www.albedotelecom.com

PRP Verification and Troubleshooting Overview

The *Parallel Redundancy Protocol* (PRP) provides failsafe operation in a LAN by duplicating the transmission media so that they work in parallel. Every piece of information is transmitted twice over mutually isolated networks (LAN A and LAN B). If one of the replicas is lost, it is very likely that the the information is still received through the alternative channel. In normal operation conditions, frames are received twice and *Doubly Attached Nodes* (DANs) implement a duplication detection algorithm designed to drop most duplicated frames before they are processed by higher protocol layers.

The PRP standard is flexible enough to accept interoperability with nodes connected to only LAN A or LAN B. These are called *Singly Attached Nodes* (SANs). There are not strict requirements for LAN A and LAN B. These are typically bridged Ethernet networks designed so that a failure in one of them does not affect to the second network.

Despite its conceptual simplicity, PRP is a quite tricky technology that requires careful assessment. Some issues are hidden behind redundancy and they are only evident in faulty conditions. There are also some special features of PRP that require special attention: LAN A and LAN B ports in a DAN share the same MAC and IP addresses. They must never be attached to the same PRP LAN. DANs must generate special data flows to report the kind of node they are. They also need to discard duplicated frames. It is not mandatory that all duplicates are detected and dropped but it is still expected that most duplicates will be correctly processed. Moreover, DANs must keep a table with other DANs attached to the network in order to detect of duplicates. For PTP there are also specific requirements and standard IEC 61850-9-3, the PTP Utility profile, specifies a special version of the *Best Master Clock Algorithm* (BMCA) for PTP over PRP. These and other conditions are potential sources for faults or service degradation and they can only be prevented or corrected through specialized testing.

1. Passive Tests

Zeus provides a complete set of PRP statistics and performance indicators. Many of them are collected from the network in a purely passive, non-intrusive mode. This information offers valuable information about interaction between nodes, including DANs and SANs in both PRP LANs. Here the Zeus ability to collect information from both LANs at the same time and correlate results from them becomes a valuable tool (see Figure 1). Zeus provides generic performance parameters about PRP, similar to the results offered for any Ethernet network but it also includes counters and statistics specific of PRP.

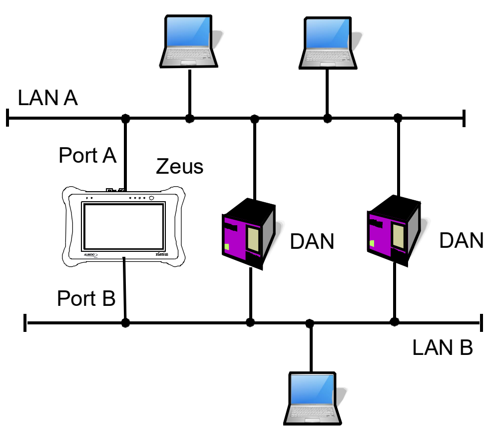

Figure 1 *Zeus interconnection to the PRP LAN A and LAN B to collect data from both networks.*

WP-PRPVERF 12/25

Traffic Statistics

The *Frame layer statistics* and *Bandwidth statistics* panels include some of the most important results collected in passive mode (see Figure 2). Both have a control that enables users to switch between Port A and Port B results collected from PRP LAN A and LAN B respectively. The Bandwidth statistics panel include a control to switch between Port A and Port B results.

Frame layer statistics

	Frames	Bytes
TX	0	0
RX	1,250,166	87,512,436
Unicast	1,250,000	
Multicast	166	
Broadcast	0	
VLAN	0	
IEEE 802.1ad	0	
Q-in-Q	0	
Control	0	
Pause	0	
Supervision frames	100	

Port A | Port B

Bandwidth statistics

	fr/s	Mbit/s	%
Eth.(current)	25003	14.0	1.8002
Eth.(min.)	2	0.0	0.0001
Eth.(max.)	25007	14.0	1.8005
Eth. unicast	25000		
Eth. multicast	3		
Eth. broadcast	0		
IPv4 (current)	0	0.0	0.0000
IPv6 (current)	0	0.0	0.0000
IPv6 (current)	0	0.0	0.0000

Port A | Port B | Units

Figure 2 *Zeus frame and traffic statistics in PRP endpoint operation mode. Independent statistics for Port A and Port B in the same screen.*

Sometimes it is useful to focus in a simple node or an small set of nodes, DAN or SAN, and collect statistics only from them. Zeus includes different kinds of filter to match the traffic from devices connected to the network:

1. *DAN* filters matches traffic from a DAN or a DAN group. You can use these filters to match traffic received from DANs connected at the same time to Port A and B by their MAC addresses, VLAN identifiers and other frame fields.
2. *SAN-A / SAN-B* filters match traffic generated from SANs that are therefore detected in LAN A or B but not in both.

Other alternatives are added to the *DAN / SAN-A* and *SAN-B* filters. For example, there is a family of filters to match traffic received from a remote loopback device with minimum configuration in the traffic generator. Users are allowed to configure up to eight traffic filters and generate custom frame and bandwidth statistics for these filters.

Traffic Capture

Sometimes it is important to verify the frame structure generated by network hosts. You may want to know if the traffic has the correct VLAN tags or if an specific field in the frame header or payload set to the correct value. The best approach to tasks related to protocol analysis is the traffic capture tool. When the capture is combined with the filters (DAN, SAN-A, SAN-B...) it becomes a powerful troubleshooting tool. If an external timing source (GNSS, 1PPS / ToD, IRIG-B) is configured, Zeus generates nanosecond accurate timestamps for each captured frame providing a powerful tool to analyze timing issues in the traffic flow.

It is worth discussing some details on the operation of the capture function with the help of an specific example (see Figure 3). When the capture is enabled, Zeus generates some statistics such as stored frames from each port and timestamps from the first and last captured frame and buffer usage. In the example, Zeus is configured to match the traffic generated from a single DAN identified by its source MAC address resulting 85 frames captured in each port. That is perfectly natural when the traffic is captured from a node that generates a copy of each transmitted frame. The capture buffer shows that frames captured in LAN A and LAN B are grouped in couples received with a time offset of a few nanoseconds. One of the members of the couple (the one that is received in last place) is detected as a duplicated frame and discarded by the PRP receiver with high likelihood. The time difference between a frame and its copy is the delay skew. If the skew is large, the duplicate detection mechanism in the PRP receiver may consider the copy a legitimate frame and forward it to the next processing block.

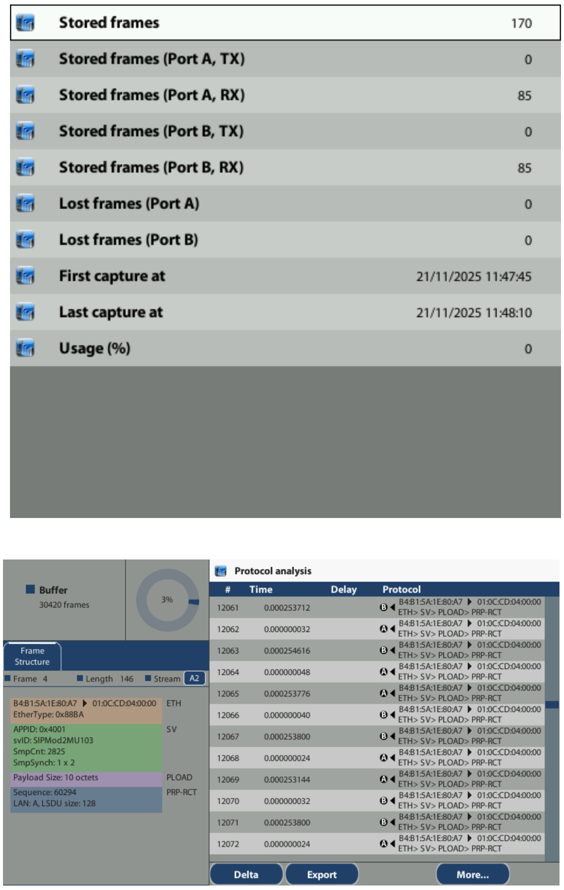

Figure 3 *Zeus capture results. Capture statistics and capture buffer showing the protocol layer structure of captured frames*

Node Population

We now start with the overview some if the most important PRP-specific results. Unlike the results described in the previous sections, these statistics wouldn't make sense in non-PRP networks and therefore they are not available in other Zeus operation modes.

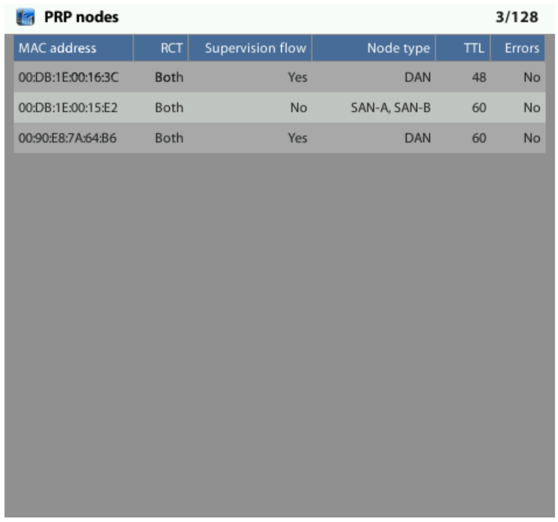

Figure 4 *Node population results. DANs and SANs are displayed in a single list. In the example, the supervision flow is missing in node 00:DB:1E:00:15:E2.*

The PRP nodes table detects DANs and SANs in the network and report information about them related to the kind of node they are, when they were last seen, or if there is any error in the frames received from them (see Figure 4). The table is populated in real time as new nodes are detected. A *time-to-live* (TTL) determines nodes to be deleted from the list. Node MAC addresses are displayed to allow users to obtain more detailed results about about specific hosts.

PRP Statistics

The PRP Statistics panels enable users to count transmitted and received PRP frames. These statistics are useful for several reasons (see Figure 5):

- To account for transmitted and received frames with and without a *Redundancy Control Trailer* (RCT). DANs must generate frames with RCT but this trailer could be lost or modified by the network. PTP is an special protocol and some PTP-aware nodes may generate PTP frames without an RCT.
- Verify that duplicates are received and dropped through the *Unique* and *Duplicate* frame counters. As explained in previous sections, PRP

receivers drop frames if they already have a copy of them. There is not noticeable effect in the service when duplicates are not received but redundancy is lost and the network is then vulnerable to faults.

PRP statistics

	A	B	Interlink
RX	2600	2600	2600
RX RCT	2,600	2,298	2600
Supervision frames	0	0	
Unique frames	1,825	775	2600
Single duplicates	775	1,523	
Multiple duplicates	0	0	
LAN mismatches	0	0	
Errors	0	0	

Figure 5 *Filtered PRP frame counts. These results focus in a specific traffic flow selected by the user. The example shows that 2600 frames are received in LAN A and LAN B. In LAN A 1825 unique frames are received and in LAN B there are 775 frames totalling 2600 frames. However, there are only 2298 duplicates (775 frames + 1523 frames), less than 2600 frames. The reason is that there are some frames received in LAN B that do not carry any RCT. Once the LAN A and LAN B traffic is merged the issue becomes undetectable.*

- Provide diagnostics by reporting different kinds of anomalies detected in PRP traffic. The tester reports if multiple duplicates are received. Single duplicates are part of the normal operation of PRP but multiple duplicates should never be found. A multiple duplicate has different consideration depending on whether it was received in the same port where the unique or the first duplicate was received. The tester also detects frames received in the wrong LAN (LAN mismatches).

2. Active Tests

In an active test, Zeus loads the network with synthetic traffic that sometimes is engineered to resemble traffic that would be transmitted in normal operation conditions and sometimes pretends to stress the network in some way to predict potential anomalies or issues.

The following sections focus on how to simulate PRP traffic, how to use this traffic to measure some key performance indicators from the network and how to tune the Zeus traffic generator to challenge the network.

DAN/SAN Traffic Generation

Zeus is equipped with a multistream traffic generation that enables users to generate up to eight independent traffic flows. Most frame fields can be edited for each of them (see Figure 6). That includes VLAN tags, source and destination MAC addresses and others. Flows can be configured to simulate DANs or SANs. In DAN flows, users can decide whether to generate the supervision flow or not. There is also the possibility to generate DAN flows from only LAN A or LAN B to simulate conditions where the traffic from one of the PRP LANs cannot reach the destination.

Frame size	70
Node type	DAN
Supervision flow	Enable
Encapsulation	None
Source MAC address from	Local profile
Source MAC address	00:00:00:00:00:00
Source MAC range	1
Destination MAC address from	Manual
Destination MAC address	00:DB:1E:00:16:3D
Destination MAC address (ARP)	00:00:00:00:00:00
Destination MAC range	1
EtherType	88B5
C-VID	1
C-VLAN priority	0

Figure 6 *Frame configuration in PRP mode. There are two settings related to PRP: The "Node type" field allows traffic generation from SANs or DANs, "Supervision flow" enables and disables the supervision from in the current stream, even if it is for a DAN.*

Synthetic traffic generated from Zeus can be analyzed by other Zeus units as explained in the previous sections (see Figure 7). The traffic statistics, capture and node population results work as if the traffic was generated from a PRP node, but synthetic traffic contains sequence numbers and timestamps that can be used to go beyond the results that can be obtained in passive tests.

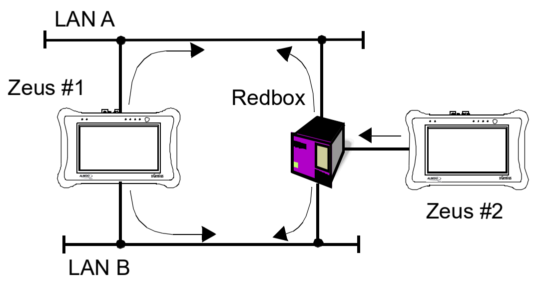

Figure 7 *Active test from two different points. Zeus #1 is generating PRP traffic that can be received by Zeus #2, which is connected to the network though the Redbox in Ethernet (non-PRP) mode.*

QoS Tests

The Ethernet *Quality of Service* (QoS) test is generalized for PRP networks. In PRP interfaces two units are normally needed for testing because each unit is equipped with two ports but four ports are required to run generation and analysis over LAN A and LAN B at the same time. This is, however, not too different to many Ethernet test setups. Particularly, two units are always required to measure delay and frame loss between two different locations.

Typical PRP QoS tests require external timing for latency measurements. Without external timing only round-trip statistics can be obtained by installing two traffic reflectors in the far end. External timing can be supplied in the same conditions than for any other Zeus test. The most common reference source is GNSS but others such as 1PPS, ToD or IRIG-B are also supported.

There is an special QoS parameter that is relevant only in PRP applications: the delay skew between LAN A and LAN B. Zeus is prepared to measure the skew generated by a DAN identified by its MAC address. Skew between LANs is important because the ability of DANs to detect and discard duplicates depends on how far apart in time they are received.

Like many other PRP tests, QoS measurement becomes more powerful when different results are combined. For example, the PRP statistics result may suggest that there is a LAN with higher delay when all duplicates are always received in the same LAN. The delay skew results can be used to confirm this point and provide quantitative results (see Figure 8).

PRP statistics

	A	B	Interlink
RX	2100	2100	2100
RX RCT	2,100	2,100	2100
Supervision frames	0	0	
Unique frames	0	2,100	2100
Single duplicates	2,100	0	
Multiple duplicates	0	0	
LAN mismatches	0	0	
Errors	0	0	

Delay statistics

	FTD	FDV	Skew
Current	8.84 µs	0.24 µs	2.61 µs
Average	8.84 µs	0.04 µs	2.58 µs
Maximum	8.94 µs	0.24 µs	2.70 µs
Minimum	8.70 µs		2.45 µs
Standard deviation	0.04 µs		0.01 µs
Range	0.24 µs		0.26 µs

Port A Port B

Figure 8 *The filtered PRP statistics for a PRP DAN show that all unique frames are received in LAN B. This result suggest that there is a lower delay in the LAN B path. The "Delay statistics" show a quite stable delay in the LAN A transmission path through the FTD and FDV statistics but the delay is around 2.6 µs larger in this LAN than in LAN B. This result is consistent with the duplicated frame statistics.*

Stress Tests

One of the most important Zeus features is that it enables users to predict network performance through different kinds of stress tests running in a controlled environment. These tests guarantee reliable transmission even in harsh operation conditions. We are going to review three different kinds of stress tests: (1) throughput, to rate the ability of a network or a network element to forward large amounts of traffic, (2) skew, by generating a large time offset between LAN A an d LAN B to check the duplicate detection algorithm in a DAN, (3) nodes test that simulates a large amount of DANs or SANs to see how the network deals with them.

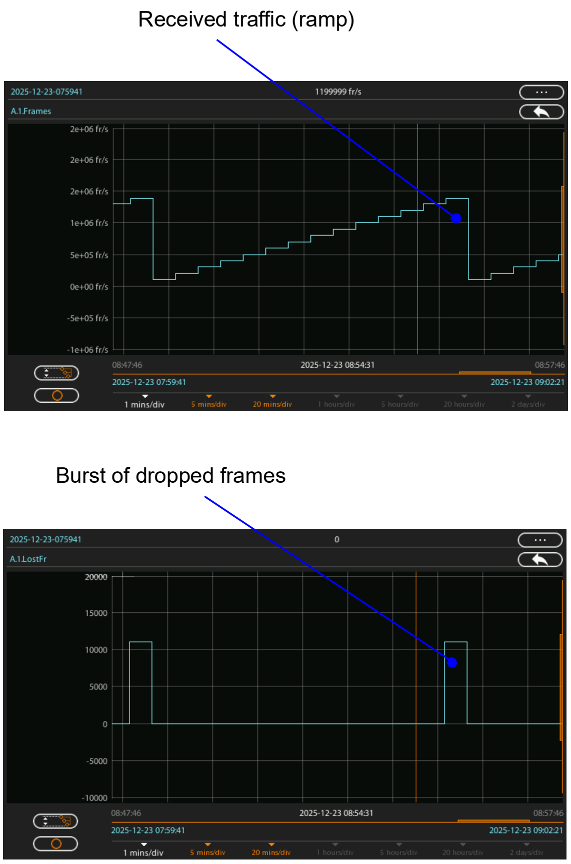

Figure 9 *Test results obtained when the interlink port of a redbox is loaded by variable traffic and measured at LAN A and LAN B (only LAN A results are shown). The received traffic has a ramp shape. If the traffic load exceeds certain limits, the DUT starts dropping packets.*

One particularly useful setup for throughput tests consists in generation of a variable traffic that increases with time. At some point, it may be difficult for the network to transmit all the frames without dropping some of them (see Figure 9). Similar tests are common in all kinds of Ethernet and IP networks. A throughput test is, for example, specified in RFC 2544 to rate network elements or services provided by packet networks. However, PRP tests have their own specific requirements. Zeus enables users to simulate large amounts of traffic from DAN or SAN devices and also to measure performance in LAN A and LAN B at the same time and to compare the results.

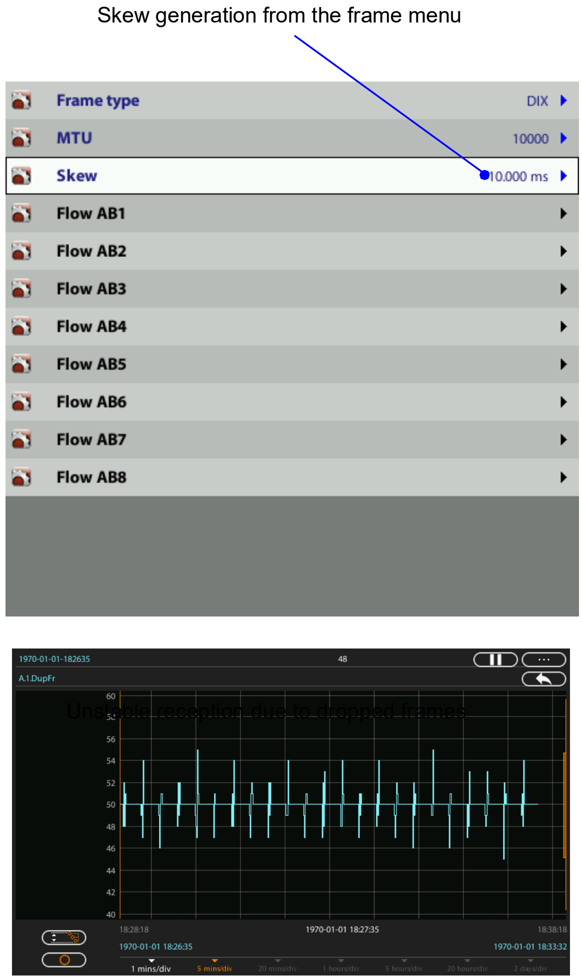

Figure 10 *Duplicated frame generation in the interlink of a redbox due to a 10 ms skew added between LAN A and LAN B. An even larger skew would cause all frames from the delayed port to be detected as duplicated or out-of-sequence frames.*

Skew is the relative delay between a packet and its copy. Since PRP receivers need to detect and drop duplicated frames, generating delay between them is a good way to stress the system. If a PRP device is not able to detect duplicates reliably, then it will forward them to an output port or to the upper protocol layers increasing the CPU load and degrading services and applications.

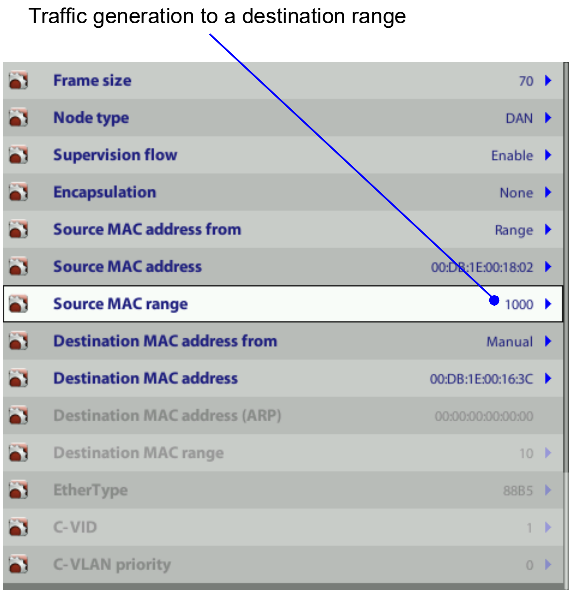

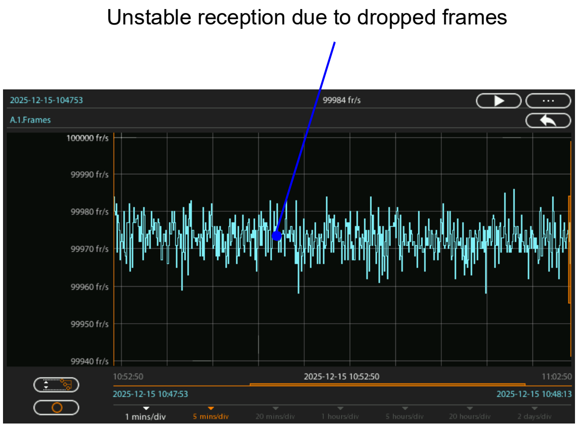

Figure 11 *Traffic results when a constant traffic of 10000 frames/s from a range of 1000 different MAC addresses is transmitted through LAN A and LAN B in a PRP redbox. The DUT is unable to process all addresses frames without dropping frames, for this reason the measured traffic smaller than 10000 frames/s*

Zeus has the ability to generate a custom amount of skew by delaying traffic from Port B from Port A. This feature, together with the duplicated and out-of-sequence frames statistics can be used to assess the effects of skew (see Figure 10). This kind of test is most useful in network elements such as redboxes because they have an input and an output interface where testers can inject traffic and analyze outgoing frames.

A different kind of stress test consists in simulating traffic from many endpoints. PRP receivers keep track of network nodes in order to detect duplicates and it may be difficult for some of them to work in large networks with many DANs (see Figure 11). Zeus can simulate traffic from a very large amount of nodes including DANs and SANs. If necessary both kinds of nodes can be simulated at the same time. This feature includes the generation of an independent PRP supervision flow for each DAN.

3. PTP over PRP tests

Timing distribution in PRP networks is a challenge because timing information is collected from two channels with two delay paths. This fact makes delay compensation unfeasible. At some point, the slave could be applying a compensation derived for LAN A in a Sync packet received in Port B. The solution to this issue is addressed in IEC 61850-9-3, the PTP Utility Profile: A DAN operating in PTP slave mode will keep only one of the ports in SLAVE state while the other remains in PASSIVE SLAVE mode. PASSIVE SLAVE and SLAVE ports are similar but PASSIVE SLAVE ports do not discipline the oscillator unless a fault is detected in the SLAVE port. In that case the PASSIVE SLAVE port becomes a regular SLAVE port. With this procedure, consistency between PTP synchronization (Sync message flow) and delay compensation (Peer Delay Request / Response flows) procedures is guaranteed.

Zeus provides special modes to generate PTP over PRP both as a PTP master and as an slave clock (see Figure 12). The main purpose of operation as a slave is to measure and compare the LAN A and LAN B delay paths. The test unit measured the *offset from master* (OFM) and *mean path delay* (MPD) individually in LAN A and LAN B. For each of them information is provided about their mean value, range and standard deviation. When the unit is synchronized by an external time source it also provides the *Time Error* (TE) between PTP and the configured reference. Again, two result sets are provided, one for each PRP LAN. In this case, also the relative TE between the LANs is computed and displayed.

As a master, xGenius can be used as an auxiliary device to provide timing to devices deployed in the substation but it can also simulate a large number of operation conditions that would be difficult to achieve without an specialized testing tool including generation of custom *clockClass*, and *accuracy* values in *Announce* messages. An important test scenario is skew generation in the traffic generated from LAN A and LAN B in a PTP grandmaster (see Figure 13). One could measure the effects on the output of the DUT, for example a *Merging Unit* (MU), synchronized over skewed timing source.

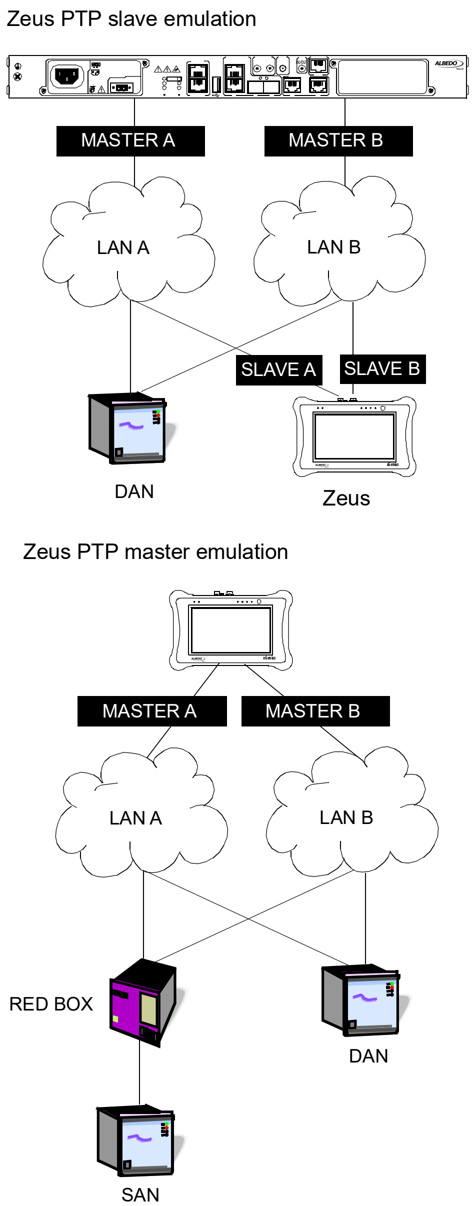

Figure 12 *Zeus PTP master and slave emulation modes*

For all PTP operation modes, Zeus provides different kinds of PTP statistics related to messages transmitted and received, delay and others. This function becomes even more powerful when they are used together with other general purpose statistics and functions such as the PRP statistics, the nodes table and the capture.

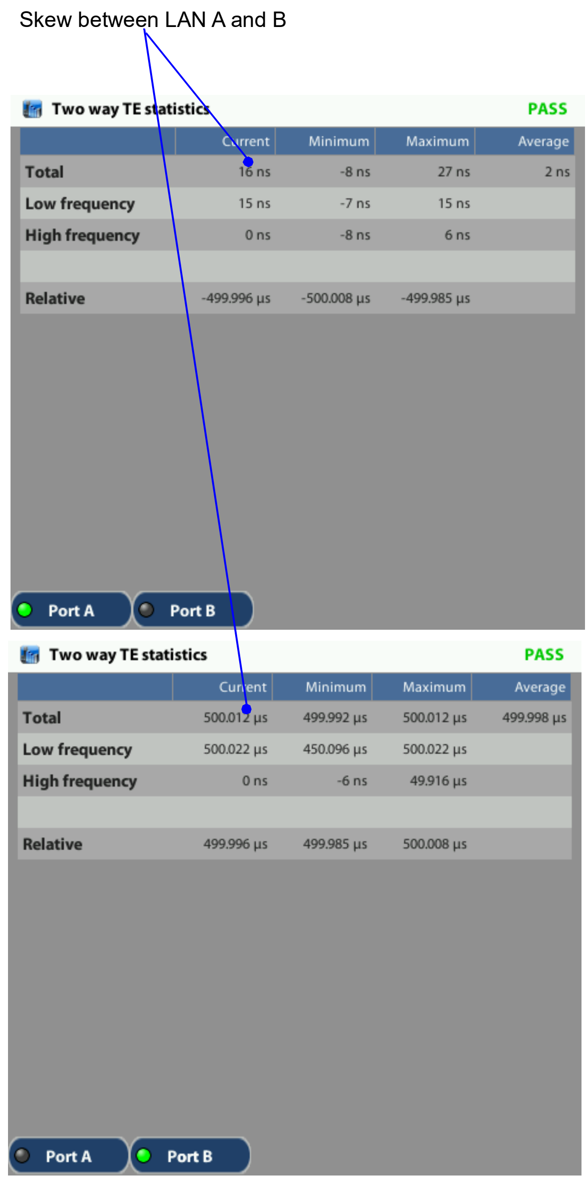

Figure 13 *PTP TE results collected in LAN A and LAN B y Zeus. These results are generated by a second Zeus unit running a PTP master emulation test with an skew added to Port B of 1 ms. We can see that the induced TE in Port B is then 0.5 ms, one half of the skew.*

4. IEC 61850 over PRP Test

The *Generic Object Oriented Substation Event* (GOOSE) and *Sampled Values* (SV) protocols carry time sensitive information and the analysis of the time required by this protocols to be transmitted and propagated through the network to the destination becomes critical. Zeus enables IEC 61850 protocol analysis with detailed time statistics required to assess the performance level of any device or network based on IEC 61850 protocols. Test setups include PRP and non-PRP-scenarios.

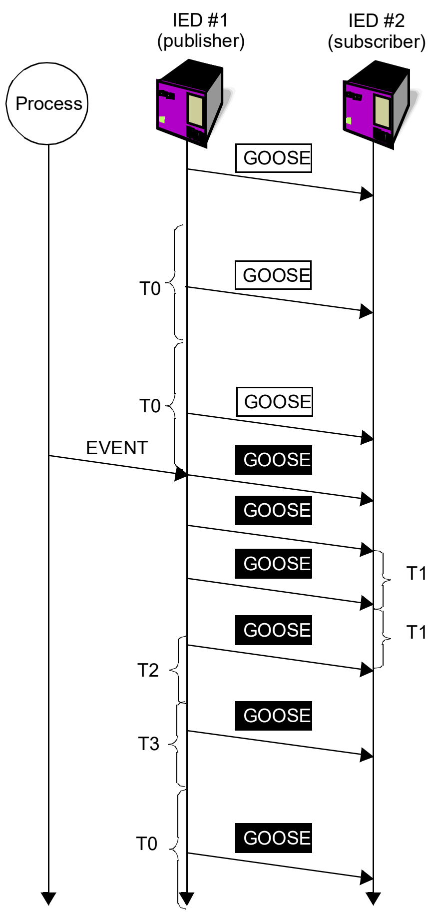

Figure 14 *(see Figure 14)Illustration of GOOSE reliable but unacknowledged transmission.*

GOOSE provides a fast and reliable system-wide distribution of input and output data values. The GOOSE protocol has an specific scheme of re-transmission to achieve the appropriate level of reliability: GOOSE publishers encode and transmit GOOSE messages on multicast associations when certain time critical events, such as the tripping of a circuit breaker, are recorded. Additional reliability is achieved by re-transmitting the same data with gradually increasing retransmission time. As the publisher-subscriber approach is not well suited for bidirectional transmission, GOOSE does not require any kind of acknowledgment mechanism (see Figure 14) . It is assumed that subscribers will receive enough copies of the message to allow for a proper operation.

SV results

	Frames	FTD	FDV
Current		3.76 ms	62.50 ms
Average		75.26 ms	920.96 µs
Maximum		783.16 ms	67.95 ms
Minimum		-430.94 ms	
Standard deviation		477.05 ms	
Range		1.21 s	
Packet number	56,839		

SV results

	Frames	FTD	FDV
Current		1.11 ms	62.50 ms
Average		45.40 ms	753.82 µs
Maximum		945.69 ms	67.95 ms
Minimum		-430.94 ms	
Standard deviation		425.33 ms	
Range		1.38 s	
Packet number	96,189		

Figure 15*Delay (FTD) and jitter (FDV) results measured by Zeus in SV transmitted over LAN A and LAN B in a PRP network.*

Transit time of GOOSE messages is important in many applications and critical in all those related to tele-protection. Zeus provides a mechanism to filter and record GOOSE messages and measure their transit delay in situations when the source and the destination are in the same or in remote substations.

The second example of peer-to-peer communications is provided by the SV protocol that implements a mechanism to distribute voltage and current samples collected from the power line through the process bus. To understand how the SV protocol works it is necessary to introduce the concept of *Merging Unit* (MU). MUs are devices with the ability to get the physical magnitudes from instrument current and voltage transformers (CTs, VTs), generate samples from these magnitudes and distribute the result of this operation in a packet switched network through an specialized protocol. MU interactions with CTs and VTs may be based on proprietary interfaces but the distribution of sampled values is standard and, potentially, devices from different vendors could decode and use the information generated by any MU.

Availability of current and voltage samples is important in protection applications because protective relays must receive accurate information about the power line status. Before the introduction of IEC 61850, relays had to be cabled directly to CTs and VTs to evaluate the line status but today all that is needed is a communications interface that is typically an optical fiber. A multicast association is established between the MU, one or several protection relays and, potentially, devices s not directly related with protection such as those involved in disturbance recording.

Time stamping of SV samples is of critical importance because it enables devices to carry out time-coherent correlation of two or more SV flows. This is important in disturbance recording applications and differential protection schemes. Zeus provides information about the Frame Transfer Delay (FTD) experienced by SV frames as they are transmitted through the network and it also provides statistics such as the range and standard variation measured in the SV flow. These results are generated in all setups, including those without PRP. In PRP tests two independent result sets are generated adding another verification layer to the protocol, which is consistency between LAN A and LAN B results (see Figure 15).