

Ericsson

Evolved IP Network Solution

Scalability, Performance, and High Availability Test Report

Introduction

Ericsson is one of the leading telecommunication equipment vendors in the world with 40 percent of the world's mobile traffic passing through Ericsson networks¹. With integrated mobile and fixed solutions, the company is in a unique position to offer true end-to-end networks to its customers. Ericsson's 'Evolved IP Network' solution targets converged and mobile operators, comprising in-house developed multiservice transport products and the Ericsson IP Operating System.

Ericsson contacted the European Advanced Networking Test Center (EANTC) to execute a test campaign in order to publicly validate Ericsson's 'Evolved IP Network' solution. Together Ericsson and EANTC agreed on the following high-level validation points:

- Test the complete IP transport portfolio with single operating system
- Demonstrate microwave integration for evolution to all-IP backhaul
- Validate Ericsson's performance monitoring to the base station

This report describes the network Ericsson created for the test, the test cases themselves and the results.



Figure 1: The Ericsson SSR 8020 Test Setup

1. http://www.ericsson.com/thecompany/company_facts

Test Highlights

- Tested end-to-end scope from base station to internet
- Measured less than 5.3 ms service outage time using IP Fast Reroute
- Confirmed solution scalability at large operator network scale
- Measured phase accuracy within +/- 325.56 ns for full path timing
- Verified suitability for mobile and multiservice deployments

Tested Solution

Ericsson provided a comprehensive and complete network environment for the tests - from multi-standard indoor macro base station (Ericsson RBS 6202) to the IP service core running on Ericsson's SSR 8010 and 8020.

In the access, the Ericsson SP router family was used to connect the radio base stations, emulated and real, to the network. The SP 420 and SP 415, equipped with 10GigabitEthernet and GigabitEthernet ports, would typically be used in a larger macro cell site where many antennas are positioned. The SP 110 provides a 1 rack unit solution for smaller macro cell sites and is also suited for small cell backhaul networks.

In addition, three Ericsson microwave transport solutions were used in the access network, these were the MINI-LINK PT 2020 (traditional spectrum from 6GHz to 42GHz), MINI-LINK PT 3060 (60GHz, also known as V-band) and MINI-LINK PT 6020 (70/80GHz, also known as E-Band).

In the access network, an extensive clock synchronization topology was set, with GrandMaster clock functions fulfilled by Microsemi TimeProvider 2700 and 5000 (supplied by Ericsson and part of the 'Evolved IP Network' solution).

The access network was then connected to the aggregation network, which was built using the Ericsson SSR 8010 routers. In the aggregation network we also connected several Ixia IxNetwork tester ports to emulate other access networks in order to mimic a typical large service provider network.

The aggregation network was attached to the IP service core network, where typically both mobile core, residential and enterprise services co-exist. For the purpose of some of the test cases we connected the biggest Ericsson router in the test, the SSR 8020, to an internal Ericsson mobile packet core.

Both the real and simulated mobile traffic in the test network were routed inside IP/MPLS layer 3 virtual private network services (L3VPN). In the access network these services were transported to the aggregation network using two MPLS labels – one was used to route the service to the aggregation network and the other was used to reach the IP service core using the transport layer. These three label-stacking principles were established by seamless MPLS methodology (<https://tools.ietf.org/html/draft-ietf-mpls-seamless-mpls-05>).

Alongside the L3VPN services, Ericsson configured virtual private wire services (VPWS) in which emulated residential and business subscriber traffic was encapsulated. Figure 2 depicts the complete test topology.

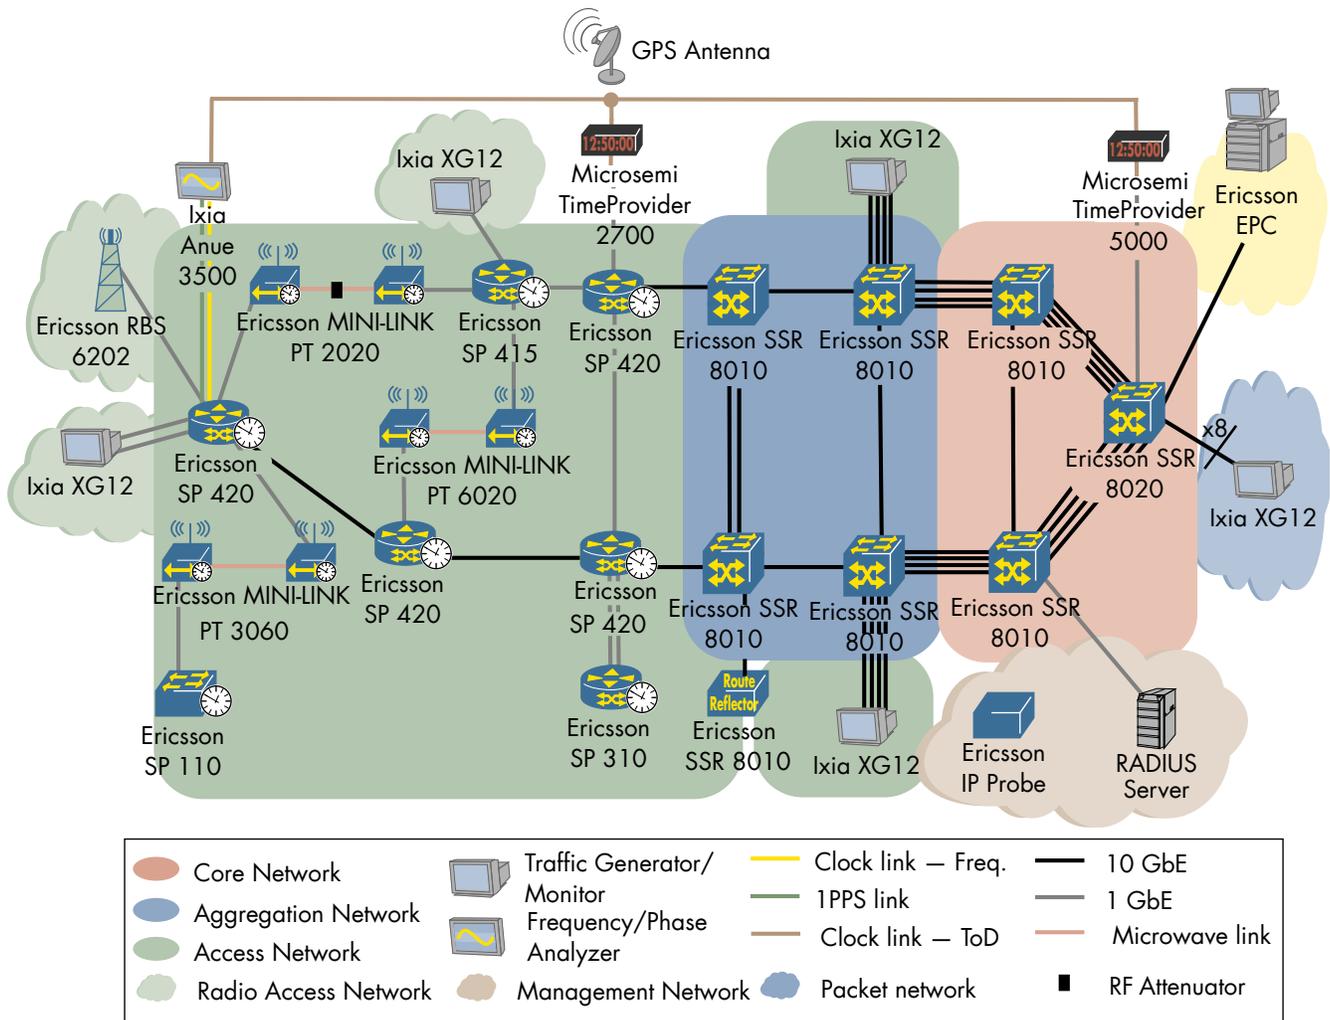


Figure 2: Ericsson Evolved IP Network Test Setup

Test Equipment

To execute the tests we used Ixia's XG12 chassis equipped with Xcellon-Flex-AP 16-port 10GigabitEthernet and LSM1000XMVDC16 cards. For our clock synchronization tests we used Ixia's Anue 3500.

We complemented the physical access network that Ericsson provided with our testers. Based on the Ericsson design, the access network would have been a single geographical location in which typically up to 1,000 base stations reside along side some 5,000 residential or business services. Based upon these guidelines, we emulated traffic sourced at base stations.

The aggregation network, however, is expected to support many more such access areas. For this purpose we used an Ixia IxNetwork tester to emulate additional 8 access networks, each operating at the same scale as the physical network. This meant that the tester had to emulate base stations as well as residential subscribers. From a tester point of view, in the aggregation network, the biggest task is to emulate the control plane, such that the real devices, in this case the Ericsson SSR 8010s, would treat the emulated traffic in the same way that they treat the physical devices.

Test Bed Traffic Configuration

In most of the tests, we used a tester configuration which enabled us to emulate 8 access networks at scale. The services, as figure 3 depicts, were all terminated on the IP service core router – the Ericsson SSR 8020.

In this configuration, the Ericsson SSR 8020 terminated 47,584 Virtual Private Wire Services as well as 8,000 Layer 3 Virtual Private Networks. The majority of the L3VPN end points were created by the emulated access ports (em_Access in the figure), however, 64 VPN Routing and Forwarding tables existed on the SP devices.

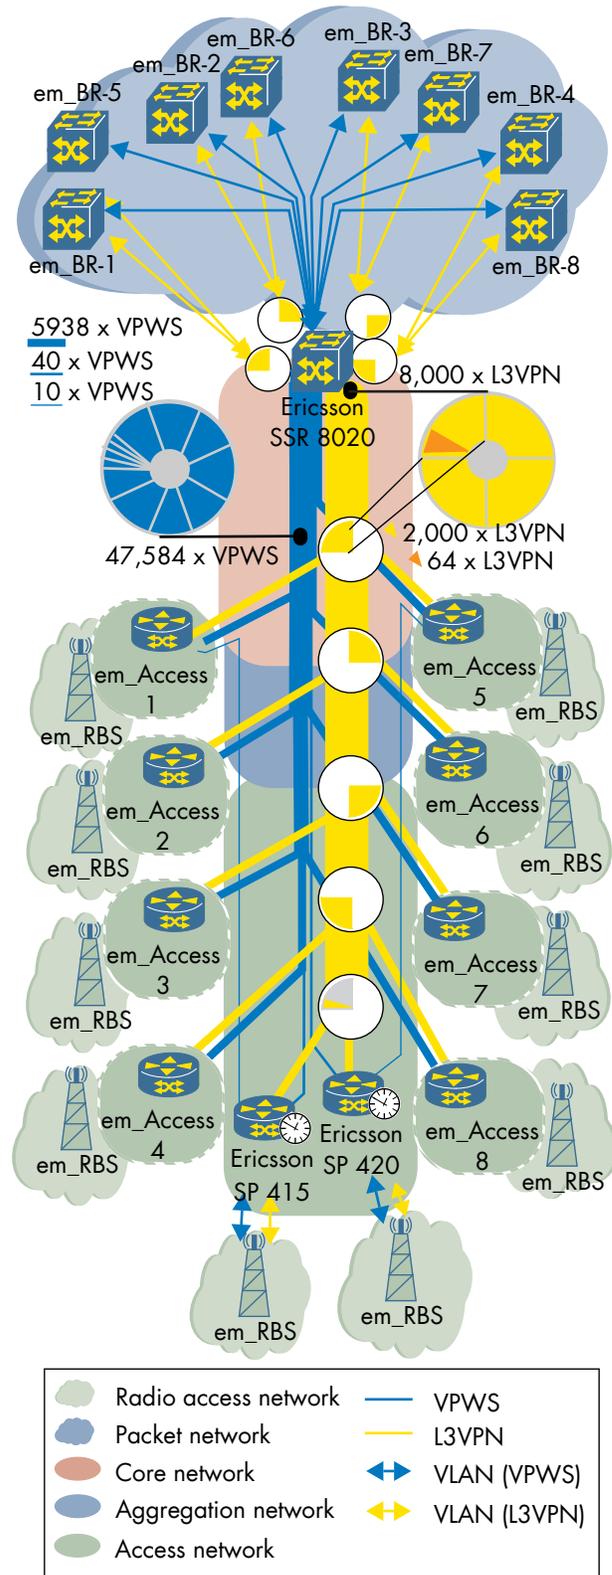


Figure 3: Ericsson Evolved IP Network Emulated Services

Intelligent Transport

When looking at today's complex network, we can see many competing and diverse requirements from the different services that run over them. Ericsson's solution is aimed at networks that carry both residential and business traffic, in addition to mobile services. Such a network should be secure, support all generations of mobile technology, and enable transport over densely populated urban areas as well as rural landscapes.

Typically we would also urge the vendor to allow us to test the Network Management System (NMS) solution, but given that Ericsson was in the process of launching¹ a new NMS just as we were in the midst of testing, we accepted that NMS testing would take place at a later date.

Multi-Standard Base Station

We started from the very edge of the network and worked our way to the core. This meant that Ericsson's multi-standard indoor macro Radio Base Station (RBS 6202) was the first observation point. Ericsson explained that the Base Station is installed with several interfaces that allow it to operate simultaneously using current mobile technologies: UMTS (3G using DUW card) and LTE (4G using DUL card).

Both UMTS and LTE ran natively over IP using the same IP/MPLS L3VPNs configured in the network to reach the mobile packet core.

EANTC does not typically perform physical layer testing. The test hardware required to perform such activities is completely different than the testers we typically use to emulate services or subscribers. So in this case, we used several mobile devices, each verified to operate using different mobile technology (i.e. UMTS and LTE) to connect to the internet. We first verified that the mobile device was really connected to the RBS 6202 in the test bed and not to another base station, by simply disconnecting the RF cables from the Ericsson RBS, one card at a time, and seeing that the mobile device we controlled lost its connectivity and that the connectivity only returned once we reconnected the cable.

1. <http://www.ericsson.com/news/1761209>

With that we were convinced that the RBS 6202 was really providing the mobile connectivity. We then browsed EANTC's web site while also streaming videos from a local content provider. We also called from the UMTS phone to another phone to verify that voice was also working. All activities were done in parallel, which verified Ericsson's point – the RBS 6202 was indeed multi-standard.

Microwave Adaptive Modulation

One of the unique challenges when using microwave as a packet transport medium is natural elements. For those times when rainfall conditions are exceptionally bad, the link capacity may be temporarily reduced due to an adaptive modulation shift that is compensating for decreased signal to noise ratio. This changes the modulation scheme on the microwave air interface to the maximum that can support the change in environmental circumstances. Ericsson explained that their MINI-LINK PT product family can support a number of traffic classes, reacting to link capacity reduction by scheduling higher-priority traffic and discarding lower-priority traffic based on configured quality of service settings traffic and prioritizing the rest.

Ericsson selected one of the three microwave devices for this test – the MINI-LINK PT 2020. They configured seven different classes for typical traffic types seen in converged networks. In the highest (most valuable) traffic class was packet clock synchronization and radio network control. The IP control plane and OAM traffic received the next level of prioritization while mobile voice, as well as gaming and video, were marked as the next priority class. These three traffic classes were configured to use a priority scheduler.

Low priority OAM traffic and mobile, non guaranteed bit rate (GBR) traffic, were being scheduled using a weighted fair queueing algorithm. The lowest class of service carried all unclassified data and was treated as best effort.

Emulating miserable weather conditions requires an RF attenuator that allows us to reduce the capacity of the microwave link. We started with link modulation of 512QAM, which at a 56MHz spectrum bandwidth provided 406Mbit/s of throughput, and reduced the modulation in regular steps all the way down to 4QAM. Using that modulation scheme the available air interface capacity was 94Mbit/s.

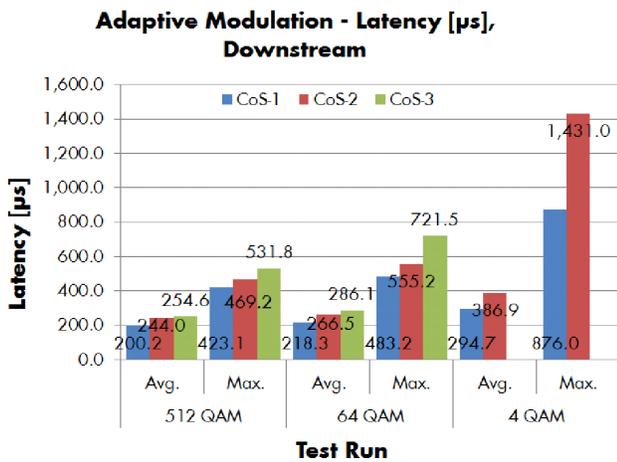


Figure 4: High Priority traffic average and Max Latency in the downstream direction

The results showed that Ericsson indeed built intelligence into the microwave devices that allowed the router/switch component of the MINILINK PT 2020 to prioritize traffic based on classes of service and to maintain healthy operations for the most valuable traffic, such as packet-based clock synchronization, even when the link capacity is reduced. We recorded no loss of high-priority traffic as we decreased the link capacity and monitored minimal change in latency: average latency was between 200 microseconds when using 512QAM. When we reduced the modulation all the way to 4QAM, the average latency for the highest priority traffic (CoS 1) was still averaging 294 microseconds with maximum latency of 876 microseconds.

Performance Monitoring Per Class of Service

Proactively identifying an issue in the network, before the customer detects any degraded performance, is a vital key to maintaining happy customers and selling premium services. The Internet Engineering Task Force (IETF) defined a mechanism in a Two Way Active Measurement Protocol (TWAMP) which Ericsson implemented in its SP 415¹ and 420 routers as well as in the Ericsson RBS 6202 base station. This mechanism allows service providers to automatically measure delay, delay variations and packet loss ratio. These are the three parameters on which our tests focused.

1. Since Ericsson explained that the SP 415 and 420 only differ in the number of ports, we only executed the test with the SP 415.

The TWAMP protocol, as implemented by Ericsson, also allows the operator to monitor specific classes of services for different values. Given that delay and delay variation could significantly differ between different traffic classes, it made sense for us to emulate delay in specific traffic classes and make sure that the reports collected by the Ericsson IP probes, were able to identify the increase delay.

We configured the Ixia Anue impairment generator to increase one class of service delay to 10 milliseconds and another class of traffic to suffer an increase delay of 20 ms. We then monitored with Ericsson's IP Probe and verified that the probe correctly reported the delay increase and later, once we removed the impairment profile, on the return of the network delay to the normal condition.

We collected measurement information from the two devices as shown in the figure below.

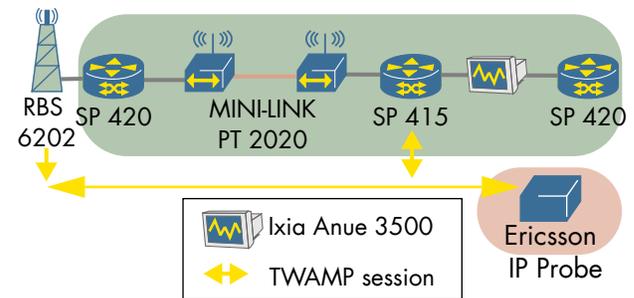


Figure 5: TWAMP Measurement Points

We also tested the ability of the same network elements to detect packet loss. We used the Ixia impairment generator to drop 5% of the traffic for one class of service and 15% of the traffic for another class of service. In all our measurement points, the monitoring devices showed the correct packet loss ratio.

We believe that these capabilities are useful for service providers especially in networks where packet-clock synchronization is used. Being able to monitor the delay and delay variations, can provide a service provider with an accurate picture of where the potential pitfalls are for their IEEE 1588-2008 deployment – where delay variation is high, the network should be looked at carefully. Of course the same tool can also monitor SLAs and provide general early warning alerts when service quality is reduced. Seeing the same performance monitoring tool being used in all network elements, from the base station itself onwards, is also an encouraging message to service providers.

Securing The Routers

There are several ways to secure routers. Control plane authentication should be used to make sure that only authorized routers could establish and exchange routes. The router itself should, as a rule of thumb, always be configured to allow only secure command line interface (CLI) connections and in the case that the router is being attacked, the router should be able to continue forwarding traffic without affecting user traffic.

We put two of the Ericsson devices through the rigor of security testing. First we successfully verified that both SSR 8010 and the SP 420 only established control plane sessions with neighbors using md5 authentication. We tested each device in the same way that it was configured in the test: the Ericsson SP 420 with OSPF and eBGP and the Ericsson SSR 8010 with IS-IS and iBGP.

We then used two well known attacks directed at the Ericsson SSR 8010: ICMP and ARP attacks. Both were originating in one of our emulated access connections such that the Ericsson SSR 8010 was the first device to detect the attack. Together with our attack traffic, we sent emulated customer traffic and monitored both CPU utilization and packet loss on the legitimate subscriber traffic.

The results showed that the Ericsson SSR 8010 was able to successfully continue operating while it was being attacked by both ICMP and ARP messages. The CPU remained stable and no packets were lost from the legitimate user streams.

Intelligent Transport Tests Summary

Security testing is of course a complete area, one on which we could spend several months. Given the scope and the tests, we wanted to make sure that the solution at least meets all the functional criteria that we believe a service provider will be looking for: transport flexibility, performance monitoring, and of course, multi-standard base station support. Our expectations were indeed met by the Ericsson 'Evolved IP Network.' Having performed these base line tests, it was prime time to move on to one of the most challenging areas in networking these days: packet-based clock synchronization.

Clock Synchronization

Mobile networks have always required frequency synchronization, but in recent years the industry is talking more and more about the need for time of day and phase synchronization. This is driven by a number of radio features such as LTE Broadcast and Small Cell co-ordination. While an operator might not today require time and phase synchronization, it is becoming important to understand how to evolve their network to be able to support it in the future. If this is not provided with sufficient accuracy and stability a base station will be degraded in terms of its capabilities, and in the worst case this may cause it to disconnect from the network

Ericsson suggested to split the clock synchronization tests into two parts - first to focus on full-path support in line with Ericsson's 'Evolved IP Network' design guidelines (based on the ITU-T G.8275.1), and then additionally test for partial-path support which is not yet standardized, but could be of interest to forward-looking service providers.

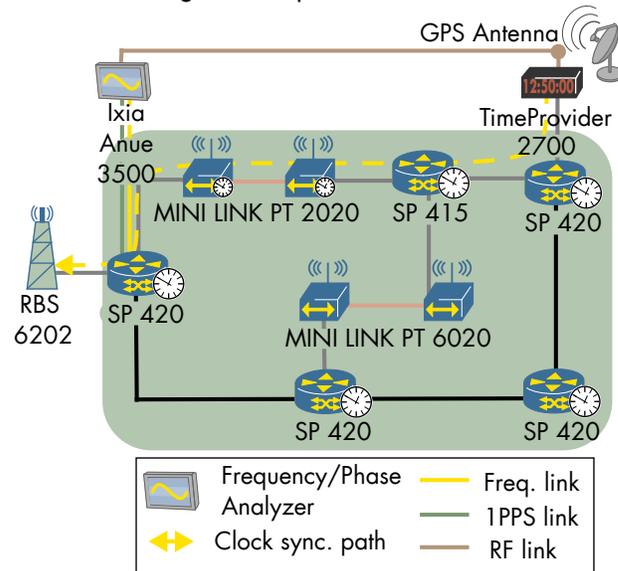


Figure 6: Full-path Clock Synchronization

Full-path support means that every node in the clock distribution path, from the primary reference clock (in this case a Microsemi TimeProvider 2700 GrandMaster) to the base stations (Ericsson RBS 6202) was in some way actively involved in maintaining the clock quality. Since Ericsson uses the standard-based IEEE 1588-2008 (also called Precision Time Protocol or PTP), the roles of the Ericsson products split between a Boundary Clock and Transparent Clock.

Between the GrandMaster and Ordinary Clock (implemented as a function inside the Ericsson RBS 6202 LTE base unit) were both Boundary Clocks in the SP routers, and Transparent Clocks in the MINI-LINK PT microwave devices. In the RBS node itself two functions existed: the Transmission Control Unit (TCU) implemented a Transparent Clock function and the air interface control cards implemented an Ordinary Clock.

Full-Path Clock Quality

We started our clock testing by measuring the clock quality in the access. In this topology, where the clock source is positioned somewhere within the network access, all base stations could synchronize to the same GrandMaster.

We generated load in the access network based on the definition in ITU-T G.8261 for Test Case 12. We used Network Traffic Model 2 which is intended to model traffic where the majority of the traffic is data. The load was equivalent to 80% of the downstream capacity (towards the base station) and 20% in the upstream direction. We then made sure that the Boundary Clock, to which we attached our tester, was locked to the GrandMaster, and captured measurements over a whole weekend.

The results showed that the LTE base station (with its stringent phase synchronization requirement) ran smoothly over the whole weekend. From a standards perspective, the phase offset should be within 1100 nanoseconds - we confirmed that the clock quality for phase was within +/- 325.56 nanoseconds. The frequency was also confirmed to pass the G.8261 EEC Option 1 mask for the complete period monitored - which in itself confirms the long-term limit requirement of 16ppb or better as well.

Now that we knew that the quality of the Ericsson clock distribution solution was compliant with both phase and frequency standards, we investigated its robustness by emulating various fail-over scenarios and checking if the clock quality remained within the accepted tolerance.

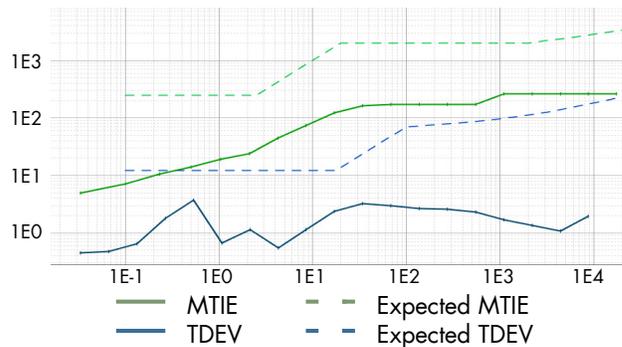


Figure 7: Full-path MTIE and TDEV measurement over a weekend G.8261 EEC Option 1

Clock Signal Robustness – Link Failure in the Access

Positioning the GrandMaster clock at the boundary between the Access and Aggregation networks has several benefits. The clock path is shorter than when the clock is positioned in the IP service core and therefore delay variation, an aspect of networks that increases with distance and the number of nodes in the path, is shorter. If a path within the access network experiences a failure, the network could re-route around the failed link or node, but the GrandMaster function will be uninterrupted.

In this scenario we disconnected the link between the Ericsson SP 415 and the MINI-LINK PT 2020. We then checked to see that in our monitoring point, the Ericsson SP 420 Boundary Clock, the clock source port changed. The SP 420 changed its state to “free-running” for 26 seconds and then locked to the GrandMaster again.

This was unexpected behavior since the natural first state to transition into, for an Ordinary and Boundary Clock should be holdover - a state in which the clock is maintained locally. Ericsson explained that phase holdover was not supported in the software release under test, but that it is planned for future release.

During those 26 seconds, the 1PPS output, which typically will be used by the base station to receive phase signal, was turned off - by design. This indicates, to the base station, that it should be using other mechanisms for phase information at that point, such as its own internal oscillator or other backup references, to maintain phase stability.

The next step was to measure clock quality during recovery - when the link which we disconnected was restored. Once we restored the link, the SP 420 changed its primary clock port back to the original port. The SP 420 reported free running state for 121 seconds and then re-locked.

These results show that the Ericsson full-path support design guideline can deliver the frequency and phase synchronization required by LTE. For those networks that are not equipped with GrandMaster clocks in every access network or use a mix of devices, some of which do not support PTP, a partial-path support model could be used. Partial-path, where the packet clock signal is crossing nodes that do not support PTP, was our next focus.

We evaluated two partial-path support scenarios: the first condition being a failure of the primary GrandMaster clock (positioned in the access network) forcing the network to use a GrandMaster in the core network. The second condition was a network which uses a GrandMaster clock positioned in the core as primary reference clock.

GrandMaster Clock Resiliency (partial path)

Clock quality may not immediately be affected by the loss of the primary reference, owing to synchronization holdover. Depending on the oscillator quality, a device can run in holdover mode for an extended period of time. It is best to return to the healthy clock operation as quickly as possible. Interestingly, the duration in which a device should maintain an accurate clock in a holdover state has not been defined by the 3GPP - this is perhaps an area for the service providers and vendors to investigate in the future.

The IEEE 1588-2008 standard describes an algorithm in which a Boundary or Ordinary Clock can select the best timing source. The algorithm compares the following parameters in this order: Priority 1, Clock Class, Accuracy, Variance, Priority 2, and Identity. In this test, we disconnected the GPS antenna from the primary GrandMaster, and verified that the Boundary Clock, to which we connected our Ixia Anue 3500 measurement device, switched to the secondary GrandMaster. We also measured the frequency and phase quality during the whole procedure in order to understand if an LTE base station will be affected by such failure condition.

Initially the Microsemi TimeProvider 2700 GrandMaster continued providing primary reference. Five minutes after we disconnected its GPS source, the GrandMaster degraded its clock class and the Ericsson SP 420 reacted by changing its state to free-running. This behavior is consistent with previous tests in the full-path support section.

Within 140 seconds of the Ericsson SP 420 changing its state to free-running, it locked to the second GrandMaster clock – a Microsemi TimeProvider 5000 positioned in the IP service core. When we calculated the phase and frequency masks over the complete time period we could see that they were not being met, upon further investigation of the data before and after the failure we could see that after the SP 420 has locked to the backup grandmaster, the MTIE was met, but not the TDEV.

Ericsson explained that since the clock is now distributed in a partial timing network, it is much more noisy. In this scenario, the healthy operation of the base station will depend on alternative clock sources or its own oscillators maintain healthy operations. Another alternative, recommended by Ericsson is to filter the additional noise caused by the partial-timing solution.

Clock Signal Robustness - Link Failure in the Aggregation

In the last failure condition, we emulated perhaps the most difficult condition for a packet clock distribution network. Here Ericsson configured a single GrandMaster clock, attached to the Ericsson SSR 8020 in the IP service core. This meant that the clock distribution network was from the get-go in a partial timing mode as the SSR 8010 and 8020 routers in the IP service core and aggregation network were not configured with any Precision Time Protocol functions.

A single VPWS service was used to transport the PTP messages from the GrandMaster to the Boundary Clock in the access network. The VPWS service was configured between Ericsson SSR 8020 where the GrandMaster clock was connected, and Ericsson SP 420 where the Ericsson SP 310 (Boundary Clock) was connected. Ericsson explained that the Boundary Clock on the SP 310 was configured such that the upstream port was configured for IP/unicast while the downstream port was configured for Ethernet/multicast.

A backup RSVP-TE tunnel was used to provide protection and path symmetry for PTP traffic in case of failure on the primary path. The PTP traffic in the access was distributed using native Ethernet encapsulation.

Once we verified that the clock was being recovered on the same Ericsson SP 420 router that was used as Boundary Clock in all our tests, we disconnected the link between the two SSR 8010s in the aggregation network. This caused the PTP traffic to reroute around the failed link.

During the whole duration of the test we monitored the 1 PPS and SyncE interfaces on the Ericsson SP 420, making sure that both phase and frequency requirements (the phase deviation within 1100 nanoseconds and the frequency within 16 ppb) were being met. The measurements showed that the absolute phase error exceeded 1100 nanoseconds on two occasions before the failure and frequency measurement did not pass ITU-T G.823 TDEV SEC mask. Ericsson explained that this was expected because the clock was distributed using partial timing path.

Even during the emulated failure condition the Ericsson SP 420 stayed time-locked to the GrandMaster clock. This was due to the fact that the failure was resolved quicker than the PTP Announce Message Interval (set to 1 packet/second). Rerouting of the PTP stream in the aggregation network added two additional network hops in the PTP path thus the

mean path delay increased from 56µs to 81µs. As a next step we removed the emulated failure on the working path and after 140 seconds Ericsson SP 420 entered in free-running state, before re-locking. This behavior is expected and is due to a change in the floor packet delay as per section I.5.1.3.2 of the ITU-T G.8260 Amendment 1, which states that the packet network limit measurement should be restarted when such a floor packet delay change occurs.

Clock Synchronization Tests Summary

The Ericsson team undertook a monumental challenge, showing not only the optimal clock conditions, but also degraded clock conditions that are likely to trouble any operator. Packet-based clock synchronization is a reality for many networks that can not or do not want to use GPS. Currently the main proven mechanism to distribute phase, Time of Day and frequency information in packet networks is the IEEE 1588-2008 Precision Time Protocol.

Aside from academic discussions about theoretical frequency and phase quality, the important question that an operator should be asking himself is whether the base stations are going to maintain operations and if the air interface is being used at optimal efficiency. The answer for the current generation for base stations is a yes as our measurements show. The other substantial take home message for operators is also that the design of a clock distribution

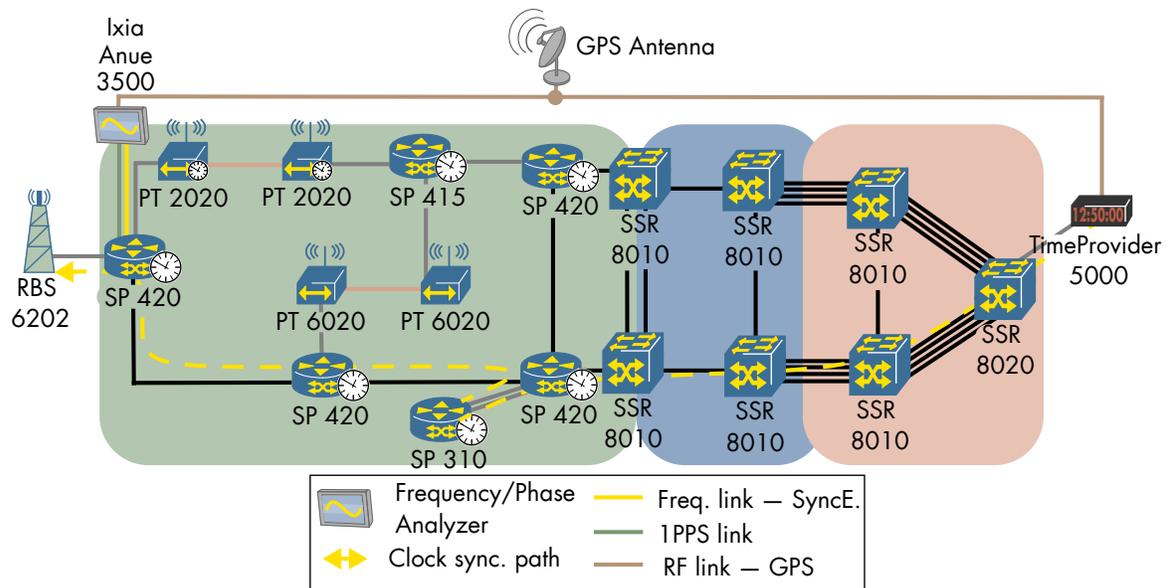


Figure 8: Partial-path clock distribution Test Setup

network must be done together with detailed understanding of the base stations capabilities and requirements, as well as with the active support of the underlying IP infrastructure.

High Availability

Once we finished the clock testing we shifted our attention to another attribute that each service provider network is very concerned with: service availability. After all, a service is often measured by well defined parameters in a service level agreement (SLA). We already verified that the Ericsson solution can measure delay and delay variations, two parameters that are often used in SLAs. Now it was time to check if services are affected when failures in the network occur.

Recovery from Link and Node Failure

Whether a link fails due to a backhoe digging where it should not have been, or a router fails due to mechanical fan issues, the safest way to operate a network is to let it react to sudden failures by itself. In the Ericsson 'Evolved IP Network' the vendor used IP Fast Reroute, as specified in RFC 5286. Ericsson explained that they favour IP Fast Reroute over MPLS Fast Reroute since the former provides fast failure recovery with LDP, and hence does not require the complexity involved with RSVP-TE.

We picked one node and one link in the network as the target for our failure condition: one of the connection points in the aggregation network. To simulate a router failure, we simply disconnected the Ericsson SSR 8010 from its power source. To simulate the link failure we disconnected the fiber link between two of the Ericsson SSR 8010s in the aggregation network. As can be seen in the figure below, none of the failover scenarios caused an out of service time larger than 5.3 milliseconds.

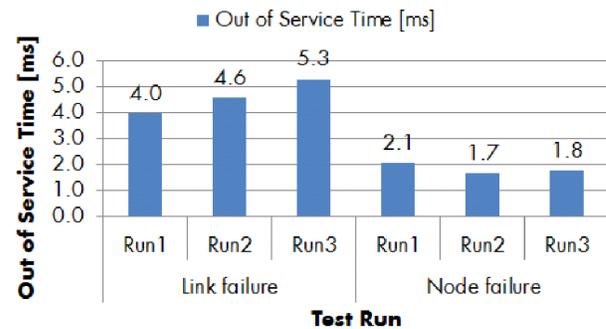


Figure 9: IP Fast Reroute – Out of Service Time

Recovery from Optical Transmission Failure

Another scenario that is particularly maddening is the failure of optical transmission units. In some of these failure scenarios, the routers terminating the links do not experience loss of signal (LOS), but traffic is no longer reaching them. Typically routing protocols will notice that their hello messages are not being answered, but the hello intervals are typically configured in the order of seconds. Prudent service providers typically look for more aggressive mechanism to detect such failures. Such a mechanism was created by the Internet Engineering Task Force (IETF) and is called Bidirectional Fault Detection or BFD.

Ericsson configured BFD on all interfaces in the test and for all routing protocols used on all devices - SP and SSR. We picked one of the links, the one between the access network SP 420 and the aggregation network SSR 8010 to insert a switch and turn off the switch's backplane. This emulated the same condition as a transmission device failure – both routers had their interfaces up, but no traffic was passing through.

Given that Ericsson configured BFD hello intervals at 100 ms and the number of missing hellos before the protocol was to raise the alert at 3, we expected that the detection time would be between 200 and 300 ms. Once the failure is detected the remaining time is due to service reconvergence. We recorded, upon triggering the failure, an out of service time of 417 ms in the upstream direction and 519 ms in the other direction (upstream direction is base station to the Internet).

The next step in the test is of course the recovery – the state in which the issue is resolved and the network needs to return to normal operations. Here the potential issue is that the Internal Gateway Protocol (IGP), the protocol that helps LDP identify the next hop in the path, will need to reconverge. This could potentially black-hole traffic until all the protocols converged to the original path. This LDP/IGP synchronization in our test was done between IS-IS and LDP and was measured to be 9ms in the upstream direction and 190ms in the downstream direction. Ericsson explained the difference in out of service time as a consequence of the different hardware architectures between the SP 420 and SSR 8010, which leads to different re-converge times in the upstream and downstream directions.

Gracefully Restarting Routing Protocols

The Internet Engineering Task Force (IETF) defined mechanisms to allow routers to restart their control plane without affecting the forwarding plane. Sometimes BGP or IS-IS need to be restarted, but no topology changes took place. In this situation, when the forwarding plane is functioning and the network topology is stable, if the process that is being restarted is able to communicate that fact, other nodes would continue forwarding traffic and no effect on traffic in the network is expected.

We restarted both BGP and ISIS process on the Ericsson SSR 8010 while sending the full traffic load. In order for the restart to really be graceful, the neighbors of the router being restarted need to operate in helper mode. This role was taken by the Ericsson SP 420. We used the CLI commands ‘process restart bgp’ and ‘process restart isis’ and measured the out of service time.

The results matched our expectations as no packets were dropped during this operation. We also verified that no re-convergence was happening in both routing protocols behind the scenes just to be sure that the routers were really using Graceful Restart mechanisms.

High Availability Tests Summary

We verified that the Ericsson ‘Evolved IP Network’ is equipped with high availability mechanisms to help service providers maintain operations during disaster scenarios. With mobile traffic moving to an all-IP framework, all customer traffic in a converged network, be it mobile, residential or enterprise, is IP-

centric. This gives credence to Ericsson’s choice of depending on IP-centric resiliency mechanisms from the large pool of options available in the various standard bodies.

Scalability and Performance

It is only logical that a service provider should be concerned with the service scalability of any solution that he or she is buying. In order to provide an overview of the service scalability and performance, we picked two of the devices – the Ericsson SP 420 and the SSR 8010, and verified that each of them can support the number of services Ericsson committed to supporting.

VPWS and L3VPN Service Scalability

For the remaining tests we disconnected the network and connected the routers directly to the Ixia tester. This meant that the Ixia tester had to emulate the roles of the other Ericsson devices, a function that the tester had no problems executing.

The service scalability tests used the same methodology regardless of the service tested. We first established control plane connectivity with the device under test and verified that all services were installed using the command line interface. Once all services came up, we sent traffic close to line rate in order to verify that the service was also usable.

Tested Network Services	# Services on SP 420	# Services on SSR 8010
L3VPN	64	8,000
VPWS	250	239,844

In the L3VPN case we also injected 128 routes to each of the VRFs on the SSR 8010 as well as 156 routes to each of the 64 VRFs on the SP 420 for a total of 1,025,792 unique routes. We then sent traffic to each of these routes to make sure that they were indeed installed in the forwarding plane.

Forwarding Performance

If a network is truly successful, one can imagine that as customers are added, more and more bandwidth will be used and the routers will be expected to deal with an increase in forwarding requirements. As a rule of thumb, service providers do not tend to run their routers at 100% utilization – on the contrary, normally routers only run at 50%-75% utilization.

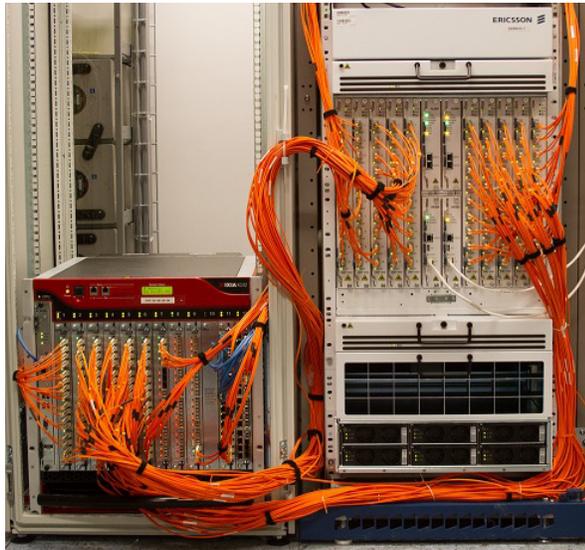


Figure 10: Ericsson SSR 8010 – Performance and Scalability Setup

To round up the testing we performed IPv4-focused forwarding performance tests on two of the devices: the Ericsson SP 420 and the SSR 8010. Both were fully loaded with line cards which meant that the SP 420 was tested in a configuration of 20 GigabitEthernet and 2 10GigabitEthernet while the SSR 8010 was loaded with 10 line cards each supporting 10 10GigabitEthernet ports.

We used the same methodology described by the IETF in RFC 2544 and selected frame sizes provide a realistic overview relevant to service providers. Typically, the smallest frame size (64 Bytes without VLAN, 68 bytes with VLAN) is used to measure packet processing performance of the solution, but are not the only packets seen in the wire. For this reason we used a packet size mix, an Imix, which is more typical of actual packets in the network. We summarize the results in the table below.

Packet Size (bytes)	Line Rate forwarding MINI-LINK PT 6020	Line Rate forwarding SP 420	Line Rate forwarding SSR 8010
68	100%	100%	69.4%
128 ^a	Not-relevant	Not-relevant	100%
373	100%	100%	100%
1518	100%	100%	100%
IMIX	100%	89%	98%

a. Frame size added to SSR test to measure 100% performance at small frame size

Summary

The networking market is dominated by vendors whose roots are in the Internet. The clear separation between mobile and Internet has been eroded since the advent of the smart phone with more and more devices accessing the Internet using radio signal as opposed to cable or wifi. We see that an independent test of a holistic converged network solution is one of the strongest statements a vendor can make in the market and the most reliable one at that.

Based on the results of the test we can confirm that Ericsson’s ‘Evolved IP Network’ solution is meeting Ericsson’s claims. The solution is IP-centric, offering multi-standard base stations and QoS-aware microwave transport. We verified that performance monitoring tools were functioning accurately and could already be used at the base stations as well as on transport nodes. We verified that Ericsson’s unique and standards-based approach to high availability provided the high availability required in modern networks, and that the solution scale to the design goals Ericsson would suggest to its customers. We must also salute Ericsson’s courage to demonstrate, along side optimal clock synchronization topologies, the more challenging partial-path setup. We expect future tests as the solution evolves.

About EANTC



The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

EANTC AG
 Salzufer 14, 10587 Berlin, Germany
info@eantc.com, <http://www.eantc.com/>
v1.0, 20140221, JG

Any marks and brands contained herein are the property of their respective owners.