

Latency Monitoring Tool on Cisco Nexus Switches: Troubleshoot Network Latency

Introduction

Networks often encounter problems related to latency. Troubleshooting such problems can be complicated. End-to-end system latency that is higher than expected may arise from the servers, applications, network design, burstiness of traffic to a single port, misconfiguration in the switch, etc.

Latency measurements can be used to identify the flows that are affected by latency. These statistics also allow administrators to plan network topologies, manage incident responses, and identify root causes of application problems in the network. The statistics also can be used to provide service-level agreements (SLAs) for latency-intensive applications.

The Cisco® latency monitoring tool on the Cisco Nexus® 5600 platform switches can help you troubleshoot latency concerns in the network with ease.

Note: All the details presented in this document apply to Cisco Nexus 6000 Series Switches as well.

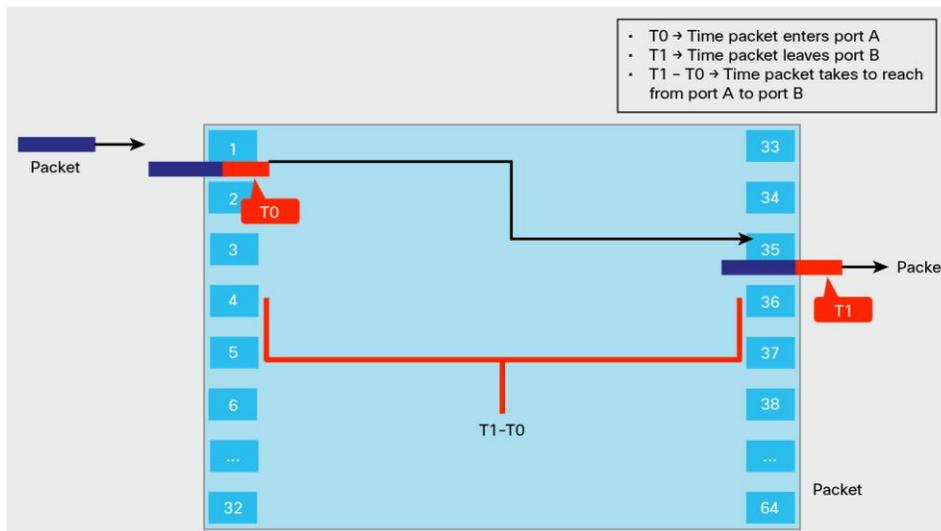
The latency monitoring tool marks each ingress and egress packet with a time-stamp value. To calculate the latency for each packet in the system, the switch compares the ingress time stamp with the egress time stamp. The tool allows you to display historical averages of latency between all pairs of ports as well as real-time latency data.

By default, Instantaneous mode latency monitoring is enabled between a pair of ports. This mode measures switch latency for each packet. The switch latency monitoring tool works at line rate; hence, packet sampling is not required.

This tool is fully implemented in hardware. Hence, it has no effect on the CPU or on data traffic.

Latency monitoring reveals the minimum, average, and maximum latency between a specified port pair and also maintains a latency histogram (accurate to several nanoseconds [ns]), as shown in Figure 1.

Figure 1. Latency Monitoring



In Figure 1, when a packet enters the source port, it is marked with a time stamp (t_0). When it exits the destination port, the time stamp is marked again (t_1). The latency monitoring tool essentially provides the time ($t_1 - t_0$).

Switch Latency Measurement Modes

The following modes can be configured between ingress and egress port pairs:

- Instantaneous mode: This is the mode enabled by default. No configuration is needed to see latency values in this mode. This mode is useful for getting a quick snapshot of port-pair latency.
- Linear Histogram mode: This mode is useful for seeing the number of packets that had the user-configured latency value.
- Exponential Histogram mode: This mode provides an exponential view of the linear histogram mode.
- Custom Histogram mode: This mode displays the number of packets that are in the low-latency range and the number of packets that are in the high-latency range.

Instantaneous Mode

In Instantaneous mode, the minimum, maximum, and average latency can be collected. This mode is enabled by default. No configuration is required.

Following is a sample of latency measurement output in Instantaneous mode. Eth 1/1 is the egress port and 1/6 is the ingress port.

```
switch# show hardware profile latency monitor interface ethernet 1/1 interface ethernet 1/6
```

```
-----
Egress Port: Ethernet1/1 Ingress Port: Ethernet1/6 Mode: Inst
-----
|           |           Minimum           |           Maximum           |           Average           |
-----
|  cnt  |           890 |           1020 |           960 |
-----
```

In the preceding output, **cnt** denotes the minimum, maximum, and average latency, in nanoseconds, of packets entering 1/6 and egressing 1/1.

Linear Histogram Mode

A linear histogram can show you the number of packets seen in a given latency range. For instance, suppose you want to count packets for the following ranges of latencies (in nanoseconds) separately: 800 to 849, 850 to 899, 900 to 949, and 950 to 999. You can do this by using a linear histogram.

You first configure a latency base value for the table to 800 ns. Then you program a step value of 50 ns. Using these two values, you partition the entry into four separate counters:

```
if base ≤ latency < base + step, increment CNT0
if base + step ≤ latency < base + 2 * step, increment CNT1
if base + 2 * step ≤ latency < base + 3 * step, increment CNT2
if base + 3 * step ≤ latency < base + 4 * step, increment CNT3
```

This histogram mode is linear because the latency ranges are increased linearly using the step value.

Following is a sample of latency measurement output in Linear Histogram mode. Eth 1/1 is the egress port and 1/6 is the ingress port.

```
switch# show hardware profile latency monitor interface ethernet 1/1 interface
ethernet 1/6
-----
Egress Port: Ethernet1/1 Ingress Port: Ethernet1/6 Mode: Linear Histogram
-----
| Range|          800-840|          840-880|          880-920|          920-960|
-----
| cnt |          123009|          556009|          640003|          450243|
-----
```

In the preceding output, **cnt** is the number of packets in the specified user-configured latency range. A base value of 800 ns and a step value of 40 ns were configured.

Exponential Histogram Mode

An exponential histogram provides the binning of latencies for ranges increasing exponentially. For instance, suppose that you want to count the packets in the following ranges of latencies: 850 to 899, 900 to 999, 1000 to 1199, and 1200 to 1599. You can achieve this by using the Exponential Histogram mode with the base value set to 800 and the step value set to 50 ns.

The exponential binning would count the packets as follows:

```
if base + step * 2^0 ≤ latency < base + step * 2^1, increment CNT0
if base + step * 2^1 ≤ latency < base + step * 2^2, increment CNT1
if base + step * 2^2 ≤ latency < base + step * 2^3, increment CNT2
if base + step * 2^3 ≤ latency < base + step * 2^4, increment CNT3
```

Following is a sample of latency measurement output in Exponential Histogram mode. Eth 1/1 is the egress port and 1/6 is the ingress port.

```
switch# show hardware profile latency monitor interface ethernet 1/1 interface ethernet 1/6
-----
```

```
Egress Port: Ethernet1/1 Ingress Port: Ethernet1/6 Mode: Exp Histogram
```

```
-----  
| Range|           840-880|           880-960|           960-1120|           1120-1440|  
-----  
| cnt |           45087|           55430|           36937|           26780|  
-----
```

In the preceding output, **cnt** is the number of packets in the specified user-configured latency range. A base value of 800 ns and a step value of 40 ns were configured.

Custom Histogram Mode

With a custom histogram, you can count the number of packets with latency within a specific range. Two counters are collected in this mode: a range specified by **latency_low** and **latency_high** boundaries, and a counter outside the first range. The counting is performed as follows:

```
if base + latency_low ≤ latency < base + latency_high, increment CNT0  
else increment CNT1
```

Following is a sample of latency measurement output in Custom Histogram mode. Eth 1/1 is the egress port and 1/6 is the ingress port.

```
switch# show hardware profile latency monitor interface ethernet 1/1 interface  
ethernet 1/6  
-----  
Egress Port: Ethernet1/1 Ingress Port: Ethernet1/6 Mode: Custom Histogram  
-----  
| Range|           840 <= Latency < 2000|           840 > Latency >= 2000|  
-----  
| cnt |           5678643|           458267|  
-----
```

In the preceding output, **cnt** is the number of packets in the specified user-configured latency range. A base value of 8 ns, a low-latency value of 832 ns, and a high-latency value of 1992 ns were configured.

Configuration

For configuration of latency monitoring on the Cisco Nexus 5600 platform, refer to the guide to [configuration details for the Cisco Nexus 5600 platform](#).

Note that you need to run the **clear hardware profile latency monitor all** command when the switch is reloaded or a new module is powered on.

Feature Guidelines

- Latency measurements are available only on physical interfaces.
- For latency measurements on Unified fabric ports, you need to add the latency of the unified port physical layer to the latency values displayed.

Conclusion

The latency measurement tool on the Cisco Nexus 5600 platform makes it easy to find out which application is experiencing problems related to high latency and to debug latency problems across the entire network. The tool can also help you determine whether the problems are caused by the switch or are a result of network problems.

This tool can be especially useful in data center environments in which low end-to-end system latency is extremely important.

For More Information

- <http://www.youtube.com/watch?v=q2ypBpo6wAY&feature=youtu.be>.



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)