



Cisco Network Assurance Engine Fundamentals Guide, Release 3.1(1)

Table of Contents

The Cisco Network Assurance Engine	3
Overview	3
Assurance Provided by the Cisco Network Assurance Engine	5
Assurance Control Modes	5
Offline Data Collection Script	5
Change Management	7
Change Management Tab	7
Real-time Change Analysis	7
Determining Configuration Issues That Prevent Network Migration	7
Verify & Diagnose	9
Verify & Diagnose Tab	9
Tenant Endpoints	9
Determining Incorrect GST Entries	9
Tenant Forwarding	9
Tenant Security	10
Determining Why Endpoint Groups Cannot Communicate	11
Epoch Analysis	12
Epoch Analysis Tab	12
Optimize	13
Optimize Tab	13
TCAM	13
Determining the Leaf Switches That Use the Most TCAM	13
Determining the Policy Distribution Across Leaf Switch TCAMs	13
Determining the Least Used TCAM Rules By Hit Count	14
Compliance	15
Compliance Tab	15
Compliance Analysis Tab	15
Manage Compliance Requirements Tab	15
Guidelines to Satisfy Compliance Requirements	17
Smart Events	19
Smart Events Tab	19
Use Cases	19
Smart Event Dashboard	19
Manage Event Rules	20
Current Epoch Event Rules Snapshot	20

First Published: 2019-03-31

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 SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017-2019 Cisco Systems, Inc. All rights reserved.

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.

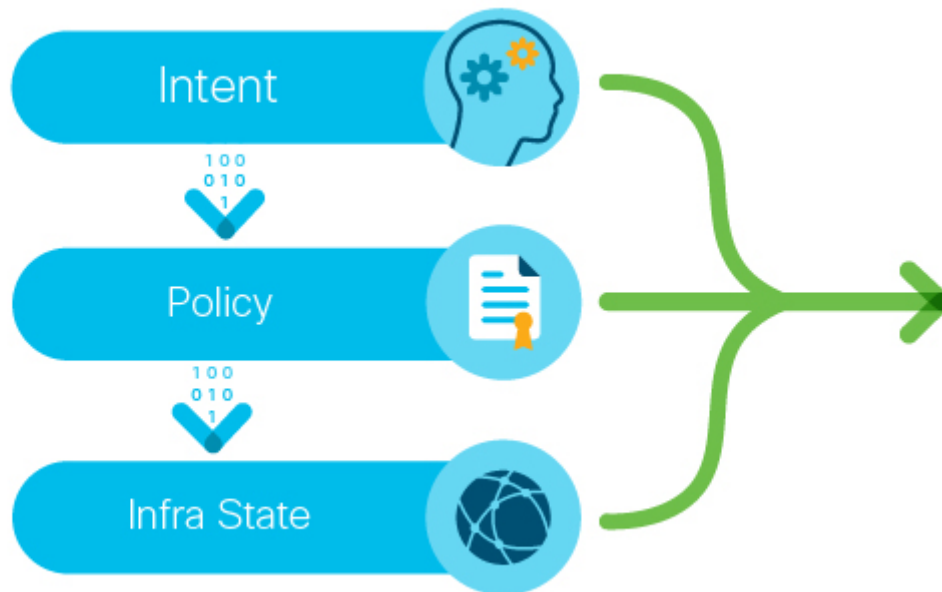


Figure 1. Taking the intent, policy, and the infra state from a device



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 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 to 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 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 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)

Additional use cases related to **Tenant Forwarding Visualization** are as follows:

In the **Prefix Communication View** visualization, you can view the state of forwarding between prefixes across the entire fabric. You can also view whether the prefixes are allowed to communicate by policy. The visualization helps operations teams determine the following:

- If a prefix communication is healthy or unhealthy.
- If the communication between prefixes is within the fabric or external to the fabric.

In the **Physical Interface View** visualization, the operations teams focus on the interface(s) that need immediate attention by getting a view of the interface status for a specific leaf switch based on the intended use of the interface and the event severity. This view helps operators determine the following interfaces:

- Used Interfaces - These interfaces are already in use by EPGs. EPGs are impacted when a Used Interface is operationally or administratively down.
- Configured Interfaces - These interfaces are ready to be used by EPGs.
- Partially Configured Interfaces - These interfaces are not ready for use by EPGs.

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 that 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.

Compliance

Compliance Tab

The **Compliance** tab provides comprehensive information about the state and health of the IT infrastructure. Assurance data helps IT operators get detailed visibility into the functioning state of various components and act on exceptions that are impacting security, performance, capacity, availability, and configuration of the IT infrastructure.

The assurance data generated for a particular time period allows you to determine if there are any issues in the network or determine if the issues generated during an earlier time period were resolved.

There are two sub-tabs under this tab in the GUI.

- Compliance Analysis
- Manage Compliance Requirements

Compliance Analysis Tab

The **Compliance Analysis** tab enables the user to verify compliance analysis results when an epoch runs for an assurance group.

When an epoch runs for a specific assurance group, all the active requirement sets that are associated with that assurance group are verified and validated.

Manage Compliance Requirements Tab

The **Manage Compliance Requirements** tab enables the user to verify Segmentation Compliance and Service Level Agreement (SLA) compliance, and traffic restriction compliance. Compliance can be used to set up regulatory compliance rules. With segmentation compliance, the user can establish walled areas around a set of entities that must not communicate with other entities. SLA compliance can also set up rules for entities that must talk with other entities. Traffic restriction compliance requirements allow the user to specify restrictions on protocols and ports for communication between EPGs.

In the NAE UI, the user specifies their compliance requirements. The NAE appliance, verifies in the subsequent epochs, whether the compliance requirements are satisfied by the policy that is configured on Cisco APIC. If satisfied, an event is raised stating that the compliance requirement is satisfied. One event per requirement per epoch is raised. For example, if an assurance group runs a compliance analysis on an epoch every 15 minutes, and there are two requirements associated with the epoch, two smart events will be raised.

Manage Compliance Requirements Compliance Requirements are managed in two areas as follows:

Define Compliance Requirements

Define compliance requirements by creating named objects such as Compliance Requirement Sets, Compliance Requirements, and EPG selectors.

EPGs are grouped in **EPG selectors** that are container objects that describe a collection of EPGs. They are configured with selectable attributes. You include and exclude the selection criteria of attributes when defining an EPG selector.

Traffic selectors are used for SLA and traffic restriction compliance. They are not used for segmentation. The SLA compliance requirement is verified only for traffic that matches criteria specified in the traffic selector.

EPG selectors are used in a **Compliance Requirement**. A compliance requirement must be configured. Compliance requirements can be verified in both directions if desired. For example, EPGSelectorA must not communicate with EPGSelectorB. And EPGSelectorB must not communicate with EPGSelectorA.

Compliance Requirements can be grouped into **Compliance Requirement Sets**.

Create Compliance Requirement Sets, and associate specific Compliance Requirements with that set. Compliance Requirement Sets must be associated with an assurance group or with multiple assurance groups. Compliance Requirement Sets contain multiple Compliance Requirements that can be activated or deactivated.

Examples of Compliance Requirement Sets

When you run a Compliance Analysis on an epoch in a particular assurance group, it runs all the Compliance Requirement Sets for that assurance group. The Compliance Requirement Sets can be activated and deactivated as required. If you modify a Compliance Requirement, the new Compliance Requirement goes into effect in the next epoch.

A **Segmentation Compliance** example is as follows:

- EPGSelectorA must not talk to EPGSelectorB. All traffic between EPGs of EPGSelectorA and EPGSelectorB must be denied.

An **SLA Compliance** example is as follows:

- EPGSelectorA must talk to EPGSelectorB on TrafficSelector1. Traffic between EPGSelectorA and EPGSelectorB that matches TrafficSelector1 must be permitted.

Traffic Restriction Compliance examples are as follows:

- EPGSelectorA must not talk to EPGSelectorB on TrafficSelector2. Traffic between EPGSelectorA and EPGSelectorB that matches TrafficSelector2 must be denied.
- EPGSelectorA may talk to EPGSelectorB only on TrafficSelector3. Traffic between EPGSelectorA and EPGSelectorB that does not match TrafficSelector3 must be denied. Traffic that matches TrafficSelector3 may be permitted or denied (It does not matter).

Analyze Compliance Results

Compliance analysis results in raising one event per Compliance Requirement per epoch. Each

event displays one of the following three results:

- Satisfied
- Violated
- Not checked (EPG selector does not match any EPG)

After a compliance analysis is completed, every epoch displays a list of the requirements that are violated and a list of the compliance requirements that are enforced.

NAE provides continuous compliance.

Use Cases

A common use case for Segmentation Compliance is as follows:

- Segmentation compliance verifies the regulatory requirement that a set of two EPGs are segmented.
- SLA compliance verifies availability of services. For example, an SLA compliance requirement can ensure connectivity of an EPG to the DNS EPG.
- Traffic restriction compliance verifies security requirements. For example, a traffic restriction compliance requirement can verify deny rules for traffic to unencrypted ports (such as a telnet port).

Network Compliance Assurance Solution Requirement Attributes

The compliance assurance solution supports the following requirement attributes:

- Comprehensively covers security
- Provides compliance for all requirements associated with a fabric when an epoch runs for that fabric.
- Understands tenant security configurations (contracts, VRFs)

Guidelines to Satisfy Compliance Requirements

For a contract to satisfy an **SLA compliance** requirement, the direction of ports in the filter should be the same as the direction of ports in the traffic selector of the SLA requirement. For example, if an EPG from EPG selector A is the consumer of the contract and EPG from EPG selector B is the provider, the ports of the filter in the filter chain from consumer to provider must be the same as those in “A to B” direction of the traffic selector. Similarly, if an EPG from EPG selector A is the provider of the contract and EPG from EPG selector B is the consumer, the ports of the filter in filter chain from provider to consumer must be same as those in “A to B” direction of the traffic selector.

For a contract to violate a **Traffic Restriction** compliance requirement, the direction of ports in the filter should be the same as the direction of ports in the traffic selector of the traffic restriction requirement. For example, if an EPG from EPG selector A is the consumer of the contract and EPG from EPG selector B is the provider, the ports of the filter in the filter chain from consumer to provider must be the same as those in “A to B” direction of the traffic selector. Similarly, if an EPG

from EPG selector A is the provider of the contract and EPG from EPG selector B is the consumer, the ports of the filter in filter chain from provider to consumer must be same as those in “A to B” direction of the traffic selector.

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** inspector page also enables you to suppress smart events in the Cisco NAE UI and view only the smart events that are relevant.

Use Cases

Some of the common use cases for smart event suppression are as follows:

- An administrator can suppress smart events that were generated for known issues so that the administrator can focus on more important smart events. The administrator can address the suppressed smart events after resolving the more important smart events.
- The network administrator can suppress all Tenant End Point Smart Events for all of the epochs during a change window while the network administrator and the server administrator both make changes to their respective infrastructure pieces.
- The network administrator can suppress all smart events in the category of Real Time Change Analysis and System Assurance, when the network administrator is undertaking a massive change in the fabric. During this time period, the network administrator has moved all of the traffic to the disaster recovery data center.
- The network administrator can suppress all smart events in a certain category except the one smart event that is of interest to the network administrator. For example, suppressing all TCAM Utilization smart events except when TCAM Utilization has reached the Critical level.
- The network administrator can view the list of all of the smart events that have been suppressed for a given epoch. The information can be used to determine if any of the smart events should be removed from the smart event suppression list.
- The network administrator can suppress all smart events containing an object, such as a VRF instance, while the network administrator makes a change to that object.

Smart Event Dashboard

In the **Smart Event Dashboard** tab ,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 smart event has on your fabric.
- **Affected Objects**—The objects in your fabric that are affected by the issue. The primary affected objects are highlighted.
- **Checks**—The Passing or Failing checks performed on the smart event and suggested steps to resolve the issue. Every passing or failing condition has a check code associated with it. The same check code may be used for a passing or failing condition and may be reused across smart events with different event codes.
- **Event ID/Code**—The ID and code associated with the smart event.



If the event suppression feature is activated, the smart event count and the smart events listed take into account the event rules.

Manage Event Rules

The **Manage Event Rules** tab enables you to create and manage event rules and event rulesets for suppressing smart events.

Current Epoch Event Rules Snapshot

The **Current Epoch Event Rules Snapshot** tab enables you to view the event rules for the current epoch selected in the timeline. Click **Smart Events Count** to view the smart events that match the event rule.