

The Network Source of Truth: How Engineering Teams Establish and Use These Critical Tools

December 2025 EMA Research Report

By Shamus McGillicuddy, VP of Research
Network Infrastructure and Operations



Table of Contents

1	Introduction
2	Rising Adoption, Modest Success
2	Research Goals: Define How You Should Build Your Network Source of Truth
3	Research Methodology
5	Defining a Network Source of Truth
6	A Network Source of Truth is...
7	What Data Lives in the NSoT?
9	Why Network Teams Need a Source of Truth
10	Data Fragmentation Bogs Down Operations
11	Establishing an NSoT Enables Efficient, Consistent Operations
11	A Powerful NSoT Data Model Enables Scalable Automation
12	The NetDevOps Angle
14	Use Cases for a Network Source of Truth
15	The Core Use Cases
17	Secondary Use Cases
18	How Network Teams Establish a Network Source of Truth
19	Step 1: Identify and Validate Key Sources of Data
19	Step 2: Establish an NSoT Tool that can Model all Relevant Data
20	Step 3: Clean Up and Import Data
21	Step 4: Ongoing Validation and Reconciliation of Data
23	Key Capabilities of a Network Source of Truth
	The Essentials
	Emerging Requirements
25	Network Source of Truth Challenges
26	Management Buy-In
26	Populating Data in NSoT: A Long and Painful Journey
27	Adoption: Network Engineers Resist Change
27	Open Source and Homegrown NSoT Comes with Technical Debt
28	What Should You Demand from NSoT Vendors?
29	Sophisticated Pre-Sales Engineering Support
29	Flexible and Extensible Data Model
30	Deep Integrations with Network Automation Toolsets
30	Event-Driven Operations
31	Example NSoT Vendors
33	BlueCat Networks
33	EfficientIP
33	Forward Networks
33	Infoblox
34	IP Fabric
34	NetBox Labs
34	Network to Code
34	OpsMill
35	EMA Perspective
36	An NSoT Action Plan
37	Case Study: Truma Accelerates Global Growth with Infoblox Universal DDI



Introduction

Over the last decade, a new class of management tool, the network source of truth (NSoT), slowly emerged in the world of network operations. Enterprise Management Associates (EMA) analysts have observed increased discussions about and rising adoption of NSoT tools as IT organizations mature their approach to network automation.

Network teams use NSoT tools to establish a centrally located, authoritative repository of operations data that network automation tools can reference programmatically when interacting with and making changes to a network. However, an NSoT tool has applications beyond network automation, as this report will demonstrate.

Rising Adoption, Modest Success

Recent EMA research found that 80% of network automation teams have an NSoT solution, but only 20% believe their NSoT is completely effective.¹ Most IT organizations are using ad hoc tools or open source software for NSoT use cases. NetBox is quite popular in this context. Released by cloud provider Digital Ocean in 2016 as an open source project, NetBox is a hybrid tool that combines IP address management (IPAM) functionality with the asset and inventory management features of data center infrastructure management (DCIM) tools. In fact, EMA research finds that the tools involved most often in establishing an NSoT include DCIM and IPAM products. Secondarily, they use configuration management databases (CMDBs) and network controllers.²

More recently, vendors started offering NSoT products. In EMA's opinion, most enterprises are still maturing their approach to NSoT tooling. While some will stay with open source tools and leverage internal personnel who have a mix of networking and software development skills to support them, many others want to engage with a vendor. In other words, EMA believes network teams are at a crossroads with NSoT solutions.

Research Goals: Define How You Should Build Your Network Source of Truth

Given the inflection point that many enterprises are approaching, EMA conducted new market research that explores NSoT tool strategy in depth. This report seeks to:

- Define the concept of a network source of truth
- Examine why IT organizations are adopting NSoT tools
- Evaluate how enterprises establish an NSoT
- Study how companies use NSoTs
- Review the challenges that are blocking progress

¹ EMA, "Enterprise Network Automation: Emerging From the Dark Ages and Reaching Toward NetDevOps," March 2024.

² Ibid.



Research Methodology

This is a qualitative research project. EMA conducted in-depth, one-on-one interviews with 17 IT professionals who are deeply involved in implementing, maintaining, and/or using a network source of truth. EMA identified common perceptions and experiences across these conversations to formulate key findings around strategies for NSoT adoption.

EMA granted 100% anonymity to these interview subjects since they were not authorized by their employers to discuss their technical programs. The following list provides insight into the people interviewed:

- 1 Global automobile manufacturer (Network engineer)
- 2 Fortune 500 bank (Senior network engineer)
- 3 Fortune 500 financial services company (Network engineer)
- 4 Fortune 500 insurance company (NetDevOps engineer)
- 5 Fortune 500 IT hardware manufacturer (Director of networking)
- 6 Fortune 500 pharmaceutical company (Network automation developer)
- 7 Fortune 500 retailer (Network engineering manager)
- 8 Fortune 500 technology company (Network automation software engineer)
- 9 Internet exchange (NetDevOps engineer)
- 10 Large American university (Network automation engineer)
- 11 Midsized renewable energy company (Network automation engineer)
- 12 Multi-billion-dollar European media company (NetDevOps architect)
- 13 Multi-billion-dollar media company (Network engineer)
- 14 Multi-billion-dollar pharmaceutical company (Network engineer)
- 15 Multi-billion-dollar gaming company (Network automation architect)
- 16 Multi-billion-dollar renewable energy company (Senior network engineer)
- 17 \$1 billion logistics company (Network operations manager)



Defining a Network Source of Truth

A network source of truth (NSoT) is fundamentally a system of documentation for a network. There is some debate about what should be documented in such a system and how. Some engineers believe it should reflect the current state of the network, but most tell EMA that it should contain the intended (declarative) state.

In other words, an NSoT should document the intent standard against which the production network (current state) can be managed. As other tools and processes reference this NSoT, engineers can ensure that operational outcomes align with network intent.

A Network Source of Truth is...

EMA defines an NSoT as:

- A system or repository that documents, models, and maintains the intended state of a production network.
- It should have both a robust user interface and open APIs that provide programmatic access so that various other tools can integrate with it, exchange data with it, and apply the data it contains to network operations



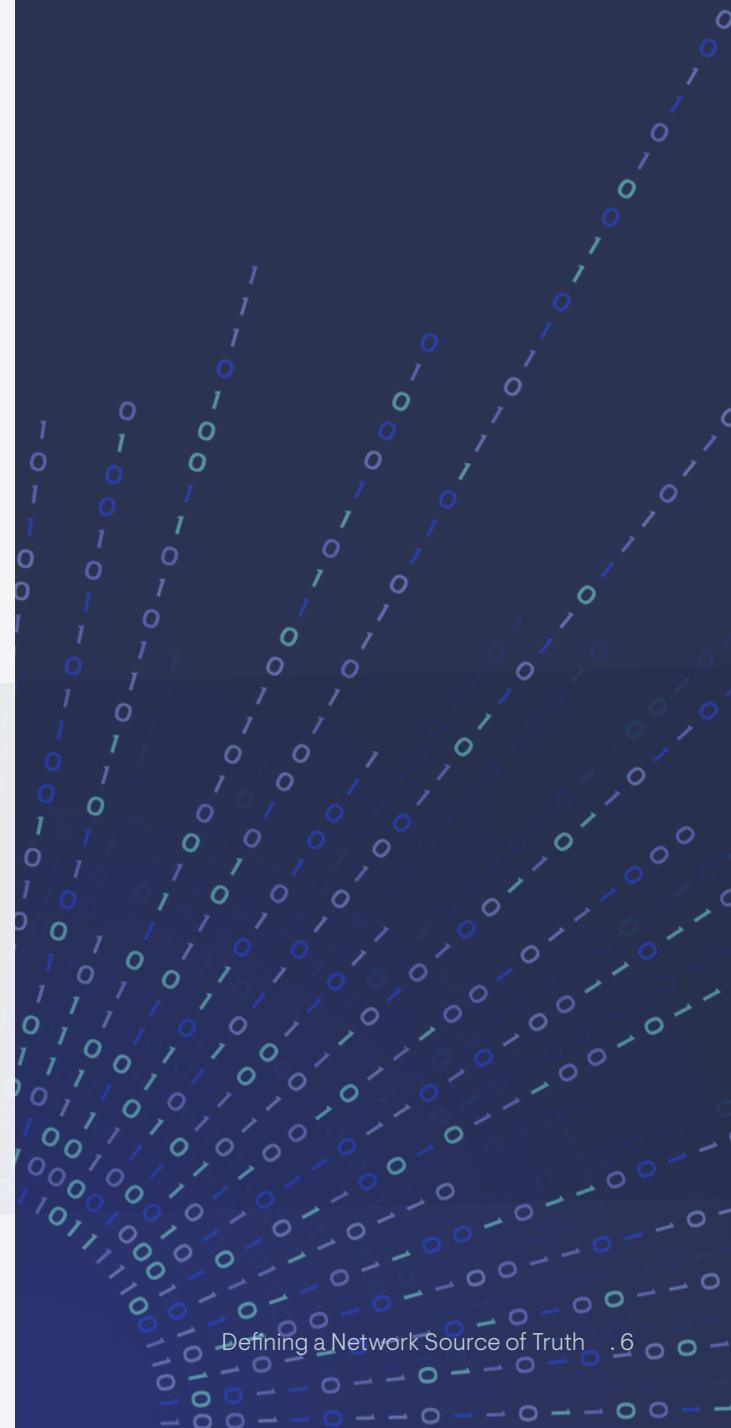
*Network automation architect,
multi-billion-dollar gaming company*

“It’s a place where you model your production network the way you intend it to work. And it’s something that your production network then needs to essentially replicate. Once you look at the source of truth, you should be able to answer any question about what a network looks like or is supposed to look like at any given point of time.”



*Network automation engineer,
large North American university*

“For us, the source of truth is our intended state for not only what devices we have, but also what their configuration should be and how they should be connected. It is the system in which we have declared our intent for the network, and everything should be based off it.”



What Data Lives in the NSoT?

Engineers have some minor differences of opinion about the data they need to centralize in an NSoT, but most agree that it should include three core types of data.

The Core Requirements

- **Inventory data**, including device make and model, serial numbers, OS versions, rack and unit information, vendor support contracts, lifecycle information, and so on.
- **IP address management data**, including IP address allocations and pools, subnets, public and private space, etc.
- **Topology and connectivity information**, including physical and logical topology, inter-device and inter-site connections, cabling, patch panels, optics, etc.

Use Cases and Business Requirements Drive Other Data Choices

Engineers told EMA that the use cases they pursue and the nature of their businesses dictate exactly what data will live in their NSoT beyond the three classes of information mentioned.

For example, WAN engineers who are focused on documenting the truth about their connectivity will want detailed information of all their WAN circuits, including provider names, bandwidth, SLA contract information, customer support contacts, and inventory information about the routers where the circuits terminate.

On the other hand, data center network engineers may want to add information about their network fabric control planes (BGP, EVPN-VXLAN), power and cooling equipment, and compute infrastructure, such as servers and hypervisors.

What About Configuration Files?

Engineers have varied opinions on where configuration files should live. In some cases, this opinion impacts their definition of an NSoT. For example, engineers who considered raw config files or golden configs to be their ultimate source of truth data will often use version control tools, like Git, as their NSoT tools. These tools are meant to store and govern the use of individual files, documents, and software code. They do not abstract configs into structured data.

Most engineers believe that an NSoT should abstract configurations into structured data, where they can make higher-level decisions about their network and translate those actions into actual configs programmatically. They typically store raw config files somewhere as a backup. This gives them more flexibility in how they design and manage their networks.



*Senior network engineer,
Fortune 500 bank*

“An [NSoT] should contain all the data that is required to build configurations.”



*Network engineer,
multi-billion-dollar media company*

“Anything that you might need to generate a config file needs to be abstracted into some form of model and stored in that source of truth.”



*Network engineer,
Fortune 500 financial services company*

“The intention with a source of truth is to create what our desired end state [for the network] should look like, and then your automation consumes that and uses it to go build that.”



Why Network Teams Need a Source of Truth

NSoT adoption is about more than powering network automation. It addresses a fundamental operational data problem. Network engineers lack good network documentation. Knowledge of what's on the network and how it is meant to operate is often fragmented, inconsistent, and difficult to find.



*Network automation software engineer,
Fortune 500 technology company*

“I’ve seen it in multiple large enterprises. As the company expands and evolves, you start losing what’s actually on your network. When I was a network engineer at [a large financial services company], we had a device that was online for 10 years with an active Verizon circuit that nobody knew was there. Everything was disconnected from this device except the active Verizon circuit, and it was sitting there for 10 years, just running and collecting a bill.”

IT organizations that lack an NSoT tool will maintain a hodgepodge of commercial and home-grown repositories, including spreadsheets, Visio diagrams, inaccurate CMDBs, and proprietary network element management systems offered by the various vendors whose infrastructure platforms populate networks (switching, routing, Wi-Fi, SD-WAN, firewalls, etc.).

Data Fragmentation Bogs Down Operations

The quality and authoritativeness of these disparate data sources vary. Also, the fragmentation of data adds complexity. Engineers must gather information from multiple documents and systems while conducting operations. After completing an operation, they must then update these various systems to ensure their changes are documented. In a fragmented data environment, network teams struggle to ensure that personnel have full access to all systems of record, and they struggle to enforce consistent maintenance of data.



Senior network engineer, multi-billion-dollar renewable energy company

“Our key primary issue was that there was so much information for operational teams. All that information was put in so many different Excel sheets and so many spaces. Since there are multiple teams involved in operations, from level one, level two, level three, each team had their own different version of the spreadsheet at any given point of time. No one had an updated sheet.”



Network engineer, global automobile manufacturer

“Network departments tend to be project-focused. You do a project, and you will produce some documentation of it. Then, you do another project, and on and on. Sometimes projects overlap. They use the same infrastructure, so it’s very difficult to know which documentation is correct. It’s not always easy to find the right information.”



Network operations manager, \$1 billion logistics company

“There was not a central repository that dictated how everything should be. That lack makes things like disaster recovery, lifecycle management, and auditing very difficult to do.”

Establishing an NSoT Enables Efficient, Consistent Operations

When organizations consolidate network data into a central location that is authoritative, engineers have consistent and easy access to the information they need to conduct operations, whether they are deploying new devices, making configuration changes, auditing the network, or troubleshooting a fault. This need to centralize network data is one of the primary drivers of adoption.

An NSoT delivers efficiency and improves the quality of operations. With the best information located in one place, network personnel can gather data quickly. In a troubleshooting scenario, for instance, an engineer can gather all the information about a problem domain and compare it to the operating state that monitoring tools reveal. Bad network documentation would cause an engineer to make incorrect assumptions, which could worsen the issue they are trying to fix.



*Network automation software engineer,
Fortune 500 technology company*

“If you’re troubleshooting an issue, actually knowing what’s on the network and how [devices] are interacting with each other really helps a lot. Rather than go into command line, a real, accurate source of truth can help with troubleshooting issues very well.”

A Powerful NSoT Data Model Enables Scalable Automation

Network automation is essentially about removing manual tasks and processes from operations. In the early stages of adoption, network teams are tactical. They might leverage one-off scripts to automate individual tasks. These automations might impact a single device. It’s easy to gather relevant data needed to execute the task.

As network teams scale automation with workflows that impact dozens, hundreds, or thousands of devices, having accurate data about intent becomes more important, since automation-driven errors can have much larger impacts on the network. Also, providing programmatic access to a comprehensive model of this data makes it easier for engineers to plan, test, commit, and validate automation. They can also quickly document any updates to the intended state of the network that are pushed out via automation.



*Network engineer,
multi-billion-dollar media company*

“An [NSoT] is about wrapping the data that’s always been around in one form or another in a nice, structured format with an API. When we try to move to an API-based approach with data consumption for automation, you have to be able to expose that data in a structured, programmatic way.”

The NetDevOps Angle

NetDevOps groups emerged in many IT organizations to apply DevOps principles to network operations. For instance, DevOps groups have adopted tools and practices that enabled them to manage application infrastructure as code, which allowed them to tightly couple IT operations with application development.

In the network operations world, many NetDevOps groups emerged with the intent to manage networks as code that is easily consumed or managed.

EMA finds that NetDevOps teams always have a mature approach to NSoT tools. Network engineers who are just getting started with NSoT should look to companies that publicly evangelize their NetDevOps philosophy for guidance on how to approach an NSoT project.



*Network engineer,
multi-billion-dollar media company*

“We’re trying to manage the network as code primarily to enforce standards. We’ve had issues in the past in which changes go through that haven’t been as well planned as they should have been. Basically, repeatable standards are what we want, so we can make a change and we know that a standard set of test suites will be run.”



*Network engineer,
Fortune 500 financial services company*

“The NetDevOps group’s whole philosophy is treating the network more like applications. Instead of doing this engineering task and that engineering task, NetDevOps engineers are more interested in analyzing workflows, then codifying those workflows. ‘How can I take this series of tasks that you’re doing and make that more into an application that you can execute?’”

The Role of NSoT in NetDevOps

If the goal of NetDevOps is to transform network operations processes into code, network data is the foundation. Network engineers reference various systems for the information they need to codify network infrastructure into applications and drive automation.

An NSoT tool provides centralized, structured data for consistent automation. NetDevOps teams need an NSoT tool to serve as a data hub that enables provisioning, change management, validation, and drift detection. Just as important, it should be a place where the NetDevOps group can collaborate as they build solutions.



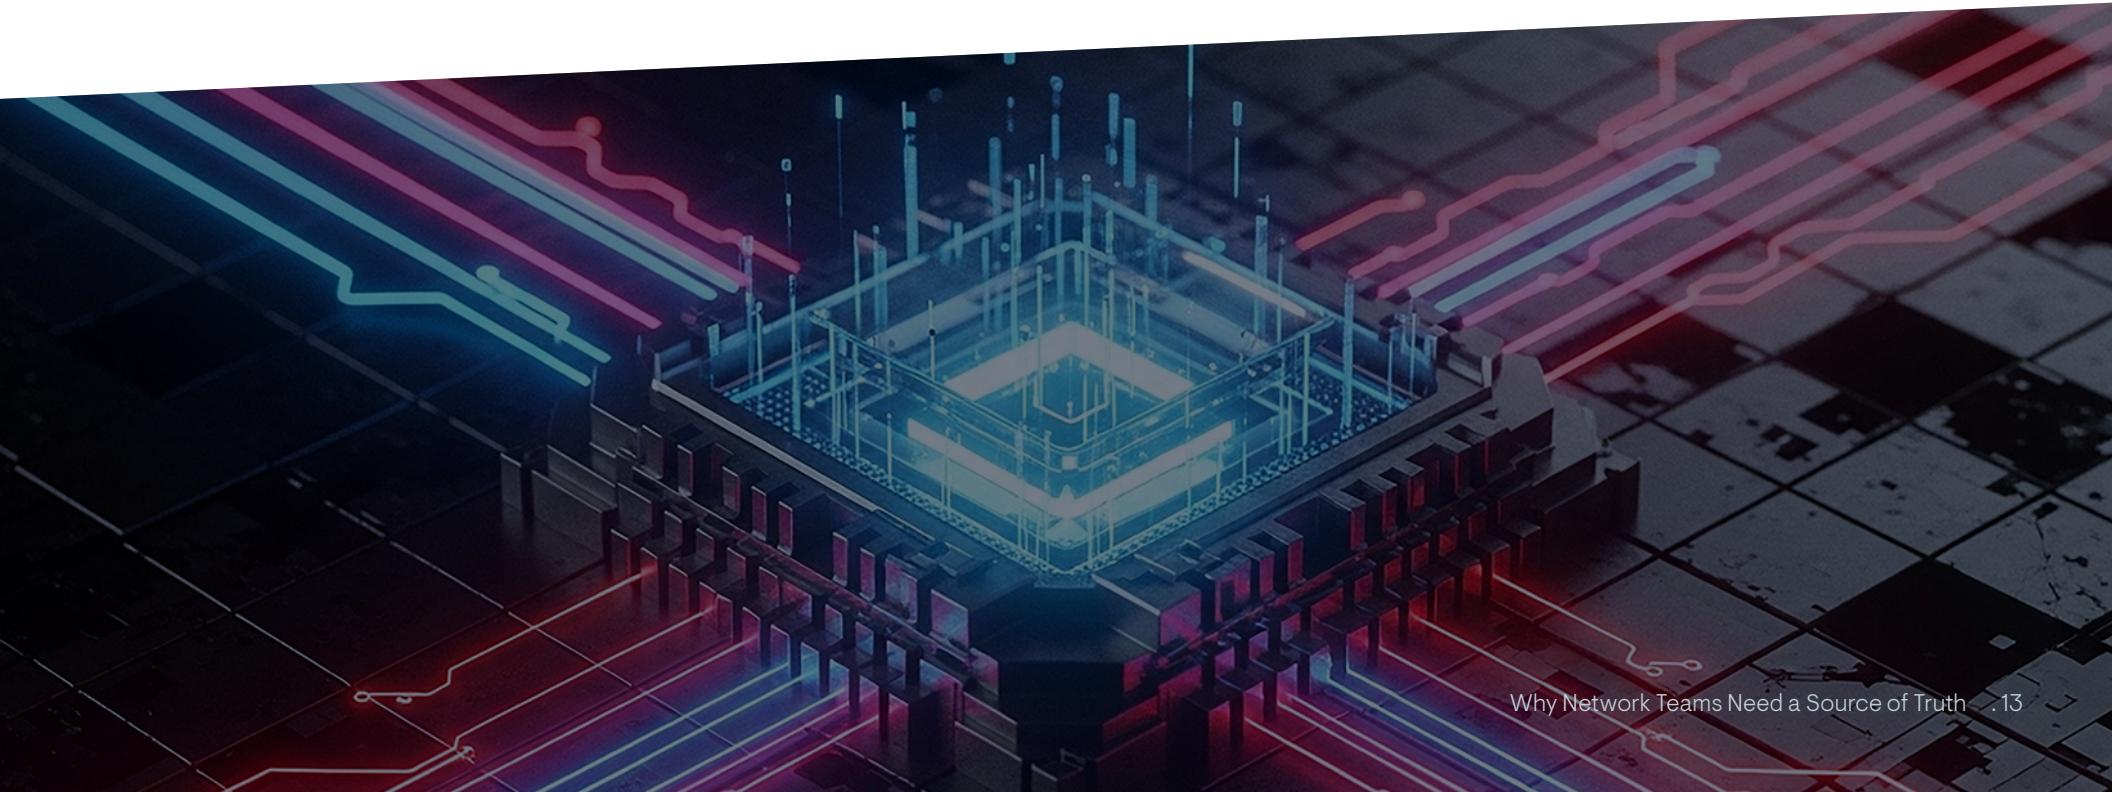
*Network engineer,
Fortune 500 financial services company*

“With NetDevOps, the goal is to take all that information and put it into a structured data format that an application or a script can then read and execute on.”



*Network automation architect,
multi-billion-dollar gaming company*

“[With an NSoT], you can have a robust automation pipeline with proper standards in place. Your network will behave in a predictable manner, and you generally are less susceptible to unexpected consequences resulting from things like config changes.”





Use Cases for a Network Source of Truth

In this section, we explore the use cases that IT organizations are pursuing with an NSoT. Network automation – a key driver of NSoT adoption – may or may not be involved in the enablement of these use cases. Some NSoT implementations simply supported manual processes for these use cases, while other organizations were fully automating some or all of them with their NSoTs at the center of their network automation solutions.

The Core Use Cases

Network Provisioning and Deployments

With all network intent data centralized in an NSoT, engineering teams can confidently pull together everything they need to execute Day 0 operations. For example, if they are provisioning a new device, the NSoT can generate a configuration for it. If they are planning to deploy a new site, the NSoT can put together a network plan for multiple devices.



*Network engineer,
global automobile manufacturer*

“I’d love to use [an NSoT] for zero touch provisioning. It can generate the configuration based on the information I have in the NSoT, then push it to Ansible. When that new switch comes online, Ansible will push the configuration to it.”

Engineering teams can implement these Day 0 operations manually with the NSoT outputs, or they can drive them automatically via integrations between the NSoT and whatever automation tools they use. In some cases, engineers extended their NSoT to support cloud deployments, too.



*Network engineering manager,
Fortune 500 retailer*

“When we onboard a new switch, the cable technician connects that switch, trunk port, or uplink port, to one of the out-of-band switches. As soon as the switch gets an IP address, NetBox comes into play to map that IP address to the specified TFTP server from where that switch can pull all its configurations.”

Change and Configuration Management

Since all information needed to generate configurations resides in an NSoT, engineers often use it for managing network changes and configurations. Many have NSoT tools that allow them to model, test, and approve proposed changes. Others told EMA that they struggle with this use case because their tools lack branching capabilities. Branching allows engineers to propose changes without impacting the central data model within the NSoT. Once changes are approved, they can merge the branch and update the model.



*Network engineer,
multi-billion-dollar media company*

“We are leveraging branching in [our NSoT] so we can do change management within it. You can create a branch and stage a change to the database. Once the change is successful, you can merge that change into the main database. We’re using that quite heavily in the automation platform that we’re building.”

Network Validation, Compliance Management, and Audits

Because an NSoT is the authoritative record of the intended state and configuration of a network, many engineers need to ensure that intent matches actual state. This helps them validate that planned changes were successful. It also enables compliance management and network audits. If the tool has features for version control and change history, network teams can track who made a change and why.

Most engineers had NSoTs that couldn't discover network state on their own. They either wrote scripts that pulled state information or they integrated an NSoT with third-party tools that discover and model production networks. These methods allowed them to compare intended network state in the NSoT with the current state of the network. Then, engineers or auditors could identify configuration drift and policy violations.



*Network engineer,
multi-billion-dollar media company*

“Having validation baked into our NSoT would be huge. To be able to push your intended state out and have a post-change report that says these reachability tests have passed automatically would be really powerful. It’s on our roadmap for next year.”



*Network engineer,
global automobile manufacturer*

“We get hit with various audits throughout the year. That gets to be a big drain on resources. It’s something that the reporting [functionality in our NSoT] can fulfill. We are generating reports in 30 minutes that would take us a week in the past.”

Network Troubleshooting

Data gathering is one of the first steps of network troubleshooting. An NSoT is a central hub for much of that data and it can accelerate the process. Some engineers told EMA that they are exploring integration of their NSoTs with their network observability tools so that they can correlate fault and performance data with intent data. This will help them identify whether intent drift or a bad change caused a problem.

Given that an NSoT is also deeply integrated into the network automation toolset, engineers can also fix the issue within the same environment by rolling back a bad change or planning and executing a fix.



*Network engineer,
Fortune 500 financial services company*

“If you have all of the information about how everything is supposed to be configured and you can compare that against the running state, that will definitely help you troubleshoot.”



*Network engineer,
multi-billion-dollar media company*

“You need to be able to understand quickly, when there’s a fault, what the network should look like. You can look at the [intent] and ask, what’s changed? At the click of a button, you can see that someone committed this change and everything’s broken. So, let’s roll back that change.”

Secondary Use Cases

What follows is a list of other examples of how network teams leverage their NSoTs. They were less popular than the cores use cases, but they still delivered tremendous value.

- **IP address management:** Most NSoT tools have an IPAM feature set, and it makes sense that many network teams manage IP address space within them. However, many other teams federate their NSoTs with a third-party IPAM tool. That third-party IPAM is the primary tool for managing addressing. The NSoT provides a place where engineers can reference the IP address plan and request changes as part of their NSoT workflows.
- **Inventory management:** As with IPAM, NSoTs typically offer inventory and asset management functionality. In this example, however, EMA found that the NSoT often becomes the primary tool for this use case. Engineers will use NSoTs to track device lifecycles, manage vendor relationships, and document the physical network.
- **Security policy design and management:** Some engineers used their NSoTs to model and track firewall policies and access control lists, for instance. Very few were using their tools to design security policies. Overall, many found NSoT tools were limited in this area.
- **Security investigations:** Some network engineers give their cybersecurity teams access to their NSoTs. Security personnel will reference it while exploring how a network was breached and how vulnerabilities occurred.



How Network Teams Establish a Network Source of Truth

Step 1: Identify and Validate Key Sources of Data

Network teams often begin an NSoT project by determining where information about the network currently lives and ensuring that this data is accurate and reliable. It may reside in various network controllers, spreadsheets, Visio diagrams, IPAM tools, CMDBs, etc. In other cases, the data may not exist at all. Engineers must extract the data from the network devices directly.



*NetDevOps architect,
multi-billion-dollar European
media company*

“The first step is to just ascertain where the truth is today. Quite honestly, it’s rarely in one place. We’ll create an inventory, then figure out where the source of truth is for each data domain actually is.”

Step 2: Establish an NSoT Tool that can Model all Relevant Data

This report will explore all the key capabilities of an NSoT in a later section. However, the first thing a network team needs in such a tool is a flexible and extensible database schema that can model disparate classes of data. Much of the data that network teams consolidate into a source of truth has traditionally resided in different systems, either as unstructured data or in specialized databases. Thus, the NSoT tool should at the very least be flexible enough to model all the data that the network team intends to put in it.



*Network automation architect,
multi-billion-dollar gaming company*

“A lot of what we did in the beginning was focused on getting the right data model in place.”

Step 3: Clean Up and Import Data

Once an NSoT with the right data model is identified, the team must begin the process of importing data from other systems. They often develop scripts to import CSV files, and they use APIs or webhooks to import data from other IT management systems and databases. In some cases, manual data entry may be necessary.



*Network engineer,
multi-billion-dollar media company*

“Once our data model was decided, it was a lot of cleaning up our spreadsheets and making sure we were mapping everything from the old ways into the new system.”

Step 4: Ongoing Validation and Reconciliation of Data

An NSoT solution requires care and feeding. Network engineers will often make network changes without updating it. Others might make mistakes while updating it. Maintaining an NSoT involves regular validation to ensure the system accurately reflects the current state of the network.



*Network operations manager,
\$1 billion logistics company*

“We do regular audits in which we compare the source of truth in NSoT with live device configs and update as necessary or fix any discrepancies.”



Key Capabilities of a Network Source of Truth

The Essentials

A Robust and Extensible Data Model

The data that lives within NSoT today is very diverse, so network engineers seek solutions that have a powerful data model that can accommodate that diversity, storing IP address space, inventory data, routing information, security policies, topology, and more in a single schema.

The types of data that live in an NSoT will evolve as network teams expand the scope of what they want to do with these tools. Thus, an NSoT must be highly extensible. The schema for modeling this data must evolve with the needs of network teams.



*Network automation architect,
multi-billion-dollar gaming company*

“Ideally, you want a system that allows you to extend the schema so you can evolve your network over time. A lot of companies are building out high-performance network fabrics for GPU clusters. So, you need to be able to model your GPU connections to switches. That’s not something that you would find available in the offerings that are out there. You should evolve [an NSoT] over time to accommodate newer topologies, new technologies.”

Integrated Workflows and Reporting Across Data Silos

An NSoT typically consolidates data from multiple toolsets and repositories into a single environment. The classic example is the combination of IPAM and inventory data. Having this data in one place gives engineers a single place to go for information. Still, it is just as important to provide integrated workflows across this data and offer the ability to produce reports and dashboards that provide visibility into these disparate classes of data.



*Network engineer,
Fortune 500 financial services company*

“[Our NSoT] has bundled it in such a way where DCIM is very tightly coupled with the IPAM system. That’s really nice for a network engineer because oftentimes, your DCIM is a completely different vendor than your IPAM. You have to go from one pane of glass to another with a completely different interface, and that makes it a little confusing for the operator.”

Platform Security and Multitenancy

Data security, fine-grained role-based access control, and multi-tenancy features are essential, ensuring that malicious actors can’t access the tool and users and teams can see or modify only data relevant to their roles. For instance, the wireless team shouldn’t see or touch data about a data center network fabric. All this ensures that sensitive information is protected and compliance requirements are met.



*Network engineer,
global automobile manufacturer*

“It’s a lot of sensitive information to put in one spot. That’s a nice target for somebody. That data needs to be encrypted when it’s in transit. When it’s at rest, that’s another thing. You need multi-tenancy and role-based access.”

APIs for Customization and Integrations

Network engineers universally agree that open and robust APIs are an essential driver of value for an NSoT. In the early stages, network teams leverage these APIs to plug their NSoT into their automation tools, including scripts and other simple tools.

As they mature their approach to automation, engineering teams leverage those APIs to enable a variety of third-party integrations. There are three basic stages of integration that engineers talk about.

Stage 1: Integration with systems of record for ingesting, reconciling, and validating network intent data

Stage 2: Integration with service management and network automation tools to drive network provisioning, deployment, changes, and compliance

Stage 3: Integration with network discovery and network monitoring to ensure NetOps monitor new networks properly, changes are detected to ensure current state aligns with intent, and workflows across monitoring and the NSoT are integrated for troubleshooting and other Day 2 operations.



*Network automation architect,
multi-billion-dollar gaming company*

“Once you have a source of truth, your entire automation ecosystem gets built around that. And it does not just facilitate Day 0. Your entire telemetry monitoring stack can also directly plug into it. So, you essentially model a network device on Day 0, then use the NSoT to provision it using things like zero touch provisioning. And as soon as the device turns active, it gets picked up by your monitoring pipeline and starts collecting metrics.”



*Network engineer,
multi-billion-dollar media company*

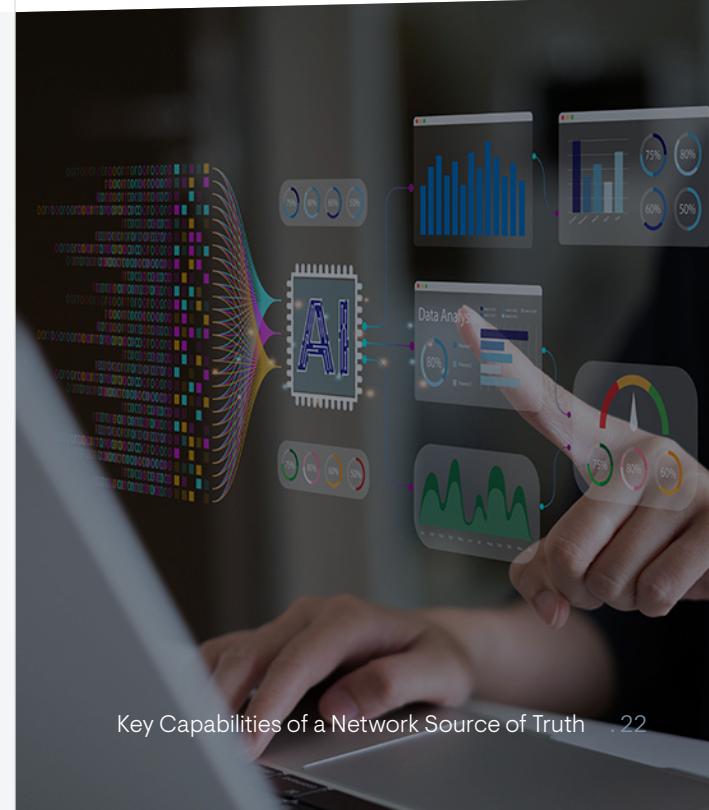
“It’s almost like a control plane for the network, where you’ve got all the data that you need to make a decision for whatever you’re working on in one tool, without you having to go all over the place to find different things. So, if I navigate to a site in source of truth, it would be good to be able to see the health of the devices at that site, maybe some metrics for the traffic level. Then, that would take you to more granular details in the other system.”

APIs are also essential for network teams that want to customize their tools to add functionality that meets their specific business requirements.



*Network automation architect,
multi-billion-dollar gaming company*

“Good APIs provide the ability to customize the product in a way that allows us to incorporate new functionality...without actually having to go back and forth with the vendor a thousand times.”



Emerging Requirements

Discovery

Early NSoT tools lacked discovery capabilities. Network teams had to import data manually via a variety of techniques. Many network teams are now seeking integrated discovery features that can find infrastructure and collect data to populate an NSoT, but also detect changes to ensure that intent and state are aligned.



*Network engineer,
multi-billion-dollar pharmaceutical company*

“I think one of the issues we’re trying to address is to find an NSoT tool that can scan the network and populate data. Also, when deployments are done, instead of forcing engineers to document a deployment, the NSoT could just automatically detect that something was deployed and update itself.”

Branching

Some network engineers have identified a need to branch the data models in their NSoTs. This allows them to plan and test changes within the tool without making changes to the NSoT data model. Essentially, they want the branching capabilities of Git so that they can apply DevOps principles to network engineering. They can plan and test the changes in a branch of the tool’s data schema. Once they’re ready to commit the change, they can merge their changes into the main system of record in the tool.



*Network automation architect,
multi-billion-dollar gaming company*

“This is about how we ensure data integrity and that changes are made in a safe and reliable manner. We want to be able to build some sort of branching capability, where you can plan a change in isolation without affecting your production database. You can get it reviewed just like you would in Git. It allows you to make changes without being concerned about how it would impact the production network.”

Expanding Beyond the Network

Many admins and users of NSoTs told EMA that they envision a future in which the word “network” drops away and it evolves into an infrastructure source of truth. Some had modest ambitions for this expansion. They simply wanted to include data about servers, hypervisors, and power and cooling equipment. This enables a broader set of users and helps integrate IT operations across teams. Others envisioned going further, modeling applications and business services.



*Network operations manager,
\$1 billion logistics company*

“It becomes more of an infrastructure source of truth, or even an environmental source of truth. It could absolutely expand to not only the network environment, but also the system environment, even potentially applications, line of business information, and third-party retail products. It becomes a one-stop shop for all the information, as it should be.”

A Foundation for AI-Driven NetOps

Some network engineering teams want to make their NSoTs more powerful and easier to use by leveraging artificial intelligence (AI). Lower-skilled admins and non-networking personnel who might struggle to work with an NSoT could interact with a virtual AI assistant that can unlock the power of an NSoT for personnel of all skill levels.



*Network automation architect,
multi-billion-dollar gaming company*

“We’ve built an AI agent that is able to answer low-hanging fruit questions. It’s also able to provide things like code snippets for querying router interfaces or circuits.”

Moreover, agentic access to an NSoT can help AI learn about a network and enable proactive and predictive management based on inferences AI can make from NSoT data.



*Network engineer,
global automobile manufacturer*

“I’d like an MCP server so we can tie our source of truth back into an internal LLM to have it start doing things like learning what the network looks like and start doing some predictive alerting.”



Network Source of Truth Challenges

This section explores the roadblocks to success that network engineering teams encounter with their NSoTs. Just finding a tool that can cover all the requirements explored in the previous section is challenging on its own.

On top of that, engineers identified several additional challenges that impact their ability to execute.

Management Buy-In

Most engineers acknowledged that NSoT tools are a new concept, and budget support is limited or nonexistent. In general, only the very largest companies that EMA spoke with were using paid products from vendors. Most were using open source tools or homegrown tools.



*Network automation engineer,
large American university*

“We have such a large network, so as soon as we get into tiered pricing, it just kind of throws it out the window. We’re usually paying Cisco an arm and a leg because we have a mandate to keep our network very up to date. So, we’re paying them a few million a year, and then we have to turn around a pay \$60,000 or \$70,000 on an [NSoT] that is not mandated? It’s very difficult to have a business justification when we also have a fairly stacked engineering team.”

Populating Data in NSoT: A Long and Painful Journey

Getting data into an NSoT is extremely painful. It can take months, even years. This isn’t necessarily the fault of the tools themselves. It’s the quality of the data they have and how it is stored. In some cases, the data simply doesn’t exist.

Finding, Cleaning up, Importing Data is Painful

Existing data might be hard to find because it lives on an engineer’s laptop in a spreadsheet, or it’s hard to extract because the data lives in a tool that lacks APIs or uses proprietary data formatting. In other cases, the data is simply bad. CMDBs, for instance, are notoriously inaccurate.



Network automation engineer, large American university

“Whenever you’re dealing with the data entry—the start of a new source of truth system—it’s always a pain because you have little spreadsheets, you’ll have documentation, you’ll have a Visio diagram in someone’s SharePoint, and someone else has a wiki.”

Some Data is Offline, Requires a Truck Roll

Some network data requires physical interaction with the network. One can only acquire it by going to the device and looking at it. This is manual labor that cannot be eliminated today.



NetDevOps engineer, internet exchange

“Capturing all the things that don’t happen online is the most difficult problem. There’s a physical cable patch between this port and this port. Somebody needs to write that down. We have incomplete data because not everyone does that consistently.”

Adoption: Network Engineers Resist Change

The effectiveness of an NSoT depends on organizational buy-in. Network engineers must incorporate the tool into their day-to-day operations. Without buy-in, engineers will go to other sources of data that are not authoritative. In other cases, they might make changes to the network without documenting them in the NSoT; over time, the NSoT drifts from the true state of the network. The quality of data in the NSoT will be unreliable.

Open Source and Homegrown NSoT Comes with Technical Debt

Most of the engineers EMA interviewed about NSoT strategy were not supported by vendors. They created their own software or adopted and customized open source tools. In a way, this fit perfectly into their tool strategy because they often integrated their NSoT into an ecosystem of open source and homegrown network automation and infrastructure as code tools.

On the other hand, building and maintaining these NSoT tools without vendor support adds technical debt. They need people with both networking and software development skills. Those people are expensive and hard to find.



*Network engineer,
Fortune 500 financial services company*



*Director of networking,
Fortune 500 IT hardware manufacturer*



*Network automation architect,
multi-billion-dollar gaming company*

“Building something in house comes at a cost. You need to have enough resources to be able to maintain it in an ongoing way. That does start building technical debt behind the scenes that does become difficult to solve over time.”



```
• 000101 0 0
• 00101000101 01 10 1
• 1101 1 1
• 01010101 1 0
• 010101010111110100001 0 1 1 0 0
• 00001010101010010100 0 0 1 0
• 010011000010 100 1 0
• 00101010101000 0 1 1 1
• 00101010101010 010101010 0 0 0 1 1
• 0010101010101 0 1 1 1 1
• 00101 0 1
• 010101010100101010 1 0 0 0 0 0
• 010101010100010010001000100 01 0 0 1 0 0
• 001010101010 00 1 0 0
• 01000010 0 1
• 010100101001000100010 0 0 0 0 0 0
• 010010010100010101010100111 1 1 0 00 1 0 010 0
• 001010100100 1 0 0 1 0
• 010010 010
• 0010101001001010 0 1 0 0 0 0 0
• 00100101010101 10 1 1
```

What Should You Demand from NSoT Vendors?

We asked interview subjects, who were mostly using open source and home-grown tools, to identify ways in which the emerging market of NSoT vendors could serve them better. Here is what they said.

Sophisticated Pre-Sales Engineering Support

NSoT tools are still new to most IT organizations. Engineers need help in demonstrating value, justifying an investment, and optimizing the overall solution. This high-touch service would also help network teams get up and running faster, since many struggle with populating their DIY NSoT with data.

Flexible and Extensible Data Model

Business requirements and technology transformation often shape the direction that network engineering teams take with their NSoT. They need solutions with flexible data models that can adapt to changes in how network teams use NSoT tools.

This flexible schema should allow network teams to add new fields of information and metadata, so that engineers can manage their networks in a way that meets their specific business requirements.



*Network operations manager,
\$1 billion logistics company*

“I need clear solutions presented by a pre-sales or solutions engineering team so we can present it to management and say these are the actual problems we’ve run into by trying to maintain this ourselves, and here is how this NSoT solves these problems.”



*Network engineer,
multi-billion-dollar media company*

“I need to have everything aggregated together. You’ve got your configuration data, but then your operational site data, so you can aggregate the business hierarchy, the regions, the sites, individual locations within a site, and then device data as well. If you had it all in the same database and you could relate those things, it would be powerful.”



*Network automation architect,
multi-billion-dollar gaming company*

“Ideally, you want a system that allows you to extend the schema so you can essentially evolve your network over time. A lot of companies are moving toward building out high-performance network fabrics for AI and GPU clusters. So, you need to be able to model your GPU connections to switches as well, things like that.”

Deep Integrations with Network Automation Toolsets

Homegrown and open source tools dominate today's network automation toolsets. Any source of truth should be ready to embrace and integrate with a variety of automation and orchestration tools. Engineers emphasized the importance of active support and engagement with Ansible and Terraform, for example.



*Network operations manager,
\$1 billion logistics company*

"We are looking for something that integrates very well with multiple automation platforms. Something that has a well-documented and up-to-date Terraform provider is a big push of ours."

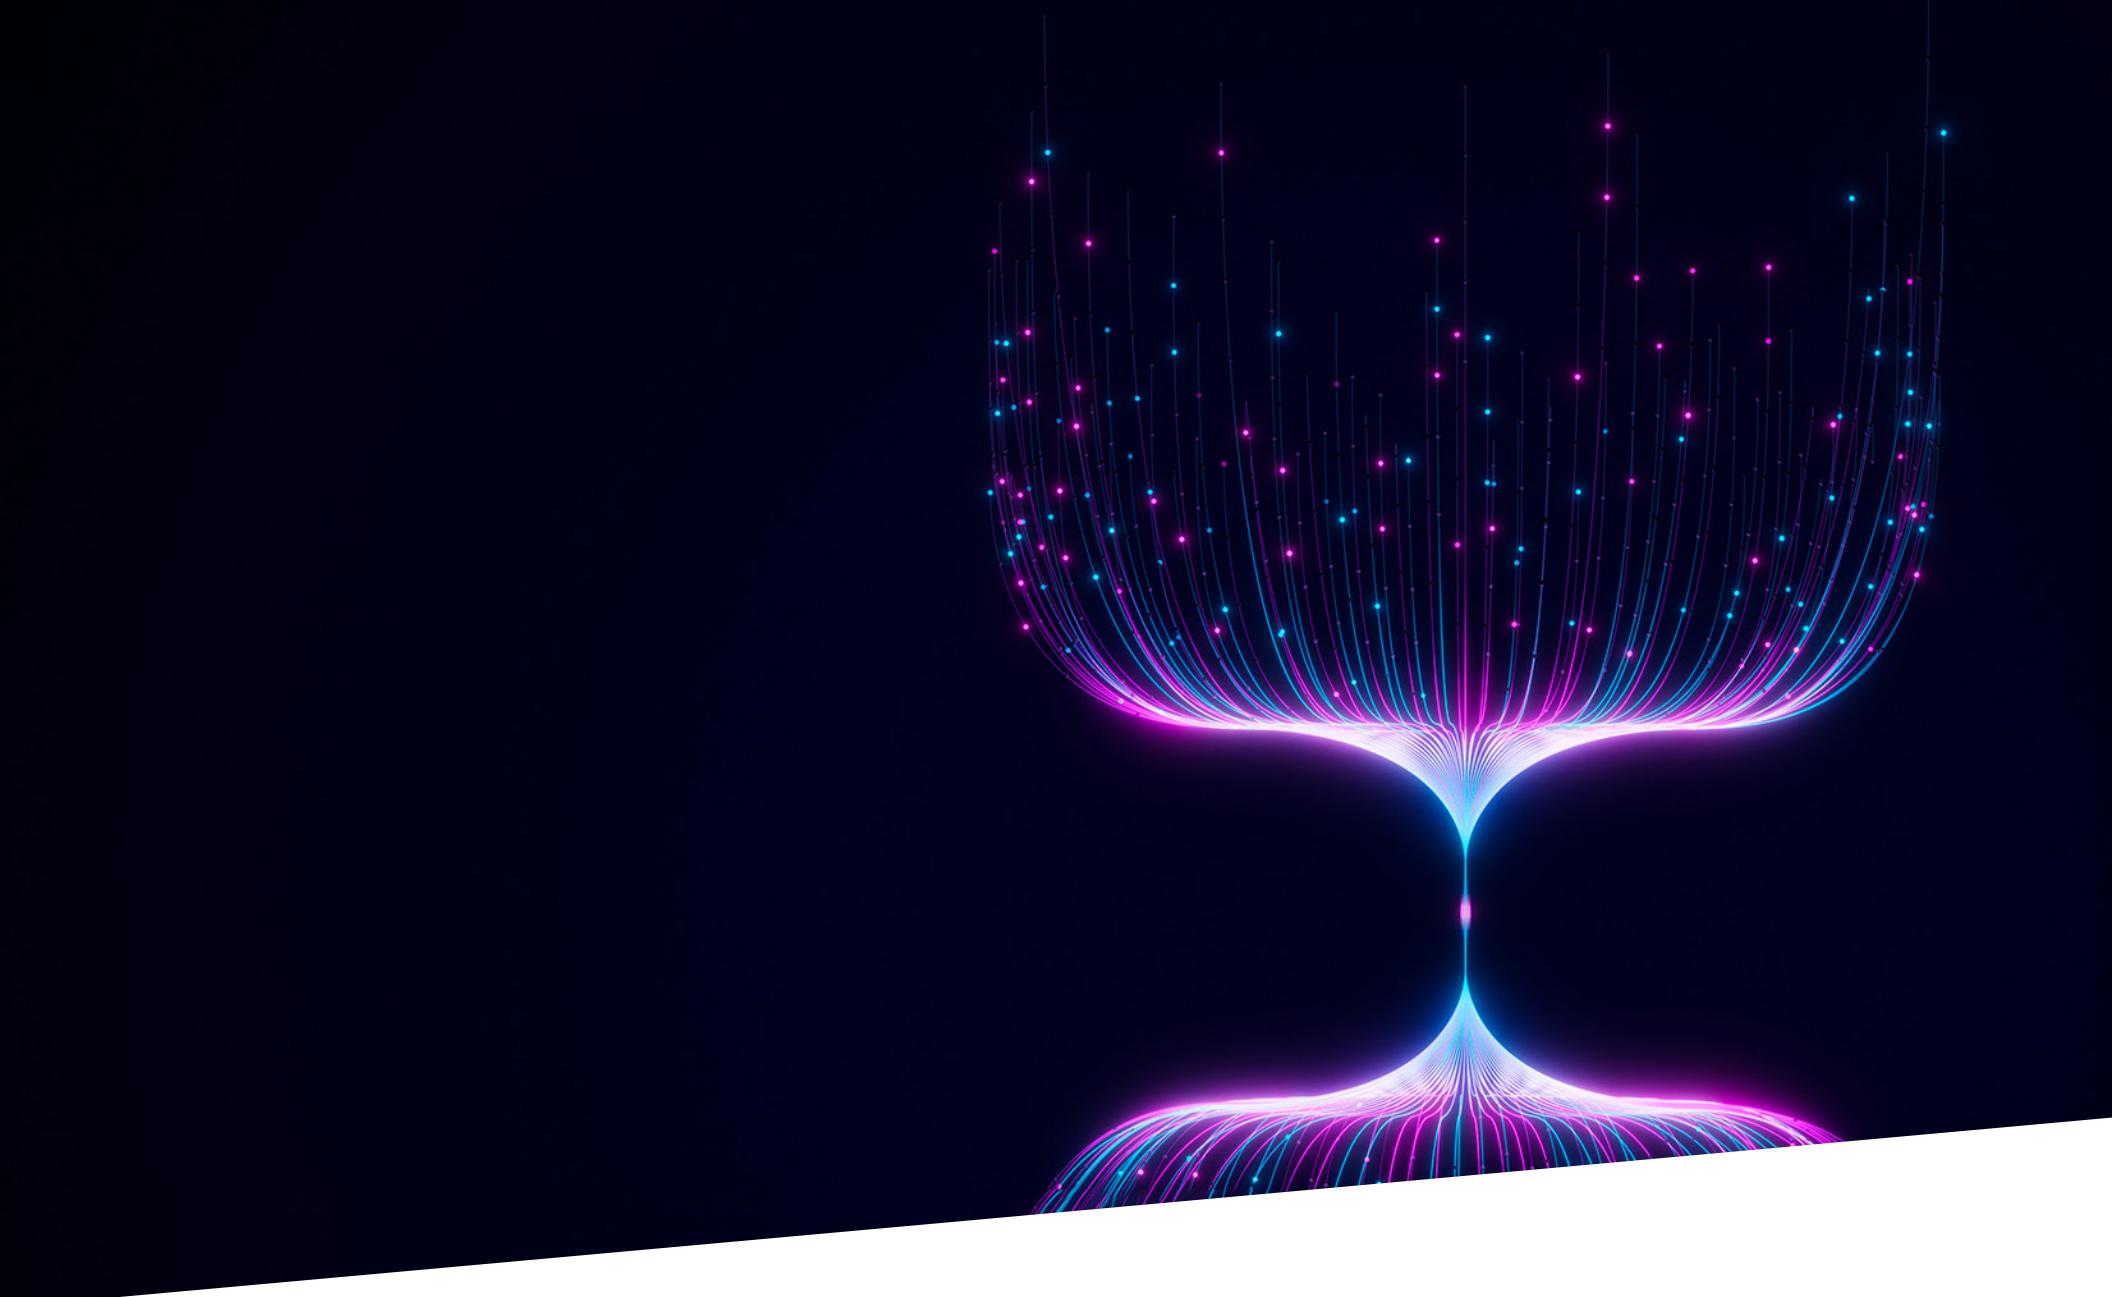
Event-Driven Operations

Network teams want an NSoT that understands what is happening on the network. This means it needs robust discovery and data ingestion features and some level of network observability for fault and performance management. This allows them to orchestrate all aspects of network operations through the NSoT.



*Network automation engineer,
large American university*

"How easy is it to do event-driven architecture? If the source of truth doesn't have web hooks, it is literally useless to us, because we need the instantaneous insight that a config changed. And if it's a source of truth and has some sort of observability via something like SNMP or even just ping, that would be nice."



Example NSoT Vendors

This section offers a non-exhaustive list of vendors that provide NSoT products or position themselves as NSoT solution providers. Notably, most full-stack network automation vendors (not listed here) also claim to offer NSoT capabilities within their overall platform, such as Gluware, Itential, and NetBrain.

BlueCat Networks

EfficientIP

Forward Networks

Infoblox

IP Fabric

NetBox Labs

Network to Code

OpsMill

BlueCat Networks

- Offers two DNS, DHCP, and IPAM (DDI) solutions, Integrity for large and complex networks and Micetro for smaller networks
- Positions its technology as a single source of truth for core network services, often integrated into an organization's network automation toolset

EfficientIP

- EfficientIP's SOLIDserver consolidates DNS, DHCP, IPAM, and all related network objects (VLANs, VRFs, devices, applications) into a single authoritative data model, delivering centralized visibility, lifecycle management, and consistent control across hybrid multi-cloud environments.
- Continuous discovery compares real-time network state with intended state to detect drift, while rich APIs integrate the NSoT into automation, security, cloud, DevOps, and ITSM workflows

Forward Networks

- Offers Forward Enterprise, a network digital twin platform that periodically discovers the current state of production networks and models it for network engineering and cybersecurity teams
- Often integrates with intent-based NSoT solutions to enable deep analysis of network intent versus current state

Infoblox



Infoblox is a network infrastructure automation company that pioneered the DDI (DNS, DHCP, IPAM) space more than 20 years ago. Universal Asset Insights extends Infoblox's authoritative IPAM foundation to deliver comprehensive network source of truth capabilities across hybrid and multi-cloud environments.

Foundation: Authoritative IPAM as Source of Truth

Infoblox's NSoT strategy uniquely leverages DNS and DHCP intelligence as the authoritative foundation for asset discovery. Unlike discovery tools that provide observational snapshots, Universal Asset Insights establishes IPAM as the single source of truth that infrastructure-as-code pipelines can trust. This prevents common automation failures from IP conflicts, stale records, and configuration drift affecting Terraform, Ansible, and ServiceNow workflows.

Comprehensive Discovery Architecture

The solution employs multi-layered discovery across:

Public Cloud. Native API integrations with AWS, Azure, and GCP for real-time asset synchronization

On-Premises Networks. NIOS-X probes with SNMP/CLI capabilities for network device discovery

Third-Party Sources. Integrations with Cisco Meraki, CrowdStrike, ServiceNow, and Microsoft Active Directory

Passive Auto-Discovery. Continuous monitoring that has inventoried 7M+ assets demonstrating enterprise scalability

Automated Asset Classification and Lifecycle Management

Universal Asset Insights categorizes assets into actionable intelligence: zombie (orphaned EC2 instances, detached volumes), idle resources, unencrypted assets with public exposure, and dangling DNS records. The platform implements a five-stage automation workflow: request, allocation (IPAM provides IP and auto-creates DNS), discovery, monitoring (continuous sync), and retirement (automatic remediation). This reduced one customer's monthly cloud waste by \$23,000 through automated cleanup.

Infoblox



Key Differentiators

Solving the Automation Challenge: Without accurate IPAM, infrastructure automation fails predictably. Universal Asset Insights demonstrated 25% IPAM accuracy improvement (73% manual to 98%+ automated), 93% reduction in IP conflicts (15-30 monthly to fewer than two), 99%+ Terraform success rates, and zero-touch cleanup of 1,200+ dangling DNS records creating security vulnerabilities.

Cross-Persona Unified View: The solution eliminates fragmented asset visibility across NetOps, CloudOps, and SecOps teams. Security investigations transform from “10.1.1.100” into actionable context: “EC2 t3.large in AWS us-east-1, DevOps team, payment processing, verified DNS, encrypted, patched 12 minutes ago.” This accelerates incident response and compliance auditing.

Enterprise Integration: Professional Services provides Getting Started Packages with documented integration guides. The platform supports active discovery (approximately 15 minutes for cloud) and continuous synchronization for accuracy.

Use Cases Addressed

- Enterprise Asset Source of Truth:** End IP conflicts across hybrid environments, eliminate manual tracking, establish reliable foundation for infrastructure automation
- Security Risk Intelligence:** Identify and remediate dangling DNS records, prevent subdomain takeover, eliminate exposures from decommissioned resources
- IT Asset Intelligence:** Continuously discover and reconcile assets, eliminate blind spots between tools and reality, maintain trustworthy compliance data

Deployment Model

Universal Asset Insights operates as part of the Infoblox Universal DDI platform, available as cloud-managed SaaS (NIOS-X as a Service) or on-premises appliances. It integrates with existing NIOS deployments or deploys greenfield. Deployment timelines range from 15 minutes for cloud discovery to several days for comprehensive on-premises environments.

IP Fabric

- Offers an eponymous platform that periodically discovers and models the current state of a network, essentially providing a digital twin
- Often integrates with intent-based NSoT solutions to enable deep analysis of network intent versus current state

NetBox Labs

- Leading contributor to open source NSoT tool NetBox, a hybrid of IPAM and DCIM with a powerful data model and open APIs for integration with automation solutions
- Offers cloud-based and on-premises enterprise versions of NetBox

Network to Code

- Creator of Nautobot, an open fork of NetBox
- Offers commercial network automation solutions and consulting services based on Nautobot

OpsMill

- Offers Infrahub, an open core NSoT product
- Combines the flexible data model of leading NSoT solutions with the version control features of Git-based solutions



EMA Perspective

NSoT tools are rising in prominence because network engineering teams have identified critical gaps in their toolsets and in how they maintain data about their networks. This report demonstrates that NSoTs are more than just repositories. They are foundational systems that enable scalable automation, consistent operations, and reliable documentation. By centralizing intent data in an NSoT and integrating it with automation and observability platforms, network teams can move beyond fragmented spreadsheets and outdated CMDBs toward a unified, programmatic approach to network management. However, success with NSoT requires more than just technology; it demands cultural change, strategic planning, and ongoing investment in data quality and tool integration.

The journey to a mature NSoT implementation is often long and complex, but the payoff is substantial: improved operational efficiency, enhanced troubleshooting, streamlined audits, and a foundation for NetDevOps and AI-driven network operations. By following a structured approach and aligning stakeholders around a shared vision, network teams can transform their infrastructure into a reliable, automated, and future-ready environment.

An NSoT Action Plan

Network teams that are just getting started with an NSoT plan or that are looking to revamp their current approach should take the following steps.

- **Identify and validate existing data sources:** Audit current repositories (e.g., spreadsheets, Visio diagrams, IPAM, CMDBs) and determine where authoritative network data resides.
- **Select or build a flexible NSoT tool:** Choose a solution with an extensible data model that can evolve with your network and business needs.
- **Clean and import data systematically:** Develop scripts and workflows to migrate data from legacy systems into the NSoT, ensuring accuracy and consistency.
- **Establish processes for ongoing validation:** Implement regular audits and reconciliation between the NSoT and live network configurations to maintain data integrity.
- **Integrate NSoT with automation and monitoring tools:** Use APIs to connect the NSoT with provisioning, configuration, and observability platforms for end-to-end automation.
- **Secure buy-in from IT executives and operations teams:** Educate stakeholders on the value of an NSoT and promote a culture that prioritizes centralized, authoritative network documentation. IT leadership may need to issue directives to keep teams aligned.



Case Study: Truma Accelerates Global Growth with Infoblox Universal DDI

Truma Gerätetechnik, a German manufacturer of RV and caravan accessories and equipment, has long focused on meeting the evolving needs of its customers. To execute, Truma needs to stay agile with a network that is adaptable, fast, reliable, and secure. This network must also align with Truma's cloud-first strategy, which enables IT to deploy services in the ways that make the most sense. The ability to securely scale mission-critical DNS and DHCP services and deliver reliable performance and visibility across the business remains a key part of Truma's cloud-first strategy.

Disparate Networking Technologies and Risk Exposure

Truma's network expanded rapidly through mergers and acquisitions. Running network operations seamlessly across new subsidiaries has become increasingly challenging for Truma's small network team, especially for foundational network services that keep these sites connected. Each newly acquired entity brings its existing DNS, DHCP, and IP address management (DDI) infrastructure and systems. Making them all speak the same language is not easy.

"We faced barriers between systems, which made it difficult to share information and collaborate across entities," said Andreas Schmidt, IT system engineer for Truma. "It became clear that we needed to centralize and standardize our networking infrastructure, so every location operates on the same foundation for IP services, DNS, and DHCP."

With limited resources and no centralized management system, maintaining consistency across so many environments posed a significant challenge for Truma's two-person central IT team. Bringing a new location online could take several days because the team manually built configurations. Each entity maintained its own local DNS and IP structure, increasing complexity and risk of configuration errors. In one instance, a DNS record tied to a production system was mistakenly deleted, halting operations for several hours until it could be restored.

The company was also running most of its DNS traffic through Microsoft DNS, Active Directory, and various domain controllers, which required considerable effort to manage both internal and external DNS. "There was no real way to streamline this; there was no dashboard to manage it," Schmidt said.

Truma also faced growing security demands. Each location operated its own legacy systems and configurations, creating inconsistencies and blind spots that made it difficult to maintain a unified security posture.

From Strong Foundations to Cloud-Ready Network Control

Truma standardized DDI management with Infoblox NIOS DDI, and the IT team gained the visibility, reliability, and automation they needed to manage services across its growing network. The NIOS deployment connected Truma's headquarters in Putzbrunn with four subsidiaries through the Infoblox Grid, providing redundancy and continuous availability of DDI services. For the first time, the IT team could manage internal and external DNS through a single interface and reduce the time spent on weekly updates from hours to minutes.

As Truma continued to grow, adding new subsidiaries across Europe and North America, the team recognized a need for a cloud-based DDI solution that could scale globally, streamline onboarding, and maintain centralized control from anywhere. It adopted the Infoblox Universal DDI™ (UDDI) platform, powered by NIOS-X virtual appliances. By adding UDDI, Truma could extend Infoblox's reliability to a cloud-managed architecture, unifying DDI services across all entities.

The hybrid design, featuring 10 NIOS-X virtual appliances running on VMware ESX and Microsoft Azure, delivers low-latency local performance alongside centralized, cloud-based control. Now, Schmidt's team can manage critical network services at any location using UDDI, through a single API and the Infoblox Portal.

Simplicity, Speed, and Scale for a Global Network

The migration to UDDI unlocked a new level of simplicity and speed. “With Microsoft DNS, even small configuration changes could take hours, sometimes days,” Schmidt said. “Now, with Universal DDI, we can make those changes in just a few minutes. Standardizing with Universal DDI let us scale globally, onboarding 10 entities across seven countries, without local partners or additional headcount,” Schmidt said. “That would have been impossible with our old systems.”

.....

Standardizing with Universal DDI let us scale globally, onboarding 10 entities across seven countries, without local partners or additional headcount,” Schmidt said. “That would have been impossible with our old systems.”

Built-in automation and rollback capabilities give Schmidt’s team the confidence to adapt fast and recover instantly, even during large-scale onboarding. “We can make changes quickly and restore configurations within minutes, without downtime or data loss.”

Universal DDI has also brought robust visibility. Every device, DNS query, and connection is visible from one portal, eliminating the blind spots that once led to downtime. “It’s much safer to make changes in production,” Schmidt explained. “We no longer have to guess whether a change might interfere with other systems.”

Truma also leverages Infoblox Threat Defense™, which is fully integrated into the UDDI platform, to protect users and devices across every location. It provides DNS-layer protection and real-time insights across all users and locations. “We can just click another button and see the DNS security view in the same portal,” Schmidt said. “With everything connected in one platform, the team can spot and respond to issues faster. We don’t need to search for hours to determine exactly what happened.”

With Universal DDI and Threat Defense working together, Truma created a modern, cloud-managed foundation that scales effortlessly while keeping its network secure and resilient. “People across Truma no longer have to think about how things work on the network; they can focus on what they do to satisfy our customers,” Schmidt said. “Infoblox works as a hidden champion in the background every day.”





About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com or follow EMA on [X](#) or [LinkedIn](#).

This report, in whole or in part, may not be duplicated, reproduced, stored in a retrieval system or retransmitted without prior written permission of Enterprise Management Associates, Inc. All opinions and estimates herein constitute our judgement as of this date and are subject to change without notice. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. "EMA" and "Enterprise Management Associates" are trademarks of Enterprise Management Associates, Inc. in the United States and other countries.

©2025 Enterprise Management Associates, Inc. All Rights Reserved. EMA™, ENTERPRISE MANAGEMENT ASSOCIATES®, and the mobius symbol are registered trademarks or common law trademarks of Enterprise Management Associates, Inc.