Carrier Ethernet World Congress 2012
Public Multi-Vendor Interoperability Event

White Paper

Editor's Note

Carsten Rossenhövel
Managing Director
EANTC

What are the *hot topics* for Carrier Ethernet business, wholesale and mobile backhaul services this year?

The agenda of the Carrier Ethernet World Congress names quite a few areas — but which of these are ready for testing?

In our interoperability tests, we aim to evaluate innovative packet transport technologies, physically available solutions, and areas that are most relevant for service provider deployments in the future. While a conference agenda can be formed as the union of these three goals, we have to find their intersection to facilitate testing.

That's why we reserved tests of software-defined networking and specifically OpenFlow for future events — it is too early for interoperability testing in the context of Carrier Ethernet services today. Cloud tests stayed out because standards-based innovation happens mostly in the data center these days, not in Carrier Ethernet networks. MPLS-TP focused on specific tests (see below) due to the limited number of actually available, interoperable implementations.

I am happy that this allowed us to focus on a couple of very relevant and factual test areas:

- Mobile backhaul remains a key driver for Carrier Ethernet network design and deployment. Indeed it is often *the* driver for the service provider to upgrade a packet network. We evaluated advances in packet-based clock synchronization for LTE-TDD scenarios including boundary clocks — with very positive outcome: Eight test cases were successfully evaluated in 17 test combinations.

  The clock synchronization tests included a new metric, called "Floor Packet Population" which enables service providers to measure their network's readiness to transport packet-based clocking easily — without the need to install the full backhaul solution in the first place.

  In addition, we were excited to test four transparent clocks as well as a clock recovery mechanism that uses a hybrid method to provide both frequency and phase signal.

- Two of our test areas are in line with the MEF's new campaign of "Carrier Ethernet 2.0."
  CE 2.0 is a group of MEF specifications and implementation agreements released in the recent past that the MEF packaged in a release. The benefit is that most of the underlying industry standards from IEEE, ITU and others have been around for a while so the industry is mostly ready for CE 2.0 today. This allowed us to implement elaborate scenarios.

  In this event, we successfully tested hierarchical Service OAM at multiple protocol levels — a topic where interoperability is specifically important, as service providers will almost always use equipment from multiple vendors in the access, aggregation and core network areas. We even included Management Information Base (MIB) objects in accordance with MEF 31 and 35 in our tests.

  As another aspect of CE 2.0, we tested the support for multiple classes of service ("multi-CoS") in our service activation test as well as in our microwave prioritization test. While Ethernet has supported CoS for ages, it may be difficult for service providers to implement it consistently across all equipment and manage CoS correctly over time. Service activation tests enable a factual check at provisioning time, for each new virtual circuit.

- Our multi-vendor resiliency testing campaign continued with more advanced scenarios as well. Dual-homing of Ethernet rings attached to an MPLS or MPLS-TP core was tested successfully: Sub-50 millisecond recovery times were achieved in optical link failure cases.

Table of Contents

The test areas and results show that Carrier Ethernet solutions are reliable these days, as expected; it is important for manufacturers, specifically in the areas of (microwave and wireline) aggregation and access, to continue innovate. Service providers aim to use Carrier Ethernet solutions to differentiate from their competition, and to provide more versatile and easier to manage services with stringent service-level agreements. To this end, the vendors participating in our showcase proved that service providers can rely on multi-vendor interoperability.

After two weeks of intense testing at EANTC's lab in Berlin, Germany, we are proud to present the results of the joint work of all participating vendors and the EANTC team in this white paper.

Interoperability Test Results

The following sections of the white paper describe the test areas and results of the interoperability event. 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.

In order to perform our tests we had to generate, measure, impair, and analyze traffic and perform synchronization analysis. We thank Ixia and Spirent Communications for their support.

Participants and Devices

Vendor	Devices
Albis Technologies	ACCEED 1104 ACCEED 1416 ACCEED 2202
Aviat Networks	Eclipse Packet Node
Ericsson	MINI-LINK TN MINI-LINK PT2010 MINI-LINK SP210 MINI-LINK SP310 SPO1410 SPO1460
HFR	CET-10216 CET-10448
Ixia	Anue 3500 Anue Hawaii ImpairNet IxNetwork
NEC	iPASOLINK 400
Omnitron	iConverter GM4
SIAE	ALCplus2e ALFOplus80
Spirent Communications	Spirent Paragon-X Spirent TestCenter
Symmetricom	TimeProvider 500 TimeProvider 1500 TimeProvider 5000 SSU 2000e

Managed Services

The topic of managed services has become a staple in our events. Here we cover test areas which should provide service providers the means to roll out services as well as monitor them efficiently. As such we looked at several new implementations of service activation testing performed on access devices as well as hierarchical service OAM and performance monitoring.

Service Activation

In last year's Carrier Ethernet World Congress Interoperability Event we tested, for the first time, a set of methodologies to verify that a service has been configured and is performing correctly. These are grouped under a standard, defined by the ITU-T, called Y.1564 – "Ethernet service activation test methodology". In 2011 one test vendor was able to demonstrate an implementation of the standard. This year, several access device vendors were ready with an implementation. The standard describes in details the test methodologies and test steps for different type of services such as color aware and color blind services, for different service parameters like Committed Information Rate (CIR), Excess Information Rate (EIR), Committed Burst Size (CBS) and Excess Burst Size (EBS).

The test methodology also evaluates the service objective performance parameters defined in Service Level Agreement (SLA). These parameters include Information Rate (IR), Frame Loss Ratio (FLR), Frame Transfer Delay (FTD), Frame Delay variation (FDV) and Availability (AVAIL). With these parameters a service provider can verify that the service is correctly configured and is functioning in accordance to the SLA before the service is handed over to the customer. Should issues with the service arise at a later time, the service provider can at least prove that the service was working correctly when it was set up for the customer.

So in order to create a test case that is as realistic as possible, we asked the participating vendors to create two Ethernet Virtual Private Line (EVPL) services. We then treated the two EVPLs as new services that a service provider would like to activate. The first service, EVPL 1, was defined to operate in a color-aware mode. The second service, EVPL 2, was configured in color blind mode.

MEF 23 service type	Color label	Color identification	SLA parameters
EVPL 1 – Color aware service			
High	PCP = 5 PCP = 4	Green Yellow	CIR = 5 Mbit/s CBS = 12,176 B EIR = 0 Mbit/s EBS = 0 B
Medium	PCP = 3 PCP = 2	Green Yellow	CIR = 10 Mbit/s CBS = 12,176 B EIR = 10 Mbit/s EBS = 12,176 B
Low	PCP = 1 PCP = 0	Green Yellow	CIR = 0 Mbit/s CBS = 0 B EIR = 15 Mbit/s EBS = 12,176 B
EVPL 2 – Color blind service			
Medium	PCP = 3	N/A[a]	CIR = 15 Mbit/s CBS = 12,176 B EIR = 0 Mbit/s EBS = 0 B
Low	PCP = 1	N/A	CIR = 20 Mbit/s CBS= 12,176 B EIR = 30 Mbit/s EBS = 12,176 B

a. Not applicable

The test setup and parameters defined for this test case were complex, however, we expect that this level of complexity will become the norm when virtual services from various customers will be expected to share the same interface. The specific service configuration is depicted in the table above.

The test service frames were generated with different Priority Code Point (PCP) values for each service type individually. One of the interesting aspects of the methodology was that the test frames were generated either by the service activation device, or by the access device under the test. The bandwidth profiles for each service were applied at the UNI of the EVPL service.

For each service type, the SLA parameters – or Service Acceptance Criteria (SAC) as Y.1564 calls a similar construct – were coordinated with the test case participants. SAC parameters including Information Rate (IR), Frame Loss Ratio (FLR), Frame Transfer Delay (FTD), Frame Delay variation (FDV) and Availability (AVAIL) are a part of Y.1564 service activation test. These are the parameters against which the test is being performed.

We performed Y.1564 service configuration test foreach service individually. If the test passed this step, we let the performance test run for 15 minutes for both services. In the final step we included the impairment generator in the test setup and introduced 20 ms delay on the service frame. We ran the performance test for one more time expecting a negative service performance test results.

The following devices participated the test: Albis ACCEED 1416, Ixia IxNetwork, and Omnitron GM4.

In the first scenario, Ixia IxNetwork participated as the device external to the network performing the service activation test. The Omnitron GM4 served as the access device at the customer site. It policed the traffic and applied bandwidth profiles on the pre-marked test traffic. The Albis ACCEED 1404 forwarded the traffic to another Ixia IxNetwork port. In this scenario both services, color aware and color blind, were tested successfully. Since the measurements were done externally to the service (using a tester) we were also able to measure Committed Information Rate (CIR), Excess Information Rate (EIR), Committed Burst Size (CBS) and Excess Burst Size (EBS) as per Y.1564.

In the second scenario, we performed the service activation test and the performance test from Albis ACCEED 1416. In this scenario, ACCEED 1416 generated the test traffic, and policed the test traffic on the UNI-C (on a port toward the service provider domain). Omnitron GM4 looped back the traffic received from ACCEED 1416 by swapping the MAC addresses. In this scenario we tested the service activation test for the Information Rates (IR) and the color blind services.

The last scenario set the Omnitron GM4 access device in the role of the service activation and service performance tester. Albis ACCEED 1416 was setup as the other end of the service was looping back the traffic. One difference between the service activation test implementation on the Omnitron device was the location in which the policer was applied. The traffic received from Albis was policed by Omnitron GM4 on the Ingress port (on a port toward customer domain). The policers were set to drop yellow traffic that did not conform to the SLA. The EBS and CBS were set to 12,000 bytes (as opposed to 12,176 bytes) bytes since these parameters' granularity in the implementation is 1,000 bytes. We ran the service activation and service performance test successfully. The service activation test measured all but Frame Loss Ratio SAC parameters correctly. The Frame Loss Ratio was calculated based on the total amount of lost frames as opposed to our expectations which were only on dropped green-marked frames. In this scenario, we tested the Information Rates (IR) for color aware services. We did not test burst sizes nor did we verify the behavior of the solution when impairment was applied.

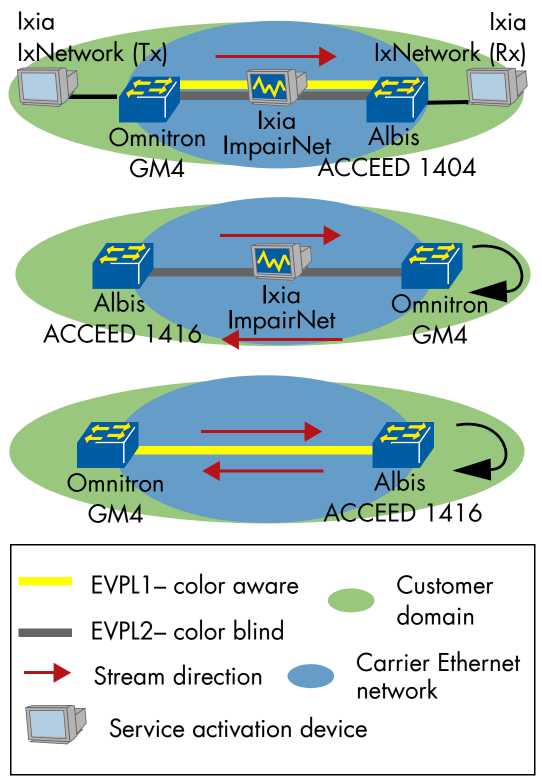

Figure 1: Service Activation Tests

We were impressed to see three new implementations of Y.1564 in the testing this year. Even though the standard is primarily written for specialized test instruments the ITU-T made specific provisions for implementations on access devices. From a service provider point of view implementations on access devices are likely to be more economical: they do not require additional devices to be sent to the customer site nor do they require an engineer to operate them. The implementations we have seen in the interoperability event will conceivably be operated from a network operations center as part of a service activation procedure. When such tests are available remotely, the cost of service roll-out could significantly be reduced – a goal that every service provider would appreciate.

Hierarchical Service OAM

It is entirely conceivable for Carrier Ethernet networks to offer native Ethernet service end-to-end. Given the recent approval of the Metro Ethernet Forum's ENNI Abstract Test Suite (MEF 37) the idea of native end-to-end services could even be tested for compliancy. This will mean that service providers will need to construct some hierarchies in their Operations Administration and Maintenance (OAM) where information and objects are separated into different maintenance levels. This will allow service providers to manage each domain independently and to provide both the customers and other operators visibility into the health of their own hierarchy.

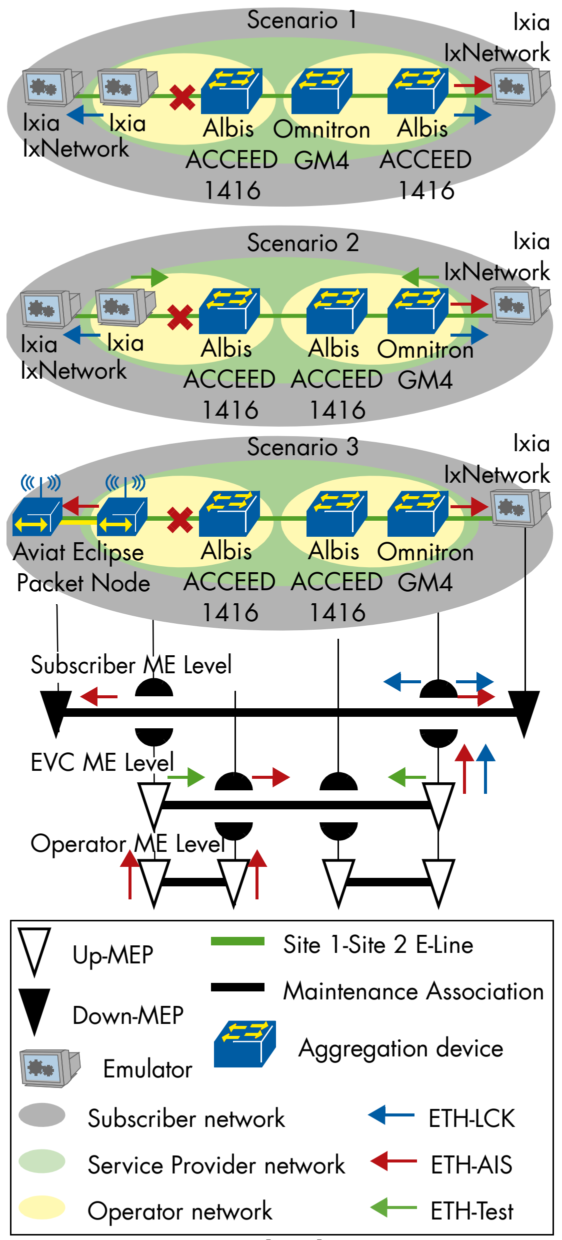

Figure 2: Hierarchical Service OAM

MEF 30 specifies the functional requirements for Fault Management in hierarchical domains. In our interoperability event we focused on Service OAM functionality between different vendor Maintenance Entity Groups (MEGs) at multiple maintenance levels.

The test was designed to distinguish between three types of OAM messages defined in MEF 30: Ethernet Alarm Indication Signal (ETH-AIS), Ethernet locked signal (ETH-LCK) and Ethernet Test Signal (ETH-Test). ETH-AIS is used to suppress alarms and to notify the failure of the service or transport path to remote Maintenance Entity Points (MEPs). ETH-LCK function is defined as a communicative mean. MEPs receiving ETH-LCK will differentiate between an administratively locked MEP and a defect condition. Peer MEPs upon receiving of ETH-LCK from a server MEP expecting interruption of data traffic. ETH-Test function is defined as an operation that runs one-way on-demand in-service or out-service and is used for diagnostics tests. This includes verifying bandwidth throughput, frame loss, bit errors, etc.

On the surface, the test scenario was straight forward from a testing perspective: build an Ethernet service in a multi-vendor environment that crosses two operator domains; remove a physical link in one of the operator network; verify that the MEPs configured at operator domain (level 2) propagate the ETH-AIS signal to service provider domain (level 4), and ultimately to subscriber domain (level 6).

To test the ETH-LCK signal we used the same topology to administratively lock a MEP at a service provider domain (level 4), and verify that the ETH-LCK signal is sent periodically to MEPs configured at subscriber domain (level 6).

The last message type, ETH-Test, was verified by generating ETH-Test frames at one of the MEPs configured at EVC level (level 4) – with specified rate, frame size and transmission patterns – and checking to see if the remote MEP received the test frames and performs the intended measurement.

As we performed the test, we created a single-tag (IEEE 802.1q) service at Customer UNI. The service was transported over the service provider network based on IEEE 802.1ad. A second tag was therefore added on top of the service tag at Service Provider's UNI. The outer tag was translated to the operator's network tagging scheme at ENNI. Down MEPs were configured at subscriber domain (MEF defined default MEG level to be 6). The devices sitting on the edge of the network configured two up MEPs: one at the service provider network (level 4), and one at the operator network (level 2). Continuity Check messages (CCMs) were exchanged between each pair of MEPs at a specific level. All maintenance associations were up and running.

As we configured the devices based on the test setup described above, we tested ETH-AIS and ETH-LCK successfully in scenario 1 in Figure 2. In the second scenario shown in the figure, scenario 2, we tested ETH-AIS, ETH-LCK and ETH-Test with success. In scenario 3, only ETH-AIS was tested.

As the link was removed in the operator network, devices next to the failure propagated AIS signal away from the failure. We observed that AIS signal was generated by MEPs configured at level 2, and received by MEPs configured at level 6.

As we tested ETH-LCK signal, we observed that upon receiving ETH-LCK signal on MEPs at level 6, the traffic was still being forwarded in one scenario, whereas in a second scenario, service frames were automatically blocked. Our expectation was that the Ethernet locked signal will interrupt data traffic forwarding towards the MEP.

As we conducted one-way, on demand, in-service ETH-Test in scenario 2, we observed that ETH-Test frames generated by MEP configured at level 4 were received by the remote peer MEP successfully. We observed that the remote MEP represented ETH-Test count frames with no interruption of data traffic.

Y.1731 explains that the tester should generate test signal with specified throughput, frame size and transmission pattern, and the remote MEP should perform the intended measurement. As the standard does not clearly define the mechanism of test signal, the vendors supporting ETH-Test signal represented only the ETH-Test count frames.

During the test, we noticed that the configuration of H-OAM was quite complex. We observed that MEPs configured at ENNI and UNI are capable of forwarding a single tagged or double tagged of OAM traffic. In order to have an end-to-end fault management module, it is required to know the capabilities of intermediate ENNIs capabilities.

Performance Monitoring

Performance monitoring tool is always interesting for service providers to monitor the performance of Ethernet Virtual Connection and to verify the Service Level Agreement (SLA). ITU-T Y.1731 provides the specification of such tool. The newer version of the standard is documented as G.8013/Y.1731 (2011) which we decided to focus in this event.

We distinguished between four measurement types for the test: two-way average frame delay and delay variation measurement, single-ended frame loss per EVC, and frame loss per EVC per CoS ID.

As a baseline reference test, we first validated performance monitoring implementations and Y.1731 protocol exchange per EVC without introducing any impairment. Once this test procedure was completed as a baseline, we added an artificial constant delay, delay variations, or packet loss using impairment tools. We expected that the delta between the impaired configuration and the reference test to be equivalent to the impairment tool settings.

In order to test average frame delay measurement, we added unidirectional constant delay of 20 milliseconds (ms) to all Ethernet Delay Measurement Message (DMM) packets, and expected that the average frame delay value would be increased by 20 ms. For average frame delay variation, we introduced 15 ms packet delay to every second DMM packet, and 25 ms delay to the remaining DMM packets. We expected the average frame delay variation would increase by 10 ms.

For loss measurement per EVC, we sent bidirectional Ethernet traffic over the network service and introduced 10% frame loss in one direction using the impairment tool. We verified whether the far-end and the near-end frame loss displayed on the device under test showed the same loss values.

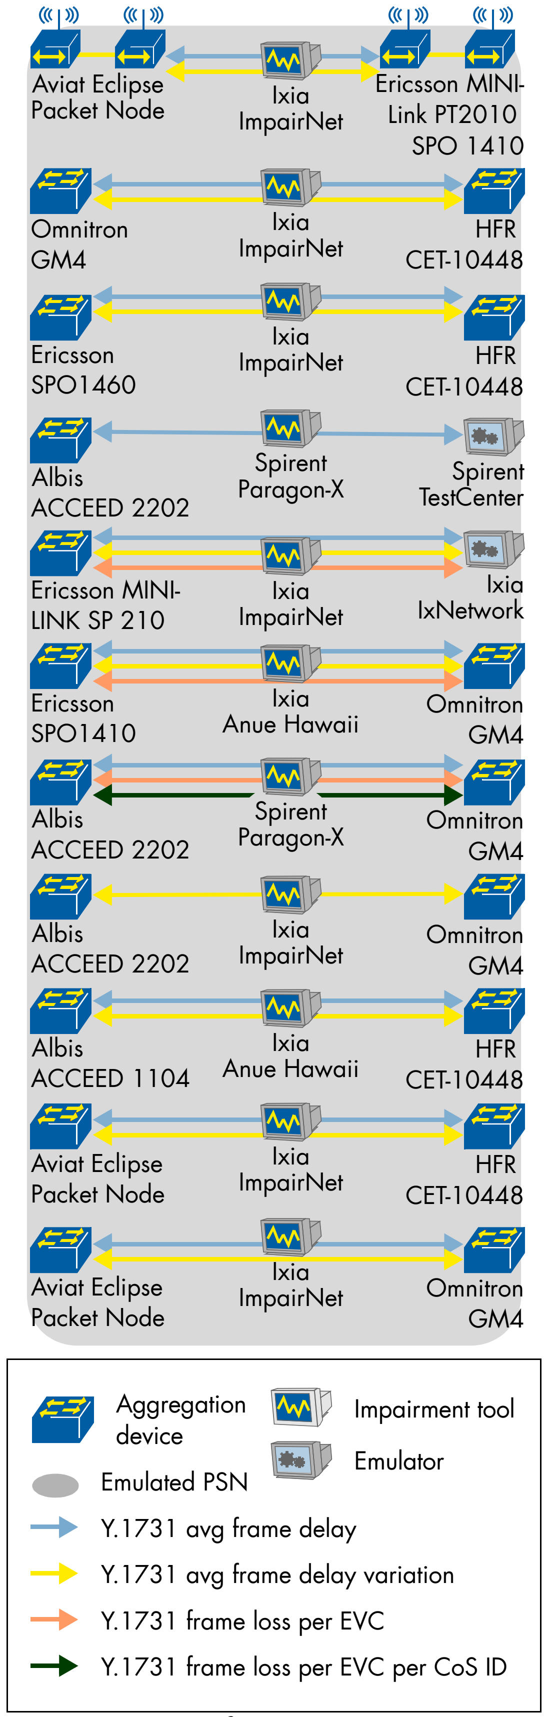

Figure 3: Performance Monitoring

As we performed per EVC per CoS ID frame loss measurement, we added a bidirectional constant frame loss of zero, and 10% per CoS ID 5 and CoS ID 3 respectively. For CoS ID 0, we added a constant frame loss of 5% in one direction, and frame loss of 15% in the opposite direction. We expected the devices to observe that the far-end and near-end frame loss stayed unchanged for CoS ID 5, increased by 10% for CoS ID 3, and increased by 5% and 15% – depends on the direction of impairment- for CoS ID 0.

All the pairs represented in Figure 3 participated and passed the frame delay and frame delay variation measurements successfully.

During testing the frame delay variation part, we observed two issues. One issue was that the impairment tool could not easily apply a delay profile of 15 ms and 25 ms to alternate packets. The impairment device could introduce a delay in the range of 15 ms and 25 ms. Due to this, it was impossible to calculate the exact expected value for the average frame delay variation. We instead observed a value in the range between 3.3 ms to 5 ms which was expected. We also found a buffer overflow bug for one implementation which was fixed in less than a day and re-tested successfully still during the event.

The following combinations were successfully tested for single-ended frame loss measurement per EVC: ACCEED 1104, Omnitron GM4; Ericsson SPO1410, Omnitron GM4; Ericsson MINI-LINK SP 210, Ixia IxNetwork.

ACCEED 1104 and Omnitron GM4 were successfully tested for frame loss measurement per EVC and per 3 CoS ID at the same time. For this pair, we observed that Loss Measurement Messages (LMM) were exchanged within a specific MEP pair for different CoS IDs. A derivative of the test was implemented and tested for Ericsson SPO1410 and Omnitron GM4. In this implementation, LMM messages were exchanged for different Class of services, but for different MEP pairs on different MD levels.

During the test, we observed MEG ID interoperability issue. ITU-T Carrier Code (ICC) defines no MD name for the MEG ID. MEF 30 follows ITU-T Y.1731 standard as an OAM tool, but the MEG ID specified in MEF 30 follows the 802.1ag string format. We observed that most vendors support both formats for MEG ID, but they needed to agree on the used format before setting up the test.

In all test scenarios, we used the Ixia Anue Hawaii, Ixia IxNetwork, Spirent Paragon-X to inject the required frame loss and delay onto the target OAM packet streams.

Carrier Ethernet Transport

This event's transport area was dominated by Ethernet Ring Protection Switching (ERPS) as well as microwave transport solutions. This is by no mean an indication that all other transport solutions have been tested and are showing interoperability. It is rather a result of the focus we gave the other test areas in this event. We see developments in the microwave solutions, now able to parse MPLS headers and prioritize traffic based on EXP bits, as well as new ERPS implementations. The next sections describe these tests and their results.

Microwave MPLS-TP QoS Support

Microwave links are a commonly-used substitute for long-haul fiber in the field, especially now that Adaptive Modulation can select the highest possible modulation scheme at a given point in time based on atmospheric conditions. Since changing conditions can decrease the overall throughput of the link, it is important to have a solution to ensure that important traffic is not dropped due when the air capacity is reduced. When the air-interface capacity is reduced it is paramount on the microwave solution to offer traffic prioritization capabilities. In this test we expected the microwave solutions to be aware of the MPLS-TP EXP bits, identifying packets marked as high-priority and prioritizing them.

MPLS-TP sessions were established between Ericsson MINI-LINK SP310 and either Ixia IxNetwork or Spirent TestCenter. Two EXP-marked classes were defined: one for high-priority traffic (Traffic Class 6), and one for best effort (Traffic Class 0). Test traffic was supplied by either Ixia IxNetwork or Spirent TestCenter.

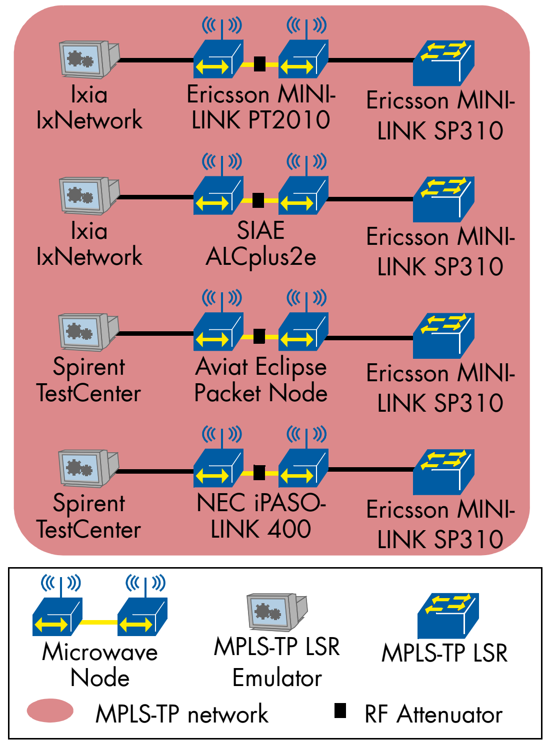

Figure 4: MPLS-TP QoS over Microwave

Physical Network Topology
Multi-Vendor Carrier Ethernet Interoperability Event 2012

EANTC

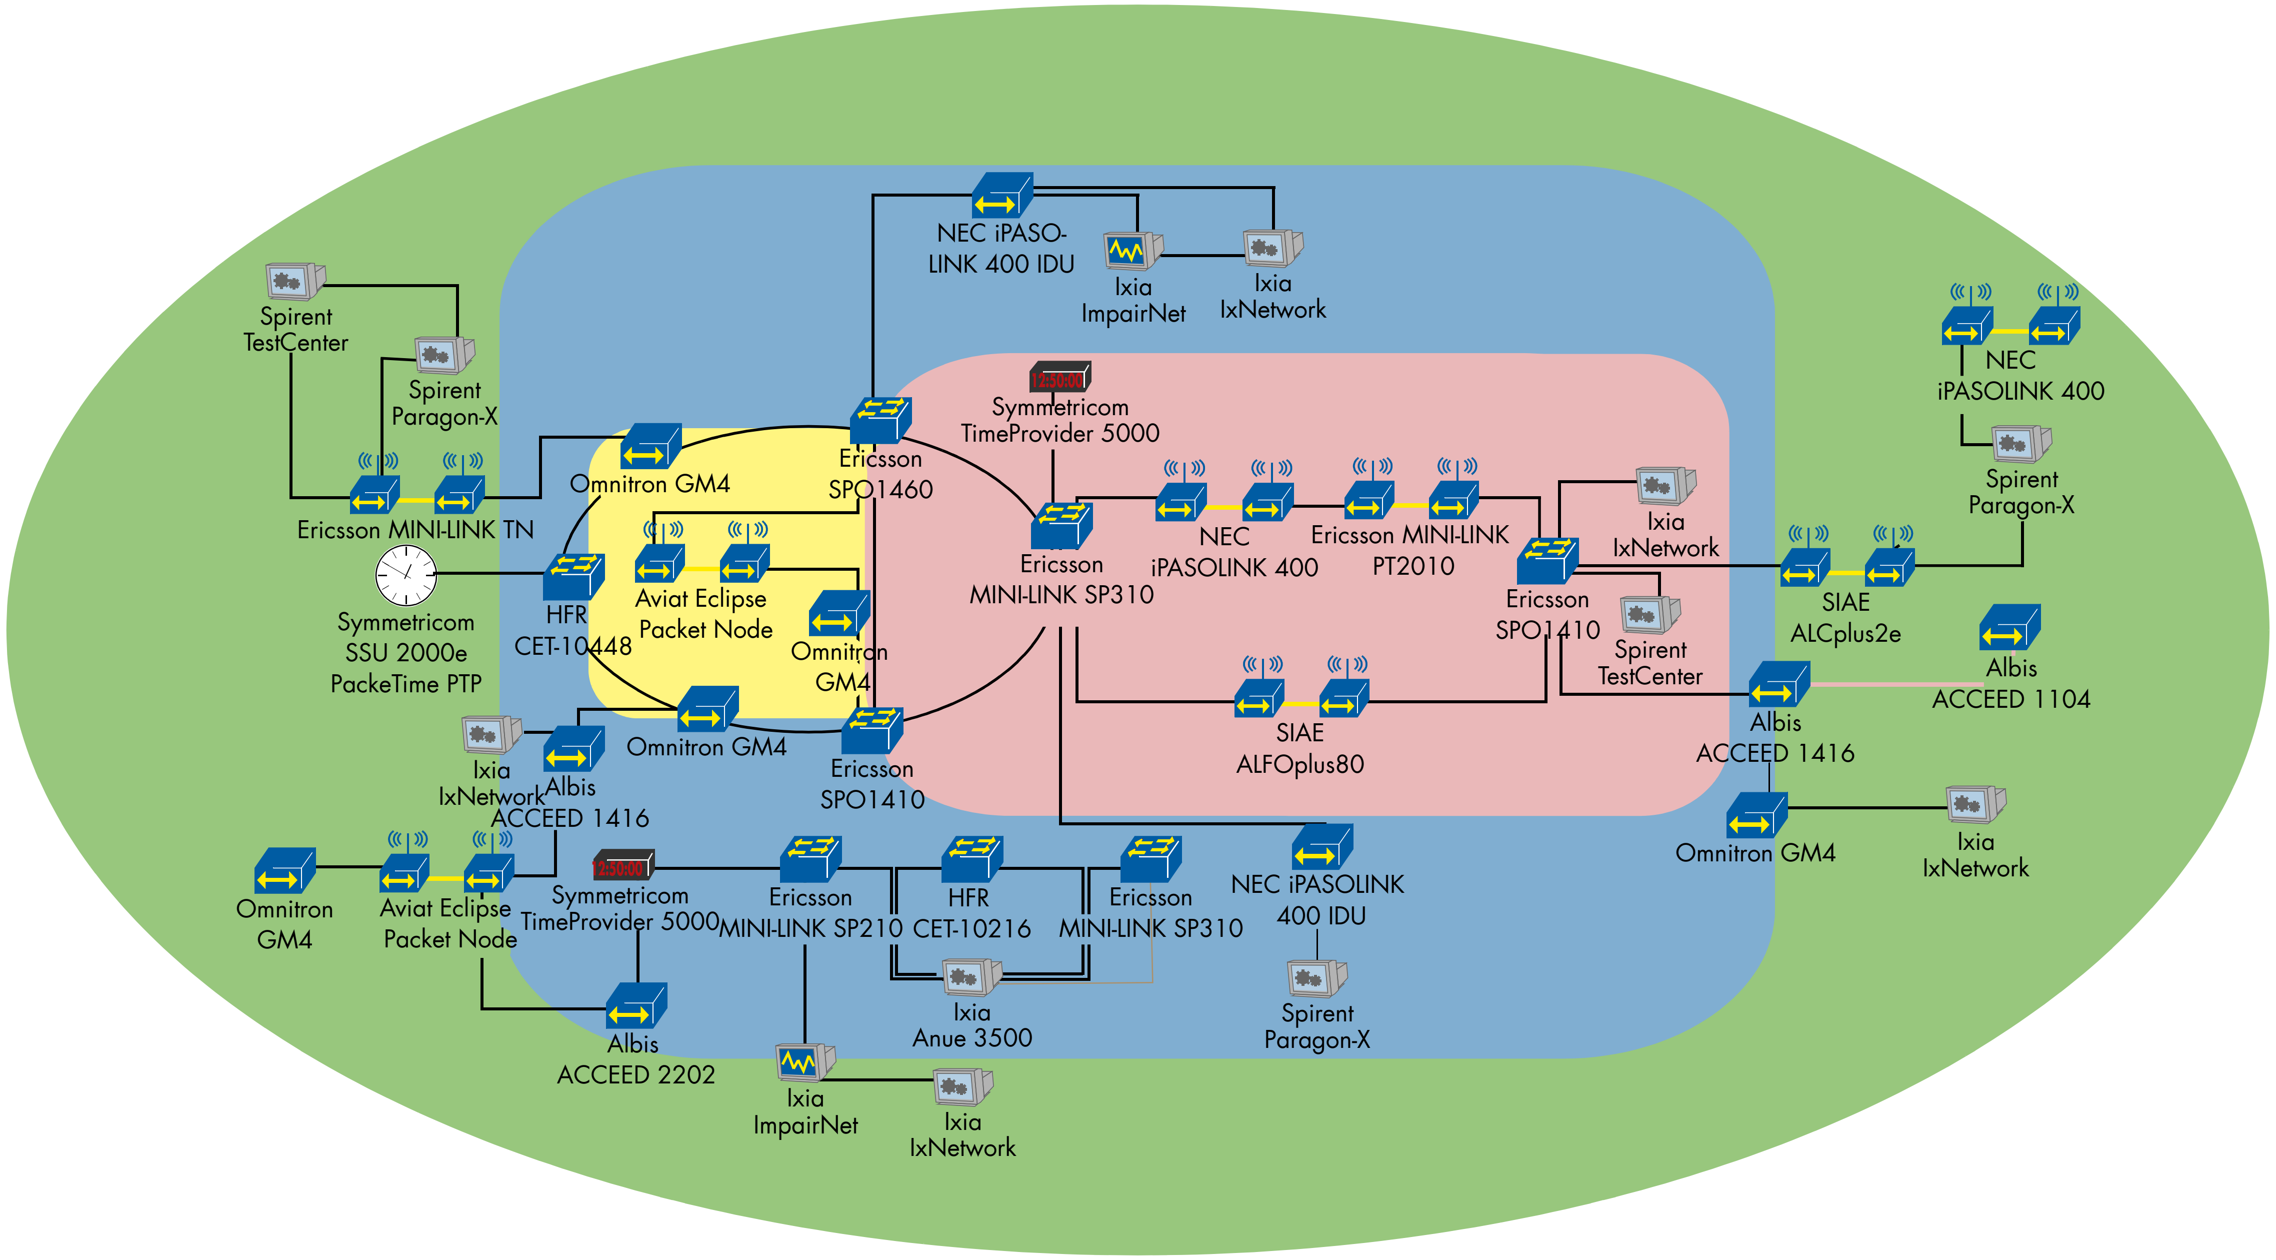

Device Types

- Access device
- Aggregation device
- Analyzer/Traffic Generator
- Emulator
- IEEE 1588v2 Grandmaster
- Impairment tool
- Reference Clock
- Microwave radio system

Connection Types

- Clock link
- Gigabit Ethernet link
- Bonded SHDSL link

Network Areas

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

Traffic was transmitted through the topology at 80% of the microwave links throughput at its maximum modulation. All test frames were 1500 bytes, and the traffic was divided up as follows: high-priority was transmitted at a rate equal to or slightly less than the throughput of the microwave link at its minimum modulation, and the rest was marked as best effort.

After we ran the 80% traffic at maximum modulation and verified that there was no frame loss, an RF attenuator was used to decrease the modulation, simulating inclement weather or other types of interference. We verified in all test scenarios that 0% of the high-priority traffic was dropped, and that 95% or more of the best-effort traffic was dropped, allowing for headroom in the predicted throughput of the minimum modulation. We also measured the maximum latency of the high-priority traffic in both modulations and observed a minimum of 122 microseconds and a maximum of 773 microseconds among the vendor spread.

The microwave devices tested in this scenario were: Aviat Eclipse Packet Node, Ericsson MINI-LINK PT2010, NEC iPASOLINK 400, and SIAE ALCplus2e.

ERPSv2 and VPLS Interworking

Rather than having a single point of failure in the attachment of an aggregation Ethernet ring to an MPLS or MPLS-TP core, it can be useful to dual-home an Ethernet ring at the attachment point, enabling end-to-end resiliency in the event of a single interworking device failure.

This test simulates that scenario by breaking a ring link connected to one of the interworking devices.

In both test scenarios that were run, an MPLS-TP VPLS was set up between three Ericsson devices: SPO1410, SPO1460, and MINI-LINK SP310. Two different ERPS sub-rings were connected to the VPLS, distinguished by the rate of CFM messages. The first ring comprised of two Aviat Eclipse Packet Nodes and an Omnitron GM4, configured with CCMs at a rate of 100 ms. The second sub-ring was made up of an HFR CET-10448, two Omnitron GM4s, and an SIAE ALCplus2e. CCMs on this second ring were transmitted at a rate of 3.33 ms. In the first ring, one Aviat Eclipse Packet Node was the RPL owner, and the second was the RPL neighbor. In the second ring, the Omnitron GM4 was RPL owner and the HFR CET-10448 served as the RPL neighbor.

Bidirectional traffic was supplied between the RPL owner and the end (non-interconnecting) VPLS node by an Ixia IxNetwork or Spirent TestCenter.

First, the primary path was verified to be fully functional by transmitting traffic through the ring to the far side of the VPLS by either an Ixia IxNetwork or Spirent TestCenter, ensuring that there was no frame loss. After that, a link in the primary traffic path was physically disconnected, causing the ring to switch to the protection path. Frame loss was again measured to ensure that there was no loss during the switchover, and while traffic was running on the protection path. Failover time was measured at 36 milliseconds in each direction on the first (100 ms CCMs) ring, and 25 milliseconds in one direction and 6 milliseconds in the other direction on the second (3.33 ms CCMs) ring. The switch over time in 100 ms CCM ring was less than 100 ms since the loss of signal (link failure) was used to trigger the fault condition instead of using impairment device on CCMs.

The link was re-connected, and after the Wait to Restore (WTR) timer expired, traffic loss and reversion time was again measured when the ring switched back to the primary path. Restoration time was measured at 47 milliseconds in one direction and 46 milliseconds in the other on the 100 ms ring, and 10 milliseconds and 1 millisecond on the 3.33 ms ring.

Afterwards, the backup path link was disconnected to ensure that no path switching occurred, as an additional confirmation that the ring indeed reverted to the primary path.

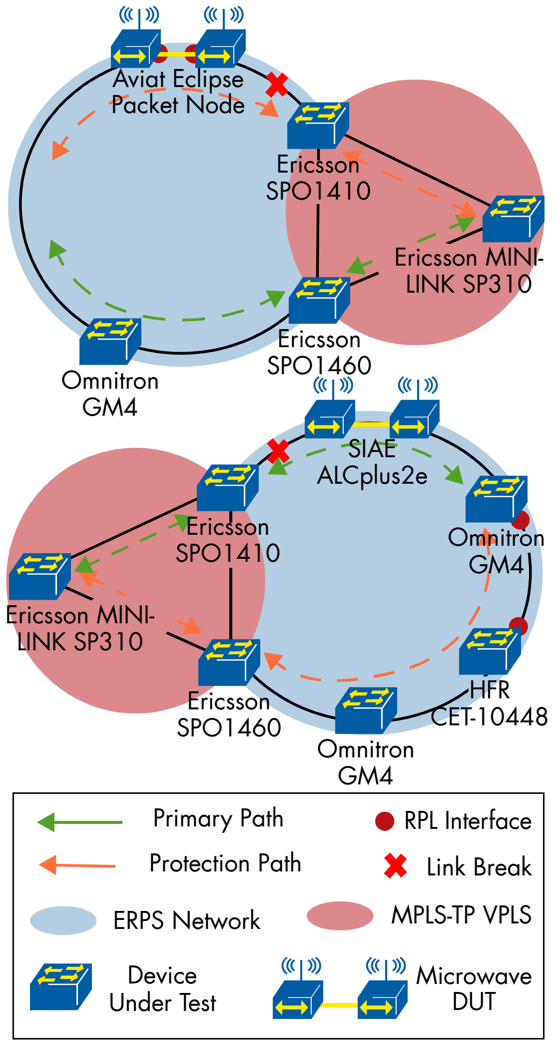

Figure 5: ERPSv2 and VPLS Interworking

MPLS-TP Fault Management

Two pairs of vendors signed up for this test. However, due to interoperability issues between participating vendors, stemming from mismatching implementations of the ACH TLV in the Generic Associated Channel Header, there were no successful results created.

MPLS-TP Pseudowire Redundancy

Due to issues encountered with interoperability scenarios in the MPLS-TP Fault Management tests, Ericsson ran a demo using a topology made entirely of their devices in the absence of compatible interop partners with enough available devices for testing.

The Ericsson MINI-LINK SP310 was set up as a UNI device. A primary MPLS-TP PW was configured to run through the SPO1410, and a backup one through the SPO1460, to another SPO1410.

Bidirectional test traffic was supplied by Ixia IxNetwork. After ensuring that it was flowing down the primary path, The link between Ericsson MINI-LINK SP310 and SPO1460 was disconnected and we measured that the service reconverged.

Clock Synchronization

What can be said about clock synchronization testing? The topic enjoys constant interest from new vendors as well as a never ending extensions from the standard bodies. This year one of our suggested test topic was for example clock synchronization for power systems applications, however, no implementations were ready to demonstrate support. Still, we were challenged by extending our testing program with new precision time protocol (PTP - IEEE 1588v2) measurement metrics, impairment profiles as well as new test cases for best master selection and hybrid-mode (PTP and synchronous Ethernet). We also performed synchronous Ethernet tests focusing on microwave transport. We used GPS as our clock source in our PTP and synchronous Ethernet tests.

Another challenge we encountered in this area was support for various configuration specified in the IEEE 1588v2 standard. These included message rates, transport mechanism (unicast or multicast) and 1-step or 2-step clock mode. The ITU-T G.8265.1 defines the telecom profile for frequency synchronization with PTP, which we intended to use in all of our tests. However, since not all vendors supported this profile (and the mechanisms for phase synchronization planned in G.8275.1 are yet to be ratified), we had to divert from it and include multicast-mode, resulting in a total of four PTP profiles.

The latest revision of the ITU-T G.8260 recommendation approved in February this year, includes a definition of "Floor Packet Population" metrics which address the measurement of PDV incurred on PTP packets throughout the network. These metrics evaluate how many synchronization packets reach the end point of a synchronization path close to (read: 150 µs) the observed floor delay with an observation duration of 200 seconds. Although PTP is agnostic to the constant delay throughout the synchronization path, it is sensitive to variable delay caused by various traffic patterns seen in production networks. The floor packet population metrics provides service providers with a mean to evaluate the readiness of an entire network path to support PTP. It enables services providers to measure their network delay and delay variation as a mean of knowing if the network could transport clock signal as opposed to measuring the clock output of the network which relies heavily on the ordinary clock output. Floor packet population is considered a neutral metric.

We included floor packet population measurements in our transparent clock, boundary clock and hybrid mode (PTP with synchronous Ethernet) test cases using the FPP (floor packet percentage) threshold described in G.8261.1 clause 8, which specifies that at the percentage packets meeting the target delay should exceed 1%.

IEEE 1588v2 Transparent Clocks

PTP enables frequency and phase synchronization between nodes over layer 2 or layer 3 transport. However, since each intermediate node introduces variable latency into the synchronization path, the traffic patterns may have adverse effects on accuracy and the slave clocks may require a longer duration to acquire a lock. One possibility to alleviate this condition is to deploy devices supporting the transparent clock function throughout the synchronization path. Each device supporting this function updates the correction field in PTP messages with an estimation of the latency it introduced. This in turn allows the slave clock to compensate for the latency and increase the precision of the synchronization.

In the first part of the test, we let the slave clock acquire the lock without any background traffic. After we observed successful results, we reset the slave clock and verified that it was in free running mode. We proceeded to introduce traffic between the transparent clock and the slave clock and allowed the slave clock to synchronize with background traffic. We used a traffic pattern based on the G.8261 VI2.2, with bursts lasting 100 ms and interburst gaps between 16 and 17 ms.

After the slave reported a synchronization lock, we measured frequency and phase against the G.823 SEC mask, as well as the floor delay packet population (discussed in ITU-T G.8260 I.5) using more than 1% FPP limit defined in G.8261.1.

An additional measurement was possible thanks to the ability of Ixia Anue 3500 and Spirent Paragon-X to compare the correction field values inserted by the transparent clock against the actual latency introduced. The testers either directly plotted the difference or subtracted the correction field from the measurements. One implementation inserted a constant value regardless of the latency. Other implementations accounted for the traffic, although we observed variations in the microsecond range. The irony of this measurement was that both approach are possibly correct. The transparent clock's job is to insert the delay through the clock's own implementation into the correction field. The task of the ordinary clock is then to subtract this value from the delay. Since both interpretations that we measured seemed to be correct, we did not actually use these measurements in the evaluation of the interoperability.

We also observed one implementation that did not recalculate the UDP checksum after updating the correction field. The vendor was able to provide an updated software image during the testing and we were able to test this fresh implementation successfully.

As in most of our clock tests Symmetricom was used to provide a clock source. In this case the Symmetricom TimeProvider 5000 was used as the grandmaster. Ericsson MINI-LINK SP210, HFR CET-10216 and Omnitron GM4 showed correct implementation of transparent clocks. The Ericsson MINI-LINK SP310 and NEC iPASOLINK 400 acted as slave clocks.

All maximum time interval error (MTIE) measurements passed the G.823 SEC mask. All measurements passed the ±1.5μs phase deviation requirements. All floor packet population measurements passed the 1% FPP threshold. We observed no loss for the test traffic.

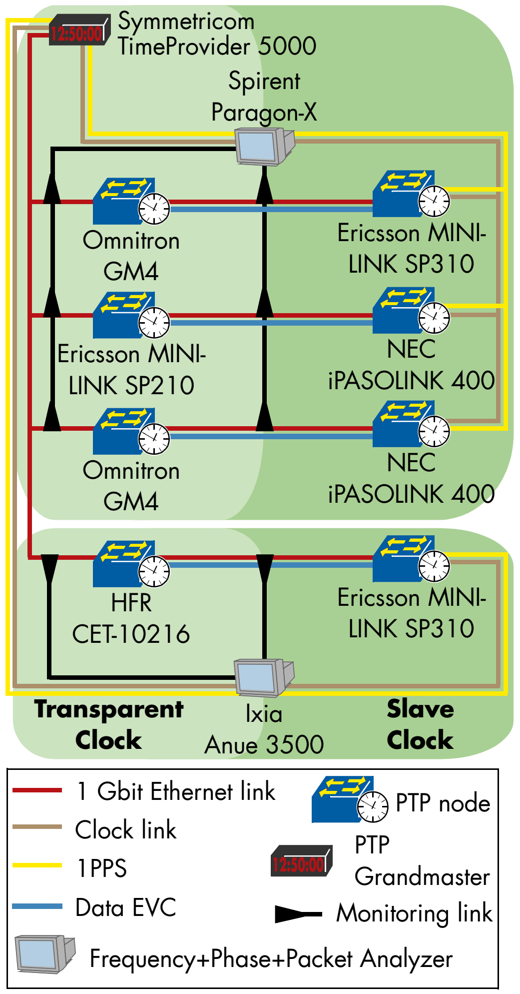

Figure 6: PTP Transparent Clock

Currently, the Internet Engineering Task Force (IETF) working group TICTOC is discussing methods for MPLS nodes to deal with PTP packets. This is certainly going to be an area of further investigation in future events.

IEEE 1588v2 Boundary Clocks

Boundary clocks provide means to distribute synchronization within a network in a scalable manner with the additional value of buffering the grandmaster from the slave clocks.

Two boundary clock implementations were included in this year's event – an increase from last year's single boundary clock test. As in last year's event, we included an impairment tool between the grandmaster and the boundary clock and another, which did not exist in last year's test setup, between the boundary clock and the slave clock. The test plan created by EANTC specified the use of G.8261 test case 13 as the impairment profile. The participating vendors, however, were interested in using the less challenging G.8261 test case 12. The impairment device in the test was simultaneously creating the impairment profile and measuring the clock output. This allowed us to verify in the first test runs that both devices locked within 1 hour, although the measurements showed that the clock output oscillated, failing the G.823 SEC mask, and stabilized only after approximately 7 hours. For a second test topology we found an incompatibility between the boundary clock, which only supported 2-step mode and other components which supported 1-step mode.

At the end, we did not obtain any successful results for this test.

Ethernet Synchronization Messaging Channel (ESMC)

ESMC provides traceability of the clock source for synchronous Ethernet and allows the receiver to lock onto the best available reference source. ESMC uses ITU-T G.781 clock source quality levels (QLs). In this test, we used 4 different clock QL levels, starting from the highest quality level QL-PRC (primary reference clock) through QL-SSU-A (synchronization supply unit, ITU-T G.812 type I or type V slave clock), QL-SEC (synchronous equipment clock) and finally QL-DNU (do not use).

In this test we concentrated on microwave devices that are designed to carry synchronous Ethernet transparently. This requires the two microwave units to maintain the semantics of a single device, recovering the clock from one port, conveying the frequency and traceability over the microwave link and utilizing it at the outputs of the peer microwave unit.

Our test topology was composed from three successive synchronous Ethernet nodes with the reference clock source attached to the Primary Reference Clock (PRC) as well as the clock slave (SSU) as depicted in the figure. We started the test with no reference clocks connected, which prompted the devices to send either QL-SEC or QL-DNU, according to the internal oscillator capability or administrative configuration of each device.

When a port is used as a clock recovery source, it should signal QL-DNU on that port unless the device has a better quality source available to it than the received quality level signal. After we connected the primary reference source, the devices propagated QL-PRC downstream and QL-DNU upstream. We then connected the secondary reference source and observed no change as expected.

We proceeded with disconnecting the primary reference clock and observed that the secondary reference clock was propagated by signalling QL-SSU-A upstream and QL-DNU downstream via ESMC, which triggered a switching event. We connected the primary reference clock once more and observed that the nodes reverted to the primary reference source by signalling QL-PRC downstream and QL-DNU upstream, again triggering a switching event.

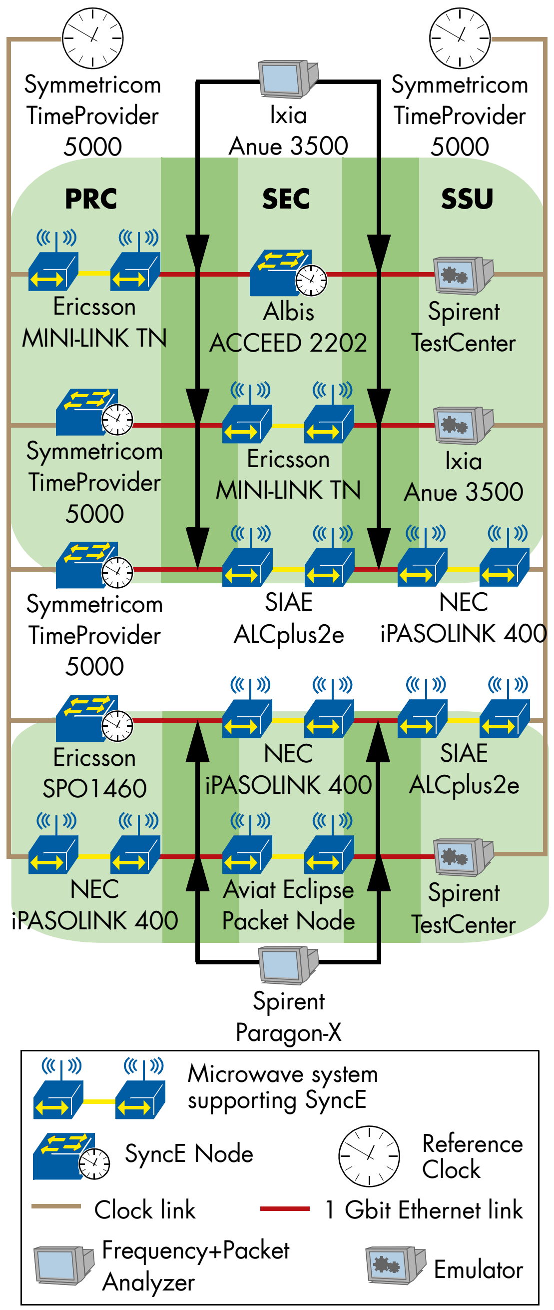

Figure 7: Ethernet Synchronization Messaging Channel

In both switching events, we verified our measurements against the G.8262 EEC Option 1 short term transient mask.Albis ACCEED 2202, Aviat Eclipse Packet Node, Ericsson SPO1460, Ericsson MINI-LINK TN, NEC iPASOLINK 400 and SIAE ALCplus2e passed the G.8262 EEC Option 1 short term transient mask.

Synchronous Ethernet over Link Aggregation

Both at network edges and at aggregation layers Link Aggregation is often used to increase link capacity at smaller increments than are available per physical interfaces capacity as well as a method for increase resiliency between nodes. If synchronous Ethernet is also used over the same path, its usage will impose new challenges in the event of a link fault – when the frequency offset between each link member is too wide, the clock signal may be impaired. Additionally, there is the possibility of a timing loop, when ESMC SSM messages on a LAG bundle do not correspond.

In order to verify that the implementations under test are able to deal with such conditions we established a 2-link LAG between the master and slave nodes. We allowed the devices' clock to settle and started measurements against the G.8262 EEC Option 1 mask to verify that the frequency is indeed stable. We then proceeded to emulate a failure on the primary link, followed by its restoration, after which we performed the same procedure for the other link. We waited at least one minute to ensure that Wait to Restore (WTR) timers expired. The failure was emulated using an impairment tool when Link Aggregation Control Protocol (LACP) was used or a physical cable disconnection when static LAG was available. We continued our measurements totalling at least 60 minutes and evaluated the results against the G.8262 EEC Option 1 mask.

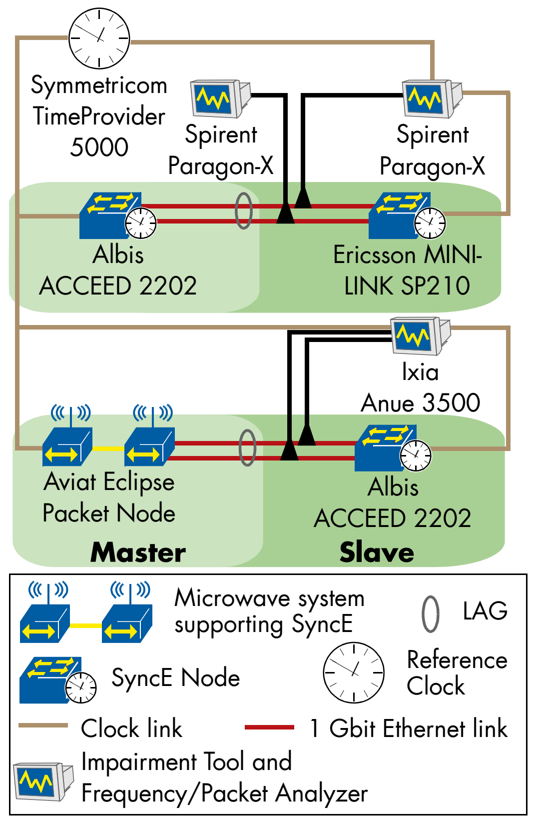

Figure 8: SyncE over LAG

In the final step we disconnected the reference source and verified that the slave node – now receiving QL-SEC instead of QL-PRC from the master node – sent QL-DNU on both LAG links, thus avoiding a timing loop.

Albis ACCEED 2202, Ericsson MINI-LINK SP210 and Aviat Networks Eclipse Packet Node passed the G.8262 EEC Option 1 mask. We did not observe timing loops in any of the test runs.

Precision Time Protocol and Synchronous Ethernet – Hybrid Mode

PTP provides both frequency and phase synchronization, whilst synchronous Ethernet provides only frequency synchronization. Frequency recovery with synchronous Ethernet is claimed to be more robust and stable, since it is independent of the traffic that traverses the devices and enables the recovery of clocking directly from the physical layer, which enables the acquisition of the frequency lock within a shorter time span.

When creating the test plan, together with the interested vendors, we received the request to show a mode which the vendors referred to as hybrid – frequency was to be received using synchronous Ethernet while Time of Day (ToD) was to be recovered from the Precision Time Protocol messages. Since the challenge was interesting and the method is potentially of interest to service providers that already have a synchronous Ethernet networks but are now required to support phase, we accepted the test.

Our setup included a grandmaster providing PTP and SyncE over separate links, a transparent clock and an ordinary clock. We used impairment tools to introduce packet delay variation (PDV) and wander. Frequency and phase analyzers were connected to the slave clock's outputs. Our goal was to verify that when packet delay variations (PDV) increases, the increase will only affect phase synchronization, while impairing frequency will affect both frequency and phase synchronization. Since potentially the frequency could also be gained from the PTP messages, we also made sure to check the influence of wander introduction on the recovered frequency signal.

First we verified the correct operation without any form of impairment, staying within the ±1.5 μs phase limit and passing the G.8262 EEC Option 1 mask for frequency. We then proceeded to impair PTP by introducing PDV and verifying that the frequency output of the slave clock was not affected and passed the G.8262 EEC Option 1 mask. Once the mask measurement was declared as pass, we stopped the PDV impairment and introduced wander with using a sinusoidal waveform. We observed that the slave clock's frequency output followed the wander we introduced as expected proving that the frequency was really recovered from the synchronous Ethernet signal. We also observed periodic oscillations in the 1PPS measurement, which confirmed that the frequency recovered from synchronous Ethernet is utilized to provide phase output. This was not for the faint hearted!

The good news was that the results of the test were positive. The complexity introduced to the test was really a product of our scepticism – after all we needed proof that the implementations were really combining both protocols to provide accurate phase and frequency. The following figure depicts the devices in each test run.

Symmetricom TimeProvider 5000 provided the PTP grandmaster functionality, Ericsson MINI-LINK SP310 and HFR CET-10216 acted as transparent clocks and slaves (in two test runs), while Ixia Anue 3500 and Spirent Paragon-X provided the impairment and measurements.

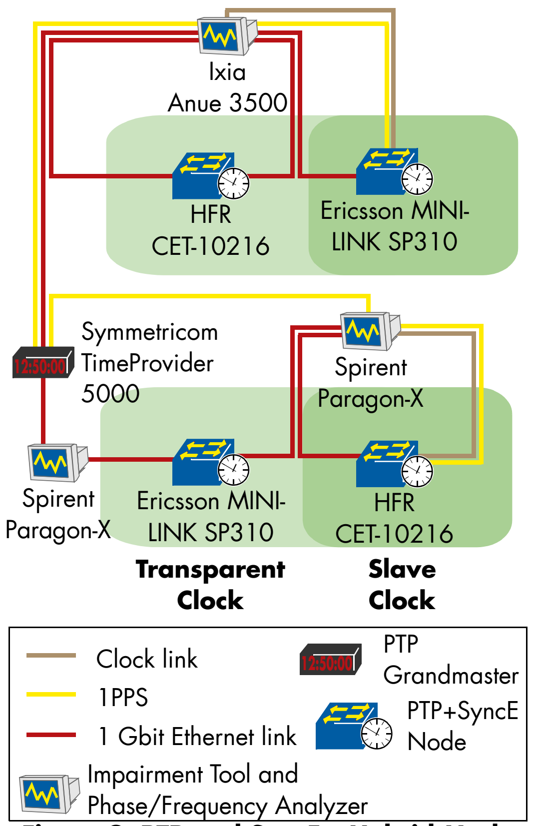

Figure 9: PTP and SyncE – Hybrid Mode

Precision Time Protocol and Synchronous Ethernet - Best Master Clock Selection

The Precision Time Protocol provides a simple mechanism for synchronization slave clocks to select the best master clock in their respective synchronization domain. The selection is performed according to the information announced by the master or by its availability. One of the parameters used in the selection algorithm is the master's clock class, which describes the reference source. When a master loses its reference source, its clock class attribute will degrade to reflect that it is in hold over mode.

The test comprised two grandmasters, identical in the configured priority values, a boundary clock and an ordinary clock and an impairment tool. We first verified that the boundary clock selects the grandmaster according to the tie-breaker in the form of the grandmaster identity. We performed the verification according to the clock class by either disconnecting the clock reference source or by emulating a clock class degradation with a tester, which caused the boundary clock to select the secondary grandmaster.

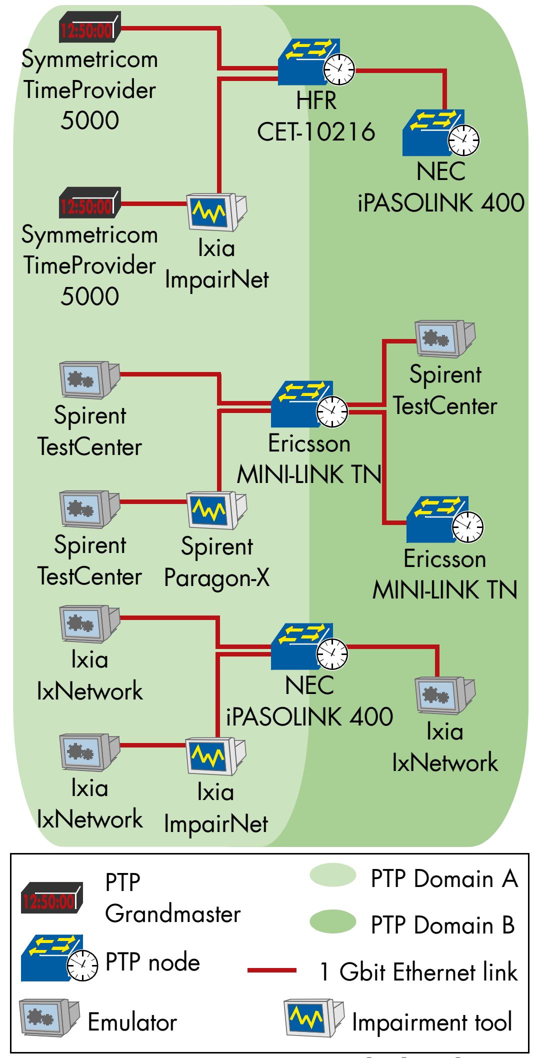

Figure 10: PTP Best Master Clock Selection

In the next step, we applied impairment on the secondary grandmaster by dropping all PTP messages, which prompted the boundary clock to revert back to the primary grandmaster. In each step we used the ordinary clock to verify that the boundary clock communicate its change of grandmaster.

Ixia IxNetwork, Spirent TestCenter and Symmetricom TimeProvider 5000 were PTP grandmasters, while Ixia ImpairNet and Spirent Paragon-X provider impairment. Ericsson MINI-LINK TN, HFR CET-10216 and NEC iPASOLINK 400 were boundary clocks. Ericsson MINI-LINK TN, Ixia IxNetwork, NEC iPASOLINK 400 and Spirent TestCenter were slave clocks.

Demonstration Network

As in every EANTC interoperability event, once testing was completed between vendor pairs (or sometimes, more than pairs), we constructed a demonstration network. The amount of the devices available in our lab enabled us to create an implementation of the ITU-T G.8261.1 Hypothetical Reference Model (HRM-2c). 6 vendors participated in this demo: Symmetricom SSU 2000e PackeTime PTP as the packet master clock. Ericsson MINI-LINK PT2010, Ericsson MINI-LINK SP310, Ericsson SPO1410, Ericsson SPO1460, HFR CET-10216, NEC iPASOLINK 400, Omnitron GM4 and SIAE ALCplus2e formed the network, with Spirent Paragon-X providing the measurements before the slave clock. Since the model includes microwave devices, devices that switch modulation schemes and thus bandwidth capacity according to weather conditions, additional prioritization to PTP packets could be supported by some of these microwave solutions to reduce packet delay variation (PDV) during modulation shifts.

In the public showcases (in Barcelona and Hong Kong), selected test results will be demonstrated to the public. We hope to see you there!

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, Shima Sajjad, Eldad Zack and Stephen Murphy.

Pulse Supply
909 Ridgebrook Road.,Sparks,Maryland 21152,USA
TEL : +1-410-583-1701 FAX : +1-410-583-1704
E-mail: sales@pulsesupply.com https://www.pulsesupply.com/siklu

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 © 2012 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.

20120910 v5