

AppManager Operator Console User Guide

NetIQ® AppManager®

March 2007



Legal Notice

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, NETIQ CORPORATION PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of NetIQ Corporation, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of NetIQ Corporation. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. NetIQ Corporation may make improvements in or changes to the software described in this document at any time.

© 1995-2007 NetIQ Corporation, all rights reserved.

U.S. Government Restricted Rights: If the software and documentation are being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), in accordance with 48 C.F.R. 227.7202-4 (for Department of Defense (DOD) acquisitions) and 48 C.F.R. 2.101 and 12.212 (for non-DOD acquisitions), the government's rights in the software and documentation, including its rights to use, modify, reproduce, release, perform, display or disclose the software or documentation, will be subject in all respects to the commercial license rights and restrictions provided in the license agreement.

Check Point, FireWall-1, VPN-1, Provider-1, and SiteManager-1 are trademarks or registered trademarks of Check Point Software Technologies Ltd.

ActiveAgent, ActiveAnalytics, ActiveAudit, ActiveReporting, ADcheck, AppAnalyzer, AppManager, the cube logo design, Directory and Resource Administrator, Directory Security Administrator, Domain Migration Administrator, Exchange Administrator, File Security Administrator, IntelliPolicy, Knowing is Everything, Knowledge Scripts, Mission Critical Software for E-Business, MP3check, NetConnect, NetIQ, the NetIQ logo, NetIQ Change Administrator, NetIQ Change Guardian, NetIQ Compliance Suite, NetIQ Group Policy Administrator, NetIQ Group Policy Guardian, NetIQ Group Policy Suite, the NetIQ Partner Network design, NetIQ Patch Manager, NetIQ Risk and Compliance Center, NetIQ Secure Configuration Manager, NetIQ Security Administration Suite, NetIQ Security Analyzer, NetIQ Security Manager, NetIQ Vulnerability Manager, PSAudit, PSDetect, PSPasswordManager, PSSecure, Server Consolidator, VigilEnt, Vivinet, Work Smarter, and XMP are trademarks or registered trademarks of NetIQ Corporation or its subsidiaries in the United States and other jurisdictions. All other company and product names mentioned are used only for identification purposes and may be trademarks or registered trademarks of their respective companies.

Contents

About This Guide

Intended Audience	vii
What's Changed?	viii
Conventions	viii
Using Online Help	viii
Other Information in the Library	ix
About Attachmate	x

Chapter 1

Welcome

Why Monitor Your Environment?	13
How Does AppManager Help?	14
What Is the AppManager Suite?	15
What Does AppManager Do?	15
How Does AppManager Work?	16

Chapter 2

Getting Started

Starting the Operator Console	27
About the Operator Console	31
Connecting to Another Repository	35

Chapter 3

Configuring the Operator Console

Customizing the Operator Console Layout	38
Selecting a View	39
Identifying the Computers to Monitor	42

Adding Computers to the TreeView Pane	43
Organizing Computers into Groups	48
Finding Objects in the TreeView Pane	53
Deleting an Object from the TreeView Pane	55
Expanding and Collapsing the Network Hierarchy	58
Viewing Detailed Information about Discovered Resources	60
Viewing AppManager System Information	61
Filtering Information Based on a Resource Object	63
Grouping and Filtering Based on Other Criteria	63
Setting Preferences	63
Displaying Toolbars	91
Displaying a Status Bar	91
Getting Information about Knowledge Scripts	92
Checking Communication with Managed Computers	92
Integrating AppManager with a Custom Web Application	93
Using the Extensions Menu	95

Chapter 4

Running Monitoring Jobs

About Knowledge Scripts and Jobs	103
Starting a New Job	107
Setting the Schedule	110
Setting the Values to Monitor	113
Specifying One or More Corrective Actions	115
Selecting the Objects to Monitor	120
Setting Advanced Job Properties	121
Customizing Knowledge Scripts	130

	Checking and Changing Job Status	137
	Initiating Maintenance on a Computer	145
	Grouping and Filtering Job Information	146
	Viewing Job Results	149
	Closing a Job	151
	Deleting a Job	152
Chapter 5	Responding to Events	
	How Events Are Raised	153
	Event Severity Levels	155
	Viewing Events	157
	Viewing Details about an Event	161
	Adding Comments to an Event	166
	Grouping and Filtering Events	168
	Changing the Status of an Event	170
	Deleting an Event	174
	Duplicate Events	174
	Automatically Closing Events When the Condition No Longer Exists	180
Chapter 6	Creating Graphs	
	Collecting Data for Graphs	185
	Viewing the List of Graph Data Streams	188
	Grouping and Filtering Data Streams	189
	Creating a Graph	193
	Viewing and Working with Graphs	198

Chapter 7	Using Snapshot Views	
	About Snapshot Views	207
	Creating a Snapshot View	208
	Updating a Snapshot View	211
Chapter 8	Using Dynamic Views	
	Understanding Dynamic Views	216
	What Kinds of Views Can You Create?	217
	Understanding How Rules Work	217
	Understanding Custom Properties	230
	Creating a Dynamic View	232
	Organizing Computers into Groups.	237
	Tips and Tricks	243
Chapter 9	Monitoring by Policy	
	About Policy-Based Monitoring	249
	Creating a Knowledge Script Group	255
	Changing a Knowledge Script Group	259
	Creating a Monitoring Policy	264
	Changing a Monitoring Policy	267

Index

About This Guide

The NetIQ AppManager Suite (AppManager Suite) is a comprehensive solution for managing, diagnosing, and analyzing performance, availability, and server health for a broad spectrum of operating environments, applications, and server hardware.

AppManager provides system administrators with a central, easy-to-use console to view critical server and application resources across the enterprise. With AppManager, administrative staffs can monitor computer and application resources, check for potential problems, initiate responsive actions, automate routine tasks, and gather performance data for real-time and historical reporting and analysis.

Intended Audience

This *AppManager Operator Console User Guide* provides information for system and network administrators working with the AppManager Operator Console.

For information about using the Control Center Console and the other AppManager console applications, including the Chart Console and the Operator Web Console, see the *Control Center User Guide*.

For information about configuring and maintaining an AppManager site, see the *Administrator Guide*.

What's Changed?

This release of AppManager includes many enhancements to functionality and new features. To learn what's new and for other late breaking information, see the `ReleaseNotes.htm` located in the `\Release Notes` folder of the AppManager installation kit.

Conventions

This guide uses consistent conventions to help you identify items throughout the documentation. The following table summarizes these conventions.

Convention	Use
Bold	<ul style="list-style-type: none">• Window and menu items• Technical terms, when introduced
<i>Italics</i>	<ul style="list-style-type: none">• Book and installation kit titles• Variable names and values• Emphasized words
Fixed Font	<ul style="list-style-type: none">• File and folder names• Commands and code examples• Text you must type• Text (output) displayed in the command-line interface

Using Online Help

AppManager provides task-based, reference, and context-sensitive online Help.

To access task-based Help or search for Help topics, click **Help Topics** on the Help menu. To view context-sensitive Help within dialog boxes, click **Help** or press **F1**.

You can get help on individual Knowledge Scripts in one of the following ways:

- On the **Values** tab of the Knowledge Script Properties dialog box, click **Help** or press **F1**.
- In the Knowledge Script pane of the Operator Console, highlight a Knowledge Script and press **F1**.

Other Information in the Library

The library provides the following information resources:

- *Installation Guide*: Provides complete information about AppManager pre-installation requirements and step-by-step installation procedures for all AppManager components.
- *Control Center User Guide*: Provides complete information about managing groups of computers, including running jobs, responding to events, creating reports, and working with the Control Center Console. A separate guide is available for the AppManager Operator Console.
- *Administrator Guide*: Provides information about maintaining an AppManager management site, managing security, using scripts to handle AppManager tasks, and leveraging advanced configuration options.
- *Upgrade and Migration Guide*: Provides complete information on how to upgrade from a previous version of AppManager.
- *Management Guides*: Provide information about installing and monitoring specific applications with AppManager.

The AppManager library is available in Adobe Acrobat (PDF) format and is located in the \Documentation folder of the AppManager installation kit.

NetIQ Online Support and Extended Support Web sites provide other resources:

- Downloads, including hotfixes, service packs, and product upgrades.
- Documentation, including white papers and the most current information about version support for the systems and applications monitored by AppManager.

Note You can access NetIQ Support without a password or registration. To access the Extended Support site, you must be a registered AppManager customer.

In addition to the AppManager documentation, you may want to consult the documentation for your Windows or UNIX operating system, or other application- or system-specific documentation for reference and conceptual information. This background information can help you get the most out of your AppManager installation.

About Attachmate

Attachmate, owned by an investment group led by Francisco Partners, Golden Gate Capital and Thoma Cressey Equity Partners, enables IT organizations to extend mission critical services and assure they are managed, secure and compliant. Attachmate's leading solutions include host connectivity, systems and security management, and PC lifecycle management. Our goal is to empower IT organizations to deliver trusted applications, manage service levels, and ensure compliance by leveraging knowledge, automation and secured connectivity. For more information, visit www.attachmate.com.

NetIQ Solutions from Attachmate

Attachmate provides a wide selection of systems and security management solutions to help you manage and secure all your essential platforms, including Windows, Unix, Linux, and iSeries. These Knowledge-Based Service Assurance products and solutions include embedded knowledge and tools to implement industry best practices and to better ensure operational integrity, manage service levels and risk, and ensure policy compliance. Our modular, best-of-breed solutions for Performance and Availability Management, Security Management, Configuration and Vulnerability Management, and Operational Change Control integrate through an open, service-oriented architecture allowing for common reporting, analytics and dashboards. Attachmate offers the following systems and security management solutions:

- **Performance and Availability Management** These products offer rapid time-to-value solutions that enable you to align your IT operations with business priorities and optimize the delivery of your IT-based business services. This solution automates the complete IT service management lifecycle: assessment of requirements, definition of Service Level Agreements, management of day-to-day operations, and review of operational metrics.
- **Security Management** These easy-to-install-and-deploy products provide effective protection from and response to security-related threats. This solution provides powerful features, such as real-time security event monitoring, mapping of threat indicators, policy violation alerts, and expedited incident forensics and resolution. These products reduce the time required to identify and resolve security threats.
- **Configuration and Vulnerability Management** These products allow you to quickly and easily assess

vulnerabilities, manage security risks, and assure policy compliance. This powerful solution measures and enforces compliance to configuration baselines based on your corporate policies, regulations, and evolving security threats. You can use the latest security knowledge, which is updated in real time, to resolve compliance and configuration issues.

- **Operational Change Control** These products enable IT organizations to control, manage, and audit operational changes to servers, Active Directory, and Group Policy with unprecedented levels of accountability. NetIQ's Operational Change Control (OCC) solutions enable enterprise customers to meet IT compliance and operational integrity needs in the most cost-effective manner, by delegating access control, managing changes according to policy, and alerting and reporting on change activities and entitlements.

Contacting NetIQ Solutions Support

Please contact us with your questions and comments. We look forward to hearing from you.

Sales Email: info@netiq.com

Telephone: 713.418.5555 (United States)
+353 (0) 91 782 677 (Europe, Middle East, and Africa)
For other locations, see our Support Contact Information Web site at www.netiq.com/support

Support Web Site: www.netiq.com/support

Welcome

This chapter provides an overview of the functionality and architecture of the NetIQ AppManager Suite. You should review this chapter to become familiar with key AppManager concepts and terminology. The following topics are covered:

- [Why Monitor Your Environment?](#)
- [How Does AppManager Help?](#)
- [What Is the AppManager Suite?](#)
- [What Does AppManager Do?](#)
- [How Does AppManager Work?](#)

Why Monitor Your Environment?

For most organizations, managing a distributed network of servers and workstations is essential to keep mission-critical servers, applications, and resources up and running. Monitoring and managing the operations on your key computer resources involves:

- **Finding out what's wrong with the systems on your network.** Regardless of the size of your network, you need to locate current and potential problems quickly so you can respond with an appropriate action. Are any servers or workstations running out of disk space? Is a printer jammed and queueing up a large number of jobs? Are key services or applications, such as Microsoft Exchange, unavailable or going down?

- **Automating system management tasks or corrective actions.** By automatically monitoring system performance and statistics at set intervals, you can ensure that certain files or processes are available, be proactive in your response to trouble spots, and establish automatic corrective actions (for example, automatically restarting a service that's gone down).
- **Gathering real-time and historical performance information.** Monitoring real-time performance information keeps you on top of current system status at a glance. By monitoring information such as CPU load or usage patterns over time, you can track bottlenecks and trends to help you do capacity and resource planning.
- **Maintaining a complete inventory of system configuration.** Keeping complete, detailed records about the configuration of the computers in your environment is an essential part of maintaining your organization's infrastructure and understanding user needs. Having easy access to information about the operating systems, CPU, memory configuration, and applications installed on the computers you are managing not only provides better visibility into the overall make-up of your network, but also allows you to better plan upgrades and evaluate problems.

By performing these tasks for your organization, you can avoid unexpected downtime, alleviate system performance problems, and ensure the resources that are crucial to your business are up and running at all times.

How Does AppManager Help?

AppManager helps you optimize performance and ensure availability of your supported Windows and UNIX systems, applications, and hardware by providing automated problem detection and correction.

AppManager also helps to reduce the support costs associated with managing a distributed heterogeneous network of clients and servers by providing a centralized console for monitoring activity while providing the flexibility of monitoring activity from a Microsoft Windows server with Microsoft Internet Explorer.

AppManager provides a scalable, multi-tier architecture that can handle small workgroups, mid-size departments, or full-enterprise deployment and delivers robust out-of-the-box functionality to make you productive right away. In addition, you can completely customize AppManager to address your company's unique configuration requirements or extend its functionality to monitor internally-developed applications.

What Is the AppManager Suite?

AppManager consists of a number of different product options that can be used together or separately, depending on the requirements of your organization. The combined set of these product options make up the NetIQ AppManager Suite. Each product option in the suite is specially focused to perform specific management tasks for a particular application or resource.

These products can be installed together as a suite on a single server or separately on distributed servers, as necessary.

What Does AppManager Do?

AppManager identifies problems in your environment; helps you assess the cause, location, and severity of these problems; and allows you to automatically correct or initiate other appropriate actions when problems occur.

With AppManager you can:

- View complete, detailed information about the configuration of the computers in your environment.
- Decide appropriate thresholds, severity, and corrective actions that suit your business environment or use the AppManager defaults to minimize the work you do to get started.
- Centralize the management of distributed and remote systems by installing a number of NetIQ AppManager Operator Consoles and Operator Web Consoles strategically across your organization for greater flexibility in how you monitor remote locations.
- Be pro-active in detecting and correcting problems before they impact your business.
- Gather real-time and historical information about your environment to help you in capacity planning.
- Continue monitoring even when network outages occur between any two components.
- Monitor internally-developed or commercial applications that run on supported Windows and UNIX operating systems.
- Add custom monitoring functions that handle the specific needs of your business and environment (see *Developing Custom Knowledge Scripts* for information about extending AppManager's monitoring capabilities).

How Does AppManager Work?

AppManager manages the availability and performance of supported Windows and UNIX systems and applications through monitoring functions called **Knowledge Scripts**.

Each Knowledge Script is a business or system management rule designed to handle one or more tasks, such as detecting whether a Windows service is down, checking UNIX CPU usage, or verifying connectivity between Exchange servers.

Depending on the task, Knowledge Scripts can collect performance data (for example, about how busy your intranet and Internet servers are), monitor systems for simple or complex **events** (for example, a computer or service detected down), and respond with one or more **actions** (for example, send an e-mail when there's a problem or restart a service automatically).

AppManager Architecture and Components

To provide the best combination of efficiency, scalability, and flexibility, AppManager uses a multi-tier architecture. This multi-tier architecture gives you flexibility in distributing process load across multiple components and allows for efficient communication between components. The AppManager architecture consists of five components:

- **AppManager consoles:** Programs for managing various aspects of your environment. The *Operator Console* allows you to configure and control the execution of Knowledge Scripts on the systems and applications you manage. The *Operator Web Console* allows you to check the status of jobs and events, create and view charts, run Knowledge Scripts, view details about the computers you are monitoring, and view reports from a Microsoft Windows server with Microsoft Internet Explorer. The *Chart Console* allows you to generate and view charts of AppManager repository data. (The Chart Console can also be displayed from the Operator Web Console). Each of these programs can be selectively installed on computers.

AppManager also offers Control Center, an enhanced version of the Operator Console that allows you to manage

multiple repositories. Refer to the *Control Center User Guide* for more information.

- **AppManager agent:** Resides on managed clients to intelligently and autonomously monitor the system and application resources (such as printers, disks, CPUs, database files, etc.) associated with those computers.

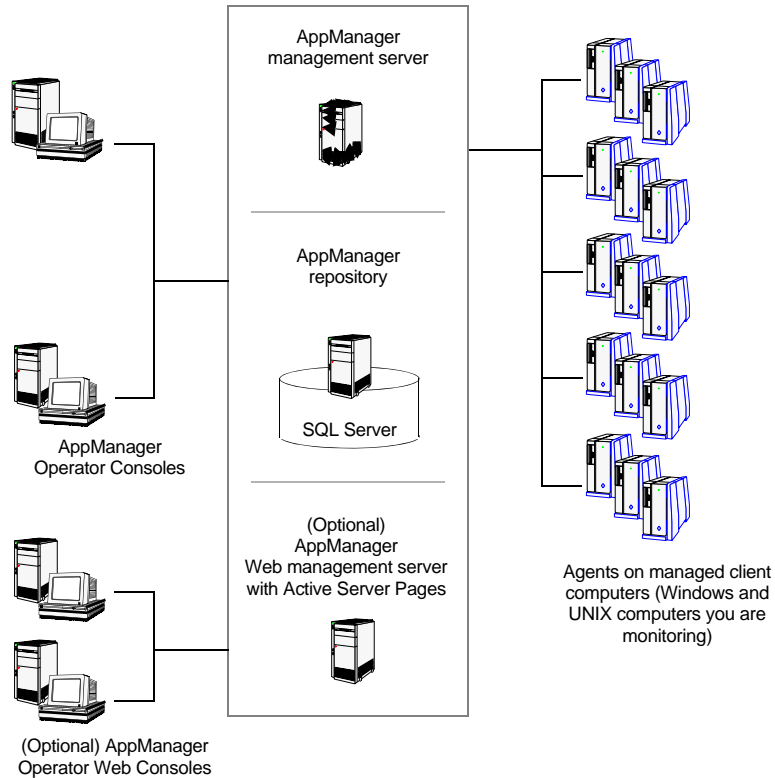
For Windows computers, the AppManager agent consists of two Windows services, the NetIQ Client Resource Monitor (netiqmc) and the NetIQ Client Communication Manager (netiqccm), a local repository database, and at least one managed object that resides locally on each computer you are managing.

For UNIX computers, the AppManager agent consists of a single daemon, nqmagt, and the supporting files and directories that provide data persistence (equivalent to the local repository) and access to system statistics (equivalent to managed objects).

- **AppManager management server:** A Windows service (NetIQms) that manages event-driven communication between the AppManager repository and AppManager agents.
- **AppManager repository:** A SQL Server database in which management data is stored.
- **AppManager Web management server (Optional):** A set of Active Server Pages that communicate with the AppManager repository. These pages make up the *Operator Web Console*.

From a Microsoft Windows server with Microsoft Internet Explorer, the Operator Web Console allows you to check the status of jobs and events, create and view charts, run Knowledge Scripts, and view details about the computers you are monitoring.

The following figure illustrates a simplified view of this architecture:



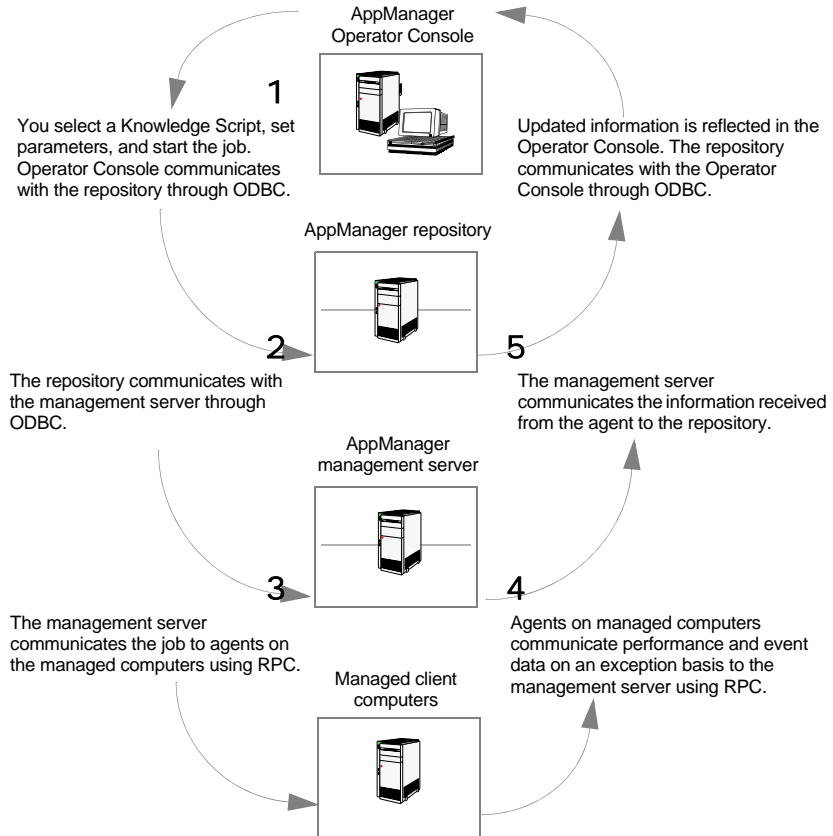
Communication among AppManager Components

Once you install the AppManager components, you use the Operator Console to select a Knowledge Script and the computers where you want it to run. You are prompted to enter scheduling information and set parameters for the Knowledge Script. Once you accept the defaults or configure the Knowledge Script parameters, AppManager starts a **job** on the selected computers. A job is an instance of a Knowledge Script running on a computer you are monitoring.

When you start a job:

- The Operator Console notifies the repository that you have requested a Knowledge Script to run.
- The AppManager repository is updated with information about the job properties and communicates with the management server.
- The AppManager management server then sends the Knowledge Script and properties you have set to the appropriate managed computers you want to monitor by contacting the AppManager agent.
- The AppManager agent runs the job on the managed computer and sends to the management server any performance or event information returned by the Knowledge Script job.
- The managed computers report information back to the management server, which inserts it into the repository. That triggers an update on the Operator Console (for example, a real-time graph gets updated with new data).
- The Web management server, using Active Server Pages, updates the Operator Web Console with information from the AppManager repository.

The following figure illustrates a simplified view of the communication flow (without the optional Web management server and Operator Web Console):



Starting a Monitoring Job

Knowledge Scripts automate management and monitoring tasks. All you need to do is start them, view their output, and then stop and restart them when needed.

To start a Knowledge Script job in the Operator Console, drag-and-drop the icon for the Knowledge Script from the Knowledge Script pane onto the resource icons that represent

your Windows systems and applications in the TreeView pane. A job is simply a particular instance of a Knowledge Script running on a server or workstation you are managing.

You can quickly customize the default properties of any Knowledge Script (for example, to change the thresholds and corrective actions defined for various events). You can also control the frequency at which a Knowledge Script performs its intended task.

Use the Operator Console or the Operator Web Console to view job status and any output created by the jobs in the form of events and data collected.

What Happens on the Managed System

When you start a job, you indicate the managed clients on which you want the job to run. A **managed client** is a supported Windows or UNIX computer on which you have installed the AppManager agent and set up to be monitored by AppManager. The actual monitoring of that computer is performed by the **AppManager agent**.

On Windows computers, the AppManager agent is composed of two Windows services: the NetIQ Client Resource Monitor (`neti qmc`); and the NetIQ Client Communication Manager (`neti qccm`). On UNIX computers, the AppManager agent is composed of a single daemon, `nqmagt`.

On an exception basis, the agent communicates back to the management server any relevant output from the Knowledge Script. (For network efficiency, the AppManager agent only communicates back to the management server when an event has occurred or data needs to be inserted into the repository database.)

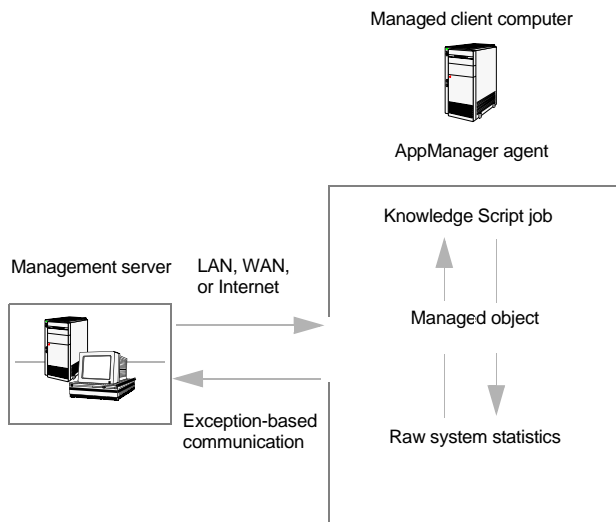
AppManager agents handle the scheduling of Knowledge Scripts and initiate corrective actions and communication with the management server. The collection of performance and

event data is facilitated through the use of software probes called **managed objects** that “plug into” the AppManager agent.

Knowledge Scripts use managed objects to access counters, event logs, queries, Application APIs, and other sources to gather statistics metrics, and other properties of specific application elements. On Windows computers, managed objects are COM/OLE objects in the form of dynamic link libraries (.DLL files). On UNIX computers, managed objects are Perl modules, in the form of dynamic shared libraries.

Using these native sources of information, managed objects collect raw statistics and information, such as current CPU utilization or database lock activity, and pass that information to the Knowledge Script jobs. Knowledge Scripts then provide the rules for what to do with this raw information. The Knowledge Scripts run under the control of the AppManager agent. On managed Windows computers, the Knowledge Scripts invoke the managed objects through the standard COM/OLE interface. On managed UNIX computers, the Knowledge Scripts invoke the managed objects through the standard Perl module interface.

The following figure illustrates a simplified view of this relationship.



How the Management Server and Repository Are Used

The management server is a Windows computer running the NetIQms service. The NetIQms service manages the event-driven communication between the repository and the managed clients.

The AppManager repository server is a central Microsoft SQL Server database that stores management data for a site.

The management server and repository can reside on the same physical computer or separate Windows servers to distribute the workload.

When you use the Operator Console or Operator Web Console to start a job, a record is inserted into the AppManager repository's database. An update of this nature is interpreted as an instruction for the management server to communicate the Knowledge Script and its properties to the appropriate managed clients.

The management server also receives communication back from the managed clients in the form of performance and event information returned by the Knowledge Script jobs. The management server inserts this information into the AppManager repository, which in turn triggers an update to occur on the Operator Console and Operator Web Console (for example, a real-time graph gets updated with new data).

Communication between the management server and the managed clients is via Microsoft RPC. Communication between the Operator Console and the repository and between the repository and the management server is through ODBC.

Reporting System Management Data

Because AppManager stores information in a SQL Server database, you can easily share, analyze and report information about the health and use of systems and applications you are monitoring. (SQL Server inherently provides the ability to handle high transaction rates and lets you collect extensive amounts of data.)

AppManager Report Scripts collect data from the AppManager repository and generate numerous standard reports that summarize system and performance data. In addition, you can customize these scripts the same way you customize standard Knowledge Scripts and tailor them to your network environment. AppManager security is set up at the database level to ensure that only authorized personnel can operate the Operator Console.

Multiple Operator Consoles can share the same data. You can use the NetIQ Analysis Center (licensed separately) to generate reports based on data from multiple AppManager repositories. You can also use AppManager Control Center to manage data from multiple repositories and to perform the same tasks you'd perform using the Operator Console.

Getting Started

This chapter provides an introduction to the NetIQ AppManager Operator Console and suggests how you can get started using AppManager. The following topics are covered:

- [Starting the Operator Console](#)
- [About the Operator Console](#)
- [Connecting to Another Repository](#)

Note After you have installed AppManager components on one or more computers, you use the Operator Console to monitor your environment. For information about installing AppManager and about system requirements, see the *Installation Guide*.

Starting the Operator Console

To start the Operator Console, you must log on to the AppManager repository with a SQL Server login account that has permission to access AppManager. See the *Installation Guide* for more information.

For information on starting the Operator Console using command-line parameters, see the *Administrator Guide*.

- 1 On a computer where you have installed the Operator Console, click the **Start** button, then **Programs > NetIQ > AppManager > Operator Console**.

At the AppManager Operator Console Logon dialog box, log into the repository you want to use. (You will see the same

dialog box when you start other AppManager programs, such as AppManager Chart Console.) If you aren't sure what to enter in any field or have problems logging on, see your system administrator.



Field	Description
Server	Type the name of the SQL server that manages the AppManager repository. When specifying a computer name, you can enter the Windows computer name or the IP address. At this time, only v4 IP addresses are supported. For example, to specify a named instance on SQL Server 2000, you can enter 10. 1. 10. 43\INST1.
Repository	<p>After you type the Server name and press the Tab key, AppManager fills in this drop-down list with the names of databases on that server. Select the name of the AppManager repository you want.</p> <p>The default AppManager repository name is QDB.</p>
Connection Information	<p>Select an option to specify the method that SQL Server uses to validate the user:</p> <ul style="list-style-type: none">• Use Windows authentication specifies that SQL Server uses the Windows user information to validate the user before access is granted. The AppManager console computer needs to be part of a Windows domain or workgroup.• Use SQL Server authentication specifies that SQL Server uses the standard SQL Server security validation. This is the the default when SQL Server is configured to use both SQL Server and Windows authentication. You must add the login name to SQL Server before a user can log in. Login name - Type the user name of the SQL Server login account used to access the AppManager repository. The default SQL Server user is the administrative user, sa. Password - Type the password for the SQL Server login account. If using the sa user, the default password is Null.

2 Click **Logon**.

Now that you're logged onto the AppManager repository, data is received dynamically.

Notes

- If you cannot connect to the AppManager repository or want to use data cached locally from a previous session, you can click **Off Line**. Working off-line displays the Operator Console but, because you are not connected to the repository, you are not able to manage systems dynamically or view dynamic information.
- Running separate instances of the Operator Console, different users can access the same repository from the same console computer.
- Once you have logged into the Operator Console, you can switch to another repository. See [“Connecting to Another Repository” on page 35](#) for more information.
- Your AppManager environment must have the same version of the AppManager repository, management server, console, and Web management server components. For example, you cannot log on to the repository with an earlier version of the Operator Console or Web Console.

Getting Help for the Operator Console

For online Help when working in the Operator Console, click **Help > Help Topics** or press the F1 key. Each of the property sheets and dialog boxes also contains a Help button, which you can click for detailed information about specific options.

Tip of the Day

Once you have logged into a database, the Tip of the Day dialog box is displayed by default. These tips provide useful hints on how to use the Operator Console.

Tip of the Day Display Options	Instructions
Dismiss the Tip of the Day dialog box	Click OK .
Disable the Tip of the Day dialog box so that it doesn't appear each time you start the Operator Console	<ul style="list-style-type: none">• Deselect Show Tips at Startup on the Tip of the Day dialog box, or• Click File > Preferences and deselect Show Tip of the Day.
To show the Tip of the Day at any time	Click Help > Tip of the Day .

Once you have dismissed the Tip of the Day dialog box, the Operator Console appears.

Registering AppManager

To qualify for early notification of upgrades, regular e-mail newsletters, and special promotions, you need to register.

When installing the AppManager repository component, you are prompted to register. You can also register from within the Operator Console.


To register AppManager:

- 1 In the Operator Console, click **Help > About**.
- 2 In the About dialog box, click **Register**.

If you have Internet access, your default Web browser is opened at the NetIQ AppManager online registration Web page (<http://www.netiq.com/register/appmanager>). Follow the instructions for filling out and submitting the registration form.

About the Operator Console

The Operator Console is where you do most of the work to manage the systems on your network. To give you the tools for viewing and managing your computer resources, the Operator Console consists of a **view** with five distinct **panes**. The example below displays the TreeView pane, the Knowledge Script pane, and the List pane.



Working with Panes

Information about the systems you are managing is grouped in panes.

Pane	Description
TreeView	<p>Shows a hierarchical list of discovered and managed objects. By default, this is one panel. You can set a preference to display two panels, with managed objects (such as SQL Server) on the left and their resources (such as Services and Databases) on the right. For more information about splitting the TreeView pane, see “Setting Operator Console Preferences” on page 64.</p> <p>Note Splitting the TreeView pane speeds up console startup and navigation in the TreeView pane. It is recommended in environments with hundreds of managed systems.</p>
Knowledge Script	<p>Shows the Knowledge Scripts you can run, organized into categories such as General, NT, and Action. Click a tab to see the Knowledge Scripts in that category.</p> <p>To see a brief description of each Knowledge Script, right-click in the pane and then click Show Description.</p> <p>Drag Knowledge Scripts from this pane to the TreeView pane to start a job.</p>
List	<p>Select a tab to show the status of events, jobs, graph data streams, and configuration details about the servers and workstations you are managing. The content in the List pane is always related to the current view (such as Master or NT).</p>
Graph	<p>Provides an area for displaying real-time or historical information in graphical form.</p> <p>Drag data streams from the Graph Data tab in the List pane to this pane. (The Graph pane appears blank until you drag and drop a data stream into it.)</p> <p>Notes</p> <ul style="list-style-type: none">• By default, the Graph pane is not displayed. To display the Graph pane, click View > Graph Pane.• You can also display graph data in the Chart Console. For information on creating charts to display graph data, see the AppManager Help.

Pane	Description
Web	Provides an area for displaying Web documents, such as a custom Web application, that integrates server or object information from the TreeView pane. For more information about the Web pane, see "Integrating AppManager with a Custom Web Application" on page 93 . Note By default, the Web pane is not displayed. To display the Web pane, click View > Web Pane .

Note To customize the Operator Console to see only the panes you are interested in working with, see ["Customizing the Operator Console Layout" on page 38](#).

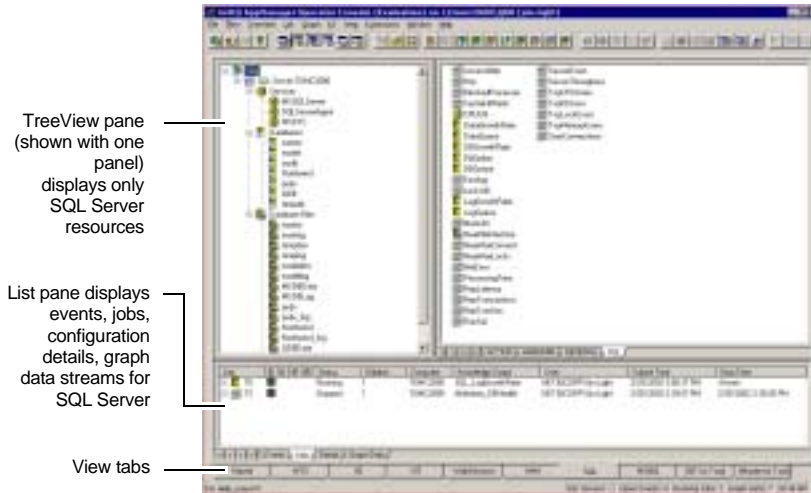
Working with Views

A view shows selected panes and filters the computers and information that is displayed in them.

The default **view** in the Operator Console is the **Master** view, which includes information about all of the Windows and UNIX computers you are managing and all Knowledge Script categories.

Standard views allow you to limit the information displayed to specific systems or applications (for example, only Exchange servers or UNIX computers). Standard views are added when you discover new types of servers or applications.

For example, the SQL view looks like this:



After you discover resources, you can create a view that displays a snapshot of the discovered resources or a view that dynamically displays resources based on one or more rules.

Topic	Where to Look for More Information
Displaying a snapshot of discovered resources	Chapter 7, "Using Snapshot Views"
Dynamically displaying resources based on one or more rules	Chapter 8, "Using Dynamic Views"

Note In the NetIQ AppManager Security Manager, AppManager users are assigned security roles, which include rights as to which views the user can see in the Operator Console and Operator Web Console. If you don't have access to a particular view, check your status in the Security Manager or contact your AppManager system administrator.

Connecting to Another Repository

Although you can only work with one repository at a time in the Operator Console, you can switch to another repository without logging out of the console.

- 1 Click **File > Connect Repository**.
- 2 Select the repository you want to log into.

Options	Steps to Take
Log into a repository for the first time during the current Operator Console session	Click New ; then click Yes when asked if you want to disconnect from the current repository. At the NetIQ AppManager Operator Console Logon dialog box, enter the user name and password, and the repository server and database names. (For more information on these options, see page 27 .)
Log into a repository that you have previously logged into during the current Operator Console session	Click the name of the repository. (You won't be prompted to disconnect from the current repository.)

Note If you decide to display the TreeView pane in two panels instead of one (the default), you can reconnect to the same repository to make the setting take effect. Follow the instructions in Step 2 above for logging into a repository for the first time. For more information on setting the TreeView display preference, see [“Customizing the Operator Console Layout” on page 38](#).

Configuring the Operator Console

This chapter describes how you can configure the NetIQ AppManager Operator Console to fit your needs. The following topics are covered:

- [Customizing the Operator Console Layout](#)
- [Selecting a View](#)
- [Identifying the Computers to Monitor](#)
- [Adding Computers to the TreeView Pane](#)
- [Organizing Computers into Groups](#)
- [Finding Objects in the TreeView Pane](#)
- [Deleting an Object from the TreeView Pane](#)
- [Expanding and Collapsing the Network Hierarchy](#)
- [Viewing Detailed Information about Discovered Resources](#)
- [Viewing AppManager System Information](#)
- [Filtering Information Based on a Resource Object](#)
- [Grouping and Filtering Based on Other Criteria](#)
- [Setting Preferences](#)
- [Displaying Toolbars](#)
- [Displaying a Status Bar](#)
- [Getting Information about Knowledge Scripts](#)
- [Checking Communication with Managed Computers](#)
- [Integrating AppManager with a Custom Web Application](#)
- [Using the Extensions Menu](#)

Customizing the Operator Console Layout

By default, the Operator Console displays three panes, each with a default size and location. For information about the Operator Console panes, see [“About the Operator Console” on page 31](#).

To improve the startup performance of the Operator Console, change the Operator Console preferences to display only the tabs and panes that you want to view. As a rule, you can improve the startup performance of the Operator Console by reducing the number of panes and tabs that are displayed on startup. See [“Setting Preferences” on page 63](#) for more information.

You can customize the console layout to display only the panes you are interested in and change the size, shape, and location of those panes. For example, you can display the Web pane (it is not displayed when you start the Operator Console) to view Web documents.

To	Do this
Display the TreeView pane as two panels, with managed objects (such as a SQL Server) appearing on the left and the managed object's resources (such as services and databases) appearing on the right	<p>Click File > Preferences. On the Console tab, check Split TreeView into two panels. To see the change, restart the console or re-connect to the repository (see “Connecting to Another Repository” on page 35).</p> <p>Notes</p> <ul style="list-style-type: none">• Splitting the TreeView pane into two panels reduces the amount of time it takes to start the Operator Console, improves the speed of navigating in the TreeView pane, and is recommended in environments with hundreds of systems.• With the TreeView split into two panels, you can hide the right panel. Right-click anywhere in the TreeView pane and then deselect Show Details Panel.

To	Do this
Show or hide a pane	Select (to show) or deselect (to hide) View > TreeView Pane, List Pane, Knowledge Script Pane, Graph Pane, and Web Pane . You can also hide a pane by right-clicking in a pane and then clicking Hide Pane .
Move or resize each pane	Use the mouse to drag pane borders in any direction.
Reset all visible panes to the default location and size	Select View > Reset Panes .

Note Throughout this guide, procedure instructions assume that the TreeView pane is displayed with one panel (the default). If you have selected to display the TreeView as two panels, select computers in the left panel and resource objects in the right.

Selecting a View

The Operator Console consists of multiple **views** that affect the details you see in all of the console panes. By default, the **Master view** is displayed. For more information, see [“About the Operator Console” on page 31](#).

To switch views, click the tab at the bottom of the Operator Console for the view you want to use.



The way the name appears in the view tab indicates whether a view is a standard, snapshot, or dynamic view.

For a	The view tab displays
Standard view	The view name only.
Snapshot view	The view name enclosed in brackets []. See Chapter 7, “Using Snapshot Views,” for more information.
Dynamic view	The view name enclosed in braces { }. See Chapter 8, “Using Dynamic Views,” for more information.

When you switch from one view to another, all of the panes reflect information for the selected view. For example, by selecting the SQL view, only jobs, events, Knowledge Scripts, and computers associated with monitoring SQL Server are displayed. To return to a more complete view of your environment, click the **Master** tab.

You can customize the console layout to display only the views in which you are interested. For example, if you are the Exchange administrator, you may want to focus strictly on Exchange Servers using the **Exchange** view, or if you are a database administrator, you may want to include a snapshot view that shows only SQL databases.

To	Do this
Show or hide a view	Click View > View Manager ; then select (to show) and deselect (to hide) the available views. Note If you plan to add computers to the TreeView pane or run Discovery Knowledge Scripts, you must perform these functions from the Master view or a snapshot view that is based on the Master view.
Rename a snapshot or a dynamic view	Click View > View Manager ; then select the view you want and click Rename . You cannot rename a standard view or the Master view.
To save these view settings for future Operator Console sessions	Click File > Preferences . On the Console tab, select Save View Manager settings on exit and reload them on startup .

Note Throughout this guide, procedure instructions assume that all available view tabs are displayed. If you have chosen to hide a view, you can open it by clicking **View > View Manager** and then selecting the view.

Showing and Hiding a View

To show or hide a view:

- 1 Click **View > View Manager**.
- 2 In View Manager, select (to show) and deselect (to hide) the available views.
- 3 Click **OK**.

Notes

- If you plan on adding computers to the TreeView pane or running Discovery Knowledge Scripts, be sure not to hide the **Master** view.
- Throughout online help, procedure instructions assume that all available view tabs are displayed.

Renaming a View

To rename a snapshot view or a dynamic view:

- 1 Click **View > View Manager**; select a view from the list. The **Based On** column indicates the view upon which the snapshot or dynamic view is based.
- 2 Click **Rename**.
- 3 In the Rename View dialog box, type a new name for the view and click **OK**.
- 4 To close the View Manager, click **OK**.

Saving View Settings

To save view settings for future Operator Console sessions:

- 1 Click **File > Preferences**.
- 2 On the **Console** tab, select **Save View Manager settings on exit and reload them on startup**.

Identifying the Computers to Monitor

To manage the servers and workstations in your environment, you first need to identify the computers and the applications on them that need to be monitored. The process by which AppManager is made aware of the computers and applications you want to monitor is called *discovery*.

Discovery typically involves:

- Installing the AppManager agent on the server or workstation you want to manage.
- Adding the server or workstation to the **Master** view. The list of managed resources appears as a network hierarchy you see displayed in the TreeView pane.
- Running one or more specialized Discovery Knowledge Scripts in the **Master** view.

Note You can automatically perform the first two scenarios when you install AppManager, if the management server and repository server are running and you have network access to them. You can also run Discovery Knowledge Scripts at any time after you add servers and workstations to a view, or as part of a monitoring policy. For information about monitoring by policy, see [Chapter 9, “Monitoring by Policy.”](#)

Adding Computers to the TreeView Pane

To add computers (servers, workstations, and cluster nodes; but not clustered virtual servers) to the TreeView pane:

- 1 Click the **Master** tab to select the **Master** view or click a snapshot view that is based on the **Master** view.

Note The **Master** view and a snapshot view based on the **Master** view are the only views in which you can run Discovery Knowledge Script jobs.

- 2 In the TreeView pane, select the top-level icon to have the computer added under the root of the hierarchy, or select a group name to have the computer added under the group.
- 3 Right-click and then click **Add Computer**.
- 4 In the Add Computer dialog box, select an option to choose the type of computer you want to add and then type the names or IP addresses of the computers you want to add in the **Computers** box (you don't need to include leading

backslashes [\\]). To add multiple computer names, separate each name by a comma.

Select an option to choose the type of computer you want to add...

...and then type the computer names or IP addresses. Or, for Windows computers, click the Browse (...) button to browse the network and select the computers.



To	Do this
Select a Computer type	<p>Select an option to choose the type of computer you want to add:</p> <ul style="list-style-type: none">• Windows computers This option allows you to add computers that run on supported Windows operating systems. This is the default.• Unix computers This option allows you to add computers that run on supported UNIX operating systems.
Select Computers	<p>Type the names or IP addresses of the computers you want to add to the TreeView pane (you do not need to include leading backslashes [\\]). At this time, only v4 IP addresses are supported.</p> <p>To add multiple computer names, separate each name by a comma.</p> <p>If the computer you want to add is not accessible from the Operator Console, such as a managed client computer behind a firewall, the Add Computer Errors dialog box lists the inaccessible computer and you must manually add the computer. See “What to Do if a Computer Is Not Added” on page 47 for more information.</p>



To	Do this
Browse (...)	<p>Click the Browse (...) button. You can then browse through the domains in your network to find the Windows NT or Windows 2000 computers you want to add. This option is not applicable for UNIX computers.</p> <p>Computers that are already in the TreeView pane will appear with a red check mark as part of their object icons. Select the computers and click OK. You can return to the Browse dialog box to select additional computers (for example, from other domains).</p> <p>Support for Microsoft Cluster Server The names of cluster nodes and the virtual servers on them appear in the list. Add only cluster nodes (not virtual servers) to the TreeView pane.</p>
Discover objects automatically	<p>This option is available when you add a Windows or UNIX computer. Click:</p> <ul style="list-style-type: none"> • Discover Windows objects automatically to discover NT resources. Selecting this option is the same as running the Discovery_NT Knowledge Script on the computer. • Discover UNIX objects automatically to discover UNIX resources. Selecting this option is the same as running the Discovery_UNIX Knowledge Script on the computer. <p>Notes</p> <ul style="list-style-type: none"> • You might want to run other discovery Knowledge Scripts on the computer after adding it to the TreeView pane. • If you discover resources in a snapshot view, after the discovery completes, you must add the discovered resources. See “Adding Discovered Resources” on page 212 for more information.

To	Do this
Immediately query the computer for time zone information	<p>Select only when:</p> <ul style="list-style-type: none"> You are adding computers that are in a different time zone from your Operator Console. It is absolutely critical that any Knowledge Script jobs that run immediately after the computer is added display correct time zone information. In the Console tab of the Preferences dialog box, you have selected the option to Display dates and times based on the local time of the managed computer. (If this preference is not selected, this option has no effect.) <p>Note When the Immediately query the computer for time zone information option is not selected, AppManager adds a computer to the TreeView without actually connecting to it. The management server then establishes a connection to the managed computer, and it receives information about the computer's time zone. Therefore, if you've set the preference to show local times in the Operator Console, for the few minutes (the actual time depends on your network configuration) after the managed computer is added to the TreeView but before the management server connects to it, any jobs you start on that computer may not reflect correct time zone information. If you do not plan to run jobs immediately, or if it is not critical that data returned in the first few minutes contains correct dates and times, you can leave this option unchecked. Selecting the Immediately query the computer for time zone information option ensures accurate data from the moment the computer is added to the TreeView, but the trade-off is that it requires AppManager to connect to each computer before adding it to the TreeView; this process takes substantially longer than simply adding the computer without this option selected.</p>

5 Click **OK**.

If you entered a valid server or workstation name that can be reached over the network, it is added to the TreeView pane. If the computer is not running or you entered the name of a computer that is not accessible from the Operator Console computer, the Add Computer details box appears. See [“What to Do if a Computer Is Not Added” on page 47](#) for more information.

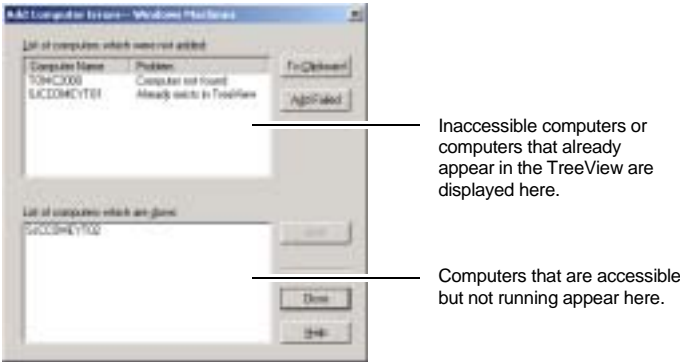
Notes

- If the AppManager client resource monitor (neti qmc.exe) is not running on the computer,  this icon appears in the TreeView. In this case, only Knowledge Scripts related to installing the AppManager agent (such as AMADMIN_AgentInstall) can be run.
 -  This icon, which only appears if the TreeView pane is split into two panels (see [page 65](#)), means the server, workstation, or cluster node hasn't been discovered. For information about Discovery Knowledge Scripts, see the online help.
-

What to Do if a Computer Is Not Added

If the computer you want to add is not running, is not accessible from the Operator Console computer, or is already added to the TreeView pane, the Add Computer Errors dialog box displays a list of computers that were not added. This dialog box only appears when a computer was not added to the TreeView.

After you close this dialog box, the results cannot be displayed again; if you need to keep a list of the computers that were not added, copy the contents of the list to the Windows Clipboard.



To	Do this
Copy all computers that were not added to the Windows Clipboard.	Click To Clipboard .
Add computers that cannot be accessed through the network from the Operator Console computer to the TreeView pane.	Click Add Failed to add all inaccessible computers in the list. If you have managed client computers that cannot communicate directly with the Operator Console, for example, managed client computers behind a firewall, click this button to add them to the TreeView pane.
Add computers that are accessible to the Operator Console computer but are not running.	Click to select the computers you want and click Add .
Close the dialog box.	Click Close . After you close this dialog box, the results cannot be displayed again.

Organizing Computers into Groups

You can organize the computers (servers, workstations, and cluster nodes, but not clustered virtual servers) on your network into logical groups. For example, you may want to

create a group for all of the West Coast computers or separate groups for computers that serve the Engineering and Finance departments in your organization.

Create groups to organize the computers in a particular view; a group only appears in the view in which it was created, and computers can be organized differently in different views.

You can also create groups within groups, giving you a great deal of flexibility in organizing your TreeView hierarchy and managing jobs.

To enable a monitoring policy for a group, use the Monitoring Policy tab. For more information, see [“About Policy-Based Monitoring” on page 249](#).

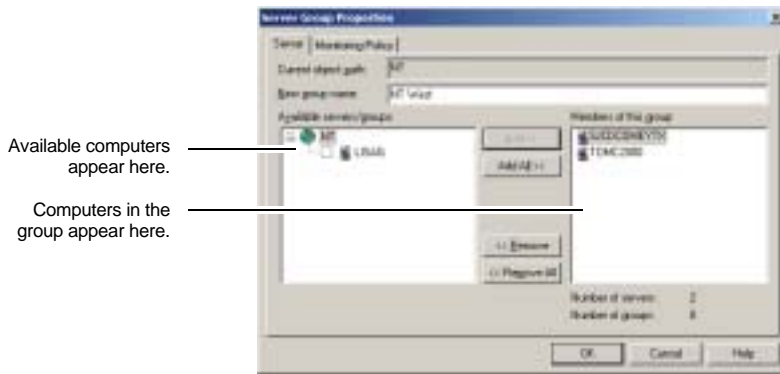
Creating a Group

Create a group to organize computers in a standard view or a snapshot view. To organize computers in a dynamic view, you must configure a group rule to select the computers you want.

After you create a group, you can run monitoring jobs on some or all of the servers in the group, and automatically monitor all servers in the group using a monitoring policy. For information about monitoring by policy, see [Chapter 9, “Monitoring by Policy.”](#)

- 1 Click a view tab to select the view in which you want to create the group.

- 2 In the TreeView pane, right-click the view name or another group name, and then click **Create Server Group**.



Note Groups can be nested, but any individual computer, server, or cluster node can appear only once in the TreeView of a particular view.

- 3 Type a name for the group in **New group name**.

The group name must be unique within the current level of the TreeView (for example, there can be only one Finance group in the top level of the NT view); however, you can use the same name in different levels of the TreeView (for example, you nest Finance groups).

- 4 From **Available servers/groups**, select the computers and existing groups you want to add.

- 5 Click **Add**.

To add all available computers and groups at once, click **Add All**. (Remember, any individual computer can appear in only one group.)

- 6 When you finish, click **OK**.

The new group is displayed in the TreeView pane. Expand the group to see its members.

Changing the Servers in a Group

Use the Server tab of the Server Group properties dialog box to:

- Organize computers in a standard or snapshot view into existing server groups. To organize computers in a dynamic view, you must configure a group rule.
- View the total number of computers in a group.

If a group is monitored by a policy, when you add a computer to the group, matching resources are automatically monitored by the policy. If you remove a computer from a policy-enabled group, any related policy-based jobs are stopped and deleted.

Tip When you add computers to or remove computers from a server group that is monitored by policy, it can take up to one minute to start new policy-based jobs or stop existing policy-based jobs. See [“Changing a Monitoring Policy” on page 267](#) for more information.

- 1 In the TreeView pane, right-click a group and click **Properties**.

Use the Current object path and group name to verify the selected group. You can add and remove groups and servers from the current group.

- 2 In the Server Group Properties dialog box, update the **Server** tab:



To	Do this
Add computers or groups	<ol style="list-style-type: none">1 In Available servers/groups, select one or more computers (servers or workstations) or groups.2 Click Add. The objects are displayed in Members of this group. <p>Note To add all of the servers and groups in the TreeView to the current group, click Add All.</p>
Remove computers or groups	<ol style="list-style-type: none">1 In Members of this group, select one or more computers (servers or workstations) or groups.2 In Available servers/groups, select a level in the TreeView hierarchy to specify where you want to move the selected computers or groups.3 Click Remove. The objects are displayed in Available servers/groups. <p>Note To move all of the members to the selected level in the TreeView hierarchy, click Remove All.</p>
Rename a group	In Group name , type a different name.

Tip You can also organize computers within existing groups from the TreeView pane of the Operator Console. Simply drag a computer or server group to another server group or the view icon.

- 4 Click **OK**.

Finding Objects in the TreeView Pane

The Operator Console provides a variety of options so you can focus your search very precisely. Using the options in the Find dialog box, you can constrain the search to find only computers or search for objects below a particular location in the TreeView hierarchy. You can also constrain the search to the current level of the TreeView hierarchy or search all sub-folders.

You can search using the explicit names of objects (Memory, SQL Server) or by using text and wildcards (M*, SQL*). The * wildcard is used for zero or more characters; the ? wildcard is used for single characters.

To search the TreeView:

- 1 Click an object in the TreeView (such as the view name or a server group) to specify where the search should begin. By default, the search finds objects in any sub-folders.

For example, if you search from the view name (at the top of the TreeView), the entire view is searched; if you search from a server name, all objects associated with that server are searched.

- 2 Click **TreeView > Find**.

- 3 In the Find in TreeView dialog box, enter search criteria in **Find what** or select a search item from the list:

Type the words you want to find here.

To search from a particular location in the TreeView, expand the list and then double-click an object.



In	Do this
Find what	Specify the words you want to find. Entering a string (such as SQL) will find any objects containing the string. You can also specify the following special characters: <ul style="list-style-type: none"> • * matches any number of characters • ? matches a single characters
Search server only	Select to only find matching server names; sub-objects (such as Exchange or SQL) are not found.
Look in subfolders	Deselect this option to only search the current level in the hierarchy. By default, this option is selected and all sub-folders are included in the search.
Look in	Click to expand the list of objects and then double-click to select a new search location. By default, the search location is the selected TreeView object. The current search location appears in the list.

- 4 Click **Find**.

The results of your search appear in the list; the top-most item is highlighted in the list and in the TreeView.

Double-click an item in the list to expand the TreeView and select the object. The search constraint is also updated in the Look in field.

- 5 To perform a new search on an object in the search results list, double-click an object in the list and then repeat Step 3 and Step 4.
- 6 You can leave the dialog box open and perform other Operator Console tasks or click **Close** to close the dialog box.

Deleting an Object from the TreeView Pane

You can delete any object from the TreeView pane of the Master view, a standard view, or a snapshot view. When you delete an object, any associated jobs are stopped and deleted. When you delete objects that are monitored by policy, it can take up to one minute to stop and delete policy-based jobs.

When you delete an object, the object may also be removed from other views:

If you delete an object from	The object is also deleted from
The Master view	The corresponding standard view and any snapshot views.
A standard view	The Master view and any snapshot views.
A snapshot view	Not applicable. In this case, the object is only deleted from the snapshot view.

Note You cannot delete objects from the TreeView pane of a dynamic view. See [“Understanding Dynamic Views” on page 216](#) for more information.

For information on automatically deleting associated events and data when a job is deleted, see [“Set Preferences for Event Management” on page 79](#) and [“Set Preferences for Data Management” on page 76](#).

Deleting an Application Server or Workstation

You can delete one or more application servers or workstations from the TreeView.

Note To delete a server that is running the AppManager management service (netiqms.exe) for the current repository, you must stop the management service on that computer before you delete the computer. If the management service is restarted, the computer automatically reappears in the TreeView.

- 1 In the TreeView pane, select a server or workstation. To delete some or all of the servers or workstations in a view or group, select the view or group.
- 2 Click **TreeView > Delete Object**.

If you selected a view or group, in the Delete Objects dialog box, select the servers and groups you want to delete and then click **Delete**.
- 3 Click **Yes** to confirm.

Deleting a Particular Resource

You can delete a resource folder (which includes all resources in the folder), for example, **SMTP Sites**, or a particular resource, for example, **Default SMTP Virtual Server**.

When you delete a resource from a standard view or the **Master** view, you can choose to prevent the deleted object from being discovered at a later time. This option is useful when you are performing discovery operations on a scheduled basis and you do not want a resource to be rediscovered.

For example, let's say that in the **IIS** view, you delete the **SMTP Sites** resource from **IIS Server: PROD1** and select **Do not rediscover**. To discover and display **SMTP Site** information for **IIS Server: PROD1**, you must delete **IIS**

Server: PROD1 from the **IIS** or **Master** view and then, in the **Master** view, run the **IIS** Discovery Knowledge Script on the **PROD1** server.

In the case of a monitoring policy, preventing an object from being discovered will also prevent the monitoring policy from monitoring the resource. See [“Creating a Monitoring Policy” on page 264](#) for more information.

To delete a particular resource:

- 1 Select an object in the TreeView and click **TreeView > Delete Object**.
- 2 If you are deleting a resource from a standard view or the **Master** view, you are prompted to prevent the object from being rediscovered. Select **Do not rediscover** to prevent the object from being rediscovered and then click **Yes**.

If you are deleting a resource from a snapshot view, click **Yes** to delete the resource.

How to Discover a Deleted Resource

If you delete a resource and select the **Do not rediscover** option, the resource does not appear when you discover. To display the resource, you must delete the corresponding system or application from the standard view or the **Master** view and then rediscover.

Deleting a View

If you have added a monitoring policy to a view, you must remove the monitoring policy to delete a view. See [“Changing a Monitoring Policy” on page 267](#) for more information.

- 1 Click **View > View Manager**; select a view from the list.
- 2 Click **Delete**.
- 3 To close the View Manager, click **OK**.

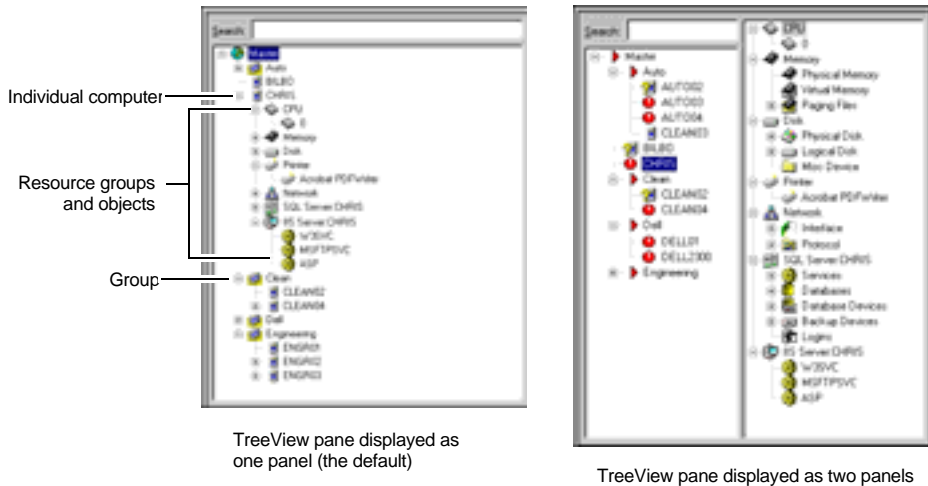
Expanding and Collapsing the Network Hierarchy

The TreeView pane is a key element of the Operator Console because it displays the computers (servers and workstations) and resources (applications like SQL or IIS Server, CPU, memory, and database objects) you are monitoring, plus indicators that tell you where jobs are currently running and where events have occurred.

The TreeView is organized as a hierarchy of managed **resource objects** and consists of:

- Servers and workstations (individual physical computers you are monitoring).
- Group folders to collect multiple computers.
- Individual resources such as applications, CPU, memory, disks, database devices, printers, network interface cards, and so on. (In some cases, folders are added to group

individual resources. For example, a “parent” Logical Disk folder is added to group individual disks.)



You can expand or collapse the network and resource hierarchy for any group, computer, or resource in the TreeView pane to see exactly where jobs are running or events have been raised, or to make starting and stopping jobs easier.

To	Do this
Expand (unfold) the hierarchy	Click the plus (+) to the left of the resource. Or select the resource, right-click, and then click Expand Branch .
Collapse the hierarchy	Click the minus (-) to the left of the resource. Or select the resource, right-click, and then click Collapse Branch .
Expand the hierarchy to the lowest level or collapse it to the highest level	Press the Ctrl key while expanding or collapsing.
Expand or collapse the hierarchy to where events are reported	Press the Shift key while expanding or collapsing. This shortcut allows you to find the source of an event quickly, without opening and closing tree branches individually.

Viewing Detailed Information about Discovered Resources

There are two ways to view detailed information about the resources in your environment.

Rest the cursor on any individual resource object in the TreeView pane to see a ToolTip. (Be sure **Show ToolTip for TreeView details** is selected in the Console tab in the Preferences dialog box.)



Or select the resource in the TreeView pane, and then click the **Details** tab in the List pane. The configuration details you see displayed vary depending on the object you have selected. For example, if you select a Memory resource object, the list includes the size (in KB) of physical and virtual memory. If you select a SQL Server database, you see the database name, owner, creation date, size, log size, and version.

Name	Owner	CreateDate	Size	LogSize	Version	CompatLevel
master	sa	Aug 15 1997 17:02	17	Shared with data	400	
model	sa	Aug 15 1997 17:02	1	Shared with data	400	
msdb	sa	Mar 15 1998 17:36	6	2	400	
pubs	sa	Aug 15 1997 17:09	3	Shared with data	400	
QC6	netiq	Apr 24 1998 12:53	100	50	400	
tempdb	sa	Oct 17 1998 8:54	12	Shared with data	400	

For Windows servers and workstations, the configuration details include the:

- Name of the server or workstation. For example, Shasta or Cobb1.
- Type of operating system installed. For example, Windows 2000 Server.
- Version of the operating system that's installed on the computer.
- Service pack version installed on the computer.


- Operating system build number.
- Date and time the operating system was installed on the computer.
- Type of processor installed on the system.
- Main directory for system resources and key files.
- Windows domain and role assigned to the computer.

Viewing AppManager System Information

To view detailed information about the AppManager configuration for a computer:

- 1 Right-click a computer in the TreeView pane and click **Properties**.
- 2 In the Computer Properties dialog box, click the **System** tab to view configuration information about AppManager.

- 3 In the list of **Available servers/groups**, expand the list and select the computer you want.



Option	Description
Managed object and Build	A list of the AppManager managed objects that are installed on this computer and the corresponding AppManager build version.
Primary management server	Indicates the name of the primary management server, if one was designated using the AMAdmin_SetPrimaryMS (for Windows computers) or the AMAdminUNIX_SetPrimaryMS (for UNIX computers) Knowledge Script.
Secondary management server	Indicates the name of the secondary (backup) management server, if one was designated using the AMAdmin_SetPrimaryMS (for Windows computers) or AMAdminUNIX_SetPrimaryMS (for UNIX computers) Knowledge Script.
Repository version	The version of the AppManager repository that is installed on this computer. If the computer is not a repository server, <not a repository server> appears in this field.
Time zone	The time zone setting for this computer.

- 4 To see the configuration information for another computer in the TreeView pane, in the list of **Available servers/groups**, select the computer you want.

Filtering Information Based on a Resource Object

In some cases you may want to display only the information associated with a particular computer, group, or resource. For example, you may want to see all jobs or events associated with a specific computer or group of computers.

To filter the information displayed in the List pane based on a particular resource:

- 1 In the TreeView pane, highlight the group, computer, or resource for which you want to view information.
- 2 Right-click; and then click **Filter List Pane by Selected Object**.
- 3 Click the tabs in the List pane to view information.

For example, click **Events** to see all of the events associated with the selected group, computer, or resource.

Once you have filtered information, you can perform additional tasks. For example, you can stop or restart all jobs or acknowledge all events associated with the selected group, computer, or resource.

Grouping and Filtering Based on Other Criteria

In addition to filtering by resource object, you can group and filter events, job information, and graph data by other criteria such as Knowledge Script name or event severity. You can also combine filtering and grouping schemes to filter by resource object to only display the information you want.

Setting Preferences

Preferences control certain aspects of how the Operator Console operates.

Setting Operator Console Preferences

To customize preferences for the Operator Console, click **File > Preferences**; then update the options in the **Console** tab. Changes you make to Operator Console preferences only apply to that console computer.



Console preference	Description
Show Tip of the Day	Select to display the Tip of the Day dialog box, which provides tips for using AppManager, each time you start the Operator Console. To disable this feature, deselect this option or deselect Show Tips at Startup in the Tip of the Day dialog box.
Flash warnings	Select to display an alert in the TreeView pane for lower severity events. If you deselect this option, only the most severe events (severity level 1 through 10) display an alert in the TreeView pane. However, event indicators are still displayed for less severe events under the Events tab in the List pane.

Console preference	Description
Flash icon on Taskbar	Select to flash the Operator Console button on the taskbar when an event is open and the application window is minimized.
Clear cache on startup	<p>AppManager uses a local cache to store information about changed records. The Operator Console foreground and background views are updated at regular intervals, using the information in the cache, to keep in sync with the repository.</p> <p>Select this option to clear the cache to free up some disk space on the computer when you start the Operator Console.</p> <p>Note It's recommended that you keep this option selected.</p>
Beep for new events	Select to hear an audible alert when events occur.
Split TreeView into two panels	<p>Select to display the TreeView pane in two panels, with computers and groups displayed in the left and resources displayed in the right. To see the change, restart the Console or re-connect to the repository (see "Connecting to Another Repository" on page 35).</p> <p>Notes</p> <ul style="list-style-type: none"> • If you are managing a large number of computers and groups in your environment, selecting this option can reduce the amount of time it takes to start the Operator Console. For more information, see "Customizing the Operator Console Layout" on page 38. • With the TreeView split into two panels, you can hide the right panel. Right-click anywhere in the TreeView pane and then deselect Show Details Panel.
Show ToolTips for TreeView details	<p>Select to display detailed configuration information about an object in the TreeView by resting the cursor on it.</p> <p>If you deselect this option, then click the Details tab in the List pane to see detailed configuration information about an object.</p>
Prompt to start, stop and close jobs	<p>Select to display a confirmation dialog box before starting, stopping, or closing a selected parent or child job. Use this option to avoid accidentally starting, stopping, or closing all of the child jobs under one or more selected parent jobs.</p> <p>If you deselect this option, no warning messages are displayed before AppManager starts, stops, or closes the selected jobs.</p>

Console preference	Description
Prompt to delete	<p>Select to display a confirmation dialog box before permanently deleting any items. Use this option to avoid accidentally deleting all of the child items under one or more selected parent items.</p> <p>If you deselect this option, no warning messages are displayed before AppManager deletes the objects, jobs, events, or graphs you have selected for deletion.</p>
Prompt to acknowledge and close events	<p>Select to display a confirmation dialog box before acknowledging or closing a selected event. Use this option to avoid accidentally acknowledging or closing an event.</p> <p>If you deselect this option, no warning messages are displayed before AppManager acknowledges or closes the selected event.</p>
Save View Manager settings on exit...	<p>Select to save the View Manager settings (in which you display only the views you are interested in) for subsequent Operator Console sessions.</p> <p>By default, all available views are displayed. For more information about using the View Manager, see "Selecting a View" on page 39.</p>
Display dates and times...	<p>Select to display all job, event, and graph data information related to date and time based on the local time of the managed client. If the managed client's local time zone setting is different from that of the computer on which the Operator Console is running, then GMT time zone bias information appears in parentheses after the date and time. For example, if the Operator Console computer is in San Jose, CA (USA) and the managed client is in Tokyo, Japan, then the time zone bias is (GMT +09:00).</p> <p>By default, date and time information is based on the local time of the computer where the Operator Console is installed.</p>
Filter events list by last number of days	<p>Select this option to specify the number of days that an event appears in the Events tab of the List pane.</p>

Console preference	Description
Foreground view	<p>To control how frequently information in the currently selected view is updated in the Operator Console, enter in seconds how often you want the information updated.</p> <p>If you change the view displayed (for example, from Master to NT), the view you switch to becomes the new foreground view and is updated accordingly.</p> <p>The default is 30 seconds.</p> <p>Note Depending on how you use the Operator Console (for example, to view activity periodically or to constantly monitor activity as it is going on), you may want to adjust the Foreground view and Background views preferences. More frequent updates refresh the Operator Console with information to give you a real-time view of your environment, but may have some impact on performance or may refresh the Operator Console unnecessarily.</p>
Background views	<p>Select to control how frequently information in background views is updated. Enter in minutes how often you want the information updated.</p> <p>If you change the view displayed (for example, from Master to NT), the Master view becomes a background view and is updated accordingly.</p> <p>The default is 5 minutes.</p>
Save current graph style as default	<p>To save the current graph style as the default style for creating graphs, click Save.</p> <p>Properties (such as the graph type, color, and fonts) of the currently selected graph are used as defaults for any new graphs you create.</p>
Save graphs on exit...	<p>Select to save the graphs currently displayed in the Graph pane and reload them the next time you start the Operator Console.</p>

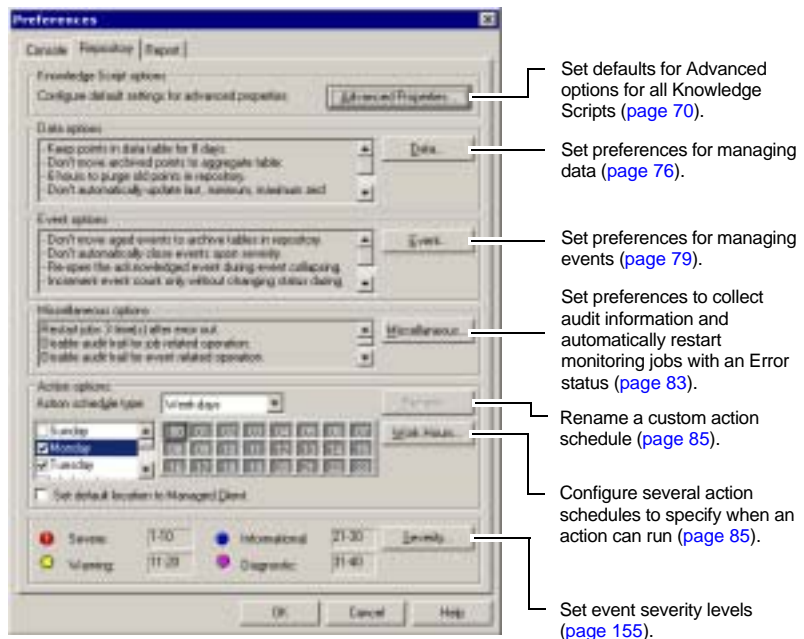
Console preference	Description
Customize the console layout	<p>Select one or more of the following options to customize the console layout. As a rule, you can improve the startup performance of the Operator Console by reducing the number of panes and tabs that are displayed on startup. You must restart the console or re-connect to the repository in order for your changes to the console layout to take effect:</p> <ul style="list-style-type: none"> • Work with Events. Deselect this option to remove the Events tab in the List pane and event-related menu commands (for example, List > Acknowledge Event). • Work with Jobs. Deselect this option to remove the Jobs tab in the List pane and remove the jobs pane, job-related menu commands (for example, List > Close Job), and the Knowledge Script pane. • Work with Details. Deselect this option to remove the Details tab in the List pane. • Work with Graphs. Deselect this option to remove the Graph Data tab in the List pane, graph data-related menu commands (for example, Graph > Delete Graph), and the Graph pane. • Customize button. Click a Work with... option and then click this button to change the layout of the corresponding tab in all views. For example, with Work with Events selected, click Customize to open the Show Event Columns dialog box. • Hide parent jobs in job pane. Select this option to hide all parent jobs in the Jobs tab of the List pane and list only the child jobs. • Hide parent events in event pane. Select this option to hide all parent events in the Events tab of the List pane and list only the child events. <p>Note For information on starting the Operator Console using command-line parameters, see the <i>Administrator Guide</i>.</p>

Setting Repository Preferences

To customize global preferences for the AppManager repository, click **File > Preferences**; then click the **Repository** tab. Changes you make to repository preferences settings apply to all repository users.

The Repository tab provides links to preferences for:

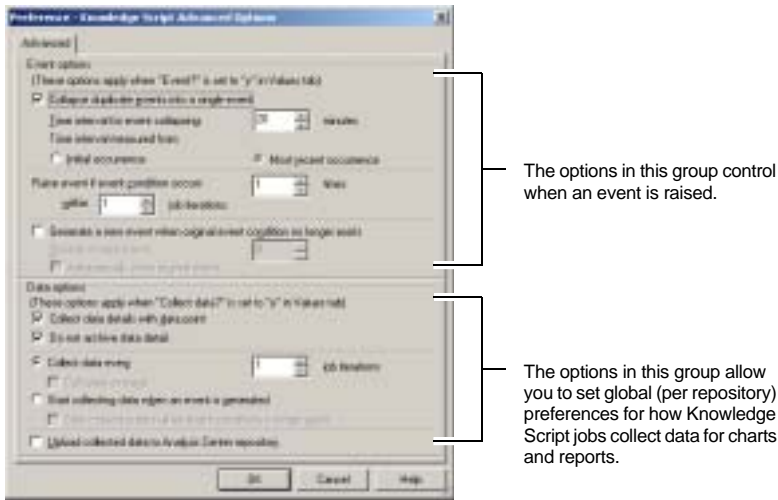
- Knowledge Script Advanced options
- Data management
- Event management
- Audit information and restart attempts
- Action options
- Event severity levels



To change the settings for a preference, click a button.

Set Preferences for Knowledge Script Advanced Options

In the Repository tab, click **Advanced Properties** to open the Preference - Knowledge Script Advanced Options dialog box.



Except for the **Generate a new event...** option, the options you set here are global across all new jobs. You can override these settings for a particular job. See [“Setting Advanced Job Properties” on page 121](#) for more information.

Knowledge Script advanced options	Description
Event options	
These options apply when Event? is set to y in the Values tab on the Job Properties dialog.	
The options in this group control when an event is raised, based upon a specified number of event occurrences during a specified number of job iterations. If a duplicate event is raised, use these options to either collapse duplicate events or automatically close the original event when the event condition no longer exists.	
Note We recommend that you see “Duplicate Events” on page 174 for a more complete explanation of event collapsing and “Automatically Closing Events When the Condition No Longer Exists” on page 180 for more information on automatically closing events.	

Knowledge Script advanced options	Description
Collapse duplicate events into a single event	<p>Select to have duplicate events collapsed into a single event.</p> <p>An event is considered a duplicate when a recent Open child event with the same object name, event message, severity, and Job ID as a previous event occurs. Just like unique events, duplicate ones are logged in the repository and, if event collapsing is not enabled, are listed as additional child event entries under the parent event. Actions, if specified, are initiated for each duplicate event.</p> <p>With event collapsing enabled, instead of creating new child event entries, the duplicate events are collapsed into the original child event and the child event count is increased. The duplicate events are logged in the repository; however, actions associated with collapsed duplicate events, whether the actions are initiated on the managed client or management server, are suppressed.</p> <p>Duplicate events are collapsed within a specified time interval (see Time interval for event collapsing).</p> <p>Note We recommend that you see “Duplicate Events” on page 174 for a more complete explanation of event collapsing.</p>
Time interval for event collapsing	<p>Set the time interval in minutes for collapsing duplicate events into a single event. The default is 20 minutes.</p> <p>This time period is fixed. For example, assuming a 20 minute time interval, events at 12:00 and 12:19 will be collapsed into a single event, while events at 12:00 and 12:21 will be raised as separate events. This option is useful to raise events at specific intervals until you are able to address the situation (in the case of the above example, one event every 20 minutes).</p>

Knowledge Script advanced options	Description
Time interval measured from...	<p>Select an option to specify when the time interval for event collapsing begins:</p> <ul style="list-style-type: none"> • Initial occurrence. This option sets the time interval to a static period of time that begins when the first event occurs. For example, using the default time of 20 minutes, the 20 minute interval is always restarted every 20 minutes. • Most recent occurrence. This option restarts the time interval for event collapsing each time an event is raised (the event collapsing interval is not a static period of time). For example, using the default time of 20 minutes, if a job generates duplicate events every five minutes, then the event is collapsed and the 20 minute interval is restarted every five minutes, which means it never effectively expires – unless you set an option to have events ignored (see “Raise event if event condition occurs N times...” on page 73). <p>This option is selected by default.</p>

Knowledge Script advanced options	Description
Raise event if event condition occurs N times...	<p data-bbox="669 245 1150 321">Enter a threshold value to specify the number of duplicate events that must occur during a specified number of job iterations before an event is raised.</p> <p data-bbox="669 334 1153 435">An event is considered a duplicate when a recent Open child event with the same object name, event message, severity, and Job ID as a previous event occurs.</p> <p data-bbox="669 448 1188 602">For example, when the threshold value is set to 3 and the number of job iterations is set to 5, if an event condition occurs each time the job runs during the interval, the first 2 event conditions are ignored and the third event condition generates an event and starts the “Time interval for event collapsing” on page 71.</p> <p data-bbox="669 615 1188 846">The subsequent events that are raised during the interval are ignored—only 2 events are raised during the interval, which falls below the threshold. The event count resets when the interval resets. In this case, to generate another event, a duplicate event must occur 3 times during the next 5 job iterations and if event collapsing is enabled, the event must be generated within the “Time interval for event collapsing” on page 71.</p> <p data-bbox="669 859 1181 959">Ignored duplicate events are not logged in the repository and do not increment the event count of the original child event. Actions associated with ignored events are also suppressed.</p> <p data-bbox="669 972 1161 1049">The default value for the number of duplicate events and the number of job iterations is 1, which means duplicate events are not ignored.</p> <p data-bbox="669 1062 727 1078">Notes</p> <ul data-bbox="669 1084 1188 1321" style="list-style-type: none"> • Use this option to filter temporary spikes in thresholds being met, improving performance by reducing the number of potentially trivial events being reported to the management server. • We recommend that you see “Duplicate Events” on page 174 for a more complete explanation of event collapsing and “Automatically Closing Events When the Condition No Longer Exists” on page 180 for more information on automatically closing events.

Knowledge Script advanced options	Description
...within M iterations	<p>Enter the number of job iterations to specify the interval for ignoring duplicate events. The interval is reset when the number of job iterations matches the interval.</p> <p>The default value for the number of duplicate events and the number of job iterations is 1, which means duplicate events are not ignored.</p>
Generate a new event when original event condition no longer exists	<p>Select this option to raise an informational event when an event condition no longer exists. This preference is not applicable for jobs running on UNIX computers.</p> <p>With this option selected, the first time an event condition is encountered, a parent event and a child event are created; no additional events are raised as long as the event condition persists, and the event count for the child event does not increment. When the event condition no longer exists, a new event is created to indicate the change in the event condition. Select the Automatically Close Original Event option to automatically close the original event.</p> <p>Note The status of the informational event that is created by AppManager in response to the changed event condition is Open.</p> <p>For example, when a job detects that physical memory usage exceeds its threshold, an event is raised. When the job detects the usage is below the threshold, the original event condition does not exist so AppManager creates an informational event (to indicate the change in the event condition) with a specified event severity.</p>
Severity of new event	<p>Specify the severity for the event generated in response to a change in the original event condition. This option is not applicable for jobs running on UNIX computers. The default severity level is 20.</p> <p>Note We recommend setting the severity level for this type of event to a lower severity than the original event.</p>

Knowledge Script advanced options	Description
Automatically close original event	<p>Select this option to automatically close the original event when the event condition no longer exists. This option is not applicable for jobs running on UNIX computers.</p> <p>For example, when physical memory usage rises above the threshold, an event is raised. When the usage falls below the threshold, the event condition no longer exists and AppManager closes the event and then generates an informational event to document the closing of the original event.</p> <p>We recommend that you see “Automatically Closing Events When the Condition No Longer Exists” on page 180 for a more complete explanation of how to raise an informational event and automatically close the original event when an event condition no longer exists.</p>
Data options	
<p>The options in this group allow you to set global (per repository) preferences for how Knowledge Script jobs collect data for charts and reports.</p> <p>These options apply when Collect data? is set to y in the Values tab.</p>	
Collect data details with data point	<p>This option only applies to data collection for AppManager graphs and charts.</p> <p>Deselect this option to only collect the value of the monitored resource; detailed information, such as server name and collection time, is not collected. This option is selected by default.</p>
Do not archive data detail	<p>This option only applies to data collection for AppManager reports that display detail data, such as ReportAM_DetailData.</p> <p>Select this option to only collect the value of the monitored resource; detailed information, such as server name and collection time, is not collected. This option is selected by default.</p>
Collect data every N job iterations	<p>Create a data point every N times that a job runs. For example, if a Knowledge Script job is set to run every 10 minutes and you set this value to 3 iterations, then a data point will be created every third time the job runs, or every 30 minutes.</p> <p>The default is 1 iteration, which means a data point is created every time the job runs.</p>

Knowledge Script advanced options	Description
Calculate average	<p>If you have configured the repository to create a data point every N iterations, you can select this option to create that data point as an average value of all data points collected during the interval. For example, if you configure the repository to collect a data point every 3rd time a job runs, and the data point for the first iteration is 2, the data point for the second iteration is 5, and the data point for the third iteration is 10, the data point is created on the third iteration and its value is 5.6666.</p> <p>This information can be displayed in the Graph Data tab of the List pane in the Operator Console by customizing the Graph Data tab to display the Average column. See “Changing the Layout of the Graph Data Tab” on page 189 for more information.</p>
Start collecting data when an event is generated	<p>Select to have a Knowledge Script job begin collecting data when an event is raised. For example, if a Knowledge Script job is set to raise an event when CPU usage exceeds 70%, then the job will also start collecting data at that point.</p> <p>Note The data header is always created when you start a job. If you select this option, the Graph Data tab displays the data stream with the current number of data points.</p>
Stop collecting data when the event condition no longer exists	<p>Select to have a Knowledge Script job stop collecting data when the event condition no longer exists. Following the previous example, the job will stop collecting data when CPU usage falls below 70%.</p>

Set Preferences for Data Management

In the Repository tab, click **Data** to open the Preference - Data Options dialog box. These options configure repository

preferences for storing, archiving, and updating data for charts and AppManager reports.



Data options	Description
Default period to keep points in data table	<p>Specifies a time period, in days, for keeping data points available for display in the Graph pane of the Operator Console and in the Chart Console. This preference does not affect the availability of data for reporting purposes.</p> <p>At the next Time interval to purge old points in repository, if the age of the data point exceeds the time period you set, the data point is deleted from the Data table of the AppManager repository and is no longer available for display in a graph or chart. The default time period is 8 days.</p> <p>If you change the repository preferences for the number of days that data points are available, the change affects new jobs only. Existing data streams must be updated to change the number of days to keep data points in the data stream, up to the maximum set for the repository. For information on configuring the number of days that points for a particular data stream can appear in a graph or chart, see “Changing Data Stream Properties” on page 192.</p> <p>To minimize the size of the repository database, consider running AppManager reports for historical charting purposes rather than increasing the time period to keep data for charts and graphs. For more information on how to manage data, see the <i>Administrator Guide</i>.</p>

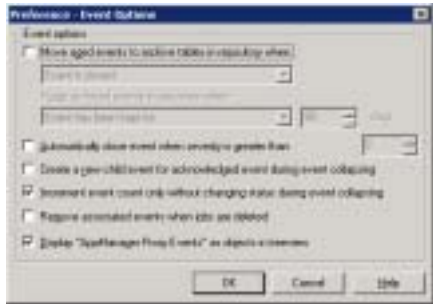
Data options	Description
Move archived points to aggregate table	<p>Select this option to specify a maximum threshold value, in months, for keeping data points in the repository available for AppManager reports.</p> <p>If the age of the data exceeds the threshold, the repository database calculates an hourly average, minimum, maximum, sum, and count value and moves the aggregate values to the ArcAvgHourlyData table.</p> <p>When data values are aggregated, the original data points are no longer available for reports. This preference does not affect the availability of data for charts and graphs. To report on aggregate values, you must run a report that uses aggregate data.</p> <p>By default, this option is disabled and the repository never deletes data points used for reporting. Changes to this preference take effect with the next interval of the AppManager repository's internal tasks.</p> <p>If you do not enable this option, the AppManager database administrator must manually manage data for reports.</p> <p>For more information on how to manage data, see the <i>Administrator Guide</i>.</p>
Time interval to purge old points in repository: <i>n</i> hours	<p>Specify the interval, in hours, to check for and delete data values that exceed the threshold for keeping data for charts and graphs. This preference does not affect the availability of data for reports. The default is 6 hours.</p> <p>Note To configure this option, you must be granted the SQL Server System Administrator server role.</p> <p>Changes to this preference take effect with the next interval of the AppManager repository's internal tasks.</p> <p>For more information on how to manage data, see the <i>Administrator Guide</i>.</p>
Automatically update last, minimum, maximum and average value for graph data	<p>Select to populate these columns in the Graph Data view of the List pane with the correct values. This setting is disabled by default.</p> <p>For information on displaying the Average column in the Graph Data tab, see “Changing the Layout of the Graph Data Tab” on page 189.</p> <p>Changes to this preference take effect with the next interval of the AppManager repository's internal tasks.</p>

Data options	Description
Remove associated graph data when jobs are deleted	<p>Select to automatically delete all graph data associated with a deleted job, including policy-based jobs. This preference does not affect the availability of data for reporting purposes. This setting is disabled by default.</p> <p>Changes to this preference take effect with the next interval of the AppManager repository's internal tasks.</p> <p>For more information on how to manage data, see the <i>Administrator Guide</i>.</p>
Collect data for physical nodes of virtual server	<p>Select to collect data about a clustered application on a physical node of a virtual server. This option only applies to AppManager modules that can collect physical node information. If this option is not selected, information about clustered applications is collected for the virtual server.</p>

Set Preferences for Event Management

In the Repository tab, click **Event** to open the Preference - Event Options dialog box. The options in this dialog box define when to archive events and when to purge archived events from the repository. By default, all options are disabled.

If you make changes to these preferences, your changes take effect with the next interval of the AppManager repository's internal tasks.



Event options	Description
Move aged events to archive tables in repository when	<p>Select one of the following options to set the schedule for archiving and purging of event data in the repository:</p> <ul style="list-style-type: none">• Event is closed This option moves events to the archive tables in the repository when the events are closed.• Event stays open status for n Days/Hours This option moves open events to the archive tables in the repository after the specified number of Days/Hours.• Event stays closed status for n Days/Hours This option moves closed events to the archive tables in the repository after the specified number of Days/Hours.• Closed event severity is greater than n severity This option moves closed events to the archive tables in the repository when the event severity is greater than the specified level. <p>For more information on managing events, see the <i>Administrator Guide</i>.</p>

Event options	Description
Purge archived events in repository when	<p>Select one of the following options to set the schedule for removing archived events from the repository:</p> <ul style="list-style-type: none"> • Event has been kept for n days This option purges events that have been kept in the repository for the specified number of days. • Oldest events exceed number of records: n events This option purges events that exceed the maximum number specified to be kept in the repository. <p>For more information on managing events, see the <i>Administrator Guide</i>.</p>
Automatically close event when severity is greater than	Select to automatically close any event with an event severity level that is higher than the specified threshold value.
Create a new child event for acknowledged event during event collapsing	<p>When event collapsing is enabled, this option changes how the Operator Console displays event information for duplicate events that are raised after you Acknowledge a child event.</p> <p>Select this option to create a new child event when a duplicate event is raised; the status of the Acknowledged child event does not change.</p> <p>With this option deselected, when a duplicate event is raised, the status of the Acknowledged child event changes to Open.</p> <p>If you select this option, you cannot select the Increment event count only... option (see below). This option is deselected by default.</p>
Increment event count only without changing status during event collapsing	<p>When event collapsing is enabled, this option changes how the Operator Console displays event information for duplicate events that are raised after you Acknowledge a child event.</p> <p>Select this option to increment the child count for an event when a duplicate event is raised; the status of the Acknowledged child event does not change.</p> <p>With this option deselected, when a duplicate event is raised, the status of the Acknowledged child event changes to Open.</p> <p>If you select this option, you cannot select the Create a new child event for... option (see above). This option is deselected by default.</p>


Event options	Description
Remove associated events when jobs are deleted	Select to automatically delete all events associated with a deleted job, including policy-based jobs. This option is disabled by default.
Display "AppManager Proxy Events" as objects in treeview	<p>If you are experiencing issues with the overall performance of the Operator Console, you may be able to improve the overall performance of the Operator Console by disabling this option. This option is disabled by default.</p> <p>Select this option to display event information for a computer that is monitored by a proxy Knowledge Script under its own parent event. For example, if you are monitoring the SJCCOMEYT01 and SJCCOMEYT03 computers with General_MachineDown, event information for each computer is organized under its own parent event.</p> <p>In addition, if the computers you are monitoring are not in the TreeView pane of the Operator Console, a new group, AppManager Proxy Events, is automatically created in the Master view to display the event information. From this group, you can view, acknowledge, close and delete proxy-based events on these computers.</p> <p>Here are some examples of Knowledge Scripts that do not require the AppManager agent to be installed on the computer you want to monitor (proxy Knowledge Scripts):</p> <ul style="list-style-type: none"> • General_MachineDown • General_MachineDownLR • General_PingMachine • NT_RemoteServiceDown • NT_RemoteServiceDownLR • UNIX_PortHealth <p>Note that you cannot run jobs on computers in the AppManager Proxy Events group. To discover the agent and run monitoring jobs, you must first delete the computer from the AppManager Proxy Events group then add the computer to the TreeView pane using the Add Computers dialog box. If necessary, stop any proxy-based Knowledge Scripts that are monitoring the computer so that you can delete the computer from the AppManager Proxy Events group.</p>

Set Preferences for Audit Information and Restart Attempts

In the Repository tab, click **Miscellaneous** to open the Preference - Miscellaneous dialog box. The options in this dialog box permit you to collect audit information on jobs, events, and actions, and to specify the number of times to restart monitoring jobs with an Error status.

Note For information on setting the number of restart attempts for a monitoring policy, see [“Creating a Monitoring Policy” on page 264](#).

If you make changes to these preferences, your changes take effect immediately.



Miscellaneous options	Description
Enable audit trail for...	<p>Select an option to collect audit information about:</p> <ul style="list-style-type: none">• ...Job related operation. Select this option to collect audit information on Knowledge Script jobs.• ...Event related operation. Select this option to collect audit information on events.• ...Action related operation. Select this option to collect audit information on actions. <p>Audit information for maintenance mode activity is automatically collected. If the computer is disconnected or the AppManager agent is not running, a status of “start-maintenance” or “stop-maintenance” can be recorded in the computer maintenance audit trail even though machine maintenance was not actually performed. For more information on maintenance mode, see “Initiating Maintenance on a Computer” on page 145.</p> <p>Note For information about viewing AppManager audit information, see “Viewing Audit Information” on page 101.</p>

Miscellaneous options	Description
Number of times to restart error jobs	<p>Select an option to specify the number of times to automatically attempt to restart a job with an Error status. If the job does not run after the specified number of retry attempts, the job status remains Error.</p> <p>The repository database attempts to restart error jobs each time SQL Server runs the Task_Minutely job. By default, the Task_Minutely SQL Server job runs every 5 minutes.</p> <p>If you attempt to restart an error job, the number of times that the job can be automatically restarted by the repository database is reset. For example, if you restart an Error job and the job does not run, AppManager attempts to restart the job the specified number of times. The default is to restart an error job 3 times.</p> <p>Note For information on setting the number of restart attempts for a monitoring policy, see “Changing a Monitoring Policy” on page 267.</p>

Set Preferences for Action Schedules

There are preferences for standard action schedules, such as **Weekdays** and **Weekends**, and custom schedules, such as **Custom1** and **Custom2**. You can configure the days and times that standard and custom actions schedule are enabled; additionally, you can rename custom action schedules. For information about using action schedules, see [“Specifying One or More Corrective Actions” on page 115](#).

There is also a preference to change the default configuration for all Knowledge Scripts in the repository to specify where a responsive action is initiated.


If you make changes to these preferences, your changes take effect immediately. To apply your changes to all existing jobs, you must stop and restart each job. Action schedules are not applicable when monitoring a UNIX computer.

- By default, Knowledge Scripts are configured to initiate an action from the Management Server computer. Select the **Set default location to Managed Client** checkbox to

change the default configuration for all Knowledge Scripts in the repository to initiate a responsive action from the managed client computer.

- To rename a custom action schedule, select a custom action schedule, such as **Custom1** or **Custom2**, from the **Action schedule type** list and click **Rename**. You cannot rename the standard action schedules.
- To view the days and times during which an action schedule is enabled, select a schedule from the **Action schedule type** list. The enabled days are indicated with a checkmark; select a day from the list to view the enabled hours which are highlighted with a gray background.
- To change the days and times during which an action schedule is enabled, click the **Work hours** button. In the

Customize Action Work Hours dialog box, configure the hours you want for an action schedule.



Action Work Hours options	Description
Schedules	<p>Configure the hours during which an action can run by defining one or more action schedules. Select a schedule to set its hours of operation:</p> <ul style="list-style-type: none">• Weekdays• Weekends• Work Hours• Non-Work Hours• Custom... <p>Notes</p> <ul style="list-style-type: none">• There are 15 Custom schedules which can be configured to meet your needs.• When configuring an action schedule for use with a job that is configured to run on a particular weekend of the month (such as first weekend, second weekend, etc.), the job runs on Sunday and the following Saturday.

Action Work Hours options	Description
Scheduled days	<p>After you select a schedule, click a checkbox to select each day of the week to be included in the time frame:</p> <ul style="list-style-type: none"> • Sunday • Monday • Tuesday • Wednesday • Thursday • Friday • Saturday
Hours of the day	<p>Select the hours during which an action can run for the selected day. Hours are specified in 24-hour format (00 to 23). For example, if you select 08 and 09, the action schedule allows an action to run from 8:00 a.m. to 9:59 a.m.</p> <p>To select an hour, click a number to change the background color to gray. To deselect an hour, click a number to change the background color to white.</p>

Set Preferences for Event Severity Levels

In the Repository tab, click Severity to open the Customize Severity dialog box and view the current event severity levels.



See [“Event Severity Levels” on page 155](#) for information about customizing event severity levels.

Setting Report Preferences

Use the **Report** tab to configure AppManager reporting preferences, including:

- The URL that the Operator Console **Extensions** menu uses to view reports. When you click **Extensions > Report Viewer**, this URL identifies the location of AppManager reports on the AppManager web management server.
- Additional AppManager repository data sources. Although you can configure the report agent to report on data in additional repositories, this is not recommended.

It is recommended that you use a separate report agent for each data source, and install each report agent on a dedicated server. Here are some considerations:

- Although you can configure a report agent to have more than one data source, at this time, an AppManager report can only include data from a particular AppManager repository. An AppManager report cannot include data from more than one AppManager repository.
- Depending on the amount of information in each report and the number and frequency of reports you run, the report agent can consume significant system resources.

See the *Installation Guide* for information on installing the AppManager report agent.

To customize preferences for the Operator Console, click **File** > **Preferences**; then update the options in the **Report** tab.



To	Do this
Configure the Operator Console Extensions menu to view reports	<p>Specify the URL that identifies the AppManager web management server and virtual directory that is linked to the reports. By default, AMReports is the virtual directory linked to the reports. Use the following syntax:</p> <p>http://Web_server/virtual_directory</p> <p>For example: http://Server1/AMReports</p> <p>Tip To view reports from the Operator Console, open the Report Viewer by clicking Extensions > Report Viewer.</p>
Data Sources	<p>This list displays the available repository data sources for AppManager report agents that belong to this AppManager site. Click a button to configure a repository data source:</p> <ul style="list-style-type: none">• Add Adds a new a repository data source.• Edit Edits an existing repository data source.• Delete Deletes an existing repository data source.

Displaying Toolbars

AppManager provides toolbars for quick access to common commands:

This toolbar	Contains commands
General	Found on the File menu.
View	For working with the View Manager and panes.
TreeView	For working with objects in the TreeView pane.
Knowledge Scripts	For working with Knowledge Scripts.
List	For working with jobs, events, and data streams.
Graph	For working with graphs.
Web	For working with documents in the Web pane.

You can selectively show or hide these toolbars.

To control which toolbars are displayed and the style of the toolbars: click **View > Toolbars**.

For information about the commands available on each toolbar, see the Operator Console online help.

Displaying a Status Bar

The status bar in the bottom portion of the Operator Console provides a summary of the number of computers being monitored, the total number of open child events, the total number of child jobs, and the total number of graph data streams being collected.



Computers: 99 | Open Events: 0 | Running Jobs: 16 | Graph Data: 441 | 1:50 PM

To show or hide the status bar: click **View > Status Bar**.

Getting Information about Knowledge Scripts

To display a short description of each Knowledge Script in the Knowledge Script pane, right-click in the pane and then click **Show Description**. (If text is truncated, pass the mouse cursor over it and the full text will appear as a ToolTip.)



To view the development history of a Knowledge Script, including the version number and any comments recorded by the developer, right-click the Knowledge Script in the Knowledge Script pane and then click **Version History**.

For detailed information about an individual Knowledge Script, select the script in the Knowledge Script pane and then press the F1 key; or click **Help** on the **Values** tab in the Knowledge Script Properties dialog box.

Checking Communication with Managed Computers

From time to time, the icons for managed Windows NT and Windows 2000 computers may appear dimmed (grayed out) in the Operator Console, which means the AppManager management server is unable to communicate with the computer. While in this state, you can't run Knowledge Scripts against the computer. In the background, the management server periodically checks whether it can communicate with the computer. Once communication is reestablished, the icon is restored to its normal state.



This icon, which only appears in the TreeView if the AppManager agent is not running on the computer, means that only Knowledge Scripts related to installing the AppManager agent can be run.

Instead of waiting for the management server's next polling interval, you can check the communication with a managed computer yourself. You might want to do this if you shut down a computer to perform some maintenance (causing it to be dimmed in the Operator Console). When you reboot the computer, you can check the connection immediately rather than wait for the management server to do it. Click **TreeView > Ping Computers**.

If a managed computer's icon remains dimmed for more than twenty minutes, check for network communication problems and the availability of the management server, AppManager repository, and the managed computer itself. If the availability of the managed computer isn't a network connection problem, you may need to use the NetIQctrl utility to track down the source of the problem. For more information about using NetIQctrl, see the *Administrator Guide*.

Integrating AppManager with a Custom Web Application

You can configure the Operator Console to open a Web document, such as a custom Web application, in the Web pane of the Operator Console. Note that by default, the Web pane is not displayed. To display the Web pane, click **View > Web Pane**.

When specifying the URL that you want to appear in the Web pane, you can include AppManager variables in order to integrate AppManager with a custom Web application. For example, if you specify the URL as `http://intranet/customweb/computerdetails.asp?Computer=%COMPUTER%`, when

you click the Prod1SQL computer in the TreeView pane, the Operator Console substitutes the computer name into the URL and updates the Web pane. The Web pane opens the Web document named `http://intranet/customweb/computerdetails.asp?Computer=Prod1SQL`.

Specifying a Document to Open

To open a document in the Web pane, the Web pane must be displayed. To display the Web pane, click **View > Web Pane**.

- 1 In the Operator Console, click **Web > Properties**.
- 2 In the Web Pane Properties dialog box, specify the URL of the document you want to open. Here's a list of AppManager variables that you can include as part of the URL.

To substitute the	Specify
Computer name of a managed client	%%COMPUTER%% For example, <code>http://intranet/computerdetails.asp?Computer=%%COMPUTER%%</code>
Object ID of a managed resource	%%OBJECTID%% For example, <code>http://intranet/computerdetails.asp?Object=%%OBJECTID%%</code>

- 3 Click **OK**.

The Web pane displays the specified document. If you have included a computer or object ID variable as part of the URL, click a corresponding object in the TreeView pane to update the Web pane.

Working with Documents in the Web Pane

Here's a list of commands that you can use when working with documents in the Web pane.

To	Do this
View pages you have recently visited	<ul style="list-style-type: none">Click Web > Previous to view the previous document.Click Web > Next to view the document you were viewing before you clicked Web > Previous.
Stop downloading the current document	<ul style="list-style-type: none">Click Web > Stop.
To reload the current document	<ul style="list-style-type: none">Click Web > Refresh.

These commands are also available from the Web toolbar. For information on displaying the Web toolbar, see [“Displaying Toolbars” on page 91](#).

Using the Extensions Menu

The Operator Console Extensions menu provides access to tools, some of which are not part of AppManager, that can be useful to perform related system management or diagnostic tasks.

You can modify the Extensions menu by adding and deleting tools, and editing a menu item.

Accessing Tools

In the AppManager Operator Console, click a command on the Extensions menu:

This command	Does this
Security Manager	Starts the NetIQ AppManager Security Manager, which allows you to update security information and create user roles.
Repository Browser	Starts the NetIQ AppManager Repository Browser, which allows you to browse the AppManager repository database. See the <i>Administrator Guide</i> for more information.
Report Viewer	Starts the Report Viewer, which allows you to view AppManager reports. The Report Viewer displays the same report information that is available from the Reports view of the Operator Web Console.
Chart Console	Starts the NetIQ Chart Console, which allows you to generate and view charts of data streams generated by Knowledge Script jobs. For more information, see the AppManager Help.
ADCheck	ADcheck quickly checks domain controllers, domains, and replication. This tool is not applicable when monitoring a UNIX computer. Note ADcheck is not included on the AppManager CD and must be installed separately.
NetIQCtrl	Starts a command-line interface for checking the status of AppManager components. For more information, see the <i>Administrator Guide</i> .
Event Viewer	Starts the Windows Event Viewer. This tool is not applicable when monitoring a UNIX computer. For more information, see your Windows documentation.
Performance Monitor	Starts the Windows Performance Monitor. This tool is not applicable when monitoring a UNIX computer. For more information, see your Windows documentation.
Computer Management	Starts the Windows Computer Management console. This tool is not applicable when monitoring a UNIX computer. For more information, see your Windows documentation.

This command	Does this
Web Recorder	Starts the NetIQ Web Recorder. The Web Recorder is a browser-based application that lets you record the customer experience of a Web site and then set thresholds that measure the performance of that site on an ongoing basis. For more information, see the <i>Monitoring Response Time for Web with AppManager</i> . This Work Smarter Guide is provided in online format (Adobe Acrobat) only and can be found on the AppManager Control Center installation kit in the \Documentation\Management Guides folder.
URL Check Recorder	Starts the NetIQ URL Check Recorder. The URL Check Recorder generates URLCheck Knowledge Scripts that validate one or more URLs for availability and performance. For more information, see the <i>Monitoring Response Time for Web with AppManager</i> . This Work Smarter Guide is provided in online format (Adobe Acrobat) only and can be found on the AppManager Control Center installation kit in the \Documentation\Management Guides folder.
Diagnostic Console	<p>Diagnostic Console expands on the problem diagnosis features of AppManager by providing you with a collection of features that enable you to remotely investigate and diagnose problems occurring on Windows servers in your environment. This tool is not applicable when monitoring a UNIX computer.</p> <p>The Diagnostic Console must be installed separately. The setup program is provided on the AppManager Suite installation kit in the \Diagnostic Console folder.</p> <p>Tip To diagnose a selected computer in the TreeView pane, right-click a Windows computer and select Diagnose Selected Computer. The Diagnostic Console displays summary information for the selected computer. Note that in order to diagnose a computer using the Diagnostic Console, the Windows user account under which the console operator is currently logged in as must have Administrator privilege on the selected computer. If the current Windows user account does not have Administrator privilege on the computer to be diagnosed, the console operator is prompted to enter a valid Windows user account.</p>
Audit Trail	<p>View audit information collected by the AppManager repository. For more information, see “Viewing Audit Information” on page 101.</p> <p>Note You must enable repository preferences to collect audit information; see “Set Preferences for Audit Information and Restart Attempts” on page 83 for more information.</p>

This command	Does this
Save SQL Password	Allows you to save the password of an SQL login account that is allowed to run SQL Knowledge Scripts on a computer. For more information, see “Saving a Password for a SQL Server Login Account” on page 99 .
Customize	Allows you to modify the Extensions menu by adding and deleting tools. For more information, see “Adding a Tool” on page 98 and “Deleting a Tool” on page 99 .
Refresh Tables	Refreshes the local cache in the AppManager Operator Console with up-to-date information (all tables).

Note If in the NetIQ AppManager Security Manager a user is assigned the Read-Only User security role, then none of the commands on the Extensions menu will be available to that user.

Adding a Tool

You don't need to know Visual Basic Script to add a menu item; AppManager creates the VB script for you.

- 1 In the AppManager Operator Console, click **Extensions > Customize**.
- 2 Click **New**.
- 3 In the New Extension dialog box:

For	Do this
Menu title	Type the name you want to appear on the Extensions menu.
Script name	The menu title, with a .VBS extension, becomes the name of the VB script file AppManager creates. The file is stored in \Netiq\AppManager\bin\plugins.
Command line	Type the name of the executable file.
Parameters	Type any parameters associated with the executable file.

- 4 Click **OK**.

Editing an Extensions Menu Item

If you are familiar with VBScript, you can edit the script. For more information on VBScript, see your Microsoft documentation.

- 1 In the AppManager Operator Console, click **Extensions > Customize**.
- 2 Select a script in the list.
- 3 Click **Edit**.

Note It's recommended that you not edit NetIQ Extensions.VBS, which is the script for the standard AppManager Extensions menu commands.

Deleting a Tool

- 1 In the AppManager Operator Console, click **Extensions > Customize**.
- 2 Select a script in the list.
- 3 Click **Delete**.
- 4 Click **Yes**.

Saving a Password for a SQL Server Login Account

Some Knowledge Scripts for SQL Server require a SQL login account to run. In the NetIQ AppManager Security Manager, you select the SQL login accounts that are allowed to run SQL Knowledge Scripts on a computer or enter the SQL login account you want and enter a password. For each SQL login account you select, you are prompted for the account password. For more information, see the *Administrator Guide*.

When working in the NetIQ AppManager Operator Console, you can save SQL login account passwords.

To save a SQL password in the repository:

- 1 In the AppManager Operator Console, click **Extensions > Save SQL Password**.



Select the SQL Server you want from the list, or...



... If the Operator Console cannot resolve the SQL Server name, type the name here.

- 2 Select the computer for which you want to save SQL login passwords.

If you first selected a computer in the TreeView pane before opening the Save SQL Password for AppManager dialog box, the selected computer's name appears in the Computer list.

- 3 In the SQL Server drop-down list, click the SQL Server you want. You can save a SQL login password for each server in a cluster.

If the Operator Console cannot resolve the SQL Server name, type the name in the SQL Server field.

- 4 Type the name of the SQL login account whose password you want to change.

- 5 Type the SQL login account password.
- 6 If changing only one password, click **Close**.

To update multiple passwords, select different computers or type different SQL logins, type the new password, and then click **Save**. When you finish saving passwords, click **Close**.

The password is verified, encrypted and saved in the repository.

Note Saving the password simply stores the information in the AppManager repository. If the password is changed on the SQL Server, then it must also be changed in the AppManager repository.

Viewing Audit Information

If you have configured the AppManager repository preferences to collect audit information for jobs, events, and actions, you can view this information, and information about maintenance mode activity, from the Operator Console. Audit information about maintenance mode is automatically collected and can also be viewed.

Note that if the computer is disconnected or the AppManager agent is not running, a status of “start-maintenance” or “stop-maintenance” can be recorded in the maintenance history even though the operation is not actually performed.

To view information about	Do this
Jobs, including when the job status changed and the name of the user who changed the job.	In the Jobs tab of the List pane, click to select a job and then click Extensions > Audit Trail > Jobs .
Actions, including the name of the action that was initiated, when it was initiated, and when the action was completed.	In the Jobs tab of the List pane, click to select a job and then click Extensions > Audit Trail > Actions .

To view information about	Do this
Events, including when the event was created, when the event status changed, and the name of the user who changed the event.	In the Events tab of the List pane, click to select an event and then click Extensions > Audit Trail > Events .
Maintenance mode, including when machine maintenance or scheduled maintenance was enabled or disabled, and the name of the user who enabled or disabled maintenance mode.	In the TreeView pane, click to select a computer and then click Extensions > Audit Trail > Computer Maintenance .

Note If there is no audit information for the selected resource, the title bar of the Audit Trail dialog box displays Error.

For information on configuring repository preferences to collect audit information for jobs, events, and actions, see [“Set Preferences for Audit Information and Restart Attempts”](#) on page 83.

Running Monitoring Jobs

This chapter discusses the basics of working with jobs. You can also configure AppManager to run jobs automatically, as part of a monitoring policy. The following topics are covered:

- [About Knowledge Scripts and Jobs](#)
- [Starting a New Job](#)
- [Setting the Schedule](#)
- [Setting the Values to Monitor](#)
- [Specifying One or More Corrective Actions](#)
- [Selecting the Objects to Monitor](#)
- [Setting Advanced Job Properties](#)
- [Customizing Knowledge Scripts](#)
- [Checking and Changing Job Status](#)
- [Initiating Maintenance on a Computer](#)
- [Grouping and Filtering Job Information](#)
- [Viewing Job Results](#)
- [Closing a Job](#)
- [Deleting a Job](#)

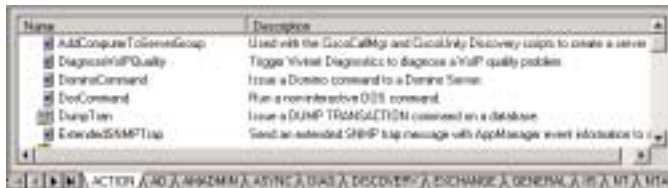
About Knowledge Scripts and Jobs

In AppManager, all management and monitoring tasks are facilitated through the use of Knowledge Scripts, which are programs that run on managed clients (servers and workstations). Knowledge Scripts can collect data, monitor for

specific events, and perform specific actions in response to events. Each script is designed to perform a certain task or set of tasks, and to work on specific types of resources. A job is an instance of a Knowledge Script running on a managed client. That is, when you run a knowledge script you are creating a job.

Knowledge Scripts are organized into categories, which are loosely connected to views in the NetIQ AppManager Operator Console. For example, SQL Knowledge Scripts that check database statistics are available in the SQL and Master views; General Knowledge Scripts can be applied to all Windows-based views for generic management of Windows NT and Windows 2000, and Action Knowledge Scripts let you initiate actions in any view. The Knowledge Script categories you see depend on the AppManager Suite products you have purchased and installed, as well as the objects and resources you have discovered in your environment.

To select a Knowledge Script category to work with, click one of the tabs displayed across the bottom of the Knowledge Script pane.



The individual Knowledge Scripts displayed in any Knowledge Script category can vary depending on the resource objects actually discovered. For example, if printers have not been discovered in the environment, then the PrinterHealth Knowledge Script won't display in the NT Knowledge Script category.

To see a brief description of each Knowledge Script, along with version and build information, right-click in the pane and then click **Show Description**. If text is truncated, pass the mouse cursor over it and the full text appears as a ToolTip.

For detailed information about an individual Knowledge Script, select the script in the Knowledge Script pane and then press the **F1** key; or click **Help** on the **Values** tab in the Knowledge Script Properties dialog box.

Planning which Knowledge Scripts to Run

Each Knowledge Script runs on certain types of objects that correspond to the types of resource objects you see in the TreeView pane. AppManager handles all of the appropriate type matching so you never have to worry about running the wrong type of Knowledge Script on a computer or resource.

In deciding which Knowledge Scripts to run, consider the critical resources in your environment, how servers and applications are distributed, potential system bottlenecks, and the breadth and depth of monitoring you want to do.

In deciding where you should run a Knowledge Script, NetIQ recommends that you only use the **Master** view to discover resources. To monitor resources, use another view such as a standard view or snapshot view.

Running Jobs to Pinpoint Problems

Most Knowledge Script jobs are run on a regular basis to keep continuous or periodic watch over system resources. In some cases, you need to run a combination of scripts to pinpoint a problem.

As an example, assume you run only the Exchange_ServerHealth Knowledge Script on a particular computer. If the Knowledge Script sends you an alert indicating the server is overloaded, you may need to run

additional Knowledge Scripts, using the Run Once schedule, to try to find the cause of the problem.

By running these additional Knowledge Scripts, you get a better sense of the activity on the computer and the potential sources of the problem you are investigating. To carry this example further, if you discover that a server's CPU is busy, you can run the NT_TopCPUProcs Knowledge Script to pinpoint the processes that are consuming the CPU time.

Once you have identified the problem source, there are several things you can do to take appropriate action. For example, in this case you might:

- Spread out the CPU load so that the top CPU processes and users run at different times to reduce system load.
- Change the computer's hardware configuration. In some cases, a computer may not have adequate or appropriate hardware for the tasks it runs. For example, there may be insufficient RAM or disk space. View the basic hardware information by examining the details on hardware resource objects.

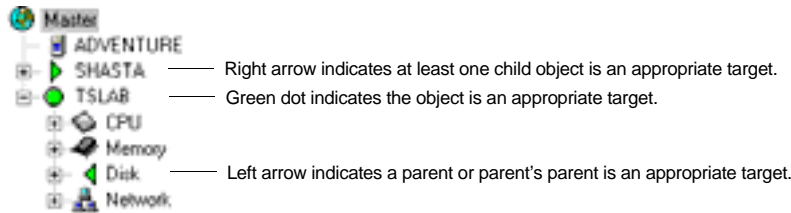
When jobs that check general system health or a number of different resources report event conditions, investigate the event detail message and consider running additional, more specialized, jobs to find the real source of the problem and the best solution.

Starting a New Job

- 1 Click the appropriate Knowledge Script tab in the Knowledge Script pane.
- 2 Select the script that performs the monitoring function you are interested in.
- 3 Drag the icon for the Knowledge Script to the TreeView pane.

To run the job on	Drop the Knowledge Script on
A specific computer	Only that computer.
More than one computer	An appropriate group, the Master level, or several computers, individually.

The icon for the resource object changes to let you know you can drop the script on the selected target.



If you have selected at least one valid target, the Properties dialog box displays.

- 4 The first time you start a Knowledge Script, you may want to review all of the default values before starting the job.

- 5 Enter properties for the Knowledge Script. In many cases, you can simply use the default schedule and values. To make changes:

Click this tab	To
Schedule	Change the default schedule for running the job. See “Setting the Schedule” on page 110 .
Values	Adjust the threshold and other parameters for the job. See “Setting the Values to Monitor” on page 113 .
Actions	Specify an action you want taken. See “Specifying One or More Corrective Actions” on page 115 .
Objects	Modify the resource objects you want to run the job on. See “Selecting the Objects to Monitor” on page 120 .
Advanced	Specify whether to filter temporary event conditions and whether to have duplicate events for this job collapsed into a single event or automatically acknowledged. Additionally, you can specify data collection options for this job. See “Setting Advanced Job Properties” on page 121 .

Any changes you make only affect this instance of the Knowledge Script job you are running. The default properties of the Knowledge Script are not affected.

Tip To permanently change the default properties of a Knowledge Script, see [“Changing Default Knowledge Script Properties” on page 130](#) or [“Copying a Knowledge Script” on page 134](#).


- 6 Once you finish reviewing or changing the parameters, click **OK** to start the jobs.

Tip You can also start a new job by selecting a resource in the TreeView pane, pressing the F2 key, and selecting the Knowledge Script you want from the Knowledge Script Browser dialog box. After you select a Knowledge Script and click **OK**, you can enter properties for the Knowledge Script as discussed in Step 5.

Parent and Child Jobs

When a Knowledge Script job is started, a parent job is created with at least one child. The number of child jobs created is based on the number of computers included in the scope of the drop. For example, if you drop the Knowledge Script on one computer, one child job is created; if you drop it on a group that contains several computers, a child job is created for each one.

In the Jobs tab in the List pane, expand the parent job by clicking the + (plus sign) next to the parent job identifier.



To	Do this
View child jobs for a parent job	Click the + (plus sign) to the left of the parent job.
View all child jobs for all parents	Ctrl-click the + (plus sign) to the left of any parent job.
Fold the child jobs back under the parent	Click the - (minus sign) to the left of the parent job.
Fold all child jobs back under all parents	Ctrl-click the - (minus sign) to the left of any parent job.

Adding a Child Job to an Existing Parent

You can use an existing parent job to monitor additional resources.

To monitor a resource using an existing parent job:

- 1 Select the parent job that you want and then click **List > Add Jobs**.

- 2 In the Add Children Jobs dialog box, select the resources that you want the existing parent job to monitor and click **OK**. A child job is created for each computer.

Note You cannot add a child job to an existing parent job if the parent job is policy-based. A monitoring policy automatically monitors resources as they are discovered. For information on adding a resource to a monitoring policy, see [“Changing a Monitoring Policy” on page 267](#).

Setting the Schedule

Although all Knowledge Scripts have a default schedule, you can modify the schedule for any Knowledge Script job as needed. You can specify when the job should start, the interval (if any) at which the job should run, and when the job should stop. You can also schedule jobs to run on specific days and at set times, or to run only once.

To set the schedule for a job:

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

2 Click the **Schedule** tab.



Note The options for setting intervals, start and stop times, and daily frequency depend on the selected schedule.

3 Select a schedule type.

If you select	Then
Run once	Set a start time for the job to run once (see Step 4).
Regular intervals	<p>Set the interval (Every) — in hours, minutes, or seconds — and start and end times for the job (see Step 4).</p> <p>The interval controls how many times the job runs between the starting and ending time. For example, if you set a 10-minute interval with a start time of 10:00:00 PM and an end time of 10:59:00 PM, the job runs every 10 minutes — for a total of six times. (If the job takes more than 10 minutes to run, then it will run fewer times. For example, if the job takes 15 minutes to run, then it will run at 10:00:00 PM, 10:20:00 PM (the next scheduled interval), and 10:40:00 — for a total of three times.)</p>
X number of times	<p>Set how often you want the job to run (Every) — in hours, minutes, or seconds. Set how many times you want the job to run (End after <u> </u> Times).</p> <p>Then set a start time (see Step 4).</p>

If you select	Then
Daily schedule	Set the daily interval (Every) you want the job to run. Then set start and end dates (see Step 4), and the frequency (see Step 5).
Weekly schedule	Set the weekly interval (Every) and the days of the week you want the job to run. Then set start and end dates (see Step 4), and the frequency (see Step 5).
Monthly schedule	Set the monthly interval (Every) and the days of the month you want the job to run (On). Depending on the option you select in the On list, you can further refine your selection. For example, if you select a day of the week (such as Monday), you can further select the weeks in the month (such as The first Monday and The last Monday of the month). If you choose Selected Day , you can then select days in calendar format (such as 3 and 17, or L for the last day of the month). Then set start and end dates (see Step 4), and the daily frequency (see Step 5). Note For an AppManager job that is configured to run monthly on a particular weekend (such as first weekend, second weekend, etc.), a weekend constitutes a consecutive Saturday and Sunday that falls within the calendar month.
Asynchronous	This option is selected for Knowledge Scripts that run asynchronously to monitor resources as changes occur.

- 4 Set start and end date/times for the job. Available options depend on the schedule type you selected in Step 3.

For	Select
Start (time or date)	Now , or a date and time.
End (time or date)	No end date , or a date and time.

Note Jobs run at the scheduled date and time using the local date and time of the managed client computer. By default, the resulting job, event, and graph data information is displayed in the Operator Console based on the local date and time of the computer where the Operator Console is

installed. To display this information using the local time of the managed client, select **Display dates and times based on the local time of the managed computer** on the Console tab of the Preferences dialog box. For more information, see [“Setting Operator Console Preferences” on page 64](#).

- 5 Set the frequency for the job.

To run the job	Select
Once during the day	Once and then set a start time.
At set intervals during the day	Every . Then set the interval — in seconds, minutes, or hours — and a start and end time.

- 6 Click another tab to change its properties or click **OK** to start the job.

Note To prevent jobs from running during scheduled maintenance periods on a Windows agent, run the AMADMIN_SchedMaint Knowledge Script. On a UNIX agent, use the AMAdminUNIX_SchedMaint Knowledge Script. For more information, select a Knowledge Script in the Knowledge Script pane and then press the F1 key or click **Help** on the **Values** tab in the Properties dialog box.

Setting the Values to Monitor

Knowledge Scripts have default threshold values or other parameters that control what they check and how they behave. The parameters associated with Knowledge Scripts varies, but typically includes:

- What the Knowledge Script does (for example, notifies you of event conditions, collects data for graphs and reports, or attempts to automatically restart services).
- The threshold values to watch for (for example, less than 10% free disk space or CPU utilization greater than 90%).

- The files, messages, and executables to monitor (for example, a list of processes that should be running, specific error messages in a log file, specific applications for which to monitor memory usage).

You can modify the parameter values, as needed.

Note that in the Operator Console, the Values tab of the Job Properties dialog box does not indicate the override value for a parameter, if one is specified. In the Operator Console, the Values column always displays the actual parameter value specified, even if it is not used. To view and change override values for a job, you must use the Control Center Console. If you change the parameter value for a job and that parameter uses an override value, the parameter value is updated but the job runs with its override value.

To modify the parameter values:

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

2 Click the **Values** tab.

List of the parameters you can set for the job.

The default for each parameter...

...Click in the **Value** column to change any parameter value.

Units associated with the parameter (if any).

Provides information about the Knowledge Script.

Description	Value	Units
Event Notification		
Create event if total system CPU is high?	<input checked="" type="checkbox"/> Yes	
Create event if any individual CPU is high?	<input type="checkbox"/> Yes	
Severity - Job failure	5	Severity
Data Collection		
Monitoring		
Threshold - Total system CPU	95	%
Threshold - Individual CPU	98	%
Threshold - Processor queue length	4	Threads

Monitor total CPU usage and queue length to determine whether the CPU is overloaded. This Knowledge Script raises an event if the CPU is overloaded (for example, both the total CPU usage and CPU queue length thresholds are crossed). Click Help for more information.

OK Cancel Help

3 Click in the **Value** column to change threshold or monitoring parameter values. (Some fields have a Browse (...) button you can use to make selections.)

For specific information about the current Knowledge Script parameters, click **Help** in the dialog box.

4 Click another tab to change its properties or click **OK** to start the job.

Specifying One or More Corrective Actions

Almost all Knowledge Scripts can initiate responsive actions, such as sending e-mail, running an external program, or sending an SNMP trap. You define the conditions under which to perform the action and what action to take. For example, if you want to receive an e-mail when CPU utilization reaches 75% on a server, create a job that checks for a CPU threshold of 75% and specify an e-mail action.

Actions appear in the Actions tab of the Knowledge Script pane. Actions that can be run on a managed UNIX computer are prefixed with “UX”, for example Action_UXCommand. All other actions can be run on a managed Windows computer (including the management server or a proxy server). For discussion purposes, actions that can be run on a managed Windows computer are referred to as **Windows actions** and actions that can be run on a managed UNIX computer are referred to as **UNIX actions**.

The Action Knowledge Scripts, such as Action_MAPIMail and Action_Messenger, have built-in severity filtering. Using the Severity Configuration parameters, you can now specify the range of severities you want the action to execute.

Understanding How Actions Work on Windows Computers

When monitoring a Windows computer, you can configure a job to initiate a different corrective action depending on when the event condition is detected. To do this, configure more than one Windows action and specify a different action schedule. Building upon the example in the previous section, when CPU utilization exceeds the threshold during business hours, an e-mail action runs; during non-business hours, a page action runs.

For information on viewing and configuring the scheduled hours of operation for an action schedule, see [“Set Preferences for Action Schedules” on page 85](#).

In addition, you can configure an action to run each time an event is raised, only when a duplicate event is created a specified number of times, or when the event condition no longer exists.

To configure a job to initiate a similar response depending on when the event condition is detected, you must make a copy

of the action Knowledge Script and configure the job to run each action with the corresponding action schedule. For example, if CPU utilization exceeds the threshold during the week, an e-mail action (Action_MapiMail) sends an e-mail notification to the IT operator; if the threshold is exceeded on the weekend, an e-mail action (Action_MapiMailOnCallIT) sends an e-mail notification to the on-call IT operator.

Note When using a custom action schedule between **Custom11** and **Custom15** on a job that runs on AppManager agents that are **earlier than version 5.0**, set the location where the action is run to the management server (**MS**). If you configure an action to use one of these custom action schedules and you configure the action to run on the managed client (**MC**), the job does not run and has a status of Error; if the action location is set to **Proxy**, the job runs but the action does not run. This problem does not occur on Version 5.0 (or later) AppManager agents.

Understanding How Actions Work on UNIX Computers

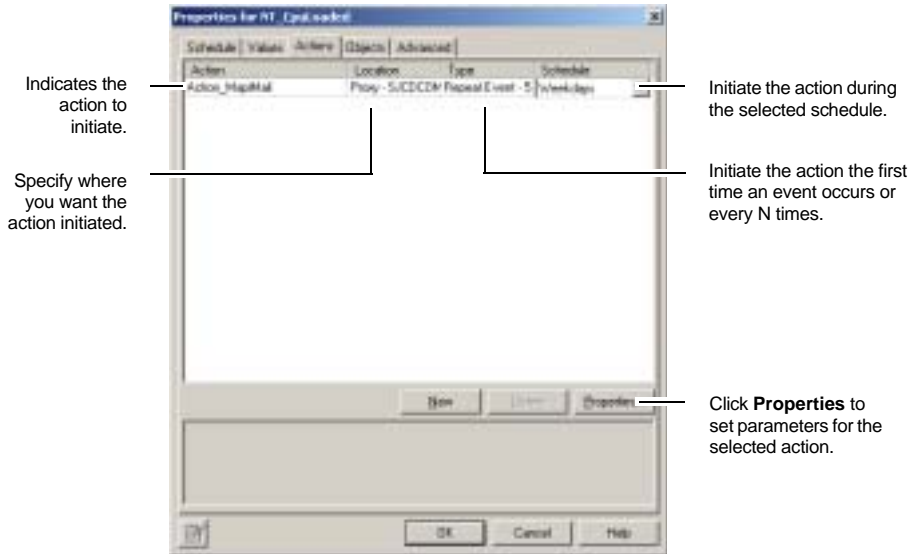
When monitoring a UNIX computer, you can configure a Windows action to run on the management server or a UNIX action to run on the managed UNIX computer. However, a UNIX action must be configured to run each time an event is raised. In addition, you cannot use action schedules to initiate a different corrective action depending on when the event condition is detected.

To set up a responsive or corrective action for a job:

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

2 Click the **Actions** tab.



3 Click **New** to create a new action and then select an Action Knowledge Script from the **Action** list. If you are monitoring a Windows computer, UNIX actions are not available.

Note To see a list of available actions, click the **Actions** tab in the Knowledge Script pane.

4 In the **Location** column, specify where the action runs:

Select	To run the action on the
MS	Management server computer Select this option to initiate a Windows action on the management server. If you are monitoring a UNIX computer, this option is required to initiate a Windows action.
MC	Managed client computer Select this option to initiate a Windows action on the managed Windows computer. If you are monitoring a UNIX computer, you must select this option to initiate a UNIX action on the managed UNIX computer.
Proxy	Proxy server Select this option to initiate a Windows action on a Windows computer with the AppManager agent (neti qmc.exe). This option is not applicable when monitoring UNIX computers.

- 5 In the **Type** column, configure an action to run the first time an event is raised (a unique event), after a duplicate child event is created a specified number of times, or when the event condition no longer exists. Note that a Unix action must be configured to run the first time an event is raised:

Select	To run an action
New event	The first time an event is raised (a unique event). This is the default. Note A Unix action must be configured to run the first time an event is raised.
Repeat Event -n	When the specified number of duplicate child events are created. The default value, 1 , permits an action to be run each time a new duplicate child event is created. Specify the number of times that a new duplicate child event must created in order to initiate an action. Note that if event collapsing is enabled, increments to the event count of a child event are not the same as a new duplicate child event. For example, if you have enabled event collapsing with the default interval of 20 minutes, when an event condition is first detected, a child event is created in the Events tab of the List pane. Subsequent events increment the event count until the 20 minute interval has passed. At this point, a new child event can be created.
Event Down	When the event condition no longer exists. To use this option, you must check the box on the Advanced tab labeled Generate a new event when original event condition no longer exists .

- 6 In the **Schedule** column, select an action schedule to specify the available hours during which the action can run. When monitoring UNIX computers, action schedules are not applicable.

For information on viewing and configuring the scheduled hours of operation for an action schedule, see [“Set Preferences for Action Schedules” on page 85](#).

- 7 To set the properties for an Action Knowledge Script, select a Knowledge Script in the Action list and then click **Properties**.

Most actions require you to set some additional properties. For example, if you select an e-mail action, you need to specify an e-mail recipient. For more information about Action Knowledge Scripts and their parameters, see the online help.

- 8 Click **OK** to close the Properties dialog box for the action Knowledge Script.
- 9 Click **New** to configure another action or click **OK** to start the job.

Selecting the Objects to Monitor

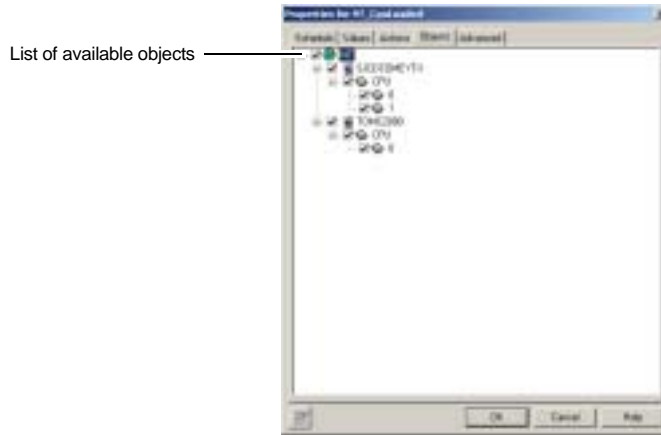
If you drop a Knowledge Script on a TreeView object that contains other objects (for example, a group that has multiple computers or a Disk folder that contains multiple disks on a single computer), the script runs by default on all the objects in the hierarchy that match the Knowledge Script type.

You can select the objects on which you want the job to run.

To select a subset of objects on which to run the Knowledge Script:

- 1 When you drag the icon for the Knowledge Script to the TreeView pane, the Properties dialog box appears.

- 2 Click the **Objects** tab.



- 3 In the list of objects, de-select the objects you do not want the Knowledge Script to monitor. By default, all objects are selected.
- 4 Click another tab to edit its properties or click **OK** to start the job.

Setting Advanced Job Properties

You can specify advanced job properties to filter temporary event conditions. If an event is raised, you can configure the advanced job properties to collapse duplicate events or automatically close the event when the event condition no longer exists. Note that on UNIX computers, you cannot configure a job to automatically close the event.

There are additional options that you can use to configure data collection. See the following sections for more information.

Filtering Temporary Event Conditions

Specify criteria to filter temporary events by configuring the number of times that a duplicate event condition must be

detected during a specified number of job iterations before an event is raised. An event condition is considered a **duplicate** when an event condition is detected on the same object name with the same event message, severity, and Job ID as a previous event condition.

For example, when the threshold for the number of duplicate event conditions is set to 3 and the number of job iterations is set to 5, if an event condition is detected each time the job runs, the first two event conditions are ignored and the third event condition raises an event and starts the [“Time interval for event collapsing” on page 71](#).

The next two event conditions that are detected during the interval are ignored—only two event conditions are detected during the interval, which falls below the threshold. The event condition count resets when the interval resets, after the fifth job iteration. In order to raise another event, a duplicate event condition must occur three times during the next five job iterations, and if event collapsing is enabled, the event must be generated within the [“Time interval for event collapsing” on page 71](#).

Ignored event conditions are not logged in the repository and do not increment the event count of the original child event. Actions associated with ignored event conditions are also suppressed.

To specify the number of times that an event condition must be detected during a specified number of job iterations to raise an event:

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

2 Click the **Advanced** tab.



3 Enter criteria to specify a threshold for raising an event:

For	Do this
Raise event if event condition occurs N times...	Enter a threshold for the number of times that a duplicate event must be detected during the specified job iterations before an event is raised. (The detected event conditions do not need to be consecutive.)
...within M iterations	Enter the number of job iterations to specify the interval. The interval is reset after the given number of job iterations. The default for the number of duplicate events and job iterations are both 1, which means that duplicate events are not ignored and an event is raised each time an event condition is detected.

4 Click **OK** to start the job.

Changing Event Collapsing Options for One Job

Sometimes a Knowledge Script job detects identical events. To suppress these duplicates and their specified actions, you can

set event collapsing options. You can set these options globally, for all jobs, or for individual jobs. This section discusses how to set event collapsing options for each job you run.

Note For more information about duplicate events and global event collapsing options, see [“Duplicate Events” on page 174](#).

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

- 2 Click the **Advanced** tab.



- 3 Filter temporary event conditions by specifying the number of event conditions that must be detected during a specified number of job iterations. If an event condition is detected the specified number of times during the interval, an event is raised. For example, to configure a job to raise an event if the event condition is detected twice during an interval of

five job iterations, set **Raise an event if event condition occurs** to **2 times** within **5 job iterations**.

See [“Filtering Temporary Event Conditions” on page 121](#) for more information.

4 Select Collapse duplicate events into a single event.

Duplicate events are logged in the repository; however, actions associated with collapsed duplicate events, whether the actions are initiated on the managed client, management server, or proxy computer, are suppressed.

5 Set the Time interval for event collapsing.

The default time interval is 20 minutes.

Any duplicate events that were created within the time interval are collapsed into the original child event. The time interval can be measured from when the first event condition is detected or when the most recent event condition was detected.

6 Click OK to start the job.

Note Event collapsing options set on the Advanced tab in the Properties dialog box affect only the job you are about to run. To set global default event collapsing options for all Knowledge Scripts, set the options on the Preferences dialog box (see [“Duplicate Events” on page 174](#) and [“Setting Global Event Collapsing Options” on page 179](#)).

Automatically Closing an Event when the Event Condition No Longer Exists

When you are monitoring a resource on a Windows computer that tends to raise events frequently, you can reduce the amount of time you spend managing these events by configuring the job to automatically close an event whenever the event condition no longer exists. This advanced option is

not available when configuring a job to run on a UNIX computer.

The first time an event condition is encountered, a parent event and a child event are created; no additional events are raised as long as the event condition persists, and the event count for the child event does not increment. When the event condition no longer exists, a new event is created to indicate the change in the event condition and optionally, the original event is closed. This option can be set globally, for all jobs, or a particular job. This section discusses how to automatically close events for each job you run.

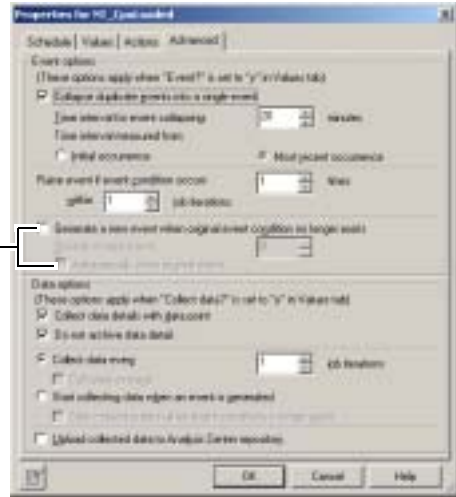
Note For more information about automatically closing events and global options, see [“Automatically Closing Events When the Condition No Longer Exists” on page 180](#).

- 1 Display the Properties dialog box for the Windows-based job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

- 2 Click the **Advanced** tab.

Select both options to automatically raise an event when the event condition no longer exists and automatically close the original event.



- 3 Select **Generate a new event when original event condition no longer exists**. With this option selected, an event is raised when the event condition that initiated the original event no longer exists. Note that this option is not available when configuring a job to run on a UNIX computer.

With this option selected, you can specify the event severity when an event is raised in response to the change in the event condition. By default, the event severity is set to 20.

- 4 To configure the Operator Console to close the original event when the event condition no longer exists, select **Automatically close original event**. Note that this option is not available when configuring a job to run on a UNIX computer.
- 5 Click **OK** to start the job.

Changing Data Collection Options for One Job

There are several data collection options that you can use to optimize how a Knowledge Script job collects data for charts and reports.

Data collection options can be set globally, for all jobs, or a particular job. This section discusses how to set data collection options for each job you run. For information on configuring global data collection options, see [page 75](#).

- 1 Display the Properties dialog box for the Knowledge Script job.

If the job is already running, double-click the job identifier in the Jobs tab; or select the job, right-click, and then click **Properties**.

- 2 Click the **Advanced** tab.

Configure data collection options to reduce the amount of data in the repository.



3 Select an option to specify whether to collect data details:

Select	To
Collect data details with data point	This option only applies to data collection for AppManager graphs and charts. Select this option to collect details for each data point (such as server name and collection time). These details are displayed when you want to view the details of a data point in a graph or chart. This option is selected by default.
Do not archive data detail	This option only applies to data collection for AppManager reports that display detail data, such as ReportAM_DetailData. Select this option to only collect the value of the monitored resource; detailed information, such as server name and collection time, is not collected. This option is selected by default.

4 To control how frequently a data point is collected, update **Collect data every N job iterations** to specify the number of times that a job must run to collect data.

For example, if a Knowledge Script job is set to run at regular intervals of 10 minutes and you set this value to 3 iterations, then a data point will be created every third time the job runs, or every 30 minutes.

The default is 1 iteration, which means a data point is created every time the job runs.

5 If you have configured the job to create a data point once during a specified interval, select **Calculate average** to create that data point as an average value of all data points measured during the interval. For example, if you configure the repository to collect a data point every 3rd time a job runs, and the data point for the first iteration is 2, the data point for the second iteration is 5, and the data point for the third iteration is 10, the data point is created on the third iteration and its value is 5.6666.

- 6 Select **Start collecting data when an event is generated** to begin collecting data when an event is raised. For example, if a Knowledge Script job is set to raise an event when CPU usage exceeds 70%, then the job will also start collecting data at that point.

Note The data header is always created when you start a job. If you select this option, the Graph Data tab displays the data stream with the current number of data points, which may be zero (0) until an event occurs.

- 7 If you have configured the job to begin collecting data when an event is raised, select **Stop collecting data when the event condition no longer exists** to stop collecting data when the event condition no longer exists. Following the previous example, the job will stop collecting data when CPU usage falls below 70%.

- 8 When you finish, click **OK**.

Customizing Knowledge Scripts

You can change the default settings of a Knowledge Script (see the next section) or create a new Knowledge Script with the settings you want (see [“Copying a Knowledge Script” on page 134](#)).

Changes you make to schedule, monitoring values, actions, and advanced options can be propagated to running jobs. See the following sections for more information.

Changing Default Knowledge Script Properties

As you gain experience in monitoring your environment, you may find that you want to change the default properties for selected Knowledge Scripts rather than changing the parameter values for individual jobs.

Before changing the default threshold values or schedule for a Knowledge Script, you might want to:

- Run any Knowledge Scripts that collect data related to the threshold you want to change to determine a typical value for your environment.
- Run the Knowledge Script you want to change with the default values to evaluate whether you are seeing too many or too few events.
- Run one or more independent jobs with the values you want to use to test that the results are what you expect before changing the defaults.

When you have determined the operational values that best suit your needs:

- 1 Click the appropriate Knowledge Script tab in the Knowledge Script pane.
- 2 Double-click the Knowledge Script you want to change.
- 3 In the Properties dialog box, click a tab and make your changes.
- 4 Click **OK**.

The changes are written to the AppManager repository, replacing the old properties, and affect all Operator Consoles and Operator Web Consoles using the same repository.

Note The version number of the Knowledge Script does not change and the history is lost. To change the version number and establish a version history, check the Knowledge Script out of the AppManager repository (**KS > Check Out Knowledge Script**) and then check it back in (**KS > Check In Knowledge Script**). For information about version history, see [“Getting Information about Knowledge Scripts” on page 92](#).

After changing the default properties for a Knowledge Script, you can propagate the Knowledge Script properties to:

- Existing jobs (more on [page 132](#))
- Knowledge Script Group members (more on [page 133](#))

Propagating Knowledge Script properties to Running Jobs or Reports

You can propagate the properties of a Knowledge Script, including changes to schedule, monitoring values, actions, and advanced options, to running jobs or reports started by that Knowledge Script. Corresponding jobs are stopped and restarted with the Knowledge Script properties.

Before you propagate the Knowledge Script properties of a report, ensure that you have specified a value for all of the required parameters, including parameter values that are not displayed in the Values tab of the Knowledge Script Properties dialog box.

For information on propagating properties to monitoring jobs started by a Knowledge Script Group, see [“Propagating Knowledge Script Group Members to Running Jobs or Reports”](#) on page 261.

- 1 In the Knowledge Script pane, click the Knowledge Script you want and then click **Properties Propagation > Ad Hoc Jobs**.

- 2 In the Properties Propagation dialog box, select the components of the Knowledge Script that you want to propagate to associated ad hoc jobs:



Select	To propagate
Script	The logic of the Knowledge Script.
Properties	The Knowledge Script properties, including schedule, monitoring values, actions, and advanced options.

- 3 When you finish, click **OK**.

Propagating Properties of a Knowledge Script or a Report Script to Knowledge Script Group Members

You can propagate the properties of a Knowledge Script or a report script, including changes to schedule, monitoring values, actions, and advanced options, to corresponding Knowledge Script Group members.

Before you propagate the Knowledge Script properties of a report, ensure that you have specified a value for all of the required parameters, including parameter values that are not displayed in the Values tab of the Knowledge Script Properties dialog box.

After you propagate Knowledge Script properties to Knowledge Script Group members, you can propagate the updated Knowledge Script Group members to monitoring jobs or reports started by the Knowledge Script Group. See [“Propagating Knowledge Script Group Members to Running Jobs or Reports” on page 261](#) for more information.

Note For information on automatically propagating changes to

policy-based jobs, see [“Changing a Knowledge Script Group” on page 259](#).

- 1 In the Knowledge Script pane, click the Knowledge Script you want and then click **Properties Propagation > Knowledge Script Group Members**.
- 2 In the Properties Propagation dialog box, select the components of the Knowledge Script that you want to propagate to associated monitoring jobs:



Select	To propagate
Script	The logic of the Knowledge Script.
Properties	The Knowledge Script properties, including schedule, monitoring values, actions, and advanced options.

- 3 When you finish, click **OK**.

Copying a Knowledge Script

You can copy a Knowledge Script and then change its properties to create a new, customized Knowledge Script. The new Knowledge Script name cannot exceed 150 characters.

We recommend limiting the length of a Knowledge Script name to 145 characters or less. Limiting the length to 145 characters prevents AppManager from truncating the Knowledge Script job name when running the Knowledge Script from a monitoring policy, or as an ad hoc job started by a Knowledge Script Group. See [“How Knowledge Script Groups Work” on page 252](#) for more information.

- 1 Click the appropriate Knowledge Script tab in the Knowledge Script pane.
- 2 Select the Knowledge Script that you want to copy.
- 3 Right-click and then click **Copy Knowledge Script**.



- 4 Enter a new name and description for the Knowledge Script. (By default, the phrase “CopyOf” has been appended to the Knowledge Script name to ensure that a new script is created.)

To	Do this
Display the new Knowledge Script in the same tab as the original script	Don't change the prefix that appears before the underscore.
Create a new tab for your customized Knowledge Scripts	Type a new name before the underscore. For example, MySi te_

- 5 Click **OK**.

The new Knowledge Script is automatically checked into the AppManager repository.

- 6 In the Knowledge Script pane, double-click the icon for the new Knowledge Script you just created.
- 7 In the Properties dialog box, make the changes you want in the **Schedule**, **Values**, **Actions**, or **Advanced** tabs.
- 8 Click **OK** to save the changes to the AppManager repository.

Note The version number of the Knowledge Script does not change and the history is lost. To change the version number

and establish a version history, check the Knowledge Script out of the AppManager repository (**KS > Check Out Knowledge Script**) and then check it back in (**KS > Check In Knowledge Script**). For information about version history, see [“Getting Information about Knowledge Scripts” on page 92](#).

Checking a Knowledge Script into the Repository

To run a new or modified Knowledge Script, you must first check the Knowledge Script into the repository.

We recommend limiting the length of a Knowledge Script name to 145 characters or less. Limiting the length to 145 characters prevents AppManager from truncating the Knowledge Script job name when running the Knowledge Script from a monitoring policy, or as an ad hoc job started by a Knowledge Script Group. See [“How Knowledge Script Groups Work” on page 252](#) for more information.






- 1 Click **KS > Check In Knowledge Script**.
- 2 Select the Knowledge Script you want, and click **Open**.

The new Knowledge Script is automatically checked into the AppManager repository.

To check a Knowledge Script out of the repository and copy it to your local computer, select the Knowledge Script you want from the Operator Console and click **KS > Check Out Knowledge Script**.

Checking and Changing Job Status


Once a job is started, it typically goes through the following states in its life cycle:

This status	Indicates
Pending	<p>A temporary state when a job is first started or when you temporarily stop and restart a job. If a job stays in this state for very long, it may indicate a network communication problem or that a managed computer has been shut down.</p> <p>Click the + (plus sign) next to a parent job to see additional status details on child jobs:</p> <p> Start - The job is in the process of starting.</p> <p> Stop - The job is in the process of stopping.</p>
Running	<p>The job has been started, but may or may not be actively executing its monitoring functions.</p> <p>Click the + (plus sign) next to a parent job to see additional status details on child jobs:</p> <p> Active - The job is currently running on the managed client, monitoring computer resources according to the schedule and parameters you have set. Any corrective actions you have defined for the job occur automatically when the event condition is met.</p> <p> Scheduled - The job has not yet started executing its monitoring functions. Monitoring takes place according to the schedule you have set.</p> <p> Inactive - The job is currently inactive due to scheduled maintenance. When the maintenance period is over, the job will resume running according to the schedule you have set. (To prevent jobs from running on a Windows agent during scheduled maintenance periods, run the AMADMIN_SchedMaint Knowledge Script. On a UNIX agent, use the AMAdminUNIX_SchedMaint Knowledge Script. For more information, select a Knowledge Script in the Knowledge Script pane and then press F1; or click Help on the Values tab in the Properties dialog box.)</p>
Stopped	<p>The job has run to completion or has been stopped manually.</p>

This status	Indicates
Closed	The job is no longer in use but is still available for historical purposes and can be re-opened in the future.
Error	A job has run into a problem (for example, an error in the managed client's runtime configuration that prevented the job from running properly).


Viewing Jobs in the TreeView Pane

The job (J) indicator next to a resource object in the TreeView pane means that one or more Knowledge Script jobs are running. Job indicators appear at each level of the hierarchy to give visibility to running jobs, regardless of how the hierarchy is expanded or collapsed.



When all the jobs on a particular branch are stopped, the job indicators disappear.

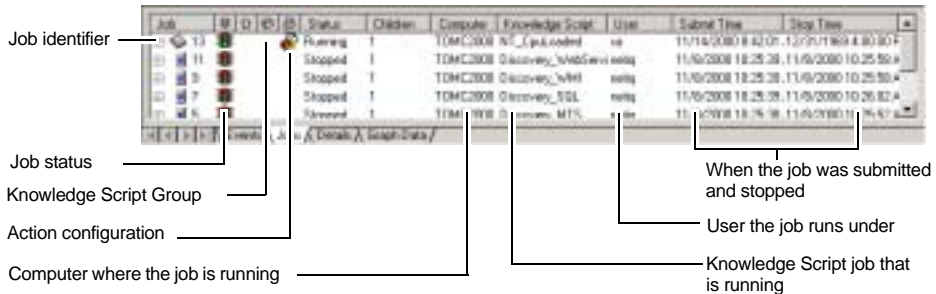
When an event condition on a resource object is detected, an alert indicator flashes and the other job indicators in the hierarchy turn into event indicators.



The event alert indicator continues to display in the TreeView until you respond by acknowledging or closing the event. See [Chapter 5, “Responding to Events”](#) for more information.

Viewing Jobs in the List Pane

Click the **Jobs** tab in the List pane to get more details on all jobs, including the current status of each job, whether the job was started by a Knowledge Script Group or was started by a monitoring policy, where the job is running, the name of the Knowledge Script being run, and the time the job was submitted.



Click the + (plus sign) next to a parent job to get more detailed status information on child jobs. For example, a Running job can be Active, Scheduled, or Inactive; a Pending job can be Starting or Stopping.


To expand all parent jobs and display child information, Ctrl-click the + (plus sign) next to any parent job.

For information on viewing policy-based jobs, see [“Viewing Policy-Based Jobs in the List Pane” on page 254](#).

Note If text is truncated, pass the mouse cursor over it and the full text appears as a ToolTip.

Viewing the Actual Script for a Running Job

To view the actual script that is running on a computer:

- 1 In the Jobs tab of the List pane, double-click a child job to view its job Properties dialog box.
- 2 Click the **View KS** button .

Hiding Parent Jobs

To improve sorting of all child jobs in the List pane, hide parent jobs. To improve Operator Console performance, it is recommended that you do not use this feature in the Master view.

To hide parent jobs in the current view, in the Jobs tab of the List pane, click **List > View** and then unselect **Parent Jobs** (or select to show parent jobs).

For information on hiding parent jobs in the Operator Console, see [“Setting Operator Console Preferences” on page 64](#).

Changing the Layout of the Jobs Tab

You can change the columns that appear in the Jobs tab and the order in which they appear. The changes are applied to the Jobs tab in the current view.

To change the layout of the Jobs tab:

- 1 In the List pane, click the Jobs tab and click **List > Customize**.

- 2 In the Show Job Columns dialog box, perform the following tasks:

To	Do this
Hide a column	Select a column name from the Show these columns list and click Remove .
Show a column	Select a column name from the Available columns list and click Add .
Change the order in which the columns are displayed	In the Show these columns list, from top to bottom, the columns are displayed in the Jobs tab from left to right. Select a column name and click: <ul style="list-style-type: none">• Move Up to move the column to the left.• Move Down to move the column to the right.

- 3 When you finish, click **OK**.

Note For information on changing the default layout of the Operator Console, see [“Setting Operator Console Preferences” on page 64](#).

Viewing Jobs Based On Status

You can select to show or hide jobs in the Jobs tab based on status. Hiding jobs removes them from the list but doesn't delete the job information. For example, you might want to hide closed jobs so they don't clutter up the list of jobs you're currently interested in. Later, if you want to view the closed jobs for historical purposes or re-open a job, you can select to show closed jobs.

- 1 Click the **Jobs** tab in the List pane.
- 2 Right-click in the List pane and then click **View**.
- 3 Select (to show) or deselect (to hide) **Running Jobs**, **Pending Jobs**, **Stopped Jobs**, **Closed Jobs** and **Error Jobs**.

Changing the Sorting Order

You can click any of the headings in the List pane to sort all parent entries in that field or toggle the sorting order preference. For example, jobs in the **Jobs** tab are normally sorted by job ID in ascending order. To sort the job list by Knowledge Script, click the **Knowledge Script** header. To reverse the sorting order from ascending to descending, click the **Knowledge Script** header again.

Changing Job Status

You can change the status of a particular job or all jobs running on an object in the TreeView pane.

For a particular job, you can click the job traffic light icon in the Jobs tab in the List pane. Clicking a running job stops it; clicking a stopped job, restarts it. You can attempt to restart a job with an Error status, but if the job still encounters an error, the status will return to Error.

Stopping a parent job stops all of its child jobs. To stop an individual child job, first click the + (plus sign) next to a parent job.

For a particular object in the TreeView, you can right-click to stop and start all jobs on that object.

You cannot change the status of a Running policy-based job. If a policy-based job has a status of Stopped, you can manually restart the job. For information on changing a monitoring policy, see [Chapter 9, “Monitoring by Policy.”](#)

To temporarily stop a particular job and restart it:

- 1 Click the **Jobs** tab in the List pane.
- 2 Click the green traffic light icon to the left of the Running status for the job you want to stop.

The status (and icon) changes to Pending (yellow) and then to Stopped (red).

- 3 When you are ready to restart the job, click the red traffic light icon to the left of the Stopped status.

The status (and icon) changes to Pending (yellow) and then to Running (green).

To temporarily stop and restart all jobs on a particular object in the TreeView pane:

- 1 Select an object in the TreeView pane.

Right-click	To stop all jobs on
An object	All resources under the object.
A computer	The computer.
A server group	All computers in the group.
The top-level icon in a view	All computers in the view.

- 2 Right-click the selected object and click **Stop Jobs**.

In the Jobs tab of the List pane, the status of the jobs (and icons) change to Pending (yellow) and then to Stopped (red).

- 3 When you are ready to restart the jobs, right-click the selected object and click **Start Jobs**.

In the Jobs tab of the List pane, the status of the jobs (and icons) change to Pending (yellow) and then to Running (green).

Changing Properties of a Running Job

When you change the properties of a parent job, all job properties of the parent are applied to its child jobs. To change the properties of individual child jobs, first click the + (plus sign) next to a parent job.

In most cases, you will want to keep the properties of the parent and all child jobs the same. Any changes to the properties of a child job will be overwritten by the parent job when you:

- Update the parent job properties
- Propagate job properties to running jobs

Note For information on changing the properties of all running jobs by propagating changes to the job properties, see [“Propagating Knowledge Script properties to Running Jobs or Reports” on page 132](#).

- 1 Click the **Jobs** tab in the List pane.
- 2 In the Job or Status columns, double-click an ad hoc job to open the Properties dialog box.

Note Clicking or double-clicking the traffic light icon changes the status of a Running job to Stopped; it doesn't open the properties dialog box.

- 3 Change the properties you want in any tab of the dialog box (for example, modify the schedule, set a new threshold value, or modify the list of objects that the job monitors).

If the dialog box is grayed out, the job is policy-based. See [“Changing a Monitoring Policy” on page 267](#) for information on updating policy-based jobs.

- 4 Click **OK**.

The job is temporarily changed to a pending state and then resumes running with the new properties.

Initiating Maintenance on a Computer

In many environments, you may need to perform unscheduled maintenance on a computer. For example, an organization may have an Apache Web server that must be shut down immediately. In this case, you can temporarily block all jobs, events, and data for a particular computer, including jobs that remotely monitor a computer, by placing the computer in **machine maintenance mode**.


Maintenance options	How it works
Machine maintenance	<ul style="list-style-type: none">• Administrator manually enables and disables machine maintenance on a Windows or UNIX agent.• Machine maintenance blocks all monitoring jobs for a computer, including jobs that remotely monitor the computer.• Machine maintenance does not require the agent to be running to enable or disable maintenance mode.• Machine maintenance does not block AMAdmin jobs, for example, the AMAdmin_DBHealth Knowledge Script.
Scheduled maintenance	<ul style="list-style-type: none">• Schedule a maintenance period on a Windows agent using the AMAdmin_SchedMaint Knowledge Script. On a UNIX agent, use the AMAdminUNIX_SchedMaint Knowledge Script.• Requires the agent to be running to configure, start, and stop maintenance for the specified period.• Blocks a particular Knowledge Script category or all Knowledge Scripts monitoring the computer, including jobs that remotely monitor the computer.

If you intend to shut down a computer that is managed by AppManager, it is always a good idea to enable machine maintenance before you shut the computer down. In some cases, as the computer is shutting down, a monitoring job may error out because the resource that the job monitors is not available.

You must manually enable and disable machine maintenance on the computer. The only exception is when you enable maintenance on a computer and replace it with a clone.

Because the clone does not have any information about its machine maintenance status, when the computer is brought online and communicates with the management server, after about 5 minutes, the management server will automatically disable machine maintenance on the computer. Alternatively, you can manually remove the machine maintenance from the clone computer.

To enable machine maintenance on a computer, in the TreeView, click to select a computer and then click **TreeView > Maintenance Mode**.

This icon , indicates that a computer is in machine maintenance mode or all application resources on a computer are in scheduled maintenance mode; AppManager has temporarily stopped monitoring the computer.



This icon appears in front of all objects on the computer when the computer is in machine maintenance mode **or** when all application resources on the computer are in scheduled maintenance mode.



This icon appears when a particular application resource is in **scheduled maintenance mode**. Note that all objects have the icon but only jobs for the specified application resource are blocked.

To disable machine maintenance and resume all monitoring jobs, in the TreeView, click to select a computer and then unselect **TreeView > Maintenance Mode**.

Grouping and Filtering Job Information





By default, AppManager displays all of the jobs in your environment in list form, with each job as a separate entry. Although this provides complete information, as you add more

jobs you may want to organize your jobs into groups, or filter the information displayed to make jobs easier to work with. For example, you may want to organize jobs by where they are running or when they were submitted, or limit the jobs displayed to those running on a specific computer or associated with a specific Knowledge Script.

Organizing jobs into Groups

- 1 Click the **Jobs** tab in the List pane.
- 2 Right-click in the List pane and then click **Group By**.
- 3 Select how you want jobs grouped: **Computer**, **Knowledge Script**, **User Name**, **Submit Time**, or **Knowledge Script Category**.

Once you make your selection, job information is folded into groups as specified. For example, if you selected Computer, the Jobs tab might display something similar to the following two groups:

  SHASTA
  TSLAB

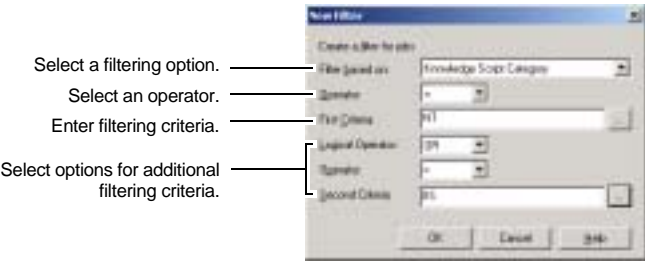
To expand the detailed job information for any group, click the + (plus sign) to the left of the group.

Filtering the Jobs Displayed

- 1 Click the **Jobs** tab in the List pane.
- 2 Right-click in the List pane. Click **Filter**; then click **New Filter** or a previously defined filter.

If you select an existing filter, the information in the Jobs tab in the List pane is immediately filtered based on the filter properties.

If you are creating a new filter, the New Filter dialog box displays.



- 3 In the **Filter based on** list, select a filtering option: **Computer**, **Knowledge Script**, **User**, **Submit Time**, or **Knowledge Script Category**.
- 4 For **Operator**:

If you select	Then
Submit Time	Select an operator (for example, greater than > or less than <).
Computer Knowledge Script User	Leave the default operator (=).

- 5 For **First Criteria**, enter the filtering criteria you want to use or click the Browse (...) button to see a list of possible values.

If you select **Submit Time**, use the format set in the Regional Settings Properties in the Windows Control Panel (the default format for English (United States) is mm/dd/yyyy h: mm: ss).

- 6 To add additional filtering criteria, select **AND** or **OR** from the **Logical Operator** list. Then select an operator and the second filtering criteria (click the Browse (...) button to see a list of possible values).

- 7 Click **OK**.

The Jobs tab is filtered according to the criteria you have specified; the new filter is added to the List menu.

Note The possible values for both the first and second criteria are based on the filter option you select. For example, if you select Computer, the possible values for both criteria are Windows computer names. If you want to create a filter that combines filtering options – for example, a filter that displays the jobs on the computer Shasta (Computer) and submitted after a certain date (Submit Time) – use a combination of the **TreeView > Filter List Pane by Selected Object** option and a **List > Filter** option.

Viewing Job Results

View Results in the TreeView Pane

To quickly see where jobs are running in your environment and whether an event has been generated, view jobs in the TreeView pane. For more information about event indicators, see [“Event Severity Levels” on page 155](#).

View Results in the List Pane

Use the tabs in the List pane to view results.

Click this tab	For information about the
Events	<p>Events generated by a Knowledge Script job.</p> <p>Double-click an event identifier to get more information about the event, to add comments, or view any detailed messages provided by the Knowledge Script job.</p> <p>To view the parent job that created an event, click the hyperlinked Job ID.</p> <p>Note To hyperlink to the child job, the parent must be expanded before you click the link.</p>
Jobs	<p>Status of jobs.</p> <p>To view the parent event created by a job, click to select the child job and then press F4.</p>
Graph Data	<p>Data streams collected by a Knowledge Script job.</p> <p>Double-click a data stream identifier to get more information about the data stream properties or drag the data stream to the Graph pane to see the data in graphical form.</p> <p>To view the associated job for a data stream, find the data stream you want and click the hyperlinked Job ID.</p> <p>Note To hyperlink to the child job, the parent must be expanded before you click the link.</p>

View Results in the Graph Pane

You can view job results in the Graph pane if you:

- Set the Knowledge Script job to collect data for graphing, and
- Create a graph from a data stream listed in the List pane

Once you have created a graph, you can view details about any point in the graph by double-clicking that point on the graph. (If there is more than one graph in the Graph pane, click the tab of the graph you want to see.)

For more information about graphs, see [Chapter 6, “Creating Graphs.”](#)

View Results in the Chart Console

You can view job results in the Chart Console if you:

- Set the Knowledge Script job to collect data for graphing, and
- Create a chart from a data stream listed in the List pane.

For more information about charts, see the AppManager Help.

Closing a Job

- 1 Click the **Jobs** tab in the List pane.
- 2 Select the job you want to close.

If the job is hidden, see [“Viewing Jobs Based On Status” on page 141.](#)

To select multiple jobs, press the Ctrl or Shift keys while making your selections.

- 3 Click the open folder to the left of the job status to close the job.

The job continues to be displayed in the Jobs tab, allowing you to access it for historical purposes. You can also re-open the job to use the same parameters at a later time (there's no need to re-enter any information). If you restart the job, it has the same job ID number it originally had.

Deleting a Job

You can delete pending, stopped, or closed jobs. When you delete a job, the job and its events and graph data are deleted from the AppManager repository.

Note For information on how to configure the repository to maintain events and graph data, see [“Setting Repository Preferences” on page 69](#).

- 1 Click the **Jobs** tab in the List pane.
- 2 Select the job you want to delete.

If the job is hidden, right-click in the List pane and then click **View**. Select (to show) the job status associated with the hidden job.

To select multiple jobs, press the Ctrl or Shift keys as you make your selections.

- 3 Press the **Delete** key; or right-click and then select **Delete Job**.

Responding to Events

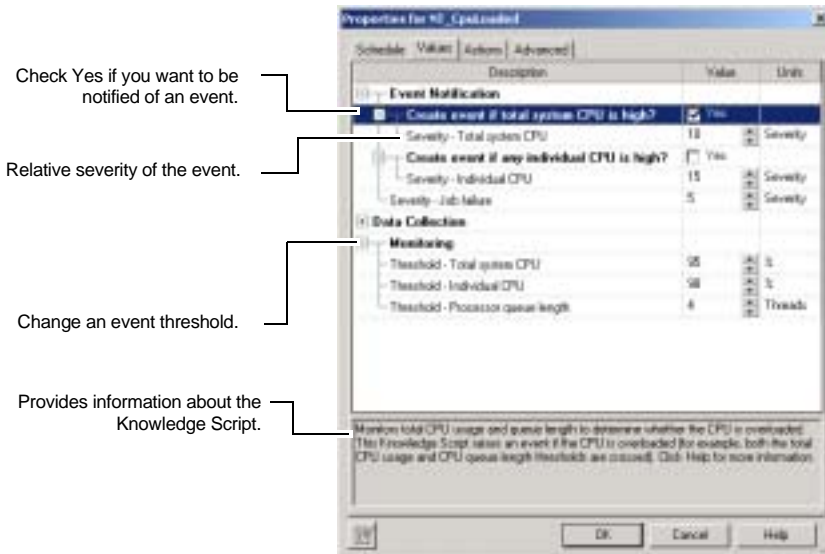
This chapter describes how events (alerts) are triggered by AppManager jobs and how to view and respond to them. The following topics are covered:

- [“How Events Are Raised” on page 153](#)
- [“Event Severity Levels” on page 155](#)
- [“Viewing Events” on page 157](#)
- [“Viewing Details about an Event” on page 161](#)
- [“Adding Comments to an Event” on page 166](#)
- [“Grouping and Filtering Events” on page 168](#)
- [“Changing the Status of an Event” on page 170](#)
- [“Deleting an Event” on page 174](#)
- [“Duplicate Events” on page 174](#)
- [“Automatically Closing Events When the Condition No Longer Exists” on page 180](#)

How Events Are Raised

When you create a Knowledge Script job, the parameters you set on the **Values** tab in the Knowledge Script Properties dialog box indicate what you want the job to do. If you want the job to send an alert or notify you in some way when a

particular condition exists, set the “Raise event. . .?” parameter to **y** (yes) or select the **Yes** check box. For example:



As a Knowledge Script job runs, it checks whether the condition you have defined exists – for example, whether a threshold value has been exceeded or a particular process is down.

If a Knowledge Script detects a problem and the Event parameter is enabled, the job triggers an event, which is:





- Logged in the AppManager repository.
- Added to the **Events** tab in the List pane.
- (In most cases) displayed in the form of an event alert in the TreeView pane in the Operator Console and in the NetIQ AppManager Operator Web Console.

In this state, the event is considered **open**, requiring your attention.

Event Severity Levels

Because some jobs check for serious problems that can severely impact your business while others simply provide informational feedback, events range in severity. The event indicators displayed in the TreeView pane in the NetIQ AppManager Operator Console and in the Operator Web Console reflect at a glance the relative severity of the problem. Also, events may trigger a sound depending on settings in the Windows Control Panel and the AppManager console preferences.

In AppManager, events are classified as follows:

Severity level	Indicator	Description
Severe Levels 1 to 10 (default)	Red  Sound: Critical Stop	A very serious problem was detected that requires your attention (for example, a computer or service is down). The Severe event indicator flashes in the TreeView pane.
Warning Levels 11 to 20 (default)	Yellow  Sound: Exclamation	Unexpected behavior or minor problems were detected that may require your attention, but they are not affecting network uptime or preventing users from doing their work (for example, a counter you want to monitor or an application you want to discover is not found on the selected computer). By default, the Warning event indicator flashes in the TreeView pane, unless you uncheck the Flash Warnings option in the Preferences dialog box.
Informational Levels 21 to 30 (default)	Blue  Sound: Asterisk	A task was successfully completed (for example, a service is successfully restarted or a discovery is completed). The Informational event indicator does not appear in the TreeView pane.
Diagnostic Levels 31 to 40 (default)	Magenta  Sound: Default Beep	Diagnostic information is being provided (for example, when debugging problems). The Diagnostic event indicator does not appear in the TreeView pane.

You can change the severity levels to better reflect your system management policies. See the following section for more information.

Note If you have trouble distinguishing any of these colors, the severity of an event is listed in the Event Properties dialog box (see [“Viewing Details about an Event” on page 161](#)) or you can rely on the sound indicator.

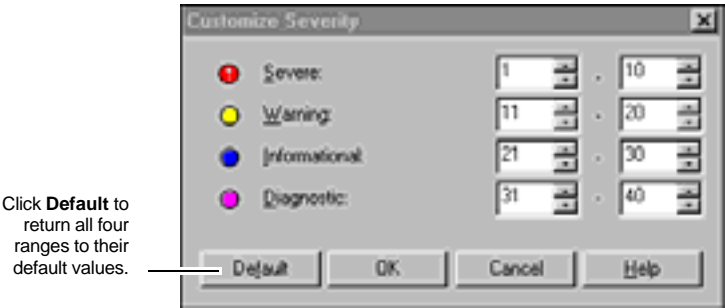
Customizing Severity Levels

By setting custom ranges, you are defining what each of the four severity levels actually mean.

Note It can take up to 24 hours for any changes you make to be reflected in the Control Center Console. The **NetIQ CC Daily Task**, which by default runs every day at midnight, updates existing events in the Control Center repository. To refresh the list of events after they are updated in the Control Center repository, you must restart the Control Center Console or press **Ctrl+F7** and then press **F5**.

- 1 Click **File > Preferences**.
- 2 Click the **Repository** tab.
- 3 Click **Severity**.
- 4 In the Customize Severity dialog box, type starting and ending values for the custom levels you want to set up. (The

fields are designed to ensure that all values are valid and do not overlap.)

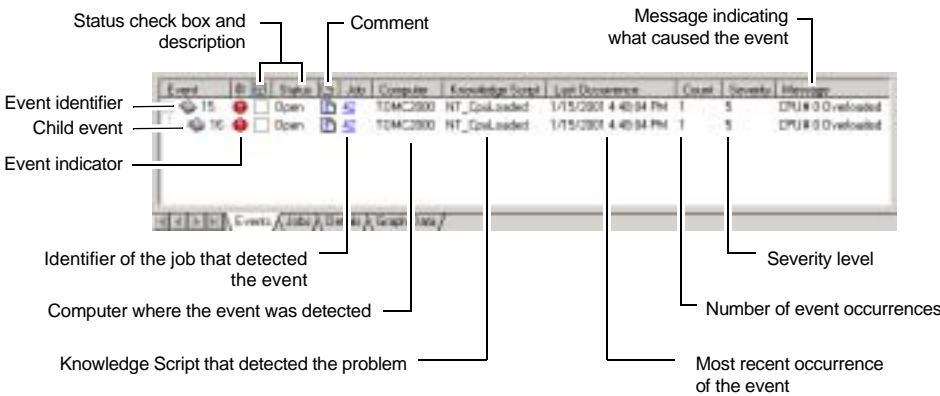


5 Click **OK**.

Viewing Events

To see a list of events generated by Knowledge Script jobs, click the **Events** tab in the List pane. If you are currently in the Master view, the **Events** tab lists all events. To see the events for a specific view—such as those associated with an Exchange or Domino Server—click the appropriate view tab.

By default, the event entry includes the following summary information:



Notes

- If text is truncated, pass the mouse cursor over it and the full text appears as a ToolTip.
- To view the parent job that created an event, click the hyperlinked Job ID. To view the child job, the parent must be expanded.
- For more information on events, see [“Viewing Details about an Event” on page 161](#).

Changing the Layout of the Events Tab

You can configure the layout of the **Events** tab to display selected columns in a particular order. Any changes are applied to the Events tab in the current view.

For information on changing the default layout of the Operator Console, see [“Setting Operator Console Preferences” on page 64](#).

To change the layout of the items on the Events tab:

- 1 In the **Events** tab of the List pane, click **List > Customize**.
- 2 In the Show Event Columns dialog box, change the layout:

Action	Steps
Show a column	Select a column name from the Available columns list and click Add .
Hide a column	Select a column name from the Show these columns list and click Remove .
Change the order in which the columns are displayed	In the Show these columns list, from top to bottom, the columns are displayed in the Jobs tab from left to right. Select a column name and click: <ul style="list-style-type: none">• Move Up to move the column to the left.• Move Down to move the column to the right.

- 3 Click **OK**.

Viewing Events Based on Status

You can select to show or hide events on the **Events** tab based on their status. Hiding events removes them from the list but does not delete the event information. For example, you might want to hide closed events so they don't clutter up the list of events you are currently interested in. Later, if you want to view the closed events for historical purposes, you can select to show them.

To determine which events are shown on the Events tab:

- 1 Click the **Events** tab in the List pane.
- 2 Right-click in the List pane and click **View**.
- 3 Select (to show) or clear (to hide) the following types of events: **Open Events**, **Acknowledged Events**, or **Closed Events**.

Note To see only a list of open events, click **TreeView > View Problems**. Be sure that **Open Events** is selected in the view list.

Expanding and Collapsing Events

Events for each job are organized in a parent-child relationship under the **Events** tab in the List pane. The first event generated by a particular job creates both a parent and a child event entry. Subsequent events from the same job are then listed under the same parent event, and an event count keeps track of how many event occurrences are generated. (The

count for the parent entry is the sum of the counts for the child entries.)

Action	Steps
View child events for a parent event	Click the + (plus sign) to the left of the parent event to examine each individual event entry.
View all child events for all parents	Ctrl+click the + (plus sign) next to any parent event.
Fold the child events back under the parent	Click the - (minus sign) to the left of the parent event.
Fold all child events back under all parents	Ctrl+click the - (minus sign) to the left of any parent event.

Changing the Sorting Order

Click any of the headings in the List pane to sort all parent entries in that field or toggle the sorting order preference. For example, events in the **Events** tab are normally sorted by event ID in ascending order. To sort the event list by severity, click the **Severity** header. To reverse the sorting order from ascending to descending, click the **Severity** header again.

Hiding Parent Events

To improve sorting of all child events in the List pane, you can hide parent events. To improve Operator Console performance, it is recommended that you do not use this feature in the Master view.


To hide parent events in the current view, in the Events tab of the List pane, click **List > View** and then unselect Parent Events (or select to show parent events).

For information on hiding parent events in the Operator Console, see [“Setting Operator Console Preferences” on page 64](#).

Viewing Details about an Event

The **Event** tab of the Event Properties dialog box includes more information (such as when the event first occurred and the user running the job) than the summary information listed in the Events tab.

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event entry; or right-click the event, and then click **Properties**. To see a child event, expand the parent entry by clicking the + (plus sign) next to it.



- 3 To acknowledge or close the event, select the appropriate check boxes. (Closing an event automatically acknowledges it. Once an event is closed, it can then be deleted.) Then click **Apply**.
- 4 Use the **Up** and **Down** buttons to browse through events as they appear in the Events tab; if the event is not displayed on the Events tab (for example, a parent event is not expanded), you cannot view the event from here.

Depending on the Knowledge Script and the event, additional tabs may be included to display even more

information.

Topic	Page Reference
Viewing event messages (F4 or the Message tab)	162
Viewing event actions (Action tab)	163
Viewing event comments (Comments tab)	164
Adding comments to an event (Comments tab)	166

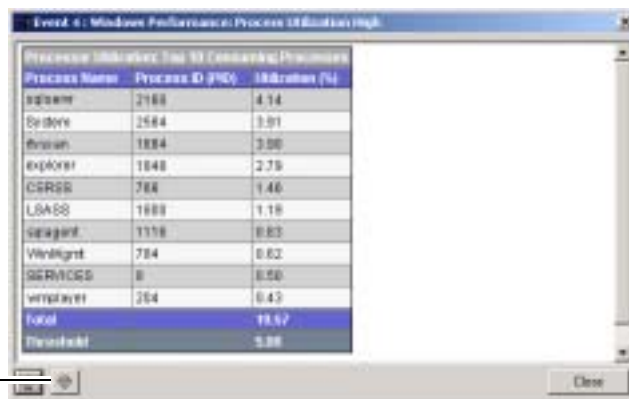
Viewing Event Messages

All child events include an **event detail message**. The message describes what caused the event or records log information that can assist you in tracking down the cause. If the event message includes a hypertext link, you can click the link to go to the corresponding URL.

To quickly view only the event message (without opening the Event Properties dialog box):

- 1 In the List pane, click the **Events** tab.
- 2 Click the + (plus sign) of a parent event to display its child events.
- 3 Click a child event entry and then press **F4**.

Click to view the properties of the next event that is visible in the Events tab.



- To view the child events for a different parent, click **Close** to close the dialog box and then repeat Steps 2 and 3.

- The new window size is retained until you change it.

Event Message Contents

Process Name	Process ID (PID)	Utilization (%)
explorer	2160	4.54
System	256	3.81
smss.exe	1684	3.80
explorer	1640	2.70
csrss.exe	768	1.86
lsass.exe	1080	1.78
explorer	1116	0.87
WinMgmt	704	0.82
SERVICES	8	0.50
smss.exe	304	0.42
Taskmgr	1912	0.30
Threshold		5.00

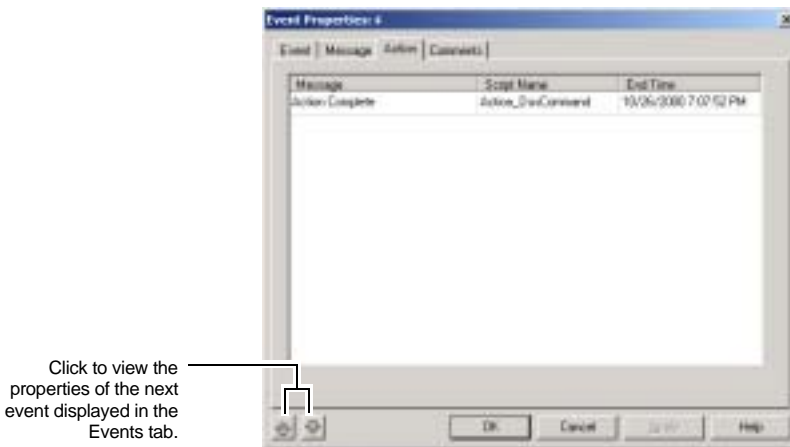
When events are detected, Knowledge Scripts can initiate a variety of actions, such as sending a page or e-mail, running an external program, or sending an SNMP trap.

If you have defined an action for a Knowledge Script job **and** the action was performed, the Event Properties dialog box includes an Action tab, which tells you when the action was performed and its current status (for example, complete or in error).

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event entry.

To see a child event, expand the parent entry by clicking the + (plus sign) next to it.

- 3 Click the **Action** tab.



- 4 Use the **Up** and **Down** buttons to browse through events as they appear in the Events tab; if the event is not displayed in the Events tab (for example, a parent event is not expanded), you cannot view the event from here.



Viewing Event Comments

Event comments are added by console operators and administrators and usually include information about what caused the event or how it was resolved. An event comment

can include a hypertext link, and you can click the link to jump to the corresponding URL.

Note If you have configured an AppManager Report Script to raise an event upon successful completion, the event comment can include a hyperlink to the actual report. For information on how to configure Report Script event comments to include a hyperlink to the actual report, see the *Reporting Guide*.

On the Events tab of the List pane, you can quickly see if an event has a comment.

Indicator	Meaning
	There is no comment information for the event entry; or in the case of a parent event, for any of its child events.
	There is comment information for this entry; or in the case of a parent event, for at least one child event.

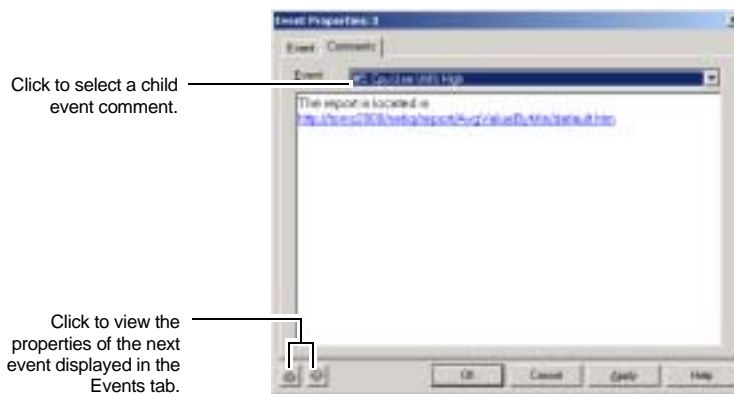
To view a comment:

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event entry; or right-click the event, and then click **Properties**.

To see a child event, expand the parent entry by clicking the + (plus sign) next to it.

- 3 Click the **Comments** tab.
- 4 Child event comments can be viewed, but not added or edited, from the parent event. On the **Comments** tab of the parent event, select the child entry from the **Event** list. Event

entries with comments appear in the list and are identified by their Event number.



- 5 Use the **Up** and **Down** buttons to browse through events as they appear in the Events tab; if the event is not displayed in the Events tab (for example, a parent event is not expanded), you cannot view the event from here.

Adding Comments to an Event

Adding comments to an event allows you to record and share information with other console operators and administrators about what caused the event or how you resolved the problem.

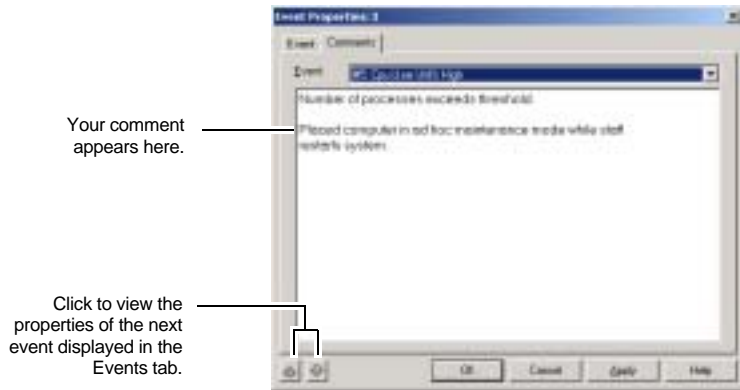
Comments can be added to parent and child events, individually.

To add a comment to an event:

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event entry.

To add a comment to a child event, expand the parent entry by clicking the + (plus sign) next to it.

- 3 Click the **Comments** tab.
- 4 Enter the text (up to 256 characters) you want saved with the event. You can also specify an URL to create a hyperlink to a web site. For example:



- 5 Click **Apply** or click **OK** to apply your changes and close the dialog box.

When you add a comment to an event, the Comment indicator in the Events tab in the List pane changes. For more information, see [“Viewing Event Comments” on page 164](#).

- 6 Use the **Up** and **Down** buttons to browse through events as they appear in the Events tab; if the event is not displayed in the Events tab (for example, a parent event is not expanded), you cannot view the event from here.


Grouping and Filtering Events

By default, AppManager displays all events in list form, with the events associated with each job as a separate parent entry. In some cases, you may want to organize events into groups or filter the information displayed based on some criteria. For example, you may want to organize events by severity, or limit the events displayed to those found on a specific computer or associated with a specific Knowledge Script.

Organizing Events into Groups

- 1 In the List pane, click the **Events** tab.
- 2 Right-click in the List pane and then click **Group By**.
- 3 Select the way you want to group events: by **Computer**, **Knowledge Script**, **Parent Severity**, or **Parent Message**.

Once you make your selection, event information is organized to match your selection. For example, if you selected Computer, the Events tab might display something similar to the following two groups:



You can expand the event information for any group by clicking the + (plus sign) to the left of the group.

Filtering the Events Displayed

You can also create filters to determine which types of events are shown in the **Events** tab.

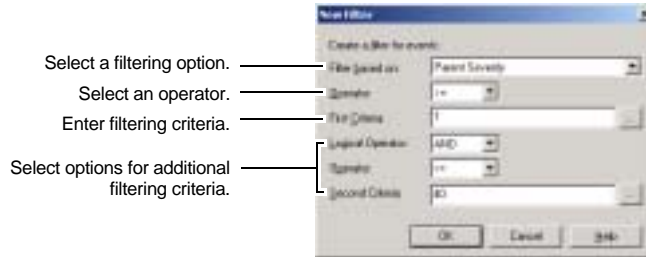
To apply event filters or create new ones:

- 1 In the List pane, click the **Events** tab.

- 2 Right-click in the List pane. Click **Filter**; then click **New Filter** or a previously defined filter.

If you select an existing filter, the information on the Events tab in the List pane is filtered based on the filter properties.

If you are creating a new filter, the New Filter dialog box displays.



- 3 In the **Filter based on** list, select a filtering option: **Computer**, **Knowledge Script**, **Parent Severity**, or **Parent Message**.
- 4 For **Operator**, choose one of the following options:

Option	Next Step
Parent Severity	Select an operator (for example, greater than > or less than <).
<ul style="list-style-type: none">• Computer• Knowledge Script• Parent Message	Leave the default operator (=).

- 5 For **First Criteria**, enter the filtering criteria you want to use or click the Browse (...) button to see a list of possible values.
- 6 To add additional filtering criteria, select **AND** or **OR** from the **Logical Operator** list. Then select an operator and the second filtering criteria (click the Browse (...) button to see a list of possible values).

7 Click **OK**.

The events displayed on the Events tab are filtered according to the criteria you have specified; the new filter is added to the List menu.

Note The possible values for both the first and second criteria are based on the filter option you select. For example, if you select **Parent Severity**, the possible values for both criteria are severity level numbers. If you want to create a filter that combines filtering options – for example, a filter that displays the events on the computer Shasta (Computer) with a severity less than five (Parent Severity) – use a combination of filtering options: TreeView > Filter List Pane by Selected Object and List > Filter.

Changing the Status of an Event

Once a Knowledge Script detects a problem and generates an event, it is added to the Events list and, depending on the severity of the event, may flash an indicator in the TreeView pane until you respond to it.

To respond to an event and turn off the event indicator, you need to either acknowledge or close the event. How you use the Acknowledged or Closed status depends on your system management policies. In general, you can acknowledge an event to indicate you are aware of the problem or that the problem is being addressed; then close it when the problem has been resolved.

Acknowledging an Event

Once a Knowledge Script raises an event, an event indicator may flash in the TreeView pane and the event is listed as Open in the List pane until you respond to it.

Acknowledging the event disables the event indicator (so that other administrators don't see alarms for events you have already responded to) and changes the status of the event in the List pane to Ack (so that other administrators know the event occurred but is being addressed).

There are three common ways to acknowledge events:

- Individually acknowledging child events in the List pane (or acknowledging all child events at once by acknowledging a parent event).
- Acknowledging all events associated with an application server, a group of servers, or all servers in a view.
- Individually acknowledging an event after viewing event details in the Event Properties dialog box.

To acknowledge an event from the List pane:

- 1 In the List pane, click the **Events** tab.
- 2 Click the event indicator of the event you want to acknowledge.



If you select a parent event, the parent and all its child events are acknowledged. If you select a child event, only the child event is acknowledged.

To acknowledge all events associated with an application server:

- 1 In the TreeView pane, select the application server, group of servers, or view icon whose events you want to acknowledge.
- 2 Right-click and then click **Acknowledge Events**.

All parent and child events are acknowledged.

To acknowledge an event from the Event Properties dialog box:

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event you want to acknowledge; or right-click the event and then click **Properties**.
- 3 In the Event Properties dialog box, select **Acknowledge**.
- 4 Click **Apply** or **OK**.

Closing an Event

There are three ways to close an event. You can:

- Individually close child events in the List pane (or close all child events at once by closing a parent event).
- Close all events associated with an application server, a group of servers, or all servers in a view.
- Individually close an event after viewing event details in the Event Properties dialog box.

When you have resolved the problem that caused an event, you can then close the event. You don't need to acknowledge an event before closing it. (However, to prevent accidental deletion of open or unresolved events, you must close an event before it can be deleted.)

Note To configure a job to automatically close an event when

the event condition no longer exists, see [page 125](#).

To close an event from the List pane:

- 1 In the List pane, click the **Events** tab.
- 2 Click the status check box for the event you want to close.



If you select a parent event, the parent and all its child events are closed. If you select a child event, only the child event is closed.

To close all events associated with an application server:

- 1 In the TreeView pane, select the application server, group of servers or view icon for which you want to close events.
- 2 Right-click and then click **Close Events**.

All parent and child events are closed.

To close an event from the Event Properties dialog box:

- 1 In the List pane, click the **Events** tab.
- 2 Double-click the event you want to close; or right-click the event and then click **Properties**.
- 3 In the Event Properties dialog box, select **Close**.
- 4 Click **Apply** or **OK**.

Deleting an Event

You don't have to acknowledge an event before you delete it. However, to prevent accidental deletion of open or unresolved events, you must close an event before it can be deleted.

To delete Closed events:

- 1 Click the **Events** tab in the List pane.
- 2 Select the event you want to delete.

To select multiple events, press the Ctrl or Shift keys as you make your selections.

- 3 Press the Delete key.

Notes

- When you select multiple events, you can delete, acknowledge, or close all of the events at once. If you attempt to delete multiple events, and there are still Open or Acknowledged events in your selection, those events are not deleted. To delete an Open or Acknowledged event, first close the event by clicking the **Status** check box in the List pane.
 - For information on configuring the AppManager repository to automatically delete an event when its associated job is deleted, see the [Remove associated events when jobs are deleted](#) option on [page 82](#).
-

Duplicate Events

The first event generated by a job creates both a parent entry and a child event entry. Subsequent events from the same job are listed as additional child event entries under the parent event. An event count keeps track of the number of child

events generated. (The count for the parent entry is the sum of the counts for each child entry.)

Event	ID	Status	Job	Computer	Knowledge Script	Last Occurrence	Count	Severity	Message
36	1	Open	08	CHRS	Discover_05	10/22/1998 1:54:05 PM	2	5	The license has expired for component 05
37	1	Open	08	CHRS	Discover_05	10/22/1998 1:54:05 PM	1	10	05 Discover 05: ver' Error = 3
38	1	Open	08	CHRS	Discover_05	10/22/1998 1:54:05 PM	1	5	The license has expired for component 05

A Knowledge Script job might detect **unique** events (for example, events with a different severity or message, as shown above), or identical events. An event is considered identical, or a **duplicate**, when a new Open child event with the same object name, event message, severity, and Job ID as a previous event occurs. Just like unique events, duplicate events are logged in the repository and, if event collapsing is disabled, are listed as additional child event entries under the parent event. Actions, if specified, are initiated for each duplicate event.

Although duplicate events are typically valid, it probably isn't useful to receive multiple alerts caused by the same condition. In addition, you don't want identical events to initiate duplicate actions, such as repeated e-mail notification or messenger dialog boxes.

AppManager offers an option to **filter temporary event conditions** (not just collapse them). You can specify the number of duplicate event conditions that must be detected during a specified number of job iterations before an event is raised. This option is useful in reducing the number of trivial events reported to the management server. Filtered event conditions are not logged in the repository and do not increment the event count of the original child event. Actions associated with ignored events are also suppressed.

After an event is raised, subsequent duplicate events can then be **collapsed** into a single event. Instead of creating new child event entries, duplicate events are collapsed into the original child event and the child event count is increased. The duplicate events are logged in the repository; however, actions

associated with collapsed duplicate events, whether the actions are initiated on the managed client or management server computer, are suppressed.

Duplicate events are collapsed within a specified time period (the default is 20 minutes). The **event collapsing time interval** can be configured to begin:

- When the first event is raised. All duplicate events within the time frame (a fixed period of time) will be collapsed into one event.
- Each time an event is generated (a dynamic period of time). For example, using the default time of 20 minutes, if a job generates duplicate events every five minutes, the 20-minute interval is restarted every five minutes, which means it never effectively expires – unless you set an option to **ignore** events.

After the original child event is closed, or after the event collapsing time interval expires, a new child event is created and an action on the management server computer, if one is specified, is initiated when an event condition is detected.

Event collapsing options can be configured globally, so that they apply to all Windows-based Knowledge Script jobs, and for individual Windows-based jobs. For more information on where to set the options for all jobs or an individual job, see the notes at the end of this section.

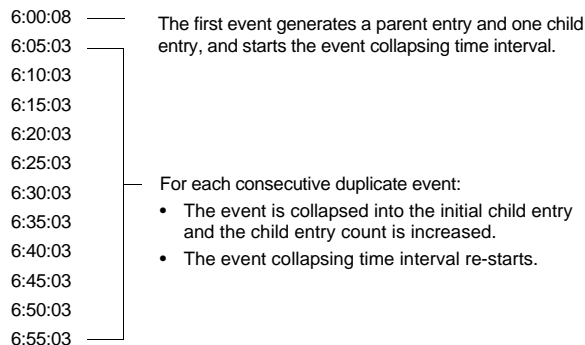
Here's an example of how event collapsing works. Assume that a job detects a problem, raises an event, and continues to detect the same condition every 5 minutes for a total of 12 identical events.

Without event collapsing enabled, a parent entry and 12 child entries are listed. The count for each child entry is 1; the count

for the parent entry is 12. All 12 events are logged in the repository and 12 instances of the specified action are initiated.

Event	Status	Job	Computer	Knowledge Script	Last Occurrence	Count
65	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	12
66	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
67	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
68	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
69	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
70	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
71	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
72	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
73	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
74	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
75	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
76	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1
77	Open	03	CHRIS	NT_CplLoaded	10/14/1999 6:55:03 PM	1

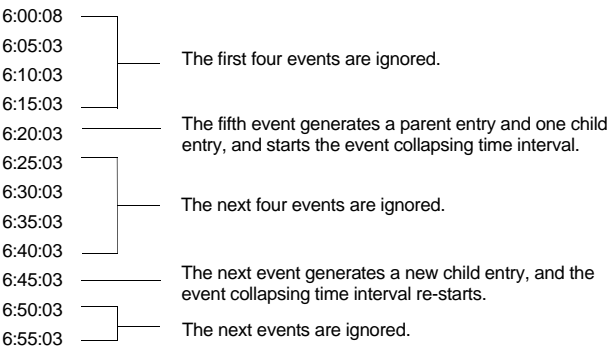
With event collapsing enabled but without ignoring events (using the default values of 20 minutes for **Time interval for event collapsing** and the event condition must occur 1 **time** during 1 **iteration**), the first event creates a parent entry and one child entry, and begins the 20-minute time interval. Because the duplicate events take place at five-minute intervals, the 20-minute event collapsing interval never expires (it restarts after each duplicate event). And, because the number of events and the number of job iterations is set to 1, duplicate events are not ignored; instead, each is collapsed into the initial child entry.



A parent entry and one child entry is listed; the count for each is 12. All 12 events are logged in the repository but only one instance of the specified action is initiated.

Event	Status	Job	Computer	Knowledge Script	Last Occurrence	Count
40	Open	62	CHRIS	NT_CplLoaded	10/14/1999 6:58:03 PM	12
41	Open	62	CHRIS	NT_CplLoaded	10/14/1999 6:58:03 PM	12

With event collapsing enabled and with events ignored (using 20 minutes for **Time interval for event collapsing** and the event condition must occur **5 times** within **5 job iterations**), the first four events are ignored. The fifth event generates a parent entry and one child entry, and begins the 20-minute time interval. The next four events (six through nine) take place during the 20-minute interval and are ignored. The next event (ten) generates a new child event and restarts the time interval and the job iteration interval. The eleventh and twelfth events are ignored.



A parent entry and two child entries are listed. The count for each child entry is 1; the count for the parent entry is 2. Only the two listed child events are logged in the repository and their actions initiated. The other events and actions are ignored.

Event	Status	Job	Computer	Knowledge Script	Last Occurrence	Count
54	Open	70	CHRIS	NT_CplLoaded	10/14/1999 6:45:03 PM	2
62	Open	70	CHRIS	NT_CplLoaded	10/14/1999 6:20:03 PM	1
55	Open	70	CHRIS	NT_CplLoaded	10/14/1999 6:45:03 PM	1

Notes

- For information on configuring repository preferences to set global, default event collapsing options for all Knowledge Script jobs, see [“Setting Global Event Collapsing Options” on page 179](#).
 - You can configure a job to either collapse duplicate events or automatically close events when the event condition no longer exists; you cannot do both. See [“Automatically Closing Events When the Condition No Longer Exists” on page 180](#) for more information.
 - For information on changing event collapsing options for individual Windows-based jobs, see [“Changing Event Collapsing Options for One Job” on page 123](#).
-

Setting Global Event Collapsing Options

Options for collapsing duplicate events can be set on a global basis, so that they automatically apply to new events.

To set global event collapsing options for all jobs:

- 1 Click **File > Preferences**.
- 2 Click the **Repository** tab, and then click **Advanced Properties**.
- 3 On the **Advanced** tab, select **Collapse duplicate events into a single event**.

Duplicate events are logged in the repository; however, actions associated with collapsed duplicate events, whether the actions are initiated on the managed client, management server, or proxy computer, are suppressed.

- 4 Set the **Time interval for event collapsing** in minutes.

The default time interval is 20 minutes.

All duplicate events within the time interval are collapsed into the original child event. The time interval begins again each time an event is generated, unless you set the option to have events ignored.

- 5 Determine when an event should actually be raised by setting the number times that an event must occur during a specified number of job iterations (job iteration interval).

For example, if set to 5 **times** within 5 **job iterations**, the first four events are ignored; the fifth occurrence generates an actual event and starts the event collapsing time interval. The next generated event will take place either:

- When five more events take place during the time interval, the first four of which will be ignored.
- After the time interval expires.

The ignored events are not logged in the repository and do not increment the event count of the original child event.

- 6 Click **OK**.

Automatically Closing Events When the Condition No Longer Exists

For some event conditions, it is more useful to raise an event when the condition is first detected and then raise a second event when the event condition stops. You can configure an AppManager job to do this; you can also choose to have the first event closed automatically.

A job may be set to raise an event and automatically close the original event if an event condition is detected. In such a case, the first time that condition occurs, a parent event and a child event are created with a status of Open. Then, if subsequent, duplicate event conditions are detected, additional child events are suppressed, and the event count for the child event is not

increased. Once the event condition no longer exists, a new event indicates the change in the event condition, and the status of the original event is changed to Closed.


For example, with these settings in place, the first time a job runs, an event condition does not exist; an event is not raised. The next time the job runs, an event condition is detected and a parent event and a child event are raised.



The job runs several times, and each time the event condition exists. Additional child events are not created, and the child event count is not increased; there is only one child event on the **Events** tab of the List pane.




On the next iteration, the job runs, and the event condition is not detected. The status of the child event is changed to Closed, and a new informational event is created (with an event severity of 40) that indicates the change in the event condition status.



If you double-click the event that indicates the change in status of the event condition (event 582), in the Event

Properties dialog box, the **Message** tab indicates that the original event (event 581) has been closed.



If the Knowledge Script job detects the event condition again, a new child event is created.

Notes

- To automatically close an event when the original event condition no longer exists, you must configure the job to raise an event with a specified event severity level each time the event condition is not detected. This event is for informational purposes only and its event severity should be set accordingly. To prevent the parent event from having an Open status when the original event is Closed, you can configure the AppManager repository preferences to automatically close all open events based on severity level. See [“Set Preferences for Event Management” on page 79](#) for more information.
- You can configure a job to either automatically close events when the event condition no longer exists or collapse duplicate events; you cannot do both. For information on event collapsing, see [“Duplicate Events” on page 174](#).

- To configure an individual job to automatically close an event when the event condition no longer exists, you can do so from the **Advanced** tab in the Knowledge Script Properties dialog box. For more information, see [“Automatically Closing Events When the Condition No Longer Exists” on page 180.](#)
 - Global, default options for automatically closing an event when the event condition no longer exists are set on the Repository tab in the Preferences dialog box. For more information, see [“Setting Global Options to Automatically Close Events When the Event Condition No Longer Exists” on page 183.](#)
-

Setting Global Options to Automatically Close Events When the Event Condition No Longer Exists

To set global options for all jobs:

- 1 Click **File > Preferences.**
- 2 Click the **Repository** tab and then click **Advanced Properties.**
- 3 On the **Advanced** tab, select **Generate a new event when original event condition no longer exists.** With this option selected, an event is raised when the condition that initiated the original event no longer exists.

With this option selected, specify the event severity when an event is raised in response to the change in the event condition. By default, the event severity is set to 20.

- 4 To configure the job to close the original event when the event condition no longer exists, select **Automatically close original event.**
- 5 Click **OK.**

Creating Graphs

This chapter describes how to collect graph data using AppManager jobs and how to view graph data in the Graph pane of the NetIQ AppManager Operator Console. The following topics are covered:

- [“Collecting Data for Graphs” on page 185](#)
- [“Viewing the List of Graph Data Streams” on page 188](#)
- [“Grouping and Filtering Data Streams” on page 189](#)
- [“Creating a Graph” on page 193](#)
- [“Viewing and Working with Graphs” on page 198](#)

Note For information about creating a chart of an AppManager data stream in the AppManager Chart Console, see the AppManager Help.

Collecting Data for Graphs

Before you can create a graph, you first need to create a Knowledge Script job that collects the data you want graphed. Most Knowledge Scripts can be used to collect data for graphs and reports.

To create a job that collects data, click the **Values** tab of the Knowledge Script Properties dialog box, and select the **Yes**

check box next to the **Collect Data** parameter (or set that parameter to **y**). For example:

Check Yes to collect data about a particular condition.



Data Points and Data Streams

The specific numeric data collected in one run of a Knowledge Script job is considered a **data point**.

A series of these data points, collected as a Knowledge Script job runs at set intervals, is called a **data stream**. Each data stream is displayed on the **Graph Data** tab in the List pane. Data streams are displayed only in the view to which they pertain, for example, data collected about components of SQL Server will not be visible in the NT or IIS view.



A single Knowledge Script job may create several data streams if it collects more than one kind of information. For example, the NT_CpuResource Knowledge Script collects four types of data points (User CPU, Number of processes, All threads, and Interrupts for all CPU), which appear as four separate data streams in the List pane.

Displaying the Data in a Graph

To display the information collected in a graphical form, drag the data stream icon from the Graph Data tab in the List pane to the Graph pane. You can combine multiple data streams in the same chart or graph each data stream separately. There are no predefined rules restricting how you can combine data streams, so you can combine them in any way you find useful.

Each data stream generated by a Knowledge Script is stored in the AppManager repository so you can view data at a later date. A graph is really just a view of one or more data streams, not data sent directly from a managed computer to the NetIQ AppManager Operator Console or Operator Web Console.

As the job runs, new data is constantly appended to the data stream. Each data stream has a maximum time period, in days, that data points are stored for graphing; and, periodically, data streams are truncated in the repository to keep them efficient while displaying the most recent information.


The AppManager Repository Preferences let you set the default time period during which the data points in a data stream are available for display in the Graph pane of the Operator Console (and the Chart Console), as well as the interval for deleting data points that exceed a specified time period. By default, the AppManager repository keeps 8 days' worth of data for charts and graphs. To view more than the most recent 8 days of data, you should consider running an AppManager report before increasing the time period to keep data for charts

and graphs. See [“Set Preferences for Data Management” on page 76](#) for more information.

If you change the repository preferences for the number of days that data points are available, the change affects new jobs only. Existing data streams must be updated to change the number of days to keep data points in the data stream, up to the maximum set for the repository. For information on configuring the number of days that points for a particular data stream can appear in a graph or chart, see [“Changing Data Stream Properties” on page 192](#).

Viewing the List of Graph Data Streams

Each job that collects data creates at least one new data stream. Each data stream has a unique ID, and is displayed in the **Graph Data** tab in the List pane. The **Graph Data** tab includes the following information:



Data Source	Job	Knowledge Script	Computer	Max days	Cur. poi.	Legend
	12	NT_Memo1	TUMC2008	9	128	Mem/Post/page 5

Labels and their corresponding table columns:

- Data stream identifier: Points to the Job column (12).
- Job identifier (job collecting the data): Points to the Job column (12).
- Knowledge Script that is running: Points to the Knowledge Script column (NT_Memo1).
- Computer where the data was collected: Points to the Computer column (TUMC2008).
- Maximum number of days that data points are stored in the repository: Points to the Max days column (9).
- Graph legend for the type of data collected: Points to the Legend column (Mem/Post/page 5).
- Data points collected: Points to the Cur. poi. column (128).

Note To view the parent job that created a data stream, click the hyperlinked Job ID. To view the child job, the parent must be expanded.

Changing the Layout of the Graph Data Tab

You can configure the layout of the **Graph Data** tab to display additional columns and to change the order in which the columns appear. Any changes are applied to the Graph Data tab in the current view.

To change the layout of items on the Graph Data tab:

- 1 In the List pane, click the Graph Data tab and click **List > Customize**.
- 2 In the Show Graph Data Columns dialog box, change the layout of the Graph Data tab:

Action	Steps to Take
Hide a column	Select a column name from the Show these columns list and click Remove .
Show a column	Select a column name from the Available columns list and click Add .
Change the order in which the columns are displayed	In the Available columns list, from top to bottom, the columns are displayed on the Jobs tab from left to right. Select a column name and click: <ul style="list-style-type: none">• Move Up to move the column to the left.• Move Down to move the column to the right.

- 3 When you finish, click **OK**.

Note For information on changing the default layout of the Operator Console, see [“Setting Operator Console Preferences” on page 64](#).

Grouping and Filtering Data Streams

By default, AppManager displays all data streams in lists. In some cases, you may want to organize data streams into groups or filter the information based on some criterion. For example, you may want to organize data streams by the type of data collected, or limit the data streams displayed to those





found on a specific computer or to those associated with a specific Knowledge Script.

Organizing Data Streams into Groups

To organize data streams based on their type or source:

- 1 Click the **Graph Data** tab in the List pane.
- 2 Right-click in the List pane and click **Group By**.
- 3 Select how you want events grouped: **Job ID**, **Knowledge Script**, **Computer**, **Maximum Days**, or **Legend**.

Once you make your selection, event information is folded into groups as specified. For example, if you selected **Computer**, the **Graph Data** tab might display something similar to the following two groups:

  SHASTA
  TSLAB

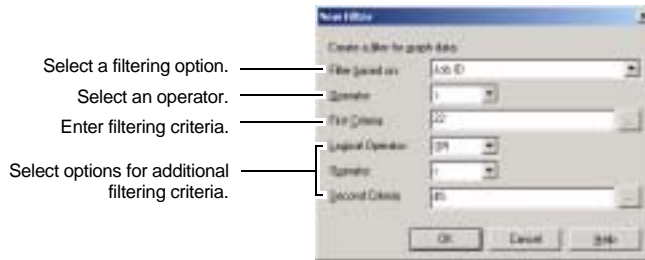
You can expand the event information for any group by clicking the + (plus sign) to the left of that group.

Filtering the Data Streams Displayed

- 1 Click the **Graph Data** tab in the List pane.
- 2 Right-click in the List pane. Click **Filter**; then click **New Filter** or a previously defined filter.

If you select an existing filter, the information in the Graph Data tab in the List pane is immediately filtered based on the filter properties.

If you are creating a new filter, the New Filter dialog box displays.



- 3 In the **Filter based on** list, select a filtering option: **Job ID**, **Knowledge Script**, **Computer**, **Maximum Days**, or **Legend**.
- 4 For **Operator**, choose from the following options:

Option	Next Step
<ul style="list-style-type: none">• Job ID• Maximum Days	Select an operator (for example, greater than > or less than <).
<ul style="list-style-type: none">• Computer• Knowledge Script• Legend	Leave the default operator (=).

- 5 For **First Criteria**, enter the filtering criteria you want to use or click the Browse (...) button to see a list of possible values.

- 6 To add additional filtering criteria, select **AND** or **OR** from the **Logical Operator** list. Then select an operator and the second filtering criteria (click the Browse (...) button to see a list of possible values).
- 7 Click **OK**.

The Graph **Data** tab is filtered according the criteria you have specified; the new filter is added to the List menu.

Note The possible values for both the first and second criteria are based on the filter option you select. For example, if you select Job ID, the possible values for both criteria are Job ID numbers. If you want to create a filter that combines filtering options – for example, a filter that displays the data streams collected on the machine Shasta (Computer) for jobs less or greater than a particular job ID (Job ID) – use a combination of the TreeView > Filter List Pane by Selected Object option and a List > Filter option.


Changing Data Stream Properties

Each data stream has properties, such as the maximum number of days that a data point is available for graphing and the legend used to describe the data. These can be edited.

To view or change the properties for individual data streams:

- 1 Click the **Graph Data** tab in the List pane.

- 2 Double-click a data stream.



Field	Description
Maximum days	Specify a time period, in days, to make data points in the current data stream available for display in the Graph pane of the Operator Console and in the Chart Console. The default time period is 8 days. Note The default value for the maximum number of days can be configured globally, for all data streams. For more information, “Set Preferences for Data Management” on page 76 .
Legend text	Type the legend text you want.
Legend unit	Type the legend unit you want (such as MB).

Note Changes you make only affect the selected data stream.

- 3 Click **OK**.

Creating a Graph

To create a graph, you first need to run a Knowledge Script job that collects data. Then you select the data stream for graphing and select graphing preferences.

To run a Knowledge Script job that collects data:

- 1 In the Knowledge Script pane, click the appropriate tab and then select a Knowledge Script capable of collecting data.

- 2 Drag the Knowledge Script icon to the TreeView pane, dropping it on the computer or resource you want to collect information about.
- 3 In the Properties dialog box, click the **Values** tab.
- 4 Enable the **Collect Data** parameter (select the **Yes** check box or set it to **y**).
- 5 Set any other options you want on any of the dialog box tabs.

Note The **Advanced** Knowledge Script Properties tab includes options for data collection. See [“Changing Data Collection Options for One Job” on page 128](#) for more information.

- 6 Click **OK** to start the job.

Once the job is running, a data stream is added to the **Graph Data** tab in the List pane for each type of information the job is set to collect. The next step is to create the graph.

To create a new graph from the data you’ve collected:

- 1 Click the **Graph Data** tab in the List pane and locate the data stream you want to graph.

You may have to wait until the **Cur Points** field indicates that some data points have been collected.

To select multiple data streams, press the Ctrl or Shift keys as you make your selections.

- 2 Drag the data stream from the List pane to the Graph pane.

- 3 At the Create Graph dialog box, make sure **Yes** is selected for **New Graph**.



- 4 Enter a name for the graph or use the default name. (The name of the graph appears on a tab in the Graph pane.)

When graphing more than one data stream, you can select the options in Step 5 through Step 8 for only the first data stream.

- 5 For the **Type** field, select one of the following options:

Option	Description
Realtime	Graphs data as it is collected.
Historical	Graphs data that was previously collected, during a particular time period.

- 6 For **Display**, select **All Data** or an average. By default, all data in the data stream is displayed.
- 7 If you are graphing data for a particular time period, move the slider bars left or right to specify the first (**Start Time**) and last (**End Time**) data points to display.

Note See [“Changing Data Stream Properties” on page 192](#) for information on configuring the number of days that a data point can be displayed in a graph or chart.

- 8 Click **OK** or **OK to All**.

Notes

- A graph can have up to eight data streams. See [“Adding Data Streams to an Existing Graph”](#) on page 199.
 - For information about creating a chart of an AppManager data stream, see the AppManager Help.
-

Changing Graph Properties

AppManager graphs provide many opportunities for customization. It's easy to change the appearance of a graph, or to change its scaling or title.

To change the properties of a graph:

- 1 In the Graph pane, right-click the graph and then click **Properties**.
- 2 Use the options on the General, Series, Scale, 3-D, and Titles tabs to change graph properties, such as the graph type, series color, X-axis and Y-axis scaling, angle of the 3-D view, and titles.

Using commands on the Graph menu, you can delete series from a graph, export and import graph data, and print and delete graphs. In the Graph pane, right-click and then select a menu item.


Here is a brief summary of the commands on the Graph menu (many of which are also available on the graphing toolbar).

Option	Description
Tile Graphs	Tiles graphs within the Graph pane. When this command is not checked, graphs are shown on different tabs.
Delete Graph	Deletes the selected graph.
Delete Series	Deletes the selected series in the graph.

Option	Description
Import Graph	Imports an AppManager graph that was previously exported as a ChartFX file (.chf). For more information, see “Importing a Graph” on page 204.
Export Graph	Exports an AppManager graph to a ChartFX (.chf) or comma-separated value (.csv) file. For more information, see “Exporting a Graph” on page 204.
Copy Graph to Clipboard	Copies the graph to the Windows Clipboard as a bitmap.
Copy Data to Clipboard	Copies the graph data to the Windows Clipboard as text.
3-D View	Toggles between 2-D and 3-D graphs.
Zoom	Provides a zoom tool to enlarge a portion of the graph. Click Zoom ; then drag a selection rectangle around the area you want to enlarge. (The cursor turns into a magnifying glass and a bounding box outlines the selected area.) To return to the normal view, deselect Zoom .
Show X-Axis Values	Shows or hides the values on the X-axis.
Show Series Legend	Shows or hides the data series legend.
Show Vertical Grid	Draws vertical lines on the graph.
Show Horizontal Grid	Draws horizontal lines on the graph.
Show Palette Bar	Shows a palette for selecting colors. Drag a color and drop it on the part of the graph you want to change (for example, a data series, the graph frame, or the pane background).
Show Pattern Bar	Shows a palette for selecting patterns. Drag a pattern and drop it on a data series.
Show Data Editor	Displays graph data in spreadsheet form.
Font	Changes the fonts, colors, and font size of the text in the graph.
Print Graph	Prints the selected graph.
Hide Pane	Hides the Graph pane.
Properties	Displays a dialog box for changing all graph properties.

Viewing and Working with Graphs

Graphs are created and displayed in the Graph pane. You may need to adjust the size of the Graph pane to see the entire graph. At the top of each graph is the graph name. The X-axis shows the time the data was taken and the Y-axis shows the value of the data. Use the scroll bar across the bottom of the Graph pane to see data points over time up to the most recent information.



Normally, each graph is displayed on its own tab. You can switch between graphs by clicking the tabs at the bottom of the Graph pane. To show multiple graphs at once in the Graph pane, right-click in the Graph pane and then click **Tile Graphs**.

Tip If you notice that a graph is missing one or more data points, you should determine if there are any AppManager events to indicate why data points are missing.

Graphing Multiple Data Streams in one Graph

Graphing multiple data streams in a single graph is useful when the data sources are related (for example, when you are collecting the same information from several different computers or collecting related data from a single computer).

The data streams don't have to be related to each other to be placed in a single graph, but your graphs will be easier to interpret if the information shares a similar scale or units.

Some examples of data from different areas that you may want to combine in one graph include:

- User CPU data from different computers. For example, because each computer is expected to maintain its User mode CPU utilization between 0 and 100%, it makes sense to have multiple CPU graphs displayed together.
- Printer queue lengths from multiple printers – connected to a single computer or multiple computers.
- CPU, disk usage, and other statistics for each computer. For example, you may want to see graphs that summarize activity on a per computer basis.

It's best to experiment with different methods and combinations of data streams to find combinations that best suit your needs. You can add data streams to an existing graph or create a new graph for multiple data streams. The following topics provide instructions.

Adding Data Streams to an Existing Graph

New data streams can easily be added to an existing graph.

To add a data stream to a graph:

- 1 In the List pane, click the **Graph Data** tab.
- 2 Locate the data streams you want to add to a graph.


To select multiple data streams, press the Ctrl or Shift keys as you make your selections.
- 3 Drag the data streams from the List pane to the Graph pane.
- 4 In the Create Graph dialog box, make sure **No** is selected for **New Graph**.

- 5 Select the name of the graph you want to add the data streams to. (The name of the currently selected graph in the Graph pane appears by default.)

When adding data streams to an existing graph, you can't set Type, Display, Start Time and End Time options

- 6 Click **OK** or **OK to All**.

The following figure shows such an example of a graph with three different data streams graphed together.



Note When graphing data from multiple data streams, use the 3-D property. Because data streams are placed in the graph according to the order in which you select or drag them to the Graph pane, one data stream may overlay another, making some difficult to read. If this happens, change the 3-D view angle or the graph type. To change graph properties, right-click in the graph and then click **Properties**.

Getting Details on Graph Data Points

To see details of a particular data point on a graph (for example, a high or low point in the graph), double-click the graph data point you are interested in.



The Graph Data Details dialog box includes the following information:

Field	Description
Job ID	A unique identifier for the job that collected the graph data.
Data ID	A unique identifier for the data source in the AppManager repository.
Graph ID	(Not currently used) A unique identifier for the graph.
Stream ID	A unique identifier for the data stream within the Knowledge Script.
Machine	The name of the computer that sent the data.
Legend	A short description of the data collected in this data stream.
Time	Time the data was collected.
Value	The numeric value of the data point.

If the Knowledge Script collects additional details or both detail and summary information (for example, a Knowledge Script that tracks CPU usage may include privileged usage, user usage, and total usage), that information is included in the Graph Data Details dialog box. Some Knowledge Scripts show

- 2 Modify the graph properties (for example, the graph type, y-axis scale, color, or other display characteristics) to reflect the style you want to use as a default.
- 3 Click **File > Preferences**.
- 4 On the **Console** tab, click **Save** in the Graph group to save the current graph style as the default style.

The graph style for the currently selected graph will be used by default for all subsequent graphs. You can still edit the graph style for any individual graph.

Saving and Reloading Current Graphs

To reload the graphs that are currently displayed the next time you start the AppManager Operator Console:

- 1 In the Graph pane, display the graphs you want loaded by default.
- 2 Click **File > Preferences**.
- 3 On the **Console** tab, check **Save graphs on exit and reload them on startup** in the Graph group.

Printing a Graph

To send a graph to a printer:

- 1 In the Graph pane, select the graph you want to print.
- 2 Right-click the graph, and then click **Print Graph**.
- 3 In the Print Graph dialog box, click **All** (or click **Pages** and then enter a page range).

When printing a page range, use the scroll bar in the Graph pane to help you decide where to begin printing. Keep in mind that the most recent information will be on the last page, and the first pages may be blank (if, for example, you

have only recently started collecting data for the graph) or contain the oldest information.

- 4 Click **OK**.

Exporting a Graph

To export a graph to a file:

- 1 In the Graph pane, select the graph to be exported.
- 2 Right-click the graph, and then click **Export Graph**.
- 3 Select the format you want: **ChartFX File** or **Comma Delimited**, and click **OK**.
- 4 Select a location to save the file and give the file a name.
- 5 Click **Save**.

Importing a Graph

AppManager supports importing graphs, but only in ChartFX (.chf) file format.

To import a ChartFX file:

- 1 In the Graph pane, select the graph to be imported, right-click on the graph, and then click **Import Graph**.
- 2 Select a ChartFX (.chf) file and click **Open**.
- 3 The graph appears in the Graph pane.

Deleting a Graph

Deleting a graph doesn't delete the data stream used to create the graph. The Knowledge Script will continue to collect data if it's still running. The data stream remains in the Graph Data tab in the List pane to provide access to historical information.

To delete a graph:

- 1 In the Graph pane, select the tab of the graph to be deleted.
- 2 Right-click the graph, and click **Delete Graph**.

Deleting a Graph Data Stream

Graphs and data streams are deleted separately.

You can delete a data stream used in a graph from the List pane or from the Graph pane. See below for instructions.

To delete a graph data stream from the List pane:

- 1 Click the **Graph Data** tab in the List pane.
- 2 Select the graph data stream you want to delete.

To select multiple data streams, press the Ctrl or Shift keys as you make your selections.

- 3 Right-click, and click **Delete Graph Data**.

To delete a graph data stream from the Graph pane:

- 1 In the Graph pane, right-click and then click **Delete Series**.
- 2 Select the graph data stream you want to delete.
- 3 To select multiple data streams, press the Ctrl or Shift keys as you make your selections.
- 4 Click **OK**.

Note For information on configuring the AppManager repository to automatically delete a graph data stream when its associated job is deleted, see the [Remove associated graph data when jobs are deleted](#) option on [page 79](#).

Using Snapshot Views

This chapter describes how you can use the NetIQ AppManager Operator Console to organize the resources you want to monitor into snapshot views. Snapshot views cannot be configured from the Operator Web Console. The following topics are covered:

- [“About Snapshot Views” on page 207](#)
- [“Creating a Snapshot View” on page 208](#)
- [“Updating a Snapshot View” on page 211](#)

About Snapshot Views

A *snapshot view* allows you to organize and display a “snapshot” of the resources in an existing standard view or the **Master** view. Unlike a standard view or the **Master** view, the resources that appear in a snapshot view are not updated automatically when you run a Discovery Knowledge Script.

When deciding how to use a snapshot view, keep the following in mind:

- If the snapshot view is based on the **Master** view, you can run Discovery Knowledge Scripts. Keep in mind that you must manually add the discovered resources you want to the snapshot view.
- The resources that appear in a snapshot view are a copy of the resources in the view upon which the snapshot view is based; any changes you make to resource objects in a snapshot view do not affect the corresponding objects in

the “based on” view. For example, you can delete a discovered resource from a snapshot view without affecting any jobs that are running that resource object.

- Typically, a snapshot view is configured to display a subset of the resources for an application or system.
- Parent information for jobs is displayed but you cannot apply changes to a parent job from a snapshot view. To modify the properties of a job and start, stop, and delete a job, you must apply your changes to the child job.

If you cannot access the **Master** view or any snapshot views based on the **Master** view, check with your AppManager administrator to verify your user account has access to the **Master** view.

As with all views, you can apply a monitoring policy to a snapshot view or one of its groups to monitor matching resources using a pre-defined set of Knowledge Scripts. For more information, see [Chapter 9, “Monitoring by Policy.”](#)

Creating a Snapshot View

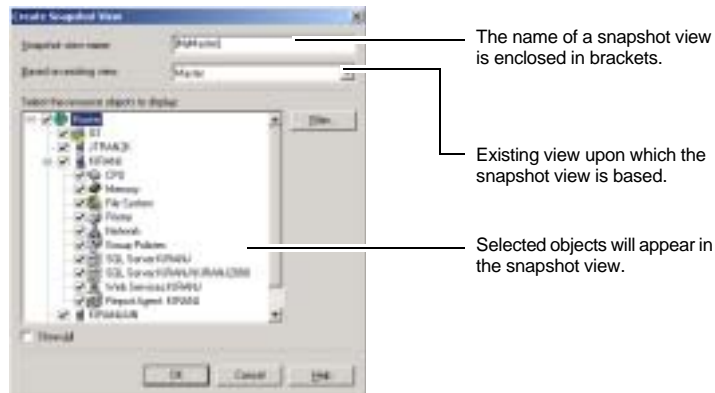
To create a snapshot view, you must specify a view upon which the snapshot view is based. Typically, the snapshot view will display a subset of the resource objects in the “based on” view.

You can add any of the resources that appear in the “based on” view to the snapshot view:

- 1 Open the View Manager by clicking **View > View Manager**.



- 2 In the View Manager, click the **Snapshot ...** button.
- 3 In the Create Snapshot View dialog box, type a name for the snapshot view and then click **Based on existing view** to select an existing view to serve as the basis for the snapshot view.



- 4 After you select an existing view upon which to base the snapshot view, a list of objects from the “based on” view is displayed. You can deselect any particular objects in the list

that you do not want and filter the types of objects that appear in the list.

- 5 Select **Show all** to view all discovered detail objects. If you have a large environment, it may take several minutes to display all detail objects.

By default, detail objects are hidden to reduce the amount of time required to display this dialog box. If you change the Show all setting, the new setting is preserved the next time you open this dialog box.

When viewing detail objects, by default, all detail objects are selected which means they will be displayed. De-select the detail objects that you do not want to include in the snapshot view.

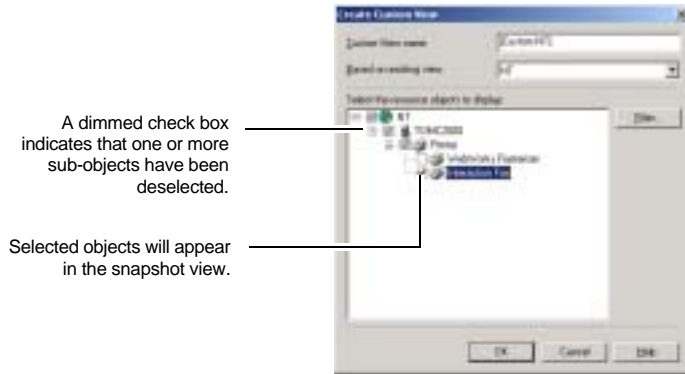
- 6 To hide some of the available objects in the list, click **Filter ...** and then deselect the objects that you do not want.

The filter for the snapshot view is configured to only display Printer objects.



When you have finished, click **OK**.

- 7 In the Create Snapshot View dialog box, click **OK** to create the snapshot view with the selected objects.



- 8 Click **OK** to close the View Manager.

In the Operator Console, you can click the snapshot view tab to display its resource objects. The name of a snapshot view is enclosed in brackets [].

Updating a Snapshot View

You can update a snapshot view to add and delete discovered resources such as systems, applications, or resource objects.

If you started a job in a snapshot view and subsequently add resource objects, you must modify the properties of an existing job to monitor the additional objects. For more information, see [“Selecting the Objects to Monitor” on page 120](#) for more information.

Note that if you delete resource objects from a snapshot view, existing jobs continue to monitor the deleted resource object.

If you have enabled a monitoring policy for a snapshot view, when you add a matching resource, the monitoring policy automatically monitors the resource. When you add resources to a monitoring policy, it can take up to one minute to monitor

the added resources. For information on enabling a monitoring policy, see [“Creating a Monitoring Policy” on page 264](#).

Adding Discovered Resources

You can add discovered resources to a snapshot view. If you cannot add the resources you want, make sure they appear in the view upon which the snapshot view is based.

- 1 Click to select a snapshot view.

The name of a snapshot view is enclosed in brackets, for example, **[NT West]**.

- 2 Click **TreeView > Add Discovered Resources**.

- 3 In the Add Discovered Resources dialog box, click to select the resources you want to add. To add a computer with no discovered resources, click to select the computer.



- 4 Click **OK**.

Deleting Objects

If you delete discovered resources from a snapshot view, the changes you make do not affect the view upon which the snapshot view is based. Any jobs running on the deleted object continue to run. See [“Deleting an Object from the](#)

[TreeView Pane](#)” on [page 55](#) for information on deleting objects.

Discovering Resources

To discover resources from a snapshot view, the snapshot view must be based on the **Master** view. To discover resources, run a Discovery Knowledge Script or configure a monitoring policy to discover resources on a scheduled basis.

After you discover resources, they are automatically displayed in the **Master** view and the corresponding standard view; to display discovered resources in the snapshot view, you must add them. See [“Adding Discovered Resources” on page 212](#) for more information.

Renaming or Deleting a Snapshot View

Use the View Manager to rename or delete a snapshot view. To open the View Manager, click **View > View Manager**. See [“Selecting a View” on page 39](#) for more information.

Using Dynamic Views

This chapter describes how you can use the AppManager Operator Console to organize the resources you want to monitor into dynamic views. Dynamic views cannot be configured from the Operator Web Console.

In general, NetIQ recommends that you do not use the Operator Console to create dynamic views and instead, use the Control Center Console to create rule-based management groups. Management groups provide an easier, more powerful, rule-based mechanism to dynamically manage the servers in your environment.

Implementing a monitoring policy on a dynamic view provides a powerful and flexible way to automatically manage the resources in your environment. Be sure to read the contents of this chapter carefully. Dynamic views are an advanced feature and you must understand how dynamic views work in order to configure them properly. The following topics are covered:

- [“Understanding Dynamic Views” on page 216](#)
- [“What Kinds of Views Can You Create?” on page 217](#)
- [“Understanding How Rules Work” on page 217](#)
- [“Understanding Custom Properties” on page 230](#)
- [“Creating a Dynamic View” on page 232](#)
- [“Organizing Computers into Groups” on page 237](#)
- [“Tips and Tricks” on page 243](#)

Understanding Dynamic Views

A *dynamic view* uses one or more rules to describe the resources that should appear in the TreeView pane. A *rule* consists of a query statement that describes the objects to include and optionally, a query statement that describes any exceptions.

As your AppManager environment changes, dynamic views automatically update to display servers that match the rules for a view.

Unlike standard views and snapshot views, a dynamic view uses rules to determine the resources that are displayed. You cannot drag-and-drop servers in the TreeView pane or delete objects from a dynamic view. To change the servers that appear in a dynamic view, you must configure the rules that define the view.

You can organize the contents of a dynamic view into one or more dynamic groups. Each dynamic group uses a single rule to determine its contents. At this time, you cannot organize groups within another group.

As with all views, you can apply a monitoring policy to a dynamic view or one of its dynamic groups to automatically monitor matching resources using a pre-defined set of Knowledge Scripts. Monitoring dynamic views by policy works the same as it does for standard views and snapshot views. For more information, see [Chapter 9, “Monitoring by Policy.”](#)

AppManager reports can also be configured to report on the resources in a dynamic view just as they do for standard views and snapshot views. For more information, see the AppManager Help.

What Kinds of Views Can You Create?

Here are some examples of the types of dynamic views you can create.

Using one rule, you can create a dynamic view that displays:

- SQL Server computers located in **London** or **Berlin**.
- SQL Server computers with a database named **SAP***
- Servers that do **not** have any **NT** objects discovered.

Using two rules, you can create a dynamic view that displays:

- SQL Server computers with **version 7.0** and **more than 512 MB** of memory.
- SQL Server computers that have a **QDB** database and **more than 512 MB** of memory.
- SQL Server computers located in **London** with **more than 512 MB** of memory.

Or you can exclude particular servers by configuring a separate exception query statement. For example, you can configure a view to display all SQL Servers except the SQL Server on which the AppManager repository database resides.

For reference information about the types of dynamic views that can be created, see [“Examples of Dynamic Views” on page 246](#).

Understanding How Rules Work

Using the Rule Wizard in the Operator Console, you can easily configure a view rule or group rule. Each rule contains a query statement that selects the objects you want to include and optionally, a different query statement that describes any objects you want to exclude from the view.

You can configure a query statement using **AND** or **OR** conditions. Use **AND** to select objects that meet all of the

specified conditions. Use **OR** to select objects that meet at least one of a set of conditions. The following condition types are available:

- **Object Name** Selects any objects that match the name you specify.
- **Object Type** Selects any objects of the type you specify.
- **Object Detail** Selects objects based on the object detail of the object type you specify.
- **Custom Property** Selects objects based on an existing custom property on at least one computer in your environment. Note that you must use the Control Center Console to add or change a custom property to a computer.

In general, you can configure a view to display the objects you want using 3 view rules or fewer. The Operator Console allows you to configure a view with up to 6 rules. Keep in mind that view rules are applied as a logical **AND**; an object on a computer must be selected by all rules in order to display that computer in the dynamic view. The order in which view rules are applied does not affect how the computers are displayed.

You can configure a dynamic view with up to 20 dynamic groups, each with its own group rule. If you create more than one group rule, the order in which group rules are applied determines the group that displays the server. At this time, you cannot organize groups within another group. See [“Organizing Computers into Groups” on page 237](#) for more information.

When configuring a rule to select objects, the query statement can include up to 4 conditions. An exception query statement can specify up to 20 conditions.

Note The computers that can be displayed in the TreeView pane are constrained by the “based on” view. For more information on the “based on” view, see [“Creating a Dynamic View” on page 232](#).

About the Object Name Condition

The Object Name condition selects objects matching the name you specify using an **IS** or **IS NOT** operator. Object names correspond to the discovered objects you see in the TreeView pane of the Operator Console.

This condition is useful when you want to select objects by name using pattern-matching characters. For example, you can specify **SJCExch*** to select any object whose name starts with **SJCExch**, such as **SJCExch01** and **SJCExch02**.

Comparison Operator	Description
IS	Selects object names that match the specified string value. For example, IS myhostname .
IS NOT	Selects object names that do not match the specified string value. For example, IS NOT SJCExch* selects any object name that does not begin with SJCExch , such as SQL Server:PROD01 .

When specified by itself, the Object Name condition selects any type of object that matches the specified name. Specify an Object Type condition and an Object Name condition (using an **AND**) to select a particular object type with the specified name.

Tip To filter object names that do **not** match a string value, you must configure separate object name conditions to select the object names from which you want to filter the object names you do not want. For example, to select object names that do not begin with **SJCExch**, configure a rule where:

```
Object Name is *  
AND Object Name is not SJCExch*
```

About the Object Type Condition

The Object Type condition selects objects of the type you specify using an **IS** or **IS NOT** operator. Discovered object

types appear as resource objects in the TreeView, for example, a SQL Server or a CPU resource. The object types you can select to configure this condition correspond to the objects in the AppManager repository—you can configure this condition to select an object type that has not been discovered but no objects will be selected until after they are discovered.

Operator	Description
IS	Selects only the specified object type. For example, IS SQLT_Server selects computers that have a discovered SQL Server.
IS NOT	Selects any object types that do not match the specified object type. For example, IS NOT SQLT_Server selects any discovered object types other than SQLT_Server, such as NT_MachineFolder. In this case, a computer that has any other discovered resources is displayed, including computers with SQL Server. Tip To avoid this problem, configure an exception query to exclude the objects you do not want.

When configuring an object type condition, select a parent or child object that corresponds to the object type you want. In the following example, the SQL Server object is selected and,

The object type for SQL Server is `SQLT_Server`.



About Selecting Object Types

When you configure an Object Type condition using the Rule Wizard, you can select a parent object type or child object type using **IS** or **IS NOT**:

- **Parent objects** represent basic resource information about an object and appear in the Rule Wizard with an “F” or “Folder” appended to the object name. For example, **Object Type is NT_MemoryFolder** describes Windows servers with a discovered memory resource.
- **Child objects** represent objects under a parent object in the TreeView pane of the Operator Console. For example, the **NT_MemoryFolder** parent object has two child objects, **NT_PhysicalMemObj** and **NT_VirtualMemObj** which represent physical memory and virtual memory, respectively.

In the Rule Wizard, the available object types have a slightly different name than objects in the Operator Console. For example, the objects for the SQL view appear under **SQLT** and the objects for the Netfinity view appear under **NETFIN**.

To select an object of a particular type and name, use two conditions to specify the Object Type condition and (using an **AND**) the Object Name condition. For example, to describe computers with at least 2 processors, create a rule that describes the **NT_CPUNumber** object type **AND** the object name is **1** (the CPU number of the first CPU on a computer is always **0**; the name of the second processor is always **1**).

To filter object types, configure an exception in the rule to select the objects from which you want to filter the object type you do not want. For example, to only display computers that are not Microsoft SQL Server computers, configure a rule where:

```
Object Name is *  
Except if Object Type is SQLT_Server
```

Query Limitations for Object Types

When configuring a query statement, you cannot use an AND to select more than one object type. For example, you **cannot** select SQL Servers with IIS using a single rule:

- Where Object Type is SQLT_Server
And Object Type is IIST_Server

This is a limitation of the query mechanism used in dynamic views.

To create a dynamic view that displays SQL Servers with IIS installed, you must create two view rules:

- Where Object Type is SQLT_Server
- Where Object Type is IIST_Server

There is no rule restriction for specifying multiple object types using an OR.

About the Object Detail Condition

The Object Detail condition compares the value of an object detail for a specified object type. An object detail corresponds to the detail columns that you see on the **Details** tab of the List pane. See [“Viewing Detailed Information about Discovered Resources” on page 60](#) for more information.

Tip You can view the details for an object by clicking the Details tab in the Operator Console and selecting the object you want from the TreeView pane.

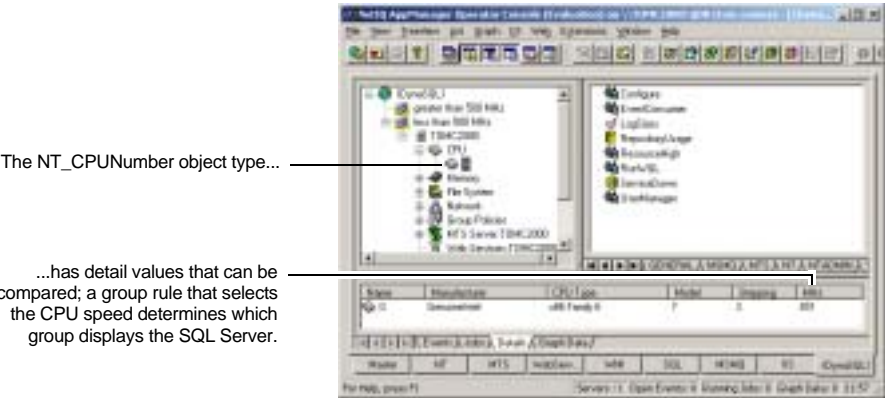
You can configure this condition to select an object by comparing its detail value using standard comparison operators:

Operator	Description
IS	Selects detail values that match the specified string value. For example, when comparing the SQL Server version (which is stored as a string value), IS 8* selects objects with a matching detail value such as 8.00.194 .
IS NOT	Selects detail values that do not match the specified string value. For example, when comparing the SQL Server version, IS NOT 8* selects computers with a discovered SQL Server version other than 8, such as 7.00.842 .
EQUAL TO	Selects detail values that match the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), EQUAL TO 451 selects objects with a detail value of 451 . Note that this operator can also be used to compare decimal values. NetIQ recommends that because decimal values are approximate, that you use comparison operators such as GREATER THAN or LESS THAN to compare decimal values.
NOT EQUAL TO	Selects detail values that do not match the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), NOT EQUAL TO 451 selects objects with a detail value that is greater than or less than 451. Note that because decimal values are approximate, NetIQ recommends that you use comparison operators such as GREATER THAN or LESS THAN to compare decimal values.
GREATER THAN	Selects detail values that are greater than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), GREATER THAN 451 selects objects with a detail value that is greater than 451 , such as 451.112 or 575.
LESS THAN	Selects detail values that are less than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), LESS THAN 451 selects objects with a detail value that is less than 451 , such as 450.112 or 375.

Operator	Description
NOT LESS THAN	Selects detail values that are not less than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), NOT LESS THAN 451 selects objects with a detail value that is not less than 451 , such as 451.112 or 575.
NOT GREATER THAN	Selects detail values that are not greater than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz), NOT GREATER THAN 451 selects objects with a detail value that is not greater than 451 , such as 450.112 or 375.

For example, you can configure a view rule to select SQL Servers with the repository database using an Object Type condition and (using an AND) an Object Name condition, and then create group rules to organize the computers into groups based on CPU speed using an Object Detail condition.

In the following example, there are two groups: **greater than 500 MHz** and **less than 500 MHz**.



In this example, the **NT_CPUNumber** object type has several object details, including Megahertz (MHz). The Object Detail condition allows you to compare the detail value of an object type. And, because the view is based on the **Master** view, all server resources on the computer are displayed.

About Selecting Object Details

When you configure an Object Detail condition using the Rule Wizard, you can select a parent object type or child object type—be sure to only select a child object type.

Do **not** select a parent object to configure an Object Detail condition; parent objects do not contain detail values. Keep in mind that parent objects usually appear in the Rule Wizard with an “F” or “Folder” appended to the object name. Note that the **NT_MachineFolder** and **UNIX_MachineFolder** objects are an exception to this rule; they include detail values and can be selected in an Object Detail condition.

To configure an Object Detail condition in the Rule Wizard, select a **child object**. For example, the **NT_MemoryFolder** parent object has two child objects, **NT_PhysicalMemObj** and **NT_VirtualMemObj** which describe the amount of physical memory and virtual memory, respectively. If, after you create a dynamic view, no objects are displayed, make sure that the correct object type is selected.

In the Rule Wizard, the available object types have a slightly different name than objects in the Operator Console. For example, the objects for the SQL view appear under **SQLT** and the objects for the Netfinity view appear under **NETFIN**.

To select detail information for a named object, use an AND condition to specify the name of the object you want. (Note that in some cases, detail information includes the object name.) For example, to describe computers with at least 2 processors and the processor is faster than 500 MHz, create a rule where:

```
Object NT_CPUNumber.Detail15[MHz] is at least 500
AND the object name is 1
```

Note that the object name of a second processor is always **1** (CPU objects are numbered beginning with **0**).

Tip To filter object details, configure separate rules to select the objects from which you want to filter the object detail you do not want. For example, to select computers with more than 512 MB of memory, configure 2 rules where:

- Object Name is *
- Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288

Query Limitations for Object Details

When configuring a query statement, you cannot use an AND to select more than one object type. This is a limitation of the query mechanism used in dynamic views.

For example, to select SQL Server computers with version 7 or version 8 and more than 512 MB of memory, create a dynamic view that is based on the **Master** view with two view rules where:

- Object SOLT_Server.Detail2[Version] is 7*
OR Object SOLT_Server.Detail2[Version] is 8*
- Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288

Note that there is no rule restriction for specifying multiple object types using an OR.

About the Custom Property Condition

The Custom Property condition compares the value of a custom property. To configure this condition, at least one computer in your environment must be configured with the custom property you want.

Note that to add a custom property to a computer, you must use the Control Center Console. It can take up to 2 minutes for a new custom property to appear in the Operator Console.

Tip You can view the custom property values for a computer by right-clicking the computer in the TreeView pane of the **Master** view or a snapshot view that is based on the **Master** view and clicking **Custom Properties**.

You can configure this condition to select a computer by comparing its custom property value using standard comparison operators:

Operator	Description
BEFORE	Selects a date value that is BEFORE the specified date of the custom property value. This data type supports date format only (datetime is not supported).
AFTER	Selects a date value that is AFTER the specified date of the custom property value. This data type supports date format only (datetime is not supported).
AT	Selects a date value that is AT (on) the specified date of the custom property value. This data type supports date format only (datetime is not supported).
IS	Selects custom property values that match the specified string value. For example, when comparing the Location custom property (which is stored as a string value), IS San Jose selects computers with a matching custom property value.
IS NOT	Selects custom property values that do not match the specified string value. For example, when comparing the Location custom property, IS NOT San Jose selects computers with a Location custom property value other than San Jose .
EQUAL TO	Selects custom property values that match the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, EQUAL TO 451 selects computers with a matching custom property value. Note that this operator can also be used to compare decimal values. NetIQ recommends that because decimal values are approximate, that you use comparison operators such as GREATER THAN or LESS THAN to compare decimal values.

Operator	Description
NOT EQUAL TO	<p>Selects custom property values that do not match the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, NOT EQUAL TO 451 selects computers with a custom property value that is greater than or less than 451.</p> <p>Note that because decimal values are approximate, NetIQ recommends that you use comparison operators such GREATER THAN or LESS THAN to compare decimal values.</p>
GREATER THAN	<p>Selects custom property values that are greater than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, GREATER THAN 451 selects objects with a detail value that is greater than 451, such as 451.112 or 575.</p>
LESS THAN	<p>Selects custom property values that are less than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, LESS THAN 451 selects objects with a detail value that is less than 451, such as 450.112 or 375.</p>
NOT LESS THAN	<p>Selects custom property values that are not less than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, NOT LESS THAN 451 selects computers with a detail value that is not less than 451, such as 451.112 or 575.</p>
NOT GREATER THAN	<p>Selects custom property values that are not greater than the specified integer or decimal value. For example, when comparing the CPU speed (in MHz) custom property, NOT GREATER THAN 451 selects computers with a detail value that is not greater than 451, such as 450.112 or 375.</p>

For example, if a **Location** custom property exists on each computer in your environment, you can organize computers into dynamic groups according to their **Location** value.

Tip To filter computers that do **not** match a custom property value, you must configure separate conditions to select the objects from which you want to filter the custom property values you do not want. For example, to select computers that are not located in San Jose, configure a rule where:

```
Object Name is *
Except if Custom Property Location is San Jose
```

To configure a condition to use a custom property, the custom property must exist on at least one computer in the TreeView pane. For more information, see [“Understanding Custom Properties” on page 230](#).

Related Topic

[Understanding How Rules Work](#)
[Examples of Dynamic Views](#)

Query Limitations for Custom Properties

When configuring a query statement, you cannot use an AND to select more than one custom property. For example, you **cannot** select **Production** computers located in **London** using a single rule:

- Custom Property Location is London
AND Custom Property SvrType is Production

This is a limitation of the query mechanism used in dynamic views.

To select **Production** computers located in **London**, create a dynamic view that is based on the **Master** view with two view rules, for example, where:

- Custom Property Location is London
- Custom Property SvrType is Production

Note The **Location** and **SvrType** custom properties must be defined on at least one server when the rule is created. Also, there is no rule restriction for selecting more than one custom property in a single rule using an OR condition.

Understanding Custom Properties

A custom property provides information about a computer. For example, you can use a custom property to identify servers by **Location**, such as **San Jose**. With this information, you can configure a dynamic view to manage servers by location.

If you want to organize and manage computers based on custom property information, NetIQ recommends that you use the Control Center Console to create a rule-based management group that selects computers based on custom property information.

To add or remove a custom property, you must use the Control Center Console. It can take up to 2 minutes for a new custom property to appear in the Operator Console. Note that any changes you make to the value of a custom property can affect any rule-based management groups or job override values that use the custom property.

In the Operator Console, you can view custom property information from the **Master** view or a snapshot view that is based on the **Master** view.

When deciding how to use custom properties, keep in mind that when you configure a view rule for a dynamic view, you can only select one custom property using an AND. With a maximum limit of 6 view rules in a view, you can configure a dynamic view to display computers that all have up to 6 matching custom properties.

To view the custom property values on a computer:

- 1 In the TreeView pane of the **Master** view or a snapshot view that is based on the **Master** view, right-click the computer you want and click the **Custom Properties...** button.
- 2 On the **Custom** tab of the Properties dialog box, a list of custom properties is displayed for the selected computer.
- 3 Click **OK** to close the Properties dialog box.

Creating a Dynamic View

A dynamic view is configured with at least one view rule. You can optionally configure up to 6 view rules to describe the contents of a dynamic view.

A dynamic view is based on a standard view or on the **Master** view; you cannot base a dynamic view on a snapshot view. The objects that can be displayed in the TreeView pane are constrained by the “based on” view. For example, if you create a dynamic view based on the **Master** view, and you create a rule that selects computers with a SQL Server database named **QDB**, the dynamic view displays resource information for SQL Servers with a **QDB** database and any other discovered resources on that computer. If the dynamic view is based on the **SQL** view, only SQL-related objects are displayed.



A dynamic view based on the **SQL** view only displays SQL-related resources for the computer in the TreeView pane...



...while a dynamic view based on the **Master** view displays all server resources for the computer in the TreeView pane.

Use the **Master** view as the “based on” view when you want to display all server resources on a computer; in the case of a dynamic view that is based on a standard view, only server information that corresponds to the “based on” view is displayed. Note that you can configure a view to select objects that do not correspond to the “based on” view but the view only displays server resources that correspond to the “based

on” view. After you create a dynamic view, you cannot change its “based on” view.

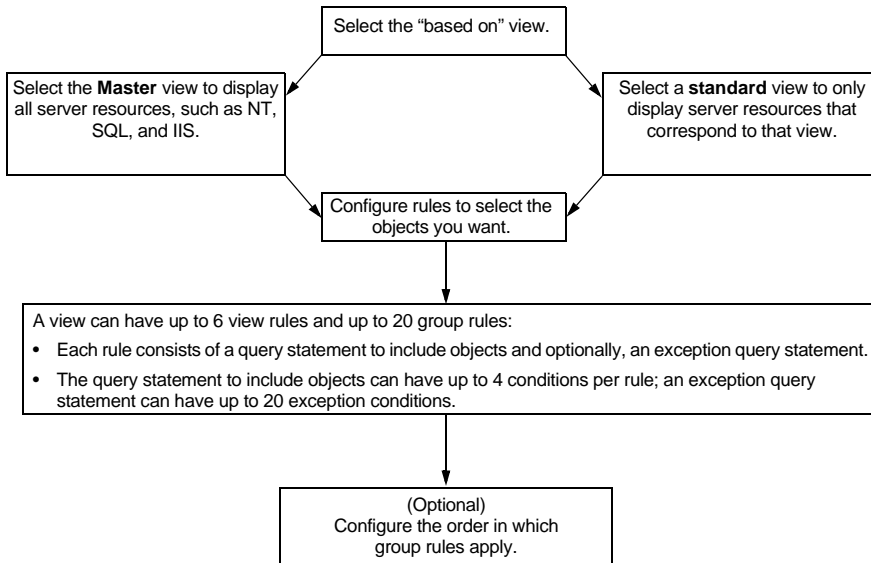
Create one or more view rules to select the objects you want to appear in the dynamic view. Keep in mind that view rules are applied as a logical AND; an object on a computer must be selected by all rules in order to display that computer in the dynamic view. The order in which view rules are applied does not affect how the computers are displayed.

After you configure one or more view rules to select the objects you want to include in the view, you can create group rules to organize the contents of the view into dynamic groups. At this time, you cannot organize groups within another group.

Using the AppManager Security Manager, you can apply AppManager permissions to enable users to view a particular dynamic view. Also, you can enable users to create or change dynamic views by configuring the **Create Custom Views** functional right. See the *Administrator Guide* for more information about configuring user permissions.

Overview: How to Create a Dynamic View

Here's an overview of the steps to create a dynamic view.



Creating a Dynamic View

When you create a dynamic view, you must create a view rule and then create additional view rules or group rules.

- 1 Open the View Manager by clicking **View > View Manager**.



- 2 In the View Manager, click the **Dynamic** button.
- 3 In the Create Dynamic View dialog box, type a name for the dynamic view.
- 4 Click the **Based on view** list to select a standard view or the **Master** view.

View	Description
The Master view	Displays all server resources for a computer.
A standard view	Only displays server resources that correspond to that view.

- 5 In the Create Dynamic View dialog box, click the **New View Rule** button to create a rule that selects the objects you want. See [“Examples of Dynamic Views” on page 246](#) for guidance on creating rules to select objects.
- 6 In the Rule Wizard: Create rule dialog box, you can create a rule based on an existing rule that you have created by selecting a rule from the **Select custom rules** list. Otherwise, you can create a new rule by clicking **Next**.
- 7 In the Rule Wizard: Select objects to include dialog box, click the **Add** button choose a condition that selects the objects you want to include. Then in the Add Rule Condition dialog box, select the type of condition you want and click **OK**.

For information about the available conditions, see [“Understanding How Rules Work” on page 217](#).

- 8 In the Rule Wizard: Select objects to include dialog box, you must click the condition to configure its value. In the Configuration dialog box, specify the value you want and click **OK**.

- 9 In the Rule Wizard: Select objects to include dialog box, you can configure additional conditions or click **Next**.

Task	Steps to Take
Add a condition using an AND statement	Click the AND button.
Add a condition using an OR statement	Click the OR button.
Change the order in which AND and OR statements apply	Select a condition and click the Move Up and Move Down buttons.
Remove a condition	Select a condition and click the Remove button.

- 10 In the Rule Wizard: Select objects to exclude dialog box, click the Add button to configure any exception to the include conditions you have previously specified, or click **Next**.

For information about the available conditions, see [“Understanding How Rules Work” on page 217](#).

- 11 In the Rule Wizard: Name the rule dialog box, specify a name for the rule and click **Finish**.

- 12 When you finish creating the view rule, the rule appears in the list of view rules. To see the rule description, make sure the rule you want is selected and view the description in the Rule Description field. Repeat these steps to add more view rules or to add and order group rules.

Task	Steps to Take
Add a view rule	Click the New View Rule button.
Add a group rule	Click the New Group Rule button.
Reorder group rules	Click a group rule and click Move Up or Move Down . Group rules apply in order from top to bottom. View rules apply as a logical AND statement.
Modify an existing rule	Click the rule you want and click Modify .
Rename or delete a rule	Click the rule you want and click Rename or Delete .

- 13 When you are finished, click **OK** to close this dialog box and display the dynamic view.

Changing a Dynamic View

- 1 Open the View Manager by clicking **View > View Manager**.
- 2 In the View Manager, select the view you want and click the **Modify...** button.
- 3 In the Modify Dynamic View dialog box, you can type a new name for the view. Note that you cannot change the “based on” view.
- 4 To see the description of a rule, select the rule you want from the view rule or group rule list. The current description of the rule is displayed in the **Rule Description** field.
- 5 To modify a view rule or group rule, click a rule in the list of view rules or group rules and click the **Modify...** button.
- 6 In the Rule Wizard, click **Next** to view and change the properties of the rule. For help at any time, click **Help**.

When you finish updating the rule, you can repeat these steps to create additional rules for the view or see the following sections for information on how to organize the contents of a dynamic view into dynamic groups and how to specify the order in which rules apply.

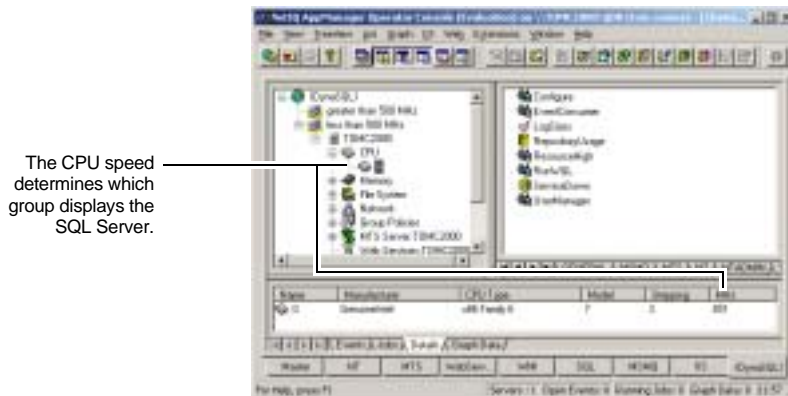
Organizing Computers into Groups

After you create at least one view rule to describe a view, you can create additional rules to organize the contents of a dynamic view into dynamic groups. A dynamic group consists of a single rule that selects objects from a dynamic view and displays corresponding computers in the group. For example, you can organize SQL Server computers with the **QDB** database into groups based on CPU speed.

You cannot manually organize the computers in a dynamic view into a standard server group. Use a group rule to select computers in the dynamic view and organize them into a dynamic group. At this time, you cannot organize groups within another group.

Tip To view the number of computers in a dynamic group, right-click the dynamic group in the TreeView pane of the Operator Console and click **Properties**. The **Server** tab displays the number of computers in the group.

In the following example, there are two groups: **greater than 500 MHz** and **less than 500 MHz**.



After you change the rules for a dynamic view or a dynamic group, it may take up to one minute to initiate an update of the view or group. Depending on the number of servers in your environment, it can take several minutes to update the entire contents of the view or group.

To simplify the process of creating groups, before you create a dynamic group, make sure a dynamic view displays the objects that you want.

When specifying more than one group rule in a view, keep in mind that a computer can only appear once within a view. If you set up more than one dynamic group, you must also

specify the order in which group rules apply. The first rule that selects an object on a computer displays the computer. You cannot organize groups within another group. See [“Understanding How Rules Work” on page 217](#) for more information.

You can apply a monitoring policy to groups in a dynamic view to automatically monitor matching resources using a pre-defined set of Knowledge Scripts. For information, see [Chapter 9, “Monitoring by Policy.”](#)

Creating a Group Rule

To create a group rule:

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify...** button.
- 2 In the Modify Dynamic View dialog box, click the **New Group Rule...** button.
- 3 In the Rule Wizard, follow the prompts to configure the rule. For help at any time, click **Help**.

When you finish creating the rule, you can repeat these steps to create additional dynamic groups or see [“Configuring the Order in which Group Rules Apply” on page 240](#) for information on how to specify the order in which group rules apply.

Changing a Group Rule

To change the objects that appear in a dynamic group, update the group rule or change the order in which group rules are applied.

Tip To view the number of computers in a dynamic group, right-click the dynamic group in the TreeView pane of the Operator Console and click **Properties**. The **Server** tab

displays the information about the selected group; you cannot use this dialog box to add or remove computers from the current group.

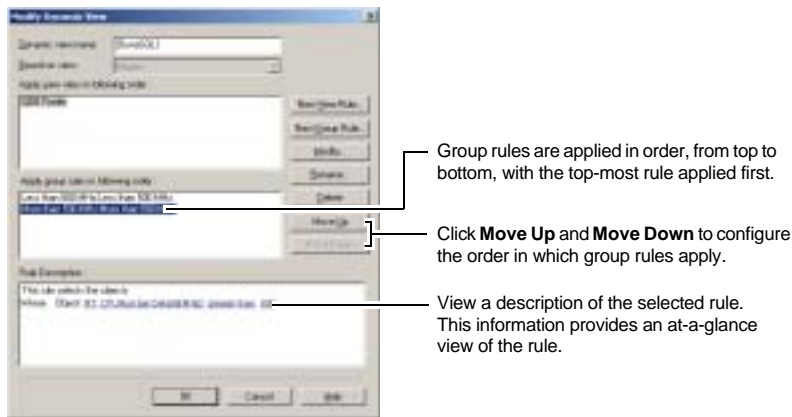
To change a group rule:

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify** button.
- 2 In the Modify Dynamic View dialog box, select a group rule from the list and click the **Modify** button.
- 3 In the Rule Wizard, follow the prompts to change the current configuration of the rule. For help at any time, click **Help**.

For information on specifying the order in which rules apply, see [“Configuring the Order in which Group Rules Apply” on page 240](#).

Configuring the Order in which Group Rules Apply

If you have created group rules, use the Dynamic View properties dialog box to configure the order in which group rules apply.



When specifying more than one group rule in a view, keep in mind that a computer can only appear once within a view. If

you configure more than one group rule, you must also specify the order in which group rules apply.

Group rules are applied in the order they are listed in the Dynamic View properties dialog box, from top to bottom. The first rule that describes an object on a computer displays the computer. At this time, you cannot organize groups within another group.

Note that after you change the rules for a dynamic view or group, it may take up to one minute to initiate an update of a dynamic view or group. Depending on the number of servers in your environment, it can take several minutes to update the entire contents of the view or group.

To change the order in which group rules apply:

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify...** button.
- 2 In the Modify Dynamic View dialog box, select a group rule from the list and click **Move Up** and **Move Down** to change the position of the rule in the list. Keep in mind that the rules are applied in the order that they appear in the list, from top to bottom.
- 3 Click **OK**.

Deleting a Rule

To delete a group rule:

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify...** button.
- 2 In the Modify Dynamic View dialog box, select a view rule or group rule and click **Delete**.

Note that a view must have at least one view rule; you cannot delete all the view rules in a dynamic view.

- 3 Click **OK** to apply your changes.

Note After you change the rules for a dynamic view, it may take up to one minute to initiate an update of a dynamic view; depending on the number of servers in your environment, it can take several minutes to update the entire contents of the view.

Renaming a Rule

You can rename a group rule to use the name you want.

Tip To rename a group, modify the group rule to change the group name. See [“Renaming a Group” on page 242](#) for more information.

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify...** button.
- 2 In the Modify Dynamic View dialog box, select a rule and click the **Rename...** button.
- 3 In the Rule Rename dialog box, type a new name and click **OK**.

Renaming a Group

To rename a dynamic group, you must modify the group rule using the Rule Wizard.

If you rename a group that is monitored by policy, the monitoring policy is removed and any associated policy-based jobs are stopped and deleted. For information about adding a monitoring policy to a group, see [“Creating a Monitoring Policy” on page 264](#).

- 1 Click **View > View Manager**, select an existing dynamic view, and click the **Modify...** button.

- 2 In the Modify Dynamic View dialog box, select a group rule and click the **Modify...** button.
- 3 In the Rule Wizard, follow the prompts to rename the group.

Tips and Tricks

Here are some tips for configuring dynamic views:

- Decide whether to base your view on a standard view (to only display server resources from that view) or the **Master** view (to display all server resources).
- When combining conditions using AND, you can only select one object type and one custom property. Note that an object type is selected when you configure an Object Type or Object Detail condition. To select more than one object type or custom property using AND, create a view rule for each object type or custom property you want to select.

There is no restriction for selecting multiple object types or custom properties using an OR. See [“Understanding How Rules Work” on page 217](#) for more information.

- Keep in mind that over time, as new objects are discovered, a dynamic view automatically identifies and displays matching objects; if a monitoring policy is associated with a dynamic view or group, the monitoring policy automatically monitors resources as the dynamic view changes.
- After you are comfortable with the dynamic view and any groups, apply a monitoring policy to a dynamic view or one of its groups to automatically monitor matching resources using a pre-defined set of Knowledge Scripts. For more information, see [Chapter 9, “Monitoring by Policy.”](#)
- When using arithmetical operators to describe a numeric value, keep in mind that some version information, such as

SQL Server version **8.00.194**, is a string value and must be described using string operators such as **IS** or **IS NOT**.

- The **NT_MachineFolder** and **UNIX_MachineFolder** objects include detail values and can be selected for an Object Detail condition.
- To describe a discovered server application, configure an Object Type or Object Detail condition using the **<category>_Server** object. For example, **SQLT_Server** describes any discovered SQL Server.
- To add a monitoring policy to a dynamic view, the view must be populated with objects that can be monitored by an available Knowledge Script Group. If the Knowledge Script Group that you want to add to a monitoring policy is not listed, make sure that the objects in the dynamic view correspond to the Knowledge Script Group you want.
- If you rename a dynamic group that is monitored by policy, the monitoring policy is removed and any associated policy-based jobs are stopped and deleted.
- To select an object type by name, be sure to select the object name that appears in the TreeView pane of the Operator Console, for example, **IIS Server:Tomc2000**. In some cases, object detail information includes the object name.

For example, the details for **SQLT_Server** object include the server name as the first detail (**Detail1[ServerName]**) while **IIS** does not include this information.

Troubleshooting a Dynamic View

Here are some tips for troubleshooting a dynamic view:

Problem	Suggestion
It takes up to one minute to display a dynamic view	<p>After you change the rules for a dynamic view, it may take some time to initiate an update of a dynamic view; depending on the number of servers in your environment, it can take several minutes to update the entire contents of the view.</p> <p>If after a minute, the objects are not displayed, verify that the SQL Server Agent service (SQLSERVERAGENT) is running on the repository database server. If SQLSERVERAGENT is not running, dynamic views will not work.</p>
After a minute or so, the dynamic view does not display any objects	<ul style="list-style-type: none">• After you change the rules for a dynamic view, it may take up to one minute to initiate an update of a dynamic view; depending on the number of servers in your environment, it can take several minutes to update the entire contents of the view.• Make sure the objects you want to display in the dynamic view appear in the view upon which the dynamic view is based.• If the objects exist in the “based on” view, verify that the query statements in each rule are defined properly. For information on viewing the configuration of a rule, see “Changing a Dynamic View” on page 237.
After 5 minutes, the dynamic view is not displayed	<p>AppManager attempts to display a dynamic view 5 times (once each minute). If a condition within a rule compares string values using an arithmetical operator (or vice versa), the view is not displayed and an event is created with a Job ID of -1. The event message is “{view_name} view failed after 5 attempts.”</p> <p>To view the event message, in the Master view, click the Events tab of the List pane.</p>
I’ve created a lot of dynamic views. Now the Operator Console is refreshing views slowly	<p>Click View > View Manager and deselect (to hide) the views you are not currently using and improve the overall refresh rate of the Operator Console.</p>

Problem	Suggestion
I created a query statement to describe a custom property, but the corresponding servers are not displayed	If the AppManager repository resides on a SQL Server that is case-sensitive, string values are case-sensitive. Check the name and value of the custom property and make sure the view rule matches.
I cannot add a custom property to a computer	To add a custom property, you must be in the Control Center Console. In general, NetIQ recommends that you use the Control Center Console to manage a dynamic group of servers.

Examples of Dynamic Views

Here are some examples of how you can use dynamic views:

View Contents	Steps to Take
SQL Server computers with version 7 or version 8 and more than 512 MB of memory	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object SQLT_Server.Detail2[Version] is 7* OR Object SQLT_Server.Detail2[Version] is 8* Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288
SQL Server computers that do not have a QDB database	Create a dynamic view that is based on either the SQL or Master view with one view rule where: <ul style="list-style-type: none"> Object Type is SQLT_DatabaseObj Exception is QDB
SQL Server computers that have a QDB database	Create a dynamic view that is based on the SQL view with one view rule where: <ul style="list-style-type: none"> Object Name is QDB
SQL Server computers that have a QDB database and more than 512 MB of memory	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object Type is SQLT_DatabaseObj And Object Name is QDB Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288
Computers that have both SQL Server and IIS	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object Type is SQLT_Server Object Type is IIST_Server

View Contents	Steps to Take
Computers that have either SQL Server or IIS	Create a dynamic view that is based on the Master view with one view rule where: <ul style="list-style-type: none"> Object Type is SOLT_Server OR Object Type is IIST_Server
SQL Server computers located in London with more than 512 MB of memory	Create a dynamic view that is based on the SQL view with one view rule where: <ul style="list-style-type: none"> Custom Property Location is London AND Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288
SQL Server computers on Netfinity servers	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object Type is NETFINI TYDIR_Server OR Object Type is NETFIN_Server Object Type is SOLT_Server
SQL Server computers named SJC* with more than 512 MB	Create a dynamic view that is based on the SQL view with 2 view rules where: <ul style="list-style-type: none"> Object Name is SJC* Object NT_PhysicalMemObj.Detail1[Size(KB)] is greater than 524288
SQL Server computers with a database named SAP*	Create a dynamic view that is based on the SQL view with one view rule where: <ul style="list-style-type: none"> Object Name is SAP* AND Object Type is SOLT_DatabaseObj
Servers that do not have any Windows objects discovered	Create a dynamic view that is based on the Master view with one view rule where: <ul style="list-style-type: none"> Object Name is * Except if Object Type is NT_CPUFolder <p>Note that this rule can display UNIX servers. To exclude UNIX servers, add another exception, for example:</p> <ul style="list-style-type: none"> Object Name is * Except if Object Type is NT_CPUFolder OR Object Type is UNIX_CPUFolder
Servers that have SQL Server discovered but not IIS Server	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object Type is SOLT_Server Object Type is not IIST_Server
Servers that have Windows discovered but not SQL Server	Create a dynamic view that is based on the Master view with two view rules where: <ul style="list-style-type: none"> Object Type is NT_CPUFolder Object Type is not SOLT_Server

Monitoring by Policy

This chapter describes how to initiate policy-based monitoring from the AppManager Operator Console.

If you plan to implement a monitoring policy, you should do so from the Control Center Console. Any monitoring policies that you implement in the Control Center Console cannot be managed from the Operator Console. The following topics are covered:

- [“About Policy-Based Monitoring,”](#) below
- [“Creating a Knowledge Script Group”](#) on page 255
- [“Changing a Knowledge Script Group”](#) on page 259
- [“Creating a Monitoring Policy”](#) on page 264
- [“Changing a Monitoring Policy”](#) on page 267

About Policy-Based Monitoring

A **monitoring policy** uses a set of pre-configured Knowledge Scripts to automatically monitor resources as they appear in the TreeView. A monitoring policy enables you to efficiently and consistently monitor all of the resources in your environment.

For example, you can create a monitoring policy to monitor physical disk resources for a particular group of computers. The monitoring policy monitors the discovered physical disk resources on each computer in the group using the same monitoring values. If the disk configuration for a managed computer changes (for example, a G: drive is added), the next

time physical disk resources are discovered, the monitoring policy automatically monitors the G: drive. If you add a computer to the group, the monitoring policy automatically monitors the discovered physical disk configuration. Of course, if you remove a physical disk from the managed client and rediscover resources (or delete the resource from the AppManager TreeView), the monitoring policy automatically stops monitoring the resource.

A monitoring policy is implemented through one or more Knowledge Script Groups. A **Knowledge Script Group** is a set of pre-configured Knowledge Scripts which you can run as part of a monitoring policy for a group or view, or as a set of **monitoring jobs** on a computer.

When using a Knowledge Script Group to run a set of monitoring jobs, you have the flexibility to stop and start a particular job. With a monitoring policy enabled, you cannot stop or change the properties of a particular policy-based job. (You can change the properties of all jobs in a policy or remove a computer from a policy.)

The advantage of using a Knowledge Script Group to monitor by policy is that matching resources are automatically monitored—a monitoring job started by Knowledge Script Group only monitors particular resources.

Additionally, a monitoring policy enforces the same monitoring values on all monitored resources. If you change the monitoring values for a monitoring policy, corresponding policy-based jobs are automatically updated to use the new monitoring values.

See the following sections for more information.

How Monitoring Policies Work

To better understand how a monitoring policy works, keep in mind the following:

- A monitoring policy is implemented at the group or view level in the TreeView pane through one or more Knowledge Script Groups. A **Knowledge Script Group** is a pre-configured set of Knowledge Script Group members. A **Knowledge Script Group member** is an instance of a Knowledge Script. For example, the NT_CPULoaded Knowledge Script can be a member of several Knowledge Script Groups, with the member in each Knowledge Script Group configured to use different monitoring values.
- A monitoring policy automatically creates policy-based jobs to monitor matching resources in the TreeView. A **policy-based job** is an instance of a Knowledge Script Group member running on a computer with a matching resource. Unlike a monitoring job, a policy-based job **cannot** be modified on the **Jobs** tab of the List pane. To stop, start, or modify a policy-based job, you must change the monitoring policy. Or you can delete or move a monitored resource from the policy-enabled view or group.
- **Properties propagation** automatically publishes changes to the monitoring policy to corresponding policy-based jobs. For example, if you change the properties of a Knowledge Script Group that is part of a monitoring policy, the changes are automatically propagated to create, remove, or change policy-based jobs.
- If a policy-based job encounters an error while attempting to monitor a resource, the monitoring policy attempts to restart the job. Only monitoring jobs can be restarted; discovery- and AppManager agent-related jobs are not automatically restarted. If a policy-based job cannot be restarted automatically, or if the policy-based job is scheduled to Run Once, when the job status is Stopped, you can manually restart the job.

- It is recommended that you implement a monitoring policy in the Control Center Console. If you implement a monitoring policy in the Operator Console, the policy is not visible from the Control Center Console. For example, if you use the Operator Console to configure a monitoring policy for the NT view, a management group with the NT view as a member does not display the monitoring policy.

Notes

- A job that is not part of a monitoring policy is a monitoring job. A **monitoring job** monitors a particular resource on a computer. Unlike policy-based jobs, you must manually propagate Knowledge Script properties to monitoring jobs. In addition, monitoring jobs do not automatically monitor resources that are discovered after the job starts.
 - The SQL Server Agent service (SQLSERVERAGENT) must be running on the repository database server. If SQLSERVERAGENT is not running, monitoring policies will not work.
-

How Knowledge Script Groups Work

If a job is started by a Knowledge Script Group, AppManager prepends the Knowledge Script job name with the identifier of the Knowledge Script Group that started the job. This information does not appear in the Jobs tab of the Operator Console but can be seen by viewing the job properties. For example, if you run

acmeNT_MemoryUtilizationByApplication as part of a monitoring policy, the Knowledge Script job name that appears in the Jobs tab of the Operator Console is truncated to appear as **acmeNT_MemoryUtilizationByApplicati**.

However, if you double-click the job to display its properties, the title bar of the Job Properties dialog box displays the Knowledge Script job name with the Knowledge Script Group

identifier, such as

1234:acmeNT_MemoryUtilizationByApplicati.

Tip To avoid truncating the Knowledge Script job name when running a custom Knowledge Script as part of a monitoring policy or as an ad hoc job started by a Knowledge Script Group, we recommend that you limit the length of the Knowledge Script name to 145 characters. The maximum length for a Knowledge Script is 150 characters.

Reporting Considerations

It is not necessary to run the same reports on more than one report agent so it is recommended that you **not** run reports as part of a monitoring policy. See [Chapter 7, “Running Reports,”](#) for more information.



Charting Considerations

To select the data streams you want, the Add Data wizard of the Chart Console lists the corresponding Knowledge Script job names. For policy-based jobs and ad hoc jobs started by a Knowledge Script Group, the Knowledge Script Group identifier is prepended to the Knowledge Script job name.

Viewing Policy-Based Jobs in the List Pane

Policy-based jobs appear on the **Jobs** tab of the List pane.



Icon	Description
	The job is part of a monitoring policy and as such, its job properties cannot be modified from the Jobs tab. Unlike monitoring jobs, you cannot start and stop a particular policy-based job from the Jobs tab of the List pane—the Job properties for a policy-based job are disabled.
	The ad hoc (monitoring) job was started by a Knowledge Script Group. You can modify its job properties from the Jobs tab.

Warning The Operator Console displays policy-based jobs created by both the Operator Console and the Control Center Console. However, you cannot use the Operator Console to remove a monitoring policy that was created by the Control Center Console. The Operator Console allows you to modify the policy and remove the appropriate Knowledge Script Group, however, the changes are restored immediately. Always use the Control Center Console to modify monitoring policies that were created in the Control Center Console.

The Jobs tab does not display the Knowledge Script Group identifier for policy-based jobs or ad hoc jobs started by a Knowledge Script Group. See [“How Knowledge Script Groups Work” on page 252](#) for more information.

The following sections discuss how to create monitoring policies and propagate refinements in your monitoring values.

Creating a Knowledge Script Group


You'll want to create a Knowledge Script Group to configure a monitoring policy or to create a group of monitoring jobs on a particular resource. By default, a Knowledge Script Group appears on the **KSG** tab of the Knowledge Script pane.

A Knowledge Script Group can be run in the **Master** view and any views that can be monitored by a Knowledge Script Group member:

View for Knowledge Script Group	Knowledge Script Categories that Can Be Included
The Master view	Any Knowledge Script category. To discover resources, you must create the Knowledge Script Group in the Master view or a snapshot view based on the Master view. To install the AppManager agent, you must create the Knowledge Script Group in the Master view or the NT view.
A standard view	Corresponding Knowledge Script categories. For example, if you create a Knowledge Script Group in the WMI view, you can add Knowledge Script Group members from the WMI and General categories.
A snapshot view or a dynamic view	The "based on" view. If the view is based on the Master view, all Knowledge Scripts are available. For example, if you create a Knowledge Script Group in a snapshot view that is based on the WMI view, you can add Knowledge Script Group members from the WMI and General categories.

- 1 In the Operator Console, click **KS > Create Knowledge Script Group**.

- 2 In the Properties for Knowledge Script Group dialog box, type a Knowledge Script Group name and description.



Note By default, the Knowledge Script Group name is prefixed with **KSG_** which means the Knowledge Script Group appears in the KSG tab of the Knowledge Script pane. To create the Knowledge Script Group in another category, prefix the name of the Knowledge Script Group using **Category_**. For example, to create a Knowledge Script Group named **NT Basics** in the NT category, name the Knowledge Script Group **NT_NT Basics**.

- 3 Click an option to select the type of Knowledge Script you want to include:

Option	Description
Regular KS	Monitors discovered resources.
Discovery KS	Discovers resources.
Install KS	Installs or updates AppManager agent software.

- 4 In the left column, click to expand the list of Knowledge Script categories; select one or more Knowledge Scripts and then click **Add**.

The Knowledge Scripts appear in the Members of the Knowledge Script Group column.

- 5 To view or update the properties for each member of the Knowledge Script Group, click a Knowledge Script in the Member column and click **Properties**. In the Job properties dialog box, view or change the monitoring values and then click **OK**.
- 6 When you finish, click **OK**.

Viewing Knowledge Script Groups

By default, a new Knowledge Script Group appears on the **KSG** tab of the Knowledge Script pane.



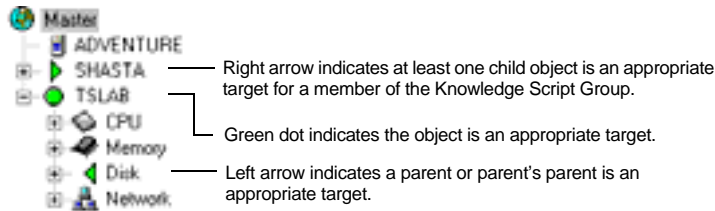
Creating a Group of Monitoring Jobs

Use a Knowledge Script Group to create a group of monitoring jobs on a selected computer. The resulting jobs are associated with the Knowledge Script Group but are not part of a monitoring policy.

To create a group of monitoring jobs:

- 1 In the Knowledge Script pane, click a tab to find the Knowledge Script Group you want and drag the Knowledge Script Group to the TreeView pane.
 - To run the set of jobs on a specific computer, drop the Knowledge Script Group only on that computer.
 - To run the set of jobs on multiple computers, drop the Knowledge Script Group on an appropriate group, on the Master level, or on several computers, individually.

The icon for the resource object changes to let you know you can drop the Knowledge Scripts on the selected target.



- 2 If you have selected at least one valid target, the Properties for Knowledge Script Group dialog box displays a list of all members in the Knowledge Script Group.

All members of the Knowledge Script Group appear here.



Note The list of members is not constrained by matching resource objects; there may be members in the list that do not match any resources.

- 3 To view or change the objects that are monitored by a Knowledge Script, click a member Knowledge Script in the list and click **Objects**.
- 4 In the Properties dialog box, a list of available objects is displayed. Note that by default, all matching objects are monitored.

Expand the list. Select (to monitor) or deselect (not monitor)


an object.



- 5 Click **OK** to close the Properties dialog box.
- 6 In the Properties for Knowledge Script Group dialog box, click **OK** to start monitoring jobs on matching resources.

Note For information on configuring the AppManager repository preferences to set the number of times to restart a monitoring job with an Error status, see [“Set Preferences for Audit Information and Restart Attempts”](#) on page 83.

Viewing Monitoring Jobs Started by a Knowledge Script Group

 This icon, which only appears on the **Jobs** tab of the List pane, means that the job is a monitoring job and was started by a Knowledge Script Group.

Changing a Knowledge Script Group

You can update a Knowledge Script Group to add or remove Knowledge Script Group members and to change the job properties of an existing group member.

When you update the job properties of a Knowledge Script Group member, any changes you make are automatically propagated to corresponding policy-based jobs. You must manually propagate changes to jobs that are not part of a monitoring policy.

If you add a member to a Knowledge Script Group, it may take up to one minute to propagate your changes to corresponding policy-based jobs.

To make changes to a Knowledge Script Group:

- 1 In the Knowledge Script pane, click a tab to find the Knowledge Script Group you want, click a Knowledge Script Group, and then click **KS > Properties**.

By default, a Knowledge Script Group appears on the **KSG** tab of the Knowledge Script pane.

- 2 In the Knowledge Script Group dialog box, you can change the description text and update the properties for each Knowledge Script Group member:

Task	Steps to Take
Remove a Knowledge Script Group member	<p>Click a Knowledge Script in the Members list and click Remove.</p> <p>When you close the dialog box, any associated policy-based jobs based on that member are stopped and deleted.</p>
Add a Knowledge Script Group member	<p>Click a Knowledge Script in the left-hand list and click Add.</p> <p>When you close the dialog box, any groups or views that use the Knowledge Script Group to monitor by policy automatically monitor corresponding resources with the new Knowledge Script.</p>

Task	Steps to Take
Change the monitoring values for a Knowledge Script Group member	<p>In the Members list, click a Knowledge Script and click Properties to display the Job Properties dialog box. Specify the monitoring values and click OK.</p> <p>For information on changing the job properties, see “Customizing Knowledge Scripts” on page 130.</p> <p>When you close the dialog box, any groups or views that use the Knowledge Script Group to monitor by policy automatically monitor corresponding resources with the updated Knowledge Script.</p>

- 3 Click **OK**.

Propagating Knowledge Script Group Members to Running Jobs or Reports

You can propagate the Knowledge Script properties of a Knowledge Script Group member to running jobs or reports that were started by that Knowledge Script Group member. Any corresponding jobs or reports are stopped and restarted with the Knowledge Script properties.

For information on propagating changes to a particular Knowledge Script to running jobs or reports, see [“Propagating Knowledge Script properties to Running Jobs or Reports” on page 132](#).

To propagate properties from a Knowledge Script Group member to running jobs or reports:

- 1 In the Knowledge Script pane, click a tab to find the Knowledge Script Group you want, click a Knowledge Script Group, and then click **KS > Properties Propagation > Ad Hoc Jobs**.

- 2 In the Properties Propagation dialog box, the selected members are to receive the transferred properties. Deselect a member to exclude it from properties propagation.

Deselect the members that you do not want to propagate.



- 3 By default, the job properties (specified in the Schedule, Values, Actions, and Advanced tabs) and the Knowledge Script itself (specified in the KPS section of the Knowledge Script) are propagated. Deselect an option to specify how to propagate Knowledge Script properties for the selected members:

Option	Description
Script	Propagates the logic of the Knowledge Script.
Properties	Propagates the changed monitoring values (including the schedule, actions, and Advanced properties).

Note You must select at least one option.

- 4 When you finish, click **OK**. Any monitoring jobs started by a Knowledge Script Group member are restarted with the job properties of the Knowledge Script Group member.

Copying a Knowledge Script Group

Copy a Knowledge Script Group to create a new Knowledge Script Group with a copy of each Knowledge Script Group member.

To avoid truncating the Knowledge Script job name when running a custom Knowledge Script as part of a monitoring policy or as an ad hoc job started by a Knowledge Script Group, we recommend that you limit the length of the Knowledge Script name to 145 characters. The maximum length for a Knowledge Script is 150 characters.

When you copy a Knowledge Script Group, a new Knowledge Script Group identifier is prepended to any Knowledge Script jobs created by the Knowledge Script Group. See [“How Knowledge Script Groups Work” on page 252](#) for more information.

To copy a Knowledge Script Group:

- 1 In the Knowledge Script pane, click a tab to find the Knowledge Script Group you want, click a Knowledge Script Group, and then click **KS > Copy Knowledge Script**.

Note By default, a Knowledge Script Group appears on the **KSG** tab of the Knowledge Script pane.

- 2 In the Copy Knowledge Script dialog box, enter the new name for the Knowledge Script Group.

Note To copy the Knowledge Script Group into another tab, change the prefix (category_) of the name. For example, to copy a Knowledge Script Group from the **KSG** tab into a new tab named **KSG 2000**, prefix the name with

KSG 2000_.



- 3 When you finish, click **OK**.

Deleting a Knowledge Script Group

Knowledge Script Groups can easily be deleted unless they are set up to run as part of a monitoring policy.

To delete a Knowledge Script Group:

- 1 In the Knowledge Script pane, click a tab to find the Knowledge Script Group you want to delete.
- 2 Click a Knowledge Script Group to delete.
- 3 Click **KS > Delete Knowledge Script**.

Note To delete a Knowledge Script Group that is configured to run as part of a monitoring policy, remove the Knowledge Script Group from any monitoring policies and then delete it. For information on removing a Knowledge Script Group from a monitoring policy, see [“Changing a Monitoring Policy” on page 267](#).

Creating a Monitoring Policy

Implementing a monitoring policy on a dynamic view provides a powerful and flexible way to monitor the resources in your environment. You can create a monitoring policy that both automatically discovers and monitor resources in a view

or a group. See [“Using Dynamic Views” on page 215](#) for more information.

When you create a monitoring policy, all matching objects to which the policy is applied are automatically monitored. See [“Deleting an Object from the TreeView Pane” on page 55](#) for information on deleting a particular resource and preventing it from being automatically rediscovered (and monitored by policy).

You can configure a monitoring policy to automatically discover and monitor resources by adding separate Knowledge Script Groups to:

- Discover resources. The Knowledge Script Group members should be set to run on a scheduled basis.
- Monitor resources. As the discovered resources change, the monitoring policy automatically monitors the updated resources.

When you create a monitoring policy, it may take up to one minute for the policy-based jobs to start. If you subsequently add or remove computers, it may take up to one minute for policy-based jobs to start or stop.



This icon, which only appears when a policy-based job is displayed, means that the job is part of a monitoring policy. As such, its job properties cannot be modified from the List pane.

Note A monitoring policy is implemented through Knowledge Script Groups. You cannot use a monitoring policy to install the AppManager agent or run reports; a Knowledge Script Group that was created with the **Install KS** option cannot be added to a monitoring policy. See [“Creating a Knowledge Script Group” on page 255](#) for more information.

To create a monitoring policy:

- 1 In the TreeView, click the top-level icon to select the view

or click to select a server group and then click **TreeView > Monitoring Policy**.



- 2 On the **Monitoring Policy** tab, click a Knowledge Script Group in the **Available** column and then click **Add**. To remove a monitoring policy, click a Knowledge Script Group in the **Selected** column and click **Remove**. Or, click **Remove All** to remove all Knowledge Script Groups.

From the Operator Console, you cannot remove a Knowledge Script Group that was added in the Control Center Console. If you select a Knowledge Script Group that was added from the Control Center Console, the **Remove** and **Remove All** buttons are enabled but does not work.

- 3 You can configure the number of times that the monitoring policy should attempt to restart policy-based monitoring jobs. Choose one of the following options:

Option	Description
Always	<p>Restarts a policy-based monitoring job until its job status changes to Running.</p> <p>Note Only monitoring jobs can be restarted by the monitoring policy. Discovery- and AppManager agent-related jobs are not restarted automatically by a monitoring policy.</p>
Number	<p>Specifies a number of retry attempts. If the specified number of retry attempts fails to restart the job, the job status remains Error.</p>

Note If a policy-based job cannot be restarted automatically,

or if the policy-based job is scheduled to Run Once, when the job status is Stopped, you can manually restart the job. See [“Changing Job Status” on page 142](#) for more information.

- 4 When you click **OK**, the monitoring policy creates policy-based jobs to monitor matching resources.

Changing a Monitoring Policy

You can change a monitoring policy by adding or removing Knowledge Script Groups. The monitoring policy automatically starts new policy-based jobs or stops existing policy-based jobs. It may take up to one minute for changes to a monitoring policy to take effect when:

- Adding a Knowledge Script Group to a monitoring policy
- Adding a Knowledge Script Group member to a Knowledge Script Group that is part of a monitoring policy

To change the job properties for an existing policy-based job, update the job properties of the corresponding Knowledge Script Group member. The changes are automatically propagated to all corresponding policy-based jobs. For more information on changing Knowledge Script Groups, see [“Changing a Knowledge Script Group” on page 259](#).

To stop policy-based jobs on a particular computer, you must move the computer out of a policy-enabled group, remove the policy, or delete the computer from a policy-enabled view.

To change a monitoring policy:

- 1 In the TreeView, select the view or server group to which the policy is applied and click **TreeView > Monitoring Policy**.

- 2 On the **Monitoring Policy** tab, add or remove a Knowledge Script Group to change the monitoring policy:

Task	Steps to Take
Remove a Knowledge Script Group	<p>Click a Knowledge Script Group in the Selected list and click Remove.</p> <p>Notes</p> <ul style="list-style-type: none">• When you close the dialog box, any associated policy-based jobs based on a member of the Knowledge Script Group are stopped and deleted.• See “Set Preferences for Data Management” on page 76 and “Set Preferences for Event Management” on page 79 for information on configuring the repository to automatically delete events and data when you delete a job.
Add a Knowledge Script Group	<p>Click a Knowledge Script Group in the Available list and click Add.</p> <p>Note When you close the dialog box, any matching resources in the view or server groups to which the policy is applied are automatically monitored by policy.</p>
Change the number of times that AppManager should attempt to restart the policy-based job	<p>Click to select an option:</p> <ul style="list-style-type: none">• Always Restart the error job until its job status changes to Running.• Number Specify the number of retry attempts. If the specified number of retry attempts fails to restart the job, the job status remains Error. <p>Note Only monitoring jobs can be restarted by the monitoring policy. Discovery- and AppManager agent-related jobs are not restarted automatically.</p>

- 3 Click **OK** to close the dialog box and update the monitoring policy.

Index

A

- acknowledge and close, preference for confirmation 66
- actions
 - automating 14
 - choosing where to run from 118
 - collecting audit information 84
 - creating schedules 87
 - event property 163
 - initiating 115
 - introduction 17
 - running based on number of duplicates events 119
 - viewing audit information 101
 - when monitoring UNIX computers 117
 - when monitoring Windows computers 116
- adding resources, to a snapshot view 212
- agent services 18
- alarms 155
- applications
 - ensuring availability 13
 - monitoring user-defined 16
 - supported 15
- AppManager
 - architecture
 - description 17
 - multi-tier 15
 - communication 19
 - extended support x
 - features 15
 - installation 15
 - introduction 14
 - monitoring your environment 13
 - operation 16
 - Operator Web Console 18
 - product suite 15
 - registering 30
 - supported products x
 - working off-line 29
- AppManager agent
 - introduction 18
 - purpose 22
- AppManager console
 - introduction 17
 - See also* Operator Console
- AppManager repository
 - data streams 187
 - how used 24
 - introduction 18
 - logging on 28
 - preferences 69
 - switching
 - Operator Console 35

asynchronous, schedule for jobs
112

audit information
collecting 84
viewing 101

automatic discovery 42

B

background view 67
blocking jobs, to perform
maintenance 145
business rules 17

C

Chart Console
selecting data 253
check in, a Knowledge Script 136
check out, a Knowledge Script 136
child events 159
closed jobs 138
closing events
automatically 180
based on event condition 125
based on event severity level 81
COM/OLE objects 23
components
communication 19
required 17
Web management server 18
configuration details
system inventory 14
viewing 60
Control Center 17, 25
conventions, documentation viii
copying, a Knowledge Script 134

Custom Property condition 227
introducing 227
query limitations 230

D

data collection
for a particular job 128
graphing 185
introduction 17
data points
averaging 76
defined 186
detail message 201
maximum number 187
data streams
combining in a graph 198
defined 186
deleting 205
displaying in graphs 187
filtering 191
grouping 189
information displayed 188
list of 188
multiple
creating a graph with 199
properties 192
view parent job 188
viewing in Chart Console 151
viewing in Graph pane 151
date and time, display 66
default threshold values 113

- deleting
 - a Knowledge Script Group 264
 - a particular resource 56
 - a view 58
 - an application server or workstation 56
 - preference for confirmation 66
 - preventing rediscovery 56
- Details tab 60
- details, dimming 68
- diagnose, selected computer 97
- Diagnostic Console, accessing from
 - Operator Console 97
- diagnostic events 155
- dimmed computer 92
- discovery
 - automatic 42
 - deleted objects with no rediscovery 57
 - displaying resources in a snapshot view 213
 - introduction 42
 - preventing 56
- discovery status, for a managed client 47
- documentation
 - conventions viii
 - online Help viii
- duplicate events 174
- dynamic view
 - changing a group rule 239
 - changing a view rule 237
 - creating 232
 - creating a group rule 239
 - creating a view rule 234
 - Custom Property condition 227
 - dynamic view *continued*
 - deleting a rule 241
 - design considerations 218
 - examples of 217
 - grouping computers 237
 - introducing condition types 218
 - introduction 216
 - limiting objects to display 232
 - Object Detail condition 223
 - Object Name condition 219
 - Object Type condition 219
 - ordering group rules 240
 - renaming a group 242
 - renaming a rule 242
 - sample rules 246
 - steps to create 234
 - tips for configuring 243
 - tips for troubleshooting 245
 - understanding how rules work 217
- E**
 - event collapsing
 - introduction 174
 - repository preference 71
 - setting options globally 179
 - setting options per job 123
 - events
 - acknowledging 170
 - actions initiated 115, 163
 - adding comments 166
 - automatically closing based on
 - event condition 125, 180, 183
 - closing 172

events *continued*

- closing based on event severity
 - level 81
- collapsing
 - global options 179
 - introduction 174
 - per job options 123
- collapsing time interval 176
- collecting audit information 84
- comments
 - adding 166
 - viewing 164
- deleting 174
- detailed message 162
- dimming 68
- duplicates 174
- expand and collapse 160
- filtering 170
- filtering by number of days 66
- filtering temporary conditions 175
- grouping 168
- how displayed 154
- how triggered 153
- ignoring 121
- information displayed 157
- initiating an action for duplicates 119
- introduction 17
- list of 157
- parent-child relationship 159
- prompting to acknowledge and close in Operator Console 66
- properties 161
- responding to 170

events *continued*

- severity levels
 - defined 155
 - flash icon on Taskbar 65
 - flash warnings 64
 - indicators 155
 - modifying 156
 - sorting 160
 - start collecting on 76
 - view parent job 150, 158
 - viewing
 - actions 163
 - comments 164
 - messages 162
 - viewing audit information 102
 - viewing by status 159
- Extensions menu
- Operator Console 95, 96

F

- filtering
 - data streams 191
 - events 170
 - job criteria 148
 - resource objects 63
 - temporary event conditions 175
 - to avoid duplicate events 175
 - user-defined criteria 63
- finding objects, in the TreeView 53
- flash warnings 64, 65
- foreground view 67

G

- graph data
 - See* data streams
- Graph Data tab, changing layout 189

- Graph pane
 - tabs 198
 - viewing graphs 198
- graphs
 - collecting data for 185
 - creating from data streams 193
 - default style 202
 - deleting 205
 - detail message 201
 - dimming 68
 - exporting 204
 - importing 204
 - modifying 196
 - preferences 197
 - printing 203
 - saving 67, 203
 - switching between 198
 - tiling 198
- grayed out computer 92
- group rule
 - changing 239
 - configuring the order to apply 240
 - creating 239
 - renaming 242
- grouping
 - data streams 189
 - events 168
 - jobs 147
 - user-defined criteria 63
- groups
 - adding 49
 - changing 51
 - computers and servers 48
 - folder objects 58
 - removing 51
- groups *continued*
 - renaming 51
 - viewing the number of servers 51
- H**
 - historical data 14
- I**
 - icons
 - grayed out computer 92
 - ignoring, events within an interval 121
 - indicators 155
 - informational events 155
 - inventory 14
- J**
 - job confirmation 65
 - jobs
 - adding a child job to an existing parent 109
 - checking status 138
 - child 109
 - closed 138
 - collecting audit information 84
 - default values 108
 - dimming 68
 - errors 138
 - expand and collapse 109
 - filtering 148
 - grouping 147
 - how to create 104
 - icon displayed 138
 - introduction 19
 - Knowledge Script Group identifier 252

jobs continued

- number created 109
- parent 109
- pending 137
- prompting to start, stop or close 65
- properties 107, 258
- restarting 143
- running
 - changing parameters 108
 - icon 138
 - multiple 109
 - status 137
 - subset of objects 120
- schedule
 - daily schedule 112
 - end 112
 - frequency 113
 - monthly schedule 112
 - regular intervals 111
 - run once 111
 - setting 110
 - start 112
 - weekly schedule 112
 - x number of times 111
- See also* policy-based jobs
- setting parameters 113
- sorting 142
- specifying corrective actions 115
- starting 21
- starting a new job 107
- status
 - changing 142
 - descriptions 137
- stopped 137

jobs continued

- triggering events 153
- view parent event 150
- viewing actual script 140
- viewing audit information 101
- viewing by status 141
- viewing results 149
- Jobs tab, changing the layout 140

K

- Knowledge Script Groups
 - changing properties 260
 - creating 255
 - deleting 264
 - identifier 252
 - propagating changes to jobs and reports 261
 - viewing in List pane 257
- Knowledge Scripts
 - categories and views 33, 104
 - checking in 136
 - checking out 136
 - copying 134
 - copying to create new 134
 - data streams 187
 - defaults
 - changing 22, 130
 - checking in 136
 - copying 134
 - unaffected 108
 - viewing 107
 - detailed messages 201
 - events 153
 - introduction 16
 - missing 104
 - parameters 154

Knowledge Scripts *continued*
 propagating changes to jobs
 and reports 132
 properties 108
 showing descriptions 92
 starting jobs 19, 21
 starting monitoring jobs 107
 type matching 105
 use of managed objects 23
 version history 92

L

List pane
 configuration details 60
 data streams 186, 188
 events 157
 filtering 63
 job status 139
 sorting 142, 160
local cache, preference setting 65
logon
 Operator Console 27

M

maintenance
 initiating on a computer 145
 viewing audit information 102
managed clients
 identifying 42
 introduction 22
managed objects, introduction 15,
 23
management server
 how used 25
 introduction 18
 Windows service 24

Microsoft SQL Server, AppManager
 repository 18, 25

Microsoft Windows
 tools

 Computer Management
 console 96
 Event Viewer 96
 Performance Monitor 96

monitoring functions 16
monitoring job status, changing
 142

monitoring jobs
 changing properties 144
 closing 151
 creating a group of 257
 deleting 152
 global preference to restart 85
 how created 107
 introduction 107
 started by a Knowledge Script
 Group 259
 starting a new job 108
monitoring policy
 changing 267
 changing a group 51
 creating 264
 how it works 250
 introduction 249

N

NetIQ AppManager Operator
 Console
 See Operator Console
NetIQccm 18, 22
NetIQCtrl, starting from the
 Operator Console 96
NetIQmc 18, 22

NetIQms 24
nqmagt 22

O

Object Detail condition
 query limitations 227
 selecting objects 226
 understanding 223
Object Name condition,
 introducing 219
Object Type condition
 query limitations 223
 selecting objects 222
 understanding 219
Objects tab 121
online Help viii
online help
 toolbar commands 91
operation management 13
Operator Console
 adding computers 43
 changing layout 38
 changing the layout of 68
 extensions 95
 indicators 155
 introduction 31
 logging on 27
 master view 33
 multiple users 29
 online help 29
 panes 32
 preferences 63, 64, 90
 running separate instances 29
 saving SQL passwords 99
 sharing data 25
 starting 27
 startup performance 68

Operator Console *continued*

 status bar 91
 supported AppManager envi-
 ronment 29
 switching to another repository
 35
 tip of the day 30
 toolbars 91
 tools
 adding 98
 Computer Management
 console 96
 editing 99
 Performance Monitor 96
 saving SQL passwords 99
 Windows Event Viewer 96
TreeView pane
 details panel 38
 grayed out computer 92
 split pane 38
updates 25
users, multiple 29
views 33
working off-line 29
Operator Web Console
 introduction 18

P

panes
 hiding 39
 moving 39
 overview 32
 reset to default 39
 resizing 39
 showing 39

- parent events
 - hiding 160
 - hiding in the List pane 68
- parent jobs
 - adding a child job 109
 - hiding 140
 - hiding in the List pane 68
- pending jobs 137
- Performance Monitor 96
- performance monitoring 14
- performance, improving Operator Console startup 68
- policy
 - See* monitoring policy
- policy-based jobs
 - introduction 251
 - viewing in List pane 254
- preferences
 - Operator Console
 - beep for new events 65
 - clear cache 65
 - customize layout 68
 - display date and times 66
 - flash icon on Taskbar 65
 - flash warnings 64
 - graph options 67
 - introduction 64, 90
 - local cache 65
 - prompt to acknowledge and close events 66
 - prompt to delete 66
 - prompt to start, stop or close jobs 65
 - tip of the day 64
 - TreeView pane
 - details panel 65
 - continued*
 - Operator Console *continued*
 - split pane 65
 - tooltip for TreeView details 65
 - View Manager settings 66
 - view updates 67
 - repository
 - archiving events 80
 - audit information 84
 - automatically closing an event based on event condition 74, 75
 - averaging data points 76
 - calculating average data stream value 78
 - closing events
 - automatically 81
 - data for graphs 77
 - data options 75
 - collecting data 75
 - deleting events
 - automatically 82
 - deleting graph data
 - automatically 79
 - displaying child events with event collapsing 81
 - event suppression 174
 - ignoring duplicate events 73
 - purging archived events 81
 - restarting monitoring jobs 85
 - prompt to start, stop or close jobs 65

- propagating
 - Knowledge Script properties to jobs and reports 132
- properties of a Knowledge
 - Script Group member to jobs and reports 261
- properties of a Knowledge
 - Script or a report script to Knowledge Script Group members 133

- properties
 - data streams 192
 - events 161
 - graphs 197
 - job 107, 258

R

- Report Viewer, starting 96
- reports
 - propagating Knowledge Script changes 132, 133
 - propagating Knowledge Script Group changes 261
 - viewing 96
- Repository Browser, starting from the Operator Console 96
- resource objects
 - defined 58
 - deleting from the TreeView 57
 - detailed information 60
 - filtering based on 63
 - inventory 14
 - type matching 105
- running jobs 137

S

- schedule
 - for a job 110
 - for an action 116
- scheduled maintenance,
 - AMADMIN_SchedMaint 145
- searching, in the TreeView 53
- services, client (NetIQmc, NetIQccm) 18
- severe events 155
- snapshot view
 - adding discovered resources 211
 - creating 209
 - introduction 208
 - viewing in List pane 211
- sorting child events, hiding parents 160
- sorting child jobs, hiding parents 140
- sound 65
- SQL password, saving 99
- status bar 91
- stopped jobs 137
- supported products x

T

- thresholds 113
- tip of the day 30, 64
- toolbars
 - available 91

tools

- adding 98
- Computer Management console 96
- editing 99
- Event Viewer 96
- Operator Console 95
- Performance Monitor 96
- saving SQL passwords 99

TreeView pane

- adding computers 43
- adding groups 49
- agent status 47, 93
- changing groups 51
- details panel 38, 65
- job indicator 138
- network hierarchy 58
- organizing the hierarchy 49
- removing computers 51
- removing groups 51
- renaming groups 51
- split pane 38, 65
- tooltips for TreeView details 65
- what's displayed 58

U

- UNIX agent daemon 22
- update interval
 - Operator Console preference for views 67

V

- version history 92
- View Manager settings 66
- view rule
 - creating 234
 - renaming 242

views

- deleting 58
- hiding 40
- introduction 33
- master view 33
- renaming 40, 41
- saving settings 40
- selecting 39
- showing 40
- updating 67

W

- warning events 155
- Web pane, displaying 93
- Windows services
 - NetIQccm 22
 - NetIQmc 22
 - NetIQms 24