Configuring Traffic Analytics

This chapter describes how to configure the Traffic Analytics feature on Cisco NX-OS devices.

About Traffic Analytics

The Traffic Analytics (TA) feature has the following capabilities:

- Provides an ability to identify services offered by servers behind a switch, delivering aggregated analytics data. To distinguish between servers and clients, TCP flags (SYN and SYN ACK) in a three-way handshake are utilized.
- Collapses multiple TCP session data traffic from a client to a server or from a server to client into a single record in the show flow cache database and exports it to the collector. During the traffic analytics aggregation, the source port of TCP is set to a value of 0.
- Supports faster export cadence for troubleshoot flows.
- Supports TA interface filter and VRF filter.

A flow is defined by the source interface, protocol, source IP address, source port, destination IP address, and destination port values. If traffic analytics is enabled, the flows of TCP sessions are aggregated based on source IP address (SIP), destination IP address (DIP), source port (SP) for server to client traffic and SIP, DIP, destination port (DP) for client to server traffic.

Aging of Traffic Database Entries

The traffic database entries will be monitored every 24 hours using a timer. If there is no traffic hitting a database entry, then within 24 to 48 hours that traffic database entry will be deleted. By default the size of the database is 5000.

Troubleshooting Rules

The Troubleshooting rules are used to debug a flow by programming an analytics ACL filter. These rules take precedence over the traffic analytics rules and can be used for capturing specific flow. Troubleshooting rules might result in two entries in the flow cache.

Troubleshooting rules should be used only for specific flows preferably host for short duration only.

Faster Export Cadence for Troubleshoot Flows

Currently, the flow records and troubleshoot records are exported at a fixed interval of one minute. To enhance the efficiency of troubleshooting analysis, a new **filter export-interval** command is introduced. This command facilitates the export of troubleshoot records at a faster interval by utilizing a dedicated hash database.

This configuration can be applied only if traffic analytics is enabled, and a filter is set up within the flow system settings. For more information on **filter export-interval** command, see Example for Traffic Analytics, on page 10.

About UDP Port support

Beginning with Cisco NX-OS Release 10.5(2)F, Traffic Analytics supports UDP port(s) configuration to mask the exported flows.

For masking, the following procedure is followed:

- If the UDP ports are configured, the flows will be masked in the TA DB and NFM flow cache.
- If the destination port is matched, then the source port is masked and vice versa.
- The NetFlow entry will be inserted first, followed by TA entries.
- If UDP port is not configured, the current functionality is not impacted.

To configure the UDP port(s), the following **[no] udp-port** *port-range* command is introduced under the flow traffic-analytics submode (under analytics).

The UDP port must be in the range of 1 to 65565. The port(s) can be entered in a comma-separated and/or range-based format (For example: `2000-3000, 400, 500`).

When the number of ports in the input exceeds the maximum number of ports that can be displayed in a single line command, they will be spilled over to a new configuration line as shown in the example below:

```
analytics
  flow traffic-analytics
    udp-port
5400,500,1024,1026,1028,1030,1032,1034,1036,1038,1040,1042,1044,1046,1048,1050,1052,1054,1056,1058,1060,1062,1064,1066,1068,1070,1072,1074,1076,1078,1080,1082,1084,1086,1088,1090,1092,1094,1096,1098,1100,1102,1104,1106,1108,1110,1112,1114,1116,1118,1120,1122

    udp-port
1124,1126,1128,1130,1132,1134,1136,1138,1140,1142,1144,1146,1148,1150,1152,1154,1156,1158,1160,1162,1164,1166,1168,1170,1172,1174,1176,1178,1180,1182,1184,1186,1188,1190,1192,1194,1196,1198,1200,2000-3000,3002,3004,3006
```

TA Interface Filter and VRF Filter

The Traffic Analytics feature is enhanced to offer more granular support to capture TCP flows using filter configuration at both the interface and VRF levels, similar to the existing FT interface configuration.

Under this TA filter configuration, you can achieve the following:

- Configure an IP address that is required for monitoring and use keywords such as
  - **permit** for IP address that requires monitoring,
  - **deny** to avoid the flow being collected, and
  - **ft-collapse** to integrate flows into a single service.

Note The **ft-collapse** keyword is not used for **troubleshoot filters**.

- Configure the VRF filter across all interfaces in a given VRF.
- Provide permit subnet rules for TCP packets (TCP SYN, SYN ACK, and without any TCP flag).
- For general TCP packets (without SYN or SYN ACK) which are considered for profile 31, the TCS flows forwarded to the collector can be stopped using the **show flow cache** command.
- The **output** option is introduced for flow filter to be applied in the egress direction only.
- Both IPv4 and IPv6 access-list are supported under the filter.

For more information on TA interface filter and VRF filter, see Example for TA Interface Filter and VRF Filter, on page 9.

Interface Traffic Analytics

Granular flow control for Traffic Analytics is provided for

- ingress interface, and
- egress interface.

This feature provides an ability to take specific action on traffic that comes in and goes out of interfaces.

Note Removal of TA is not allowed if it is enabled at interface level.

Granular control for TA feature is provided on ingress and egress interfaces by allowing filters at interface and VRF level, like FT interface configuration.

The table displays the interfaces that are supported on ingress and egress interfaces through releases.

Interfaces	Supported on Ingress and Egress from Release
SVI interface	• Ingress from Release 10.5(2)F • Egress from Release 10.5(3)F
sub interface	Ingress and Egress from Release 10.5(3)F

Interfaces	Supported on Ingress and Egress from Release
port-channel interface	• Ingress from Release 10.5(2)F • Egress from Release 10.5(3)F
VRF interface	• Ingress from Release 10.5(2)F • Egress from Release 10.5(3)F
VNI interface	Ingress and Egress from Release 10.5(3)F

VNI interface

Flow filters can be applied under Layer 3 VNI interfaces in a VXLAN fabric like any other interface filters. The flow filters are supported in both ingress and egress directions. The filters can be either IPv4 or IPv6 filters.

The limitations for TA on VNI interface include:

- Bridged traffic or Layer 2 forwarded traffic cannot be filtered using flow filters applied under L3VNI interface in BGW of a VXLAN fabric.
- A deny flow filter in one direction blocks the traffic in the opposite direction too as VNI interfaces are not direction aware in BGWs of a VXLAN fabric.

Figure 1: Traffic Analytics in VNI interface

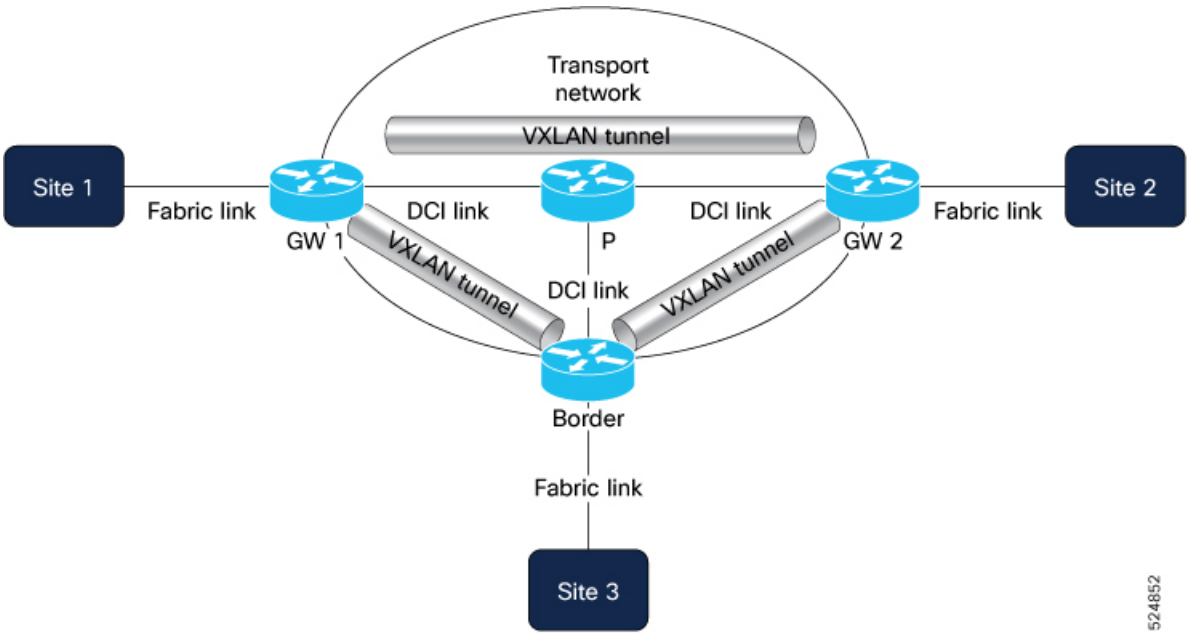

The image depicts what happens during Traffic Analytics in the VNI interface in ingress and egress directions.

- Ingress - All traffic coming from DCI link gets decapsulated using VTEP and goes through the VNI interface where policy is applied. It is then forwarded to either fabric or host interface.

- Egress – All traffic coming from either fabric or host interface goes through VNI interface where policy is applied and gets encapsulated using VTEP. It is then forwarded to the DCI link.

Example configuration for VNI interface

```
vrf context TENANT-VRF
vni 70000 l3

interface nve1
  member vni 70000 associate-vrf


interface vni70000
  flow filter v4_vni_filter_input
  flow filter v4_vni_filter_output output
```

ECN Detection for Traffic Analytics

Explicit Congestion Notification (ECN) helps network devices signal congestion without losing packets. It focuses on the CE (Congestion Experienced) notification, which indicates that a packet has encountered congestion on its path. The enhancement in Traffic Analytics allows the system to find and report ECN bits in the IP header. This feature is designed for use with switches managed by Network Insights, where records are exported to Network Insights Resources (NIR) for consumption and further analysis.

This feature is crucial for closely monitoring and managing congestion across network traffic. It is particularly beneficial for real-time applications, such as VoIP calls and video streaming, where maintaining consistent quality is vital. By focusing on CE notifications and leveraging Network Insights Resources (NIR), network managers gain insights into congestion patterns, helping maintain performance stability in environments sensitive to delays.

- **Enhanced Network Management:** Accurate detection of ECN bits provides administrators with the necessary information to effectively manage congestion, such as rerouting traffic or adjusting bandwidth.
- **Optimized Quality of Service:** By focusing on CE notifications, this feature helps keep real-time applications running smoothly, allowing for proactive management of congestion.
- **Better Troubleshooting:** Monitoring ECN bits provides detailed insights into the network's health, aiding in quick fixes and long-term planning.

How ECN Detection Works

These stages describe how the Traffic Analytics system detects and reports ECN bits in IP traffic:

1. The Traffic Analytics system continuously monitors IP traffic.
2. For each packet, the system examines the IP header to detect ECN bits, specifically looking for the CE (Congestion Experienced) notification.
3. When ECN bits are detected, the system records this information, identifying instances of congestion.
4. The collected data is used to generate reports or alerts for network administrators, highlighting congestion areas, and is further analyzed using NIR.

This process ensures that network administrators receive timely and accurate information about congestion in IP traffic, enabling effective management and optimization of network performance.

Disable global Traffic Analytics

Configure **mode interface** under **flow traffic-analytics** to disable global Traffic Analytics.

```
switch(config)# analytics
switch(config)# flow traffic-analytics
switch(config)# mode interface
```

Use the **no** form of the **mode interface** command to disable the interface mode.

Guidelines and Limitations Traffic Analytics

The following guidelines and limitations are applicable to Traffic Analytics:

- If the Traffic Analytics feature is enabled, other than TCP all other IP protocols get 3 tuple information.
- The Traffic Analytics feature is supported only on Mixed mode in standalone devices.
- Before enabling the Traffic Analytics feature, ensure to remove the flow filters else an error message will be displayed.
- When a system flow filter is configured, the traffic flow behavior is as follows:
  - If a traffic analytics database has information, two flows are seen in the cache.
  - If a traffic analytics database does not have information, only one flow is seen in the cache.
- If the traffic analytics database size is reduced, new entries will happen only after removing the old entries.
- When both NetFlow and traffic analytics are enabled, the profiles used for both functions in a scaled NetFlow configuration are:
  - 29-31 until Cisco NX-OS Release 10.5(2)F
  - 26-31 from Cisco NX-OS Release 10.5(3)F

  When neighbor discovery or special packets hit these profiles, it is not possible to distinguish whether the record created is for traffic analytics or NetFlow. Consequently, the record gets processed twice, resulting in the appearance of two packets with one AN profile.
- Netflow and Flow Telemetry are not supported in N9K-C9364C-H1 platform SFP+ ports, Ethernet1/65, and Ethernet1/66.
- Beginning with Cisco NX-OS Release 10.5(2)F, ingress traffic analytics is supported.
- Beginning with Cisco NX-OS Release 10.5(3)F, the traffic analytics features supported are
  - egress traffic analytics,
  - in ingress traffic analytics:
    - sub-interface,
    - VNI level interfaces, and
    - collapse action.

- explicit congestion notification for flows,
- global traffic analytics is also supported on Cisco Nexus 9500 switches with GX and FX line cards, and
- mode interface disables global traffic analytics only on 9300-FX3, -GX, -GX2, -H2R, and -H1 switches.

Platform support

The table lists the supported platforms for TA features through releases.

Features	Platforms	Release
Support for Traffic Analytics	9300-FX, -FX2, -FX3, -GX, and -GX2	10.4(2)F
Support for Traffic Analytics	9300-H2R and -H1	10.4(4)M
Ingress Traffic Analytics	9300-FX, -FX2, -FX3, -GX, -GX2, -H2R, and -H1	10.5(2)F
Egress Traffic Analytics	9300-FX3, -GX, -GX2, -H2R, and -H1	10.5(3)F
Global Traffic Analytics	9500 with -GX and -FX line cards	10.5(3)F

Note For more information about supported platforms for features through releases, refer to Nexus Switch Platform Support Matrix.

Guidelines and Limitations for TA Troubleshooting Rules

- when upgrading to Cisco NX-OS Release 10.5(1)F using a nondisruptive upgrade, the default value of **filter export-interval** is derived from the NetFlow **flow timeout** value.

Guidelines and Limitations for TA Interface Filter and VRF Filter

- The TA interface filter is not supported for loopback, tunnel interfaces (such as NVE), and management interfaces.
- The TA interface filter is not supported for L3 subinterfaces and L3 port-channel (PO) subinterfaces.
- The VRF filter is not supported for default and management VRFs.
- If TA interface filters and VRF filters are configured, TA interface filters take precedence.

Guidelines and Limitations for ECN Detection for Traffic Analytics

- Beginning with Cisco NX-OS Release 10.5(3)F, the ECN Detection for Traffic Analytics feature is supported on:

- Cisco Nexus 9300-FX3/GX/GX2/H2R/H1 platform switches.
- Cisco Nexus 9700-FX/GX Line Cards.
- Cisco Nexus 9500 EOR switches with GX and FX line cards

- This feature is designed for networks that require detailed congestion monitoring, especially for real-time applications. Configure Traffic Analytics to focus on detecting ECN bits.
- ECN detection is only supported on switches managed by Network Insights Resources (NIR).

Configuring Traffic Analytics

You can configure traffic analytics feature only on mixed mode.

Beginning with Cisco NX-OS Release 10.5(1)F, Traffic Analytics flows can be marked as troubleshoot flows for debugging purposes, and TA flows are exported to the Nexus Dashboard at a faster interval rate.

In the following example, the troubleshoot flows are defined in both IPv4 and IPv6 ACL lists and are attached to a flow filter. The flow filter has been enabled system-wide under the flow system configuration.

Before you begin

Ensure that you are in mixed mode before enabling the traffic analytics feature. To enable the mixed mode, use the following commands. For more information on mixed mode, see Configuring Mixed Mode:

```
(Config)#feature netflow
(Config)#feature analytics
```

Procedure

Step 1 Configure traffic analytics feature with higher cadence support as follows:

Example:

```
ip access-list ipv4-global_filter
  statistics per-entry
  1 permit ip 10.1.1.2/32 11.1.1.2/32
  2 permit ip 11.1.1.2/32 10.1.1.2/32
  3 permit ip 101.1.1.2/32 111.1.1.2/32
  4 permit ip 111.1.1.2/32 101.1.1.2/32

ipv6 access-list ipv6-global_filter
  statistics per-entry
  1 permit ipv6 10::2/128 11::2/128
  2 permit ipv6 11::2/128 10::2/128
  3 permit ipv6 101::2/128 111::2/128
  4 permit ipv6 111::2/128 101::2/128

flow filter global_filter
  ipv4 ipv4-global_filter
  ipv6 ipv6-global_filter


switch(config)# feature netflow
switch(config)# feature analytics
```

```
switch(config)# analytics
switch(config-analytics)#

switch(config-analytics)#  flow traffic-analytics
switch(config-analytics-traffic-analytics)#  db-size 200
switch(config-analytics-traffic-analytics)#  filter export-interval 30
switch(config-analytics-traffic-analytics)#  flow system config
switch(config-analytics-system)#  traffic-analytics
switch(config-analytics-system)#  monitor monitor input
switch(config-analytics-system)#  profile profile
switch(config-analytics-system)#  event event
switch(config-analytics-system)#  filter global_filter
```

Step 2 Use the **flow filter** <*filter*> command to configure traffic analytics for ingress interface.

Example:

```
switch(config)# interface Ethernet1/1
switch(config-if)# flow filter test
```

Step 3 Use the f**flow filter** <*filter*> **output** command to configure traffic analytics for egress interface.

Note
Before using egress filters, ensure that the **egress netflow tcam** region is carved.

Example:

```
switch(config)# interface Ethernet1/1
switch(config-if)# flow filter test output
```

Example for TA Interface Filter and VRF Filter

Interface Filter Configuration

The following example shows how the interface filter configuration is performed:

```
ip access-list ipv4-l3_intf_filter
  statistics per-entry
  1 permit tcp 10.1.1.7/32 11.1.1.7/32 syn
  2 permit ip 10.1.1.7/32 11.1.1.7/32

ipv6 access-list ipv6-l3_intf_filter
  statistics per-entry
  1 permit tcp 10::7/128 11::7/128 syn
  2 permit ipv6 10::7/128 11::7/128

flow filter l3_filter
  ipv4 ipv4-l3_intf_filter
  ipv6 ipv6-l3_intf_filter

analytics

  flow traffic-analytics
    db-size 200
    filter export-interval 30
  flow system config
    traffic-analytics
    monitor monitor input
    profile profile
```

```
    event event

interface Ethernet1/63/1
  flow filter l3_filter
flow filter l3_filter output

switch(config-analytics)# show running-config inter e 1/63/1

interface Ethernet1/63/1
  vrf member vrf1
  flow filter l3_filter
  ip address 10.1.1.1/24
  ipv6 address 10::1/64
  no shutdown
```

VRF Filter Configuration

The following example shows how the VRF filter configuration is performed:

```
ip access-list ipv4-vrf1_filter
  statistics per-entry
  1 permit tcp 10.1.1.9/32 11.1.1.9/32 syn
  2 permit tcp 11.1.1.9/32 10.1.1.9/32 ack syn

ipv6 access-list ipv6-vrf1_filter
  statistics per-entry
  1 permit tcp 10::9/128 11::9/128 syn
  2 permit tcp 11::9/128 10::9/128 ack syn

flow filter vrf1_filter
  ipv4 ipv4-vrf1_filter
  ipv6 ipv6-vrf1_filter

analytics

  flow traffic-analytics
    db-size 200
    filter export-interval 30
  flow system config
    traffic-analytics
    monitor monitor input
    profile profile
    event event

vrf context vrf1

  flow filter vrf1_filter
flow filter vrf1_filter output
```

Example for Traffic Analytics

The following example displays the output of the troubleshoot flows export interval:

```
switch(config-analytics-traffic-analytics)# show flow traffic-analytics
Traffic Analytics:
    Service DB Size: 200
    Troubleshoot Export Interval: 30
```

The **filter export-interval** command allows setting the troubleshoot timer with a range of 10 to 60 seconds. The default value for this timer is set to 10 seconds.

The **no filter export-interval** will reset the troubleshoot timer range to default value of 60 seconds.