Arista Networks 7508 Switch

384-Port 10 Gigabit Ethernet Performance Test Report

Introduction

In late 2009 EANTC tested the performance and latency of Arista's 7148SX switch. As Arista prepared to launch a new high-density switch EANTC was requested to return and verify the performance of the new Arista 7508 switch. EANTC's task was to validate that the 384 10 Gigabit Ethernet ports device, neatly packed into 11 rack units, is indeed a high performance modular switch.

Arista positions the DCS-7508 switch for applications in the data center, for High Performance Computing (HPC), Web farms, Storage and financial. The testing challenge was therefore to show that the switch can really deliver the performance Arista claimed it has under conditions that will meet all of the target users.

Arista Networks DCS-7508 Switch

- ☑ 384-Port 10GbE
  Wire-speed performance
- ☑ Low Latency
  Consistent across line cards
- ☑ High Availability
  N+1 fabric redundancy

Test Period: March 2010
Arista DCS-7508 with EOS-4.4.0

Tested by EANTC 2010

EANTC Validation Highlights

- ➜ **7.68 Tbit/s full-duplex throughput**
- ➜ **Zero frame loss throughput performance**
- ➜ **Average latency of 4.5 µsec**
- ➜ **Best measured maximum latency at 8.6 µsec[a]**
- ➜ **Wirespeed performance with reduced fabrics**

a. For frame size up to 1518 bytes

Tested Devices & Test Equipment

The test was set up in March of 2010 at Ixia's iSimCity lab in Santa Clara, California. All 384 10 Gigabit Ethernet ports on the Arista DCS-7508 fabric modules were connected to ports on Ixia's IxYukon load modules housed in four XM12 chassis. This large installation was not without its own challenges. As expected from tests of this scale, cabling and preparation itself took time, however, the team's dedication prevailed and the tests were prepared extremely well.

Figure 1: The Arista Networks DCS-7508 Switch

Both the device under test (DUT) and the tester used pre-release software – Arista provided an early version of EOS-4.4.0 while Ixia used IxNetwork version 5.60.

Ethernet Throughput Test Results

We performed two throughput tests on the DCS-7508 switch. The tests differed from each other in one fundamental aspect: the distribution of traffic across the device.

In both tests we used all 384 ports on the DUT to send bidirectional traffic for a number of frame sizes: 64, 65, 73, 128, 256, 512, 1024, 1280, 1518 and 9216 bytes. Every port in the test had 16 MAC addresses configured on the tester in order to emulate a realistic scenario. Both tests shared the same goal: reach the theoretical maximum rate at which none of the offered frames are dropped by the switch.

In modular switches traffic is normally forwarded through the fabric. In some cases, performance is better when switched locally on the same line card, as compared to switching across line cards. Hence, to identify any corner cases, we used two test case configurations. In the first test case we used a full mesh traffic configuration: every port was sending to, and receiving from, all of the other 383 ports. The second test case setup had a different traffic configuration. Every port was sending and receiving traffic from its neighboring port only.

Both tests results were in effect identical: the switch forwarded frames, at line rate, for all frame sizes, without dropping a single frame.

In this test, when using 64 bytes frames, the switch processed more than 5.71 billion frames per second. This is no easy task and we were impressed to see that the switch did not loose any frames at wire-speed. With these results we can then confirm that Arista's claim, for the two traffic configurations used in the test, were accurate. The DCS-7508 can indeed utilize the full capacity of its ports and therefore switch 7.68 Tbit/s worth of Ethernet traffic with the 48 port 7548S line card modules.

Latency Test Results

The second test area focused on latency through the DCS-7508. These days, latency has become a key metric in several industries. Low network latency is critical in some applications such as financial trading environments where a millisecond advantage plays a big financial role. Even in other applications, where the disk or the operating system contributes to latency, reducing network latency will result in increased overall performance.

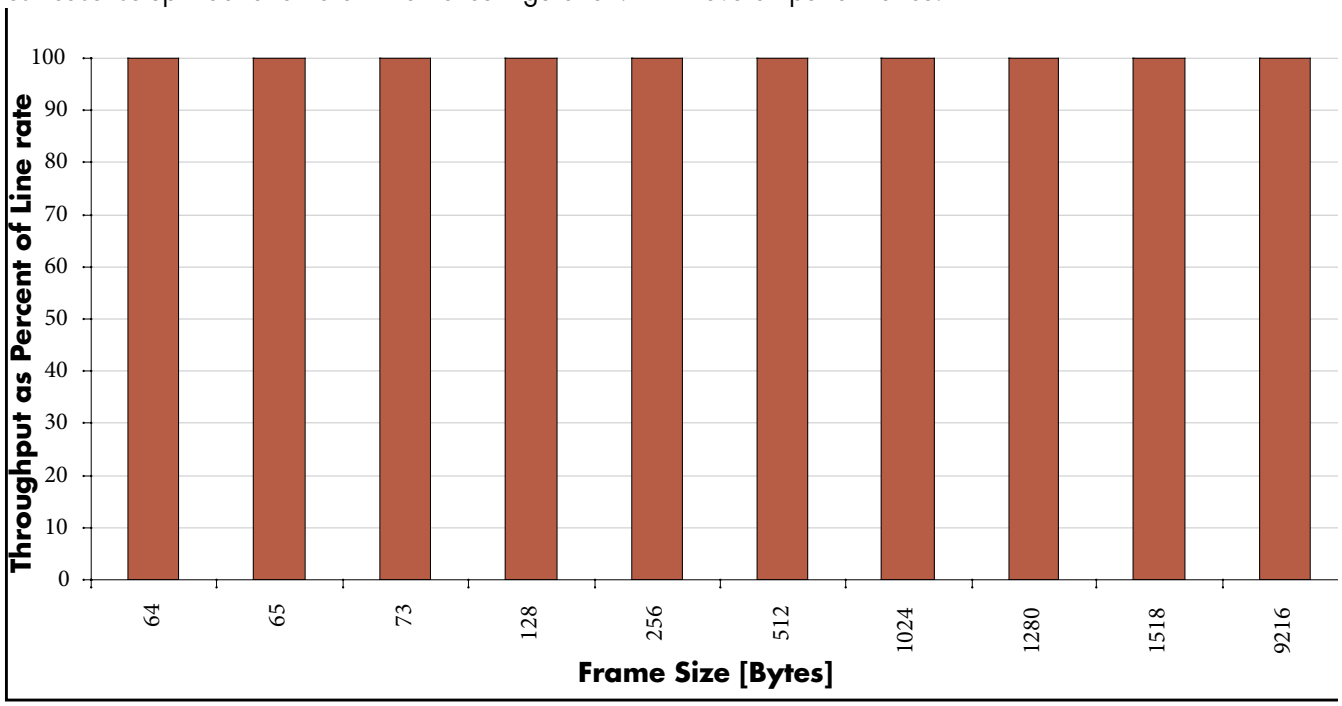

Figure 2: Throughput Test Results For Full-Mesh Configuration

Several traffic configurations could have been used for the latency tests, however, we focused on one scenario - port to port latency. According to Arista, their customers worry about latency values in such configuration more than any other. The configuration we used for the test was to measure port-to-port latency in a fully loaded system using all 384 ports on the DUT.

We repeated the test with the same frame size sweep as was used for the throughput tests, using both standard frame sizes (such as 64 bytes) and adding non-word aligned frames such as 65 and 73 bytes. We ran the test for a duration of 3 minutes for each frame size this time using 70% of line rate on all 384 ports.

The results of the test confirmed Arista's claims that the new device, under typical usage conditions, is a low-latency switch. The *maximum latency* for frame sizes up to 1518 bytes was 8.6 microseconds, while latency increased, as expected, for the 9216 byte frames up to 14 microseconds. The highest *average latency* we recorded for frame sizes up to and including 1518 bytes did not exceed 4.5 microseconds. The results are presented in figure 3 below.

One of the side effects of the test configuration is that traffic remains on a single slot (each slot hosting 48 ports). In our analysis of the results we looked at the latency breakdown as a function of the line card position. The biggest maximum latency difference between one slot and another was 1.2 microseconds for 1518 byte frames.

The test showed that in a realistic scenario, in which the ports are well used, the latency recorded by the switch remains in the single digit microsecond range.

We performed an additional test to validate that indeed no significant difference exists between traffic on a single slot and traffic traversing the whole switch. Four ports, split into two groups, were used in this test. The first group included the first and the last port on the switch. The second group included two ports that were on the same slot. Again we ran through all 10 frame sizes in the test and recorded the latency for the port pairs.

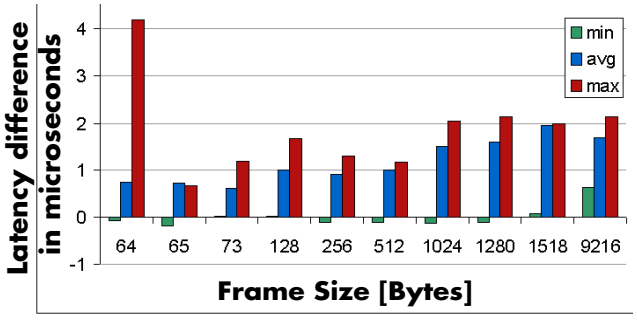

Figure 3: Latency Differences Results Between Port Controllers

As seen in figure 3 some latency differences did exist between the two port groups, however, for most frame sizes the average difference was minimal. The reason for the latency differences, as explained by Arista, was that the test traffic sent between the first and last port was switched on different port controllers while the other stream was switched locally.

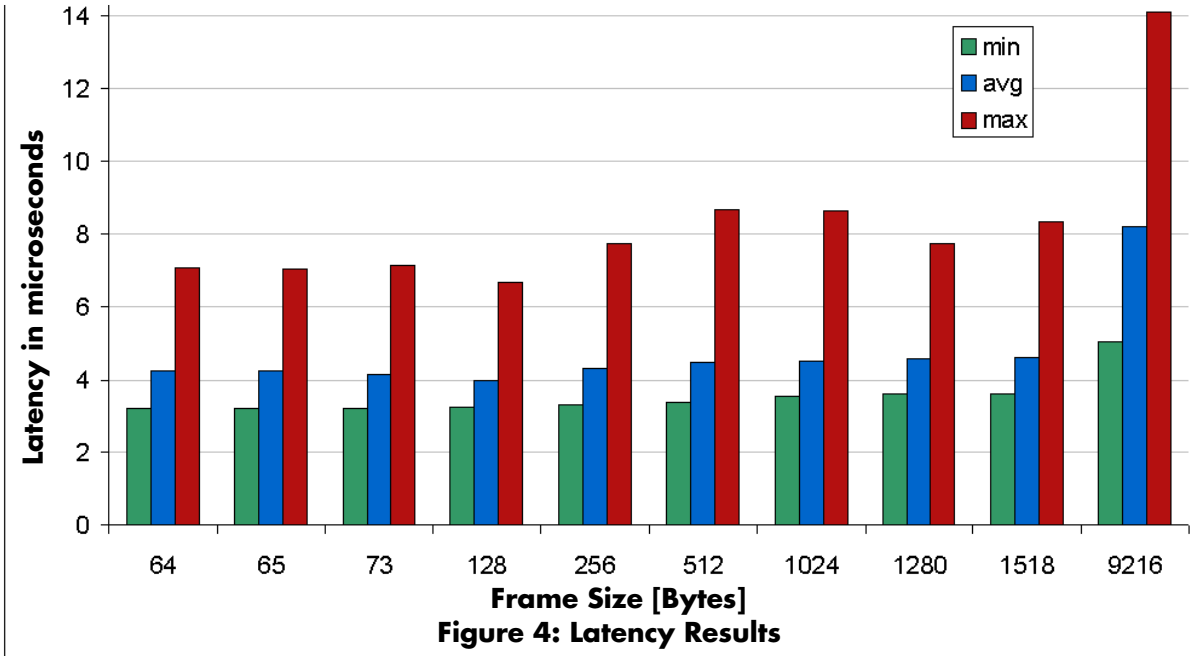

Figure 4: Latency Results

High Availability Test Results

Arista asked us to validate that the device can continue to perform as expected when only five of the six fabrics are installed in the device. In order to validate this claim we simply removed one of the fabrics (Fabric Module 1) and re-ran the throughput test using 1518 byte frames. Since we already had the throughput results from the previous tests, we could compare the two test results.

Both throughput results for 1518 byte frames were identical - the DCS-7508 could indeed forward at 100% line rate with only 5 fabrics. Thus we were able to validate that the fabric modules are indeed N+1 redundant and the switch had the same throughout of 7.68 Tbit/s with just 5 fabric modules.

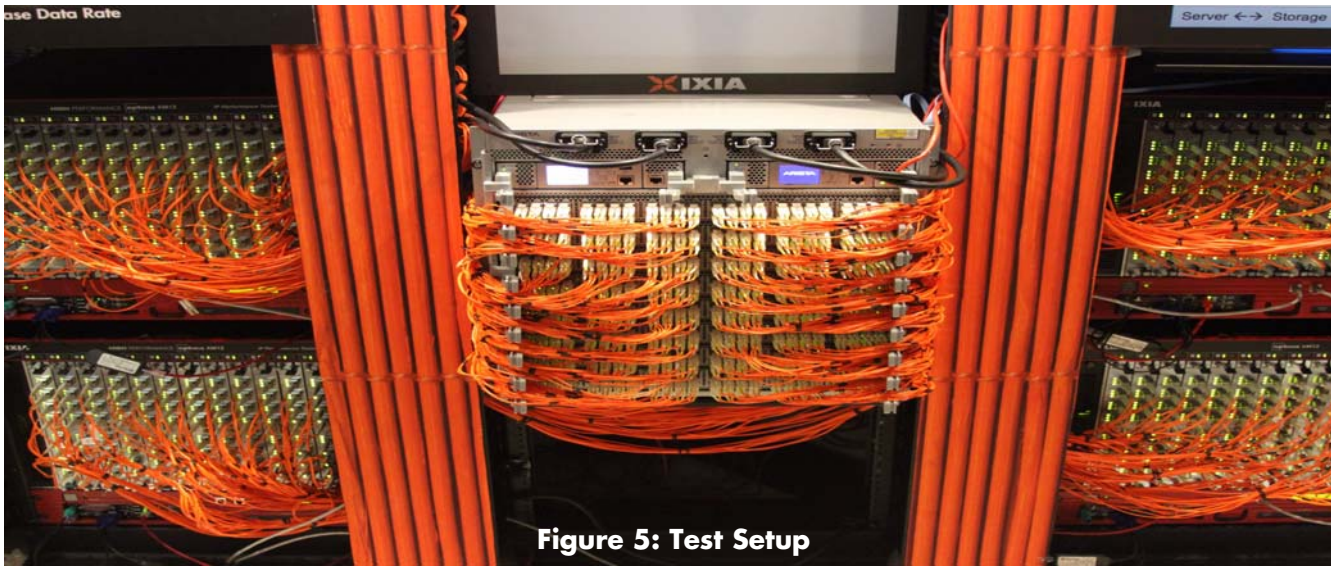

Figure 5: Test Setup

Conclusion

Once again we were impressed with the performance and latency of the Arista DCS-7508 switch and its ability to perform at wirespeed with reduced number of fabrics. Packing such high density in such a small foot-print must have been an engineering challenge that has fulfilled its goals - we were able to find no compromise in the throughput and latency results - the system performed as advertised by Arista.

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/

V1.1 20100417