Experimenting and Deploying IEEE 1588v2

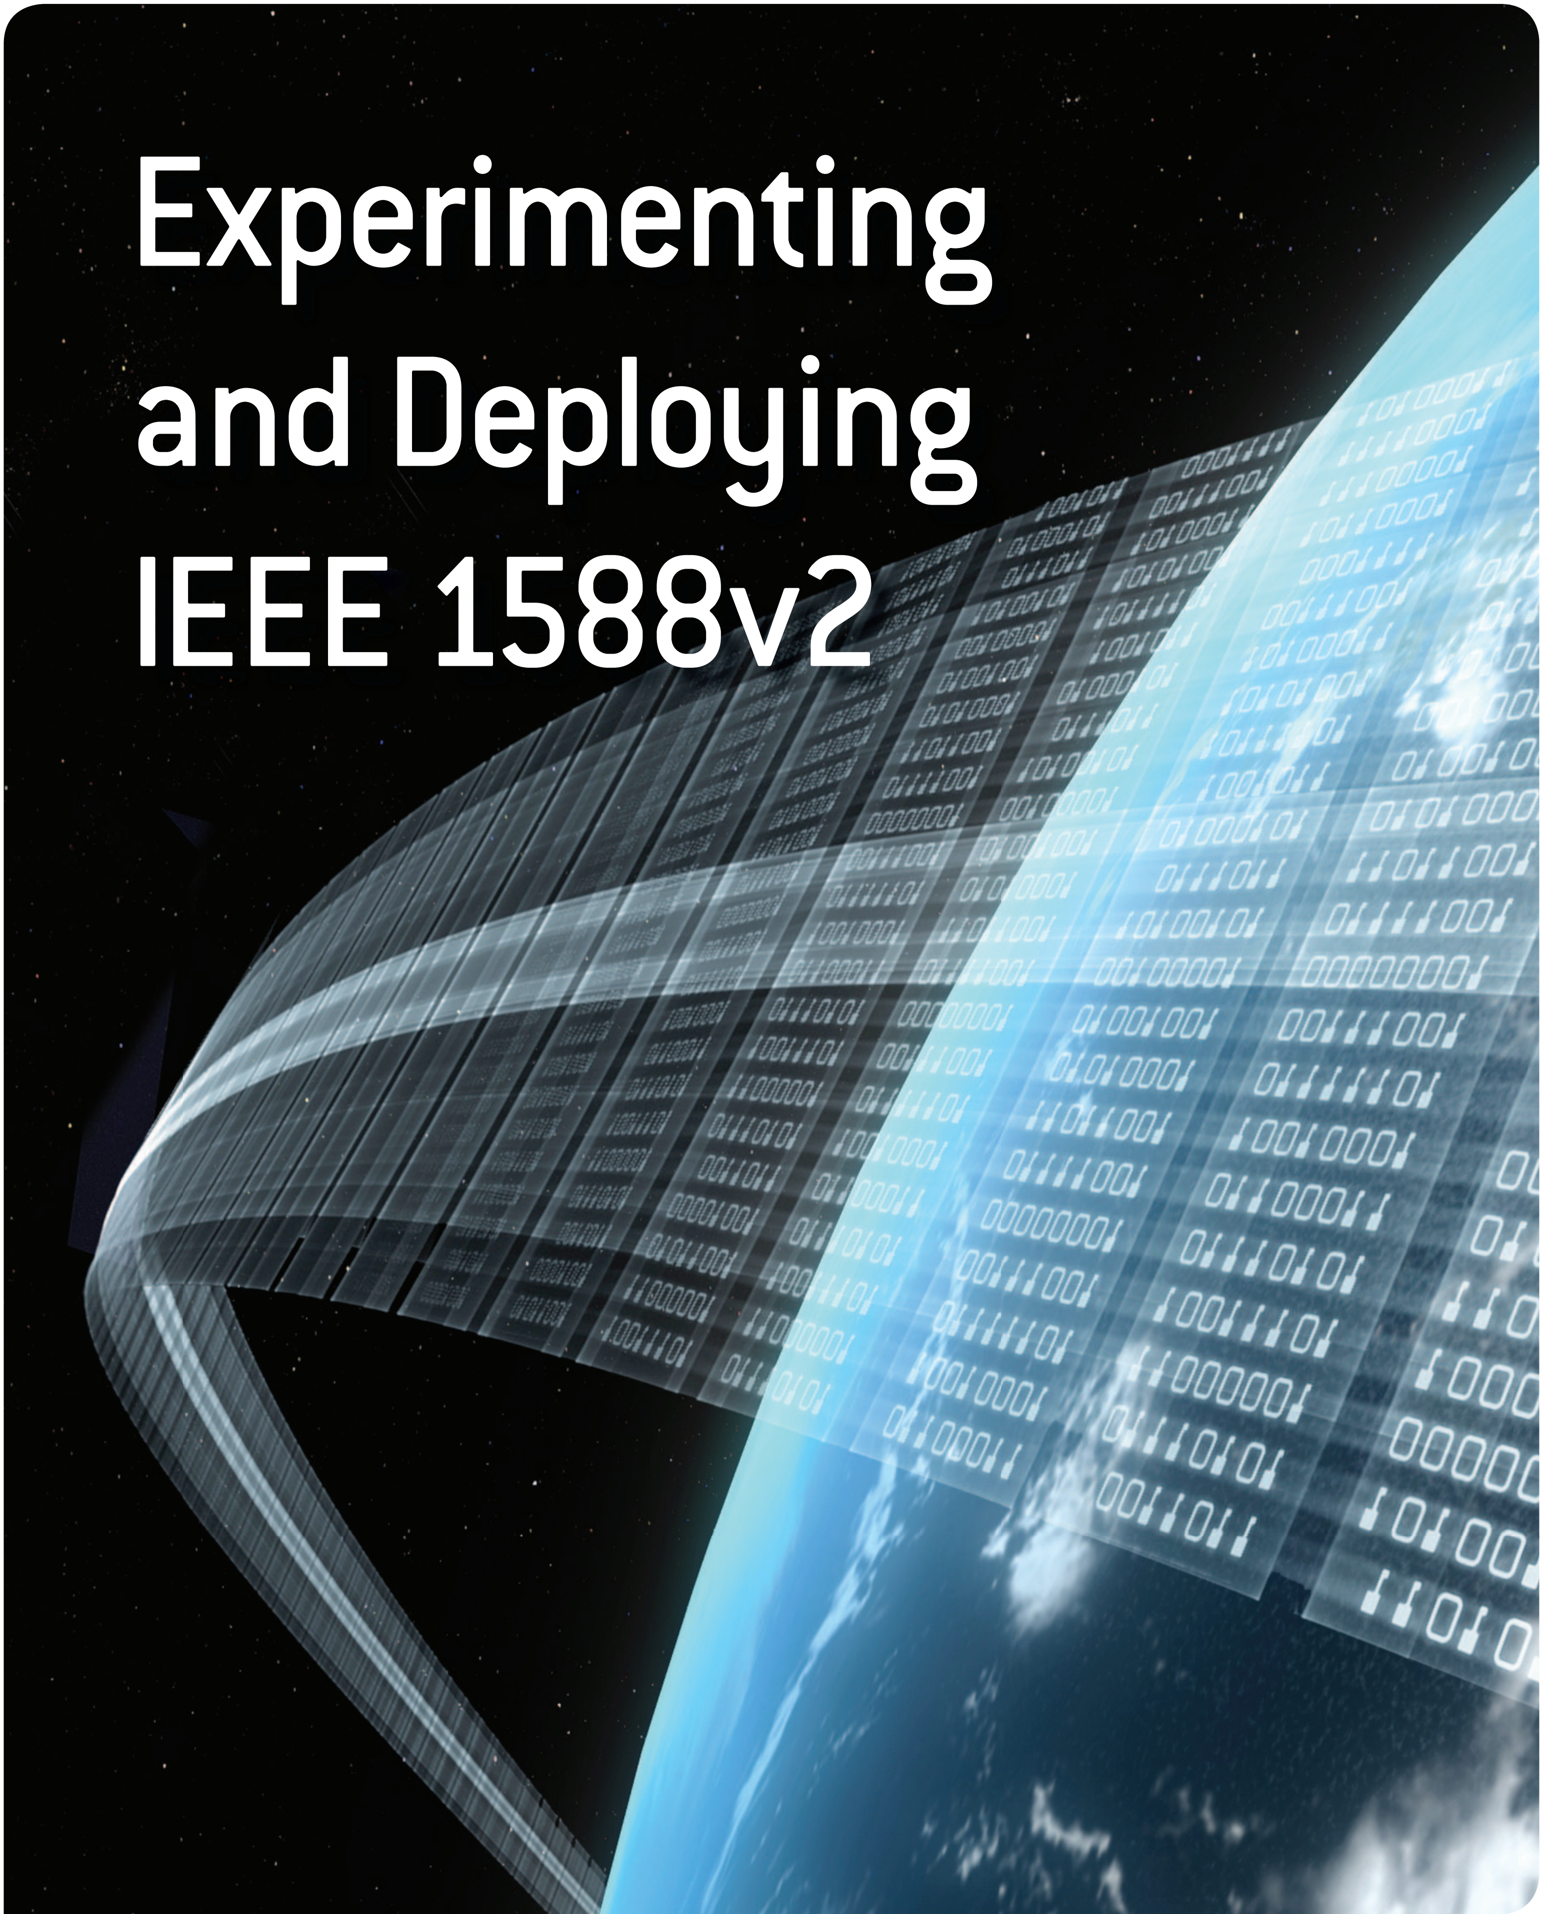

EXFO | Assessing Next-Gen Networks

EXECUTIVE SUMMARY

THE CONTEXT

Mobile operators are upgrading the backhaul network to a packet-based network in order to answer the bandwidth explosion resulting from the increase in subscribers and the enhanced mobile data services.

A major challenge of this transition is maintaining synchronization, a mandatory requirement for mobile stability and handset handoff. As Ethernet is an asynchronous technology, newer methods and standards have been introduced to implement synchronization over Ethernet.

IEEE 1588v2 PTP is becoming the solution of choice for synchronization distribution, as it uses an exchange of timestamp between master and slave devices, thus providing frequency, phase and time sync over the existing Ethernet infrastructure.

The Problem

Network operators deploying IEEE 1588v2 PTP face numerous challenges:

› The complexity of the IEEE 1588v2 PTP standard, whose performance is clearly affected by the one-way delay and one-way delay variation characteristic of the network

› A lack of experience with this relatively new standard and its one-way performance requirements

› The lack of proper tools to troubleshoot the technologies, as traditional performance metrics are focused on round-trip results while IEEE 1588v2 PTP is affected by one-way performance

› The proprietary aspect of the PTP algorithm, which can result in performance differences between competing solutions for the same network conditions–adding complexity in the deployment

The Solution

A five-step methodology designed to help operators select and deploy a PTP solution.

This five-step methodology follows the network lifecycle approach and consists of the following phases:

1. Network survey
2. Benchmark
3. Sync service turn-up/rollout
4. Sync monitoring
5. Troubleshooting

The Benefit

› Simple five-step approach based on tier-1 and tier-2 expertise, resulting in time and OPEX benefits

› Methodology with clear objectives and deliverables for each phase

› Early-stage identification of performance concerns that could affect PTP performance

› Benchmark report that can be used to simplify the tender process by providing key relevant metrics for PTP performance

› Benchmarking and selection of vendor solution based on important, factual data such as network performance

› Methodology designed around tools available today, resulting in strong CAPEX benefits

Experimenting and Deploying IEEE 1588v2

Thierno Diallo, Product Specialist, Transport and Datacom

The introduction of enhanced mobile services has increased the bandwidth usage on mobile network bandwidth, making for a better user experience. Long-gone are the days of a mobile phone with voice service only. Today's typical subscribers are bandwidth consumers, navigating the Internet, viewing videos and downloading applications directly from their mobile handset.

Mobile operators had to adapt to this capacity crunch and newer standards such as LTE promise to deliver the bandwidth, quality of service and range required to enable these new enhanced mobile services. Such new performance requirements have had a great impact on the mobile backhaul infrastructure, the network carrying traffic from the data centers to the mobile towers, which is now evolving from a TDM-based network to the more efficient and higher bandwidth Ethernet-based network, supported by a host of new carrier technologies such as OAM and MPLS-TP.

A key aspect of mobile technologies is synchronization. Basically, synchronization refers to the coordinated and simultaneous relationship between clocks or time-keeping applications. For mobile transmissions, synchronization requirements are critical performance parameters that must be met in order to ensure a trouble-free, stable operation and enable user handoff between serving towers. In essence, without synchronization, mobile freedom and usefulness as we know them would not exist.

The original TDM-based backhaul networks provided synchronization via the use of physical-layer frequency distribution. Simply put, the deployed SONET/SDH and lower rate DSn/PDH transport technologies were based on synchronous operation and provided deterministic and stable timing. As the mobile backhaul is transitioning to a packet-based network, operators face the challenge of maintaining synchronization over the inherently asynchronous Ethernet technology.

Over the last few years, the packet world has seen great advancements on synchronization, with two major technologies emerging as the methods of choice for sync distribution: ITU-T G.8261 Synchronous Ethernet, typically refered to as SyncE, and IEEE 1558v2 Precise Time Protocol, or PTP. As both technologies provided frequency distribution capabilities, PTP demarked itself by using the efficient aspect of packet distribution to also enable phase and time synchronization, making it the most advanced and applicable solution for newer mobile technologies such as LTE/4G.

Mobility Air-Interface Stability Needs		
Mobility Standard	Frequency (ppb)	Time/Phase
CDMA2000	50	<3 to <10 μ/sec
GSM	50	
WCDMA (UMTS)	50	
TD-SCDMA (China Mobile)	50	3 μ/sec intercell phase change
LTE (FDD)	50	
LTE (TDD)	50	2.5 μ/sec intercell phase change
LTE MBMS	50	5 μ/sec intercell phase change
WiMAX	20	1 μ/sec

Unfortunately, PTP is a challenging technology. This timing method requires strict conditions in terms of network performance for proper operation, which adds to the difficulty due to the lack of experience with this method. Many operators are looking at deploying PTP but are still wondering how to approach it in order to limit its impact on the existing investment.

The following pages build upon EXFO's experience in PTP and proposes a comprehensive and efficient five-step method for operators and carriers interested in PTP. This guide starts with a brief introduction on PTP, then describing the challenges of this timing solution and presenting the five steps with added operational and implementation details.

AN INTRODUCTION TO IEEE 1588v2

IEEE 1588v2 PTP is specifically designed to provide high clock accuracy through a packet network via a continuous exchange of packets with appropriate timestamps. In this protocol, a highly precise clock source, referred to as the grand-master clock, generates timestamp announcements and responds to timestamp requests from boundary clocks, thus ensuring that the boundary clocks and the slave clocks are precisely synchronized with the grand-master clocks. By relying on the holdover capability and the precision of the integrated clocks in combination with the continuous exchange of timestamps between PTP-enabled devices, frequency and phase accuracy can be maintained at a sub-microsecond range, thus ensuring synchronization within the network.

The objective of the PTP deployment is really simple: by exchanging timestamps, the slave clock can determine its offset from the master clock and thus adjust itself. This provides frequency and phase synchronization through packet distribution.

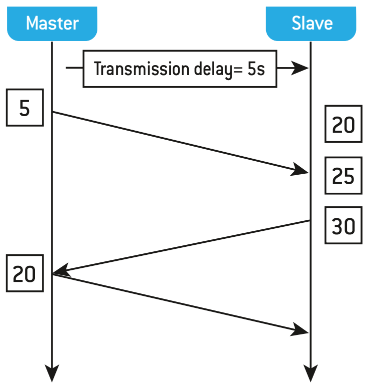

Scenario 1: perfectly symmetrical one-way delay

Transmission delay = 5 s
Offset between slave clock and master clock = 15 s

t1 = 5, t2 =25,
t3 = 30, t4 = 20

Unidirectional transmission delay =
Delay (M->S) + Delay (S->M)= (t2-t1) + (t4-t3)
------------------------------------ --------------------
2 2
= (25-5) + (20-30) / 2 = (20 -10)/2 = 5

This accurately reflects the one-way transmission delay, but only if we assume that the delay in both directions is exactly the same.

The offset between master and slave clock is deduced by:
Offset = slave clock – master clock – transmission delay
Offset = t2-t1-UTD = 25-5- 5 = 15

Figure 1. PTP scenario: best case.

The PTP challenge

The great advantage of PTP is that as a packed-based technology, only the boundary and slave clocks need to be aware of the nature of the packets; therefore, synchronization packets are forwarded as any other data packets within the network.

The major weakness of PTP is also due to its packet-based nature. As the synchronization packets used by PTP are forwarded in the network between the grand master and hosts, they are subject to all network events such as delay (latency), delay variation (packet jitter) and frame loss. Even with the best practice of applying high priority to synchronization flows, these synchronization packets will still experience congestion and possible routing and forwarding issues such as out-of-sequence packets and route flaps. The host clock's holdover circuit must be stable enough to maintain synchronization in the event that synchronization packets experience these network events.

The impact of delay asymmetry

The PTP protocol requires near-perfect latency symmetry between the master-to-slave direction and the slave-to-master direction, and it uses this assumption to estimate the phase difference between master and slave clocks. When the assumption is correct, the PTP protocol can easily identify the transmission delay between master and slave via packet exchange.

However, this is rarely the case as PTP packets can often experience different levels of delay per direction, typically due to the congestion level of the paths used. This asymmetry can cause significant errors in the offset estimation at the slave clocks and therefore introduce frequency and phase difference in the output clock compared to the master clock. The following example shows the impact of asymmetry on offset estimation:

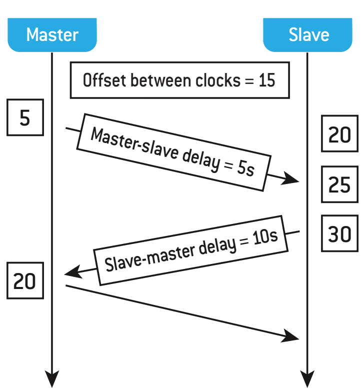

Scenario 2: asymmetrical one-way delay

Transmission delay = 5
Offset between slave clock and master clock = 15 s

t1 = 5, t2 =25,
t3 = 30, t4 = 25

Unidirectional transmission delay =
Delay (M->S) + Delay (S->M) = (t2-t1) + (t4-t3)
------------------------------------ --------------------
2 2
= (25-5) + (25-30) / 2 = (20 -5)/2 = 7.5 != to 5

The one way delay is incorrect due to the different delay per direction.

The offset between master and slave clock is deduced by:
Offset = slave clock – master clock – transmission delay
Offset = t2-t1-UTD = 25-5- 7.5 = 12.5 != 15.

The offset measured by the slave is incorrect and it will not adjust itself properly to the master clock.

Figure 2. PTP scenario: the impact of asymmetrical delay.

The impact of delay variation

As PTP packets are exchanged, the latency experienced by each packet can vary, essentially due to the various processing time as the packets traverse the network. Fortunately, PTP clients typically implement proprietary algorithms to detect packets with abnormally high delay and will implement smoothing algorithms to limit the delay variation on the phase adjustment of the slave clock's output. However, packet delay variation is still an important source of phase error and must be controlled in order to provide a clock with frequency and phase characteristics near those of the master clock's reference signal.

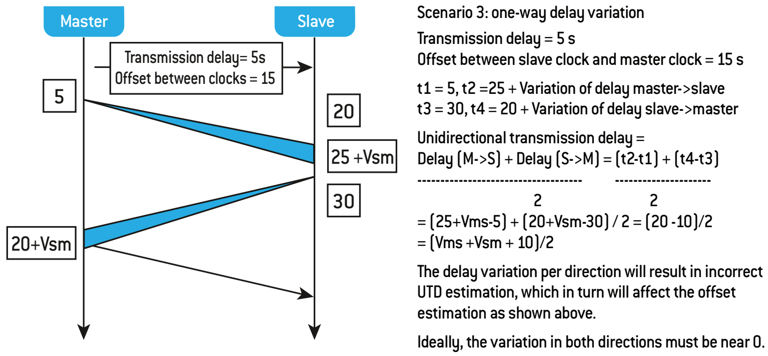

Figure 3. PTP scenario: the impact of one-way delay variation.

Delay asymmetry and delay variation are the two most important PTP packet metrics. They will determine the success of a PTP deployment, and most importantly the sync performance of the network using PTP. For network engineers, controlling these parameters is a significant yet attainable challenge:

Delay asymmetry is essentially a result of the path taken by the traffic in each direction. Network designers must ensure that the path distance in both directions is as similar as possible to eliminate any form of extra delay due to excess distance. This assessment must be performed for the active pipes and the alternate pipes that PTP traffic can take in case of protection switchover.

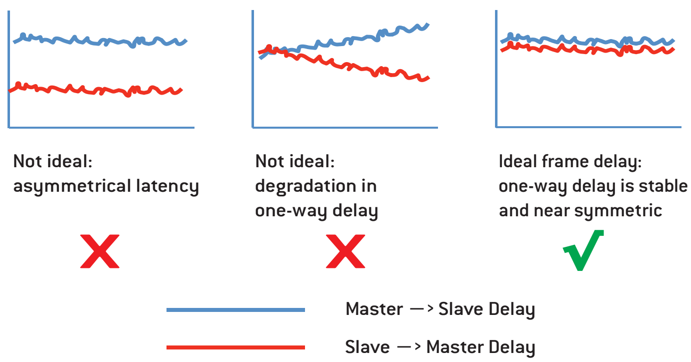

Figure 4. Interpreting one-way delay results.

Delay variation will essentially be due to the number of devices traversed by the PTP traffic and the processing time in each device. With each hop, a variable component of delay due to processing will ultimately create delay variation at the destination. In order to reduce this variation, a typical approach is to configure network devices such as they recognize PTP traffic and process the PTP packets as fast and with the highest priority possible.

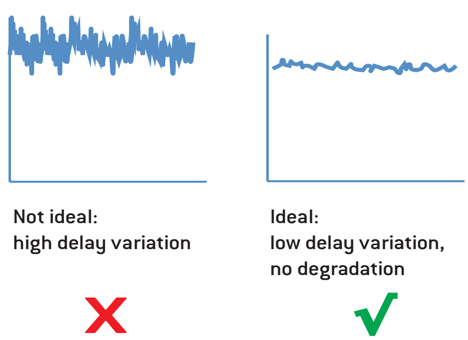

Figure 5. Interpreting delay variation results.

Even with these practices in place, it is impossible to avoid asymmetry and delay variation. The key to success is to control these parameters and maintain them as stable as possible. In fact, we can even consider that only the vendors who qualified their systems and have set maximum operational limits can guarantee at what levels or asymmetry and delay variation their system will work, hence the need to qualify the network first.

The PTP ecosystem

The grand-master clocks are the core of the PTP architecture, and the holders of the reference in the network. These devices are typically synchronized to a high-quality, high-precision reference signal–typically a GPS-based signal or a reference derived from high-stability primary reference clocks, traceable cesium clocks and output packets with timestamps to the edge slave clocks.

The ordinary clock, or slave clock, is typically found at the end node where synchronization is needed. The ordinary clock typically communicates to higher hierarchy devices and typically outputs a clock signal whose frequency and phase characteristics are similar to the grand-master reference clock.

As PTP is sensitive to packet-delay variation, intermediate devices known as boundary clocks and transparent clocks are used to provide jitter correction capabilities and act as master clocks for attached PTP slave clocks. Boundary and transparent clocks have different characteristics in terms of intrusiveness:

- Boundary clocks typically serve multiple PTP ordinary clocks or boundary clocks and are typically connected and synchronized to grand-master clocks. They act as interface between separate PTP domains and terminate PTP flows, and as master to the connected clocks. These devices have the unique capability of intercepting and processing all PTP messages and correcting the timestamps in the PTP message in order to reduce the jitter component as PTP traffic flows through multiple nodes.
- Jitter is mostly due to the buffering and transfer time of packets as they traverse network devices. Transparent clocks act as bridge between PTP domains and remove the jitter component by adding a packet residence time in PTP messages, thus allowing the slave clocks to remove the residence time and only measure physical latency.

DEPLOYING PTP IN FIVE PHASES

Anyone considering PTP deployment faces a steep challenge, from the strict requirements on packet delay and packet-delay variation to simply the lack of information and expertise on this new technology. With proper methodology though, deploying PTP can be well within the reach of most operators and entities interested in it.

The objective of this document is to discuss a method that efficiently enables operators to implement and deploy PTP within their existing network, making use of the tools already available.

The proposed method consists of five phases, each with a key element to address and a specific objective to be met before jumping to the next phase, with the goal of having a fully functional and stable synchronization network based on PTP when reaching the last phase.

The five phases

1. **Network survey**: In this phase, the operators answer the critical questions related to the one-way performance that can be experienced by the PTP flow in the exisiting networks and paths. This first step is of double importance. First it acts as a network audit to identify weaknesses that can be addressed before investing in a PTP solution. Second, it produces a powerful document which we will refer to as the PTP stress report, that provides the one-way metrics that must be handled by the vendors' solutions.

   This PTP stress report is the end result of this phase. This document will benefit everyone as it enables the vendors to propose the most effective solution from the start of the tender process, in turn helping operators make better use of available time and resources throughout their lab trials.

2. **Vendor benchmarking**: In the tender process, vendors must comply with the requirements of the PTP stress report. However, due to the nature of the PTP process, the quality and the stability of the output clock for a PTP system is also dependent on the algorithm implementation, a vendor-proprietary aspect of the design. This means that two solutions can have different clock performances, even under the same network conditions.

   The objective of this phase is to submit the selected vendors to a benchmarking comparison, enabling the operators to identify the strengths and weaknesses of each solution according to their network reality. This last point is important, as specifications do not necessarily lead to the same performance in an operating network.

3. **Service turn-up and PTP rollout**: The vendor has been selected and it is finally time to deploy. Operators can face a great challenge in this phase: PTP packets are transmitted just like another network service that must be properly configured and deployed across the network and through externally controlled networks.

   Once the service has been verified as properly configured, the clocks must be put in service and performance must be tested to provide a first baseline and correct as many errors as possible. This type of sync testing is typically known as the sync audit and closes with the acceptance of the sync deployment by the operator.

4. **Monitoring**: A key challenge for operators is ensuring that the network is performing at the expected levels of quality and stability. Unfortunately, PTP synchronization is unpredictable and continuously evolves with the network, since it is sensitive to packet delay and packet-delay variation, which in turn are also affected by network conditions.

   Operators must monitor the two layers associated with the PTP synchronization network. Monitoring the packet layer mostly consists of continuously assessing the one-way performance and ensuring that the network's metrics are within acceptable levels; synchronization monitoring ensures that the all-important clock signal generated by the slave clocks is still within the frequency and phase requirements of the network.

   Combining these two layers of monitoring provides powerful information for network operators, integrating the packet data aspect and the sync performance into a unified view and control. The benefits are outstanding, from faster resolution of issues to proactive reaction to sync degradation and more trust in the sync network.

5. **Troubleshooting**: Even with proper monitoring, failures can still occur—it is simply inevitable. The whole purpose of the monitoring system is to alert before degradation reaches a dangerous stage and starts to affect the stability of the network. Troubleshooting enables operators to find, identify and isolate the source of sync issues in order to complete corrections and restore stability. However, PTP deployment introduces new troubleshooting challenges that we address in this guide through a simple method for time-efficient troubleshooting.

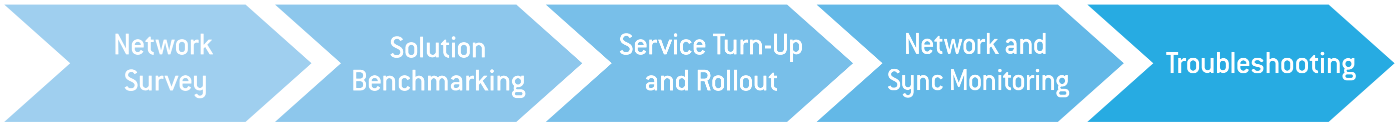

Figure 6. The five phases of PTP deployment.

The proposed tools

In order to implement this methodology, we made it a requirement that the measurement tools are readily available today. From a PTP deployment perspective, the selected tools are:

- FTB-8130NGE: This multiservice test tool offers 10 Base-T to 10G Ethernet testing capabilities in one module that uses two key applications during the whole process:
  - EtherSAM (Y.1564) enables test personnel to quickly and efficiently perform service turn-up and is used as a performance measuring tool.
  - One-way delay option: This critical option enables users to measure unidirectional delay using an innovative reference based on a globally available timescale such as GPS or CDMA signals. In this approach, one-way delay is combined with the EtherSAM suite to provide a complete and powerful testing solution and produce the PTP stress report.
- SyncWatch-110: This powerful unit is designed for network synchronization testing and monitoring in legacy, SyncE and PTP environments. The SyncWatch-110 provides stable and accurate results for PTP and clock sync signals through the use of GPS or rubidium-based references. It can also be used for monitoring sync through the managed option, enabling operators to maximize their investment.
- Brix Monitoring: The Brix Service Assurance platform provides continuous monitoring of Ethernet performance via its combination of network devices and powerful data storage and correlation engine. Using devices synchronized with GPS or CDMA signals, performance dashboards and custom SLAs, the Brix system enables operators to perform continuous assessment of the one-way performance of PTP services.

Phase 1: network survey

A key question for anyone considering PTP is the viability of this technology in their existing network. As discussed in the previous section, PTP is extremely sensitive to packet delay and packet-delay variation, and these two parameters must be well controlled to enable proper operation.

It is therefore critical to measure these two key network parameters in the path used by PTP flows and determine their levels before performing any other step, as it delivers key elements for the PTP challenge:

1. Testing and gathering results provides real-world values that can be used in the tender process, establishing minimum network conditions that the vendors must meet to achieve a fully functional and reliable PTP deployment.
2. Determining these values will determine if the network operator can use the current paths or must dedicate alternate paths for PTP flows, a question that can have serious OPEX implications.

Y.1564 is an ideal tool to make this critical assessment as it provides fast and reliable methodology to measure performance metrics on a per-service flow in a multiservice environment. Since one-way performance is critical, EXFO's implementation of this standard, EtherSAM, is the ideal test method as it provides Y.1564 methodology with one-way performance assessment via the dual test set method and the one-way latency measurement capabilities.

Using this application, network operators can qualify the path used by the PTP flow, the PTP flow service configuration, as well as PTP multiservice performance using a single and unified test method.

- **Per-direction performance assessment**: EtherSAM can be executed in a dual test set mode where two units collaborate and provide one-way performance assessment of the network KPIs of throughput, frame loss, packet delay and packet-delay variation. In the case of PTP, the last two KPIs are the most critical as they directly influence the performance of the PTP flow.
- **PTP flow validation**: EtherSAM generates synthetic test traffic and measures performance on this flow. In a typical PTP deployment, PTP traffic is configured and treated as another network service and protected by an SLA that guarantees its performance. EtherSAM uses this SLA concept to measure the performance of the flow under committed information rate or in the allotted bandwidth for PTP performance in the tested pipe.
- **PTP validation under multiservice conditions**: Since the PTP flow will be protected and serviced under an SLA, PTP services are also given the highest priority and must be treated as such by the network devices. But with the high number of devices on the forwarding path and the equally high possibility of error or mismatched configuration, the PTP flow must be validated in multiservice conditions; it is therefore mandatory for network operators to verify the performance of the PTP flow under multiservice conditions and ensure that traffic flows are not affecting negatively the PTP flow.

PTP stress report generation via EtherSAM

Using EtherSAM, network operators can perform data gathering over the PTP service pipes through two methods:

1. Using EtherSAM with a single PTP service configured, we can proceed to test a PTP flow inside a path that is already carrying customer traffic. This in-service method provides the benefit of measuring real-world metrics of a possible PTP upgrade.
2. The second method is to simulate a multiservice scenario via EtherSAM, by simulating a PTP service and adding background streams that consume bandwidth. This scenario is ideal for service turn-up of for dark fibers that could eventually be used for PTP service. This method also has the benefit of providing metrics on background traffic, thus proving the PTP flow but also the effect of the PTP flow on these background services during congestion phases.

In both scenarios, the testing begins by deploying two EXFO FTB-8130NGE modules at each end of the path where PTP traffic will eventually originate and terminate. This could be at the central location, where the grand-master clock could be located, and at a base station, where the PTP flow will terminate. Between these two points, a specific path should be configured with the proper identifiers for the PTP flow, either through VLANs or IP TOS/Diffserv or any other differentiation method, in order for the network to differentiate between existing traffic and the PTP flow, and then enforce the proper quality of service.

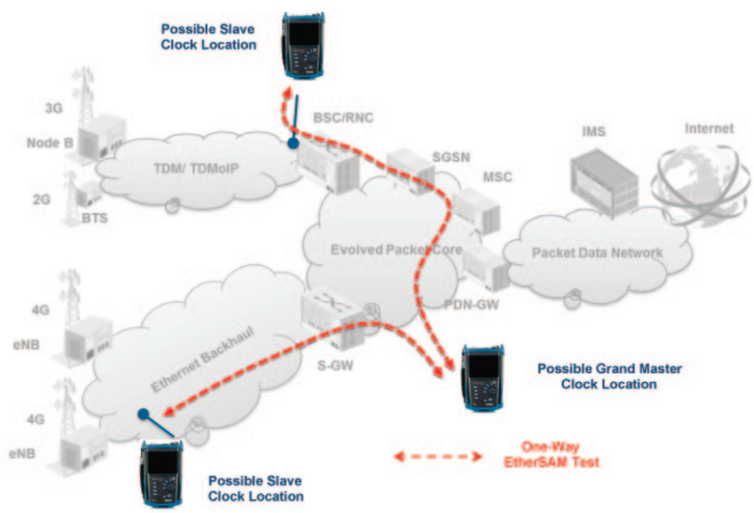

EtherSAM implements the two phases described by the Y.1564 standard, namely the service configuration and service performance tests. Both phases should be executed as they each provide benefits.

The service configuration phase tests each service individually, ramps up traffic from a low level to the committed information rate, and then tests the rate-limiting enforcements. This provides the benefit of measuring performance at low traffic utilization and at maximum traffic utilization (CIR rate). We can then assess if there is a performance degradation between this range of throughput and make a proper performance assessment. The rate-limiting test determines if the service has been properly configured in the network.

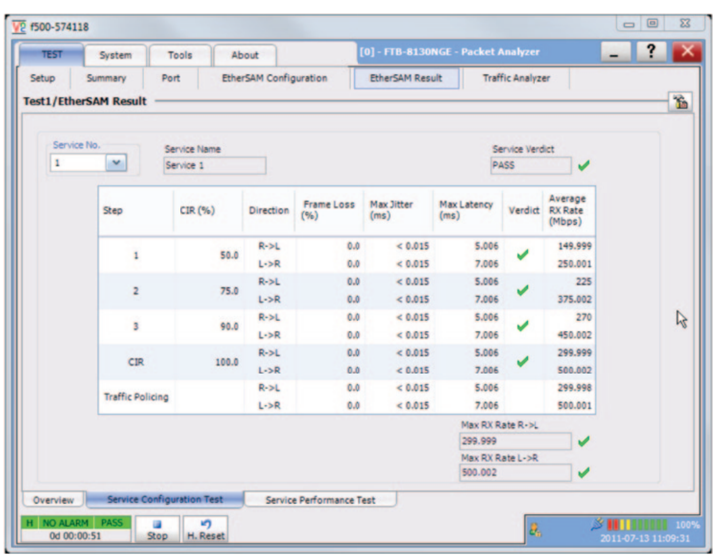

Figure 7. EtherSAM service configuration test with one-way delay.

The **service performance phase** is designed for long-term soaking and multiservice assessment. The enabled services are generated for a defined soaking period and all KPIs are assessed for all services simultaneously, allowing to measure the stability of the PTP flow and the effect of background flows on its performance. This is the key phase for PTP as it provides minimum, maximum and average metrics for a longer period of test, allowing for proper network assessment.

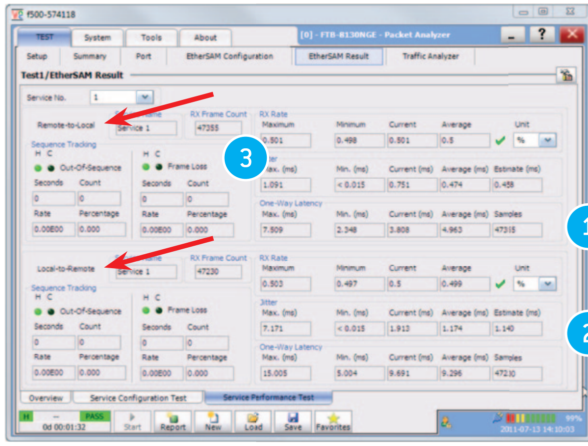

Figure 8. EtherSAM service performance results with one-way delay and delay variation.

As each direction is tested individually, performance metrics are reported for each direction:

- Packet-delay variation measurements displayed in the jitter section
- One-way delay measurements

For each metric, the following values are reported:

- **Per-direction metrics**: Information must be displayed per direction in order to assess the asymmetry of the delay and its variation; this is the most critical aspect of this testing phase.
- **Minimum and maximum jitter**: The peak delay variation per direction, determining the floor and ceiling reached during the test period.
- **Average**: Delay distribution can consist of a high number of low packet-delay variations and a low number of high packet-delay variations, usually some outliers from a normal distribution. Relying simply on minimum and maximum analysis does not provide a complete assessment; the average delay or delay variation determines the typical performance experienced by the PTP flow.

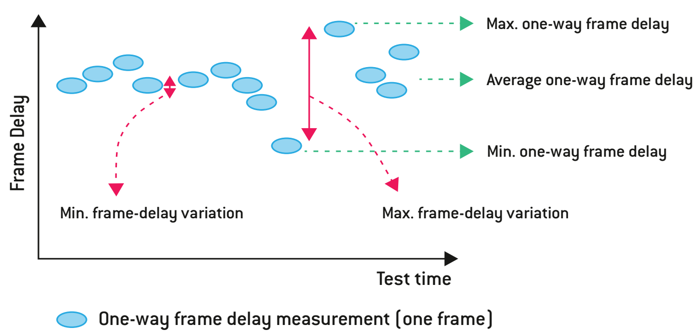

- **Frame loss ratio**: IEEE 1588v2 PTP is designed to recover from levels of frame loss; as such, frame loss therefore is not a major problem. However, for network designers, it is important to measure the frame loss ratio. If frame loss is detected or a high ratio is detected, it is important for network designers to determine the cause as this can be an indication of incorrect settings or network configuration issues.

Once the test is complete, the user should save a complete report. This will provide traceability for future testing but most importantly serve as a network performance document for the tender process. Vendors will appreciate this document as it enables them to offer the proper products and solution to fit the customer's reality, while network operators will benefit from proper solutions and reduced time and capital spent on testing.

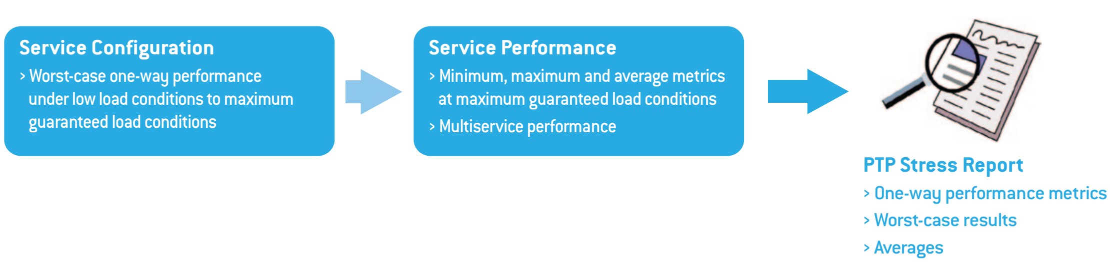

Phase 2: benchmarking the selected grand-master clock and slave clocks

Once the tender process is complete, hopefully you will have received offers and retained vendors for your PTP deployment. Now comes the next stage, thatis, benchmarking the different vendors and comparing performance.

In the case of PTP deployments, this can be a very important phase as each PTP vendor implements a proprietary algorithm in order to produce a stable clock signal. It is important to benchmark and compare the performance of each vendor's system under your own network conditions and determine which system produces the most stable and highest-performing clock.

The objective of all PTP deployments is to transfer a clock through a packet network and to output to the client or edge device a clock with a frequency and phase characteristics traceable to the grand-master reference clock. Therefore, the benchmarking must be performed on the synthesized clock from the slave's clock output. Ultimately, this is the most important signal in the PTP system.

Benchmarking these different implementations can be facilitated by the use of performance masks such as the maximum time-interval error (MTIE) clock masks defined by the different standards. The purpose of these masks is to provide a clear pass/fail assessment of a clock's performance based on MTIE data. Thus, two different PTP implementations can be tested in the same network for the same period of time, and the performance and stability of their output clock can be measured and compared to ensure that network planners make the proper decisions and justify them with real-world performance data.

Testing sync performance with the SyncWatch-110

The EXFO SyncWatch-110 network performance tool combined with the SyncSMART application enables network engineers to measure clock performance via TIE and MTIE assessment. The SyncWatch-110 is a powerful test solution that supports both PTP and legacy sync metrics.

In the case of PTP, the slave clock will typically output a timing signal to the edge client that needs synchronization. This timing signal is critical and must be measured for MTIE performance to determine that it meets the clock timing and accuracy levels required for network stability.

The SyncWatch-110 enables operators to carry out such performance measurements against a host of stable references for the highest possible accuracy. The typical approach to testing sync performance requires a stable reference and a test time that can take between 24 and 48 hours of continuous data measurements, ensuring that the long-term stability of the clock under various network loads can be estimated.

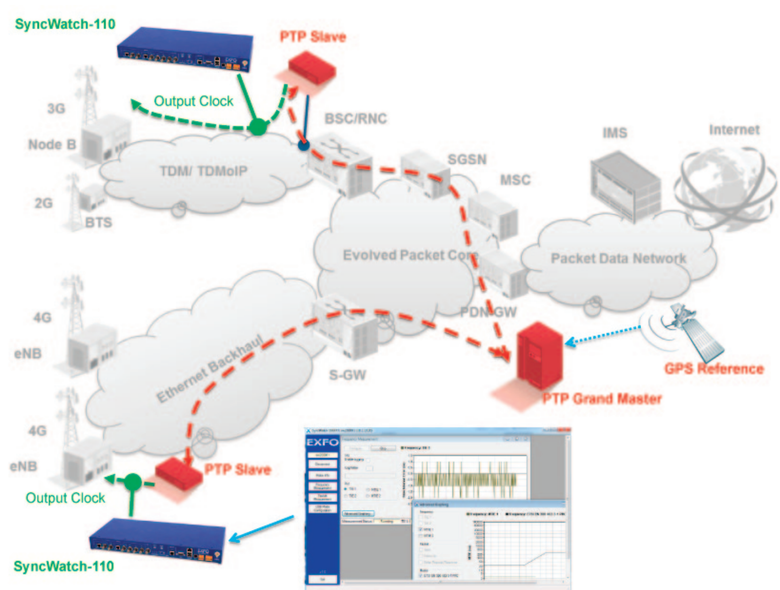

The testing phase will focus on MTIE assessment, and ensure that the output clock of the PTP slave device is in the pass region of the expected performance mask. This output clock can be either the signal sent to the towers, as is the case for most 2G and 3G deployments, or be obtained from a clock monitoring port from the terminating node when the client is embedded in the edge device, a typical occurrence for the modern eNode B in 4G architectures.

The core document at this phase is the MTIE report, representing the stability and accuracy of the output clock and showing if the signal passes the performance mask:

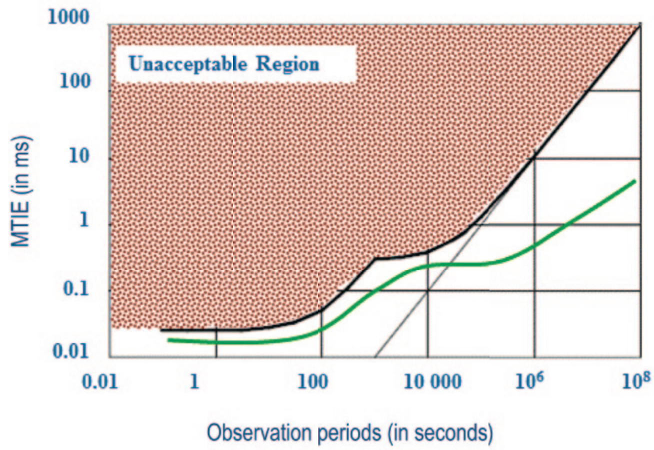

Figure 9. MTIE masks pass/fail.

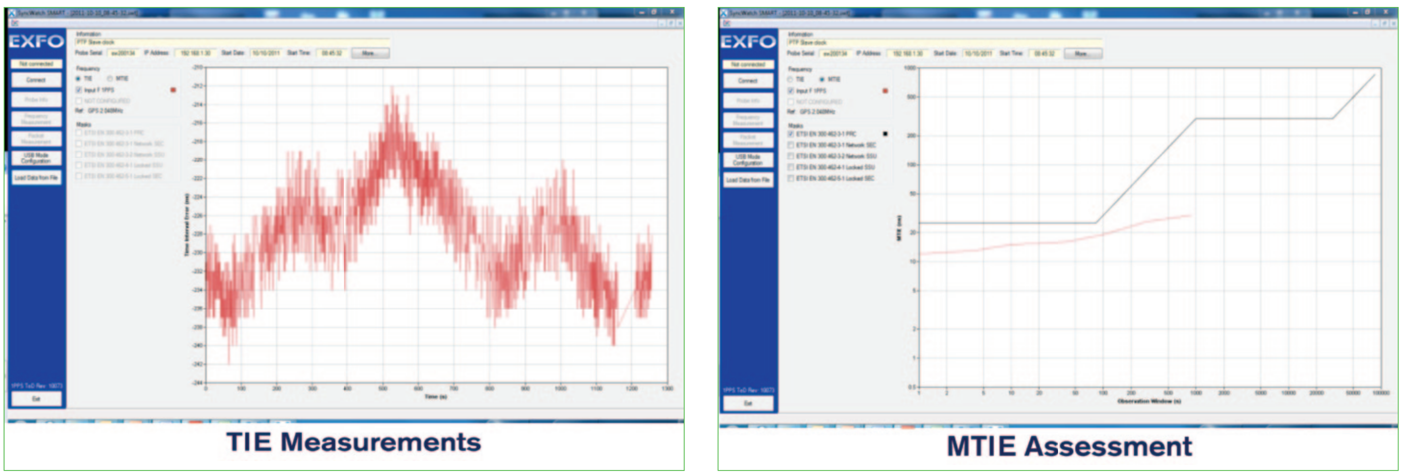

Figure 10. PTP slave clock output results during benchmark.

Phase 3: service turn-up and referencing

A solution has finally been selected that fits all network requirements and provides the needed sync performance. However, a major part of the work effort is now deploying all the PTP pipes for this new system as well as the associated protection channels, and ensuring that these services are well configured in the network.

A PTP flow will typically be protected and serviced under an SLA that will guarantee minimum performance objectives. Network operations must now include deploying the proper configurations, dispatching teams and ensuring that the services deployed do not affect current customers.

This type of action, called the service turn-up, can easily be completed via the EtherSAM test application. Designed from the beginning as a turn-up and troubleshooting tool, EtherSAM (EXFO's implementation of the Y.1564 methodology) provides operators with a fast and efficient method to turn up services in a multiservice environnement.

Using traditional methods such as RFC 2544 could transform this into a logistical nightmare due to the amount of resources and time required to perform such testing. EtherSAM and its innovative simultaneous bidirectional testing method dramatically decrease the testing effort and reduce the test time from hours to a few minutes per service, enabling operators to limit the impact on customers.

The result of this phase is a report called the birth certificate, which certifies that the PTP services have been properly implemented in the network and provides a benchmark of the operating metrics on the network. In the case of PTP deployment, this birth certificate should focus on the end-to-end connectivity between the network location of the edge clock and the network location of the grand-master clock, the one-way characteristics between these two points, as well as the impact of the current services on the PTP flow.

Once the birth certificate is issued for each service and approved by the proper personnel, PTP services can be activated. A good practice at this point is to perform a similar turn-up test for all PTP clients and benchmark the performance of the end clocks. It is a good moment to involve vendors in the deployment process and make use of their experience to iron out any outstanding issue.

Once everything is validated, the PTP sync network is ready to be put in service.

Phase 4: monitoring

The PTP system is now deployed on your network–network synchronization is achieved and customers and top management are happy. However, the challenge does not stop here. How do you maintain synchronization? Since synchronization degrades over time, as network loading increases and equipment failure occurs, how do you ensure that you react before a catastrophic sync failure affects the quality of experience and ultimately costs you business?

For these reasons, it is imperative, once synchronization is deployed, to continuously monitor its performance in order to ensure a stable network. This continuous monitoring can occur at two key layers.

› **At the sync layer**: The objective of the PTP system is to transfer a signal with the same frequency and phase requirements as the grand-master reference clock to the edge device through the packet cloud. The performance of the slave clock is not only affected by the packet layer, but also by the algorithm used to synthesize the output clock. Other events such as mechanical failures can cause sync loss. It is important to monitor synchronization at all times to guarantee that it stays within the acceptable levels.

› **At the packet layer**: Since one-way delay and delay variation, which highly influence PTP performance, are continuously evolving with different load conditions on the PTP path, network events such as congestion, path changes and network errors can cause changes in the levels of latency and the spread of packet-delay variation which can immediately affect the performance of clock synchronization.

EXFO's ecosystem provides monitoring capabilities at both layers through the Brix Service Assurance platform and SyncWatch-110 Synchronization Testing Unit.

Synchronization monitoring via SyncWatch

The EXFO SyncWatch can be configured in Managed mode where multiple probes perform real-time sync measurements and store data on a centralized server. This provides a global view of the sync performance across the network and enables proactive reaction to sync degradation–a key benefit for operators.

In this mode, SyncWatch probes are distributed at key points across the network. In the case of PTP, since the flow terminates at the slave clock, it would be costly to deploy probes at all client sites. Instead, a valid approach is to deploy sync probes at a certain percentage of towers serviced by the same boundary clocks or the same MSC. In this approach, the sync health at these locations is considered to represent the sync performance of the region, and any failure detected by the probes would limit the troubleshooting within this closed region.

SyncWatch units can also be deployed at intermediate points such as boundary clocks, where a clock monitoring port is available for testing purposes. This provides greater visibility into the synchronization network and provides better isolation of faulty regions for troubleshooting.

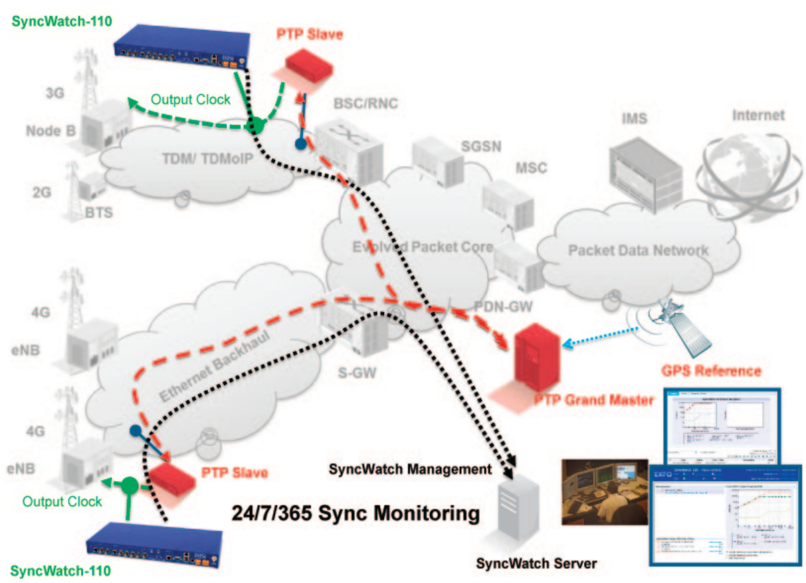

Here are some advantages of the SyncWatch's Managed mode:

› **Distribution of probes**: Any number of probes can be deployed as the server and bandwidth requirements are efficiently managed. All probes collect and transfer data and alarms/errors to the server with instant notification of failures for quick reaction.

› **User-defined thresholds**: Operators can define warning thresholds below the "fail" region of the various performance masks, which ensures that operators are notified of sync degradation before actual failures occur.

› **Simple management**: Global view and focused view of results allows operators to quickly pinpoint regions at fault through a simple probe inspection within a virtual map and then focus on a specific probe for current fault analysis and historical analysis of the sync performance.

PTP flow monitoring with Brix Service Assurance

EXFO's Brix System provides 24/7 monitoring of the network performance via a centralized monitoring system. Consisting of a central data repository and a series of data collection agents called Verifiers, the Brix System provides a complete view of the network's health and performance. In the case of PTP monitoring, the Brix System can perform one-way delay and delay variation measurements between Verifiers equipped with either GPS or CDMA receivers for increased accuracy.

Monitoring of these KPIs can be performed via a variety of active tests such as the one-way active performance, Ethernet Delay and Y.1731 OAM DMM/DMR tests. In each case, the monitoring is achieved by generating synthetic traffic and measuring one-way performance. All results are stored in the central server, allowing for a clear view of all one-way performance metrics across the network.

One-way performance is compared to SLA criteria for pass/fail assessment, enabling the operator to quickly determine if failure occurred and pinpoint the span where one-way failure has been detected. This provides automatic failure notification, resulting in more efficient use of troubleshooting resources. Another benefit is historical tracing of results, as a circuit's performance can be analyzed for anomalies or corrections.

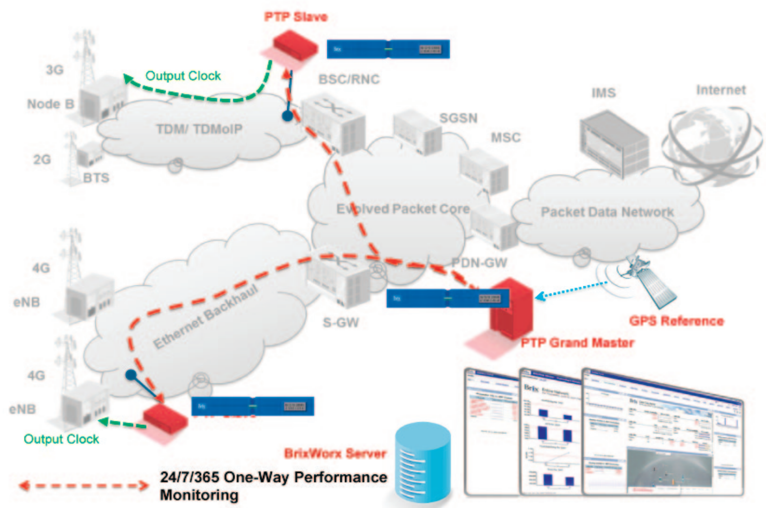

The one-way delay results are mostly similar to those described above, allowing operators to measure delay asymmetry and monitor the spread of one-way delay or packet-delay variation:

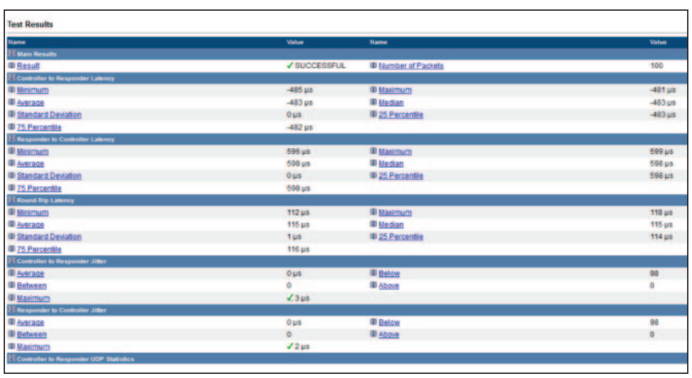

Test Results

Name	Value	Name	Value
Main Results			
Result	✓ SUCCESSFUL	Number of Packets	100
Controller to Responder Latency			
Minimum	-485 µs	Maximum	-481 µs
Average	-483 µs	Median	-483 µs
Standard Deviation	0 µs	25 Percentile	-483 µs
75 Percentile	-482 µs		
Responder to Controller Latency			
Minimum	596 µs	Maximum	599 µs
Average	598 µs	Median	598 µs
Standard Deviation	0 µs	25 Percentile	598 µs
75 Percentile	598 µs		
Round-Trip Latency			
Minimum	112 µs	Maximum	118 µs
Average	115 µs	Median	115 µs
Standard Deviation	1 µs	25 Percentile	114 µs
75 Percentile	116 µs		
Controller to Responder Jitter			
Average	0 µs	Below	98
Between	0	Above	0
Maximum	✓ 3 µs		
Responder to Controller Jitter			
Average	0 µs	Below	98
Between	0	Above	0
Maximum	✓ 2 µs		
Controller to Responder UDP Statistics			

Figure 11: Sample one-way delay monitoring result.

Performance Matrices

Maximum Jitter To Responder of One-Way Performance Active Test in Silver One-way Performance (milliseconds)

Controller \ Responder	BER-BV1000-5212	IAD-BV1000-6864	LAX-BV1000-6608	LHR-BV1000-5403	ORD-BV1000-5497	PEK-BV1000-6827	SIN-BV1000-9114	YVR-BV1000-9127
BER-BV1000-5212		0.00875	0.0094	0.097	0.0018	0.03	No Data	0.00875
IAD-BV1000-6864	0.006		0.0113333	0.1314	0.0066	0.0218	No Data	0.005
LAX-BV1000-6608	0.0064	0.021		0.0688	0.0042	0.02725	No Data	0.014
LHR-BV1000-5403	0.0516	0.0334	0.0462		0.006	0.0575	No Data	0.0425
ORD-BV1000-5497	0.0022	0.0226	0.004	0.0648333		0.0535	No Data	0.015
PEK-BV1000-6827	0.0093333	0.0098	0.01375	0.0095	0.009		No Data	0.0238
SIN-BV1000-9114	No Data	No Data	No Data	No Data	No Data	No Data		No Data
YVR-BV1000-9127	0.00775	0.0053333	0.0095	0.0615	0.0074	0.0256	No Data	

Figure 12. One-way delay monitoring matrix.

Perfect PTP monitoring: combining sync and network monitoring

When monitoring both layers, operators gain a tremendous benefit: the capability to continuously monitor both PTP and sync performance and ensure that these layers are under proper operational conditions. Therefore, they can:

- React faster in case of failures and simplify troubleshooting; network engineers can quickly determine if sync failure is due to packet performance or concentrate their efforts on the sync layer.
- Prevent failures by detecting degradation at any layer.
- Have better assurance in the performance of the network.

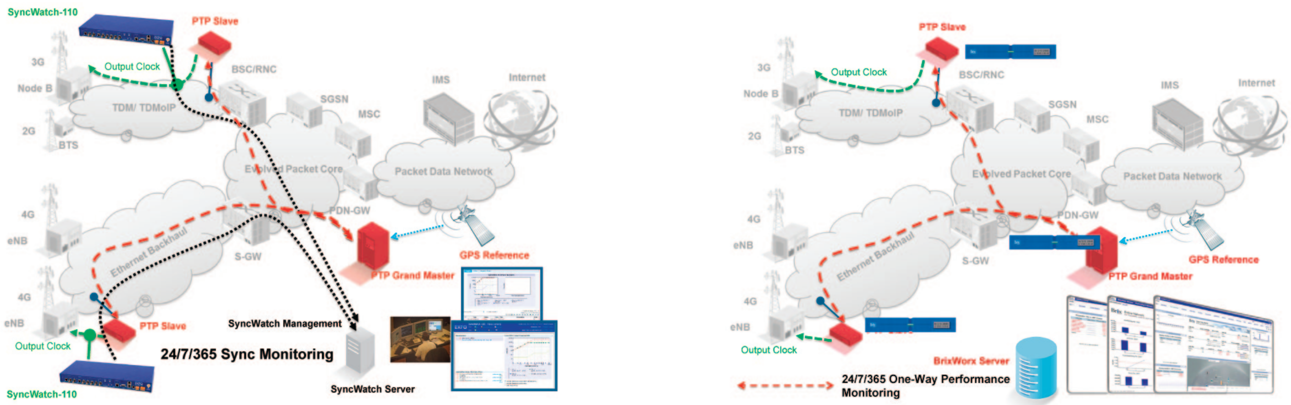

Figure 13. Best-in-class monitoring: sync and network performance monitoring.

Phase 5: troubleshooting failures

When a sync failure is detected, how can we quickly troubleshoot and determine the possible cause of problem and solution? In the case of PTP, the impact of packet performance and the reliance on the client algorithm make troubleshooting more complex.

In a typical PTP scenario, the possible causes of failure are numerous, ranging from the equipment clock to reference loss. For each item, the SyncWatch can be used to perform measurements and isolate faults.

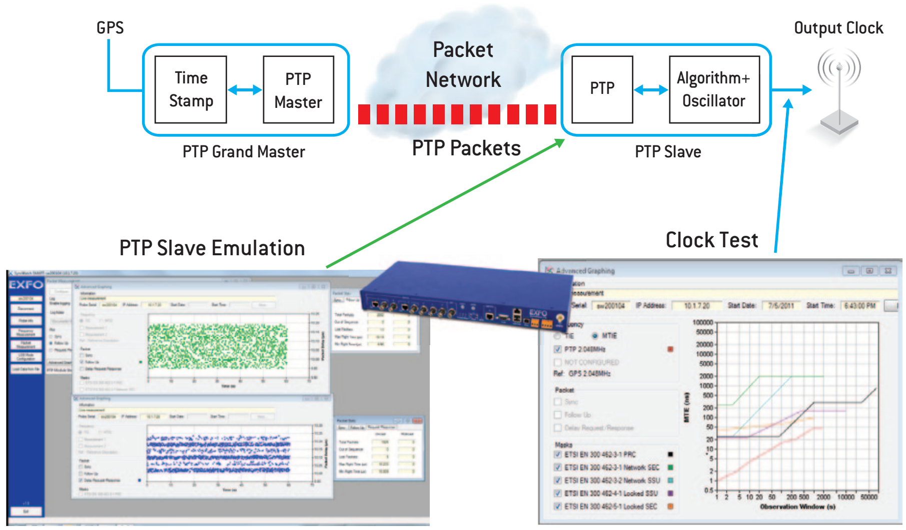

Figure 14. PTP troubleshooting via SyncWatch.

Possible cause no. 1: PTP flow asymmetry or high packet-delay variation

The SyncWatch should be connected either through the same Ethernet link as the slave under investigation or to the same Ethernet switch. It will then simulate a slave clock, connect to the grand-master clock and perform one-way delay and delay variation measurements using the PTP protocol.

The one-way performance measurements can be analyzed or compared to the measurement performed during the service turn-up phase, and any anomaly or important discrepancy can be identified as the possible cause.

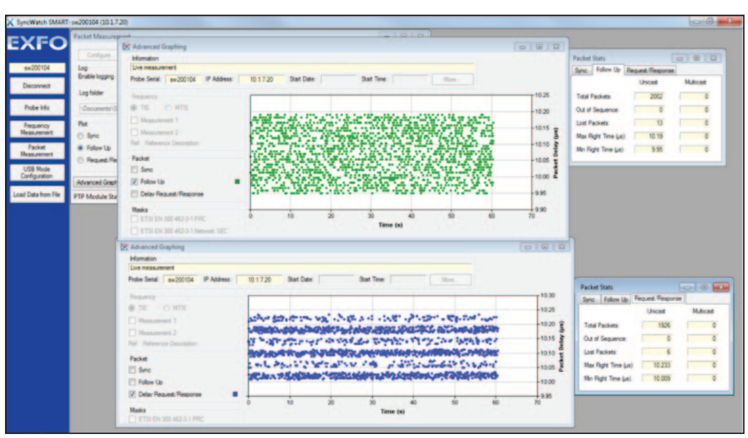

Figure 15. PTP slave emulation results.

Possible cause no. 2: network degradation

If PTP flow asymmetry or high packet-delay variation is identified as the possible cause, a troubleshooting step would be to isolate if a segment is at fault. To perform this simple troubleshooting, SyncWatch can be used to test at intermediate points between the grand-master clock and slave clock under investigation, with the objective of measuring packet performance at the nodes and identifying where the performance is degrading.

Then, more investigation will typically be needed at the faulty node to determine if the errors are caused by a configuration or loading issue. By using EtherSAM, operators can perform this type of assessment and carry out the appropriate corrective actions as needed.

Similarly, this type of testing can also be performed via the Brix System by monitoring the one-way performance between the key verifiers within the segment under investigation.

Possible cause no. 3: grand-master clock/PTP slave equipment failure

Of course, if packet delay and delay variation is well within the accepted performance agreed with the synchronization vendor, the grand-master clock or the slave clock can be considered at fault. In this case, the SyncWatch can be either connected directly to the edge slave clock, and the output clock can be measured and compared to a stable reference.

Possible cause no. 4: algorithm issues

A major dependency for PTP slave performance is the algorithm used to synthesize the output clock. Typically, slave clocks will selectively use some packets for clock adjustments in order to avoid high variations in the clock output. Since each algorithm is proprietary, performance can differ for the same network conditions.

The SyncWatch provides a simple way to troubleshoot the algorithm by offering the capability to internally synthesize a clock signal based on the PTP packets exchanged with the grand-master clock. This signal can then be compared to a stable reference for TIE and MTIE measurements.

- If performance is better with the onboard SyncWatch PTP client, then a doubt can be placed on the slave clock's implementation and should be brought up to the vendor for discussion.
- If performance is similar or not meeting the MTIE performance masks, then more investigation is needed and focus should be placed on network performance.

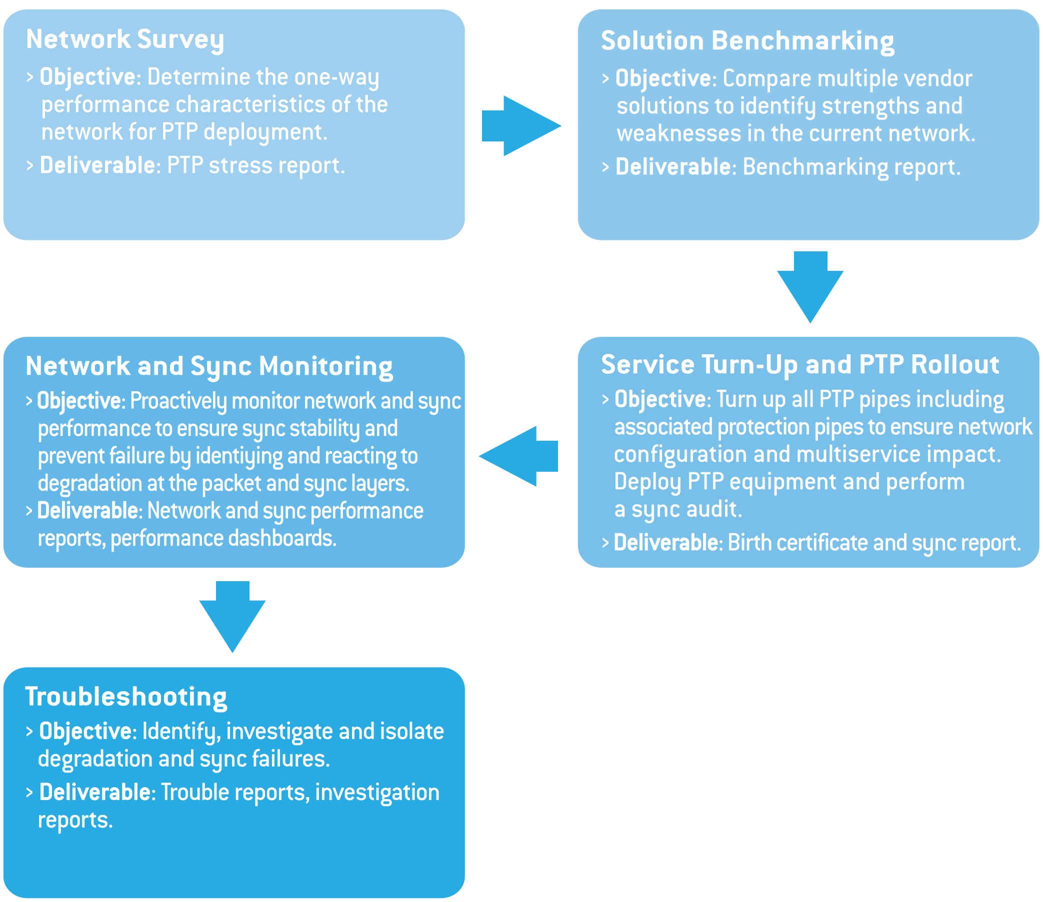

Conclusion

IEEE 1588v2 PTP is a promising technology which now delivers the synchronization capabilities required to enable the new mobile technologies. Although IEEE 1588v2 is complex to implement, we can reflect on this document and notice that we have the knowledge and the tools to understand, master and deploy IEEE 1588v2 PTP in our networks.

There has been quite a surge in interest in PTP over the last few years as the industry is looking to deploy all-packet networks. Experts have studied the PTP challenge and new ideas are being put forth to facilitate the deployment and operation of PTP as a sync technology. A few examples:

- PTP devices are now becoming mainstream applications instead of niche products, as more and more vendors are entering this space.
- There is growing interest and activity within sync-focused groups such as the International Telecom Sync Forum (ITSF) and IEEE 1588 Conformance Alliance Program (ICAP), who aim to regroup vendors, operators and sync experts with the objective of promoting packet-based sync technologies.
- Study groups within the Metro Ethernet Forum (MEF) committee and the International Telecommunication Union (ITU-T) are actively looking at new metrics to help better translate packet performance into sync performance. Metrics such as minTDEV, BanTDEV and others have been proposed by sync experts and are currently being debated by the various committees.

All these advances and many more point to PTP as an increasingly promising technology of choice for the future, one filled with challenges but that also provides key advantages as we strive for better quality and more bandwidth.

It is in this spirit that this document is published, to promote discussion and provide a solid basis for future deployments based on EXFO's proven experience in the packet domain and growing presence in the synchronization space. This document also aims to show our commitment to developing synchronization and pursuing the advancement of packet-based sync as a mainstream technology, as we recognize the current challenge and the future opportunities in enhanced mobile communication.

ANNEX 1: PROPOSED TOOLS

MULTISERVICE TESTING

FTB-8130NGE

Features:

- 10 Base-T to 10GigE
- Rugged platform (field application)
- Powerful test capabilities:
  - EtherSAM (Y.1564)
  - Traffic generation
  - RFC 2544
  - BERT
  - Smart Loopback
  - TCP throughput
  - Packet capture
  - Traffic scan

Main application used in PTP deployment:

- EtherSAM Y.1564
- One-way delay
- RFC 2544

SYNCHRONIZATION TESTING

SyncWatch-110

Features:

- 1U, lightweight chassis
- Complete sync capabilities of legacy, SyncE and PTP
- Capable of dual test signal versus reference
- Flexible reference options:
  - External
  - Internal GPS card
  - GPS + rubidium
- SyncWatch SMART: easy-to-use control software

Main application used in PTP deployment:

- Legacy signal testing
- PTP slave emulation

NETWORK MONITORING

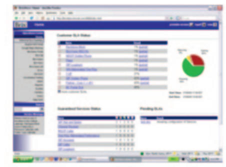

BrixNGN Service Assurance Software

Features:

- Powerful monitoring and service assurance capabilities
- Extensive reporting capabilities
- Seamless integration in network and management system

Main application used in PTP deployment:

- One-way delay monitoring
- One-way delay variation monitoring

ANNEX 2: RELEVANT LINKS

TEST METHODOLOGIES AND TECHNOLOGIES

One-way delay application note	http://documents.exfo.com/appnotes/anote255-ang.pdf
EtherSAM application note	http://documents.exfo.com/appnotes/anote230-ang.pdf
SyncWatch application Note	http://documents.exfo.com/appnotes/anote243-ang.pdf
Sync testing short guide	http://documents.exfo.com/Misc/SyncWatchReference%20Laminate-ang.pdf

TEST TOOLS

FTB-8130NGE Multiservice Test Module

Product page	http://www.exfo.com/en/Products/Products.aspx?Id=370
Spec sheets	http://documents.exfo.com/specsheets/FTB-8120NGE-8130NGE-angHR.pdf

SyncWatch-110 Synchronization Testing Unit

Product page	http://www.exfo.com/en/Products/Products.aspx?Id=472
Spec sheets	http://documents.exfo.com/specsheets/SyncWatch110-angHR.pdf

BrixNGN Next-Generation Service Assurance Software

Product page	http://www.exfo.com/en/Products/Products.aspx?Id=410
Solution overview	http://documents.exfo.com/specsheets/BrixNGN-angHR.pdf
The Brix System	http://documents.exfo.com/specsheets/BrixSystem-angHR.pdf

EXFO Corporate Headquarters > 400 Godin Avenue, Quebec City (Quebec) G1M 2K2 CANADA | Tel.: +1 418 683-0211 | Fax: +1 418 683-2170 | info@EXFO.com

Toll-free: +1 800 663-3936 (USA and Canada) | www.EXFO.com

EXFO America	3400 Waterview Parkway, Suite 100	Richardson, TX 75080 USA	Tel.: +1 972 761-9271	Fax: +1 972 761-9067
EXFO Asia	100 Beach Road, #22-01/03 Shaw Tower	SINGAPORE 189702	Tel.: +65 6333 8241	Fax: +65 6333 8242
EXFO China	36 North, 3rd Ring Road East, Dongcheng District Room 1207, Tower C, Global Trade Center	Beijing 100013 P. R. CHINA	Tel.: + 86 10 5825 7755	Fax: +86 10 5825 7722
EXFO Europe	Omega Enterprise Park, Electron Way	Chandlers Ford, Hampshire S053 4SE ENGLAND	Tel.: +44 23 8024 6810	Fax: +44 23 8024 6801
EXFO NetHawk	Elektroniikkatie 2	FI-90590 Oulu, FINLAND	Tel.: +358 (0)403 010 300	Fax: +358 (0)8 564 5203
EXFO Service Assurance	270 Billerica Road	Chelmsford, MA 01824 USA	Tel.: +1 978 367-5600	Fax: +1 978 367-5700

Printed in Canada 11/11

Assessing Next-Gen Networks