

LANDesk® System Manager 8.7

User's Guide



»»»
LANDesk®



Nothing in this document constitutes a guaranty, warranty, or license, express or implied. LANDesk disclaims all liability for all such guaranties, warranties, and licenses, including but not limited to: Fitness for a particular purpose; merchantability; non infringement of intellectual property or other rights of any third party or of LANDesk; indemnity; and all others. LANDesk products are not intended for use in medical, life saving, or life sustaining applications. The reader is advised that third parties can have intellectual property rights that can be relevant to this document and the technologies discussed herein, and is advised to seek the advice of competent legal counsel, without obligation of LANDesk.

LANDesk retains the right to make changes to this document or related product specifications and descriptions at any time, without notice. LANDesk makes no warranty for the use of this document and assume no responsibility for any errors that can appear in the document nor does it make a commitment to update the information contained herein.

Copyright © 2002-2006, LANDesk Software Ltd. or its affiliated companies. All rights reserved.

LANDesk, Autobahn, NewRoad, Peer Download, and Targeted Multicast are either registered trademarks or trademarks of LANDesk Software, Ltd. or its controlled subsidiaries in the United States and/or other countries.

*Other brands and names are the property of their respective owners.

Contents

Cover	1
Contents	3
Overview	5
About LANDesk® System Manager	5
Getting started	9
Licensing	21
Adding licenses	21
The console	22
Starting the console.....	22
Using the console	22
Targeting devices	26
Filtering the display list.....	27
Using groups	28
Using the Actions tab	29
Custom columns.....	31
Custom attributes	33
Page settings.....	33
Viewing the Server information console.....	34
Role-based administration	46
About role-based administration.....	46
Adding product users.....	50
Creating scopes.....	51
Assigning rights and scope to users.....	52
Device discovery	54
Using device discovery.....	54
Creating discovery configurations	56
Scheduling and running discovery	58
Viewing discovered devices	59
Moving discovered devices to the My devices list.....	60
Device agent installation and configuration	64
Agent installation and configuration overview	64
Configuring agents	66
Deploying agents to managed devices	69
Installing agents.....	70
Installing agents with an installation package	70
Pulling the agents.....	71
Installing Linux server agents.....	74
Device monitoring	80
About monitoring	80
Setting performance counters	83
Monitoring performance	84
Monitoring configuration changes	85
Monitoring for connectivity.....	85
Alert configuration	87
Using alerts.....	87
Configuring alert actions.....	90
Configuring an alert ruleset	91
Deploying rulesets.....	93

Viewing alert rulesets for a device.....	93
Viewing the alert log	94
Software updates.....	96
Scripts.....	107
Managing scripts	107
Scheduling tasks	110
Reports	113
About reports	113
Viewing reports.....	113
Queries	115
Using queries.....	115
Understanding custom queries.....	117
Creating custom queries	118
Viewing query results	121
Viewing drill-down query results.....	121
Exporting query results to CSV files.....	121
Changing query column headings.....	122
Exporting and importing queries.....	122
Inventory management	124
Managing inventory	124
Inventory scanning overview	124
Viewing inventory data	126
Customizing inventory options	128
Editing the LDAPPL3.TEMPLATE file	128
Hardware configuration	132
Intel* AMT support.....	132
Configuring Intel* AMT devices	133
Changing the username and password for Intel* AMT devices	137
Configuring System Defense policies.....	138
Intel* AMT Agent Presence configuration	140
IPMI support	141
IPMI BMC configuration	143
Managing Dell* DRAC devices.....	148
Core database installation and maintenance	150
Core database installation	150
Microsoft SQL Server 2005 configuration	150
Appendix A: System requirements and port usage.....	153
Appendix B: Activating the core server	157
Appendix C: Configuring services.....	160
Configuring services tabs	160
Configuring preferred server credentials.....	165
Configuring the custom jobs service	165
Appendix D: Agent security and trusted certificates.....	169
Troubleshooting tips.....	171

Overview

About LANDesk® System Manager

Welcome to LANDesk® System Manager 8.70, a stand-alone management application that lets you maintain the availability of your servers including those running Windows, Linux, HP-UX and AIX. It can also be installed and used concurrently with LANDeskManagement Suite, using the same core database as Management Suite to ease IT-wide reporting.

Designed with an emphasis on low resource impact, this product has several "on-demand" agents and services that run only when they are needed, thus freeing memory and CPU cycles for other tasks. LANDesk knows device availability is critical for your company, so the product is designed for stability, running in 24/7 environments. It leaves you in control of the software running on your devices. You can install the full agent, select specific components, or move devices to your device list without installing any agents.

For dialogs and windows to display properly, the System Manager website must be added to the allowed list of the browser's popup blocker.

What's New to Version 8.70

The following features have been added or upgraded from the previous version of System Manager:

Agentless device management: Manage devices in the **My devices** view without installing a management agent on them, when they are enabled with an out-of-band management technology such as Intel* AMT, IPMI, or DRAC.

Off-core Inventory Service option: Place the cumulative inventory collection process onto another server to lessen the resource requirement on the core server

Custom groups in Scheduled tasks: You can group tasks into custom groups for execution.

Beginner mode: You can display labels to the buttons on toolbars, thus making it easier for new users to see what the buttons do. Likewise, you can choose not to display labels (tooltips still display on mouseover).

Hardware configuration: This new tool lets you configure options for devices with Intel* AMT capabilities. You can generate IDs for provisioning Intel AMT devices, view the generated IDs, and change configuration options related to provisioning your Intel AMT devices. You can also define circuit breaker policies, which detect and block suspicious network activity on the devices.

Enhanced support for Active Management Technology: This product now supports Intel* Active Management Technology 2 (in addition to 1). AMT version 2 also supports agentless management and automatic discovery of Intel* AMT 2 devices.

Product features

System Manager enables you to choose your level of management coverage, from simple information gathering to extended performance analysis, security and configuration control. System Manager includes the following:

Easy-to-use Web console: Run the product anytime, anywhere using a Web-based console designed to deliver rich data in an easy-to-use interface. You can run it from your primary workstation or from a workstation in the server room with no install. Simply browse to the product URL, <http://coreserver/LDSM>. You "target" specific devices for actions such as software distribution by selecting them for placement in the **Targeted devices** list, similar to the "shopping cart" model in many Web applications.

Expanded operating system support: Manage your diverse server environment from one integrated console. In addition to managing Windows 2000 and 2003 servers, System Manager provides support for several flavors of Linux and Unix:

- Red Hat Enterprise Linux v3 (ES 32-bit - U6)
- Red Hat Enterprise Linux v3 (ES EM64t - U6)
- Red Hat Enterprise Linux v3 WS 32-bit - U6
- Red Hat Enterprise Linux v3 WS EM64t - U6
- Red Hat Enterprise Linux v3 (AS 32-bit - U6)
- Red Hat Enterprise Linux v3 (AS EM64t - U6)
- Red Hat Enterprise Linux v4 (ES 32-bit - U3)
- Red Hat Enterprise Linux v4 (ES EM64t - U3)
- Red Hat Enterprise Linux v4 (AS 32-bit - U3)
- Red Hat Enterprise Linux v4 (AS EM64t - U3)
- Red Hat Enterprise Linux v4 WS 32-bit - U3
- Red Hat Enterprise Linux v4 WS EM64t - U3
- SUSE* Linux Server 9 ES 32-bit SP3
- SUSE Linux Server 9 EM64t SP3
- SUSE Linux Professional 9.3
- SUSE Linux Professional 9.3 EM64t
- HP-UX 11.1
- Unix AIX

Scheduled task view: View all scheduled or completed agent deployment, discovery, software update, and custom script tasks from one location. You can reschedule the task, modify it, or make it a recurring event.

Intel* AMT support: Support for Intel* Active Management Technology versions 1 and 2. Intel AMT lets you remotely manage networked devices in any system state through out-of-band (OOB) communication, even when the OS is unresponsive or the device is turned off. The only requirements on the device side are that it be connected to a corporate network and have stand-by power.

IPMI support: The product provides support for Intelligent Platform Management Interface (IPMI) enabled servers (versions 1.5 or 2.0, allowing out-of-band remote recovery of downed servers and viewing of autonomous management data even when the OS or processor isn't running.

Scripting tool: You can execute custom tasks on devices by creating local scheduler scripts.

Performance monitoring: You can monitor the real-time performance of your managed enterprise or blade servers, using a wide range of attributes. You can even track these attributes and see historical performance data reported over several days. You can monitor devices that have the monitoring agent installed, and you can also monitor out-of-band IPMI-enabled servers without an agent.

Blade server support: IBM blade chassis and server blades are supported, including discovery, chassis detection, inventory, and patch management capabilities. Product tools enable you to group blades by function, chassis, rack or other criteria for more effective data-gathering.

Reporting: You can run reports on any device in the database showing usage statistics, resource allocation, and many other measurements. This product includes several canned (pre-formulated) reports. These reports run quickly by directly accessing the database to gather information and represent data in two- or three-dimensional pie and bar graphs. You can create additional reports by creating custom queries.

Health monitoring / Alerts: Monitoring the overall health of a device is easy. You can set thresholds for measurements such as disk space or CPU usage, and configure how you want to be alerted in the event a threshold is exceeded. You can see health concerns of the selected device and initiate action to resolve the problem before users experience sluggish performance or downtime because of the problem.

Software updates: You can receive software updates for System Manager and for Intel* hardware. You can manually deploy the selected updates using the software distribution capabilities.

Role-based administration: Add users and configure their access to tools and other devices based on their administrative role. With role-based administration, you assign scope to determine the devices a user can view and manage, and rights to determine the tasks they can perform, such as reporting only users.

Inventory: Using the inventory scanning tool, the product compiles a wealth of hardware and software information into the core database. You can then view, print, and export this data.

Device discovery: Be sure of what is on your network. Device discovery gathers basic information on all devices and other devices in your environment, enabling greater control and speeding agent deployment to targeted devices.

Support for Active System Console: Support for the Active System Console, which provides a quick overview of system health when you install the Active System Console agent on a device. You can see at a glance whether selected hardware elements are functioning correctly and whether there are potential problems that may need to be addressed. You can also view detailed system performance metrics and view a list of system components, including hardware, software, logs, and information about Intel* AMT and IPMI (if the device is enabled with either).

Executive dashboard: A group of tools (informative charts, diagrams, dials, and meters) that enables executives to monitor the health or status of their business.

Help: This product includes a [Getting Started Guide](#), as well as context-sensitive help topics.

Product Terms

- **Core server:** The center of a management domain. All the product's key files and services are on the core server. A management domain has only one core server. A core server can be a new server or a repurposed server.
- **Console:** The browser-based console that is the main product interface.
- **Core database:** The product creates an MSDE database on the core server to store management data.
- **Managed devices:** Devices in your network that have product agents installed. "Devices" include desktop computers, servers, laptop/mobile computers, blade chassis, and so forth. A core server can manage thousands of devices.
- **Public:** Items (such as groups, distribution packages, or tasks) that are visible to all users. When a user modifies a Public item, the modification remains Public. Public groups are created by a user with Administrator rights.
- **Private or User:** Items that are created by the currently logged-in user. They are not visible to other users. Private or User items appear under the **My delivery methods**, **My packages**, and **My tasks** trees. Users with Administrator rights can see Private groups and User packages and tasks.
- **Common:** An item visible to other users. When a user assumes ownership of a Common item (by modifying it), the item branches into two items: the Common item remains, and a User item is saved in the Users folder. The User instance of the item is no longer visible to other users. A user can mark any task that is visible to them as Common, thus sharing it with other users. Once a user clears the Common option of the item's properties, the task is only visible in the user's User tasks group.

Getting started

- [Overview](#)
- [Running the installation program](#)
- [Activating the core server](#)
- [Adding users](#)
- [Configuring services and credentials](#)
- [Running the console](#)
- [Discovering devices](#)
- [Scheduling and running the discovery](#)
- [Viewing discovered devices](#)
- [Moving devices to the My devices list](#)
- [Grouping devices for actions](#)
- [Configuring devices for management](#)
- [What's next?](#)

Overview

Welcome to LANDesk® System Manager, a stand-alone device management application that maximizes your valuable time by letting you quickly and efficiently manage your devices, thus saving you and your organization time and money. System Manager lets you manage your devices in a central location, group them for actions (such as power cycling, vulnerability assessments, or configuring alerts), remotely troubleshoot any problems, keep your network secure, and keep your devices updated with the latest patches.

This guide's purpose is to help you start using System Manager quickly by configuring services, running the console, discovering devices, moving the devices into the **My devices** list, and configuring the managed devices for actions.

System Manager is a Web application, allowing you to access it using your browser so you can manage your servers from a remote workstation. It behaves like many of the Web applications which you are accustomed to, but it also contains several advanced Windows-type controls to enhance your usability experience. For example, you can hover the mouse pointer over a control and then double-click it or right-click it (just as you would in a Windows application). For example, in the **My devices** list, you can double-click a device name to access its specific information, or right-click to see available actions.

The steps below guide you through getting System Manager up and running, discovering devices on your network, selecting the servers to move to your **My devices** list, deploying agents, and then targeting those devices for various tasks.

Running the installation program

During the install, on the Autorun page, select LANDesk® System Manager. Specific installation instructions can be found in Phase 2 of the *Installation and Deployment Guide*.

After you have installed System Manager, you are ready to start using it. The sections below tell you how to complete several required tasks: running the core activation utility, configuring

services, discovering computers, specifying which devices to actively manage by moving the devices the **My devices** list, grouping devices, adding users, and deploying agents. Once these tasks are completed, you are ready to begin exploring how the robust feature set of System Manager can help you manage your devices.

Activating the core server

You won't be able to run the product until you have activated the core server.

Use the Core Server Activation utility to:

- Activate a new System Manager core server for the first time
- Update an existing System Manager core server or upgrade to Management Suite or System Manager

Each core server must have a unique authorization certificate.

This utility runs automatically on the first reboot.

With your core server connected to the Internet,

1. Click **Start | All Programs | Core Server Activation**.
2. Type in the unique user name and password provided by when you purchased your licenses.
3. Click **Activate**.

The core communicates with the Software licensing server via HTTP. If you use a proxy server, click the utility's **Proxy** tab and enter your proxy information. If your core has an Internet connection, communication with the license server is automatic and won't require any intervention by you. If the core is not connected, click Close on reboot, and email the authorization file into licensing@landesk.com.

Periodically, the core server generates node count verification information in the "Program Files\LANDesk\Authorization Files\LANDesk.usage" file. This file gets sent periodically to the LANDesk Software licensing server. This file is in XML format and is digitally signed and encrypted. Any changes made manually to this file will invalidate the contents and the next usage report to the Software licensing server.

- The Core Server Activation utility won't automatically launch a dial-up Internet connection, but if you launch the dial-up connection manually and run the activation utility, the utility can use the dial-up connection to report usage data.
- You can also activate the core server by e-mail. Send the file with the .TXT extension located under Program Files\LANDesk\Authorization to licensing@landesk.com. LANDesk customer support will reply to the e-mail with a file and instructions on copying the file to the core server to complete the activation process.

Adding users

System Manager users are users who can log in to the console and perform specific tasks for specific devices on the network. You manage users through the role-based administration feature. Role-based administration lets you assign product users special administrative roles

based on their rights and scope. *Rights* determine the product tools and features a user can see and utilize. *Scope* determines the range of devices a user can see and manage. You can create a variety of users and customize their rights and scope to fit your management requirements. For example, you can create a user who fills the Help Desk role by giving this user the rights necessary for this role. More detail is available in the Role-based Administration chapter of the System Manager *Users Guide*.

When you install the product, two user accounts are automatically created (see below). If you want to add more users, you can do so manually. Users are not actually created in the console. Instead, users appear in the Users group (click Users in the left navigation pane) after they have been added to the LANDesk Management Suite group in the Windows NT users environment on the core server. The Users group shows all of the users currently residing in the LANDesk Management Suite group on the core server.

There are two default users in the Users group. One of the users is the Default Administrator. This is the administrative user who was logged in to the server when the product was installed.

The other default user is the Default Template User. This user contains a template of user properties (rights and scope) that is used to configure new users when they are added to the Management Suite group. In other words, when you add a user to that group in the Windows NT environment, the user inherits the rights and scope currently defined in the Default Template User properties. If the Default Template User has all rights selected and the Default All Machines Scope selected, any new user placed in the LANDesk Management Suite group will be added to the Users group with rights to all of the product tools and access to all devices.

You can change the property settings for the Default Template User by selecting it and clicking Edit. For example, if you want to add a large number of users at once, but do not want them to have access to all of the tools or devices, change the settings for the Default Template User first, then add the users to the LANDesk Management Suite group (see steps below). The Default Template User cannot be removed.

When you add a user to the LANDesk Management Suite group in Windows NT, the user is automatically read into the Users group in the Users window, inheriting the same rights and scope as the current Default Template User. The user's name, scope, and rights are displayed. Additionally, new user subgroups, named by the user's unique login ID, are created in the User Devices, User Queries, User Reports, and User Scripts groups (note that ONLY an Administrator can view User groups).

Conversely, if you remove a user from the LANDesk Management Suite group, the user no longer appears in the Users list. The user's account still exists on the core server and can be added back to LANDesk Management Suite group at any time. Also, the user's subgroups under User Devices, User Queries, User Reports, and User Scripts are preserved so that you can restore the user without losing their data, and so that you can copy data to other users.

Refresh the Users frame in the System Manager console by pressing **F5**. To learn how to add a user or domain group to the LANDesk Management Suite group or how to create a new user account, please see "Adding Product Users" in the Role-based Administration chapter of the System Manager *User's Guide*.

To add a user or domain group to the LANDesk Management Suite group

USERS GUIDE

1. Navigate to the server's **Administrative Tools | Computer Management | Local Users and Groups | Groups** utility.
2. Right-click the **LANDesk Management Suite group**, and then click **Add to group**.
3. Click **Add**, then type or select a user (or users) from the list.
4. Click **Add**, and then **OK**.

Note: You can also add a user to the LANDesk Management Suite group by right-clicking the user account in the **Users** list, clicking **Properties | Member Of**, and then clicking **Add** to select the group and add the user.

If user accounts do not already exist on the server, you must first create them on the server.

To create a new user account

1. Navigate to the server's **Administrative Tools | Computer Management | Local Users and Group | Users** utility.
2. Right-click **Users**, and then click **New User**.
3. In the **New User** dialog, enter a name and password.
4. Specify password settings.
5. Click **Create**. The **New User** dialog remains open so that you can create additional users.
6. Click **Close** to exit the dialog.

Add the user to the LANDesk Management Suite group to have them appear in the Users group in the console.

Configuring services and credentials

Before you can manage devices on your network, you must provide System Manager with the necessary device credentials. Use the Configure Services utility on the core (SVCCFG.EXE) to specify the required operating system, Intel* AMT, and IPMI BMC credentials. You can also specify additional settings, such as inventory defaults, PXE holding queue settings, and LANDesk database settings.

Use Configure Services to configure:

- The database name, username, and password. (Set at installation time.)
 - Credentials for scheduling jobs to the managed devices. (You can enter more than one set of administrator credentials.)
 - Credentials for configuring IPMI BMCs. (You can enter only one set of BMC credentials.)
 - Credentials for configuring Intel AMT-enabled devices. (You can enter only one set of Intel AMT credentials.)
 - Server software scan interval, maintenance, days to keep inventory scans, and login history length.
 - Duplicate device ID handling.
 - Scheduler configuration, including scheduled job and query evaluation intervals.
 - Custom job configuration, including remote execute timeout.
1. At the core server, click **Start | All Programs | LANDesk | LANDesk Configure Services**.
 2. Click the **Scheduler** tab.
 3. Click the **Change Login** button.

4. Enter the credentials you want the service to use on the managed devices, typically a domain administrator account.
5. Click **Add**. Add additional credentials as necessary, if the managed devices do not all have the same administrator user name accounts enabled.
6. Click **Apply**.
7. If you have IPMI-enabled servers in your environment, click the **BMC Password** tab. Type a password in the **Password** text box, retype the password in the **Confirm password** text box, then click **OK**. (All managed IPMI servers must share the same BMC user name and password.)
8. If you have Intel AMT-enabled devices, click the **Intel AMT Configuration** tab. Type the currently configured Intel AMT user name in the **User name** text box and the currently configured password in the **Password** text box. Retype the password in the **Confirm password** text box, then click **OK**.
9. Set any other settings as desired, such as software scan intervals.
10. Click **OK** to save the changes.

Click **Help** on each Configure Services tab for more information.

Running the console

System Manager includes a full range of tools that let you view, configure, manage, and protect the devices on your network. The console is the entry point by which you can use these tools.

The top pane in the console displays the server you are logged in to and the user you are logged in as. The **My devices** list is the main window of the console and is the starting point for most functions. The left-hand pane shows available tools. The right-hand pane in the console displays dialogs and screens which allow you to complete management tasks.

The convenience of the console is that you can perform all of its functions from a remote location, such as your workstation, freeing you from the need to take additional trips to the server room or to go to each managed device individually to perform routine maintenance or troubleshoot problems.

Launch the console one of three ways:

- On the core server, click **Start | All Programs | LANDesk | System Manager**.
- In a browser at a remote workstation, type the URL <http://coreserver/LDSM>.

Discovering devices

Use the **Discovery configurations** tab to create new discovery configurations, edit and delete existing configurations, and schedule a configuration for discovery. Each discovery configuration consists of a descriptive name, the IP ranges to scan, and the discovery type.

Once you create a configuration, use the **Schedule discovery** dialog to configure when it will run.

1. In the left navigation pane, click **Device discovery**.
2. In the **Discovery configurations** tab, click the **New** button.
3. Fill in the fields described below. When you have finished, click the **Add** button, and click **OK**.

The text below describes the parts of the **Discovery configuration** dialog box.

- **Configuration name:** Type a name for this configuration. Give the configuration a meaningful name so you can easily remember the configuration. The configuration can be up to 255 characters long, and should not contain the following characters: ", +, #, & or %. The configuration name will not display after the use of any of these characters.
- **Standard network scan:** Looks for devices by sending ICMP packets to IP addresses in the range you specify. This is the most thorough search, but also the slowest. By default, this option uses NetBIOS to gather information about the device.

The network scan option has an **IP fingerprint** option where device discovery tries to discover the OS type through TCP packet responses. The IP fingerprint option slows down the discovery somewhat.

The network scan option also has a **Use SNMP** option, where you can configure the scan to use SNMP. Click **Configure** to enter information about your SNMP configuration.

- **LANDesk CBA discovery:** Looks for the standard management agent (formerly known as the common base agent [CBA] in Management Suite) on devices. The standard management agent allows the core server to discover and communicate with clients on the network. This option discovers devices that have product agents on them. Routers block standard management agent and PDS2 traffic. In order to run a standard CBA discovery across multiple subnets, the router must be configured to allow directed broadcast across multiple subnets.

The CBA discovery option also has a **LANDesk PDS2 discovery** option, where device discovery looks for the LANDesk Ping Discovery Service (PDS2) on devices. LANDesk Software products such as LANDeskSystem Manager, Server Manager, and LANDesk Client Manager use the PDS2 agent. Select this option if you have devices on your network with these products installed. CBA discovery is not supported for Linux machines, but if you choose PDS2, Linux machines with an agent installed can be discovered.

- **IPMI:** Looks for IPMI-enabled servers. IPMI is a specification developed by Intel,* H-P,* NEC,* and Dell* to define the message and system interface for management-enabled hardware. IPMI contains monitoring and recovery features that let you access these features regardless of whether the device is turned on or not, or what state the OS may be in. Please keep in mind that if the Baseboard Management Controller is not configured, it will not respond to ASF pings which the product uses to discover IPMI. This means that you will have to discover it as a normal computer. When you push the client, ServerConfig will scan the system and detect it is IPMI and configure the BMC.
- **Server chassis:** Looks for blade server chassis management modules (CMMs). The blades in the server chassis are detected as normal servers.
- **Intel* AMT:** Looks for devices with Intel Active Management Technology support.
- **Starting IP:** Enter the starting IP address for the range of addresses you want to scan.
- **Ending IP:** Enter the ending IP address for the range of addresses you want to scan.
- **Subnet mask:** Enter the subnet mask for the IP address range you want to scan.
- **Add:** Adds the IP address ranges to the work queue at the bottom of the dialog.
- **Clear:** Clears the IP address range fields.
- **Edit:** Select an IP address range in the work queue and click **Edit**. The range appears in the text boxes above the work queue, where you can edit the range and add the new range to the work queue.
- **Remove:** Removes the selected IP address range from the work queue.
- **Remove all:** Removes all IP address ranges from the work queue.

Now that you have configured a discovery task, you are prepared to discover the devices connected to your network by scheduling when the discovery task is to run.

Scheduling and running the discovery task

Use the **Schedule** button on the **Discover devices** tab to display the **Schedule discovery** dialog. Use this dialog to schedule when a discovery will run. You can schedule a discovery task to run immediately, at some point in the future, make it a recurring schedule, or run it just once and never worry about doing it again.

Once you schedule a discovery task, see the Discovery tasks tab for discovery status. Scheduling a recurring discovery task assists you by automatically discovering new devices that come up on the network.

The **Schedule discovery** dialog has these options.

- **Leave unscheduled:** Leaves the task unscheduled but keeps it in the Discovery configurations list for future use.
- **Start now:** Runs the task as soon as possible. It may take up to a minute for the task to start.
- **Start at scheduled time:** Starts the task at the time you specify. If you click this option, you must enter the following:
 - **Time:** The time you want the task to start
 - **Date:** The date you want the task to start. Depending on your locale, the date order will be day-month-year or month-day-year.
 - **Repeat every:** If you want the task to repeat, select whether you want it to repeat Daily, Weekly, or Monthly. If you pick Monthly and the date doesn't exist in all months (for example, 31), the task will only run in months that have that date.

To schedule a discovery task

1. In the left navigation pane, click **Discovered devices**.
2. On the **Discovery configurations** tab, select the configuration you want and click **Schedule**. Configure the discovery schedule and click **Save**.
3. Monitor the discovery progress in the **Discovery tasks** tab. Click **Refresh** to update the status.
4. When the discovery completes, click **Unmanaged** to view all discovered devices in the upper **Discovered devices** pane (the pane does not refresh automatically).

Viewing discovered devices

Discovered devices are categorized by device type in the **Discovered devices** pane. The **Computers** folder is displayed by default. Click the folders in the left pane to view devices in different categories. Click **Unmanaged** to view all devices returned by the discovery.

- Blade server chassis appear in the **Chassis** folder.
- Standard enterprise devices appear in the **Computers** folder.
- Routers and other devices appear in the **Infrastructure** folder.
- Intel AMT-enabled devices appear in the **Intel AMT** folder.
- IPMI-enabled servers appear in the **IPMI** folder.

- Non-categorized devices appear in the **Other** folder.
- Printers appear in the **Printers** folder.

Note: Some Linux servers appear with the generic "Unix" as the operating system name (or even sometimes show as Other). When the standard management agent is deployed, these servers will update their OS name entry in the **My devices** list and display a full inventory.

To view discovered servers

1. In the Device discovery page, in the left pane, click **Computers** or another type of device you want to view. The results are displayed in the right pane.
2. To filter the results, click the **Filter** icon, type at least a portion of what you are searching for, and click **Find**.

Assigning names

When doing a network scan discovery, some servers return with blank node name (or host name). This occurs most frequently with servers running Linux. You must assign a name to the device before you can use Manage to move it to the My devices list.

1. In the Device discovery page, click the device with a blank name. (You must click the blank area in the node name column.)
2. Click **Assign name** on the toolbar.
3. Type in the name and click **OK**.

When you install a product agent on a device, it automatically scans the host name and updates the core database with the correct information.

Moving devices to the My devices list

Once discovered, you must manually target the devices you wish to manage and move them to the My devices list. Moving the device does not install any software to the device. It only makes the device available for querying, grouping, and sorting in the My devices list. You "target" specific devices for specific action, a model similar to the "shopping cart" model in many Web applications.

1. In the **Device discovery** view, click the device you want to move to the **My devices** list. You can select multiple devices by pressing SHIFT+click or CTRL+ click.
2. Click the **Target** button. Click the **Target list** tab.
3. Select the devices, then click the **Manage** tab.
4. Select **Move targeted devices**.
5. Select **Manage out-of-band-enabled devices agentless (without using an OS-specific agent)** to perform an out-of-band scan on an IPMI device. This scan will return IPMI-specific information only, such as FRU and SDR information.

The BMC Password you configured in Configure Services is used to communicate with the BMC. You must set the same password in Configure Services and the BMC of the target machine before the move will succeed.

6. Click **Move**.

The devices are moved to the **My devices** list and their information is moved to the database. Once the information is in the database, you can run queries and reports on it, configure thresholds for alerts, and many other vital management tasks.

After you select devices in the **Device discovery** view you can also go directly to the **Manage** tab (without targeting the devices). If you do this, select **Move selected devices** in step 4 above.

Grouping devices for actions

You may want to organize your devices into groups, such as by geographic location or function, so you can perform actions on them more quickly. For example, you may want to see the processor speeds of all the devices in a specific location.

1. In the **My devices** list, click **Private groups** or **Public groups**, then click **Add group**.
 1. In the **Device discovery** view, click the device you want to move to the **My devices** list. You can select multiple devices by pressing SHIFT+click or CTRL+ click.
 2. Click the **Target** button. Click the **Target list** tab.
 3. Select the devices, then click the **Manage** tab.
 4. Select **Move targeted devices**.
 5. Select **Manage out-of-band-enabled devices agentless (without using an OS-specific agent)** to perform an out-of-band scan on an IPMI device. This scan will return IPMI-specific information only, such as FRU and SDR information.

The BMC Password you configured in Configure Services is used to communicate with the BMC. You must set the same password in Configure Services and the BMC of the target machine before the move will succeed.

6. Click **Move**.

The devices are moved to the **My devices** list and their information is moved to the database. Once the information is in the database, you can run queries and reports on it, configure thresholds for alerts, and many other vital management tasks.

After you select devices in the **Device discovery** view you can also go directly to the **Manage** tab (without targeting the devices). If you do this, select **Move selected devices** in step 4 above.

2. Type a name for the group in the **Group name** box.
3. Click the type of group you want to create.
 - **Static:** Devices that have been added to the group. They remain in the group until they are removed or until you no longer manage them.
 - **Dynamic:** Devices that meet one or more criteria as defined by a query. For example, a group may contain all servers that are currently in a Warning state. They remain in the group as long as they match the criteria defined for the group. Devices are automatically added to dynamic groups when they meet the group query criteria.
4. When you are finished, click **OK**.
5. To add devices to a static group, click devices in the right pane of the **My devices** list, click **Move/Copy**, select the group, and click **OK**.

Configuring devices for management

Discovering devices alone does not put them under the management umbrella. Before you can fully manage devices with the console and receive health alerts, you need to install management agents on them. You can choose to install the default agent configuration (which installs all management agents) or customize your own agent configuration to install on your devices. (The agent configuration must include the monitoring agent to receive health alerts.)

You can install management agents in any of the following ways:

- Target devices in the **My devices** list, then schedule an agent configuration task to remotely install agents on the devices. (steps below)
- Map to the core's LDlogon share (//coreserver/ldlogon) and run SERVERCONFIG.EXE. (steps are in "Pulling the agents" in the Device Agent Installation and Configuration chapter of the *System Manager User's Guide*)
- Create a self-extracting device installation package. Run this package locally on the device to install the agents. This must be done while logged in with administrative privileges. (steps are in "Installing Agent with an Installation Package" in the Device Agent Installation and Configuration chapter of the *System Manager User's Guide*)

To push the agent:

1. Target devices in the **My devices** list (as explained above in Moving devices to the My devices list)
2. In the left navigation pane, click **Agent configuration**, right-click the configuration you want to push, and click **Schedule task**.
3. In the left pane, click **Target devices**, and click the **Add target list** button.
4. Click **Schedule task**, click **Start now** to start the task immediately or **Start later** and set the task's start date and time, and click **Save**.

You can view the status of the task in the **Configuration tasks** tab.

Installing Linux server agents

You can remotely deploy and install Linux agents and RPMs on Linux servers. Your Linux server must be configured correctly for this to work. The instructions to properly configure a Linux server found in "Installing Server Agents" in the Device Agent Installation and Configuration chapter of the *System Manager User's Guide*.

Setting alerts

When a problem or other event occurs on a device (for example, the device is running low on disk space). System Manager can send an alert. You can customize these alerts by choosing the severity level or threshold that will trigger the alert. Alerts are sent to the console and can be configured to perform specific actions. You can set alerts for many events or potential problems. The product comes with a default alert ruleset, and this ruleset is installed to a managed device when the monitoring component is installed. This ruleset of alerts provides health status feedback to the console. This default ruleset includes alerts such as:

- Disk added or removed
- Drive space
- Memory usage
- Temperature, fans, and voltages
- Performance monitoring
- IPMI events (on applicable hardware)

To learn more about alerting, please see the Alert configuration chapter in the *System Manager User's Guide*.

What's next?

You now have gotten Server Manager up and running. You have used only a fraction of the features available in Server Manager, and only a portion of the features you did use (like device discovery and agent configuration). The companion guides (the *Installation and Deployment Guide* and the *Users Guide*) can provide more in-depth information on all the product's features. Some of these features include:

Software updates: Establish ongoing patch-level security on the managed devices across your network. You can automate the repetitive processes of maintaining current vulnerability information, assessing vulnerabilities for the various operating systems running on your managed devices, downloading the appropriate patch executable files, remediating vulnerabilities by deploying and installing the necessary patches on affected devices, and verifying successful patch installation.

Alerting: Ensure that you are alerted if any of your devices reach a particular threshold. Related to the Monitoring feature, Alerting can notify you in many different ways. For example, if you need to know when the storage on your devices reaches 95% of capacity, you can choose how you want to be alerted (the agent can send e-mail or pager messages, reboot or shut down a device, or add information to the alert log).

Queries: Manage your network by searching for and organizing devices in the core database based on specific system or user criteria. You can query the list of managed devices for those which match the criteria you specify (such as all located in the corporate office or all with 256K of RAM) and group them for actions. These groups can be static (the members of the group can only be changed manually) or dynamic (the members change when devices meet or fail to meet specified criteria).

Software distribution: Create tasks to distribute software packages (one or more MSI files, an executable, a batch file, RPM files (Linux), or a package created with LANDesk package builder) to target devices.

Monitoring: Monitor a device's health status by using one of the supported types of monitoring (direct ASIC monitoring, in-band IPMI, out-of-band IPMI, CIM, and so forth). Monitoring lets you keep track of many pieces of data on your devices, such as usage levels, OS events, processes and services, historical performance, and hardware sensors (fans, voltages, temperatures, etc.). Alerting is a related feature that uses the monitoring agent to initiate alerting actions.

Reporting: Generate a wide variety of specialized reports that provide critical information about the managed devices on your network. Server Manager uses an inventory scanning utility to add devices (and collected hardware and software data about those devices) to the core database.

USERS GUIDE

You can view and print this inventory data from a device's inventory view, as well as use it to define queries and group devices together. The reporting tool takes further advantage of this scanned inventory data by collecting and organizing that data in useful report formats, which can be helpful in gathering and formatting data for regulatory reports.

Unmanaged device discovery: Find devices that aren't being managed by the console. Discovery is the first step to getting new machines into management quickly. You can set up a discovery task to scan for new machines every month.

Software license monitoring: Track overall license compliance. The software license monitoring agent gathers data (such as the total minutes of usage, the number of launches and the last launch date of all installed applications on a device) and stores this data in the device's registry. You can use the data to monitor product usage and denial trends. The agent passively monitors product usage on devices, using minimal network bandwidth. The agent continues to monitor usage for mobile devices that are disconnected from the network.

OS Deployment: Deploy OS images to devices on your network by using the PXE-based deployment tool . This allows you to image devices with empty hard drives or unusable operating systems. Lightweight PXE representatives eliminate the need for a dedicated PXE server on each subnet. OS deployment streamlines new device provisioning without requiring additional end user or IT interaction once the process starts.

Licensing

The license process helps keep your organization in compliance with its licensed node agreements by running an ongoing authorization process. This approach also enables you to use multiple core servers under a defined user account. The license process uses a backend database to create and manage user accounts. The license process is a simple request and reply from the core server to the backend process allowing the core to renew its activity for another period.

When you run the product (or any add-on) after an installation, you can activate an evaluation license for a trial period or enter a username and password to activate a purchased license obtained from LANDesk Sales. The same username and password is used to activate all core servers for the existing account.

The activation process is essentially the same for evaluation and purchase products. When the device has an Internet connection, the process is a simple information exchange. When the device is not connected, the manual process of e-mailing a file to LANDesk and then saving a returning file to the core server must be followed. The activation process works like this:

1. User runs [Activate Core utility](#)
2. A file is created containing both server and usage information. It is signed by the core's private key and encrypted with the LANDesk public key.
3. If an Internet connection is available, the core and LANDesk servers communicate and the core uploads the activation file. The backend processes the information and sends back the activation information which is written directly to the database.
4. If an Internet connection is not available, you may e-mail the file in the \Program Files\LANDesk\Authorization Files folder to licensing@landesk.com.

Adding licenses

The functionality available to you through the console is dependent on a license key. You can add a new license key to access additional functionality or update the number of users. During installation, a trial 45-day license is generated. When you add a valid license using the console, the temporary license is deleted.

To add a license key

1. In the left navigation pane, click **Preferences**.
2. Click the **License** tab.
3. At the bottom of the screen, click the link <http://www.landesk.com/contactus/>.

If the above link does not work, it could be because the browser security level is not set to Medium. You should change the default Internet Security Level to Medium in Internet Explorer (**Tools > Internet Options > Security > Internet > Default Level**).

The console

Starting the console

To start the console

1. On the core server, click **Start | All Programs | LANDesk | LANDeskSystem Manager**.

or

On a remote workstation, open a browser and type the address of the console. This will be in the format of `http://corename/ldsm`.

2. Enter a valid user name and password.

If you're connecting to a remote core server, follow the normal Windows rules for remote login (i.e., if the user is local to that core server, just enter the user name; if the user is a domain user, enter the domain name\user name).

3. Click **OK**.

If the device list and buttons do not appear when you start the console, you may need to [activate the core server](#).

About the System Manager login dialog

Use this dialog to launch the console and connect to a core server.

- **Username:** Identifies a user. This might be an administrator user or some other type of product user with restricted access (for more information, see [Role-based administration](#)). The user must be a member of the LANDesk Management Suite group on the core server. If you're connecting to a remote core server, enter the domain name and user name.
- **Password:** The user's password.

Using the console

You can use tools to view, configure, manage, and protect the devices on your network—all from a single console. You can update software or configuration settings, diagnose hardware and software issues, and use role-based administration to control users' access to features and devices. Additionally, if you are also using other LANDesk products, you can connect to them directly from the console.

The top pane in the console displays the server you are logged in to and the user you are logged in as. The **My devices** list is the main window of the console and is the starting point for most functions. The left-hand pane shows available tools. The right-hand pane in the console displays dialogs and screens which allow you to manage devices and users, view reports, run discoveries, create and modify queries, and so on. You can resize the panes and columns of the **My devices**

list. When no agents are installed on a device, the name and IP address are the only columns that contain information. In some instances, the operating system will also display.

System Manager provides some Windows application-like functionality in the convenience and accessibility of your Web browser.

- Right-click a device in the **My devices** list to view available options for that device, such as Ping and Target.
- To select multiple consecutive entries in a list, click the first item, press and hold down the **Shift** key, and then click the last item.
- To select multiple nonconsecutive entries in a list, press and hold down the **Ctrl** key, and then click each item.

For dialogs and windows to display properly, the System Manager website must be added to the allowed list of the browser's popup blocker.

Role-based administration

As a user, the devices you can view and manage in the **My devices** list and the management tools you can use are determined by the access rights and device scope assigned to you by the Administrator. For more information, see "[Role-based administration](#)."

This section provides information about:

- [My devices list](#)
- [Device icons](#)
- [Using shortcut menus](#)
- [Using tools](#)
- [Viewing device properties](#)

My devices list

The **My devices** list contains the following groups and sub-groups. In addition, depending on your access rights and device scope, you can [create your own groups](#) for easier management of devices.

All devices

The **All devices** list shows the devices for the currently logged-in user, based on the user's scope, in a flat list (no sub-groups). When connected to a particular core server, the administrator can see every device managed by that core server. Product users, on the other hand, are restricted and can only see the devices that reside within their assigned scope (a scope is based on either a database query or a directory location).

Devices running product agents (Standard management agent and Inventory) automatically appear in the **All devices** list when they are scanned into the core database by the inventory scanner. Typically, this scan takes place for the first time during initial device configuration. Once a device is scanned into the core database it is considered to be a managed device—it can now be managed by that core server. For more information on setting up devices, see "[Configuring client agents](#)."

Because the **All devices** group is populated automatically via an inventory scan, you may never need to manually discover devices. However, to discover devices not already in the core database (or to move unmanaged devices to the servers group), you can use the device discovery tool to scan the network for devices. For more information, see "[Using discovery](#)."

The **All devices** group provides the following information for each device. Double-click **All devices** to open the list.

- **Name:** The device host name, such as the Windows* computer name.
- **IP address:** The IP address of the device.
- **Health:** The health and availability status of the device. This can be Normal, Warning, or Critical.
- **Agent:** The current agent running on the device.
- **Device type:** Displays the kind of hardware on the machine (Intel AMT, IPMI, ASIC, or IPMI Advanced).
- **Operating system:** The type of operating system the device is running.
- **Up since:** The date and time the computer has been operating without interruption (in the time zone of the database).

When you select a device, the device's properties are displayed in the **Properties** pane below the device list. The **Properties** pane shows many important device attributes:

- **ID:** The identification number of the device. This number is determined by the sequence in which the device was added to the **All devices** list.
- **IP address:** The IP address of the device.
- **Manufacturer:** The device's manufacturer.
- **Model:** The model of the device.
- **Processor speed:** The speed of the device's CPU.
- **Processor type:** The type of the device's CPU.

From the console, you can view a detailed inventory and target the device for an action, such as running a report.

Double-clicking a device in the **All devices** list takes you to the [Server information console](#), which contains device summary information, configuration, remote control options, and alert ruleset information.

Public groups

The **Public groups** list shows groups of devices that have been created by a user with Administrator rights. They are visible to other users.

This list also shows blade chassis groups that are automatically created when a chassis management module (CMM) is added to the list of managed devices. The group lists the CMM and each associated blade server that you manage. You cannot edit a chassis group in the same way you edit a group you have created.

Groups can be static or dynamic. Dynamic groups contain devices that meet predefined filter criteria, such as processor speed, device OS, or a custom attribute such as a device type. Static groups include a set list of devices, other static groups, or dynamic groups.





Private groups

The **Private groups** list shows groups of devices created by the currently logged-in user. Private groups are not visible to other users, so they can't be used by other users.

Device icons

Device icons display in the **All devices** list and show the current health status of each device. You can update the health status for devices one at a time as you select them in devices in the **My devices** list by clicking the **Refresh** toolbar button.

The following table lists the possible device and status icons and what they mean:

Icon	Description
	Device with Normal status
	Device with Warning status
	Device with Critical status
	Device with Unknown status

Using shortcut menus

Shortcut (context) menus are available for all items in the console, including groups, devices, queries, scheduled tasks, scripts, and so on. Shortcut menus provide quick access to an item's common tasks and critical information.

To view an item's shortcut menu, right-click the item. For example, when you right-click a managed device in the **My devices** list, its shortcut menu will typically display the following options:

- **Remove from group:** Removes the item from a user-defined group.
- **Target:** Moves the selected device to the [Targeted devices](#) list. **Note:** If targeted devices do not appear in the **Target list**, click **Refresh** on the **Targeted devices** tab.
- **Ping device:** Verifies the device is awake.
- **Tracert device:** Sends a trace route command to view a network packet being sent and received and the amount of hops required for the packet to reach its destination.

The help does not cover every console item's shortcut menu, but it is recommended that you right-click any item to see the options that are available.

Using tools

Tools are available through the left pane. Use the arrow keys at the top of the pane to view all the tools.

An administrator sees all of the tools in the left navigation pane. Other users will see only the tools (features) that are allowed by their assigned rights. For example, if a user doesn't have the Reports right, the Reports tool does not appear in the left navigation pane.

Here is a complete list of tools:

- **Software updates:** Download applicable update packages
- **Scripts:** Create and manage scripts.
- **Scheduled tasks:** View all tasks (originating in Agent configuration, Vulnerabilities, Device discovery, or Scripts) in the Scheduler.
- **Monitoring:** Monitor the real-time performance of your managed devices using a wide range of attributes.
- **Alerting:** Configure alerts, setting thresholds and the response the product will use if a threshold should be exceeded.
- **Agent configuration:** Configure an IPMI (Baseboard Management Controller), Linux, or Windows agent configuration.
- **Device discovery:** Find devices on the network that aren't scanned into the core database.
- **Logs:** Displays the Alert log, which shows the alerts you have flagged as ones you want to see on your managed devices.
- **Reports:** Manage predefined service reports.
- **Queries:** Create and modify queries to the database to isolate specific devices that meet your criteria.
- **Users:** Control user access to tools and devices based on user rights and scope.
- **Preferences:** Create custom inventory attributes and view licensing information.
- **Hardware configuration:** Open a separate window showing configuration options for Intel* AMT devices.

When you click a tool name, the tool's window opens in the right pane.

Viewing device properties

In the **My devices** view, you can quickly view information about a device by clicking the device in the list and selecting **Properties** in the bottom pane.

More detailed information about the device is available in its inventory data. You can view inventory data in the **All devices** view by clicking the device and selecting the **View inventory** tab in the bottom pane to open the full **Inventory** window.

Targeting devices

The **Targeted devices** list helps you complete tasks on selected devices, such as deploying agents or scanning for software updates to a select group of devices.

The recommended number of devices that you should add to the list is 250 or fewer. The devices will stay in the list until your console session times out (after 20 minutes of inactivity).

Add devices to the **Targeted devices** list by selecting them from any list of devices. If you don't see the devices you want, use the **Find** button on the toolbar. Search for one particular device, or search for several using the wildcard characters % or *. Click the **Target** toolbar button to add the device to the **Targeted devices** list. If the button is not visible, click the << button.

If several devices are found, select the ones you want to add to the list, then click **Target**. If the returned device list spans multiple pages, you must click **Target** for each page. You can't select devices on multiple pages and click the buttons just once for all of the pages. You can click the down arrow below the toolbar on the far right to set how many devices you want to display per page. You can display up to 500 devices per page. To change the number of devices displayed in a list, see [Page settings](#) under **Preferences**.

With one or more devices in the **Targeted devices** list, you can complete a task such as deploying an agent configuration to each of the targeted devices, or moving unmanaged devices to the **My devices** list.

To target devices


1. In the **My devices** list or the **Discovered devices** view, click the device you want to target for an action. You can select multiple devices by using the standard methods of multiple selection (SHIFT+click or CTRL+click).
2. Click the **Target** button. If it is not visible, click << on the toolbar. The button is on the far right.

In the lower pane the selected devices are listed under the **Targeted devices** tab. Once they are listed under this tab, you can open a tool (such as Agent deployment) and schedule a task that can be applied to the targeted devices. If you have targeted unmanaged devices, you can click the **Manage** tab and move them to the **My devices** list.

Filtering the display list

The **My devices** list has a filter icon you can use to determine which devices appear in the list. You can filter by only one of the criteria (by device name or IP address), or you can combine the criteria to focus on a subset of computers.

To filter the display list

1. From the **My devices** list, double-click **All devices** or navigate to a group.
2. Click **Filter**  on the toolbar.
3. In the drop-down list, select **Device name** or **IP address**.

4. Set the parameters of the specified criteria by typing in the text box. In the **Find** box, the following extended characters are not supported: < , > , " , ' , !.

If you filter by device name, type the host name or range of computer names. You can enter wildcard characters to find certain computer names (such as *srv .

5. Click **Find**.

Using groups

You can organize devices in groups for easier management. You can create groups to organize devices based on function, geographic location, department, device attribute or any other category that meets your needs. For example, you could create a Web server group for all servers configured as Web servers, or create a group that includes all devices running a specific OS. You can right-click a group to open it, delete it, or target all of the devices it contains for actions such as alert rulesets and agent deployment.

The main **My devices** view contains the following groups:

- **All devices:** Lists all devices that can be seen by the currently logged-in user, based on the user's scope, in a flat list (no subgroups). For an administrator, **All devices** lists all devices that have been scanned or moved into the core database. Devices configured with the standard management agent automatically appear in the **All devices** group/folder when they are scanned into the core database by the inventory scanner. Users, including administrators, cannot create groups under **All devices**.
- **Public groups:** Lists groups/devices an administrator has added from the **All devices** group, as well as blade chassis groups. An administrator (a user with the Administrator right) sees all of the devices in this group, while other users see only the devices allowed by their scope. Only administrators can create groups under **Public groups**.
- **Private groups:** Lists groups/devices for the currently logged-in user, based on the user's scope. A user can create device subgroups only under **Private groups**. Users can add devices to their **Private groups** group, or any of its subgroups, by moving or copying them from the **Public groups** and **All devices** groups. All users can create groups under **Private groups**.

For more information on which devices you can view and manage in the device view, and the management tools you can use, see "[Role-based administration](#)."

Group types

You can create and manage two types of groups:

- **Static groups.** A *static group* is composed of devices that you have added manually to that group. Static groups can only be changed by manually adding or removing devices.
- **Dynamic groups.** A *dynamic group* is composed of computers that meet filter or query definition. Each time the group is expanded, the query is resolved and the results are displayed. For example, a dynamic group may contain all devices currently in a Warning state. Machines would move in and out of the group as their statuses change.

To create a static group

1. In the console's device view, double-click the parent group (such as **Private groups**), and then click **Add group**.
2. Type a name for the new group.
3. Select **Static** and then click **OK**.

After you have created a static group, you can move or copy devices into the group by selecting them from a list and clicking **Move/copy** on the toolbar. You can copy devices from the **All devices** list into a group or move/copy from other groups.

To create a dynamic group

1. In the console's device view, double-click the parent group (such as **Private groups**), and then click **Add group**.
2. Type a name for the new group.
3. Select **Dynamic** and then click **OK**.

After you have created a dynamic group, you must create a filter for it to determine which computers will appear in that group. You can specify a new filter or base the filter on an existing query.

To create a new filter

1. Select the dynamic group you created (this displays **Group properties** in the bottom pane).
2. From **Group properties**, select **Create a new filter** and click **New filter**.
3. Select the filter criteria you want to use, then click **OK**.

To create a filter based on an existing query

1. Select the dynamic group you created (this displays **Group properties** in the bottom pane).
2. From **Group properties**, select **Create a filter based on an existing query**
3. Select the existing query you want to use to filter the group and click **New filter**.
4. Add any additional filter criteria you want to use, then click **OK**.

If you base a filter on an existing query and that query is later modified by you or another user, the filter based on that query will not dynamically change to match the modified query.

Using the Actions tab

Use the **Actions** tab to execute operations on selected and targeted devices. You can delete devices from the list of managed computers, power on, power off, and reboot devices, and monitor connections with managed devices.

- [Delete devices](#)
- [Power options](#)

- [Device monitor](#)

Delete devices

Delete devices lets you delete selected or targeted devices from the list of managed computers. The delete function can delete single or multiple devices from any group (either a default group or user-created group in System Manager). Once a device has been deleted from a group, it is completely removed from all lists of managed/inventoried devices, including the default **All devices** group.

If you are deleting a large number of devices, the operation may time out. If the operation times out, try breaking up the operation into smaller operations.

Power options

Power options lets you power off, reboot and, in the case of managed IPMI machines, power on remote devices. In the case of non-IPMI servers, the device must have the LANDesk agent deployed to it in order to execute the reboot and power off functions. With IPMI machines, you must have the correct IPMI credentials to execute the power on/power off and reboot features. If an IPMI box has the LANDesk agent deployed to it, then you can execute the power off and reboot features without the IPMI credentials. Use the [Configure Services utility](#) to set the IPMI BMC password to use for managing IPMI servers.

To use power options

1. In the **My devices** list, click a device or [target](#) a list of devices.
2. In the bottom pane, click the **Actions** tab.
3. Click **Power options**.
4. Select whether to perform the action on devices in the [Targeted devices](#) or only selected devices.
5. Select from the following options:
 - Reboot
 - Power off
 - Power on (works on IPMI-enabled and Wake on LAN-enabled devices)

6. Click **Show console redirection window** to launch an ultra-slim container with a launcher (the launcher will be much easier to recompile for EM64T than the TTY control would be).

When powering on or rebooting a managed IPMI server, you can open a console redirection window that displays the server's boot information. This can be useful if you want to verify that the server is rebooting. You can also use the console window to pause the boot process and change BIOS settings on the managed server.

To view a console redirection window, the server must have Console redirection over serial port enabled in its BIOS setup. The console data is sent to the serial port. If there is a serial cable connecting the server and the administrator console, the console redirection is carried over that cable. If not, System Manager initiates a serial over LAN (SOL) connection to redirect the data from the serial port to the LAN connection. The SOL connection stays open as long as the console window is open. When the console data has finished displaying you should close the window.

When the console window opens, a second message window opens. You can close the message window. After the console redirection window opens but before the console displays the boot sequence, you may see random characters in the window. These appear because the server's BMC is sending heartbeat messages, which are transmitted on the connection to the administrator console. The characters do not display while the console shows the boot display, but may reappear after the boot process has completed.

Device monitor

Use Device monitor to check the connectivity of the selected devices. If a device loses its network connectivity, it cannot send an alert to the core server. Device monitor checks to see whether devices are still able to communicate on the network.

1. In the **All devices** list, click a device or [target](#) a list of devices.
2. In the bottom pane, click the **Actions** tab, then click **Device monitor**.
3. To view a list of devices currently being monitored, click **Show monitored devices**.
4. Type numbers for the minutes between ping sweeps and number of times the product will attempt to communicate with a device.
5. Select whether to perform the action on devices in the [Targeted devices list](#) or on all devices in the **All devices** group.
6. To stop monitoring all devices, select **Never ping devices**.
7. Click **Apply**.

Only the last group of targeted devices are monitored. For example, if you target device A and device B and apply device monitoring to them, only device A and device B will be pinged by the core server. If you then target device C and device D and apply device monitoring to those devices, only device C and device D will be monitored; devices A and B will no longer be monitored.

Custom columns

Use **Custom columns** to modify column names and fields. A name is the name of the column, and a field contains the attribute(s) that can appear in the column (if the attribute is present). Any

column changes you make will not be seen by other users. Custom column changes will be seen in the **My devices** view.

This product includes a default column set with seven columns. You cannot edit the default set, but you can define a custom column set to use as your default.

It is not advisable to create custom columns in which there can be multiple field names. For example, if you were to create a Computer.Software.Package.Name field and the device had multiple packages installed, System Manager will list only one package name per line, even if the different package names are on the same device, making the **All devices** list have multiple entries for the same device.

To create a custom column set

1. In the left navigation pane, click **Preferences**.
2. Click the **Custom columns** tab.
3. Click **New**.
4. Type a name for the column set.
5. In the top box, select each column heading you want in the column set and click **Add**.

The box shows a list that represents all of the inventory data currently in the database. Drill down this list to select an attribute to display in the query results list. Remember to select attributes that will help you identify the clients returned in the query. If you cannot find attributes you want to display, you can add them in the [Custom attributes](#) dialog. However, these attributes must be assigned to machines before they appear in the query dialog.

Note: If you select an attribute in the database that has a 1:* relationship, you will get duplicate entries for the device. When you select attributes with a 1:1 relationship (only one possible attribute, like Computer.System.Asset Tag), you will not receive duplicate entries.

6. To change the order of the columns, select a column heading and click **Move up** or **Move down**.
7. To remove a column, select it in the lower box and click **Remove**.
8. To change the heading displayed for a column, select the heading in the lower box, click **Edit**, make your modifications, and press **Enter**. The following extended characters are not supported: < , > , ' , " , !.
9. Click **OK** to save the column set.
10. To use the custom column set when you view the **All devices** list, select it and click **Set as current column set** on the toolbar.

To edit a custom column set

1. In the left navigation pane, click **Preferences**.
2. Click the **Custom columns** tab.
3. Select the custom column set and click **Edit**.
4. In the top box, select a column heading and click **Add** to add the column (see notes under step 5 above).
5. To remove a column, select it in the lower box and click **Remove**.

6. To change the heading displayed for a column, select the heading in the lower box, click **Edit**, make your modifications, and press **Enter**. The following extended characters are not supported: < , > , ' , " , !.
7. To change the order of the columns, select a column heading and click **Move up** or **Move down**.
8. Click **OK** to save your changes.

Custom attributes

Attributes are characteristics or properties that belong to a device. The more attributes a device has in the database, the easier it becomes to uniquely identify the device. You can create custom attributes only if you are using LANDesk® Server Manager with the Administrator right. If custom attributes have been created and added to the core database, you can assign values for those attributes to a managed device. If no custom attributes have been added to the core database, the **Assign attributes** option does not appear under the **Actions** tab..

To assign custom attributes to devices

1. In the **All devices** list, select a device or multiple devices.
2. In the bottom pane, click the **Actions** tab.
3. Select **Assign attributes** from the left pane.
4. Each Attribute Name has a drop-down list of values. Select a value from the drop-down list for the attribute name, and repeat as necessary. Click **Selected devices**.
5. Click **Assign**, then click **OK**.

You can also assign custom attributes to multiple devices that you have targeted. If you have devices in the Target list, click **Targeted devices** in step 4 above.

Page settings

Use **Page settings** to set display preferences for pages listing devices or displaying graphics.

1. In the left navigation pane, click **Preferences**.
2. Click the **Page settings** tab.
3. In the **Graph type** drop-down, select the type of graph you want to display in **Reports**.
4. In the **Items/page** box, type the maximum number of items you want to display in each page that uses pagination. The value must be 500 items or less.

Beginner mode

You can show text alongside the buttons on toolbars, thus helping new users with feature identification. If this option is not selected, only icons appear on the toolbars. The icons will display text on mouseover.

1. To show text alongside the buttons on toolbars, click the **Show text on toolbars** checkbox.
2. Click **Update**.

Viewing the Server information console

Use the Server information console to view top-level summary information about a device, view system information like CPU or fan information, monitor the health status and thresholds of key components of a device, manage vulnerabilities, and power on, power off, or reboot a device. The Server information console has the following sections listed under the left navigation pane.

- [System information](#)
- [Software updates](#)
- [Monitoring](#)
- [Rulesets](#)
- [Power options](#)
- [Hardware configuration](#)

In order to view the server information console for a device, you must first deploy the Standard management agent on that device (see [Configuring agents](#)). Also, the device must be rebooted after the agent is installed for the server information console to function correctly. This reboot is required when you install the agent on the core server as well as on managed devices.

To view the Server information console

1. In the **My devices** view, double-click the device name.

The console opens in a new browser window and shows the **Health summary** page by default.

2. Click buttons in the left navigation pane to view server information and use the available tools.

System information

System information contains summary data about the health of the device, as well as information about hardware and software, system logs, and other data such as asset and network information.

Health summary

The **Health summary** page provides a quick overview of system health for this device. You can see at a glance whether selected hardware elements are functioning correctly and whether there are potential problems that may need to be addressed.

When any of the health elements are in a warning or critical state, the corresponding button contains a yellow (warning) or red (critical) icon indicating that a problem exists. Click the button to view a description of the event that caused the warning or critical alert.

System summary

Use the **System summary** page to view important information about the selected device. The information listed on the page can include the following, depending on the type of hardware and software configured on the device.

- **Health:** The overall health of the device as defined by the conditions and parameters you set.
- **Type:** The type of the device, such as print, application, or database.
- **Manufacturer:** The maker of the device.
- **Model:** The model of the device.
- **BIOS version:** The version of the device's BIOS.
- **Operating system:** The device's operating system.
- **OS version:** The version number of the operating system.
- **CPU:** The manufacturer, model, and speed of the device's processor.
- **Vulnerability scanner:** The version of the vulnerability scanner.
- **Remote control:** The version of the remote control agent.
- **Software distribution:** The version of the software distribution agent.
- **Inventory scanner:** The version of the inventory scanner.
- **IPMI type, IPMI version:** The type and version number of IPMI the device is using.
- **SDR version:** The version of the Sensor Data Record in the device's BMC.
- **BMC version:** The version of the device's Baseboard Management Controller.
- **Kernel:** For Linux devices, the version number of the kernel installed.
- **Monitoring:** The version number of the monitoring agent on the device.
- **CPU usage:** The percentage of the processor currently being used.
- **Physical memory used*:** The percentage of total physical memory used on the device.
- **Virtual memory used*:** The percentage of total virtual memory used on the device.
- **Last reboot*:** The date and time the device was last rebooted (in the time zone of the database).
- **Drive:** The drives on the device with the total size of the drive and percentage of space used.

This information is taken from the registry in Windows or from configuration files in Linux.

*This information appears when an agent has been installed on the device.

Hardware

Use the **Hardware** page to view details about the device's hardware configuration. Items in the **Hardware** list are grouped in the following categories. Note that not all categories appear for all devices. For example, if the device does not have fan and temperature sensors, the **Cooling** category does not appear in this list.

- **CPU:** processors and cache
- **Storage:** logical drives, physical drives, removable media, and storage adapters
- **Memory:** usage information and memory modules
- **Chassis:** the server's chassis; view whether the case is open or closed
- **Input devices:** keyboard, mouse, and other devices
- **Motherboard:** motherboard, expansion slots, and BIOS

- **Cooling:** fans and temperature sensors
- **Power:** power supplies and voltage

Setting alerting thresholds for hardware items

Some items in the **Hardware** list represent data from sensors in the device, such as temperature sensors. If a managed device contains components with supported sensors, you can change the sensor readings that will trigger an alert. For example, a CPU temperature sensor can have lower and upper temperature readings that trigger warning and critical alerts. Thresholds are generally based on manufacturer's recommended settings, but you can change the upper and lower settings using the **Thresholds** dialog box.

1. In the server information console, click **System information**.
2. Expand the **Hardware** folder and drill down to find the hardware element you want (such as **Cooling | Temperatures**).
3. In the list of sensors, double-click the sensor you want to set thresholds for.
4. Type values in the lower and/or upper threshold text boxes, or drag the sliders on the trackbar to the left or right to change the values.
5. Click **Update** to save your changes.
6. To return to the original values for the thresholds, click **Restore defaults**.

Logs

The Logs page displays local system logs, the System Events Log (SEL) for IPMI devices, and an alert log.

Local logs, such as the Application, Security, and System logs, don't have buttons to clear the log from the console, but you can view and clear the logs using Windows Computer Management.

If this device's BIOS has the ability to clear the SMBIOS log, click the **Clear log** button to remove all log entries. This button is unavailable if the BIOS does not support this action.

Software

The **Software** page shows summary information about processes, services, and packages on this device, as well as a list of current environment variables.

- **Processes:** displays processes that are running; select a process and click **Kill process** to terminate it
- **Services:** displays services available on the device and their status; select a service and click **Stop**, **Start**, or **Restart** to make changes
- **Packages:** lists installed packages with version numbers and vendor name
- **Environment:** lists the environment variables currently set on the device

Other

The **Other** page shows asset information and a summary of network hardware and connections.

- **Asset information:** view and edit asset management information, such as location and asset tag number; you can also view system information such as serial number, manufacturer, and chassis type
- **Network information:** view a list of network hardware installed, statistics of network activity, a configuration summary (including IP address, default gateway address, and WINS, DHCP, and DNS server information), and a list of current network connections (mapped drives)

Software updates

Use the **Software updates** page to scan for detected vulnerabilities on the selected device.

To check for detected vulnerabilities

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in a new browser window.
2. In the left navigation pane, click **Software updates**.

Column descriptions

- **ID:** Identifies the vulnerability with a unique, vendor-defined alphanumeric code.
- **Severity:** Indicates the severity level of the vulnerability. Possible severity levels include: Service Pack, Critical, High, Medium, Low, Not Applicable, and Unknown.
- **Title:** Describes the nature or target of the vulnerability in a brief text string.
- **Language:** Indicates the language of the OS affected by the vulnerability.
- **Date published:** Indicates the date the vulnerability was published by the vendor.
- **Silent install:** Indicates whether the vulnerability's associated patch file installs silently (without user interaction). Some vulnerabilities may have more than one patch. If any of a vulnerability's patches don't install silently, the vulnerability's **Silent install** attribute is **No**.
- **Repairable:** Indicates whether the vulnerability can be repaired through patch file deployment and installation. Possible values are: Yes, No, and Some (for a vulnerability that includes multiple detection rules and not all detected vulnerabilities can be repaired).

Monitoring

Use **Monitoring** to view performance counters and graphs, and to set thresholds for device components. For more information about this feature, see the [Device monitoring](#) section.

To select a performance counter to monitor

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in a new browser window.
2. In the left navigation pane, click **Monitoring**.
3. Click the **Performance counter settings** tab.
4. From the **Objects** column, select the object you want to monitor.
5. From the **Instances** column, select the instance of the object you want to monitor, if applicable.

6. From the **Counters** column, select the specific counter you want to monitor.
7. Specify the polling frequency and the number of days to keep the counter history.
8. In the **Alert after counter is out of range** text box, specify the number of times the counter will be allowed to cross the thresholds before an alert is generated.
9. Specify upper and/or lower thresholds.
10. Click **Apply**.

To view a performance graph for a monitored counter

1. Click the **Active performance counters** tab.
2. Select a counter from the list.
3. In the **Counters** list, select the counter for which you want to see a performance graph.
4. Select **View real-time data** to display a graph of current performance, or select **View historical data** to display a graph showing performance over the period you specified (Keep history when selecting the counter.)

On the performance graph, the horizontal axis represents time that has passed. The vertical axis represents the units you are measuring, such as bytes per second (when monitoring file transfers, for example), percentage (when monitoring percentage of the CPU that is in use), or bytes available (when monitoring hard drive space).

Rulesets

Use the **Rulesets** page to view a list of the alert and monitoring ruleset configurations assigned to the selected device, and to view the details of each alert.

To view alert rulesets

1. In the **My devices** view, double-click the device you want to configure. The console opens in a new browser window.
2. In the left navigation pane, click **Rulesets**.
3. Click the **Alerting rulesets** tab.

The following text describes the details provided about each alert. For more information on modifying these details, see [Using alerts](#).

- **When state reaches:** When the state of the alert reaches the displayed state, an alert will be generated.
- **Affects health:** If the alert state reaches the specified threshold, the state affects the overall health state of the device. The selection of an alert affecting health is determined in the Alert rulesets dialog.
- **Ruleset name:** The name of the alert ruleset, as defined in the [Alert rulesets](#) dialog.
- **Alert type:** The kind of alert to be generated, such as an e-mail, an SNMP trap, or executing a program.
- **Action configuration:** The action that occurs when the alert is generated, as defined in the [Action rulesets](#) dialog.
- **Alert handler:** The handler associated with the alert, such as an e-mail handler.
- **Instance:** Indicates the specific source of the alert.

To view monitoring rulesets

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in a new browser window.
2. In the left navigation pane, click **Rulesets**.
3. Click the **Monitoring rulesets** tab.

The following text describes the details provided about each monitoring ruleset. For more information on modifying these details see [About monitoring](#).

- **Name:** The name of the ruleset configuration, as defined in the [Monitoring](#) page.
- **Ruleset name:** Whether the ruleset is a default ruleset or not.
- **Enabled:** The ruleset has or has not been enabled to execute on the device.
- **Warning threshold:** The threshold at which, if exceeded, the device will send a warning message to the core.
- **Critical threshold:** The threshold at which, if exceeded, the device will send a critical message to the core.
- **Check every:** The frequency at which the item is monitored.




Power options

Power options lets you power off, reboot and, in the case of managed IPMI and Intel AMT devices, power on remote devices. In the case of non-IPMI servers, the server must have the LANDesk agent deployed to it in order to execute the reboot and power off functions.

With IPMI and Intel AMT devices, you must have configured the correct credentials to execute the power on/power off and reboot features. If IPMI or Intel AMT devices have the LANDesk agent deployed, then you can execute the power off and reboot features without the IPMI or Intel AMT credentials. To configure BMC credentials for IPMI devices, or Intel AMT device credentials, use the Configure Services utility (see [Configuring services and credentials](#)).

To use power options on the selected device

1. In the **My devices** view, double-click the device you want to configure. The console opens in a new browser window.
2. In the left navigation pane, click **Power options**.
3. Select from the following options:

-  Reboot
-  Power off
-  Power on

Hardware configuration

The **Hardware configuration** tool lets you configure options for devices with IPMI or Intel* AMT capabilities. This tool and any options are only displayed on devices with the corresponding

hardware (for example, IPMI options are displayed only if the device is recognized as an IPMI device).

You can generate IDs for provisioning Intel AMT devices, view the generated IDs, and change configuration options related to provisioning your Intel AMT devices. You can also define circuit breaker policies, which detect and block suspicious network activity on the devices, and you can enable Agent Presence monitoring to ensure that management agents on your devices are continually running. (For more information, see [Intel AMT support](#).)

For IPMI devices, you can customize configuration options such as the watchdog timer, power options, and BMC user settings. You can also configure the use of a LAN channel or serial over LAN to maintain out-of-band communication with an IPMI device. (For more information, see [IPMI BMC configuration](#).)

For devices that have a Dell* DRAC (Remote Access Controller), you can view Dell DRAC logs and edit usernames for access to the OpenManage Server Administrator. (For more information, see [Managing Dell DRAC devices](#).)

Managing Intel* AMT devices

After an Intel* AMT device has been discovered and added to the core database to be managed, it can be managed in limited ways even if the device does not have a LANDesk agent installed. (See [Discovering Intel* AMT devices](#) for information on discovering devices and moving them to the core database).

The following table lists the management options available for a device that has Intel AMT only compared with a device that has Intel AMT and a System Manager management agent installed.

	Intel AMT only	Intel AMT and agent	Agent only
Inventory	summary	X	X
Event log	X	X	X
Remote boot manager	X	X	
Disable OS network		X	
Enable OS network		X	
Force vulscan on reboot		X	
Inventory history		X	X

	Intel AMT only	Intel AMT and agent	Agent only
Remote control		X	X
Chat		X	X
File transfer		X	X
Remote execute		X	X
Wake up		X	X
Shut down		X	X
Reboot		X	X
Inventory scan		X	X
Scheduled tasks and policies	limited	X	X
Group options		X	X
Run inventory report		X	X
Intel AMT alerting		X	X

To view the Intel AMT inventory summary for a device

1. Double-click the device in the **All devices** list.
2. In the server information console, click **Intel AMT options**.
3. Click **Inventory summary**.

The summary shows the device's GUID, product and manufacturer, serial number, and BIOS, processor, and memory summaries, and the Intel AMT version number. If any information is missing you can refresh the data by clicking **Update inventory**.

Accessing devices provisioned with Enterprise mode

When you provision an Intel AMT device in Enterprise mode, the core server installs a certificate on the device for secure communication. If the device is to be managed by another core server, it must be unprovisioned and then re-provisioned by the new core server. If not, the device's Intel AMT access will not respond because the new core server does not have a matching certificate. Similarly, if any other computer attempts to access the Intel AMT functionality on the device, it will not succeed because it does not have a matching certificate. (See [Intel* AMT support](#) for information about provisioning modes).

Intel AMT event log

System Manager provides a view of the event log that Intel AMT devices generate. The settings determine what events are captured in this log. You can view the date/time of the event, the source of the event (Entity column), a description, and the severity as determined by the Intel AMT settings (Critical or Non-Critical). You can also export the log data in comma-separated value (CSV format).

To view the Intel AMT event log

1. Double-click the device in the **All devices** list.
2. In the server information console, click **System information**.
3. Expand **Logs** and click **Intel AMT log**.
4. To export the log to a CSV format file, click the **Export** button on the toolbar and specify a location to save the file to.
5. To clear all data in the log, click the **Purge log** button on the toolbar.
6. To update the log entries, click the **Refresh log** button on the toolbar.

Intel AMT power options

System Manager contains options to power on and off Intel AMT devices. These options can be used even when a device's operating system is not responding, as long as the device is connected to the network and has standby power.

When System Manager initiates power option commands, in some cases it is not possible to verify that the commands are supported on the hardware receiving the command. Some devices with Intel AMT may not support all power option features (for example, a device may support IDE-R reboot from CD but not from a floppy). Consult the hardware vendor's documentation if it appears that a power option is not working with a particular device. You may also check for any firmware or BIOS upgrades from Intel for the device if power options do not work as expected.

You can simply turn on or off the device's power, or you can reboot and specify how the device is rebooted. The options are described in the table below.

Power off	Shuts down the power on the device
Power on	Turns on the power on the device

Reboot	Cycles the power off and on again on the device
Normal boot	Starts up the device using whatever boot sequence is set as the default on the device
Boot from local hard drive	Forces a boot from the device's hard drive regardless of the default boot mode on the device
Boot from local CD/DVD drive	Forces a boot from the device's CD or DVD drive regardless of the default boot mode on the device
PXE boot	When restarted, the PXE-enabled device searches for a PXE server on the network; if found, a PXE boot session is initiated on the device
IDE-R boot	Reboots the device using the IDE redirection option selected (see below)
Enter BIOS setup	When the device is booted, it allows the user to enter the BIOS setup
Show console redirection window	When the device is booted, it starts in serial over LAN mode to display a console redirection window
IDE redirection: Reboot from floppy	When the device is booted, it starts from the floppy disk drive or image that are specified (floppy image files must be in .img format; see note below)
IDE redirection: Reboot from CD/DVD	When the device is booted, it starts from the CD drive or image that are specified (CD image files must be in .iso format; see note below)
IDE redirection: Reboot from specified image file	When the device is booted, it starts from the image file that is specified (see note below)

To use Intel AMT power options

1. Double-click the device in the **All devices** list.
2. In the server information console, click **Power options**.
3. Select a power command. If you select **Reboot**, select a boot option.
4. Click **Send** to initiate the command.

Notes on using IDE redirection options

To use IDE redirection options, both a boot floppy or floppy image file and a boot CD/DVD or CD/DVD image file must be specified. Floppy image files must be in .img format, and CD image files must be in .iso format. Some BIOSes may require the CD image to be located on a hard drive.

Intel AMT normally remembers the last IDE-R settings, but System Manager clears the settings after 45 seconds, so on subsequent boots it will not restart the IDE-R feature. The IDE-R session on an Intel AMT device lasts 6 hours or until the System Manager console is turned off. Any IDE-R operation still in progress after 6 hours will be terminated.

Forcing a vulnerability scan and disabling network access on Intel AMT machines

When an Intel AMT-configured device has the LANDesk agent installed, the agent includes functionality that can help resolve problems with malicious software or other issues that prevent you from accessing the device.

The amtmon.exe service is installed with the LANDesk agent. When this service is running on a device, you can force a vulnerability scan at the next reboot to attempt to identify any malicious software on the device. If communication with the device fails, you can disable the device's network connection even if the OS is not functional, such as when malicious software has disabled the OS by consuming all CPU cycles. By disabling the network connection you can prevent the device from sending unwanted packets through the network.

When the LANDesk agent is installed on an Intel AMT device, the following options are available on the **Intel AMT options** page:

- **Operating system network connection:** click **Disable** to disable the OS network stack to stop network access; click **Enable** to enable OS network access if it has been disabled.
- **Scan for vulnerabilities after reboot:** forces the vulnerability scanner to run the next time the device reboots.

When a device is not responding or may have malicious software running on it, the recommended use case is to first run a vulnerability scan on the next reboot to attempt to identify the problem. If the problem continues and the machine is infecting/attacking the network, or if you can't access the device, you have the option to disable the OS NIC.

To force a vulnerability scan after a reboot

1. Double-click the device in the **All devices** list.

2. In the device console window, click **Intel AMT options**.
3. Click **Configuration options**, then click **Scan**. A message appears on the device stating that a scan will be run the next time it reboots.
4. To shut down or reboot the device, use the Intel AMT remote boot manager features above.

To disable or enable the network connection on an unresponsive device

1. Double-click the device in the **All devices** list.
2. In the device console window, click **Intel AMT options**.
3. To disable the device's network card to stop communication with other devices on the network, click **Disable**. When the network connection is disabled, a message appears on the device stating that the network card has been disabled.
4. If the device is safe to connect to the network again, click **Enable**. When the connection is restored, a message appears on the device stating that the network card is enabled again.

Opening the Intel AMT Configuration Screen

System Manager includes a link that lets you open the Intel AMT Configuration Screen. This is an interface provided by Intel to view device status, hardware information, the Intel AMT event log, remote boot settings, and network settings. It also lets you add and edit Intel AMT user accounts for the device. The window that displays this screen is separate from the System Manager console, and any questions you may have about your use of this interface should be addressed to the device manufacturer's technical support.

To open the Intel AMT Configuration Screen

1. Double-click the device in the **All devices** list.
2. In the device console window, click **Intel AMT options**.
3. Click **Intel AMT console**, then click **Launch Intel AMT web console**.

Role-based administration

About role-based administration

Use role-based administration to configure user access to product tools and other devices based on their administrative role in your system. With role-based administration, you assign scope to determine the devices a user can view and manage, and rights to determine the tasks they can perform.

Administrators (users with the Administrator right) can access the role-based administration tools by clicking **Users** in the left navigation pane.

Role-based administration lets you assign product users special administrative roles based on their rights and scope. *Rights* determine the product tools and features a user can see and utilize. *Scope* determines the range of devices a user can see and manage.

You can create roles based on users' responsibilities, the management tasks you want them to be able to perform, and the devices you want them to be able to see, access, and manage. Access to devices can be restricted to a geographic location like a country, region, state, city or even a single office or department. Or, access can be restricted to a particular platform, processor type, or some other device hardware or software attribute. With role-based administration, it's completely up to you how many different roles you want to create, which users can act in those roles, and how big or small their scope of device access should be.

Example administrative roles

The table below lists some of the possible administrative roles you might want to implement, the common tasks that user would perform, and the rights that user would need in order to function effectively in that role.

Role	Tasks	Required rights
Administrator	Configure core servers, manage users, configure alerts, integrate other company products, etc. (Of course, administrators with full rights can perform any management tasks.)	Administrator (all rights implied)
Asset manager	Discover devices, configure devices, run the inventory scanner, enable inventory history tracking, etc.	Device discovery, Software distribution, and Public query management
Reporting manager	Run predefined reports, print reports, etc.	Reports (required for all reports)

These are just example roles. Role-based administration is flexible enough to let you create as many custom roles as you need. You can assign the same few rights to different users but restrict their access to a limited set of devices with a narrow scope. Even an administrator can be restricted by scope, essentially making them an administrator over a specific geographic region or type of managed device. How you take advantage of role-based administration depends on your network and staffing resources, as well as your particular needs.

To implement and enforce role-based administration, simply designate current local Windows users, or create and add new local Windows users, as product users, add the users to the Management Suite user group, and then assign the necessary rights (to product features) and scope (to managed devices). Follow the procedures below:

Understanding rights

Rights provide access to specific tools and features. Users must have the necessary right (or rights) to perform corresponding tasks. For example, in order to remote control devices in their scope, a user must have the remote control right. If you have multiple LANDesk management products installed, rights can be assigned to users from any console and are effective in all consoles.

When a right is not assigned to a user, tools associated with that right are not visible to that user in the product console. For example, if a user is not given the reports right, the reports item doesn't appear in the left navigation pane. The table below shows which rights are required for the tool to display for a user.

Tool	Rights needed to display in left navigation pane
My devices	Basic Web console
Agent configuration	Administrator
Alerting	Alerting and monitoring
Device discovery	Device discovery
Monitoring	Alerting and monitoring
Queries	Basic Web console, Public query management, Reports
Reports	Reports, Patch management
Scheduled tasks	Device discovery

Tool	Rights needed to display in left navigation pane
Scripts	Patch management
Users	Administrator
Software updates	Patch management
Preferences	Basic Web console
Hardware configuration	Administrator

See the descriptions below to learn more about each product right and how rights can be used to create administrative roles.

Scope controls access to devices

When using the features allowed by these rights, users will always be limited by their scope (the devices they can see and manipulate).

Administrator

The Administrator right provides full access to all of the product tools (however, use of these tools is still limited to the devices included in the administrator's scope).

This is the default right for a newly added user, unless you've modified the settings for the Default Template User.

The Administrator right provides users the ability to:

- See and access the **Users** tool in the left navigation pane
- See product licensing in **Preferences** in the left navigation pane.
- Perform all of the product tasks allowed by the other rights listed below

Deleting the Administrator user is not recommended. If you are the last administrator to log into a particular LDSM console, and go into Windows Computer Management and remove the Administrator user from the Management Suite group, you may encounter problems upon re-entering the console. You will still be logged in as Administrator for next 20 minutes or so (which is the default session timeout), but when you initiate any action or refresh the console browser (F5 key) in which you are logged in as Administrator, any rights belonging solely to the Administrator will no longer be accessible. It is recommended that you do not delete the last Administrator user under any circumstance.

Note on rights and tools

The Administrator right is exclusively associated with the **Users** tool. If a user does not have the Administrator right, this tool will not appear in the console.

All of the tools in the product console are associated with a corresponding right (as described below).

Device discovery

The Device discovery right provides users the ability to:

- Find devices on the network that haven't submitted an inventory scan to the product core database through many ways, such as a network scan, Standard management agent discovery, and IPMI discovery
- Schedule periodic discoveries
- Move devices from Discovered to Managed

Public query management

The Public query management right provides users the ability to:

- Create queries available to all users
- Ability to create or delete public queries
- Ability to modify/edit existing public queries

Reports

The reports right provides users the ability to:

- See and access the **Reports** tool in the left navigation pane
- Run predefined reports

Patch management

The Patch management right is specific to the vulnerability scanning feature. For more information, see "Using the software updates tool."

Basic Web console

The Basic Web console right provides users with the ability to use the features associated with the right. The features are listed below, along with any exceptions within the feature.

- **My devices** (the right doesn't allow for the updating of public groups, or deleting devices under the **Actions** tab)
- Change preferences (but not custom attributes)

Alerting and monitoring

The Alerting and monitoring right provides users with the ability to:

- Monitor the performance of various system and OS components, such as drives, processors, memory, processes, bytes/sec transferred by the system's Web server, and so forth
- Track the exact health of all managed devices
- Customize alerts to be sent by severity level (Critical, Warning, Informational, OK, Unknown) or threshold (for example, if the hard disk usage exceeds 90% of hard disk capacity)
- Choose the action to be taken if an alert exceeds a threshold (by adding information to the log, e-mailing a notice, running a program on the core or an individual device, or sending an SNMP trap to an SNMP management console on the network)

Adding product users

Product users are users who can log in to the product console and perform specific tasks for specific devices on the network.

Product users are not actually created in the console. Instead, users appear in the **Users** tab (in the left navigation pane, click **Users** after they have been added to the LANDesk Management Suite group in the Windows users environment on the core server). The **Users** group shows all of the users currently residing in the LANDesk Management Suite group on the core server.

There are two default users in the **Users** group:

- **Default Template User:** This user is basically a template of user properties (rights and scope) that is used to configure new users when they are added to the LANDesk Management Suite group. In other words, when you add a user to that group in the Windows environment, the user inherits the rights and scope currently defined in the Default Template User properties. If the Default Template User has all rights selected and the Default All Machines Scope selected, any new user placed in the LANDesk Management Suite group will be added to the **Users** group with rights to all of the product tools and access to all devices.

You can change the property settings for the Default Template User by right-clicking it and clicking **Edit rights**. For example, if you want to add a large number of users at once, but do not want them to have access to all of the tools or devices, change the settings for the Default Template User first, then add the users to the LANDesk Management Suite group (see steps below).

The Default Template User cannot be removed.

- **Default Administrator:** This is the administrative user who was logged in to the server when this product's core was installed.

When you add a user to the LANDesk Management Suite group in Windows, the user is automatically read into the **All Users** group in the **Users** window, inheriting the same rights and scope as the current Default Template User. The user's name, scope, and rights are displayed.

If you remove a user from the LANDesk Management Suite group in the Windows users environment, the user is no longer active and can be deleted from the **Users** group. The user's account still exists on your server and can be added back to LANDesk Management Suite group at any time. Also, the user's subgroups under **User Devices**, **User Queries**, **User Reports**, and **User Scripts** are preserved so that you can restore the user without losing their data, and so that you can copy data to other users.

To refresh the **Users** list to display any newly added users, click **Users** and click the **Refresh** button on your browser.

To add a user or domain group to the LANDesk Management Suite group

1. Navigate to the server's **Administrative Tools | Computer Management | Local Users and Groups | Groups** utility.
2. Right-click the **LANDesk Management Suite** group, and then click **Add to group**.
3. Click **Add**, then type or select a user (or users) from the list.
4. Click **Add**, and then **OK**.

Note: You can also add a user to the LANDesk Management Suite group by right-clicking the user account in the Users list, clicking **Properties | Member Of**, and then clicking **Add** to select the group and add the user.

If user accounts do not already exist in Windows, you must first create them on the server.

To create a new user account

1. Navigate to the server's **Administrative Tools | Computer Management | Local Users and Group | Users** utility.
2. Right-click **Users**, and then click **New User**.
3. In the New User dialog, enter a name and password.
4. Specify password settings.
5. Click **Create**. The New User dialog remains open so that you can create additional users.
6. Click **Close** to exit the dialog.
7. Add the user to the LANDesk Management Suite group to have them appear in the Users group in the console.

You can now assign your product users rights and scope.

Creating scopes

A scope defines the devices that can be viewed and managed by a product user. If you have multiple LANDesk management products installed, scopes can be assigned to users from any console and are effective in all consoles.

A scope can be as large or small as you want, encompassing all of the managed devices scanned into a core database, just a single device, or no devices. This flexibility, combined with modularized tool access, is what makes role-based administration such a versatile management feature.

Default scopes

Role-based administration includes two default scopes. These two predefined scopes can be useful when configuring the user properties of the default template user.

- **(Default) No Machines Scope:** Excludes all devices in the database.
- **(Default) All Machines Scope:** Includes all devices in the database.

You can't edit or remove the default scopes.

Custom scopes

You can create the following types of custom scopes and assign them to users:

- **Query-based:** Controls access to only those devices that match a custom query search. You can select an existing query, or create new queries from the **Query** dialog, to define a scope. For more information on creating queries, see "[Creating database queries](#)."
- **Group-based:** Controls access to only those devices located in the selected group. You can select groups from the **Group scope properties** dialog to define a scope.

You can assign more than one scope to any of the users. When multiple scopes are assigned to a user, the cumulative effective scope (i.e., the complete range of devices that can be accessed and managed as a result of the combination of assigned scopes is a simple composite.)

You can customize a user's effective scope by adding and removing scopes at any time. All types of scopes can be used together.

To create a scope

1. In the left navigation pane, click **Users**.
2. On the **Scope** tab, click the **New query scope** or the **New group scope** toolbar button.
3. Enter a name for the new scope.
4. If you selected query-based, choose an existing query, or click **Define** to create a new query. Click **OK**.
5. If you selected group-based, choose a group, then click **OK**.
6. Click **OK** to save the scope and close the dialog.

Assigning rights and scope to users

- [About the User rights/scope dialog](#)
- [About the Remote control settings dialog](#)

Once you've added product users, learned about rights and how they control access to features and tools, and created device scopes to allow or restrict access to managed devices, the next step in establishing role-based administration is to assign the appropriate rights and a scope to each user.

You can modify a user's rights and scope at any time.

If you modify a user's rights or scope, those changes will take effect the next time that user logs into the console.

To assign rights and scope to a user

1. In the left navigation pane, click **Users**.
2. Expand the list of users to view all of the users that are currently a member of the LANDesk Management Suite group in the core server's Windows NT environment.

This list displays user names and assigned rights (a check character indicates the right is enabled or active).

3. Right-click a user, and then click **Edit**.
4. In the **User rights/scope** dialog, check or clear rights as desired.
5. Click the **Scope** tab and select a scope from the **Assigned scopes list**.
6. Click **Apply**.

The new rights display next to the user's name in the list and will take effect the next time the user connects to the core server.

To delete a scope

1. In the left navigation pane, click **Users**.
2. In the **Scope** tab, click the scope you want to delete and click **Delete**. Click **OK**.

Exercise caution when deleting scopes. The users assigned to them will be able to access rights previously prohibited by the scope.

About the User rights/scopes dialog

Use this dialog to view and modify a user's assigned rights and scope. Open the dialog by selecting a user and clicking **Edit**.

Rights tab: Lists the rights assigned to the user.

- **Administrator**
- **Device discovery**
- **Public query management**
- **Reports**
- **Patch management**
- **Basic Web console**
- **Alerting and monitoring**

Scope tab: Lists the scopes assigned to the user.

- **Assigned scopes:** Identifies the user's current scopes.
- **Add:** Opens the **Add scope** dialog where you can select a scope to add to the user.
- **Remove:** Deletes the selected scope.
- **Cancel:** Closes the dialog without saving changes.

Device discovery

Using device discovery

Device discovery finds devices on your network that do not have agents of the discovering core installed and have not submitted an inventory scan to the same core database. Device discovery has multiple ways of finding devices on your network.

- **Network scan:** Looks for computers by doing an ICMP ping sweep. This is the most thorough search, but it can take longer (if IP fingerprinting is used) . You can limit the search to certain IP and subnet ranges. By default this option uses NetBIOS to gather information about the device. You can also select IP Fingerprinting, which also provides the OS type (in most cases) . The network scan option also has a **Use SNMP** option, where you can configure the scan to use SNMP for SNMP devices such as some printers.
- **CBA discovery:** Looks for the Standard management agent (formerly known as the common base agent [CBA] in Management Suite) on computers. This option discovers computers that have been managed by Server Manager, System Manager, and so on. You can select the PDS2 option to discover devices using the older LANDesk PDS2 agent. CBA discovery is not supported for Linux machines, but if you choose PDS2, Linux machines with an agent installed can be discovered.
- **IPMI:** Looks for servers enabled with [Intelligent Platform Management Interface](#), which allows you to access the baseboard management controller (BMC) regardless of whether the server is turned on or not, or what state the OS may be in.
- **Server chassis:** Looks for blade server chassis management modules (CMMs). The blades in the server chassis are detected as normal servers.
- **Intel* AMT:** Looks for devices with Intel Active Management Technology (version 1) , which allows you limited management features regardless of whether the server is turned on or not, or what state the OS may be in.

Device discovery tries to discover basic information about each device. Not all of the information below is available for every device.

- **Node name:** The discovered device name, if available.
- **IP address:** The discovered IP address.
- **Subnet mask:** The discovered subnet mask.
- **Category:** The device discovery group the device belongs to.
- **OS name:** The discovered operating system description, if available.

When device discovery finds a device for the first time, it looks in the core database to see if that device's IP address and name are already in the database in the **My devices** list. A device in the **Unmanaged** list will be rediscovered and potentially yield more data. If there's a match, device discovery ignores the device. If there isn't a match, device discovery adds the device to the **Unmanaged** device table. Devices in the **Unmanaged** table don't use a System Manager license. A device is considered managed once it sends an inventory scan to the core database. Once you move a device to the **All devices** group, it no longer displays in the **Discovered devices** list.

IPMI devices must have a BMC (baseboard management controller) that is configured in order to be discovered as IPMI devices and to use full IPMI functionality. If the BMC is not configured, the device can be discovered as a computer. You can then add the device to the list of managed devices and run the Hardware configuration feature to configure the BMC password. The device's IPMI functionality will then be recognized by this product. Note that the BMC's IP address is not necessarily the same as the OS's IP address, and therefore might not be able to take a direct agent push to the BMC IP address. A rediscovery of standard IPs may be required to get a standard agent to push down to the BMC's IP. The BMC IP should be able to take an IP agent push..

Devices enabled with Intel* AMT (version 1) should be configured with an Intel AMT username and password to be recognized and discovered as Intel AMT devices. After being discovered, you can run the Hardware configuration feature to configure the Intel AMT settings and provision the device in Small business mode or in the secure Enterprise mode.

To automate device discovery, you can schedule discoveries to occur periodically. For example, you could divide your network by subnets and schedule a ping sweep for a different subnet each night. In all discoveries, the core server does the discovering.

To discover and manage devices on your network, complete the following tasks:

- Create discovery configurations
- Schedule and run the discovery
- View discovered devices
- Move discovered devices to the **My devices** list

Using unmanaged device discovery with firewalled devices

Be aware that unmanaged device discovery usually can't discover devices that use a firewall, such as Windows Firewall, unless you manually configure the firewall. You must open the following ports. To change these settings, access the Windows Firewall through the Windows Control Panel.

Managed servers:

- File and Printer Sharing: TCP 139, 445; UDP 137,138 (Push won't work without this)
- Software distribution: TCP 9595 (Push won't work without this)
- Advanced - ICMP: "Allow incoming echo request" (cannot be discovered if this is not enabled)

Core server:

- Inventory: 5007
- Remote control: 9535

Creating discovery configurations

Use the **Discovery configurations** tab to create new discovery configurations, edit and delete existing configurations, and schedule a configuration for discovery. Each discovery configuration consists of a descriptive name, the IP ranges to scan, and the discovery type.

Once you create a configuration, use the **Schedule discovery** dialog to configure when it will run.

1. In the left navigation pane, click **Device discovery**.
2. In the **Discovery configurations** tab, click the **New** button.
3. Fill in the fields described below. When you have finished, click the **Add** button, and click **OK**.

The text below describes the parts of the **Discovery configuration** dialog box.

- **Configuration name:** Type a name for this configuration. Give the configuration a meaningful name so you can easily remember the configuration. The configuration can be up to 255 characters long, and should not contain the following characters: ", +, #, & or %. The configuration name will not display after the use of any of these characters.
- **Standard network scan:** Looks for devices by sending ICMP packets to IP addresses in the range you specify. This is the most thorough search, but also the slowest. By default, this option uses NetBIOS to gather information about the device.

The network scan option has an **IP fingerprint** option where device discovery tries to discover the OS type through TCP packet responses. The IP fingerprint option slows down the discovery somewhat.

The network scan option also has a **Use SNMP** option, where you can configure the scan to use SNMP. Click **Configure** to enter information about your SNMP configuration. For more information, see [Configuring SNMP scans](#).

- **LANDesk CBA discovery:** Looks for the standard management agent (formerly known as the common base agent [CBA] in Management Suite on devices). The standard management agent allows the core server to discover and communicate with clients on the network. This option discovers devices that have product agents on them. Routers block standard management agent and PDS2 traffic. In order to run a standard CBA discovery across multiple subnets, the router must be configured to allow directed broadcast across multiple subnets.

The CBA discovery option also has a **LANDesk PDS2 discovery** option, where device discovery looks for the LANDesk Ping Discovery Service (PDS2) on devices. LANDesk Software products such as LANDesk[®] System Manager, Server Manager, and LANDesk Client Manager use the PDS2 agent. Select this option if you have devices on your network with these products installed. CBA discovery is not supported for Linux machines, but if you choose PDS2, Linux machines with an agent installed can be discovered.

- **IPMI:** Looks for IPMI-enabled servers. IPMI is a specification developed by Intel,* H-P,* NEC,* and Dell* to define the message and system interface for management-enabled hardware. IPMI contains monitoring and recovery features that let you access these features regardless of whether the device is turned on or not, or what state the OS may be in. Please keep in mind that if the Baseboard Management Controller is not configured, it will not respond to ASF pings which the product uses to discover IPMI. This means that you will have to discover it as a normal computer. When you push the client, ServerConfig will scan the system and detect it is IPMI and configure the BMC. For an overview of IPMI, see [IPMI support](#).
- **Server chassis:** Looks for blade server chassis management modules (CMMs). The blades in the server chassis are detected as normal servers.
- **Intel* AMT:** Looks for devices with Intel Active Management Technology support.
- **Starting IP:** Enter the starting IP address for the range of addresses you want to scan.
- **Ending IP:** Enter the ending IP address for the range of addresses you want to scan.
- **Subnet mask:** Enter the subnet mask for the IP address range you want to scan.
- **Add:** Adds the IP address ranges to the work queue at the bottom of the dialog.
- **Clear:** Clears the IP address range fields.
- **Edit:** Select an IP address range in the work queue and click **Edit**. The range appears in the text boxes above the work queue, where you can edit the range and add the new range to the work queue.
- **Remove:** Removes the selected IP address range from the work queue.
- **Remove all:** Removes all IP address ranges from the work queue.

To edit or delete a configuration

- On the **Discovery configurations** tab, click the configuration you want and click **Edit** or **Delete**.

Configuring SNMP scans

Network scan discoveries can use SNMP. Depending on your network's SNMP configuration, you may need to enter additional SNMP information in the discovery configuration. Clicking **Configure** next to the **SNMP** option shows the **SNMP configuration** dialog, which has these options:

- **Retries:** How many times device discovery retries the SNMP connection.
- **Wait for response in seconds:** How long device discovery should wait for an SNMP response.
- **Port:** What port device discovery should send SNMP queries to.
- **Community name:** The SNMP community name device discovery should use.
- **Configure SNMP V3:** Device discovery also supports SNMP V3. Click this button to configure SNMP V3 options in the **SNMP V3 configuration** dialog.

The **SNMP V3 configuration** dialog has these options:

- **User name:** The username device discovery should use to authenticate with the remote SNMP service.
- **Password:** The password for the remote SNMP service.
- **Authentication type:** The authentication type SNMP is using. Can be **MD5**, **SHA**, or **None**.

- **Privacy Type:** The encryption method the SNMP service is using. Can be **DES**, **AES128**, or **None**.
- **Privacy Password:** The password to use with the specified privacy type. Not available if you selected a privacy type of **None**.

Scheduling and running discovery

Use the **Schedule** button on the **Discovery configuration** tab to display the **Scheduled task** dialog. Use this dialog to schedule when discovery configurations run. You can schedule a discovery configuration to run immediately, run at some point in the future, make it a recurring schedule, or run it just once.

Discovery tasks can be rescheduled or deleted from the **Discovery tasks** tab. Once you schedule a discovery, see the **Discovery tasks** tab for discovery status. You can also access discovery task status from the **Scheduled tasks** tool. Once a discovery task completes, new devices not already in the core database will be added to the discovered device categories.

The **Scheduled task** dialog has these options.

- **Show in common tasks:** Lets other users see the task. When another user edits or runs the task, the user becomes the owner of an instance of the task.
- **Owner:** The task's owner.
- **Leave unscheduled:** (default) Leaves the task in the Task list for future scheduling.
- **Start now:** Runs the task as soon as possible. It may take up to a minute for the task to start.
- **Start at scheduled time:** Starts the task at the time you specify. If you click this option, you must enter the following:
 - **Date:** The date you want the task to start. Depending on your locale, the date order will be day-month-year or month-day-year.
 - **Time:** The time you want the task to start.
 - **Repeat every:** If you want the task to repeat, select a frequency (every Day, Week, or Month). If you pick Month and the date doesn't exist in all months (for example, 31), the task will only run in months that have that date.

To schedule a discovery

1. In the left navigation pane, click **Device discovery**.
2. On the **Discovery configurations** tab, select the configuration you want and click **Schedule**. Configure the discovery schedule. When finished, click **Save**.
3. Monitor the discovery progress in the **Discovery tasks** tab.
4. When the discovery completes, view all discovery results in the upper **Discovered devices** pane. If you double-click the discovery task, the quantity and percentage of devices are zero because these numbers are tied to targeted devices, and discovery tasks have no targets.

The **Discovery tasks** tab shows discovery job status. The status includes the following:

- The discovery configuration's name.
- The task status, which can be Working, All Completed, None Completed, or Failed.

- The last time the task ran.
- The type of task run.

To delete or reschedule a discovery

If you want to delete a task from the list, whether it has run already or not, click the task and click **Delete**. If the task hasn't run yet or is a recurring task, deleting it prevents it from running in the future.

You can also reschedule a discovery task in the list to run again or at a different time by clicking the task, clicking **Edit**, choosing **Schedule**, and reset the schedule. To rerun the task immediately, select the task and click **Start now**.

To view discovery task status

1. In the left navigation pane, click **Discovered devices**.
2. Click the **Discovery tasks** tab or click **Refresh** on that panel's toolbar.

Viewing discovered devices

View all discovered devices in the **Device discovery** top pane. This pane shows the results of all discoveries you have run. When you run a new discovery, any devices found are added to this list.

When device discovery finds a device, it tries to identify the device type so it can add the device to one of these categories:

- **Chassis:** Contains blade server chassis management modules (CMMs).
- **Computers:** Contains computers. Linux systems may be labeled as Unix systems in the **OS Name** column.
- **Infrastructure:** Contains routers and other network hardware.
- **Intel AMT:** Contains devices with Intel* Active Management Technology support.
- **IPMI:** Contains IPMI-enabled devices.
- **Other:** Contains unidentified devices.
- **Printers:** Contains printers.

These categories help keep the **Device discovery** list organized so you can more easily find the devices you're interested in. You can sort the device lists by any column heading when you click on a heading. Device discovery may not categorize devices correctly every time. You can easily move misidentified devices to the correct group by right-clicking the device you want to move, clicking **Move**, selecting the correct category, then clicking **OK**.

Occasionally, the core server is listed twice. This occurs because the same machine is found via different discovery mechanisms (for example, CBA, IPMI, CMM, PDS1, and PDS2) and the machine's information is added to the database under that mechanism.

Once you deploy agents to a discovered device and the device sends an inventory scan to the core, the discovered device will be removed from the discovered devices list.

Filtering the device list

Find devices matching search criteria you specify by using the **Filter by** field in the toolbar. You can filter on Node name, IP address, Subnet mask, Category, or OS Name. When you use a filter, devices are sorted alphabetically by the attribute you filtered on.

To filter the device list

1. In the left navigation pane, click **Device discovery**.
2. In the **Unmanaged** tree, click the group you want to filter.
3. In **Filter by**, click the attribute you want to filter on. (If **Filter by** is not visible, click **>>** to expand it.)
4. In the box beside the attribute, enter the text you want to filter on.
5. Click **Find**.

Adding categories

You can create device categories to group unmanaged devices. If you move a device to another category, it will appear in that group again if device discovery detects the device later. By moving devices you know you won't be managing with the console into other groups, you can more easily see new devices in the **Computers** group.

If you delete a group that contains devices, device discovery moves the devices to the **Other** group.

To add a device category

1. In the **Device discovery** view, click **Add category**.
2. Type a name for the group in the **Category name** box, then click **OK**.
3. To delete a category you have added, select the category, click **Delete category**, and click **OK** to confirm the deletion.

Moving discovered devices to the My devices list

After you discover devices you can move them to the **My devices** list. In the process of moving the devices their information is added to the database. Once the information is in the database, you can deploy the agent configuration, run queries and reports on it, and perform many other management tasks.

For devices that can be managed out of band (those with IPMI, Intel AMT, or DRAC functionality) you also have the option of managing the devices without deploying an agent. When you do this, the device information is saved in the database and the device's BMC is configured so you can use the management features allowed by the device's out-of-band management hardware.

To move discovered devices to the My devices list

1. In the **Device discovery** view, click the device you want to move to the **My devices** list. You can select multiple devices by using SHIFT+click or CTRL+click.
2. Click the **Target** button. The selected devices are then listed under the **Target list** tab.
3. Click the **Manage** tab.
4. Select **Move targeted devices**.
5. If the devices can be managed out of band, and you do not want to deploy a management agent to them, select **Manage out-of-band-enabled devices agentless**.
6. Click **Move**.

Devices are removed from the list of unmanaged devices and appear in the **My devices** list.

If you selected the out-of-band management option, you can view the status of the move process by clicking the **Move status** tab in the lower pane. Any errors in the configuration are noted here. To move an IPMI-enabled device to the **My devices** list, you must have provided the proper BMC credentials in the [Configure Services utility](#) to allow the core server to successfully authenticate to the device.

When you move a chassis management module (CMM) to the **My devices** list, it is displayed in the **All devices** list and also as a group in the **Public groups** list. The group details show the CMM and a list of available bays in the chassis with the names of blade servers in the bays. The blade servers are also detected and managed as individual servers.

Discovering Intel* AMT devices

System Manager includes the option to discover devices that are configured with Intel* Active Management Technology (Intel* AMT version 1). Devices can be discovered as Intel AMT devices only after you have accessed the Intel AMT Configuration Screen on the device and changed the manufacturer's default password to a secure password. (Refer to the manufacturer's documentation for information on accessing the Intel AMT Configuration Screen). If you haven't done this, the devices will be discovered but not identified as Intel AMT devices, and you won't be able to view the same inventory summary information as you otherwise would.

AMT discovery can take longer due to a double check of the secure and unsecure ports.

Devices with Intel AMT version 2 are not discovered using this process. After provisioning IDs are entered at the Intel AMT Configuration Screen with the IP address of the core server, the device is automatically discovered. See [Configuring Intel AMT devices](#) for details about working with version 2.

To discover Intel AMT devices

1. In the left navigation pane, click **Device discovery**.
2. Click **New** to create a new configuration, and type a name for the configuration. Or click an existing configuration and click **Edit** to modify it.
3. Check **Discover Intel AMT devices**.
4. Enter starting and ending IP addresses to scan a range of addresses, and enter a subnet mask.
5. Click **Add**, then click **OK**.

USERS GUIDE

6. Select the configuration and click **Schedule**. Set scheduling options, or click **Start now**, then click **Save**.
7. To view the progress of the scan, click the **Discovery tasks** tab.

Intel AMT-configured devices are displayed in a folder labeled **Intel AMT**. From this folder you can select the device and move it to the list of managed devices.

To add the device to the core database in order to manage it, the username/password for the device must match the username/password saved in the Configure Services utility, which allows System Manager to authenticate to the device. When you save the password configuration in the Configure Services utility, it stores the information in the core database so System Manager can authenticate to Intel AMT devices.

If you have Intel AMT devices with different credentials, you will need to make sure the credentials for each device match those in the Configure Services utility before you can manage them.

When an Intel AMT device is discovered and moved to the **My devices** list, it is automatically provisioned using the mode that you select using the Configure Services utility. Small Business mode provides basic management without network infrastructure services and is non-secure, while Enterprise mode is designed for large enterprises and provides security based on network services such as DHCP, DNS, and a TLS certificate authority service.

If your core is using a proxy server, the proxy server must support Digest Access Authentication to be able to discover Intel AMT devices.

To configure the Intel AMT password

1. Click **Start | All Programs | LANDesk | Configure Services**. Click the **Intel AMT Configuration** tab.
2. Type the current user name and password. These must match the user name and password as configured in the Intel AMT Configuration Screen (which is accessed in the computer BIOS settings) in order to manage Intel AMT devices.
3. To change the user name and password, complete the **New Intel AMT password** section.
4. Select the mode (**Small Business** or **Enterprise**) that you want to use for provisioning devices when you add them to the core database to manage them.
5. Click **OK**. This change will be made when the client configuration is run.

To move discovered Intel AMT devices to the list of managed devices

1. Click one or more device names in the list of unmanaged devices.
2. Click the **Target** button on the toolbar.
3. Click the **Manage** button on the lower pane, select **Move targeted devices**, then click **Move**.

If all devices you want to manage are visible in the list, you can select them, click the **Manage** button on the lower pane, select **Move selected devices**, then click **Move**.

The device is removed from the list of unmanaged devices and appears in the **All devices** list. Note that when you move devices to the **My devices** list, the Intel AMT provisioning runs in a

separate background process. You can continue with other discovery or management tasks while this occurs.

For more information about managing Intel AMT devices, see [Managing Intel* AMT devices](#) and [Intel* AMT support](#).

Device agent installation and configuration

Agent installation and configuration overview

In order to fully manage devices with the console, you need to install management agents on them. You can choose to install the default agent configuration (which installs all product agents) or customize your own agent configuration to install on your devices. Installation of System Manager does not install agents on the core automatically; you must also install agents to the core, then manually reboot the core. The agent configuration must include the monitoring agent to receive health alerts.

You can install management agents using one of the following methods:

- [Deploying agents](#). Target devices in the **My devices** list, then schedule an agent configuration task to remotely install agents on the devices.
- [Installing agents with an installation package](#). Create a self-extracting device installation package. Run this package locally on the device to install the agents. This must be done while logged in with administrative privileges.
- [Pulling the agent](#). Map to the core's Idlogon share (`//servername/Idlogon`) and run `SERVERCONFIG.EXE`.
- Manually on a device with a portable USB drive (see [Installing agents with an installation package](#))

Another resource for agent installation and configuration is the [Device discovery](#) chapter in the *User's Guide*.

Note: You can make a device configuration the default configuration by selecting the configuration in the **Agent configuration** page and clicking **Set as default**. An IPMI BMC-only configuration cannot be a default configuration. Default configurations cannot be deleted.

For Windows systems, the following port settings below require manual configuration of the firewall for full product functionality. To change these settings, access the Windows Firewall through the Windows Control Panel.

Managed Servers:

- File and Printer Sharing: TCP 139, 445; UDP 137,138 (Push won't work without this)
- Software distribution: TCP 9595 (Push won't work without this)
- Advanced: ICMP - "Allow incoming echo request" (Cannot be discovered if this is not enabled).

Core Server:

- Inventory: 5007

- Remote Control: 9535

To do so, click **Start | Control Panel | Security**.

Updating existing agents

You can push agent configurations to your devices, even if the Standard management or Remote control agents are not yet present. See [Configuring services and credentials](#) in the *User's Guide* for information on configuring credentials.

Once you've installed an agent package, installing removes the previous installation and installs the new one. You can uninstall an agent by creating a new agent package that doesn't include the agent you want removed.

Uninstalling agents

If you need to uninstall agents from servers, follow this procedure.

Warning: By default, Uninstallwinclient.exe reboots the device after uninstalling the agents unless you use the **/noreboot** switch to the command line. The reboot is necessary to complete the uninstall. If a reboot is initiated, the server will reboot without notice and all other applications are forced to quit. The **/noreboot** switch lets the server continue without a reboot.

To uninstall agents from a server

1. Log in at the server with administrative rights.
2. Map a drive to the core server's **ldmain** share.
3. Open a command prompt, change to the ldmain folder's drive letter, and enter the following:

```
uninstallwinclient.exe /noreboot
```

The uninstall will run silently, removing all agents.

You can also select **Start > Run > \\core name\ldmain\uninstallwinclient.exe /noreboot**.

To uninstall agents from a Linux server

1. Copy the linuxuninstall.tar.gz file to a temporary directory on the Linux device. This can be found in the core server's ManagementSuite shared folder.

The Linux device may not have Samba installed/configured, so you won't be able to copy it directly; you might either either pscp it from the core, copy it to the ldlogon folder, or copy it to removable media.

2. From a shell prompt (on the Linux machine), unpack this file using tar and the x, z, and f options.

```
tar -xzf linuxuninstall.tar.gz
```

3. After the file is unpacked, from a shell prompt run the linuxuninstall script from the current directory:

```
./linuxuninstall.sh
```

Configuring agents

In order to fully manage devices with the console, you need to install management agents on them. Installation of System Manager does not install agents on the core automatically; you must also install agents to the core, then manually reboot the core. Whether you use one of the default agent configurations or create an agent configuration in the console, you can install it on Windows or Linux devices in one of three ways:

- Create an agent configuration, target devices in the **My devices** list, then schedule an agent configuration task to remotely install agents on the devices.
- Create a self-extracting installation package. Run this package locally on the device to install the agents. This must be done while logged in with administrative privileges. For more information, see "[Installing agents with an installation package.](#)"
- From a Windows device, map to the core's Idlogon share (`\\myserver\ldlogon`) and run SERVERCONFIG.EXE.

To create an agent configuration

1. In the left navigation pane, click **Agent configuration**.
2. Click **New**.
3. Type a name for the new configuration in the **Configuration name** box.

Type a name that describes the configuration you're working on. This cannot be an existing configuration name.

4. Select the install type of the configuration (user selected or IPMI BMC-only). Select **IPMI BMC-only** under **Configuration** to configure the baseboard management controller (BMC on IPMI-enabled devices).

An IPMI BMC-only configuration configures the Baseboard Management Controller for out-of-band access, performs a full inventory scan, and removes itself. An IPMI BMC-only configuration cannot be a default configuration. When you create an IPMI BMC-only configuration, note that most of the editing options described in the following steps are not available.

5. Select the configuration you just created, and click **Edit**.

In the tabs, some options are dimmed because they are not configurable for the configuration you chose.

6. In the **Agent** tab, select the agents you want to deploy.
 - **All:** Installs all agents on the selected device.

- **Standard management agent:** Forms the basis of communication between devices and the core server. This is a required agent (except for BMC-only configurations). Most of this agent's processes are on-demand.
 - **Software updates:** Installs the software updates scanner. With this agent installed, you can configure how the scanner runs. This is not an on-demand agent.
 - **Monitoring:** Installs the monitoring agent on the selected server. The monitoring agent allows for many types of monitoring, including direct ASIC monitoring, In band IPMI, Out of band IPMI, and CIM. This is not an on-demand agent.
 - **Active System Console:** Installs the agent that allows for the Active System Console to be accessed from System Manager through the interface or through the menus. This agent is supported only on devices with Intel motherboards.
7. In the **Configuration** system type boxes, select the type. If this is dimmed, it is because you have already selected the type.
8. Select a **Reboot** option.

Rebooting manually means that devices will not reboot after installation. A device reboot is not required after agent configuration. You must manually reboot the device.

Rebooting if necessary causes a reboot for agent updates when updated files are locked.

9. In the Inventory tab, set the Inventory Scanner configuration settings. These are explained below.
- **Automatic update:** Remote devices read the software list from the core server during software scans. If this option is set, each device must have a drive mapped to the LDLOGON directory on the core server so they can access the software list. Changes to the software list are immediately available to devices.
 - **Manual update:** The software list used to exclude titles during software scans is loaded down to each remote device. Each time the software list is changed from the console, you must manually resend it to remote devices.
 - **Inventory scanner settings:** The time the inventory will run. You can select frequency, and you can specify that it always runs on startup. You can run the scanner manually from the managed server; you can launch it from Start | Programs | LANDesk Management | Inventory Scan. In Linux, you should be logged in as root, and run the following from the command line: `/usr/LANDesk/ldms/ldiscan - ntt`
 - **Always run on startup:** Runs the Inventory scanner on startup of the device. If you are creating an HP-UX configuration, this button is dimmed because the HP-UX scanner is set to run as a cron job, which will run either daily, weekly, or monthly. This cannot be modified.
 - **Start time:** Specify an hour range that the scanner can run between. If a device logs in during the time range you specify, the inventory scan runs automatically. If the device is already logged in, once the starting hour arrives the inventory scan starts automatically. This option is useful if you want to stagger inventory scans on devices so they don't send scans all at once.
 - **Repeat every:** Enter a number that represents the increment (like 1, 2, or 3), and the measurement (Minutes, Hours, or Days).

- **Restrictions:** Limits the available days and times the inventory scanner can be run. Click **Time of day**, **Day of week**, or **Day of month**, and enter inclusive parameters. For example, enter 10 for **Day of month** and 1:00 AM and 3:00 AM for **Time of day** to allow for the inventory scanner to be run on the 10th of each month between the hours of 1:00 AM and 3:00 AM.
10. In the **Software updates** tab, set the days and times you want the Software updates scanner to run. The scanner will run automatically, without requiring a scheduled task.
 - **Always run on startup:** Runs the Software updates scanner on startup of the device.
 - **Start time:** Specify an hour range that the scanner can run between. If a device logs in during the time range you specify, the software updates scan runs automatically. If the device is already logged in, once the starting hour arrives the software updates scan starts automatically. This option is useful if you want to stagger scans on devices so they don't send scans all at once.
 - **Repeat every:** Enter a number that represents the increment (like 1, 2, or 3), and the measurement (Minutes, Hours, or Days).
 - **Restrictions:** Limits the available days and times the software updates scanner can be run. Click **Time of day**, **Day of week**, or **Day of month**, and enter inclusive parameters. For example, enter 10 for **Day of month** and 1:00 AM and 3:00 AM for **Time of day** to allow for the inventory scanner to be run on the 10th of each month between the hours of 1:00 AM and 3:00 AM.
 11. In the **Rulesets** tab, select any monitoring and/or alerting rulesets you want included with the configuration. These rulesets are stored in the `Idlogon/alertrules` folder. New rulesets can be created in **Monitoring** or **Alerting**. In order for newly-created rulesets to display in the drop-down lists, you must generate the XML for the custom ruleset.
 12. Click **Save changes** to save the information to the database. Click **Save as file** to save the configuration as a distributable package.

Note: You can make an agent configuration the default configuration by selecting the configuration in the **Agent configurations** page and clicking **Set as default**. Default configurations cannot be deleted.

To schedule an agent configuration task

1. In the left navigation pane, click **Agent configuration**.
2. Click the agent configuration and click **Schedule task**.
3. Edit the list of [targeted devices](#) and the task schedule.
4. Click **Save**.

When you click **Schedule task**, a task is created (it has no targeted devices, and it is unscheduled). If you cancel this agent configuration task without saving it, be aware that it has still been created and appears in the **Task** list with a status of Not scheduled. You can delete it from the **My tasks** list.

After the agent configuration task is complete, you must reboot the device in order to view details about the device in the console (see [Viewing the Server information console](#)). This reboot is required when you install the agent on the core server as well as on managed devices. The agent configuration process lets you choose when to reboot so the reboot doesn't interfere with the server's use.

Deploying agents to managed devices

Once you have discovered devices, you can deploy agents to them. You can only deploy agents to supported Windows, Linux, and HP-UX devices. You must have Administrator rights to deploy agents to Windows devices, and the root privilege to configure Linux and HP-UX devices.

You can deploy agents to unmanaged devices in one of these ways:

- Push-based deployments using a discovery job and a domain administrative account you've configured for the scheduler service, which processes discovery jobs. The domain administrative account gives the scheduler services the rights it needs to install the server agents. This works for Windows NT family servers.
- Push-based deployments using the Standard management agent. If the servers have the Standard management agent, which is used by many LANDesk Software products, you can deploy to them without requiring a domain administrative account.

When deploying to discovered devices, use the **Unmanaged** tree's **Filter by** option. You can filter on IP address to isolate devices.

For Windows systems, the following port settings below require manual configuration of the firewall for full product functionality. To change these settings, access the Windows Firewall through the Windows Control Panel.

Managed Servers:

- File and Printer Sharing : TCP 139, 445; UDP 137,138 (Push won't work without this)
- Software distribution: TCP 9594, 9595 (Push won't work without this)
- Advanced - ICMP: "Allow incoming echo request" (Cannot be discovered if this is not enabled.)

Core Server:

- Inventory: 5007
- Remote control: 9535

Configuring device authentication credentials

Unmanaged devices with the Standard management agent on them don't require authentication credentials for agent deployment. To install agents on Windows OS servers that don't have the Standard management agent, you must specify the credentials that the scheduler service on the console device will use to get the required rights.

To install device agents on unmanaged devices, the scheduler service needs to be able to connect to devices with an administrative account. The default account the scheduler service uses is LocalSystem. The LocalSystem credentials generally work for devices that aren't in a domain.

If devices are in a domain, you must specify a domain administrator account. If you're configuring unmanaged devices in multiple domains, you must configure them one domain at a time, since

the scheduler service authenticates with one set of credentials, and each domain will require a different domain administrator account.

The core server includes Configure Services utility, which you can use to customize inventory options. This utility can only be run at the core server.

To configure the scheduler service login credentials

1. Launch the Configure Services utility on the core server by clicking **Start | Program Files | LANDesk | Configure Services**.
2. Click the **Scheduler** tab.
3. Click the **Change Login** button.
4. Enter the credentials you want the service to use on clients, typically a domain administrator account.

Installing agents

Once you've created an agent configuration in the console, you need to install it on devices. Installation of System Manager does not install agents on the core automatically; you must also install agents to the core, then manually reboot the core.

Client agent packages are a single self-extracting executable file. By default they're stored in the \Program Files\LANDesk\ManagementSuite\ldlogon folder on the core server. Running the executable installs the client agents silently, without requiring any user interaction. A browser is not required on a targeted device for successful agent installation.

Installing agents

You can update the agents by creating a new client configuration and distributing it from the console, or install agents directly to unmanaged devices.

Once you've installed a client agent package, installing other client agent packages removes all agents and installs the agents specifically selected. You can uninstall an agent by creating a new client agent package that doesn't include the agent you want removed.

Uninstalling agents

If you need to uninstall agents from devices, please see "[Agent installation and configuration overview](#)."

Installing agents with an installation package

One of the ways you can install agents is with a self-extracting device agent package. This allows you to copy the file to a CD or USB drive to install agents manually. You can create these packages by clicking **Save as file** in the bottom of the **Configuration** dialog.

1. Click **Agent configuration**, then double-click a configuration name.
2. In the **Agent configuration** dialog, click **Save as file**, then click **Close**.

Clicking **Save as file** creates a self-extracting executable package with the filename matching the configuration name you specified. It might be a few minutes before the package is available in the \Program Files\LANDesk\ManagementSuite\ldlogon\ConfigPackages folder on the core server.

Running the executable installs the agents, without requiring any user interaction. You must log in with administrative privileges.

If your users can't log in with administrative privileges to install the package, you can deploy the packages via e-mail, Web download, login scripts, or from a share.

Pulling the agents

This section contains details on deploying agents from the command line. You can control what components are installed on devices by using SERVERCONFIG.EXE command-line parameters. You can launch SERVERCONFIG.EXE in standalone mode. It's located in the *http:\coreserver\LDLogon* share, which is readable from any Windows server.

SERVERCONFIG.EXE uses SERVERCONFIG.INI for configuring devices.

Understanding SERVERCONFIG.EXE

SERVERCONFIG.EXE configures Windows NT family servers for management using the following process:

1. SERVERCONFIG determines whether the computer has been previously configured with a management agent. If it has, SERVERCONFIG removes all components and reinstalls selected components.
2. SERVERCONFIG loads the appropriate initialization file (SERVERCONFIG.INI) and executes the instructions contained in it.

The following command-line parameters are available for SERVERCONFIG.EXE:

Parameter	Description
/i	Components to include (quotation marks included): "Common Base Agent" "Inventory Scanner" "Alerting" "Vulnerability Scanner" "Server Monitor"

Parameter	Description
	You can combine these on the same command line. For example: <code>SERVERCONFIG.EXE /I="Alerting" /I="Vulnerability Scanner"</code>
/L or /Log=	Path to the CFG_YES and CFG_NO log files that log which servers were and were not configured
/LOGON	Execute [LOGON] prefixed commands
/N or /NOUI	Do not display the user interface
/NOREBOOT	Don't reboot server when done (default)
/REBOOT	Force reboot after running
/X=	Components to exclude. For example: <code>SERVERCONFIG.EXE /X=SD</code>
/CONFIG= /[CONFIG]=	Specifies a server configuration file to use in place of the default SERVERCONFIG.INI files. For example, if you've created configuration files called NTTEST.INI, then use this syntax: <code>SERVERCONFIG.EXE /CONFIG=TEST.INI</code> The custom .INI files should be in the same directory as SERVERCONFIG.EXE and note that the /config parameter uses the filename without the NT prefix.
/? or /H	Display help menu

Creating an agent configuration

Use **Agent configuration** to create and update server agent configurations (such as what agents are installed on managed). You can create different configurations for a group's specific needs. For example, you could create a configuration for Web servers and another one for application servers.

To push a configuration to a server, you need to:

- **Create the agent configuration:** Set up specific configurations for your servers.

- **Schedule the agent configuration:** Push the configuration to servers or, from the server, run SERVERCONFIG.EXE from the core server's LDLogon share.

To create an agent configuration

1. In the console, click **Agent configuration**.
2. Click the **New** toolbar button.
3. Enter a **Configuration name** and select the operating system, then click **OK**.
4. Click the new configuration name, then click **Edit**.
5. Select the agents you want to deploy.
6. Use the tabs on the top of the dialog to navigate to options relating to the components you selected. Customize the options you selected as necessary.
7. Click **Save changes**, then close the dialog.
8. If you want the configuration to be the default, click **Set as default**.

Pulling a Linux agent configuration

To pull a Linux agent configuration

1. Create a temporary directory on your Linux device (for example, /tmp/ldcfg), and copy the following into the directory:
 1. All files from the LDLOGON\unix\linux directory.
 2. Copy the shell script named after the configuration (<configuration name>.sh) into the temporary directory.
 3. Copy the *.0 file named after the configuration into the temporary directory. The * represents eight characters (0-9, a-f).
 4. Copy all files listed in the <configuration name>.ini file into the temporary directory. To identify these files, search the .INI file for "FILExx", where xx is a number. Most of the entries you find will have been copied to the client in step 1, but you will find .XML files that must be copied. The filenames should be left intact, with the following exceptions:
 - alertrules\<any text>.ruleset.xml should be renamed to internal.ruleset.xml
 - monitorrules\<any text>.ruleset.monitor.xml should be renamed to masterconfig.ruleset.monitor.xml
2. If the device has IPMI and a BMC (with Monitoring included in the installation), type the following on a command line:

```
export BMCPW="(bmc password)"
```

3. Running as root, execute the shell script for the configuration. For example, if you named the script "pull," use the full path used below:

```
/tmp/ldcfg/pull.sh
```

4. Remove the temporary directory and all its contents.

Note: Be aware that if you push or pull an agent out to a Linux device, then run

```
./linuxuninstall.sh -f ALL
```

to clean it and then push or pull again, the file with the GUID is the only file left on the device after this operation is completed.

The -f option deletes all directories owned by the product. See the [Linux uninstall documentation](#) for additional information.

Creating standalone agent configuration packages

Normally the agent configuration utility, SERVERCONFIG.EXE, configures agents on managed devices. If you want, you can have the **Agent configuration** window create a self-extracting single-file executable that installs an agent configuration on the server it's run on. This is helpful if you want to install agents from a CD or portable USB drive.

Pushing an agent configuration to devices

To push an agent configuration

1. In the console, select the devices you want to deploy the agent to, then click **Target**.
2. In the left navigation pane, click **Agent configuration**.
3. Right-click the agent configuration you want to push, then click **Schedule task**.
4. Click **Target devices** in the **Schedule task properties** dialog box, then click **Add target list**.
5. Click **Schedule task**.
6. Specify the time to deploy the agent, then click **Save**.

Installing Linux server agents

You can remotely deploy and install Linux agents and RPMs on Linux servers. Your Linux server must be configured correctly for this to work. To install an agent on a Linux server, you must have root privileges.

The default Linux install (Red Hat 3 and 4, and SUSE) includes the RPMs that the Linux standard management agent requires. If you select the monitoring agent in **Agent configuration**, you need an additional RPM, sysstat. For the complete list of RPMs that the product requires, see the *System Manager Deployment Guide*.

For an initial Linux agent configuration, the core server uses an SSH connection to target Linux servers. You must have a working SSH connection with username/password authentication. This product doesn't support public key/private key authentication. Any firewalls between the core and Linux servers need to allow the SSH port. Consider testing your SSH connection from the core server with a 3rd-party SSH application.

The Linux agent installation package consists of a shell script, agent tarball(s), .INI agent configuration, and agent authentication certificates. These files are stored in the core server's LDLogon share. The shell script extracts files from the tarball(s), installs the RPMs, and configures the server to load the agents and run the inventory scanner periodically at the interval you specified in the agent configuration. Files are placed under /usr/landesk.

You must also configure the scheduler service on the core to use the SSH authentication credentials (username/password) on your Linux server. The scheduler service uses these credentials to install the agents on your servers. Use the [Configure services utility](#) to enter the SSH credentials you want the scheduler service to use as alternate credentials. You should be prompted to restart the scheduler service. If you aren't, click **Stop** and then **Start** on the **Scheduler** tab to restart the service. This activates your changes.

Deploying the Linux agents

After you've configured your Linux servers and added Linux credentials to the core server, you must add servers to the **My devices** list so you can deploy the Linux agents. Before you can deploy to a server, you must add it to the **My devices** list. Do this by discovering your Linux server with **Device discovery**.

To discover your Linux servers

1. In **Device discovery**, create a discovery job for each Linux server. Use a standard network scan and enter the Linux server's IP address for the starting and ending IP ranges. If you have many Linux servers, enter a range of IP addresses. Click **OK** once you've added your discovery IP ranges.
2. Schedule the discovery task that you just created by clicking it and clicking **Schedule**. When the task finishes, verify that the discovery process found the Linux servers you want to manage.
3. In **Device discovery**, select the servers you want to manage, and click **Target** to add the selected devices to the Target List. Click the **Manage** tab in the lower half of the window. Click **Move selected devices** and click **Move**. This adds the servers to the **My devices** list, so you can target them for deployment.

To create a Linux agent configuration

1. In **Agent configuration**, click **New**.
2. Enter a configuration name, click **HP-UX** or **Linux Server Edition**, select the install type (server or desktop), and click **OK**.
3. Select the configuration you just created and click **Edit**.
4. Select the agents you want.
5. In the **Inventory** tab, select the options and the scanner frequency interval that you want. The installation script will add a cron job that runs the scanner at the interval you select.
6. In the **Rulesets** tab, select any monitoring and/or alerting rulesets you want included with the configuration. These rulesets are stored in the `ldlogon/alertrules` folder.
7. Click **Save changes**.

To deploy your agent configuration, select it in **Agent configurations** and click **Schedule task**. Configure the task and monitor the task progress in **Configuration tasks**.

Note: You will not receive any health information on a Linux device until after the inventory scanner completes its first scan after installation.

To pull a Linux agent configuration

1. Create a temporary directory on your Linux device (for example, /tmp/ldcfg), and copy the following into this temporary directory:
 - All files from the LDLOGON\unix\linux directory.
 - The shell script named after the configuration (<configuration name>.sh).
 - The *.0 file named after the configuration. The * represents eight characters (0-9, a-f) .
 - All files listed in the <configuration name>.ini file. To identify these files, search the .INI file for "FILExx", where xx is a number. Most of the entries you find will have been copied to the client in step 1, but you will find .XML files that must be copied. The filenames should be left intact, with the following exceptions:
 - alertrules\<any text>.ruleset.xml should be renamed to internal.ruleset.xml
 - monitorrules\<any text>.ruleset.monitor.xml should be renamed to masterconfig.ruleset.monitor.xml
2. If the device has IPMI and a BMC (with Monitoring included in the installation), type the following on a command line:

```
export BMCPW="(bmc password)"
```

3. Running as root, execute the shell script for the configuration using the full path below:

```
/tmp/ldcfg/lsminstall.sh
```

4. Remove the temporary directory and all its contents.

Note: Be aware that if you push or pull an agent out to a Linux device, then run

```
./linuxuninstall.sh -f ALL
```

to clean it and then push or pull again, the file with the GUID is the only file left on the device after this operation is completed.

The -f option deletes all directories owned by the product. See the [Linux uninstall documentation](#) for additional information.

Inventory scanner command-line parameters

The inventory scanner, ldiscan, has several command-line parameters that specify how it should run. See "ldiscan -h" or "man ldiscan" for a detailed description of each. Each option can be preceded by either '-' or '/'.

Parameter	Description
-d=Dir	Starts the software scan in the Dir directory instead of the root. By default, the scan starts in the root directory.

Parameter	Description
-f	Forces a software scan. If you don't specify -f, the scanner does software scans on the day interval (every day by default) specified in the console under Configure Services Inventory Scanner Settings .
-f-	Disables the software scan.
-i=ConfName	Specifies the configuration filename. Default is /etc/ldappl.conf.
-ntt=address:port	Host name or IP address of core server. Port is optional.
-o=File	Writes inventory information to the specified output file.
-s=Server	Specifies the core server. This command is optional, and only exists for backward compatibility.
-stdout	Writes inventory information to the standard output.
-v	Enables verbose status messages during the scan.
-h or -?	Displays the help screen.

Examples

To output data to a text file, type:

```
ldiscan -o=data.out -v
```

To send data to the core server, type:

```
ldiscan -ntt=ServerIPName -v
```

Linux inventory scanner files

File	Description
ldiscan	The executable that is run with command-line parameters to indicate the action to take. All users that will run the scanner need sufficient rights to execute the file.

File	Description
	There is a different version of this file for each platform supported above.
/etc/ldiscan.conf	<p>This file always resides in /etc and contains the following information:</p> <ul style="list-style-type: none"> • Inventory assigned unique ID • Last hardware scan • Last software scan <p>All users who run the scanner need read and write attributes for this file. The unique ID in /etc/ldiscan.conf is a unique number assigned to a computer the first time the inventory scanner runs. This number is used to identify the computer. If it ever changes, the core server will treat it as a different computer, which could result in a duplicate entry in the database.</p> <p>Warning: Do not change the unique ID number or remove the ldiscan.conf file after it has been created.</p>
/etc/ldappl.conf	<p>This file is where you customize the list of executables that the inventory scanner will report when running a software scan. The file includes some examples, and you'll need to add entries for software packages that you use. The search criteria are based on filename and file size. Though this file will typically reside in /etc, the scanner can use an alternative file by using the -i= command-line parameter.</p>
ldiscan.8	Man page for ldiscan.

Console integration

Once a Linux computer is scanned into the core database, you can:

- Query on any of the attributes returned by the Linux inventory scanner to the core database.
- Use the reporting features to generate reports that include information that the Linux scanner gathers. For example, Linux will appear as an OS type in the Operating Systems Summary Report.
- View inventory information for Linux computers.

Queries on "System Uptime" sort alphabetically, returning unexpected results

If you want to do a query to find out how many computers have been running longer than a certain number of days (for example, 10 days), query on "System Start" rather than "System Uptime." Queries on System Uptime may return unexpected results, because the system uptime is simply a string formatted as "x days, y hours, z minutes, and j seconds." Sorting is done alphabetically and not on time intervals.

Path to config files referenced in Idappl.conf doesn't appear in the console

ConfFile entries in Idappl.conf file need to include a path.

Device monitoring

About monitoring

System Manager provides several methods for monitoring a device's health status. Monitoring features gather data from various sources to help you keep track of many pieces of data on your devices, such as:

- Usage levels
- OS events
- Processes and services
- Historical performance
- Hardware sensors (fans, voltages, temperatures, etc).

This chapter includes information about the different features that monitor your managed devices:

- [Installing a monitoring agent](#) on devices and creating monitoring rulesets that can be deployed to devices
- [Setting performance counters](#) on devices and monitoring the performance data
- [Monitoring configuration changes](#) with alerts when changes occur
- Pinging devices regularly to [monitor their connectivity](#), using the **Device monitor** feature

Alerting is a related feature that uses the monitoring agent to initiate alerting actions such as sending e-mail or pager messages, rebooting or shutting down a device, or adding information to the alert log. You can generate alerts from any of the device events that can be monitored. See "[Using alerts](#)" for more information.

Notes

- Communications to the monitoring agent are via HTTP over TCP/IP in the form of GET or POST requests. Responses to requests are in XML documents.
- To run and store a query on the health status of devices (Computer.Health.State), you should be aware that the state in the database is represented by a number. The numbers correspond to the following states: 4=Critical, 3=Warning, 2=Normal, 1=Informational, null or 0=unknown.
- Hardware monitoring is dependent on the capabilities of the hardware installed on a device, as well as on the correct configuration of the hardware. For example, if a hard drive with S.M.A.R.T. monitoring capabilities is installed on a device but S.M.A.R.T. detection is not enabled in the device's BIOS settings, or if the device's BIOS does not support S.M.A.R.T. drives, monitoring data will not be available.
- If reporting from a specific machine appears to have stopped, you can use restartmon.exe in the LDCLIENT folder to restart the collector and all monitoring providers. This utility is for machines on which reporting has been installed, and reporting has stopped. Use this utility to restart the collector and providers without having to reboot the device.

Deploying the monitoring agent to devices

System Manager provides an immediate summary of a device's health when the monitoring agent is installed on the device. The monitoring agent is one of six agents that can be installed on managed devices. It checks the device's hardware and configuration on a regular, periodic basis and reflects any changes in the device's health status. This is shown by the status icon in the **My devices** list, and details are displayed in log entries (shown in the **System information** summary) for the device and in graphs (shown in the **Monitoring** summary page for the device).

For example, a monitored device with a disk drive that is filling up can display a warning status icon when the disk is 90% full, changing to a critical status icon when the disk is 95% full. You may also receive alerts for the same disk drive status if the device has an alert ruleset that includes rules for a drive space alert.

You can deploy the default monitoring ruleset to devices. Or, if you prefer, you can create a custom ruleset that includes only the health items you're concerned with.

Creating a monitoring ruleset

You can choose what is monitored on a device by creating a monitoring ruleset, which defines what the monitoring agent checks on the device. You can deploy a ruleset to one device or a targeted group of devices. For example, you may define one ruleset for servers dedicated to storage and use a different ruleset for Web servers.

The default monitoring ruleset includes 16 items. When you create a ruleset you can turn any of these items on or off, specify how frequently to check them, and, for some items, set performance thresholds. You can also select services running on the devices that you want to monitor.

The overall process for creating and deploying an alert ruleset is as follows:

1. Select the devices to which you will deploy the ruleset and click **Target** to add them to the **Targeted devices** list.
2. Create or edit a monitoring ruleset. Note that you must check the box to turn monitoring on for each event you want monitored in the ruleset. Not all are set to be monitored by default. Some events such as services also require that you select each service you want monitored. (See detailed steps below).
3. Deploy the ruleset to the targeted devices. You can target additional devices before deploying the ruleset if needed. (See detailed steps below).

To create a monitoring ruleset

1. In the left navigation pane, click **Monitoring**.
2. Click **New**, type a name and description for the configuration, and click **OK**.
3. Select the configuration in the left column.
4. In the list of items, click an item you want to change and click **Edit**.
5. To turn off monitoring of the item, clear the checkbox and click **Update**.
6. To change the frequency at which the item is monitored, select **Seconds** or **Minutes** and specify a number in the text box.
7. If applicable, set the threshold percentages for warning and critical status.

8. For **Services** monitoring, select the OS from the drop-down list. Select one or more services to monitor (use CTRL + click to select more than one) and click **>>** to add the services to the list on the right.
9. For each item that you modify, click **Update** to apply your changes to the configuration. If you modify an item and then decide not to change it, click **Revert** to restore the original settings.

When you edit services in a monitoring configuration from the core server, the **Available services** list displays known services from the inventory database. No services are displayed in the **Available services** list box until a LANDesk agent has been deployed to one or more devices and an inventory scan is returned to the core. For example, to select any Linux services from the list, you must have first deployed an agent to a Linux device.

To deploy a monitoring ruleset

1. In the left navigation pane, click **My devices**, then click the **All devices** group.
2. Select the devices to which you want to deploy the ruleset, then click **Target** to place the devices in the **Targeted devices** list.
3. In the left navigation pane, click **Monitoring**, then click the **Deploy ruleset** tab.
4. In the **Monitoring rulesets** box, select the ruleset you want to deploy.
5. Click the link to view the **Targeted devices** list. To remove a device from this list, right-click it, then click **Remove**. (To add devices, you must add them to the targeted list as described in step 2).
6. Click **Deploy** to deploy the selected ruleset to the targeted devices.

As part of the deployment process, an XML page is created that lists the deployed ruleset and devices the ruleset was deployed to. This report is saved on the core server in the LDLOGON directory, and is named with a sequential number assigned by the database. If you want to view this XML page separately from deploying a ruleset, click the **Generate XML** button and then click the link to view the XML file. Generating a ruleset as XML also allows it to be displayed in the list of available rulesets in the **Agent configuration settings**.

Turning off the ModemView service

The ModemView service is the service/driver that monitors modem calls (both incoming and outgoing) and generates an alert if it sees one. This service uses about 10 Mb of memory because it uses MFC. You may not want it to be running, especially if you don't have a modem on the device.

To turn the ModemView service off

1. On the device (either directly or via remote control) click **Start > Control Panel > Administrative Tools > Services**.
2. Double-click **LANDesk Message Handler Service**.
3. Under **Startup Type**, select **Manual**, and click **OK**.

You can also click **Stop** under **Service Status**.

Setting performance counters

System Manager lets you select performance items (counters that you want to monitor) on a managed device. You can monitor many different kinds of items including hardware components (such as drives, processors, and memory), or OS components (such as processes), or application components (such as bytes per second transferred by the system's Web server). When you select a performance counter you also specify the frequency for polling the item, as well as specify the performance thresholds and number of violations that are allowed before an alert is generated.

When a performance counter has been selected, you can then monitor performance on the **Monitoring** page by viewing a graph with real time or historical data. See "[Monitoring performance](#)" for more details.

To select a performance counter to monitor

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in another browser window.
2. In the left navigation pane, click **Monitoring**.
3. Click the **Performance counter settings** tab.
4. From the **Objects** column, select the object you want to monitor.
5. From the **Instances** column, select the instance of the object you want to monitor, if applicable.
6. From the **Counters** column, select the specific counter you want to monitor.

If the counter you want doesn't appear in the list, click **Reload counters** to refresh the list with any new objects, instances, or counters.

7. Specify the polling frequency (**Check every n seconds**) and the number of days to keep the counter history.
8. In the **Alert after counter is out of range** text box, specify the number of times the counter will be allowed to cross the thresholds before an alert is generated.
9. Specify upper and/or lower thresholds.
10. Click **Apply**.

Notes

- Performance log files can quickly grow in size; polling a single counter at a two second interval adds 2.5 MB of information to the performance log daily.
- A warning alert is generated when a performance counter drops below a lower threshold on either a Windows or a Linux device. When a performance counter exceeds an upper threshold on a Linux device, a warning alert is generated. When a performance counter exceeds an upper threshold on a Windows device, a critical alert is generated.
- When setting thresholds, bear in mind that alerts will be generated regardless of whether an upper or lower threshold is crossed. In the case of something like disk space, you may want to be alerted only if the device is running low. In this case, you would want to set the upper threshold to a high enough number that you would not be alerted if a lot of disk space became available on the device.

- Changing the **Alert after counter is out of range** number lets you choose to focus on an issue when it is a persistent problem or when it is an isolated event. For example, if you are monitoring the bytes sent from a Web server, System Manager can alert you when the bytes/sec consistently runs high. Or, you can specify a low number such as 1 or 2 to receive an alert whenever your anonymous FTP connections exceed a certain number of users.

Monitoring performance

The **Monitoring** page lets you monitor the performance of various system objects. You can monitor specific hardware components, such as drives, processors, and memory, or you can monitor OS components, such as processes or bytes per second transferred by the system's Web server. The **Monitoring** page includes a graph that displays real-time or historical data for counters.

In order to monitor a performance counter you must first select the counter, which adds it to the list of monitored counters. When you do this you also specify the frequency for polling the item and set performance thresholds and the number of violations that are allowed before an alert is generated. See "[Setting performance counters](#)" for details on selecting counters.

To view a performance graph for a monitored counter


1. In the **My devices** view, double-click the device you want to configure. The server information console opens in another browser window.
2. In the left navigation pane, click **Monitoring**.
3. Click the **Active performance counters** tab, if necessary.
4. In the **Counters** drop-down list, select the counter you for which want to see a performance graph.
5. Select **View real-time data** to display a graph of real-time performance.

or

Select **View historical data** to display a graph showing performance over the period you specified (Keep history) when selecting the counter.

On the performance graph, the horizontal axis represents time that has passed. The vertical axis represents the units you are measuring, such as bytes per second (when monitoring file transfers, for example), percentage (when monitoring percentage of the CPU that is in use), or bytes available (when monitoring hard drive space). The line height is not a fixed unit. The height of the line changes relative to the extremes in the data; for one counter the vertical axis might represent 1 to 100 and for another it might represent 1 to 500,000. When the data varies across a wide extreme, minimal changes can appear as a flat line.

Notes

- Selecting another counter refreshes the graph and resets the units of measurement.
- Click **Refresh**  to clear and restart the graph.
- If you receive an alert generated by a counter in the list, right-click the counter and click **Acknowledge** to clear the alert.

To stop monitoring a performance counter

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in another browser window.
2. In the left navigation pane, click **Monitoring**.
3. Click the **Active performance counters** tab, if necessary.
4. Under **Monitored performance counters**, right-click the counter and click **Delete**.

Monitoring configuration changes

This product can [generate alerts](#) if a device's hardware or software configuration changes and the monitoring agent is installed on the device. These changes may affect a device's performance and stability or cause problems with a standard installation. By monitoring vital pieces of the device, this product can reduce total cost of ownership (TCO).

The device configuration changes that will generate alerts are:

- **Application installed or uninstalled:** You can see which users installed or removed applications. This may be helpful in tracking licenses or employee productivity. Applications registered in the Windows Add/Remove Programs area of the Control Panel are monitored. Other applications are ignored. The application name that is used in the Windows Add/Remove Programs is the application name that appears in the notification log or alert pop-ups window.
- **Memory added or removed:** This product detects and monitors the quantity and type of memory installed. If the configuration changes, an alert is generated.
- **Hard drives added or removed:** This product detects and monitors the type and size of the drives installed on devices. If the configuration changes, an alert is generated.
- **Processor (or processors added, removed, or modified):** This product detects and monitors the number, type, and speed of the processor(s). If the configuration changes, an alert is generated.
- **Network card added or removed:** This product detects and monitors the number and type of network interface cards on devices, and generates alerts when the configuration changes.

To view a record of alerts for configuration changes, review the alert log on the server information console. See "[Viewing the alert log](#)" for details.

Monitoring for connectivity

In most cases, devices can alert you when critical situations arise, such as when the hard disk is filling up or a fan has stopped. However, in some situations, the device can go offline before it can send an alert. For example, a switch or router may disrupt network traffic, or the device may experience power failure.

In these situations, this product can check on devices periodically to determine whether they are available on the network. If the device does not respond to the ping, its health status is changed to critical the next time you refresh the **My devices** list.

You must set up the device monitor to ping targeted devices or all devices in the **All devices** group.

To set up the device monitor

1. In the **My devices** list, select the devices you wish to monitor. You can select them from **All devices** or from a public or private group.
2. Click **Target**.
3. In the bottom pane, click **Actions**, then click **Device monitor**.
4. To view a list of devices currently being monitored, click **Show monitored devices**.
5. Type numbers for the minutes between ping sweeps and number of times the product will attempt to communicate with a device.
6. Select whether to perform the action on devices in the [Targeted devices list](#) or on all devices in the **All devices** group.
7. To stop monitoring all devices, select **Never ping devices**.
8. Click **Apply**.

Only the last group of targeted devices are monitored. For example, if you target device A and device B and apply device monitoring to them, only device A and device B will be pinged by the core server. If you then target device C and device D and apply device monitoring to those devices, only device C and device D will be monitored; devices A and B will no longer be monitored.

Alert configuration

Using alerts

When a problem or other event occurs on a device (for example, the device is running low on disk space). System Manager can send an alert. You can customize these alerts by choosing the severity level or threshold that will trigger the alert. Alerts are sent to the console and can be configured to perform specific actions. Use this chapter to understand how alerts work.

- [How do I see alerts?](#)
- [What kinds of device problems can generate alerts?](#)
- [Configuring severity levels for events](#)
- [Process for configuring custom alert rulesets](#)
- [Example: Configuring an alert ruleset for a disk space problem](#)

How do I see alerts?

This product can notify you of problems or other computer events by:

- Adding information to the log
- E-mailing a notice or sending a message to a pager
- Running a program on the core or an individual device
- Sending an SNMP trap to an SNMP management console on the network
- Rebooting or shutting down a device

Be aware that certain alerts assigned to groups of machines can simultaneously generate a large number of responses. For example, you can set the alert "Computer configuration change" and associate it with an e-mail action. If a software distribution patch is applied to those machines with this alert setting, it would generate a number of e-mails from the core server equal to the number of machines to which the patch was applied, potentially "flooding" your e-mail server. In this case, an option might be to handle this alert by simply writing it to the core log rather than send e-mail.

What kinds of device problems can generate alerts?

This product has an extensive list of events that can generate alerts. Some are problems that need immediate attention; others are configuration changes that may or may not be a problem but that provide useful information to a system administrator. (See "[Monitoring configuration changes](#)" for related information). Alerts can only be generated when devices are equipped with the appropriate hardware. For example, alerts generated from sensor readings are only applicable to devices equipped with the correct sensors.

The types of events that you can potentially monitor include the following:

- **Hardware change:** A component such as a processor, memory, a drive, or a card has been added or removed.
- **Application added or removed:** An application has been installed or uninstalled on a device.
- **Service event:** A service has started or stopped on the device.

- **Performance:** A performance threshold has been crossed, such as for drive capacity, available memory, etc.
- **IPMI event:** An event detectable on IPMI devices has occurred, including changes to controllers, sensors, logs, etc.
- **Modem usage:** The system modem has been used, or a modem has been added or removed.
- **Physical security:** Chassis intrusion detection, power cycling, or another physical change has occurred.
- **Package installation:** A package has been installed on the target computer.
- **Remote control activity:** Remote control session activity has occurred, including starting, stopping, or failures.

Hardware monitoring that generates alerts is dependent on the capabilities of the hardware installed on a device, as well as on the correct configuration of the hardware. For example, if a hard drive with S.M.A.R.T. monitoring capabilities is installed on a device but S.M.A.R.T. detection is not enabled in the device's BIOS settings, or if the device's BIOS does not support S.M.A.R.T. drives, alerts will not be generated from S.M.A.R.T. drive monitoring.

Configuring severity levels for events

Device problems or events can be associated with some or all of the severity levels shown below.

- **Informational:** Supports configuration changes or events that manufacturers may include with their systems. This severity level does not affect device health.
- **OK:** Indicates that the status is at an acceptable level.
- **Warning:** Provides some advance warning of a problem before it reaches a critical point.
- **Critical:** Indicates that the problem needs immediate attention.
- **Unknown:** The alert status cannot be determined or the monitoring agent has not been installed on the device.

Depending on the nature of the event or server problem, some severity levels don't apply and aren't included. For example, with the Intrusion detection event, the device's chassis is either open or closed. If it is open, this can trigger an alert action with a severity of Warning. Other events, such as Disk space and Virtual memory, include three severity levels (OK, Warning, and Critical).

You can choose the severity level or threshold that will trigger some alerts. For example, you can select different actions as a result of a Warning or Critical status for an alert. The Unknown status can't be selected as an alert trigger but simply indicates that the status cannot be determined.

Process for configuring custom alert rulesets

You can configure an alert ruleset to deploy to an individual device or to a group of targeted devices. Each managed device must have the product monitoring component installed before it can send alerts to the core server. (See "[Configuring agents](#)" for more information).

When the monitoring component is installed to a managed device, a default ruleset of alerts is included to provide health status feedback to the console. This default ruleset includes alerts such as:

- Disk added or removed
- Drive space
- Memory usage
- Temperature, fans, and voltages
- Performance monitoring
- IPMI events (on applicable hardware)

In addition to the default ruleset, you can configure and deploy custom alert rulesets. You can include custom alert actions to respond a particular event. For example, if a fan stops, it can trigger an alert and send an e-mail to your hardware support group.

The overall process for creating and deploying an alert ruleset is as follows:

1. Select the devices to which you will deploy the ruleset and click **Target** to add them to the **Targeted devices** list.
2. Create the alert action rulesets you will use. These action rulesets define the kinds of actions that can be triggered by alerts. (See "[Configuring alert actions](#)" for more information).
3. Create your custom alert ruleset. When you do this you can select the actions that you previously defined. (See "[Configuring an alert ruleset](#)" for more information).
4. Deploy the ruleset to the targeted devices. You can target additional devices before deploying the ruleset. (See "[Deploying rulesets](#)" for more information).

The following is a simple example of this process.

Example: Configuring an alert ruleset for a disk space problem

1. In the left navigation pane, click **My devices**, then double-click the **All devices** group.
2. Select the devices for which you want to set the alert, then click **Target** to place the devices in the **Targeted devices** list.
3. Click **Alerting**, then click the **Action rulesets** tab.
4. In the **Actions** drop-down list, select the action you want to configure (such as **Send e-mail/page**). Click **New**, type a name in the **Name** field, then click **OK**.
5. Back in the **Action rulesets** page, select the ruleset you just named, and click **Edit actions**. Specify the data as needed in the text boxes. When you have finished, click **Save**.
6. Click the **Alert rulesets** tab.
7. Click **New**, type something like "Disk Space Problem" in the **Name** field, type a description in the **Description** field, and click **OK**.
8. Click the alert ruleset you just named, and click **Edit ruleset**.
9. Click the **New** button.
10. In the **Alert type** drop-down list, click **Drive space**.
11. Check the status you wish to alert on: **OK**, **Warning**, or **Critical**. (If you want the same action for multiple statuses, select more than one. If you want a different action for each status, create a separate configuration for each status so you can trigger different actions for different status levels).

12. In the **Action** drop-down list, select the action you want to occur if the conditions specified in steps 6 and 7 are met. If the action you want is not on the list, you can [create](#) it using the **Action rulesets** page. (If you have not created an alert action ruleset it will not be in the list).
13. In the **Alert action** drop-down, select the configuration you want.
14. Check **Affects device health** if you want the alert to apply to the server's health state when it is displayed in the **All devices** list. If the severity level for the alert is **Informational** only, the alert will not affect device health.
15. Click **Add**.
16. Repeat steps 6-12 if you want to add additional alerts to the ruleset.
17. When finished, click **Close**.
18. If you want to change any alert types in the ruleset, select the alert type and click **Edit**, make the changes, click **Update**, then click **Close**.
19. When the alert ruleset is defined, apply the ruleset to the targeted devices: click **Deploy ruleset**, select the ruleset, and click **Deploy**.

Configuring alert actions

Use the **Action rulesets** page to provide additional information for how you want the actions to behave when they are selected. When a threshold is crossed, an alert is generated. The alert can have an action associated with it, such as sending an e-mail. Each action has its own configurations, and must be set up individually.

To create an action ruleset

1. In the left navigation pane, click **Alerting**, then click the **Action rulesets** tab.
2. In the **Actions** drop-down list, select the action you want to configure. Each action has its own list of unique configurations.
3. Click **New**, type a name in the **Name** field, then click **OK**.
4. Back in the **Action rulesets** page, select the ruleset you just named, and click **Edit actions**.
5. If you selected **Execute program on core** or **Execute program on client**, type or paste the path to the program you want to execute on the alert, then click **Save**. When you select either **Execute program** action, note that programs may not display as expected on the desktop. When the program is run, it is started as a service in Windows and so is not displayed as a regular application would be. Programs that are run in this way should not contain a user interface that requires interaction. To definitively determine if the program executed, check the processes in the Windows Task Manager.

If you selected **Send e-mail/page**, type the full e-mail address of the person you want to receive the e-mail in the **To** field; type a valid e-mail address in the **From** field; type a subject in the **Subject** field; type a message in the **Body** field; select the day or time you want to have the message sent; and type the location of an SMTP server in the **SMTP server** field. Click the **Help** box to learn how to send messages to multiple recipients and how to use variables in your messages. When you have finished, click **Save**.

If you selected **Send an SNMP trap**, type the host name, select a version, type the community string in the **Community string** box, then click **Save**.

Notes

- Certain alerts assigned to groups of machines can simultaneously generate a large number of responses. For example, you can set the alert "Computer configuration change" and associate it with an e-mail action. If a software distribution patch is applied to those machines with this alert setting, it would generate a number of e-mails from the core server equal to the number of machines to which the patch was applied, potentially "flooding" your e-mail server. In this case, an option might be to handle this alert by simply writing it to the core log rather than send e-mail.
- Some alert actions do not affect device health. These include actions such as Run Program on Client, Shut Down/Restart, and any alert that is Informational only. However, if any of these actions are combined with other alert actions that do affect device health, then any alerts generated will affect device health and will appear in the alert log.
- The **From** field in an e-mail must contain a valid e-mail address in order for SMTP alerting to work.
- SNMP traps identified as version 1 are processed, while those identified as version 3 are only forwarded.
- For SNMP traps, the Severity levels are reported in the Specific Trap Type field of the trap. Values are 1 = unknown, 2 = info, 3 = OK, 4 = warning, 5 = critical.

Configuring an alert ruleset

Use the **Alert rulesets** page to create a new alert ruleset. Before you can configure alerts, you must configure actions. (See "[Configuring alert actions](#)" for more information).

There are two alert rulesets that appear by default on the **Alert rulesets** page:

- **Core alert ruleset:** this ruleset ensures that alerts are sent to the core server when the **Device monitor** feature is enabled (see [Monitoring for connectivity](#)). This ruleset contains a predefined group of alert types, including Device monitor, AMT Circuit Breaker Alert, and Serial Over LAN Session alert types. You can edit the status, action, alert action, and health settings for the core alert types, but if you attempt to make any other changes they will be ignored.
- **Default ruleset:** this ruleset is deployed to all managed devices and contains a number of alert types that are of general use for most network administrators. You can edit this ruleset to add other alert types and change the settings for the default alert types. Any time you edit this ruleset, the changes are deployed to all managed devices even if you don't explicitly redeploy the ruleset.

In addition to these rulesets you can create custom rulesets to apply to targeted groups of managed devices. These rulesets must be generated in XML format to display in **Agent configuration**.

When you create a custom ruleset for a device, be aware that if a default ruleset has already been deployed to the device you may have overlapping or conflicting alerting rules. If you deploy the default ruleset when you configure the managed device, and then deploy a custom ruleset, both rulesets will be executed on the device. For example, if both rulesets generate alerts for the same alert type but take different actions, you may have duplicate or unpredictable alert actions as a result. While you can't remove the default ruleset once it has been deployed, you can edit the default ruleset if you want to change any part of it.

To create an alert ruleset

1. In the left navigation pane, click **Alerting**, then click the **Alert rulesets** tab (if necessary).
2. Click **New**, type a name in the **Name** field, type a description of the alert in the **Description** field, then click **OK**.
3. Click the ruleset you just named, and click **Edit ruleset**.
4. Click **New**.
5. In the **Alert type** drop-down list, select the component, action, or event type you want to receive alerts on.
6. Check each status you wish to alert on: **Informational**, **OK**, **Warning**, or **Critical**. For example, to receive an alert if the type you selected in step 5 crosses a critical threshold, check **Critical**.
7. In the **Action** drop-down list, select the action you want to occur if the conditions specified in steps 5 and 6 are met. These actions are defined beforehand; if you want an action that is not in the list, you can [create](#) one using the **Action rulesets** page.

Note: Certain alerts assigned to groups of machines can simultaneously generate a large number of responses. For example, you can set the alert "Computer configuration change" and associate it with an e-mail action. If a software distribution patch is applied to those machines with this alert setting, it would generate a number of e-mails from the core server equal to the number of machines to which the patch was applied, potentially "flooding" your e-mail server. In this case, an option might be to handle this alert by simply writing it to the core log rather than send e-mail.

8. In the **Alert action** drop-down list, select a configuration. There may only be one configuration available (the contents of this list change depending on what you selected in step 7).
9. Check **Affects device health** if you want the alert to apply to the server's health state when it is displayed in the **All devices** list. If the severity level for the alert is **Informational** only, the alert will not affect device health.
10. Click **Add**.
11. Repeat steps 5-10 to add additional alerts to the ruleset.
12. When finished, click **Close**.

To edit an alert ruleset, select the ruleset (step 3 and click **Edit ruleset**, then continue with the steps above).

A few minutes after you create or edit a ruleset, the ruleset deployment service automatically attempts to update all computers that previously had that ruleset deployed to them. Or, if you want to deploy the ruleset immediately, click the **Deploy ruleset** tab and click **Deploy**.

Deploying rulesets

Use the **Deploy ruleset** page to move the selected alert ruleset to targeted devices.

In order to deploy a ruleset to a managed device, you must first have a management agent installed on that device. When you deploy the standard management agent, the Default ruleset is deployed. After the agent setup is complete you can update the Default ruleset or deploy new rulesets. To begin you should target the devices you want to deploy the ruleset to.

To deploy an alert ruleset

1. In the left navigation pane, click **My devices**, then click the **All devices** group.
2. Select the devices to which you want to deploy the alert ruleset, then click **Target** to place the devices in the **Targeted devices** list.
3. In the left navigation pane, click **Alerting**, then click the **Deploy ruleset** tab.
4. In the **Alert rulesets** box, select the ruleset you want to deploy.
5. Click the link to view the targeted devices list. To remove a device from this list, right-click it, then click **Delete**. To remove all devices, right-click any device name and click **Reset**. To add devices, you must add them to the [targeted list](#) (steps 1-2 above).
6. Close the **Target list** window, then click **Deploy** to deploy the selected configuration to the targeted devices.

As part of the deployment process, an XML page is created that lists the deployed ruleset and devices the ruleset was deployed to. This report is saved on the core server in the `\dlogon\alertrules` directory, and is named with a sequential number assigned by the database. If you want to view this XML page separately from deploying a ruleset, click the **Generate XML** button and then click the link to view the XML file.

Note that only one custom ruleset can be in effect on a managed device at any time. If you have deployed a custom ruleset and then deploy a second custom ruleset to the same device, the first custom ruleset is overwritten and the second is in effect.

Viewing alert rulesets for a device

Use the **Alerting rulesets** page to view a list of the alert rulesets assigned to the selected device, and to view the details of each alert.

To view alert rulesets

1. In the **My devices** view, double-click the device you want to configure. The server information console opens in another browser window.
2. In the left navigation pane, click **Rulesets**.
3. Click the **Alerting rulesets** tab.

The following details are provided about each ruleset. For more information on modifying these details, see [Using alerts](#).

- **When state reaches:** When the state of the alert reaches the displayed state, an alert will be generated.
- **Affects health:** Indicates whether the alert status will apply to the server's health state when it is displayed in the **All devices** list.
- **Ruleset name:** The name of the alert ruleset, as defined in the [Alert rulesets](#) dialog.
- **Alert type:** A description of the source of the alert (hardware, software, event, etc.).
- **Action configuration:** The action that occurs when the alert is generated, as defined in the [Action configurations](#) dialog.
- **Alert handler:** The kind of alert to be generated, such as an e-mail, an SNMP trap, or executing a program.
- **Instance:** Indicates the specific source of the alert.

You can also click the **Alert log** button to go to the device's alert log and view details about alerts. (See [Viewing the alert log](#) for more information.)

Viewing the alert log

Use the **Alert log** page to view alerts sent to the core (the global alert log) or to managed devices. The log is sorted by Time (GMT), the most recent being at the top of the log.

The Alert log contains the following columns:

- **Alert name:** The name associated with the alert, as defined in the **Alert configurations** page.
- **Time:** The date and time the alert was generated (GMT).
- **Status:** The status of the alert, which can be one of the following:
 - **Unknown:** The status cannot be determined.
 - **Informational:** Supports configuration changes or events that manufacturers may include with their systems.
 - **OK:** Indicates that the status is at an acceptable level.
 - **Warning:** Provides some advance warning of a problem before it reaches a critical point.
 - **Critical:** Indicates that the problem needs your immediate attention.
- **Instance:** Indicates the specific source of the alert.
- **Device name:** The name of the device on which the alert was generated. This should be a fully qualified domain name. (Global alert log only).
- **IP address:** The IP address of the device on which the alert was generated. (Global alert log only).

If the device name does not appear as a fully qualified domain name, it is because this product was unable to resolve the fully qualified domain name for the device.

To view the global alert log

1. In the left navigation pane, click **Logs**.
2. To sort entries by time, name, status, or instance, click a column heading.
3. To view a more detailed description of the alert, double-click the entry in the **Alert name** column.

4. To list log entries by name, status, or instance, select the filter criteria in the Filter drop-down list. For example, select **Alert name** and type a complete name (such as Performance) or a partial name with the * wildcard (such as Remote*). To search by date, select **Enable date filtering**, enter a range with a start date and end date, and click **Find**.
5. To clear the health status of an alert, select the alert by clicking the number in the **Alert name** column, click **Clear alert**, and click **OK**. To delete a log entry, select the alert and click **Delete entry**.
6. To delete all entries in the log, click **Purge log**.

To view the alert log for a specific device

1. Double-click the device in the **My devices** list.
2. In the left navigation pane, click **System information**.
3. Click **Logs**, then double-click **Alert log**.
4. To sort entries by time, name, status, or instance, click a column heading.
5. To view a more detailed description of the alert, click the entry in the **Alert name** column.
6. To list log entries by name, status, or instance, click the **Filter** button on the toolbar and select the filter criteria. For example, select **Alert name** and type a complete name (such as Performance) or a partial name with the * wildcard (such as Remote*). Then click Find on the toolbar to view the alerts associated with the filter options you chose.
7. To view log entries for a range of dates, clear the **Show events for all dates** check box and select a range of dates. Click **Refresh** to view the entries for that range of dates.

Software updates

System Manager includes a software updates tool that lets you search for updates to management software, operating system software, and device drivers. You can download these types of updates and remediate affected devices by deploying and installing the appropriate updates (also referred to as patches).

Read this chapter to learn about:

- [Software updates overview](#)
- [About the Software updates window](#)
- [Configuring devices for software update scanning](#)
- [Updating vulnerability definitions](#)
- [Scheduling software update downloads](#)
- [Viewing software update and detection rule information](#)
- [Purging software updates information](#)
- [Scanning devices for software updates](#)
- [Viewing detected updates](#)
- [Downloading patches](#)
- [Remediating software updates](#)

Software updates overview

The software updates tool helps you maintain up-to-date software on the managed devices across your network. You can automate the repetitive processes of maintaining current software, downloading the appropriate update files, and deploying and installing the necessary updates on affected devices.

This product uses standard role-based administration to allow users access to the software updates tool. Role-based administration is the product's access and security model that lets administrators restrict access to tools and devices. Each user is assigned specific rights and scope that determine which features they can use and which devices they can manage. An administrator assigns these rights to other users (see [About role-based administration](#) for more information). To use the software updates tool, a user must be logged in with the Patch management, Basic Web console, and Reports rights.

Supported server platforms

Software updates supports most standard server platforms, enabling you to scan for updates and deploy them to managed servers running the following operating systems:

- Windows 2000 Server SP4
- Windows 2000 Advanced Server SP4
- Windows 2000 Professional SP4
- Windows 2003 Standard Edition SP1
- Windows 2003 Enterprise Edition SP1
- Windows XP Pro SP2

- RedHat Enterprise Linux ES/AS 3
- SUSE Linux Server 9 (Professional, Enterprise, and Advanced)

About the Software updates window

Users with the Patch management right will see the **Software updates** tool in the console's left navigation pane. When you click **Software updates**, you will see a toolbar and two panes in the right side of the window. The left pane shows a hierarchical tree view of software update groups. Click on a group to view its contents in the right pane. The right pane displays software update definition details in a column list. At the top, it contains a **Find** button to quickly search for the specified criteria. In the **Find** box, the following extended characters are not supported: <, >, ', ", !.

Toolbar buttons

- **Update:** Opens the **Update vulnerabilities settings** dialog where you can specify the platforms and languages whose software update information you want to update. You can also configure whether to place updates in the **Scan** group, whether to download associated patches concurrently, the location where patches are downloaded, and proxy server settings.
- **Schedule download:** Opens the download task in the **Scheduled task** dialog where you can configure task options. When you click **Save**, the download task is placed in the **Scheduled tasks window** and under the **Vulnerability tasks** tab.
- **Schedule patch tasks:** Opens the **Schedule vulnerability scan** dialog where you can provide a name and configure scanner options.
- **Refresh:** Updates the list in the right pane with the latest downloaded update information.
- **Purge:** Opens the **Purge security and patch definitions** dialog where you can specify the platforms and languages whose vulnerability information you want to remove from the core database.

Left pane (tree view)

The left pane of the window shows the following groups:

- **Scan:** Lists all of the updates that are searched for when the software updates tool runs on managed devices. In other words, if an update is included in this group, it will be part of the next scan operation; otherwise, it won't be part of the scan.

Scan can be considered one of three vulnerability states, along with Don't scan and Unassigned. As such, a software update can reside in only one of these three groups at a time. An update is identified by a unique icon for each state (question mark (?) icon for Unassigned, red X icon for Don't Scan, and the regular vulnerability icon for Scan). Moving an update from one group to another automatically changes its state.

To move a software update from one group to another, right-click the update and select the group to move it to.

By moving updates into the Scan group, you can control the specific nature and size of the next software update scan.

New updates can also be automatically added to the Scan group during an update by checking the **Put new definitions in the Scan group** option on the **Update Vulnerabilities Settings** dialog.

Caution about moving software updates from the Scan group

When you move software updates from the Scan to the Don't Scan group, the current information in the core database about which scanned devices detected those updates is removed from the database and is no longer available in either the Software Update Properties dialog or in the scanned Server Information dialog. To restore that assessment information, you would have to move the software updates back into the Scan group and run the scan again.

- **Do not Scan:** Lists the software updates that aren't searched for the next time the scanner runs on devices. As mentioned above, if an update is in this group, it can't be in the Scan or Unassigned group. You can move updates into this group to remove them from a software update scan.
- **Detected:** Lists all of the software updates detected by the previous scan, for all of the target devices included in that scan job. The contents of this group are always determined by the last software update scan, whether one device was scanned or many devices.

The Detected list is a composite of all detected software updates found by the most recent scan. The Scanned and Detected columns are useful in showing how many devices were scanned, and on how many of those devices the software update was detected. To see specifically which servers have a detected update, right-click the definition and then select **View affected computers**. Note that you can also view update information for a specific server in its [Server information console](#) dialog.

You can only move software updates from the Detected group into either the Unassigned or Don't Scan groups.

- **Unassigned:** Lists all of the software updates that do not belong to either the Scan or Don't Scan groups. The Unassigned group is essentially a holding area for collected updates until you decide whether you want to scan for them or not.

By default, collected software updates are added to the Scan group during an update.

You can move software updates from the Unassigned group into either the Scan or Don't Scan groups.

- **View by OS:** Lists all of the downloaded software updates organized into specific device operating system subgroups. These subgroups help you identify updates by OS category. You can use these OS subgroups to copy a set of updates into the Scan group for OS-specific scanning.

Software updates can be copied from an OS group into the Scan, Don't Scan, or Unassigned group. Updates can reside in more than one platform and/or product group simultaneously.

- **View by Product:** Lists all of the downloaded software updates organized into specific product subgroups. These subgroups help you identify updates by product category. You can use these product subgroups to copy updates into the Scan group for product-specific scanning.

Right pane (list view)

The right pane of the window displays the following software update details, listed in sortable columns:

- **ID:** Identifies the update with a unique, vendor-defined alphanumeric code.
- **Severity:** Indicates the severity level of the update. Possible severity levels include: Service Pack, Critical, High, Medium, Low, Not Applicable, and Unknown.
- **Title:** Describes the nature or target of the update in a brief text string.
- **Language:** Indicates the language of the OS affected by the update.
- **Date Published:** Indicates the date the update was published by the vendor.
- **Silent Install:** Indicates whether the update's associated patch file installs silently (without user interaction). Some updates may have more than one patch. If any of an update's patches don't install silently, the update's Silent Install attribute is No.
- **Repairable:** Indicates whether the update can be repaired through patch file deployment and installation. Possible values are: Yes, No, and Some (for an update that includes multiple detection rules and not all detected updates can be repaired).

Double-click the update ID to view more detailed information on its properties dialog. From a software update properties dialog, you can see the detection rules for the update, download associated patch files, and click the rule to view its detailed properties dialog.

Configuring devices for software update scanning

Before managed devices can be scanned for vulnerabilities, and receive patch deployments, they must have the Software updates agent installed.

The easiest way to deploy the Software updates agent to multiple managed devices is to create a new agent configuration, with the Software updates agent selected (default setting), and then schedule the configuration for the desired target devices with **Scheduled tasks**.

When you configure a device to support software updates, the necessary files for software update scanning and remediation (i.e., patch deployment and installation) are installed on the target device.

Updating software update definitions

Your network is continuously vulnerable to maintenance issues like software updates and bug fixes. The software updates tool makes the process of gathering the latest known patch information quick and easy by letting you update software via a LANDesk-hosted database. This service consolidates known updates from trusted, industry/vendor sources.

By establishing and maintaining up-to-date patch information, you can better understand the nature and extent of needed software updates for each server operating system you support. The first step is to keep up with the latest known update information.

You can configure and perform software updates at once, or create a scheduled update task to occur at a set time or as a recurring task.

To update software update information

1. In the left navigation pane, click **Software updates**. (For a description of the dialog, see [About the Software updates window](#).)
2. Click the **Update** toolbar button.
3. Select the download source site from the list of available content servers.
4. Select the platforms whose software update information you want to update. You can select one or more platforms in the list. The more platforms you select, the longer the update will take.
5. Select the languages whose software update information you want to update for the platforms you've specified. You can select one or more languages in the list. The more languages you select, the longer the update will take.
6. If you want new software update definitions (i.e., those that do not already exist in the database) to automatically be placed in the Unassigned group instead of the default location (which is the Scan group), uncheck the **Put new definitions in the Scan group** check box.
7. If you want to automatically download the actual patch executable files, check the **Download patches for definitions selected above** check box, and then click one of the download options.
 - **For detected definitions only:** Downloads only the patches that are associated with software updates detected by the last software update scan (i.e., the updates that are currently residing in the Detected group).
 - **For all referenced definitions:** Downloads ALL of the patches that are associated with software updates currently residing in the Scan group. This will take a long time.

Patches are downloaded to the location specified in the Patch Settings section of the dialog (see procedure below).

8. If you have a proxy server on your network that is used for external Internet transmissions (required to update software update information and download patches), click the **Proxy Settings** tab and check the **Use Proxy Server** box. Specify the server's address, port number, and authentication credentials if a login is required to access the proxy server.
9. Click **Apply** at any time to save your settings.
10. Click **Update Now** to run the software update. The **Update security and patch definitions** dialog displays the current operation and status.
11. When the update has completed, click **Close**. Note that if you click **Cancel** before the update is finished, only the software update information that has been processed to that point is downloaded to the core database. You would need to run the update again in order to obtain all of the remaining information.

Note: Do not close the console while an update process is running or the process will be terminated. This does not apply to a scheduled download task.

If you have System Manager and LANDesk® Management Suite installed on the same core server, both products use the same settings file to determine what types of vulnerabilities are updated. In some cases you may see updates in System Manager that are only configurable from Management Suite when you run the update. For example, if you have selected Security threats as an update option in Management Suite, and then choose to update Software updates in System Manager, when you run the update in System Manager you will see both software updates and security threats in the updated items.

To configure the patch download location

1. On the **Update vulnerabilities settings** dialog, click the **Patch Settings** tab.
2. Enter a UNC path where you want the patch files copied. The default location is the core server's \LDLogon\Patch directory.
3. If the UNC path entered above is to a location other than the core server, enter a valid username and password to authenticate to that location.

The folder must have file and web sharing enabled and Anonymous access must be enabled.

4. Enter a Web URL where servers can access the downloaded patches for deployment. The Web URL should match the UNC path above.
5. You can click **Test Settings** to check to see if a connection can be made to the Web address specified above.
6. If you want to restore the UNC path and Web URL to their default locations, click **Reset patch settings**. The default location is the core server's \LDLogon\Patch directory.

Scheduling software update downloads

You can also configure software updates as a scheduled task to automatically occur at a set time in the future, or as a recurring task. To do this, click the **Schedule download** toolbar button to open the **Scheduled task properties** dialog where you can name the task and configure its options. When you click **Save**, the task appears in the Scheduled tasks window.

All scheduled software update tasks will use the current settings found in the **Update vulnerabilities settings** dialog. So, if you want to change the source site, platforms, languages,

patch download site, or proxy server settings for a particular update job, you must first change those settings in the **Update vulnerabilities settings** dialog before the task is scheduled to run.

To configure a Schedule download task

1. In the left navigation pane, click **Software updates**.
2. Click **Schedule download**.
3. On the **Schedule task** page, configure the [schedule](#).
4. Click **Save**.

When you click **Schedule download**, a task is created (it has no targeted devices, and it is unscheduled). If you cancel this **Scheduled tasks** procedure, be aware that it has still been created and appears in the **My tasks** list.

Viewing software update and detection rule information

After software updates have been updated with the latest information from the LANDesk security service, you can view software update lists in the console, view them by platform and product, and move the updates into different status groups. For information on the different groups in the window and how to use them, see [About the software updates window](#) earlier in this chapter.

To view software update details, double-click a Software updates ID to open its properties dialog. From this dialog you can also access detection rule details by double-clicking a patch file name in the **Detection rules** list to open the patch properties dialog (see [About the Patch properties dialog](#)).

This information can help you determine which updates are relevant to your network's supported server platforms, how an update's detection rules check for the presence of a vulnerability, what patches are available, and how you want to configure and perform remediation for affected devices.

You can also view software update definition and detection rule information specific to scanned devices directly from the console by accessing the server information console from **My devices**, and clicking **Software updates** in the left navigation pane.

Purging software updates information

You can purge software update information from the software updates window (and subsequently from the core database) if you determine that it isn't relevant to your environment.

When you purge software update information, associated detection rule information is also removed from the database. However, the actual patch executable files aren't removed by this process. Patch files must be removed manually from the local repository, which is typically on the core server.

To purge software updates information

1. Click the **Purge** toolbar button. (For a description of the dialog, see [About the Purge security and patch definitions dialog](#).)

2. Select the platforms whose software update information you want to remove. You can select one or more platforms in the list.

If an update is associated with more than one platform, you must select all of its associated platforms in order for the update's information to be removed.

3. Select the languages whose update information you want to remove (associated with the platform specified above).

If you select a Windows platform above, you should specify which language update information you want to remove. If you select a UNIX platform above, you must specify the Language-neutral option in order to remove cross-language update information.

4. Click **Remove**.

Scanning devices for software updates

Software update assessment means checking the currently installed versions of operating system-specific files and registry keys on a device against the most current known software updates in order to identify update needs on your servers. After reviewing known software update information (updated from industry sources) and deciding which updates you want to scan for, you can perform customized assessment on managed devices that have the Software updates agent installed. (For information on configuring devices for scanning and patch deployment, see "[Configuring devices for software update scanning](#)" earlier in this chapter.)

When the software update scanner runs, it always reads the contents of the Scan group and scans for those specific updates. Before scanning servers for updates, you should always make sure only the software updates you want to scan for are included in that group. You can move software updates into and out of the Scan group to customize the size and nature of a scan.

Running the software updates scanner

The software updates scanner can be pushed to devices as a scheduled scan task from the console.

To create a software updates scan task

1. In the left navigation pane, click **Software updates**.
2. Make sure software update definitions have been updated recently.
3. Make sure the Scan group contains only those updates you want to scan for.
4. Click the **Schedule patch tasks** toolbar button. (For a description of the dialog, see "About the Schedule Vulnerability Scan dialog.")
5. Enter a unique name for the scan. If the task script already exists, you can select whether to overwrite the existing script.
6. Specify whether you want the software updates scanner to display a progress dialog on the target device. You can also specify whether you want a Cancel button to appear with the scanner dialog, so that the end user has the option of canceling the scan.
7. Specify how you want the software updates scanner dialog to close when it is done running on target devices. You can require end user input, or you can set the dialog to close after a specified timeout period.

8. Click **OK**.
9. Select the task in the bottom pane (under **Vulnerability tasks** and click **Edit**). Set targeting and [scheduling](#) parameters, and click **Save**.

Viewing detected updates

If the software updates scanner discovers updates for any of the enabled software updates on any of the target devices, this information is reported to the core server and added to the **Detected** list.

You can use any of the following methods to view detected updates after running a software updates scan:

By the Detected group

Select the **Detected** group in the software updates window to view a complete listing of all updates detected by the most recent scan.

By an individual device

Double-click a device name in **My devices**, and then click **Software updates** to view detailed software updates assessment information for that device.

Downloading patches

In order to deploy patches to devices with detected software updates, the patch executable file must first be downloaded to a local patch repository on your network. The default location for patch file downloads is the core server's /LDLogon directory. You can change this location in the **Patch settings** tab of the **Update vulnerabilities settings** dialog.

Patch download location and proxy server settings

Patch downloads always use the download location settings currently found in the **Patch settings** tab of the **Update vulnerabilities settings** dialog. Also note that if your network uses a proxy server for Internet access, you must first configure the proxy server's settings in the **Proxy settings** tab of the **Update vulnerabilities settings** dialog before you can download patch files.

The product first attempts to download a patch file from the URL shown on the Patch Properties dialog. If a connection can't be made, or if the patch is unavailable for some reason, the product downloads the patch from the LANDesk security service, which is a company-hosted database containing patches from trusted industry sources.

You can download one patch at a time, or a set of patches together at the same time.

To download single patches

1. Double-click a software update name to open its **Properties** dialog.
2. In the **Detection rules** section, select the detection rule patch files you want to download, and then click **Download Selected Patches**.

3. The download operation and status displays in the **Downloading patches** dialog. You can click **Cancel** at any time to stop the entire download process.
4. When the download is finished, click the **Close** button.

To download multiple patches

All scheduled software update tasks will use the current settings found in the **Update vulnerabilities settings** dialog. So, if you want to change the source site, platforms, languages, patch download site, or proxy server settings for a particular update job, you must first change those settings in the **Update vulnerabilities settings** dialog before the task is scheduled to run.

1. In the left navigation pane, click **Software updates**.
2. Click **Schedule download**.
3. On the **Schedule task** page, configure the schedule.
4. Click **Save**.

Removing patch files

To remove patch files, you must delete the files manually from the patch repository, which is typically the core server's LDLogon directory.

Remediating software updates

Once you've updated software update definitions, put the updates you want to scan for in the Scan group, run a scan on managed devices, determined which software updates require attention, and downloaded the necessary patches, the next step is to perform software updates remediation by deploying and installing the necessary patches on affected devices.

Software updates remediation is done on an individual software update basis. In other words, you create a remediation task for a specific software update that deploys and installs the necessary patch files.

Note that remediation, like software update scanning, only works on devices that have been configured with the Software updates agent. For more information, see [Configuring devices for software update scanning](#) earlier in this chapter.

Linux remediation is supported. You can use the Software updates tool to discover vulnerabilities on Linux devices, then decide if you want to remediate the updates. If you want to do so, you can use a support subscription for your Linux vendor to download the necessary RPMs, then deploy the RPMs to devices.

Warning: Many patches will automatically reboot the device upon completion.

To create a custom remediation script

1. In the left navigation pane, click **Software updates**.

USERS GUIDE

2. Select the **Detected** group to view software updates detected by the most recent scan. (You don't have to select this group. If you want to create a custom remediation script for updates that haven't been scanned for or haven't been detected yet, click any of the other vulnerability groups to view their contents and select a specific vulnerability.)
3. Right-click the definition, then select **View affected devices** to view devices that are affected by this software update.
4. Right-click the definition, then select **Create remediation task**.
5. (optional) Modify the name in the **Task name** text box.
6. Select from the options, then click **OK**.
 - **Copy affected computers to the target cart:** Copies computers affected by the software update to the target cart for remediation.
 - **Show progress when running:** Enables the scanner to display information on end user devices while it is running. Click this option if you want to show scanner activity, and if you want to configure other display and interaction options in this dialog. If you don't click this option, none of the other options on this dialog are available to configure, and the scanner runs transparently on devices.
 - **Require user input before closing vulnerability scan dialog:** Click this option if you want the scanner to prompt the end user before its display dialog closes on the device. If you select this option and the end user does not respond, the dialog remains open, which could cause other scheduled tasks to timeout.
 - **Automatically close dialog after a timeout:** Click this option if you want the scanner's display dialog to close after the duration you specify.

Scripts

Managing scripts

This product uses scripts to execute custom tasks on devices. Completing the script creation dialogs generates an ASCII text file in the Windows INI format with an .INI extension. These scripts are stored on the core server in the \Program Files\LANDesk\ManagementSuite\Scripts folder. The script filename becomes the script name in the console. You can create local scheduler scripts for Windows devices using the **Scripts** window (click **Scripts** in the left navigation pane) or write your own script files and save them to the Scripts folder.

The **Scripts** window divides scripts into the following categories:

- **My scripts:** Scripts that you associate with this group.
- **All other scripts:** All scripts on the core server.
- **User scripts** (only visible to administrators): Scripts created by all product users. These are sorted by creator.

You can create groups under the **My scripts** item to further categorize your scripts. To create a new local scheduler script, click the **Local** button.

Once you've created a script, you can click **Schedule** on the script's shortcut menu. From the **My devices** window, you can target devices the task should run on, and you can schedule when the task should run from the **Scheduled tasks** window. See the Scheduling tasks section for more information on scheduling tasks.

Changes to script and task ownership for users of previous Management Suite versions

With Management Suite versions prior to 8.70, all scripts were global and all users could see them. Now scripts are only visible to the script creator and to administrators.

The **Scripts** window has a State column. The State column shows Public if all users can see the script, or Private if only the user that created the script or administrators can see it. Users can right-click scripts they have created and click Private or Public to change a script's state. Administrators can change the state of any script.

DOS PE is the default. If you select any other PE, you will not be able to create a command script. For Windows and Linux PEs, you can only generate a capture or deploy script.

If you choose Linux PE, you only have the LANDesk or Other imaging tool options. If you choose Windows PE, you have the LANDesk, Other and Microsoft* XImage.

Creating a local scheduler script

The local scheduler is a service that runs on devices. It's installed when you deploy an agent configuration as part of the standard management agent. Usually the local scheduler handles

product tasks, such as running the inventory scanner periodically. Other tasks that you schedule are handled by the core server rather than the local scheduler. You can use the local scheduler to schedule your own tasks to run periodically on devices. Once you create a local scheduler script, you can deploy it to your managed devices, just as you would any other script.

The local scheduler assigns each task an ID number. Local scheduler scripts have an ID range that is different from the default local scheduler scripts that the product uses. You can only have one custom scheduler script active on each device. If you create a new script and deploy it to devices, it will replace the old script (any script in the custom local scheduler ID range) without affecting the default local scheduler scripts, such as the local inventory scan schedule.

When selecting schedule options for the script, keep in mind that the restrictive qualities of the various options. For example, if you select Monday as the day of the week and 17 as the day of the month, the task will only execute on a Monday that's also the 17th of the month, which happens very infrequently.

You can create a script to run `restartmon.exe` on a local machine, immediately or any time you prefer. If reporting from a specific machine appears to have stopped, you can use `restartmon.exe` in the `LDClient` folder to restart the collector and all monitoring providers. This utility is for machines on which reporting has been installed, and reporting has stopped. Use this utility to restart the collector and providers without having to reboot the device.

1. In the left navigation pane, click **Scripts**.
2. Click **Local**.
3. Enter a script name.
4. Click **Add** to define the script options.
5. Configure the local scheduler options as described earlier. When finished, click **Save**.
6. Click **Save** to save your script.
7. Select the script in the **My scripts** group, then click **Schedule** to deploy the script you created to devices.

Understanding bandwidth options

When configuring local scheduler commands, you can specify the minimum bandwidth the managed device must have for the task to execute. When the time comes for the task to execute, each device running the local scheduler task will send a small amount of ICMP network traffic to the computer you specify and evaluate the transfer performance. If the test target computer isn't available, the task won't execute.

You can select these bandwidth options:

- **RAS:** The task executes if the device's network connection to the target computer is at least RAS or dialup speed. Selecting this option generally means the task will always run if the device has a network connection of any sort.
- **WAN:** The task executes if the device's connection to the target computer is at least WAN speed. WAN speed is defined as a non-RAS connection that's slower than the LAN threshold.
- **LAN:** The task executes when the device's connection to the target computer exceeds the LAN speed setting. LAN speed is defined as 262,144 bps by default.

Scheduling scripting tasks

The **Scheduled tasks** window shows scheduled task status and while the task is running and upon completion. The scheduler service has two ways of communicating with devices:

- Through the standard management agent (must already be installed on devices).
- Through a domain-level system account. The account you choose must have the login as a service privilege and you must have specified credentials in the Configure Services utility. For more information on configuring the scheduler account, see "[Configuring the scheduler service](#)."

LANDesk installs several standard scripts that you can schedule to perform routine maintenance tasks such as running inventory scans on selected devices. Click **Scripts** in the left navigation pane, then click **All other scripts** to view and schedule these scripts.

To schedule a task

1. In the left navigation pane, click **Scripts**.
2. Click to navigate to the script group.
3. Click a script, and click **Schedule**.
4. Type a name for the task, and click **OK**.
5. In the **Custom script tasks** tab, click **All tasks**, click the task you named in step 3, and click **Edit**.
6. Fill out the pages of the custom script task. Click the Help button for help on any page, or see the [Task scheduler](#) help.

When you click **Schedule**, a task is created (it has no targeted devices, and it is unscheduled). If you cancel this Scheduled task procedure, please be aware that the task has still been created and will appear in the Task list.

Using the default scripts

This product ships with two default scripts. You can use them to help you complete some typical tasks. These scripts are available under the **All other scripts** tree in the **Scripts** window (left navigation pane | **Scripts**).

- **inventory scanner:** Runs the inventory scanner on the selected devices. This script contains documentation that describes how to write a script file; read or print this script file for more information about correct use of commands and parameters.
- **Restore client records:** Runs the inventory scanner on selected devices, but the scanner reports to the core the device was configured from. If you have to reset the database, this task helps you add devices back to the proper core database in a multi-core environment.

Scheduling tasks

- [Custom task groups](#)
- [Target devices page](#)
- [Schedule task page](#)
- [Custom scripts page](#)
- [Moving tasks](#)

The **Scheduled tasks** tool is common to Agent configuration, Software updates, Scripts, and Device discovery. The tasks are filtered in the lower pane of the specific feature pages to show only related tasks. For example, if you open the **Device discovery** tool, discovery tasks are displayed in the **Discovery tasks** tab in the lower pane. All tasks are still visible through the **Scheduled tasks** tool. Here you can schedule configurations to run immediately, at some point in the future, on a recurring schedule, or run just once.

The left pane of the **Scheduled tasks** page shows these task groups:

- **My tasks:** Tasks that you have scheduled. Only you and administrative users can see these tasks.
- **All tasks:** Both your tasks and tasks marked as public.
- **Common tasks:** Tasks that users have marked common. Anyone who edits or schedules a task from this group will become the owner of that task. The task remains in the Common tasks group and will also be visible in the User tasks group for that user.
- **User tasks** (administrative users only) : Tasks users have created.

When you click **My tasks**, **Common tasks**, or **All tasks**, the right pane shows this information:

- **Task:** The task names.
- **Start On:** When the task is scheduled to run. Click a task name and click **Edit** to edit the start time or to reschedule it.
- **Status:** The overall task status. View the right pane Status column for more details. The right pane column shows the task status, which can be Working, All Completed, None Completed, or Failed.
- **Distribution package:** The package name the task distributes. This field applies to software distribution.
- **Delivery method:** The delivery method the task uses. This field applies to software distribution.
- **Owner:** The name of the person who originally created the script this task is using.

When you double-click a scheduled task, the right pane shows this summary information:

- **Name:** The task state name.
- **Quantity:** The number of devices in each task state.
- **Percentage:** The percentage of devices in each task state.

Before you can schedule tasks for a device, it must have the appropriate agent and be in the inventory database. Server configurations are an exception. They can target a device that doesn't

have the standard management agent. Tasks can be rescheduled (edited or deleted from the Tasks tabs). Once you schedule a task, see the Tasks tab for task status.

You can edit a task by selecting the task you want to edit and clicking **Edit**. The task opens with editing options applicable to the task.

Custom task groups

You can create custom groups for the task types **My tasks**, **All tasks**, and **Common tasks**. With custom groups, you can group related tasks such as scanning for vulnerabilities and running a script. Groups and subgroups can be 20 layers deep.

To create a custom task group

1. In the left navigation pane, click **Scheduled tasks**.
2. In the left pane, click the Task type in which you want to create the group.
3. Click **New group** on the toolbar.
4. Type a name in the **Group name** text box, and click **OK**

After you have created a custom group, you can move or copy tasks or other groups into the group by selecting them from a list and clicking **Move** on the toolbar.

About the Target devices page

Use this page to add device targets for the task you're configuring. You can also see the targeted devices, queries, and device groups for the task on this tab. If you have multiple LANDesk management products installed, device groups created in one product's console can be viewed in all consoles. This page is not needed for device discovery tasks.

- **Add target list:** Add the devices previously put in the target list from **My devices**.
- **Add query:** Targets the results of a query that you've previously created.
- **Remove:** Removes the selected targets.

While this page displays targeted device groups, note that groups will be displayed only if LANDesk Management Suite is installed on the core server. If you are running Server Manager, System Manager, or are running the Web console in Management Suite, device groups are not targeted as groups. Instead, if you select a group and target it, the individual devices in the group are added to the targeted devices list and will be displayed under **Targeted devices** rather than under **Targeted groups**.

About the Schedule task page

The Scheduler contains a **Scheduled task properties** tab that includes these options.

- **Leave unscheduled:** (default Leaves the task in the Task list for future scheduling).
- **Start now:** Runs the task as soon as possible. It may take up to a minute for the task to start, depending on other setting.

- **Start at scheduled time:** Starts the task at the time you specify. If you click this option, you must enter the following:
 - **Date:** The date you want the task to start. Depending on your locale, the date order will be day-month-year or month-day-year.
 - **Time:** The time you want the task to start.
 - **Repeat every:** If you want the task to repeat, click whether you want it to repeat every **Hour**, **Day**, **Week**, or **Month**. If you pick **Month** and the date doesn't exist in all months (for example, 31), the task will only run in months that have that date.
- **Schedule these devices:** For the first time a task runs, you should leave the default of Waiting or currently working. For subsequent runs, choose from All, Devices that did not succeed, or Devices that did not try to run the task. These options are explained in more detail below.
 - **Devices that did not succeed:** Select this if you only want the task to run on all devices that didn't complete the task the first time. This excludes devices that have a Successful state. The task will run on devices in all other states, including Waiting or Active. Consider using this option if you need the task to run on as many unsuccessful devices as possible, but you only need the task to complete successfully once per device.
 - **Waiting or currently working:** Select this if you want the task to run on devices that are waiting to be processed or are currently being processed.
 - **All:** Select this if you want the task to run on all devices, regardless of state. Consider using this option if you have a task, especially a repeating one, that needs to run on as many devices as possible.
 - **Devices that did not try to run the task:** Select this if you only want the task to run on devices that didn't complete the task and didn't fail the task. This excludes devices that were in an Off, Busy, Failed, or Canceled state. Consider using this option if there were a lot of target devices that failed the task that aren't important as targets.

About the Custom scripts page

- **Currently selected custom script:** Select the script you want to schedule.

Moving tasks

Use the **Move** dialog to move a task from one Task group to another, thus making the task visible to users who can view the task group you are moving the task to.

1. In the Scheduled tasks tab, select a task under **My tasks**, **All tasks**, **Common tasks**, or **User tasks**.
2. Click **Move**.
3. Select the task group you want to move the task to, then click **OK**.

Reports

About reports

System Manager includes a reporting tool you can use to generate a wide variety of specialized reports that provide critical information about the managed devices on your network.

System Manager uses an inventory scanning utility to add devices (and collected hardware and software data about those devices) to the core database. You can view and print this inventory data from a device's inventory view, as well as use it to define queries and group devices together. The reporting tool takes further advantage of this scanned inventory data by collecting and organizing that data in useful report formats.

You can use the predefined service reports and inventory asset reports. After running a report, you can view it from the console.

If you have Server Manager and Management Suite installed together, the reports you run in Server Manager will only include servers. If you run a query you will get both servers and other devices, unless the query has been configured to exclude other devices.

If reporting from a specific machine appears to have stopped, you can use `restartmon.exe` in the `LDClient` folder to restart the collector and all monitoring providers. This utility is for machines on which reporting has been installed, and reporting has stopped. Use this utility to restart the collector and providers without having to reboot the device.

Understanding report groups and predefined reports

Reports are organized in groups in the **Reports** window (left navigation pane | **Reports**). Administrators can view the contents of all of the report groups. System Manager includes a specific role, called Reports, to allow others to view reports without providing them access to other management capabilities. (For more information, see "[Role-based administration](#).") Users with the Reports right can also see and run reports, but only on the devices included in their scope.

The **Reports** window has the following groups of reports:

- Hardware
- Software
- Security and patch manager

Viewing reports

You can run any report from the **Reports** window.

From the **Reports** window, click a report group, then click the report you want to run. The report data displays in the **Report view**.

About the Report view window

Reports allow you to quickly access a graphical representation of the assets on your client computers. The reports are created from data the scanner stores in the database. You can view reports or print them through your browser.

To view a report

1. In the left navigation pane, click **Reports**. Report categories are listed in the right pane. Click a category heading to view the list of reports. An icon next to each report indicates the report type.



A report with a chart icon next to it displays as a pie or bar chart (two- or three-dimensional). In a chart, you can click on any colored bar or pie section to drill down to a summary.



A report with a document icon next to it displays as text.

2. Click the report name to view the report.
3. For the hardware or software scan date summaries, click the start and end dates to set the time frame, then click **Run**.

The Disk Space Summary report contains data for Windows-based devices only.

To print a report, right-click the page and click **Print**. On the Print dialog, click **Print**. If a report spans multiple pages, you must right-click in each page to print it.

To distribute a report

- To e-mail a report, the recommended method is to print the report to a .PDF file, then attach it to the e-mail.

The console displays report charts as pie or bar charts. To set the chart type, click the drop-down list in the report chart, then change the chart type.

In order to view the interactive bar and pie charts displayed in many reports, you must have Macromedia Flash Player* 7 installed.

Queries

Using queries

Queries are customized searches of core databases. This product provides tools that let you create database queries for devices in your core database. You create core database queries in the console's **Query** view. System Manager public queries are visible in LANDesk® Management Suite, and vice-versa, if both are being used.

Read this section to learn about:

- [Queries overview](#)
- [Query groups](#)
- [Creating database queries](#)
- [Running queries](#)
- [Importing and exporting queries](#)

Queries overview

Queries help you manage your network by allowing you to search for and organize devices in the core database based on specific system or user criteria.

For example, you can create and run a query that captures only devices with a processor clock speed of less than 1 GHz, or with less than 512 MB of RAM, or with a hard drive of less than 80 GB. Create one or more query statements that represent those conditions and relate statements to each other using standard logical operators. When the queries are run, you can print the results of the query, and access and manage the matching devices.

Query groups

Queries can be associated with groups in the **My devices** view. These are called dynamic groups, and the contents of a dynamic group are the result of the query associated with that dynamic group. For example, a group comprising all the devices in a geographic area can be associated with a query on memory, hard disk size, and so forth.

For more information on how query groups and queries display in the **All devices** view, and what you can do with them, see "[Grouping devices for actions.](#)"

Creating database queries

Use the **New query** dialog to build a query by selecting from attributes, relational operators, and attribute values. Build a query statement by choosing an inventory attribute and relating it to an acceptable value. Logically relate the query statements to each other to ensure they're evaluated as a group before relating them to other statements or groups.

To create a database query

1. In the console's **Queries** view, click **New**.
2. Select a **component** from the inventory attributes list.
3. Under **Step 1: Search conditions**, click **Edit**.
 1. Drill down this list to select the attributes that will be your search condition. For example, to locate all clients running a particular type of software, you would select `Computer.Software.Package.Name`.
 2. After selecting the attributes, you'll notice that a series of fields appear in the right side of the window. From these fields, select an operator and value to complete the search condition. For example, to locate all clients running Internet Explorer 6.0, the attributes would be "`Computer.Software.Package.Name`," the operator "=", and the value "Internet Explorer 6."
 3. At the bottom of the window, click **Add** to fill in the empty field with your search condition.
 4. You can continue to refine the query by creating another search condition, then adding it to the first with a boolean operator (AND or OR) . Also use the buttons to add, delete, replace, group, or ungroup the conditions you create.
 5. When you're finished, click **OK**.
4. Under **Step 2: Attributes to display**, click **Edit**.
 1. Drill down this list to select an attribute to display in the query results list. Remember to select attributes that will help you identify the clients returned in the query. If you cannot find attributes you want to display, you can add them in the [Custom attributes](#) dialog. However, these attributes must be assigned to machines before they appear in the query dialog.
 2. After you've selected an attribute, click **Add** to move it into the empty field on the bottom of the window. If you want to enumerate your query results list, click **Include count**.
 3. Repeat the process if you want to add more attributes. Use the **Remove** button to remove attributes, and click **Move up/Move down** to change the order of attributes.
 4. Click **Make results targetable** to enable the results of the query to be targetable for any actions you specify.
 5. When you're finished, click **OK**.
5. (optional) Under **Step 3: Sort results by attribute**, click **Edit** to customize the order of query results.
6. If you want to run the query additional times, click **Save query**, and enter a unique name for the query. If you run the query prior to saving it, the query parameters are lost and must be reconstructed to run the same query again.
7. Under **Step 4: Run query**, click **Run query**.

Query statements are executed in the order shown

If no groupings are made, the query statements listed in this dialog are executed in order from the bottom up. Be sure to group related query items so they're evaluated as a group; otherwise, the results of your query may be different than you expect.

Running queries

To run a query

1. In the left navigation pane, click **Queries**.
2. Select the query and click **Run**.

or

To make changes to the query before running it, double-click the query, click **Edit**, modify steps 1-3, and then click **Run query**.

Note: If you have modified the query and want to save your changes, click **Save query** to save the changes or **Save query as** to give the modified query a new name. Do this before running the query. If you do not save your changes before running the query, the changes will not be saved with the query.

3. The results (matching devices display in the right-hand pane of the **All devices** view).

Importing and exporting queries

You can use import and export to transfer queries from one core database to another. Exported queries are saved as .XML files.

To import a query

1. Right-click the query group where you want to place the imported query.
2. Select **Import** from the shortcut menu.
3. Navigate to the query you want to import and select it.
4. Click **Open** to add the query to the selected query group in the **All devices** view.

To export a query

1. Right-click the query you want to export.
2. Select **Export** from the shortcut menu.
3. Navigate to the location where you want to save the query (as an .XML file).
4. Type a name for the query.
5. Click **Save** to export the query.

Understanding custom queries

Custom queries are useful when you want inventory details about hardware and software installed on your devices. Use a custom query to build a list of computers that have similar inventory. Custom queries are also used to define groups and scopes.

The **Custom queries** page (click **Queries** in the left navigation pane) displays a list of queries that you have saved. To run a saved query, select the query, then select **Run**.

If the query list spans multiple pages, use the arrows at the top of the page to navigate between pages. Enter the number of items to display per page and click **Set**.

Creating custom queries

Custom queries are useful when you want inventory details about hardware and software installed on your devices. Use a custom query to build a list of devices with similar inventory. For example, if you want to upgrade all devices to at least a 2 GHz processor, you can query for all devices in your database with processor speeds of less than 2 GHz. Custom queries are also used to define groups and scopes.

You can query on any of the inventory items (known as "attributes") that the inventory scanner stores in the database, as well as any custom attributes.

Managing queries

Manage queries in the **Queries** view. Use this view to create, edit, or delete queries:

- To run an existing query, select it and click **Run**.
- To create a new query, click **New**. Once you have created and saved that query, its name will appear in the list on this page.
- To edit a query in the list, double-click it. The **Edit query** page appears with query parameters you can edit.
- To edit the most recent query, click **Edit current query**.
- To delete a query, select the query and click **Delete**.

Creating a query is a four-step process:

1. **Create a search condition:** Specify a set of inventory attributes that will be the basis of your query.
2. **Select attributes to display:** Refine or "filter" the query so that the results display the attributes most useful to you, such as IP addresses or computer device names.
3. **Sort results by attributes (optional):** Specify how you want the query results sorted. (Only applies if, in Step 2, you selected to display more than one type of attribute in the query results.)
4. **Run the query:** Run the query you just created. You can also save it for later use, or clear all of the query information to begin again.

Step 1: Creating a search condition (required)

A search condition is a set of inventory attributes and associated values that you query for. You can use one search condition or group several together to form the basis of a query.

The following steps take place on the **Edit query** page. From the **Run queries** view, click **New**, or select an existing query and click **Edit**.

To create a search condition

1. Under **Step 1**, click **Edit**. A window appears showing a list that represents all of the inventory data currently in the database.
2. Drill down this list to select the attributes that will be your search condition. For example, to locate all clients running a particular type of software, you would select `Computer.Software.Package.Name`.
3. After selecting the attributes, you'll notice that a series of fields appear in the right side of the window. From these fields, select an operator and value to complete the search condition. For example, to locate all clients running Internet Explorer 6.0, the attributes would be "`Computer.Software.Package.Name`," the operator "=", and the value "Internet Explorer 6."
4. At the bottom of the window, click **Add** to fill in the empty field with your search condition.
5. You can continue to refine the query by creating another search condition, then adding it to the first with a boolean operator (AND or OR). Also use the buttons to add, delete, replace, group, or ungroup the conditions you create.
6. When you're finished, click **OK**.

To run and store a query on the health status of servers (`Computer.Health.State`), you should be aware that the state in the database is represented by a number. Use the table below to create search conditions. For example, to create a search condition for machines with "Unknown" health, use the operator "NOT EXIST".

Health condition	Operator
Unknown	NOT EXIST
Normal	2
Warning	3
Critical	4

Step 2: Selecting attributes to display (required)

For Step 2, select the attributes that will be most useful for identifying computers returned in the query results. For example, if you want results that help you physically locate each computer matching the search condition set in Step 1, you would specify attributes such as each computer's display name (`Computer.DisplayName`) or IP address (`Computer.Network.TCPIP.Address`).

The following steps take place on the **Edit query** page.

To select attributes to display

1. Under **Step 2**, click **Edit**. A window appears showing a list that represents all of the inventory data currently in the database.

2. Drill down this list to select an attribute to display in the query results list. Remember to select attributes that will help you identify the clients returned in the query. If you cannot find attributes you want to display, you can add them in the [Custom attributes](#) dialog. However, these attributes must be assigned to machines before they appear in the query dialog.

Note: If you're using an Oracle database, make sure you select at least one attribute that is natively defined by the inventory scanner (for example, Computer.Display Name, Computer.Device Name, Computer.Device ID, Computer.Login Name, and so on).

3. After you've selected an attribute, click **Add** to move it into the field on the bottom part of the window. If you want to enumerate your query results list, click **Include count**.
4. Repeat the process if you want to add more attributes. Use the arrow buttons to add or remove attributes, and click **Move up/Move down** to change the order of attributes.
5. Click **Make results targetable** to enable the results of the query to be targetable for any actions you specify.
6. When you're finished, click **OK**.

You can also add column heading(s) to your query results list.

To change column headings (optional)

1. Under **Step 2**, click **Edit**.
2. In the bottom box, click a column heading and click **Edit**. Edit the heading and press **Enter**. Repeat as necessary.
3. Click **OK**.

At this point, you may want to save your query; the next procedure in the query-creation process is optional and applies only to query results that contain two or more columns. To save your query, click **Save Query** at the top of the page. A window appears prompting you to type a name for this query. Type a name, then click **Save** in the top right corner of the window.

Step 3: Sorting results by attribute (optional)

This procedure is necessary only if you defined more than one attribute and column heading in Step 2 and now want to sort the results alphabetically or numerically within one of those columns.

For example, let's say you specified two different attributes to display in the query results: the IP address and the processor type of each returned computer. In Step 3, you could sort alphabetically by processor type in the results.

If you skip this step, the query will automatically sort by the first attribute selected in Step 2.

To sort results by attribute

1. Under **Step 3**, click **Edit**. A window appears showing the attributes you selected in **Step 2**.
2. Select which attribute you want to sort by, then click **>>** to move it over to the empty text box.
3. Click **OK**.

Step 4: Running the query

After creating your query, you can run, save, or clear it to start over.

To save the query for future use, click the **Save** toolbar button. The query now appears in the list on the **Custom queries** page. If your query is a modified version of another, click the **Save as** toolbar button to give it a new name.

By default, saved queries are only visible by the person who saved them. If you check **Public query** before saving, the saved query will be visible to all users. Only administrators with the public query management right can make a query public.

If you have multiple product family products installed, queries are shared between products. If you save a query in one product's console, it will also be visible in other product consoles.

To view the results of this query, click the **Run** toolbar button.

To clear the query parameters from the **Edit query** page, click the **Clear** toolbar button. If the query has already been saved, it's cleared from this page but remains in the **Custom queries** list.

Viewing query results

Query results match the search criteria you specified in the query-building process. If the results aren't what you expected, go back to the **Edit query** page and refine the information.

To drill down to more information about one of the devices in the list of query results, double-click the query data or right-click and click **View computer** in the resulting menu.

From the **Query results** page, you can click the **Save as CSV** toolbar button to export the results into a format compatible with spreadsheet or other applications.

To print the query results, click **Print view** in the query results page.

Viewing drill-down query results

Query results match the search criteria you specified in the query-building process. If the results aren't what you expected, go back to the **Edit query** page and refine the information.

To drill down to more information about one of the devices in the list of query results, double-click the query data or right-click and click **View computer** in the resulting menu.

Exporting query results to CSV files

To view your query results data in a spreadsheet application, export the data as a comma-separated values (CSV file). From the **Query results** page, click the **Save as CSV** toolbar icon to save your information as a CSV file. You can then use an application like Microsoft Excel* to import and work with the CSV file.

Changing query column headings

1. Open an existing query or create a new query.
2. In the bottom box, click a column heading and click **Edit**. Edit the heading and press **Enter**. Repeat as necessary.
3. Click **OK**.

Exporting and importing queries

You can export and import any queries you create. All queries export as XML files. If you export the same query filename more than once, it will overwrite the existing file. To avoid this, you may want to copy the file to another location once it's exported.

The export and import features are useful in two scenarios:

- If you need to reinstall your database, use the export/import features to save your existing queries for use in a new database.

For example, you could export the queries, then move them to a directory unaffected by a database reinstall. After reinstalling the database, you could move the queries back into the queries directory on your Web server, then import them into the new database.

- You can use the export/import features to copy queries to other databases.

For example, you could export a query to a queries directory on your Web server, then e-mail or FTP it to someone. That person could then place the queries into the queries directory on another Web server, then import them into a different database. You could also map a drive and directly copy queries into the queries directory on another Web server.

To export a query

Complete these steps while connected to a database that has a query you want to export.

1. In the left navigation pane, click **Queries**.
2. On the **Custom queries** page, click the query name you want to export. Click **Edit**.
3. On the **Edit query** page, click the **Export** toolbar button to export the query to disk.
4. On the **Query exported** page, right-click the query to download it as an XML file to a selected directory. The query becomes the XML file.

Note that if you export the same query filename more than once, it will overwrite the existing file. To avoid this, you may want to copy the file to another location once it's exported.

If you want to eventually import the query back into a database, you must move it to the queries directory recognized by the Web server, by default c:\inetpub\wwwroot\LANDesk\LDSM\queries.

To import a query

Complete these steps while connected to a database to which you want to import a query.

1. In the left navigation pane, click **Queries**.
2. On the **Custom queries** page, click **New**.
3. On the **Edit query** page, click the **Import** toolbar button.
4. Select the query you want to import. If you want to verify the parameters of this query before importing it, click **View**.
5. Click **Import** to load the query in the **Edit query** page.
6. Once the query is loaded, scroll down and click **Save query** to save it into this database.

Inventory management

Managing inventory

You can use the inventory scanning utility to add devices to the core database and to collect devices' hardware and software data. You can view, print, and export inventory data. You can also use it to define queries, group devices together, and generate specialized reports.

Read this section to learn about:

- [Inventory scanning overview](#)
- [Viewing inventory data](#)

Inventory scanning overview

When you configure a device with the device setup feature, the inventory scanner is one of the components that gets installed on the device. When creating a client configuration, you can specify when the inventory scanner runs on the device.

The inventory scanner runs automatically when the device is initially configured. The scanner executable is named LDISCN32.EXE for Windows and LDISCAN for Linux. The inventory scanner collects hardware and software data and enters it into the core database. After that, the hardware scan runs at the interval you specify; on Windows devices it also runs each time the device is booted. To configure the software scan settings, on the core server, click **Start | Program Files | LANDesk | LANDesk Configure Services**.

For more information on configuring the inventory service, see "[Configuring the Inventory service](#)" in Appendix C.

After the initial scan, the inventory scanner can be run from the console as a scheduled task. The standard management agent must be running on remote devices to schedule an inventory scan to them.

Note: A device added to the core database using the discovery feature has not yet scanned its inventory data into the core database. You must run an inventory scan on each device for full inventory data to appear for that device.

You can view inventory data and use it to:

- Customize the **All devices** list columns to display specific inventory attributes
- Query the core database for servers with specific inventory attributes
- Group devices together to expedite management tasks
- Generate specialized reports based on inventory attributes
- Keep track of hardware and software changes on devices, and generate alerts or log file entries when such changes occur

Read the sections below to learn more about how the inventory scanner works.

Delta scanning

After the initial full scan is run on a Windows device, subsequent running of the inventory scanner only captures delta changes and sends them to the core database. Use the scanner option /RSS to gather software information from the Windows registry. (This option does not apply to Linux devices.)

Forcing a full scan

If you want to force a full scan of a Windows device's hardware and software data, you can delete the existing delta scan file and change a setting in the **Configure Services** applet. (This option does not apply to Linux devices.)

1. Delete the **invdelta.dat** file from the server. A copy of the latest inventory scan is stored locally as a hidden file named **invdelta.dat** on the root of the hard drive. (The **LDMS_LOCAL_DIR** environment variable sets the location for this file.)
2. Add the **/sync** option to the inventory scanner utility's command line. To edit the command line, click **Start | All Programs | LANDesk Management**, right-click the **Inventory Scan** shortcut icon, select **Properties | Shortcut**, then edit the **Target** path.
3. At the core server, click **Start | All Programs | LANDesk | LANDesk Configure Services**.
4. Click the **Inventory** tab, then click **Advanced settings**.
5. Click the **Do Delta** setting. In the **Value** box type **0**.
6. Click **OK** twice, then click **Yes** at the prompt to restart the service.

Scan compression

Inventory scans performed by the Windows inventory scanner (LDISCN32.EXE) are compressed by default. The scanner compresses full scans and delta scans with approximately an 8:1 compression ratio. Scans are first built completely in memory, then compressed and sent to the core server using a larger packet size. Scan compression requires fewer packets and reduces bandwidth usage. (Scan compression is not used on Linux devices.)

Scan encryption

Inventory scans on Windows devices are encrypted (TCP/IP scans only). You can disable inventory scan encryption by changing a setting in the LANDesk Configure Services applet.

1. At the core server, click **Start | All Programs | LANDesk | LANDesk Configure Services**.
2. Click the **Inventory** tab, then click **Advanced settings**.
3. Click the **Disable Encryption** setting. In the **Value** box type **1**.
4. Click **Set**, then click **OK**.
5. Click **OK**, then click **Yes** at the prompt to restart the service.

Offloading the core server

You can offload inventory scans onto a machine that is not your core server, provided that the machine has an inventory service installed. This can offload traffic and storage space requirements from your core server.

1. At the core server, click **Start | All Programs | LANDesk | LANDesk Configure Services**.
2. Click the **Inventory** tab, then click **Advanced settings**.
3. Click the **Off-core Inventory Server** setting. In the **Value** box, type the name of the machine on which you have installed an inventory service.
4. Click **Set**, then click **OK**.
5. Click **OK**, then click **Yes** at the prompt to restart the service.

Viewing inventory data

Once a device has been scanned by the inventory scanner, you can view its system information in the console.

Device inventories are stored in the core database, and include hardware, device driver, software, memory, and environment information. You can use the inventory to help manage and configure devices, and to quickly identify system problems.

You can view inventory data in the following ways:

- [Summary inventory](#)
- [Full inventory](#)
- [Viewing attribute properties](#)
- [System information](#)

You can also view inventory data in reports that you generate. For more information, see "[Reports overview](#)."

Viewing summary inventory from the server information console

Summary inventory is found on the **Summary** page in the server information console and provides a quick look at the device's basic OS configuration and system information.

Note: If you added a device to the core database using the discovery tool, its inventory data isn't yet scanned into the core database. You must run an inventory scan on the server for the summary inventory feature to complete successfully.

To view summary inventory

1. In the console's **All devices** view, double-click a device.
2. In the left navigation pane, click **System information**, then click **System summary**.

Windows 2000/2003 server summary data

This information appears when you view summary inventory for a Windows 2000/2003 server.

- **Health:** The current health state of the server.
- **Type:** The type of server, such as application, file, e-mail, and so forth.
- **Manufacturer:** The manufacturer of the server.
- **Model:** The server's model type.
- **BIOS version:** The version of the ROM BIOS.
- **Operating system:** Windows or Linux OS running on the server: 2000, 2003, or Red Hat.
- **OS Version:** Version number of the Windows 2000/2003 or Linux OS running on the server.
- **CPU:** Type of processor or processors running on the server.
- **Vulnerability scanner:** The version of the agent installed.
- **Inventory scanner:** The version of the agent installed.
- **Monitoring:** The version of the monitoring scanner installed
- **Last reboot:** The last time the server was rebooted.
- **CPU usage:** The percentage of the processor currently in use.
- **Physical memory used:** Amount of RAM available on the server.
- **Virtual memory used:** Amount of memory available to the server, including RAM and swap file memory.
- **Drive space used:** The percentage of drive space currently used. If you have more than one hard drive, each drive will be listed.

Servers that are IPMI-enabled display additional IPMI-specific data. Linux servers also display similar information in the **Summary** view.

Viewing a full inventory

A full inventory provides a complete listing of a device's detailed hardware and software components. The listing contains objects and object attributes.

To view a full inventory

1. In the console's **All devices** view, click a device.
2. In the **Properties** tab, click **View inventory**.

Viewing attribute properties

You can view attribute properties for a device's inventory objects from the inventory listing. Attribute properties tell you the characteristics and values for an inventory object. You can also create new custom attributes and edit user-defined attributes.

To view an attribute's properties, click the attribute in the left pane.

To print this information in Internet Explorer, right-click in the frame and click **Print**. To print in Mozilla, right-click in the frame, click **This Frame | Save Frame As**, click **Save**, then open the file in an application and click **Print**.

System information

From the server information console, you can view and modify the device's system information. Information in the **Hardware**, **Software**, **Logs** and **Other** categories is either stored data or real-time data. When you click an information link you can view detailed information about the selected component and, in appropriate cases, set thresholds and enter information.

1. In the console's **All devices** view, double-click a device.
2. In the left navigation pane of the server information console, click **System information**.
3. Expand the group and click the information link you want to view.

Customizing inventory options

The console includes a Configure Services utility that you can use to customize inventory options. The defaults for most options should be fine, but if you need to change them, you can run this utility. To launch the Configure Services applet, on the core server, click **Start | Program Files | LANDesk | LANDesk Configure Services**. (The filename for this utility is svccfg.exe.)

Use the Configure Services utility to configure:

- The database name, username, and password
- Device software scan interval, maintenance, days to keep inventory scans, and client login history length
- Duplicate device ID handling
- Scheduler configuration, including scheduled job and query evaluation intervals
- Custom job configuration, including remote execute timeout

Click **Help** on each tab in the Configure Services utility for more information.

Editing the LDAPPL3.TEMPLATE file

Information relating specifically to the scanner's inventory parameters is contained in the LDAPPL3.TEMPLATE file. This template file works with the LDAPPL3.INI file to identify a device's software inventory. This file is placed on managed Windows devices as part of agent configuration. Its parameters are set in the Inventory tab of [Agent configuration](#).

On Linux devices, a similar configuration file (/etc/ldappl.conf) contains information about the scanner's parameters. You can edit this file to change how the scanner operates. The file contains instructions for how to modify the Linux scanner's operation.

You can edit the template file's [LANDesk Inventory] section to configure the parameters that determine how the scanner identifies software inventory. By default, LDAPPL3.TEMPLATE is located in the core server's LDLogon share.

Use this table as a guide to help you edit the [LANDesk Inventory] section in a text editor.

Option	Description
Mode	<p>Determines how the scanner scans for software on devices. The default is Listed. Here are the settings:</p> <ul style="list-style-type: none"> • Listed: Records the files listed in LDAPPL3. • Unlisted: Records the names and dates of all files that have the extensions listed on the ScanExtensions line but that are not defined in the LDAPPL3. This mode helps discover unauthorized software on the network. • All: Discovers listed and unlisted files.
Duplicate	<p>Records multiple instances of files. Set the value to OFF to record only the first instance, or ON to record all detected instances. The default is ON.</p>
ScanExtensions	<p>Sets the file extensions (.EXE, .COM, .CFG, etc.) that will be scanned. Use a space to separate the file extensions. By default, only .EXEs are scanned.</p>
Version	<p>The version number of the LDAPPL3 file.</p>
Revision	<p>The revision number of the LDAPPL3 file; helps ensure future compatibility.</p>
CfgFiles 1-4	<p>Records the date, time, file size, and contents of the specified files. You can leave out the drive letter (for example, c:) if you want to search all local drives. You can specify more than one file on each of the four lines, but the line length is limited to 80 characters.</p> <p>Separate path names on the same line by a space.</p> <p>The scanner compares the date and size of the current file with that of the previous scan. If the date and size don't match, the scan records the contents of the file as a new revision.</p>
ExcludeDir 1-3	<p>Excludes specific directories from a scan. You can leave out the drive letter (for example, c:) if you want to exclude all local drives. Enumeration must start at 1 and be continuous. You must end each line with "\".</p>
MifPath	<p>Specifies where MIF files are stored on a client's local drive. The default location is c:\DMINDOS\MIFS.</p>

Option	Description
UseDefaultVersion	If set to TRUE, the scanner reports a match when a file matches an exact filename and file size entry in LDAPPL3 on filename only (the version will be reported as EXISTS). This can cause some false positives for applications that share a common filename with an unknown application. In the as-delivered LDAPPL3.TEMPLATE file, this parameter is set FALSE; that is, only add an entry if the match is exact. If the parameter is missing, it defaults to TRUE.
SendExtraFileData	If set to TRUE, sends extra file data to the core server. The default is FALSE. This means that by default, only path, name, and version are entered into the core database.

To edit the LDAPPL3.TEMPLATE file

1. From your core server, go to the \Program Files\LANDesk\ManagementSuite\LDLogon directory and open LDAPPL3.TEMPLATE in Notepad or another text editor.
2. Scroll down to the parameter you're interested in updating and make your changes.
3. Save the file.

Updating the application list

The data from the applications list, DEFAULTS.XML, is stored on the core database. Because the names and version numbers of commonly-used software applications change fairly often, LANDesk publishes a new DEFAULTS.XML several times a year (in previous versions of LANDesk software, this file was named LDAPPL.INI).

To update the application list

1. Download a new DEFAULTS.XML or LDAPPL3.TEMPLATE file from <http://www.landesk.com/support/downloads>. Select a product and click **Software update** to download the file.
2. Save the file to the LDLOGON directory.
3. Publish a new LDAPPL3.INI by following the steps in "[Publishing the application list](#)."

Publishing the application list

Publishing the Application list involves importing the most current application list in DEFAULTS.XML into the database, and then combining the application list with the contents of LDAPPL3.TEMPLATE to generate an updated LDAPPL3.INI file. There is a standalone utility COREDBUTIL.EXE in the \Program Files\LANDesk\ManagementSuite directory which is used to automatically perform both of these steps.

To publish the application list

1. Start CoreDBUtil.exe

2. Click the **Publish App List** button.

You should publish the application list after modifying or downloading an updated version of LDAPPL3.TEMPLATE or DEFAULTS.XML.

Hardware configuration

Intel* AMT support

System Manager supports devices using Intel* Active Management Technology (Intel* AMT), a hardware and firmware functionality that enables remote device management. Intel AMT uses out-of-band (OOB) communication for access to devices regardless of the state of the operating system or power to the device.

Intel AMT support in this product includes versions 1 and 2. The process for provisioning Intel AMT 2 devices includes some new features not found in version 1. See [Configuring Intel AMT devices](#) for details about provisioning with version 2. The information in this section applies to both versions except as noted.

The Hardware configuration tool includes the following features for managing Intel AMT devices:

- [Automatic generation of provisioning IDs \(PID\) and PPS \(version 2\)](#)
- [Changing the username and password for managed devices](#)
- [Configuring and enabling System Defense policies \(version 2\)](#)
- [Configuring and enabling Agent Presence monitoring \(version 2\)](#)

Managing devices with or without management agents

When devices are configured with Intel AMT, a limited number of management features are available even if the device does not have a LANDesk agent installed. As long as devices are connected to the network and have standby power, they can be discovered and can be added to inventory to be managed with other devices on the network.

If a device has Intel AMT but no management agent installed, it can be discovered with unmanaged device discovery, moved to the inventory database, then viewed in the **My devices** list. However, many System Manager management options are unavailable. These options are only made available when the LANDesk agent is installed. Management features that are available for Intel AMT-configured devices include:

- **Inventory summary:** a subset of the normal inventory data can be queried and viewed in real time for the device even if the device is powered off.
- **Event log:** a log with Intel AMT-specific events, showing severity and description of the events, can be viewed in real time.
- **Remote boot manager:** power cycling and several boot options can be initiated from the remote management console, regardless of the state of the device's OS or power. The options available are based on the support for the options on the device. Some devices may not support all boot options.
- **Force vulnerability scan and disable OS network:** if a device appears to have malicious software running, a vulnerability scan can be run at the next reboot; if necessary, the device's OS-level network access can be disabled to prevent unwanted packets from being spread on the network.

For more information about management options, see [Managing Intel AMT devices](#).

Intel AMT version 1 provisioning requirements

Devices can be discovered as Intel AMT devices only after you have accessed the Intel AMT Configuration Screen on the device and changed the manufacturer's default password to a secure password. (Refer to the manufacturer's documentation for information on accessing the Intel AMT Configuration Screen). If you have not done this, the devices will be discovered but not identified as Intel AMT devices, and you will not be able to view the same inventory summary information as you otherwise would.

In order for the core server to authenticate with discovered Intel AMT devices, the username/password credentials must match the credentials that you configure using the Configure Services utility. You can change the credentials using the Intel AMT Configuration Screen.

When an Intel AMT device is added to the core database to be managed, System Manager automatically provisions it in the mode you select in the Configure Services utility, regardless of whether it has already been provisioned. Small business mode provides basic management without network infrastructure services and is non-secure, while Enterprise mode is designed for large enterprises and uses DHCP, DNS, and a TLS certificate authority service to ensure secure communication between the managed device and the core server.

When you provision an Intel AMT device in Enterprise mode, the core server installs a certificate on the device for secure communication. If the device is to be managed by another core server, it must be unprovisioned and then re-provisioned by the new core server. If not, the device's Intel AMT access will not respond because the new core server does not have a matching certificate. Similarly, if any other computer attempts to access the Intel AMT functionality on the device, it will not succeed because it does not have a matching certificate.

Configuring Intel* AMT devices

Devices equipped with Intel AMT functionality should be configured when they are first set up and powered on. The provisioning process includes several security measures to ensure that only authorized users have access to the Intel AMT management features.

Intel AMT devices communicate with a provisioning server on the network. This provisioning server listens for messages from Intel AMT devices on the network and allows IT staff to manage servers through out-of-band communication regardless of the state the device's OS is in. System Manager acts as a provisioning server for Intel AMT devices and includes features that help you provision devices when you set them up. You can then manage the devices with or without additional System Manager management agents.

This section outlines a recommended process for configuring new Intel AMT (version 2) devices. During this process you will use System Manager to generate a set of provisioning IDs (PID and PPS). These IDs, when entered on the device's Intel AMT Configuration Screen, ensure a secure connection with the provisioning server that is aware of these IDs so the Intel AMT device can complete its initial provisioning process.

Devices with Intel AMT version 1 use a similar process but don't use the PID and PPS keys. See the notes at the end of this section for details.

Provisioning for Intel AMT 2 devices

When an Intel AMT 2 device is received, the IT technician assembles the computer and powers it on. After powering on the device, the technician logs in to the BIOS-based Intel ME (Management Engine) Configuration Screen and changes the default password (admin) to a strong password. This allows access to the Intel AMT Configuration Screen.

In Intel AMT Configuration Screen, the following pre-provisioning information is entered:

- A provisioning ID (PID)
- A pre-provisioning passkey (PPS) , also known as a pre-shared key (PSK)
- The IP address of the provisioning server
- Port 9982 as the port for communicating with the provisioning server
- Enterprise mode should be selected
- The host name of the Intel AMT device

The PPS must be known by the provisioning server and the managed device, but cannot be transmitted on the network for security purposes. It needs to be entered manually on the device (at the Intel AMT Configuration Screen) and stored on the provisioning server, which in this case is also the core server for System Manager. PID/PPS pairs are generated by System Manager and stored in the database. You can print a list of generated ID pairs for use in provisioning.

The IT technician should enter the IP address of the System Manager core server as the provisioning server and specify port 9982. Otherwise, by default, the Intel AMT device sends a general broadcast that can be received if the configuration server is listening on port 9971.

The default username and password for accessing the Intel AMT Configuration Screen are "admin" and "admin". These are changed during the provisioning process. The username can stay the same but the password must be changed to a strong password. The new username/password combination is entered in the Configure Services utility that is included with System Manager, as described in the procedural steps below. After each device is configured you can change the username/password individually per device, but for provisioning purposes you use the username/password that is found in Configure Services.

After the above information is entered in the Intel AMT Configuration Screen, the device sends "hello" messages when it is first connected to the network, attempting to communicate with the provisioning server. If this message is received by the provisioning server, the provisioning process will begin as the server establishes a connection with the managed device.

When the core server receives the hello message and verifies the PID/PPS keys, it provisions the Intel AMT device to TLS mode. TLS (Transport Layer Security) mode establishes a secure channel of communications between the core server and the managed server while the provisioning is completed. This process includes creating a record in the database with the device's UUID and encrypted credentials. When the device's data is in the database, the device appears in the list of unmanaged devices.

When an Intel AMT device has been provisioned by the core server, it can be managed using only Intel AMT functionality. You can select it in the list of unmanaged devices and add it to your managed devices. You can also deploy System Manager management agents to the device to use a larger set of management features.

The recommended process for using System Manager to provision Intel AMT 2 devices is as follows. Specific instructions for items 1 and 2 are given in the following procedural steps.

1. Run the Configure Services utility to specify a new, strong password for provisioning Intel AMT devices. (See detailed steps below).
2. Use System Manager to generate a batch of Intel AMT provisioning IDs (PID and PPS), and print the list of keys. (See detailed steps below).
3. Log in to the device's Intel ME Configuration Screen from the BIOS and change the default password to a strong password.
4. Log in to the Intel AMT Configuration Screen. Enter a PID/PPS key pair from the list of provisioning IDs that you printed. Enter the IP address of the core server (provisioning server), and specify port 9982. Make sure Enterprise mode is selected for provisioning. Enter the host name of the Intel AMT device.
5. After you exit the BIOS screen the device will begin sending "hello" messages.
6. The core server receives a "hello" message and checks the PID/PPS against the list of generated keys. If there is a match, it provisions the device to TLS mode.
7. The device is added to the unmanaged device discovery list.
8. Select the device and add it to your managed devices. It will be managed as an agentless device by default, and you can also deploy management agents to it.

To set the Intel AMT username and password in Configure Services

1. On the core server, click **Start | LANDesk | Configure Services**.
2. Click the **Intel AMT Configuration** tab.
3. Type **admin** as the username and password under **Current Intel AMT Credentials**.
4. Type a new username (optional) and a strong password under **Provision with new Intel AMT Credentials**.
5. Click **OK**.

These username and password fields must be entered here before you can generate a batch of provisioning IDs.

To generate a batch of Intel AMT provisioning IDs

1. At the core server, click **Hardware configuration** in the left navigation pane.
2. Expand **AMT** and drill down through **Provisioning** to **Generate AMT IDs**.
3. Type the number of IDs to generate (generally the number of devices you plan to provision).
4. If you want to use a different prefix for the PIDs, type it in the **PID prefix** text box. This prefix can only contain uppercase alphabetic characters and numerals in the ASCII character set. You can enter a maximum of 7 characters for a prefix.
5. Type a batch name to identify this group of generated IDs.
6. Check **Show generated AMT IDs** to display the generated IDs in the list. If you don't check this box, the IDs are generated and saved in the database but are not displayed here.
7. Click **Generate IDs**.

USERS GUIDE

8. After the IDs have been generated, click **Print ID list** to open a new window with the list of IDs. (Only the IDs currently shown in the list are displayed in the new window). Use your browser's Print feature to print the list.
9. To view all IDs that have been previously generated, leave the **Batch name** box empty and click **View batch IDs**.
10. To view one batch of generated IDs, type the batch name in the **Batch name** text box and click **View batch IDs**.

You can generate any number of provisioning keys at once. The keys are stored in the database for future reference as you provision new Intel AMT devices. As the devices are provisioned and the provisioning keys are consumed, the **Generate AMT IDs** page will display shading for the IDs that have been consumed, so you can track which IDs have been used.

A PID prefix is added for your convenience in identifying the IDs as PIDs, but you are not required to use a prefix. We recommend using 0-4 characters; you can use a maximum of 7 characters for the prefix.

To identify batches of provisioning keys, specify a batch name. This should be a descriptive name that indicates which devices the IDs apply to. For example, you could generate batches for each organization in your company and name the batches Development, Marketing, Finance, and so forth. If you later want to view the generated IDs, you type the batch name and click **View batch IDs** to see a list with only those IDs.

Strong passwords

Intel AMT requires the use of a strong password to enable secure communications. Passwords should meet these requirements:

- At least 8 characters long
- Includes at least one number character (0-9)
- Includes at least one non-alphanumeric ASCII character (such as !, &, %)
- Contains both upper- and lowercase Latin characters, or non-ASCII characters (UTF+00800 and above)

Errors in the provisioning process

If you enter a PID and PPS that are not paired correctly (i.e., the PPS should be paired with a different PID), you will see an error message in the alert log and provisioning will not continue with that device. You will need to restart the device and re-enter a correct PID/PPS pair in the Intel AMT Configuration Screen.

If, as you type a PID, the Intel AMT Configuration Screen displays an error message, you have mis-typed the PID. A checksum is performed to ensure that the PID is correct.

Discovering Intel AMT 1.0 devices

When you run a device discovery scan, Intel AMT version 1 devices are discovered and added to the Intel AMT folder in the **Unmanaged** devices list. The devices are recognized as Intel AMT

devices if they have been configured with a secure username and password, replacing the defaults set by the manufacturer.

When you add a secure username and password at the Intel AMT Configuration Screen, you can also enter the IP address of the provisioning server and specify port 9982, as is done with Intel AMT 2 devices. However, no PID/PPS pairs are used in provisioning Intel AMT 1 devices. If you specify a provisioning server IP address, the core server acts as a provisioning server and you can manage the device as an agentless device.

Note that Intel AMT version 1 does not use the same level of security as version 2. Intel recommends that devices with version 1 be configured on an isolated, secure network. After configuration is complete they can be moved to a less secure network for management.

Changing the username and password for Intel* AMT devices

A secure username and password are required for provisioning new Intel AMT (version 1) devices. For devices that you will manage with System Manager, the username and password you enter in the Intel AMT Configuration Screen should be the same as the username and password that you entered in the System Manager Configure Services utility. The username and password in the Configure Services utility is saved in the database and applied globally for provisioning Intel AMT devices.

Intel AMT requires the use of a strong password to enable secure communications. Passwords should meet these requirements:

- At least 8 characters long
- Includes at least one number character (0-9)
- Includes at least one non-alphanumeric ASCII character (such as !, &, %)
- Contains both upper- and lowercase Latin characters, or non-ASCII characters (UTF+00800 and above)

After provisioning, you should regularly change usernames and passwords as part of your IT maintenance. You can use a different username/password combination for each Intel AMT device or apply a username/password combination to multiple devices. The new username/password combinations you enter in the Hardware configuration page are stored in the database and used by System Manager to communicate securely with managed Intel AMT devices.

To change the username and password for Intel AMT devices

1. At the core server, click **Hardware configuration** in the left navigation pane.
2. Expand **AMT** and drill down to **Configuration**.
3. In the **All devices** list, select one or more devices for which you want to change the username and password. Click **Target** on the toolbar.
4. In the lower pane, type the new username, then type and confirm the new password.
5. Click **Targeted devices**, then click **Apply**.

For a single device or multiple devices in the same list, you can select the devices and click **Selected devices**, then click **Apply**.

Configuring System Defense policies

Intel AMT* 2.0 includes a System Defense feature, which enforces network security policies on devices with Intel AMT 2.0 functionality. You can select and apply System Defense policies for managed devices using the **Hardware configuration** tool.

When a System Defense policy is applied on an Intel AMT device, the device filters incoming and outgoing network packets according to the defined policies. When network traffic matches the alert conditions defined in a filter, an alert is generated and the device's network access is blocked. The device is then isolated from the network until you complete the remediation steps for that policy.

System Manager contains predefined System Defense policies that you can apply to your Intel AMT devices. Each policy contains a set of filters that define what kind of network traffic is not allowed and what the resulting actions are when traffic meets the criteria of the filter. The process for selecting and applying policies is as follows:

1. Target one or more managed devices
2. Select the System Defense policy to apply; if needed, edit the policy
3. Apply the policy to the targeted devices

When a System Defense policy is active on a managed device, the device monitors all incoming and outgoing network traffic. If a filter's conditions are detected, the following occurs:

1. The managed device sends an ASF alert to the core server and an entry is added to the alert log
2. The core server determines which policy has been violated and shuts down network access on the managed device
3. The device is listed in the System Defense remediation queue (in the **Hardware configuration** tool)
4. To restore network access on the device, the administrator follows the appropriate remediation steps and then removes the device from the remediation queue; this restores the original System Defense policy on the device

This process is described in more detail in the following sections.

Selecting and applying System Defense policies

System Manager contains the following predefined System Defense policies that can be applied to Intel AMT 2.0 devices. Policies are defined with parameters such as port number, packet type, and number of packets within a specific amount of time. When you enable a policy, it is registered with Intel AMT on the devices you have selected. Policies are saved as XML files on the managed device, in the CircuitBreakerConfig folder.

- **BlockFTPSvr:** This policy prevents traffic through an FTP port. When packets are sent or received on FTP port 21, the packets are dropped and network access is suspended.
- **LDCBKillNics:** This policy blocks traffic on all network ports except for the following management ports:

Port description	Number range	Traffic direction	Protocol
LANDesk management	9593-9595	Send/receive	TCP, UDP
Intel AMT management	16992-16993	Send/receive	TCP only
DNS	53	Send/receive	UDP only
DHCP	67-68	Send/receive	UDP only

When the core server shuts down network access on a managed device, it actually applies this policy to the device. Then, when the device is removed from the remediation queue, the original policy is re-applied to the device.

- **LDCBSYNFlood:** This policy detects a SYN flood denial-of-service attack: it allows no more than 10,000 TCP packets with the SYN flag turned on, in one minute. When that number is exceeded, network access is suspended.
- **UDPFloodPolicy:** This policy detects a UDP flood denial-of-service attack: it allows no more than 20,000 UDP packets per minute on ports numbered between 0 and 1023. When that number is exceeded, network access is suspended.
- **RemoveAllPolicy:** Select this to remove all policies, unregistering them with Intel AMT on the selected devices.

To select a System Defense policy

1. Click **Hardware configuration** in the left navigation pane.
2. Click **AMT** and drill down through the tree to **Policies**.
3. In the list of devices, select the devices you want to apply the policy to (use Ctrl+click or Shift+click to select multiple devices).
4. Click **Target** on the toolbar to add the devices to the **Targeted devices** list.
5. In the lower pane, select a policy from the drop-down list.
6. Click **Targeted devices**, then click **Apply**.

Restoring network access to devices in the remediation queue

If a device's network access is suspended because of a System Defense policy, the device is listed in the remediation queue. It remains there until you remove it from the list, which reinstates the active policy on that device. Before you do that, you need to resolve the issue that placed the device in the queue. For example, if FTP traffic was detected, you need to verify that appropriate actions are taken to prevent further FTP traffic on the device.

To remove a device from the remediation queue

1. Click **Hardware configuration** in the left navigation pane.
2. Click **AMT** and drill down through the tree to **Remediation**.
3. Select the devices that can have their original System Defense policy restored and click **Remove**.

Intel* AMT Agent Presence configuration

Intel* AMT 2.0 includes an Agent Presence security tool that can monitor the presence of software agents on managed devices. You can enable Agent Presence monitoring to ensure that management agents on your devices are continually running, and be alerted when an agent stops even when other, software-based, agents can't detect the problem.

System Manager uses the Intel AMT Agent Presence to monitor two agents: the standard management agent and monitoring service. It is useful in situations where normal monitoring communications are not available. For example, a device's communication layer may not be functioning or the monitoring agent itself may have stopped running. By default, Agent Presence also monitors its own monitoring process so you are alerted if it has stopped running.

Agent Presence monitoring is done by configuring a timer that listens for "heartbeat" messages from management agents on the device, to verify that the agents are running. If a timer expires because it has not received a heartbeat message, Intel AMT sends an alert to the core server.

When you set up Agent Presence configuration, the agent on the device registers with Intel AMT to send the heartbeats directly to Intel AMT; if the heartbeats stop, Intel AMT can then alert the core server through out-of-band communication that the device agent is not responding. Intel AMT sends a platform event trap (PET) alert to the core server with a description of the changed state. By default, this alert is logged with device health. You can configure other alert actions to be initiated when this alert is received (for information about configuring alert actions, see [Configuring alert actions](#)).

When you configure Agent Presence monitoring, you can enable or disable monitoring for two agents and set the following values:

- **Heartbeat:** the maximum amount of time (in seconds) that can pass between heartbeat signals. If this time limit is exceeded without a new heartbeat being received, the agent is considered to be not responding. The default value is 120 seconds for the standard management agent and 180 seconds for the monitoring service; the minimum value for both is 30 seconds.
- **Startup time:** the maximum amount of time (in seconds) that can pass after the operating system starts before a heartbeat must be received from the agent. If this time limit is exceeded the agent is considered to be not responding. Agent Presence is configured on Intel AMT when the agent is installed, so this should allow for enough time for the agent to start running and send its first heartbeat. The default value is 360 seconds; the minimum value is 30 seconds.

To edit the Intel AMT Agent Presence configuration

1. Click **Hardware configuration** in the left navigation pane.
2. Expand **AMT** and drill down through the tree to **AP configuration**.
3. To disable Agent Presence monitoring on Intel AMT 2.0 devices, clear the **Enable Agent Presence monitoring** checkbox.
4. To disable monitoring for a specific agent, clear the checkbox next to the agent name. (Even if both these checkboxes are cleared, Agent Presence will continue to monitor its own monitoring process as long as it is enabled).

5. Type a new value in the **Heartbeat** text box to change the maximum allowed time between heartbeats (minimum 30 seconds).
6. Type a new value in the **Startup time** text box to change the maximum allowed time for the agent to send its first heartbeat after the operating system starts on the device (minimum 30 seconds; 120 seconds is recommended).

IPMI support

System Manager includes support for Intelligent Platform Management Interface (IPMI) 1.5 and 2.0. IPMI is a specification developed by Intel,* H-P,* NEC,* and Dell* to define the message and system interface for management-enabled hardware. IPMI contains monitoring and recovery features that let you access many features regardless of whether or not the machine is powered on, or what state the OS may be in. For more details on IPMI, visit Intel's Web site.

IPMI monitoring is handled by the BMC (baseboard management controller). The BMC operates on standby power and autonomously polls system health status. If the BMC detects that any elements are out of range, you can configure the resulting IPMI actions, such as logging the event, generating alerts, or performing automatic recovery actions such as system power-down or reset.

You must have SMBIOS 2.3.1 or higher installed in order for the BMC to be detected on the system. If the BMC is not detected, you may not see some IPMI information in reports, exports, and so forth.

IPMI defines common interfaces to the hardware used to monitor physical health characteristics, such as temperature, voltage, fans, power supplies, and chassis intrusion. In addition to health monitoring, IPMI includes other system management capabilities including automatic alerting, automatic system shutdown and restart, remote restart and power control capabilities, and asset tracking.

The System Manager menu choices vary slightly for an IPMI-enabled device, depending on the state of the operating system.

Management features for IPMI-enabled devices

Monitoring capabilities depend on what has been installed on the device being monitored, as well as the state of the device. Any IPMI-enabled device with a baseboard management controller (BMC) can be monitored by the administrator console in limited ways with no additional management agents after the BMC has been configured. This includes out-of-band management when the device is powered down or the OS isn't functional. Full-featured management is available when the management agent is installed, a BMC is present, the device is powered on, and the OS is functional. The table below compares the functionality available with these different configurations.

	BMC only*	BMC + agent	Agent (no IPMI)
Out-of-band management enabled	X	X	

USERS GUIDE

	BMC only*	BMC + agent	Agent (no IPMI)
In-band management enabled		X	X
Device can be discovered**	X	X	X
Read environment sensors	X	X	Hardware dependent
Power on/off remotely	X	X	X
Read & clear event log	X	X	
Configure alerts	X	X	X
Read OS information		X	X
Graceful shutdown		X	X
Read SMBIOS information (processor, slots, memory)		X	X
IP syncing (OS to BMC)		X	
Watchdog timer		X	
BMC communicates with core server	X	X	
Local System Manager components communicate with core server		X	X
Full range of System Manager management features		X	

*Standard BMC. The mini BMC is a scaled-down version of a baseboard management controller. It has the functionality listed above, with limitations such as the following:

- Does not support serial over LAN (SOL redirection)
- Has only one username for BMC management
- Uses only one channel for communicating with the BMC
- Has a smaller system event log (SEL) repository

**If the BMC is not configured, it will not respond to ASF pings which the product uses to discover IPMI. This means that you will have to discover it as a normal computer. When you deploy a management agent, the server configuration executable will scan the system and detect it is IPMI and configure the BMC.

Conflicts with other IPMI drivers

If you have installed other management software that includes IPMI drivers on devices you want to manage with System Manager, you need to uninstall those products before you can deploy Management Suite agents with IPMI management features.

For example, Microsoft* Windows* Server 2003 includes IPMI support through the installation of Windows Remote Management (WinRM), which includes a Windows Management Instrumentation (WMI) provider and an IPMI driver. However, System Manager does not support the installation of that IPMI driver and installs its own IPMI driver. If WinRM has been installed on a device that you want to manage with System Manager, you must first uninstall WinRM through Windows Add/Remove Programs (**Start | Control Panel | Add or Remove Programs | Add/Remove Windows Components | Management and Monitoring Tools | clear the Hardware Management check box | click OK**).

IPMI BMC configuration

Use the **IPMI BMC configuration** page to customize settings for communicating with [IPMI](#)-enabled devices. The features described below are available for in-band devices; if a device is out of band, only the power configuration and BMC user settings are available.

CAUTION: It is strongly recommended that you do not change IPMI settings unless you are familiar with the [IPMI specification](#) and understand the related technologies involved in these settings. Improper use of these configuration options may prevent System Manager from successfully communicating with IPMI-enabled devices.

The following configuration options are available:

- [Watchdog timer](#)
- [Power configuration](#)
- [User settings](#)
- [BMC password](#)
- [LAN configuration](#)
- [SOL configuration](#)
- [IMM configuration](#)

Changing watchdog timer settings

IPMI provides an interface for the BMC watchdog timer. This timer can be set to expire periodically, and is configured to initiate certain actions if it expires (such as power cycling). System Manager is configured to reset the timer periodically so it does not expire; if the device becomes unavailable (for example, it is powered down or hangs), the timer is not reset and it then expires, which initiates the action.

You can specify how much time to allow before the timer expires and select an action to perform if it does expire. You can choose to take no action, do a hard reset (shut down and restart of the device, power down the device gracefully, or run a power cycle (power down gracefully and then start up again).

You can also set the BMC to stop broadcasting ARP (Address Resolution Protocol) messages while the watchdog timer is enabled, which can reduce the amount of network traffic being generated. If you suspend ARPs, they will automatically resume if the watchdog timer expires.

To change watchdog timer settings

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **Watchdog timer**.
4. Check **Turn on the watchdog timer** to enable the timer.
5. Specify the frequency of checking the timer (number of minutes or seconds).
6. Select an action to initiate when the watchdog timer expires.
7. If you want the BMC to stop broadcasting ARP messages while the watchdog timer is enabled, check **Suspend BMC ARPs**.
8. Click **Apply**.
9. If you have changed the watchdog timer settings, you can revert to the default settings by clicking **Restore defaults**.

Changing power configuration settings

When power is lost on an IPMI-enabled computer, you can specify what action should be taken when power is restored. We recommend that you restore the computer to whatever state it was in at the time power was lost, but you can also choose to keep it powered off or always power up the computer.

To change power configuration settings

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **Power configuration**.
4. Select an option for when power is restored.
5. Click **Apply**.
6. If you have changed the power configuration settings, you can revert to the default settings by clicking **Restore defaults**.

Changing BMC user settings

System Manager authenticates to a BMC with a user name/password combination that is unique to the BMC (separate from any other System Manager user names). System Manager reserves the first user name so it can always communicate with the BMC. If the BMC allows other user names to be defined, you can define user names with passwords for BMC authentication.

You can also specify privilege levels for each user. For advanced IMMs, you can specify protocol privilege levels (telnet, http, and https) for each channel.

CAUTION: Use extreme caution when making changes to these settings. Erroneous settings can disable the device's BMC communication with this product.

To change BMC user settings

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **User settings**.
4. To clear the data for a user name, click the index number and click **Clear**.
5. To add or change a username, click the index number and click **Edit**.
6. Type a user name.
7. To set a password, check the **Set password** checkbox, then type the password and confirm it.
8. Select the privilege levels for LAN and serial access.
9. Click **Save changes**.

Changing the BMC password

System Manager authenticates to a device's BMC using the default user name (user 1 and password). You can't change the user name but you can change its password. When you change this password setting, the change is saved in the database and on the BMC.

To change the default BMC password

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **Password**.
4. Type the new password and then confirm it.
5. Click **Apply**.

Changing LAN configurations

IPMI messages can be carried directly from the BMC over a LAN interface in addition to the device's system interface. Enabling LAN communication allows the core server to receive IPMI-specific alerts even if the device is powered down. The core server maintains this communication as long as the device has a physical network connection with a valid network address, and as long as the device's main power remains connected.

CAUTION: If you choose to set custom configuration for LAN or serial communication to the BMC, use extreme caution when making changes to the settings. Erroneous settings can disable the device's BMC communication with this product.

If you have a LAN channel defined, you can use the default settings for the device's BMC, or you can change the IP address and gateway settings. Use these options to configure destinations for the SNMP traps sent by the BMC for each platform event trap (PET) event.

You can also change SNMP community string settings for sending alerts over LAN. When configuring these settings, you must specify the SNMP community string used for SNMP authentication. For each configuration, you can edit the trap destination information to specify where and how traps are sent, and whether they are acknowledged.

To set properties for LAN channel configuration

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **LAN configuration**.
4. Select **Always available** in the LAN communication drop-down list to keep access to the BMC open. If you select **Disabled** you will not have LAN access to the BMC when the device is out of band.
5. Select the user privilege level for the channel: **Administrator-level** has access to all commands, while **User-level** is limited to read-only access (you will have a restricted feature set if you select User-level).
6. Check **Permanently turn off BMC ARPs** to turn off Address Resolution Protocol messages from the BMC. This reduces network traffic but can prevent communication with the BMC when the device is out of band.
7. Check **Turn off ARP responses** to stop the BMC from sending ARP message responses when the OS is not available. If you enable this setting you might prevent communication with the BMC when the device is out of band.
8. IP settings for the LAN channel are automatically set if the BMC is in sync with the OS channel. If not, the checkbox under the **IP settings** tab is enabled. You can leave the box checked to use DHCP settings provided automatically, or you can uncheck the box and edit the text fields with static settings. It is generally preferable to use automatic settings.
9. Click the **Send alerts over LAN** tab to configure SNMP community string settings (see details below).
10. Click **Apply** to save your changes.

To change Send alerts over LAN properties

1. Open the **LAN configuration** page (steps 1-3 above).
2. Click the **Send alerts over LAN** tab.
3. Check the **Enabled** checkbox to enable sending SNMP alerts.
4. Specify the **SNMP community string** to be used for SNMP authentication.
5. To configure the trap destinations, double-click the index number to open a **Properties** dialog box.
6. Specify the IP address to which the BMC will send alerts, as well as the corresponding MAC address.

7. Specify the number of times to retry, how frequently to retry, and the preferred gateway to use.
8. If you want the alerts to be acknowledged (which increases the amount of network traffic that is generated), check the **Acknowledge alerts** checkbox.
9. Click **OK**.
10. At the LAN configuration page, click **Apply** when all settings are complete.

Changing Serial Over LAN (SOL configurations)

Use SOL (Serial Over LAN) configuration options to customize serial modem settings for special uses, such as redirecting BIOS POST messages to the serial port. If the BMC is required to dial out over a modem connection, specific modem settings such as initialization strings and dial strings must also be specified.

For serial modem operation, you may need to configure the device board's BIOS and jumper settings. See the documentation for the particular device for details.

CAUTION: If you choose to set custom configuration for LAN or serial communication to the BMC, use extreme caution when making changes to the settings. Erroneous settings can disable the device's BMC communication with this product.

To change SOL configuration settings

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **SOL configuration**.
4. Check **Turn on Serial-over-LAN communications** to enable SOL.
5. Select the minimum **User level required to activate SOL**.
6. Select the **Baud rate for SOL sessions** that is appropriate to the device's hardware configuration.
7. Click **Apply**.

Changing IMM configurations

The **IMM configuration** page is displayed only for IPMI devices that are equipped with an advanced IMM add-in card. The options on this page let you enable or disable protocols and features for use with the IMM-enabled device. Consult the manufacturer's documentation for the IMM before you make changes to these settings.

To change IMM configuration settings

1. In the **My devices** view, double-click the device you want to configure.
2. In the left navigation pane of the server information console, click **Hardware configuration**.
3. Expand **IPMI BMC configuration** and click **IMM configuration**.
4. Check the boxes for protocols and features that you want to enable, and add any settings that are required. Available options include:
 - KVM

- SNMP
 - telnet
 - SMTP alerting
 - HTTP
 - HTTPS
5. Click **Apply**.

Managing Dell* DRAC devices

This product includes management integration with devices that have a Dell* DRAC (Remote Access Controller). The DRAC is a remote hardware controller that provides an interface to the IPMI-compliant server management hardware on the Dell device. The DRAC has an IP address assigned to it, which is used to identify the DRAC device in device discovery and in managing the device.

Devices that contain a Dell DRAC can be managed with the same functionality as other IPMI-compliant devices. When the device has been discovered and added to the list of managed devices, it is managed as any other IPMI device. In addition, System Manager also has unique Dell DRAC features.

The OpenManage Server Administrator is a Web-based console provided by Dell for managing the Dell DRAC device. Normally it is accessed by typing the IP address of the DRAC in a browser and logging in with a username and password. When a Dell DRAC device is managed with System Manager you can also open this utility directly from the System Manager interface.

In addition, System Manager lets you manage usernames and passwords for accessing the OpenManage Server Administrator, and displays three logs from this utility in the server information console.

To open the OpenManage Server Administrator for a Dell DRAC device

1. Double-click the device in the **All devices** list.
2. In the server information console, expand **Hardware** and click **Dell DRAC**. The device's IP address and other identifying information is displayed.
3. Click **Launch Dell DRAC utility** to open the device's OpenManage Server Administrator in a new window.

Dell DRAC logs available in System Manager

Three logs from the OpenManage Server Administrator utility are displayed in the System Manager server information console.

- **Dell DRAC log:** tracks all events recorded by the Server Administrator, such as login activity, session status, firmware update status, and interaction between the DRAC and other device components. Information displayed in System Manager includes event severity, description, and suggested corrective actions for errors.
- **Dell DRAC command log:** tracks all commands issued to the Server Administrator. It shows what commands were performed, by whom, and when, including attempts to log in and out and access errors.

- **Dell DRAC trace log:** useful for tracing details about network communication events, such as alerting, paging, or network connections from the DRAC.

To view logs for a Dell DRAC device

1. Double-click the device in the **All devices** list.
2. In the server information console, expand **Logs**.
3. Click **Dell DRAC log**, **Dell DRAC Command log**, or **Dell DRAC trace log**.

Managing usernames for Dell DRAC-enabled devices

To access the OpenManage Server Administrator interface you log in with a username and password that is defined for the device. The default **root** user is the first user in the list and cannot be deleted, but its password can be changed. Up to 15 additional users can be added. While DRAC usernames can have different access levels, System Manager only defines usernames at the Administrator level.

To add or edit usernames and passwords for a DRAC-enabled device

1. Double-click the device in the **All devices** list.
2. In the server information console, click **Hardware configuration**.
3. In the hardware configuration console, expand **Dell DRAC configuration** and click **Dell DRAC users**. A list of currently defined users appears.
4. To change the password for a user, click the user number and click **Change password**. Type and confirm the new password, then click **Apply**. (To assign the same password to multiple users, select them using Ctrl +Click or Shift+click.)
5. To add a user, click **Add user**. Type a username and password and confirm the password, then click **Apply**. The user is added to the list.

Note: if you type a username that is already in the list, the new password you specify will overwrite the existing password for that user name; a second user with that name is not added to the list.

6. To delete a user, click the user number and click **Delete user**, then click **OK**. (To delete multiple users, select them using Ctrl +Click or Shift+click.)

All users in this list have Administrator level access to the OpenManage Server Administrator.

Core database installation and maintenance

Core database installation

The default installation for this product installs a Microsoft MSDE database on your core server. This is the only database option available for System Manager, and only one core database can be installed. The database should be installed only on a standalone server.

The database schema supports Microsoft SQL Server 2000 with SP4. All database servers need to have MDAC 2.8 on them.

The database installed on your core server must be a completely new database. If you are installing System Manager on a server that previously had an installation of LANDesk[®] Management Suite or Server Manager, you cannot use the existing database structure for your installation of System Manager.

The LANDesk Configure Services utility includes an interface that allows you to configure various services. The General tab in this utility displays the current server name, database name, and user name/password required to access the core database. These credentials are used by all services that access the core database. Because System Manager can use only one core database, you will not need to change the server or database names. You can change the credentials, if needed. For more information, see [Appendix C: Configuring services](#).

Microsoft SQL Server 2005 configuration

This product needs the following parameters. These parameters will be set by default if you use a typical install for SQL 2005:

SQL server configuration parameters

- Microsoft SQL 2005 performs self-tuning. You shouldn't need to tune any parameters manually.

Database parameters

- Use the defaults.

Other settings

- Use another user aliased into the database as DBO when creating the database.
- Set up database maintenance.

To install the product so that it uses your SQL 2005 database

1. Install the product to the point where you need to choose a database.
2. In the **Choose a Database** page, click **Add LANdesk tables to a database** and then click **Next**.
3. Enter the Server and Database names, and enter the User and Password that the product should use to authenticate to the database. You **MUST** use a user who is aliased into the database as DBO. Don't use "sa" for the login name. If you're using an Oracle database, check **Use an Oracle database**.
4. Click **Next** and finish the product install.

SQL maintenance

You must regularly perform maintenance on a Microsoft SQL Server database. Over time, the indexes become very inefficient. If your database queries seem to be running more slowly than normal, updating statistics on all tables within the database can substantially improve query performance. On very large databases, you might want to update statistics daily.

Microsoft SQL maintenance requires the SQLServerAgent service to be running on the SQL server. You may need to set the service to Automatic in the Control Panel Services applet. SQL maintenance won't run unless the SQLServerAgent service is started.

To set up a maintenance task

1. Click **Start | Programs | Microsoft SQL Server 2005 | SQL Server Management Studio**.
2. Click the + next to these folders: **Microsoft SQL Servers**, **SQL Server Group**, the name of your server, and **Management**.
3. Right-click **Database Maintenance** and click **New Maintenance Plan**.
4. In the **Database Maintenance Plan** dialog, click **Next**.
5. In the **Select Databases** dialog, select **These databases** and select the checkbox for your database. Click **Next**.
6. In the **Update Data Optimization Information** dialog, click **Reorganize data and index pages**.
7. Set the **Change free space per page percentage to** option to **10**.
8. Click the **Change** button next to the **Schedule** window.
9. In the **Edit Recurring Job Schedule** dialog, select the schedule you want for maintenance. We suggest you perform the maintenance at least weekly at a time when there will be minimal database activity.
10. Click **OK**.
11. In the **Database Integrity Check** dialog, select these options: **Check database integrity** and **Include indexes**, and click **Next**.
12. In the **Specify the Database Backup Plan** dialog, specify your own backup schedule and click **Next**.
13. In the **Specify the Transaction Log Backup Plan** dialog, specify your own backup schedule and click **Next**.
14. In the **Reports to Generate** dialog, select the **Write report to a text file in directory** option and click **Next**.
15. In the **Maintenance Plan History** dialog, select the **Write history to the msdb.dbo.sysdbmaintplan_history table on this server** option.
16. Set the **Limit rows in the table to** option to **1000**.

USERS GUIDE

17. Click **Next**.
18. In the **Completing the Database Maintenance Plan** dialog, enter a **Plan name** and click **Finish**.

Appendix A: System requirements and port usage

The core must have a static IP address.

- [Administrative Core](#)
- [Server Support \(agents\)](#)
- [Browsers](#)
- [Databases](#)
- [Microsoft Data Access Components](#)
- [Port usage](#)

Administrative Core

The administrative core supports the following operating systems:

- Microsoft Windows 2000 Server (with SP4)
- Microsoft Windows 2000 Advanced Server (with SP4)
- Microsoft Windows 2003 Server R2 Standard Edition
- Microsoft Windows 2003 Server R2 Enterprise Edition

Server Support (agents)

- Microsoft Windows 2000 Server (with SP4)
- Microsoft Windows 2000 Advanced Server (with SP4)
- Microsoft Windows 2000 Professional (with SP4)
- Microsoft Windows 2003 Server R2 Standard Edition x86
- Microsoft Windows 2003 Server R2 Standard x64 Edition
- Microsoft Windows 2003 Server R2 Enterprise Edition x86
- Microsoft Windows 2003 Server R2 Enterprise x64 Edition
- Microsoft Windows XP Professional (with SP2)
- Microsoft Windows XP Professional x64 (with SP2)
- Windows Small Business Server 2000 (with SP4)
- Windows Small Business Server 2003 (with SP1)
- Red Hat Enterprise Linux v3 (ES 32-bit - U6)
- Red Hat Enterprise Linux v3 (ES EM64t - U6)
- Red Hat Enterprise Linux v3 WS 32-bit - U6
- Red Hat Enterprise Linux v3 WS EM64t - U6
- Red Hat Enterprise Linux v3 (AS 32-bit - U6)
- Red Hat Enterprise Linux v3 (AS EM64t - U6)
- Red Hat Enterprise Linux v4 (ES 32-bit - U3)
- Red Hat Enterprise Linux v4 (ES EM64t - U3)
- Red Hat Enterprise Linux v4 (AS 32-bit - U3)
- Red Hat Enterprise Linux v4 (AS EM64t - U3)
- Red Hat Enterprise Linux v4 WS 32-bit - U3
- Red Hat Enterprise Linux v4 WS EM64t - U3

USERS GUIDE

- SUSE* Linux Server 9 ES 32-bit SP3
- SUSE Linux Server 9 EM64t SP3
- SUSE Linux Server 10 ES 32-bit
- SUSE Linux Server 10 EM64t
- SUSE Linux Professional 9.3
- SUSE Linux Professional 9.3 EM64t
- HP-UX 11.1
- Unix AIX

Browsers

- Microsoft Internet Explorer 6.x (with SP1)
- Mozilla 1.7 and later
- Firefox 1.5 and later

Databases

- MSDE (with SP4)

Microsoft Data Access Components

- MDAC 2.8 or later

If you want to have more than one LANDesk management product using the same database, you must install both products on the same "core" machine. Accordingly, if you want to install multiple products on the same "core" machine, you must use the same database. If both products use the same database, then both products must also be version 8.70.

Port usage

Introduction

When using this product in an environment that includes firewalls (or routers that filter traffic), you may need to adjust firewall or router configurations to allow the product to operate. This section describes the ports used by the various product components. The information here focuses on information you need to configure routers and firewalls, leaving out ports only used locally (within individual subnets).

Background information on firewall rules

This information applies to setting up firewall rules. If you aren't familiar with the subject, this section provides some generic background information on the main concepts.

Firewall rules

"Opening a port" is not a precise term. You can't just go to a firewall and "open port x." Opening a port is shorthand for setting up a firewall rule. Firewall rules describe what traffic will or will not be allowed through the firewall. Firewall rules don't filter traffic on port number only. Rules can be based on protocols, source and destination port numbers, direction (inbound / outbound), source and destination IP addresses, and other things.

A typical firewall rule looks like this: "allow inbound traffic on TCP port 9535." For using this product, this rule is needed to support remote control. The rule is based on three elements:

1. The protocol (TCP or UDP)
2. The port number
3. The direction (inbound or outbound)

These three elements are required to set up firewall rules.

Source and destination ports, dynamic ports

There are always two ports involved in TCP or UDP communication. Any TCP or UDP packet is from a source port to a destination port. Firewall rules can be based on the source port, the destination port, or both. Ports listed in documents such as this one are always destination ports.

Well-known ports such as 5007 (used by the inventory service) refer to only one side of the communication. The other side of the communication is using a dynamic port. Dynamic ports are assigned automatically by the operating system in the range 1024-5000.

Firewalls and UDP traffic

To allow TCP traffic through a firewall, a single rule is sufficient, such as to allow inbound TCP connections to port 5007. Once the TCP connection is established, data can flow both ways through the connection.

UDP traffic is different because it is connectionless. For example, by default the core server will "ping" devices at UDP port 38293 before starting a task. A firewall rule that allows outgoing UDP packets to port 38293 will allow packets from the core server to a device outside the firewall, but not the device's response packets.

A rule that allows both outgoing and incoming packets to port 38293 won't work either because only one side of the communication is listening on the well-known port. The other side is using a dynamic port. Because the core server's outgoing packets are from a dynamic port to port 38293, the device's response packets are from port 38293 to the same dynamic port, not to port 38293. To allow two-way communication, a rule is needed that allows UDP packets with source port or destination port = 38293. Such a rule is usually acceptable on the intranet, but not on an external firewall (because it would allow inbound packets to all UDP ports).

For this reason, UDP traffic is usually not considered "firewall friendly". Coming back to the example, there is an alternative to UDP port 38293: TCP port 9595. When managing devices across a firewall, you probably want to configure the product to use the TCP port.

Ports used

Port	Direction	Protocol	Service
31770	console to device, device to core	TCP	communication between console and device
9595, 9594	console to device	TCP	Server configuration
9595	console to device	UDP	discovery
623	console to device	UDP	ASF, IPMI discovery
5007	console to device	TCP	inventory
9535	console to device	TCP	remote control
139, 145	console to device	TCP	file and printer sharing
137, 138	console to device	UDP	file and printer sharing

This product needs to discover nodes with the standard management agent installed before it can manage them. UDP port 9595 is used for discovery. You can also manually add individual devices to the console, but this still requires the device to respond to a "ping" on UDP port 9595. Communication between the console and the device uses TCP ports 31770 and 6787. Traffic on the latter port is HTTP-based. UDP port 623 is used for ASF (alert standard forum) discovery. In addition, this product uses TCP port 9535 for remote control. IPMI discovery is a linked with ASF discovery and uses the same port (udp/623).

Appendix B: Activating the core server

Before you can use the console, you must first activate your core server using the Core Server Activation utility. This is normally a one-time procedure that only needs to be repeated if you purchase additional licenses. Use the Core Server Activation utility to:

- Activate a new server for the first time.
- Update an existing core server or upgrade to Management Suite or Server Manager.
- Activate a new server with a 45-day trial-use license.

Start the utility by clicking **Start | All Programs | LANDesk | Core Server Activation**. If your core server doesn't have an Internet connection, see "[Manually activating a core or verifying the node count data](#)" later in this section.

Each core server must have a unique authorized certificate. Multiple core servers can't share the same authorization certificate, though they can verify node counts to the same LANDesk account. This utility runs automatically on the first reboot after installing System Manager.

Periodically, the core server generates node count verification information in the "\Program Files\LANDesk\Authorization Files\LANDesk.usage" file. This file is periodically sent to the LANDesk Software licensing server. This file is in XML format and is digitally signed and encrypted. Any manual changes to this file will invalidate the contents and the next usage report to the LANDesk Software licensing server.

The core communicates with the LANDesk Software licensing server via HTTP. If you use a proxy server, click the utility's **Proxy** tab and enter your proxy information. If your core has an Internet connection, communication with the license server is automatic and won't require you to perform any manual steps. If the core is not connected, click **Close** on reboot, and email the authorization file into licensing@landesk.com.

The Core Server Activation utility won't automatically launch a dial-up Internet connection, but if you manually launch the dial-up connection and run the activation utility, the utility can use the dial-up connection to report usage data.

If your core server doesn't have an Internet connection, you can verify and send the node count manually, as described later in this section.

Activating a server with a LANDesk Software account

Before you can activate a new server with a full-use license, you must have an account set up with LANDesk Software that licenses you for the LANDesk Software products and number of nodes you purchased. You will need the account information (contact name and password) to activate your server. If you don't have this information, contact your LANDesk Software sales representative.

Don't change the core server's date or time between installing the product and activating the core. The activation will fail. You will have to uninstall and reinstall the product.

To activate a server

1. Click **Start | All Programs | LANDesk | Core Server Activation**.
2. Click **Activate**.

Activating a server with a trial-use license

The 45-day trial-use license activates your server with the LANDesk Software licensing server. Once the 45-day evaluation period expires, you won't be able to log in to the core server, and it will stop accepting inventory scans, but you won't lose any existing data in the software or database. During or after the 45-day trial use license, you can rerun the Core Server Activation utility and switch to a full activation that uses a LANDesk Software account. If the trial-use license has expired, switching to a full-use license will reactivate the core.

To activate a 45-day evaluation

1. Click **Start | All Programs | LANDesk | Core Server Activation**.
2. Click **Activate this core for a 45-day evaluation**.
3. Click **Evaluate**.

Updating an existing account

The update option sends usage information to the LANDesk Software licensing server. Usage data is sent automatically if you have an Internet connection, so you normally shouldn't need to use this option to send node count verification. You can also use this option to change the core server associated with the LANDesk Software account. This option can also change a core server from a trial-use license to a full-use license.

To update an existing account

1. Click **Start | All Programs | LANDesk | Core Server Activation**.
2. Click **Update this core server using your LANDesk contact name and password**.
3. Enter the **Contact name** and **Password** you want the core to use. If you enter a name and password that's different than the one used to originally activate the core, this switches the core to the new account.
4. Click **Activate**.

Manually activating a core or verifying the node count data

If the core server doesn't have an Internet connection, the Core Server Activation utility won't be able to send node count data. You'll see a message prompting you to send activation and node count verification data manually through e-mail. E-mail activation is a quick and simple process. When you see the manual activation message on the core, or if you use the Core Server Activation utility and see the manual activation message, follow these steps.

To manually activate a core or verify the node count data

1. When the core prompts you to manually verify the node count data, it creates a data file called ACTIVATE.TXT in the \Program Files\LANDesk\Authorization Files folder. Attach this file to an e-mail message and send it to licensing@landesk.com. The message subject and body don't matter.
2. LANDesk Software will process the message attachment and reply to the mail address you sent the message from. The LANDesk Software message provides instructions and a new attached authorization file.
3. Save the attached authorization file to the \Program Files\LANDesk\Authorization Files folder. The core server immediately processes the file and updates its activation status.

If the manual activation fails or the core can't process the attached activation file, the authorization file you copied is renamed with a .rejected extension and the utility logs an event with more details in the Windows Event Viewer's Application Log.

Appendix C: Configuring services

You can use the Configure Services applet to configure the following services for any of your core servers and databases:

- [Selecting a core server and database](#)
- [Configuring the Inventory service](#)
- [Configuring duplicate device name handling](#)
- [Configuring duplicate device ID handling](#)
- [Configuring the scheduler service](#)
- [Configuring preferred server credentials](#)
- [Configuring the custom jobs service](#)
- [Configuring the multicast service](#)
- [Configuring the BMC password](#)
- [Configuring the Intel AMT password](#)

To launch the Configure Services applet, on the core server, click **Start | Program Files | LANDesk | LANDesk Configure Services**.

Configuring services tabs

Before configuring a service, use the **General** tab to specify the core server and database you want to configure the service for.

Note: Any service configuration changes you make for a core server and database will not take effect until you restart the service on that core server.

Selecting a core server and database

The **General** tab lets you select a core server and database and provide authentication credentials so that you can configure services for that core server.

About the Configure Services dialog: General tab

Use this dialog to select the core server and database you want to configure a specific service for. Then, select the service tab and specify the settings for that service.

- **Server name:** Displays the name of the core server you're currently connected to.
- **Server:** Lets you enter the name of a different core server and its database directory.
- **Database:** Lets you enter the name of the core database.
- **Username:** Identifies a user with authentication credentials to the core database (specified during Setup).
- **Password:** Identifies the user's password required to access the core database (specified during Setup).

- **This is an Oracle database:** Indicates that the core database specified above is an Oracle database. (Does not apply to System Manager).
- **Refresh settings:** Restores the settings that were present when you opened the Service Configuration dialog.

Configuring the Inventory service

Use the **Inventory** tab to configure the Inventory service for the core server and database you selected using the General tab.

About the Configure Services dialog: Inventory tab

Use this tab to specify the following inventory options:

- **Server name:** Displays the name of the core server you're currently connected to.
- **Log statistics:** Keeps a log of core database actions and statistics.
- **Encrypted data transport:** Enables the inventory scanner to send device inventory data from the scanned device back to the core server as encrypted data through SSL.
- **Scan server at:** Specifies the time to scan the core server.
- **Perform maintenance at:** Specifies the time to perform standard core database maintenance.
- **Days to keep inventory scans:** Sets the number of days before the inventory scan record is deleted.
- **Primary owner logins:** Sets the number of times the inventory scanner tracks logins to determine the primary owner of a device. The primary owner is the user who has logged in the most times within this specified number of logins. The default value is 5 and the minimum and maximum values are 1 and 16, respectively. If all of the logins are unique, the last user to log in is considered the primary owner. A device can have only one primary owner associated with it at a time. Primary user login data includes the user's fully qualified name in either ADS, NDS, domain name, or local name format (in that order), as well as the date of the last login.
- **Advanced settings:** Opens the **Advanced settings** dialog where you can set many different advanced settings related to the inventory scanner. To change a setting, click the setting, change the setting in the **Value** text box, then click **Set**. To view a description of a setting, click the setting and view the details in the **Description** box.
- **Software:** Opens the **Software Scan Settings** dialog where you can configure server software scanning time and history settings.
- **Attributes:** Opens the Select attributes to store dialog where you can select the inventory scan attributes that get stored in the database.
- **Manage duplicates: Devices:** Opens the [Configuring duplicate device name handling](#) dialog, where you can choose an option to remove devices with duplicate device names, MAC addresses, or both (see **Duplicate Devices** below).
- **Manage duplicates: Device IDs:** Opens the **Duplicate Device ID** dialog, where you can select attributes that uniquely identify devices. You can use this option to avoid having duplicate device IDs scanned into the core database (see [Configuring duplicate device ID handling](#) below).
- **Inventory service status:** Indicates whether the service is started or stopped on the core server.

- **Start:** Starts the service on the core server.
- **Stop:** Stops the service on the core server.

About the Software Scan Settings dialog

Use this dialog to configure the frequency of software scans. A device's hardware is scanned each time the inventory scanner is run on the device, but the device's software is scanned only at the interval you specify here.

- **Every login:** Scans all of the software installed on the device every time the user logs on.
- **Once every (days):** Scans the device's software only on the specified daily interval, as an automatic scan.
- **Save history (days):** Specifies how long the device's inventory history is saved.

Configuring duplicate device name handling

Use the Duplicate Devices dialog to delete duplicate devices from the database.

1. In the Inventory tab, click **Devices**.
2. In the Duplicate Devices dialog, click the option you want to use when deleting duplicate devices, then click **OK**.

Remove duplicate when:

- **Device names match:** Removes the older record when two or more device names in the database match.
- **MAC addresses match:** Removes the older record when two or more MAC addresses in the database match.
- **Both device names and MAC addresses match:** Removes the older record ONLY when two or more device names and MAC addresses (for the same record match.)

Configuring duplicate device ID handling

Because imaging is often used to configure devices in a network, the possibility of duplicate device IDs among devices is increased. You can avoid this problem by specifying other device attributes that, combined with the device ID, create a unique identifier for your devices. Examples of these other attributes include device name, domain name, BIOS, bus, coprocessor, and so on.

The duplicate ID feature lets you select device attributes that can be used to uniquely identify the server. You specify what these attributes are and how many of them must be missed before the device is designated as a duplicate of another device. If the inventory scanner detects a duplicate device, it writes an event in the applications event log to indicate the device ID of the duplicate device. The Duplicate Device ID dialog contains the following options:

- **Attributes list:** Lists all of the attributes you can choose from to uniquely identify a device.
- **Identity attributes:** Displays the attributes you've selected to uniquely identify a device.
- **Duplicate device ID triggers:**

- **Identity attributes:** Identifies the number of attributes that a device must fail to match before it's designated as a duplicate of another device.
- **Hardware attributes:** Identifies the number of hardware attributes that a device must fail to match before it's designated as a duplicate of another device.
- **Reject duplicate identities:** Causes the inventory scanner to record the device ID of the duplicate device and reject any subsequent attempts to scan that device ID. Then, the inventory scanner generates a new device ID.

To configure duplicate ID handling

1. In the Configure Services dialog, click the **Inventory** tab, then click **Device IDs**.
2. Select attributes from the **Attributes List** that you want to use to uniquely identify a device, and then click the right-arrow button to add the attribute to the **Identity Attributes** list. You can add as many attributes as you like.
3. Select the number of identity attributes (and hardware attributes) that a device must fail to match before it's designated as a duplicate of another device.
4. If you want the inventory scanner to reject duplicate device IDs, check the **Reject duplicate identities** option.

Configuring the scheduler service

Use the **Scheduler** tab to configure the scheduler service for the core server and database you selected using the **General** tab. You must have the appropriate rights to perform these tasks, including full administrator privileges to the managed devices, allowing them to receive package distributions from System Manager. You can specify multiple login credentials to use on devices by clicking **Change login**.

About the Configure Services dialog: Scheduler tab

Use this tab to see the name of the core server and the database that you selected earlier, and to specify the following scheduled task options:

- **Username:** The username under which the scheduled tasks service will be run. This can be changed by clicking the **Change login** button.
- **Number of seconds between retries:** When a scheduled task is configured with multiple retries, this setting controls the number of seconds the Scheduled Tasks will wait before retrying the task.
- **Number of seconds to attempt wake up:** When a scheduled task is configured to use Wake On LAN, this setting controls the number of seconds that the scheduled tasks service will wait for a device to wake up.
- **Interval between query evaluations:** A number that indicates the amount of time between query evaluations, and a unit of measure for the number (minutes, hours, days, or weeks).
- **Wake on LAN settings:** The **IP port** that will be used by the Wake On LAN packet set by the scheduled tasks to wake up devices.
- **Schedule service status:** Indicates whether the service is started or stopped on the core server.
- **Start:** Starts the service on the core server.

- **Stop:** Stops the service on the core server.
- **Restart:** Restarts the service on the core server.
- **Advanced:** Opens the **Advanced scheduler settings** dialog, in which you can modify settings controlling how the scheduler will operate. To change a setting, click it, click **Edit**, change the setting, and click **OK**.

About the Configure Services dialog: Change login dialog

Use the **Change login** dialog (click **Change login** on the **Scheduler** tab) to change the default scheduler login. You can also specify alternate credentials the scheduler service should try when it needs to execute a task on unmanaged devices.

To install System Manager agents on unmanaged devices, the scheduler service needs to be able to connect to devices with an administrative account. The default account the scheduler service uses is LocalSystem. The LocalSystem credentials generally work for devices that aren't in a domain. If devices are in a domain, you must specify a domain administrator account.

If you want to change the scheduler service login credentials, you can specify a different domain-level administrative account to use on devices. If you're managing devices across multiple domains, you can add additional credentials the scheduler service can try. If you want to use an account other than LocalSystem for the scheduler service, or if you want to provide alternate credentials, you must specify a primary scheduler service login that has core server administrative rights. Alternate credentials don't require core server administrative rights, but they must have administrative rights on devices.

The scheduler service will try the default credentials and then use each credential you've specified in the **Alternate credentials** list until it's successful or runs out of credentials to try. Credentials you specify are securely encrypted and stored in the core server's registry.

You can set these options for the default scheduler credentials:

- **Username:** Enter the default domain\username or username you want the scheduler to use.
- **Password:** Enter the password for the credentials you specified.
- **Confirm password:** Retype the password to confirm it.

You can set these options for additional scheduler credentials:

- **Add:** Click to add the username and password you specified to the Alternate Credentials list.
- **Remove:** Click to remove the selected credentials from the list.
- **Modify:** Click to change the selected credentials.

When adding alternate credentials, specify the following:

- **Username:** Enter the username you want the scheduler to use.
- **Domain:** Enter the domain for the username you specified.
- **Password:** Enter the password for the credentials you specified.
- **Confirm password:** Retype the password to confirm it.

Configuring preferred server credentials

There is a **Credentials** button at the bottom of the **Configure LANDesk Software services** dialog. This button launches the **Server credentials** dialog, where you can specify the preferred servers that devices will check for software distribution packages. These preferred servers offload demand on the core server and help you distribute network traffic in low-speed WAN environments where you don't want devices downloading packages from off-site servers. Preferred servers work for every delivery method except for multicast. UNC package shares work with all packages. HTTP package shares only work with MSI and SWD packages.

The **User name and password** dialog (click **Add** in the **Server credentials** dialog) has the following options:

- **Description:** A description for this preferred server. The description appears in the **Server credentials** dialog.
- **Server name:** The name of the server that will host packages.
- **User name:** The user name devices will use to log into the server. This user name should allow only read access for security reasons.
- **Password and Confirm password:** The password for the user name you specified.
- **Limit preferred server usage by these IP address ranges:** If you only want devices within a specified IP range to use the preferred server you're configuring, you can specify the **Starting IP address** and **Ending IP address**, and click the **Add** button to add it to the list.
- **Test credentials:** Click this button to make sure the server name and credentials you entered work correctly.

When controlling preferred server access through IP address ranges, note that devices within the same multicast domain share their configuration files and may use the same servers, even if some of those devices aren't in a particular preferred server's IP address range.

Configuring the custom jobs service

Use the **Custom jobs** tab to configure the custom jobs service for the core server and database you selected using the General tab. Examples of custom jobs include inventory scans or software distributions.

When you disable TCP remote execute as the remote execute protocol, custom jobs uses the Standard management agent protocol by default, whether it's marked disabled or not. Also, if both TCP remote execute and Standard management agent are enabled, custom jobs tries to use TCP remote execute first, and if it's not present, uses Standard product remote execute.

The **Custom jobs** tab also enables you to choose options for server discovery. Before the custom jobs service can process a job, it needs to discover each server's current IP address. This tab allows you to configure how the service contacts servers.

About the Configure Services dialog: Custom jobs tab

Use this tab to set the following custom jobs options:

Remote execute options:

- **Disable TCP execute:** Disables TCP as the remote execute protocol, and thereby uses the CBA protocol by default.
- **Disable CBA execute / file transfer:** Disables the Standard management agent as the remote execute protocol. If Standard management agent is disabled and TCP remote execute protocol is not found on the device, the remote execution will fail.
- **Enable remote execute timeout:** Enables a remote execute timeout and specifies the number of seconds after which the timeout will occur. Remote execute timeouts trigger when the device is sending heartbeats, but the job on the device is hung or in a loop. This setting applies to both protocols (TCP or Standard management agent). This value can be between 300 seconds (5 minutes) and 86400 seconds (1 day).
- **Enable client timeout:** Enables a device timeout and specifies the number of seconds after which the timeout will occur. By default, TCP remote execute sends a heartbeat from device to device in intervals of 45 seconds until the remote execute completes or times out. Client timeouts are triggered when the device doesn't send a heartbeat to the device.
- **Remote execute port (default is 12174):** The port over which the TCP remote execute occurs. If this port is changed, it must also be changed in the client configuration.

Distribution options:

- **Distribute to <nn> servers simultaneously:** The maximum number of devices to which the custom job will be distributed simultaneously.

Discovery options:

- **UDP:** Selecting UDP uses a standard management agent ping via UDP. Most System Manager components depend on the Standard management agent, so your managed devices should have the Standard management agent on them. This is the fastest discovery method and the default. With UDP, you can also select the UDP ping **Retries** and **Timeout**.
- **TCP:** Selecting TCP uses an HTTP connection to the server on port 9595. This discovery method has the benefit of being able to work through a firewall if you open port 9595, but it's subject to HTTP connection timeouts if devices aren't there. These timeouts can take 20 seconds or more. If a lot of target devices don't respond to the TCP connection, your job will take a while before it can start.
- **Both:** Selecting Both has the service attempt discovery with UDP first, then TCP, and lastly DNS/WINS if it's selected.
- **Disable subnet broadcast:** When selected, disables discovery via a subnet broadcast.
- **Disable DNS/WINS lookup:** When selected, disables a name service lookup for each device if the selected TCP/UDP discovery method fails.

Configuring the multicast service

Use the **Multicast** tab to configure the multicast domain representative discovery options for the core server and database you selected using the **General** tab.

About the Configure Services dialog: Multicast tab

Use this tab to set the following multicast options:

- **Use multicast domain representative:** Uses the list of multicast domain representatives stored in the network view's **Configuration > Multicast domain representatives** group.
- **Use cached file:** Queries each multicast domain to find out who might already have the file cached. The cached file can then be used instead of downloading the file to a representative.
- **Use cached file before preferred domain representative:** Changes the order of discovery to make **Use cached file** the first option attempted.
- **Use broadcast:** Sends a subnet-directed broadcast to find any device in that subnet that could be a multicast domain representative.
- **Log discard period (days):** Specifies the number of days that entries in the log will be retained before being deleted.

Configuring the BMC password

Use the **BMC password** tab to create a password for the IPMI Baseboard Management Controller (BMC).

- In the **BMC password** tab, type a password in the **Password** text box, retype the password in the **Confirm password** text box, then click **OK**.

The password cannot be longer than 15 characters, each of which must be numbers 0-9 or upper/lower case letters a-z.

Configuring Intel AMT options

Use the **Intel AMT Configuration** tab to create or change the password on Intel Active Management Technology-enabled devices, and to view instructions for discovering AMT devices.

To configure an Intel AMT password

1. Type the current user name and password. These must match the user name and password as configured in the Intel AMT Configuration Screen (which is accessed in the computer BIOS settings).
2. To change the user name and password, complete the **New Intel AMT Password** section.
3. Click **OK**. This change will be made when the client configuration is run.

Note: The new password must be a strong password, meaning that the password

USERS GUIDE

- Is at least seven characters long
- Contains letters, numbers and symbols
- Has at least one symbol character in the second through sixth positions
- Is significantly different from prior passwords
- Doesn't contain names or user names
- Isn't a common word or name

Discovering and provisioning Intel AMT devices

To discover AMT devices, enter the Core server's IP address in the Provisioning Server field of the AMT BIOS, and use port 9982. Press **Help** in **Configure services** for more information. When an Intel AMT device is discovered and moved to the **My devices** list, it is automatically provisioned using the TLS mode.

Providing validation for OSD imaging environments

In order to create Windows and DOS Use the OSD validation

Appendix D: Agent security and trusted certificates

Each core server has a unique certificate and private key that Setup creates when you first install the core server on a device. Devices will only communicate with core servers for which they have a matching trusted certificate file.

These are the private key and certificate files that are installed:

- **<keyname>.key:** The .KEY file is the private key for the core server, and it only resides on the core server. If this key is compromised, the core server and server communications won't be secure. Keep this key secure. For example, don't use e-mail to move it around.
- **<keyname>.crt:** The .CRT file contains the public key for the core server. The .CRT file is a viewer-friendly version of the public key that you can view to see more information about the key.
- **<hash>.0:** The .0 file is a trusted certificate file and has content identical to the .CRT file. However, it's named so that the computer can quickly find the certificate file in a directory that contains many different certificates. The name is a hash (checksum) of the certificate's subject information. To determine the hash filename for a particular certificate, view the <keyname>.CRT file. There is a .INI file section [LDMS] in the file. The hash=value pair indicates the <hash> value.

All keys are stored on the core server in \Program Files\LANDesk\Shared Files\Keys. The <hash>.0 public key is also in the LDLOGON directory and needs to be there by default. <Keyname> is the certificate name you provided during core server setup. During setup, it's helpful to provide a descriptive key name, such as the core server's name (or even its fully qualified name) as the key name (example: Idcore or Idcore.org.com). This will make it easier to identify the certificate/private key files in a multi-core environment.

Backing up and restoring certificate/private key files among core servers

When you install a core server, Setup creates a new certificate. If you reinstall over an existing core server, Setup still creates a new certificate. If you install devices with a certificate that doesn't match your new core server certificate, the core server won't be able to communicate with them. If you need to reinstall your core server, you have two options:

1. Reinstall the agents manually with a configuration built on your new core server. You can't use software distribution to update the agents, because the core server and devices don't have a matching certificate and key.
2. Before reinstalling a core server, back up the existing certificate and key files to a safe place. After you have reinstalled, copy the old keys to the new core installation. The new and old keys can coexist. The core will use the appropriate key automatically.

Cores can contain multiple certificate/private key files. As long as a client can authenticate with one of the keys on a core, it can communicate with that core.

USERS GUIDE

A utility is included in this product that performs the second option listed above. The core data migration utility (CoreDataMigration.exe) is installed in the \ProgramFiles\LANDesk\ManagementSuite folder. It handles the backing up and copying of data such as keys and certificates when you install a new core.

To save and restore a certificate/private key set

1. At the source core server, go to the \Program Files\LANDesk\Shared Files\Keys folder.
2. Copy the source server's <keyname>.key, <keyname>.crt, and <hash>.0 files to a floppy disk or other secure place.
3. At the destination core server, copy the files from the source core server to the same folder (\Program Files\LANDesk\Shared Files\Keys). The keys take effect immediately.

Warning: Keep the private key file secure

Make sure that the private key <keyname>.key isn't compromised. Don't transfer it using an insecure method, like e-mail or a public file share. The core server uses this file to authenticate devices, and any core with the appropriate <keyname>.key file can perform remote executions and file transfer to a managed device.

Troubleshooting tips

The following troubleshooting tips are for issues that most frequently occur with the console.

I can't activate the core.

If you installed a core, then changed the device time, you will not be able to activate. You must reinstall the product in order to activate the core.

When I tried to activate the core I got a failure message that the core server database could not be read.

Check to make sure the core server is physically connected to the network and has a valid internet connection. If a cable is unplugged or the core server's internet connection is not valid, the activation process can't be completed.

I don't know the URL to the console pages.

Contact the person who installed the core server, most likely the network administrator for your site. However, typically the URL for Server Manager and System Manager is `http://core server machine name/ldsm`. The URL for Management Suite is `http://core server machine name/remote`.

Who am I logged in as?

Look above the bar below the name LANDeskSystem Manager, at the **Connected As** section.

What machine am I logged in to?

Look above the bar below the name LANDeskSystem Manager, at the **Connected To** section.

I launch System Manager and I get a "Session Timed Out" message immediately.

If you open System Manager from the Favorites or Bookmark menu with the `/frameset.aspx` extension at the end of the URL, it will not launch correctly. To fix this, edit your Bookmarks or Favorites link to remove this extension, or paste the URL (without the extension) directly into the browser window.

I don't see some of the left navigation pane links.

This occurs because your network administrator is using LANDeskSystem Manager's role-based administration or feature-level security option that limits you to performing certain tasks that you have the rights to do.

The scanner can't connect to the device.

If the scanner can't connect to the device, verify that the Web application directory is configured correctly. If you're using https, you must have a valid certificate. Verify that you have a valid certificate.

I get a "permission denied" error when I try to access the console.

To use feature-level security on Windows 2000 and 2003, you must disable anonymous authentication. Verify the authentication settings on the Web site and the `..\LANDesk\ldsm` folder under the Web site.

1. On the server that hosts the Web console, click **Start | Administrative Tools | Internet Information Services (IIS Manager.)**
2. From the **Default Web Site** shortcut menu, click **Properties.**

3. On the **Directory Security** tab, in the **Anonymous access and authentication control** area, click **Edit**. Clear the **Enable anonymous access** option and check **Integrated Windows authentication** option.
4. Click **OK** to exit the dialogs.
5. From the Default Web Site's .\LANDesk\ldsm subfolder, click **Properties**. Repeat steps 3-4.

I get an invalid session when viewing the console.

It's possible the browser session has timed out. Use your browser's **Refresh** button to start a new session.

An ASP.NET error appears when I try to start the Web console.

If you see an ASP.NET error message when attempting to log in to the Web console, ASP and the ASP directory permissions may not be configured correctly. Reset the ASP.NET configuration by running the following command:

```
ASPNET_REGIIS.EXE -i
```

The number of items per page is different from the number I specified.

When you specify how many items to display per page, that setting is stored in the Web browser's cookies directory and expires when the console session times out.

The console times out too frequently.

You can change the default session timeout for the console's Web pages. The IIS default is 20 minutes of inactivity before a login expires. To change the IIS session timeout:

1. On the Web server, open the IIS Internet Service Manager.
2. Expand the default Web site.
3. Right-click the **LDSM** folder, then click **Properties**.
4. Under the **Virtual Directory** tab, click **Configuration**.
5. Click the **Application Options** tab, then change the session timeout to the value you want.

Note:LANDeskSystem Manager8.70 is a session-based product. Do not disable the session state.

Report charts don't display properly.

In order to view the interactive bar and pie charts displayed in many reports, you must have the Macromedia Flash Player* installed. Verify that Flash is installed, then run the report again.

Why am I seeing two instances of the same device in my database?

Have you deleted a device from the core database and reinstalled it using UninstallWinClient.exe?

UninstallWinClient.exe is in the LDMain share, which is the main ManagementSuite program folder. Only administrators have access to this share. This program uninstalls LANDesk agents on any device it runs on. You can move it to any folder you want or add it to a login script. It's a Windows application that runs silently without displaying an interface. You may see two instances of the device in the database you just deleted. One of these instances would contain historical data only, while the other would contain data going forward. See the *Deployment Guide* for more information on UninstallWinClient.exe.

When I try to discover an IPMI device, it is not listed in the IPMI folder in the Unmanaged devices page.

IPMI devices must have a BMC (baseboard management controller) that is configured in order to be discovered as IPMI devices and to use full IPMI functionality. If the BMC is not configured, the device can be discovered as a computer. You can then add the device to the list of managed devices and run the Configure Services utility to configure the BMC password. The device's IPMI functionality will then be recognized by this product.

I added a S.M.A.R.T. drive on a server, but I don't see S.M.A.R.T. drive monitoring in the inventory list for that server.

Hardware monitoring is dependent on the capabilities of the hardware installed on a device, as well as on the correct configuration of the hardware. If a hard drive with S.M.A.R.T. monitoring capabilities is installed on a device but S.M.A.R.T. detection is not enabled in the device's BIOS settings, or if the device's BIOS does not support S.M.A.R.T. drives, monitoring data will not be available, and resulting alerts will not be generated.

A USB disk device is not listed in Inventory list until inventory scan has been run.

When a disk device is connected with a USB cable to a managed device, it is not immediately listed under Hard drives in the device's inventory. It is listed under Logical drives after being connected to the device. However, it will not appear under Hard drives until an inventory scan has been run on the device.

On managed Linux devices, a USB disk device must be mounted in order to be listed in inventory. If it is mounted but an inventory scan has not been run, it will appear under Logical drives; after the inventory scan it will also be listed under Hard drives. When the device is disconnected it should be dismounted from the system. On some Linux systems running an older kernel, the device may also stay in the inventory list even after it has been disconnected and dismounted. In this case the managed device needs to be rebooted before the device will be removed from the inventory list.

When I view the Web console help index, it's blank.

The Web console's HTML online help has a full-text search feature that relies on the Windows Indexing Service. Normally, this is enabled by default. If you need to enable indexing on your Web server, do the following:

1. Click **Start | Programs | Administrative Tools | Services**.
2. Double-click **Indexing Service** and click **Start**.
3. Click **OK** to exit out of the dialogs.

You may have to wait a while (up to several hours) for the indexing service to index your server.