



Configuring Traffic Analytics

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

- [About Traffic Analytics, on page 1](#)
- [Guidelines and Limitations Traffic Analytics, on page 2](#)
- [Configuring Traffic Analytics, on page 3](#)
- [Example for TA Interface Filter and VRF Filter, on page 4](#)
- [Example for Traffic Analytics, on page 5](#)

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 5](#).

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.
- Configure an IP address which does not need flow collection using a **deny** keyword.
- 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.

For more information on TA interface filter and VRF filter, see [Example for TA Interface Filter and VRF Filter, on page 4](#).

Guidelines and Limitations Traffic Analytics

The following guidelines and limitations are applicable to Traffic Analytics:

- Beginning with Cisco NX-OS Release 10.4(2)F, the Traffic Analytics feature is supported on Cisco Nexus 9300-FX/FX2/FX3/GX/GX2 platform switches.
- 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 NetFlow and traffic analytics are enabled, profiles 29–31 will be used for both functions if we have a scaled NetFlow configuration that is using those profiles. When neighbor discovery or special packets hit these profiles, it is not possible to differentiate whether the record created is traffic analytics or NetFlow. As a result, it gets processed twice, leading to the appearance of two packets with an AN profile.
- Netflow and Flow Telemetry are not supported in N9K-C9364C-H1 platform SFP+ ports, Ethernet1/65, and Ethernet1/66.

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.

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
```

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

```

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

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

```

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.