Cisco Network Assurance Engine Fundamentals Guide, Release 4.1(1)

Table of Contents

First Published: 2020-02-10

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

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 figuresincluded in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbersin illustrative content is unintentional and coincidental.

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.

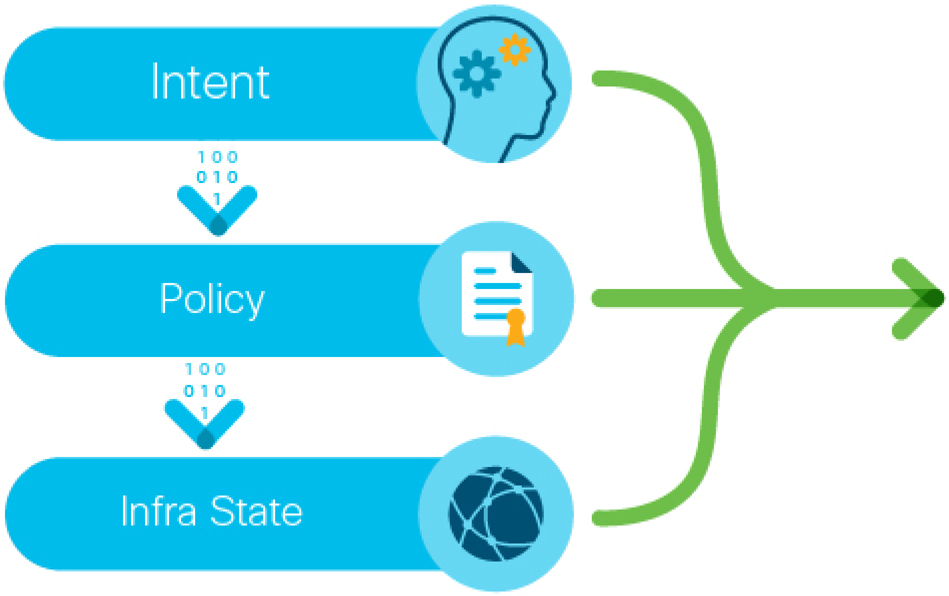

Figure 1. Taking the intent, policy, and the infra state from a device

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 policies in 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 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.
- Policy CAM optimization—You can analyze the amount of policy content-addressable memory (Policy CAM) 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 Policy CAM 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 is comprised of the entire ACI fabric. An ACI fabric is made up of the APIC host and all leaf switches and spine switches controlled by the APIC controller. All the network nodes (APIC controller, leaf switches and spine switches) are analyzed together as part of the Assurance Group. Optionally, an Assurance Entity (if it exists in the ACI fabric) can be included in the Assurance Group. An Assurance Entity is an ancillary item in the ACI fabric that provides support to the overall fabric. For example, a load balancer would be considered an ancillary item in an ACI fabric and could be included in an Assurance Group.

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 a later time. Offline analysis enables you to collect the data during change management windows and then perform the analysis.

Beginning with release 4.0(1), you can analyze multiple Assurance Groups. See the *Cisco NAE Getting Started Guide* for more information.

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.

Dashboard

Dashboard Tab

The **Dashboard** tab screen provides a high-level overview of the health of the assurance group.

Smart Events by Severity

In this area, the total number of Smart Events by severity is displayed.

You can view, at a glance, all the smart events in the assurance group based on their severity. Typically, there are more info smart events and smaller numbers of critical, major, minor, and warning Smart Events. An endpoint with a Critical, Major, or Minor smart event is considered as unhealthy. The total figure is the sum of all of these smart events.

- Critical—An icon that indicates critical smart events.
- Major—An icon that indicates major smart events.
- Minor—An icon that indicates minor smart events.
- Warning—An icon that indicates warning smart events.
- Info—An icon that indicates information smart events.
- Total—This indicates the total number of smart events.

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

When you expand the arrow at the bottom of this area, all smart events for various categories by their severity are displayed.

The categories are as follows:

- Policy Analysis
- Tenant Endpoint
- Tenant Forwarding
- Tenant Security
- Resource Utilization
- System
- Compliance

From the Dashboard screen, you can also navigate for more detailed information to other parts of the UI. Click a specific non-zero smart event count for further details about that category in the **All Smart Events** table. Following the smart event that you click in the table, the relevant filters are applied and the filtered results are displayed in the **Global Search** page.

Unhealthy Count by Resources

This **Unhealthy Count by Resource** area displays the Smart Event counts by unhealthy resources. Resources are policy constructs. Resources that have smart events are called unhealthy policy or forwarding resources. Based on this information, you can root cause the resources that are unhealthy. For each resource the number of unhealthy interfaces is listed alongside the total number of interfaces.

The Smart Event counts for these resources are not mutually exclusive. For example, if a particular Smart Event affects a Tenant and a VRF, the Tenant and the VRF will each display a Smart Event count for that resource issue.

The following resources are displayed:

- Tenants—A secure and exclusive virtual computing environment. In Cisco ACI, a tenant is a unit of isolation from a policy perspective, but it does not represent a private network.
- VRFs—A virtual routing and forwarding instance defines a Layer 3 address domain that allows multiple instances of a routing table to exist and work simultaneously.
- Bridge Domains—These are a set of logical ports that share the same flooding or broadcast characteristics.
- Application Profiles—These define the policies, services, and relationships between endpoint groups (EPGs).
- Endpoint Groups—These are the APIC EPGs.
- Contracts—These specify how communications between EPGs take place.
- External Routed Networks—These are L3 Outs.
- Internal Subnets—These are all the subnets manually defined by the user, excluding L3 Out static routes.
- External Routes—These are L3 Out static routes and externally learned routes.
- Interfaces—These are Layer 1 ports.
- Endpoints—Endpoints are devices connected to the network directly or indirectly.
- Leafs—These are the APIC leaf switches.

Hot Topics

This tab also displays hot topics where you can obtain detailed information about the following:

- Objects by Endpoint Counts
- Objects by Route Count
- Leafs by Policy CAM Usage

Explorer

About Explorer

The **Explorer** feature in NAE analyses a policy snapshot from the Cisco APIC to enable data center operators and architects to:

- Explore the ACI object models and associations
- Verify connectivity and segmentation between network assets

The **Explorer** feature allows network operators to discover assets and their object associations in an easy-to-consume natural language query format. Operators can quickly get visibility into their infrastructure and connectivity or segmentation between assets. The **Explorer** feature allows operators to easily discover associations between traditional networking constructs such as VRFs, subnets, VLANs to the ACI object model.

The Explorer feature is based on natural language query interface. The types of queries supported by the feature include:

- **What Query**: Answers how the different ACI networking entities are related to each other.

Example:

1. What EPGS are associated with VRF: */uni/tn-secure/ctx-secure*
2. What EPs are associated with INF: *topology/pod-1/paths-101/pathep-[eth1/3]* or *VRF:uni/tn-secure/ctx-ctx1*
3. What EPGs are associated with BD: *uni/tn-secure/BD-BD1* and *LEAF: :topology/pod-1/node-103*

- **Can Query**: Answers whether the entities in the ACI policy can communicate with each other. Can queries can also be used to determine if the entities in the ACI policy can communicate using protocols such as TCP, UDP, or ICMP and the source and destination ports used for communication.

Example:

1. Can entity *A* talk to entity *B*.
2. Can EPG: *uni/tn-secure/ap-AP0/epg-B* talk to EPG: *uni/tn-secure/ap-AP0/epg-A* on tcp dport: *80* sport: *10*

- **How Query**: Provides details on the communication between the entities in the ACI policy.

Example: How does EPG *X* talk to EPG *Y*.

- **View Query**: Provides the visual indication of the interface status for any leaf switch in the assurance group.

Example: View interfaces on leaf *X*.

Use Cases

- **Design verification**: Ad-hoc query model enables operators to quickly understand and reason about their infrastructure. The natural language query model returns search results and associations in an easy to understand tabular format. In a single concise view, operators are able to answer design verification questions or discover deviations from organizational best practices.
- **Lightweight book-keeping**: Administration and maintenance teams can provide on demand visibility into the current state of their policy and networking infrastructure allowing inventory, book-keeping, and asset tracking procedures to be lightweight.
- **Connectivity and Segmentation**: Easily answer connectivity questions between a pair of assets or containers of assets. For example, if a group of EPGs needs to be quarantined, the Can query can quickly answer if policy has been correctly setup.

Change Management

Change Management Tab

The **Change Management** tab provides information about the assurance on changes in your fabric. The sub-tabs available here are:

- Policy Analysis
- Manage Pre-Change Analysis

Policy Analysis

The **Policy 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 **Policy 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 **Policy 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.

Manage Pre-Change Analysis

The **Manage Pre-Change Analysis** screen is under the **Change Management** tab.

When you want to change a configuration in an assurance group, a pre-change analysis allows you to model the intended changes in Cisco NAE, perform a pre-change analysis against an existing base epoch in the assurance group, and verify if the changes generate the desired results.

After you model the changes, you can initiate an analysis by saving and running the job. This operation results in smart events being raised as part of the analysis.

When you save and run the job, the changes are applied to the selected base epoch, the analysis is performed, and results are generated. When the analysis is completed, the pre-change analysis instance is listed in the **Manage Pre-Change Analysis** table. For every pre-change analysis listed in the **Manage Pre-Change Analysis** table, a delta analysis is generated between the pre-change analysis and the base epoch. When you click **View Epoch Delta**, it displays the delta analysis screen where you can view the differences in the smart events.

The pre-change analysis is also displayed in the Timeline as an icon above the associated base epoch. When you hover over this icon, you see a list of all pre-change analysis that have been run on that epoch.

If there are smart events raised in the analysis, you can fix them and run a pre-change analysis again and again until you are satisfied with the results.

After you are satisfied with the pre-change analysis, you can download the configuration and upload the JSON file in Cisco APIC.

Additional details about Pre-Change Analysis are as follows:

- Once the analysis starts, the status of the job will be shown as **Running**. During this time, the specified changes will be modeled on top of the base epoch, and complete logical checks will be run, including Policy Analysis and Compliance. No switch software or TCAM checks will be performed. The status of the Pre-Change Analysis job is marked **Completed** when the entire analysis including Epoch Delta completes. The Epoch Delta is automatically triggered and the associated Pre-Change Analysis job us displayed as running during that time. The Epoch Delta is performed only on checks supported in Pre-Change Analysis job.
- You can view the changes applied by a user to a specific Pre-Change Analysis job by clicking the Pre-Change Analysis job in the **Manage Pre-Change Analysis** table. This action displays the Pre-

Change Analysis details for that job. If the changes in the Pre-Change Analysis job were applied manually, you can view the different changes selected by the user. If the Pre-Change Analysis job is created using a JSON file, the **Change Definition** field displays the name of the JSON file from where the changes were imported.

- A Pre-Change Analysis job can be cloned if the same set of changes are required to be verified on different base epochs.
- The resulting data of an analysis are stored as a Pre-Change Analysis epoch, displayed as an annotation on top of the base epoch, in the epoch timeline. Clicking this annotation takes you to the Delta Analysis screen.

Pre-Change Analysis Use Cases

The Pre-Change Analysis feature supports the following use cases:

- The user can create and run a new pre-change analysis job with some changes modeled.
    - The user selects a base epoch on which to make changes.
    - The user specifies one or more changes to add/remove/modify a configured object.
    - The user selects **Save & Run** to analyze this change.
    - When you attempt to run a pre-change analysis, if there is already an epoch analysis running, a dialog box displays where you can stop the analysis that is running and start the new pre-change analysis. If you stop an epoch analysis that is running, you must manually restart it after your pre-change analysis job is complete. Alternatively, you can wait for the analysis that is currently running to complete, and then run the new pre-change analysis.
- The user can create a new pre-change analysis job by cloning a previous job. The user can clone a previous pre-change analysis job to start with the specified changes that exist on top of the specified base epoch, and then create the new job with additional changes.
- The user can delete old pre-change analysis jobs. This will delete all the associated data for the Pre-Change Analysis epoch such as modeled changes, analysis, events and delta jobs.
- The user can export the modeled changes as a JSON file that can be imported into Cisco APIC.

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

Policy CAM

Policy CAM

The **Policy CAM** inspector page provides information about the resource utilization in the network, and the amount of policy content-addressable memory (Policy CAM) 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 Policy CAM
- Determining the policy distribution across leaf switch Policy CAMs
- Determining the least used Policy CAM rules (security policies), based on their hit count

Determining the Leaf Switches That Use the Most Policy CAM

The **Policy CAM** inspector page visualization area enables you to view the leaf switches that have the most Policy CAM being used (top leafs by Policy CAM 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 Policy CAM 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 Policy CAM being used. In such a case, you can modify your Policy CAM rules to balance the Policy CAM usage.

Determining the Policy Distribution Across Leaf Switch Policy CAMs

The **Policy CAM** inspector page visualization area enables you to view the policy distribution across leaf switch Policy CAMs. This data shows you which tenants and endpoint groups are consuming the most contracts, which contracts are consuming the most Policy CAM, 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 Policy CAM 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 Policy CAM Rules By Hit Count

The **Policy CAM** inspector page summary area provides the least used Policy CAM 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 Policy CAM utilization of each Policy CAM rule.

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

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

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 Tab

The **Manage Compliance** tab enables the user to verify Segmentation Compliance and Service Level Agreement (SLA) compliance, and traffic restriction compliance, and configuration 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 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 Object Selectors.

Objects, such as EPGs, VRFs, BDs, are grouped under **Object Selectors** that are container objects that describe a collection of a specific objects. They are configured with selectable attributes. You include and exclude the selection criteria of attributes when defining an object 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 Policy CAM Utilization smart events except when Policy CAM 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.

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

- Smart Events Dashboard
- Manage Event Rules

Smart Event Dashboard

In the **Smart Event Dashboard** sub-tab, the smart events are organized by severity, category, subcategory, and name.

You can manage the display of the smart events with the following:

- Clicking one of the icons in the Smart Events by Severity bar displays only the smart events of the specified severity. All the smart events are displayed when clicking Total in the Smart Events by Severity bar.
- Clicking one of the buttons in the row above the Smart Events by Severity bar displays only the smart events in the selected category. Clicking an additional button adds smart events of the additional category to the display. Clicking a previously selected button removes the smart events of that category from the display. When none of the buttons are clicked, smart events of all categories are displayed. Clicking the filter icon at the beginning of the row de-selects all filter selections.
- Clicking **Aggregated** displays groups of smart events identified by event name.
- Clicking **Individual** displays every instance of the smart event.

In addition, the view of the smart events can be specified with the View pull-down menu:

- All events
- Any rule matches
- Never suppressed
- Suppressed
- Unsuppressed

When selecting and expanding a smart event in the Smart Event Dashboard, the lifecycle of a smart event and details of the individual smart event are displayed.

Lifecycle of a Smart Event

The lifecycle of a smart event appears as an overview summary timeline when displaying the details of an individual smart event. The lifecycle of a smart event is a graphical representation of the individual smart event occurrences in the epochs on the timeline. The color of the epoch icons signify the severity of the smart event. The magnification of the lifecycle can be controlled with the Zoom Level controls below the timeline.

Zoom Level: Lifecycle

Selecting the Lifecycle Zoom Level (default magnification) displays the time when the smart event reached a certain state. (A gray color icon indicates the smart event did not reach the threshold for the state.)

Lifecycle State	Description
⬆ First Raised	Initial occurrence of the smart event and the affected object.

Lifecycle State	Description
Last Raised	Last occurrence of the smart event and the affected object.
Clearing	Smart event and affected object did not occur in one subsequent epoch.
Cleared	Smart event and affected object did not occur in two subsequent epochs.

In addition to these four smart event states, the CONNECTED_EP_LEARNING_ERROR and the FABRIC_EP_LEARNING_ERROR smart events have a Raising state that precedes the First Raised state. If a CONNECTED_EP_LEARNING_ERROR and the FABRIC_EP_LEARNING_ERROR smart event occurs in an epoch, it is in the Raising state. If it occurs in the next consecutive epoch, it is then considered to be in the First Raised state.

Zoom Level: Magnified

Selecting a more granular Zoom Level (or clicking one of the icons on the timeline) increases the magnification of the timeline and displays epochs where the smart event occurred.

With the increased timeline magnification, you can navigate among the epochs by clicking one of the epochs or by using the navigation controls (located below the timeline and to the right of the Zoom Level). The selected epoch displays the date/time for the epoch as well as detailed information for the smart event occurrence.

Clicking the settings icon ⚙ in the Action column opens a menu so that you can edit one of the following pieces of information for the smart event:

- Operation Status: [New, In Progress, or Closed]
- Assign to: Assign to a UserId
- Comment: Add a comment
- Tags: Add a metadata tag

Upon completion of editing these pieces of information, the updated information is displayed in the detailed information for the smart event. You can search the individual smart events by Operation Status, UserID, or Tag value by making these columns a visible attribute to the individual smart event by setting **Column Customization** ☰ and typing the desired search term in the filter field of the appropriate column.

For example, if you edit an individual smart event to have a metadata tag of "Rare event" and add the Tags column to the table of individual smart events using **Column Customization**; then you can search for the smart event by typing "Rare event" in the filter field of the Tags column.

Previous Occurrence and Next Occurrence

If NAE has access to information about the lifecycle of the smart event that is contained in earlier or later epochs, you can click on **Previous Occurrence** or **Next Occurrence** to display these lifecycles.

Lifecycle does not support epochs from Cisco NAE release 3.1(1) and earlier releases.

Smart Event Details

- 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. You can use the View drop down menu to specify other ways to view the smart events, such as Unsuppressed or All events.

Also on the Smart Event Dashboard is the Current Epoch Event Rules Snapshot table.

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

Manage Event Rules

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