



Cisco Nexus Dashboard Insights Flows,
Release 6.5.1 - For Cisco NDFC or
Standalone NX-OS

Table of Contents

New and Changed Information	2
Flows	3
Flows	3
Flows Hardware Requirements	3
Flows Guidelines and Limitations	3
View Flows	5
L4-L7 Traffic Path Visibility	6
Flow Telemetry Events	7
Multi-Site Traffic Path - Beta Feature	9
Multi-Site Traffic Path Trace and Fault Correlation	9
Configure Multi-Site Traffic Path Trace and Fault Correlation	9
Copyright	11

First Published: 2024-06-28

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

New and Changed Information

The following table provides an overview of the significant changes up to the current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

New Features and Changed Behavior in the Cisco Nexus Dashboard Insights

Feature	Description	Release	Where Documented
Terminology change	The term " sites" is renamed to " fabrics" .	6.5.1	Entire document

This document is available from your Nexus Insights GUI as well as online at www.cisco.com. For the latest version of this document, visit [Cisco Nexus Dashboard Insights Documentation](#).

Flows

Flows

Flows provides deep insights at a flow level giving details such as average latency and packet drop indicator. It also raises anomalies when the latency of the flows increase or when packets get dropped because of congestion or forwarding errors.

Each flow has a packet counter representing the number of packets entering the ASIC for that flow over a period of time. This period of time is called aggregation interval. There are several points where flow statistics for a given flow can be aggregated. Aggregation can happen in the ASIC, switch software, and server software.

The Flows section of Nexus Dashboard Insights displays the telemetry information collected from various devices in the fabric that were added to the fabric.

For details on Flow Telemetry support for Cisco Nexus series switches and line cards, see the **Compatibility Information** section in the Nexus Dashboard Insights [Release Notes](#).

Flows Hardware Requirements

For details on Flows Telemetry support for Cisco Nexus platform switches, see [Nexus Dashboard Insights Release Notes](#) Compatibility Information section.

Flows Guidelines and Limitations

For details on Flow Telemetry hardware support, see the Compatibility Information section in the [Nexus Dashboard Insights Release Notes](#).

- Output port information for outgoing traffic from N9K-C93180YC-EX, N9K-C93108TC-EX, and N9K-C93180LC-EX line cards will not be displayed.
- Flows does not support multicast traffic. The access list must be provisioned to exclude the multicast traffic flows.
- A maximum of 63 VRFs are supported on flow telemetry nodes.
- The number of anomalies in the **Fabric Overview** dashboard will not match the number of anomalies in the flow browse page. The Fabric Dashboard contains the total anomaly count for the time range you selected. The flow records are not aggregated in the flow browse view, where multiple flow records can point to the same anomaly entry.
- The L3-VNI flows show as L2-VNI flows when the VXLAN flow is dropped in the ingress node. When VXLAN packets are dropped in the first-hop, the exported VXLAN flow telemetry records will indicate the drop. However, they don't carry the VNI information in it. The Ingress interface from the flow telemetry export along with the VRF associated with the interface, does not deduce if the flow is L2-VNI or L3-VNI. In this case Nexus Dashboard Insights associates the L2-VNI for the flow.
- When a VXLAN encapsulated packet enters an Cisco Nexus 9500-EX switch and feature overlay (EVPN) is configured, the packet will be treated like a VXLAN transit node packet. Also the ingress interface and egress interface are set as zeros in the flow telemetry export. The ingress and egress interfaces are needed to consider this record for flows. The limitation on these switches

results in the Cisco Nexus 9500-EX switch not being considered in the path stitching and correlation if the switch is in ingress, transit, or egress direction. Cisco Nexus 9500-EX switches will be treated like a transit node for an overlay packet.

- For Nexus Dashboard Insights to work in VXLAN deployments, you must have symmetric configuration on the switches involved in the overlay. This enables Nexus Dashboard Insights to correlate and stitch the overlay flows. When such a symmetric configuration is not present, the VXLAN feature and forwarding will work, but Nexus Dashboard Insights will not stitch the flows correctly. See the following examples to understand what is meant by symmetric configuration on switches:
 - For Layer 2 VXLAN VNI cases: If vlan-x is mapped VNI-A in PE1, then the same vlan-x must be mapped to VNI-A in PE2, where PE1 and PE2 are VTEP endpoints for the Layer 2 overlay.
 - For Layer 3 VXLAN VNI cases: If SVI-x is mapped to VRF-A mapped VNI-P on PE1, then the same SVI-x must be mapped to VRF-A mapped VNI-P in PE2, where PE1 and PE2 are VTEP end points for the Layer 3 overlay.
- The ingress and VRF information will not be shown for all interfaces which use the flow telemetry 'tenant-id' for encoding the logical interface ID, as this ID will be used for 'overlay-id'. It's not possible to derive the logical interface (SVI with trunk port, sub interface, SVI with trunk and port channel) and get the VRF associated with it. This results in the flow browse page and details page not showing the ingress and egress VRFs.
- On the Cisco Nexus 9500-EX switches connected to VPC pair, the current design limits in identifying the ingress leaf nodes between VPC pairs causing the loss of flow in Nexus Dashboard Insights.
- When there are 29 million anomalies in the indices, flow database writes are too slow, which causes KAFKA lag on 350 nodes supported for software telemetry and flow telemetry. The KAFKA lag results in partial data in Nexus Dashboard Insights user interface.
- Flows information is retained for 7 days or until flow database reaches 80%, which ever happens first, then older flows information is deleted from the database.
- Flow telemetry and flow telemetry events will not export **drop bit** if there is an egress ACL drop in Cisco Nexus FX switches.
- For Nexus Dashboard Insights to receive Flow Telemetry data, the TCAM region for **ing-netflow** must be set to 512. See [Nexus 9000 TCAM Carving](#).
- For flows, if the time range you have selected is greater than 6 hours, the data may not get displayed. Select a time range that is less than or equal to 6 hours.
- When the scale limit for anomalies is reached, some of the unhealthy flows may not be displayed as anomalies in the Flow Record Details page. A system issue is raised when this condition happens. Navigate to **Admin > System Settings > System Issues** to view the system issue.
- For Layer 4 to Layer 7 Services only intra VRF is supported for flows. For Layer 4 to Layer 7 Services inter VRF is not supported for flows.
- To enable Flow Telemetry on a NDFC fabric with Netflow configured and Traffic Analytics enabled, you must perform the following steps:
 1. Delete the netflow policy on the NDFC.
 2. Disable Traffic Analytics. See [Traffic Analytics](#).

3. Enable Flow Telemetry. See [Getting Started](#).

- Multicast is not supported for Flow Telemetry.
- When traffic flows through sub-interfaces, in the **Flow Record Details** page, sub-interface is displayed only in the ingress direction in the Flow Path area. In the egress direction the parent interface is displayed.
- For unicast routes in Flow Telemetry, Layer 4 to Layer 7 Services traffic path visibility is supported on Cisco Nexus 9300 -GX2 and -FX3 platform switches.
- Configuring Netflow on a switch before onboarding a Standalone NX-OS fabric on Nexus Dashboard Insights and then enabling Flow Telemetry on the corresponding fabric from Nexus Dashboard Insights could lead to anomalies being generated for incorrect configurations being pushed to the switch.

To resolve the issue, perform the following steps:

1. Remove switch from the Standalone NX-OS fabric.
 2. Use the command **no feature netflow** to remove Netflow from the switch.
 3. Add the switch back to the Standalone NX-OS fabric.
 4. Enable Flow Telemetry.
- When Flow Telemetry is disabled while one of the switches is unreachable, the fabric goes into the **Disable Failed** state. This is expected behavior. Following this condition, when the switch becomes available and you enable Flow Telemetry, the ACL configurations get corrupted.

We recommend either of these workarounds:

- When this condition occurs and the fabric is in a **Disable Failed** state, retry the switch that is in this state. This will trigger the ACL to unconfigure successfully. Thereafter, when you enable the fabric, start a **disable acl job** to clean any existing ACLs in the switch.
- If you have already enabled Flow Telemetry without cleaning the existing ACLs in the switch, then perform these actions:
 1. Disable Flow Telemetry and then enable it.
 2. Remove the switch that has the issue from the DCNM fabric and add it back in the fabric.

View Flows

The Flows page displays telemetry information collected from various devices in an online fabric. The Flows records let the user visualize the flows per fabric. For a particular fabric, you can view flows by Anomaly Score, Packet Drop Indicator, and Average Latency.

The flows engine also runs machine-learning algorithms on the behavior of the flows to raise anomalies in the behavior, such as average latency and packet drop indicator. The graph represents the anomalies in the behavior over a period of time.

Flow telemetry and analytics gives in-depth visibility of the data plane. Flows collects the flow records streamed from the nodes and converts to understandable EPG-based flow records. Top nodes by flow anomalies displays the nodes in the network with the most anomalies.



To view flow details, you must first enable flows. See [Getting Started](#).

1. Navigate to **Manage > Fabrics**.
2. Select **Online Fabrics** from the drop-down list.
3. Click a fabric name to view the fabric details.
4. Navigate to **Connectivity > Flows**.
5. Select a time range.
6. In the Fabric Flows by area select an option from the drop-down list to view flows by Anomaly Score, Packet Drop Indicator, and Average Latency. The graph displays a time series plot for flows properties recorded in the entire fabric. The flows recorded for Top Sources and Top Destinations are also displayed.

Anomaly Score—The score is based on the number of detected anomalies logged in the database.

Packet Drop Indicator—The flow records are analyzed for drops. The primary method of detecting drops is based on the drop bit received from the switch (flow records).

Latency—The time taken by a packet to traverse from source to destination in the fabric. A prerequisite for fabric latency measurement is that all the nodes shall be synchronized with uniform time.

7. The Flows table displays information such as anomaly score, flow record time, nodes, flow type, protocol, latency, packet drop indicator.
8. Use the search bar to filter the flows. The Flows table displays the filtered flows. Click the column heading to sort the flows in the table.
9. Click **Record Time** to view the flow record details. The details include record time, flow type, aggregated flow information, ingress and egress information, flow path, anomalies, and trends for average latency, traffic, packet drop indicator, bursts.

L4-L7 Traffic Path Visibility

Starting with Nexus Dashboard Insights release 6.1.1, you now have expanded visibility in the Flow Path to L4-L7 external devices such as firewalls. Nexus Dashboard Insights tracks the end-to-end flow across the service chain in real-time and helps locate data plane issues across the device silos. In the current release, a non-NAT environment across all third-party vendors is supported.

For L4-L7 traffic path visibility, your flow telemetry must be enabled and the appropriate rules must be configured. See [Getting Started](#). Based on your rules, if the flow is passing through Policy Based Redirects (for example, a firewall), it will display that information in the flow path.

To view traffic path visibility in the GUI,

1. Navigate to **Manage > Fabrics**.
2. Select **Online Fabrics** from the drop-down list.
3. Click a fabric name to view the fabric details.

4. Navigate to **Connectivity > Flows**.
5. Select a time range.
6. In the flows table, click Record Time to view the flow record details.

In the **Path** area, a graphical flow path will display the end-to-end information, from source to destination, and it will also identify the firewall in the path if a firewall is present. The graph also captures the end-to-end flow path network latency that is occurring. In the graph, if there are any anomalies, a red dot is displayed next to the symbol for the leaf switch or the spine switch.

7. Click **Anomalies** in the **Flow Details** page to view further details related to the anomaly.



In the current release firewalls are not supported for anomalies.

Guidelines and Limitations for L4-L7 Traffic Path Visibility

- This feature is currently recommended only if Policy Based Redirect can be configured using L4-L7 Service for NDFC.
- Service node types will be detected if the service node is directly connected. However, if multiple service nodes are connected on the same physical port, Nexus Dashboard Insights will not identify the exact service node type information. As a result, it will be identified as an unknown service node.
- In the current release, firewalls are not supported for anomalies.
- In the current release, the latency information that is being displayed is the network latency, and it does not capture the latency that is occurring in the firewall.
- In the current release, NAT is not supported.
- This feature is currently supported if you use the following switches:
 - Cisco Nexus 9300-FX Platform Switches
 - Cisco Nexus 9300-FX2 Platform Switches
 - Cisco Nexus 9300-GX Platform Switches.
- Policy Based Redirect destinations on L3Out are not supported because such configurations use an internal VRF because of which only a partial flow path would be available.
- Without a service graph for L4-L7, if the client > service node is VRF_A and the service node > server is VRF_B, the paths will be recorded as separate flows as there is no common or single contract to stitch the flows.
- Load Balancers are not supported.

Flow Telemetry Events

Flow telemetry events are enabled implicitly when flow telemetry is enabled and flow rules are configured. The flow telemetry enables triggering events when a configured rule is met, where packets are exported to the collector for analysis.

Flow telemetry events enhance and complement current flows in {PlatformFullName}. They enrich anomaly generation for flow telemetry and flow telemetry events.

It monitors security, performance, and troubleshooting. This is achieved using the periodic flow table

event records exported every second.

The data export to {PlatformFullName} is done directly from the hardware without control plane needing to handle the data. Statistics are assembled as a packet with a configurable MTU size and a defined header. These packets are sent as in-band traffic from NDFC fabric. Headers are configured by software, and packets streamed are UDP packets.

When flow telemetry is available for a triggered flow telemetry event, then you can navigate to flow details page for aggregated information. These events are based on the following drop events:

- **Cisco ACL Drop**—When packet hits **sup-tcam** rules and the rule is to drop the packet, the dropped packet is counted as ACL_Drop and it will increment the forward drop counter. When this occurred, it usually means the packet is about to be forwarded against basic Cisco ACL forwarding principals. The **sup-tcam** rules are mainly to handle some exceptions or some of control plane traffic and not intended to be checked or monitored by users.
- **Buffer Drop**—When the switch receives a frame and there are no buffer credits available for either ingress or egress interface, the frame is dropped with buffer. This typically hints at a congestion in the network. The link that is showing the fault could be full or the link containing the destination may be congested. In this case a buffer drop is reported in flow telemetry events.
- **Forward Drop**—The packets that are dropped on the LookUp block (LU) of the Cisco ASIC. In a LU block a packet forwarding decision is made based on the packet header information. If the packet is dropped, forward drop is counted. There may be a variety of reasons when forward drop is counted.

TCP packet RTO anomaly is not supported on NDFC.

Flow Telemetry Events Vs Flow Telemetry

- The flow telemetry event packets are exported only when configured events occur, where as flow telemetry packets are streamed continuously.
- The flow telemetry events are captured for all traffic, where as flow telemetry is captured for filtered traffic.
- The total number of collectors between flow telemetry and flow telemetry events is 256.

Guidelines and Limitations for Flow Telemetry Events

- Flow telemetry event anomalies are aggregated. For example, a packet drop anomaly occurred from time T0 to T1. No packet drop anomaly occurred from time T1 to T2. Another packet drop anomaly occurred from time T2 to T3. Although there is no anomaly from T1 to T2, the time stamp for the aggregated packet drop anomalies is from T0 to T3.
- The flow telemetry events do not report policing drop anomalies in {PlatformFullName}, when the egress data plane policer is configured on front-panel ports and there is traffic drop.
- To export flow telemetry events on FX platform switches, you must configure flow telemetry filters.

Multi-Site Traffic Path - Beta Feature

Multi-Site Traffic Path Trace and Fault Correlation



This is a **Beta** feature. We recommend that you use features marked as **Beta** in your test environments but not in production deployments.

To monitor flows, you can stitch together flows from across two different fabrics into one single view. This enables you to have end-to-end views for the paths and end-to-end details for the particular flow and the latency information for that flow.

Use cases for Multi-Site Traffic Path Trace and Fault Correlation:

- You can correlate flows across fabrics and display flow details with stitched paths.
- You can monitor flows across fabrics and generate inter-fabric anomalies that are trigger based.
- You can monitor flows across fabrics and provide the end-to-end latency.

Configure Multi-Site Traffic Path Trace and Fault Correlation

In the **Explore** area, you can view the flow path between two ports, their IP addresses, and their VRFs.

1. In **Flows** table, click Record Time to view the flow record details.
2. In the **Flow Path** area, in the flow path, click the **Multi-Site Flow - View in Flow Explore** tab to go to the **Explore** page.

In the **Explore Flows** page, the flow information filters in the Search field auto-populate, and you can see in which fabrics the flows exist. The **View** query area displays the information with the source IP address, source port information and the destination IP address, destination port information. **Explore** finds and returns all the fabrics where this flow was discovered for the specified VRF.

Next, select the appropriate source and destination fabrics to view their aggregated information, path summary, and anomalies. You must specify which fabric you want as your source and which fabric you want as your destination. Cisco Nexus Dashboard Insights will stitch the information based on your input. You can only choose one source and one destination at a time to stitch together this information. Based upon your selection of source and destination fabrics, Cisco Nexus Dashboard Insights returns the names of the fabrics that it finds.

In the **Flow Path Summary** area, details for the two fabrics are displayed in the **Explore** page as a graphical flow path that displays the end-to-end information, from source to destination. You can see the first fabric with the endpoint and a set of nodes and it is connected to the second fabric with the second set of nodes followed by the endpoint. It will also identify the firewall in the path if a firewall is present. The graph also captures the end-to-end flow path network latency.

Specific details for the source and the destination fabrics are displayed in each of the flows tables. In the **Anomalies** table, choose **Grouped** to view the grouped anomalies for your selections flows.



If you enter different flow details in the Search field in the **Explore** page, you can view in which fabrics those flow exists. Alternatively, in the **Explore** page, you can directly begin your search by entering details in the Search field for details about flows and the Flow Path across more than one fabric.

Copyright

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

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

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

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

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

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

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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

© 2017–2024 Cisco Systems, Inc. All rights reserved.