European Advanced Networking Test Center

Competitive Test
Cisco Catalyst 6500 — Force10 TeraScale E1200
Executive Summary

Introduction

In June 2005, Cisco Systems commissioned the European Advanced Networking Test Center (EANTC) to independently compare Force10's TeraScale E-Series with Cisco's Catalyst 6500 switches. EANTC and Cisco systems developed a detailed test plan of almost 60 test cases and invited Force10 to participate in the testing. Force10 declined the invitation, but offered to help with potential questions. During the months of July and August 2005, testing were performed by EANTC engineers with support from Cisco and on a few occasions test configuration provided by Force10. The high level findings are presented here.

Force10's TeraScale E-Series switches support high port density, offering up to 1,260 Gigabit Ethernet interfaces in a single switch. Force10 encourages network designs where large numbers of ports are "collapsed" into very high-density single-switch solutions. Cisco typically suggests distributed solutions spread across multiple switches in a redundant configuration.

Force10 contracted the test lab Tolly Group to conduct a report highlighting the performance, scalability and high availability features of the TeraScale solution. The reports seemed to support Force10's marketing claims, however, lacked the details to facilitate the reproduction of the results by another testing body.

EANTC and Cisco conducted comprehensive tests to reproduce the Tolly Group test results and compare Force10's TeraScale E-Series against Cisco's Catalyst 6500. The tests looked beyond pure speeds and feeds, which were the focus of Force10's Tolly Group tests; examining other critical areas of switch performance. EANTC and Cisco focused on providing users with a complete stable solution with the understanding that slower network speeds do not prevent users from working while an unstable network will stop users from performing their tasks.

More details can be found in the full documentation of all 59 test cases, covering more than 100 pages.

Overall Test Highlights

In-depth independent comparison of Cisco's Catalyst 6500 and Force10's TeraScale switch families

- More than five weeks of testing
- 100+ pages of public test report

EANTC independently verified Force10's performance claims and Tolly Group findings about Force10's TeraScale architecture and compared them to Cisco's Catalyst performance:

- Force10's TeraScale switch does not support the advertised zero-loss throughput with 672 ports for any packet sizes except 64 and 1518 bytes.
- Force10's TeraScale high availability claims are misleading and at times incorrect.
- EANTC found major heat flow issues leading to a total switch failure when removing one line card for 10 minutes.
- The Cisco Catalyst BGP routing implementation converged two times faster than Force10's. In contrast to the Tolly report, EANTC found a Force10 TeraScale hardware scalability limit at 200,000 unique BGP routes.
- Force10's TeraScale Bandwidth Aggregation implementation tended to polarize traffic, usually using only 50 % of the aggregated capacity, sometimes even only 25 %.
- Cisco's Catalyst multicast implementation outperformed the Force10 TeraScale switch in every area tested.

The Test Areas

Our tests were divided into the following broad test categories:

Investigate the Force10 commissioned Tolly Group test reports:

We compared the performance, availability and scalability claims made by Force10 to our own findings. Force10's Tolly Group tests can be found at *http://www.tolly.com/Search.aspx?VendorID=27* and are contained in Test Reports 204147 and 204148.

Verify Force10's High Availability Claims:

Force10 has published a number of white papers on their web site to explain their high availability and resiliency stance. The publications are titled "High Availability in the Force10 Networks E-Series" (Version 1.3) and "Guaranteed Access to System Management even During Processor Overload." It stands to reason that a single box solution will be designed with the utmost resiliency and redundancy.

EANTC and Cisco used these documents as guidance for High Availability testing. In addition, the Tolly Group "Hitless Failover" tests were reproduced and later repeated with multiple TeraScale E-Series switches in a realistic network topology.

Routing Protocol Scalability, Stability & Resilience Tests:

In this section of tests, we focused on the scalability, stability and resilience of the control plane, the brains of the switch.

A high port-density switch must match its forwarding plane (the ports) with a strong and capable control plane. We investigated what we considered to be logical expectations from a switch with 1,260 ports: Routing neighbors equal to the number of ports; routing table learning time; convergence time; reaction to changes in routing information and more.

IP Multicast Scalability, Stability & Resilience Tests:

IP multicast traffic is becoming an increasingly important requirement in many networks. Applications range from remote learning, distribution of financial data in trading rooms, IPTV and the control streams used in grid computing applications. All have recognized the benefits of multicast distribution.

Our tests investigated:

- Multicast performance in various failure conditions. We investigated whether various resilience features adversely affect IP Multicast.
- Multicast routing and group scalability. We compared multicast features offered by Force10 to those on the Catalyst 6500.

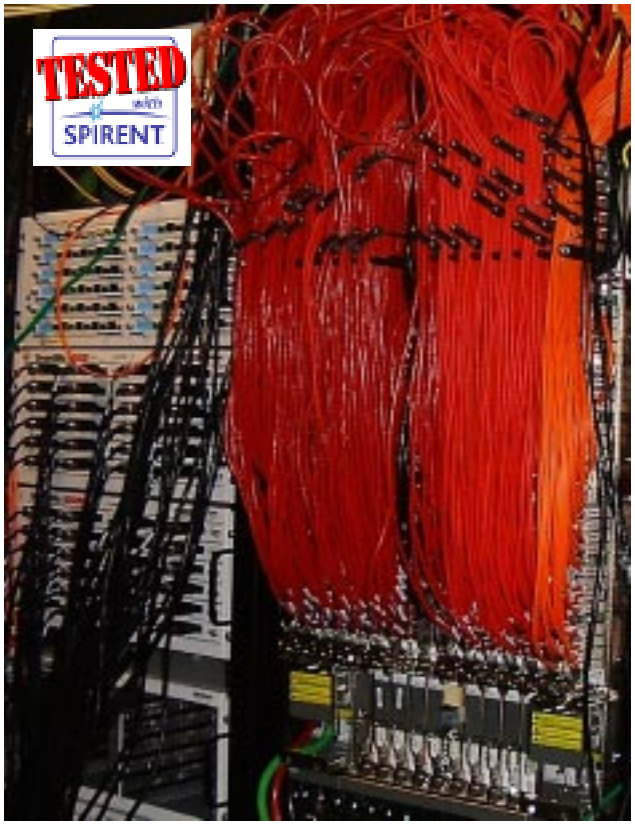

Test setup for Force10 TeraScale E1200 data plane scalability »snake« test with 672 Gigabit Ethernet ports

Critical Findings

The following is a high-level summary of the major tests findings:

Force10 Tolly Group Tests:

- Various Tolly Group report findings were not reproducible by EANTC. Through correspondence with the Tolly Group and Force10 Networks (available on EANTC's web site) we concluded that several results were reached using software and potentially hardware that are not available for the general public.
- Force10's "1-Million ACL" test was confirmed by the Tolly Group to be an access control list (ACL) with 20,993 entries, not the 1-Million entries implied by the original test report. Tolly confirmed that the 1-Million ACL claim was a result of applying the same 20,993-entry ACL to 56 10-Gigabit Ethernet ports in parallel, hence allowing Force10 to claim "greater than 1-Million ACLs." (56 x 20,993 = 1,175,608)

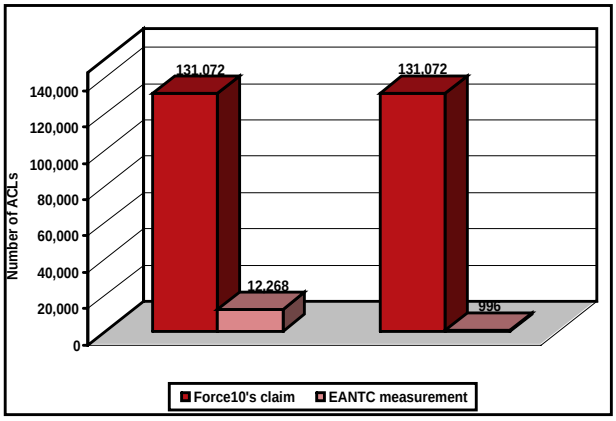

Force10 E1200 Ingress Access Control List Capacity

EANTC also verified that when running the latest FTOS production software the switch cannot apply an access control list with 20,993 entries to an interface, it runs out of hardware table space. It is apparent from correspondence with Force10 that "specially modified" FTOS software was used to re-allocate internal hardware table space to accommodate the ACL for this test, something not available to the average customer. EANTC confirmed it is impossible to reach 1 Million ACL entries in a single list and therefore Force10's claim is highly misleading.

- Force10's "6-Million BGP Paths" scalability test was conducted using 2 ports with 260,000 BGP routes learned over 25 equal cost paths (260,000 x 25 = 6,500,000). Tolly's report did not publish details of this test's configuration.

EANTC could not reproduce the Tolly Group test and asked Tolly Group to provide the Force10 switch configuration file. The configuration file showed that 25 equal cost paths were used to achieve these results, while the production code available to the public supports only 16 equal cost paths.

- EANTC was unable to reproduce Force10's IPv6 forwarding tests:
  - Force10's current production software does not support any configuration commands involving IPv6.
  - Force10 support a total of 8 (eight) hardware IPv6 routing entries. This finding explains the Tolly report using only 8 ports for IPv6 forwarding testing as opposed to the higher port count in the IPv4 tests.
  - Force10 demonstrated IPv6 forwarding between directly connected IPv6 hosts. No IPv6 routing protocols are supported.
- The Tolly Group tests focused exclusively on unicast traffic performance. In a real network, there is always a mix of traffic types, with a small amount of background broadcasts. We ran a similar L2 Snake test to confirm the TeraScale E-Series unicast performance, and then ran the same test, but this time adding 10% background broadcast traffic in a separated, isolated VLAN. The capacity of the switch was not exceeded and we expected no loss on either traffic type. Our test found that when the broadcast traffic was present on the TeraScale's switch fabric, the packet loss experienced by the unicast traffic increased from virtually zero to 20% loss. This suggests that the TeraScale's switch fabric has difficulty coping with the mixed traffic types found in most networks.

High Availability Tests:

- Force10's white paper "Guaranteed Access to System Management even During Processor Overload" states that during CPU overload conditions their 3-CPU control plane design would allow an operator to "always get precise information on switch status and overcome the problem without rebooting".
  - Our tests found that the Control Processor (CP) relies on the other two CPUs presented in the

system to retrieve status. Hence, when one of the other CPUs is under heavy load the Control Processor is effectively blocked and is unable to provide the operator with any information, debugging or otherwise.

- While the Control Processor is waiting for one of the other processors to respond with information, the command line interface (CLI) is "locked up" until such time that the commands times out. During one test the CLI was unusable for 20 minutes.

- Force10's E1200 switch overheats under common usage conditions. During one of the tests, a single card slot was left uncovered for 10 minutes. The switch fabric cards overheated and shut down the switch. The incident is documented in a video that can be found at: *http://tools.cisco.com/cmn/jsp/index.jsp?id=46123*

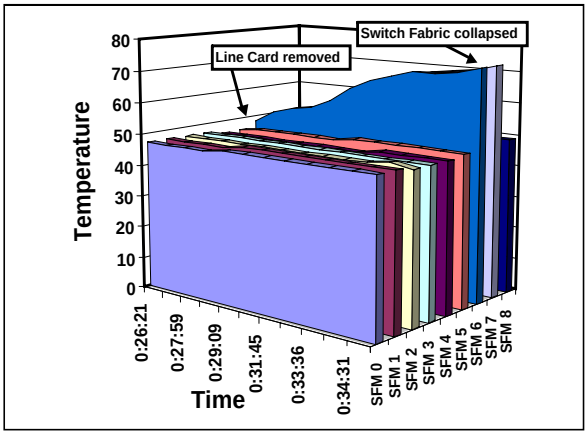

Force10 E1200 Switch Fabric Module (SFM) temperature after removing one line card

- Hitless Route Processor Module (RPM) failover does not work in a real-world network topology. We found prolonged service disruption when measuring RPM failover times in a realistic, multi-switch network design that employed OSPF, Multicast, VRRP and other common features.

- Hitless Switch Fabric Module (SFM) failover causes all nine switch-fabrics to fail when an unrelated software feature (sFlow) is enabled on interfaces. This was discovered under FToS Version 6.2.1.1. In Version 6.3.1.2 this behavior was not observed, however, EANTC found that in this version of FToS, sFlow had been severely rate-limited and no longer functioned correctly. Force10's sFlow workaround doesn't fix the underlying problem. See video evidence at: *http://www.eantc.de/F10-SFM*

- Force10's operating system (FToS) is not modular and does not support hitless restarts and upgrades as advertised by Force10. In comparison, Cisco's Modular IOS, which is available for the Catalyst 6500, offers individual sub-system upgrades and restarts. EANTC's report on Cisco's Modular IOS is available under: *http://www.eantc.com/press/pressreleases/aug05/EANTC-Summary-Report-Cisco-ION.final3.pdf*

Bandwidth Aggregation Tests:

- Force10's load-sharing in both OSPF Equal Cost Multi-Path Routing and in Link-Aggregation proved unable to use all available links when these load-sharing techniques were extended beyond a single switch-hop.

- Force10 distributes traffic over a subset of the available links leaving a number of links empty of traffic. We tested all available load-sharing algorithms, but could find no solution for an even distribution of traffic across all link members.

- During one bandwidth aggregation test we found that only 2 out of 8 configured links were being used.

- Cisco's Catalyst 6500 used all configured links, but did not achieve completely even traffic distribution.

Route Scalability & Network Stability:

- Simple OSPF routing instabilities caused the Force10 switch's CPU to overload. The switch was unable to respond to routing updates and continued forwarding traffic based on outdated forwarding information.

- In some tests, the E1200 did not populate its routing table until the routing process was manually reset. The OSPF neighbors showed "full" adjacencies, but no OSPF routes were installed in the routing table.

- During BGP route scalability tests we were unable to reproduce the 6-million BGP paths test results reported by the Tolly Group. We verified that the Tolly Group BGP scalability test did not in fact represent a test with 6-Million unique BGP routes but was limited by hardware table sizes of only 260,000 route entries. In addition, we found that even with traffic rates of only a few hundred packets per second the TeraScale switch suffered frame loss when traffic was sent to all BGP routes in the test.

- The TeraScale E-Series took longer to converge BGP routes than the Catalyst 6500, even when offered 50% less routes to learn.
- The Catalyst 6500 maintained its advertised performance levels and scaled to 500,000 unique BGP routes while Force10's verified BGP route limitation was 262,000.
- We verified that Force10's route scalability limitation is grounded in its hardware CAM size. At 262,000 BGP routes, CAM memory is exhausted and no further routes are learned.

IP Multicast Scalability & Resilience:

- The TeraScale E-Series imposed a limit on the number of Protocol Independent Multicast (PIM-SM) routing interfaces that could be configured. EANTC configured 31 out of the 1,260 Gigabit Ethernet interfaces for PIM-SM before receiving an error messaging indicating the limit had been reached. The TeraScale E-Series can only support a maximum of 31 PIM-SM neighbors, whereas the Cisco Catalyst 6500 was able to establish PIM-SM peering with over 550 neighbors.
- The TeraScale E-Series failed to load-share multicast traffic across multiple inter-switch links. This would severely limit end-to-end multicast bandwidth.

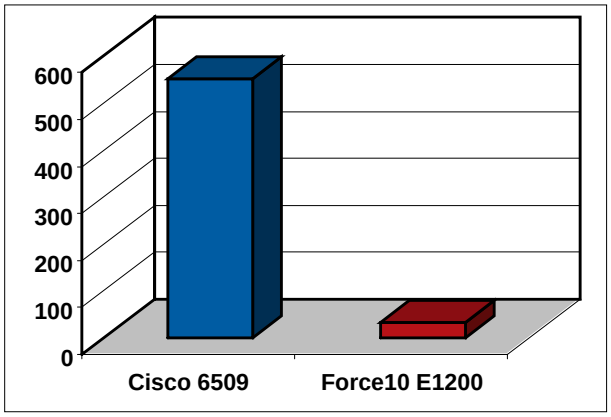

Maximum number of multicast (PIM-SM) neighbors

- Force10's TeraScale E-Series did not recover from a PIM Sparse Mode Rendezvous Point failure, even with just 170 multicast groups in the test. In the same test configuration the Catalyst recovered within 8 seconds.
- PIM-SM route convergence required over 10 minutes for 10,000 multicast routes to be learned on the Force10 solution while the Cisco Catalyst learned 5 times as many routes in under 3 minutes.

Summary

This executive report summarizes over five weeks of intense testing. We highlighted what we consider to be the most critical deficiencies in the Force10 TeraScale solution.

As is the case with every network device, the individual customer requirements should be matched to product capabilities. It might be the case that as a layer 2 Ethernet switch the TeraScale offers companies the port density they need.

In all other tests; Routing, Multicast, Performance and Resiliency, we found alarming deficiencies.

From EANTC's point of view, Force10 published product literature and Tolly Group test results highlight maximum performance in artificial conditions. We suggest readers should not rely on these figures for real-life network design.

EANTC AG
Einsteinufer 17, 10587 Berlin, Germany

Phone: +49.30.3180595-0
E-Mail: info@eantc.com
Web: http://www.eantc.com/

About EANTC

EANTC (European Advanced Networking Test Center) is an internationally recognized test lab based in Berlin, Germany. We offer vendor neutral network test facilities for manufacturers, service providers and enterprise customers. Our business areas include:

➜ Test and Certification Services for Network Equipment Manufacturers

EANTC provides external quality assurance by testing conformance, performance and robustness of single systems as well as the interoperability of multiple components. More than 14 years of testing experience, the use of state of the art analysis equipment and active participation in industry forums enable us to conduct tests even for the most advanced customer requirements.

EANTC tests cover the full range of Carrier Ethernet, IP, and MPLS networks, IP services including firewalls and VPN gateways, ATM, and voice over packet networks.

➜ Test Services for Service Providers and Enterprises

EANTC supports network operators during all stages of the infrastructure life cycle: From network design and RfP creation through prototype testing during the tender process and acceptance testing up to trouble-shooting and future capacity planning. Our tests ensure network performance, availability and resilience. They reduce the risk of network failure and cost of operation.

➜ Philosophy

In all our activities — whether we conduct tests, provide consultancy or training — only the highest quality serves our purpose. EANTC invests much in research, education and development of new test methods to meet the demands of our customers. Our active participation in industry forums is a must to stay up to date.

We often find ourselves mediating between network operators, vendors and systems integrators in a cross-cultural environment. EANTC test services result in neutral, unbiased facts. Our detailed test reports provide the basis for rational decisions and actions. We ensure confidentiality in any required way.

➜ History

In 1991, EANTC was founded at the Technical University of Berlin. EANTC soon started to test and certify network components for high speed network technologies, and continued the close cooperation with the university after its spin-off in 1999.