Getting in Sync with Open Source - Addressing RAN Requirements

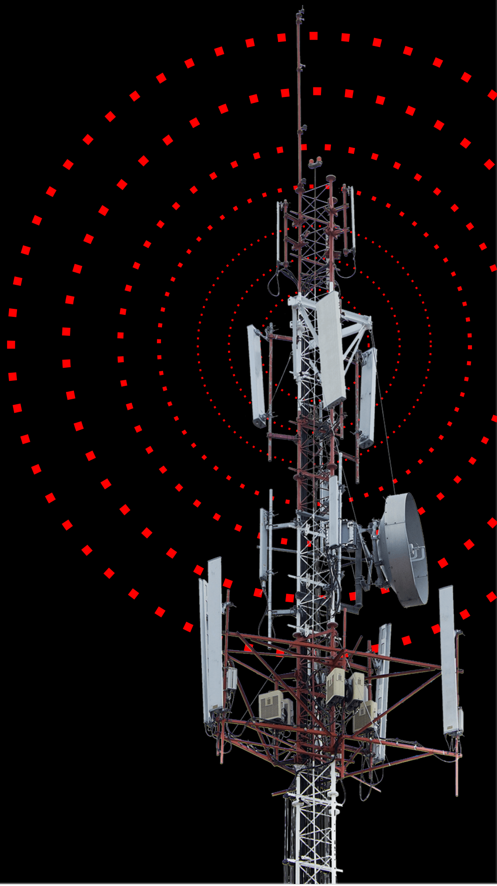

Workshop on Synchronization and Timing Systems
March 14th, 2023
Vancouver, Canada

WSTS

Pasi Vaananen
Systems Architect, Office of the CTO
Telco Enablement & Solutions

WORKSHOP ON SYNCHRONIZATION AND TIMING SYSTEMS

Red Hat

Agenda

- ➢ (O-)RAN synchronization network topologies and clock types
- ➢ Server node synchronization HW & SW implementation
- ➢ ITU-T Performance requirements summary
- ➢ Our Test Environment & Capabilities Overview
- ➢ Test results examples for different clock types
- ➢ Impact of SyncE on time/phase Synchronization performance

Red Hat

O-RAN "LLS-Cx" vs. Target Clock Types

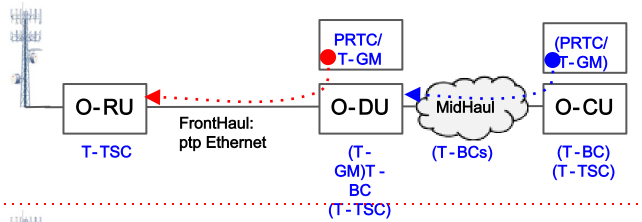

WG4 Configuration LLS-C1

FH Direct connection to O-RU(s) from O-DU; sync source in O-DU

WG9 C1, Option A: GNSS/T-GM Embedded in ~~O-DU~~ Cloud node)

WG9 C1, Option B: T-GM Directly Connected to O-DU

WG9 C1, Option C: T-GM connected to O-DU via chain of network nodes

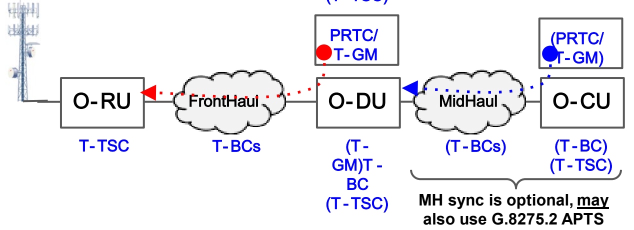

WG4 Configuration LLS-C2

FH Network connection to O-RU(s) from O-DU; sync source in O-DU

WG9 C2, Opt. A: ~~O-DU~~ Cloud node is the nearest common T-BC

WG9 C2, Opt. B: nearest common T-BC not ~~O-DU~~ Cloud Node

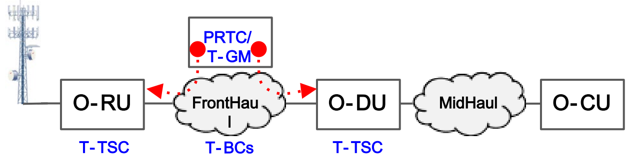

WG4 Configuration LLS-C3

Network connection to O-RU from O-DU & sync source in FH network

WG9 C3, Option A: T-GM is the nearest common master

WG9 C3, Option B: nearest common master is not T-GM

WG9 C3, Option C: T-GM in Mid/Back-haul

WG9 C3, Option D: T-GM in Mid/Back-haul with T-BC chain

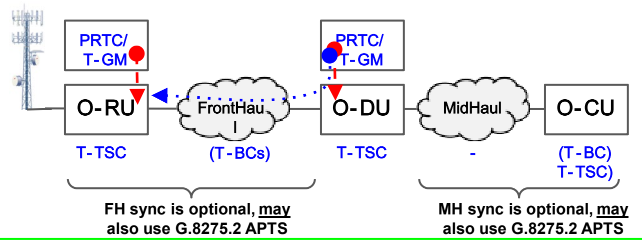

WG4 Configuration LLS-C4

Network connection to O-RU from O-DU & local sync sources

WG9 C4, Option 1: GNSS at Cell Site (e.g. in O-RU / xNB)

WG9 C4, Option 2: GNSS at Cell & Edge + APTS network

(O-)RAN Market priority list

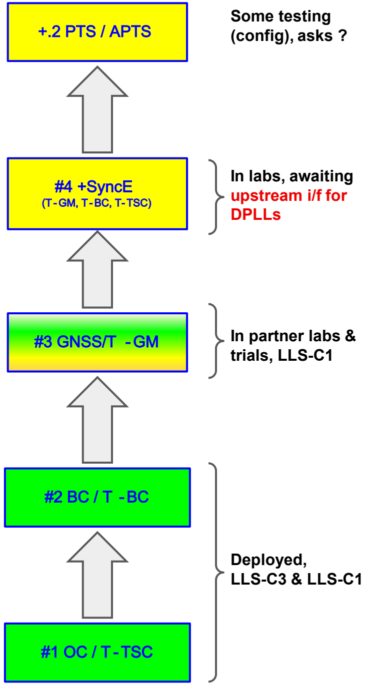

Note:
In clouds, synchronization is part of cloud infrastructure, and decoupled from RAN instances

Synchronization Components in Linux / k8s nodes

Each Physical Node

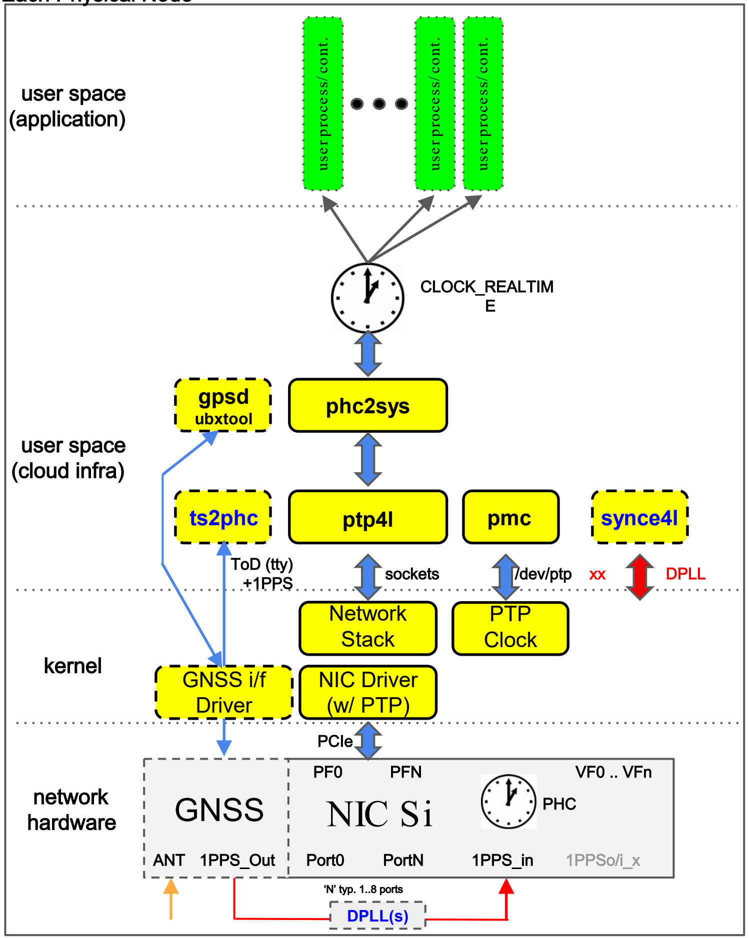

Key Components of the node Synchronization implementations

- HW specific synchronization SW support features are implemented in HW device drivers
- HW Clock (PHC) support in NIC Si is required for high accuracy; SyncE requires DPLLs
- **Linuxptp** is an Open Source project implementation of the PTP SW stack for Linux
- **ptp4l** implements Boundary Clock (BC) and Ordinary Clock (OC), it synchronizes PTP hardware clock (PHC) to remote master clock
- ptp4l is very flexible, and can be configured to support specific profiles, assuming that HW & driver supports associated features (e.g. PHC, L3 vs. L2 transport, accuracy targets)
- **phc2sys** synchronizes two or more clocks in the system, typically used to synchronize the system clock from PTP / PHC
- **pmc** - PTP management client; 1588 basic management access for ptp4l
- ts2phc synchronizes PHC(s) from external reference signals, such as 1PPS_in and ToD messages- used in multi-card T-BC, and GNSS driven T-GM configurations
- synce4l : implements ESMC protocol and associated state transition controls
- gpsd/ubxtool can be used to interface w/ GNSS receivers / u-blox specifics for configuration management and status monitoring
- In k8s clusters, synchronization functions are configured and monitored with k8s, and associated general CM, PM and FM event and metrics tools.
- O-RAN WG6 specifies a notification interface to inform synchronization users about node synchronization state changes

intel NVIDIA XILINX ... And many other Si & HW Partners

BROADCOM ublox Red Hat

2 kinds of typical HW from Sync. & SW Perspective

- Synchronization optimized designs add various functions, typically DPLL and PHY clock recovery/generation, as well as phase/f req physical inputs and outputs, and GNSS options all which need SW enablement for configuration, monitoring and control purposes
- Various design specific ways to achieve equivalent things, none is “right” or “wrong”, they are just different; goal is to ab stract to the “right” level through Linux interfaces and device drivers
- FPGA cards and SoCs with embedded NI functionality are generally similar in design from sync perspective (i.e. incorporate “N I Si” block below)

“Basic” PTP HW-timestamping capable HW Without Synchronization Enhancements Features

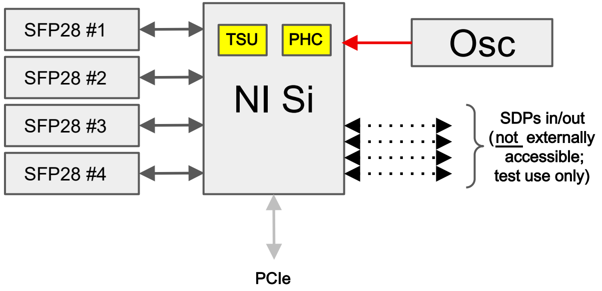

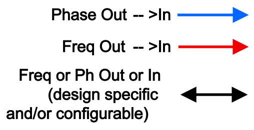

PHC: PTP HW Clock
TSU: TimeStamping Unit
1PPS: 1 Pulse Per Second
GNSS: Global Navigation Satellite System
DPLL: Digital Phase Locked Loop
OCXO: Oven Controlled Crystal (Xtal) Oscillator
ToD: Time of Day
NMEA: National Marine Electronics Association
PCIe: Peripheral Component Interconnect express

HW With Synchronization Enhancements Features

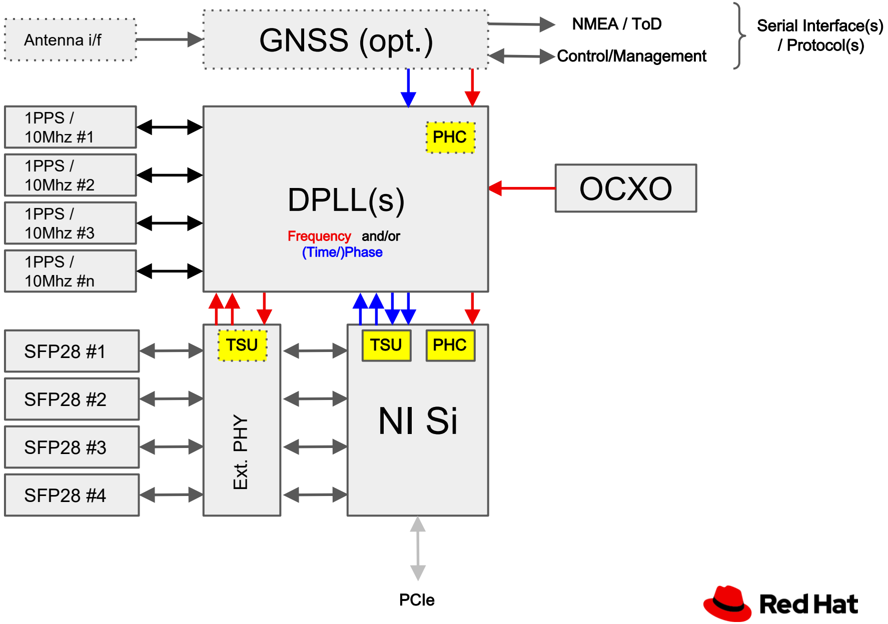

Red Hat

The Test applicability - non-SyncE and SyncE capable HW

Test Cases	Basic HW	Enhanced HW (w/ SyncE)
G.8273.2 7.1.x: TE Noise Generation	Yes, all; 7.1.4.x TE_R defined for Class-C only	
G.8273.2 7.2.x: Noise Tolerance	Yes 7.2.1 (A/B) vs. 7.2.2 (C)	Yes, 7.2.2 (Class C)
G.8273.2 7.3.x: Noise Transfer	7.3.1 PTP-PTP only	Yes, 7.3.1 PTP-PTP & 7.3.3 phy. layer freq to PTP (C/D)
G.8273.2 7.4.1 Transient Response	(7.4.1.2 PTP only; perf is FFS)	Yes
G.8273.2 7.4.2 Holdover Performance	(7.4.2.1 PTP only; perf is FFS)	Yes; Class-C mask FFS

- We do also conduct the applicable tests for 1PPS and/or 10Mhz phase/freq. Reference outputs if those outputs are supported as physical interface with connectors
- In “standard” cards without connectorization, we may also use these for testing if available e.g. through pin headers, but they are expected to **not** be used/usable/supported in end-application
- Also, G.8262.1 eEEC test sequence typically applies to SyncE capable cards, not shown here
- We view eEEC tests more of DVT tests, minimum SW impact; primarily FW/driver/DPLL configs

G.8273.2 T-TSC & T-BC Noise Gen Reqs Summary

Parameter	Class-A	Class-B	Class-C	(Class-D); Still WIP in ITU-T	Notes
7.1 Max. Absolute Time Error; max|TE|	≤ 100 ns	≤ 70 ns	≤ 30 ns	(≤ 15 ns)	Unfiltered measurement, absolute value
7.1 Max. Absolute Time Error; max|TE_L|	-	-	-	≤ 5 ns	0.1Hz low-pass filtered, 1000s measurement, absolute value
7.1.1 Max. Constant Time Error; cTE	≤ ±50 ns	≤ ±20 ns	≤ ±10 ns	(≤ ± 4 ns)	cTE Averaged over 1000s
7.1.2 Max. dynamic Time Error, 0.1Hz Low-Pass Filtered; dTE_L (MTIE)	≤ 40 ns		≤ 10 ns	(≤ 3 ns)	MTIE Mask, 1000s observation interval constant temp., (10000s for A/B variable temp.)
7.1.2 Max. dynamic Time Error, 0.1Hz Low-Pass Filtered; dTE_L (TDEV)	4 ns		2 ns	(≤ 1 ns)	TDEV Mask, 1000s observation interval at constant temp.
7.1.3 Max. dynamic Time Error, 0.1Hz High Pass Filtered; dTE_H	70 ns p-p		FFS (30ns p-p?)	(15 ns p-p)	Peak-to-peak value, 1000s measurement
7.1.4.1 Relative constant Time Error Noise Generation; cTE_R	FFS		≤ ±12 ns	FFS	cTE averaged over 1000s
7.1.4.2 Relative dynamic Time Error Low-Pass Filtered Noise Generation; dTE_{RL} (MTIE)	FFS		≤ 14 ns	FFS	MTIE Mask; 1000s observation interval at constant temp.

Note: Accuracy required is primarily determined by specific Use Case requirements & number of elements on the synchronization transfer path

Red Hat

Synchronization Test Setup at Red Hat TelcoLab

Path-delay calibrated Optical single-mode Switch for remote/automated test reconfiguration

Spirent TestCenter for Bulk Traffic testing (≤1.2Tbps)

Spirent TC for Detailed Perf. Testing (≤800Gbps), w/ PTP/SyncE slave emulation

Calnex Attero-100G Network Emulation System

Calnex Paragon Neo PTP/SyncE analyzer w/ comprehensive SW feature set for Telco PTP & SyncE

Keysight 8ch Scope w/ time&freq analysis SW

Cisco&Dell PTP/SyncE capable Switches & routers (back)

Primary GNSS signal splitters (roof antenna + simulator)

1PPS & 10Mhz reference distribution amplifiers

Primary ePRTC (TP4100 + 5071A w/ hi-performance CBT)

Secondary ePRTC (TP4100 + 5071A w/ hi-performance CBT)

Spirent GNSS Constellation Simulation System

22x servers w/ various NICs, FPGA, SoC etc. PCIe cards w/ Sync. Support ("DUT" servers)

1PPS/10M signal switch (rear)

Not Shown: O-RU & UE emulation systems, Test O-RU's, 50Ghz Spectrum Analyzer, very high-end scope and OTA performance test gear etc.

Test Target Adjacencies for Black-Box DUTs Overview

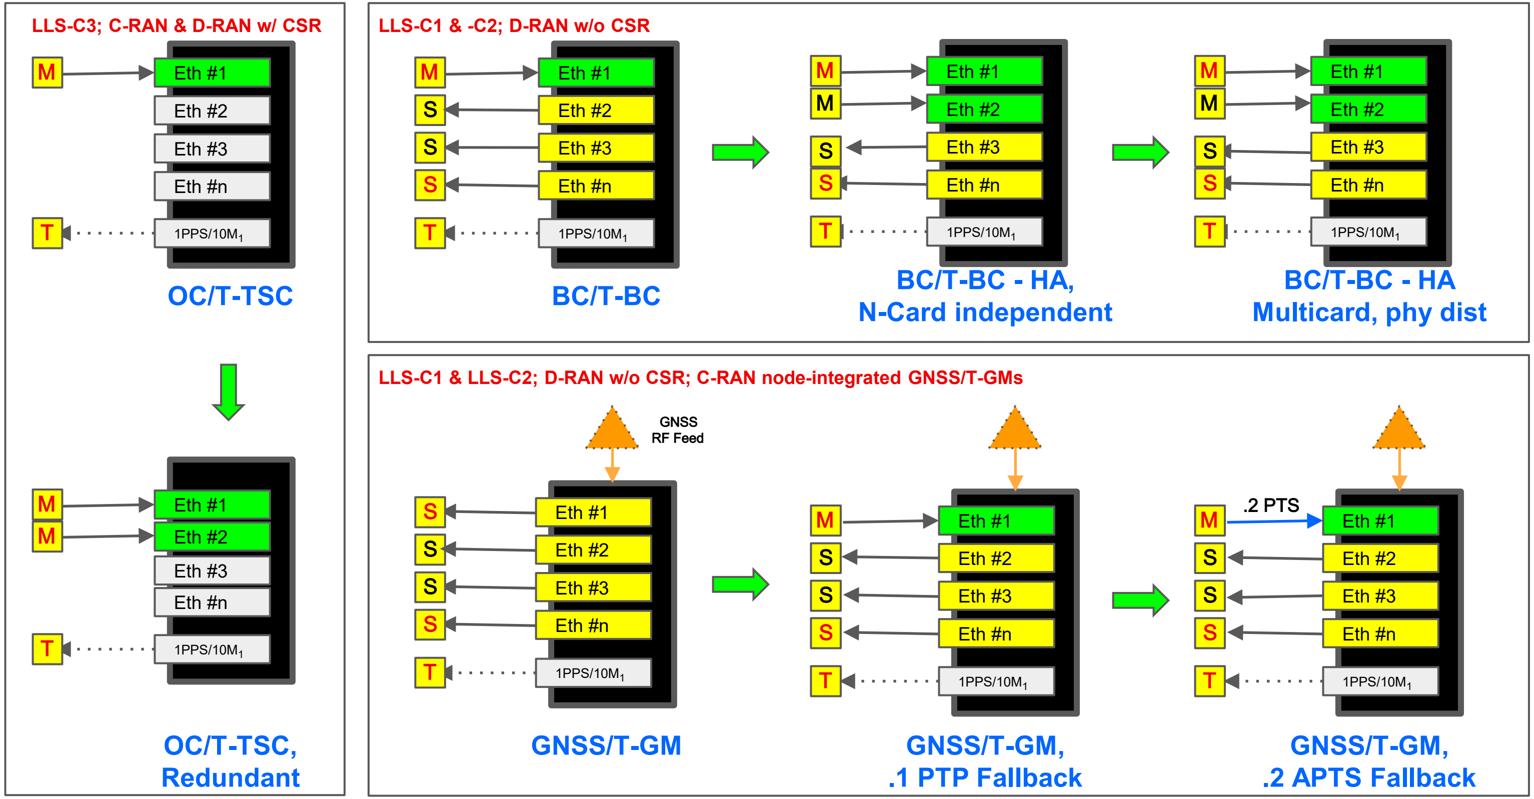

G.8275.1 FTS (+opt SyncE subject to HW support); M →S
G.8275.2 PTS (+opt SyncE) can be substituted to any i/f

G.8275.1 PTS (+ opt optional SyncE subject to HW support); M →S

1PPS | 10Mhz test signal, if supported by HW; as needed (opt.)

M PTP (+optionally SyncE) Master Port; external to DUT

S PTP (+optionally SyncE) Slave Port; external to DUT

T 1PPS or 10Mhz Test input port; external to DUT; optional

Red M/ S/ T main test ports

Red Hat

Server DUTs as "Grey Boxes", multi-card (simplified)

OC/BC (1..N "basic" NICs)

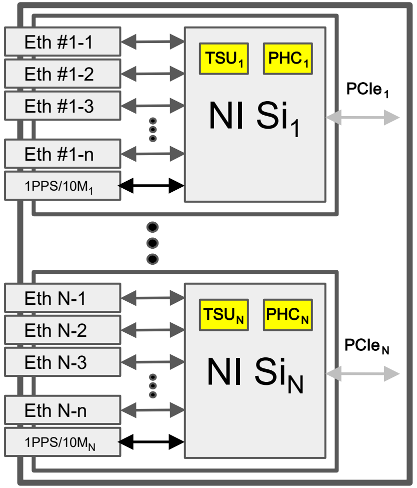

T-TSC/T-BC (1..N "enhanced" NICs)

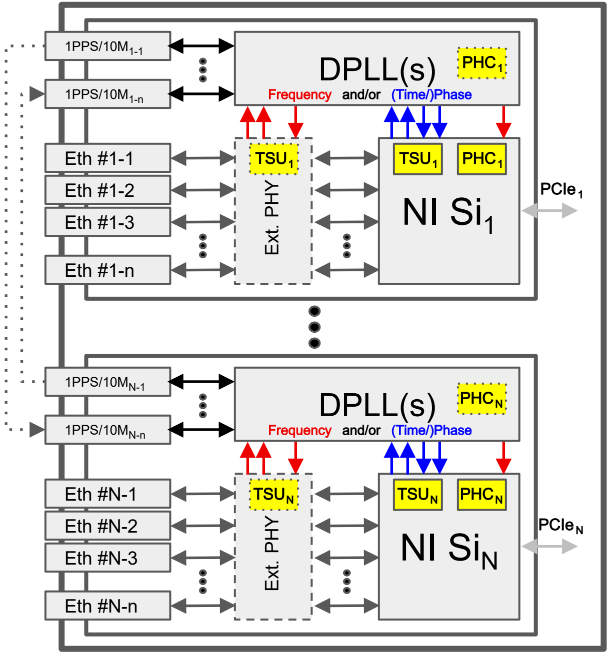

GNSS/T-GM (1..N "enhanced" NICs)

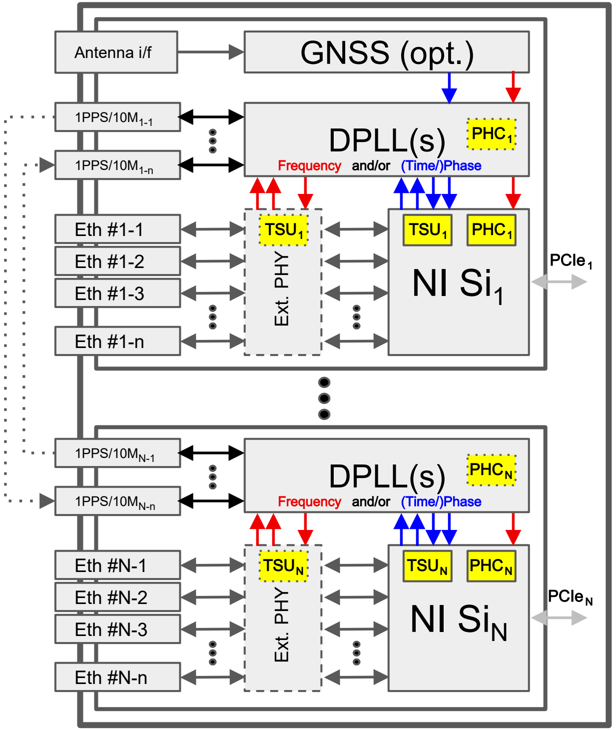

OC/BC, Target Configurations:

- Up to N*n PTP Interfaces
- OC: typ 1 or 2 (HA, 1 or 2-card) active i/f's
- BC: typ N*(1*slave + n-1*Master)
- OC: 1*ptp4l + 1*phc2sys
- BC: N*ptp4l + 1*phc2sys
- 1PPS/10M: out only, lab / field test

T-TSC/T-BC, Target Configurations:

- Up to N*n PTP Interfaces
- T-TSC: typ 1..2 (HA, 1 or 2-card) active i/f's
- T-BC: typ N*(1*slave + n-1*Master) i/f's, possibly less slaves than cards when cards synchronized w/ physical signals
- T-TSC: 1*ptp4l + 1*phc2sys + 1*synce4l
- T-BC: 1..N*ptp4l + 1*phc2sys + 1*ts2phc +1*synce4l
- 1PPS/10Mhz monitoring (out) + in (multicard sync with physical signals, using ts2phc to sync PHCs)

GNSS/T-GM Target Configurations:

- Basics as with T-TSC/T-BC
- Plus 1xGNSS (using ts2phc, different config)
- +gpsd and/or +ubxtool, dep. on use case
- T-GM: up to N*n Master i/f's

Set-up for a G.8275.1&G.8273.2 T-BC (T-TSC) w/ Neo

Cs Frequency Standard (+ ePRTC)

Notes:

- Only a single port-pair + 1PPS & 10Mhz can be measured at the time
- Measurement of the System Clock (i.e. OS / CPU clock) is not possible in normal CPU SKUs, and those that have some time enabled pins are not commonly used
- No under-load testing, or if enabled only port pair, not sufficient to exercise entire NIC and/or PCIe related issues (latter is esp. relevant to phc2sys / time transfer over PCIe)

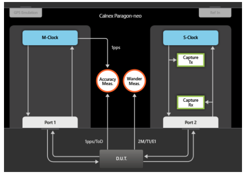

10Mhz
(opt. 1PPS+ToD)

Physical 1PPS and/or 10Mhz, as supported / required

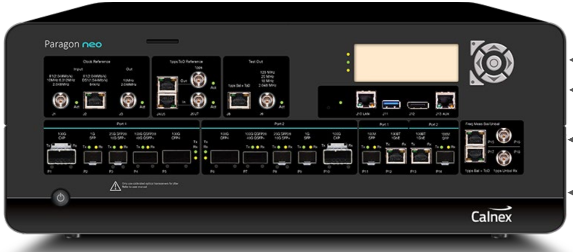

T 1PPS (Coax)

T 10Mhz (Coax)

M 1/10/**25**/100G Ethernet

S 1/10/**25**/100G Ethernet

S

M

DUT PTP Test Ports In Different Cards (worst Case, max # of 1PPS 'hops')

DUT: **T-BC** w/ N cards (2 shown)

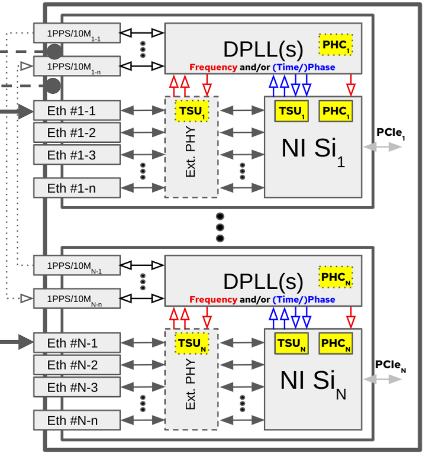

Paragon Neo PTP Tester

G.8275.1 PTP FTS (+SyncE opt.)

OC/T-TSC Clock Tests:

- OC/T-TSC configuration is subset, at least one Neo Master to at least one DUT Slave port
- OC/T-TSC tests require measurement signal back to tested, i.e. 1PPS
- In Practice, if it meets specific class as (T-)BC, it will do at least as well as T-TSC/OC, and (T-)BC test does not mandate need for 1PPS physical signals

Red Hat

Using the Linux / PTP metrics and events

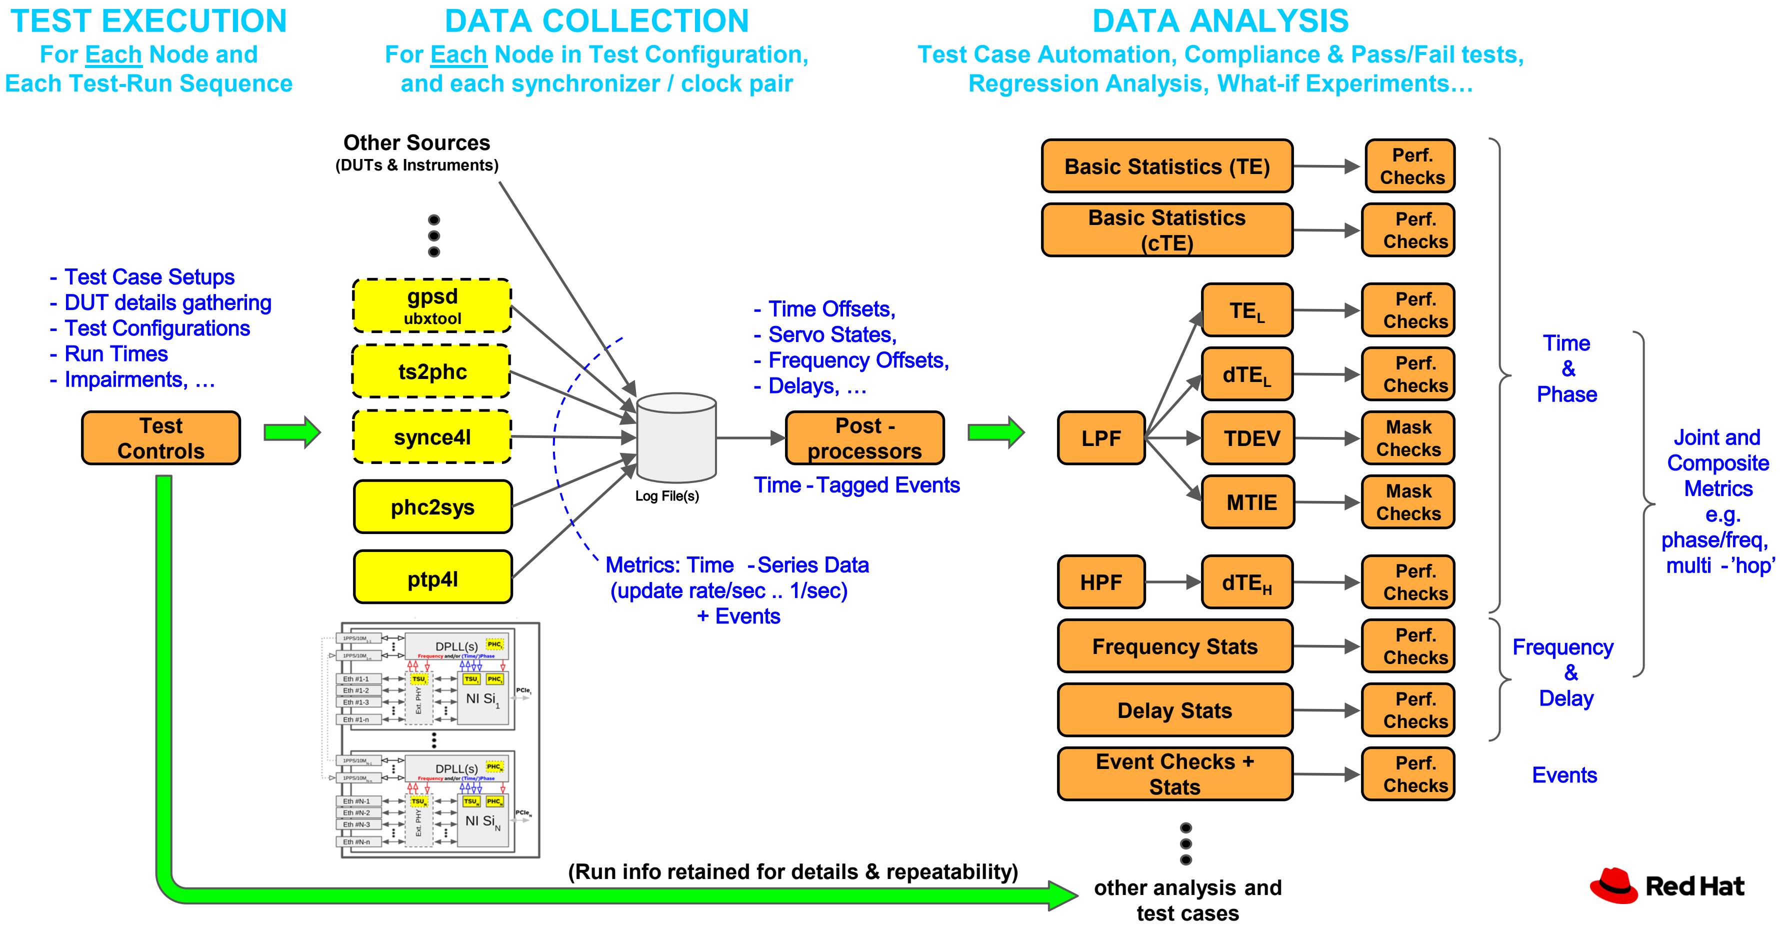

Linux / PTP metrics analysis G.8275.1, ptp41

Basic HW

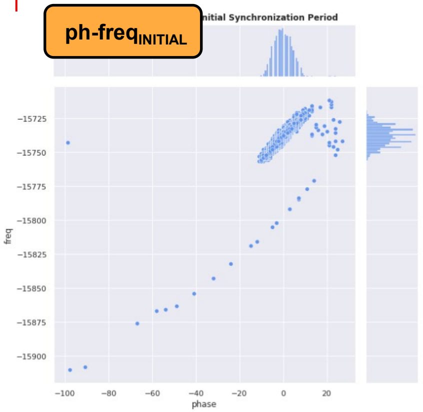

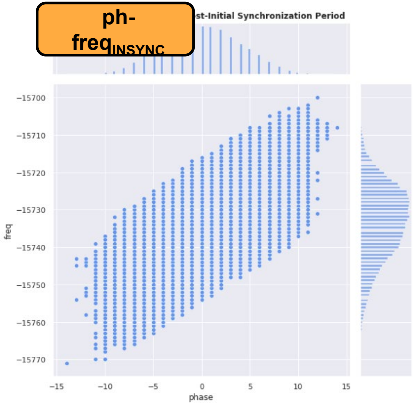

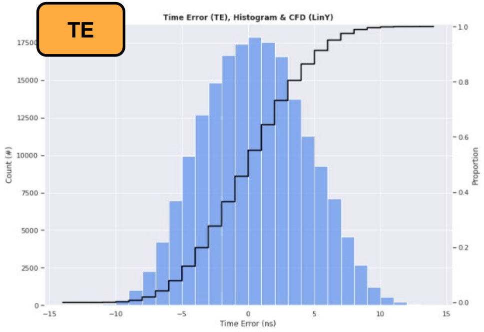

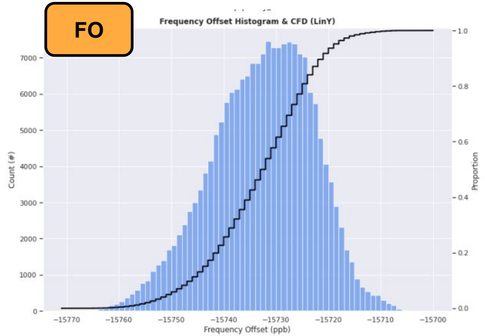

Time from 1st S0 to 1st S2: 0.188 s

Min phase: -14.000 ns
Max phase: 14.000 ns
Phase stddev: 3.891 ns

Min Freq: -15771 ppb
Max Freq: -15700 ppb
Freq stddev: 9.708 ppb

Enhanced HW; NO SyncE

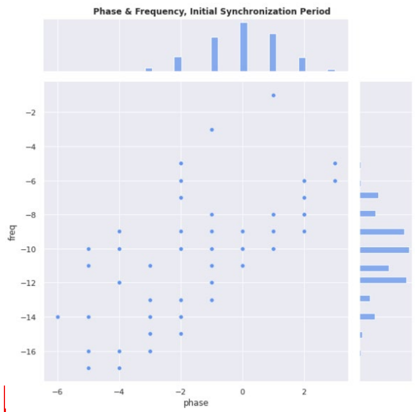

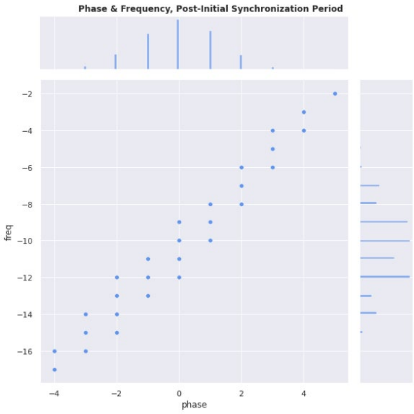

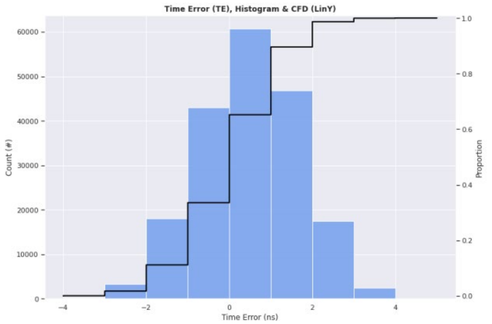

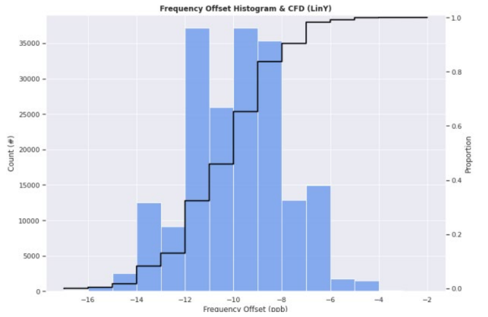

Time from 1st S0 to 1st S2: 0.251 s

Min phase: -4.000 ns
Max phase: 5.000 ns
Phase stddev: 1.222 ns

Min Freq: -17 ppb
Max Freq: -2 ppb
Freq stddev: 2.061 ppb

Linux / PTP metrics analysis; G.8275.1 ptp4l basic HW

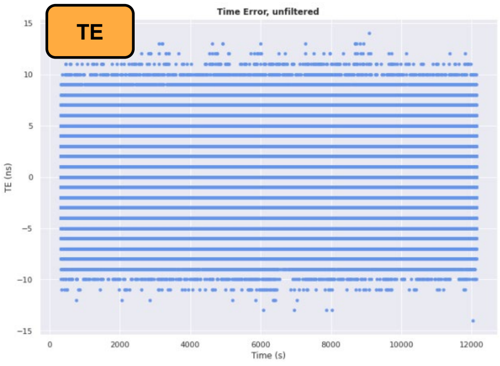

pk-pk phase: 28.000 ns
Mean phase: -0.000 ns
Min phase: -14.000 ns
Max phase: 14.000 ns
Phase stddev: 3.891 ns
Max |TE|: 14.000 ns

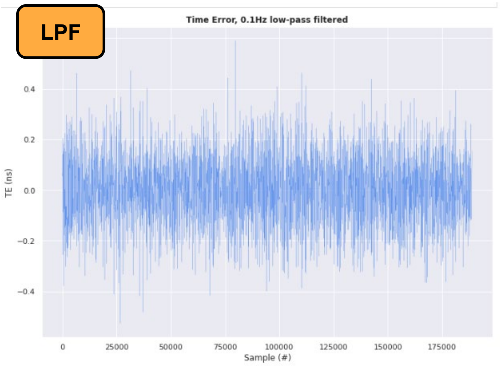

pk-pk lpf phase: 1.115 ns
Mean lpf phase: -0.000 ns
Min lpf phase: -0.526 ns
Max lpf phase: 0.589 ns
Lpf phase stddev: 0.127 ns
Max lpf |TE|: 0.589 ns

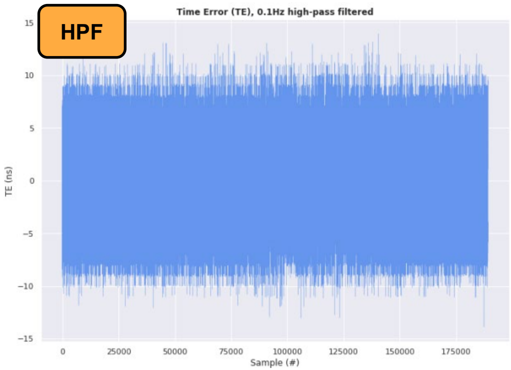

pk-pk hpf phase: 27.806 ns
Mean hpf phase: 0.000 ns
Min hpf phase: -13.900 ns
Max hpf phase: 13.905 ns
hpf phase stddev: 3.887 ns
Max hpf |TE|: 13.905 ns

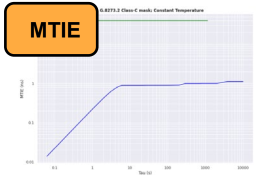

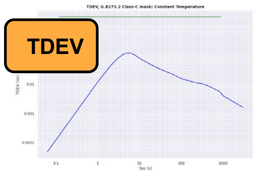

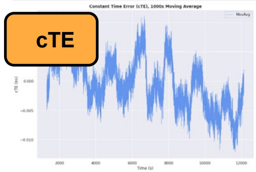

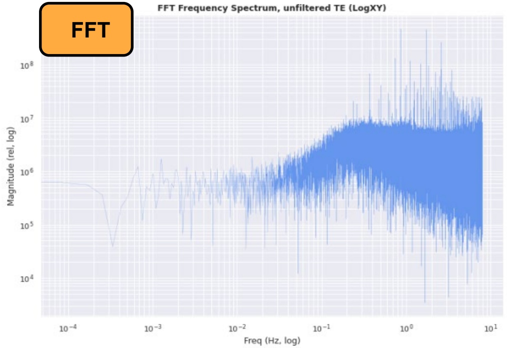

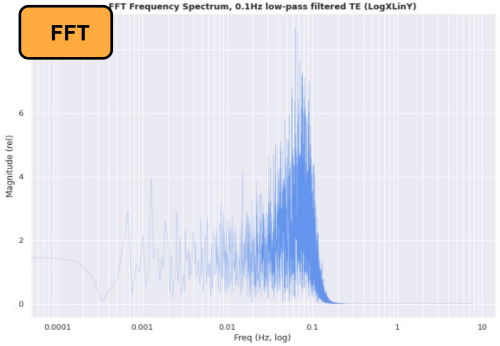

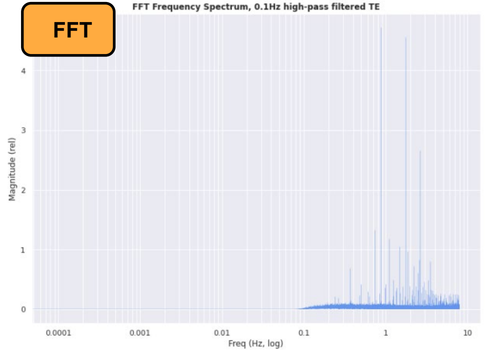

Red Hat

7.1T-BC Noise Generation - Unfiltered TE

Basic HW

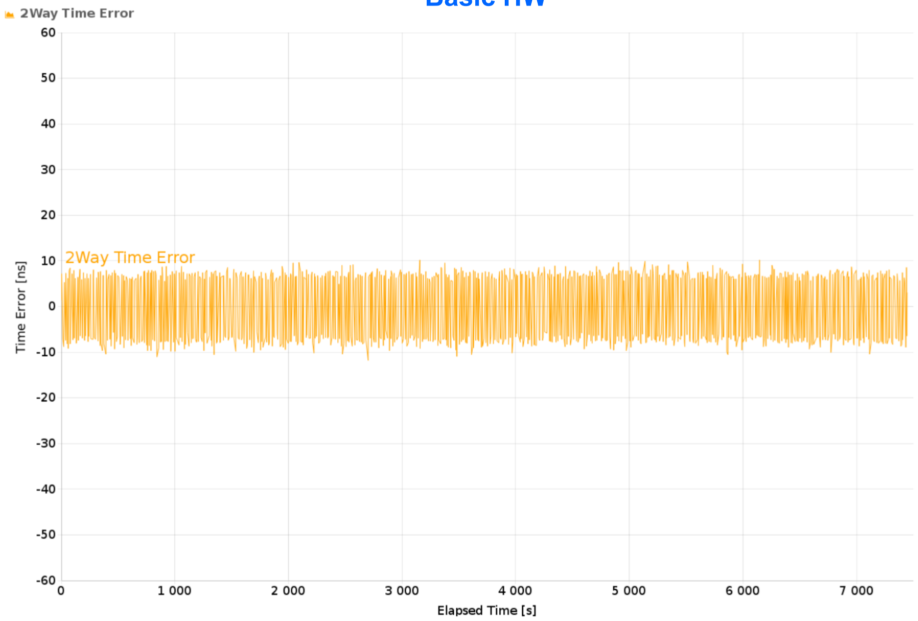

Pk-Pk	22 ns
Mean	-0.374 ns
Min	-11.805 ns
Max	10.195 ns
Max-Min	22 ns

Enhanced HW With SyncE

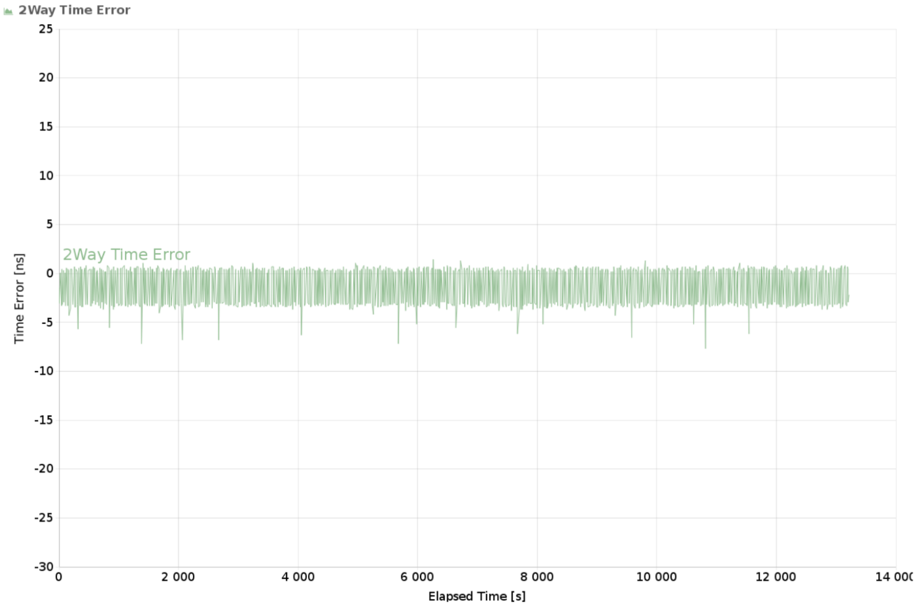

Pk-Pk	9.125 ns
Mean	-1.461 ns
Min	-7.68 ns
Max	1.445 ns
Max-Min	9.125 ns

- Both DUTs do pass Class C for TE with lots of margin (Class-C requires 30ns for Unfiltered |TE|)
- The performance with SyncE in this test is ~2x better then without it in the same configuration

Red Hat

7.1.2 T-BC Noise Generation - Dyn. MTIE LF Class-C mask

Basic HW

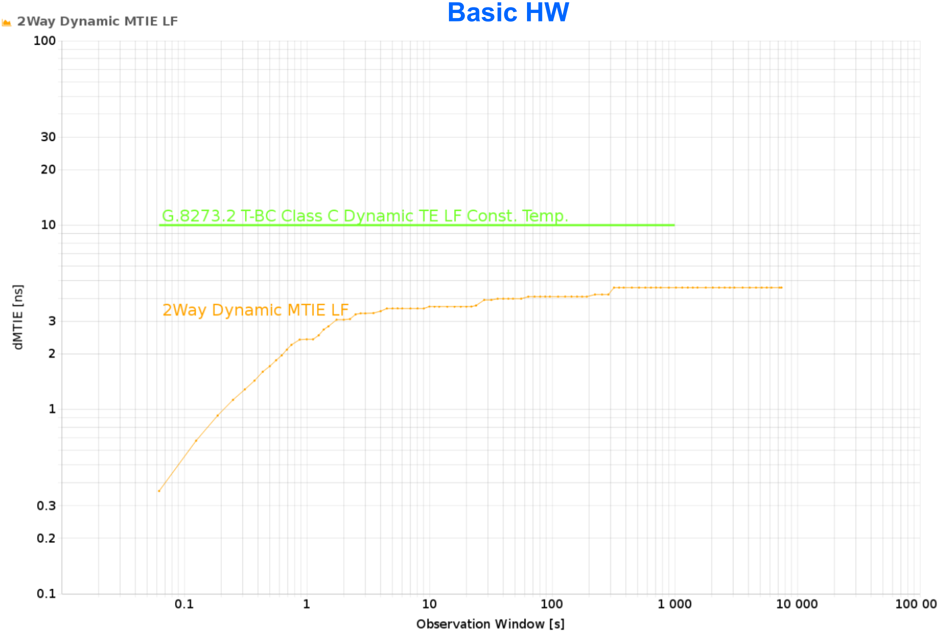

Min	0.361 ns
Max	4.58 ns
Max-Min	4.219 ns

Enhanced HW With SyncE

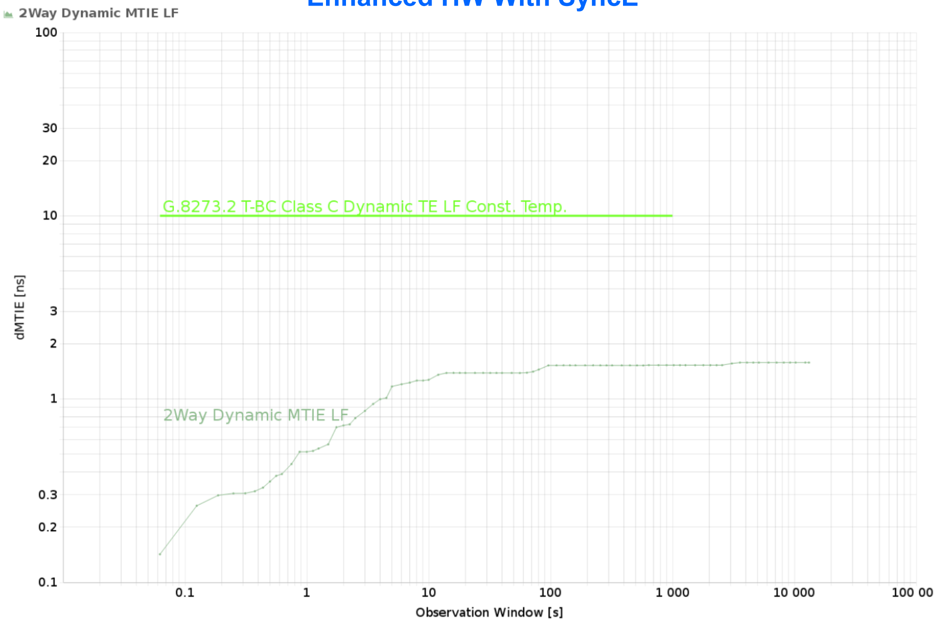

Min	0.142 ns
Max	1.582 ns
Max-Min	1.439 ns

- Both DUTs do pass Class C for MTIE LF 10ns mask (as well as TDEV mask) with margin
- The performance with SyncE in this test is >2x better then without it in the same configuration

Red Hat

7.1T-BC Noise Generation - 1vs. 2-Card with SyncE

2-Way TE, Single Card

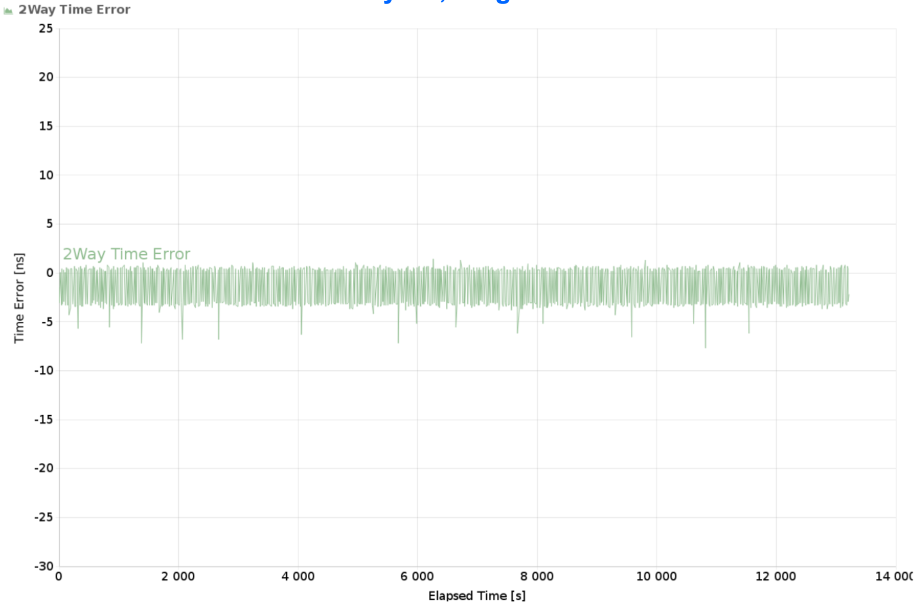

Pk-Pk	9.125 ns
Mean	-1.461 ns
Min	-7.68 ns
Max	1.445 ns
Max-Min	9.125 ns

2-Way TE, Dual Card

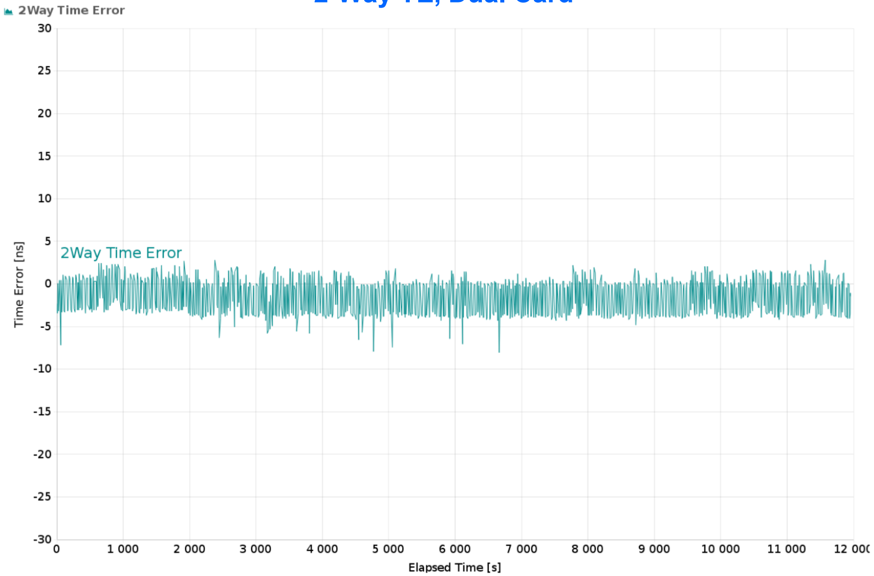

Pk-Pk	10.875 ns
Mean	-1.6 ns
Min	-8.063 ns
Max	2.812 ns
Max-Min	10.875 ns

- In this test, the BC slave port is in 1st card, and master port is in 2nd card
- Combination still passes Class C for TE with margin (important, multi-card configs are common !)
- This is primarily enabled by physical phase/freq signals to sync the 2nd card from the 1st

Red Hat

Holdover Performance - SyncE HW ("good" OSC+DPLL)

A: Holdover (dMTIE), w/ SyncE

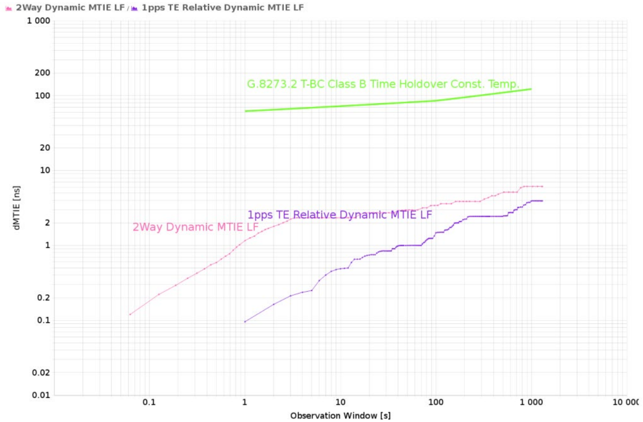

B: Holdover (dMTIE), SyncE lost

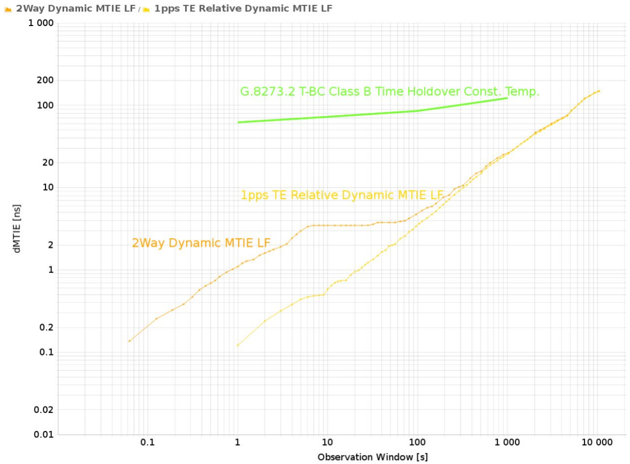

C: Holdover (dMTIE), Freq FreeRUN - no SyncE lock/use at all

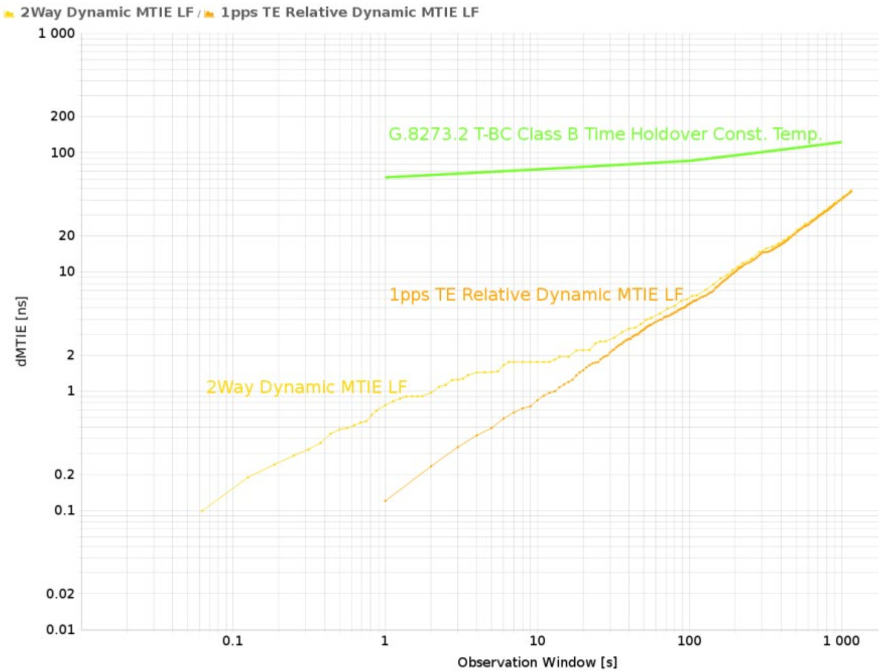

- Mask (above) is for Class-B, Class C/D is FFS
- ~constant temperature (lab !) tests ONLY done at Red Hat
- Case A: SyncE assisted holdover: ~6ns @1000sec
- Case B: SyncE lost holdover: ~25ns @1000 sec
- Case C: FreeRun, No SyncE: ~47ns @1000 sec

Red Hat

Get Involved & Get In Sync

Upstream projects - SW & open HW

- Chrony (primarily NTP)
- LinuxPTP project
- Linux kernel - common sync if/s
- Linux HW dev. Drivers- sync features
- OpenCompute TAP
- OpenCompute Networking
- OpenCompute Telco
- OpenCompute Telco Edge
- TIP OOPT
- TIP RAN projects
- Jupyter, MatplotLib, Scipy/Numpy, pandas, …
- Allantools

Key Standards / Spec. Organizations

- O-RAN WG4,5,6,7,8 and 9…
- ITU SG15 Q13/15
- IEEE P802 / 802.3cx (TS accuracy)
- IEEE P1588
- 3gpp

How to contact me

Pasi Vaananen: pvaanane@redhat.com

Special Thanks to

- LinuxPTP community, Linux Kernel Community, our NIC Si/card & FPGA Si/card HW partners, vRAN SW partners, Calnex & Spirent & Keysight