CA Nimsoft Monitor

Discovery User Guide

v7.0

September 2013

Legal Notices

Warranty

The material contained in this document is provided "as is," and is subject to being changed, without notice, in future editions. Further, to the maximum extent permitted by applicable law, Nimsoft LLC disclaims all warranties, either express or implied, with regard to this manual and any information contained herein, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. Nimsoft LLC shall not be liable for errors or for incidental or consequential damages in connection with the furnishing, use, or performance of this document or of any information contained herein. Should Nimsoft LLC and the user have a separate written agreement with warranty terms covering the material in this document that conflict with these terms, the warranty terms in the separate agreement shall control.

Technology Licenses

The hardware and/or software described in this document are furnished under a license and may be used or copied only in accordance with the terms of such license.

No part of this manual may be reproduced in any form or by any means (including electronic storage and retrieval or translation into a foreign language) without prior agreement and written consent from Nimsoft LLC as governed by United States and international copyright laws.

Restricted Rights Legend

If software is for use in the performance of a U.S. Government prime contract or subcontract, Software is delivered and licensed as "Commercial computer software" as defined in DFAR 252.227-7014 (June 1995), or as a "commercial item" as defined in FAR 2.101(a) or as "Restricted computer software" as defined in FAR 52.227-19 (June 1987) or any equivalent agency regulation or contract clause. Use, duplication or disclosure of Software is subject to Nimsoft LLC's standard commercial license terms, and non-DOD Departments and Agencies of the U.S. Government will receive no greater than Restricted Rights as defined in FAR 52.227-19(c)(1-2) (June 1987). U.S. Government users will receive no greater than Limited Rights as defined in FAR 52.227-14 (June 1987) or DFAR 252.227-7015 (b)(2) (November 1995), as applicable in any technical data.

Trademarks

Nimsoft is a trademark of CA.

Adobe®, Acrobat®, Acrobat Reader®, and Acrobat Exchange® are registered trademarks of Adobe Systems Incorporated.

Intel® and Pentium® are U.S. registered trademarks of Intel Corporation.

Java(TM) is a U.S. trademark of Sun Microsystems, Inc.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Netscape(TM) is a U.S. trademark of Netscape Communications Corporation.

Oracle® is a U.S. registered trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of the Open Group.

ITIL® is a Registered Trade Mark of the Office of Government Commerce in the United Kingdom and other countries.

All other trademarks, trade names, service marks and logos referenced herein belong to their respective companies.

For information on licensed and public domain software, see the *Nimsoft Monitor Third-Party Licenses and Terms of Use* document at: http://docs.nimsoft.com/prodhelp/en_US/Library/index.htm?toc.htm?1981724.html.

Contact CA Nimsoft

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

Send comments or questions about CA Technologies Nimsoft product documentation to nimsoft.techpubs@ca.com.

To provide feedback about general CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Document Revision History

Document Version	Date	Changes
7.0	Sept 2013	Revised for NMS 7.0 ■ Addition of device correlation. ■ New probe_discovery queue. ■ Changes to the Discovery Wizard. ■ Elimination of discovery probe user interfaces. ■ Content improvement.
6.5	April 2013	First edition of the guide, covering Nimsoft Discovery as implemented in NMS v6.5.

Contents

Chapter 1: Discovery Architecture 7

Discovery Components 8
Discovery Considerations 9
    Prerequisites and Supported Platforms 10

Chapter 2: Configuring Discovery 11

Deploy NMS 12
Configure probe_discovery Queues 13
Launch the Discovery Wizard 16
    Create Authentication Profiles 17
    Define Ranges 21
    Schedule Discovery 22
Run File-based Import 23
View Discovered Systems 24

Appendix A: Advanced Configuration 27

Running discovery_server on a Robot Other Than the Primary Hub 27
Setting Maximum Java Heap Size 28
    Discovery Server 28
    Discovery Agent 29
File-based Import Reference 30
    XML File Schema 31

Chapter 1: Discovery Architecture

Finding and recording all addressable devices and computers within a managed IT environment is the job of automated *discovery*. When multiple discovery records correspond to a single device, this is recognized by device *correlation*. The device record list can be augmented by XML *file-based* device import, or with manual device entry.

When CA Nimsoft Monitor is first installed, the Discovery Wizard starts and you are prompted to configure and run automated discovery. This process can be initiated any number of times afterwards, or scheduled to run on regular intervals.

The Discovery Wizard allows you to set authentication credentials and define IP address ranges to scan. Discovery finds virtually all connected resources on the network and provides detailed information on device type, configuration, and asset/inventory data. By using ICMP, ARP, DNS, SNMP (v1, v2, and v3), WMI, SSH, and NetBIOS, discovery finds a wide range of devices and device information.

This diagram illustrates the flow of data among the key components of automated discovery:

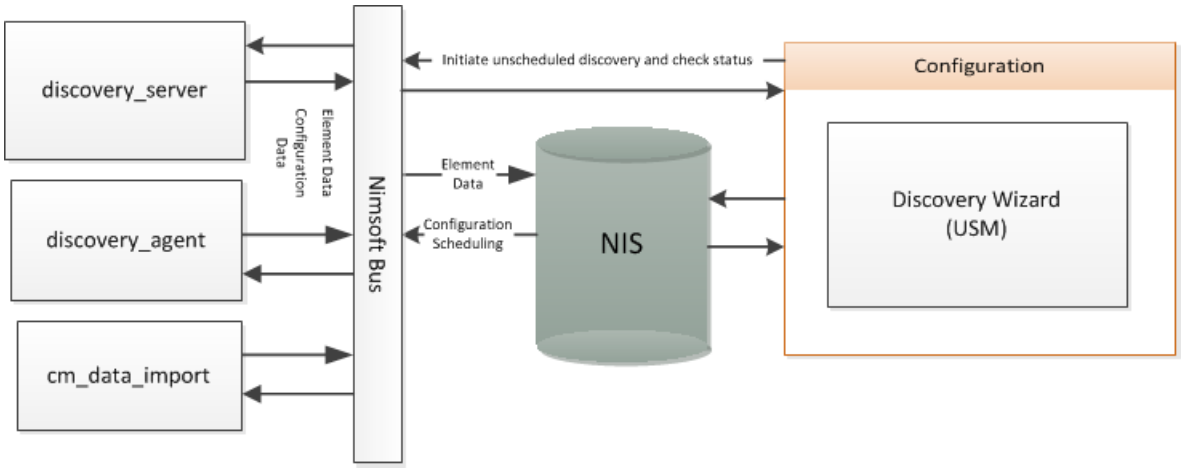

This section contains the following topics:

Discovery Components (see page 8)
Discovery Considerations (see page 9)

Discovery Components

All discovery components are included in a basic installation of CA Nimsoft Monitor Server.

Discovery Wizard

The Discovery Wizard is where you configure discovery scans. Launch the wizard from the **Actions** pop-up menu within USM (Unified Service Manager). USM is the GUI that Nimsoft Monitor opens first after initial installation.

You can run the Discovery Wizard from any discovery agent node in the Discovery tree of USM. In the wizard, where you can set up either of the following discovery methods:

- *Automated discovery* -- populates the device inventory by scanning the network according to authentication profiles and scanning scopes that you configure.
- *File-based import* -- imports a list of one or more hosts or network devices from an XML file into the device inventory.

These topics are covered fully in the section on the Discovery Wizard (see page 16).

Discovery is implemented through these NMS probe components (all deployed on the primary Hub in an initial installation):

Discovery Server

Discovery requires one discovery_server probe that, in typical NMS configurations, runs on the primary Hub. The discovery_server probe performs these major tasks:

- Configures discovery agents and collects status from them.
- Collects information about the Nimsoft infrastructure: hubs, robots, probes, packages, monitored systems or devices, monitored subsystems or items and monitored metrics.
- Collects element data from probes that publish discovery information.
- Applies correlation rules to associate new device records, where appropriate, with any already-existing master device records. One example is to represent multi-homed devices (devices with multiple network interfaces) accurately.

The information that is collected by the discovery_server probe is saved into the NiS database and used by other components in the Nimsoft Monitor solution. The discovery_server probe also helps maintain the NiS database by expiring inactive systems.

Note: Even without any discovery_agent probes deployed, the discovery_server probe is still needed to generate the data required by other components in Nimsoft Monitor.

Discovery Agent

The agent performs IT network scans, pinging and querying devices according to subnet masks/ranges, credential profiles, and selected profiles. These scanning parameters are configured within the Discovery Wizard that runs within the USM portlet in UMP. More information is available in the section on the Discovery Wizard (see page 16).

CM Data Import

This probe processes an XML file that lists information describing hosts and devices, and adds this information to the device inventory. This probe is usually co-located with the discovery server. When you run file-based import from Discovery Wizard, CM Data Import carries out the work. More information about file-based import is in the section on File-Based Import.

Additional components that play a role in discovery:

probe_discovery queue

This queue on the primary hub collects discovery data that is processed by the discovery server. On secondary hubs, queues configured to handle the probe_discovery subject are used to collect and relay up to the primary hub. For instructions on setting up this queue, see the section on Configure the discovery_probe Queue (see page 13).

Nimsoft Information Store (NIS)

The NIS is the database that holds all persistent data in Nimsoft Monitor, including discovery data.

Discovery Considerations

- Discovery works together with the Topology and Root Cause Analysis features in Nimsoft Monitor. Using the data that discovery collects, Topology and Root Cause Analysis can deduce the structure of the network and model it. The model is viewable in the Relationship Viewer portlet in the Unified Management Portal (UMP). More information about the topics of topology and root cause analysis is available in the Topology and Root Cause Analysis User Guide.
- Devices that are imported into Nimsoft via file-based import are not reflected in Topology or in Root Cause Analysis. Topology depends on SNMP information gathered by the discovery agent about the devices.

Prerequisites and Supported Platforms

- Discovery v7.0 requires NMS v7.0.
- Discovery Server v7.0 only works with v7.0 Discovery Agents. The Discovery Server raises an alarm for any pre-v7.0 Discovery Agent it finds.
- Discovery Server v7.0 does not collect any discovery results from any pre-v7.0 discovery agent.

For supported NMS system platforms, see the Nimsoft Monitor Compatibility Support Matrix for details.

Chapter 2: Configuring Discovery

Here is a high-level overview of the entire process.

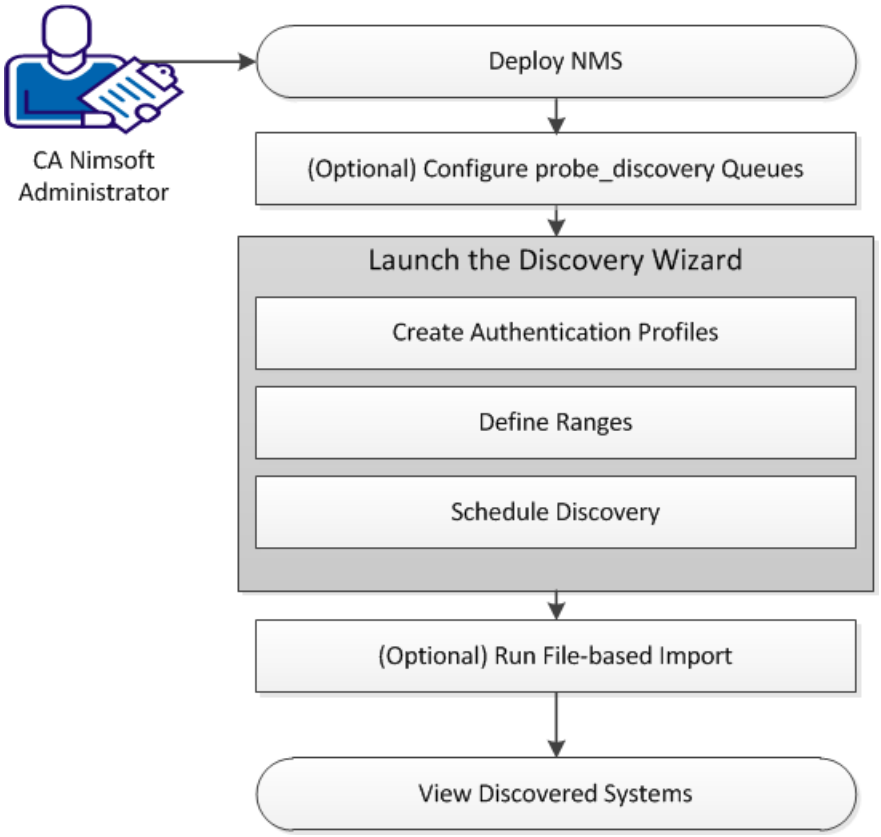

The basic steps are as follows:

1. Deploy NMS (see page 12). The components required for discovery are put in place with a basic install of Nimsoft Monitor Server.
2. (Optional) Configure probe_discovery queues (see page 13). When discovery components are distributed over two or more hubs, queues for probe_discovery messages must be configured between the hubs.

3. Launch the Discovery Wizard (see page 16). The Discovery Wizard in USM leads you through the process of configuring discovery, and allows you to schedule discovery scans to run now or in the future. Note that with a new or upgraded Nimsoft Monitor installation, Discovery Wizard launches automatically the first time UMP is opened.
    1. Create authentication profiles (see page 17)
    2. Define ranges (see page 21) (sets or ranges of IP addresses and IP masks that define and bound the scope of discovery)
    3. Schedule discovery (see page 22).
4. (Optional) Run file-based import (see page 23). To augment automated discovery, you can prepare an XML file with device information and import this information into the device inventory.
5. View discovered systems (see page 24). View computers and devices that have been discovered on your network.

Deploy NMS

The components (probes) required for discovery are deployed on the primary hub with a basic install of CA Nimsoft Monitor:

- Discovery Server
- Discovery Agent
- CM Data Import.

Keep the following in mind if you wish to modify the default discovery probe deployment:

- For minimal discovery, only the discovery_server probe is required. No network scanning is performed
- To add network scanning, deploy the discovery_agent probe on the NMS primary hub or elsewhere
- For optimal discovery in larger environments, more than one discovery agent can be deployed. Some users, particularly service providers and those with very large networks, find it useful to deploy multiple discovery agents in various locations.

  Discovery of a large network can be divided across administrative boundaries for these reasons:

  - To provide different users with access to different parts of the network
  - In situations where there is no direct connectivity to devices at a remote site because of firewall constraints or network-address translation (NAT). For efficient discovery, deploy discovery agents such that each one discovers an exclusive part of the network.

- Note that the WMI protocol is only supported for discovery_agent probes running on Windows systems.

Tip: Discovery Agent requires read-only SNMP access to network devices. To simplify discovery configuration, consider setting up as many network devices as possible to use a "universal" read-only community string (SNMP v3 recommended over v1 or v2c). For example, you could define read-only (get-only) credentials to be "**nms_get_only**". Set up every device possible to allow read-only SNMP access via those credentials. This minimizes the number of SNMP authentication credentials that must be attempted on network nodes, and vastly simplifies your discovery configuration.

Configure probe_discovery Queues

When probes that participate in the discovery process are deployed on a single hub, communication of discovery data is automatically configured and no additional action on the part of the end user is required.

When discovery probes are deployed under hubs *other* than the hub that hosts the discovery_server probe, you must ensure that discovery data can flow from those hubs up to the primary hub (NM Server).

Discovery data between hubs and their probes is carried by queues handling the *probe_discovery* subject. Queues for probe_discovery must be configured on all secondary hubs that host probes which collect discovery data:

- discovery_agent v7.0 (in large environments multiple instances of this probe can be deployed one per network segment/subnet).
- vmware (v5.10 and later).
- cm_data_import (typically this probe is deployed with discovery_server on the primary hub).

Queues for probe_discovery, for secondary hubs, consist of both *attach* queue and *get* queue components.

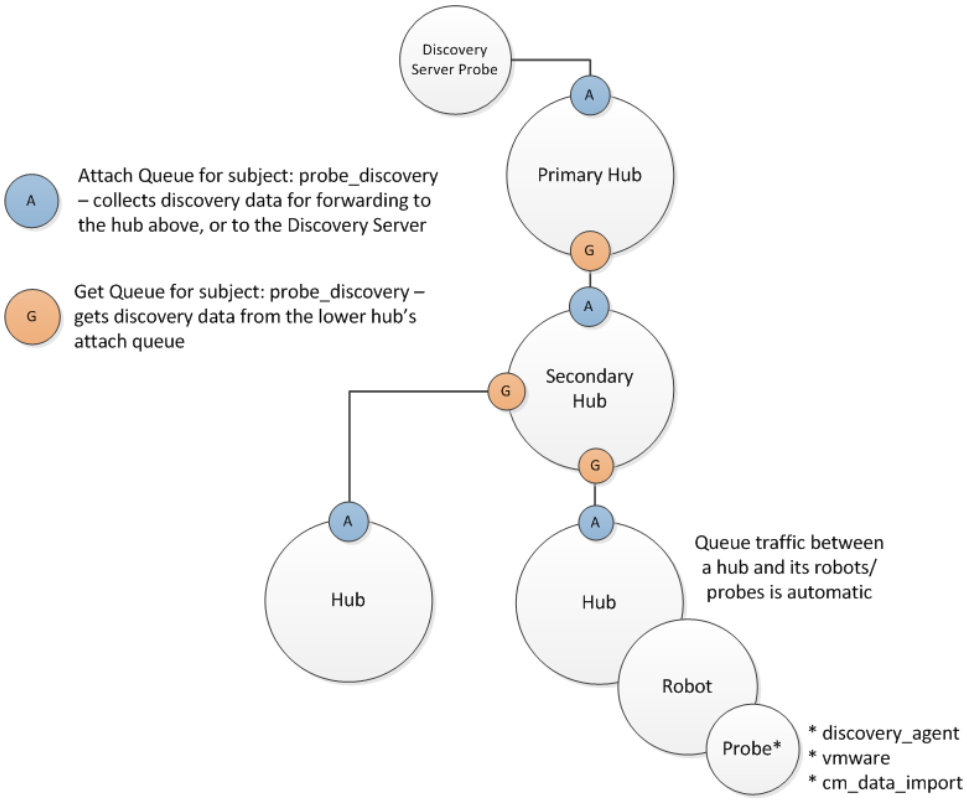

To set up and configure probe_discovery queues, follow these steps:

1. Identify hubs which host any of the following probes
   - discovery_agent
   - vmware (v5.10 and later)
   - cm_data_import.

   You will set up *attach* and *get* queues on the identified hubs.
2. Launch the hub configuration GUI on the primary hub. This can be done in Infrastructure Manager by double-clicking on the hub node. In Admin Console, open the the web-based probe configuration GUI by clicking on the icon next to the hub, then selecting **Configure** from the pop-up menu.
3. Navigate to the **Queues** tab.

4. Set up the *attach* queue.

   a. Click the **New** button.

   b. In the **New Queue** dialog, ensure that **Type** is set to **attach**, and that **Active** is checked

   c. Choose **probe_discovery** under the **Subject** drop-down menu.

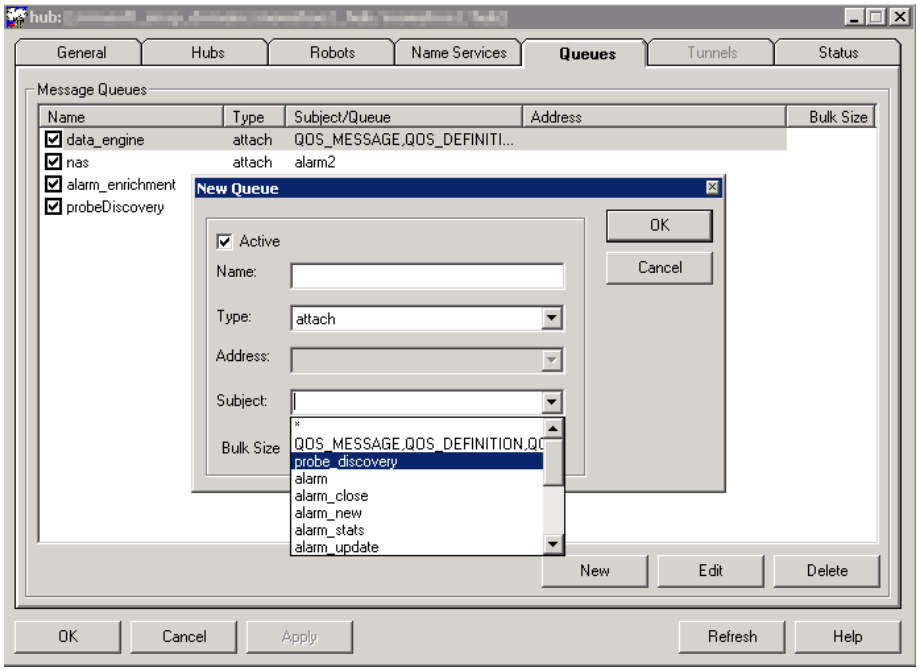

   d. Click **OK**.

5. Next, set up the *get* queue.

   a. Click the **New** button.

   b. In the **New Queue** dialog, ensure that **Type** is set to **get**, and that **Active** is checked

   c. Choose **probe_discovery** under the **Subject** drop-down menu.

   d. Click **OK**.

6. Repeat steps 2 through 5 on the secondary hub or hubs you identified in the first step.

   Note: It is not necessary to set up the *get* queue on the secondary hub(s) lowest in the NMS hierarchy--only an *attach* queue is required to make discovery data available to hubs above.

Confirm the queues are operational by running an automated discovery scan. Scans are initiated using the Discovery Wizard in USM. Check the list of discovered devices that is returned. In addition to local devices, it should contain devices that are only addressable from the secondary hubs in your infrastructure.

Note: Setting up other queues for alarms, QoS, and baseline data is a similar procedure of configuring attach and get queues. The subject of the queue changes as required by the type of data to be carried.

Tip: In small to medium NMS deployments, a wildcard (*) subject, which carries any message, can simplify queue configuration. Use of a wildcard subject in large NMS installations is not recommended.

Launch the Discovery Wizard

The first time you open the Unified Management Portal (UMP) it opens to the Unified Service Manager portlet and the Discovery Wizard is automatically launched.

After the first time you open UMP, you can launch the Discovery Wizard manually if you want to run discovery or change your discovery settings. You can launch the Discovery Wizard from the Inventory node or from the **Actions** menu.

Note: The Discovery Wizard will not run after an update of CA Nimsoft Monitor if there are existing ranges that define *excluded* IP addresses. You must either choose to accept the system prompt to delete excluded ranges, or remove them manually from the database before discovery will run.

Follow these steps:

1. Hover the cursor over or click the name of a discovery agent or range in the tree.

   Discovery agents are indicated by the magnifying glass icon (), and ranges are indicated by the network icon ().

2. Click the gear icon () to the right of the discovery agent or range name in the tree, or choose **Discovery Wizard** from the **Actions** menu.

   Note: The **Discovery Wizard** menu option is enabled only when you click on a discovery agent or range in the tree.

Create Authentication Profiles

The WMI, Linux/Unix, and SNMP tabs allow you to create, edit, view, and delete authentication profiles for discovery. An authentication profile contains credential information necessary for discovery to access and gather information about computer systems and devices in your network.

You can create one or more authentication profiles under each of the WMI, Linux/Unix, and SNMP tabs.

Note: Creating authentication profiles is not required for discovery. However, only IP discovery is used if no authentication profiles exist, and information about discovered systems may be limited.

Select the WMI, Linux/Unix, or SNMP tab, and click the name of an authentication profile in the left pane to view its properties in the pane to the right.

To modify an existing authentication profile, select it and edit the fields as necessary, then click **Save**. To delete an authentication profile, click the trash can icon () next to the name of the profile in the left pane, and click **Save**.

Follow these steps to create an authentication profile:

1. Click **New credentials** in the left pane.
2. Enter information in all of the required fields.

   Required fields are outlined in red.
3. Click **Next**.

   The information you enter is saved when you click **Next** and move through the Discovery Wizard.

Linux/Unix

Linux/Unix authentication profiles use SSH or Telnet to access and discover Linux and Unix systems.

Description

Name for the authentication profile.

ID

This read-only field is the Nimsoft system ID for this authentication profile, assigned when the profile is saved. It identifies the profile uniquely for re-use in other areas of USM that reference authentication profiles.

User

User name.

Password

The user password. Check the **Show new passwords** check box to verify the text as you enter it.

SSH or Telnet

Select the communication protocol to use, SSH (Secure Shell) or Telnet (no secure authentication or encryption).

Note: The Discovery Agent uses *password* authentication to connect to a target device over SSH. The Discovery Agent cannot communicate with a device where SSH is configured for other authentication methods, such as keyboard-interactive. Discovery Agent does not support public key authentication.

SNMP

SNMP (Simple Network Management Protocol) is a widely-used standard protocol for managing devices on IP networks.

CA Nimsoft Discovery supports SNMP versions 1, 2c, and 3. SNMP v3 adds security features which v1 and v2c lack. As a result, authentication profile configuration fields in the Discovery Wizard that deal with security and privacy (encryption) are only active when **3** is selected in the **Version** pull-down menu.

We recommend the following best practices:

- Create a minimal set of SNMP Authentication Profiles that will, in aggregate, provide SNMP access to all your network devices and hosts that support SNMP.
- Set up as many of your network devices as possible to use "universal" read-only credentials. For example, you could define a read-only (get-only) credential to be **nms_get_only**. Create an authentication profile that uses that. Then set up every device possible to allow read-only SNMP access via this universal credential. This minimizes the number of SNMP authentication credentials that must be attempted on network nodes, and simplifies your discovery configuration.
- If there are devices that accept unique SNMP credentials, create one authentication profile for each of those.

Field (SNMP v1 or v2)	Required	Description
Description	Yes	Name for the authentication profile
ID		This read-only field is the Nimsoft system ID for this authentication profile, assigned when the profile is saved. It identifies the profile uniquely for re-use in other areas of USM that reference authentication profiles.

Version	Yes	The SNMP version supported by the monitored device. When version 1 or 2 is selected, only the Community field is active.
Community	Yes	The SNMP community string. Check **Show new passwords** to verify the text as you enter it. Be aware that this string is sent across the network in clear text as part of SNMP v1 or v2c requests, which may pose a security risk.

Field (SNMP v3)	**Required**	**Description**
Description	Yes	Name for the authentication profile
ID		This read-only field is the Nimsoft system ID for this authentication profile, assigned when the profile is saved. It identifies the profile uniquely for re-use in other areas of USM that reference authentication profiles.
Version	Yes	SNMP version supported by the monitored device. Versions 1, 2c, and 3 are supported. When v3 is selected, other fields for security and privacy are enabled.
Password	See note	The password associated with the SNMP v1/v2c device or SNMP v3 user. Check **Show new passwords** to verify the text as you enter it. **Note:** This field is enabled and required if either **AuthNoPriv** or **AuthPriv** security is selected. See the description for the Security field below.
User	Yes	SNMP v3 user name used to access the monitored device. Required for all SNMP v3 security levels. See the description for the Security field below.
Method	Yes	SNMP v3 method of encryption, when **AuthPriv** security is selected (see the description for the Security field below): ■ **None** ■ **MD5** - MD5 Message-Digest Algorithm (HMAC-MD5-96) ■ **SHA** - Secure Hash Algorithm (HMAC-SHA-96)

Security	Yes	SNMP v3 security level of the user. Depending on what level of security is selected, other security fields are enabled or disabled. ■ **NoAuthNoPriv** - messages sent unauthenticated and unencrypted ■ **AuthNoPriv** - messages sent authenticated but unencrypted ■ **AuthPriv** - messages sent authenticated and encrypted
Priv.Password	See note	SNMP v3 privacy password to use if **AuthPriv** security level is selected. Must be at least eight characters. Do not confuse with the user password (authentication). **Note:** This field is enabled and required if **AuthPriv** security is selected
Priv.Protocol	See note	SNMP v3 privacy (encryption) protocol to use. ■ **DES** - Data Encryption Standard ■ **AES** - Advanced Encryption Standard **Note:** Enabled and required if **AuthPriv** is selected.

WMI

WMI (Windows Management Interface) discovery scans servers and hosts running Windows to gather system information. WMI discovery runs only on discovery agents hosted on Windows systems.

Description

Name for the authentication profile.

ID

This read-only field is the Nimsoft system ID for this authentication profile, assigned when the profile is saved. It identifies the profile uniquely for re-use in other areas of USM that reference authentication profiles.

User

User name, in the form of **Domain\user name**. **user_name** and **IP_address\user_name** are also allowable.

Password

User password. Check the **Show new passwords** check box to view the text as you enter it.

Define Ranges

Use the Ranges tab of the Discovery Wizard to define network addresses, ranges, or masks where devices are to be discovered. At least one network range must be entered for discovery to run.

You can assign any combination of SNMP, Linux/Unix, and WMI authentication profiles to a range. The discovery process records *any* device within a range that responds to a request on any protocol, including a simple ICMP ping. This means you can include end nodes (such as servers, network printers, network storage systems, or workstations) in a range, even if they don't respond to requests via SNMP or other management protocols.

If no authentication profile is assigned to a range, basic discovery is performed using protocols that do not require authentication, but discovery may not be complete and information about discovered systems is limited.

Best Practices for Creating Ranges

For each discovery agent, review the assigned ranges to minimize predictable timeouts. To optimize performance and avoid duplicate entries, each discovery agent should discover an exclusive part of the network.

Tips to decrease discovery run time:

- The discovery agent tries each credential on each IP address and waits for a timeout (or success) with each attempt. Use a single credential in a range that has a high probability of immediate success on the nodes in that range to speed up discovery.
- When you apply an authentication profile to a range, make sure that most, if not all, devices defined by that range will accept the authentication profile.
- If you include devices that do not respond to requests on any management protocol, place them in a discovery range with no authentication profiles assigned to the range.
- If you use SNMP for a device that accepts only a unique SNMP community string, create a **Single** type range and specify the device's IP address. Assign the corresponding authentication profile to the range.
- When using SNMP, to avoid unnecessary authentication traps/alerts, assign only one SNMP authentication credential per discovery range.

Create a Range

Follow these steps:

1. Click **New range** in the left pane of the Ranges tab.
2. Enter a name for the range.

3. In the Range definition section, enter one of the following to specify the area(s) of your network where you want to perform discovery:

   - **Mask** - Bitmask for a subnet using Classless Inter-Domain Routing (CIDR) notation with a base IP address and a routing prefix. For example, 195.51.100.0/24. The value /24 refers to a Class C subnet of 256 addresses. Other values for reference: /30 (4 addresses) and /16 (65,536 addresses, or a Class B subnet).
   - **Range** - Range of IP addresses.
   - **Single** - Single IP address.

   An additional field is automatically added if you type the hyphen (-) or slash character (/) to enter a range or mask. The pull-down menu also dynamically updates if you enter a range or mask.

4. To add another IP range, address, or mask, click **New IP range or single IP address** and repeat the previous step.

5. In the Credentials section, you can assign authentication profiles to the selected range. By default, all of the authentication profiles are selected.

   If you have a large number of authentication profiles in the list, you can enter the name of a profile to filter the list.

   To view only the profiles that are selected, select the **Hide unused credentials** check box.

6. When you have finished defining ranges, click **Next**.

Schedule Discovery

In the Schedule tab, you can schedule discovery to run in the future, and/or you can run discovery immediately. You can schedule either a single discovery run or recurring runs.

A scheduled discovery does not interrupt a discovery that is already running. If at the time a discovery run is scheduled another discovery run is in progress, the scheduled discovery is ignored.

If you select **Run discovery now** and discovery is in progress, the current discovery run is terminated and the new run is executed.

Follow these steps to start and/or schedule discovery:

1. Leave the **Run discovery now** check box selected unless you do not want to run discovery when you complete the Discovery Wizard.
2. To schedule discovery, select the **Schedule discovery** check box.
3. Enter information in the date and time fields.

   The time field is in 24-hour format. The time is the local time of the user.
4. To schedule recurring discovery runs, select the **Recurring every** check box, and enter the number of hours for the recurrence interval.
5. Click **Finish** to complete the Discovery Wizard.

Navigating in the Discovery Wizard

There are a few things to be aware of when using the Discovery Wizard:

- If you click the **Close** button or the **X** icon in the title bar before completing the Discovery Wizard, you see a prompt asking whether you want to save your changes. If you execute discovery by clicking **Finish** on the final screen of the Discovery Wizard, changes are retained.
- If valid information is entered in the required fields of an authentication profile or network range, the information is automatically saved when you click **Next**. Required fields are outlined in red.
- Passwords for authentication profiles are displayed as asterisks. If you want to see a password as you enter it, click the Show password icon () next to the **Password** field. After you click **Next**, the password is displayed as asterisks again.

Run File-based Import

Using file-based import, CA Nimsoft administrators can import device and host information into CA Nimsoft Monitor without network scans or manual entry. Because it is not necessary to scan the IT environment, file-based import of devices causes fewer security alerts, and can be faster than automated discovery using the Discovery Wizard.

Note: If a system is discovered by an automated scan of the network and is also included in a file-based import, the file-based import takes precedence. If information about the system differs, the information in the XML file for file-based import is the information that is stored in the database.

Follow these steps:

1. Create an XML file containing information about computers or network devices.

   For details about the contents of the XML file, see the help topic XML File Schema (see page 31).

2. Expand the **Discovery** node in the tree view in the Unified Service Manager.
3. Hover over the **External** node in the tree and click the import icon (), or click the **External** node and choose **Discovery Import** from the **Actions** menu.
4. Navigate to the XML file in the file browser, then click **OK**.

   The device information is imported into the Nimsoft database. Processing by discovery_server starts, and can take several minutes or more to finish.
5. To view imported devices, click the **External** node.

   The devices are displayed in the table to the right.

Alternative import method:

The cm_data_import probe monitors a directory for valid XML files, and if it finds one, it automatically imports the information into the database. Here is how the process works:

1. Copy the XML file you prepared to <Nimsoft install directory>\Probes\Service\cm_data_import**import** directory on the system that hosts the cm_data_import probe.
2. The cm_data_import probe scans this directory at regular intervals (the default is 60 seconds).
3. If the probe finds a valid import file, it imports the device information in the file into the Nimsoft database.
4. The probe moves the file to a timestamped subfolder in the <Nimsoft install directory>\Probes\Service\cm_data_import**processed** directory, also on the probe host, and logs the results of the process.

View Discovered Systems

The **Discovery** node in the tree view of the Unified Service Manager allows you to view computers and devices that have been discovered on your network.

The Discovery section of the tree contains discovery agents, with network ranges under each discovery agent. The tree also has an Automatic and an External node.

Icons next to the tree nodes help identify the type of node and provide additional information:

 - Top-level Discovery node or discovery agent.

 - Network range.

 - Automatic. Some probes automatically discover systems, and those systems are displayed under this node.

 - External. Systems listed under this node were imported using file-based discovery.

 - A discovery is scheduled. Hover over the icon to see the next scheduled time in the tool tip.

 - Discovery in progress. The proportion of blue indicates the progress of discovery.

 - No discovery scheduled.

Click a node in the tree to view associated systems and their properties in the table to the right. To view properties for all discovered systems, click the **Discovery** node.

A pie chart above the table displays information about discovered systems for the selected node. Choose a different criterion (**Device Type**, **Operating System,** etc.) from the pull-down menu to change the data displayed in the pie chart.

Click a slice in the pie chart or an item in the chart legend to filter for those systems. Only the systems represented in the slice are displayed in the table and reflected in the response links to the right. Click the slice or legend item again to clear the filter.

The response links to the right of the pie chart list systems according to how recently they responded to a request from the discovery agent. Click one of these links, such as **Recent (last day)**, to filter for those systems. Only those systems are displayed in the pie chart and in the table. Click the link again to clear the filter.

Note: Systems that do not respond are eventually purged from the database. By default, 30 days after the last response from a system, the system is deleted from the database.

A Quick Filter field below the response links allows you to filter for text in the **Name**, **IP Address**, **Domain**, **OS Name**, and **Origin** columns of the table.

Click a column header to sort the table by the column.

A key icon () in the table indicates a discovery agent was able to authenticate with the system using one of the defined authentication profiles. Hover over the key icon to view the type and name of the authentication profile used.

Click the reload icon () to refresh discovery data. This may take a few moments; the time depends on the number of systems in your environment and the speed of your network connections. After you click it, the icon is disabled until discovery is finished, then the icon is enabled again.

You can export data for a discovery agent or network range. The data includes more columns than are displayed in the Inventory table. Data is exported to a .csv file, which is saved in a location you choose. To export data, click a discovery agent or network range in the tree, then select **Export Group** from the **Actions** menu.

Note: When you choose **Export Group**, all systems for the selected discovery agent, or selected network range, are exported, regardless of whether you filtered the display in the Inventory view.

Appendix A: Advanced Configuration

Note: Automated discovery scan settings, such as network ranges and authentication credential profiles, are configured within the Discovery Wizard that runs within the USM portlet in UMP. For information, see the section on the Discovery Wizard (see page 16).

This section contains the following topics:

Running discovery_server on a Robot Other Than the Primary Hub (see page 27)
Setting Maximum Java Heap Size (see page 28)
File-based Import Reference (see page 30)

Running discovery_server on a Robot Other Than the Primary Hub

By default, the discovery server runs on the primary hub, which is the same robot where the data_engine is running. The discovery server can run on a different robot as long as the discovery server can communicate with the data_engine probe, and the database server, from its new location. To run the discovery server on a different robot other than the primary hub, follow these steps:

1. Deactivate or delete the discovery server on the primary hub--only one instance of the discovery server can be deployed.
2. In Infrastructure Manager, right click on the discovery_server probe on the secondary hub. In Admin Console, click on the icon next to the discovery_server on the secondary hub.
3. Select Raw Configure.

4. In the content window navigate to the setup > data_engine key and click the Edit Key button. In Admin Console, click on the value field to edit it.

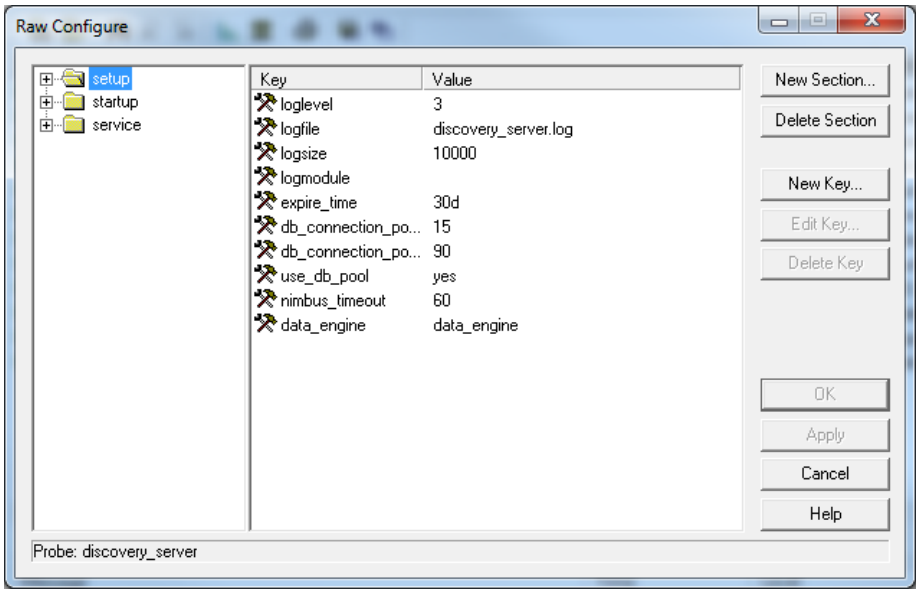

5. Specify the full data_engine probe address. You can look up the data_engine address in Infrastructure Manager under the primary hub's SLM category. For example: /my_domain/my_primary_hub/my_primary_robot/data_engine.

6. Activate or restart the discovery_server in its new location.

Setting Maximum Java Heap Size

The default maximum Java heap size for the discovery_server and discovery_agent probes is set using the Raw Configure option.

Discovery Server

The default maximum Java heap size is set to 1 GB and is intended to support up to 5000 robots. For systems with more than 5000 robots, the maximum Java heap size should be increased by 1 GB for every 5000 additional robots.

Example: 2 GB for 5001 to 10,000 robots; 3 GB for 10,001 to 15,000 robots.

To set the maximum Java heap size:

1. In Infrastructure Manager, hold the Shift key and right-click on the discovery_agent probe, and select **Raw Configure**. In Admin Console, click on the icon next to the discovery_agent and select **Raw Configure**.

2. Navigate to **startup**, then open **opt**.

3. In the content window select **java_mem_max**. In Infrastructure Manger, click the **Edit Key** button. In Admin Console, click on the value field to edit it.

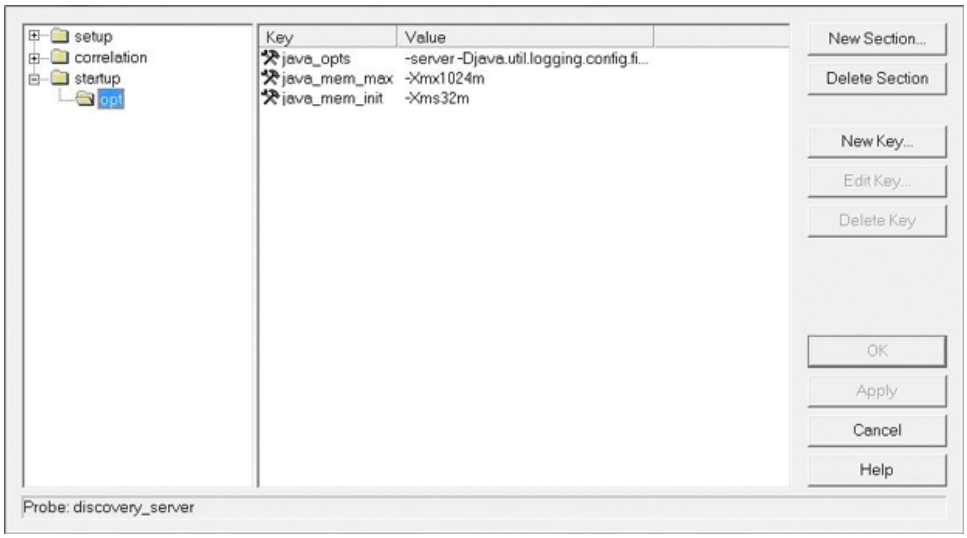

4. Enter the new value using increments of 1024 MB.
   - 1 GB = -Xmx1024m
   - 2 GB = -Xmx2048m

Discovery Agent

The default maximum Java heap size is set to 256 MB. For very large discovery ranges (equivalent to a class B subnet, or in excess of 30,000 addressable devices), increasing the maximum heap allocation to 512 MB or 1024 MB is recommended.

To set the maximum Java heap size:

1. In Infrastructure Manager, hold the Shift key and right-click on the discovery_agent probe, and select **Raw Configure**. In Admin Console, click on the icon next to the discovery_agent and select **Raw Configure**.
2. Navigate to **startup**, the open **opt**.

3. In the content window select **java_mem_max**. In Infrastructure Manger, click the **Edit Key** button. In Admin Console, click on the value field to edit it.

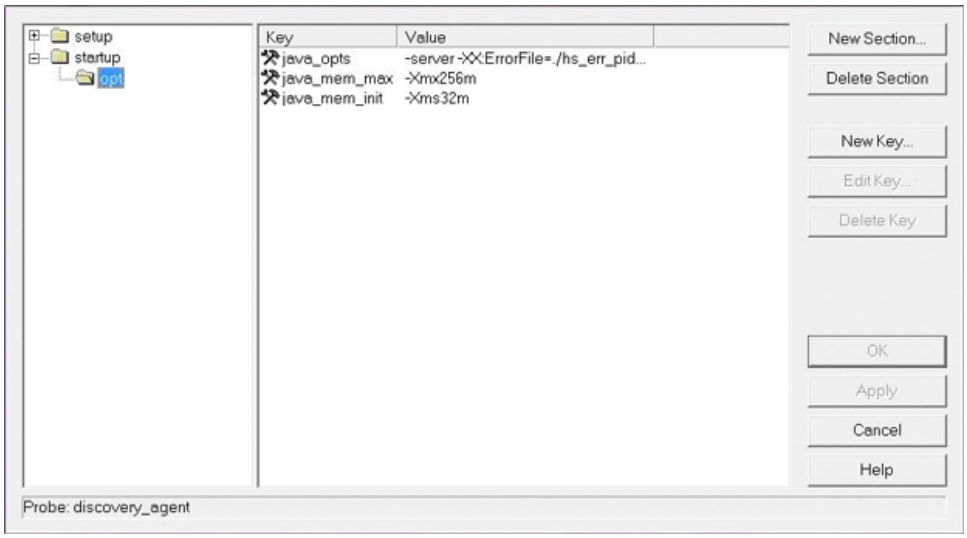

4. Enter the new value.
   - 512 MB = -Xmx512m
   - 1 GB = -Xmx1024m

File-based Import Reference

The discovery function called *file-based import* provides a convenient way to import device description data into the discovery database. File-based import provides an alternative to automated discovery for populating the discovery inventory, without incurring the overhead of scanning the IT environment.

The cm_data_import probe processes device data in an XML file that you prepare. It can be initiated in either of two ways:

Option 1: Open a file browser from the External node of the Discovery Wizard in USM, and navigate to the prepared XML file on your local file system to process. ■ Click **OK.** ■ The file is processed. ■ The devices are published to the Nimsoft bus. The discovery_server receives this information and adds the devices to the device database. ■ The devices are displayed in USM.	**Option 2**: Copy an XML file into the Nimsoft\Probes\Service\cm_data_import\import directory on the machine hosting the cm_data_import probe. ■ The cm_data_import probe scans this directory at a configurable interval (default is 60 seconds). If it finds a file it will process it, then moves it to a time-stamped subfolder within the directory <Nimsoft>\Probes\Service\cm_data_import\processed ■ The result of the process is logged. ■ The probe publishes the devices to the Nimsoft bus. The discovery_server receives this information and adds the devices to the device database. ■ The devices are displayed in USM.

Devices are visible in the USM interface either under Groups, or listed under the **External** branch of the Discovery tree.

Note: Devices imported via file-based import are not reflected in Nimsoft Topology.

XML File Schema

This section describes how to create an XML file for use with file-base discovery.

The XML file must include these required properties for each host or device:

- PrimaryIPv4Address - List the IPv4 address. Although the PrimaryIPv6Address tag exists, IPv6 addresses are not currently supported in discovery.
- Origin - Setting the origin correctly is important. See details on the Origin property in the table below.

Here is an example of XML that illustrates how to import one device with IP address 1.2.3.4 and origin "MyOrigin" in the database.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<DevicesToImport xmlns="http://nimsoft.com/2012/11/cm-data-import"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
        <Device>
                <PrimaryIPV4Address>1.2.3.4</PrimaryIPV4Address>
                <Origin>myOrigin</Origin>
        </Device>
</DevicesToImport>
```

Additional optional properties can be included, as shown in the example below. You can also find this example file, named example1MaximalDevice.xml, in the <Nimsoft install directory>\Probes\Service\cm_data_import**schema** directory, located on the system that hosts the cm_data_import probe--typically the primary hub.

```
  <?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
  <DevicesToImport xmlns="http://nimsoft.com/2012/11/cm-data-import"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
        <Device>
              <Origin>myOrigin</Origin>
              <Label>myComputer</Label>
              <Description>myComputer description goes here</Description>
              <PrimaryDnsName>myComputer.myCompany.com</PrimaryDnsName>
              <PrimaryIPV4Address>1.2.3.4</PrimaryIPV4Address>
              <PrimaryIPV6Address>fe80::223:ebff:fe06:9d40%10</PrimaryIPV6Address>
              <PrimaryMacAddress>F0-4D-A2-25-5B-7A</PrimaryMacAddress>
              <PrimaryOSType>WindowsServer-2008</PrimaryOSType>
              <PrimaryOSVersion>6.1.7601</PrimaryOSVersion>
              <ProcessorType>x86-64</ProcessorType>
              <Vendor>Dell Inc.</Vendor>
              <Model>PowerEdge T620</Model>
              <PhysSerialNumber>123-456-789-ABCD</PhysSerialNumber>
              <PrimaryDeviceRole>VirtualMachine</PrimaryDeviceRole>
              <PrimarySoftwareRole>DatabaseServer</PrimarySoftwareRole>
              <DBServerType>MSSQLServer</DBServerType>
              <WmiAuthId>3</WmiAuthId>
              <ShellAuthId>5</ShellAuthId>
              <SnmpAuthId>7</SnmpAuthId>
              <AppServerType>Unknown</AppServerType>
              <VirtualizationEnvironment>VMware</VirtualizationEnvironment>
              <MonitorFrom>monitoringRobotHostName</MonitorFrom>
       </Device>
</DevicesToImport>
```

The following table describes the XML properties. For properties that refer to open enumerations, navigate to <Nimsoft installation path>\Probes\Service\cm_data_import**schema** and open either **usm-openenums.xml** or **cm-data-import-openenums.xml** to view the defined values for each enumeration instance. It is strongly recommended you use values defined by the open enumerations, though not strictly required.

To deploy a robot to an imported system using USM and ADE, some additional properties beyond IP address and origin are required. These are noted in the table below.

Property	Required?	Description
Origin	Yes	QoS data from probes are tagged with an origin name to identify the origin of the data. The origin name defaults to the Nimsoft hub name but can be overridden at the hub or robot (controller) in order to separate data in a multi-tenancy environment. To ensure that QoS data from probes is correlated to this device, the origin name specified here should match the origin name you intend to use in your Nimsoft infrastructure of hubs and robots.
Label	No	A short description or caption.
Description	No	Text description of the device.
PrimaryDnsName	No	The entity's Domain System Name, which may be used for correlation.
PrimaryIPV4Address	An IPv4 address is required	An IPv4 address for the entity that may be used for correlation and identity.
PrimaryIPV6Address	No	An IPv6 address for the entity that may be used for correlation and identity. The address is expressed using the formal, complete IPv6 notation (8 groups of up to 4 hex digits, using only uppercase where applicable, separated by colons).
OtherIPAddresses	No	An entity may have multiple IP addresses. This element captures the values of those addresses, whereas the PrimaryIPAddress elements are designed to be used for correlation and identity. The various values are comma-separated. Either IPv4 or IPv6 values may be specified but the addresses should be formatted following the regex patterns defined by usm-core:IPV4AddressFormat or usm-core:IPV6AddressFormat.

Property	Required?	Description
PrimaryMacAddress	No	A MAC address for the entity that may be used for correlation and identity. The address is expressed as 6 groups of 2 hex digits (using only uppercase), separated by dashes.
OtherMacAddress	No	An entity may have multiple MAC addresses. This element captures the values of those addresses, whereas the PrimaryMacAddress element is designed to be used for correlation. The various values are comma-separated and are formatted following the regex pattern defined by usm-core:MacAddressFormat.
PrimaryOSType	Required by ADE for robot deployment	OS type, defined by the open enumeration OSTypeEnum. For Linux, the Linux distribution name is required by ADE (for example, **Linux-RedHat**).
PrimaryOSVersion	No	OS version details.
ProcessorType	Required by ADE for robot deployment	Processor environment/type (such as "x86") as defined by the open enumeration ProcessorEnvironmentEnum.
Vendor	No	The hardware vendor/manufacturer's name, as defined by the open enumeration VendorEnum.
Model	No	The hardware model name/number.
PhysSerialNumber	No	An identifying string assigned by the hardware manufacturer and printed on a tag attached to the component. The data for this element should be input directly from the manufacturer's tag on the component (which may be an RFID tag), or read from the entPhysicalSerialNum field of SNMP's Entity-MIB. Note that a virtual entity would NOT have a PhysSerialNumber.
PrimaryDeviceRole	No	The device role as defined by the open enumeration DeviceRoleEnum.
PrimarySoftwareRole	No	The software role as defined by the open enumeration SoftwareRoleEnum.

Property	Required?	Description
DBServerType	No	The type of database server of which this is an instance, defined by the open enumeration DBServerTypeEnum.
AppServerType	No	The type of application server, as defined by the open enumeration AppServerTypeEnum.
VirtualizationEnvironment	No	Value indicating the specific virtualization environment (hypervisor manager) of a hypervisor or virtual system. Values are defined in the open enumeration VirtualizationTypeEnum.
WmiAuthId	Either WmiAuthId or ShellAuthID is required by ADE for robot deployment	A Nimsoft defined authentication profile ID to use for WMI access. This is the ID field in the WMI authentication profile.
ShellAuthId	Either WmiAuthId or ShellAuthID is required by ADE for robot deployment	A Nimsoft defined authentication profile ID to use for SSH or telnet access. This is the ID field in the Shell authentication profile.
SnmpAuthId	No	A Nimsoft defined authentication profile ID to use for SNMP access. This is the ID field in the SNMP authentication profile.

Property	Required?	Description
MonitorFrom	No	If the device will be remotely monitored, this specifies the system to monitor this device from. The value can be specified as an IP address, simple host name, fully qualified domain name or Nimsoft address (/NimsoftDomain/HubName/RobotName). A Nimsoft robot should be installed on the system specified here. If the robot is not installed, this device will not be imported. The origin name used by the robot should match the origin specified for this device to ensure that QoS data from probes is correlated with this device.