18/21 MAR. MARRIOTT PARIS

MPLS SDN WORLD CONGRESS 2014

Public Multi-Vendor Interoperability Event 2014

White Paper

EANTC

Editor's Note

Carsten Rossenhövel
Managing Director, EANTC

For the annual MPLS & SDN World Congress interoperability event, our job at EANTC is simply to identify the latest technologies relevant to service providers, to invite vendors and to test their solutions — *et voilà!*

At least that is the theory. Sometimes, however, we find ourselves a little too far ahead on the technology development cycle. This year, our grand scheme for the showcase was end-to-end multi-vendor cloud datacenter connectivity with elastic network services. This would have involved combining new OpenFlow and Open DayLight solutions with IP/MPLS advances such as Segment Routing, Ethernet VPNs, RSVP-TE signalled bi-directional LSPs, BGP flow specifications and BGP-LS. We were pretty excited — vendors market their advances in these relevant areas quite a bit.

It was not meant to happen, though. As we started hearing back from our vendor customers our excitement subdued. Even where seemingly mature IP/MPLS products were involved, large manufacturers told us the task at hand was simply too big and implementations of the above were just on the road map.

In fact, these days manufacturers' resources are stretched between legacy and new technologies. SDN developments are time-consuming and lots of brains are required for Network Functions Virtualization (NFV) and OpenDaylight activities as well. In this situation, MPLS and IPv6 are technologies where development may be slowed down and where interoperability can be de-prioritized. At EANTC's service provider proof of concept tests, we notice an increasing ratio of functional and performance software issues in related implementations.

Is this part of the new world? The term "Open" is pounded more than ever, but in fact service providers are getting locked in to "ecosystems", often deliberately, praising cost savings and time to market.

I find these approaches surprising, to be diplomatic. Experience has shown that new technologies are more successful if customers have many standards-based choices to select from.

And in the existing IP/MPLS transport market, vendors who continue development are more likely to win the inevitable network upgrades that are yet to come. Let's not forget that network resource and performance management, provisioning efficiency and high availability at scale could still be improved in many MPLS implementations. Plus, some large enterprise customer groups only now discovering MPLS…

The deployment of interoperable, well-supported and better manageable packet transport solutions will even increase in the mid-term, as mobile and fixed broadband services push new scale boundaries every day.

This is a wake-up call for network operators to decide:

- How long will MPLS remain in use in your networks? Do you even foresee a complete migration any time soon? How much should vendors go into maintenance mode from now on and focus their engineering on new markets?
- What should the new world of SDN & NFV look like? How should the market be balanced between *friends & family* eco-system solutions (faster to deploy with less pain, allow single RfPs) and standards-based, fully interoperable solutions (enabling best-of-breed, economic selection of products and less dependence on a single vendor's roadmap and support)?

These questions are open *right now*. There is a great opportunity for network operators to decide about future network designs in 2014. I hope that service providers will help vendors to steer their investments into the right directions.

Let's focus the on the **Test Areas Covered**. Packet clock synchronization in mobile backhaul networks is an area where a group of vendors has tirelessly worked together to improve interoperability continuously. This time, vendors were ready for very advanced and performance oriented tests. We provided a platform for vendors to test the interoperability of Long Term Evolution (LTE) phase clock quality under rather extreme conditions. The tests were very successful — please see details in this report.

We were also elated to hear that several vendors were ready to take the opportunity to show real world, interoperable SDN applications, controllers and orchestrators. The participants worked hard on getting demos working, resolving protocol interoperability issues on the way that service providers could expect to encounter in these early days of technology adaptation.

At the end of our two weeks hot staging phase in our lab in Berlin, Germany, we were proud to have a working multi-vendor SDN-network and to have completed all clock synchronization test cases we had planned for.

Test Equipment. With the help of participating test equipment vendors, we generated and measured traffic, emulated control and management protocols and performed clock sync analysis. We thank Ixia, Microsemi and Spirent Communications for test equipment and support. In addition, thanks to QualiSystems for providing their orchestrator solution to facilitate the SDN tests.

Table of Contents

Participants and Devices 3
Software-Defined Networking 3
Clock Synchronization 5
Topology.................................... 6
Demonstration Network 11

Participants and Devices

Vendor	Devices
Adva Optical Networking	FSP 150SP-100
Ericsson	MINI-LINK PT 2020 SP 110 SP 210 SP 310 SP 415 SP 420 SSR 8010
Huawei	SN-640 SOX (Smart OpenFlow Controller)
Ixia	Anue 3500 Anue Network Emulator ImpairNet IxNetwork RackSim
Metaswitch	SDN Controller
Microsemi	TimeProvider 2300 TimeProvider 2700 TimeProvider 5000
OE Solutions — AimValley	Chronos Smart SFP OAM Smart SFP TWAMP Smart SFP
Pica8	P-3922
QualiSystems	CloudShell
Spirent Communications	Spirent TestCenter

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 evaluated with equipment from a single vendor only. In any case, demonstrations were permitted only when the topic had been covered in the previously agreed test plan; when a test area had only one vendor, or when multi-vendor combinations failed, vendors performed demonstrations.

Software-Defined Networking

In recent years, the network industry increasingly focused its attention on Software Defined Networking (SDN). SDN is a paradigm shift – from distributed to a centralized control plane; from proprietary controller interfaces to standardized protocol defined between controller and network elements. The promise of SDN is to separate between network components that are responsible for packet forwarding from the component that are responsible for the network control, though enabling service providers more flexibility in choosing their suppliers and fast provisioning. Currently the main focus of SDN, from a protocol perspective, is OpenFlow – the interface between the control and data planes.

In our interoperability testing this year we focused on OpenFlow version 1.3. The following features and use cases were highlighted specifically:

- Rate Limiting
- Interworking between non-OpenFlow and OpenFlow devices
- 1:1 Protection

We were excited to also demonstrate the value of utilizing an SDN orchestrator to enable service delivery.

OpenFlow: Rate Limiting

Providing appropriate QoS to the user traffic is essential and part of managing Service Level Agreements (SLA). To this end the OpenFlow specification introduced the concept of meter tables and meter bands, which can be used to limit the transmission rate of an output port. While meter tables are used to store a collection of meter bands, meter bands specify transmission rate and actions to be performed when the specified rate is exceeded. The specification defines three meter band types:

DROP defines a simple rate that drops packets that exceed the band rate value.

DSCP REMARK defines a simple DIFFSERV policer that remark the drop precedence of the DSCP field in the IP header of the packets that exceed the band rate value.

EXPERIMENTER allows additional functionality in future OpenFlow message types.

In this test setup an OpenFlow (OF) controller was connected to an OF switch (OF Forwarder) across an IP network, over which the OF channel was established. Either Ixia IxNetwork or Spirent TestCenter was connected to the OF Forwarder, sending traffic to validate the data plane.

During the test we configured the OF controller with the band type DROP for Low and DSCP Remark for High traffic class as described in the following table.

We generated IP traffic for each traffic class at their corresponding bandwidth and verified that controller successfully installed the meter table into the OF switch table. We did not observe traffic drop.

Traffic Class	DSCP Value	Band Rate [Mbit/s]	Band Type
High	48	100	DSCP Remark
Low	0	250	Drop

We then doubled the traffic rate for the High traffic class and monitored that half of the traffic was remarked with DSCP value for the low traffic class (DSCP 0).

In order to make sure that the Low traffic class was also being metered, we increased its rate and observed that the Low traffic class was rate limited to the bandwidth we defined in the test plan.

Ixia IxNetwork and Huawei SOX successfully participated as OF Controller. Huawei SN-640 successfully participated as OF switch.

During this test we initially encountered an issue between two participating vendors. One implementation was looking for the meter band type OFPMBT_DROP with length 12 bytes, which according to the standard should be 16 bytes and hence was not installing the meter band value correctly. The vendor fixed the issue by updated the code version.

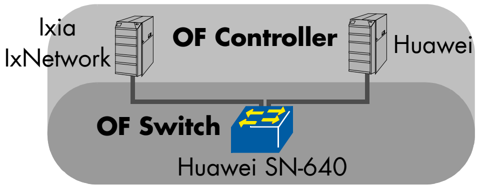

Figure 1: OpenFlow: Rate Limiting

Interoperability between OpenFlow and Non-OpenFlow Switches

It is likely that for the foreseeable future, OpenFlow based networks and traditional networks, will be operating side by side. As in our test a year ago, we wanted to provide both type of devices a platform to create a blue-print deployment scenario that could be shared. We had three roles defined in the test. As OpenFlow controllers Huawei SOX and Metaswitch both played their part. We had Huawei SN-640 and Pica8 P-3922 as switches in the OpenFlow environment. Ixia IxNetwork played the part of the non-OpenFlow switch. QualiSystems' CloudShell functioned as the orchestrator, sending service requests to the Metaswitch controller.

In all test combinations, we used a test setup consisting of two switches: an OpenFlow (OF) switch and non-OF switch. The OF switch was connected to the OF controller using an IP network, over which the OF channel was established. Likewise, OF switch and non-OF switch were connected over an IP network. To perform this test, participating vendors configured Resource Reservation Protocol with Traffic Engineering (RSVP-TE) to setup and tear down the Label Switch Path (LSP) between OF controller and non-OF switch. The vendors used Open Shortest Path First with Traffic Engineering (OSPF-TE) extensions as Interior Gateway Protocol (IGP) to build topology information about the network.

Once control plane sessions were established we sent traffic, between both ends of the network. We also made sure that the OF switches successfully installed flow entries to push and pop MPLS headers. We successfully validated OSPF-TE and RSVP-TE interoperability between non-OpenFlow and OpenFlow switches using the Metaswitch and Huawei OpenFlow controllers and Pica8 and Huawei OpenFlow switches.

We discovered an issue between two vendors participating in this test. Messages sent by the controller to the OF switch contain cookie and cookie mask fields. One vendor was not handling the cookie and cookie mask fields in strict conformance with the OpenFlow 1.3.1 specification, causing all rules programmed by the OF controller and sent to the OF switch to be deleted. Within the course of the two weeks hot staging phase, the vendor successfully updated the code and was able to perform the test.

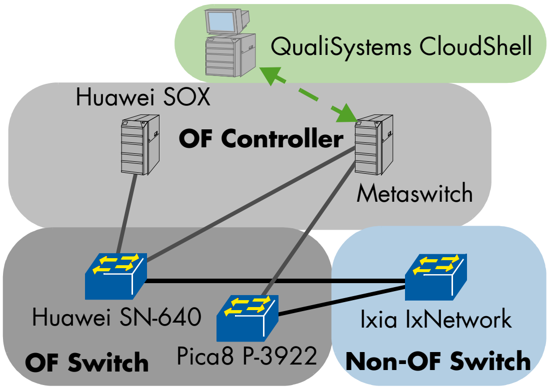

Figure 2: Interoperability between OpenFlow and Non-OpenFlow Switches

OpenFlow: 1:1 Protection

Resiliency mechanisms are crucial to the healthy operation of modern network. There are several different types of protection mechanisms that are commonly used such as 1:1, 1+1, 1:N and M:N.

The OpenFlow specification addresses resiliency by introducing Fast Failover group type in OpenFlow 1.2. The type allows fast failover since it does not require a round trip communication to the controller. In a 1:1 protection the OF controller installs two disjoint path in the OpenFlow network.

Our test setup consisted of four OF switches: three Huawei SN-640 and one Pica8 P-3922, each of which was connected to the Huawei SOX controller. The controller was configured to install two disjoints paths much like any other network: working path and protected path. While sending traffic at 1,000 frame/seconds using Spirent TestCenter we triggered a failover condition by pulling the link in the working path. We did not observed any impact on the traffic and the measured failover time was 0 milliseconds. After setting the network back to its original state, traffic reverted back to the working path without any impact.

Two-Way Active Measurement Using TWAMP

The quality, the performance, and the reliability of a network are essential to the user experience and customer satisfaction. Operators rely on performance measurement tools to monitor performance metrics such as packet delay, packet delay variation, packets loss and availability of their networks. In a multi-vendor environment, Layer 3 Operation Administration and Maintenance (OAM) solution based on Two-Way Active Monitoring Protocol (TWAMP) is one way to measure performance.

The Two-Way Measurement Protocol is specified in RFC 5357 and provide standard-based methods for measuring round-trip IP performance, such as packet loss, packet delay and packet delay variation between any two devices that support the standard. TWAMP uses the methodology and architecture of One-Way Active Measurement Protocol (OWAMP) defined in RFC 4656 to define a way to measure round-trip metrics.

TWAMP includes two protocols: the TWAMP control protocol and the TWAMP test protocol. The TWAMP control protocol is used to initiate, start and stop the TWAMP sessions, while the TWAMP test protocol is used to exchange TWAMP test packets between TWAMP endpoints.

The TWAMP standard also specifies a lighter version called TWAMP Light. In the TWAMP Light implementation, the role of Server, Control-Client and Session Sender are performed by the sending host and the role of the Session-Reflector is performed by the responding host thus eliminating the TWAMP control protocol. TWAMP Light provides a simple architecture for responders where their role will be to simply act as light test points in the network, thereby enabling the measurement of two-way IP performance from anywhere in the network.

In our event we focused on the TWAMP Light testing. In the test topology the TWAMP Light implementation consisted of two hosts: the controller and the session-reflector. The control-client, server and session-sender were setup on a laptop. The controller connected to the OE Solutions – AimValley Smart TWAMP SFP which was inserted into the Ericsson SP 110, acting as a session-reflector. The controller initiated the two-way measurement and the server accepted the incoming TWAMP test packets and reflected them back to the controller, which then performed the measurement.

We used Ixia IxNetwork between both controller and session-reflector to introduce impairment, such as packet loss, packet delay, packet delay variation, packet reordering and packet duplication. For each type of impairment we compared the two-way measurement results with the emulated impairment.

In all cases the measurement was correct.

Service Activation

When setting up and handing over Ethernet Service to the customer, Service Providers often require tools to check if the provisioned service complies with the Service Level Agreement (SLA). Service Activation helps service providers to verify and validate the correct configuration and performance of the service at the time of their deployment. There are major standards in this area: ITU-T Y.1564 "Ethernet Service Activation Test Methodology" and Carrier Ethernet Service Activation Testing (SAT), a work in progress from the Metro Ethernet Forum (MEF). The operation of the Service Activation can be facilitated by using test Protocol Data Unit (PDU). It provides the ability to configure and control the Service Activation Testing (SAT) steps and to fetch test result at the completion of the test without the need of a loopback topology, which may not appropriate when testing for configuration of Ingress Bandwidth Profile. SAT test PDU is being defined by the MEF in "Service Activation Testing Test and Control Protocol Data Units and Control Protocol" document and enables Network operator and Service Provider to perform SAT with interoperable device from diverse test equipment vendors. The specification defined Frame Loss PDU (FL-PDU) and Frame Delay PDU (FD-PDU) to support the service activation as approved in the Y.1564 recommendation.

SAT test PDU, analogue to the test methodology defined in ITUT-T Y.1564, is designed to test Ethernet-based service attributes, including bandwidth profile parameters: Committed Information Rate (CIR), Excess Information Rate (EIR), Committed Burst Size (CBS), Excess Burst Size (EBS), Color Mode (color-blind and color-aware) and Coupling Flag. It also covers performance attributes: Frame Loss, Frame Delay and Frame Delay Variation.

We conducted the test known as "Service Under Test" by configuring two Ethernet Services, EVPL 1 and EVPL2, that were being activated.

In our test we used an Ethernet Test Support System (ETSS), which was connected to the Media Converter electrical port. The ETSS commands were then transferred to the SFP port where the Control End (CE) OAM Smart SFP was inserted. Two OE Solutions – AimValley OAM Smart SFPs were employed, one acting as the CE, another as Responder End (RE). SAT control protocol messages between the CE and RE were used to configure each test, followed by SAT test frames generated between the two Smart SFPs. Multiple EVC services were tested in parallel, and each test was carried out as 2 uni-directional tests, allowing for asymmetric EVC parameters in each direction of traffic.

The ETSS running on Personal Computer (PC). We used Ixia Anue to impair the service in order to demonstrate violation of the Service Acceptance Criteria (SAC).

We ran the test in two phases: In the first phase we validated that each service under test is correctly configured. As soon as the first phase was successfully completed, participating vendor initiated the performance test from the ETSS for 10 minutes. The performance test evaluated the service against performance parameters know as Service Acceptance Criteria (SAC), which is a subset of a SLA. The SAC was agreed with the participant vendor prior to test execution. During the test the CE successfully retrieved the test results from the Responder End (RE) using control protocol, combined them and returned these to the ETSS.

OE Solutions – AimValley Smart SFP successfully participated in the test as CE and RE. No major issue was observed during this test.

Clock Synchronization

Over the past six years we have tested clock synchronization mechanisms from ways to transport TDM signal over packet (SAToP), through Synchronous Ethernet and IEEE 1588-2008.

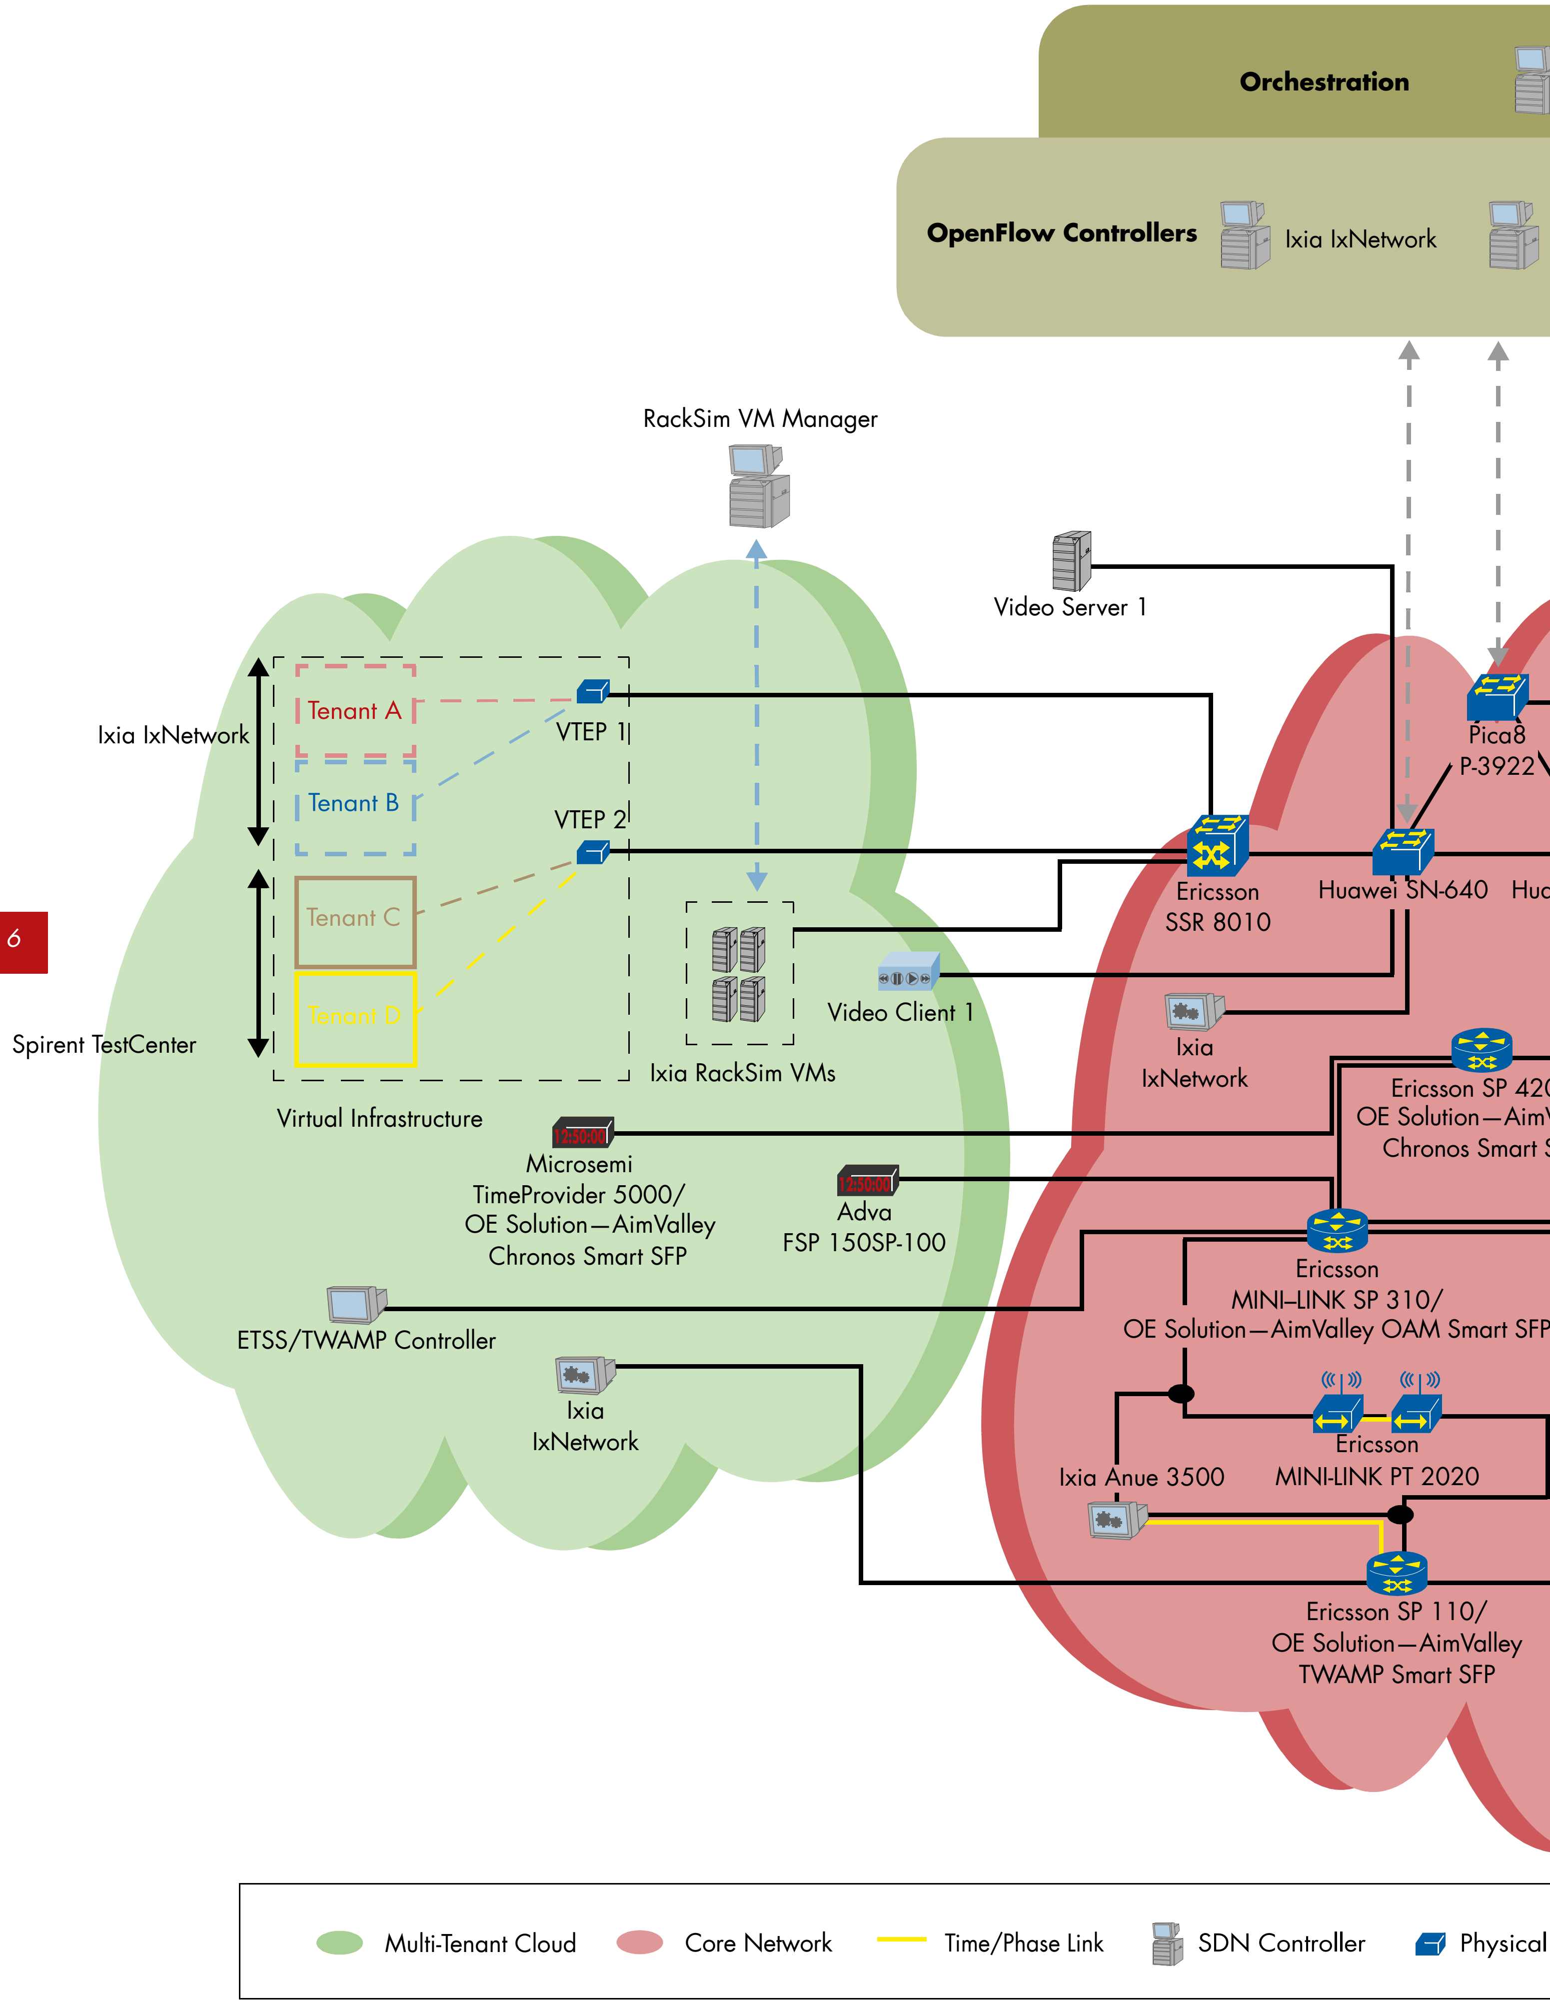
Orchestration
OpenFlow Controllers
Ixia IxNetwork
RackSim VM Manager
Video Server 1
Ixia IxNetwork
Tenant A
Tenant B
Tenant C
Tenant D
VTEP 1
VTEP 2
Spirent TestCenter
Virtual Infrastructure
Ixia RackSim VMs
Video Client 1
Pica8
P-3922
Ericsson
SSR 8010
Huawei SN-640
Ixia
IxNetwork
Microsemi
TimeProvider 5000/
OE Solution—AimValley
Chronos Smart SFP
Adva
FSP 150SP-100
Ericsson
MINI-LINK SP 310/
OE Solution—AimValley OAM Smart SFP
ETSS/TWAMP Controller
Ixia
IxNetwork
Ixia Anue 3500
Ericsson
MINI-LINK PT 2020
Ericsson SP 110/
OE Solution—AimValley
TWAMP Smart SFP
Multi-Tenant Cloud
Core Network
Time/Phase Link
SDN Controller
Physical

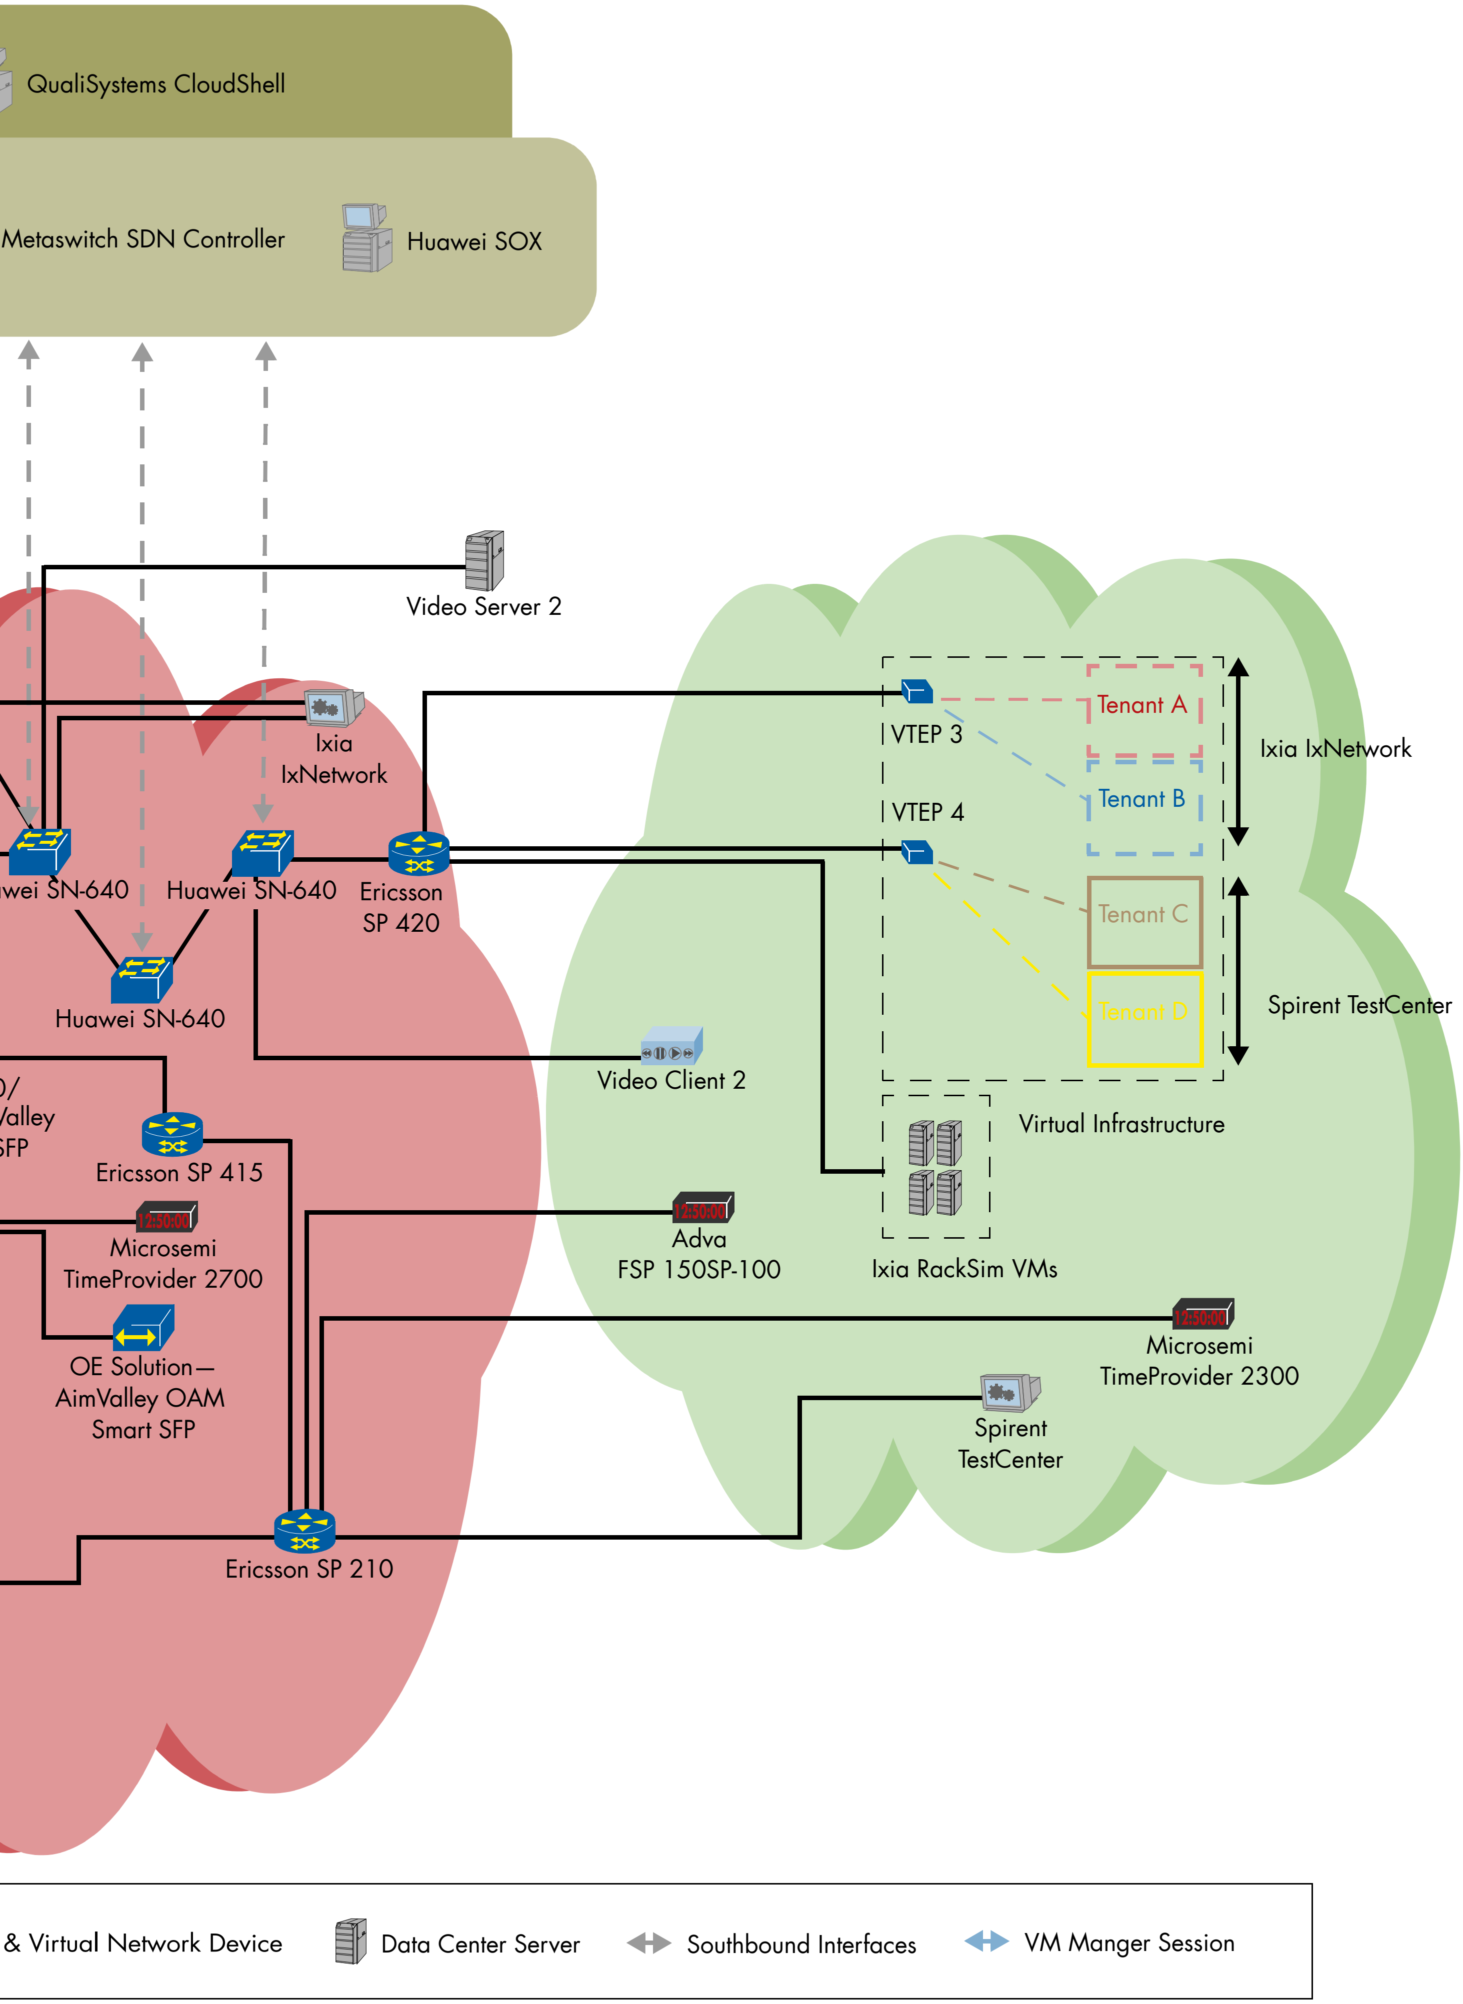
QualiSystems CloudShell
Metaswitch SDN Controller
Huawei SOX
Video Server 2
Ixia
IxNetwork
VTEP 3
VTEP 4
Tenant A
Tenant B
Tenant C
Tenant D
Ixia IxNetwork
Spirent TestCenter
wei SN-640
Huawei SN-640
Ericsson
SP 420
Huawei SN-640
Video Client 2
Virtual Infrastructure
0/
Valley
SFP
Ericsson SP 415
Microsemi
TimeProvider 2700
Adva
FSP 150SP-100
Ixia RackSim VMs
OE Solution—
AimValley OAM
Smart SFP
Microsemi
TimeProvider 2300
Spirent
TestCenter
Ericsson SP 210
& Virtual Network Device
Data Center Server
Southbound Interfaces
VM Manger Session

In this period, we worked with a group of vendors that not only return to our events again and again, but also help us in setting challenging goals.

The tests executed in this area focused on phase synchronization. Modern mobile networks require this technology for Time Division Duplex (TDD), enhanced inter-cell interference coordination (eICIC) and LTE Broadcast. These solutions promise higher bandwidth and spectral efficiency and wider service coverage, but require certain phase accuracy.

We borrowed the accuracy level of ±1.5 μs from the ITU-T recommendation G.8271 – accuracy level 4 as an initial starting point for the testing goals. We defined 0.4 μs as the phase budget for the air interface, which meant that the network phase accuracy level had to be ±1.1 μs. All of our tests in this event used this level of accuracy as a condition to passing a test. Measurement of phase was performed using a 1PPS interface or a Time of Day (ToD) and 1PPS composite interface. For frequency measurements, we used either an E1, a 10MHz or a SyncE signal. Frequency measurements were evaluated using the G.823 SEC requirements.

Since our tests imposed high accuracy requirements, in the order of single nanosecond, measurements of phase were given special care. We measured the length of the physical cables used to carry the 1PPS signal and accounted for the constant offset caused by the propagation speed of signals through the physical medium. This applies to the connection between the reference source and the analyzer, as well as the measured clocks and the analyzer.

The primary time reference clock was GPS using an L1 antenna located on the roof of EANTC's lab.

Precision Time Protocol as GPS Backup

The Global Positioning System (GPS) is an optimal choice for phase synchronization as it can deliver — under normal working conditions — a maximum absolute time error in the range of ±0.1μs. This allows deployment of accurate phase distribution, however GPS is subject to jamming, which could bring severe operational risks. Since GPS provides phase, frequency and time of day information, currently the only protocol that could serve as an alternative to delivering this information is the IEEE 1588-2008 or Precision Time Protocol (PTP).

The test started with both grandmaster and slave clocks locked onto GPS and PTPv2 was active. We then impaired PTP by dropping all PTP messages and verified that no transients occur, indicating that GPS was the primary source, while also verifying that the slave clock detected PTP failure. We then re-enabled PTP packet flow and introduced packet delay variation (PDV) based on G.8261 Test Case 12 to simulate a network of 10 nodes without on-path support. Following, we took baseline phase and frequency measurements from the slave clock. Afterwards, we disconnected the GPS antenna, simulating an outage. We restarted the measurements and evaluated the results according to the phase requirement of ±1.1 μs and G.823 SEC MTIE mask.

The measured phase accuracy with GPS was less noisy than PTP without on-path support. We measured a maximum of 45 ns (0.045 μs) time error with GPS in our tests, well below our set measurement threshold. We still managed to measure a maximum time error of 1 μs with PTP — also within our set goals.

The diagram depicts the results that passed the phase accuracy requirement of ±1.1 μs and frequency accuracy requirements of G.823 SEC.

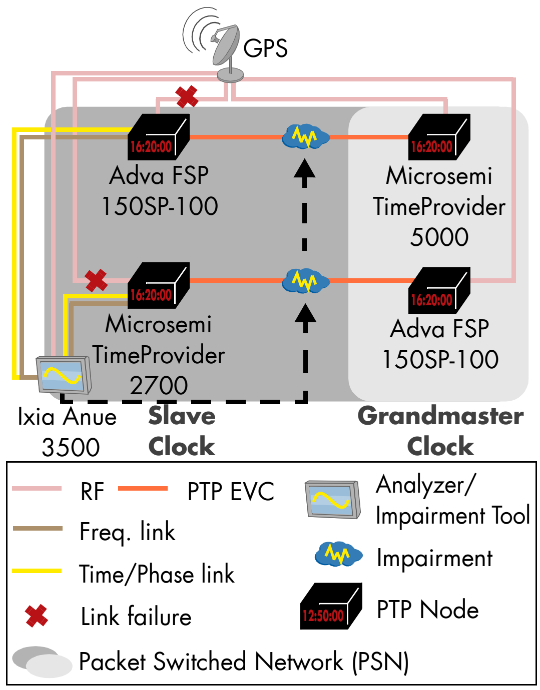

Figure 3: Precision Time Protocol as GPS Backup Results

Phase/Time Hold-Over Performance

Hold-over time is a crucial metric for mobile Service Providers. It is a major factor in the decision whether to send a field technician to a cell site to perform urgent maintenance or delay it for more cost-effective scheduling of operations. In case of a prolonged outage, a slave clock in a radio controller which exceeds its hold-over period will most likely result in major failure in hand-over from (and to) neighboring cell sites. Equipment vendors design their frequency hold-over oscillator performance accordingly. But what about time/phase hold-over performance?

We started the test with the slave clock in free-running mode and allowed it to lock onto the grand-master clock. We then performed baseline measurements. After passing the masks we set for the test, we used an impairment generator to drop all PTP packets, simulating a PTP outage. We then verified that the slave clock is in hold-over mode and started the measurements, letting them run over night.

We observed that with SyncE providing frequency reference while PTP is impaired, phase accuracy hold over was stable, exceeding 14 hours (the test was stopped at this point due to time considerations). In one test run we performed the test with no SyncE frequency reference and measured a hold-over time of approximately 3.5 hours while still providing a phase accuracy of up to ±1.1 μs.

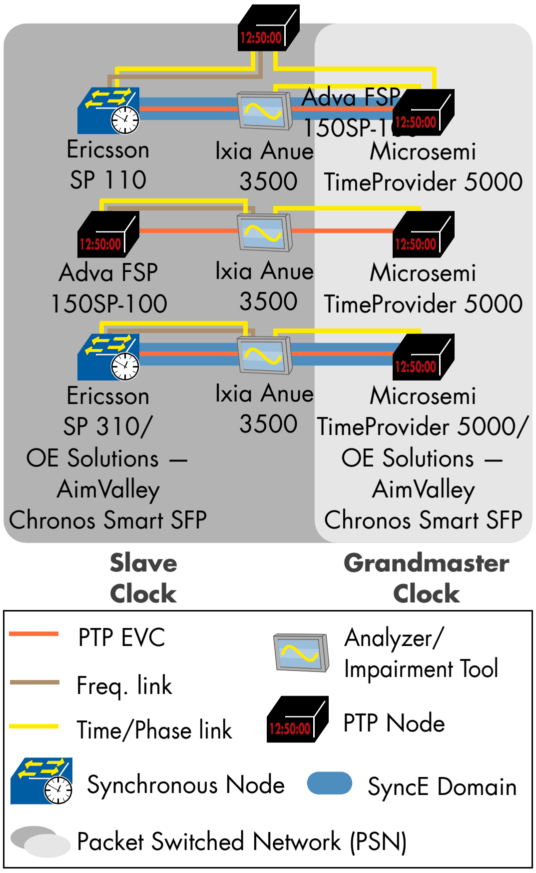

Figure 4: Phase/Time Hold-Over Performance Results

Adva FSP 150SP-100, Ericsson SP 110 and Ericsson SP 310 passed the phase accuracy requirement of ±1.1 μs and frequency requirements of G.823 SEC as slave clocks.

We executed one additional test where the optical link between the grandmaster clock and slave clock was replaced with Copper SFPs that support SyncE master/slave mode and provide symmetric delay for PTP. This test run also passed the frequency and phase requirements using OE Solutions — AimValley Chronos Smart SFPs. The diagram depicts the tested combinations we successfully executed.

Precision Time Protocol: Boundary Clock Noise Generation (Time/Phase)

When considering the phase budget to reach the required accuracy level, several factors come into play. One of them is the internal noise generated by each boundary clock, currently under study by the ITU-T in the upcoming recommendation called G.8273.2. There are two forms of noise — one is constant time error (cTE), estimated by averaging the measured time error, while the other is dynamic time error (dTE), estimated by calculating the MTIE of the phase measurements. We used the preliminary quality targets of 50 ns constant time error and 40 ns MTIE (over the whole period) as our goal.

We measured the time error of PTP packets at the ingress of the boundary clock on the packets originating from the grandmaster to estimate the inbound constant and dynamic noise. At the same time we measured the time error at the egress of the boundary clock. As an additional control, we also measured the physical phase output via 1PPS interface. In this test, we also measured the cable lengths and accounted for the physical medium latency for the PTP packets, to estimate the time error at the boundary clock itself.

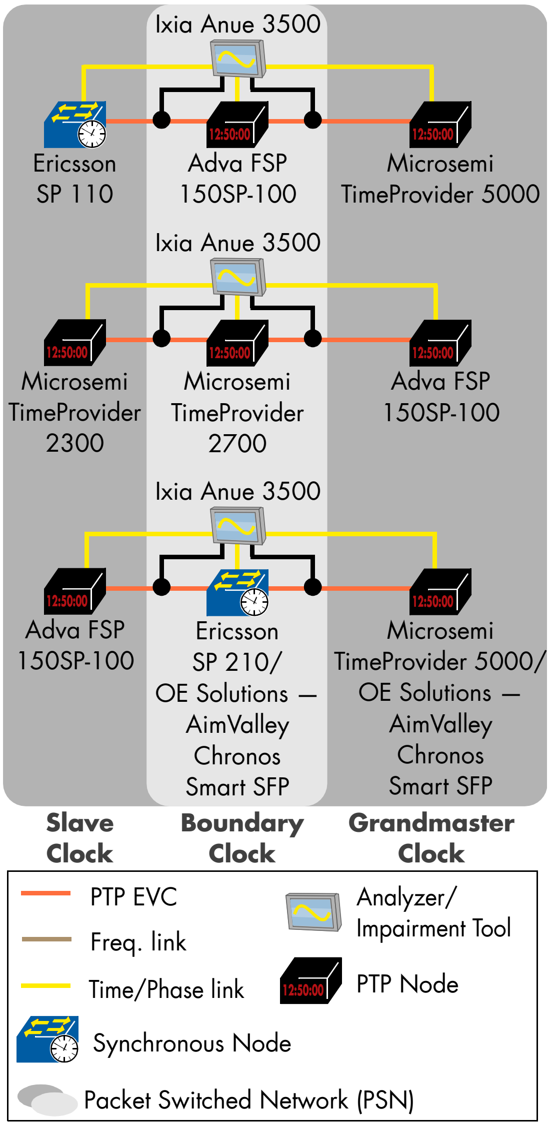

Figure 5: PTP — Boundary Clock Noise Generation (Time/Phase) Results

Adva FSP 150SP-100, Ericsson SP 210 and Microsemi TimeProvider 2700 passed as boundary clocks with the requirements of 50 ns constant time error (cTE) and 40 ns MTIE dynamic time error (dTE).

Precision Time Protocol over Adaptive Modulation Microwave System

In some deployment scenarios, such as rural areas access, a microwave transport is the most cost-effective solution for mobile backhaul. Microwave radios use Adaptive Coding and Modulation (ACM) to adapt the radio modulation to changing transmission conditions. Variation in modulation changes the link capacity, making it challenging to control packet delay variation (PDV) for packet clock protocols as well as guarantee its transport under severe weather conditions.

We designed this test case to verify that when the microwave link is 100% utilized, accuracy of the phase synchronization does not degrade — in normal and emulated severe weather conditions. To emulate such severe weather conditions, we used an attenuator to reduce the RF signal to the lowest modulation scheme available.

We started the test with the slave clock in free-running mode and generated traffic according to G.8261 VI2.2 at the maximum line rate for the maximum modulation scheme and expected no traffic loss. We took baseline measured for phase and frequency from the slave clock. After passing the requirements, we attenuated the signal to the lowest modulation scheme. Since the bandwidth decreased accordingly, we expected to observe verified that data packets were dropped according to the now available bandwidth. We restarted the measurements on the slave clock, evaluated them with the requirements and compared them to the baseline measurements.

We performed a single test run for this test with Microsemi TimeProvider 5000 as grandmaster clock; Ericsson SP 310 as the boundary clock; Ericsson MINI-LINK PT 2020 as the microwave system and transparent clock; Ericsson SP 210 as the slave clock. Measurements were taken using Ixia Anue 3500. We measured up to 18.4 ns time error in the highest modulation scheme (512QAM) and up to 19.2 ns time error in the lowest modulation scheme (4QAM).

Precision Time Protocol: Transparent Clock Scalability

As is the case for boundary and grandmaster clocks, an important characteristic of a transparent clock is the amount of clients it supports – but since a transparent clock does not require a context for each client, the governing factor is the total PTP message rate per second. We designed a test to verify that with maximum client utilization of the grandmaster, PTP accuracy quality remains within the requirements for phase. We performed this test with each client configured for a message rate of 64 packets per second (sync, delay request and delay response).

In this test we measured the dynamic error of the transparent clock, by comparing the correction field accuracy in the ingress and egress of the transparent clock. We also measured the phase output from the non-emulated slave clock.We started the test with one PTP client and performed baseline measurements. We then started the emulated clients, and repeated the same measurement, comparing the results.

We observed a maximum dynamic accuracy of the transparent clock up to 46 ns peak to peak with one client and 48 ns peak to peak with 500 clients. In both test runs, we observed an absolute maximum time error of 39 ns on the slave clock.

In one run, after establishing 500 clients through the transparent clocks, we observed occurrences of outliers up to 50 ns. All observed outliers were non-contiguous. We did not observe any outliers during the baseline measurements. No transients were observed on the slave clock output. The diagram depicts the test combinations we executed.

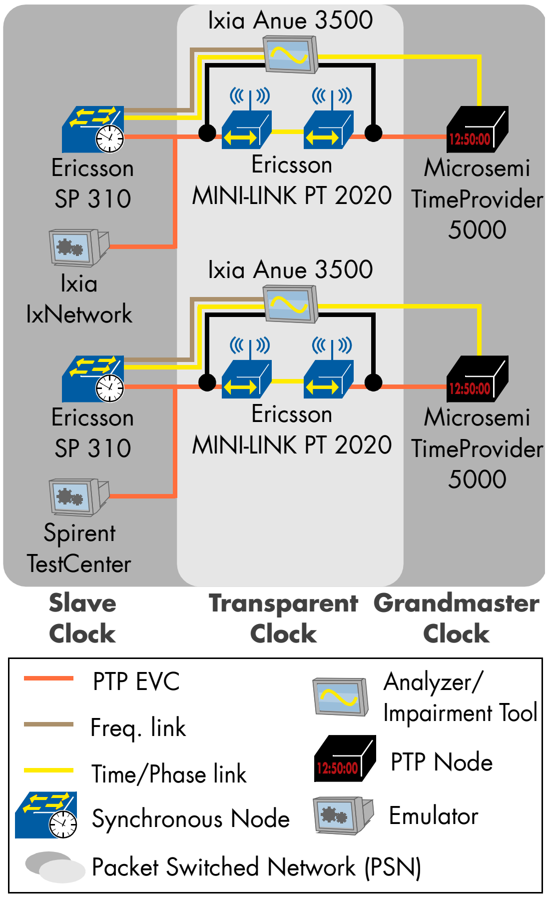

Figure 6: PTP — Transparent Clock Scalability Results

Precision Time Protocol: Master Clock Scalability

An important characteristic of a master-capable PTP node (either a grandmaster or a boundary clock) is the amount of clients it supports. We designed a test to verify that with the maximum client utilization, PTP accuracy quality meets the requirements for phase.

We started with one non-emulated slave clock in free-running mode and allowed it to lock to the master clock — either a boundary clock or a grandmaster clock. We then performed baseline measurements. After passing the requirements, we restarted the measurements and started the emulated clients.

The number of emulated clients was set according to the vendor's specifications of supported client to match the maximum together with the non-emulated slave clock. We verified that no transients occurred when we started the emulated clients. We then evaluated the results of the slave clock according to the phase and frequency requirements and also compared it with the baseline measurements.

We tested all devices with a message rate of 64 packets per second (sync, delay request and delay response) for each client. The following devices were tested for their PTP client scalability: Adva FSP 150SP-100 with 32 clients; Ericsson SP 110 with 8 multicast master ports; Ericsson SP 310 with 7 unicast master ports and a multicast slave port back to the Grandmaster; Microsemi TimeProvider 5000 with 500 clients. The results are depict on page 11.

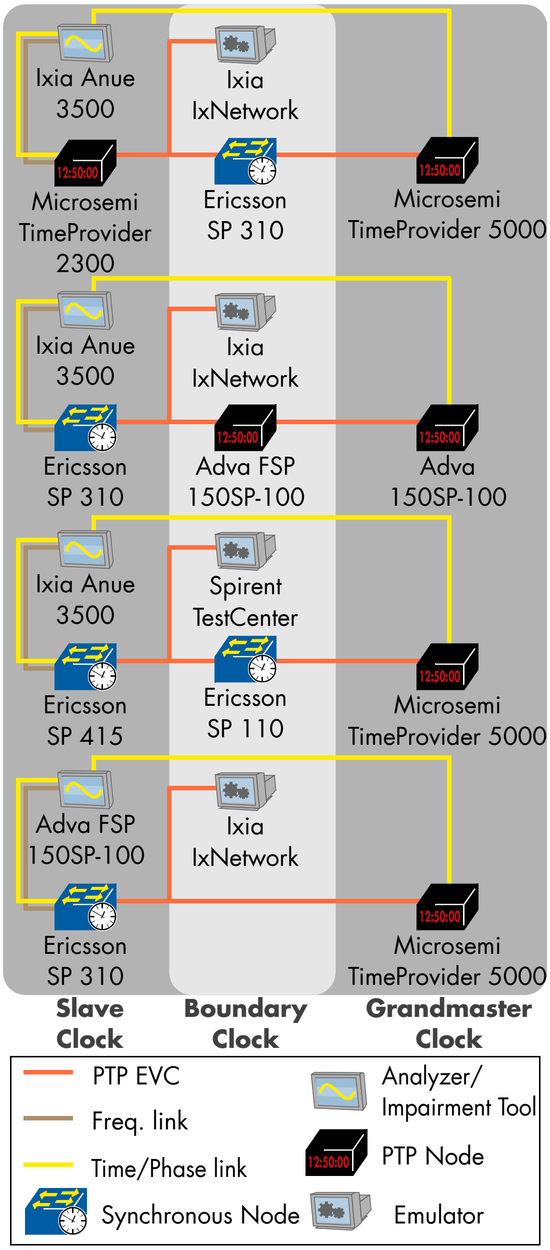

Figure 7: PTP — Master Clock Scalability Results

Demonstration Network

The first advanced use case aimed to demonstrate video content delivery over an SDN network with quality assurance. This is a great use case for a chain that a Service Provider is likely to see: a service orchestrator, OpenFlow controller (or two) and a network of OpenFlow switches as well as video servers and clients. QualiSystems CloudShell served as the orchestrator for providing the overall service life cycle of the demo – provisioning the controllers, the video hosts, and the Ixia IxNetwork tester.

Huawei, Metaswitch and QualiSystems agreed that the orchestrator would use a combination of Secure Shell (SSH) and REST APIs for provisioning the controllers. During the limited hot staging time, we verified the interaction between the QualiSystems CloudShell orchestrator and the Metaswitch OpenFlow controller. We successfully tested that the OpenFlow controllers were able to control the switches in the test network.

The second use case aimed to emulate data center workload mobility. In this scenario, we used Ixia's RackSim solution to create two data center sites with a large number of VMs in a multi-tenancy environment. Using a built-in VM Manager, Ixia emulated forward and reverse migration of virtual machines (VMs) between both data center sites across a core SDN network using Ericsson SP 420 and Ericsson SSR 8010 routers as data center gateways. Here we were challenged by the various transport network components, specifically the required support for IP localization. Therefore, we were not able to measure the out-of-service time during the VM migration process.

In the SDN area, we integrated two scenarios. The first was OpenFlow rate limiting with Huawei SOX as controller and Huawei SN-640 switches. We also showcased interoperability between OpenFlow and non-OpenFlow switches with Metaswitch SDN Controller as the OpenFlow controller, Pica8 P-3922 as OpenFlow switch, Ixia IxNetwork emulating a non-OpenFlow switch and QualiSystems CloudShell functioning as orchestrator.

In the transport area we demonstrated SAT measurement by placing the OE Solutions–AimValley OAM Smart SFP into Ericsson SP 310 and performing measurements. TWAMP measurements in the network are show-cased with an OE Solutions–AimValley TWAMP Smart SFP, positioned in the Ericsson SP 110, while the TWAMP sender client connects via Ericsson SP 310.

Furthermore, Ixia IxNetwork and Spirent TestCenter emulated VTEPs in each simulated data center. Both data centers were interconnected over VPWS circuit and Layer 3 VPNs. We sent bidirectional traffic between emulated IPv4/IPv6 and DHCP hosts across data center sites using VXLAN encapsulation.

In the clock synchronization area, we constructed a transport network with full on-path support, where every device is either an IEEE 1588-2008 boundary or transparent clock.

The devices for the transport part of the network were Ericsson MINI-LINK PT 2020, Ericsson SP 110, Ericsson SP 210, Ericsson SP 310, Ericsson SP 415 and Ericsson SP 420.

Adva FSP 150SP-100 and Microsemi TimeProvider 5000 were integrated as grandmaster clocks hosted in the data center portion of the network.

OE Solutions — AimValley Chronos Smart SFP were inserted into Microsemi TimeProvider 5000 and Ericsson SP 420, providing synchronization links over copper SFPs.

Microsemi TimeProvider 2700 acted as a slave clock within the transport network with GPS backup, while Microsemi TimeProvider 2300 acted as slave clock located in a data center. An additional Adva FSP 150SP-100 acted as a slave clock within the data center with GPS backup while performing time error measurements on the raw PTP stream.

Ixia Anue 3500 provided measurements for the transparent clock correction field accuracy and measuring the slave clock output from Ericsson SP 110, while Ixia IxNetwork and Spirent TestCenter emulated slave clocks.

EANTC AG European Advanced Networking Test Center	Upperside Conferences
Salzufer 14 10587 Berlin, Germany Tel: +49 30 3180595-0 Fax: +49 30 3180595-10 info@eantc.de http://www.eantc.com	54 rue du Faubourg Saint Antoine 75012 Paris - France Tel: +33 1 53 46 63 80 Fax: + 33 1 53 46 63 85 info@upperside.fr http://www.upperside.fr

20140310 v05