White Paper

SPAN-on-Latency Feature on Cisco Nexus Switches: Troubleshoot Network Latency

Introduction

In a data center environment, often certain applications have high latency, but you can't identify which applications they are. The lack of visibility can make debugging difficult and hinder you from identifying why the network is experiencing latency higher than expected. In latency sensitive environments latency can be crucial. You must be able to determine which application is experiencing high latency so that you can easily fix the problem.

With the new Cisco® SPAN-on-Latency feature in the Cisco Nexus® 5600 switches, debugging latency problems is easy.

Note: The information in this document applies to Cisco Nexus 6000 Series of Switches as well.

The SPAN-on-Latency feature allows the system to apply the Cisco Switched Port Analyzer (SPAN) feature to packets that exceed a preconfigured latency threshold.

For high-latency flows, the system can be configured to send a copy to any preconfigured SPAN destination. This configuration creates a data set for analytics that can be used to identify which applications are affected by increased latency in the network. This feature can also be used to identify traffic flows that experience congestion. SPAN copies can be transported to a local analyzer port or to a remote analyzer using Encapsulated Remote SPAN (ERSPAN). The SPAN copies can be truncated to save bandwidth. Latency threshold configurations are set per port.

The Cisco Nexus 5600 platform has one buffer for data and a separate buffer for SPAN. Replicated traffic uses the SPAN buffer so that it doesn't affect the production traffic. Production traffic is always given higher priority than SPAN traffic in times of congestion.

© 2014 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

How SPAN-on-Latency Works

The SPAN-on-Latency source port is the egress port on which you monitor the latency. The SPAN-on-Latency session makes a copy of all high-latency packets that egress this port from any ingress port.

Figure 1 shows the SPAN-on-Latency source and destination ports.

Figure 1. SPAN-on-Latency Source and Destination Ports

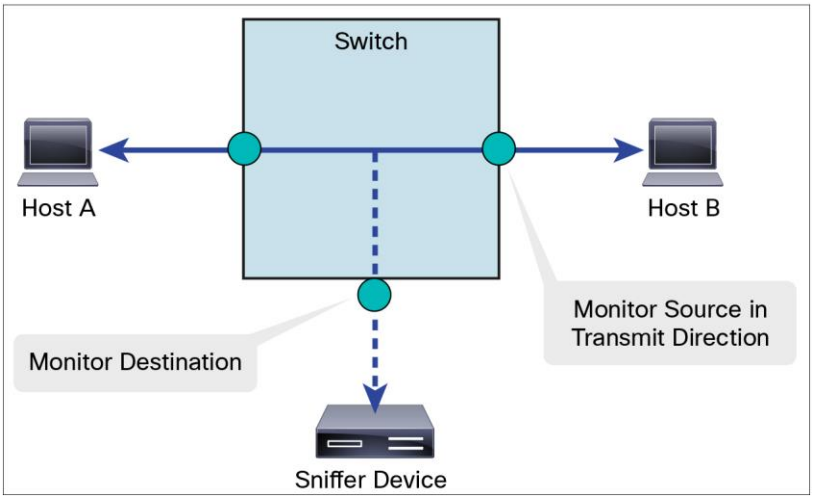

In Figure 1, suppose that the user expects data to be transferred from Host A to Host B in 850 nanoseconds (ns). However, because the buffers are nearly full at Host B's end (because many hosts simultaneously are sending data to Host B), this data transfer takes more than 2 milliseconds (ms). With the SPAN-on-Latency feature, the user can determine that the data is not traveling from Host A to Host B in the expected time frame. In this example, Host B is the SPAN-on-Latency source port, and the destination port is either a local or remote host port with a packet analyzer running to identify which flows took a long time to reach Host B. The latency threshold is set at 850 ns.

Configuration

The command syntax for the configuration is shown here:

```
[no] packet latency threshold <threshold_value>
```

- Run the command in the command-line interface (CLI) configuration mode.
- Configure the threshold_value in multiples of 8; otherwise, the value will be truncated to nearest multiple of 8.
- If the same command is entered a second time with different values, the configuration will be updated.
- The **no** form of the command does not requires the user to enter all the configured parameters.
- Threshold value configurations are per port for 40-Gbps interfaces; however, four 10-Gbps interfaces share the same configuration.

© 2014 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Here is an example:

```
switch(config)# int ethernet 3/7/1
switch(config-if)# packet latency threshold 1000


Interfaces Eth3/7/1, Eth3/7/2, Eth3/7/3, and Eth3/7/4 are configured with a
latency threshold of 1000ns.
```

You must issue the **clear hardware profile latency monitor all** command when the switch is reloaded or when a module is powered on. Until this command is issued, no packets are spanned.

Configuration of Span-on-Latency

The following example monitors the traffic to identify which flows reaching e2/1 and e2/2 took more than 1 us and 2 us, respectively. All these flows are spanned to local port e2/3. In addition, the maximum transmission unit (MTU) is truncated to save bandwidth.

```
switch#conf t
switch(config)#monitor session 1 type span-on-latency
switch(config-span-on-latency)#source interface ethernet 2/1
switch(config-span-on-latency)#source interface ethernet 2/2
switch(config-span-on-latency)#destination interface ethernet 2/3
switch(config-span-on-latency)#mtu 100
switch(config-span-on-latency)#no shut
```

Configure Latency Threshold on Interface 2/1

```
switch(config)#interface ethernet 2/1
switch(config-if)#packet latency threshold 1000
```

Configure Latency Threshold on Interface 2/2

```
switch(config)#interface ethernet 2/2
switch(config-if)#packet latency threshold 2000
```

Configure SPAN Destination Interface

```
switch(config)#interface ethernet 2/3
switch(config-if)#switchport monitor
```

Configuration of SPAN-on-Latency with ERSPAN

The following example monitors the traffic to identify which flows reaching e2/1 took more than 1 us. All these flows are spanned to a remote destination.

```
switch#conf t
switch(config)#monitor session 1 type span-on-latency-erspan
switch(config-span-on-latency-erspan)#erspan-id 20
switch(config-span-on-latency-erspan)#source interface ethernet 2/1
switch(config-span-on-latency-erspan)#destination ip 65.65.65.2
switch(config-span-on-latency-erspan)#vrf default
switch(config-span-on-latency)#no shut
switch(config)#monitor erspan origin ip-address 10.10.10.1 global.
```

© 2014 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Configure Latency Threshold on Interface 2/1

```
switch(config)#interface ethernet 2/1
switch(config-if)#packet latency threshold 1000
```

Configure Local Interface to Connect to Destination IP Address

```
switch(config)#interface ethernet 2/3
Switch(config-if)#no switchport
Switch(config-if)#ip address 10.10.10.1/24
```

Note: Connectivity to destination IP address 65.65.65.2 through Virtual Routing and Forwarding (VRF) should be set by default.

Feature Guidelines

Note the following guidelines when using the SPAN-on-Latency feature:

- Only Ethernet sources are supported (PortChannels are not supported). The source cannot be part of any other session: local SPAN, SPAN-on-Drop, etc. Only a single destination can be specified. However, multiple sources can be configured.
- Direction on the source interface is not supported. A packet egressing the source will be spanned.
- Fabric extender interfaces are not supported as sources; however, fabric ports are supported.
- Access control list (ACL)-based SPAN-on-Latency is not supported.
- Only one SPAN-on-Latency or SPAN-on-Latency with ERSPAN session can be active at a time.

Conclusion

SPAN-on-Latency can be especially useful in identifying applications that are experiencing congestion. If numerous applications are running and a few of them are experiencing high latency, you can apply SPAN to packets from each application and find out which application is causing latency problems in the network.

Especially in HFT environments, in which latency is extremely important, this feature can help diagnose any misconfigurations in the system, traffic patterns, port-speed mismatches, buffer discrepancies, or network problems that are leading to high latency.

For More Information

- http://www.youtube.com/watch?v=hCdBPbQje6s

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 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: 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)

Printed in USA C11-733021-00 11/14

© 2014 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.