

Net.Time: Deploying the PTP Utility Profile over PRP

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 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.

and LAN B. These are typically bridged Ethernet networks designed so that a failure event in one of them does not affect to the second network.

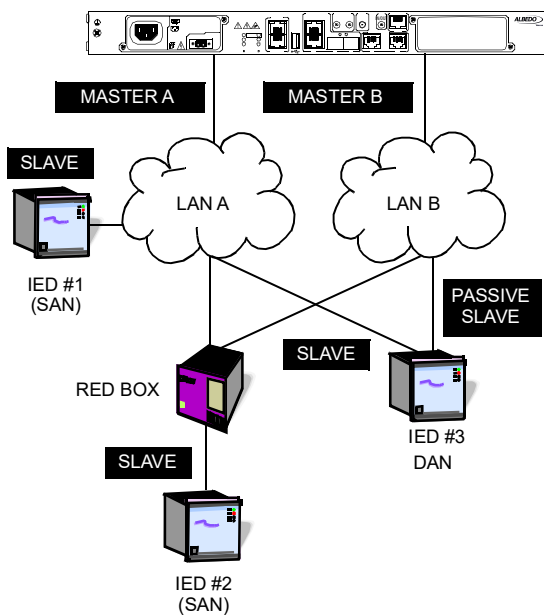


Figure 1 PRP provides redundant communications to IED #2 and #3. IED#1 is sensitive to network faults.

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




Figure 2 IED #2 and #3 keep their normal operating conditions when LAN A stops working but IED #1 loses its connection to the PTP grandmaster and it goes to holdover.

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



AN.INTIME.PRP.12/22

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.

It is next described the procedure to configure the IEC 61850-9-3 profile in a Net.Time PRP interface to achieve failsafe communications in a critical environment such as in a power substation.

1. SETTING UP MAIN REFERENCE

In this note we will consider that the main reference is GNSS. It is assumed that the Net.Time unit is restored to its factory values before starting the configuration and the following configuration is set in the unit:

- The device host name is set to “nettime”. The domain name “nettime” is also assigned to the clock.
- SSH management is enabled.

These are the steps to follow to enable and configure the GNSS reference in the unit using the CLI:

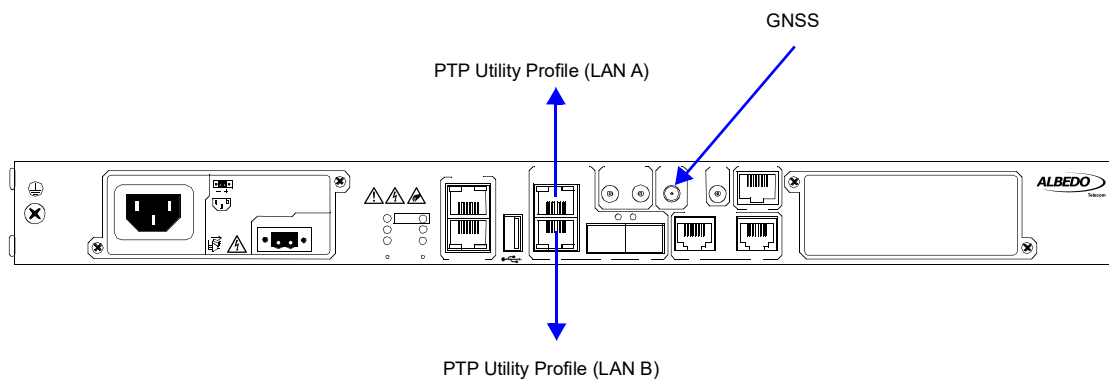


Figure 3 .Location of clock reference inputs (GNSS) and outputs (LAN A and B) in the current setup.

1. Attach the antenna to the unit. Make sure that the antenna sees as much of the sky as possible. The unit may fail to achieve synchronization if there are not enough satellites on sight. Details about how to install the antenna are omitted in this document and it is assumed that installation tasks are finished.

2. Log in the system using SSH and a controller account

```
ssh ctrl@nettime
ctrl@nettime's password:
```

3. Map the GNSS interface to the GNSS input

```
ctrl@nettime> set map GNSS(G) add gnss01
```

The locking process to the GNSS reference starts.

4. Optionally, enable the fixed position mode to improve the time accuracy

```
ctrl@nettime> set input gnss gnss01 fixed-position mode enable
```

If the Net.Time unit doesn't have a stored position, then it will retrieve its latitude, longitude and ellipsoidal height coordinates by means an averaging process that lasts for 24 hours in the default configuration. Fixed position mode will start at the end of this averaging process.

5. Wait for the locking process to finish. You can optionally display the oscillator sync status with

```
ctrl@nettime> show oscillator

Oscillator sync status
-----
Oscillator type:                OCXOE
Oscillator lock status:        Time synced
Elapsed time:                   0/00:06:25
Holdover remaining time:       -
```

6. Optionally, provision the GNSS input with

```
ctrl@nettime> set admin-status GNSS(G) provision
```

This action locks the configuration to avoid unwanted port misconfiguration in the future

2. ENABLING PRP

To enable the PRP in the unit, we need to modify the default packet service and configure the parameters of the resulting interface, PRP(J). This is the detailed procedure:

1. Set the packet service

```
ctrl@nettime> set packet-service prp-grandmaster
```

2. Provide an IP address (assumed here 10.0.0.1) to the ETH(A) port with

```
ctrl@nettime> set network PRP(J) address 10.0.0.1
```

3. Provide a subnet mask (assumed 255.0.0.0) to ETH(A)

```
ctrl@nettime> set network PRP(J) netmask 255.0.0.0
```

This and the previous steps are not strictly necessary when the PTP profile is going to be the *Utility profile* but it is still useful to avoid the INIP alarm and to allow the port to be reached by ping.

4. In many implementations, PTP is provisioned in a VLAN tagged interface. The following command sets the VLAN ID 101 with priority code-point 4:

```
ctrl@nettime> set frame PRP(J) encapsulation vlan vid 101 pcp 4
```

Unplugging one of the PRP LAN A or LAN B ports generates now an alarm that can be displayed with the following command:

```
ctrl@nettime> show alarm

Alarm      Subst.  Cat.  Origin  Status  Description
-----
(705)-PLAD ETH     minor PRP(J)  current PRP LAN A down
(704)-INIP ETH     major -        history  Invalid IP profile in Ethernet port
```

If both ports are unplugged at the same time, an individual alarm is generated for each port. PLAD for LAN A and PLBD for LAN B. An additional EDWN major alarm is generated to report total connectivity loss.

```
ctrl@nettime> show alarm
```

Alarm	Subsyst.	Cat.	Origin	Status	Description
(703)-EDWN	ETH	major	PRP(J)	current	Ethernet port link down
(705)-PLAD	ETH	minor	PRP(J)	current	PRP LAN A down
(706)-PLBD	ETH	minor	PRP(J)	current	PRP LAN B down
(704)-INIP	ETH	major	-	history	Invalid IP profile in Ethernet port

3. VERIFYING PRP

When Net.Time PRP interface is connected both to LAN A and LAN B, alarms related with PRP are disabled and the system becomes ready for operation. Users could easily get statistics about PRP with a single command:

```
ctrl@nettime> show analysis prp
```

PRP	PRP(J) frames	PRP(J).A frames	PRP(J).B frames
TX	44	44	44
TX RCT	-	44	44
RX	0	0	0
RX RCT	0	0	0
LAN mismatches	-	0	0
Errors	-	0	0
Unique	0	0	0
Single duplicate	-	0	0
Multiple duplicate	-	0	0

In the previous example we can see that frames generated from Net.Time are replicated in the PRP physical ports, *PRP(J).A* and *PRP(J).B*. These frames carry an special *Redundancy Check Trailer (RCT)*, that enables duplicate detection by other nodes in the Network. There are no timing interfaces configured in the system yet. The transmitted frames correspond with the PRP supervision flow. One of these frames is generated every second. It is likely that frames from other nodes, including supervision frames and other, are received in the PRP port. The following example illustrates the scenario of a DAN connected to the network:

```
ctrl@nettime> show analysis prp
```

PRP	PRP(J) frames	PRP(J).A frames	PRP(J).B frames
TX	214	214	214
TX RCT	-	214	214
RX	145	145	145
RX RCT	145	145	145
LAN mismatches	-	0	0
Errors	-	0	0
Unique	145	122	23
Single duplicate	-	23	122
Multiple duplicate	-	0	0

Received frames also carry the RCT. This is typical from DANs. Despite the fact that 145 frames are received in each port, totaling 290 frames, half of these frames were duplicates. The system detects and drops redundant information. That is the reason why only 145 unique frames are received. The first copy of each frame was received in *PRP(J).A* for 122 times and only 22 times in *PRP(J).B* with duplicates reported by the *Single*

duplicate row. That is probably because the delay in *PRP(J).B* is slightly larger than in the complementary port. No multiple duplicates are received. Unlike single duplicates, multiple duplicates correspond with an error condition and they should never be detected under normal operation.

The previous example was obtained by pinging Net.Time from a remote DAN attached to the PRP network. Net.Time replies ping requests and registers the DAN MAC address in the PRP nodes table:

```
ctrl@nettime> show prp PRP(J) nodes-table
```

MAC address	Type	TTL
XX:XX:XX:00:16:3C	DAN	59

The remote host is classified as a DAN. Any frame from the DAN can be used to populate the nodes table. If this host stops generating traffic to Net.Time, the system will remove the MAC address from the list. The *Time-To-Live* (TTL) field counts the time left before the address is deleted. Usually, the PRP supervision flow keeps entries alive in the nodes table when they are not exchanging other kinds of information, but in the example, the supervision flow was absent and, for this reason, the entry corresponding with this DAN is removed one minute after the last ping request message is received.

Non-PRP devices are identified as SANs, usually referred as SAN-A or SAN-B, depending on the LAN where they are connected. Obviously, it is not possible for these nodes to become DANs under any operating condition. Traffic from one of these SANs looks like in the following example:

```
ctrl@nettime> show analysis prp
```

PRP	PRP(J) frames	PRP(J).A frames	PRP(J).B frames
TX	195	195	195
TX RCT	-	195	195
RX	133	133	0
RX RCT	0	0	0
LAN mismatches	-	0	0
Errors	-	0	0
Unique	0	0	0
Single duplicate	-	0	0
Multiple duplicate	-	0	0

No traffic is received from *PRP(J).B* this time. We also see that the received traffic does not carry RCTs, which seems logical because the remote host doesn't have any knowledge of the PRP protocol. Traffic generated by Net.Time, is safely decoded by the SAN, however. This is because standard Ethernet nodes are designed to ignore the PRP trailer Net.Time is adding to the traffic. The SAN is correctly identified as a SAN-A in the nodes table:

```
ctrl@nettime> show prp PRP(J) nodes-table
```

MAC address	Type	TTL
XX:XX:XX:00:16:3C	SAN-A	59

4. ENABLING THE PTP UTILITY PROFILE

Once the clock inputs and the Ethernet links are ready we still need to configure the PTP output to provide accurate timing to the IEDs deployed in the substation. We will assume that these IEDs are compatible with the standard IEC 61850-9-3, which describes the PTP Utility profile.

1. Map the PTP (*ptp01*) output interface to the *PRP(J)* port

```
ctrl@nettime> set map PRP(J) add ptp01
```

2. Configure the Utility profile in *PRP(J)*

```
ctrl@nettime> set output ptp ptp01 profile utility
```

3. Optionally, verify the PTP output

```
ctrl@nettime> show ptp ptp01 status
```

```
ptp01 current status
-----
Status:                      ready
Clock state:                  started
Port0 state:                  master
Current clockClass:          6
Master IP address:           -
Master Ethernet address:     -
Master identity:              00:DB:1E:FF:FE:00:0E:1E
Master port number:          0
Timescale:                    TAI
Grandmaster identity:        00:DB:1E:FF:FE:00:0E:1E
Grandmaster priority 1:      128
Grandmaster priority 2:      128
Grandmaster clock class:     6
Grandmaster clock class desc: Synchronized to PRC
Grandmaster clock accuracy:  1.00e-07 s
Grandmaster clock variance:  -
Grandmaster time source:     GPS
```

4. Optionally, provision the Ethernet output with

```
ctrl@nettime> set admin-status PRP(J) provision
```

At this point, the unit will start distributing time in the Ethernet network through the PTP protocol. If the oscillator has already finished the locking process the IEDs will receive accurate timing from the clock. The address from the slave is included in the nodes table:

```
ctrl@nettime> show prp PRP(J) nodes-table
```

MAC address	Type	TTL
YY:YY:YY:7A:64:B6	DAN	52

In some situations, slaves may not be listed in the nodes-table. This is because PTP slaves running the IEC 61850-9-3 profile do not need to exchange traffic with the master. They only need to generate peer-delay-request messages to their peers, usually represented by switches located between the master and the slave. These messages are processed by the peer and dropped. However, if the slave is a DAN and it is generating the PRP supervision flow, it will be transmitted through the switch and it will be received by the master.

PTP traffic collected by Net.Time looks like in the following example:

```
ctrl@nettime> show analysis prp
```

PRP	PRP(J) frames	PRP(J).A frames	PRP(J).B frames
TX	623	323	325
TX RCT	-	25	25
RX	293	156	157
RX RCT	0	0	0
LAN mismatches	-	0	0
Errors	-	0	0
Unique	0	0	0
Single duplicate	-	0	0
Multiple duplicate	-	0	0

Duplicate detection does not apply to PTP traffic. Instead, PTP requests received from LAN A and B are independently processed and replied. The RCT, if present, is ignored by the receiver because it plays no role. If LAN B is disconnected from the master, operation continues through the alternative LAN. The frame counters in this case evolve as follows:

```
ctrl@nettime> show analysis prp
```

PRP	PRP(J) frames	PRP(J).A frames	PRP(J).B frames
TX	2166	1211	1048
TX RCT	-	93	93
RX	861	589	333
RX RCT	0	0	0
LAN mismatches	-	0	0
Errors	-	0	0
Unique	0	0	0
Single duplicate	-	0	0
Multiple duplicate	-	0	0

The RX frame counter for *PRP(J).B* is not increased because no frames are received through the broken link. Frame counts for the transmitter still increase but they do so more slowly. This is because some PTP traffic is still generated for LAN B (*Sync*, *Announce* and *Peer delay request* messages). Peer delay requests are not received and therefore they don't need to be replied.