Cisco Network Assurance Engine Fundamentals Guide, Release 2.1(1)

Table of Contents

First Published: 2018-07-25

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

The Cisco Network Assurance Engine

Overview

The Cisco Network Assurance Engine (NAE) software provides operators with a new approach to manage SDN-based data centers confidently. The Cisco NAE software is built on a comprehensive formal model of the network, combined with deep domain knowledge of networking. The Cisco NAE software provides operations teams with continuous and proactive network verification and intent assurance.

Business drivers such as cloud, mobile, and digitization trends are demanding more from modern data centers, rapidly increasing their scale, rate of change, and complexity. With the Cisco Application Centric Infrastructure (ACI) and other SDN technologies, network infrastructures have evolved to provide programmable interfaces, automation, agility, and virtualization. However, operational tools still center around traditional approaches, such as probe tools, packet sniffers, and the command line interface (CLI) to reason about the network. These are inherently reactive-after-the-fact, manual, and rely on the tribal knowledge of a handful of experts to reasonably reconstruct a network state.

The Cisco NAE software takes the intent from the controller as a logical policy, as well as configurations and the data plane (infra) state from each switch device, to build a network-wide model of the underlay, overlay, and virtualization layers.

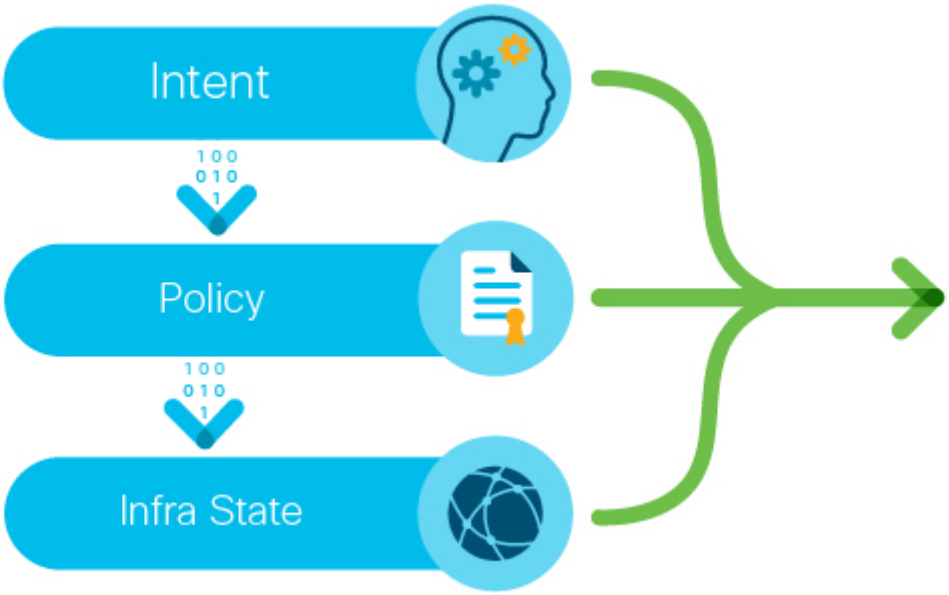

Figure 1. Taking the intent, policy, and the infra state from a device

Figure 2. Network-wide model of the underlay, overlay, and virtualization layers

Leveraging formal mathematical techniques and a deep understanding of the networking domain, the Cisco NAE software is able to answer three fundamental questions about the network:

1. How do I guarantee that I have not introduced errors into the fabric while specifying my policy and configuration?
    - In SDN networks, the impact of a misconfiguration is amplified with centralized automation
    - With increased frequency of changes, misconfigurations are much more common
2. How do I understand the actual current state of the network?
    - Dynamic events (protocol learning, VM mobility, and so on) can make the network deviate from the intended state
    - Central intent has to be propagated to multiple nodes in a large distributed system; consistency issues are unavoidable in such systems
    - With multiple abstraction layers in SDN networks, manually inspecting and reconstructing the network state and configuration has become prohibitively challenging
3. How do I rapidly diagnose the network for the root cause when issues arise?
    - How do I identify these lingering issues before they impact the application?
    - How do I reduce the cost of downtime?

By proactively running this broad set of checks on the network model and providing deep visibility across the fabric, the Cisco NAE software transforms the operating mode from reactive to proactive. The Cisco NAE software enables operators to predict network outages and vulnerabilities before they impact business, reduce risk while accelerating changes and migrations, and rapidly find the root cause of problems. With a complete diagnostic record and compliance rules, operators can ensure continuous compliance and easily satisfy audits.

Assurance Provided by the Cisco Network Assurance Engine

The Cisco Network Assurance Engine provides assurance for the following areas:

- Change management—You can analyze real-time changes to your fabric with the online analysis mode, or one-time changes with the offline analysis mode. In either case, you can determine when something changed and what changed in your fabric during the specified time interval.
- Tenant endpoints—You can analyze connectivity issues with tenant endpoints, including the number of endpoints with issues and which endpoints that have the most severe issues.
- Tenant forwarding—You can analyze issues in your fabric regarding tenant forwarding. The information includes the quantity of tenants with forwarding issues and the severity of the issues.
- Tenant security—You can analyze issues with tenant security, which includes determining if deny or permit security policies have been violated and suggestions for resolving the violations.
- TCAM optimization—You can analyze the amount of ternary content-addressable memory (TCAM) utilization in your fabric. You can determine where your fabrics resources are being used the most and least, which can help you to balance the resource load or determine if you can remove a TCAM rule from a leaf switch.

Assurance Control Modes

The Cisco Network Assurance Engine assurance control capability enables you to analyze the assurance group in two modes: online analysis and offline analysis. An assurance group contains of all the network nodes that should be analyzed together. Assurance control involves collecting data from the assurance group, running the analysis to create a model with the collected data, and generating the results. The results are then displayed on the **Dashboard**.

Online analysis provides assurance on the assurance group in real-time. With online analysis, data collection, model generation, and results generation are carried out simultaneously. The collected data is analyzed immediately after collection, followed by result generation. This is repeated after at a fixed time interval as specified by the operator.

Offline analysis provides a one-time assurance of the assurance group. Offline analysis offers the flexibility of decoupling the data collection stage from the analysis stage. Data is collected using a Python script and the collected data is then uploaded to the Cisco NAE to provide one-time assurance. The collected data can also be analyzed at any later point in time. Offline analysis enables you to collect the data during change management windows and then perform the analysis.

Offline Data Collection Script

The Cisco Network Assurance Engine offline data collection script is a Python script that polls the Cisco Application Policy Infrastructure Controllers (APICs), spine switches, and leaf switches for a series of REST API and CLI calls. For information about the REST API calls and CLI calls, see the

readme.md file that is included with the script.

The script has the following dependencies:

- Python 2.7.11+
- Ubuntu/OS X /Cent OS
- Python dependencies
  - Requests (Python REST library)
  - Paramiko (Python SSH library)
  - Setuptools (Python packaging library)

See the readme.md file for information on the Python dependencies and the process to install the dependencies in a virtual environment. The readme.md file provides the complete list of objects and show commands collected from the Cisco APIC, spine switches, and leaf switches. The readme.md file is available inside the same zip file with the offline analysis script file. The offline analysis script is downloadable directly from the Cisco NAE appliance from the settings menu.

The workstation on which the script is being launched must have out-of-band management connectivity to the Cisco APICs, leaf switches, and spine switches. Make sure that every node in the Cisco ACI fabric has an out-of-band management IP address configured. Make sure that the firewall does not block HTTPS (for using the REST API) and SSH (for connecting t the leaf switches and spine switches). Make sure that the proxy settings are properly set to allow HTTPS connections.

The readme.md file provides the syntax for using the script. By default, the script will run 3 iterations of the data collection at a 5 minute interval between iterations, although you can specify the number of iterations by using the **-iterations** option. The total expected collection time ranges between 18 to 20 minutes from start to finish for 3 epochs for a fabric with around 20 leaf switches. Larger fabrics will take longer time depending on complexity of the configuration and scale of the fabric.

Change Management

Change Management Tab

The **Change Management** tab provides information about the assurance on changes in your fabric.

Real-time Change Analysis

The **Real-time Change Analysis** inspector page provides information about issues that are caused by changes made to the network. On this page, you can quickly see how many policy violations of all types exist.

One common use case for this inspector page is to resolve the following network issue:

- Determining Configuration Issues That Prevent Network Migration

Determining Configuration Issues That Prevent Network Migration

When you migrate to a Cisco Application-Centric Infrastructure (ACI) environment, typically you pre-provision all of the required configurations that relate to tenants, VRF instances, bridge domains, and endpoint groups. You then program the correct policies required for Layer 1 and Layer 2 connectivity for server ports. This process includes the following steps:

1. Configure the interface policies, interface policy groups, interface profiles, leaf switch profiles, attachable access entity profiles, VLAN pools, physical and virtual domains, and so on.
2. Create the tenants, VRF instances, bridge domains, application profiles, endpoint groups, static path bindings, and so on.
3. Extend Layer 2 connectivity from the old infrastructure to the new Cisco ACI environment and begin moving workloads into the Cisco ACI fabric. Subsequently, during a cutover, the default gateways for multiple VLANs are moved from the old infrastructure to Cisco ACI.

As Cisco ACI starts to take over the default gateway functionality, several Cisco ACI configurations take effect, such as DHCP relay, contracts, and inter-subnet routing. The accurate configuration of contracts are key to providing connectivity between endpoint groups in the fabric and to entities outside of the fabric.

In several cases, migrations have been been delayed, postponed, or canceled due to human errors even after multiple reviews with the change advisory boards. The Cisco Network Assurance Engine allows operators quickly to pinpoint configuration issues, both within the ACI fabric and outside the fabric, by using the **Real-time Change Analysis** inspector page. This helps reduce the time taken to perform migrations and the costs associated with multiple change windows in the event that things do not work as expected.

The following list specifies some of the common network migration issues that the **Real-time Change Analysis** inspector page can help you to troubleshoot:

- After you extend the Layer 2 connectivity from the old infrastructure to the new Cisco ACI

environment, you attach a server in the new Cisco ACI fabric. However, the server cannot ping its gateway in the old infrastructure through the Layer 2 extension.

- After you migrate a gateway from the old infrastructure to the new Cisco ACI environment, you observe intermittent ping loss to hosts in a few subnets, while other hosts in other bridge domains and subnets have no connectivity.

In the first case, you can discover the cause of the issue by performing the following actions on the **Real-time Change Analysis** inspector page:

1. In the Hot Topics area, change the right dashlet to **Top EPGs by Policy Issue Subcategory**.
2. In the dashlet, click the forwarding issue count for the endpoint group that contains the endpoint (the server) with the issue.
3. In the Smart Events area, view the smart events that relate to the appropriate endpoint group and perform the suggested next steps.

Verify & Diagnose

Verify & Diagnose Tab

The **Verify & Diagnose** tab provides assurance information for tenant endpoints, tenant forwarding, and tenant security.

Tenant Endpoints

The **Tenant Endpoints** inspector page provides information about connectivity issues with tenant endpoints. On this page, you can quickly see how many endpoints have issues and which endpoints have the most severe issues. Endpoints can have multiple issues; they are not limited to one issue each.

Some of the common use cases for this inspector page are as follows:

- Diagnosing endpoint learning errors, such as:
  - Local Station Table (LST) or Council of Oracle Protocol (COOP) mismatches
  - Incorrect Global Station Table (GST) entries
  - VPC peer synchronization issues
- Finding endpoint IP addresses that are not in the bridge domain subnet
- Finding endpoints with a duplicate IP addresses
- Determining if an endpoint is unreachable due to an endpoint learning error

Determining Incorrect GST Entries

In Cisco APIC, the Global Station Table (GST) contains information about remote endpoints on the leaf switch. Remote endpoint learning helps Cisco APIC forward packets more efficiently by allowing leaf switches to send packets directly to a destination leaf switch without using the resources on the spine switch to lookup endpoints on the COOP database.

Starting from Cisco NAE Release 2.1(1), Cisco NAE checks for all the GST entries in the leaf switches and determines the incorrect GST entries. Cisco NAE checks if the GST entry is pointing to the correct leaf switch where the end point resides and if there any stale GST entries. An incorrect GST entry may cause momentary traffic disruption. It also accounts for bounce entries before determining the incorrect GST entries.

You can use the **Tenant Endpoint** inspector page to determine the cause of the issue. In the Smart Events area, view the `CONNECTED_EP_LEARNING_ERROR` smart event that lists the leaf switches where the EP is present on the GST table and perform the suggested next steps.

Tenant Forwarding

The **Tenant Forwarding** inspector page provides information about issues with tenant forwarding. On this page, you can quickly see how many tenants have forwarding issues and which tenants

have the most severe issues.

Some of the common use cases for this inspector page are to resolve the following network issues:

- Overlapping subnets across external interfaces for a bridge domain in a VRF instance
- Overlapping subnets across external learned routes for a bridge domain in a VRF instance
- Overlapping subnets across external endpoint groups for a bridge domain in a VRF instance
- Overlapping subnets across VRF instances due to a contract
- Tenants that have a mismatch in the internal subnet routing information base (RIB) and forwarding information base (FIB)

Tenant Security

The **Tenant Security** inspector page provides information about issues with tenant security. On this page, you can quickly see how many tenants have policy violations and which tenants have the most severe issues.

Some of the common use cases for this inspector page are to resolve the following network issues:

- Resolving deny policy violations
- Resolving permit policy violations
- Resolving deny log traffic policy violations
- Resolving permit log policy violations
- Determining why endpoint groups cannot communicate

Determining Why Endpoint Groups Cannot Communicate

In this use case, two endpoint groups cannot communicate. You can use the **Tenant Security** inspector page to determine the cause of the issue. Take as an example tenant "Finance" with endpoint groups "App" and "DB" that cannot communicate with one another.

In the visualization area, use the filters and view control to narrow down the radial display to make it easier to find the endpoint groups with the issue. See if the endpoint groups have an arrow pointing from one to the other. If they do not, then the application profile is not configured to allow the endpoint groups to communicate. Reconfigure the application profile to allow the endpoint groups to communicate.

Otherwise, click on the arrow the connects the endpoint groups. The radial display changes to show you the contracts that are intended to enable the endpoint groups to communicate. In this example, App is the consumer while DB is the provider. This view gives you additional information about possible communication issues between the endpoint groups.

In the smart events area, find the smart events that pertain to the endpoint groups with the issue—App and DB in the example. Expand the events to get more information about the issue and suggested resolutions.

Epoch Analysis

Epoch Analysis Tab

The **Epoch Analysis** tab provides information about the state of the fabric between two epochs.

An epoch is a period of time in your network's history during which the Cisco NAE collected and analyzed data. The size of the epoch gives a rough indication of the quantity of smart events at that time, with a larger epoch indicating more smart events. Epoch data is collected in 15 minute intervals.

As you make changes to your network, **Epoch Delta Analysis** enables you to compare two epochs to determine what changes occurred, where the changes occurred, and the resources that were affected by the changes.

In the **Health Delta** view you can determine the changes in the health of the fabric. For example, when you are making changes to the fabric in your Change Control Window, you can compare the epoch after each change to the epoch prior to the change. The **Smart Event Count** provides a summary view of the changes. The **Health Delta By Resources** indicates the Cisco APIC resources that were impacted by the change. The **All Smart Events** indicates the impact of the changes. You can determine if new events were raised or if the same event has been raised with a different failing check.

In the **Policy Delta** view you can visualize the changes made to a policy. For example, when you add a new tenant, or make changes to BDs or EPGs in a policy you can verify the changes in the policy using the **Policy Delta Visualization** before you implementing the changes.

Some of the common use cases for this inspector page are as follows:

- Site preparation and migration of workload
- Change control
- Maintenance upgrades
- Capacity management
- Fabric improvement

Optimize

Optimize Tab

The **Optimize** tab provides information about the resource utilization in the network.

TCAM

The **TCAM** inspector page provides information about the amount of ternary content-addressable memory (TCAM) utilization.

For switches and line cards that use hash tables, if a filter has a large number of port ranges specified, the number of concrete actrl rule objects and zoning rules match, but the Hardware Abstraction Layer (HAL) output does not have all of the hardware entries that are present in zoning rules. Because of this, smart events show the actual HAL usage, but HardwareStats show the zoning rule usage. The N9K-C93180LC-EX, N9K-C93108TC-EX, and N9K-C93180YC-EX top-of-rack switches and the N9K-X97160YC-EX, N9K-X9732C-EX, N9K-X9732C-EXM, and N9K-X9736C-EX line cards use hash tables.

Some of the common use cases for this inspector page are as follows:

- Determining the leaf switches that use the most TCAM
- Determining the policy distribution across leaf switch TCAMs
- Determining the least used TCAM rules (security policies), based on their hit count

Determining the Leaf Switches That Use the Most TCAM

The **TCAM** inspector page's visualization area enables you to view the leaf switches that have the most TCAM being used (top leafs by TCAM usage). This data is provided in an x-y graph. The leaf switches are listed on the x-axis in descending order of usage from left to right. The y-axis indicates the quantity of TCAM being used on the switches. You can use the filter field to view the data for a subset of the switches.

You can use this information to determine if any switches have too much TCAM being used. In such a case, you can modify your TCAM rules to balance the TCAM usage.

Determining the Policy Distribution Across Leaf Switch TCAMs

The **TCAM** inspector page's visualization area enables you to view the policy distribution across leaf switch TCAMs. This data shows you which tenants and endpoint groups are consuming the most contracts, which contracts are consuming the most TCAM, and which filters are constantly being hit.

You can use this information to determine where you can optimize your network's security posture and where can you recover the most TCAM resources. The data enables you to determine the top

consumers of the network policies, after which you can optimize the appropriate contracts in the Cisco Application Policy Infrastructure Controller (APIC).

Determining the Least Used TCAM Rules By Hit Count

The **TCAM** inspector page's summary area provides the least used TCAM rules (security policies) based on how many hits the rules had over time. The table of this information is sorted in the order of the least number of hits, then by the highest TCAM utilization of each TCAM rule.

You can use this information to determine if there are any TCAM rules that are never hit, meaning that they are unused. In such a case, you can close any ports that are associated with those TCAM rules to improve the security posture of your network.

The summary area also specifies the TCAM utilization of each TCAM rule, which enables you to know how much TCAM resources you can regain on the relevant leaf switch by removing a TCAM rule.

Smart Events

Smart Events Tab

The **Smart Events** inspector page provides a table of all of the smart events that have been triggered in the currently-chosen epoch. A smart event provides information about the state of your network at the time represented by the epoch. A smart event can be informational, meaning that the smart event is only informing you that something happened, such as that the Cisco Network Assurance Engine successfully logged into a Cisco Application Policy Infrastructure Controller (APIC) cluster. A smart event can also warn you of an issue with your network in varying degrees of severity, such as the critical issue of a bridge domain subnet that is missing from a leaf switch. The smart events are organized by category, subcategory, severity, and name.

A smart event contains the following information:

- Description—A description of the smart event.
- Impact—The negative impact that the event has on your fabric.
- Affected Objects—The objects in your fabric that are affected by the issue.
- Checks—The reason for the event and suggested steps to resolve the issue.
- Event ID/Code—The ID and code of the smart event.