

Arista Networks 7148SX Switch

Unicast and Multicast Performance and Latency Tests

Introduction

The European Advanced Networking Test Center (EANTC) was commissioned by Arista Networks to test the performance of Arista's 7148SX 48-port, 10GbE switch. The goal of the tests was to verify the performance characteristics advertised by Arista, namely the non-blocking throughput capacity and low latency of the switch. The latency was expected to remain low throughout line-rate unicast and multicast traffic tests.

There are several market segments that are extremely sensitive to throughput and latency. Examples of these markets are data centers that must connect a multitude of servers to each other. Applications in the financial sector for example require incoming information to be received as soon as possible - every nanosecond counts.

Test Highlights

- **Extraordinarily low-latency switching**
- **Line-rate throughput for unicast and multicast**
- **RFC 2889 throughput with zero loss**

Tested Devices & Test Equipment

Arista 7148SX is a 48-port Layer 2/3 switch. The 7148SX offers density of 48 ports in a single rack unit. The switch uses Fulcrum FM4000 ASICs to provide non-blocking performance and low latency for both layer 2 and 3. Despite the small form factor, the device offers datacenter-class resiliency with fully redundant & individually hot-swappable fans and power supplies.



Figure 1: Arista 7148SX Switch

To perform the tests we used an Ixia XM12 chassis running IxNetwork 5.40.200.6. The tester was configured with 48 10-GbE ports. For all tests we used the following frame sizes based on RFC 2544: 64, 128, 256, 512, 1024, 1280, 1518, 9216 bytes.

IP Unicast Throughput and Latency

In order to verify the switching performance for unicast IP traffic, we employed the methodology specified in Internet Engineering Task Force (IETF) Request For Comments (RFC) 2544 - "Benchmarking Methodology for Network Interconnect Devices". The RFC specifies a methodology where traffic is transmitted to a series of port pairs. We followed this methodology; however, when deciding which ports would be paired, we took into account information provided

Arista Networks 7148SX®

✓ **Ethernet Performance**
At line rate for Unicast and Multicast

✓ **Low Latency**
In nanosecond range

Test Period: October 2009
7148SX with EOS 4.3.0
© 2009 EANTC AG



regarding the backplane architecture of the device, thus increasing the stress induced on the device under test (DUT).

After creating the mesh of 48-to-48 10 Gigabit Ethernet ports, we ran a line-rate traffic test for each of the eight frame sizes. Despite putting maximum load onto each interface, all traffic was successfully passed without losing any frames. We recorded the maximum latency between any two ports to be 680 nanoseconds. The results of the first test allowed us to validate that the switch is capable of forwarding traffic at line rate for all tested frame sizes.

IP Multicast Performance

Accumulated Multicast Throughput and Latency

IP video applications often use an accumulated IP multicast distribution model. A centralized server or video head-end hosts all broadcast channels that are then joined by all downstream nodes that deliver them to the end user. In order to verify the multicast replication and switching capabilities of the device, we performed tests according to RFC 3918, which specifies a benchmarking methodology for multicast performance.

In this scenario, all ports joined all groups. One port sourced 256 multicast groups while the other 47 ports connected to emulated users which sent IGMP joins to all groups. The test was conducted for all frame sizes as listed in the introduction, each at full line rate (10 Gbit/s per port).

Test Highlights - Line Rate Multicast Traffic	
Minimum Latency	600 nanoseconds
Average Latency	1273 nanoseconds
Maximum Latency	1520 nanoseconds

While many switches can handle this load without loss, many would incur significant latency given the high number of multicast replication required in this test setup. However, in our test with the Arista switch, the absolute maximum latency observed on a single port was 1,520 *nanoseconds*, and the minimum 600 nanoseconds. The device successfully switched all frames without any loss, keeping a low latency on all frame sizes. The results meet the requirements of applications that need extremely low-latency multicast performance.

Distributed Multicast Throughput and Latency

The next test followed a realistic scenario often seen in financial applications. Traders are typically only interested in a fraction of information streamed from the central servers. The server nevertheless must provide information to all its users. Therefore, a high number of multicast groups are forwarded to the switch, which then distributes the multicast groups to the subscribers requesting them.

To test this scenario, we conducted an RFC 3918 test with some alterations to the scenario. A single 10 Gigabit Ethernet port on the device under test (DUT) was connected to an Ixia tester port sourcing a total of 846 multicast groups which added up to 10 Gbit/s. Each of the other 47 10 Gigabit Ethernet ports in the test was emulating multicast subscribers each of which joining 18 of the 846 groups.

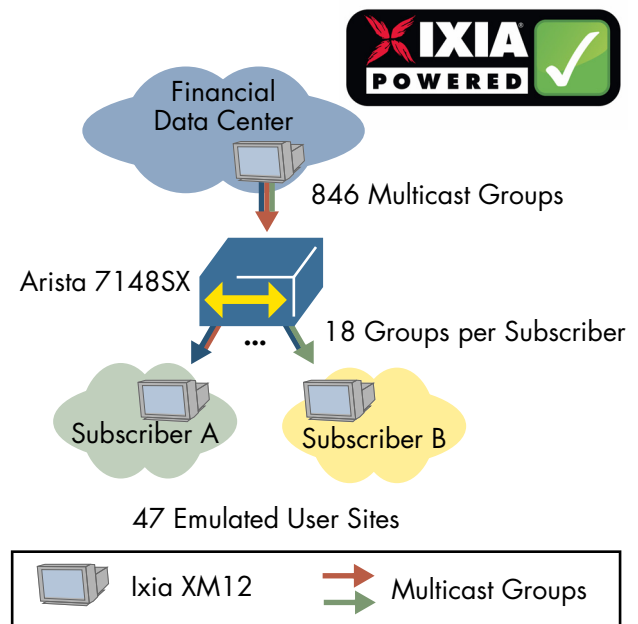


Figure 2: Distributed Multicast Throughput and Latency Test Setup

The test ran successfully for all frame sizes. All traffic was received as expected and no frames were dropped. This proves that the Arista switch is a viable option for financial applications as described above.

Mixed Class Throughput

RFC 3918 testing methodology was also used for the mixed class performance scenario. The test investigates the throughput performance when a mix of multicast and unicast packets is being sent through the device under test. This mix is more realistic for converged networks serving multiple customers and applications, some of which are unicast based, and other multicast.

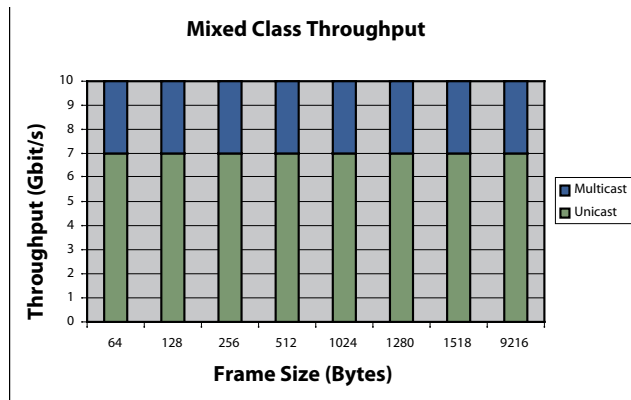


Figure 3: Mixed Class Throughput Results

In order to test the full capacity of the switch, the test setup was configured using a single port sourcing multicast traffic which was then replicated to the other 47 ports. In addition, all 47 ports also carried a full-mesh of unicast traffic. A multicast to unicast ratio of 3:7 was used. The test ran flawlessly, observing 0 frame loss. The test closely reflects real world scenarios where both multicast and unicast traffic is used as is true of today's converged packet-based networks. The throughput results are therefore promising, showing that the switch will perform well in a real-world scenario.

Ethernet Performance

Unicast and Broadcast Throughput

Since the 7148SX functions as a Layer 2 switch as well as Layer 3, we conducted two tests of Ethernet performance. The methodology used was according to RFC 2889 which is specifically written for LAN devices' benchmarking. The first test was using unicast traffic, and the second broadcast.

In the unicast test, a full-mesh traffic configuration was used. Each of the 48 10-Gigabit Ethernet ports sent Ethernet frames to each of the other 47 ports adding to a total of bi-directional line rate traffic for each port. In order to create a more realistic scenario, the traffic

reflected the eight unique hosts sitting behind each port. The test was performed for all eight frame sizes.

The second test used purely broadcast traffic. The Ixia tester transmitted line rate broadcast traffic into a single port of the 7148SX. This test was also performed for the eight frame sizes.

Both tests passed according to the expectations. We recorded no frame loss. In fact, no stream of the full mesh unicast traffic dropped a single frame. In the broadcast case, all 47 ports egressed the line rate traffic incoming from the ingress port as expected.

Summary

While network technologies progress to solve a wider breadth of problems, many providers simply need a straightforward solution that performs well. The testing performed on the Arista 7148SX put the switch to exactly that test. The results reflect that the switch meets these needs - a simple, high performance switch that will deliver what many administrators are looking for without promising more than it can deliver. With latency values lower than typical products on the market, the results further prove wrong those who believed that architectures for high performance Ethernet devices had hit their peak.

About EANTC



The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier

Ethernet, Triple Play, and IP applications.

EANTC AG
Einsteinufer 17, 10587 Berlin, Germany
info@eantc.com, <http://www.eantc.com/>

V1.0 20091113