CAST

— Center for Alternative —
Synchronization and Timing

Technical Bulletin #001

Summary: This CAST technical bulletin highlights the effects of network path routing automation on time synchronization precision, network performance visibility through synchronization data. Utilizing Precision Time Protocol (PTP) over DOE's Energy Sciences Network (ESnet), we show that as ESNet reroutes network traffic to maintain/improve overall network latency can result in significant inaccuracies for time synchronization between network devices.

This bulletin presents some of the effects and observations of the "state of network" on PTP accuracy. "State of network" here refers to anomalies include traffic on the network, network congestion, network routes, differences in paths taken in communication between two entities, etc. The network under consideration is the Department of ESnet[1] provisioned through OSCARS (On-demand Secure Circuits and Advance Reservation System[2]).

The CAST team has been testing PTP performance over an encrypted Layer 2 network path using an ESnet circuit rated to 400 gigabit/sec. For many months, the PTP accuracy was objectively sufficient (within 250ns). In early 2023, we identified a sudden instability in the PTP accuracy. After discovering no issues with lab and deployed components, it became clear there were anomalies within the network.

Our testing configuration included a Grand Master Clock (GMC) running at ORNL, distributing time to a Boundary Clock (BC) installed at Savannah River National Laboratory (SRNL), a project partner for this experiment. The network path between the two laboratories is as shown in Figure 1. Outside of site-specific internal network segments symbolized in green and orange, the remainder of the circuit is all ESnet (symbolized in gray). The GMC is part of an Enhanced Primary Reference Time Clock (ePRTC) system. A 10 MHz signal from a magnetic Cesium atomic clock fed to the GMC and used as primary clock reference. Time on the BC at SRNL is synchronized with PTP messages from the GMC at ORNL. The BC also has a GPS connection, solely for the purpose of measuring the time offset (also known as phase offset) between PTP and Coordinated University Time (UTC), and the one-way network latency. Table 1 provides details GMC, BC, and router components used in this test configuration.

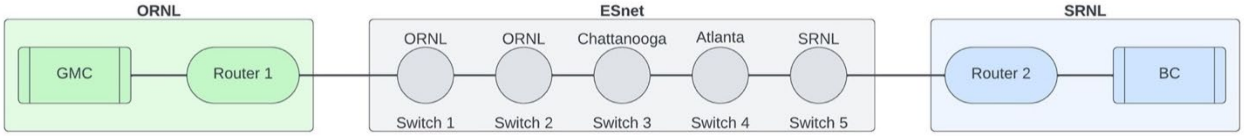

Figure 1: Network topology for PTP testing over ESnet.

---

[1] https://www.es.net/

[2] https://www.es.net/engineering-services/oscars/

The circuit is a Layer 2 circuit with IEEE 802.1Q[3] VLAN Tags used at the endpoint sites of the network path and Multiprotocol Label Switching (MPLS) used on the ESnet portion. IEEE 802.1AE (MACSec)[4] is enabled on the end points for data confidentiality, data integrity, and data origin authentication. Time synchronization between the GMC and BC is done using PTP version 2, described by IEEE 1588-2008[5]. The standard enables a hierarchical master-slave architecture for time distribution.

You will observe in Figure 2 the time offset ("phase offset") between the clock time on BC at SRNL (derived from PTP, coursed from the ORNL GMC) and GPS time at the same location is fairly consistent for much of the year, averaging around 250ns, with a noticeable shift around 13 December 2022. Figure 3, a graph of the network latency for between the two sites, also reflects this shift, but not at the same time/date.

Figure 2: Two-month plot of phase offset data between SRNL BC and GPS.

---

[3] "IEEE Standard for Local and Metropolitan Area Network--Bridges and Bridged Networks," in IEEE Std 802.1Q-2018 (Revision of IEEE Std 802.1Q-2014), vol., no., pp.1-1993, 6 July 2018, doi: 10.1109/IEEESTD.2018.8403927.

[4] "IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security," in IEEE Std 802.1AE-2018 (Revision of IEEE Std 802.1AE-2006), vol., no., pp.1-239, 26 Dec. 2018, doi: 10.1109/IEEESTD.2018.8585421.

[5] "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," in IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002), vol., no., pp.1-269, 24 July 2008, doi: 10.1109/IEEESTD.2008.4579760.

Figure 3: Two-month plot of one-way network latency between SRNL BC and ORNL GMC.

Figure 4 and Figure 5 show these plot areas zoomed in, highlighting both the significant shifts as well as the calendar offset of these two events. Next, we will discuss possible explanations of this, how to interpret these data, and the larger implications for time synchronization.

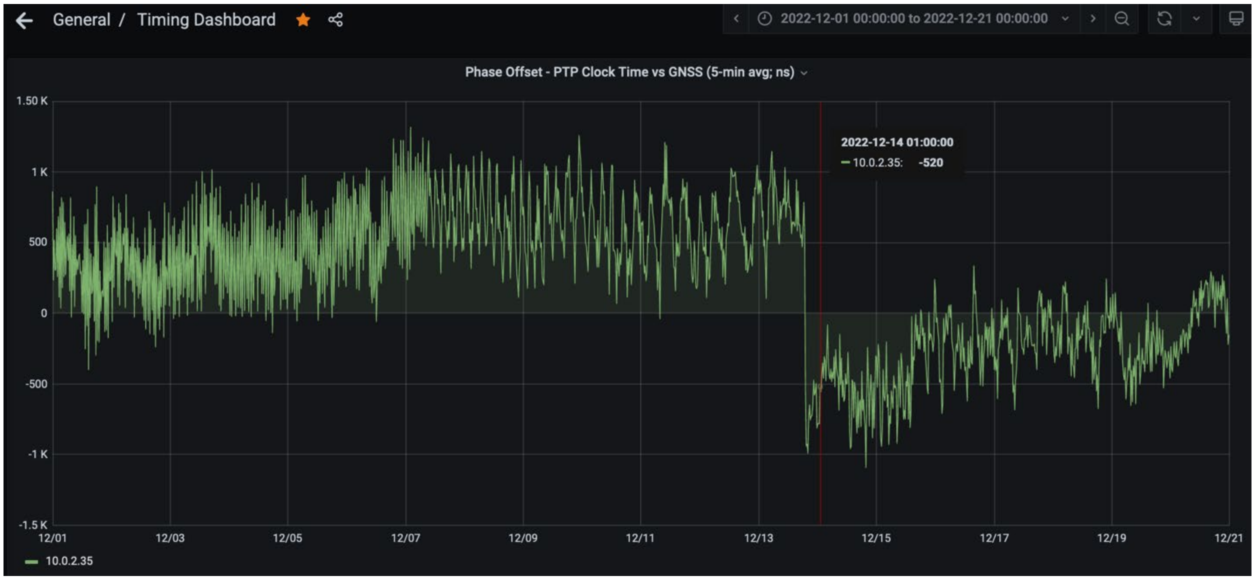

Figure 4: Zoomed in view of the dip in phase offset on 13 December 2023.

Figure 5: Zoomed in view of the network delay jump on 15 December 2022.

As seen in Figure 1, the network path from GMC to BC involves many network segments, maintained, and operated by different entities. The exact reasons for the instability cannot be precisely ascertained. However, there are some plausible explanations. Between 13 December and 15 December, it is likely that the PTP algorithms on the clocks detected delay asymmetry in their messaging back and forth to each other (see PTP standard documentation for more details on this). If the algorithm detected the BC was ahead of the GMC, the BC time is adjusted backwards, with an immediate trend of adjustment toward zero (shown in Figure 4). This trend toward zero indicates that the PTP algorithm, from subsequent PTP messages, recognized that the BC was now behind the GMC and is trying to adjust the time on the BC by moving it forward. This is followed by another drop and another movement up toward zero.

Querying the ESnet team, we learned that for OSCARS-provisioned circuits, if anything along the path fails to meet a quality of service (QoS) threshold, the circuit will automatically reroute to any other available path. The link through Atlanta had a significant issue from December to February, which led to an automated re-routing (fail over) of our data through Nashville on a different portion of the ESnet circuit, resulting in the changes we observed. As shown in Figure 5, we observed a spike in latency around 15 December, which was when the fail over tripped.

Because of our ability to monitor synchronization performance metrics across the network, we were able to understand network anomalies that we would otherwise not have known, in particular the automatic re-routing and other QoS metrics. No doubt, with further investigation, we will likely be able to identify other network artifacts through anomaly detection on the timing data. The power of the PTP protocol and this “visibility” into the networks we are using lies in the bi-directional communication between the clocks and their goal of calibrating out network effects to achieve symmetry in timing between the two nodes. By monitoring these signals and their performance, we have a new lens into the quality, and potentially even the security, of the communications networks supporting the grid.

Table 1: Pertinent timing and routing components used in PTP test.

Equipment	ORNL	SRNL
Clock	ADVA/Oscilloquartz OSA 5422	ADVA/Oscilloquartz OSA 5422
Cesium Reference	ADVA/Oscilloquartz OSA 3230B	N/A
Router	Juniper MX 10003	Juniper ACX 5448-M

For additional details, contact CAST@ornl.gov.

The Center for Alternative Synchronization and Timing (CAST) at Oak Ridge National Laboratory (ORNL) performs research, development, testing, evaluation, and technical assistance to enable resilient timing and synchronization for the power grid. Working closely with power utilities, timing hardware and software vendors, network operators, and federal stakeholders, CAST helps develop and validate alternative timing architectures to augment GPS time. CAST also translates and transfers ORNL's research and development (R&D) advances in secure timing and grid communications to power sector applications, and engages across the broader timing community to develop best practices to ensure the resilience of US critical infrastructure. CAST is sponsored by DOE's Office of Electricity. Visit https://cast.ornl.gov for more information.