CA Nimsoft Monitor

Discovery User Guide

v8.0

September 2014

Copyright Notice

Legal information on third-party and public domain software used in this product is documented in the *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 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 support website at http://ca.com/docs.

Document Revision History

Version	Date	Changes
7.6	June 2014	Revised for NMS 7.6: ■ File-based import of SNMP authentication profiles
7.5	March 2014	Minor updates for NMS 7.5
7.1	December 2013	Revised for NMS 7.1: ■ Dscovery of IPV6 devices
7.0	September 2013	Revised for NMS 7.0: ■ Device correlation ■ Probe_discovery queue ■ Changes to the Discovery Wizard ■ Elimination of discovery probe configuration GUIs ■ Content improvement
6.5	April 2013	First edition, documents Discovery as implemented in NMS v6.5

Contents

Chapter 1: Introduction 7

Chapter 2: Configuring Discovery 11

Appendix A: Advanced Configuration 41

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 is the job of automated *discovery*.

When the CA Unified Infrastructure Management 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:

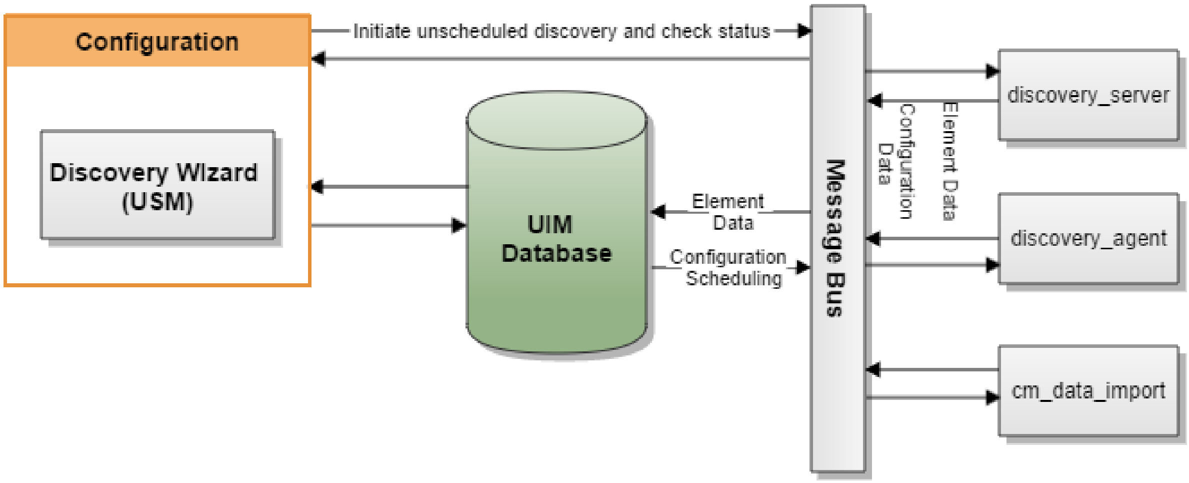

Discovery Components

All discovery components are included in a basic installation of CA Unified Infrastructure Management 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 UIM Server 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 stored in the UIM Server database and used by other UIM Server components. The discovery_server probe also helps maintain the 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.

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 devices and authentication profiles. Device information is added to your inventory. Authentication profiles are accessible in the Discovery Wizard. The cm_data_import probe is usually co-located with discovery_server, typically on the primary hub. When you run file-based import from the Discovery Wizard, cm_data_import does 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).

UIM Database

The UIM Database is the database that holds all persistent data, including discovery data.

Other monitoring probes

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

Discovery Considerations

- 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.
- The *Topology and Root Cause Analysis* feature uses data provided by discovery to deduce the network structure and create a model that you can view in the Relationship Viewer portlet in UMP. However, devices that are imported through file-based import are not reflected in topology or in root cause analysis, because topology depends on SNMP information gathered by the discovery agent about the 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 system platforms, see the Compatibility Support Matrix.

Chapter 2: Configuring Discovery

Here is how the discovery process works.

1. The components required for discovery are deployed when you install UIM Server. See Discovery Probe Deployment (see page 12) for deployment considerations.
2. If your installation includes secondary hubs, you configure *probeDiscovery* queues so that messages with the probe_discovery subject 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 17).

    b. Define ranges (see page 21) of IP addresses and IP masks that define and bound the scope of discovery.

    c. Schedule discovery (see page 24).

    Note: If you don't want to create your inventory at this time, or if you want to create it solely with file-based import, cancel the wizard. You can run discovery or import devices at any time.

5. To augment automated discovery, you can prepare an XML file with device information and import this information into the device inventory. See Run File-based Import (see page 25).

    Note: If desired, you can create your inventory solely with file-based import.

6. When discovery is complete, you can view computers and devices that have been discovered on your network. See View discovered systems (see page 38).

Discovery Probe Deployment

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

- 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.
- A domain should have a single instance of discovery_server. Deploying multiple discovery_server probes is not supported and will result in adverse behavior.
- To add network scanning, configure the discovery_agent probe on the UIM Server 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 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 "**uim_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 can use a combination of *attach* and *get* queues. An *attach* queue creates a permanent queue on a downstream hub. A corresponding *get* queue set up on the upstream hub is paired with each *attach* queue to retrieve messages from the downstream hub.
- Alternatively, you can use *post* queues. A *post* queue set up on a downstream hub sends a directed stream of messages to the upstream hub.

An *attach* queue is automatically set up on the primary hub to collect discovery data. You need to set up additional queues to collect discovery data from downstream 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 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.

Navigate

2. Navigate to **Queue List** or **Queue**.
3. If using *attach*/*get* queues, set up the queues with the corresponding values:
   - *Attach* queue on the downstream hub:
     - **Active:** enabled
     - **Name:** probeDiscovery (or other name of your choice)
     - **Type:** attach
     - **Subject:** probe_discovery
   - *Get* queue on the upstream hub:
     - **Active:** enabled
     - **Name:** probeDiscovery (or other name of your choice)
     - **Type:** get
     - **Address:** address of the hub that has the *attach* queue
     - **Queue:** name of the corresponding *attach* queue
     - **Bulk size:** number of messages to be transferred together (optional; if you expect the queue to carry a significant number of messages, sending them in bulk can improve performance)
4. If using *post* queues, set up the queue with the corresponding values:
   - *Post* queue on the downstream hub:
     - **Active:** enabled
     - **Name:** probeDiscovery (or other name of your choice)
     - **Type:** post
     - **Subject:** probe_discovery
     - **Address:** address of the upstream hub
     - **Bulk size:** number of messages to be transferred together

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

   For queue setup details, click the question mark or **Help** button in the configuration GUI.
5. Repeat the previous steps on all hubs that require a queue.

The following illustration shows discovery queue configuration using *attach*/*get* pairs. If you choose to use a *post* queue configuration, the flow in the illustration is similar. The difference is that you set up *post* queues instead of *attach* queues on all downstream hubs, and you do not need to set up *get* queues.

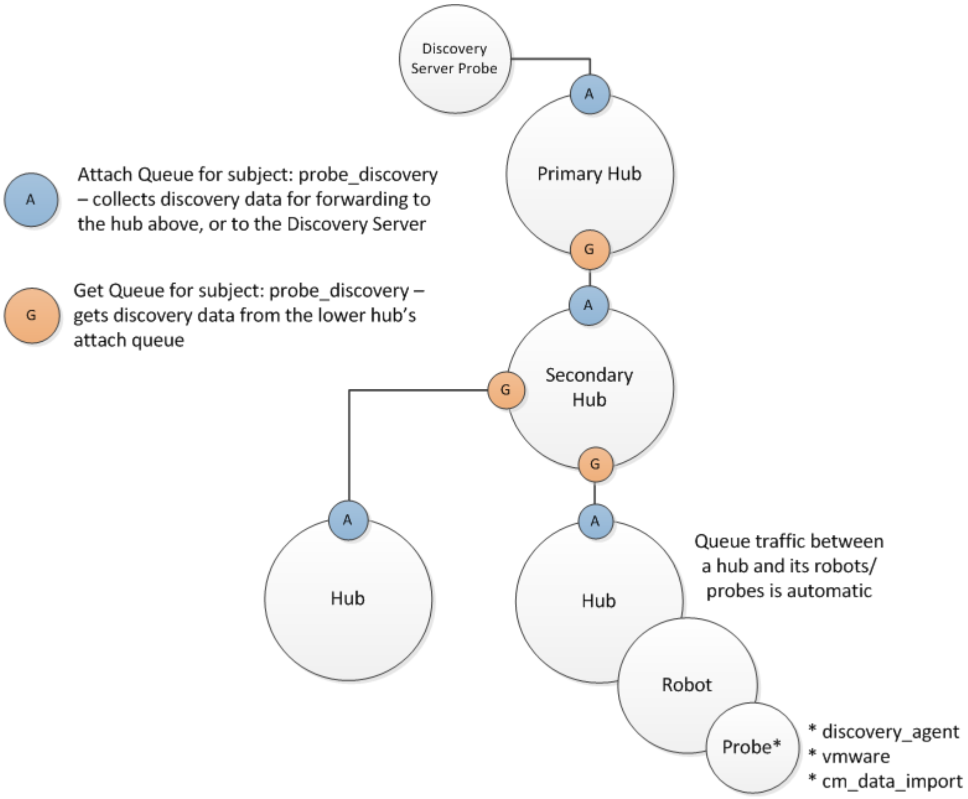
Discovery Server Probe
A
Primary Hub
G
A
Attach Queue for subject: probe_discovery – collects discovery data for forwarding to the hub above, or to the Discovery Server
G
Get Queue for subject: probe_discovery – gets discovery data from the lower hub's attach queue
A
Secondary Hub
G
G
A
Hub
A
Hub
Queue traffic between a hub and its robots/ probes is automatic
Robot
Probe*
* discovery_agent
* vmware
* cm_data_import

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, this list should also 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.

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 Unified Infrastructure Management if there are existing scopes 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 (🖧).

   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.

Navigating in the Discovery Wizard

Be aware of the following 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 are prompted to save your changes. If you execute discovery by clicking **Finish** on the final screen, changes are saved.
- 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. When you click **Next**, the password is displayed as asterisks.

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.

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.

To view the properties of an existing profile, select the **WMI**, **Linux/Unix**, or **SNMP** tab, and select a profile in the left pane.

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**.

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 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

Discovery supports SNMP versions 1, 2c, and 3. SNMP v3 provides security features that are not available in v1 and v2c. As a result, authentication profile configuration fields in the Discovery Wizard that deal with security and privacy (encryption) are only active when you select **3** in the **Version** pull-down menu.

Note: SNMP authentication profiles can also be imported from an XML file. See Run File-based Import (see page 25) for details.

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 **uim_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.
- For network devices such as routers and switches, SNMP is the sole source for detailed discovery information. For host systems such as Windows, Unix, or Linux servers, it is recommended that you use WMI or SSH discovery instead of, or in addition to SNMP. While SNMP provides detailed network interface information for devices and systems, the host system information available from SNMP, such as processor attributes, is less complete than the information obtained from WMI or SSH discovery.

Field **(SNMP v1 or v2)**	**Required**	**Description**
Description	Yes	Name for the authentication profile
ID		This read-only field is the 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 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	Enabled and required if **AuthNoPriv** or **AuthPriv** is selected (see *Security* description)	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 **AuthNoPriv** or **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	Enabled and required if **AuthPriv** is selected	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).

Priv.Protocol	Enabled and required if **AuthPriv** is selected	SNMP v3 privacy (encryption) protocol to use. ■ **DES** - Data Encryption Standard ■ **AES** - Advanced Encryption Standard

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 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 Scopes

Use the **Scopes** tab of the Discovery Wizard to define network seed devices, 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 scope. 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 using SNMP or other management protocols.

You can also add network seed devices to act as an initial communication point for discovering your network topology. Seed devices are core devices that Discovery can use as a starting point for determining the routed subnets in your network and help accelerate the discovery of known devices on the network. Use the **New Seed Scope** node to add individual seed device IPs. By default, the **LAN** check box is selected. This allows discovery of all the devices on the local subnets determined from the seed device. If you want to limit seed device discovery to a previously defined range scope, clear the **LAN** check box.

If no authentication profile is assigned to either a range or seed scope, basic discovery is performed using protocols that do not require authentication, but discovery may not be complete and information about discovered systems is limited. A seed scope requires at least one SNMP credential. The same authentication credentials are used for both range and seed scopes.

Best Practices for Creating Scopes

For each discovery agent, review the assigned range scopes 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 scope that has a high probability of immediate success on the nodes in that scope to speed up discovery.
- When you apply an authentication profile to a scope, make sure that most, if not all, devices defined by that scope 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 scope with no authentication profiles assigned to the scope.
- If you use SNMP for a device that accepts only a unique SNMP community string, create a **Single** type range scope and specify the device's IP address. Assign the corresponding authentication profile to the range scope.
- When using SNMP, to avoid unnecessary authentication traps/alerts, assign only one SNMP authentication credential per discovery range.

Create Scopes

Both range and seed scopes are defined on the Scopes tab of the discovery wizard.

To create range scopes:

1. Click **New range scope** in the left pane of the Scopes tab.
2. Enter a name for the range scope.
3. In the Range Scope definition section, specify the area(s) of your network where you want to perform discovery.

    - **Mask** - Defines 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, anycast, multicast, link-local, and loopback addresses are *not* supported.

    You can also click the Add multiple IPs icon ( ) above the Range Scope 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.

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

To create seed scopes:

1. Click **New seed scope** in the left pane of the **Scopes** tab.
2. Enter a name for the seed scope.

3. In the Seed Scope definition section, enter the IP address of your seed device.
4. (Optional) If you want to limit seed device discovery to a previously defined range scope, clear the **LAN** check box.

   Note: If the LAN check box is cleared and no range scope is defined, only the seed device and no additional devices are discovered.
5. Click **New seed device IP** to add another seed device.

In the Credentials section, you can assign authentication profiles to the selected range. By default, all of the authentication profiles are selected. Seed scopes require at least one SNMP credential. If the LAN checkbox is selected, you should assign the authentication profiles that are applicable to all devices in the local subnets covered by the seed device.

When you have finished defining scopes, 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.

Run File-based Import

In many IT environments, device and host configuration information is maintained in a configuration management database (CMDB). If you have this data, you may prefer to import your device information and SNMP authentication profiles through file-based import. With this method, you import an XML file that contains your device and profile data. Scanning the IT environment is not necessary.

This method offers several benefits:

- Speed. Automated discovery can take several hours or longer an enterprise deployment. An XML file can be imported in minutes.
- More control over your inventory. File based import helps ensure that your inventory includes all the devices you want to monitor, and no others. Automated discovery could potentially add devices that you are not interested in (such as printers or personal computers) or fail to include devices that are temporarily non-responsive.
- Fewer security alerts.

The **cm_data_import** probe imports device and SNMP authentication data. If a system is discovered by automated discovery and is also included in file-based import, the information from file-based import takes precedence for most properties, but depends on device reconciliation rules and heuristics. Device reconciliation is the process of aggregating data from multiple views of the same device and resolving it into a single device view.

Note: The XML schema was updated for cm_data_import version 7.6 to support the import of SNMP authentication profiles. While the probe will import XML files using the old schema, we recommend you migrate to the new schema.

You can launch file- based import by navigating to the file in USM or by placing the file in a directory where cm_data_import will find it.

Method 1: Navigate to the XML File

1. Prepare the XML file. See XML File Schema (see page 27) for details.
2. In the USM navigation tree, expand the **Discovery** node.
3. Hover over the **External** node and click the import icon (⇪), or click the **External** node and choose **Actions > Discovery Import**.
4. Navigate to the XML file and click **OK**. Cm_data_import publishes the data to the message bus, and discovery_server imports it.

Method 2: Automatic file import

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

1. Copy the prepared XML file to *<UIM_install_directory>*\probes\service\cm_data_import\import on the system that hosts the cm_data_import probe (typically the primary hub).
2. Cm_data_import scans the import directory at regular intervals (default is 60 seconds).
3. When the probe finds the XML file you copied to the directory, it publishes the data to the discovery_server to be imported.

The result of both methods:

- The XML file is placed in a time-stamped subfolder in *<UIM_install_directory>*\probes\service\cm_data_import\processed.
- The activity of the process is logged.
- The probe publishes the imported information to the discovery_server.
- Processing by discovery_server can take several minutes or more to finish. Once complete, the devices and/or authentication profiles are stored in the UIM database.
- Imported devices are listed in the **Inventory** node in USM.
- Authentication profiles are viewable in the Discovery Wizard in USM.
- When importing has completed, you can deploy and/or configure probes to monitor the imported devices.

XML File Schema

An XML file can import SNMP authentication profiles, device information, or both. This section describes how to create an XML file for use with file-based discovery.

The XML schema was updated for cm_data_import version 7.6 to support the import of SNMP authentication profiles. While the probe will import XML files using the old schema, we recommend you migrate to the new schema. Refer to:

- *<UIM_install_folder>*/probes/service/cm_data_import/**schema** – new XML schema definition file and example XML files
- *<UIM_install_folder>*/probes/service/cm_data_import/**schema_old_201211** – previous schema and example files for reference

Note the following:

- The schema allows you to define the following sub-sections:
  - Devices
  - SmpV1Profiles
  - SnmpV2Profiles (for SNMP V2c profiles)
  - SnmpV3Profiles
- The definition for V1 and V2 SNMP profiles is the same, but they are placed in the XML file in different sections in order to identify which are V1 and which are V2.
- Properties in the schema that contain **minOccurs="1"** are required.
- For properties that refer to openenumerations, go to *<UIM_install_folder>*\probes\service\cm_data_import\schema and open either **usm-openenums.xml** or **cm-data-import-openenums.xml** to view the defined values. Although it is not strictly required, we strongly recommended you use values defined by the open enumerations.
- In addition to the sub-sections listed above, the following top-level properties are available in the CmData section:

Property	Description
DefaultOrigin *Optional*	Top-level property that specifies the origin to be assigned to any device that does not have a specified origin.
PasswordsEncrypted *Optional*	Top-level property that specifies whether the passwords in SNMP profiles are encrypted. If the value is unspecified or false and the XML file includes passwords, the passwords will be encrypted and *PasswordsEncrypted* will be set to true in the XML file that gets saved in the processed directory. The encryption method is UIM-defined.

For sub-section examples and property descriptions, refer to:

Device Examples and Properties (see page 28)
SNMP V1 and V2c Profile Examples and Properties (see page 33)
SNMP V3 Profile Examples and Properties (see page 34)
Example File with Devices and Profiles (see page 36)

Device Examples and Properties

A minimal device XML file must include these properties for each host or device:

- At least one property that enables the device to be correlated: *PrimaryIPV4Address*, *PrimaryIPV6Address*, *PrimaryDnsName*, *PrimaryMacAddress* or *VirtualID*.
- *Origin* or *DefaultOrigin*. This value identifies the hub from which an entity's QoS messages originate. Setting the origin correctly is important; see the table below for details.

The following example shows how to create a *Devices* section that imports one device with IP address **1.2.3.4** and origin **MyOrigin**. A complete XML file is shown in Example File with Devices and Profiles.

```
- <Devices>
  - <Device>
      <PrimaryIPV4Address>1.2.3.4</PrimaryIPV4Address>
      <Origin>myOrigin</Origin>
    </Device>
  </Devices>
```

Additional optional properties can be included, as shown in the example below. You can find example files in <*UIM_install_folder*>\probes\service\cm_data_import\schema on the system that hosts cm_data_import (typically the primary hub).

```
 - <Devices>
  - <Device>
      <ElementUUID>550e8400-e29b-41d4-a716-446655440000</ElementUUID>
      <Origin>myOrigin</Origin>
      <Label>myComputer</Label>
      <Description>myComputerIsFast</Description>
      <PrimaryDnsName>myComputer.myCompany.com</PrimaryDnsName>
      <OtherDnsNames>name2,name3</OtherDnsNames>
      <PrimaryIPV4Address>1.2.3.4</PrimaryIPV4Address>
      <PrimaryIPV6Address>fe80::223:ebff:fe06:9d40</PrimaryIPV6Address>
       <OtherIPAddresses>2.2.2.2,3.3.3.3</OtherIPAddresses>
      <PrimaryMacAddress>F0-4D-A2-25-5B-7A</PrimaryMacAddress>
      <OtherMacAddresses>22-22-22-22-22-22,33-33-33-33-33-33</OtherMacAddress
es>
      <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>ComputerSystem</PrimaryDeviceRole>
      <PrimarySoftwareRole>DatabaseServer</PrimarySoftwareRole>
      <DBServerType>MSSQLServer</DBServerType>
      <WmiAuthId>3</WmiAuthId>
      <ShellAuthId>5</ShellAuthId>
      <SnmpAuthId>7</SnmpAuthId>
      <AppServerType>Unknown</AppServerType>
      <VirtualizationEnvironment>Unknown</VirtualizationEnvironment>
      <VirtualID>550e8400-e29b-41d4-a716-446655440000</VirtualID>
      <MonitorFrom>1.2.33.44</MonitorFrom>
    </Device>
  </Devices>
```

The following table describes the XML properties. Note the following:

- For properties that refer to open enumerations, navigate to <*UIM_install_folder*>\probes\service\cm_data_import\schema and open either **usm-openenums.xml** or **cm-data-import-openenums.xml** to view the defined values. Although it is not strictly required, we strongly recommended you use values defined by the open enumerations.
- To deploy a robot to an imported system using USM and the Automated Deployment Engine (ADE), some additional properties beyond IP address and origin are required. These are noted in the table below.
- Optional values allow you to add detail to your inventory.

Property	Description
ElementUUID *Recommended; automatically created if not specified*	The universally unique identifier for the device. The UUID must be in standard string representation (for example: 550e8400-e29b-41d4-a716-446655440000). If the UUID is not specified, one is automatically generated for the element. If an existing element is found with a matching UUID, the existing element is updated. Otherwise, a new element is inserted and associated with the UUID. To avoid inserting duplicate records, the same UUID per element should be maintained and used on updates. The XML file saved in the processed directory includes any automatically generated UUIDs for later reuse.
SnmpProfileUUID *Optional*	The universally unique identifier of the authentication profile to use for SNMP access. The specified profile can either be defined in this XML file or have been previously imported.
Origin *Required if DefaultOrigin is not specified*	Identifies the hub from which QoS messages originate. The default origin is the name of the primary hub, but this can be overridden at the hub or robot (controller) in order to separate data in a multi-tenancy environment. To ensure that QoS probe data is correlated to the discovered device, the origin specified here should match the origin you intend to use. You can avoid specifying this value for each device by specifying the top-level *DefaultOrigin* property.
Label *Optional*	A short description or caption.
Description *Optional*	Text description of the device.
PrimaryDnsName *Optional*	The entity's Domain System Name, which may be used for correlation.
OtherDnsName *Optional*	If an entity has multiple DNS names, this property captures those names. (*PrimaryDnsName* is used for correlation.) Multiple names must be are comma-separated.
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).

Property	Description
OtherIPAddresses *Optional*	If an entity has multiple IP addresses, this property captures those addresses for correlation and identity. Multiple addresses must be are comma-separated. IPv4 or IPv6 values may be specified. Addresses should be formatted following the regex patterns defined by usm-core:IPV4AddressFormat or usm-core:IPV6AddressFormat.
PrimaryMacAddress *Optional*	A MAC address that can be used for correlation and identity. The address is expressed as 6 groups of 2 hex digits (uppercase only), separated by dashes.
OtherMacAddress *Optional*	If an entity has multiple MAC addresses, this property captures those addresses. The *PrimaryMacAddress* property is used for correlation. Multiple addresses must be comma-separated and 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 *Optional*	OS version details.
ProcessorType *Required by ADE for robot deployment*	Processor environment/type (such as "x86") as defined by the open enumeration *ProcessorEnvironmentEnum.*
Vendor *Optional*	The hardware vendor/manufacturer's name, as defined by the open enumeration *VendorEnum*.
Model *Optional*	The hardware model name/number.
PhysSerialNumber *Optional*	ID string assigned by the hardware manufacturer and attached to the component. Enter the number directly from the manufacturer's tag on the component (which may be an RFID tag), or read the value from the *entPhysicalSerialNum* field of SNMP's Entity-MIB. Virtual entities do NOT have serial numbers.
PrimaryDeviceRole *Optional*	The device role as defined by the open enumeration *DeviceRoleEnum*.
PrimarySoftwareRol e *Optional*	The software role as defined by the open enumeration *SoftwareRoleEnum*.
DBServerType *Optional*	The type of database server of which this is an instance, defined by the open enumeration *DBServerTypeEnum*.

Property	Description
WmiAuthId *ADE requires WmiAuthId or ShellAuthID for robot deployment*	Authentication profile ID to use for WMI access. This number is generated by (and displayed in) the Discovery Wizard when you create a WMI authentication profile.
ShellAuthId *ADE requires WmiAuthId or ShellAuthID for robot deployment*	Authentication profile ID to use for SSH or telnet access. Authentication profile ID to use for WMI access. This number is generated by (and displayed in) the Discovery Wizard when you create a shell authentication profile.
SnmpAuthId *Optional*	Authentication profile ID to use for WMI access. This number is generated by (and displayed in) the Discovery Wizard when you create an SNMP authentication profile.
AppServerType *Optional*	The type of application server, as defined by the open enumeration *AppServerTypeEnum*.
VirtualizationEnvironment *Optional*	Value that specifies the virtualization environment (hypervisor manager) of a hypervisor or virtual system. Values are defined in the open enumeration *VirtualizationTypeEnum*.
VirtualID *Optional*	Identifier for a VirtualSystem assigned by the virtualization solution (such as VMware or Microsoft Hyper-V).
MonitorFrom *Optional*	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 CA Unified Infrastructure Management address (/UIMDomain/HubName/RobotName). A 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.

SNMP V1 and V2c Profile Examples and Properties

The following example shows how to create *SnmpV1Profiles* and *SnmpV2Profiles* sections. A complete XML file is shown in Example File with Devices and Profiles.

```
- <SnmpV1Profiles>
    - <SnmpV1Profile>

<SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440001</SnmpProfileUUID>
        <Description>SnmpV1Profile</Description>
        <GetCommunityString>public</GetCommunityString>
    </SnmpV1Profile>
  </SnmpV1Profiles>
  - <SnmpV2Profiles>
    - <SnmpV2Profile>

<SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440002</SnmpProfileUUID>
        <Description>SnmpV2Profile</Description>
        <GetCommunityString>public</GetCommunityString>
    </SnmpV2Profile>
  </SnmpV2Profiles>
```

SNMP V1 and V2 authentication profile import uses the following properties.

Property	Description
SnmpProfileUUID *Recommended*	The universally unique identifier of the authentication profile to use for SNMP access. The specified profile can either be defined in this XML file or have been previously imported.
Description *Optional*	Description for the profile.
Port *Optional*	The SNMP port to use. If not specified, port 161 (default) is used.
GetCommunityString *Required*	Community string to be used for *Get*, *GetNext* and *GetBulk* SNMP requests.

SNMP V3 Profile Examples and Properties

The following example shows how to create the *SnmpV3Profiles* section. A complete XML file is shown in Example File with Devices and Profiles.

```
- <SnmpV3Profiles>
    <SnmpV3Profile>
      <SnmpProfileUUID>6f0bf0b7-89cf-416c-a4c3-c70ab84d3483</SnmpProfileUUID>
      <Description>SnmpV3Profile NoAuthNoPriv</Description>
      CA Portal161</Port>
      <UserName>NoAuthNoPrivUser</UserName>
      <AuthenticationProtocol>None</AuthenticationProtocol>
      <PrivacyProtocol>None</PrivacyProtocol>
    </SnmpV3Profile>
    <SnmpV3Profile>
      <SnmpProfileUUID>543af64d-50d0-46b5-a81e-4bef93005259</SnmpProfileUUID>
      <Description>SnmpV3Profile AuthPriv</Description>
      CA Portal161</Port>
      <UserName>AuthPrivUser</UserName>
      <AuthenticationProtocol>MD5</AuthenticationProtocol>
      <AuthenticationKey>authKey</AuthenticationKey>
      <PrivacyProtocol>AES</PrivacyProtocol>
      <PrivacyKey>privKey</PrivacyKey>
    </SnmpV3Profile>
   </SnmpV3Profiles>
```

SNMP V3 authentication profile import uses the following properties.

Property	Description
SnmpProfileUUID *Recommended*	The universally unique identifier of the authentication profile to use for SNMP access. The specified profile can either be defined in this XML file or have been previously imported. If the UUID is not specified, one is automatically generated for the element. If an existing element is found with a matching UUID, the existing element is updated. Otherwise, a new element is inserted and associated with the UUID. To avoid inserting duplicate records, the same UUID per element should be maintained and used on updates. The XML file saved in the processed directory includes any automatically generated UUIDs for later reuse.
Description *Optional*	Description for the profile.
Port *Optional*	The SNMP port to use. If not specified, port 161 (default) is used.
UserName *Required*	SNMP user name.

Property	Description
AuthenticationProtocol *Required*	Type of authentication used for messages (if any). The values are defined by *SnmpV3AuthenticationProtocolEnum*.
AuthenticationKey	Specific key used by the *AuthenticationProtocol* for authenticating messages.
PrivacyProtocol *Required*	Type of encryption used for messages (if any). The values are defined by *SnmpV3PrivacyProtocolEnum*.
PrivacyKey	Specific key used by the *PrivacyProtocol* for encrypting and decrypting messages.

Example File with Devices and Profiles

Example files are located in *<UIM_install_folder>*/probes/service/cm_data_import/schema.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
- <CmData xmlns="http://nimsoft.com/2014/05/cm-data-import2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  - <Devices>
    - <Device>
        <PrimaryIPV4Address>10.10.10.1</PrimaryIPV4Address>
        <Origin>origin</Origin>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440001</SnmpProfileUUID>
      </Device>
    - <Device>
        <PrimaryIPV4Address>10.10.10.2</PrimaryIPV4Address>
        <Origin>origin</Origin>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440002</SnmpProfileUUID>
      </Device>
    - <Device>
        <PrimaryIPV4Address>10.10.10.3</PrimaryIPV4Address>
        <Origin>origin</Origin>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440003</SnmpProfileUUID>
      </Device>
     - <Device>
        <PrimaryIPV4Address>10.10.10.4</PrimaryIPV4Address>
        <Origin>origin</Origin>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440004</SnmpProfileUUID>
      </Device>
    - <Device>
        <PrimaryIPV4Address>10.10.10.5</PrimaryIPV4Address>
        <Origin>origin</Origin>
      </Device>
    </Devices>
  - <SnmpV1Profiles>
    - <SnmpV1Profile>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440001</SnmpProfileUUID>
        <Description>SnmpV1Profile</Description>
        <GetCommunityString>public</GetCommunityString>
      </SnmpV1Profile>
    </SnmpV1Profiles>
  - <SnmpV2Profiles>
    - <SnmpV2Profile>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440002</SnmpProfileUUID>
        <Description>SnmpV2Profile</Description>
        <GetCommunityString>public</GetCommunityString>
      </SnmpV2Profile>
    </SnmpV2Profiles>
  - <SnmpV3Profiles>
    - <SnmpV3Profile>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440003</SnmpProfileUUID>
        <Description>SnmpV3Profile NoAuthNoPriv</Description>
        <UserName>NoAuthNoPrivUser</UserName>
        <AuthenticationProtocol>None</AuthenticationProtocol>
        <PrivacyProtocol>None</PrivacyProtocol>
      </SnmpV3Profile>
  - <SnmpV3Profile>
        <SnmpProfileUUID>550e8400-e29b-41d4-a716-446655440004</SnmpProfileUUID>
        <Description>SnmpV3Profile AuthPriv</Description>
        <UserName>AuthPrivUser</UserName>
        <AuthenticationProtocol>MD5</AuthenticationProtocol>
        <AuthenticationKey>authKey</AuthenticationKey>
        <PrivacyProtocol>AES</PrivacyProtocol>
        <PrivacyKey>privKey</PrivacyKey>
      </SnmpV3Profile>
    </SnmpV3Profiles>
  </CmData>
```

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 range scopes 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.

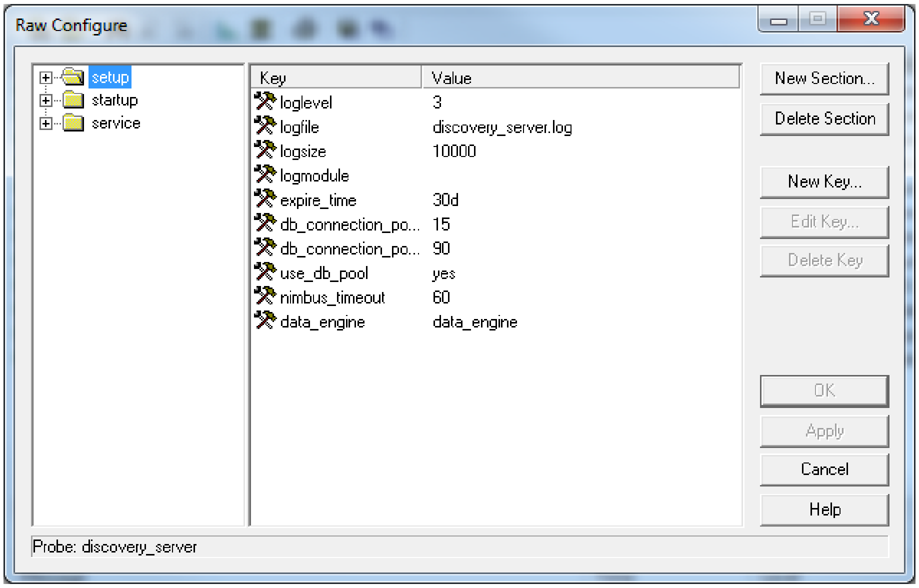

Key	Value
loglevel	3
logfile	discovery_server.log
logsize	10000
logmodule	
expire_time	30d
db_connection_po...	15
db_connection_po...	90
use_db_pool	yes
nimbus_timeout	60
data_engine	data_engine

Probe: discovery_server

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