Alcatel-Lucent
1850 Transport Service Switch®
MPLS Transport Profile Ring Protection Test

Introduction

As more transport networks are evolving from traditional pure SONET/SDH toward packet-based technologies, vendors are proposing solutions that provide packet-based transport services with the same level of quality, reliability and operational simplicity to which carriers are accustomed. Resiliency is therefore an essential element for next-generation packet transport networks. This test report explores Alcatel-Lucent's flagship Transport Service Switch (TSS) ability to provide highly resilient and scalable services. The test was performed by EANTC engineers in Alcatel-Lucent's Vimercate, Italy laboratory based on a test plan created by EANTC.

Alcatel-Lucent 1850 TSS®

- ☑ MPLS-TP Ring Protection
  In-line with SDH networks
- ☑ Service Scalability
  High-availability at scale

Test Period: September 2012
1850 TSS-320, Version 6.0.0 (beta)

Tested by EANTC 2012

Test Highlights

- ➔ **Support for up to 32,000 protected MPLS-TP bidirectional circuits**
- ➔ **Protection switching in less than 50 ms for all circuits in all test cases**
- ➔ **Node failure protection switching in 15 ms for all circuits**

Tested Devices & Test Equipment

Alcatel-Lucent engineers setup a test bed constructed of four 1850 TSS-320 switches. The devices under test were interconnected in two 10 Gigabit Ethernet rings as is shown in Figure 1. The switches used MPLS-TP to transport services with MPLS-TP Ring Protection Switching (MRPS) for service high-availability. MPLS-TP CCM OAM (link-livelihood mechanism) was employed for section fault detection.

Alcatel-Lucent configured two types of services:

- 10,000 Ethernet pseudowires each mapped to a transport label switched path (LSP)
- 22,000 Multi-segment pseudowires without a tunnel

In total, for the duration of the whole test, 32,000 bidirectional services were sending test traffic. In our evaluation of the out of service time results we monitored all services and report on the highest out of service time. From the tester perspective, 64,000 test streams existed and were active in all tests.

The User to Network Interface (UNI) was created using double VLAN tag (IEEE 802.1ad). Two switches were defined as customer facing, each providing eight Gigabit Ethernet interfaces to the Ixia tester. We employed Ixia's IxNetwork test solution, which enabled us to configure 64,000 test streams and track the traffic flows we generated for each service. Ixia XM2 chassis was used with XMV16 line cards. For one of the test cases we also used Ixia's ImpairNet - an in-line impairment module.

Test Goal

Alcatel-Lucent contacted EANTC to verify the major new feature in its next R6.0 code release - section-based ring protection mechanism. The goal of the test was rather straight forward: demonstrate that regardless of the number of services or transport tunnels, when an Alcatel-Lucent 1850 TSS based network is setup in rings, the protection switching times, in any imaginable failure condition, will always remain under the industry standard benchmark of 50 milliseconds.

Section Protection Background

The protection mechanism used in this test is specifically designed for MPLS-TP based ring topologies. Ring topologies are widely deployed, as SDH rings, in metro/aggregation networks, so a technology aimed to address this market segment should provide a solution for protection switching in ring topology as similar as possible to what transport operators are familiar with (i.e., SDH MS-SPRing).

Currently a standard solution for MPLS-TP Ring Protection is under discussion in ITU-T and IETF. Nevertheless, it is impressive to already see and test an implementation that is based on draft-helvoort-mpls-tp-ring-protection-switching.

Two mechanisms are used by the MPLS-TP Ring Protection Switching (MRPS) to protect tunnels or services in the ring: MPLS-TP CCM OAM for fault detection and MRPS Automatic Protection Switching (APS) to coordinate the protection switching actions. When a node detects a failure, all traffic is switched to the protection path away from the failure. Traffic is then switched along the ring until it reaches the node adjacent to the failure, at which point traffic is switched back to the working path. Once the failure is cleared, the node that performed the protection switching reverts back to the working path.

The effect of this mechanism is that capacity in the ring must be reserved for the protection traffic in order to guarantee Service Level Agreements (SLAs) during failure conditions. Much like on SDH rings, the ring topology allows sharing the protection bandwidth between different paths and therefore only 50% of the ring capacity should be allocated to protection traffic regardless of the working path traffic configuration. Moreover, statistical multiplexing, which is inherent to MPLS QoS mechanism, allows using the whole ring bandwidth.

Test Execution

We created five different failure scenarios that are likely to be seen in a transmission network. They ranged from complete failure of a link or a node to logical failures and single ended fiber failures. We configured the tester to send the same traffic profile in all test cases for all 32,000 services. The packet size we used was 100 bytes and the data rate in each service was 200 packets per second. This profile meant that our test traffic remained below the available capacity in the ring such that when the failure was emulated we were not causing packet loss in the network. The data rate also meant that the margin of error in our measurements was +/- 5 ms. As you can later see in the results, that error margin was not an issue. The switches in the test network were configured with 3.3 ms CCM intervals.

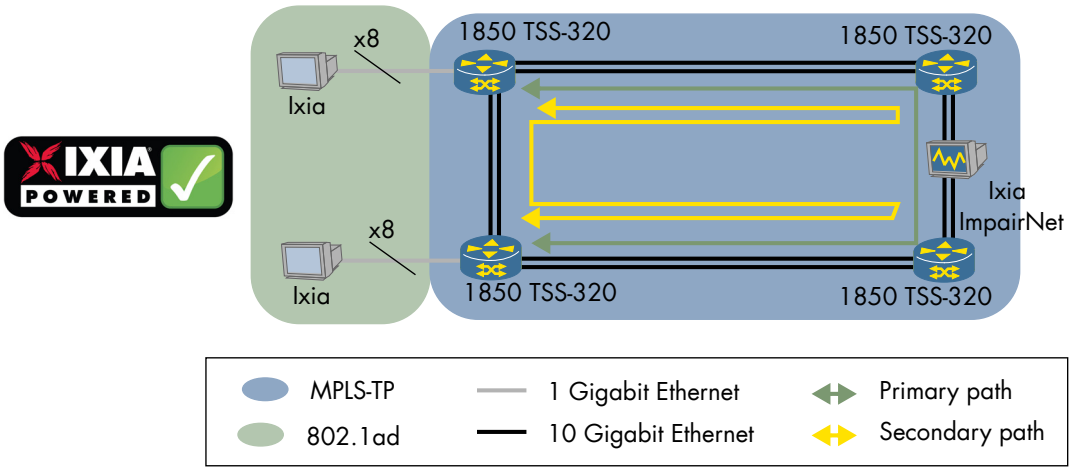

Figure 1: Test Topology

Test Highlights - Max Out of Service Times

- ➔ Unidirectional Link Failure - 20 ms
- ➔ Bidirectional Link Failure - 40 ms
- ➔ Logical Link Protocol Failure - 5 ms
- ➔ Physical Node Failure - 15 ms
- ➔ APS Controller Hot Standby - 5 ms

Test Results - Unidirectional Link Failure

The first condition that we emulated was a failure of the Tx laser on one of the devices. In this failure the neighboring switch looses its light source, but the optical port remains in up state. After we started sending traffic on all 32,000 services we created the failure condition and measured the amount of frames that were lost in all services. Our goal was to measure an out of service duration of less than 50 ms and indeed these were the results. In three test runs the worst out of service duration was 20 ms and the best 5 ms. 82% of the streams converged in 10 ms, while only 0.01% of the 64,000 actually incurred 20 ms service interruption. Since the network was configured in the only mode available - revertive, we also measured how fixing the issue effect the services. Here we measured consistently, in all three test runs, 5 ms.

Figure 2: Percent of Out of Service Time Distribution Per Service

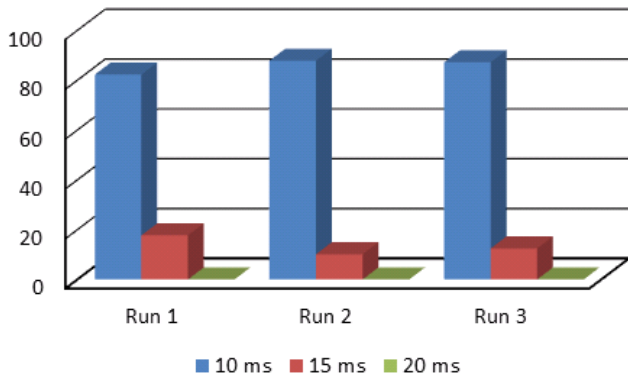

Test Results - Bidirectional Link Failure

In the next test we simulated a complete laser failure by disconnecting, simultaneously, both Tx and Rx fiber pairs from one of the devices. This failure condition is analogous to the headline grabbing "Fiber Cut" news which seems to plague the industry. The network wrapped around the troubled link loosing up to 40 ms of traffic in the worst test run for one circuit only. As is seen in the histogram Figure 3, the majority of the circuits (between 45% and 65%) converged within 10 ms.

The impact on the services during the process of reverting to the working path was measured, as in the previous test case, at 5 ms.

Figure 3: Out of Service Time as Percent of Services Histogrm

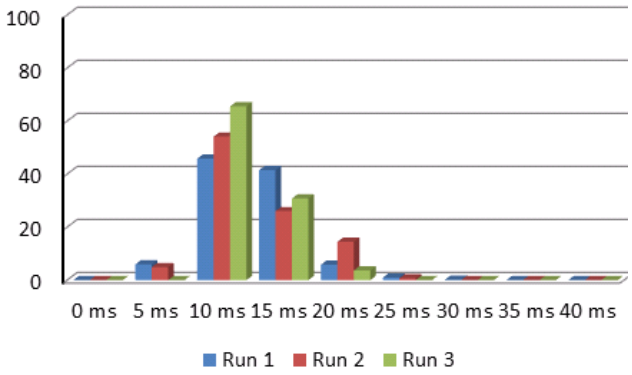

Test Results - Logical Link Protocol Failure

The third failure condition we simulated was aimed at demonstrating that the MPLS-TP section CCM is actually used for fault detection. After we first verified that traffic was traversing the ring in its working path, we used Ixia's ImpairNet to drop all Connectivity Check Messages (CCM) in the path. The two nodes never lost the laser light and traffic could still traverse the link, but the link liveliness protocol was no longer exchanged between the two nodes.

The results were a consistent 5 ms out of service time in all three test runs. We had to scratch our heads a little and discuss with the Alcatel-Lucent engineers the conditions which led us to these values. The explanation we came to was quiet elegant though.

CCM was configured with 3.3 ms intervals with a multiplier of three. This meant that at best we could expect an out of service time of 6.6 ms and at worst 9.9 ms. The explanation to the expectation is in figure 2. So how did we explain 5 ms out of service time?

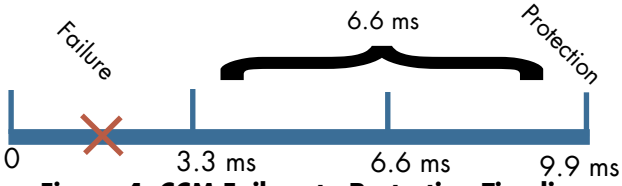

Figure 4: CCM Failure to Protection Timeline

Since traffic was still passing in the network while CCM was no longer being exchanged between the two nodes, we only measured the out of service duration caused by the protection switching action. That is the time the switch adjacent to the failure took to signal the protection event and switch to the protected path. Investigating further, we noted that 99.99% of the circuits were not affected by the loss of the CCM messages at all. From all three test runs, the highest number of affected circuits was 3 out of 64,000.

When we deactivated the filter blocking all CCM messages traffic converged to the protection path in the same time that we have been measuring all along: 5 ms in all three test runs.

Test Results - Physical Node Failure

In this test, to emulate a catastrophic failure of a complete node, we pulled the rag under a node by disconnecting its power supplies at once. This node immediately was taken offline leaving both its neighbors with a coordinated efforts of switching the traffic to the protecting path.

Figure 5: Percent of Out of Service Time Distribution Per Service

We were positively surprised at the results. In all three test runs we recorded the same out of service time of 15 ms.

Test Results - APS Controller Hot Standby

In the last test case we simulated a failure of the active APS controller, residing on the Matrix card, in a network element that had two matrix cards - active and standby. The failure of the active matrix card did not require detection using any messages since it happened internal to the switch, but since the APS state machine is also residing on the matrix card we expected some hit to traffic due to equipment protection and no MRPS protection switching action.

The two matrix cards were also configured in a non-revertive way. This meant that when an active matrix card failed, the standby one became active. When the failed card was brought back online it did not assume active state, but remained standby.

The effect this matrix card failure had on the network was consistent and minimal - we recorded 5 ms out of service time on all three test runs.

Summary

As more and more SDH devices are reaching their end of life, service providers could be assured that packets based solutions are being created to provide the same level of high-availability to which they have grown accustomed. Our tests show that the industry standard 50 ms failure recovery, so heavily established in SDH networks, could also be reliably supported by modern packet-based devices.

The ability to recovery a high number of services within 50 ms is only one of the messages we could deduce from this test. We also see consistency in the results. Why is consistency important? In the service provider world, where SLAs often involve dire financial consequences, knowing what could really be guaranteed in an SLA affects the bottom line at the end of the day. The solution performed as expected, without exceptions - exactly as transport equipment should.

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
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/

v5.0 20121018 JG