# CA Nimsoft Monitor

## Discovery User Guide

**v7.5**


March 2014

# Legal Notices

This online help system (the "System") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This System may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This System is confidential and proprietary information of CA and protected by the copyright laws of the United States and international treaties. This System may not be disclosed by you or used for any purpose other than as may be permitted in a separate agreement between you and CA governing your use of the CA software to which the System relates (the “CA Software”). Such agreement is not modified in any way by the terms of this notice.

Notwithstanding the foregoing, if you are a licensed user of the CA Software you may make one copy of the System for internal use by you and your employees, provided that all CA copyright notices and legends are affixed to the reproduced copy.

The right to make a copy of the System is limited to the period during which the license for the CA Software remains in full force and effect. Should the license terminate for any reason, it shall be your responsibility to certify in writing to CA that all copies and partial copies of the System have been destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS SYSTEM “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS SYSTEM, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The manufacturer of this System is CA.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.



Legal information on third-party and public domain software used in the Nimsoft Monitor solution is documented in *Nimsoft Monitor Third-Party Licenses and Terms of Use (*http://docs.nimsoft.com/prodhelp/en_US/Library/Legal.html*).*

# Contact CA

**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

| Version | Date | Changes |
|---|---|---|
| 7.5 | March 2014 | Minor updates for NMS 7.5. |
| 7.1 | December 2013 | Revised for NMS 7.1:<br>■ Updated to include discovery of IPV6 devices. |
| 7.0 | September 2013 | Revised for NMS 7.0:<br>■ Addition of device correlation.<br>■ New probe_discovery queue.<br>■ Changes to the Discovery Wizard.<br>■ Elimination of discovery probe user interfaces.<br>■ Content improvement. |
| 6.5 | April 2013 | First edition of the guide, covering Nimsoft Discovery as implemented in NMS v6.5. |

# Contents

# Chapter 1: Introduction

## Discovery Architecture

A critical part of IT monitoring is creating and maintaining an accurate list of the devices in your IT environment. Finding and listing all addressable devices and computers within a managed IT environment the job of automated *discovery*.

When the CA Nimsoft Unified Management Portal (UMP) is installed, the Discovery Wizard starts automatically and prompts you to configure and run discovery. The wizard allows you to specify 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.

The list of devices, referred to as your *Inventory*, can be augmented by XML file-based device import. When multiple discovery records correspond to a single device, this is recognized by device correlation.

To maintain your inventory, you can re-run discovery at any time, modifying the credentials and ranges as needed. You also can schedule discovery to run on regular intervals. This diagram illustrates the flow of data among the key components of discovery:

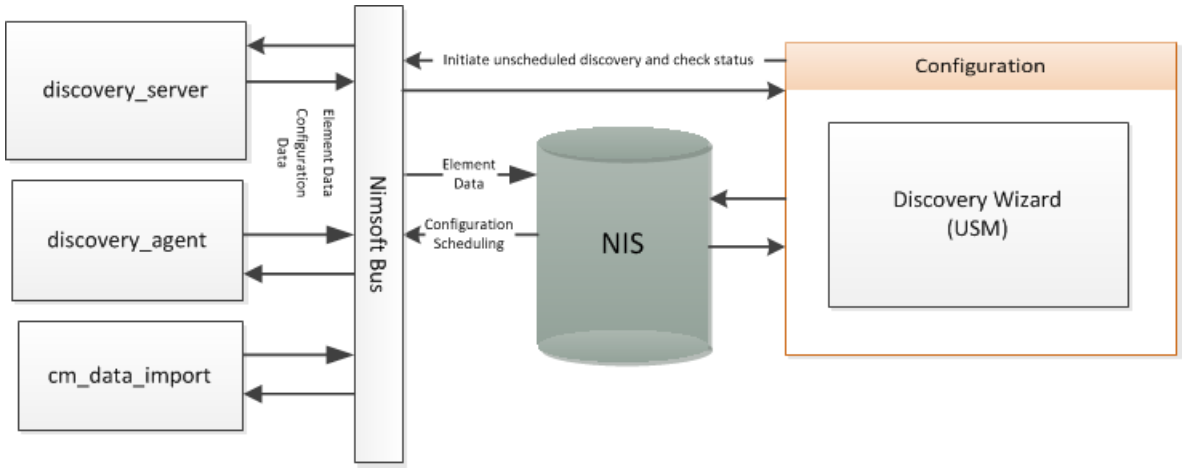

# Discovery Components

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

**Discovery Wizard**

The Discovery Wizard lets you easily configure discovery scans. To launch the wizard, open the Unified Service Manager (USM) portlet in the Unified Management Portal (UMP) and select **Actions**. You also can run the wizard from any discovery agent node in the Discovery tree of USM. The wizard lets you specify authentication profiles and the range of addresses you want to search. Discovery then uses this information to scan the network and populate the device inventory.

**Discovery Server probe**

In most installations, the discovery_server probe runs on the primary hub. The 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 device 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 that don't have any associated QOS data.

**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 probe**

The discovery_agent probe scans the IT network, pinging and querying devices according to subnet masks/ranges, credential profiles, and selected profiles. These scanning parameters are configured within the Discovery Wizard.

**CM Data Import**

This probe processes an XML file that describes 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.

Additional components that play a role in discovery:

**probeDiscovery queue**

This queue on the primary hub collects discovery data that is processed by the discovery server. On secondary hubs, you will configure probe_discovery queues to collect data and route it to the primary hub. See 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.

**Other Nimsoft monitoring probes**

All monitoring probes provide information about systems that are monitored to the Discovery Server. Several of these probes publish directly to the probeDiscovery queue. These monitoring probes help supplement auto-discovery.

# Discovery Considerations

- The Nimsoft Monitor *Topology and Root Cause Analysis* feature uses data provided by discovery to 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.
- Raw discovery data is correlated so that a master record is created when a single device responds to discovery in multiple ways. For example, when a multi-homed device responds to discovery pings on multiple IP addresses, it is reported correctly as a single system, rather than multiple devices.

# Prerequisites and Supported Platforms

- Discovery 7.x requires NMS 7.x.
- Discovery Server 7.x only works with 7.x Discovery Agents. The Discovery Server raises an alarm for any pre-v7.0 Discovery Agent it finds.
- Discovery Server 7.x does not collect any discovery results from pre-7.0 discovery agents.

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

# Chapter 2: Configuring Discovery

Here is how the discovery process works.

1. You install NMS, which includes the components required for discovery.
2. If your NMS installation includes secondary hubs, you must configure *probeDiscovery* queues so that probe_discovery messages reach the primary hub. See Configure Discovery Queues (see page 13).
3. You install UMP, which includes the Discovery Wizard.
4. After you install UMP, the Discovery Wizard launches in USM and leads you through the process of configuring discovery. You will:
   a. Create authentication profiles (see page 16)
   b. Define ranges (see page 21) (sets or ranges of IP addresses and IP masks that define and bound the scope of discovery)
   c. Schedule discovery (see page 24).
5. To augment automated discovery, you can prepare an XML file with device information and import this information into the device inventory (optional). See Run File-based Import (see page 25).
6. When discovery is complete, you can view computers and devices that have been discovered on your network. See View discovered systems (see page 26).

# Discovery Probe Deployment

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, configure the discovery_agent probe on the NMS primary hub or deploy and configure a discovery agent at another location.
- 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 Discovery Queues

If all of your discovery process probes are deployed on a single hub, communication of discovery data is automatically configured. However, if discovery probes are deployed to 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.

This is accomplished by setting up queues that handle the *probe_discovery* subject. You will set up *attach* queues (which collect messages) on the primary hub and on hubs that host discovery_agent or any CTD-publishing probes. This list includes (but is not limited to):

- discovery_agent
- vmware 5.10 or later
- cm_data_import (typically deployed with discovery_server on the primary hub)
- snmpcollector
- vcloud
- rhev (Red Hat Virtualization)

You will set up a corresponding *get* queue (which retrieves messages from the attach queue) on the primary hub and on any hub that needs to transfer the messages to another hub.

The following illustration shows where queues are required.

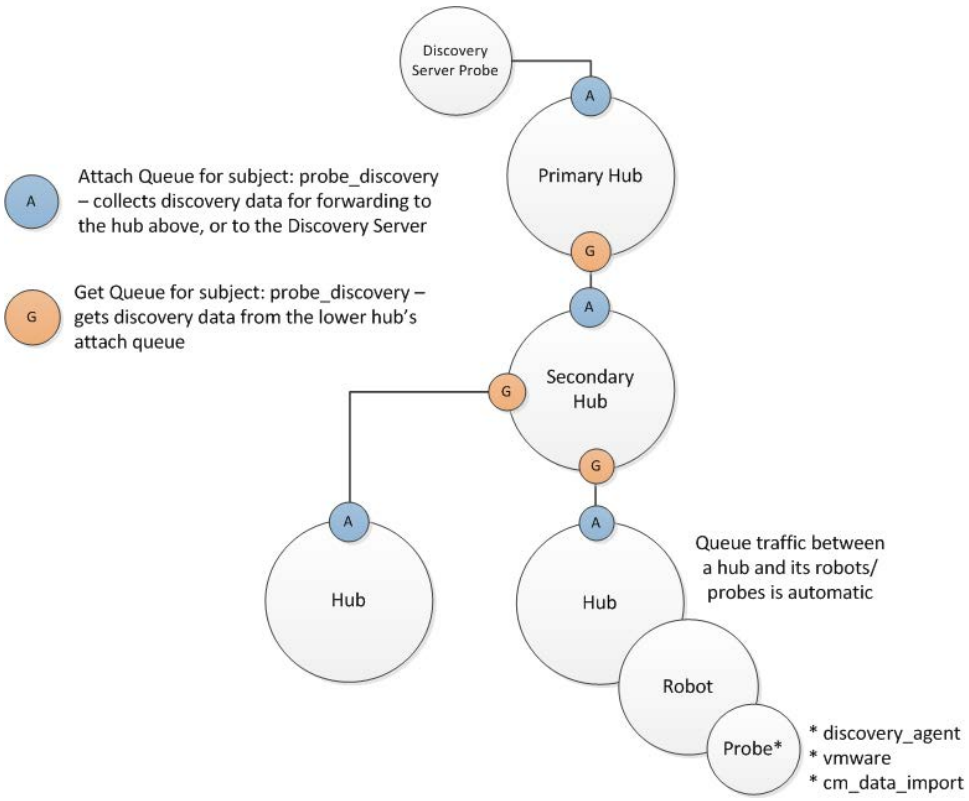


You can set up discovery queues in either Admin Console or Infrastructure Manager.

1. Identify the hub on which you want to create a queue and open the hub configuration GUI:
   - *Admin Console:* expand the hub in the navigation tree and select its robot. Click the arrow next to the hub probe and select **Configure**.
   - *Infrastructure Manager:* expand the hub's node and double-click the hub probe.
2. Navigate to **Queue List** or **Queue**.
3. Set up the required attach and/or get queues with the following values:
   - Attach queue:
     - **Active:** enabled
     - **Name:** probeDiscovery
     - **Type:** attach
     - **Subject:** probe_discovery

- Get queue:
  - **Active:** enabled
  - **Name:** probeDiscovery
  - **Type:** get
  - **Address:** address of the hub that has the attach queue
  - **Bulk size:** number of messages to be sent together (optional; if you expect the queue to carry a significant amount of messages, sending them in bulk can improve performance)

  For queue setup details, click the question mark or **Help** button in the configuration GUI.

4. Repeat the previous steps on all hubs that require a queue.

When you have set up all required queues, run an automated discovery scan to confirm the queues are operational. Review the list of discovered devices. In addition to local devices, it also 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:** Discovery Agent uses password authentication to connect to a target device over SSH. Discovery Agent cannot communicate with a device where SSH is configured for other authentication methods, such as keyboard-interactive or challenge-response authentication.

## SNMP

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**. 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. You can specify a unique port within the range of 1 to 65535 for the profile. If no port is specified, the default port 161 is used.

| 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.<br>**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 **AuthNoPriv** or **AuthPriv** security is selected (see the description for the Security field below):<br>■ **None**<br>■ **MD5** - MD5 Message-Digest Algorithm (HMAC-MD5-96)<br>■ **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.<br>■ **NoAuthNoPriv** - messages sent unauthenticated and unencrypted<br>■ **AuthNoPriv** - **messages** sent authenticated but unencrypted<br>■ **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).<br>**Note:** This field is enabled and required if **AuthPriv** security is selected |
| Priv.Protocol | See note | SNMP v3 privacy (encryption) protocol to use.<br>■ **DES** - Data **Encryption** Standard<br>■ **AES** - **Advanced** Encryption Standard<br>**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, 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 IPv4 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).

     **Note**: When you enter a subnet mask, the number of IP addresses the mask represents is displayed (the number of effective hosts minus two). Only /16 subnets or smaller are supported.
   - **Range** - Range of IPv4 addresses.
   - **Single** - Single IPv4 or IPv6 address. You can use abbreviated IPv6 address forms, and IPv6 addresses that refer to IPv4 addresses. However, IPv6 anycast, multicast, and loopback addresses, and default routes are *not* supported.

   You can also click the Add multiple IPs icon (  ) above the Range definition section. Copy and paste the IP addresses into the Import IPs dialog, one entry per line. After you click **OK**, any errors are highlighted in red.
4. Click **New IP range or single IP address** to add another IP range, address, or mask if desired.
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, click 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. 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 33).

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.

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 15).

This section contains the following topics:

# 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.

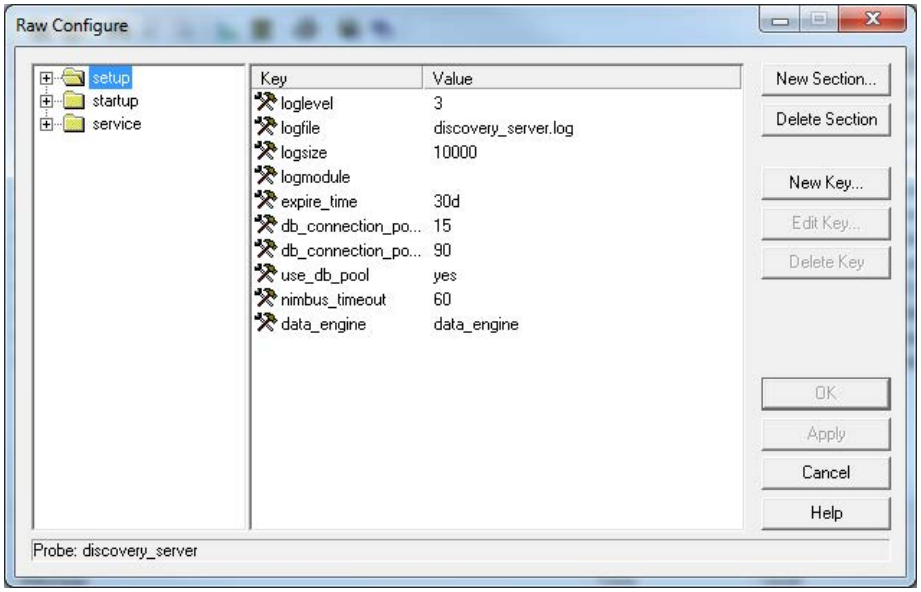


5. Specify the full data_engine probe address (*/domain/primary_hub/primary_robot*/data_engine). You can look up the data_engine address in Infrastructure Manager under the primary hub's SLM category.
6. Activate or restart the discovery_server in its new location.

# 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 1 GB and is intended to support up to 5000 robots. For deployments with more than 5000 robots, we recommend you increase the maximum Java heap size by 1 GB for every 5000 additional robots (2 GB for 5001 to 10,000 robots; 3 GB for 10,001 to 15,000 robots).

1. Open the discovery_server probe in Raw Configure:
   - *Admin Console:* click the icon next to the probe and select **Raw Configure**.
   - *Infrastructure Manager:* shift+right-click the probe and select **Raw Configure**.
2. Navigate to **startup > opt**.
3. Enter the desired value for **java_mem_max** using increments of 1024 MB:
   - 1 GB = -Xmx1024m
   - 2 GB = -Xmx2048m

## Discovery Agent

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

1. Open the discovery_agent probe in Raw Configure:
   - *Admin Console:* click the icon next to the probe and select **Raw Configure**.
   - *Infrastructure Manager:* shift+right-click the probe and select **Raw Configure**.
2. Navigate to **startup > opt**.
3. Enter the desired value for **java_mem_max**:
   - 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 the device data included in an XML file.Initiate the process with one of these methods.

**Method 1**

1. Open a file browser from the External node of the Discovery Wizard in USM.
2. Navigate to the prepared XML file on your local file system to process and click **OK**.
3. The file is processed. When processing is complete:
   - 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.

**Method 2**

1. Copy an XML file into the *<Nimsoft>\Probes\Service\cm_data_import\import* directory on the system hosting the cm_data_import probe.
2. The cm_data_import probe recognizes the new files and processes it. This probe scans the directory at a configurable interval (default is 60 seconds). When processing is complete:
   - cm_data_import moves the file to a time-stamped subfolder in *<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-based 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 XML example 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 | Either IPv4 or IPv6 is required | An IPv4 address for the entity that may be used for correlation and identity. |
| PrimaryIPV6Address | Either IPv4 or IPv6 is required | 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 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. |
| 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**). |

| Property | Required? | Description |
|---|---|---|
| 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. |
| 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 | ADE requires WmiAuthId or ShellAuthID 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 | ADE requires WmiAuthId or ShellAuthID 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. |