

Configuring Traffic Analytics

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

- Traffic Analytics, on page 1
- Guidelines and limitations for Traffic Analytics, on page 6
- Configure Traffic Analytics, on page 8
- Example for Traffic Analytics interface filter and VRF filter, on page 9
- Example for Traffic Analytics, on page 11

Traffic Analytics

Traffic Analytics is a feature that

- provides the ability to identify services offered by servers behind a switch, aggregates analytics data, and exports summarized flow records for analysis
- distinguishes between servers and clients using TCP flags (SYN and SYN ACK) in a three-way handshake
- collapses multiple TCP session data traffic into a single record in the show flow cache database and exports it to the collector; during aggregation, the source port of TCP is set to 0
- supports faster export cadence for troubleshoot flows, and
- supports traffic analytics 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 flow records

Currently, the flow records and troubleshoot records are exported at a fixed interval of one minute. 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 11.

UDP port support and configuration

UDP port support in Traffic Analytics allows masking of exported flows based on configured UDP ports.

- When UDP ports are configured, flows are masked in the TA DB and NFM flow cache.
- If the destination port matches, the source port is masked, and if the source port matches, the destination port is masked.
- If UDP port is not configured, the current functionality is not impacted.



Note First enter the NetFlow entry, and then make the TA entries.

UDP port configuration

Configure UDP port(s) for masking exported flows, use the [**no**] **udp-port** *port-range* command under the flow traffic-analytics submode (under analytics).

- The UDP port must be in the range of 1 to 65565.
- Ports 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 are spilled over to a new configuration line as shown in this example.

```
analytics
  flow traffic-analytics
    udp-port
5,40,50,102,104,106,108,110,112,204,206,302,304,306,308,310,312,404,406,408,410,412,414,416,418,420,422,424,426,428,430,432,434,436,438,440,442,444,446,448,450,452,454,456,458,460,462,464,466,468,470,472,474,476,478,480,482,484,486,488,490,492,494,496,498,500,502,504,506,508,510,512,514,516,518,520,522

    udp-port
524,526,528,530,532,534,536,538,540,542,544,546,548,550,552,554,556,558,560,562,564,566,568,570,572,574,576,578,580,582,584,586,588,590,592,594,596,598,600,200-300,302,304,306
```

Traffic Analytics interface filter and VRF filter

The Traffic Analytics (TA) 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 Traffic Analytics interface filter and VRF filter, on page 9.

Interface Traffic Analytics

Interface Traffic Analytics is a network analytics feature that provides

- granular flow control for ingress and egress interfaces, enabling specific actions on traffic that comes in and goes out of interfaces
- filters can be applied at both the interface and VRF level, similar to FT interface configuration, and
- removal of Traffic Analytics is not allowed if it is enabled at the interface level.

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.

Table 1: Supported interfaces for ingress and egress Traffic Analytics

Interface	Ingress Support	Egress Support
SVI interface	From Release 10.5(2)F	From Release 10.5(3)F
sub interface	From Release 10.5(3)F	From Release 10.5(3)F
port-channel interface	From Release 10.5(2)F	From Release 10.5(3)F
VRF interface	From Release 10.5(2)F	From Release 10.5(3)F
VNI interface	From Release 10.5(3)F	From Release 10.5(3)F

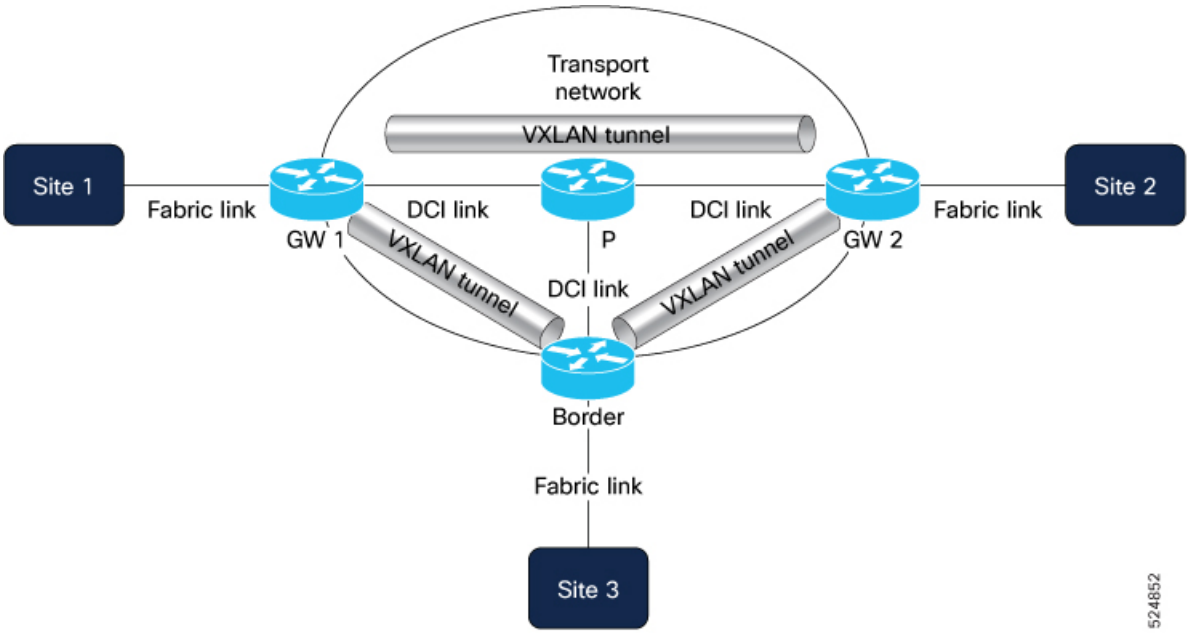
VNI interface

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) is a mechanism that provides

- 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.

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.

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

Procedure

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

Example:

```
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 for Traffic Analytics

The guidelines and limitations are applicable to Traffic Analytics are:

- 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 is displayed.
- If the traffic analytics database size is reduced, new entries happen only after removing the old entries.

  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.

- When both NetFlow and traffic analytics are enabled, the profiles used for both functions in a scaled NetFlow configuration are:

  - 29-31 until NX-OS Release 10.5(2)F
  - 26-31 from 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 NX-OS Release 10.5(2)F, ingress traffic analytics is supported.
- Beginning with 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 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, GX2, and H2R	10.4(2)F
Support for Traffic Analytics	9300-H1	10.4(3)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
Egress Traffic Analytics	9300-FX and FX2	10.6(1)F



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

Guidelines and limitations for Traffic Analytics troubleshooting rules

- When upgrading to 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 Traffic Analytics 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 NX-OS Release 10.5(3)F, the ECN Detection for Traffic Analytics feature is supported on:
  - Nexus 9300-FX3, GX, GX2, H2R, and H1 platform switches.
  - Nexus 9700-FX and GX Line Cards.
  - 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).

Configure Traffic Analytics

Enable and configure the traffic analytics feature to monitor and analyze network flows.

You can configure the traffic analytics feature only in mixed mode.

Beginning with 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 Mixed Mode:

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

Procedure

Step 1 Configure traffic analytics feature with higher cadence support.

Example:

```
ip access-list ipv4-global_filter
  statistics per-entry
  1 permit ip 192.0.2.1/32 198.51.100.1/32
```

```
  2 permit ip 198.51.100.1/32 192.0.2.1/32
  3 permit ip 203.0.113.1/32 192.0.2.2/32
  4 permit ip 192.0.2.2/32 203.0.113.1/32

ipv6 access-list ipv6-global_filter
  statistics per-entry
  1 permit ipv6 2001:DB8::2/128 2001:DB8:1::2/128
  2 permit ipv6 2001:DB8:1::2/128 2001:DB8::2/128
  3 permit ipv6 2001:DB8:2::2/128 2001:DB8:3::2/128
  4 permit ipv6 2001:DB8:3::2/128 2001:DB8:2::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 **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 Traffic Analytics interface filter and VRF filter

This topic provides examples for configuring Traffic Analytics interface filters and VRF filters.

Interface filter configuration

This example shows how the interface filter configuration is performed.

```
ip access-list ipv4-l3_intf_filter
  statistics per-entry
  1 permit tcp 10.0.0.0/8 172.16.0.0/12 syn
```

```
  2 permit ip 10.0.0.0/8 172.16.0.0/12

ipv6 access-list ipv6-l3_intf_filter
  statistics per-entry
  1 permit tcp 2001:DB8::7/128 2001:DB8:1::7/128 syn
  2 permit ipv6 2001:DB8::7/128 2001:DB8:1::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 10.0.0.1/24
  ipv6 address 2001:DB8::1/64
  no shutdown
```

VRF filter configuration

This example shows how the VRF filter configuration is performed.

```
ip access-list ipv4-vrf1_filter
  statistics per-entry
  1 permit tcp 192.0.2.1/32 198.51.100.1/32 syn
  2 permit tcp 198.51.100.1/32 192.0.2.1/32 ack syn

ipv6 access-list ipv6-vrf1_filter
  statistics per-entry
  1 permit tcp 2001:DB8::9/128 2001:DB8::A/128 syn
  2 permit tcp 2001:DB8::A/128 2001:DB8::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

This section provides an example and explanation for configuring and troubleshooting the traffic analytics export interval.

This 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** resets the troubleshoot timer range to default value of 60 seconds.