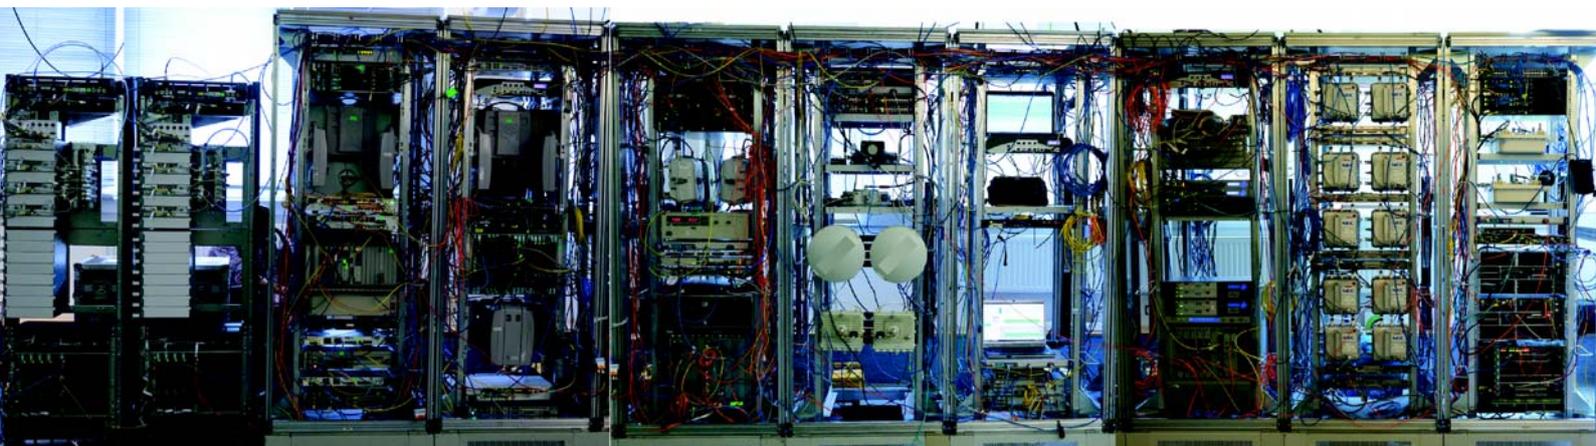


Carrier Ethernet World Congress 2011

Public Multi-Vendor Interoperability Event

White Paper



EDITOR'S NOTE



2011 marks the seventh year of the Carrier Ethernet World Congress and its EANTC multi-vendor Carrier Ethernet interoperability tests. In preparation of our showcase, I wondered how much innovation we would be able to witness year. Typically, vendors who were early adopters aim to capitalize on their investment when the market matures, as the absence of some large backbone equipment vendors shows this time.

Substantial innovation happens in other areas these days, closing some gaps for Carrier Ethernet services:

- The active integration of packet-based microwave solutions into Carrier Ethernet networks is driven by fierce competition. We tested full-blown switching functionality with clock synchronization, Ethernet Ring Protection Switching and more.
- End-to-end service activation quality assurance procedures is progressing to reduce the carriers' huge provisioning cost and helping to ensure service levels.
- Protection in provider bridging- and MPLS-TP-based access and aggregation networks, bringing them up to speed with MPLS in the access.

We have tested MPLS-TP since 2009. The industry is still undecided about two alternative, incompatible ways to implement protection. As usual we asked all vendors to bring innovative solutions with new aspects, as opposed to repeat previous showcases. This reduced the field to two equipment vendors committed to BFD-based solutions. Sometimes the difference between political statements and technology reality can be surprising.

Packet-based clock synchronization continues to be an active test area. Our event provides a unique opportunity to test multi-vendor solutions in a realistic

end-to-end network scenario. This time we focuses on boundary clocks, but had to deal with a number of functional multi-vendor interoperability points. Although progress has been made, IEEE 1588:2008 is still clearly a non-trivial technology.

Overall the Carrier Ethernet paradigm has gained overwhelming support in the industry. Basic E-Line services are a staple worldwide. The more advanced concepts which enable service providers to build value-added services, permit competitive differentiation and reduce operational cost, however, still require care in multi-vendor environments.

INTRODUCTION

Sometimes diplomatic skills are required when we organize interoperability events at EANTC. We poll vendors' technical interests, author documents describing what will be tested, provide a platform to execute tests and when things do not work as planned, bring the participants to the table, mitigate and help solve problems.

One of the major cornerstones of our interoperability showcases is that they focus on new technology solutions and/or new products. This is why the reader will find most test cases updated and refined compared to previous years' events, or new products being put to test.

In 2011, our focus was on pure Ethernet transport tests. 12 out of the 14 test cases were concerned with Ethernet testing from clock synchronization to service activation.

After months of preparation and two intense weeks of hot staging in our new offices in Berlin in which we created 106 different test permutations, we concluded a successful event. As in each interoperability event we had long debugging sessions and meetings to discuss test expectations and results, but the incremental increase in implementations and advanced features does show that Carrier Ethernet is moving forward.

INTEROPERABILITY TEST RESULTS

The following sections of the white paper describe the test areas and results of the interoperability event. The document generally follows the structure of the test plan – a document edited by EANTC and reviewed together with vendors in preparation for the event.

Terminology. We use the term "tested" when reporting on multi-vendor interoperability tests. The term "demonstrated" refers to scenarios where a service or protocol was terminated by equipment from a single vendor on both ends.

Test Equipment. In order to perform our tests we had to generate, measure, impair, and analyze Ethernet traffic and perform synchronization analysis. We thank Calnex Solutions, Ixia, Spirent Communications, Sunrise Telecom, Symmetricom and Veryx for their test equipment and support throughout the hot staging.

TABLE OF CONTENTS

Participants and Devices	3
Service Activation Test.....	3
Resiliency and Redundancy	6
Topology	8
Clock Synchronisation.....	10
Ethernet Microwave Tests.....	12
Demonstration Network.....	14
References.....	15

PARTICIPANTS AND DEVICES

Vendor	Devices
Albis Technologies	ACCEED 1416 ACCEED 2202
Aviat Networks	Eclipse Packet Node
Brocade	MLXe-4 NetIron CER
Calnex	Paragon-X
ECI	SR9705 SR9608 BG9310
Ericsson	MINI-LINK TN MINI-LINK LH MINI-LINK PT2010 MINI-LINK CN510 MINI-LINK PT6010 MINI-LINK SP110 MINI-LINK SP210 MINI-LINK SP310 SPO1460 SPO1410
Extreme Networks	E4G-200 E4G-400
Hitachi	AMN1710
Ixia	IxNetwork ImpairNet
Linkra	WIDHOP
NEC	iPASOLINK 400
Nokia Siemens Networks (NSN)	FlexiPacket MultiRadio FlexiPacket FirstMile 200 FlexiPacket Hub 800
Rohde & Schwarz	SITLine ETH 1G
Siklu	EtherHaul-1200
Spirent Communications	Spirent TestCenter Spirent GEM Spirent XGEM Spirent Anue 3500
Sunrise Telecom	RxT SL GigE RSPDR

Vendor	Devices
Symmetricom	Cesium Reference CsIII TimeProvider 1500 TimeProvider 5000 TimeProvider 5000 Evolution SSU 2000e
Veryx	XC1000

SERVICE ACTIVATION TEST

A brand new test area this year was Ethernet service activation. The test followed the ITU-T standard Y.1564 "Ethernet Service Activation Test Methodology," a standard approved in January 2011. The standard provides methodologies and specific tests to allow service providers to validate that new Carrier Ethernet services coming online meet Service Level Agreements (SLA) sold to the customer.

An informative appendix of the standard defines a number of tests that evaluate Frame Loss Ratio, Frame Transfer Delay, Frame Delay Variation, and Availability service performance parameters. These parameters are commonly known as objective service performance indicators and are also often described in SLAs. In order to meet the performance parameters defined for a service, the provider has to configure certain Quality of Service parameters on the devices used to construct the service — this is normally where the difficulties lie.

In order to support multiple service classes, a service provider has to configure parameters such as Committed Information Rate (CIR), Excess Information Rate (EIR), Committed Burst Size (CBS), Excess Burst Size (EBS) as well as color mode (color-aware or color-blind). As our testing showed, configuring all these parameters is a challenging task since MEF standard nomenclature does not necessarily align with different vendors' configuration language.

The test was conducted over two Ethernet Virtual Private Line (EVPL) services that were defined as "Services Under Tests". In essence these were the new services that were being activated. The First service, EVPL1, was defined to operate in color-aware mode. The service was provisioned with three Bandwidth Profiles (BWP) per EVC per CoS as specified in the MEF 10.2 standard and the MEF 23.1 draft standard. The profiles were applied on the UNIs of EVPL1, with no policing on subsequent ports in the network:

BWP1 – CIR 5 Mbps, CBS 16,000 bytes, EIR 0 Mbps, EBS 0 bytes; The service was defined as very important (we called this service Metro High) and had one color: Green.

The second bandwidth profile was used for medium priority traffic. Its characteristics were: CIR 10 Mbps, CBS 16,000 bytes, EIR 10 Mbps, EBS 16,000 Bytes. Here, two colors were defined: Green and Yellow.

The third bandwidth profile, *BWP3*, was defined for

low priority traffic and included CIR 0 Mbps, CBS 0 bytes, EIR 15 Mbps, EBS 16,000 Bytes. Here only Yellow color was defined.

In order to make sure that the tests are thorough, each bandwidth profile at each end of the network was first separately tested to verify its configuration and performance prior to connecting the service and testing the entire service. This procedure separately confirmed the proper operation of each bandwidth profiler. In EVPL1, the Albis ACCEED 2202 was defined as the customer premises equipment and the ECI 9608 was defined as the Aggregation or Provider Edge device.

Once the configuration of each bandwidth profile was confirmed, a loopback test confirmed the performance of the EVC to Y.1564.

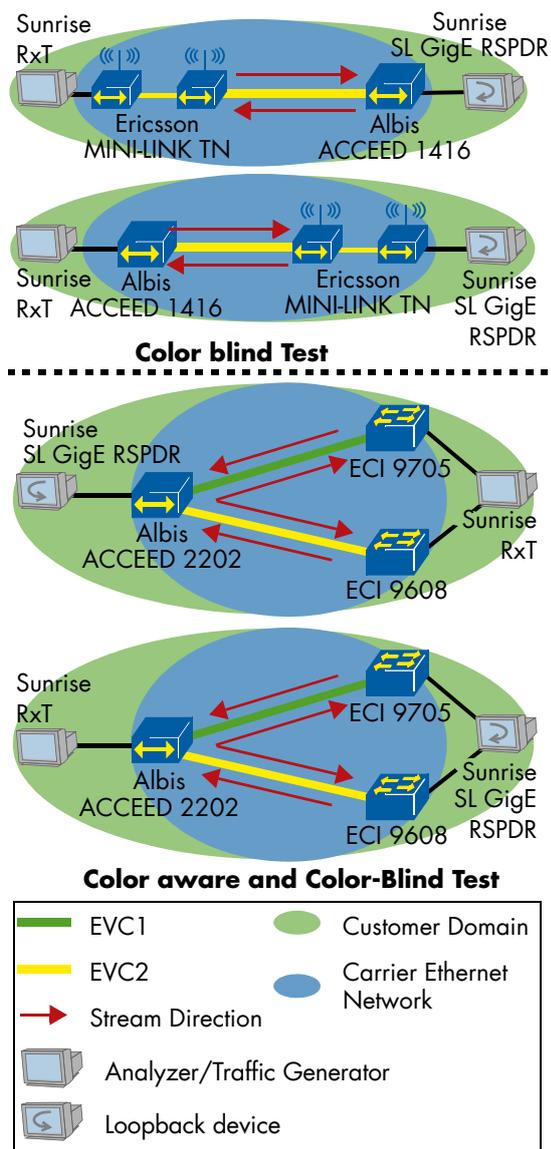


Figure 1: Service Activation Tests

In the second, color-blind scenario, we defined a second EVPL service with slightly different per EVC per CoS bandwidth profiles:

BWP1 – CIR 15 Mbps, CBS 16,000 bytes, EIR 0 Mbps, EBS 0 bytes and *BWP2* – CIR 20 Mbps, CBS 16,000 bytes, EIR 30 Mbps, EBS 16,000 Bytes.

We repeated the same procedure, first verifying that the participating vendors devices', which included Albis ACCEED 1416, ECI 9608 and Ericsson MINI-LINK TN, had the correct configuration applied and then verified that the service parameters such as Frame Loss Ratio, Frame Transfer Delay, Frame Delay Variation, and Availability reported the expected values within the Service Acceptance Criteria defined by Y.1564 and the Metro Class of Service Performance Tier of MEF 23.1 (under development).

One of the biggest challenges of the test was to achieve a common understanding of QoS configuration between vendors. Since QoS architecture is not standardized well, every vendor has its own interpretation of QoS. For example, the committed information rate parameter on one device was configured with an Information Rate value while another device configured the same parameter using Utilized Line Rate (ULR). We also found that some implementations did not allow the operator to configure burst size at the granularity we had originally specified in the test plan. We actually asked for 12,176 bytes of burst size, but due to some limitations of some of the devices we converged on a value of 16,000 bytes. In the same configuration area (burst size) we also found that some of the implementations allowed a very granular configuration, however, the tests showed that the configuration was actually not enforced by the hardware.

For service providers we see these initial results as a positive sign. The instrumentation and standards to verify service activation now exist and both solutions and services can be tested before an Ethernet service is handed over to a customer. Especially now that Ethernet services are becoming more and more mission critical and are replacing legacy TDM solutions, having the standards to verify that the services are really working as advertised.

Performance Monitoring

Once an Ethernet service has been activated, monitoring its performance in-service is one of the most desirable requirements for service providers. After all service providers sign contracts with their business customers to deliver a certain quality of service (defined in Service Level Agreements, SLAs). These SLAs often stipulate financial implications in the case that the service is "out of SLA".

The metrics to perform monitoring of services is defined in the ITU-T specification Y.1731. The specification introduces techniques to measure service performance objectives such as frame loss, frame delay, and frame delay variation on point-to point Ethernet services.

We started interoperability testing of Y.1731 implementations when the standard was young in 2008. But just around now, three years later, service providers really begin to deploy performance monitoring in production-grade Carrier Ethernet services. Vendors brought new or updated products to our test this time. The level of interest was so high

that we were able to create a list of 25 vendor pairs for the test.

We distinguished between two measurement type for the tests: Two-way frame delay and two-way frame delay variation.

We used a procedure that allowed us first to validate that the protocol exchange between two implementations worked as expected focusing on Y.1731 Delay Measurement Messages and Replies (DMMs and DMRs). Once this test procedure was completed and a baseline was reached, we added an impairment generator for all further tests. We used Calnex Paragon-X, Ixia ImpairNet and Spirent GEM impairment generators to introduce controlled delay and delay variations so we could verify that not only the implementations worked with each other, but also worked correctly.

For the frame delay measurement tests we used the impairment generators to add constant delay of 20 milliseconds (ms) to all DMM packets. We expected that the implementations will report an average frame delay value of 20 ms. It was not an unrealistic expectation given that the propagation delay in our lab approached zero. In order to verify that the two implementations measured frame delay variation correctly we asked the impairment generator vendors to introduce packet delay of 15 ms to every second DMM packet and 25 ms to the remaining packets. With this procedure we expected to observe average frame delay of 20 ms, and frame delay variation of 10ms.

The following devices successfully participated in the Y.1731 delay/delay variation tests (see Figure 2): Albis ACCEED 2202, Albis ACCEED 1416, Brocade MLXe-4, ECI 9608, ECI 9705, Ericsson SPO1460, Ericsson MINI-LINK SP310, Ixia IxNetwork, Spirent TestCenter.

Each pair in the figure represents a vendor pair which constructed a Carrier Ethernet service using VLANs (IEEE 802.1Q standard) and that have passed the test.

After we completed performance measurement tests between Brocade MLXe-4 and Ericsson SPO1460, as well as between Albis ACCEED 2202 and ECI 9608, we inserted the Rohde & Schwarz SITLine ETH 1G encryption devices and repeated the test with the same measurements. Our goal was to verify if the measurements could still be executed between two devices over an encrypted EVPL service.

We then verified that the performance measurements were still working over the encrypted service and also, using the delay measurements themselves, that the service's delay and delay variation was only marginally impacted.

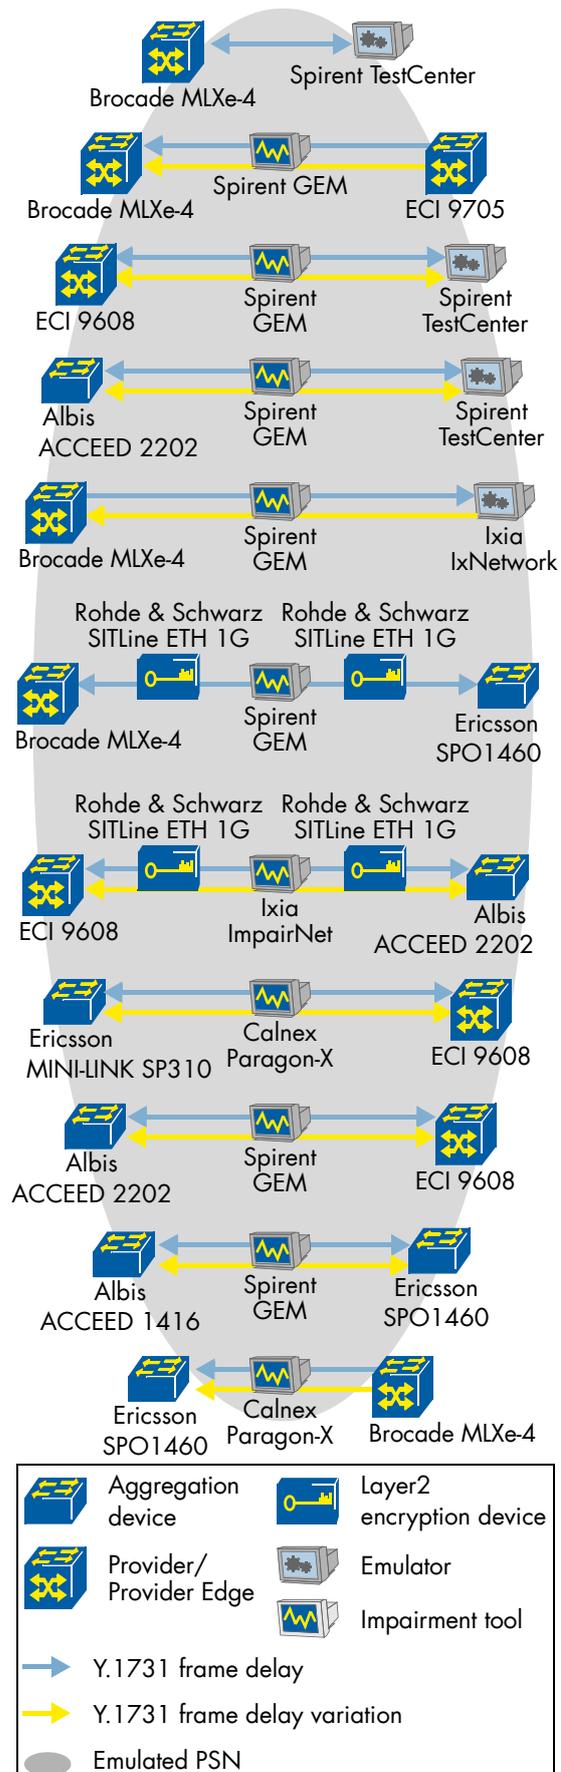


Figure 2: Performance Monitoring

Brocade MLXe-4 and Ixia IxNetwork also performed the baseline performance monitoring test over 100 Gigabit Ethernet (GbE) interface. As we could not instrument the event with a 100GbE impairment generator, we had to skip the accuracy verification.

Another performance measurements test that was executed using the Brocade MLXe-4 over 100GbE interface used a Virtual Private LAN Service (VPLS) service with Spirent TestCenter. Again, only the basic protocol interoperability was verified in this case and no delay and delay variation accuracy was measured.

During the performance monitoring test, we observed three implementations for delay measurement. One vendor supported only on-demand delay measurement. This implementation is useful where the delay measurement function is triggered by an operator or by an external management system.

Two vendors supported scheduled measurement. In this case, the values of delay and delay variation are calculated over a specific time interval. The counters used for the calculation were cleared at the end of the time interval, and the measurement continued over the next interval continuously.

There was also one implementation where the mean value of delay and delay variation could be observed only once the DMM session was stopped manually. In order to follow the procedure in the test plan, a restart was required at each step. It was possible to read the live run delay values however.

We were pleased with the results we collected for this test and believe that some of the implementations benefited from the opportunity to test against so many other implementations. We see that the protocol gets used in a growing number of implementations. This is a positive signal to the industry that services could be delivered with objective and ongoing quality measurements.

Hierarchical Service OAM

The intention of Hierarchical Service OAM was to verify the interoperability of Service OAM functions between different vendor Maintenance Association End Points (MEPs) at multiple maintenance levels. Since no other vendor participated in this test case, Albis created a demo to show the support of Ethernet Alarm Indication (ETH-AIS), and Ethernet Locked Signal (ETH-LCK) functionality in an OAM hierarchy.

We verified that the ETH-AIS messages were received successfully at the subscriber level in case that the transport path and the configured service failed. We also verified that the ETH-LCK information was received periodically at the subscriber level with no impact on the traffic, until the administrative lock condition was removed. The demo was presented successfully with no issues observed.

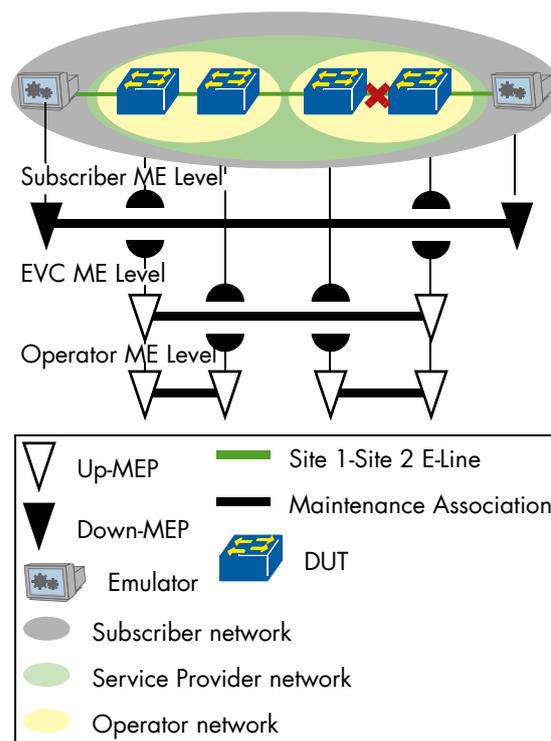


Figure 3: Hierarchical Service OAM

RESILIENCY AND REDUNDANCY

Some assumptions could be made these days with respect to the way carrier packet networks are designed: The core of the network is likely to be based on MPLS (with or without transport profile) and the access will be built based on available physical layer infrastructure. The test scenarios planned for this area were meant to provide an overview on the available tools for service provider to offer reliable Ethernet services to customers. Resiliency is after all one of the five attributes that describe Carrier Ethernet so we set out to investigate the options available in the market.

Ethernet Ring Protection Switching

Ethernet has come a long way since the days in which resiliency was performed using spanning tree protocols. These days the target set in the ITU-T G.8032 standard for resiliency is 50 ms — a target as aggressive as the *gold* standard set by SDH.

G.8032 or Ethernet Ring Protection Switching (ERPS) is a standard for fast recovery in Ethernet based ring topologies using automatic protection switching (APS) protocol.

ERPS-enabled Ethernet nodes send ring automatic protection switching signal failure (R-APS(SF)) messages as soon as they detect a failure on a ring. The failure signal causes the switch on that ring that has one of its ports in blocking state (RPL owner) to unblock its RPL port therefore forcing a change to the ring in order to bypass the failed ring segment. The nodes adjacent to the failed link isolated the failed

network segment by blocking their ports and traffic is then forwarded through the open RPL port.

In revertive operation mode the nodes detect that the failure is resolved and are expected to send ring automatic protection switching no request (R-APS(NR)) message causing the RPL owner to start the Wait to Restore (WTR) timer. Upon expiration of the WTR timer, the RPL owner blocks its RPL port. Nodes adjacent to the former failed link unblock their ports and the ring then resumes its original operation. At this point all nodes flush their Forwarding DataBases (FDBs). This whole procedure should theoretically occur within 50 milliseconds. The standard allows for non-revertive operations, however, we did not test this mode.

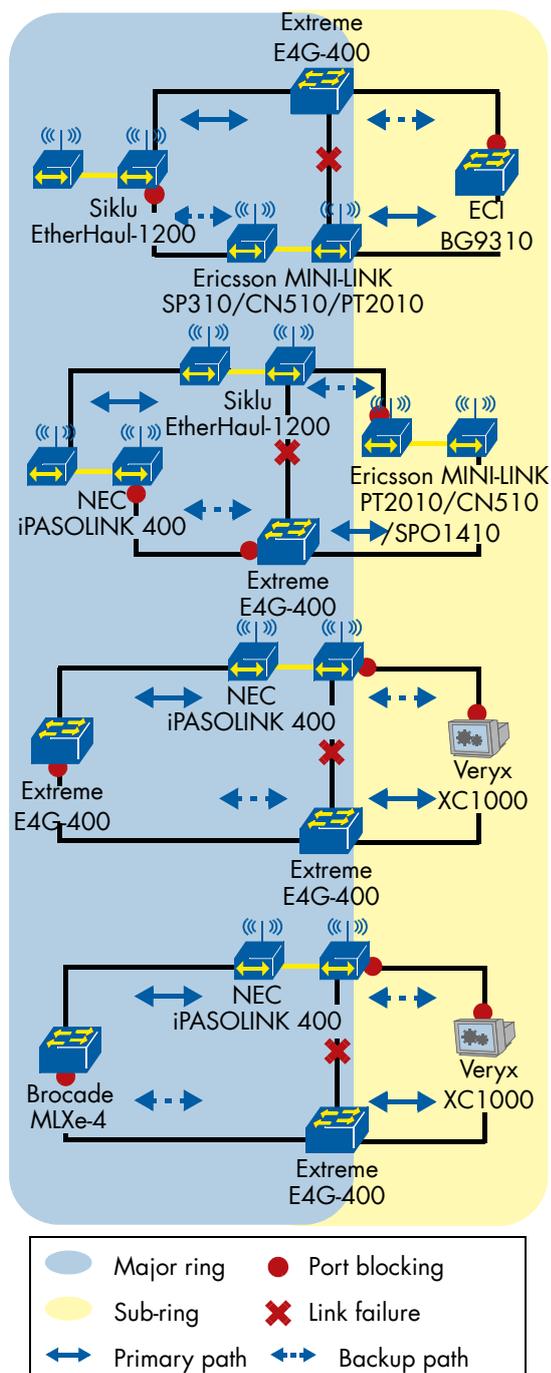


Figure 4: Ethernet Ring Protection Switching

Our goal was to verify that a ring, constructed by a

number of vendor solutions, will behave as described above. The assumption we made was that the access is exactly where carriers would require a functioning multi-vendor rings. After all, access is often the network area that includes switches that were positioned in the core of the network or new solutions that were added at a second network rollout phase.

The following vendors participated in this test area (see Figure 4): Brocade MLXe-4, ECI BG9310, Ericsson MINI-LINK PT2010/ MINI-LINK CN510/ SPO1410, Ericsson MINI-LINK SP310/MINI-LINK CN510/MINI-LINK PT2010, Extreme E4G-400, NEC iPASOLINK 400, Siklu EtherHaul-1200, and Veryx XC1000.

In all test scenarios two rings were constructed: a major ring and a sub-ring. Both were interconnected via a shared link. In each ring one dedicated node was provisioned as RPL owner. In most of the tests a RPL neighbor was configured in both major and sub-ring to block the other end of the RPL link.

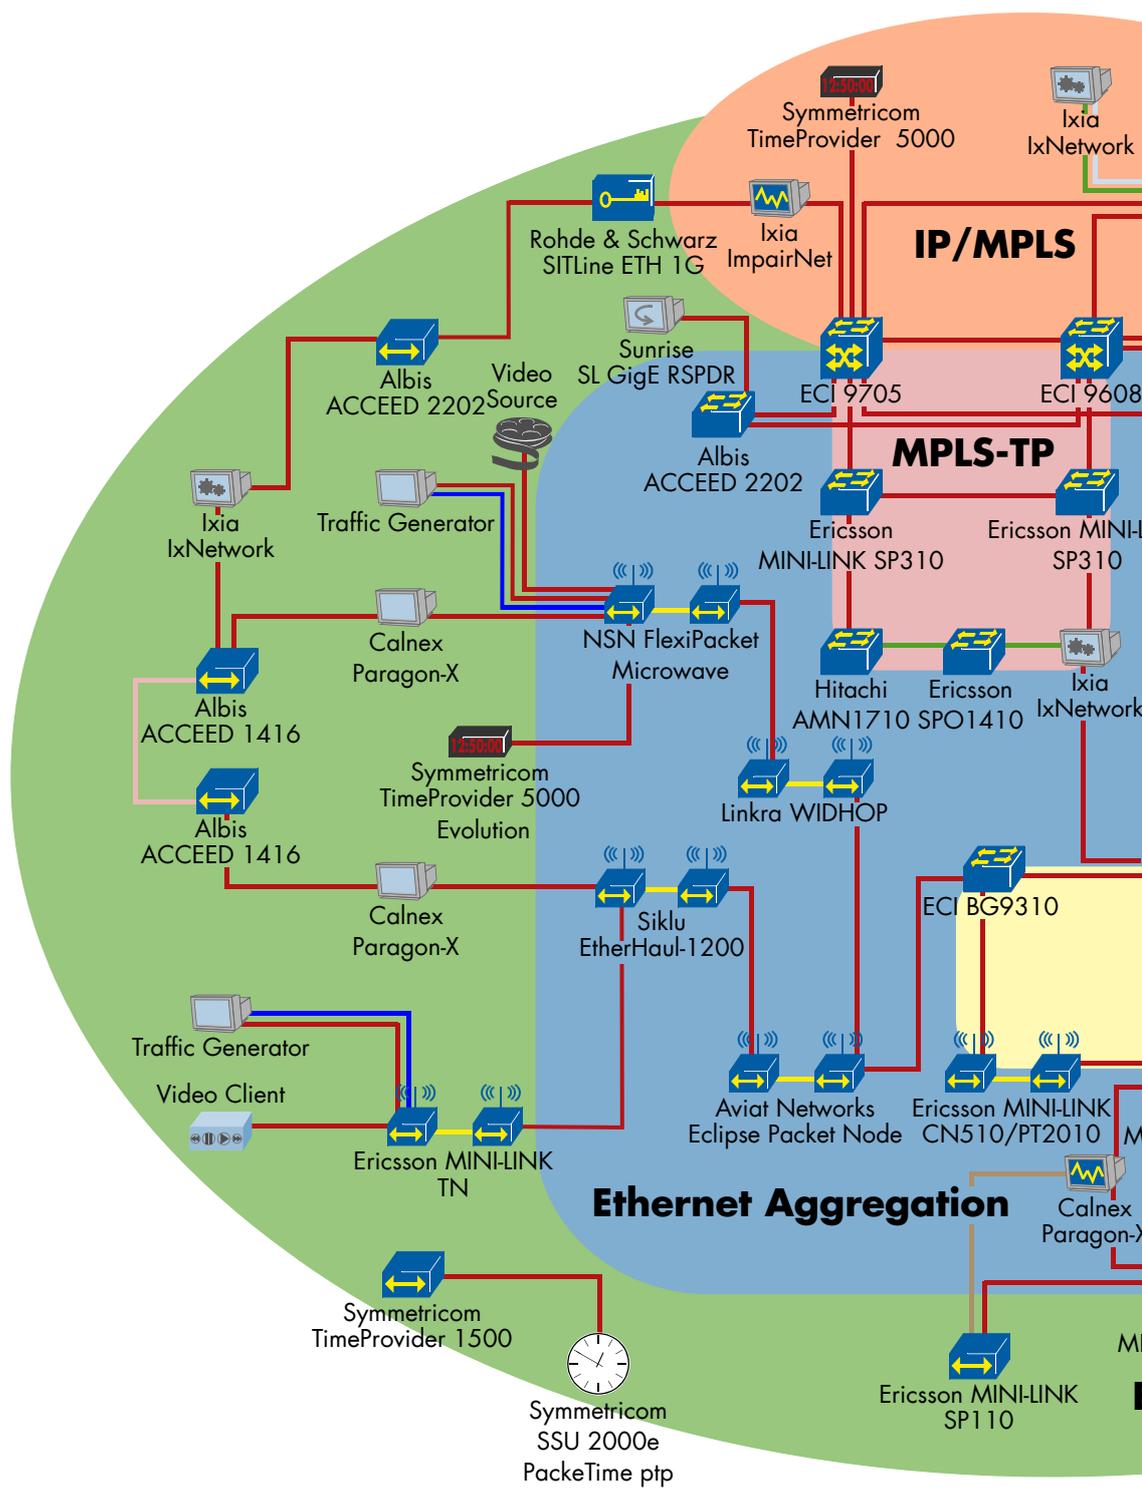
In order to perform the test we defined two profiles to accommodate different implementations: The first profile (Profile 1) was designed for vendors allocating different VLAN ID in both ring for R-APS communication. The second profile was intended for vendors that use the last octet of the MAC address as RingID for the R-APS communication to distinguish the R-APS channel in between the rings.

Initially, Connectivity Fault Management (CFM) was requested to be used to monitor the liveness of the ring. Since vendors take different approaches to realize CFM - Some vendors used a separate VLAN ID for CFM and R-APS communication whereas other vendors used the same VLAN ID for both CFM and R-APS, it was difficult to find a common link liveliness. After a long debugging session we could not converge on a single method and in order to move on with the ERPS testing we decided to run the test without CFM. Instead we triggered the failover by pulling the cable from one of the devices.

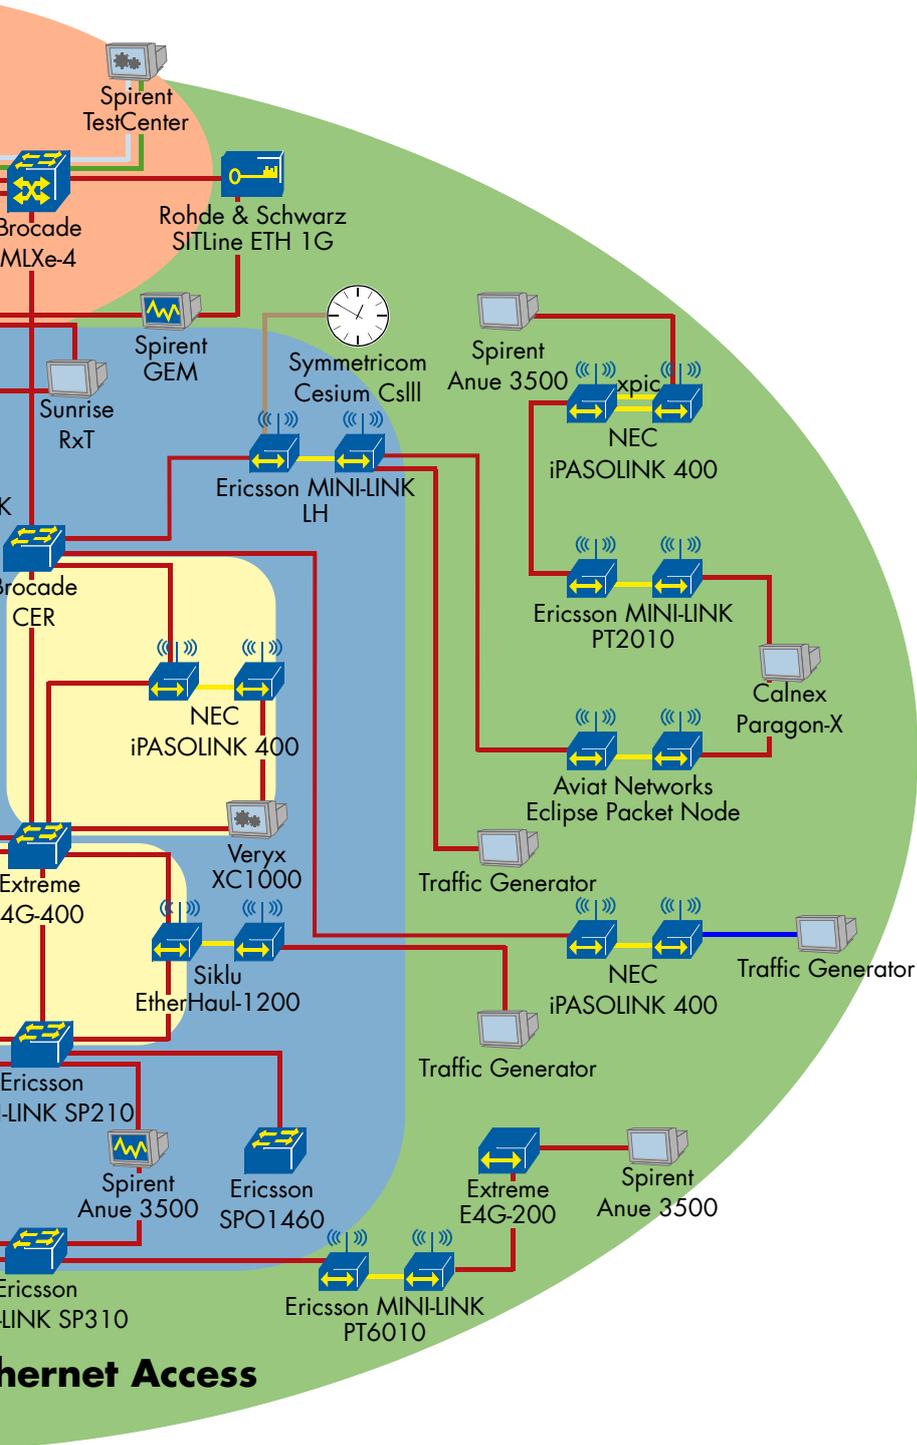
In all test scenarios, traffic was sent at 1,000 frames/seconds on the service configured in the ring using either Ixia IxNetwork, Spirent TestCenter or Veryx. We emulated link failure between the interconnection nodes and verified the proper reaction of the ring.

In the first test scenario, ECI BG9310 participated in the sub-ring as ERPSv1 and all other devices participated in this test scenario and in all subsequent test scenarios with their ERPSv2 implementation.

The Veryx solution acted as a ring node and was also able to generate traffic in the ring. When we wanted to verify that the Veryx XC 1000 had the RPL port block, we realized, together with the Veryx engineers that there are no commands available on the solution to show such port state. None the less we could verify, using traffic flow directionality in the ring that the Veryx solution was indeed blocking the port as expected.



Connection Types		Device Types	
 Clock link	 Bonded SHDSL link	 Provider/ Provider Edge device	 Emulator
 1Gbit Ethernet link	 10Gbit Ethernet link	 Aggregation device	 Access device
 TDM link	 100Gbit Ethernet link	 Analyzer/Traffic Generator	



Internet Access

Device Types

-  IEEE 1588v2 Grandmaster
-  Layer2 encryption device
-  Microwave radio system
-  Reference Clock
-  Impairment tool
-  Loopback device

Network Areas

-  IP/MPLS network
-  MPLS-TP network
-  Ethernet Aggregation network
-  Access network
-  ERPS Rings

In all test scenario after pulling the cable between interconnection ring nodes, we successfully verified that the protection switching was triggered - RPL node in the major ring unblocked the RPL port and the port connected to the failed link was blocked. The RPL port in the sub-ring remained blocked. The failover and restoration time was asymmetric and ranged from 51 to 290 milliseconds and 7 to 96 milliseconds respectively.

We also successfully tested the following administrative commands on all major ring RPL owner devices: "Manual Switch" and "Force Switch" both of which move the blocked port as desired around the ring. We also tested the "Clear" command, which removes both "Manual Switch" and "Force Switch" commands. The failover time during these set of tests ranged from 0 to 103 milliseconds and the restoration ranged from 0 to 91 milliseconds.

During the test we encountered several issues ranging from configuration to interoperability. Specially we were unable to perform the test using profile 2. Using said profile, and while emulating the link failure on the shared link, nodes connected to the failed link (interconnection nodes) started to send Signal Fail (SF). After receiving this message, the sub-ring node interprets the messages and removes the RPL block, which causes the traffic to flow on a different direction than expected.

MPLS-TP OAM 1:1 Protection

One of the areas of interest for MPLS-TP vendors has been protection scenarios. For this event three vendors were ready to support MPLS-TP 1:1 protection tests. Two impairment generators — Calnex Paragon-X and Spirent Anue 3500 — offered to provide impairment functions. In preparation for the tests we asked vendors to bring new implementations as usual. As a result of the poll, we focused on BFD implementations only this time since all available Y.1731 implementations have been tested in our interoperability events frequently in the past.

The test scenario was straightforward from a testing perspective: Build an MPLS-TP based network using 1:1 protection, fail the active link and verify that the network converges in under 50 ms. This test was performed by statically configuring two bidirectional Label Switched Paths (LSP) - primary and secondary - using a single physical connection between Ericsson SPO1410 and Ixia IxNetwork. BFD Continuity Check (CC) sessions were running on both primary and secondary LSPs to monitor the liveness of the LSP and the pseudowire service configured on top of the LSPs. BFD-CC was transmitted over the Generic Associated Channel (G-ACh) and Generic Associated Label (GAL). Both BFD sessions were running in a slow start mode of operation. Once the BFD sessions moved into the UP state the transmission and receive intervals changed to 3.33ms. Once both primary and secondary LSPs were established we offered load to verify that the traffic was indeed using the primary path. We then emulated an

unidirectional link failure on the primary path by dropping all traffic and BFD packets. We verified that traffic moved to the secondary path. The recorded failover when only the Worker BFD-CC flow was impaired was 0 milliseconds. The recorded failover when both the Worker BFD-CC and Worker Pseudowire flows were impaired was 12.4 milliseconds. After disabling the link failure emulation the traffic reverted back to the primary path without any loss. The restoration time was then zero milliseconds.

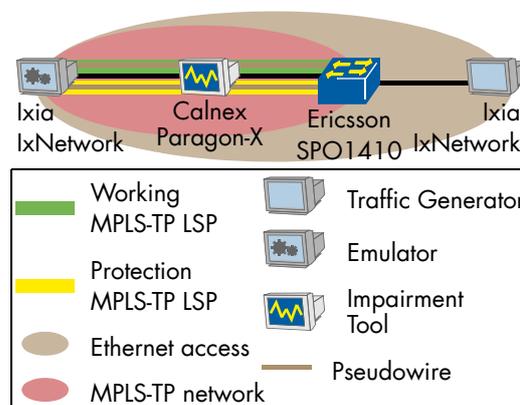


Figure 5: MPLS-TP 1:1 Protection

Besides protection switching due to the link failure, Administrative commands "Lockout of Protection", "force Switch normal traffic signal-to-protection", "manual switch normal traffic signal-to-protection" were all successfully tested based on draft-ietf-mpls-tp-linear-protection-03.

CLOCK SYNCHRONISATION

EANTC started testing clock synchronization interoperability already in early 2008. Over time we have seen a constant increase in the number of implementations and more and more positive results.

Participating vendors were still interested to continue clock synchronisation tests so we obliged. We also hear from our service provider customers that the level of trust in the multi-vendor interoperability of the associated protocols (specifically IEEE 1588v2) is still not as high as it should be. With this mandate, we set two areas of testing for the IEEE Precision Time Protocol (PTP or IEEE 1588v2): Ordinary Clocks and boundary clocks.

IEEE 1588 Master/Slave – Phase and Frequency Synchronization

In certain conditions transporting clock signal over a packet based solution is the only option available for service providers. For example, Synchronous Ethernet can not be used over Ethernet services leased from other providers as the majority do not offer such services. At this point a service provider requiring synchronization services, for mobile backhaul or other services, could use the IEEE 1588v2 Precision Time Protocol (PTP). The protocol is also attractive when Time of Day and phase synchronization is needed.

Our methodology for verifying 1588v2 focuses on protocol interoperability as well as on clock accuracy. The latter is only used in order to verify that the interoperable implementations are also delivering useful clock information to the clients of the service.

The measurements were performed by connecting each slave device to a wander analyzer for frequency, via either E1 or 2048 KHz interfaces, and a frequency counter via 1 PPS (pulse per second) for phase (time of day, or TOD) deviation measurement. All test scenarios were required to pass the 15ppb mask for MTIE, and a maximum time error of 3 microseconds for the time of day deviation.

Impairment was introduced into the network according to ITU-T G.8261 Test Case 12 using a Calnex Paragon-X or Spirent Anue 3500 impairment generators. Frequency wander measurement was performed using a Spirent Anue 3500 and another measurement device over a period of 4 hours or more. Phase measurements were planned, but at the end were not tested since the slave clocks participating in the test only supported frequency output.

Configuration of all slave clocks was done over unicast IP/UDP on a specified VLAN, with 64 Sync messages per second in Two-Step mode.

We successfully tested two PTP slaves (see Figure 6) under impairment, the Ericsson SPO1460 and NSN FlexiPacket Hub 800, and one slave without impairment emulated by Ixia IxNetwork.

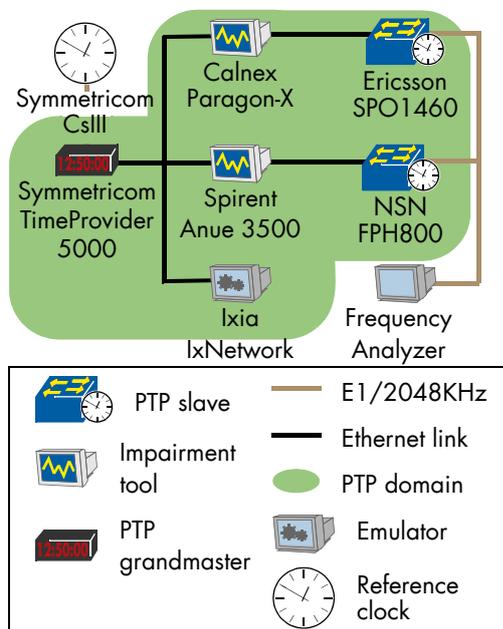


Figure 6: PTP Master/Slave Clock

In addition to successful tests, we found several interop issues among some vendors. These were mostly due to differing features (such as supported sync message rates or One-Step only vs. Two-Step only).

An additional reason for the lower than expected results for this test was the fact that the test duration was so long. This meant that synchronization was

lost during the measurement period, we did not have enough time to repeat the tests.

IEEE 1588 Boundary Clocks

There were two Boundary Clock scenarios run, as only one vendor attending this event brought a Boundary Clock to be tested.

Like in the Master/Slave synchronization test cases, the Symmetricom TimeProvider 5000 served as the PTP Grandmaster. Impairment was once again based on G.8261 Test Case 12, and provided by either an Calnex Paragon-X and Spirent Anue 3500.

The Ericsson MINI-LINK SP310 served as the Boundary Clock in both scenarios tested. The Slave Clocks tested were in a second PTP domain, as a slave to the Boundary Clock. This role was supported by an Ericsson MINI-LINK SP110 and an Extreme E4G-200 (see Figure 7). In the case of the Ericsson-Extreme interop, Ericsson used the MINI-LINK PT6010 microwave as a fiber replacement in between the Master and Slave Clock devices.

Both Boundary Clock and Slave Clock were configured with 64 sync packets per second in Two-Step mode, but the VLAN and PTP Domain were different on either side of the Boundary Clock to represent a realistic real-world scenario.

Frequency measurement was done with the Calnex Paragon-X or Spirent Anue 3500. 1PPS (Time of Day) measurement was done with a Calnex Paragon-X or a frequency counter. Like in the Master/Slave test case, all tests had to pass the G.823 15ppb mask or better for frequency, and have a maximum Time of Day deviation of 3 microseconds or less.

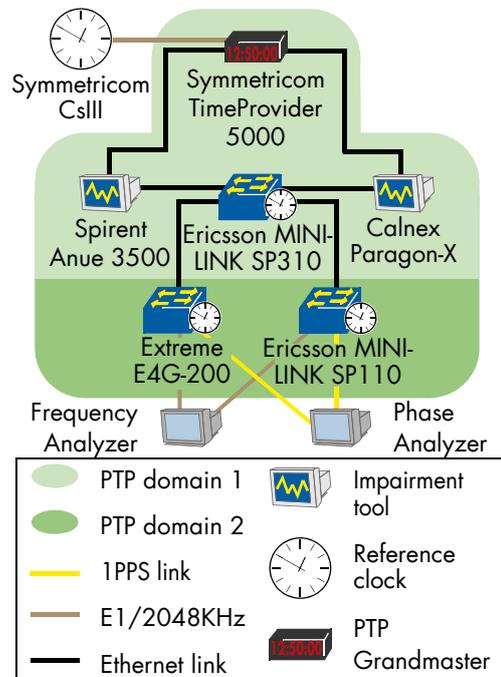


Figure 7: PTP Boundary Clock

Ethernet Synchronization Messaging Channel (ESMC)

To test ESMC, three devices were configured in a linear fashion to have varying reference clock qualities. Supplied externally to the devices are Primary Reference Clock (PRC) and Synchronization Supply Unit (SSU) signals. A device that does not have a clock source connected is expected to send DNU (Do Not Use) messages or SEC (SDH Equipment Clock), since the internal clock on most devices are expected to comply to SEC specifications at best. Due to the clock sources attached to devices in specific locations, we will refer to those nodes as PRC, SEC, and SSU.

Clock sources were connected and disconnected from each node, allowing them to adapt while we monitored the links between them to be sure that the SSMs (Synchronization Status Messages) being sent were changed accordingly.

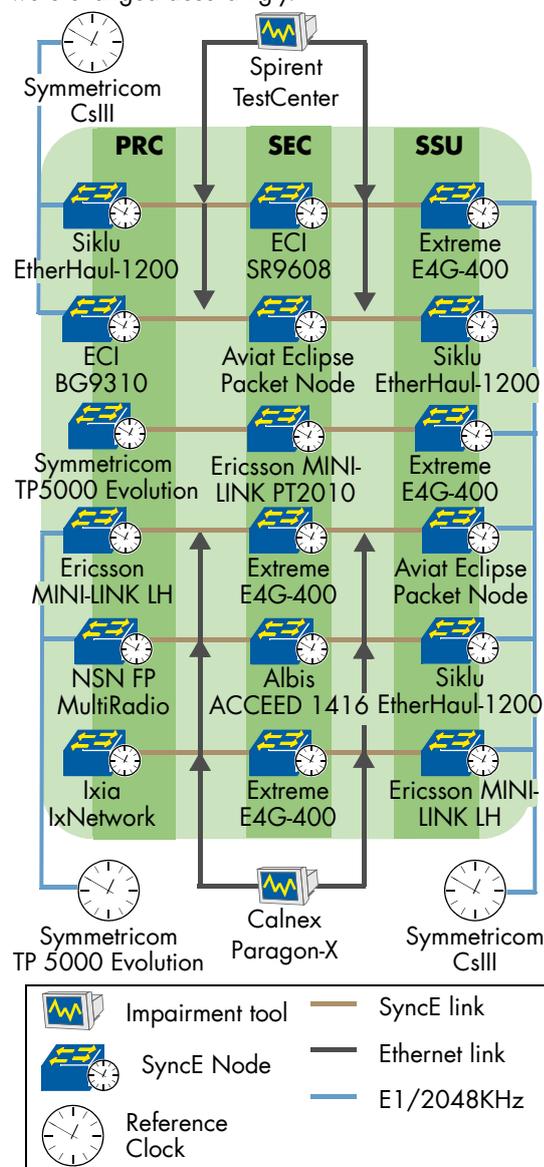


Figure 8: ESMC

Six ESMC scenarios were tested. The following vendors all took on various roles (see Figure 8) in each given configuration: Albis ACCEED 1416, Aviat Networks Eclipse Packet Node, Ericsson MINI-LINK LH, Ericsson MINK-LINK PT2010, Extreme

E4G-400, Ixia IxNetwork, NSN FlexiPacket Microwave, Siklu EtherHaul-1200, and Symmetricon TimeProvider 5000 Evolution.

The Spirent Anue 3500, Calnex Paragon-X, and Spirent TestCenter were used to evaluate the Synchronous Ethernet clock Quality Levels on the line, which were sourced from the Symmetricon CsIII.

One vendor found a bug in which they were not sending messages when the device reported they were, and one had an issue with switching to a lower-priority clock source. Both vendors tested successfully later with other products or bug fixes.

ETHERNET MICROWAVE TESTS

Our test ideas in this area focused on Ethernet functionality and not on microwave physical layer functionality. The microwave vendors were more than ready to demonstrate interoperability with other switches and routers.

We understand the reason behind the inclusion of more and more advanced features in microwave solutions: The microwave solutions are an integral part of many mobile backhaul solutions and as such have to satisfy an extensive set of requirements.

These requirements on microwave solutions could clearly be seen in the tests in which microwave vendors exclusively participated as well as in the tests that were opened to all.

We requested microwave vendors to participate as active components in all the tests we describe below and, as a policy, we did not focus on air interface testing (including capacity).

Microwave QoS Support

As microwave systems are evolving from pure transport to a more intelligent and complex solutions, one transport capability required is support for quality of service (QoS). The nature of microwaves link is its capacity variation that depends on weather conditions. It is therefore often desirable, that the bandwidth is preserved for more critical data, if the link capacity decreases.

In this test we configured two classes for which we generated test traffic: Priority and Best Effort. With the BestEffort test traffic we emulated Internet traffic, with the Priority test traffic we emulated legacy TDM voice traffic encapsulated into packets by means of Circuit Emulation Service (CES).

For each test we used two microwave systems of two different vendors that were connected via an Ethernet wire. Both systems were required to configure QoS and differentiate traffic based on Ethernet Priority Code Points (PCP) - a part of IEEE 802.1Q VLAN tag. We used PCP 0 for BestEffort and PCP 6 for priority class. During the test we generated 60 Mbit/s bidirectional IMIX (Internet mix) BestEffort traffic and 2 Mbit/s bidirectional Priority traffic with packet size equal to 256 Bytes. During the test we changed modulation of

microwave links on both participating systems first manually then via attenuating of the microwave link - emulating bad weather condition and triggering adaptive modulation on the systems. Our expectation was that we do not lose any packets from the Priority class and do lose packets of BestEffort class when link capacity was reduced.

11 Microwave System pair combinations successfully passed this test (see Figure 9): Aviat Networks Eclipse Packet Node and Ericsson MINI-LINK TN, Aviat Networks Eclipse Packet Node and Siklu EtherHaul-1200, Aviat Networks Eclipse Packet Node and Linkra WIDHOP, Ericsson MINI-LINK TN and Siklu EtherHaul-1200, Ericsson MINI-LINK TN and NSN FlexiPacket Microwave, Ericsson MINI-LINK TN and NEC iPASOLINK 400, Linkra WIDHOP and Siklu EtherHaul-1200, Linkra WIDHOP and NSN FlexiPacket Microwave, Linkra WIDHOP and NEC iPASOLINK 400, NEC iPASOLINK 400 and Siklu EtherHaul-1200.

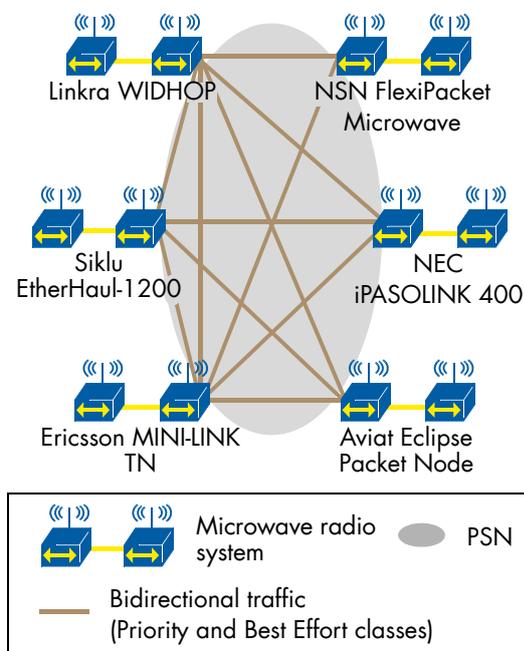


Figure 9: Microwave QoS Support

On the microwave links we used different modulation per vendor. When we emulated normal weather condition we used the following modulations: on the Aviat Eclipse Packet Node 256QAM with 28MHz channel bandwidth; on the Ericsson MINI-LINK TN 1024QAM with 28Mhz channel bandwidth; on the Linkra WIDHOP 256QAM with 28MHz channel bandwidth; on the NEC iPASOLINK 400 512QAM with 28MHz channel bandwidth; on the NSN FlexiPacket Microwave 256QAM with 28MHz channel bandwidth; on the Siklu EtherHaul-1200 64QAM with 500MHz channel bandwidth.

When we emulated bad weather conditions, all vendors were stepping back to a most robust modulation which was 4QAM with 256 MHz channel capacity for Siklu and 4QAM with 28 MHz channel capacity for all other vendors.

On the Ericsson MINI-LINK TN and NSN FlexiPacket

Microwave combination we configured SAToP CESoIP according to RFC4553, and on the Ericsson MINI-LINK TN and NEC iPASOLINK 400 combination we configured CESoETH (Circuit Emulation Service over Ethernet) according to MEF8. The microwave systems successfully converted the E1 signal from and to Ethernet packets and transported it over microwave in the priority class.

Synchronous Ethernet over Microwave

One of the other capabilities that we looked at for Microwave vendors was transporting Synchronous Ethernet, a physical layer technology, over microwave links. In that respect, the challenge for a microwave vendor is to transform the Synchronous Ethernet clock received on an Ethernet interface to a format that could be transported across the air interface and then recover the clock signal on the other microwave base station. At this point the solution will convert the signal back to Synchronous Ethernet and provide it further in the clock chain.

As is seen in Figure 10 we defined one microwave as a clock master, starting the clock chain and then transporting the clock over its own microwave link, attaching to a second vendor solution using an Ethernet link and then repeating the story again. This provided a way to both verify the ability of the solution to act as a clock master, slave and transport the clock signal over two different air interfaces.

Five Vendors were tested in a Synchronous Ethernet Master/Slave scenario.

A PRC clock source was provided via a Symmeticom SSU 2000e, which also provided a reference for measurements on a Calnex Paragon-X and Spirent Anue 3500.

We waited for the master and slave clocks to synchronize monitoring the duration the clocks took in moving from free-running to synchronize state and then ran measurements for a minimum of two hours. All test measurements were required to pass the G.8262 SSU Option 1 or Option 2 Mask and have a PPB (parts per billion) of under 16 nanoseconds at a sample rate of 1 second.

Eight vendor pairings were successfully tested (see Figure 10) consisting of: Aviat EclipsePacket Node, Ericsson MINI-LINK LH, Ericsson MINI-LINK PT2010/SPO1460, Ericsson MINI-LINK PT2010, NEC iPASOLINK 400, NSN FlexiPacket Microwave, and Siklu EtherHaul-1200.

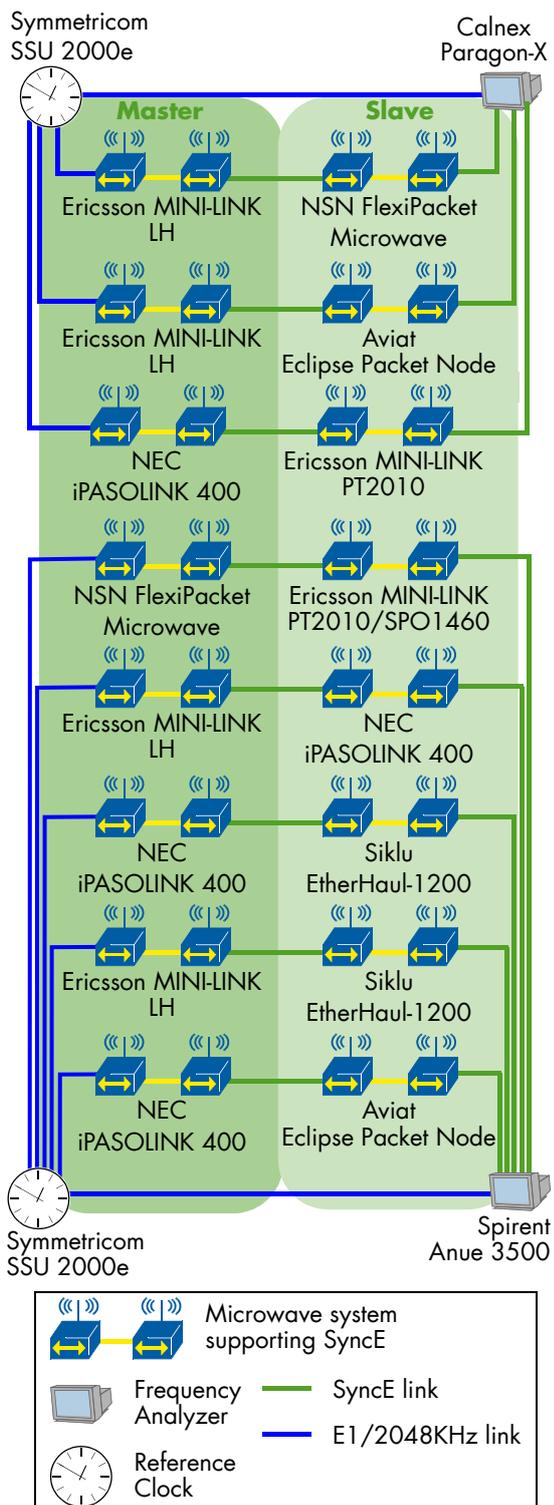


Figure 10: SyncE Test Results

DEMONSTRATION NETWORK

The demonstration network was built based on selected successful results and vendor demonstrations. Since we could not show all results in the same network topology the detailed test results are described in the test sections.

In the demonstration network we reflected the different network technologies tested in the event according to different by network areas. We tried to mimic a real provider network as possible. We created Ethernet Access, Aggregation Ethernet, MPLS-TP aggregation and IP/MPLS core network areas.

In the IP/MPLS area, Brocade MLXe-4, Ixia IxNetwork, and Spirent TestCenter configured IP/MPLS over 100Gbit/s Ethernet interfaces. Ixia and Spirent emulated multiple PE devices connecting them to a VPLS service that was built with the Brocade MLXe-4 router.

In the MPLS-TP area, Ericsson SPO1410, Hitachi AMN1710 and Ixia IxNetwork all established MPLS-TP LSPs and pseudowire services as well as successfully established BFD-CC and PSC sessions. The MPLS-TP services were transporting test traffic generated from the emulated devices.

In the aggregation area we showed two ERPS rings, synchronous Ethernet over microwave, transport of PTP sessions from a PTP Grandmaster towards a PTP Boundary Clock over the IP/MPLS core and Ethernet aggregation network areas. We also demonstrated redistribution of the Precision Time Protocol signal from a boundary clock towards ordinary clocks. Another demonstration included service activation for a service configured between ECI and Albis devices.

Based on QoS over microwave test results we created a chain of microwave systems that transport Circuit Emulation Services (CES), video signal, and Internet-like data in the aggregation area over 5 microwave hops.

We also demonstrating an encrypted EVPL transported over IP/MPLS network and performance monitoring for this service. The performance monitoring results of Albis and ECI devices demonstrate that the encryption of the EVPL did not negatively influence the service.

Acknowledgements

We would like to thank Stephen Murphy from the University of New Hampshire IOL for his support during the testing.

Editors. This report has been written and edited by Jambi Ganbar, Sergej Kaelberer, Stephen Murphy, Ronsard Pene, Carsten Rossenhoevel and Shima Sajjad.

Acronyms

Term	Definition
APS	Automatic Protection Switching
BC	Boundary Clock
BFD	Bidirectional Forwarding Detection
BFD-CC	Bidirectional Forwarding Detection - Continuity Check
BWP	Bandwidth Profiles
CBS	Committed Burst Size
CC	Continuity Check
CCM	Continuity Check Message
CFM	Connectivity Fault Management
CIR	Committed Burst Size
CoS	Class of Service
CPE	Customer Premise Equipment
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DNU	Do Not Use
EBS	Excess Burst Size
EIR	Excess Information Rate
ERPS	Ethernet Ring Protection Switching
ESMC	Ethernet Synchronization Messaging Channel
ETH-AIS	Ethernet Alarm Indication
ETH-LCK	Ethernet Locked Signal
EVC	Ethernet Virtual Connection
EVPL	Ethernet Virtual Private Line
FDB	Forwarding DataBase
G-ACh	Generic Associated Channel
GAL	G-ACh Label
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
LAN	Local Area Network
LSP	Label Switched Path
MA	Maintenance Association
ME	Maintenance Entity
MEP	Maintenance Entity Point
MIP	Maintenance Association Intermediate Point
MTIE	Maximum Time Interval Error
OAM	Operations, Administration and Management
OC	Ordinary Clock
PRC	Primary Reference Clock
PTP	Precision Time Protocol
R-APS	Ring Automatic Protection Switching

Term	Definition
R-APS(NR)	ring automatic protection switching no request
R-APS(SF)	ring automatic protection switching signal failure
RPL	Ring Protection Link
SEC	SONET/SDH Equipment Clock
SLA	Service Level Agreement
SSM	Synchronization Status Messages
SSU	Synchronization Supply Unit
SyncE	Synchronous Ethernet
TDEV	Time DEVIation
TDM	Time Division Multiplexing
TIE	Time Interval Error
TOD	Time of Day
VPLS	Virtual Private LAN Service
WTR	Wait to Restore timer

REFERENCES

"Ethernet ring protection switching", ITU-T G.8032 version 2

"MPLS Transport Profile Data Plane Architecture", RFC 5960

"A Framework for MPLS in Transport Networks", RFC 5921

"Proactive Connection Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile", draft-ietf-mpls-tp-cc-cv-rdi

"MPLS-TP Linear Protection", draft-ietf-mpls-tp-linear-protection, work in progress

"Precision Time Protocol (PTP)", IEEE 1588-2008

"Mobile Backhaul Implementation Agreement Phase 2", MEF technical specification, work in progress

"Timing and Synchronization Aspects in Packet Networks", ITU-T G.8261/Y.1361

"Precision Time Protocol Telecom Profile for frequency synchronization", ITU-T G.8265.1

"Distribution of timing information through packet networks", ITU-T G.8264

"Synchronization layer functions", ITU-T G.781

"OAM Functions and Mechanisms for Ethernet Based Networks", ITU-T Y.1731

"Transport Performance Metrics MIB", RFC 4150

"Service OAM Fault Management Implementation Agreement", MEF 30 technical specification

"Ethernet ring protection switching", ITU-T G.8032 version 2

"Ethernet service activation test methodology", ITU-T Y.1564

"Timing characteristics of synchronous Ethernet equipment slave clock (EEC)", ITU-T G.8262/Y.1362

"Distribution of timing through packet networks", ITU-T G.8264/Y.1364



EANTC AG
European Advanced Networking Test Center

Salzufer 14
10587 Berlin, Germany
Tel: +49 30 3180595-0
Fax: +49 30 3180595-10
info@eantc.de
<http://www.eantc.com>



Hosted by:
IIR Telecoms

29 Bressenden Place
London SW1E 5DR, UK
Tel: +44 20 7017 7483
Fax: +44 20 7017 7825
registration@iir-conferences.com
www.carrierethernetworld.com

This report is copyright © 2011 EANTC AG. While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein.

All brand names and logos mentioned here are registered trademarks of their respective companies in the United States and other countries.

20111004 v1