

Media Flow Analytics

Contents

Every packet matters in video delivery over IP	3
The move to IP	3
Operational challenges	3
Telemetry with the Cisco Nexus 9000 Series Switches and Cisco NX-OS	4
Flow-health monitoring	4
Feature in action: remote production use case	5
Configuring media flow analytics on Cisco Nexus 9000 Series Switches	6
Configuration steps	6
Displaying RTP flows and errors	6
New architectures require new tools	8

Using Cisco Nexus 9000 Series Switches for Media Flow Health Monitoring

Every packet matters in video delivery over IP

The media and entertainment industry is undergoing a massive transformation. Content production is rapidly moving from standard definition to high definition and ultra-high definition and beyond. To adapt to these changes, media companies have moved their production infrastructures from serial digital interfaces (SDI) to IP.

Although IP offers advantages and opportunities for the industry, the nature of the architecture makes it more difficult to isolate problems when they arise. When video flow is affected, such as when a packet is lost, the viewer sees a disturbance on the screen, but doesn't know why. However, using a combination of hardware and software telemetry, the Cisco Nexus® 9000 Series Switches can help you isolate traffic problems on an IP network, so you can resolve them in seconds instead of hours.

The move to IP

In the past, the broadcast industry used an SDI router and SDI cables to transport video and audio signals. SDI cables carry only a single unidirectional signal, so a large number of cables, frequently stretched over long distances, are required. However, with an IP-based infrastructure, a single cable has the capacity to carry multiple bidirectional traffic flows and can support different flow sizes without requiring changes to the physical infrastructure. Media companies are moving to IP-based infrastructures to meet the demands for more content and rich media experiences, including more camera feeds, higher resolutions, and virtual reality capabilities. An IP network architecture makes production options such as live production in studios, stadiums, and remote production locations feasible.

Because SDI architectures offered one-to-one connections, it was relatively easy to detect the source of a problem. But with an IP infrastructure the cable is common, which affects problem detection. The lack of visibility into flow health in IP networks is a serious concern in the industry and is one of the reasons some media and entertainment organizations are reluctant to move their production workflows to IP. Media businesses need to be able to efficiently operate their IP networks and ensure reliability. If video flow is degraded, you need to be notified proactively or in real time to minimize the effect on viewers.

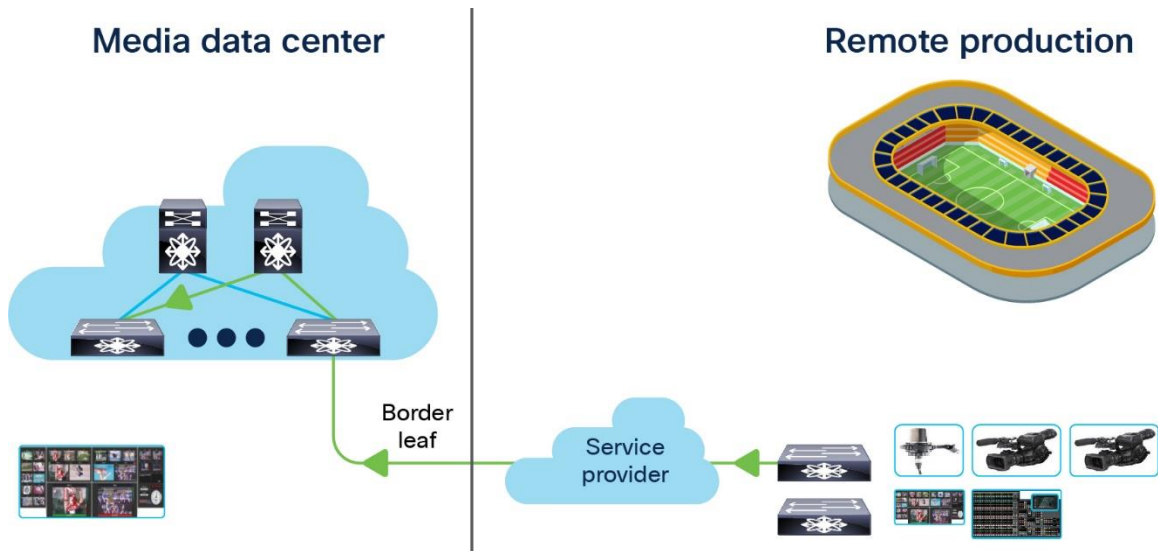
Operational challenges

Bad video on air is an operator's nightmare. Although IP networks are designed to ensure resiliency with multiple paths from source to destination, interface errors and drops are beyond an operator's control.

For example, suppose content that is captured at a stadium is transported to a production facility across a service provider network (see Figure 1). When a problem arises, it's difficult for the network operations team to pin-point the source of the problem. The network operator has complete control over the network in the production facility but may not have any visibility over the IP network in the stadium or the service provider.

If the video on screen goes bad because of packet loss, the loss could be happening at the stadium, at the service provider, or at the production facility itself. Isolating where the loss is taking place involves long hours of troubleshooting. Problems like these can lead to a loss of revenue and affect the entire business. Sometimes the issue may disappear before troubleshooting is complete, which makes it impossible to determine the cause of the failure.

Figure 1. Remote production example



Telemetry with the Cisco Nexus 9000 Series Switches and Cisco NX-OS

The traditional means of checking network health involved polling each network element for link errors, usage, and resource utilization. Because of the nature of periodical polling, often the information isn't captured in real time, so you can miss events that occur between polling intervals. The information returned from the network could indicate issues such as packet loss on an interface, but it never provided deeper insights into what traffic or application was affected.

Rather than polling the network, the Cisco Nexus 9000 solution uses hardware and software telemetry to proactively notify operators of traffic problems (see Figure 2). Because you find out about problems more quickly, this proactive notification can reduce the mean time for resolution from hours to seconds.

Figure 2. Hardware and software telemetry with Cisco Nexus 9000 and Cisco NX-OS

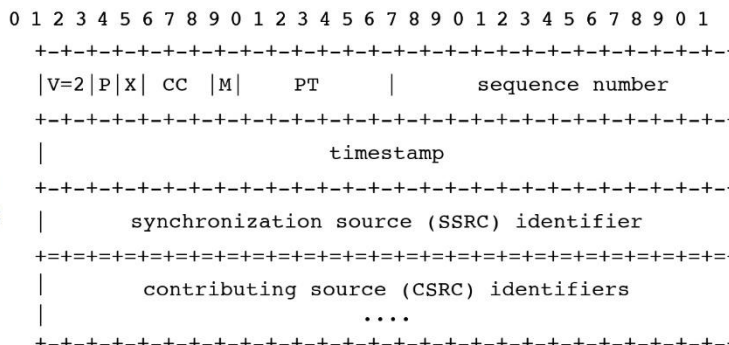


Flow-health monitoring

Almost all of the applications where video is delivered over IP use Real-Time Transport Protocol (RTP) over User Datagram Protocol (UDP). These applications include uncompressed video such as ST2110, Aspen, or compressed video such as MPEG and H.261. Each packet in an RTP flow has a sequence number from 0 to 65535 (see Figure 3). When the Cisco Nexus 9000 switch is used for video transport, it inspects every single packet of every single RTP flow that is traversing the switch. The flow table on the Cisco Cloud Scale application-specific integrated circuit (ASIC) powers the Cisco Nexus 9000. The Cisco Nexus 9000 switch is the first in the industry that can examine unsampled flow information. Using this feature, the Cisco Nexus 9000 can detect a gap in an RTP flow and generate an alert to notify the operator that packet loss has occurred. The alert also includes the flow that was impacted, its IP address, and its UDP port. The alert can be sent to a syslog or be streamed using [streaming telemetry](#) to an external collector.

Figure 3. RTP header

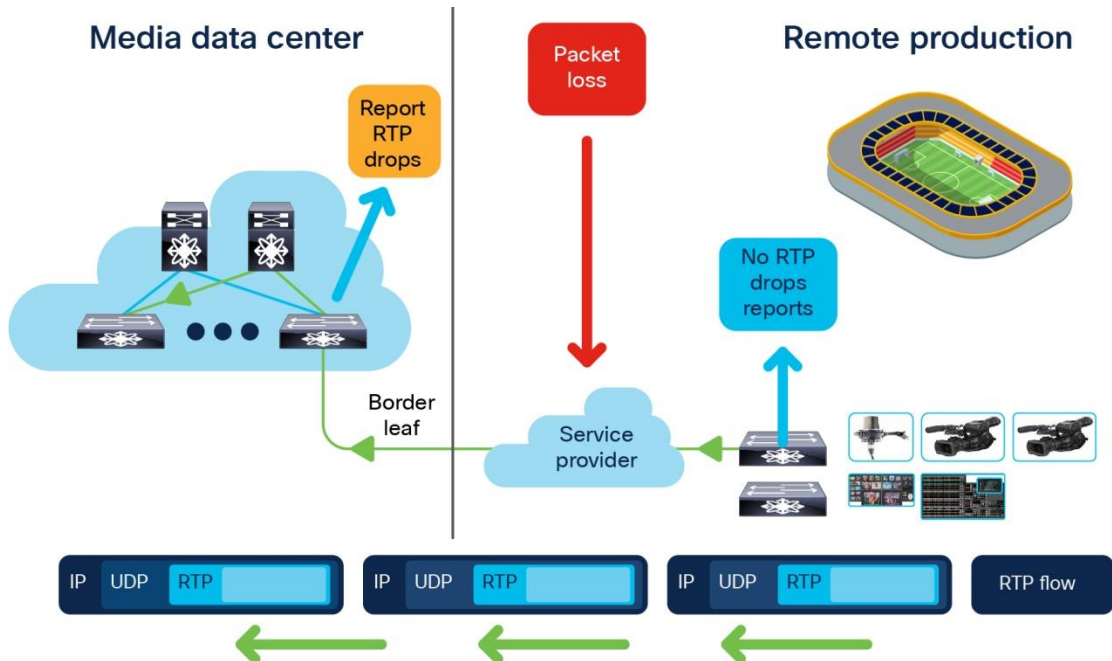
RTP headers include a sequence number which can be used to track loss



Feature in action: remote production use case

In the remote production example described earlier, suppose the IP network has one Cisco Nexus 9000 switch at the stadium and another at the production network. When bad video appears on screen, the RTP flow analytics feature is activated and tells the operator where the issue is taking place. If the Cisco Nexus 9000 switch at the stadium reports no issues, but the switch at the production edge toward the service provider does report an issue, you can conclude that the service provider is dropping the packet. Tracking down and isolating a problem like this one no longer takes several hours of troubleshooting; it only takes a matter of seconds (Figure 4 and 5). This capability is available on Cisco Nexus 9300-FX/FX2/FXP switches running Cisco NX-OS 9.3(1).

Figure 4. Detecting packet loss in real time



Configuring media flow analytics on Cisco Nexus 9000 Series Switches

Media flow analytics uses the NetFlow TCAM region on the Cisco Nexus 9000 switch. The region is carved by default. If, for any reason, the region has not been carved, ensure that it is allocated before enabling the feature.

For details on TCAM carving on Cisco Nexus 9000, please refer the [TCAM configuration guide](#).

Configuration steps

```
show hardware access-list tcam region | i ing-netflow
```

```
Ingress Netflow/Analytics [ing-netflow] size = 512
```

Step 1: Enable feature netflow.

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

Step 2: Enable User Defined Filter to match RTP traffic. (This requires reloading, saving the configuration, and then reloading the box.)

```
Switch(config)# udf netflow_rtp netflow-rtp
```

Step 3: Enable RTP flow monitoring.

```
Switch(config)# ip|ipv6 flow rtp <ACL>
```

Step 4: (Optional) Modify the ACL to match RTP traffic.

By default, the ACL matches all RTP traffic.

IP access list nfm-rtp-ipv4-acl

```
ignore routable
10 permit udp any range 16384 32767
```

If your application uses RTP but a different port range, this ACL can be modified to include the application, or a user-defined ACL can be included as shown in Step 3. Also, if monitoring is required for a few flows, then only those flows need to be included in the ACL. The switch can monitor up to 24,000 flows.

The feature works for unicast as well as multicast RTP flows (both IPv4 and IPv6).

Displaying RTP flows and errors

To display all RTP flows traversing the switch, we can use the “show flow rtp details” CLI.

```
show flow rtp details
```

RTP Flow timeout is 1440 minutes

IPV4 Entries

SIP	DIP	BD ID	S-Port	D-Port	Intf/Vlan Name	Packet Count	BytesPerSec
50.1.1.2	20.1.1.2	4151	16385	17999	Ethernet1/49/1	269207033	594468000
00:21:16	PST	Apr 07 2019					
20.1.1.2	50.1.1.2	4100	16385	18999	port-channel500	2844253	199000
00:21:59	PST	Apr 07 2019					

IPv6 Entries

SIP FlowStart	DIP	BD ID	S-Port	D-Port	Intf/Vlan Name	Packet Count	BytesPerSec
20::2 00:22:04 PST Apr 07 2019	50::2	4100	30000	31999	port-channel500	2820074	199000
50::2 00:21:16 PST Apr 07 2019	20::2	4151	30000	31999	Ethernet1/49/1	3058232	199000

As soon as the switch detects loss in RTP traffic, a notification is generated indicating a loss has taken place, and the list of all flows currently impacted can be seen with “**show flow rtp errors active**” CLI.

%NFM-1-RTP_FLOW_ERROR_DETECTED: Flow SIP: 30.30.1.2 DIP: 20.20.1.2 Interface: Ethernet1/98 loss detected

show flow rtp errors active

RTP Flow timeout is 1440 minutes

IPv4 Entries

SIP Packet Count Loss End	DIP BytesPerSec	BD ID FlowStart	S-Port	D-Port	Intf/Vlan Name Packet Loss Start
30.30.1.2 200993031 May 31 2019	20.20.1.2 10935633	4197 20:23:15 UTC	30000 May 30 2019	20392	Ethernet1/98 1558 03:48:32 UTC
20.20.1.2 204288988 May 31 2019	30.30.1.2 11114959	4196 20:23:15 UTC	30000 May 30 2019	20392	Ethernet1/97 222 03:48:30 UTC

The switch continues to monitor the flow for loss. When loss is no longer detected for a period of 10 seconds, the switch declares the loss has stopped and creates a notification:

%NFM-1-RTP_FLOW_ERROR_STOP: Flow SIP: 30.30.1.2 DIP: 20.20.1.2 Interface: Ethernet1/98 loss no longer detected

The flow now is placed in “show flow rtp error history”. It remains in the history for a period of 24 hours on the switch.

The following example shows a sample output for the **show flow rtp errors history** command:

Show flow rtp errors history

RTP Flow timeout is 1440 minutes

IPv4 Entries

SIP Packet Count Loss End	DIP BytesPerSec	BD ID FlowStart	S-Port	D-Port	Intf/Vlan Name Packet Loss Start
---------------------------------	--------------------	--------------------	--------	--------	-------------------------------------

```
20.20.1.2      30.30.1.2      4196      30000      20392      Ethernet1/97
204187441     11122753      20:23:15 UTC May 30 2019 2061      03:47:57 UTC
May 31 2019 03:47:57 UTC May 31 2019
```

```
30.30.1.2      20.20.1.2      4197      30000      20392      Ethernet1/98
199495510     10937237      20:23:15 UTC May 30 2019 1882      03:45:06
```

New architectures require new tools

As you transition your workflows to IP, you will need new solutions. In the past, networks lacked the ability to provide granular information on flow health or the ability to see problems in real time. Cisco offers the media industry a reliable, scalable, and flexible network and provides the tools you need to efficiently operate that network. Cisco is uniquely positioned in the networking industry with the ability to create intelligent ASICs that are complemented with innovative software.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

 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: <https://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)