

WHITEPAPER

Authoritative DDI

Insights and Best Practices for Improving Your DDI Deployment



Table of Contents

Introduction	3
Definitions.....	4
IPAM.....	4
IPAM structure	4
Metadata	5
APIs	5
Unmanaged networks	5
Unused networks.....	6
Discovery timestamps	6
Filters.....	7
Prioritization of which networks to remove.....	7
Validate unused networks.....	8
DHCP	8
Unused DHCP ranges.....	8
Unused fixed addresses.....	9
API approach	9
Logging approach	10
DNS.....	10
Automatic DNS scavenging.....	10
Assisted DNS scavenging.....	11
Default.....	11
Large datasets.....	12
Conclusion.....	12



Introduction

If your DNS, DHCP and IPAM (DDI) deployment has been running for even a relatively short time, it is natural for configuration drift to occur. For example, provisioned networks often never get decommissioned or a host changes names, but nobody updates the DNS record. This whitepaper will provide some practical insights and best practices for establishing authoritative IPAM and improving your DDI deployment for better visibility, automation and control.

This whitepaper assumes you have a background in Infoblox NIOS and managing NIOS deployments. The features mentioned require named applications or dedicated appliances, such as Reporting and Analytics or Network Insight, and assume that your network is correctly configured, available and running as intended.

Definitions

Authoritative DDI is the automated discovery and visibility of all network data assets, real time updates, status, ownership and attributes gathered into a centralized database to provide an accurate, reliable reflection of the true network state.

Discovery — Authoritative IPAM automatically discovers all network types (e.g., on-prem, private/virtual, public, hybrid and multi-cloud, wired, wireless, SDN), assets (e.g., IP and MAC addresses, hostnames, topologies, subnets, VLANs, device and end hosts), and contextual data (e.g., device availability, vendor, model, OS, location and timestamp attributes, user context, and more) and syncs it into a centralized database.

Accuracy — Authoritative IPAM ensures accuracy by comparing IPAM database records with the actual network state to detect discrepancies, provide notification, reports and automated, policy-based remediation. Rather than manual user inputs and tracking, authoritative IPAM uses end-to-end workload automation to identify which assets are in use, and which are available for allocation.

Automation — When you possess an accurate database of record, opportunities for automation are extensive. It enables automated end-to-end workload provisioning, deprovisioning and DNS, DHCP and IPAM (DDI) component syncing, NAC end-host identification and quarantine, end-host vulnerability scanning, ecosystem intel and threat sharing, alerting, reporting, analytics and more. This delivers greater accuracy, reliability, processing speed and cross-team collaboration while enabling skilled workers to be redeployed to higher-value assignments.

Result— Authoritative DDI delivers a clear, real-time summary and detailed forensic visibility into everything on your network regardless of complexity, diversity, vendor or geographic distribution so you can see what you're managing. It provides up-to-date insights on your infrastructure, vendors, users and ecosystem activity to avoid conflicts and outages and ensure ultimate availability, performance and efficiency.

IPAM

IPAM structure

It is best to begin building the IPAM structure using a top-down approach. Your IP spaces should represent your network correctly. Make sure you have the top-level networks and network containers in place and validate the Extensible Attributes that are present.

Because these attributes are inherited to lower-level objects, you wield considerable power at this level. If you want to make large-scale changes, it is good to know that any objects can be exported in CSV format, modified and imported with the “override” setting. You can also consider using the API if you require more programmatic control or if you want to leverage external data repositories.

The most efficient way of dealing with DDI data at scale in a human readable format is by leveraging CSV import and export. You can easily manipulate CSV files in a text editor or your choice of spreadsheet handler. You also can readily transform them through scripting and load them in other databases for further integration. For all information regarding CSV import, check out the [CSV Import Reference](#), which provides syntax examples and all available fields.

Metadata

Planning and implementing a metadata strategy and how to assign it can significantly improve automation and save considerable time, headaches and rework later. The strategy should cover how you want to organize, access and use your data in production, especially regarding automation. Knowing how and where to automate are critical decisions, and assessing them in advance provides the greatest return on investment. Most IPAM strategies consider data in various ways, including hierarchical, functional and overlay methodologies. Start by planning your hierarchy in advance, whether simply by defining “building:floor,” or with greater complexity, like “national:regional:local:building:floor.” This approach allows you to predefine metatags for easy programmatic assignments, improve efficiency, minimize errors, avoid rework and support an easier, more fluid automation process.

APIs

Another approach to defining IPAM structure is via the API. Check the API documentation available on your Grid Manager. When structuring metadata, API automation workflows depend on consistent metadata. Infoblox IPAM can act as an enforcement mechanism to guarantee automation workflows with a standards-based approach to the IPAM service. Standardization helps to establish, organize and maintain a dynamic and authoritative repository of IP address assignments.

Unmanaged networks

After deploying Network Insight (NI) you may see that unmanaged networks are highlighted in yellow. These are networks that NI discovered on infrastructure devices in your network that were not defined in your IPAM structure. You have two options for resolving unmanaged networks. Either they are unsanctioned networks that you should remove from the configuration of the relevant devices, or they are networks that were never created in IPAM that you can convert to managed networks.



There are a variety of resources available to simplify tasks like performing mass conversions when managing a large number of networks, using CSV imports for bulk changes, or using APIs and even the Ansible Playbook:

- [Infoblox REST API Deployment Guide](#)
- [Importing and Exporting Data using CSV Import](#)
- [GitHub Infoblox-Client](#)
- [Ansible Infoblox Playbook](#)

Because unmanaged networks can emerge at any time, it's good practice to check every few weeks to ensure that no new unmanaged networks have emerged. You can either create a saved filter or leverage smart folders to display any unmanaged networks.

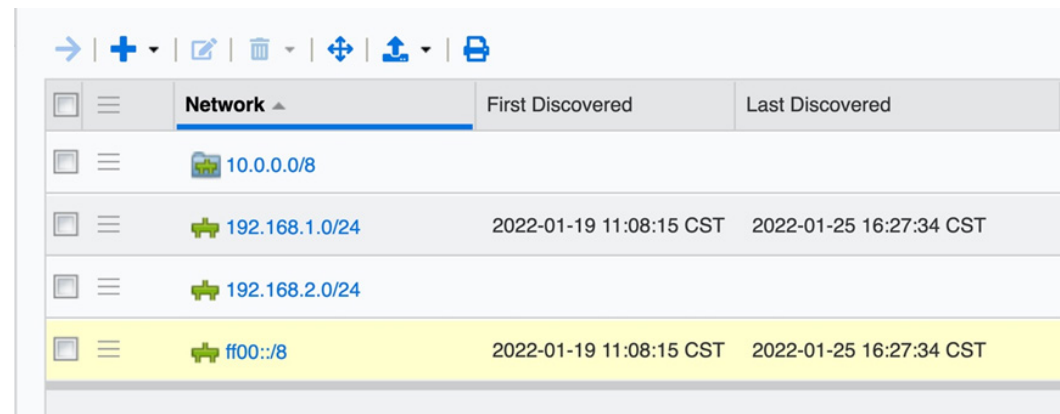
Unused networks

Once you are satisfied with the overall IPAM structure and once you convert all unmanaged networks, review the existing networks to determine if they are still in use.

Discovery timestamps

Navigate to Data Management > IPAM. Verify that discovery is enabled on all networks and network containers. All networks should have a “First Discovered” and “Last Discovered” timestamp (see Figure 1).

If the “First Discovered” and “Last Discovered” columns are not showing, edit the column settings. If no timestamps are showing, then NI has not detected any network device where this network was configured since the initial deployment.



	Network	First Discovered	Last Discovered
	10.0.0.0/8		
	192.168.1.0/24	2022-01-19 11:08:15 CST	2022-01-25 16:27:34 CST
	192.168.2.0/24		
	ff00::/8	2022-01-19 11:08:15 CST	2022-01-25 16:27:34 CST

Figure 1: Discovery Timestamps

If you have been running discovery for a while, you will be able to spot the networks that are in use because they will have a current “Last Discovered” timestamp.

Active networks should have a “Last Discovered” timestamp within the last 48 hours. To remove stale networks, export them from the UI in NIOS CSV format and import that file with the “Delete” action.



Filters

As in other parts of the GUI, you can create a filter to show a subset of the data. This step allows you to look for networks that are not discoverable as of a given date (see Figure 2). For example, use filters to find all networks related to a site after it has been decommissioned.

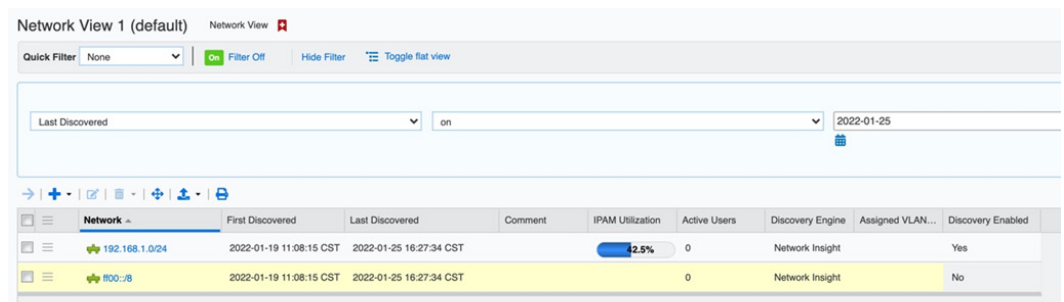


Figure 2: Discovery Filters

Prioritization of which networks to remove

When dealing with large environments, you need to prioritize the cleanup where it has the greatest impact. Aside from business drivers that determine network priorities, you can leverage IPAM utilization statistics to further decide which networks to remove. Note that IPAM utilization is not equivalent to DHCP utilization. IPAM utilization shows how much of a given IP space has been allocated (i.e., exists). This is contrary to DHCP utilization, which shows how many IP addresses have been assigned within a given range in the DHCP section of the Data Management tab. A second method for prioritizing which networks can be removed involves leveraging and considering timestamps as discussed earlier in the “Discovery timestamps” section.

Validate unused networks

DNS Query Monitoring

Before removing networks from IPAM, first determine that the network is not in use. A quick way to do so is by leveraging the DNS Query Monitoring feature, which tracks existing zones and records and when they were last queried. To validate unused IP space, check the in-addr.arpa zones for the respective IP networks that you plan to decommission.

In the example that Figure 3 illustrates, the query monitoring feature is displayed with a filter “Zone ends with in-addr.arpa”. While not completely foolproof, this is a very reliable way to validate the use or disuse of a network space because many software applications use reverse lookups when they connect to an IP. You can also use this method by itself to find unused forward and reverse DNS entries.

ZONE	NAME	RECORD TYPE	RECORD DATA	MONITORED SINCE	LAST QUERIED
0.0.10.in-addr.arpa	0	CNAME Record	0.test.0.0.10.in-addr.arpa	2019-05-08	Not Queried since 2019-05-08
0.0.10.in-addr.arpa	2	CNAME Record	2.test.0.0.10.in-addr.arpa	2019-05-08	Not Queried since 2019-05-08
0.0.10.in-addr.arpa	3	CNAME Record	3.test.0.0.10.in-addr.arpa	2019-05-08	Not Queried since 2019-05-08
0.0.10.in-addr.arpa	4	CNAME Record	4.test.0.0.10.in-addr.arpa	2019-05-08	Not Queried since 2019-05-08

Figure 3: DNS Query Monitoring

DHCP

Removing the unused networks in the previous steps will reduce the DHCP-only focus of this part of the whitepaper.

Unused DHCP ranges

(Aka DHCP scopes)

To find unused DHCP ranges, navigate to reporting and enter the following search:

```
sourcetype=ib:dhcp:network index=ib_dhcp | eval dedup_key=view."/"address."/"cidr | dedup dedup_key | msservers ms_servers | sort 0 -num(dhcp_utilization) | eval Free=address_totaldhcp_hosts | rename timestamp as Timestamp, view as "Network View", address as Network, cidr as CIDR, dhcp_utilization as "DHCPv4 Utilization %", ranges as "Total # of Ranges", address_total as "Available IP", dhcp_hosts as "Used total", static_hosts as "Statically assigned", dynamic_hosts as "Dynamically assigned" | table Timestamp, "Network View", Network, CIDR, "DHCPv4 Utilization %", "Total # of Ranges", "Available IP", "Dynamically assigned", "Statically assigned", Free, "Used total"
```

This search provides you with a report of all networks with a display of DHCP ranges, reservations or fixed addresses. While these are absolute values, they supply a good indication of the IPAM networks that contain only fixed addresses or DHCP ranges that are too small for practical use.

Events	Patterns	Statistics (4,273)	Visualization							
100 Per Page <input type="checkbox"/> Format <input type="checkbox"/> Preview <input type="checkbox"/>										
<div>< Prev 1 2 3 4 5 6 7 8 9 ... Next ></div>										
Timestamp <input type="checkbox"/>	Network View <input type="checkbox"/>	Network <input type="checkbox"/>	CIDR <input type="checkbox"/>	DHCPv4 Utilization % <input type="checkbox"/>	Total # of Ranges <input type="checkbox"/>	Available IP <input type="checkbox"/>	Dynamically assigned <input type="checkbox"/>	Statically assigned <input type="checkbox"/>	Free <input type="checkbox"/>	Used total <input type="checkbox"/>
2020-11-30 18:27:51	Company 1	6.6.6.0	29	100.0	1	6	0	6	0	6
2020-11-30 18:27:51	Company 1	5.6.7.96	28	100.0	0	1	0	1	0	1
2020-11-30 18:27:51	Company 1	34.34.1.240	28	100.0	0	1	0	1	0	1
2020-11-30 18:27:51	Company 1	128.242.99.0	25	100.0	0	2	0	2	0	2
2020-11-30 18:27:51	Company 1	1.0.179.0	24	100.0	0	1	0	1	0	1
2020-11-30 18:27:51	Company 1	10.70.144.0	20	100.0	0	1	0	1	0	1

Figure 4: Discovering Unused DHCP Ranges

In this example report in Figure 4, you can see that the /29 CIDR is impractical as a DHCP range and would qualify as a good candidate for removal. Review ranges with no available IPs because they may need to be expanded. Note that DHCP will designate addresses as abandoned, which marks them “in use” until the range is exhausted. Only at that stage will DHCP try to reclaim the IP.

Unused fixed addresses

API approach

Full list

The following curl command will return a full listing of all fixed addresses. If these addresses do not have a first or last discovered value, then you can consider deleting them. The absence of a value in either field indicates that NI has never encountered them as active in any ARP or CAM table.

```
curl --location --request GET 'https://<IP address or FQDN>/wapi/v2.10/fixedaddress?_return_fields=discovered_data.mac_address,discovered_data.last_discovered,discovered_data.first_discovered&_max_results=2000'--header 'Authorization: Basic *****'
```

Before removing these addresses, validate the MAC addresses against the lease history report to ensure that these clients did not request a lease.



Last discovered

By making a small modification to the previous curl command, you can find fixed addresses that have not been discovered since a specific time.

To leverage the timestamp values, decide what time will be your designated cutoff point. Any IPs that have not been discovered since then can be marked for deletion. To do so, use the “epoch” value from your system clock or an online [calculator](#) for the cutoff value.

The following example uses the epoch value of 1575170577, which is the equivalent of Saturday, November 30, 2019 7:22:57 PM GMT-8. Note the percent notation of the lesser than sign “%3C”. “%3C=1575170577” is the equivalent of “<=1575170577” and directs this curl command to return all fixed addresses that have a last discovered timestamp prior to November 30, 2019 7:22:57.

```
curl --location --request GET 'https://<IP address or FQDN>/wapi/v2.10/
fixedaddress?discovered_data.last_discovered%3C=1575170577&_return_fields=discovered_
data.mac_address' --header 'Authorization: Basic *****'
```

Bogus MAC

Often you can find fixed address placeholder entries that are not using actual MAC addresses. To find all such entries for MACs with “00:00:00:00:00:00”, use the following example. Note that NIOS also uses this all-zero address to designate reservations:

```
curl --location --request GET 'https://<IP address or FQDN>/wapi/v2.10/
fixedaddress?mac=00:00:00:00:00:00' --header 'Authorization: Basic *****'
```

Logging approach

Fixed addresses can be on the network without obtaining a lease. Further, if logging was not enabled, they will not stand out and it can be hard to differentiate them from other hosts when obtaining their lease.

You can use the CSV export feature to export all fixed addresses, and then cross reference the MAC addresses and IPs from that file with your SIEM for other events associated with those IPs. You can also leverage your authentication logs because these will usually track the source IP of the authentication.

DNS

Automatic DNS scavenging

Infoblox NIOS includes a helpful built-in feature for DNS scavenging, which means that a record, if it is not queried, can be marked for deletion. Additionally, because Infoblox DNS scavenging is so flexible, you can mark records as reclaimable based on many different attributes that extend beyond the last queried time. You also can define policies regarding the record types, record source or user-defined Extensible Attributes. The relevant documentation is available [here](#).

DNS Scavenging also allows for a semi-automated process, where records are flagged as reclaimable and then can be manually deleted, or where the policy can automatically delete dynamic records when they become reclaimable.



Assisted DNS scavenging

Default

Due to the risk of enabling automated processes to perform DNS record reclamation, you can manage this functionality manually for greater control.

1. Ensure query monitoring is enabled
2. Allow sufficient time to pass before moving to the next step *
3. Filter and export the entries from the query monitoring viewer that have not been queried
4. Craft a CSV import file for the records that were exported
5. Import the CSV file with the delete option

**Note: The waiting time is subjective and left to user context and discretion due to breadth of impacting variables including but not limited to workflow scheduling, related processes, dependent projects, other time intervals, and so on. It could be 30 minutes, an hour or tens of hours subject to the needs of the user and organization.*

If records are accidentally deleted, you can either revert to a previous state using your last known backup or import the CSV file with the create option.

Large datasets

Sometimes the dataset of monitored records is too large for the query monitoring viewer to export. Should this occur, you can leverage the backup file.

1. Ensure query monitoring is enabled
2. Allow sufficient time to pass before moving to the next step
3. Take a backup of the Grid
4. Rename the file extension from .bak to .tgz
5. Extract the archive file
6. Open Onedb.xml (an xml dump of the database)
7. Note that each DNS record is an .xml object with an epoch format timestamp. This is the “last queried” value.
8. Filter and extract stale records based on your requirements.
 - a. Under normal circumstances, you can assume a host is no longer on the network if it has not been around for a year.
 - b. This approach accounts for events and workloads that only take place once a year that might have dedicated equipment associated with them. Because these timeframes are unique to your network, it’s always good to practice due diligence.
9. Craft a CSV import file
10. Import the CSV file with the delete option

As noted above in the case where records are accidentally deleted, you can either restore your last known backup, import the CSV file with the create option or use the snapshot feature.

Conclusion

This whitepaper has provided some practical insights, best practices and processes for improving your DNS, DHCP and IPAM infrastructure. With proper attention and management, authoritative DDI can help discover and deliver visibility into everything on your network. It enables you to establish and maintain a single-source-of-truth database so you can automate workflows with confidence for greater cost savings and efficiency. Most importantly, it provides a proven way to keep services available and performing as expected, while keeping your users, partners, customers and stakeholders connected even under adverse scenarios.

If you’d like more information, please contact your Infoblox Solutions Architect or Account Team, or visit www.infoblox.com.



Infoblox unites networking and security to deliver unmatched performance and protection. Trusted by Fortune 100 companies and emerging innovators, we provide real-time visibility and control over who and what connects to your network, so your organization runs faster and stops threats earlier.

Corporate Headquarters
2390 Mission College Blvd, Ste. 501
Santa Clara, CA 95054

+1.408.986.4000
www.infoblox.com