

Implementing Buffer Utilization Histogram on Cisco Nexus 5600 and Nexus 6000 Switch Families

White Paper

January 2015

Contents

What You Will Learn	3
Introduction	3
Buffer Utilization Histogram Concept	3
Buffer Utilization Histogram Configuration	4
Buffer Utilization Histogram Verification	5
Conclusion	10

What You Will Learn

This white paper describes how to configure the Buffer Utilization Histogram feature on Cisco Nexus® 5600 and Nexus 6000 Switch Families, and explains the benefits that this feature offers. Traffic drops happen in every network, and it can be difficult to discover the cause of the drops. Sometimes the buffer overflow is the cause of the traffic drop. After a drop occurs it can be difficult to see the buffer state before and during traffic overflow.

Nexus 5600 and Nexus 6000 Switch Families introduced the Buffer Utilization Histogram feature. It provides a way to look in the past and see the state of buffer. The feature can show buffer state with time granularity for up to 250 milliseconds.

Introduction

A micro-burst of traffic is an example where buffer plays a role of preserving traffic not to be dropped. Buffer should not be shallow because it can easily overflow in the event of traffic microbursts. Furthermore, buffer that is too deep can affect traffic latency. This is not a desirable behavior in a network. Before a network administrator connects a new device to a switch, the admin should first gain data of buffer usage on the particular device. To do so, the admin needs to see how buffer has been used in the past.

With common tools on the switch, it is hard to see how buffer is used. Unfortunately, many of the regular tools in the network do not provide a history of usage. As a result, a user must connect a device to the switch (without previous knowledge about buffer usage) and wait to see how that new device will affect the traffic. The same situation may occur simply by installing a new server with a “bursty” application.


The Buffer Utilization Histogram feature on Cisco Nexus 5600 and Nexus 6000 switches provides per-port, instantaneous, and histogram buffer utilization. Buffer Utilization Histogram provides data of buffer utilization in both ingress and egress directions. For the egress direction, data is separated on multicast and unicast traffic. The results can be extracted into an XML file to help network administrators analyze and compare data with other processes in the network on the remote system.

Buffer Utilization Histogram Concept

The Buffer Utilization Histogram feature gives admins the ability to analyze the maximum queue depth and buffer utilization on Cisco Nexus 5600 and Nexus 6000 switches in real time. Every three ports of 40 Gigabit Ethernet or every 12 ports of 10 Gigabit Ethernet share a 25 MB buffer space. The devices maintain separate buffer on the ingress and egress directions. 15.6 MB are reserved for ingress and 8.6 MB are reserved for egress. The remaining space is used for switched port analyzer (SPAN) and control packets buffer. Egress buffer is divided into unicast and multicast buffer pools.

The Buffer Utilization Histogram feature brings instantaneous and histogram statistics on a per-port basis. The instantaneous mode gives current buffer states. To obtain the histogram of buffer utilization, software pools usage values from the hardware. Hardware pooling happens every 1 second for slow pooling mode and every 250 milliseconds for fast pooling. After execution show command, switch uses software pooling to make and display a histogram (Figure 1).

Figure 1. Buffer Utilization Histogram



A traffic microburst occurrence in the network is common, can fill buffer in a short period of time, and can increase latency. To correlate the microburst occurrence and buffer usage, the time between two pooling intervals needs to be shorter, due to the very short duration of a microburst. Default time between two buffer samples of one second can be reduced to 250 milliseconds and give higher granularity of collected data. In a steady environment, a sampling time of one second is ideal for tracking regular network behavior and spotting irregular traffic patterns.

To maintain a history of statistics, the switch collects data every hour and stores statistic in a "buffer_util_stats" file in the switch memory. Collecting data in this way provides the network administrator with the ability to look at data in the past. Furthermore, the data can be extracted to an XML file, which gives users the freedom to import buffer usage data into some other application. As a result, correlation of buffer usage and latency can be implemented using data from the XML file and Encapsulated Remote Switched Port Analyzer (ERSPAN) with a RFC 1588 PTP timestamp to calculate latency on a traffic analyzer or on a server with the traffic analyzer application.

Buffer Utilization Histogram Configuration

On Cisco Nexus 5600 and Nexus 6000 Series Switches instantaneous mode is enabled by default, and cannot be disabled. It gives buffer statistics in real time. By default, histogram mode is disabled, and can be enabled on a per-port basis. The Buffer Utilization Histogram feature can be enabled only on physical ports. To enable the histogram mode, on the particular interface on which buffer wants to be observed, execute the command:

```
switch(config-if)# hardware profile buffer monitor
```

The Buffer Utilization Histogram default sampling time is one second. When the administrator is familiar with traffic patterns, and without or with minimal numbers of traffic microbursts, this sampling period can provide a comprehensive view of usage buffer resources and provide information on how buffer was used in the past. Where a network administrator is not familiar with the network or traffic patterns, or where traffic microbursts are common, the recommended sampling period is 250 milliseconds, or fast sampling mode. This mode gives results with higher granularity, and can offer a better idea of how traffic microbursts or newly connected devices affect buffer. To enable fast sampling on the device, execute the following command in global configuration mode:

```
switch(config)# hardware profile buffer monitor sampling fast
```

The buffer usage histogram will be saved after changing the sampling mode. Once the network is stable and free of unexpected traffic behavior, the sampling period can be restored to the default sampling period of one second. To return the default sampling period to one second execute the following command:

```
switch(config)# no hardware profile buffer monitor sampling fast
```

When attaching a new device to the interface, and lower load to buffer is expected, the statistic can be erased. The commands below can erase the statistic on a particular interface or on all interfaces in which Buffer Utilization Histogram is enabled:

```
switch# clear hardware profile buffer monitor interface ethernet 1/21
switch# clear hardware profile buffer monitor
```

Buffer Utilization Histogram Verification

On Cisco Nexus 5600 and Nexus 6000 Series Switches the instantaneous mode is enabled by default and cannot be disabled. The instantaneous mode shows the buffer utilization at the moment of executing the command. To see how much buffer is utilized in a given moment, execute the following command:

```
switch# show hardware profile buffer monitor interface ethernet 1/1
```

```

+-----+
| Instant Ingress Buffer utilization per class per port. Every line |
| displays the number of cells utilized for a given port for each class |
| One cell represents 320 bytes |
+-----+
-----
Interface :          Eth1/1
-----
Total Port Instant Usage           1 (0.000320MB)
Remaining Asic Instant Usage      48836 (15.627520MB)
Per asic ingress cell count       48840 (15.628800MB)
+-----+-----+-----+-----+-----+-----+-----+-----+
port| class0| class1| class2| class3| class4| class5| class6| class7|
+-----+-----+-----+-----+-----+-----+-----+-----+
Eth1/1|      0|      0|      0|      1|      0|      0|      0|      0|
+-----+-----+-----+-----+-----+-----+-----+-----+
| Instant Egress Buffer utilization per unicast/multicast pool per port |
| Each line displays the number of cells utilized. |
| One cell represents 320 bytes |
+-----+

```

```

-----
                        Interface :          Eth1/1
-----
                        Total Port Instant Usage          85 (0.027200MB)
                        Remaining Asic Instant Usage      26772 (8.567040MB)
                        Per asic egress cell count        26912 (8.611840MB)
-----+-----+-----+-----+-----+-----+-----+
      port|   upool0|   upool1|   upool2|   upool3|   mcast pools|
-----+-----+-----+-----+-----+-----+
      Eth1/1|         0|         0|         0|         80|         5|
-----+-----+-----+-----+-----+

```

The output shows the buffer usage per port for the ingress and egress traffic directions. The table in the first line displays the interface on which the given state is shown. Buffer usage is shown as a number of used cells (320 bytes); the total number of used megabytes is shown in brackets. The same data is displayed in the table for both ingress and egress directions. In the ingress direction, all results are presented in the table. On the far left, the port on which buffer is observed is shown. The next seven columns describe classes of service for the traffic that are in the buffer, and how many cells are occupied in each of the classes from all occupied cells.

In the egress direction, buffer is divided for unicast and multicast, and shown in a table. The table on the left specifies the port on which buffer is observed. The unicast buffer is divided into four pools. The fifth column shows the cells that are occupied by multicast traffic.

Histogram mode can work in two modes: brief and detailed. In the brief mode only average values are shown over time. To display the brief buffer utilization histogram, the key words "history brief" need to be specified. The brief histogram is displayed below.

```

switch# show hardware profile buffer monitor interface ethernet 1/1 history brief
-----
                        Interface : Eth1/1
-----
                        Sampling Mode : Slow (1 second)
-----
                        Ingress Buffer Utilization Detected(in KB)
                        Per asic Ingress Total Usage (15.628800MB)
-----
      1 sec   |   5 sec   |   1 min   |   5 min   |   1 hour   |
-----
                        1.0|         0.4|         0.8|         0.8|         0.7|
-----
                        Egress Buffer Utilization Detected(Unicast|Multicast) (in KB)
                        Per asic Egress Total Usage (8.611850MB)
-----
      1 sec   |   5 sec   |   1 min   |   5 min   |   1 hour   |
-----

```

3.2| 12.8| 17.3| 12.2| 12.2| 8.7| 11.9| 6.9| 12.4| 7.1|

Results are displayed in a table. The first row reveals the interface in which the buffer is observed. The second row shows in which sampling mode the system is running. In the example, sampling mode is slow, and software pools statistics are recorded every second. For the brief histogram, fast sampling mode will give the same results. The header in third part of the table shows the total available space for the ingress buffer. Average buffer usage in the ingress direction over time is shown in the next two rows, starting from one second and going up to one hour. In the egress direction, unicast and multicast buffer states are displayed separately in the same row. The value on the left shows the unicast buffer usage, and the value on the right shows the multicast buffer usage. The values are displayed for the same time periods as on the ingress buffer.

The second mode of displaying the histogram is detail mode, which shows minimum, maximum, and average buffer usage. All statistics of buffer usage are grouped and displayed over time. In the histogram mode, the tables for slow and fast sampling are slightly different. For fast sampling, the column for one-second contains minimum, maximum, and average values. This is possible because, in fast sampling mode system, the switch can obtain four samples per second, and can display the lowest value as a minimum; the highest values as maximum; and can calculate an average value using all four samples. In the slow sampling mode, the one-second column does not have a minimum, maximum, or average value. In slow mode, only four samples are displayed in one column. To display the detailed histogram mode key words, the command "history detail" needs to be specified. The outputs for detailed a buffer usage histogram for slow and fast are displayed below.

```
switch# show hardware profile buffer monitor interface ethernet 1/1 history
detail
```

```
-----
Interface : Eth1/1
-----
```

```
-----
Sampling Mode : Slow (1 second)
-----
```

```
-----
Ingress Buffer Utilization Detected(Min|Max|Avg) (in KB)
Per ASIC Ingress Total Usage (15.628800MB)
-----
```

1 sec		5 sec		1 min		5 min		1 hour	
0.0	- -	0.0 0.6	0.2 0.0	1.6 0.7	0.0 1.6	0.8 0.0	1.6 0.7	0.0 1.6	0.7
0.0	- -	0.0 1.6	0.7 0.0	1.6 0.8	0.0 1.6	0.7	N/A		
0.0	- -	0.0 1.6	0.7 0.0	1.6 0.7	0.0 1.6	0.7	N/A		
0.0	- -	0.0 1.6	0.7 0.0	1.6 0.8	0.0 1.6	0.7	N/A		
0.6	- -	0.0 1.6	0.6 0.0	1.6 0.7	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.9	N/A	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.6	N/A	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.6	N/A	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.7	N/A	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.6	N/A	0.0 1.6	0.7	N/A		
N/A		0.0 1.6	0.7	N/A	0.0 1.6	0.8	N/A		

N/A		0.0	1.6	0.6	N/A		0.0	1.6	0.8	N/A		

Egress Unicast Buffer Utilization Detected (Min Max Avg) (in KB)												
Per ASIC Egress Total Usage (8.611850MB)												

1 sec		5 sec		1 min		5 min		1 hour				
1.6	-	1.6	12.8	6.7	0.0	41.6	11.8	0.0	48.0	11.9	0.0	49.6
3.2	-	1.6	33.6	12.0	0.0	41.6	12.0	0.0	46.4	12.3		N/A
12.8	-	0.0	46.4	12.6	0.0	48.0	12.5	0.0	46.4	11.9		N/A
8.0	-	1.6	28.8	10.8	0.0	44.8	12.4	0.0	46.4	11.8		N/A
8.0	-	0.0	40.0	16.3	0.0	48.0	12.0	0.0	46.4	11.2		N/A
N/A		0.0	33.6	13.4		N/A		0.0	46.4	12.2		N/A
N/A		0.0	22.4	11.1		N/A		0.0	48.0	12.4		N/A
N/A		0.0	28.8	9.0		N/A		0.0	48.0	12.6		N/A
N/A		1.6	20.8	7.3		N/A		0.0	49.6	12.9		N/A
N/A		0.0	22.4	10.8		N/A		0.0	49.6	13.2		N/A
N/A		1.6	41.6	16.2		N/A		0.0	49.6	13.5		N/A
N/A		0.0	25.6	8.4		N/A		0.0	49.6	12.8		N/A

Egress Multicast Buffer Utilization Detected (Min Max Avg) (in KB)												
Per ASIC Egress Total Usage (8.611850MB)												

1 sec		5 sec		1 min		5 min		1 hour				
14.4	-	1.6	12.8	5.8	0.0	22.4	7.1	0.0	24.0	6.9	0.0	25.6
3.2	-	0.0	20.8	8.0	0.0	22.4	6.4	0.0	25.6	7.5		N/A
1.6	-	1.6	20.8	9.3	0.0	24.0	6.8	0.0	25.6	7.7		N/A
4.8	-	1.6	20.8	9.5	0.0	22.4	6.6	0.0	25.6	7.7		N/A
6.4	-	1.6	20.8	9.2	0.0	22.4	7.7	0.0	27.2	7.3		N/A
N/A		0.0	20.8	6.8		N/A		0.0	25.6	7.8		N/A
N/A		0.0	16.0	5.4		N/A		0.0	25.6	7.6		N/A
N/A		1.6	20.8	7.1		N/A		0.0	25.6	6.7		N/A
N/A		0.0	20.8	6.4		N/A		0.0	24.0	6.4		N/A
N/A		0.0	22.4	8.2		N/A		0.0	22.4	6.3		N/A
N/A		0.0	20.8	7.0		N/A		0.0	22.4	6.9		N/A
N/A		0.0	20.8	6.6		N/A		0.0	24.0	6.4		N/A

switch# show hardware profile buffer monitor interface ethernet 1/1 history detail

Interface : Eth1/1

Sampling Mode : Fast (250 milliseconds)

Ingress Buffer Utilization Detected(Min|Max|Avg) (in KB)
Per asic Ingress Total Usage (15.628800MB)

1 sec	5 sec	1 min	5 min	1 hour
0.0 0.6 0.2	0.0 1.6 0.8	0.0 1.6 0.8	0.0 1.6 0.7	0.0 1.6 0.7
0.0 1.6 1.1	0.0 1.6 0.7	0.0 1.6 0.7	0.0 1.6 0.7	N/A
0.0 1.6 1.0	0.0 1.6 0.8	0.0 1.6 0.8	0.0 1.6 0.7	N/A
0.0 1.6 0.6	0.0 1.6 0.6	0.0 1.6 0.7	0.0 1.6 0.7	N/A
0.6 1.6 1.3	0.0 1.6 0.8	0.0 1.6 0.7	0.0 1.6 0.7	N/A
N/A	0.0 1.6 0.7	N/A	0.0 1.6 0.7	N/A
N/A	0.0 1.6 0.5	N/A	0.0 1.6 0.7	N/A
N/A	0.0 1.6 0.8	N/A	0.0 1.6 0.7	N/A
N/A	0.0 1.6 0.5	N/A	0.0 1.6 0.7	N/A
N/A	0.0 1.6 0.8	N/A	0.0 1.6 0.8	N/A
N/A	0.0 1.6 0.8	N/A	0.0 1.6 0.8	N/A
N/A	0.0 1.3 0.6	N/A	0.0 1.6 0.8	N/A

Egress Unicast Buffer Utilization Detected(Min|Max|Avg) (in KB)
Per asic Egress Total Usage (8.611850MB)

1 sec	5 sec	1 min	5 min	1 hour
1.6 40.0 16.4	0.0 30.4 11.9	0.0 44.8 12.4	0.0 46.4 12.3	0.0 49.6 12.4
0.0 32.0 14.8	3.2 36.8 14.4	0.0 48.0 12.0	0.0 46.4 11.9	N/A
0.0 32.0 12.4	0.0 48.0 14.6	0.0 46.4 10.8	0.0 46.4 11.8	N/A
0.0 9.6 4.4	0.0 28.8 10.6	0.0 43.2 13.1	0.0 46.4 11.2	N/A
3.2 27.2 15.2	0.0 32.0 12.2	0.0 46.4 11.9	0.0 46.4 12.2	N/A
N/A	0.0 38.4 12.3	N/A	0.0 48.0 12.4	N/A
N/A	0.0 30.4 14.0	N/A	0.0 48.0 12.6	N/A
N/A	0.0 38.4 9.3	N/A	0.0 49.6 12.9	N/A
N/A	1.6 28.8 14.0	N/A	0.0 49.6 13.2	N/A
N/A	0.0 40.0 15.8	N/A	0.0 49.6 13.5	N/A
N/A	0.0 33.6 7.4	N/A	0.0 49.6 12.8	N/A
N/A	1.6 33.6 12.4	N/A	0.0 48.0 11.9	N/A

Egress Multicast Buffer Utilization Detected(Min|Max|Avg) (in KB)
Per asic Egress Total Usage (8.611850MB)

```

-----
      1 sec      |      5 sec      |      1 min      |      5 min      |      1 hour      |
-----
0.0| 8.0|  3.2| 0.0|19.2|  7.4| 0.0|22.4|  6.6| 0.0|25.6|  7.5| 0.0|25.6|  7.1|
0.0|14.4|  4.4| 0.0|24.0|  9.3| 0.0|22.4|  7.7| 0.0|25.6|  7.7|      N/A      |
0.0|12.8|  6.4| 0.0|24.0|  8.6| 0.0|22.4|  6.7| 0.0|25.6|  7.7|      N/A      |
0.0|16.0|  7.6| 0.0|19.2|  7.1| 0.0|24.0|  8.1| 0.0|27.2|  7.3|      N/A      |
4.8|19.2| 12.4| 0.0|19.2|  6.4| 0.0|24.0|  6.7| 0.0|25.6|  7.8|      N/A      |
      N/A      | 0.0|19.2|  6.1|      N/A      | 0.0|25.6|  7.6|      N/A      |
      N/A      | 0.0|19.2|  6.3|      N/A      | 0.0|25.6|  6.7|      N/A      |
      N/A      | 0.0|20.8|  5.1|      N/A      | 0.0|24.0|  6.4|      N/A      |
      N/A      | 0.0|17.6|  5.8|      N/A      | 0.0|22.4|  6.3|      N/A      |
      N/A      | 0.0|20.8|  7.4|      N/A      | 0.0|22.4|  6.9|      N/A      |
      N/A      | 0.0|20.8|  6.6|      N/A      | 0.0|24.0|  6.4|      N/A      |
      N/A      | 1.6|20.8|  7.1|      N/A      | 0.0|22.4|  6.6|      N/A      |
-----

```

The output displays the buffer usage histogram as a table. The first line shows which interface buffer utilization histogram is displayed. The second line specifies the sampling mode. The first example is obtained for slow sampling mode and the second example is for fast sampling mode. The buffer utilization histogram states are divided into three sections. The first table displays ingress buffer utilization histogram data. The second displays egress unicast data. The last table belongs to the egress multicast buffer usage histogram.

Each ingress and egress table is divided by time periods from one second to one hour. For each time period, three columns are presented. The first column for each time period shows a minimum value of buffer usage. The second column shows maximum values. The third column shows average usage of buffer over time. All values are measured in kilobytes (KB).

For the one-second column, the system generates values by pooling samples every 250 milliseconds or one second, depending on the sampling period. Five samples from the one-second column make the first entry in the five-second column. Twelve samples from the five-second column make up the first entry in the one-minute column. Five samples from the one-minute column create the first entry in the five-minute column. Twelve samples from the five-minute column create the first entry in the one-hour column. All information is copied to a "buffer_util_stats" file every hour.

Conclusion

The Buffer Utilization Histogram feature on Cisco Nexus 5600 and Nexus 6000 switches reveals real-time buffer usage per port. The switches can obtain histograms of per-port buffer usage for up to an hour's duration. After an hour all data are saved in the switch memory. Furthermore, the states can be extracted into an XML file. All of these benefits give network administrators a comprehensive view of buffer usage in a particular network. Administrators can use this functionality to correlate latency in the network and buffer usage as a means of discovering traffic microbursts in the network.

For More Information

- For more information about Buffer Utilization Histogram, see the Buffer Utilization Histogram configuration guide chapter in [Cisco NX-OS Quality of Service Configuration Guide](#).
- For more information about Cisco Nexus 5600 platform switches, see <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 Series Switches, see <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)