



**Lab Test Report
DR100401G**

**Cisco Nexus 5000
Arista 7100 Switches**

26 April 2010

Miercom
www.miercom.com

Contents

Executive Summary	3
Overview	4
Burst Rate and Packet Loss Summary by Test Cases.....	5
Key Findings	7
Test Bed Diagram	9
How We Did It	9
2-to-Many Multicast Traffic – 24 Ports.....	11
128-byte Frame Size	11
<i>Figures 1 and 1a: Two-to-many multicast / 24 ports</i>	12
2-to-Many Multicast Traffic – 48 Ports.....	13
<i>Figures 2 and 2a: 2-to-many multicast – 48 ports</i>	14
4-to-Many Multicast Traffic – 48 Ports.....	15
<i>Figures 3 and 3a: 4-to-many multicast – 48 ports</i>	16
8-to-Many Multicast Traffic – 48 Ports.....	17
<i>Figures 4 and 4a: 8-to-many multicast – 48 ports</i>	18
23-to-1 Unicast Traffic – 24 Ports	19
24 Ports with 128-byte Frame Sizes	19
<i>Figures 5 and 5a: 23-to-1 unicast – 24 ports</i>	20
Full-Mesh Multicast Traffic	21
12 Ports with 128-byte Frame Size.....	21
<i>Figures 6 and 6a: Full-mesh multicast – 12 ports</i>	22
24 Ports with 128-byte Frame Size.....	23
<i>Figures 7 and 7a: Full-mesh multicast – 24 ports</i>	24
48 Ports with 128-byte Frame Size.....	25
<i>Figures 8 and 8a: Full-mesh multicast – 48 ports</i>	26
47-to-1 Unicast Traffic – 48 Ports	27
<i>Figures 9 and 9a: 47-to-1 unicast – 48 ports</i>	28
Full-Mesh Unicast Traffic	29
12 Ports with 128-byte Frame Size.....	29
<i>Figures 10 and 10a: Full-mesh unicast – 12 ports</i>	30
24 Ports with 128-byte Frame Size.....	31
<i>Figures 11 and 11a: Full-mesh unicast – 24 ports</i>	32
48 Ports with 128-byte Frame Size.....	33
<i>Figures 12 and 12a: Full-mesh unicast – 48 ports</i>	34
Arista 7148SX Flooding Issue.....	35

Executive Summary

Cisco commissioned Miercom to conduct an independent third-party performance test of the Cisco Nexus 5000 models 5010 and 5020 and Arista 7124S and 7148SX switches. The testing focused on evaluating the switches with bursty traffic profiles and included measurements of latency and packet loss with traffic for 128-byte frames. Typical real world application scenarios of full-mesh multicast, full-mesh unicast, two-to-many multicast and many-to-one multicast were used in the testing featured in this report.

In bursty traffic conditions the Arista switches dropped packets at low burst rates for 12-, 24- and 24-port tests. When the tests increased from 12 to 24 ports, the performance of the Arista switches degraded as well, dropping frames at a lower burst rate when the number of ports increased. The Cisco Nexus 5000 products did not drop packets at these bursts rates and demonstrated consistent results for all port tests.

Cisco Nexus 5010 and Arista 7124S were tested with 128-byte sized frames. Because the test results were strikingly similar between the two frame sizes, all test results that are shown are based on 128-byte sized packets. There was less than 0.1% difference in percentage of packet loss and latency delay between the two packet sizes. Based on the those results the Cisco Nexus 5020 and Arista were tested with 128-byte sized frames.

It is understandable that different traffic profiles used in testing can produce different performance results. Traffic profiles should include more bursty characterized traffic for testing of switching products that will be employed in environments, such as financial markets, that will have surges of high volume, of short duration, and small transaction type traffic.

Detailed test results follow and demonstrate the advantages in using the Cisco Nexus 5000 in a network environment that consists of high, bursty traffic. The tests in this report are intended to be reproducible for customers who wish to recreate them with the appropriate test and measurement equipment. Contact reviews@miercom.com for additional details on the configurations applied to the system under test and test tools used in this evaluation. Miercom recommends customers conduct their own needs analysis study, and test specifically for the expected environment for product deployment before making a product selection.

Arista was notified regarding this competitive testing analysis in accordance with the Miercom fair testing policy and was afforded the opportunity to participate in this evaluation. Representatives from Arista have contacted Miercom and Arista has offered a testing demonstration in the near future of their product with a challenge to these test results or to alternatively clarify Arista's position. There have been no identified inaccuracies in the test findings at this time of update to this report. Additional tests were conducted which revealed product anomalies, but are not included in this report as we are waiting on Arista Technical Support to confirm these findings.

The Cisco Nexus switches performed exceptionally well and demonstrated advantages over other competitive products we have evaluated, particularly in environments where surges of bursty transaction based traffic occur. The Cisco Nexus 5000 platform has proven it can support from 2.6 up to 26.7 times the level of burst traffic without packet loss compared to the Arista 7100 family switches. In review of the overall testing methodology there was no unfair advantage afforded to either vendor or product evaluated. We will nonetheless afford all vendors tested the opportunity to demonstrate their own product in retesting if requested.

Rob Smithers
CEO
Miercom

Overview

Real world networking environments are designed to handle high volumes of traffic and have the flexibility to handle sudden increases in smaller frame bursty traffic. Bursty traffic can cause network congestion or loss of data if the switch is not able to handle large volumes of traffic when these bursts occur.

Standard benchmark tests do not take into account architectural differences in buffering, congestion management and priority queuing. Two switches equally matched for handling line rate traffic of various traffic sizes can perform very differently when employed in a real network with peak demand loads causing oversubscription of ports and contention for resources.

Since there is no congestion with benchmark tests, there is no validation of how the queuing, scheduling and other internal switch advanced capabilities are implemented. In the real world, there is always short-lived congestion and how the platform handles traffic in such network environments is important.

Real world data center traffic characteristics include burst conditions with packet arrivals exhibiting an on/off pattern. This makes an existing steady state traffic engineering scheme not entirely applicable. To accurately evaluate the performance of a data center switch, it needs to be tested in network traffic conditions that consist of bursts of frames.

High bursty, fast data transmissions can lead to short lived network congestion and overfill the Ethernet switch buffers, causing packet loss that leads to TCP timeouts of 200 milliseconds or more. Such timeouts and the resulting delay can reduce application throughput by 90% and harm latency-sensitive applications such as those used on transactional based financial transactions, as well as voice and video applications. To overcome this problem and successfully handle bursty traffic, datacenter switches need large size buffers.

To evaluate the ability of the switch to handle traffic bursts, we ran tests for real world burst rates and used a bracketing technique, where the high and low initiating points eventually converge at a threshold burst rate. This is the point at where the switch begins to drop frames in multicast full-mesh, multicast two-to-many full-mesh and unicast full-mesh traffic profiles for 12, 24 and 48 ports under test.

We ran tests for 128-byte packet sizes as these reflect the smaller packet sizes used in transactional workloads such as high frequency trading.

The following page is a detailed chart which clearly shows the tests that were performed on each platform and where the detailed results can be found.

Burst Rate and Packet Loss Summary by Test Cases


Test	12 Ports				24 Ports				48 Ports			
	Cisco Nexus 5010		Arista 7124S		Cisco Nexus 5010		Arista 7124S		Cisco Nexus 5020		Arista 7148SX	
	Burst Rate / % Loss or Latency		Burst Rate / % Loss		Burst Rate / % Loss or Latency		Burst Rate / % Loss		Burst Rate / % Loss or Latency		Burst Rate / % Loss	
2-to-Many Multicast	N/A		N/A		1250	No Loss	1250	0.5%	1250	No Loss	1250	0.1%
					68.8	μsec delay	---	---	77.7	μsec delay	---	---
					3250	2.5%	3250	31.0%	3250	1.0%	3250	20.9%
See details:					Pages 11 - 12				Pages 13 - 14			
4-to-Many Multicast	N/A		N/A		N/A		N/A		425	No Loss	425	0.4%
									79.2	μsec delay	---	---
									2100	1.1%	2100	48.3%
See details:									Pages 15 - 16			
8-to-Many Multicast	N/A		N/A		N/A		N/A		180	No Loss	180	0.2%
									78.3	μsec delay	---	---
									1750	0.2%	1750	77.1%
See details:									Pages 17 - 18			
23-to-1 Unicast	N/A		N/A		60	No Loss	60	5.0%	N/A		N/A	
					81.8	μsec delay	---	---				
					1600	3.9%	1600	92.3%				
See details:					Pages 19 - 20							
Full-Mesh Multicast	150	No Loss	150	15.7%	60	No Loss	60	5.0%	35	No Loss	35	5.1%
	95.8	μsec delay	---	---	83.3	μsec delay	---	---	100	μsec delay	---	---
	1600	0.6%	1600	83.9%	1600	4.2	1600	92.3%	1600	3.3%	1600	94.3%
See details:	Pages 21 - 22				Pages 23 - 24				Pages 25 - 26			
47-to-1 Unicast	N/A		N/A		N/A		N/A		30	No Loss	30	1.3%
									44.4	μsec delay	---	---
									1600	3.1%	1600	91.1%
See details:									Pages 27 - 28			
Full-Mesh Unicast	60	No Loss	60	6.8%	15	No Loss	15	9.6%	15	No Loss	15	9.5%
	30.7	μsec delay	---	---	18.5	μsec delay	---	---	43.8	μsec delay	---	---
	275	0.8%	275	50.7%	120	2.5%	120	64.6%	55	1.7%	55	54.4%
See details:	Pages 29 - 30				Pages 31 - 32				Pages 33 - 34			

Note: Burst Rate is measured in Packets per Burst

Chart details results of test cases showing at what point burst rates start to cause packets to be lost. In the event that the burst rate does not cause a packet loss, but instead causes latency, that delay is noted in microseconds (μsecs).


Empty cells (N/A) indicate the test was not run for that scenario/product. All results shown were tested with 128 bytes per frame.

Max Burst without Loss (Full Mesh Multicast)



Comparison of maximum burst rate when packet loss is observed with full-mesh multicast.

Max Burst without Loss (Full Mesh Unicast)



Comparison of maximum burst rate when packet loss is observed with full-mesh unicast.

Key Findings

Cisco Nexus 5010 and Arista 7124S

2-to-Many Multicast

Arista 7124S starts dropping packets at bursts of 1250 frames with a 0.05% traffic loss observed. 1250 frames equates to 148.0 μ sec at 128 bytes. The Nexus 5010 continues to forward frames with increasing average latency until the burst size reaches 3250 frames. At this burst rate, the Arista switch drops 31.0% of traffic and the Nexus 5010 drops 2.5% of traffic. For this traffic pattern the Nexus can support 2.6 (3250/1250) times the level of burst traffic versus Arista without loss.

23-to-One Unicast

The Arista switch starts dropping packets at bursts of 60 frames with a 5.0% traffic loss observed. 60 frames equates to 7.1 μ sec at 128 bytes. The Nexus 5010 continues to forward frames with increasing average latency until the burst size reaches 1600 frames. At this burst rate, the Arista switch drops 92.2% of traffic and the Nexus 5010 drops 3.9% of traffic. For this traffic pattern the Nexus can support 26.7 (1600/60) times the level of burst traffic versus Arista without loss.

Full-Mesh Multicast

12 Ports

The Arista switch starts dropping packets at bursts of 150 frames with a 15.7% traffic loss observed. 150 frames equates to 17.8 μ sec at 128 bytes. Cisco Nexus continues to forward frames with increasing average latency until the burst size reaches 1600 frames. At this burst rate, the Arista switch drops 83.9% of traffic and the Nexus 5010 drops 0.6% of traffic. For this traffic pattern the Nexus 5010 can support 10.7 (1600/150) times the level of burst traffic versus Arista without loss.

24 Ports

Arista 7124S starts dropping packets at bursts of 60 frames with a 5.0% traffic loss observed. 60 frames equates to 7.1 μ sec at 128 bytes. The Nexus 5010 continues to forward frames with increasing average latency until the burst size reaches 1600 frames. At this burst rate, the Arista switch drops 92.3% of traffic and the Nexus 5010 drops 4.2% of traffic. For this traffic pattern the Nexus can support 26.7 (1600/60) times the level of burst traffic versus Arista without loss.

Full-Mesh Unicast

12 Ports

The Arista 7124S starts dropping packets at bursts of 60 frames with a 6.8% traffic loss observed. 60 frames equates to 7.1 μ sec at 128 bytes. Cisco Nexus continues to forward frames with increasing average latency until the burst size reaches 275 frames. At this burst rate, the Arista switch drops 50.7% of traffic and the Nexus 5010 drops 0.6% of traffic. For this traffic pattern the Nexus 5010 can support 4.6 (275/60) times the level of burst traffic versus Arista without loss.

24 Ports

The Arista switch starts dropping packets at bursts of 15 frames with a 9.6% traffic loss observed. 15 frames equates to 1.8 μ sec at 128 bytes. The Nexus 5010 continues to forward frames with increasing average latency until the burst size reaches 120 frames. At this burst rate, the Arista switch drops 64.6% of traffic and the Nexus 5010 drops 1.9% of traffic. For this traffic pattern the Nexus can support 8.0 (120/15) times the level of burst traffic versus Arista without loss.

Cisco Nexus 5020 and Arista 7148SX

2-to-Many Multicast

Arista 7148SX starts dropping packets at bursts of 1250 frames with a 0.1% traffic loss. 1250 frames equates to 148.0 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing average latency until the burst size reaches 3250 frames. At this rate, the Arista switch drops 20.9% of traffic and the Nexus 5020 drops 1% of traffic. With this traffic pattern the Nexus can support 2.6 (3250/1250) times the level of burst traffic versus Arista without loss.

4-to-Many Multicast

Arista 7148SX starts dropping packets at bursts of 425 frames with a 0.4% traffic loss. 425 frames equates to 50.3 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing latency until the burst size reaches 2100 frames. At this rate, the Arista switch drops 48.3% of traffic and the Nexus 5020 drops 1.1% of traffic. In this traffic pattern the Nexus can support 4.9 (2100/425) times the level of burst traffic versus Arista without loss.

8-to-Many Multicast

Arista 7148SX starts dropping packets at bursts of 180 frames with a 0.2% traffic loss. 180 frames equates to 21.3 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing latency until the burst size reaches 1750 frames. At this rate, the Arista switch drops 77.1% of traffic and the Nexus 5020 drops 0.2% of traffic. With this traffic pattern the Nexus can support 9.7 (1750/180) times the level of burst traffic versus Arista without loss.

Full Mesh Multicast

Arista 7148SX starts dropping packets at bursts of 35 frames with a 5.1% traffic loss. 35 frames equates to 4.14 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing latency until the burst size reaches 1600 frames. At this rate, the Arista switch drops 94.3% of traffic and the Nexus 5020 drops 3.3% of traffic. In this traffic pattern the Nexus can support 45.7 (1600/35) times the level of burst traffic versus Arista without loss.


47-to-1 Unicast

Arista 7148SX starts dropping packets at bursts of 30 frames with a 1.3% traffic loss. 30 frames equates to 3.5 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing latency until the burst size reaches 1600 frames. At this rate, the Arista switch drops 91.1% of traffic and the Nexus 5020 drops 3.1% of traffic. With this traffic pattern the Nexus can support 53.3 (1600/30) times the level of burst traffic versus Arista without loss.

Full-Mesh Unicast

Arista 7148SX starts dropping packets at bursts of 15 frames with a 9.5% traffic loss. 15 frames equates to 1.78 μ sec at 128 bytes. The Nexus 5020 continues to forward frames with increasing latency until the burst size reaches 55 frames. At this rate, the Arista switch drops 54.4% of traffic and the Nexus 5020 drops 1.7% of traffic. In this traffic pattern the Nexus can support 3.6 (55/15) times the level of burst traffic versus Arista without loss.

Test Bed Diagram



How We Did It

The tests we used measured the latency and loss characteristics of datacenter traffic patterns traversing through a network platform. Seven traffic profiles were simulated: 2-to-many multicast, 4-to-many multicast, 8-to-many multicast, 23-to-1 unicast, full-mesh multicast, 47-to-many unicast and full-mesh unicast. These profiles were selected as they are prevalent traffic flow models in high performance networks.

All tests were conducted for 128-byte frame sizes. The term packet and frame are used interchangeably in this report. Note that throughput and overall performance numbers were reported using full Ethernet “frame” size (including header information) and the Ixia test equipment was applied with configuration settings relative to packets, as the control and delivery is IP packet manipulation.

The setup used to conduct burst testing is shown in the above test bed diagram. Bursty traffic conditions were simulated for 128-byte frame sizes utilizing binary search with incremental increases in packets per burst until packet loss was recorded.

The systems under test included:

Cisco Nexus 5010 switch running Cisco NX-OS v 4.1(3)N1(1). The switch is a 1 RU, 10 Gigabit Ethernet switch supporting 20 fixed 10 Gigabit Ethernet ports and one expansion module slot fitted

with 6 port 10 GE module with only four used for testing to ensure the products compared had the same port count for fair comparison.

Cisco Nexus 5020 switch running Cisco NX-OS v 4.1(3) N2 (1a). The switch is a 2 RU, 10 Gigabit Ethernet switch supporting 40 fixed 10 Gigabit Ethernet ports and two expansion module slot fitted with 8 port 10 GE module with only four used on each for testing to ensure the products compared had the same port count for fair comparison.

Arista 7124S switch running EOS v 4.3.0. The switch is a 1 RU supporting (24) 10-Gigabit Ethernet ports.

Arista 7148SX switch running EOS v 4.3.2. The switch is a 1 RU supporting 48 fixed 10-Gigabit Ethernet ports.

Ixia: The traffic generator used in this evaluation was the Ixia (www.ixiacom.com) XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with test application IxExplorer. Eight port 10GE LSM cards on the chassis were used to drive Layer 2 traffic streams simulating bursty traffic conditions.

For example, to run a test with 100 frames burst for 128-byte frame size, Ixia generates bursts at a burst time of 11.84 microseconds using the following calculations:

- Time to transmit 1 byte at 10GE = 0.8ns
- Total Frame Size = (128-Bytes) + MAC Preamble (8 Bytes) + Inter Frame Gap (12 bytes) = 148 bytes
- 148 bytes * 0.8ns/bytes*100 frames burst = 11.84 microsecond burst time

The tests in this report are intended to be reproducible for customers who wish to recreate them with the appropriate test and measurement equipment. Contact reviews@miercom.com for additional details on the configurations applied to the System Under Test and test tools used in this evaluation. Miercom recommends customers conduct their own needs analysis study and test specifically for the expected environment for product deployment before making a product selection.

2-to-Many Multicast Traffic – 24 Ports

Objective

To evaluate the ability of the switch to handle multicast bursts in a two-to-many traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5010 and the Arista 7124S with 24 ports connected to Ixia traffic generator. Testing was conducted for a 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts in a binary search fashion. The metrics to be recorded were packets sent, packets received, packet loss and latency with latency calculated only at burst rates with no packet loss. Since dropped packets equate to infinite latency, latency is not shown once a platform starts dropping packets.

Test Setup

Two Ixia ports were configured to send traffic to unique multicast addresses per port. The remaining 22 ports would send an IGMP join for each multicast group address to simulate a two-to-many traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations

128-byte Frame Size


- **Cisco Nexus 5010**
Packet loss occurred when the burst rate exceeded 3000 packets per burst. Up to 2.5% packet loss was observed when 3250 packets per burst were applied.
- **Arista 7124S**
Packet drops were observed at burst rates of 1000 packets per burst and up to 0.5% packet loss at 1250 packets/burst and 31% at 3250 packets per burst was applied.
- It was noted that the Arista 7124S at 1250 packets per burst dropped 0.5% of its packets, while Cisco Nexus 5010 experienced a 68.8 μ sec delay and continued to forward packets. Latency is acceptable rather than packet loss.

Lengths of Bursts (Microseconds)

Frames / Burst	128-bytes
1250	148 μ secs
3250	384.80 μ secs

Figures 1 and 1a: Two-to-many multicast / 24 ports


Loss versus Latency



The average latency and packet loss percentage at different ingress burst rates for 128-byte frames are compared for the Cisco Nexus 5010 and the Arista 7124S. Packet drops for Arista occur at rates exceeding 1000 packets per burst while Cisco starts to drop packets at rates exceeding 3000 packets per burst.

See below enlargement

Enlargement



An enlarged view of the x-axis, packet loss and latency.

2-to-Many Multicast Traffic – 48 Ports

Objective

To evaluate the ability of the switch to handle multicast bursts in a two-to-many traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-bytes frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts in a binary search fashion. The metrics to be recorded were packets sent, packets received, packet loss and latency, calculated only at burst rates without packet loss. Since dropped packets equate to infinite latency, latency is not shown once a platform starts dropping packets.

Test Setup

Two Ixia ports were configured to send traffic to unique multicast addresses per port. The remaining 46 ports would send an IGMP join for each multicast group address to simulate a two-to-many traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
We observed packet loss when the burst rate exceeded 3000 packets per burst, with 1% packet loss for 3250 packets per burst.
- **Arista 7148SX**
Packet drops were observed at rates exceeding 1000 packets per burst. Up to 0.1% packet loss for 1250 packets per burst and 20.9% for 3250 packets per burst was observed.
- When the Arista 7148SX reached 1250 packets per burst, 0.1% of packets were dropped. Cisco Nexus 5020 experienced a 77.7 μ sec latency delay at 1250 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows the burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
1250	148 μ secs
3250	384.80 μ secs

Figures 2 and 2a: 2-to-many multicast – 48 ports



Average latency and packet loss percentages at different burst rates for 128-byte frames are compared for the Cisco Nexus 5020 and the Arista 7148SX. Once the bursts exceed 1000 packets per burst, Arista drops packets and continues to drop as burst rates increase. The Nexus 5020 switch starts to drop packets when bursts exceed 3000 packets per burst.

Enlargement



An enlarged view of the x-axis, packet loss and latency.

4-to-Many Multicast Traffic – 48 Ports

Objective

To evaluate the ability of the switch to handle multicast bursts in a four-to-many traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to Ixia traffic generator. Testing was conducted with 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts in a binary search fashion. The metrics recorded were packets sent, packets received, packet loss and latency, calculated only at burst rates without packet loss.

Test Setup

Four Ixia ports were configured to send traffic to unique multicast addresses per port. The remaining 44 ports would send an IGMP join for each multicast group address to simulate a four-to-many traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
Packet loss occurred when the burst rate exceed 2000 packets per burst. Up to 1.1% packet loss was observed when 2100 packets per burst were applied.
- **Arista 7148SX**
Packet drops were observed when burst rate exceeded 400 packets per burst. We noted 0.4% packet loss at 425 packets per burst and 48.3% at 2100 packets per burst.
- When the Arista 7148SX reached 425 packets per burst, 0.4% of packets were dropped. Cisco Nexus 5020 experienced a 79.2 μ sec latency delay at 425 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
425	50.3 μ secs
2100	248.6 μ secs

Figures 3 and 3a: 4-to-many multicast – 48 ports



Average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7148SX and Cisco Nexus 5020. Once the burst rates exceed 400 packets per burst, Arista drops packets and continues to drop as burst rates are increased. The Nexus 5020 starts to drop packets when rate exceeds 2000 packets per burst.

Enlargement



An enlarged view of the x-axis, packet loss and latency.

8-to-Many Multicast Traffic – 48 Ports

Objective

To evaluate the ability of the switch to handle multicast bursts in a eight-to-many traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts. The metrics recorded were packets sent, packets received, packet loss and latency, calculated at burst rates without packet loss.

Test Setup

Eight Ixia ports were configured to send traffic to unique multicast addresses per port. The remaining 40 ports would send an IGMP join for each multicast group address to simulate an eight-to-many traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
Packet loss occurred when the burst rate exceeded 1700 packets per burst. Up to 0.2% packet loss was observed when 1750 packets per burst were applied.
- **Arista 7148SX**
Packet drops were observed when burst rate exceeded 175 packets per burst. We noted 0.2% packet loss at 180 packets per burst and 77.1% at 1750 packets per burst.
- When the Arista 7148SX reached 180 packets per burst, 0.2% of packets were dropped. Cisco Nexus 5020 experienced a 78.3 μ sec latency delay at 180 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.


Lengths of Bursts (Microseconds)

Frames / Burst	128-bytes
180	21.3 μ secs
1750	207.2 μ secs

Figures 4 and 4a: 8-to-many multicast – 48 ports



Average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7148SX and Cisco Nexus 5020. Once the burst rates exceed 175 packets per burst, the Arista 7148SX drops packets and continues to drop as burst rates are increased. The Nexus 5020 starts to drop packets when rate exceeds 1700 packets per burst.



An enlarged view of the x-axis, packet loss and latency.

23-to-1 Unicast Traffic – 24 Ports

Objective

To evaluate the ability of the switch to handle unicast bursts in a 23-to-1 traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5010 and the Arista 7124S with 24 ports connected to the Ixia traffic generator. Testing was conducted for a frame size of 128 bytes. The test was configured to step through multiple burst rates representing real world datacenter frame bursts in a binary search fashion. The metrics to be recorded were packets sent, packets received, packet loss and latency. Latency was calculated only at burst rates with no packet loss. Since dropped packets equate to infinite latency, latency is not shown once a platform starts dropping packets.

Test Setup

On the Ixia traffic generator port 1 through port 23 were set to send unicast traffic to port 24, and port 24 set up to send traffic to port 1, simulating a 23-to-1 unicast traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations

24 Ports with 128-byte Frame Sizes


- **Cisco Nexus 5010**
We observed packet loss when burst rate exceeded 1500 packets per burst with 3.9% packet loss at 1600 packets per burst.
- **Arista 7124S**
Packet drops were observed when burst rate exceeds 50 packets per burst. There is a 5% packet loss at 60 packets per burst and 92.3% at 1600 packets per burst.
- It was noted that the Arista 7124S at 60 packets per burst dropped 5% of its packets, while Cisco Nexus 5010 experienced an 81.8 μ sec delay and continued to forward packets. Latency is acceptable rather than packet loss.

Lengths of Bursts (Microseconds)

Frames / Burst	128-bytes
60	7.1 μ secs
1600	189.44 μ secs

Figures 5 and 5a: 23-to-1 unicast – 24 ports


Loss versus Latency



Average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7124S and Cisco Nexus 5010. Once the rates exceed 50 packets per burst, the Arista 7124S drops packets and continues to drop as burst rates are increased. The Nexus 5010 starts to drop packets when rate exceeds 1500 packets per burst.

See below enlargement

Enlargement



An enlarged view of the x-axis, packet loss and latency.

Full-Mesh Multicast Traffic

Objective

To evaluate the ability of the switch to handle continuous bursts of frames and to record packet loss and latency when the switch is subjected to multicast bursts in a full-mesh traffic profile.

Description

This test was conducted for the Cisco Nexus 5010 and the Arista 7124S with 24 ports connected to an Ixia traffic generator, and then repeated using 12 ports. The Ixia traffic generator was set up to send traffic to a single multicast address from each connected port. Each port sent an IGMP join to each multicast group address, except the group it sent to, and thus simulated a full-mesh traffic profile. Traffic was offered to each interface and destined to all other interfaces. Testing was conducted with a 128-byte frame size. The test was configured to step through multiple burst rates representing real world frame bursts in a binary search fashion. The metrics recorded were packets sent, packets received, packet loss and latency. Latency was calculated only for burst rates without packet loss. There isn't any method to calculate latency for dropped packets, and dropped packets equate to infinite latency.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


12 Ports with 128-byte Frame Size

- **Cisco Nexus 5010**
Packet loss was observed when the burst rate exceeded 1500 packets per burst, with 0.6% packet loss when 1600 packets per burst was realized.
- **Arista 7124S**
Packet drops were observed at burst rates of 150 packets per burst with a 15.7% traffic loss, and an 83.9% packet loss at 1600 packets per burst.
- It was noted that the Arista 7124S at 150 packets per burst dropped 15.7% of its packets, while Cisco Nexus 5010 experienced a 95.8 μ sec latency delay and continued to forward packets. Latency is acceptable rather than packet loss.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
150	17.76 μ secs
1600	189.44 μ secs

Figures 6 and 6a: Full-mesh multicast – 12 ports



Average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7124S and Cisco Nexus 5010. Once the rates exceed 150 packets per burst, the Arista 7124S drops packets and continues to drop as burst rates are increased. The Nexus 5010 starts to drop packets when rate exceeds 1500 packets per burst.

Enlargement



An enlarged view of the x-axis, packet loss and latency.


24 Ports with 128-byte Frame Size

- **Cisco Nexus 5010**
We observed packet loss when the burst rate exceeded 1500 packets per burst, with 4.2% packet loss when 1600 packets per burst was realized.
- **Arista 7124S**
Packet drops were observed at burst rates exceeding 50 packets per burst and up to 92.3% packet loss at 1600 packets per burst.
- It was noted that the Arista 7124S at 60 packets per burst dropped 5% of its packets, while Cisco Nexus 5010 experienced an 83.3 μ sec delay and continued to forward packets. Latency is more readily acceptable in data transmission than packet loss.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
60	7.1 μ secs
1600	189.44 μ secs

Figures 7 and 7a: Full-mesh multicast – 24 ports



Average latency and packet loss percentages at different burst rates for 128-byte frames are compared for the Cisco Nexus 5010 and the Arista 7124S. Packet drops for the 7124S start to occur at burst values that exceed 50 packets per burst, and packet drops for the Nexus 5010 switch happen at burst values greater than 1500 packets per burst.

Enlargement



An enlarged view of the x-axis, packet loss and latency.

48 Ports with 128-byte Frame Size

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts in a binary search fashion. The metrics recorded were packets sent, packets received, packet loss and latency, calculated at burst rates without packet loss.

Test Setup

The Ixia traffic generator was set to send traffic to a single multicast address from each connected port. Each port sent an IGMP join to each multicast group address, except the group it sent to, and therefore simulated a full-mesh traffic profile. Traffic was offered to each interface and destined to all other interfaces. There were a total of 48 multicast groups for this test.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
Packet loss occurred when the burst rate exceeded 1500 packets per burst. Up to 3.3% packet loss was observed when 1600 packets per burst were applied.
- **Arista 7148SX**
Packet drops were observed when burst rate exceeded 30 packets per burst. There is a 5.1% packet loss at 35 packets per burst and 94.3% at 1600 packets per burst.
- When the Arista 7148SX reached 35 packets per burst, 5.1% of packets were dropped. Cisco Nexus 5020 experienced a 100 µsec latency delay at 35 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
35	4.14 µsecs
1600	189.4 µsecs

Figures 8 and 8a: Full-mesh multicast – 48 ports



Average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Cisco Nexus 5020 and the Arista 7148SX. Once the bursts exceed 30 packets per burst, the Arista 7148SX starts to drop packets and continues to drop as burst rates are increased. The Cisco Nexus 5020 starts to drop packets when bursts exceed 1500 packets per burst.

Enlargement



An enlarged view of the x-axis, the packet loss and latency.

47-to-1 Unicast Traffic – 48 Ports

Objective

To evaluate the ability of the switch to handle unicast bursts in a 47-to-1 traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame bursts. The metrics to be recorded were packets sent, packets received, packet loss and latency, calculated at burst rates without packet loss.

Test Setup

On the Ixia traffic generator port 1 through port 47 were set to send unicast traffic to port 48, and port 48 set up to send traffic to port 1, simulating a 47-to-1 unicast traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
Packet loss occurred when the burst rate exceeded 1500 packets per burst. Up to 3.1% packet loss was observed when 1600 packets per burst were applied.
- **Arista 7148SX**
Packet drops were observed when burst rate exceeds 25 packets per burst. There is a 1.3% packet loss at 30 packets per burst and 91.1% at 1600 packets per burst.
- When the Arista 7148SX reached 30 packets per burst, 1.3% of packets were dropped. Cisco Nexus 5020 experienced a 44.4 μ sec latency delay at 30 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
30	3.5 μ secs
1600	189.4 μ secs

Figures 9 and 9a: 47-to-1 unicast – 48 ports



The average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7148SX and Cisco Nexus 5020. Once the burst rates exceed 25 packets per burst, the Arista 7148SX starts to drop packets and continues to drop as burst rates are increased. The Cisco Nexus 5020 starts to drop packets when burst rate exceeds 1500 packets per burst.

Enlargement



An enlarged view of the x-axis, the packet loss and latency.

Full-Mesh Unicast Traffic

Objective

To evaluate the ability of the switch to handle unicast bursts in a full-mesh traffic profile and record packet loss and latency of traffic.

Description

This test was conducted on the Nexus 5010 and the 7124S with 24 ports and then repeated with 12 ports. Testing was conducted with a 128-byte frame size. The test was configured to step through multiple burst rates representing real world frame bursts in a binary search fashion. The metrics to be recorded were packets sent, packets received, packet loss and latency. Latency was calculated only at burst rates with no packet loss. Since dropped packets equate to infinite latency, latency is not shown once a platform starts dropping packets.

Test Setup

A fully meshed traffic configuration was used with each port on the Ixia traffic generator set up to send traffic to every other port. However this creates a single port that cannot send traffic during each burst run, since the port cannot send traffic to itself. For this reason a floating port was allocated. Port 24 was configured as the floating port with port 1 sending traffic to port 24 and all other ports sending traffic to port 1 during the first burst run. Similarly for the second burst run, port 2 sends traffic to port 24 and all other ports send traffic to port 2 and so on for other burst runs. With a burst size of 10 packets per burst, a total of 240 packets are sent from each port except port 24 which only sends 230 packets since there is no floating port defined for it. See calculations:

Burst Size = 10 packets per burst

10 Packets Burst x 24 Ports(Floating Port + 23 Other Ports) = 240 Total Packets sent per port.

Test Tools


Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


12 Ports with 128-byte Frame Size

- **Cisco Nexus 5010**
Packet loss was observed when the burst rate exceeded 250 packets per burst, with 0.8% packet loss when 275 packets per burst was realized.
- **Arista 7124S**
Packet drops were observed when burst rate exceeds 40 packets per burst. There is a 6.8% packet loss at 60 packets per burst and 50.7% with 275 packets per burst.
- It was noted that the Arista 7124S at 60 packets per burst dropped 6.8% of its packets, while Cisco Nexus 5010 experienced a 30.7 μ sec latency delay and continued to forward packets. Latency is acceptable rather than packet loss.

Figures 10 and 10a: Full-mesh unicast – 12 ports



The average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7124S and Cisco Nexus 5010. Once the burst rates exceed 40 packets per burst, the Arista 7124S starts to drop packets and continues to drop as burst rates are increased. The Cisco Nexus 5010 starts to drop packets when rate exceeds 250 packets per burst.



An enlarged view of the x-axis, the packet loss and latency.


24 Ports with 128-byte Frame Size

- **Cisco Nexus 5010**
We observed packet loss when burst rate exceeded 100 packets per burst with 1.9% packet loss at 120 packets per burst.
- **Arista 7124S**
Packet drops were observed when burst rate exceeds 10 packets per burst. There is a 9.6% packet loss at 15 packets per burst and 64.6% with 120 packets per burst.
- It was noted that the Arista 7124S at 15 packets per burst dropped 9.6% of its packets, while Cisco Nexus 5010 experienced an 18.5 μ sec delay and continued to forward packets. Latency is acceptable rather than packet loss.

Lengths of Bursts (Microseconds)


Frames / Burst	128-bytes
15	1.78 μ secs
120	14.21 μ secs

Figures 11 and 11a: Full-mesh unicast – 24 ports



The average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7124S and Cisco Nexus 5010. Once the rate exceeds 10 packets per burst, the Arista 7124S starts to drop packets and continues to drop as burst rates are increased. The Cisco Nexus 5010 starts to drop packets when rate exceeds 100 packets per burst.

Enlargement



An enlarged view of the x-axis, the packet loss and latency.

48 Ports with 128-byte Frame Size

Description

This test was conducted on the Cisco Nexus 5020 and the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-byte frame size. The test was configured to step through multiple burst rates representing real world datacenter frame. The metrics to be recorded were packets sent, packets received, packet loss and latency, calculated at burst rates without packet loss.

Test Setup

A fully meshed traffic configuration was used with each port on the Ixia traffic generator set to send traffic to every other port. This creates a single port that cannot send traffic during each burst run, since the port cannot send traffic to itself. For this reason, a floating port was allocated. Port 48 was configured as the floating port with port 1 sending traffic to port 48 and all other ports sending traffic to port 1 during the first burst run. Similarly for the second burst run, port 2 sends traffic to port 48 and all other ports send traffic to port 2 and so on for other burst runs. With a burst size of 10 packets per burst, a total of 480 packets are sent from each port except port 48 which only sends 470 packets since there is no floating port defined for it. See calculations:

Burst Size = 10 Packets per Burst

10 Packets per Burst x 48 Ports (Floating Port + 47 Other Ports) = 480 Total Packets sent per port

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations


- **Cisco Nexus 5020**
Packet loss occurred when the burst rate exceeds 50 packets per burst. Up to 1.7% packet loss was observed when 55 packets per burst were applied.
- **Arista 7148SX**
Packet drops were observed when burst rate exceeds 10 packets per burst. There is a 9.5% packet loss at 15 packets per burst and 54.4% at 55 packets per burst.
- When the Arista 7148SX reached 15 packets per burst, 9.5% of packets were dropped. Cisco Nexus 5020 experienced a 43.8 μ sec delay at 15 packets per burst and continued to forward packets. Latency is acceptable rather than packet loss.

The table shows burst rate when the switches start to drop traffic. We used 128-byte sized frames for testing.


Lengths of Bursts (Microseconds)

Frames / Burst	128-bytes
15	1.78 μ secs
55	6.51 μ secs

Figures 12 and 12a: Full-mesh unicast – 48 ports



The average latency and packet loss percentage at different burst rates for 128-byte frames are compared for the Arista 7148SX and Cisco Nexus 5020. Once the rate exceeds 10 packets per burst, the Arista 7148SX starts to drop packets and continues to drop as burst rates are increased. The Cisco Nexus 5010 starts to drop packets when rate exceeds 50 packets per burst.



An enlarged view of the x-axis, the packet loss and latency.

Arista 7148SX Flooding Issue

While testing with unicast bursts in a 47-to-1 traffic profile, we observed an anomaly with the Arista 7148SX. The device started to flood unexpectedly at 30 frames per burst after 5 - 6 minutes of running the test. Deeper analysis and multiple iterations of this test revealed that the issue only arises for burst rates at which the Arista 7148SX drops frames and after 5 - 6 minutes of running the test.

In the 47-to-1 unicast test, ports 1 to 47 send traffic to port 48 and port 48 sends only to port 1. At a burst rate of 30, the Arista 7148SX starts to drop frames. After 5 - 6 minutes of running 47-to-1 unicast burst traffic, we observed flooded packets destined for port 48 were arriving on all other ports. We verified that the MAC address at port 48 was learned.

Objective

To conduct the unicast burst in a 47-to-one traffic profile and recreate the scenario that triggers the flooding anomaly.

Description

This test was conducted on the Arista 7148SX with 48 ports connected to the Ixia traffic generator. Testing was conducted with 128-byte frame sizes. The test was configured to step through multiple burst rates that testing had revealed would lead to the Arista 7148SX dropping packets.

Test Setup

On the Ixia traffic generator port 1 through port 47 were set to send unicast traffic to port 48 and port 48 set up to send traffic to port 1, simulating a 47-to-1 unicast traffic profile.

Test Tools

Ixia XM12 chassis running IxOS 5.60.550.3 EA-Patch1 with Ixia test application IxExplorer.

Observations

- The Arista 7148SX starts to flood traffic at 30 frames per burst after running the test for 5 - 6 minutes. Flooded packets destined for port 48 arrive on all ports. This should not have happened since a *show mac-address-table* command revealed port 48 still has MAC learned.
- This flooding continued and only stopped when port 48 stopped transmitting traffic for 300 seconds (MAC aging timer) or the MAC address table was cleared.
- We conducted multiple iterations of this test and observed the flooding on the Arista 7148SX only in this scenario.

Discussion

Additional tests conducted after conferring with Arista technical support regarding this test included increasing the learning address table time out, however this did not alleviate the problem. But we noted that reducing the traffic load substantially does stop the flooding.