table Property File Format	266
NLS File Format.....	271
Inventory Extension Property File Format.....	274
Static MIF Data Collection	277
Server Initialization and Table Property Files.....	280
Examples	282
E. Agent-Server Security	287
How IBM Director Agent-Server Security Is Implemented	287
Installing IBM Director Agents in a Secure State	290
Determining the Origin of a Public or Private Key	291
Recovering Lost Public and Private Keys Files.....	291
Index	293

Preface

The *IBM Director: User's Guide* provides the installation and start-up instructions for the IBM Director product. It also describes the IBM 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"
Describes how IBM Director works and introduces the various tasks available to the network administrator.
- Chapter 2, "Planning"
Discusses considerations for network setup and management that should be addressed before installation and network administration tasks are performed with IBM Director.
- Chapter 3, "Installation and Configuration"
Lists the prerequisites and restrictions that apply to IBM Director and provides step-by-step instructions for installing and configuring IBM Director component and agent software. Also, the procedure for uninstalling IBM Director is provided.
- Chapter 5, "Using the Management Console"
Describes the IBM Director Management Console graphical user interface (GUI).

- Chapters “Chapter 6. Inventory Management,” on page 111 through “Chapter 20. Troubleshooting,” on page 227
Describe the various administrative tasks available in IBM Director for managing the hardware and software in your network.
- Chapter 20, “Troubleshooting”
Describes some typical problems you may encounter and possible ways to resolve them.
- Appendix A., “Resource Monitor Attributes”
Contains a list of all the attributes that can be monitored by IBM Director’s Resource Monitor task.
- Appendix B., “Creating the ODBC Entry for the Default Database”
Contains procedures for manually creating the default Microsoft Jet database.
- Appendix C., “Converting to Other Supported Databases”
Contains procedures for switching from the default Jet database to the DB2 Universal Server database support.
- Appendix D., “Defining Table Property Files”
Contains information on setting up your server to inventory CIM and DMI information.
- Appendix E., “Agent-Server Security”
Contains information on the process used to establish trust relationships between the IBM 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.

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.

Conventions Used in This Guide

This book contains information for installing and using IBM Director.

This Guide uses several typeface conventions for special terms and actions. These conventions have the following meaning:

Bold	Commands, keywords, file names, or other information that you must use literally appear in bold .
<i>Italics</i>	Variables and values that you must provide appear in <i>italics</i> .
<i>Bold Italics</i>	New terms appear in <i>bold italics</i> the first time they are used.
Monospace	Code examples appear in a monospace font.

Platform-Specific Information

The following table identifies text used to indicate platform-specific information:

Text	Supported Platform
NetWare	Novell Netware Versions 3.12 (and all applicable service patches), 4.1, 4.11 or 4.2 (with Service Pack 5 or higher), 5.0 (with Service Pack 1 or higher), or 5.1
OS/2	IBM OS/2 Version 4.0, 4.5, and IBM OS/2 Warp Server for eBusiness Version 4.5
Windows	Microsoft Windows 95, Microsoft Windows 98, Microsoft Windows NT 4.0 (Workstation or Server) with Service Pack 4 or higher, Windows 2000 Server, Windows 2000 Professional, Windows 2000 Advanced Server
Unix	SCO UnixWare 7.1, or higher, SCO Open Server 5.0, HP UX Versions 10.2 or 11.0
Linux	Linux Red Hat Version 6.1 or 6.2

1

Introduction

IBM Director is a management product for the systems administrator in small- to medium-sized growth companies. IBM 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 IBM Director functions enable you to manage all aspects of the computing infrastructure, including software applications, network operating system (NOS) services, networks, and hardware.

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

The following sections use these terms:

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

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

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

- IBM Director Management Console

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

The IBM 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.

- IBM Director Server

The IBM Director server is the heart of the IBM Director product. This is where management data, the server engine, and the management application logic reside. It is a Java and native C++ application. IBM 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 IBM Director server manages systems and devices in your network by communicating with the UM Services 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 IBM Director. IBM Director recognizes two types of managed systems:

- Native Managed Systems

These are managed systems that have the UM Services 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 IBM 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 IBM Director.

What Is New in This Release

This section discusses the new features this release of IBM Director contains.

Expanded Managed System Support

IBM Director now supports the Director management agent on several platforms, including Novell NetWare, IBM OS/2, Microsoft Windows, SCO UnixWare, and Red Hat Linux.

Refer to the product README file for information supported versions, managed system requirements, and the requirements for each Director management agent.

Expanded Database Support

In addition to Microsoft Jet, IBM Director supports IBM DB2 Universal Database Versions 5.2, 6.1, and 7.1, Oracle Server Versions 7.3.4 through 8.1.6, Microsoft SQL Server Versions 6.5 and 7.0 with Service Pack 1, and Microsoft Data Engine (MSDE) Version 1.0 with Service Pack 1. Refer to “Database Support” on page 13 for information on configuring IBM Director to use these databases.

Expanded Transport Support

IBM Director transport now allows dynamic loading of protocols for Windows. This support enables users to modify their protocols without having to start and stop IBM Director.

Transports are specified during installation. Refer to “Installation and Configuration” on page 49 for information on installing all IBM Director components.

New methods have been added for discovering managed systems over the TCP/IP and IPX protocols. The new broadcast relay function allows the server to discover agents that, due to network configuration, are not directly accessible by broadcast packets. 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.

The unicast discovery function enables the server to discover managed systems on networks where both broadcasts and multicasts are filtered. In many cases, for example, Remote Access Servers (RAS) do not forward broadcast traffic.

In the Discovery Preferences configuration panel, the IP tab has been enhanced so that you can now enable and disable the unicast function. The IP and IPX tabs have been modified to enable use of the broadcast relay function.

Expanded Associations Function

The Associations popup menu on the IBM Director management console has been enhanced to include task view association items for the following:

- Software Packages
- Jobs
- Activations
- Resource Monitors
- Event Action Plans

Refer to the online help for more information.

Expanded Support for Groups

The Group pane context menu in the IBM Director management console has been enhanced to include a Group Category Editor, Dynamic Group

Editor, Static Group Editor, and Task Based Group Editor, all of which allow you to create different types of groups. A Group Import/Export function has been added to the Management Console as well. This function enables a user to export group definitions from one server and import them to another server.

For more information on using these functions, refer to the online help.

Event Management Enhancements

IBM Director provides distributed event action plans that enable you to:

- Distribute existing rule bases to strategic servers throughout your network.
- Create and distribute event action plans to control the flow of events to the TME 10 event server.
- Track events as problem reports and manage the resolution of the problems.

Inventory Enhancements

The Inventory Query Browser now caches the customized query view each time you change column order, width, and so on. It is no longer necessary to modify the inventory query in order to save a customized query view.

You can now export inventory query results in the XML (extensible markup language) format as well as the CSV (spreadsheet) and HTML (hypertext markup language) formats.

The Inventory Database tables previously documented in Appendix A are being shipped as HTML on the product CD for this release. You can access them through the online help index (under Inventory Database Tables).

Resource Monitoring Enhancements

File monitors for all systems are now available. Individual file items that can be tracked include file size, file changes, file existence, directory changes, and directory existence. Registry monitors for Windows platforms are available now as well.

You can now create threshold plans that can be applied to multiple groups and systems, and each plan can contain threshold settings for multiple resources. These plans can then be exported to a file and imported from files. The ability to enable and disable individual and group thresholds has been enhanced.

Another new feature is a user preference view, which allows you to specify the levels of attributes in a path to show in the Selected Resources pane.

The number of NetWare monitors has increased. See Appendix A for more information. Also, monitor lists have been added to Appendix A for the new platforms supported this release (Unix and Linux).

For more information, see Chapter 8, “Resource Monitoring” or refer to the online help.

Process Management Enhancements

Windows NT device and service monitoring is now supported. For more information, see Chapter 18, “Process Management” or refer to the online help.

Remote Control Enhancements

The remote control task now supports two new options: recording a remote control session and changing the refresh rate for current remote control sessions. For more information on these options, see Chapter 7 or refer to the online help.

Two new options on the Network Driver Configuration dialog have also been added. And a new dialog has been added to the Server Preferences window to allow for automatic timeout of remote control sessions. For more information on these options, see Chapter 3. Installation and Configuration or refer to the online help.

File Transfer Enhancements

The file transfer task now supports the ability to synchronize directories, drives, and files with expanded wildcard search and transfers.

Task Scheduler Enhancements

The task scheduler feature now supports selecting 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.

On the Options page, there is a new Special Executions Option: Execute in client time zone. Selecting this option causes tasks to execute according to the time zone in which the target system resides.

The main Scheduler console now contains four tabs: one for Jobs and three relating to the previous Calendar tab. Now, you not only have the ability to display the Month view, but you also have the ability to display Week and Day views. 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. On the Week Calendar page, the current week appears. And, on the Day Calendar page, the current day appears in calendar format.

In addition, a job can be executed again on selected groups and individual systems. A selected job's execution history results can now be exported to a CSV or HTML file as well.

For more information, see Chapter 19, "Task Scheduler" or refer to the online help.

Removed Multiplatform Manager (MPM) Support

IBM Director no longer supports systems that the MPM API has discovered and that are being managed by MPM providers or any MPM-compliant application that has an MPM provider installed. The MPM providers that IBM Director previously shipped are: Microsoft Systems Management Server (SMS) 1.2; Intel LANDesk Management Suite 2.5.1 and 2.5.2; and IBM NetFinity Manager 5.0, 5.1 and 5.2.

Supported Communication Protocols

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

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

IBM Director uses the following protocols to communicate between the IBM 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 IBM Director management agents and servers on Unix systems.

Modem Connections

For IBM 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

IBM Director supports a comprehensive set of tasks for managed systems operating as full-function IBM Director management agents.

These agents communicate directly with the IBM Director server, enabling the following tasks to be performed.

Asset ID

IBM 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, IBM Director displays this information as a client-based task. See Chapter X, for more information.

Inventory Management

IBM 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” for details.

Remote Control

Enables you to provide 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 IBM Director Management Console. It is useful for training and educating new network administrators as well. See Chapter 7, “Remote Control” for details.

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” 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” for details.

Software Distribution

Enables you to collect software distribution packages that can then be applied to one or more managed systems for immediate or scheduled delivery. See Chapter 10, “Software Distribution” 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” for details.

Process Management

Enables you to start, stop, and monitor applications and processes on remote native systems. You can have IBM Director watch a particular process or application and generate an event if the application or process is started or terminated. See Chapter 18, “Process Management” 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 19, “Task Scheduler” for details.

Managing DMI-Enabled and CIM-Enabled Native Systems

IBM 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 13, “DMI Management” and to Chapter 14, “CIM Management” for information on tasks you can perform on DMI and CIM data.

Managing Cluster-Enabled Native Systems

IBM 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:

- Cluster Browser
- Resource Monitors
- Event Management

See Chapter 17, “Cluster Management” for information on viewing cluster data.

Managing SNMP Devices with IBM Director

IBM 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

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.

- SNMP Browser

See Chapter 12, “SNMP Management” 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 12, “SNMP Management” for more information.

Additional Features in IBM Director

Security

The IBM Director server uses the security subsystem of the operating system for validating user IDs and passwords. Each user has a unique login profile. This enables different users to log in to the IBM Director Management Console. Refer to the section “Security” on page 29 for more information.

Database Management

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

2

Planning

This chapter provides information you should consider before you begin installing and configuring IBM Director.

IBM Director Server

The IBM Director server is where most of the IBM Director processing occurs and requires more computer resources than the IBM 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 IBM Director server.

Management Server Prerequisites

Refer to the product README file for the minimum hardware and software requirements for the IBM Director server.

The TCP/IP networking transport and a network adapter that supports the TCP/IP networking protocol is also required. The adapter must also support NetBIOS, IPX, or SNA, depending on which transport is needed to communicate with the managed systems. See “Transport Support” on page 24 for information on supported versions of transports.

Database Support

IBM Director supports the following databases and versions:

-
- Microsoft Jet
 - IBM DB2 Versions 5.2, 6.1, and 7.1
 - Microsoft SQL Server Versions 6.5 and 7.0 with Service Pack 1
 - Oracle Server Versions 7.3.4 through 8.1.6
 - Microsoft Data Engine (MSDE) Version 1.0 with Service Pack 1.
If you plan to use MSDE, the MSDE database engine **must** be installed first.

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 IBM 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.

Planning to Use the Jet Database

IBM Director ships 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 IBM Director server. The Jet database has a maximum size of 1 GB.

Relocating the Jet Database

The Jet database cannot be split. After the IBM Director server is installed, it is possible to move the Jet database to another subdirectory besides *\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 you selected when you installed the IBM Director server (either the default or another name that you selected). Refer to the Windows NT online help for ODBC, 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. DB2 Server has additional storage capability and is more impervious to unwanted access attempts. Before installing the IBM Director server software, your network must be configured to use the DB2 Server database. Install the DB2 Client Application Enabler to access the DB2 server. Make sure that the DB2 Java Enablement option is installed and that CLASSPATH points to the directory 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:
 1. Set up a trusted connection or give proper login access to the IBM Director server database user ID.
 2. Create a node entry for the remote DB2 Server.
- Make sure you have sufficient licenses for DB2 Server, as this is a separate product from IBM Director and is not included in IBM Director licensing requirements.

Setting up Trusted Connections

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

DB2 Server Login Access for Windows NT

Your database administrator and your system administrator must configure security so that the IBM 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 IBM Director management server and the DB2 Server are on different domains. The IBM 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 DB2 Server Database

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

- The IBM Director Management Server Database User ID must be given user access to the database server.
- To create the database automatically, the IBM Director Management Server Database User ID must be given Create Database permission on DB2 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 IBM 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 - 500 managed systems. More space may 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 name of the server where the database is located, and the name of the database itself. You are now ready to proceed with IBM Director Management Server installation.

Note: DB2 has size restrictions on items such as user ID, table names, etc. 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 IBM 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 IBM Director server database user ID.
- Creating the SQL Server database manually or during IBM Director server installation

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

Setting up Trusted Connections

IBM Director uses trusted connections when logging in to the SQL Server. 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 IBM Director

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

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

You may need to set up a trusted relationship between domains if the IBM Director server and the SQL Server are on different domains, and then the IBM 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 IBM

Director server installation. Your database administrator should consider the following:

- The IBM Director management server user ID must be given user access to the master database.
- To create the database automatically, the IBM 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 IBM 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 - 500 managed systems. You may 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 IBM Director server installation.

You are now ready to proceed with IBM 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 IBM Director server software, your network must be configured to use the Oracle Server database. The following may require preliminary action:

1. If you do not have a User ID, one is created during the Database Configuration process.
2. The JDBC Thin client-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 client-side driver does not require Oracle client software to be installed. However, it does require the server to be configured with a TCP/IP Listener.
3. Make sure you have sufficient licenses for Oracle Server, as this is a separate product from IBM Director and is not included in IBM 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 tablespaces that are created and TWG_ROLE for security.

Configuring the Oracle TCP/IP Listener

The Oracle TCP/IP Listener must be configured and started prior to running the Database Configuration dialog.

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 **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.

Creating the Oracle Server Database

Your database administrator may choose to create the tablespaces manually, or allow the tablespaces to be created automatically during IBM 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 IBM 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 IBM 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 12. SNMP Management,” on page 167.

Web Server

If you are using IBM 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 and have access to the file system of the IBM Director server before the IBM Director server software is installed.

News and Mail Servers

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

Wake-On-LAN

IBM 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, IBM Director can send a "magic packet" to a managed system that is powered off. The packet is decoded by the system's 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 IBM 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.

IBM Director Management Console

The IBM Director Management Console (Management Console) is installed when you install the IBM 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 IBM Director Management Console.

Tivoli Management Agent

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

Refer to the product README file for the hardware requirements for each Director management agent.

To enable communication with the IBM Director server, the managed system must have one of the following network transports installed. (See “Transport Support” on page 24 for information on supported versions of transports.)

- TCP/IP
- NetBIOS
- IPX
- SNA

Note: For Unix TCP/IP is the only protocol supported.

Operating System Platform Support

The following table shows which operating system platforms are supported by the three main components of IBM Director.

Note: This table applies only to Version 2.2 levels of IBM Director components. If you are upgrading from an earlier version of IBM Director, refer to “Migration Support” on page 28 to determine which combinations of component versions are supported.

Operating System	Server	Console	Agent
Windows NT 4.0	Yes	Yes	Yes
Windows 95 (OSR2) or 98	No	Yes	Yes
Windows Millennium Edition (ME)	No	No	Yes
Windows 2000 Server, Professional, or Advanced Server	Yes	Yes	Yes
NetWare 3.12, 4.1, 4.11, 4.2, 5.0, and 5.1	No	No	Yes
OS/2 4.0 and Warp Server for eBusiness Version 4.5	No	No	Yes
SCO UnixWare 7.1, SCO Open Server 5.0	No	No	Yes
Linux Red Hat 6.1, 6.2	No	No	Yes

The agent running on Unix or NetWare does not support CIM, DMI, or Remote Control.

Transport Support

The IBM Director server communicates with the IBM Director Management Console using TCP/IP only. You can use TCP/IP,

NetBIOS, SNA, or IPX to establish communication between the IBM Director server and a managed system.

Supported transport software is *not* included as part of IBM 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 95, Windows 98, Windows NT Server 4.0, Windows NT Workstation 4.0, OS/2 4.0, Warp Server for eBusiness, NetWare 3.12, 4.1, 4.11, 5.0, and 5.1
NetBIOS	Native NetBIOS versions supported by Windows 95, Windows 98, Windows NT Server 4.0, Windows NT Workstation 4.0, OS/2 4.0 and Warp Server for eBusiness 4.5
IPX	IPX versions supported by NetWare 3.12, 4.1, 4.11, 5.0, and 5.1, Windows 95, Windows 98, Windows NT Server 4.0, and Windows NT Workstation 4.0, OS/2 4.0
SNA	<p>Windows NT: Microsoft SNA 4.0 with Service Pack 1</p> <p>Microsoft SNA 3.0 with Service Pack 2</p> <p>IBM Communication Server 5.0 for Windows NT</p> <p>IBM Personal Communications (PCOMM) 4.2 or later for Windows NT</p> <p>OS/2: IBM Communications Server 5.0 for OS/2</p> <p>IBM Personal Communications (PCOMM) 4.2 or later for OS/2</p> <p>Windows 95 and 98:</p> <p>IBM Personal Communications (PCOMM) 4.2 or later for Windows 95 and Windows 98</p>

Discovery

IBM Director discovery operates by sending out a discovery request from the server and then listening for responses from any IBM 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 accessible 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, IBM 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. IBM 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 once 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.

License Management

The License Administration window, available in the Options pulldown in the Management Console, enables you to:

- View current IBM Director license information
- Add new license key
- Remove existing license keys.

The Add License Key selection in the License Administration window enables you to:

-
- Increase the number of agent licenses
 - Increase the total number of agents the IBM Director server can manage
 - Upgrade the level of server functionality
 - Upgrade the number of agents supported by a particular version of IBM Director.

Refer to the online help for procedures on adding or removing a license. For more information on license support levels, refer to the International Program License Agreement (IPLA) for IBM Director licenses and the License Information booklet included in the product packaging.

Migration Support

If you are upgrading your network to IBM Director Version 2.2 and you intend to continue using previous version(s) of IBM Director components (server, console, or agent), use the listing in this section to determine what versions of these components are supported for use with IBM Director Version 2.2.

- Only 2.2 IBM Director consoles are supported on 2.2 IBM Director servers.
- The 2.2 IBM Director console supports both 2.2 and 2.12 IBM Director servers.
- The 2.2 IBM Director server supports 1.2. IBM Director agents (must be upgraded using Software Distribution).
- The 2.2 IBM Director server supports 2.12 IBM Director agents at a 2.12 level of function; however, most version 2.2 functions will be unavailable.
- The 2.2 IBM Director server supports version 2.2 and higher IBM Director agents at a 2.2 level of function.
- The 2.2 IBM Director agents support being managed from 2.12 IBM Director servers or higher.
- The 2.12 IBM Director agents are supported by 2.2 IBM Director servers (most new functions may not be used until upgraded).

Note: All existing configuration, accumulated data, installed software packages, and AMPs on a version 2.12 server will be fully-preserved and migrated when version 2.2 is installed as an upgrade. Once upgraded, the server and its data do **not** support being downgraded (this includes uninstall-but-keep data, followed by reinstall of version 2.12).

Security

To protect your network from unauthorized access, IBM 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 IBM Director server and IBM Director agents when the network is brought up. This section describes user-logon security, which you need to establish immediately after installing the IBM Director server and Management Console. Agent-server security is described in “Appendix E. Agent-Server Security,” on page 287.

IBM Director User-Logon Security

IBM 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 IBM 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, select **Options** → **Console Security** on the IBM Director Management Console.

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

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

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

IBM Director Accounts

Because the IBM Director server runs on an operating system which 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 IBM Director, but not administered by IBM Director. This means that IBM Director does not edit the user information for that account (such as changing the password or the user description) but you can modify the IBM Director-specific information. To add or remove those specific accounts or to change the password, use the specific operating system user administration function.

IBM Director also provides the capability to create accounts for which the IBM 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 IBM Director. All administration of these accounts is done through the IBM Director Console Security task.

Listing IBM Director Users

When you launch the IBM Director Console Security task, the window presented shows a list of all users that are authorized to login to the IBM 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 IBM 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 IBM Director.

Creating a New User

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

- Select the **New User** option from the main menu or right-click on the User Information table to bring up a context menu and select the **New User** option. Because IBM 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 on a user listed in the Unauthorized Server Users table and select the **Authorize User** option. This option creates an IBM Director account for this user using the server information and the current set of user defaults. Because IBM 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.

IBM 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** checkbox if you wish to make a user a superuser, which grants the user full authority on IBM 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 IBM 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 IBM 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** checkbox. If you wish to limit which groups the user can access, select the **Limit user access only to the groups listed**

checkbox. 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** checkbox. 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 IBM 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 31. 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

IBM Director provides a default template of attributes that is used to set up new user accounts when created. You can modify this template by selecting **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 31.

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 IBM Director user, right-click on 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 which 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.

Investment Protection and Integration

The IBM Director server includes a TMR gateway function that enables it to recognize Tivoli management agents on the network. If the TMR gateway function is enabled, the IBM Director server can discover Tivoli management agents on network systems and automatically install IBM Director agent software on these systems. Select **Options** → **Discovery Preferences** → **TME Gateway Parameters** on the IBM Director Management Console to set up this function.

Java Classes

All Java classes are installed locally on both the IBM Director Server system and the system where the IBM Director Management Console is

installed. This increases performance given the large number of classes needed for full console operation and all management tasks.

Planning for IBM Director Tasks

This section describes concepts, setup and usage considerations, and usage restrictions for IBM 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.

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

Redirected distribution

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 on 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 IBM 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, IBM 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 supports UNC-based and FTP-based file distribution servers and does not require the installation of the IBM 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 only receives the bare minimum of installation code needed to access the share and install the software from the IBM Director server.

This method, known as a *redirected install*, 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 IBM Director server and allow read access to all potential target systems. Refer to "Configuring IBM Director to Use File Distribution Servers" on page 69 for information on configuring the IBM 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, IBM 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, IBM Director tries to use each share before streaming the package directly to the

managed system. If a managed system can access the share and you have configured IBM 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 IBM 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 streamed installation, IBM 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 39 for more information on specifying streaming to IBM 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 IBM Director server. A software package is only cached 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

package was stored. When the package is distributed again, it may be cached on the same share or on a different share. If a software package is deleted it is removed from any share where it was cached.

During a redirected distribution, the IBM Director server first determines if the package is already cached on one of the file distribution server shares. If the package is not cached, the IBM Director server's list of file distribution server shares is searched to determine which share has enough free space to hold the package. Generally, the amount of free space used on a share is the maximum disk space specified for the share or the amount of free space available on the share, whichever is less. If a share is not found with enough free disk space to cache the package, the least recently used package may be deleted from a share if the deletion makes enough room for the new package. If a share cannot be found to hold the software package, the software package may be streamed to the managed systems.

Note: Do not attempt to manage software packages on the shares outside of the IBM Director Management Console. Doing so may adversely impact your software distributions. Use File Distribution Servers Manager to assist you in managing your software packages.

Determining Which Share Is Used on Redirected Installs

The list of shares that can potentially be used in a redirected distribution is determined by the IBM Director server and managed system configurations, and the server's and managed system's ability to access the shares on the list. The server configuration suggests which shares a managed system may use, but the managed system's configuration determines the shares that it prefers. Refer to “Defining Server Preferences” on page 81 for instructions on defining server shares. Refer to “Configuring Distribution Preferences for Managed Systems” on page 82 for instructions on setting distribution preferences for individual managed systems.

By default, the IBM Director server sends its suggested list of shares to a managed system or set of managed systems. The managed system evaluates the server's list of shares based upon the ordering of its share preference list, unless the managed system allows only streaming. The

managed system only evaluates shares if you choose to restrict distribution to only the shares in the managed system's list, and one or more of those shares are in the server's list. If you do not restrict the managed system's share preference list, it can evaluate the shares in the server's list 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 IBM 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

You have several options for forcing the streaming of software distribution packages:

- For an individual managed system or group, you can select to always stream a package from the IBM Director server. Refer to "Configuring Distribution Preferences for Managed Systems" on page 82 for instructions on accessing the appropriate option through the Management Console.

Specifying the Transport for Server Shares

If a server on 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 IBM Director server to the share.

Note: For OS/2, FTP is supported only for file transfer between the IBM 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 82 for instructions on specifying the FTP protocol to IBM 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, reboot 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 95 or Windows 98 managed system *that does not have a logon session* (no one is logged on to the target machine), you must first run TWGSHARE on the file distribution server. Refer to “Enabling UNC-based Share Access to Windows Managed Systems” on page 70 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 is 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 70 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 49.

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 bandwidth used to stream (copy) packages from the IBM Director server to managed systems
- Limiting the bandwidth used to stream (copy) packages from the IBM Director server to file distribution servers
- Limiting the bandwidth used to stream (copy) packages from file distribution servers to managed systems

Redirected software distribution is designed to minimize the network bandwidth dedicated to a package installation. If the IBM 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 83 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 distributions 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 IBM 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 84 for instructions on limiting concurrent streamed distributions. Refer to “Defining the Maximum Number of Concurrent Redirected Distributions” on page 83 for instructions on limiting the bandwidth for all managed systems. Refer to “Configuring Distribution Preferences for Managed Systems” on page 82 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 125 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 95 and 98
 - Windows 2000
 - OS/2 4.0 and Warp Server for eBusiness

-
- You cannot perform the remote control task on the following network nodes:
 - Unix systems
 - Native managed systems running under NetWare
 - SNMP devices

Note: Some nodes that can be detected as SNMP devices can also have a Tivoli management agent installed, which enables remote control to be performed.
 - You can concurrently monitor or control two or more remote systems from a single IBM Director Management Console.
 - If multiple IBM 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 IBM Director Management Consoles can remotely control multiple managed systems concurrently; however, the overhead load generated can cause system response to degrade significantly.
 - Only one IBM Director server can communicate with a remote system through remote control. If more than one IBM Director server attempts remote control communication, the communication is rejected and an error message is displayed on the IBM 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's desktop.
 - Remote control does not support full-screen graphic modes, including Win-OS/2 full screen graphics mode. You cannot use

remote control for 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 129.
- 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` 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, reboot the agent system or stop and restart the IBM Director agent.

- If TME 10 Remote Control or IBM Director Remote Control has already been installed on a system, the IBM Director server or management agent software can be installed on that system if, during installation of IBM Director, the option to install remote control is disabled.
- If the NetWare IPX agent software has been installed on an OS/2 system, the IBM Director management agent software can be installed on that system if, during installation of IBM 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, IBM 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 IBM 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 IBM Director events. When the mapping file is defined, IBM Director can detect and present CIM events for filtering.

SNMP Trap Support

IBM Director recognizes SNMP traps and generates a corresponding SNMP event if an SNMP trap is sent to the IBM Director server. The Event Type field in the Event Filter Builder window is updated to include the SNMP filtering category if the IBM 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 IBM Director for SNMP trap recognition, configure the network devices that generate SNMP traps to specify the IP address of the IBM Director server as a trap destination.

Following is an example of an SNMP trap event (cold start) entry in the IBM 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

IBM Director collects inventory information from managed objects 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 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 265.

The inventory database tables are HTML files (one each for each table). The HTML 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 the 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 string values is shown immediately after the table.

DB2 Version 5 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 names.

3

Installation and Configuration

IBM Director is divided into the following components:

- IBM Director server
- IBM Director management console
- UM Services Agent

See “Chapter 2. Planning,” on page 13 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.

Before You Begin

Before you install UM Services on your client or server, consider the following:

- Supported Systems by Component
- Hardware Requirements
- Workgroup/Enterprise Integration

Note: If you intend to use the Microsoft SQL Server, the IBM DB2 Universal Database Server, the Oracle Server, or the Microsoft Data Engine (MSDE) instead of the Microsoft Jet Database for IBM Director database support, refer to Chapter 2 for

information on steps to set up these servers before installing the IBM Director server.

Supported Systems by Component

The IBM Director Management Server, which installs all IBM Director components, is supported on the following operating systems:

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

The IBM Director Console component supports the same operating systems as IBM Director Management Server, including:

- Windows 2000 Professional
- Windows NT Workstation 4 (with Service Pack 4 or later)
- Windows 98
- Windows 95 (with OEM Service Release 2 (OSR2) or later)

The UM Services Client component is supported on the following Windows operating systems:

- 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 95 (with OEM Service Release 2 (OSR2) or later)
- Windows ME

Note: The UM Services Client component is installed directly on supported Microsoft operating systems through the main installation program. UM Services Client is also supported on other operating systems through a direct installation from each supported system directory on the CD. See “Installing the UM Services Client for IBM Director” on page 65.

Hardware Requirements

IBM Director Management Server and Console require the following hardware, memory, and disk space:

- Pentium II class processor, 300 MHz or faster
- 64 MB of random access memory (RAM)
Note: This memory requirement is for the operational capabilities of IBM Director only. Operating systems such as Windows 2000 Advanced Server have a higher memory requirement for installation and operation.
- 150 MB of virtual storage
- 85 MB of free disk space
- A network adapter that supports the TCP/IP protocol. The adapter must support also NetBIOS, IPX, or SNA, depending on which transport is needed to communicate with the managed systems.

The UM Services Client for IBM Director in a Windows operating system requires the hardware, memory, and disk space:

- An IBM Netfinity server or @server, IBM Desktop, IBM IntelliStation computer, or IBM ThinkPad mobile computer.
Note: Client systems must support SMBIOS version 2.0 or higher.
- An Intel Pentium 200 MHz or faster processor.
- 75 MB of hard disk space on the client systems.
- A minimum of 32 MB RAM.

Workgroup/Enterprise Integration

IBM Director installation allows you to install UM Services as an integrated systems-management solution for a supported systems-management console application. Detailed information about the appropriate environments for integration can be found in the *UM Services Installation and User's Guide*, included on the *IBM Director with UM Server Extensions* CD.

Installing IBM Director 2.2 Instructions

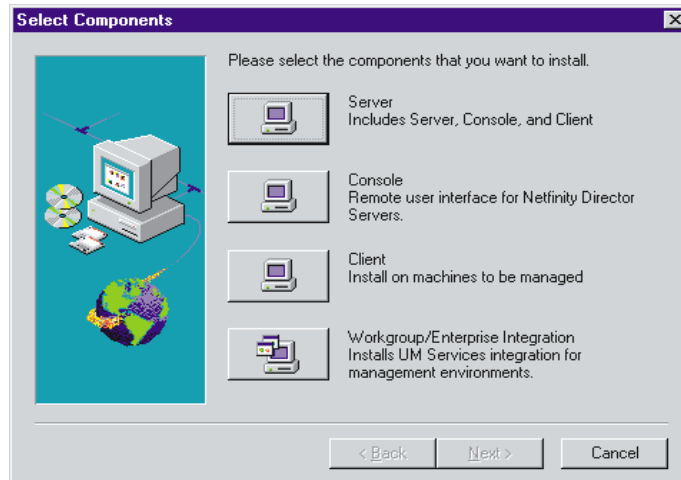
Take the following steps to install the IBM Director components (Server, Console, and Client for supported Microsoft Windows operating systems).

Notes:

- Pressing **Cancel** on the database configuration dialogs only cancels the database configuration process, not the installation process.
 - Pressing **Cancel** during an upgrade, either while files are being copied or after they are copied **will** corrupt the currently-installed copy.
 - Selecting the close window button (the **X** in the upper right-hand corner of the dialog) does not always have the same results as pressing the **Cancel** button. For example, selecting the close window button during the database configuration process cancels the entire installation (not just the database configuration process).
1. Place the *IBM Director* CD in the CD-ROM drive of the machine to which you will be installing.
 2. The Autostart program will automatically start the IBM Director installation program in a Windows environment. However, if Autostart is disabled, you can initiate installation from the command line.
 - a. Click **Start** → **Run**.
 - b. In the **Open:** field, type
X:\Director\win32\install\Ibmsetup.exe
where *X* is the location of the CD-ROM drive.
 3. Click the **Install IBM Director** button. Clicking the Install UM Server Extensions button will start the UM Server Extensions installation application. For more information about this product, refer to the *UM Server Extensions User's Guide* included in the \DOCS directory on this CD.

-
- Click through the **Welcome** window and accept the License Agreement.

The **Select Components** window opens.



There are four different installation choices from the **Select Components** window:

Server Install the files for the Server, Console, and Client.

Console Install the remote user interface for IBM Director Servers.

Client Install the Client files on Microsoft operating systems only. For other managed operating system environments, see “Installing the Windows UM Services Client” on page 66.

Workgroup/Enterprise Integration

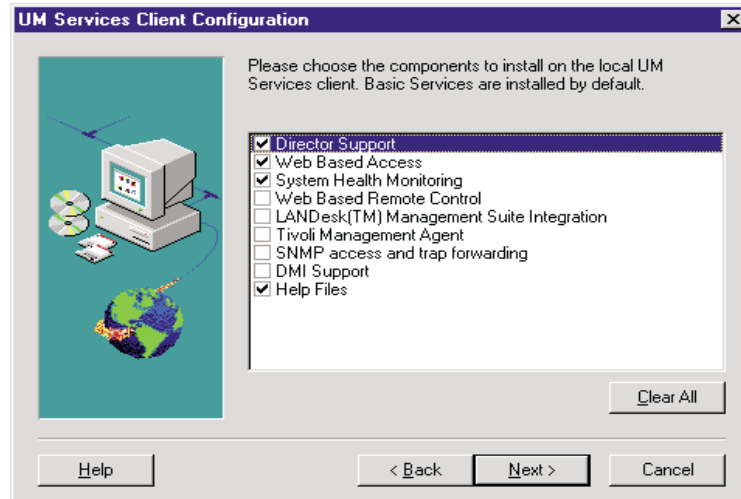
Install UM Services integration for other management environments. See the *UM Services Installation and User’s Guide*, provided on the *IBM Director with UM Server Extensions CD*, for complete information.

Installing IBM Director Server

Selecting **Server** from the **Select Components** window installs the files for the Server, Client, and Console. Take the following steps:

- Click **Server**.

The **UM Services Client Configuration** window opens.



2. Select the check box beside any of the components that you want to install on the client 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 client installation only. IBM Director is an advanced Intel processor-based workgroup hardware manager, with centralized client and group management console and server services. Selecting this feature enables the client system to be managed in a Director environment by installing UM Services on this system.

Web Based Access (default)

Web Based Access offers a convenient Java-based tool for managing a client 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 client system, the client 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 and SMS status message (Microsoft SMS 2.0 only). Critical problems also cause a message to be displayed on the monitor of the client 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 client 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 client system. This enables the system administrator to use UM Services with LANDesk Management Suite.

Tivoli Management Agent

Tivoli Management Agent installs support on the client 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 UM Services installation.

DMI Support

Selecting this component installs the Desktop Management Interface (DMI) compliant Service Provider. When enabled, this feature maps a managed system CIM data and events to DMI.

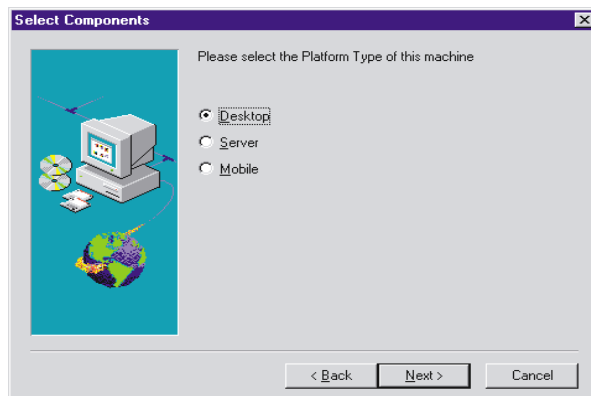
Help Files (default)

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 client system.

3. After you select the components you want to install, click **Next** to continue.

The following steps are dependent on the selections that are made in the UM Services Client Configuration menu.

4. If you selected **DMI Support** from the UM Services Client Configuration menu, the **DMI Platform** window opens.



Select the type of platform you are installing the DMI support to. Click **Next** to continue.

5. If Microsoft IIS is installed, a prompt will ask if you want to use IIS as the UM Services web server. Click **Yes** or **No** according to your preferences.
6. The **Add icons for UM Services** prompt asks if you want to add UM Services icons to the Start menu. Click **Yes** or **No** according to your preferences.
7. The **User ID** window opens. Use this window to set the user ID and password for the client system and to specify the TCP/IP port that is used to access the client.

You must provide a unique user ID and password for the client system. To use the UM Services console to manage this system, you must first provide a valid user ID and password before being allowed access to the system. Type in the **User ID** field the user ID. Then type the password in the **Password** field, and type the password again in the **Confirm Password** field.

Note: The user ID and password are case sensitive.

Then, select a TCP/IP port that is used to access the UM Services console. The default port is 411. If this port is not available, you can select port number 6411, 6500, 6600, or 6611. Make sure that other TCP/IP applications do not use the selected port. Click **Next** to continue.

-
8. If you selected **SNMP access and trap forwarding** from the UM Services Client Configuration menu and do not have the SNMP network service installed, IBM Director installation will prompt you with an SNMP installation query.
 - Click **No** to continue with the IBM Director installation without installing the SNMP network service.
 - Click **Yes** to install the SNMP network service on the server. The Installing SNMP window and your Network window opens. Follow the directions for installing SNMP. When the Windows operating system prompts you to restart, click **No**. In the Installing SNMP window, click **Next** to return to the Nefinity Director installation program.
 9. The system asks if you want to add a IBM Director icon to the **Start** menu. Click **Yes** or **No** depending on your preference.
 10. 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.

The directory that is used to store software distribution packages you create (see “Software Distribution” on page 35 for details). You can use the default drive and directory name provided or you can change it to something else. For large packages, select a drive with sufficient free space for large packages.
 11. 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.

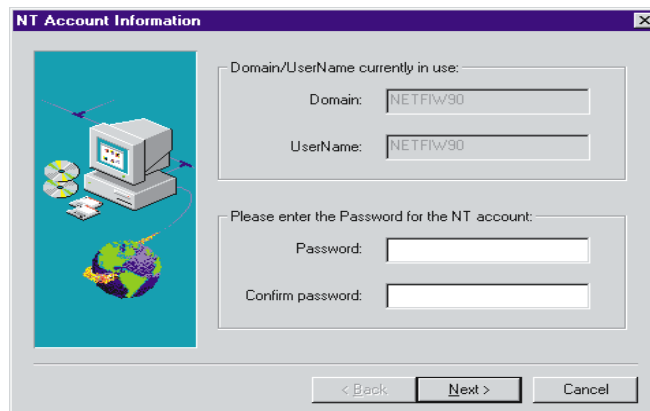
See “Software Distribution” on page 35 for details. You can use the default drive and directory name provided or you can change it to something else. For large packages, select a drive with sufficient free space for large packages.

12. Click **Next** to accept the default directory (**C:\Program Files\Director\SwDistPk**), or click **Browse** to choose a different directory.
13. The system asks you if you want to install files for remote control. Select **Yes** or **No**.

Selecting **Yes** enables IBM Director to perform remote control operations (remember that this system also acts as a managed system that can be remotely controlled; refer to Chapter 7 on page 125 for more information). If you select **No**, then remote control is disabled on this system.

Note: There are certain restrictions regarding the installation of remote control capability. Refer to “Remote Control” on page 43 for details.


14. The installation continues, installing software according to your selections. The **Account Information** window opens.



The Domain and User Name for the machine you are using are displayed.

15. Enter your account password in the **Password** field and enter it again in the **Confirm Password** field.
16. Click **Next** to continue.

The **Director Database Configuration** window opens.



17. Select the database of your choice (refer to “Database Support” on page 13 for information on database support considerations):

- **Microsoft Access:** Selecting this option creates a Microsoft Jet database. If a Jet database already exists which uses the default database name of IBM Director, you are given the option to delete it and re-create it, or specify a new database name.

- **Microsoft SQL Server:** Select this option if you want to use Microsoft SQL Server database support. The Database Configuration window is displayed, and you are prompted for the ODBC Data Source, Server Name, Database Name, User ID, and Password information. If you are using SQL 6.5, a second window will be displayed, prompting you for device name information for the database and transaction log. Your database administrator should give you the device names to use.

For detailed information on the fields displayed on the Database Configuration window for this option, refer to the online help.

- **IBM DB2 Universal Database:** Select this option if you want to use IBM DB2 Universal Database support. The Database Configuration window is displayed, and you are prompted for Database Name, User ID, and Password

information. A second panel is displayed, and you are prompted for node name information for the DB2 Server.

For detailed information on the fields displayed on the Database Configuration window for this option, refer to the online help.

- **Oracle:** Select this option if you want to use Oracle Server database support. The Database Configuration window is displayed, and you are prompted for TCP/IP Listener Port, Oracle Host Name, System Identifier, User ID, Password, and Oracle Administrator Account and Password information. A second panel is displayed, and you are prompted for Tablespace Name, Data File, and Size information.

For detailed information on the fields displayed on the Database Configuration window for this option, refer to the online help.

Note: Prior to running the Database Configuration dialog, you must ensure that the Oracle TCP/IP Listener is configured and started. If you are running Oracle 7.3.4, you must also make several other modifications. See “Planning to Use the Oracle Server Database” on page 19 for information on how to make these changes.

- **MSDE:** Select this option if you want to use Microsoft Data Engine (MSDE) database support. The Database Configuration window is displayed, and you are prompted for the ODBC Data Source, Server Name, and Database Name information.

Note: The MSDE database engine **must** be installed prior to selecting this option.

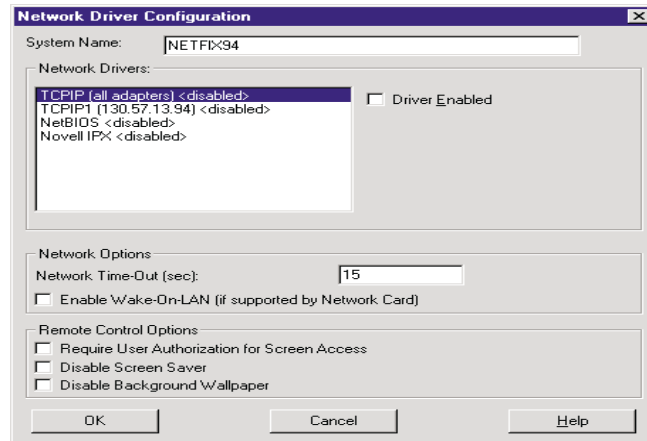
For detailed information on the fields displayed on the Database Configuration window for this option, refer to the online help.

- **Use existing ODBC entry:** Selecting this option displays a list of existing Jet, SQL, and MSDE data sources on the system. You can select one of these to use as your database.

Note: If you want to use an existing DB2 database, start the DB2 Configuration Database dialog and enter the name of the database.

Use Back and Next to move between these options as desired.

18. The **Network Driver Configuration** window opens.



This window defines the network transport options for an IBM 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 IBM Director, click on the driver name and check the **Driver Enabled** box.
- **Network Address (for NetBIOS only)** - This is the NetBIOS network name.
- **Network Time-out (sec)** - 15 seconds is the default time-out.

-
- **Require User Authorization for Screen Access** - Check this box if you want users to be able to control remote access to their systems.
 - **Enable Wake On-LAN** - Enables IBM Director to wake up a managed system that may be in “sleep” mode before performing tasks on that system. The managed system must have a network card that supports this feature. See “Additional Considerations” on page 21 for more details..
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 **Require User Authorization for Screen Access** checkbox if you want to give client users the authority to deny the system administrator remote control access to their machines. This option allows users to control who accesses their machines.
 22. Select the **Enable Wake On-LAN** checkbox if the server has Wake On-LAN capability.
 23. To disable screen saver programs on monitored systems, use the **Disable Screen Saver** option. To minimize system overhead, this option temporarily disables screen saver programs when the server detects active remote control sessions. When the remote control session ends, the previous screen saver is restored on the monitored system.
 24. To disable screen wallpaper graphics on monitored systems, use the **Disable Background Wallpaper** option. To minimize system overhead, this option temporarily disables screen background wallpaper when the server detects active remote control sessions. When the remote control session ends, the previous screen wallpaper is restored on the monitored system.
 25. When you have finished, select **OK** to save your settings. You can change these values at a later time by selecting **Start → Programs → Director → Network Driver Configuration** to bring up the Network Driver Configuration dialog once more.

When the installation is complete, the **Setup is Complete** window opens.

26. 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 IBM Director.

The IBM Director server and supporting programs run as Windows NT services and starts automatically when Windows NT is started. You can disable the automatic startup (select **Start** → **Settings** → **Control Panel**, then select **Services**, highlight **Director Support Program**, select **Startup** and change the **Startup Type** to **Manual**) and start the server manually using the **Services' Start** button, or by issuing the command: **NET START TWGIPC**.

Installing the IBM Director Console

The IBM Director Management Console performs all IBM Director tasks and should be installed on the network administrator's system. The IBM Director Management Console can be installed on multiple systems.

Note: For Windows only, if you have already installed the IBM Director server on your system, the IBM Director Management Console has already been installed, and you can skip this procedure.

To install the IBM Director console, take the following steps:

1. From the **Select Components** window, select **Console**.
The **Choose Destination Location** window opens.
2. Click **Next** to accept the default directory (**C:\Program Files\Director**), or click **Browse** to choose a different directory.
The necessary files are installed on the system.
The **Setup is Complete** window opens.
3. 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 IBM Director.

4. Click **Finish**.

Installing the UM Services Client for IBM Director

IBM Director, as a highly integrated workgroup hardware manager, allows you to manage a heterogeneous environment through the use of the UM Services Client.

You can install the Client on a Microsoft Windows system as part of the IBM Director installation. For OS/2, Netware, SCO UnixWare and Red Hat systems that are managed by IBM Director, use the *IBM Director with UM Services* CD-ROM directly to install the Client.

Supported Operating Systems

The following operating systems and versions are supported:

- Microsoft ME, Windows 95, Windows 98, Windows NT 4.0 (Workstation or Server), and Windows 2000 Professional and Server
- OS/2 3.0 and 4.0 with Service Pack 5 and Warp Server for eBusiness
- NetWare 4.x (with Service Pack 5 or higher), NetWare 5.x (with Service Pack 1 or higher)
- SCO UnixWare 7.1 (with SCO UnixWare Patch ptf7441a) or SCO UnixWare 7.1.1 with these patches:
 - ptf7045
 - ptf7410
 - ptf7441
 - ptf7601
 - ptf7602
 - ptf7603
 - ptf7608

-
- ptf7616
 - Linux Red Hat 6.1 and 6.2 with Linux kernel version 2.2 or later. The following libraries must be included:
 - ld-linux.so
 - libc.so
 - libm.so
 - libdl.so
 - libpthread.so

Hardware Requirements

Hardware requirements for Windows operating systems are detailed on page 51. The other operating systems have these requirements:

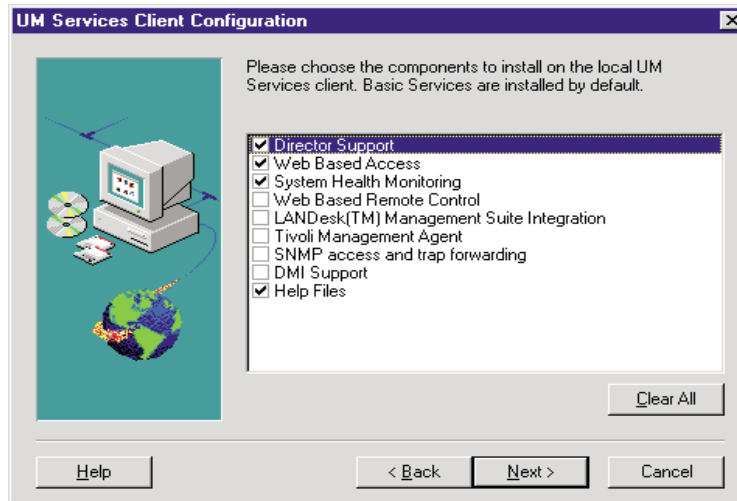
- OS/2 requires 16 MB RAM
- SCO UnixWare 7.1.1 and Linux Red Hat require a minimum of 64 MB RAM and 16 MB RAM available hard drive space on the mount point

Installing the Windows UM Services Client

From the **Select Components** window, do the following.

1. Click **Client**.

The **UM Services Client Configuration** window opens.



2. Select the check box beside any of the components that you want to install on the client 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 client installation only. IBM Director is an advanced Intel processor-based workgroup hardware manager, with centralized client and group management console and server services. Selecting this feature enables the client system to be managed in a Director environment by installing UM Services on this system.

Web Based Access (default)

Web Based Access offers a convenient Java-based tool for managing a client 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 client system, the client

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 and SMS status message (Microsoft SMS 2.0 only). Critical problems also cause a message to be displayed on the monitor of the client 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 client 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 client system. This enables the system administrator to use UM Services with LANDesk Management Suite.

Tivoli Management Agent

Tivoli Management Agent installs support on the client 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 UM Services installation.

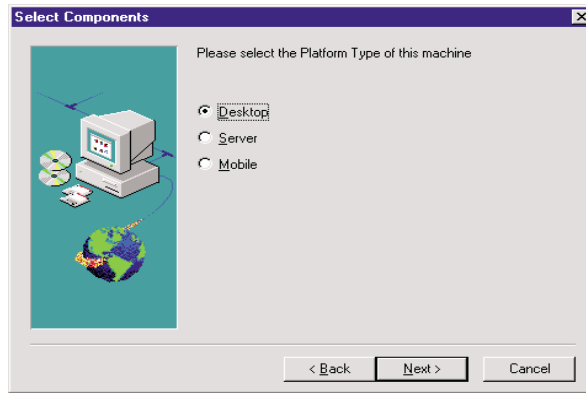
DMI Support

Selecting this component installs the Desktop Management Interface (DMI) compliant Service Provider. When enabled, this feature maps a managed system CIM data and events to DMI.

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 client system.

3. Click **Next** to continue.
The **Choose Destination Location** window opens.
4. Click **Next** to accept the default directory (**c:\Program Files\UMS**), or click **Browse** to select a different directory.
5. If you did not select to install **DMI Support** in step 10, go to step 14. If you selected **DMI Support**, the **DMI Platform** window opens.



Select the type of platform that you are installing the DMI support to. Click **Next** to continue.

6. To use the UM Services console to manage this system, you must first provide a valid user ID and password before being allowed access to the system. In the **User ID** field, type the user ID. Then type the password in the **Password** field, and type the password again in the **Confirm Password** field.

Note: The user ID and password are case sensitive.

Then, select a TCP/IP port that is used to access the UM Services console. The default port is 411. If this port is not available, you can select port number 6411, 6500, 6600, or 6611. Make sure that other TCP/IP applications do not use the selected port.

7. Click **Next** to continue.
8. If you selected **SNMP access and trap forwarding** from the UM Services Client Configuration menu and do not have the SNMP network service installed, the UM Services installation program will prompt you with an SNMP installation query.
 - Click **No** to continue with the UM Services installation without installing the SNMP network service.
 - Click **Yes** to install the SNMP network service on the server. The **Installing SNMP** window and the **Network** window opens. Follow the instructions for installing SNMP. When the Windows operating system prompts you to restart the

system, click **No**. In the **Installing SNMP** window, click **Next** to return to the UM Services installation program.

9. When the system asks you if you want to place a UM Services icon on the start menu, click **Yes** or **No**.
10. When the system asks you if you want to install files for remote control, click **Yes** or **No**.
11. When the system asks you if you require user authorization for remote control window access, click **Yes** or **No**.
The installation program begins installing the necessary files.
12. When the Director Installation Complete window opens, Click **Finish**.
13. Restart the computer now or later. If you click **Yes**, the system shuts down and restarts immediately. If you click **No** the UM Services installation program closes. However, you must restart and log in to the system to begin using UM Services.

Using the UM Services Client Attended Installation for OS/2

To install the UM Services Client software on an OS/2 3.0, 4.0, or Warp Server for eBusiness system, do the following.

1. Insert the *IBM Director with UM Services* CD-ROM into the CD-ROM drive of the system.
2. Change the directory to the **x:\Director\os2** subdirectory, where *x*: is the drive letter of the CD-ROM drive.
3. Run **setup.cmd** to start the installation utility.
4. The default location of the UM Services Client files is 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 Client 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 Client 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. Refer to the Tivoli IT Director section, “Remote Control,” for more information.
 7. Select **Install**. The files are copied to 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 IBM Director administrator uses this name to identify this system on the network.
 9. Click on one of the available network drivers the managed system uses to communicate with the IBM 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 IBM 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 local user’s permission. If this option is enabled and a IBM 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. Refer to the IBM Director section, “Remote Control” for more information on using this service.

13. Select **Enable Wake on LAN** option to enable IBM Director to wake up a managed system that may be in “sleep” mode before performing tasks on that system. The managed system must have a network card that supports this feature.
14. When you have finished, click **OK** to save the settings.
15. The installation utility displays the changes that must be made to the **config.sys** and **startup.cmd** 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** and **startup.new** instead.

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

16. Installation is now complete, click **OK** to save the settings.

Using the UM Services Client Unattended Installation for OS/2

IBM Director supports unattended installations, meaning that 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 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, DirServ.rsp, is located in the **x:\Director\win32\install\files\Nfd\Agent\OS2\en** subdirectory, where **x:** 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 that detail the usage of each entry.

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

1. Copy and modify the sample response file (**DirAgent.rsp**).
2. Change the directory to the **x:\win32\install\files\Nfd\Agent\OS2** subdirectory, where **x:** is the drive letter of the CD-ROM drive.
3. Run **setup.cmd** to start the installation utility:
install.exe /R:filename (where *filename* is the fully-qualified response file)

Installing the UM Services Client for NetWare

Notes:

- Do not install the UM Services Client on a NetWare system running Netfinity Manager. Netfinity Manager has to be commented out of **Autoexec.ref**.
- UM Services Client is only supported in NetWare 4.10, 4.11, 4.2 and 5.0.
- The Netware client system for the UM Services Client must be currently logged into the Novell NetWare Server.

To install the UM Services Client software on a Novell NetWare system:

1. Insert the *IBM Director with UM Services* CD-ROM into the CD-ROM drive of the system.
2. Change the directory to the **x:\Netware** subdirectory, where **x:** is the drive letter of the CD-ROM drive.
3. Run **setup.bat**.
4. Click **Next** on the **Welcome** window.
5. The End User License Agreement window opens. Click **Yes**.
6. From the **Choose Destination Location**, select the appropriate drive that is mapped to the **sys volume** of the targeted Novell server. Click **Next**.
7. The target location for the client files is the **\tivoliwg** directory.

-
8. It copies the necessary files, and adds the following lines to the Autoexec.ncf file:

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

9. The final window is displayed, listing several manual tasks the user must perform before this application is used:
 - Type the following:
Search add sys:tivoliwg
 - Configure the UM Services Client by loading **twgipccf**
 - Start the client by loading **twgipc**
10. Installation is now complete. The client automatically runs on the next Novell server boot.

Installing UM Services Client for SCO UnixWare

To install the UM Services Client on a SCO UnixWare system, you must install both the Server Extensions and the UnixWare Lightweight Client.

Before installing the UM Services Client for SCO UnixWare, these conditions must first be met:

- SCO UnixWare 7.1 or 7.1.1 is installed on the client system.
- SCOUNixWare 7.1 Patch ptf7441a must be applied to the client system.

To install the UnixWare Lightweight Agent for each SCO UnixWare Client, do the following.

1. At the SCO UnixWare system you want to manage, open the command line program.
2. Type **#mount /cdrom** and then hit **Enter**.
3. Type **#cp /cdrom/sco7.1/uagent.pkg /temp** and then hit **Enter**.
4. Type **#pkgadd -d /temp/uagent.pkg** and then hit **Enter**.

SCO UnixWare installs the IBM Director UnixWare Lightweight Agent.

5. Type **#umount /cdrom** and then hit **Enter** to finish the installation.

IBM Director will recognize this SCO UnixWare system as a client on the next discovery. Unix Agents are secure by default. Consequently, when they are first discovered by the IBM Director Management Server, the managed client name appears with a lock icon.

To gain access, right-click on the locked system and select, **Request Access**, from the pop-up menu. Enter the appropriate user name and password. Click **OK**.

Installing UM Services Client for Linux Red Hat

To install the Director management agent software on a Red Hat Linux 6.1 system:

1. Insert the *IBM Director with UM Server Extensions* CD into the CD-ROM drive.
2. Change the directory to *X:/Director/Linux* where *X* is the CD-ROM drive label
3. Copy the *ITDAgent-3.10-1.i386.RPM* package to a temporary directory.
4. Logged in as **root**, from the temporary directory, type:

```
/bin/rpm -i ITDAgent-3.10-1.i386.rpm
```

The agent is installed in the */opt/tivoliwg* directory.
5. After the installation is complete, the agent starts automatically, and the system is ready for discovery by a IBM Director server.

Reserving Agent Disk Space During Software Distribution (Windows and Unix)

Streamed software distributions take double the amount of disk space on a managed system in order to distribute a software package. For this

reason, you may want to reserve space on disk volumes to prevent the overwriting of critical directories.

To do this, you need to create a text file named **swdstcfg.prp** in the agent in the “data” directory. For volumes with reserved disk space,

- Entry identifiers must be numbered consecutively, starting with 0.
- For defined Unix volumes, the path separator character must be appropriate for the system type.
- If no values are specified, the default reserved space for all volumes is 10 megabytes. This default can be overridden.

The volume identifier is specified by the following line, where *n* is the entry identifier, and *vol id* is the volume identifier.

```
volume.id.n=vol id
```

The amount of reserved space on volume identifier *n* is specified by the following line, where the value of *x* is specified in megabytes.

```
volume.reserve.n=x
```

The default override for undefined volumes is specified by the following line, where the value of *x* is specified in megabytes.

```
volume.reserve.default=x
```

The following is a Windows-specific example:

```
# Volume identifier is a drive letter. volume.id.0=c:  
# Reserved space is in megabytes.  
# volume.reserve.0=35  
# Default override for all undefined volumes.  
volume.reserve.default=25
```

The following is a Unix-specific example:

```
# Volume identifiers are directory mount points.  
volume.id.0=/home  
Reserved space is in megabytes. volume.reserve.0=75  
volume.id.1=/usr  
volume.reserve.1=100  
# Default override for all undefined volumes.
```

Installing Oracle Server or DB2 Universal Databases Using the Command Line (Unix)

It is recommended that you use the graphical interface database installation process. However, you may use the command line to install the DB2 or Oracle Servers on Unix.

Installing the DB2 Universal Database

To install the DB2 database from the command line, do the following:

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

```
twg.database.odbc.name=test20
twg.database.jdbc.driver.name=COM.ibm.db2.jdbc.app.DB
2Driver
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:

```
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
```

Installing the Oracle Server

To install the Oracle Server from the command line, do the following:

1. In the IBM Director **/data** directory, edit or create the TWGServer.prop file. Add the following lines (where *goth* 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:

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

Defining Server Preferences for Database Properties

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

On the IBM Director Management Console, select **Options** → **Server Preferences** and then from the Server Preferences window, select the **Database** tab. This tab 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 IBM 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 “Software Distribution” on page 35 before you attempt to use the file distribution servers (server shares). Perform the steps described in this section to enable server access and configure IBM Director to use server shares for software distribution.

Enabling UNC-based Share Access to the IBM Director Server

The user ID under which the IBM 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 IBM 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 Windows systems 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 Windows 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 IBM 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 IBM Director server to use it.

On the IBM 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 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 notation:

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 bandwidth 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 managed 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 managed 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 IBM 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 IBM Director server to the systems 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 IBM 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 IBM Director server, IBM 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 managed 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 IBM 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 IBM 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 IBM 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 IBM Director server and a managed system's 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 42 for more information on limiting the bandwidth of a distribution.

Restricting Access Check

If you select **Restrict server access check**, the IBM 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 NetWare systems, access the NetWare server console either locally or by remote control (rconsole). 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 (Unix only)

Before uninstalling or at other times, you can back up the program files for the IBM 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 IBM 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.saven*, 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.saven) 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 IBM 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.

Uninstalling IBM Director

Before you remove IBM Director, if you have the UM Server Extensions installed, you must first uninstall the extension tools. For information on the UM Server Extensions, refer to the *UM Server Extensions User's Guide*.

To remove IBM Director:

1. Click **Start**→ **Settings**→ **Control Panel**→ **Add/Remove Programs**. Then select **Director**.
2. The system displays a message that reminds you that IBM Director must be closed before you can continue.
3. The system displays a message that verifies that you want to delete the configuration data and database content.

The uninstallation program is automated and prompts you when the process is finished.

4. Restart the computer now or Restart later. If you choose **Restart Now**, the system shuts down and restarts immediately.

Uninstalling IBM Director Components on OS/2

To uninstall IBM Director components from an OS/2 system:

1. Change to the Director management agent directory (typically `c:\tivoliwg`).
2. From a command prompt, type **bmunist** 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** (substitute the appropriate directory where you installed the Director management agent) directory and all of its contents.

Uninstalling IBM Director Components on NetWare

To uninstall IBM Director components from a Novell NetWare system:

1. Unload IBM Director by entering **unload twgipc.nlm**.
2. From a Director management agent logged in to the Novell server, delete the `\tivoliwg` directory.
3. Edit the `autoexec.ncf` file on the Novell server and remove the IBM Director section.

4

Upgrading to IBM Director 2.2


You can upgrade previous versions of Netfinity Director to IBM Director, version 2.2. The installation program checks for a previous version of IBM Director and, depending on the type of installation, upgrades the necessary IBM Director components.

Upgrading the IBM Director Server

You follow the same steps to upgrading the Server as you would with installation. You begin by inserting the *IBM Director with UM Server Extensions* CD into the CD-ROM of the Director system you are upgrading. Follow steps 2 on page 52 through 4 on page 53 of the installation section.


To upgrade the Director Server to version 2.2, do the following.

1. Select **Server** from the **Select Components** window.
Installation detects a previous installation. Click **Yes** from the Question pop-up to proceed with the upgrade.
2. The **Current Install** window opens.



Items that were installed with the previous version are listed here. These items are upgraded with the new version. Click **Next** to continue.

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




Features that are either new to Director or were not installed in the previous version are listed here. Check the box beside each feature you wish to add.

Click **Next** to continue.

4. Click **Yes** to the Add UM Services Icons pop-up for the additional menu choices provided with the upgrade.
5. The upgrade program detects the previous installation path. Click **Yes** to continue.
6. The installation continues with steps 9 through 16 on page 59. The upgrade program detects the current Director database. Click **Yes** to use the database with the upgrade.
7. Select **Yes, I want to restart my computer now**. Click **Finish** to complete the upgrade process and to restart the system.

Upgrading the IBM Director Console

After the Welcome and End User License Agreement windows, select **Console** from the **Select Components** window. Installation detects the previous version of the IBM Director Console and this prompt appears.




Click **Yes** to upgrade the console. 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 choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the UM Services installation program closes. However, you must restart and log in to the system to begin using IBM Director.

Upgrading the IBM Director Client


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

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



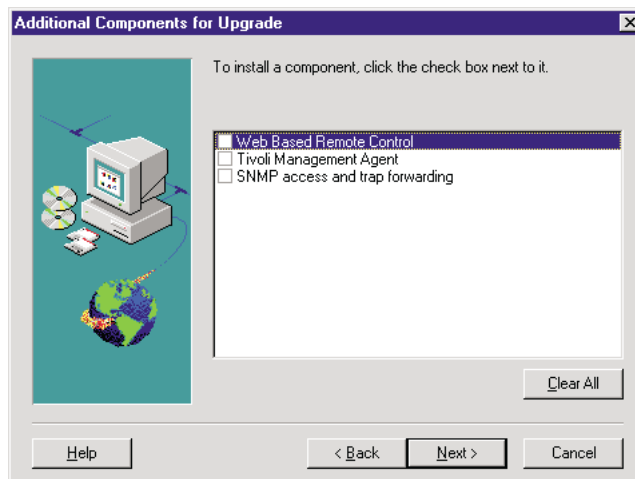
Click **Yes** to begin.

2. Installation detects the client components from the previous version of the installed UM Services client. The Current Install window opens.



Installation upgrades the client components listed in this window. Click **Next** to continue.

3. The Additional Components for Upgrade window opens.



- Select additional components to add to the upgraded components of the UM Services Client. Click **Next** to continue.
4. Depending on which additional components you add, the upgrade program will prompt you for additional responses. For more information about these prompts, see “Installing the UM Services Client for IBM Director” on page 65.
 5. A prompt to add UM Services icons appears. If you have icons from the previous version, clicking **No** will not remove the icons. Clicking **Yes** will not add additional icons to those previously installed.
 6. The program now upgrades the existing components and installs the selected new components. Restart the computer now or later. If you choose **Restart Now**, the system shuts down and restarts immediately. If you choose **Restart Later**, the UM Services upgrade program closes. However, you must restart and log in to the system to begin using IBM Director.

5

Using the Management Console

The IBM Director Management Console is your interface into the IBM 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 IBM 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

IBM Director's operation 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. IBM Director recognizes two types of managed systems:

Native systems

Systems that have the UM Services or Tivoli management agent code installed

SNMP devices

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

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

Starting the Management Console

The IBM Director Management Console is a Java application.

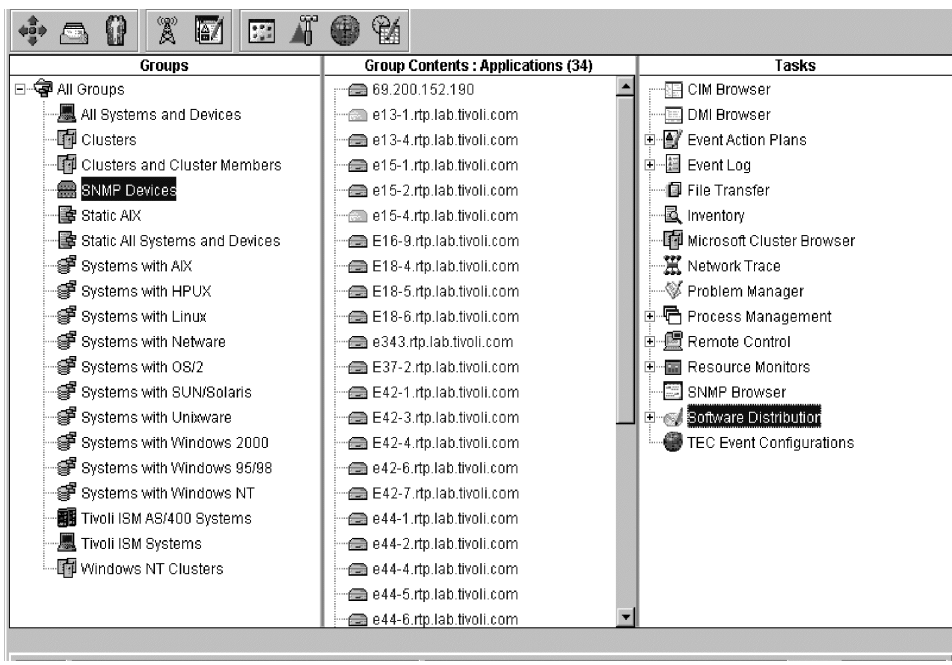
1. To start the IBM Director Management Console, select **Start** → **Director** → **Management Console**.
After the console starts, the IBM Director Login dialog appears.

2. Enter the name of your IBM Director server, your user ID, and your password.

The server name is the TCP/IP host name or the address of the IBM Director server. The user ID and password must be an authorized user account on the IBM Director management server. Your user account maintains your IBM Director Management Console configuration preferences since the last time you were logged in, including status and security settings.

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

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



Getting Around in IBM Director

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

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 IBM Director. Use the method that is most convenient for you.

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.

Using Drag and Drop

Several windows displayed in IBM 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 IBM Director windows.

To execute a task on a single managed system on the IBM Director Management Console (see page 97):

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:

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 IBM Director. Examples include the following:

-
- The file transfer task (see Chapter 11 on page 161). This action enables you to drag files and subdirectories from one system to another.
 - The event management task (see Chapter 9 on page 143). 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 on page 157), 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 Your Mouse's Double-Click Function

You can double-click specific tasks such as Inventory. This action performs an *untargeted* activation of the Inventory task (see Chapter 6 on page 111). 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 almost any task icon or system icon 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 IBM Director product.

Using Add and Remove Buttons

Some windows in IBM Director contain Add and Remove buttons, such as the Inventory Query Builder window (see the figure on page 115).

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

Managing Columns of Information

Many panes of information in IBM 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:
 - Hide a column: click **Hide** in the context menu.
 - Restore a hidden column: Position the mouse pointer over Show Columns in the context menu and click on the column you want to restore.
 - Sort data in a column: click **Sort** 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 IBM 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 IBM 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 Management Console

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

Group Contents

The middle pane in the IBM 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 103).

You can use the drag and drop methods described earlier to perform tasks on managed systems, or select an option from the system's context menu. 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 102), 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 IBM 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 IBM 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.

By default, IBM Director uses no particular association. Based on the presence of AMS modules and LAN tools using the MPM API, other associations are available.

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 IBM Director server with the IBM 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.

If IBM Director detects the presence of Application Management Specification (AMS) enabled applications through an AMS module, then a default group for that particular application would also be created. Furthermore, if IBM Director detects the presence of one or more LAN management tools utilizing the MPM API, a default group for that particular LAN tool would also be created. The IBM Director server default groups are based on the contents of your system.

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 group's criteria.

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, IBM Director automatically updates the group when your network changes. IBM Director adds and deletes managed systems from the group 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 a group's systems. These groups are then designated as *static*. This term means that the IBM Director server does not automatically update the contents of the group.

You can copy a dynamic group into a new static group. IBM 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, follow these steps:

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.

IBM Director dynamically populates the group with all managed systems that meet the specified criteria. When the IBM 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 IBM 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 IBM 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 allows 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. For example, in the Tasks pane, the Software Distribution task can have a secondary task underneath it for a particular software distribution package, such as Lotus Software Distribution.


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

Using the Scheduler feature of IBM 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 Management Console Features

A menu bar appears across the top of the IBM 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 Toolbar

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

- **Discover all systems:** Initiates a discovery of all IBM 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 IBM 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 on page 143 for more details.
- **Scheduler:** Enables you to schedule any non-interactive task to occur at another time, such as software distribution. See Chapter 19 on page 215 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 IBM 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.



Using the Ticker Tape

One feature of the IBM 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 IBM Director Management Console. See Chapter 8 on page 133 for details.



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.

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 IBM 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 “Getting Around in IBM Director” on page 97 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 IBM 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 265 for information on setting up CIM, DMI, and Static MIF files.

Using 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 untargeted mode (for all systems and devices) by double-clicking the Inventory icon. This window is divided into two panes:

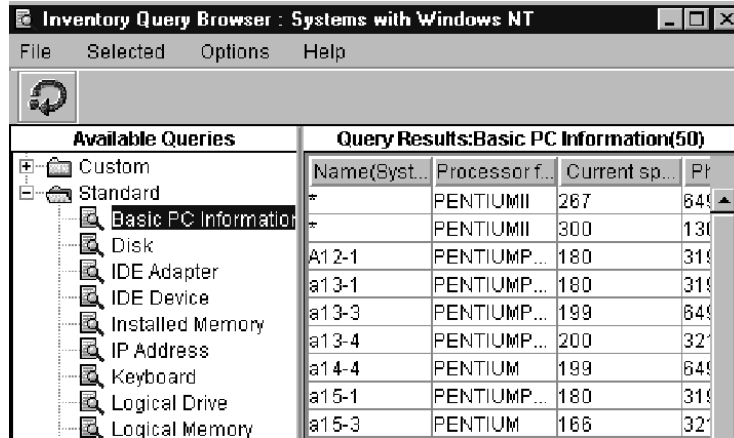
Available Queries

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.



Additional Inventory Query Browser Features

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

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

Custom queries are created by selecting **Build Custom Query** (see “Building a Customized Query” on page 114 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 100. You can save your inventory query results using the standard techniques described in “Saving Files” on page 101.

Using Menu Bar Options

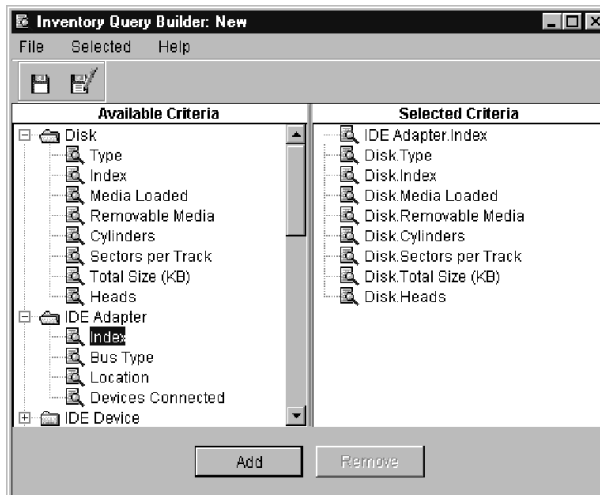
Many of the operations included in the menu bar selections File, Selected, Options, and Help have already been described, such as Perform Query, Refresh, Modify, Copy, Rename, Delete, and Build Custom Query.

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 101. 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 115 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 use the **Add** and **Remove** buttons to create your query in the Selected Criteria pane (see “Using Add and Remove Buttons” on page 100). 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

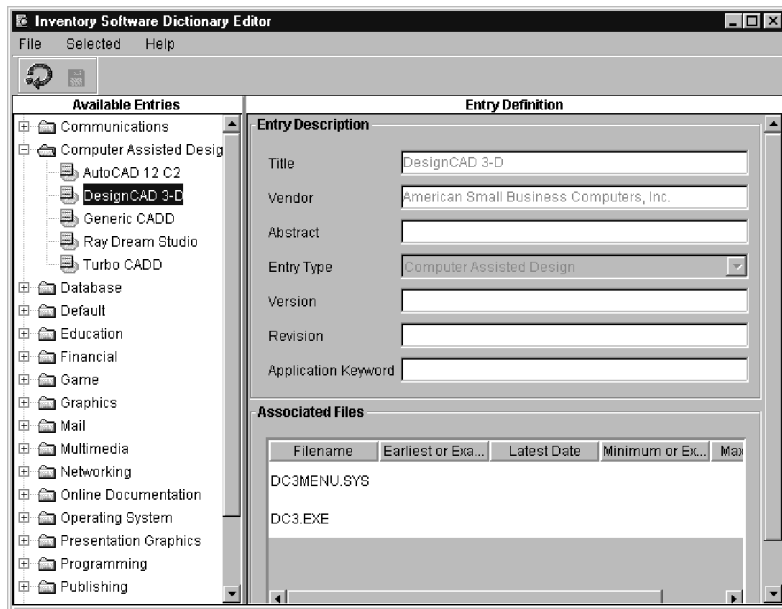
You can use the Inventory Software Dictionary editor to track software on your PC-managed systems.

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.



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

When a **TWGCLI** operation is finished, restart the IBM 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

TWGCLI SWDictionaryReader[*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 **TWGCLI SWDictionaryReader** and associated options.

-dict file

Specifies the path and name of the software dictionary file to be read from. This file must have a filetype of **.sid**.

c:\TivoliWg\Classes\com\tivoli\twg\inventory\default.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 IBM 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

TWGCLI SWDictionaryWriter *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 **TWGCLI SWDictionaryWriter** 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
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), 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\Dictionaryes\Other.sid
```

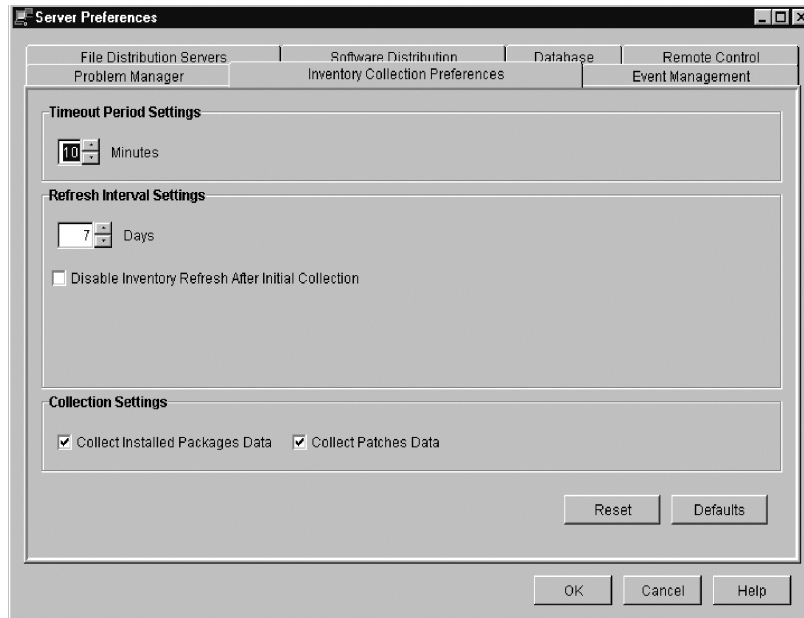
Reads the specified software inventory dictionary (update.sid) and merges its entries with the specified software inventory dictionary (D:\Data\Dictionaryes\Other.sid).

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.



You should select the **Disable Inventory Refresh After Initial Collection** checkbox 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.

7

Remote Control

Remote Control enables you to manage a remote system by displaying the desktop of a remote managed system within a IBM 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.



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 IBM Director Management Console at a remote location. When a IBM Director Management Console assumes control of a managed system in the active state, the screen image of the managed system is displayed within the IBM 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 system's display. The screen image is updated automatically at the active console when a change occurs on the remote system's display.

Monitor state

View-only mode. A managed system in the monitor state is not under the control of a IBM Director Management Console. Consoles attached to the managed system only display the screen image of the managed system and the local user has control of his or her desktop. If a change occurs on the managed system's display, 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 IBM 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 managed system's screen image 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 IBM 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 IBM 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

Managed-system users can change the remote control state of the managed system by pressing **Alt+T** at any time to interrupt the controlling console, terminate or suspend control, and return control to the managed system. When a managed-system user changes the system's control state, the change affects all remote control sessions that are established with the managed system at the time the system's control state is changed.

Control State Scenarios

Scenario 1

Assume that a native managed system is in an active mode and multiple IBM 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 managed system's state automatically changes to 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 IBM 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 IBM Director Management Consoles have remote control sessions with the managed system. If the managed system's user changes the system's state 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 43. Please refer to this section before attempting to perform remote control on your managed systems.

Remote Access Security

During configuration of network drivers, either during the process of installing the Tivoli management agent or by bringing up the Network Driver Configuration window (Start → Programs → IBM Director → Network Driver Configuration, or using the icon in OS/2), the Remote User Authorization for Screen Access option can be enabled. 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 IBM 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 95, the remote system would lock up. An inventory of the managed system tells IBM 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 IBM 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 remote system's cursor movement, but will see screen changes as they occur on the remote managed system's desktop.

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 IBM 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

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 IBM 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 managed system's context menu. See "Getting Around in IBM Director" on page 97 for tips on navigating your way through this task, or see the online help for detailed assistance.

Stopping a Remote Control Session

In addition to using the Alt+T key combination, you can end a remote control session by:

- Closing the remote managed system window.
- Selecting **File** → **Close** from the top of the window.

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 IBM 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 IBM 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 IBM 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.

8

Resource Monitoring

The IBM Director resource monitoring task enables you to view statistics on critical system resources, for example, CPU, disk, file, and memory usage.

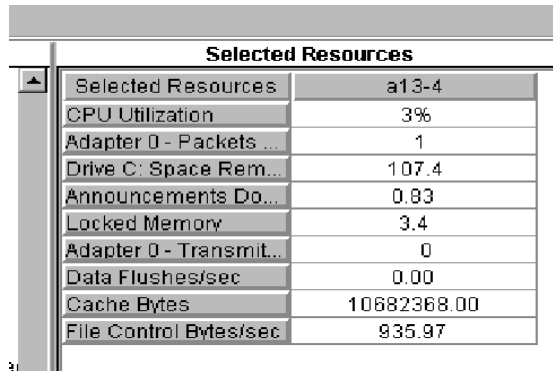
When monitor data indicates a problem or potential problem with network resources, you can set thresholds and trigger events according to the requirements of your site. You can respond to resource monitor events by specifying event action plans. See Chapter 9 on page 143 for more information on event action plans.

Other monitors can be set up to monitor specific processes and system applications. See Chapter 18 on page 207 for details.

Understanding Monitors

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

Monitoring categories are called *attributes*. For example, the performance monitor function is an expandable attribute with subcategories, while CPU utilization is a single attribute with associated data.



Selected Resources	
Selected Resources	a13-4
CPU Utilization	3%
Adapter 0 - Packets ...	1
Drive C: Space Rem...	107.4
Announcements Do...	0.83
Locked Memory	3.4
Adapter 0 - Transmit...	0
Data Flushes/sec	0.00
Cache Bytes	10682368.00
File Control Bytes/sec	935.97

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 machine or application.

The IBM Director server can monitor data from native managed systems, SNMP devices, or Windows NT devices and services.

Monitoring Data on Native Managed Systems

You can monitor data for native managed systems running on remote machines 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 IBM 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 IBM 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. To provide DMI data, managed systems must be running under Windows 95, Windows 98, Windows 2000, or Windows NT 4.0 and must have the Intel V2.0 or V2.0s Service Layer 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 95, 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 machine must be using either IP or Internet Packet Exchange (IPX) transport protocols to communicate with the IBM 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 12 on page 167 for more details.

Monitoring Data on Windows NT Devices and Services

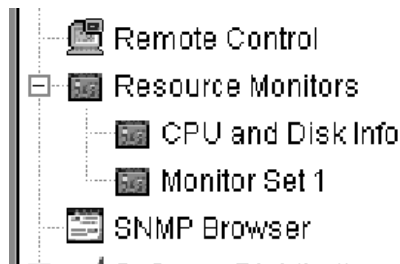
You can also monitor the status of a Win32 device or device service by setting an individual threshold.

Starting Resource Monitors

The Resource Monitors task is started from the IBM Director Management Console by using the standard drag-and-drop methods or by selecting **Resource Monitors** from the managed systems' context menu. (See "Getting Around in IBM Director" on page 97 for tips on navigating your way through this task, or see the online help for detailed assistance.)

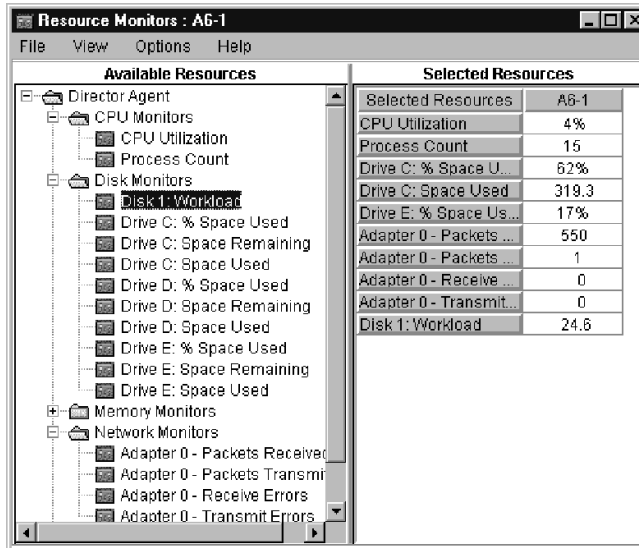
The Resource Monitors task has two subtasks: All Available Recordings and All Available Thresholds. These subtasks provide you with a quick overview of the data recordings and thresholds that have been set, and enable you to perform operations, such as ending a recording or removing a threshold. Refer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



The main part of 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 “Getting Around in IBM Director” on page 97).

Viewing Monitor Data on the Ticker Tape

You can monitor your managed systems from the IBM Director Management Console using the ticker tape display feature (see “Using the Ticker Tape” on page 109 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 on page 143 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.

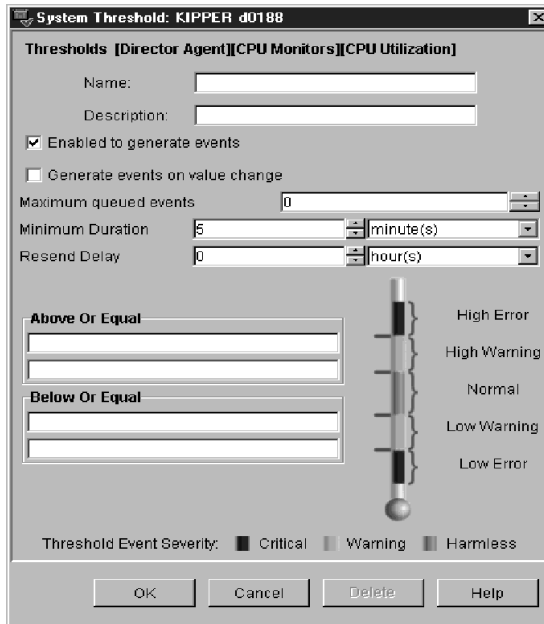
You can set thresholds for a specific managed system. You can also create threshold plans, which are a collection of thresholds. A threshold plan can then be exported to a file that can be imported at a later time for use on other systems or for archival purposes. A threshold plan task allows you to 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, shown below.



Refer to the online help for details on setting thresholds.

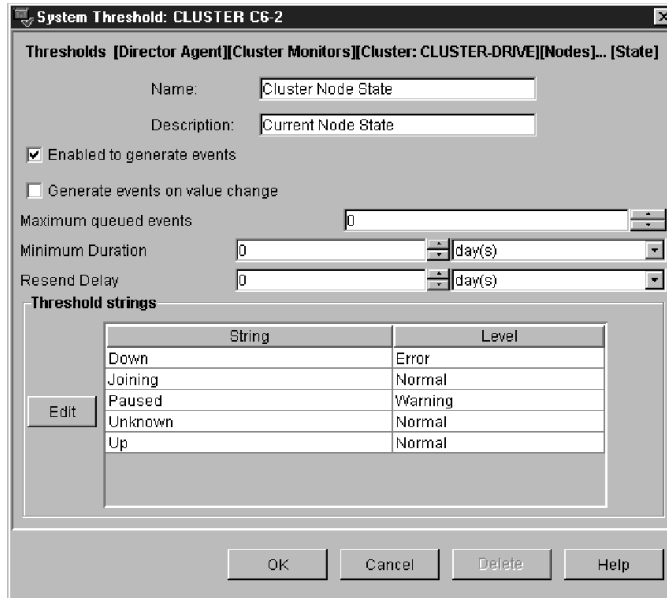
The event type generated is listed at the top of the System Threshold window. In the example shown above, the event type is set to:

[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, shown below.



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

Once you have created a set of monitor attributes in the Selected Resources window, you can save them and re-apply them 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.

9

Event Management

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

For example, you may have used the resource monitor task (see “Chapter 8. Resource Monitoring,” on page 133) to configure a threshold on your file server to generate an event when the remaining free space on the main data drive drops below 100Mb. 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 about your file server’s hard drive approaching its capacity and can take corrective action before your users are impacted.

New Terms in This Chapter

You will see the following new terms 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 is used to select a single event. IBM 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. IBM Director provides a set of predefined actions which 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 IBM 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 management task enables you to:

- Create and apply new event action plans.

Using the Event Action Plan Builder, you can create new event action plans and 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 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.

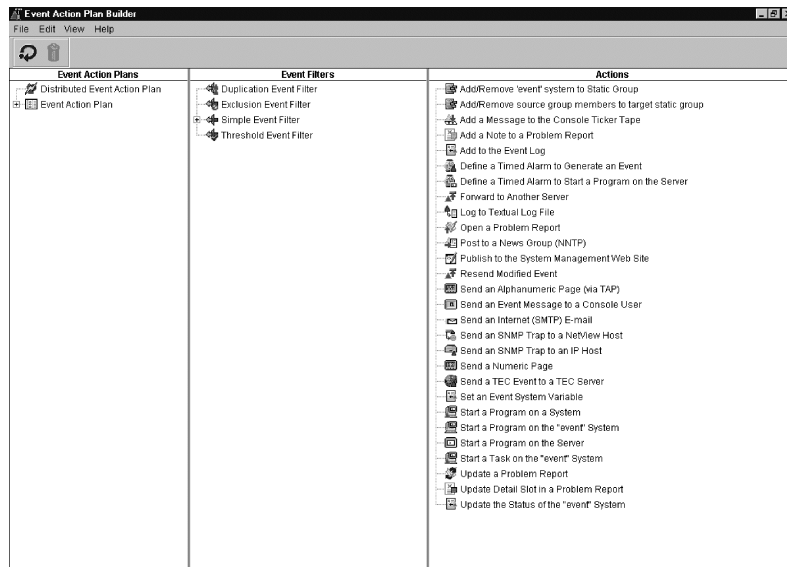
- Log and view event details.

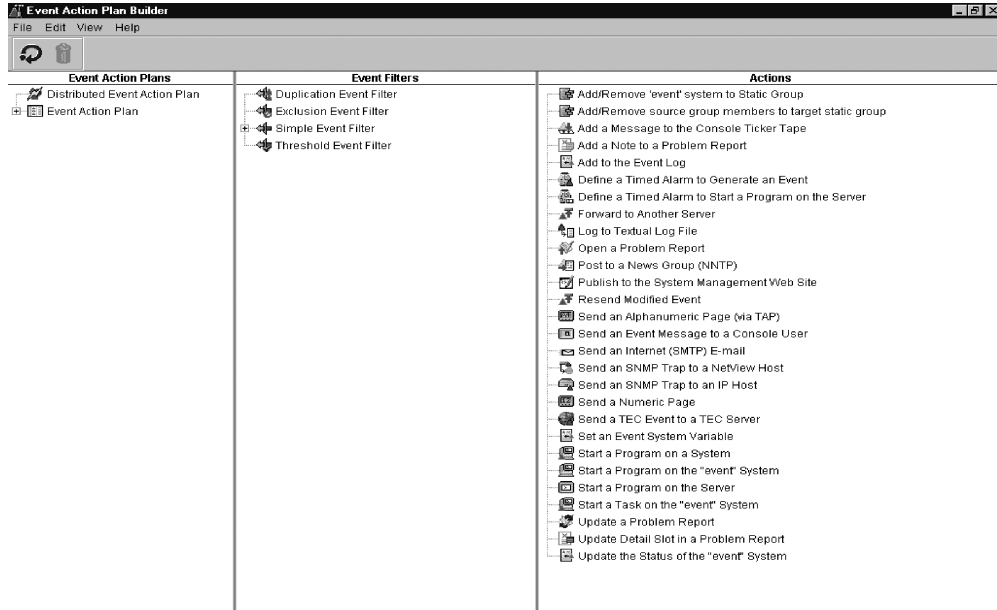
Events are recorded in the Event Log and you can view, sort, and delete these log entries as desired. You can also tailor the view to show only those events that occur on specific systems, or limit the view to only show predefined event action plans.

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 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 IBM 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 (via 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 IBM 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. You can select event filters for viewing the event log and for creating event action plans.

See “Assigning an Event Filter to an Event Action Plan” on page 148 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 from the context menu **New** → . 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 select **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 “Getting Around in IBM Director” on page 97 for tips on navigating your way through this task, or see the online help for detailed assistance).

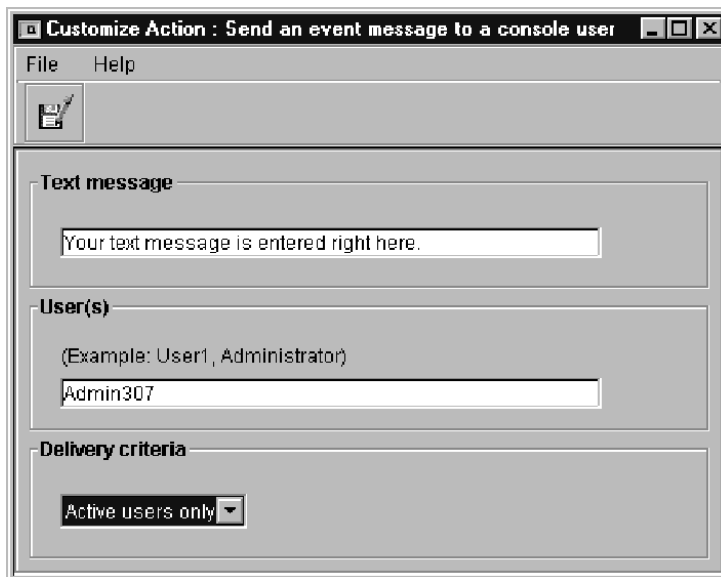
When you add a filter to a plan, the filter icon appears under the event action plan icon in the Event Action Plans pane of the Event Action Plan Builder window.

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.

IBM 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 select 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.



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 on 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 Management Console to the appropriate managed systems in the Management Console.

Displaying Applied Event Action Plans

The Event Action Plans association in the 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, 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 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 IBM Director server.

The event log is started from the Event Log icon in the Tasks pane of the IBM Director Management Console.

The screenshot shows the 'Event Log' window with a menu bar (File, Edit, View, Options, Help) and a toolbar with refresh and delete icons. The main area is titled 'Events (51) - Last 24 Hours' and contains a table of event entries. The right side of the window shows 'Event Details' for the selected entry.

Events (51) - Last 24 Hours						Event Details	
Date	Time	Event Type	Event Text	System No.		Keywords	Values
5/25/1998	12:08 AM	Director Dir...	Monitor 'he...	J2-1		Date	24-May-1998
5/24/1998	11:09 PM	Director Dir...	Monitor 'he...	J2-1		Time	4:43:46 PM
5/24/1998	10:09 PM	Director Dir...	Monitor 'he...	J2-1		Event Type	Director Topology Online
5/24/1998	9:08 PM	Director Dir...	Monitor 'he...	J2-1		Event Text	System 'FLETCHER' is online
5/24/1998	8:08 PM	Director Dir...	Monitor 'he...	J2-1		System Name	FLETCHER
5/24/1998	7:08 PM	Director Dir...	Monitor 'he...	J2-1		Severity	Harmless
5/24/1998	6:08 PM	Director Dir...	Monitor 'he...	J2-1		Category	Resolution
5/24/1998	5:29 PM	Director To...	System 'Tr...	Trantor, H...		Group Name	
5/24/1998	5:08 PM	Director Dir...	Monitor 'he...	J2-1		Sender Path	
5/24/1998	4:43 PM	Director To...	System 'FL...	FLETCHER		Sender Name	
5/24/1998	4:31 PM	Director To...	System 'Mit...	Mitscher-			
5/24/1998	4:10 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	3:36 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 'Gin...	Ginny			
5/24/1998	2:47 PM	Director To...	System 'Gin...	Ginny			

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 IBM 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 20,000; however, the log holds up to 100,000 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				Action Details	
Start Date	Start Time	Action Name	Start Status	Keywords	Values
7/20/2000	2:40 PM	Sample Event Action	Successful	Start Date	7/20/2000
7/20/2000	2:40 PM	Add to the Event Log	Successful	Start Time	2:33 PM
7/20/2000	2:38 PM	Sample Event Action	Successful	Action Name	Add to the Event Log
7/20/2000	2:38 PM	Add to the Event Log	Successful	System Name	EWBROWN
7/20/2000	2:38 PM	Add to the Event Log	Successful	End Status	
7/20/2000	2:37 PM	Add to the Event Log	Successful	Additional Information	
7/20/2000	2:36 PM	Add to the Event Log	Successful	Start Status	Successful
7/20/2000	2:36 PM	Add to the Event Log	Successful	End Date	
7/20/2000	2:36 PM	Add to the Event Log	Successful	End Time	
7/20/2000	2:35 PM	Add to the Event Log	Successful		
7/20/2000	2:33 PM	Add to the Event Log	Successful	Event Type	Director.Topology.Online
7/20/2000	2:33 PM	Add to the Event Log	Successful	Event Severity	Harmless
7/20/2000	2:32 PM	Add to the Event Log	Successful	Event Category	Resolution
7/20/2000	2:22 PM	Sample Event Action	Successful	Event Sender Name	ITSERVER
7/20/2000	1:21 PM	Sample Event Action	Successful	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 “Getting Around in IBM Director” on page 97.

Generating Your Own Events

The IBM 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. **Genevent** must be used from a command prompt on the IBM Director server or a managed system; it is not available through the IBM 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 IBM Director server to which this event will be sent and *name* is the name of the targeted IBM Director management server used by the specified protocol, for example, NETBIOS::TWGSRV1. 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.

TCP/IP is used between this managed system and the IBM Director server to which this event will be sent.

10

Software Distribution

The software distribution task enables you to distribute popular applications and install them on your network's native systems.

Using the software distribution task, you can remotely install applications that use the popular InstallShield installer. You can also remotely install applications which provide a Microsoft Package Definition File (PDF).

This chapter describes how to import and distribute a software distribution package using the IBM Director interface. Before you attempt to distribute a package, make sure you read the guidelines for software distribution in “Software Distribution” on page 35.

Importing a File Package

To import a file package that has been exported, you must use the IBM 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 Management Console and drop it on the desired system icon or group of systems. Only IBM Director managed systems are valid targets for software distribution file packages. Refer to “Chapter 20. Troubleshooting,” on page 227 for more help on distributing software distribution packages.

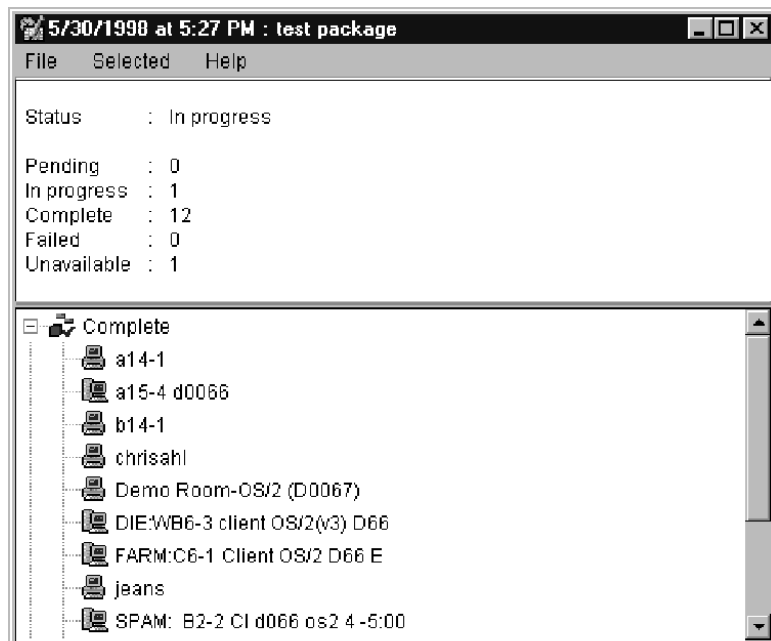
Scheduled Distributions

When you initiate a software distribution task, you are asked if you want 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 “Scheduling Tasks” on page 215 for more information about scheduling software distribution packages.

Immediate Distributions

When you perform an immediate software distribution, the following window is displayed:



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.

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 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 Management Console, right-click on **Software Distribution**, then select **File Distribution Servers Manager** in the context menu. For more information, refer to the online help.

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 IBM Director Management Console local system and the IBM Director server system
- The IBM Director Management Console local system and a native managed system
- The IBM 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 IBM Director Management Console or IBM Director server, and then send that file to a different managed system.

Using the File Transfer Task

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. It is used 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

Bring up the File Transfer console from the IBM Director Management Console by double-clicking the task or by using normal drag-and-drop techniques. Refer to “Getting Around in IBM Director” on page 97 for tips on navigating your way through this task, or see the online help for detailed assistance.

IBM 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 IBM 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 IBM 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 IBM 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 (e.g. .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 for 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. Highlight a file or subdirectory in the source pane (local system or IBM Director server).
2. Highlight 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 IBM Director management server to the target drive or subdirectory.

Transfer File(s) to Source

1. Highlight a drive or subdirectory in the source pane (the remote system or server).
2. Highlight 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 IBM 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. Only release 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 IBM 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 dialog box 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, do the following:

1. Select a source object as explained in “Starting a File Transfer Session” on page 162.
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 equal (synchronized).

Notes:

- 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.
- 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 IBM Director server or managed system are mapped using a different username or password than the user name/password specified for the IBM Director service during installation (that is, the username/password of the IBM 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 an entire drive's contents by dragging the drive icon. 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 *and* 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 IBM 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 IBM 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.

12

SNMP Management

IBM Director includes Simple Network Management Protocol (SNMP) support that enables you to isolate SNMP devices for the event management, inventory, and resource monitor services. For information on using the IBM Director Management Console to specify SNMP devices for these tasks, see the following chapters:

- “Chapter 6. Inventory Management,” on page 111
- “Chapter 8. Resource Monitoring,” on page 133
- “Chapter 9. Event Management,” on page 143.

IBM 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, IBM Director provides a MIB compiler that enables you to specify and compile MIBs that are not supplied by IBM 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 IBM Director Services

IBM Director recognizes MIBs in the System Management Information (SMI) Version 1 format. IBM 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 IBM Director compile the first time the IBM Director management server starts. Additional MIBs may be compiled as needed from the IBM Director Management Console.

Performing SNMP Tasks

From the IBM 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 IBM Director server
- Invoke the SNMP browser to view SNMP-formatted data

Understanding SNMP Discovery

IBM 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 IBM Director. Refer to “Installing the IBM Director Server on Windows” on page 49 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 your network’s Domain Name Server, the address of the machine that acts as your network’s router, 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 IBM 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 IBM 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 IBM 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.

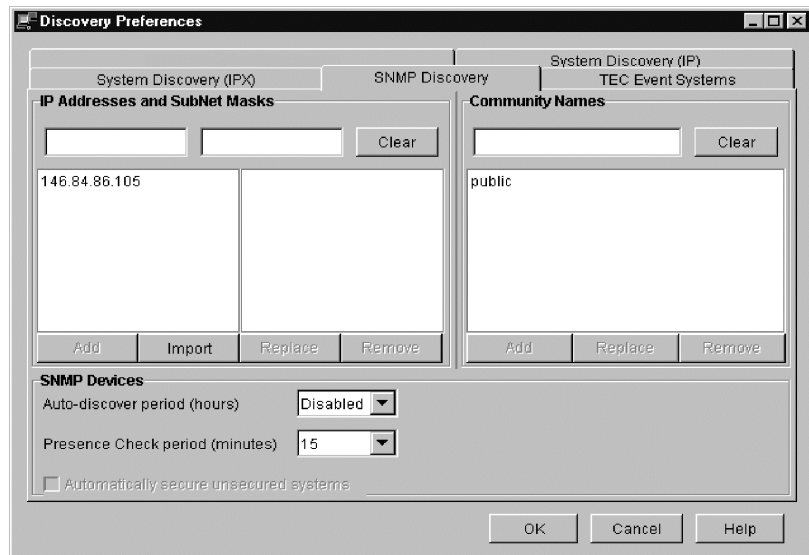
Using your lists of IP addresses, community names, and subnet masks, a series of SNMP GET statements are performed against port 161 of the IP address to determine if the address is a valid SNMP device of some kind. If it is determined to be a valid SNMP device, another series of

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 IBM Director, applying the community names as appropriate.

Setting SNMP Discovery Parameters

From the menu bar of the IBM Director Management Console, select **Options** → **Discovery Preferences**. When the Discovery Preferences window is displayed, select the **SNMP Discovery** tab.



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 your network's router, 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 IBM Director.

In the Group Contents pane of the IBM Director Management Console, select **New** → **SNMP Devices** from the context menu. The Add SNMP Devices window is displayed.



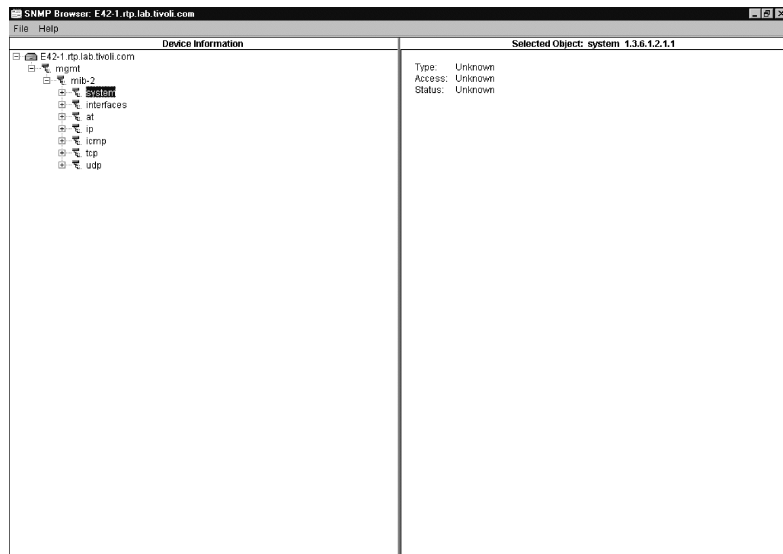
Select either the IP or IPX network transport, and then enter the network address. For IP, the dotted decimal address must be specified. Specify a

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.

Using the SNMP Browser

You can view the attributes of SNMP and RMON devices using the SNMP Browser.



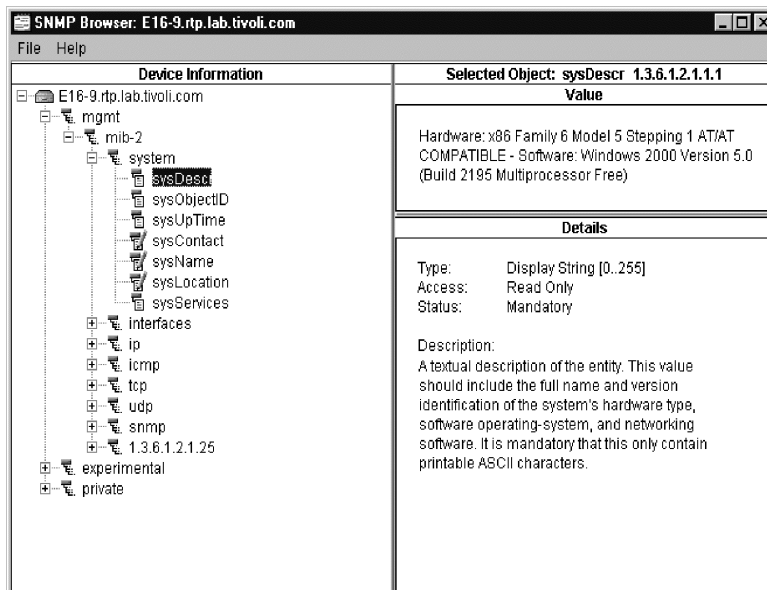
Starting the SNMP Browser

You can start the SNMP Browser from the IBM 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 IBM 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.



In the Device Information pane, information is displayed in a tree view of the devices. 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 of the device's network connections.

To open a multi-homed device, right-click on the device on the **Group Contents** pane and then select **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.

13

DMI Management

IBM 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. For information on DMI, see information on the Web at <http://www.dmtf.org>.

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 111
- “Chapter 8. Resource Monitoring,” on page 133
- “Chapter 9. Event Management,” on page 143

The 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” on page 176. 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 95, Windows 98, Windows 2000, or Windows NT 4.0. They must have the Intel V2.0 or V2.0s Service Layer installed. The Service Layer does not have to be present when the IBM Director management agent is installed. The Service Layer can be added to a managed system after IBM 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, follow these steps:

1. Right-click the **Groups** pane of the IBM 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 IBM 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, follow these steps:

- 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, IBM 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, follow these steps:

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 system's DMI tree 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 grayed out and its DMI tree cannot be expanded.
2. Double-click a system to display the components of the system, then click on a component to display descriptive information in the right-hand pane.
 3. To view the group classes of a component, double-click on the component name.
 4. To view the attributes of a group class, click on the group class name. A description of the group class appears in the upper right-hand 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-click on a line item and select **Sort** → **Descending**.
 6. When you have finished viewing information, select **File** → **Close** to close the dialog.

Setting an Attribute Value for a DMI Group

It is strongly recommended that you do not change an attribute's value unless you are thoroughly familiar with the structure and manipulation of DMI data. Improperly setting a system's value can cause unpredictable results on the target system.

To change an attribute value, follow these steps:

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 177.
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 IBM Director to attempt to change the value, cancel the window.

If IBM 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. Once 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, follow these steps:

1. In the IBM Director Management Console, apply the DMI Browser task to a managed system to display the DMI Browser window.
2. Double-click on the managed system to display the associated components.
3. Double-click on a component to display the contained group classes.
4. Right-click on the group class name to display the context menu, and select **Create task for group class**. A dialog 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, select **OK**. The new task is entered as a subtask under DMI Browser in the IBM 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:

- 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 ComponentId 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.

-
- 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.

14

CIM Management

IBM 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 111
- “Chapter 8. Resource Monitoring,” on page 133

Unlike DMI events, CIM events are not automatically detected by IBM Director. The IBM Director Software Development Kit provides information on how to set up managed systems to map CIM events to IBM Director events. When the mapping file is defined, IBM Director can detect and present CIM events for filtering.

The Management Console window does not automatically display CIM-enabled systems as a separate group of systems. To create a dynamic group of CIM-enabled systems, follow the procedure described in “Creating a CIM Dynamic Group” on page 182. You can also use the Static Group Editor to create a group of systems that are CIM-enabled in your network.

CIM Requirements

To provide CIM data, managed systems must be running under Windows ME, Windows 95, 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 IBM Director management agent is installed. You can add WMI to a managed system after IBM Director is installed. When the managed system restarts, it is enabled for CIM operations.

Creating a CIM 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 CIM-enabled.

To create a dynamic group for CIM-enabled systems, follow these steps:

1. Right-click in the Groups pane of the IBM Director Management Console window to display the context menu.
2. Select **New Task Based** to display the Task Group Editor.
3. In the Available Resources pane, select **CIM Browser** and click **Add** to add the selection to the Selected Criteria pane. Selecting CIM Browser creates a filtering criteria for managed systems that are CIM-enabled.
4. Click **Save As** to save the new group with a name of your choosing. A dialog appears to name the group.
5. Enter a descriptive name for the group, for example, “CIM-enabled systems.”
6. Select **Close Group Editor** to save your group and exit the dialog.
7. Refresh the IBM Director Management Console window with a discovery operation. The new group appears in the Groups pane.
8. Select your new group to see which managed systems match the CIM criteria. The Group Contents pane lists any discovered CIM-enabled systems.

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, IBM 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, follow these steps:

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 on 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 a property's value unless you are thoroughly familiar with the structure and manipulation of CIM data. Improperly setting a system's value can cause unpredictable results on the target system.

To change a property's value, follow these steps:

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 183.

2. Right-click the value on the property 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 IBM Director to attempt to change the value, cancel the window.

If IBM 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, follow these steps:

1. Using the procedure described in “Starting the CIM Browser and Viewing Information” on page 183, navigate to the class that has the method you want to execute. The associated methods appear under the Methods tab in the lower right-hand pane.
2. Right-click a method and select **Execute** from the context menu. The Execute Method window appears.
3. If the method receives any input arguments, one or more Input fields appear. Enter the arguments in these fields.
4. Select **Execute** at the bottom of the Execute Method window to execute the method. If you do not want to execute the method, cancel the window. If IBM Director is unable to execute 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, follow these steps:

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 183.
2. Right-click anywhere on the class name and select **Create browser task for class**. A dialog appears 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, select **OK**. The new subtask is entered under CIM Browser in the IBM 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, follow these steps:

1. Use the procedure described in “Starting the CIM Browser and Viewing Information” on page 183 to navigate to the CIM class that has the method for which you want to create a subtask. The associated methods appear under the Methods tab in the lower right-hand pane.
2. Right-click on a method and select **Execute** from the context menu. The Execute Method window appears.
3. If the method receives any input arguments, one or more Input fields appear. Enter the arguments in these fields.
4. Select **Save** at the bottom of the Execute Method window. A dialog appears 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, select **OK**. The new subtask is entered under CIM Browser in the IBM 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 execute.

Because method subtasks are non-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 “Scheduling Tasks” on page 215 for information on task scheduling.

15

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.

You retrieve Asset ID information from the UM Services agent installed on any Director-managed system. The UM Services agent polls Asset ID data from IBM systems that have the Enhanced Asset Information Area EEPROM.

The screenshot shows a window titled 'Asset ID' with several tabs: 'Serialization', 'System', 'User', 'Lease', 'Asset', 'Personalization', and 'Warranty'. The 'Asset' tab is selected, displaying a table with the following data:

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	19658810U	IBM
Motherboard	NDB70300052	IBM

At the bottom of the window, it says 'Data space remaining: 596' and there is an 'Apply' button.

Note: UM Services 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.

The Asset ID Interface

The Asset ID interface is activated as a Director task. Upon activation, the Asset ID window opens to display the data polled from an IBM Asset ID equipped system or an other DMI-enabled system.

The Asset ID interface contains the following tabbed interfaces:

Tab	Description
Serialization	Click the Serialization tab to display serial numbers for the client system hardware.
System	Displays the current client system characteristics: system name, MAC address, user login name, operating system, GUID address, IBM LAN Client Control Manager (LCCM) 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 client system hardware.
Asset	Displays the inventory factors that are related to the client 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 client system hardware.

At the bottom of the Asset ID window is the **Data space remaining:** *nnn* information line. This information is an indicator of the amount of remaining available data storage on the EEPROM. *nnn* represents this

storage as a number of characters that can be entered. Once the limit has been reached, the Data space remaining line turns red. At this point, any further information entered is discarded.

Click the **Apply** button 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 on-line help window.

Serialization

Click the **Serialization** tab to display the Serialization interface. The Serialization interface displays the serial numbers for the various components in the client 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** tab to display the System interface. The System interface displays information about the client system.

The screenshot shows a software interface with a tabbed menu at the top containing 'Serialization', 'System', 'User', 'Lease', 'Asset', 'Personalization', and 'Warranty'. The 'System' tab is active, displaying the following information:

- 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: A text input field containing 'YY'.

At the bottom left, it says 'Data space remaining: 1017'. At the bottom right, there is an 'Apply' button.

The System interface provides information about the items that are described in the following table.

Item	Description
System Name	<p>The NetBEUI name of the client system (the computer name as it appears under Network Properties).</p> <p>NetBEUI is NetBIOS extended user interface, and NetBIOS is network basic input/output system.</p>
MAC Address	The unique hexadecimal character string that identifies the network adapter in the client 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 UM Services resides).

Item	Description
System GUID	The client 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** tab to display the User interface. The User interface displays information about the logged-in user.

The screenshot shows a web-based user interface. At the top, there is a navigation bar with several tabs: 'Serialization', 'System', 'User' (which is currently selected), 'Lease', 'Asset', 'Personalization', and 'Warranty'. Below the navigation bar, there are five input fields, each with a label to its left: 'Name', 'Phone', 'Location', 'Department', and 'Position'. At the bottom left of the form area, there is a red text message that reads 'Data space remaining: -21'. At the bottom right, there is an 'Apply' button.

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.

Item	Description
Location	The user office location.
Department	The user department name or number.
Position	The user job title.

Lease

Click the **Lease** tab to display the Lease interface. The Lease interface displays lease information for the client system.

The screenshot shows a software interface with several tabs: 'Serialization', 'System', 'User', 'Lease', 'Asset', 'Personalization', and 'Warranty'. The 'Lease' tab is selected. Below the tabs, there are several input fields:

- Start Date:** A date picker showing 'September', '27', and '2000'.
- End Date:** A date picker showing 'September', '27', and '2000'.
- Term (Months):** A text input field containing the number '0'.
- Amount:** An empty text input field.
- Lessor:** An empty text input field.

 At the bottom of the interface, it displays 'Data space remaining: 975' and an 'Apply' button.

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.

Item	Description
Term (months)	The number of months for which the client system is leased.
Amount	The total price of the lease agreement.
Lessor	The name of the company that leased the client system.

Asset

Click the **Asset** tab to display the Asset interface. The Asset interface displays inventory information about the client system.

The screenshot shows a software interface with several tabs: 'Serialization', 'System', 'User', 'Lease', 'Asset', 'Personalization', and 'Warranty'. The 'Asset' tab is selected. Below the tabs, there are two rows of date pickers. The first row is for 'Purchase Date', with a dropdown menu set to 'January', a spinner box set to '1', and another spinner box set to '1999'. The second row is for 'Last Inventoried', also with a dropdown menu set to 'January', a spinner box set to '1', and another spinner box set to '1999'. Below these are two text input fields: 'Asset Number' and 'RF-ID'. At the bottom left, there is a red text message: 'Data space remaining: -21'. At the bottom right, there is an 'Apply' button.

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 client system was purchased.
Last Inventoried (mm/dd/yy)	The date of the last inventory check.

Item	Description
Asset Number	A unique number that is assigned to the client system for inventory purposes.
RF-ID	The radio-frequency identification (RF-ID) number that was encoded in the client system by the manufacturer. Not all computers have RF-ID capabilities. This is a fixed field and cannot be changed.

Personalization

Click the **Personalization** tab 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.

The screenshot shows a software interface with a tabbed menu at the top containing 'Serialization', 'System', 'User', 'Lease', 'Asset', 'Personalization', and 'Warranty'. The 'Personalization' tab is active. Below the menu is a form with two columns: 'Label' and 'Value'. Each column contains five empty text input fields. At the bottom left of the form, there is a red text indicator that reads 'Data space remaining: -21'. At the bottom right, there is an 'Apply' button.

Warranty

Click the **Warranty** tab to display the Warranty interface. The Warranty interface displays information about the warranty on the client system.

The Warranty interface provides information about the items that are described in the following table.

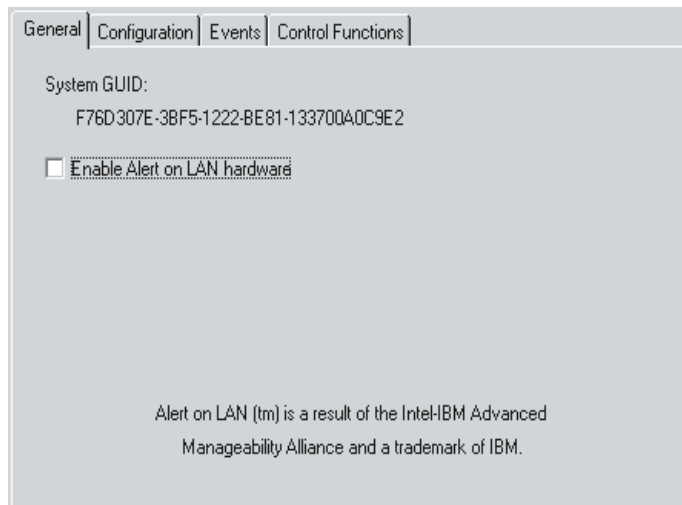
Item	Description
Duration (months)	The duration of the warranty agreement.
Cost	The total cost of the warranty.
End Date (mm/dd/yy)	The date that the warranty ends.

16

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.

When you select **Alert on LAN** from the IBM Director Tasks pane and apply it to an applicable system, the system displays the following screen.



The following items are available on the Alert on LAN screen.

Item	Description
General Tab	
System GUID	<p>A Global Unique ID (GUID) is assigned to each system board for system-management purposes.</p> <p>The GUID is stored in the BIOS on the system board.</p>
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 Tab	
Proxy server (IP address port)	The internet protocol address for the server you use to communicate with the client systems. The IP address is assigned by the system administrator. (default port is 5500.)
Heartbeat timer period	<p>The Alert on LAN proxy computer verifies that the client system is running. This is the number of seconds between system checks. The default value is 32.</p> <p>The enabled heartbeat timer period values range from 43 to 5461 seconds and can be set in intervals of 43 seconds.</p>
Watchdog Timer Period	<p>If the watchdog timer indicates that a client 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.</p> <p>The watchdog timer period values range from 86 to 5461 seconds and can be set in intervals of 86 seconds.</p>

Item	Description
Transmission attempts	The number of retries for transmission after the client 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 client system is disconnected from the LAN, even when the computer is off. If a client 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 client system that is an event message is generated.
Auto-clear events	If this option is enabled, the client 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.

Item	Description
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.

17

Cluster Management

IBM 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,” on page 133
- “Chapter 9. Event Management,” on page 143

Understanding Cluster Management

In IBM 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.

IBM Director supports only the clustering implementation of Microsoft Clustering Service (MSCS). For Windows NT systems configured with MSCS, a IBM 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 IBM 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 4.0 with Service Pack 3 or 4 and must have Microsoft Clustering Service installed.
- Each node in a cluster should have the IBM 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 IBM Director Management Console displays clusters as both the Clusters and Windows NT Clusters groups. Because IBM 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 IBM 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 IBM Director by using the guidelines in the IBM 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.

18

Process Management

IBM 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.

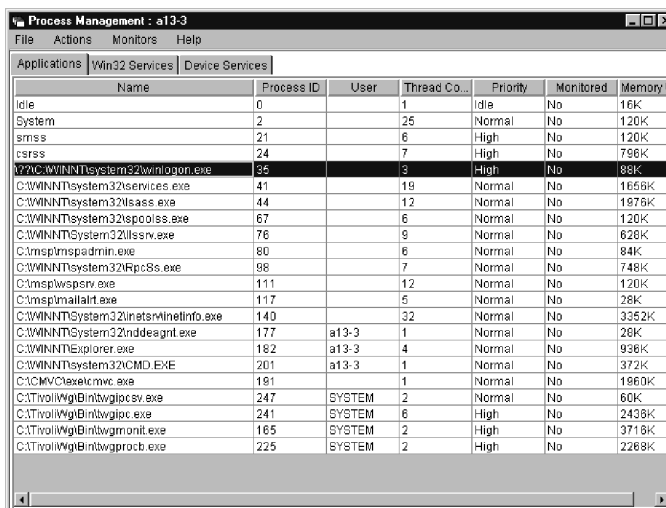
The process management task is an interactive task that applies only to native managed systems. SNMP devices do not have the capability to be monitored and managed to this level of detail.

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 systems

Starting the Process Management Task

You can start the main Process Management window from the IBM Director Management Console using drag and drop and context menu techniques (refer to “Getting Around in IBM Director” on page 97).



The screenshot shows the 'Process Management' window for host 'a13-3'. It has a menu bar with 'File', 'Actions', 'Monitors', and 'Help'. Below the menu bar are three tabs: 'Applications', 'Win32 Services', and 'Device Services'. The 'Applications' tab is active, displaying a table of running processes. The table has columns for Name, Process ID, User, Thread Co., Priority, Monitored, and Memory. The processes listed include system processes like 'Idle', 'System', 'smss', 'csrss', and various system services like 'W3C:\WINNT\system32\winlogon.exe', 'C:\WINNT\system32\services.exe', 'C:\WINNT\system32\lsass.exe', 'C:\WINNT\system32\spoolss.exe', 'C:\WINNT\system32\lsrv.exe', 'C:\msp\mspadmin.exe', 'C:\WINNT\system32\Fpc8s.exe', 'C:\msp\wspcvr.exe', 'C:\msp\mallat.exe', 'C:\WINNT\System32\inetser\inetinfo.exe', 'C:\WINNT\System32\nddeagnt.exe', 'C:\WINNT\Explorer.exe', 'C:\WINNT\system32\CMD.EXE', 'C:\CMV\Chasek.msc.exe', 'C:\Trolli\qgl\Bint\wgipc sv.exe', 'C:\Trolli\qgl\Bint\wgipc.exe', 'C:\Trolli\qgl\Bint\wgmomil.exe', and 'C:\Trolli\qgl\Bint\wgprocb.exe'.

Name	Process ID	User	Thread Co.	Priority	Monitored	Memory
Idle	0		1	Idle	No	16K
System	2		25	Normal	No	120K
smss	21		6	High	No	120K
csrss	24		7	High	No	796K
W3C:\WINNT\system32\winlogon.exe	36		3	High	No	88K
C:\WINNT\system32\services.exe	41		19	Normal	No	1659K
C:\WINNT\system32\lsass.exe	44		12	Normal	No	1976K
C:\WINNT\system32\spoolss.exe	67		6	Normal	No	120K
C:\WINNT\system32\lsrv.exe	76		9	Normal	No	628K
C:\msp\mspadmin.exe	80		6	Normal	No	84K
C:\WINNT\system32\Fpc8s.exe	98		7	Normal	No	748K
C:\msp\wspcvr.exe	111		12	Normal	No	120K
C:\msp\mallat.exe	117		5	Normal	No	28K
C:\WINNT\System32\inetser\inetinfo.exe	140		32	Normal	No	3352K
C:\WINNT\System32\nddeagnt.exe	177	a13-3	1	Normal	No	26K
C:\WINNT\Explorer.exe	182	a13-3	4	Normal	No	936K
C:\WINNT\system32\CMD.EXE	201	a13-3	1	Normal	No	372K
C:\CMV\Chasek.msc.exe	191		1	Normal	No	1860K
C:\Trolli\qgl\Bint\wgipc sv.exe	247	SYSTEM	2	Normal	No	60K
C:\Trolli\qgl\Bint\wgipc.exe	241	SYSTEM	6	High	No	2439K
C:\Trolli\qgl\Bint\wgmomil.exe	165	SYSTEM	2	High	No	3716K
C:\Trolli\qgl\Bint\wgprocb.exe	225	SYSTEM	2	High	No	2289K

All operating systems' Process Management windows contain an Applications tab. Windows NT contains 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 and 2000 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 NT System and Device Services” on page 213 for more information.

Device Services

(Windows NT and 2000 only) Enables you to interact with Windows NT 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 NT System and Device Services” on page 213 for more information.

Notes:

- Not all services can be controlled in this manner.
- 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 on page 133 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 kilobytes (KB), for the selected system.

Subsystem

Identifies the subsystem in which a job is running.

Viewing Windows NT Services Information

For Windows NT 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 machine as either administrator or root. There are provisions on Windows NT and Unix 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 version of IBM Director do not use this feature.

For Windows NT systems

To enable or disable this feature for Windows NT, modify the Windows NT 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 **RestrictAnonCmdExe**.
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**. Windows NT saves your registry entry.

For Unix systems

To require users to enter ID and password on UNIX systems, do the following in a UNIX 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.



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 on page 133 for details on viewing your process monitor.

Controlling NT System and Device Services

You can use the Process Management task to start, stop, pause, and resume system services on Windows NT 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 IBM 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” for more information on setting thresholds.

19

Task Scheduler

The task scheduler feature of IBM 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 IBM Director server.

Interactive tasks require direct interaction with a user and cannot be scheduled. Examples of interactive tasks include remote control and file transfer.

Scheduling Tasks

To schedule a task from the IBM 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/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 “Scheduling Tasks” on page 215) 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 server time matches the IBM 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

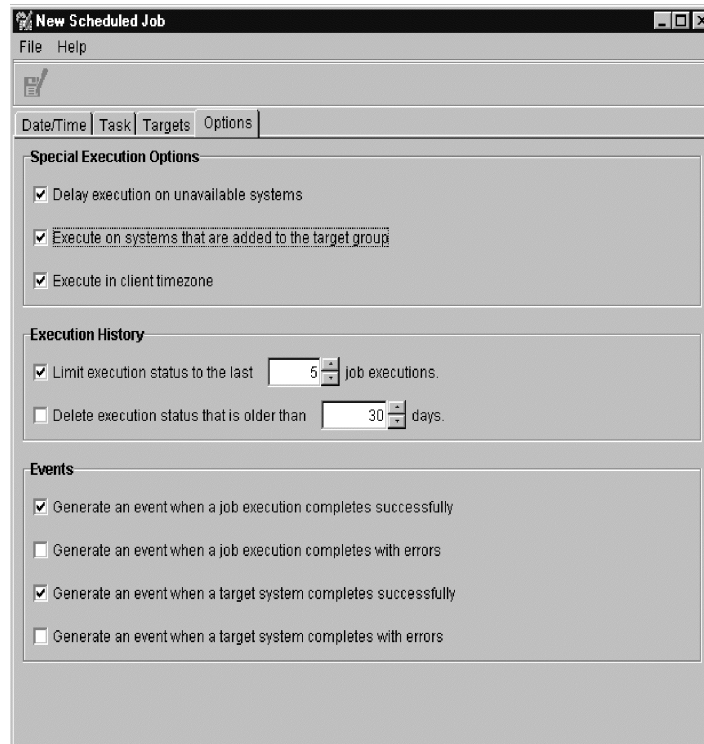
This option enables you to select a group from a list of all groups. You can select only one group. If you select a second group, it replaces the first group.

- 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.



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 client time zone

Delay execution on unavailable systems

When You Do Not Check This Option

When this box is not checked, 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 box is checked, 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 check 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 check 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 client 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:

- You cannot create a job to repeat hourly and be executed in the managed system's time zone.

-
- One job activation record per 24 hours is created when the Execute in client time zone option is selected. The job activation dynamically updates as managed-system clients move from pending to active when their time-zone window occurs.
 - If the first scheduled time zone start date occurs before the server date, the job cannot be created.
 - 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 toolbar. 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 IBM Director Management Console using the Scheduler task in the menu bar or the Scheduler icon in the toolbar. 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 216 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 216 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 216 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.

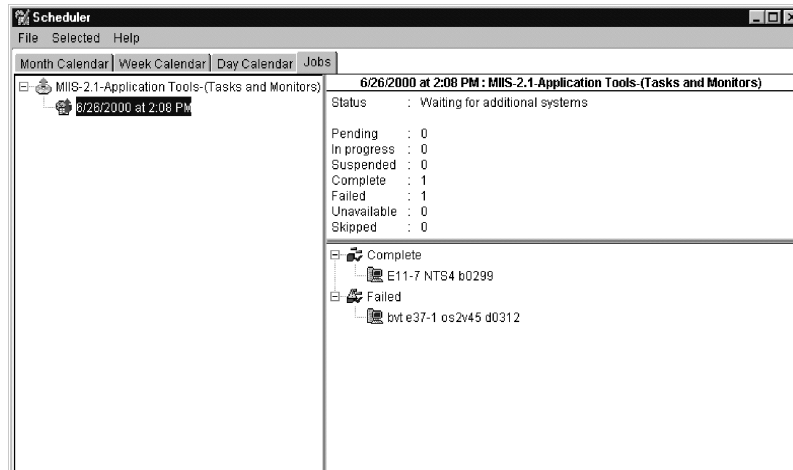
When viewing either log, you can control the level of detail displayed by using the menu bar options. By default, the log displays the lowest level of detail.

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.



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 223 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 “Scheduling Tasks” on page 215), and an Execution History window shows information about the execution. See “Viewing Scheduled Jobs” on page 223 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.

20

Troubleshooting

While every effort has been made to provide you with a simple and easy to use interface, you might find problems when running IBM Director.

Your IBM 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 IBM 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 Microsoft SQL Server database. For information on switching database support from Jet to SQL, see:

- Chapter 2 on page 13
- Chapter 3 on page 49.

Q: Why are my component installations failing even though I have verified that sufficient space is available?

A: IBM 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	Install Space Required (in bytes)	Uninstall Space Required (in bytes)
Novell Agent	1,420,331	N/A
Windows 98 Agent	4,999,233 with TMA; 3,727,506 without TMA	72,192
Windows NT or Windows 2000 Agent	5,073,303 with TMA; 3,727,506 without TMA	72,192
Windows 98 Console	1,420,331	72,192
Windows NT or Windows 2000 Console	1,420,331	72,192
Windows NT or Windows 2000 Server	3,727,506	2,706,431

Q: Why are some SNMP devices not being discovered?

A: Verify that the IBM 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 IBM 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 IBM Director's Discovery Preferences window allow IBM 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 IBM 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 IBM 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 won't IBM Director allow me to change a MIB value?

A: Check the following:

- Verify that IBM Director is using a community name that allows write access to the MIB you wish to set.
- Verify that the MIB is writable. IBM 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: IBM Director describes setting a particular MIB value to a hexadecimal/octal/binary value, but it will not accept my number. Why?

A: IBM Director expects all values to be added in decimal. You must convert the number from hexadecimal/octal/binary to decimal.

Q: What protocols does IBM Director use for sending and receiving SNMP traps?

A: This version of IBM 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 IBM Director server resides, seed devices must be configured. Note that:

- You should use only one system for each subnet.

-
- The IBM Director server must be able to ping the seed devices.
 - The seed devices must be able to ping the IBM Director server.

Configure these from the IBM Director Management Console window. Select **Options** → **Discovery Preferences** → **IBM Director System Discovery (IP)**.

In addition, discovery requires that any routers or bridges between the IBM 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 IBM 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 IBM Director server. Another method would be that seed devices are configured. Note that:

- You should use only one system for each network.
- The IBM 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 IBM 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 IBM Director server is running under Windows NT Service Pack 4, the symc810.sys device driver is probably causing the blue screen. Reinstall 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 IBM Director management server is completely up. Check the IBM 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 IBM 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 IBM 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 IBM Director management server. By default, you cannot manage other IBM 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 IBM 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: I cannot unzip log files that were archived using the Logfile Management AMS tasks. Why?

A: Most unzip utilities cannot handle long file names. Rename the zip files so that they have no more than eight characters in the file name.

Q: Why do some of my managed systems become unavailable on the console?

A: Perhaps the timeout value for IBM Director to access the system needs to be increased. Modify your Network Time-Out value in the Network Driver Configuration window (select **Start Programs** → **IBM 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 IBM 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 TCP/IP causes data to flow through a different port than the one that IBM 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 IBM 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 IBM 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 IBM Director Management Console. When a package is deleted and the package has been cached on a share, then IBM 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 IBM Director server and should not be modified by a user. If a software package directory is deleted through a means other than the IBM 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 IBM Director agent, specifying a different drive and directory.

Q: Why am I getting flooded with AMS events, originating from thresholds?

A: AMS events do not rearm. If a monitored attribute is continuously meeting a threshold value, each time the attribute is monitored, an event is generated. To remedy this, configure a threshold response that will resolve the condition.

Q: Why is IBM Director not starting up?

A: This is usually due to security issues. If the administrator's password has changed, then you must change the password for the IBM Director Support Program, in the Services section of the Control Panel.

You must have logged in with an administrator's ID when installing IBM 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 IBM Director was installed.

Q: Why are the groups created by AMS empty?

A: When an AMS component is committed, the software dictionary is updated with an entry for that component. The new entry won't be discovered on the client until the next time inventory is run. To populate the groups, run an inventory collection, either manually or through the Scheduler, after committing AMS components.

Q: Why are my redistributed installs not working properly?

A: If IBM AntiVirus is installed on the IBM 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 SC_{SI}_DE_{VICE} table that also have a value of TA_{PE} in the DE_{VICE}_TY_{PE} column will be returned.

Likewise, if you select Inventory (PC) / SC_{SI} Device / Device Type ^= TA_{PE} as a criterion, only managed systems that appear in at least one row of the SC_{SI}_DE_{VICE} table, of which none of those rows have a value of TA_{PE} in the DE_{VICE}_TY_{PE} column, will be returned. **It does not necessarily return all managed systems that do not have SC_{SI} 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 / SC_{SI} Device / Device Type ^= TA_{PE}
- Inventory / Operating System / Type = WINDO_{WS} NT

Using these criteria, a Windows NT managed system with no SC_{SI} devices would not be returned, because such a managed system would not appear in the SC_{SI}_DE_{VIC}ES table. However, if a Windows NT managed system had a SC_{SI} hard drive but no SC_{SI} tape drive, it would be returned, because such a system would appear in the SC_{SI}_DE_{VIC}ES 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.



Resource Monitor Attributes

This appendix contains a list of the attributes that can be monitored by the IBM Director resource monitors task on managed systems that have the Tivoli management agent installed.

Monitor collection rates are every 5 seconds, unless otherwise noted.

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:

- The disk drive monitors will repeat for each local non-removable logical drive found.

-
- 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:

- Only for IBM systems
- 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:

- For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
- 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

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:

- Each registry entry is an attribute.
- 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:

- The number of applications or executables monitored by the process monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.
- 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:

- The disk drive monitors will repeat for each local non-removable logical drive found.
- 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:

- Only for IBM systems
- 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:

- For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
- 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

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:

- Each registry entry is an attribute.
- 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:

- The number of applications or executables monitored by the process monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.
 - 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

Sentry Monitors

Note: The Sentry Monitors depend upon which AMS packages have been installed and have variable data collection refresh rates.

Windows 95 Operating System

CPU Monitors

- CPU Utilization
- Process Count (1 minute refresh rate)

Disk Monitors

Notes:

- The disk drive monitors will repeat for each local non-removable logical drive found.
- Monitor data collection rate is every 60 seconds.
- Disk 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:

- For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
- 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

Performance Statistics

Note: The Windows 95 Performance statistics are dynamic and may be different on each machine.

File System

- Bytes read/second
- Bytes written/second
- Dirty data
- Reads/second
- Writes/second

IPX/SPX compatible protocol

- IPX packets lost/second
- IPX packets received/second
- IPX packets sent/second
- Open sockets
- Routing Table entries
- SAP Table entries
- SPX packets received/second
- SPX packets sent/second

Kernel

- Process usage (%)
- Threads
- Virtual Machines

Memory Manager

- Allocated memory
- Discards
- Disk cache size
- Free memory
- Instance faults
- Locked memory
- Maximum disk cache size
- Minimum disk cache size
- Other memory
- Page faults
- Page-ins
- Page-outs
- Swapfile defective
- Swapfile in use
- Swapfile size
- Swappable memory

Microsoft Client for NetWare

- BURST packets dropped
- Burst received gap time
- Burst send gap time
- Burst send gap time
- Bytes in cache

-
- Bytes read/second
 - Bytes written/second
 - Dirty bytes in cache
 - NCP packets dropped
 - Requests pending

Microsoft Network Client

- Bytes read/second
- Bytes written/second
- Number of nets
- Open files
- Resources
- Sessions
- Transactions/second

Process Monitors

Notes:

- The number of applications or executables monitored by the Process Monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each one of the attributes under Process Monitors will be present for each executable being monitored.
 - 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

Registry Monitors

Notes:

- Each registry entry is an attribute.
- Monitor data collection rate is every 60 seconds.

Sentry Monitors

Note: The Sentry Monitors depend upon which AMS packages have been installed and have variable data collection refresh rates.

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 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:

- The disk drive monitors will repeat for each local non-removable logical drive found.
 - 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:

- For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
- 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 Swapfile Monitors

- Swap File Size
- Swap Space Remaining

Process Monitors

Notes:

- The number of applications or executables monitored by the Process Monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each one of the attributes under Process Monitors will be present for each executable being monitored.

-
- 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

Sentry Monitors

Note: The Sentry Monitors depend upon which AMS packages have been installed and have variable data collection refresh rates.

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:

- The disk volume monitors will repeat for each volume detected on a NetWare Server.
- 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:

- For compatible file system types, the Directory Exists or File Exists attribute (depending on which is applicable) should always be a valid datapoint.
- 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:

- The number of applications or executables monitored by the Process Monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each

one of the attributes under Process Monitors will be present for each executable being monitored.

- 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:

- The list of file systems will appear first; the following attributes will appear under each file system.
- 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:

- The monitor attributes listed below 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.
- 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:

- If there are additional directories, additional subelements will be present.
- It is possible that directories that contain a large number (more than several hundred) of subelements 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:

- The number of applications or executables monitored by the process monitors is variable and configured by the IBM Director administrator from the Process Manager console. Each of the attributes under Process Monitors will be present for each executable being monitored.
 - 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

Sentry Monitors


Note: The Sentry Monitors depend upon which AMS packages have been installed and have variable data collection refresh rates.

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




Creating the ODBC Entry for the Default Database

Use the following steps to manually create the default Microsoft Jet database:

1. Shut down the IBM Director server and ensure that you are logged on with the IBM 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 IBM 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 IBM Director installation directory (for example, **c:\Tivoliwg\Data**) with the following entry:**twg.database.odbc.name=Director**.
 13. Restart the IBM Director server and perform an inventory collection to fill the database.



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 IBM Director server you should have specified that you wanted to use the default Microsoft Jet database that ships 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 261.

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 13, 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 49, and the online help.

- **cfgmssql** - Microsoft SQL Server Database
- **cfgdb2** - IBM DB2 Universal Database
- **cfgoracle** - Oracle Server Database
- **cfgmsde** - Microsoft Data Engine (MSDE) Database.



Defining Table Property Files

This appendix contains information on setting up your server to inventory CIM and DMI information.

Setting up the Server to Inventory CIM and DMI Information

IBM 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" below.

Because the created tables may be viewed in any locale supported by IBM 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."

In addition to table property files, you must supply files that specify associations between IBM 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.” Without these files, IBM 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 IBM 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.”

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.

Table Property File Format

A Table Property File can be created and edited using any ASCII editor. These files are placed in the UserTables subdirectory of the server's data directory; this path will usually be C:\TivoliWG\Data\UserTables. 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 = IBM 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:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789-_.)
```

This restriction is based on restrictions set by the supported DBMSs.

table.token: name of the table used internally in the IBM 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; displaynames 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 IBM 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 IBM 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 IBM 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 IBM 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; pt_br_win for Brazilian Portuguese, 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 IBM 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 indices 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 filename (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 IBM 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.c  
olumnVa  
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.TWGDdbUserTable?tableTokenName =
translated string
for “all true” for this table
```

To specify the filter prompt string for “Any True:”

```
FilterTablePrompt.AnyTrue.TWGDdbUserTable?tableTokenName =
translated string
for “any true” for this table
```

To specify the filter prompt string for “All True For Same:”

```
FilterTablePrompt.AllTrueForSame.TWGDdbUserTable?tableToken
Name =
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.TWGDdbUserTable?tab
leTok
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:\TivoliWG\Data\InvExtension. 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 must 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 only 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. 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.

Note: CIM arrays will be handled using the following method: If the CIM property is a STRING, all STRINGS in the array will be appended together, space separated, into a single string. If the CIM property is any other type, only the first value will be stored.

CIM Type	Default Database Type
EMPTY STRING	CHAR
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	CHAR
DISPLAYSTRING	
DATATYPE_0	INT
COUNTER	
COUNTER64	
GAUGE	
DATATYPE_4	
INTEGER	
INTEGER64	
DATATYPE_9	
DATATYPE_10	
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 client 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 clients 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:\TivoliWG\bin). 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 IBM 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.”

Server Initialization and Table Property Files

When the IBM Director server starts, it searches the UserTables subdirectory of the server's data directory (usually C:\TivoliWG\Data\UserTables) and loads all of the user table files, which have the extension.TWGDbt, that it finds. It is important to know that IBM 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 IBM 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 IBM 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 the IBM Director 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.DESCRPTION=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=M
anufac
```

```
turer English
```

```
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.PRODUCT=Produc
```

```
t
```

```
English
```

```
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.VERSION=Versio
```

```
n
```

```
English
```

```
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.SERIAL_NUMBER=
```

```
Serial
```

```
Number English
```

```
ColumnName.TWGDbUserTable?DMI_WIN32_COMP_ID.INSTALL_DATE=I
```

```
nstall
```

```
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



Agent-Server Security

Agent-server security is an authentication process used to establish trust relationships between the IBM 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, and provides guidelines for:

- Initializing managed systems securely
- Determining the origin of public or private keys
- Recovering lost public or private key files

How IBM Director Agent-Server Security Is Implemented

IBM Director provides a means of security by which a managed system configured with Director management agent (agent) can authenticate an IBM 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 IBM Director server, which controls the administrator's ability to issue requests to the IBM Director server and agents through the IBM 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 IBM Director server. If an agent is *unsecure*, *any* IBM Director server is allowed to manage the system. If the agent is *secure*, only IBM Director servers that pass authentication are allowed to manage the system.

Public-private authentication is the method used by an IBM Director agent to authenticate an IBM Director server once the agent is secure. IBM Director's 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 IBM 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 IBM 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 9x, 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 IBM 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 superset of the agent) are stored in files named **dsaxxxx.pub**, where **xxxx** 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 **dsaxxxx.pvt**. When an IBM 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 IBM 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 IBM 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 IBM 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 IBM Director Management Console.

Installing IBM Director Agents in a Secure State

The IBM Director management console supports a "request access" function to initiate an access request from the IBM 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 IBM Director server to a Director management agent in a secure state. Refer to the online help for more information.

To install IBM Director agents in a secure state, use the following procedure:

1. Install and start any IBM 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\UMS\Director\data for Windows Clients`, `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 3.1, Windows 9x, 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, **twgipc** on Windows 3.1, **start twgipc** on Windows 9x, **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
- `DSAPstring` for public key files
- `DSAPstring` 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 wouldn't 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.

Index

A

- Account Information window 59
- Action History window
 - using 153
- actions
 - assigning to filters 149
 - definition 144
- activating event action plans 150
- active state 126
- add and remove buttons
 - using 100
- agent software
 - installing on Netware 74
 - installing on OS/2 71
- Alert on LAN 199
- application information
 - viewing 209
- applications
 - closing process management tasks 213
- applying event action plans 148, 150
- Asset ID
 - Asset tab 195
 - Lease tab 194
 - Personalization tab 196
 - Serialization tab 191
 - System tab 191
 - User tab 193
 - Warranty tab 196
- Asset tab 195
- associations 102
- attributes
 - resource monitor 239

B

- basic services, IBM Director 67

C

- calendar tabs
 - using 222
- columns of information
 - managing 100
- communication protocols 8
- configuration
 - configuring to use file distribution servers 79
- configuration settings
 - changing 85
- configuration tasks
 - Alert on LAN 199
- context menus
 - using 100
- control states
 - overriding and changing 127
 - remote control 126
- controlling
 - device services 213
 - NT system 213
- customizing
 - your scheduled job 216
 - understanding the special execution options 219
 - using the Date/Time tab 216

D

- Data/Time tab 216
- database
 - converting from Jet to SQL server 263, 265, 287
 - creating the ODBC entry 261
 - database, selecting a 59

-
- default database, ODBC entry 261
 - definitions
 - actions 144
 - event 143
 - event action plans 144
 - event filters 144
 - event management 9
 - file transfer 10
 - inventory management 9
 - native 1
 - process management 10
 - remote workstation control 9
 - resource monitoring 9
 - RMON 1
 - SNMP 1
 - software distribution 10
 - task scheduling 10
 - deleting
 - file packages 160
 - Director Database Configuration window 59
 - Director management console
 - associations 102
 - group contents 101
 - groups 103
 - hardware and software requirements 23
 - installing 64
 - overriding and changing control states 127
 - performing SNMP tasks from 168, 177, 183
 - remote control tasks 125
 - starting 96
 - starting resource monitors 136
 - tasks 106
 - Director management server
 - configuring distribution preferences 82
 - database support 13
 - hardware and software prerequisites 13
 - Director transport
 - software requirements 24
 - directories
 - synchronizing 164
 - distributing
 - file packages 157
 - DMI Mapper 56, 69
 - drag and drop 98
 - drives
 - synchronizing 164
- ## E
- event
 - definition 143
 - event action plan builder
 - creating new plans 145
 - event action plans
 - activating 150
 - applying 148, 150
 - assigning actions to event filters 149
 - assigning event filters to 148
 - creating new 145
 - customizing an action 148
 - definition 144
 - managing 151
 - saving 150
 - using predefined filters 147
 - event filters
 - assigning actions to 149
 - assigning to event action plans 148
 - customizing an action 148
 - definition 144
 - using predefined 147
 - event log
 - using 151
 - viewing all logged events 152
 - viewing events by filter characteristics 153
 - event management 143
 - definition 9
 - understanding 144
 - events
 - generating your own 154
 - SNMP traps 46
-

executing commands 51
 creating non-interactive tasks 211
 on selected systems 211

F

FAQs 227
features of
 IT Director management console 108
file distribution servers
 configuring IT Director to use 79
file transfer
 additional features 165
 between managed systems 163
 definition 10
 selecting files for 162
 starting a session 162
 task 161
 using 161
files
 saving 101
 synchronizing 164
filter characteristics
 viewing by 153

G

generating
 your own events 154
group contents 101
groups 103

H

hardware requirements, IBM Director
 Console 51
hardware requirements, IBM Director
 Management Server 51
hardware requirements, UM Services Client

I

IBM Director basic services 67
IBM Director Console, supported systems
 50
IBM Director Installation
 server files 53
IBM Director Management Server, supported
 systems 50
importing
 file packages 157
initiating a resource monitor 137
installation
 information to consider beforehand 13
 prerequisites xix
 redirected 36
 restrictions xix
 streamed 36
 tasks 49
installing agent software on Netware 74
installing agent software on OS/2
 attended installation 71
 unattended installation 73
installing IBM Director
 Server 53
Inventory
 and remote control 129
inventory
 collecting 111
Inventory database tables 47
inventory management 111
 definition 9
Inventory Query Browser
 building a customized query 114
 manging your query results 113
 menu bar options 114
 updating the list of available queries 113
 using 112
Inventory Query Builder 114

-
- Inventory Software Dictionary Editor 115 129
- IT Director
- additional features of
 - database management 12
 - security 12
 - concept of managed systems 95
 - configuring to use file distribution servers 79
 - how it works 2
 - introduction 1
 - monitoring tasks or services 101
 - navigating 97
 - uninstalling components on OS/2 88
- IT Director management console
- additional features 108
 - using 95, 101
- IT Director management server
- additional considerations 21
- IT Director services
- MIB requirements 168
- J**
- Java environment
- restrictions with remote control 130
- Jet
- converting to SQL server database 263, 265, 287
- jobs
- managing 222
 - saving 222
 - scheduling 216
- jobs tab
- using 224
- K**
- keyboard arrow keys
- using 101
- keyboard information to a remote system
- L**
- LANDesk Management Suite Integration 55, 68
- Lease tab 194
- logged events
- viewing 152
- M**
- main components 2
- management agent 2
 - management console 2
 - management server 2
- managed object
- native clients 2
 - SNMP devices 3
- managed systems
- concept in IT Director 95
- managing
- columns of information 100
 - event action plans 151
 - hardware and software inventory 111
- managing scheduled jobs
- using the calendar tabs 222
 - using the jobs tab 224
- MIB requirements
- for IT Director services 168
 - for SNMP browser 168, 176, 182, 204
- monitor console
- using 137
- monitor data
- recording 140
 - viewing datas on the ticker tape 138
- monitor resources
- managing 141
- monitor state 126

monitor thresholds
 setting 138

monitoring
 IT Director tasks or services 101

monitoring data
 on native managed systems 134
 on SNMP devices 135

monitors 133
 adding new 213

mouse functions
 double-click function 99

N

native
 definition 1

native clients
 managed object 2
 managing your network with 8

native managed systems
 monitoring data 134

navigation 97

NetWare xxi

network
 preparing 13

Network Driver Configuration window 62

network drivers, configuring 62

O

ODBC entry 261

operating systems
 platforms supported 24

OS/2 xxi

P

password 57

password, setting 57

Personalization tab 196

planning tasks 13

process management
 definition 10
 task 207

process management task
 adding new process monitors 213
 closing applications 213
 controlling NT system and device
 services 213
 creating non-interactive tasks to execute
 commands 211
 executing commands on selected
 systems 211
 removing process monitors 214
 starting 208
 viewing application information 209
 viewing Windows NT services 211

process monitors
 adding new 213
 removing 214

Q

queries
 building customized queries 114
 managing your results 113
 updating list of available 113
 using menu bar options 114

R

redirected install 36

remote access security 128

Remote Control
 active state 126
 and inventory function 129
 control states 126
 monitor state 126
 restrictions 130

-
- sending keyboard information 129
 - suspend state 126
 - task 125
 - tasks 130
 - usage restrictions 128
 - remote control
 - definition 9
 - recording a session 131
 - refresh rate 132
 - remote network monitor 1
 - removing
 - process monitors
 - monitors
 - removing 214
 - resource monitor attributes 239
 - resource monitoring 133
 - definition 9
 - resource monitors
 - initiating a resource monitor 137
 - managing your monitor resources 141
 - recording monitor data 140
 - setting monitor thresholds 138
 - starting 136
 - using the monitor console 137
 - viewing monitor data 138
 - restrictions
 - usage restrictions of remote control 128
 - RMON
 - definition 1
- S**
- saving
 - event action plan 150
 - files 101
 - your scheduled job 222
 - scheduling jobs 216, 222
 - security
 - remote access 128
 - selecting files for transfer 162
 - Serialization tab 191
 - server files
 - basic services 67
 - installing 53
 - Simple Network Management Protocol 1
 - SNMP
 - definition 1
 - SNMP access and trap forwarding 56, 69
 - SNMP browser
 - MIB requirements 168, 176, 182, 204
 - starting 172
 - using 172, 185
 - viewing information 173
 - SNMP clients
 - managing 11
 - SNMP devices
 - creating new 171, 184
 - monitoring data 135
 - SNMP discovery
 - understanding 168
 - SNMP discovery parameters
 - setting 170
 - SNMP management
 - MIB requirements 168, 176, 182, 204
 - performing tasks 168, 177, 183
 - task 167, 175, 181, 203
 - understanding 167, 175, 181, 203
 - SNMP tasks
 - creating a new SNMP device 171, 184
 - setting SNMP discovery parameters 170
 - understanding SNMP discovery 168
 - using the browser 172, 185
 - SNMP trap 46
 - Software Distribution
 - definition 10
 - deleting file packages 160
 - distributing file packages 157
 - importing file packages 157
 - redirected install 36
 - streamed install 36
 - task 157
-

-
- special execution options
 - when customizing a scheduled job 219
 - SQL Server database
 - converting from Jet 263, 265, 287
 - starting
 - file transfer session 162
 - process management task 208
 - SNMP browser 172
 - startup command 64
 - streamed install 36
 - suspend state 126
 - synchronizing files, directories, or drives 164
 - System Health
 - temperature out of specification 201
 - voltage out of specification 201
 - System Health Monitoring 55, 68
 - System tab 191
 - systems
 - executing commands on 211
- T**
- tables
 - inventory database 47
 - task scheduler
 - customizing your scheduled job 216
 - managing your scheduled job 222
 - saving your scheduled job 222
 - task 215
 - viewing immediate execution information 225
 - task scheduling
 - definition 10
 - tasks
 - Asset ID
 - Asset tab 195
 - Lease tab 194
 - Personalization tab 196
 - Serialization tab 191
 - System tab 191
 - User tab 193
 - Warranty tab 196
 - event management 143
 - file transfer 161
 - installation and configuration 49
 - inventory management 111
 - planning 13
 - process management 207
 - product introduction 1
 - remote control 125, 130
 - resource monitoring 133
 - SNMP management 167, 175, 181, 203
 - software distribution 157
 - task scheduler 215
 - troubleshooting 227
 - using the IT Director management console 95
- tasks pane 106
 - ticker tape
 - using 109
 - Tivoli Management Agent 55, 68
 - Tivoli management agent
 - hardware and software requirements 23
 - transferring files between managed systems 163
 - troubleshooting 227
 - TWGIPC startup command 64
- U**
- UM Services Client, supported systems 50
 - uninstalling IBM Director 87
 - user ID, setting 57
 - User tab 193
 - using
 - action history window 153
 - add and remove buttons 100
 - context menus 100
 - drag and drop 98
 - event log 151

file transfer task 161
Inventory Query Browser 112
Inventory Software Dictionary Editor
115
IT Director management console 101
keyboard arrow keys 101
ticker tape 109
your mouse's double-click function 99

V

viewing
all logged events 152
application information 209
by filter characteristics 153
immediate execution information 225
SNMP information 173
Windows NT services information 211

W

Warranty tab 196
Web-based Access 54, 67
Web-based Remote Control 55, 68
Windows xxi
Windows NT
controlling NT system and device
services 213
viewing services information 211