

Monitor Microbursts on Cisco Nexus 5600 Platform and Cisco Nexus 6000 Series Switches

White Paper

October 2014

Contents

What You Will Learn	3
Introduction	3
Microbursts	3
Microburst Monitoring Feature Concept	4
Microburst Monitoring Configuration	5
Microburst Monitoring Verification	6
Conclusion	6
For More Information	7

What You Will Learn

This document provides an overview of the microburst monitoring feature on the Cisco Nexus® 5600 platform and Cisco Nexus 6000 Series Switches. With the proliferation of virtual machines in the data center, the likelihood that multiple devices will send traffic to the same destination increases dramatically. Because Ethernet is a serial transport medium, these burst events cause temporary congestion that can go undetected. This microburst congestion can lead to buffer exhaustion, which results in tail drops and lost packets. The Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches introduce a way to mark microbursts so that they can be observed, tracked, and understood so that effective mitigation actions can be implemented.

Introduction

A microburst is a short burst of traffic with a duration tending toward zero and an intensity tending toward infinite. Microbursts usually lead to temporary traffic buffering, increasing latency momentarily, and can cause traffic drops if the buffering capacity is exceeded. Microbursts in typical data center networks are actually very common, but they are often undetected. Numerous virtual machines, physical machines, and data storage devices communicate at the same time, which leads to intense traffic for extremely short periods of time, or microbursts. It is hard to predict where microbursts will occur and how traffic on a particular device will be affected.


Common data center switches provide basic per-port bandwidth monitoring and can show the current bandwidth on a port. However, these measurement tools are unable to detect microbursts because they are monitoring very large time intervals.

The solution is a feature that can detect unexpected bursts of data in extremely short periods of time, on the order of microseconds. The Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches support a microburst monitoring feature that detects microbursts that occur in microsecond time periods.

Microbursts

A microburst is very intense traffic in a very short period of time. Microbursts in a network most commonly occur at the moment when two or more devices send traffic to the same destination. For example, in Figure 1 two servers are sending traffic to the same destination and causing a microburst at the egress port on the switch.

Figure 1. Example of a Microburst in the Network



In Figure 1, three servers are connected to a switch. Two servers are sending and one server is receiving the traffic. The arrows from the servers on the left to the switch represent the traffic flow, and the graph next to the servers shows the traffic in each flow. The arrow on the right represents traffic flow from the egress port on the switch to the destination server. The graph above the switch shows traffic in the flow from the switch to the destination server, which is sum of the two signals from the source servers. At the marked point, two ingress signals create a microburst at the destination, and if the switch does not have enough resources to buffer the part of the destination traffic shown in orange, that traffic will be dropped.

The Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches introduce a microburst monitoring feature that observes this behavior and notifies the network administrator about its occurrence. The network administrator can then use the Switched Port Analyzer (SPAN) feature to collect traffic and mitigate the microburst.

Microburst Monitoring Feature Concept

Microburst monitoring on the Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches can detect microbursts on a per-port basis in both the ingress and egress traffic directions. In the case in which a server with a bursty application is connected to a port, microbursts typically occur in the ingress direction. When servers simultaneously access a storage array, microbursts occur in the egress direction on the port going to the storage array.

Microburst monitoring detects microbursts when a specific number of packets or a specific amount of data (in bytes) is exceeded during a defined time interval. The switch counts every occurrence and tallies the information in a table. This table allows the network administrator to quickly correlate application behavior with a server and network port. When this information is used in combination with logging information, the exact time of the event can be known.

By combining the microburst monitoring and SPAN features, the network administrator can correlate application events with network events. For example, in a typical use case the application team reports a performance problem. The network administrator can use microburst monitoring to see whether bursts are occurring and whether traffic is being dropped; the administrator can then use SPAN-on-drop to send any dropped packets to the monitor. This process allows the administrator to know whether those dropped packets belong to the reported application. If there are traffic microbursts that do not lead to traffic drop, SPAN can be used to identify the source of the microburst. This data allows direct correlation of network behavior with application behavior.

Microburst Monitoring Configuration

The Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches recognize a microburst when a specific amount of data exceeds a given threshold in a given time interval. The microburst threshold can be specified in two ways: as an amount of data in bytes, and as a percentage of link speed, in a given amount of time. Furthermore, the maximum number of microbursts can be defined for a port in each direction, and after that threshold is reached, the system generates a syslog message. The feature can be implemented only on physical interfaces on the Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches.

To configure microburst monitoring, on the interface on which microbursts are expected you need to be set the threshold values and the direction in which traffic will be observed, using the keywords **ingress** and **egress**. You can also monitor traffic in both directions at the same time on the same switch port, and if you want, the feature can be configured on all ports on the system.

The threshold parameter is used to define the number of bytes or percentage of link use. When the threshold is defined as a number of bytes, the range is from 1 byte to 68.7 GB (68719476735 bytes). To define the threshold in bytes, you use the keyword **size**. If the threshold is defined as a percentage of link use, you use the keyword **limit**; the limit percentage can range from 1 percent to 100 percent of link use.

The other parameter you need to specify is the time interval, in microseconds. The interval can range from 1 microsecond to 16.8 seconds (16777215 microseconds).

The following command configures an ingress burst size of 1 byte with a 1-microsecond time interval:

```
switch(config-if)# burst threshold ingress size 1 interval 1
```

The command shown here configures an egress burst limit of a 1 percent of link speed with a 1-microsecond time interval:

```
switch(config-if)# burst threshold egress limit 1 interval 1
```

The second step in configuring the microburst monitoring feature is to specify the maximum number of microbursts allowed within a time interval before the system generates a syslog message. The time interval is equal to 10 times the microburst threshold interval. This setting allows the system administrator to refine the notification intervals.

When microbursts are very common, the notification interval can be increased to avoid multiple notifications and save log entries. When microbursts are unexpected and notification is critical, then a smaller number can be set to notify network administrators quickly after microbursts occur. This setting can also be used to help correlate microbursts with time-based events that occur on the network. The maximum number of microbursts can also be defined in both the ingress and egress directions.

The following command configures the ingress burst maximum. After one microburst occurs, the system will generate a syslog message.

```
switch(config-if)# burst maximum ingress burst-count 1
```

The command shown here configures the egress burst maximum. After one microburst occurs, the system will generate syslog message.

```
switch(config-if)# burst maximum egress burst-count 1
```

Microburst Monitoring Verification

After you have configured all the parameters, the microburst monitoring feature on the Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches will monitor traffic behavior and count each microburst that is recognized. The display will show the port, burst count, and burst direction. To display burst counters, run the **show interface burst-counters** command. To display the count for only one interface, specifying the interface number after the keyword **interface**.

```
switch# show interface burst-counters
-----
| Interface | Ingress Bursts | Egress Bursts | Total Bursts |
-----
| Ethernet1/1 | 10 | N/A | 10 |
| fc2/1 | 15 | 0 | 15 |
```

The first column in the table indicates the interfaces on which microburst monitoring is configured. In the example, microburst monitoring is configured on two interfaces: one Ethernet interface and one Fibre Channel interface. The second column displays the number of microbursts that occurred in the ingress direction, going into the system, for particular interfaces. In the example, bursts occurred on both interfaces going into the device. The third column shows the number of bursts in the egress direction, going out of the system. In the example, Ethernet 1/1 shows N/A in the Egress Bursts column, indicating that microburst monitoring in the egress direction is not configured on this interface. The column for Interface FC2/1 shows that zero bursts have occurred in the egress direction, out of the system. The last column shows the total number of bursts per interface in both directions.

The following is a sample log message showing a microburst event:

```
2011 Mar 5 01:13:23 %US switch %$ VDC-1 %$ R-2-SYSTEM_MSG: Micro Burst has been
detected on ingress side on Ethernet1/1 - bigsurusd
```

Conclusion

The microburst monitoring feature on the Cisco Nexus 5600 platform and Cisco Nexus 6000 Series Switches lets you see the microbursts that have been silently occurring in the network. This real-time capability gives network administrators a better understanding of the application demands on the network over time. Through this real-time data analysis feature, congestion points and times can be understood so that the network architecture can be optimized to meet the increasing demands of next-generation applications.

For More Information

For more information about Microburst monitoring future configuration:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus6000/sw/qos/7x/b_6k_QoS_Config_7x/b_6k_QoS_Config_7x_chapter_01101.html

For more information about Cisco Nexus 5600 switches:

<http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/datasheet-c78-730760.html>

For more information about Cisco Nexus 6000 switches: <http://www.cisco.com/c/en/us/products/switches/nexus-6000-series-switches/datasheet-listing.html>




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)