



Infoblox Discovery Best Practices Guide

Copyright Statements

© 2018, Infoblox Inc.— All rights reserved. The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Infoblox, Inc.

The information in this document is subject to change without notice. Infoblox, Inc. shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

This document is an unpublished work protected by the United States copyright laws and is proprietary to Infoblox, Inc. Disclosure, copying, reproduction, merger, translation, modification, enhancement, or use of this document by anyone other than authorized employees, authorized users, or licensees of Infoblox, Inc. without the prior written consent of Infoblox, Inc. is prohibited.

Trademark Statements

Infoblox, the Infoblox logo, Grid, NIOS, and NetMRI are trademarks or registered trademarks of Infoblox Inc.

All other trademarked names used herein are the properties of their respective owners and are used for identification purposes only.

Company Information

<http://www.infoblox.com/contact/>

Document Updated: January 17, 2019

Contents

1.	Introduction	1
2.	Discovery Overview	2
2.1	Active Polling	2
2.1.1	Path Collection	2
2.1.2	Complete Ping Sweep	2
2.1.3	Smart IPv4 Subnet Ping Sweep	2
3.	Data Collection	4
4.	Best Practices	12
5.	Q&A	16

1. Introduction

This document describes the Infoblox recommended best practices aimed at comprehensive, accurate, and efficient discovery of network devices using the NIOS Network Insight feature. It includes the following:

- [Discovery Overview](#)
- [Data Collection](#)
- [Best Practices](#)
- [Q&A](#)

2. Discovery Overview

Network Insight discovers new devices on network based on the following sources:

- [Active polling](#)
- [Data collection](#)

The next sections discuss these sources in detail.

2.1 Active Polling

Active polling includes the following methods to identify addresses where devices are located:

- [Path collection](#)
- [Complete ping sweep](#)
- [Smart IPv4 subnet ping sweep](#)

2.1.1 Path Collection

Path collector tries to identify potential routers within a network (defined as a network object in Network Insight) by using traceroute to first, second, middle, and last addresses in it. For smaller networks that do not have that many addresses, a traceroute is performed to the manageable addresses in them.

If no routers are found, path collector splits the network in two networks and repeats the process for smaller networks (up to four times). Path collection frequency is not configurable and traceroutes run every 30 minutes.

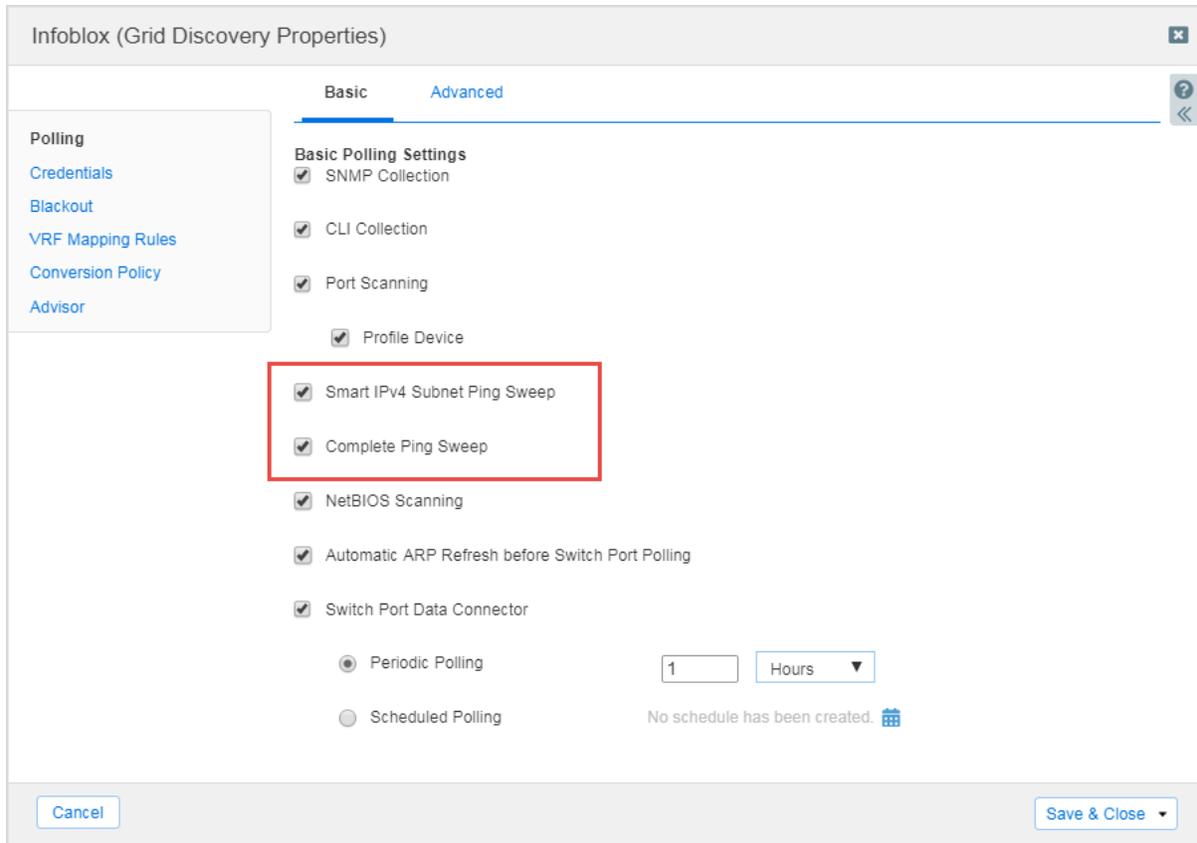
2.1.2 Complete Ping Sweep

Performs brute-force sweep on the network using a range of packets to detect the presence of a system on each IP in the network, using ports that are generally open.

2.1.3 Smart IPv4 Subnet Ping Sweep

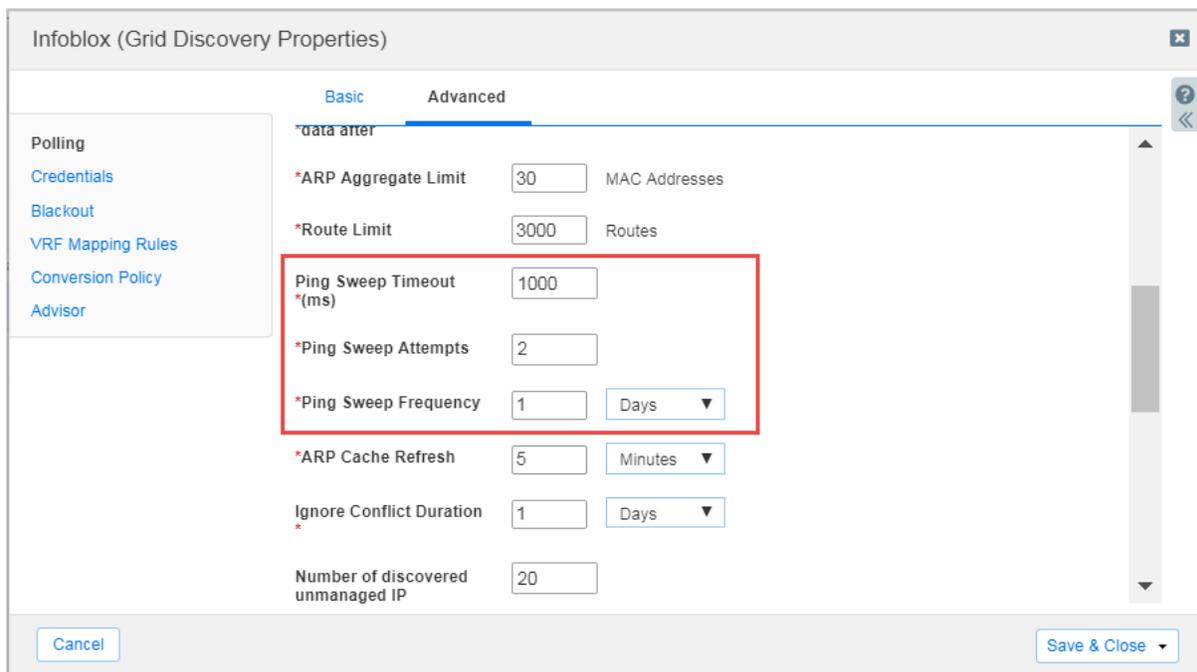
Performs ping sweeps only on networks that are known to exist but no IPs can be found through ARP tables or other means.

Note: Ping sweeps are a tool to help in discovery. It is especially useful for complete discovery of end-host network segments, but most discovery operations take place through ARP tables and routing tables are collected from infrastructure devices. Avoid using ping sweeps on a large number of networks or networks that are too large (more than a /22 in size) as devices discovered through this method may expire from Network Insight's database before they are refreshed by another discovery cycle. Ping sweeps are not used on IPv6 networks because of the dramatically greater scale of network addresses in the IPv6 realm. Smart ping sweeps are not performed on subnets larger than /22.



Complete and Smart Ping Sweeps are enabled

The Ping Sweep Frequency can be configured in the Advanced Grid Discovery Properties and affects both the Smart Ping Sweep and Complete Ping Sweep features.



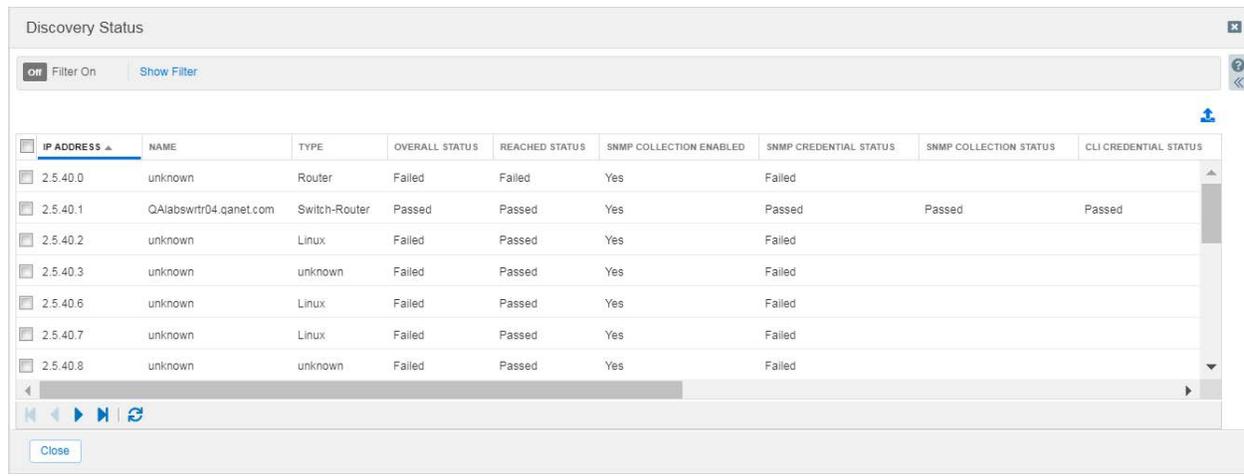
Ping Sweep configuration

3. Data Collection

Information about devices that exist on the network can be retrieved from data collected from devices that are already discovered:

- ARP Table—Discovered from collected ARP data.
- ACI—Discovered fabric nodes from configured APIC controller.
- Call Server—Discovered IP phone and VoIP devices from call server data.
- CDP—Discovered from CDP neighbor data.
- HSRP—Discovered from HSRP data.
- IP Phone—Discovered call server from IP phone.
- LLDP—Discovered from LLDP neighbor data.
- Route Table—Discovered as a route next hop or /32 or /128 route destination.
- VPN Table—Netscreen nsVPNMonTable and nsAddrStatusTable data.
- VRRP—Discovered from VRRP data.
- Wireless AP—Wireless forwarding data where IP address information is included.
- Wireless Controller—Discovered Wireless AP from wireless controller data.

After Network Insight identifies the address by which the device can be found using the above means, it performs the following steps to collect all required information from the device:



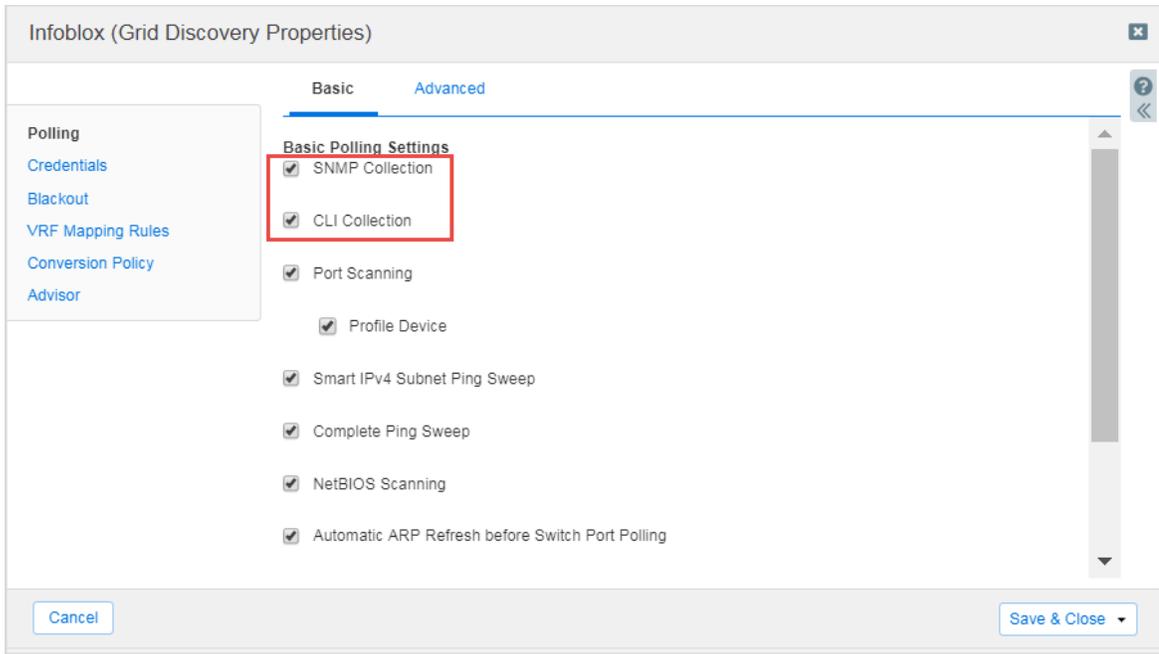
IP ADDRESS	NAME	TYPE	OVERALL STATUS	REACHED STATUS	SNMP COLLECTION ENABLED	SNMP CREDENTIAL STATUS	SNMP COLLECTION STATUS	CLI CREDENTIAL STATUS
2.5.40.0	unknown	Router	Failed	Failed	Yes	Failed		
2.5.40.1	QAAlabswrtr04.qanet.com	Switch-Router	Passed	Passed	Yes	Passed	Passed	Passed
2.5.40.2	unknown	Linux	Failed	Passed	Yes	Failed		
2.5.40.3	unknown	unknown	Failed	Passed	Yes	Failed		
2.5.40.6	unknown	Linux	Failed	Passed	Yes	Failed		
2.5.40.7	unknown	Linux	Failed	Passed	Yes	Failed		
2.5.40.8	unknown	unknown	Failed	Passed	Yes	Failed		

Discovery Status

- **Reachability Check**

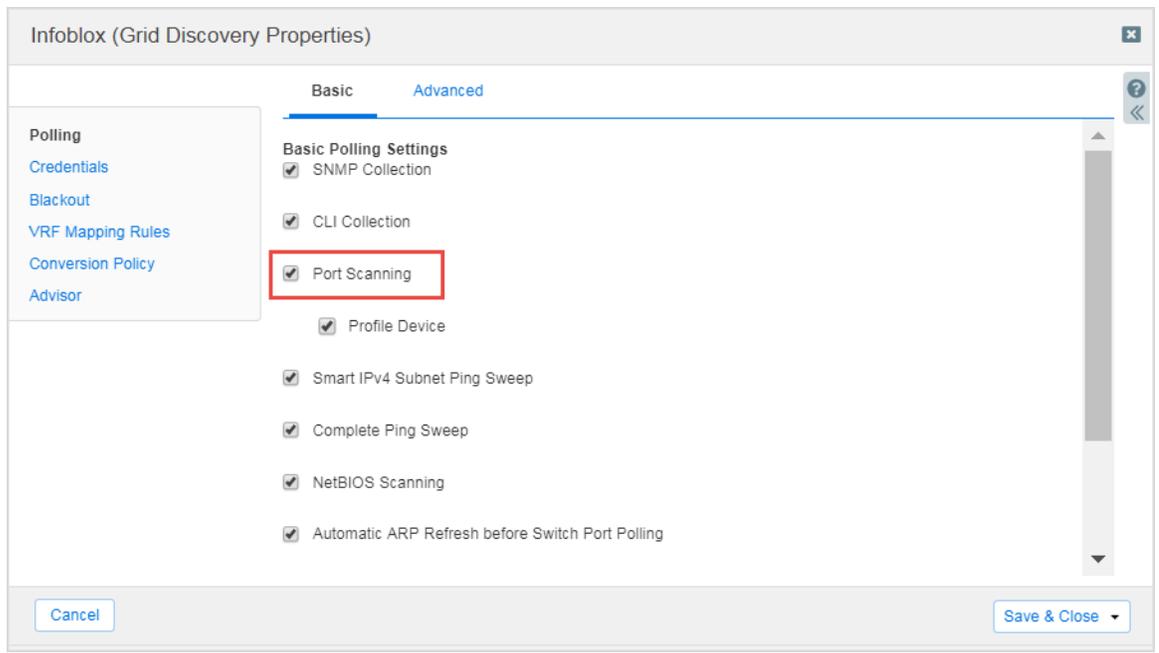
Indicates the reachability of the discovered device (Figure “Discovery Status”). Typically, devices will be reported Passed for Reached Status if they are reachable through SNMP, CLI (Figure “SNMP and CLI collection are enabled”), a path trace through ICMP, or UDP-based path tracing for an IPv6 address. If a value of Failed appears, you will likely see a Last Action reading of Reachable: Failed to Reach.

You may see a Reached Status of Passed and still receive an Overall Status of Failed. This often occurs because either the CLI credentials or SNMP credentials provided for discovering the device do not work, or another problem occurs in some part of the discovery process.

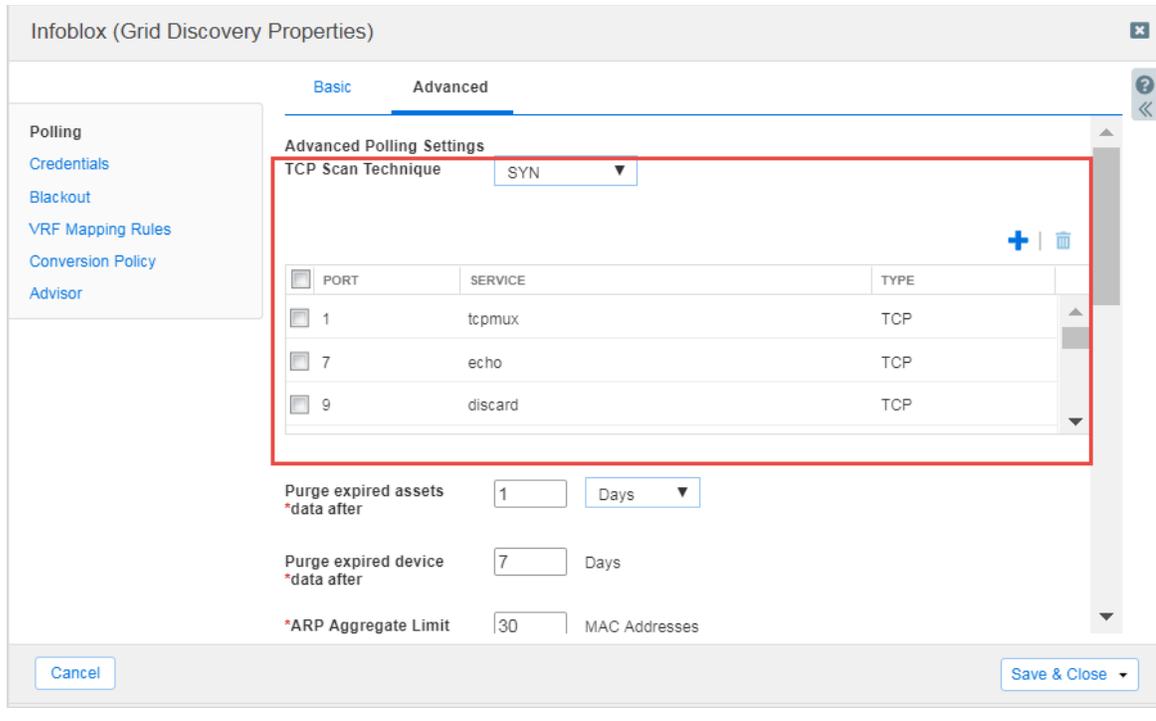


SNMP and CLI collection are enabled

- **Port Scanning** (Figure “Port scanning is enabled”)
If enabled, Network Insight probes the TCP and UDP ports on device to determine whether they are open (listed in **Discovery Properties -> Polling -> Advanced**).



Port scanning is enabled



TCP Scan Technique

NOTE: Advanced SNMP polling settings consist of choosing the TCP Scan Technique (Figure “TCP Scan Technique”), along with a number of specialized settings for Ping Sweeps and other operations.

- **TCP Scan Technique:** Select the TCP technique you want to use for the discovery. The default is SYN.

- **SYN:** Select this to quickly perform scans on thousands of TCP ports per system, never completing connections across any well-known port. SYN packets are sent and the poller waits for a response while continuing to scan other ports. A SYN/ACK response indicates the protocol port is listening while an RST indicates it is not listening. The SYN option presents less impact on the network.

- **CONNECT:** Select this to scan IPv6 networks. Unlike the SYN option, complete connections are attempted on the scanned system and each successive TCP protocol port being scanned.

In the port table, select the checkboxes of the TCP ports you want to discover. You can select all ports by clicking the checkbox in the header.

Optionally, you can click the Add icon and complete the following to add a new port to the list.

- **Port:** Enter the port number you want to add to the list. You must enter a number between 1 and 65535.

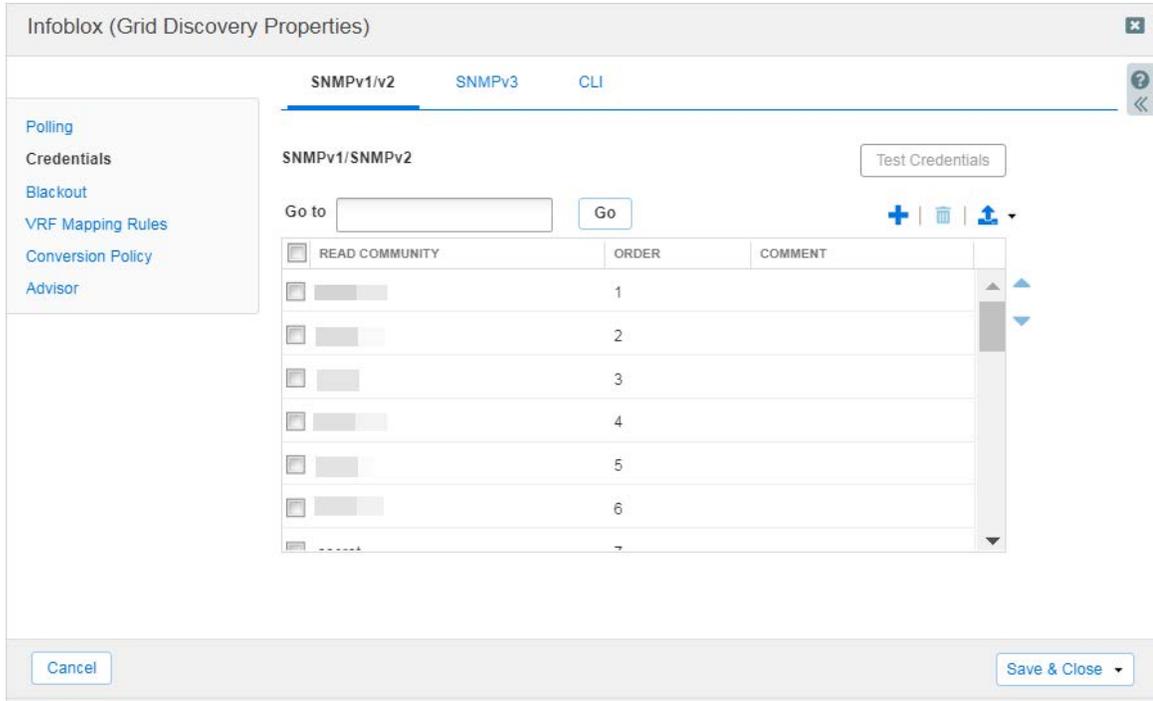
- **Service:** Enter the name of the service.

You can also delete a specific TCP port in the list, or select multiple ports for deletion.

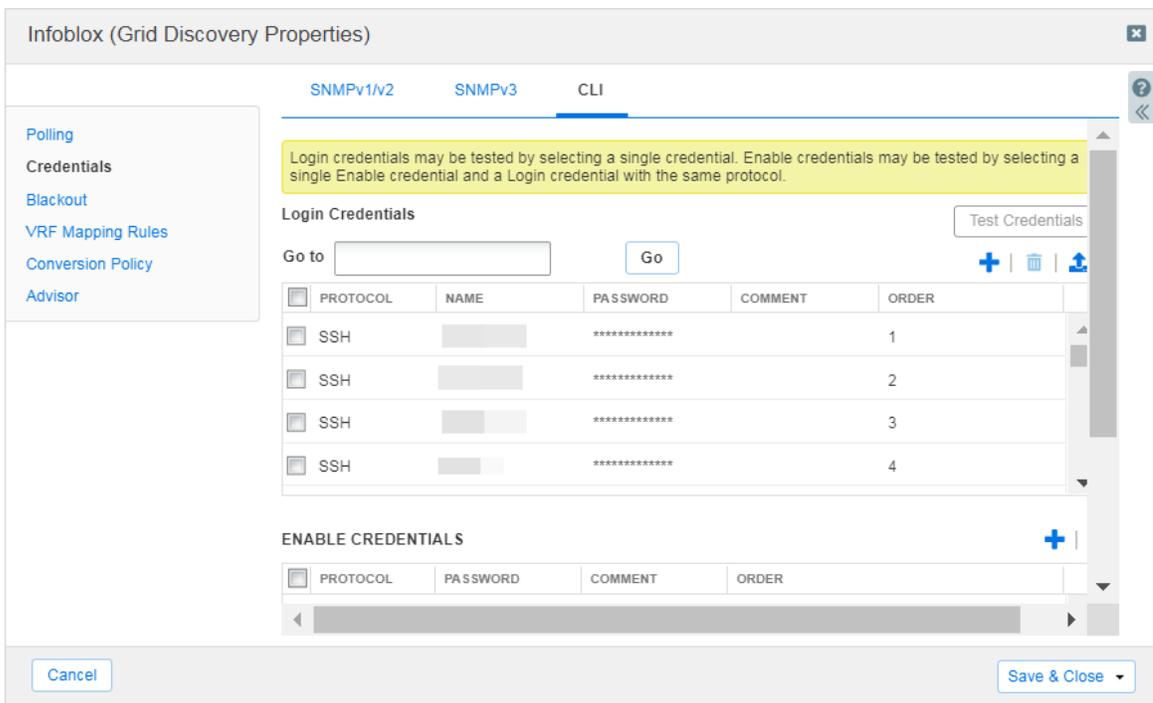
- **Fingerprinting/Profile Device**—If enabled, Network Insight attempts to identify each network device based on the response characteristics of its TCP stack. This information is used to determine the device type. In the absence of SNMP access, fingerprinting is

usually the only way to identify non-network devices. If disabled, devices accessible via SNMP are identified correctly; all other devices are assigned the Unknown device type.

- **Credential guessing**—Network Insight iterates through credentials provided by user to find the right ones for each device (Figures “SNMPv1/v2 credentials”, “CLI credentials”).

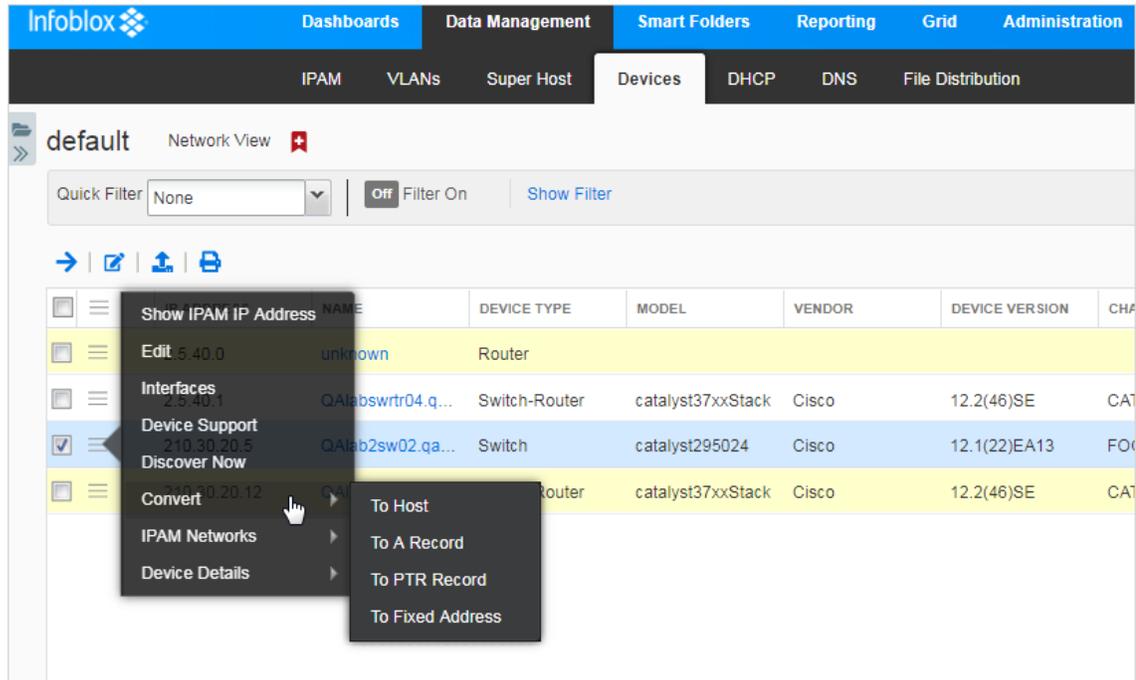


SNMPv1/v2 credentials

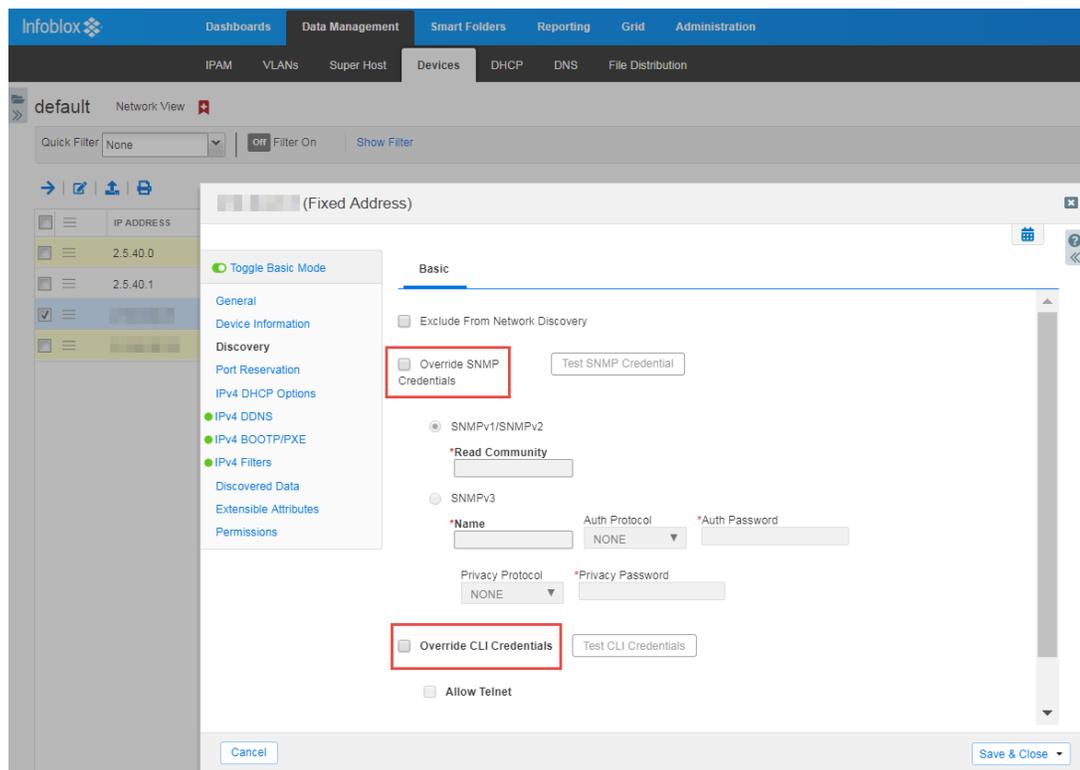


CLI credentials

Credentials can be overridden at the device level. This option is available when the IP address is converted to a managed object (Host, A Record, PTR Record, Fixed Address). Figures “Unmanaged device is converted to managed”, “Overriding SNMP and CLI credentials”.



Unmanaged device is converted to managed



Overriding SNMP and CLI credentials

- **Collection**—Data is collected from devices using SNMP and CLI.
- **NetBIOS Scanning**—Collect the NetBIOS name for endpoint devices in the network.

Clarification on discovery parameters:

- **ARP Aggregate Limit**—Defines the maximum number of ARP records per one MAC address. If there are more, it is considered to be an invalid configuration and all records for this MAC address are discarded.
- **ARP Cache Refresh**—Defines the time period between ARP refreshes by Network Insight across all switch ports. ARP Cache refreshes are used to improve the accuracy of end-device discovery.
- **Disable discovery for networks not in IPAM**—Prevents Network Insight from executing discovery on any infrastructure networks that are not present in the Infoblox IPAM, and prevents it from creating unmanaged networks found in devices.

Check the NIOS Admin Guide for details on other discovery settings.

When multiple IP addresses exist on a device, Network Insight picks one of those IP addresses to be the management IP address of the device. Network Insight only chooses IP addresses that are in defined networks. Network Insight picks the management IP address based on the following priority:

1. User-specified management IP address (available since 8.3).
2. IP address matching any IP address of seed router configured during setup.
3. IP address of an interface with “softwareLoopback” port type and with the lowest ifIndex value (it is an interface property collected by SNMP; check SNMP documentation for details). If multiple IPs are on the same ifIndex, the lowest numerical value IP address is chosen.
4. Any interface named "mgmt" or "management" with the lowest ifIndex value. If multiple IPs are on the same ifIndex, the lowest numerical value IP address is chosen.
5. IP address of an interface with “ethernet-csmacd” port type and with the lowest ifIndex value. If multiple IPs are on the same ifIndex, the lowest numerical value IP address is chosen.
6. Interface with the lowest ifIndex value. If multiple IPs are on the same ifIndex, the lowest numerical value IP address is chosen.

The same device can be discovered from different sources and with different addresses. To ensure there are no multiple records for the same device, the discovery engine performs a deduplication. A specific property called `snmpEngineID` is generated for each device. The `snmpEngineID` is an MD5 sum of the following SNMP data (check SNMP documentation for details on each item):

- `sysDescr`
- `sysLocation`
- `sysName`
- `sysObjectID`
- `sysServices`
- `ipForwarding`
- `ifAddrCheckSum`—MD5 sum of interface information (interface index, address, and network mask).

If two devices with the same `snmpEngineID` are discovered, the last one is removed as duplicate.

If a device is an end host, e.g. a Linux workstation, but with several network adapters, each address is discovered as a separate device. Network Insight does not try to deduplicate end hosts.

The table below presents information collected, frequency, and source:

Information Type	Frequency	Available Sources
Device Chassis Information	1 hour	SNMP/CLI
Device Environment Information: fans speed, temperature and power supply	1 hour	SNMP
Information of CPU, memory and disk usage on device	10 min	SNMP
Information about interfaces, their configuration, addresses, and performance	1.5 hours	SNMP/CLI
Forwarding Tables	Configurable in Switch Port Data Collection in Grid Discovery Settings	SNMP/CLI
Switch Port Information		SNMP
VLAN Information		SNMP/CLI
ARP Tables		SNMP/CLI
Route Table	2 hours	SNMP/CLI
Routing Protocols Information	2 hours	SNMP
Routing Counter Information	2 hours	SNMP
VRF Information	2 hours	CLI only for Juniper, Cisco, Arista SNMP only for Fortinet
Firewall statistics	1 hour	SNMP/CLI
Wireless Information	1 hour	SNMP
Cisco ACI fabric information and each data category from this table which applies	1 hour	HTTP (is controlled by the SNMP Collection setting)
Port Scan	24 hours	nmap

Note: These are the most common data types with the most common intervals for reference. Devices for some vendors may have a different polling frequency.

4. Best Practices

Do configuration on the Grid level first and override only when necessary.

It is simpler in management and more efficient in resources.

Perform discovery in phases for simplicity.

In order to catch issues early, start with a small set of devices and ensure credentials, firewall filters, etc. are all correct.

Configure settings and credentials before starting discovery.

This ensures that they are applied and tested beforehand.

Avoid large networks such as /8 or /16.

Larger networks will cause longer discovery times when using active polling.

This restriction is not about network containers but about *networks*. It is quite obvious why we should not create too large networks: Network Insight will need to discover a huge amount of IP addresses that can appear just as empty records in IPAM. Better way is to create several small networks that most likely contain important devices. It is ok to enable discovery for large containers, because Network Insight does not discover the whole IP range of a container, it discovers only networks that are within the container. These networks will inherit the discovery settings of the container.

Avoid more than 1,000 networks or exclusions.

Each time an address is found, the discovery engine checks to see if it is in an included network, so the more networks, the longer the discovery cycle and the more impact on performance. It is better to merge several networks into larger ones. If they become too large, disable ping sweep and add relevant seed routers.

It is better to have a large network and seed routers than thousands of small networks.

With seed routers, discovery engine will not have to rely on active polling to discover most of the devices so a large network becomes a less of an issue.

Use seed routers.

Seed routers are discovered quickly and after routing data is collected from them, the discovery engine finds new devices. It is a much quicker way than using active polling.

IMPORTANT: Seed routers are a good way to speed up the discovery, but adding too many seed routers is not a good idea (a seed router is counted towards the included networks and exclusion count).

BE CAREFUL with the use of seed from DHCP routers.

Pros:

1. Use the default gateways for associated DHCP ranges and networks as seed routers to more quickly discover and catalogue all devices (such as endpoint hosts, printers and other devices).
2. All such default gateways are automatically leveraged by discovery, and no further configuration is necessary unless you wish to exclude a device from usage.

Cons:

1. If you enable using DHCP seed routers, it will instruct Network Insight to use **all** of them as seed routers, including those that do not belong to the networks you would like to manage with Network Insight. As they are still discovered and processed, it will add some extra load on the system.

Select all discovery services (SNMP, CLI, etc.).

All the settings are important for discovery and having one disabled may cause incomplete data. Disable one only if you are sure that this data will not be needed.

SNMP Collection is disabled, for example, when you use NIOS for offline assessments. By disabling SNMP Collection before removing the NIOS Poller member performing discovery from the Grid, data can be examined later without any data expiring.

You also may want to disable both SNMP and CLI if the network contains only end hosts.

You can see what kind of information is collected via SNMP and CLI in the polling frequency table in the [Data Collection](#) section.

Note: SNMP is used to get basic information about a device (vendor, model, OS, etc.) that is required to do proper CLI polling, so if you disable SNMP before this information is collected, CLI collection will not happen.

Consider discovery blackouts.

If you want to release some resources within the network, you can stop discovery for a while by setting blackout periods. That means you can establish times when Network Insight does not talk to devices or networks for discovery. You can disable discovery Grid wide or, if you know some traffic- or latency-sensitive networks in IPAM, it will be useful to set blackout periods specifically for them.

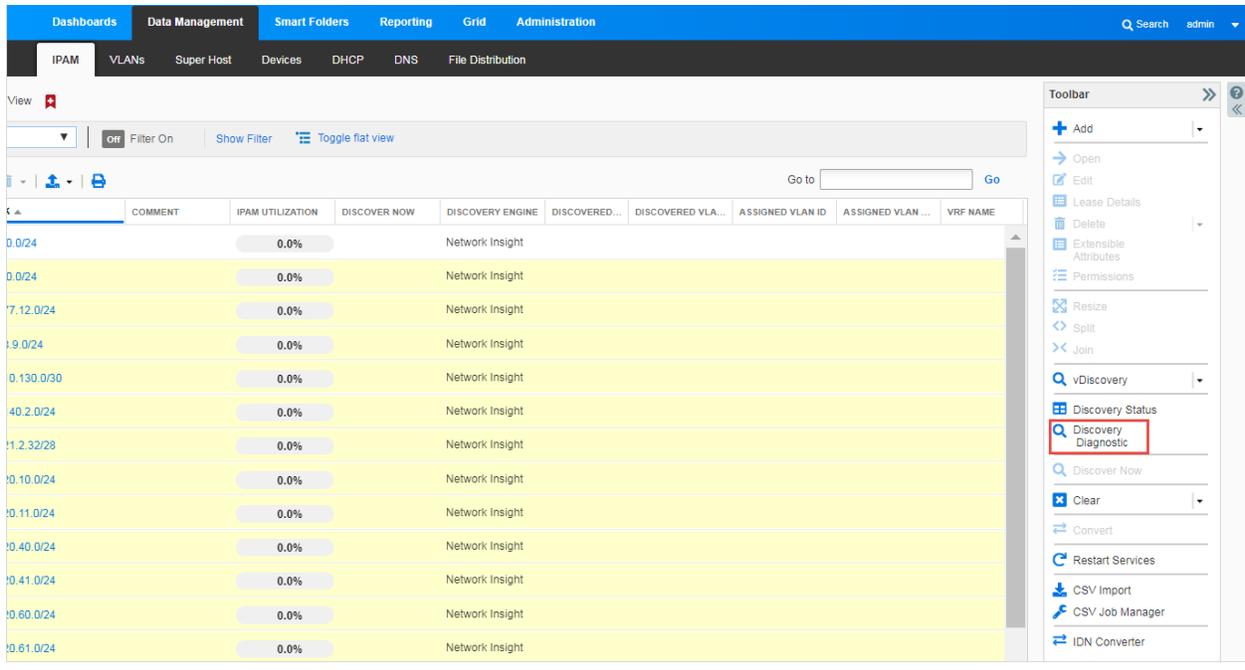
Discovery tasks may already be running when a blackout period takes effect. Current tasks are not interrupted and will complete within their time. Network Insight does not activate new discovery tasks during the blackout period, however.

Avoid polling same devices by different probes.

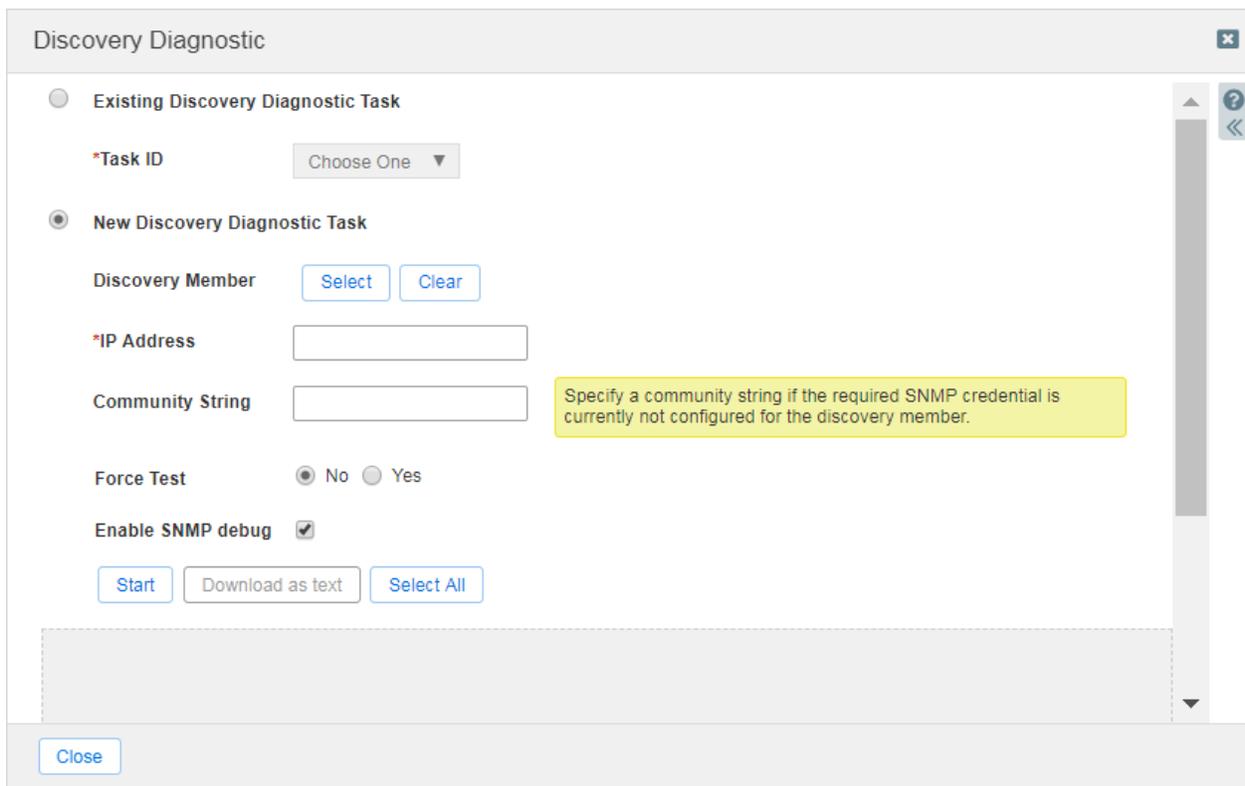
The consolidator will take care of duplicates, but as each probe does device polling and data processing separately, it will increase the load on device and waste processing resources on probes.

Use Discovery Diagnostic to troubleshoot discovery issues.

The following figures show how to access the Discovery Diagnostics feature and which options are available.



Discovery Diagnostic in the toolbar



Discovery Diagnostic dialog

Map one VRF to one Network View.

This will ensure that IP addresses are not overlapping.

If you have different networks with overlapping address space, ensure that they are handled by different network views.

Ensure you use a proper unit to discover your network.

If the Network Insight member serves as a probe or single consolidator:

Appliance Model	Maximum IP Addresses or Interfaces	Maximum Switches & Routers
ND-800	11K	300
ND-805	15K	400
ND-1400	80K	3000
ND-1405	110K	4000
ND-2200	180K	6000
ND-2205	250K	8000
ND-4000	700K	15000
ND-4005	900K	16500

Caution: Network Insight stops SNMP collection for new devices when discovered switches/routers exceed 120% of the maximum.

If the Network Insight member serves as a consolidator with multiple probes:

Appliance Model	Maximum IP Addresses or Interfaces
ND-1400	160K
ND-1405	210K
ND-2200	360K
ND-2205	500K
ND-4000	1.6M and no more than 36,000 devices
ND-4005	2.6M and no more than 40,000 devices

Note: Network Insight collects information about interfaces of infrastructure devices even if there is nothing connected to them. This is why the limitation is set not only on maximum IP addresses but also on interfaces.

5. Q&A

Q. How long will discovery take? Is there an indication when the first pass is complete?

A. Discovery is an iterative process when new devices are found based on active polling or data collected from already discovered devices. Discovery is constantly running (unless disabled by user) in order to keep the discovered data relevant, and it is not possible to predict or tell when all devices are discovered. As in the beginning of discovery there is no collected data and active polling is usually not finished yet, the first iteration either discovers no devices or “discovers” only those devices that user explicitly added as seed routers.

You can monitor the whole discovery process in **Administration -> Logs -> Syslog**. Many interesting logs from various discovery services can be found on discovery members. On the Grid Master, you can monitor messages about new discovered (unmanaged) networks and newly found devices.

Q. How to avoid polling end host devices? In other words, I want to discover the end hosts but I don't want them polled directly. Related to this, in Network Insight if I want to discover end hosts, do I need to enable discovery on a subnet of end hosts when I already have discovery enabled on the management subnet that manages the routers/switches the end hosts are attached to?

A. Discovery should be enabled for both subnet of end hosts and management subnet. To avoid polling of end hosts subnet, SNMP and CLI Collection need to be disabled in network properties; end hosts will be found and identified even without polling.

In the case of when end hosts and management devices exist in the same subnet, it is not possible to avoid polling the end hosts and the only opportunity here is to split this network into subnetworks to override its discovery properties.

Q. How do I poll only specific network devices? I don't want to discover by network.

A. You can add necessary devices as seed routers. To do that, you should open Member Discovery Properties for appropriate member of the type Probe and add device IP address in the Seed tab.

You can also specify such devices as /32 networks.

Note: If a device is added as a seed router, its type is considered as Router by default. After successful SNMP collection, fingerprinting, etc. Network Insight will change the device type to the correct one. However, if the device is not reachable or its type cannot be identified, its type may remain as Router.

Q. How do I automate discovery of new network devices? In other words, a new device is on the network and I want it discovered ASAP.

A. You can increase discovery priority of particular device by adding it as Fixed Address or Host. From the NIOS Admin Guide: “*You can optionally **Enable Immediate Discovery**. If you choose not to perform immediate discovery, but do **Enable Discovery**, the new network or other object is discovered at a normal time determined by Network Insight ... Devices matching IP addresses selected for immediate discovery are given one-time priority over other discovered devices, for data collection and counting toward any device found matching the license limits.*” You can also convert the corresponding IP address from the IPAM tab to Fixed Address or Host.

There is another way to discover a device immediately. In IPAM tab, you can create a network with prefix /32 with IP address of your device and click **Discover now** on it. This network will be treated as direct path to the device and it will be discovered very soon.

Q. How do I stop automatic creation of IPAM objects via Network Insight?

A. You can check an option in **Grid Discovery Properties -> Advanced Polling Settings -> Disable discovery for networks not in IPAM**. If new discovered networks are not contained in any existing network, new “unmanaged” objects (networks that were collected from devices’ routing tables) will not be created in IPAM. Otherwise, yellow networks will appear even if you have deleted them before.

If you want to disable the discovery of a network, you can add it in IPAM or click **Edit** for an existing one and clear the **Enable Discovery** check box. The network will never be discovered unless you change this setting again.

You can also explicitly exclude specific IPs or IP ranges from discovery in **Network Editor -> Discovery Exclusions**. Another way to exclude specific IP addresses is to select them in the IP map or IP list and click **Exclusion -> Enable Exclusion**. Discovery will never take place on these IPs unless you specifically change their exclusion setting.

Q. Are both SNMP and CLI required for interrogating devices? If so, what is the order of their operation?

A. SNMP is used at least to collect basic information about the device (name, model, vendor, etc.). This information is required to enable CLI collection on the device. To find which of SNMP/CLI is used to collect specific data, see the table in the [Data Collection](#) section.

Q. How is the network discovery database cleaned up? Like if I want to wipe all data and start again.

A. It is possible and necessary to do from Infoblox Admin CLI on both consolidator and probe members using the following command: `reset database net-automation`. It will take some time on these members for the discovery service to restart again and respond to the Grid Master. No additional actions like stopping services are required.

Q. How do I remove unwanted devices and device information?

A. We have no special button in the UI to remove specific devices. If you do not want to see some device in the Devices tab, you can clear discovered data from the IP addresses of its interfaces and add this IPs to the exclusion in the IPAM tab.

Note: No need to exclude interfaces with gray background, only managed and unmanaged (white and yellow) interfaces are discovered by Network Insight.

To remove an unwanted device, do one of the following:

1. Delete the network that contains this device in IPAM.
2. Disable discovery of the network that contains this device in IPAM.
3. Add the device IP address in Discovery Exclusions of the corresponding network in IPAM.
4. Add the IP range that contains the device IP address in Discovery Exclusions of the corresponding network in IPAM.

Q. What is the best way to deploy VRF mapping in our environment?

A. Per best practices listed in this document, map one VRF to one network view to ensure that IP addresses are not overlapping.

Q. How many devices are scanned simultaneously?

A. It is not possible to tell the exact number. Network Insight engine does its best to make a parallel polling of devices with a good balance between keeping the data most up-to-date and polling non-aggressively. For information about frequencies of device polling, see the [Data Collection](#) section.

Q. Once devices are registered, how often does discovery hit the devices?

A. Devices can be hit by different means—by active polling, collection, etc., and it is not possible to tell exact time. Their frequencies are listed in the [Data Collection](#) section.

Q. How often do new scans occur?

A. See the [Data Collection](#) section for details.

Q. Which log files can I consult for the discovery process?

A. Logs are visible in **Administration -> Logs -> Syslog**. In the dropdown menu, you can select a member for which you want to see the logs.

The following types of discovery events can be logged to the syslog on the Grid Master:

1. “Discovery”—Sends notifications about the discovery status. For example, “The grid member is connected to the grid master” or “Discovery Collector Service is working.”
2. “Discovery Conflict”—Sends notifications about conflicts between the DHCP address and the existing IP address.
3. “Discovery Unmanaged”—Sends notifications related to the discovery of unmanaged devices and networks. For example, “New unmanaged devices/networks were found during network discovery process. New unmanaged devices in '10.40.16.0/20' network in 'default' network view.”

Some error messages appear in the syslog. For example, connection errors of vDiscovery (“Communication with vSphere server failed with error”), parsing errors of native discovery (“CSV import error ... Please contact the support team”). Also, some start/stop events during member upgrade are logged: “Stopped the discovery operation. A member has not completed its upgrade”. Network Insight, if enabled, can log discovery process details such as “Some object has been created/deleted” or “Cannot find the network in NIOS for IP ... “. Also, there are some Conversion Policy messages, e.g. “%name% created from the discovered IP address” or “conversion type have been changed since the last run”.

On discovery members, you can see mostly logs from discovery core utilities (PathCollector, netautoctl, DiscoveryEngine, etc.) and kernel messages.

Q. Why doesn't my container have a First Discovered/Last Discovered value when networks in the container have been discovered?

A. Network containers can be created manually or automatically. In the first case, the First Discovered/Last Discovered timestamps are not displayed, because by design the information about discovered networks inside is not propagated to the container level. If the container was automatically converted from a network, it preserves the value from this network.

Q. Why do some networks have a gray background?

A. Gray networks, also called “non-NIOS networks”. These networks are displayed with a blank value in the Managed column. This indicates that the network is discovered but

available network information is not sufficient to identify and catalog the network in IPAM at the present time. This can be due to the following reasons:

- The admin or operation status of the corresponding interface is "down". That is, the interface is either disconnected physically or disabled by the administrator.
- The prefix length for the network is /32 (ipv4) or /128 (ipv6). Network Insight treats this as a route to a specific device rather than a subnet, therefore it does not create such network in IPAM.
- The route for this interface is configured incorrectly.
- The route is learned from a remote source via BGP, OSPF and so on (i.e., indirect next hop), or it comes from a static route using the netmgmt protocol. Network Insight creates networks in IPAM for only direct and local routes from routing tables.
- The network is within a VRF and the VRF is not mapped to a network view. VRF mapping is required in this case for the network to appear in IPAM. Some time after the VRF is mapped, the network turns from non-NIOS to unmanaged, or managed, if the network is already present in IPAM.

Yellow networks. These networks are unmanaged. It indicates that the network is not managed under IPAM, but enough network information is catalogued so that the network can be converted to managed status. You can provision these networks onto devices.

White networks. These networks are currently managed under IPAM, converted to an IPAM network. You can provision and de-provision managed networks. When you manually create or convert an existing unmanaged network in IPAM tab, this network is considered as managed. Otherwise, it considered as unmanaged. Managed networks are always white, and unmanaged are yellow. (Figure “Managed, unmanaged, and gray networks”)

	NETWORK	VRF NAME	VRF DESCRIPTION	VRF RD	MANAGED
		(default)IOS		0:0	No
		(default)IOS		0:0	No
		(default)IOS		0:0	No
	2.5.40.0/24	(default)IOS		0:0	Yes
		(default)IOS		0:0	
		(default)IOS		0:0	
		(default)IOS		0:0	No
	10.220.171.0/29	(default)IOS		0:0	
	2001:dc0:c2:a6...	(default)IOS		0:0	No

Managed, unmanaged, and gray networks

From this screenshot, you can see that we created network 2.5.40.0/24, and it becomes white and “managed”. Gray and yellow networks were later collected from discovered devices.

Q. How is the management interface on an end host determined?

A. As end hosts normally have one interface, it is considered to be the management interface. If they are several, e.g. in the case of multiple network cards in a PC, they are treated as separate devices with their own addresses.

Q. What does this message mean when you hover over a network object in NIOS: “This network is currently not available as a NIOS Network object.”

A. This message appears over the gray networks described above. These networks are not handled by NIOS and not discovered by Network Insight. From the NIOS Admin Guide: *“When a network is in this state, you are limited to de-provisioning discovered networks of this type from their host device and viewing device details.”*