ORNL/TM-2021/2331

Effects of the State of Network on PTP accuracy – An ESnet Case Study

Joseph Olatt, Chris Cooper & Carter Christopher

May 2023

OAK RIDGE National Laboratory

ORNL/TM-2021/2331

DarkNet, Electrification and Energy Infrastructures Division

Effects of the State of Network on PTP accuracy – An ESnet Case Study

Joseph Olatt (olattjv@ornl.gov)
Chris Cooper (coopercs@ornl.gov)
Carter Christopher (christophesc@ornl.gov)

May, 2023

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-Battelle LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

1. Introduction

This document presents some of the effects and observations of the "state of network" on Precision Time Protocol (PTP)[*] accuracy. "State of network" here refers to things like traffic on the network, network congestion, network routes, differences in paths taken in communication between two entities, etc. The network under consideration is ESnet[†].

This document will, hopefully, help future adopters of PTP plan their deployments such that they can derive the benefits of PTP without the pitfalls mentioned herein.

2. Background

The Department of Energy, specifically, Office of Electricity has partnered with Oak Ridge National Laboratory (ORNL) to lead an effort to find an alternative precision timing service for the nation's power grid. This is in response to the vulnerabilities identified in the Global Navigation Satellite Systme (GNSS), which Global Positioning System (GPS) is a part of.

Additionally, Executive Order 13905[‡] has highlighted the need for some sort of an alternative or backup timing solution.

ORNL has established a Timing Lab and has been testing various technologies and timing devices as part of this effort. PTP and off-the-shelf timing devices and network connections that support it are part of the testing and research being conducted.

One of the research items involved testing PTP performance over an encrypted Layer 2 network path using an ESnet circuit. For many months, the PTP accuracy was pretty good. But then, suddenly, instability in the PTP accuracy started occurring. In researching it, many observations were made and the dependence of PTP accuracy on the "state" of the network paths became quite clear.

3. Short Description of the PTP Protocol

The PTP protocol is described in the IEEE 1588 standard. PTP Version 2 is described in the IEEE 1588-2008 standard. The standard describes a hierarchical master-slave architecture for time distribution. In this architecture there can be multiple clocks at a given "network level" in the hierarchy. All clocks in a given "network level" are able to send and receive network broadcast messages from the other clocks at the same level.

A clock that has a single network interface (typically referred to as an ordinary clock) can be either a master clock or a slave clock. Master clocks distribute time. Slave clocks are at the end of the hierarchy, receiving time. There are clocks that have multiple network interfaces. One of those interfaces can be configured to be a slave clock receiving time from an upstream master. One or more other network interfaces may be configured as master clocks, distributing time to slave clocks further downstream. This type of a clock is called a boundary clock (BC). BCs adjust the time received from its master by calculating the network delay and applying it before relaying the time to its downstream clocks.

---

[*] https://en.wikipedia.org/wiki/Precision_Time_Protocol

[†] https://en.wikipedia.org/wiki/Energy_Sciences_Network

[‡] https://www.federalregister.gov/documents/2020/02/18/2020-03337/strengthening-national-resilience-through-responsible-use-of-positioning-navigation-and-timing

The clock at the root of this hierarchy is called a grandmaster clock (GMC). A GMC is elected by all the other boundary and ordinary clocks on that network segment, if configured. GMCs can also be manually configured by adjusting priority values.

Time synchronization using the PTP protocol is achieved as shown in figure 1[§]. In the figure "Time server" denotes the master clock. "n/w element" denotes the slave clock.

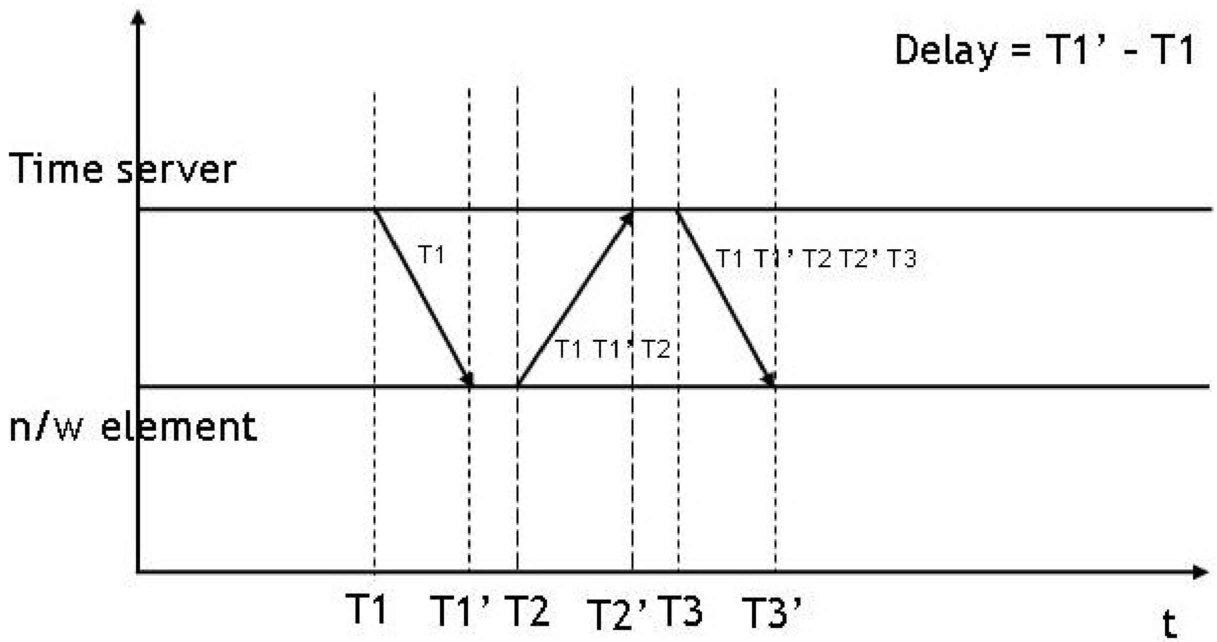

Figure 1. IEEE 1588 synchronisation mechanism and delay calculation

At time T1, a master clock broadcasts a *Sync* message to all the other clocks. That message reaches the slave clock at T1'. At time T2, the slave clock sends a PTP *Delay_Req* message to the master clock. The master clock receives that message at time T2'. The master clock now prepares a PTP *Delay_Resp* message that includes the timestamp T2' and sends it back to the slave clock. Using T1, T1', T2 and T2' the slave clock is able to calculate the average one-way delay, d, and its clock offset with the master clock, $\tilde{o}$, using:

$$\tilde{o} = \frac{1}{2}(T_1' - T_1 - T_2' + T_2) \quad (1)$$

$$d = \frac{1}{2}(T_1' - T_1 + T_2' - T_2) \quad (2)$$

An assumption made in the protocol specification is that the time taken for PTP messages to travel from master to slave and vice-versa is the same. If this is not true, it contributes to PTP inaccuracy.

[§]RadhaKrishna.Arvapally, Public domain, via Wikimedia Commons - https://commons.wikimedia.org/w/index.php?curid=31269989

4. PTP Setup and Network Path between the Clocks

A GMC, running at ORNL, distributes time to a boundary clock (BC) installed at a partner laboratory, Savannah River National Laboratory (SRNL). The network path between the two laboratories is as shown in figure 2. Outside of site-specific internal network segments, the remainder of the circuit is all ESnet.

The GMC is part of an Enhanced Primary Reference Time Clock (ePRTC) system. The GMC is disciplined by a multi-band GNSS receiver. A 10 MHz signal from an Magnetic Cesium atomic clock is also fed to the GMC and is used as a frequency reference.

Time on the BC at SRNL is synchronized with PTP messages from the GMC at ORNL. BC also has a GNSS feed for purposes of measuring the time offset (Phase Offset) between PTP and GNSS, and the one-way network latency.

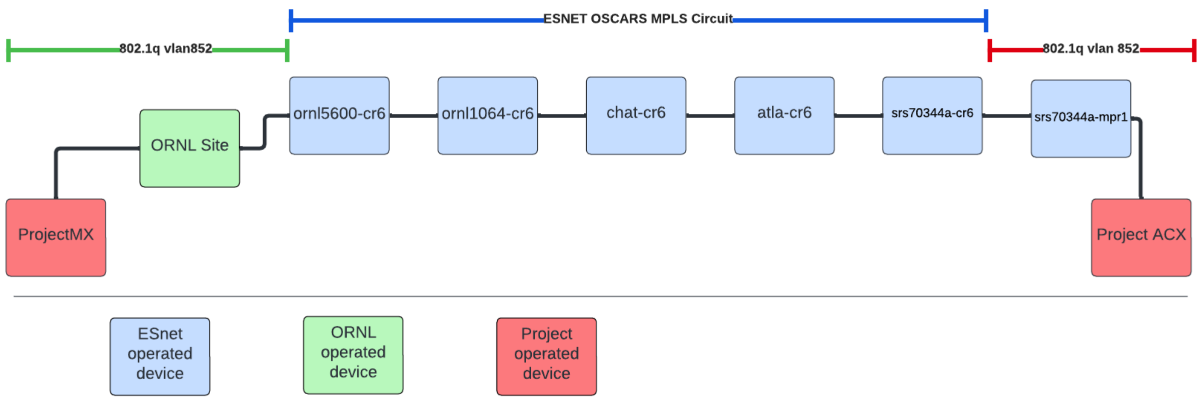

Figure 2. Network path between ORNL and SRNL

The entire circuit is a Layer 2 circuit with IEEE 802.1Q[¶] VLAN Tags used at the remote site portions of the network path and Multiprotocol Label Switching (MPLS)[‖] used on the ESnet portion. IEEE 802.1AE (MACSec)[**] is enabled on the end points for data confidentiality, data integrity and data origin authentication.

5. Observations

Figure 3 shows the time offset (Phase Offset) between the clock time on BC at SRNL vs. GNSS time at the same location. The mean and standard deviation for that period is shown in Table 1, row 1. One-way network delay/latency (BC to GMC) is shown in Figure 4. The mean and standard deviation of the network latency is shown in Table 1, row 2.

The one-way network delay/latency values were got from ADVA/Oscilloquartz's built-in measurement software called Syncjack Suite™. The specific probe/check is called "Slave to Master Probe". For this probe to work, the boundary clock has to have a local GNSS reference. The check/probe is calculated on

[¶] https://en.wikipedia.org/wiki/IEEE_802.1Q
[‖] https://en.wikipedia.org/wiki/Multiprotocol_Label_Switching
[**] https://en.wikipedia.org/wiki/IEEE_802.1AE

the master clock when the slave clock sends the PTP *Delay_Req* message alongwith the local GNSS timestamp at the time of the request. The probe software on the master then computes the actual time taken for that message to travel from the boundary clock to it, by using its own GNSS reference timestamp.

No.	Date Range	Measurement	Mean (ns)	Std. Dev
1	2022/10/10 - 2022/12/21	Phase Offset	245.108	910.192
2	2022/10/10 - 2022/12/21	One-way Network Latency	8273397.202	780322.821
3	2022/12/01 - 2022/12/21	Phase Offset	248.287	486.927
4	2022/12/01 - 2022/12/21	One-way Network Latency	2324370.085	3873261.612
5	2023/01/15 - 2023/02/24	Phase Offset	2312.116	7143.287
6	2023/01/15 - 2023/02/24	One-way Network Latency	2324370.085	3873261.612
7	2022/09/24 - 2023/03/21	Phase Offset	248.287	486.927
8	2022/09/24 - 2023/03/21	One-way Network Latency	3951785.707	4160714.853

Table 1. Phase Offset – Mean & Standard Deviation

Since this probe/check could not supply the actual time taken for the PTP *Sync* message to travel from master to boundary, the real network delay asymmetry could not be determined.

Figure 3. Phase Offset - PTP Clock Time vs GNSS (2022/10/10-2022/12/21)

Figure 3 and row 1 in Table 1 show that the PTP performance over a fairly long run on ESnet (see figure 2) for the observation time period was a mean value of sub 250 ns (nanoseconds). That is pretty good, especially when compared to the accuracies afforded by NTP[††].

In Figure 4, just before 12/16 on the X-axis, a large jump can be seen in network latency from BC to GMC. Figures 5 and 6 show a zoomed-in view of that area. In figure 5, between the X-axis values of 12/13 and 12/15 we see a rather large drop in the phase offset values and in figure 6, a large rise can be seen in the

[††]Network Time Protocol

Figure 4. Network Latency - SRNL vs ORNL GMC (2022/10/10-2022/12/21)

one-way network delay/latency. However, the drop and rise do not correlate. Upon reaching out to ESnet support, the following information was provided:

> "...the circuit is provisioned with a working path like below, if anything along that path fails, the circuit will automatically reroute to any available path. The link between atla-cr6 and srs70344a-cr6 is a circuit we procure from a telecom vendor which has a pretty significant issue from late December to early February which is why you saw the changes. Of note the working path being up actually corresponds with the higher metric in your graph. There is an overall lower latency path from Nashville to SRS which is what your failing over to, historically that path was limited to 2G and caused some significant jitter ..."

Figures 7 and 8 appear to confirm the information received from ESnet support. Between 02/07 and 02/20, the phase offset started hovering around 255 ns.

6. Reasons for Instability in PTP Accuracy

As seen in figure 2, the network path from GMC to BC involves many network segments, maintained and operated by various entities. The exact reasons for the instability cannot probably be ascertained. However, some plausible explanations can be posited.

In figure 5, between X-axis values of 12/13 and 12/15, it is possible that the *Sync* messages from GMC to BC took inordinately longer than the *Delay_Req* messages from BC to GMC. This is called a PTP delay asymmetry. As a result the value from equation 1 was a large positive number indicating that the BC was ahead of the GMC. So, the BC time was adjusted backwards. Immediately after that, the phase offset moving up to the zero on the Y-axis can be seen. This indicates that the PTP algorithm, from subsequent PTP messages, recognized that the BC was now behind the GMC and tries to adjust the time on the BC, by moving it forward. This is followed by another drop and another movement up to zero on the Y-axis. This indicates general instability in the network, due to some reason or the other.

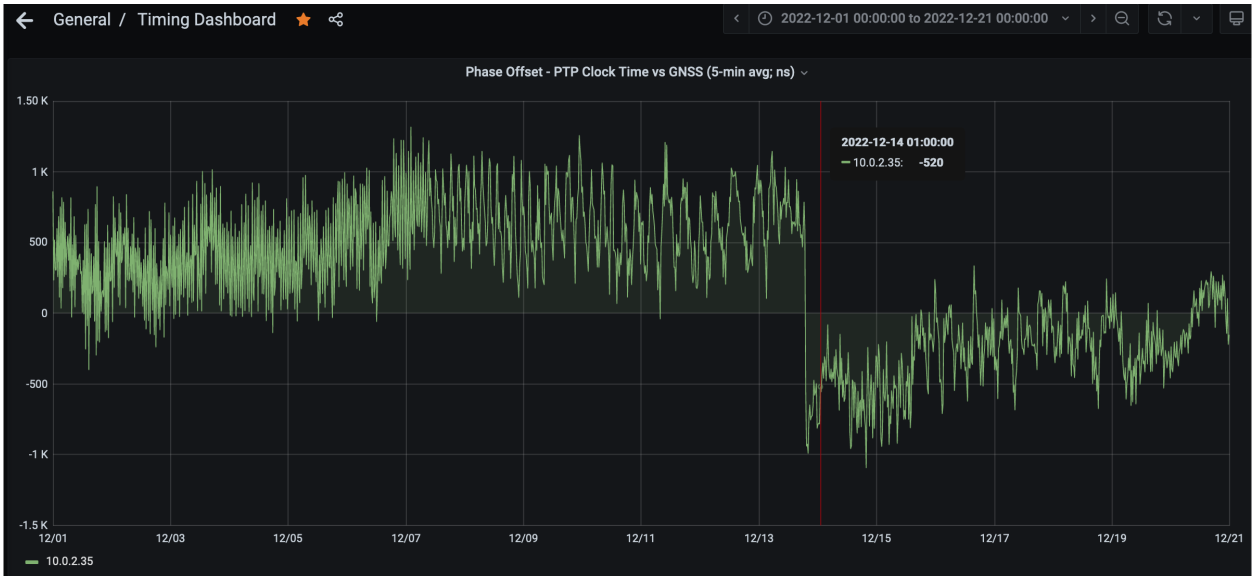

Figure 5. Phase Offset - PTP Clock Time vs GNSS (2022/12/01-2022/12/21)

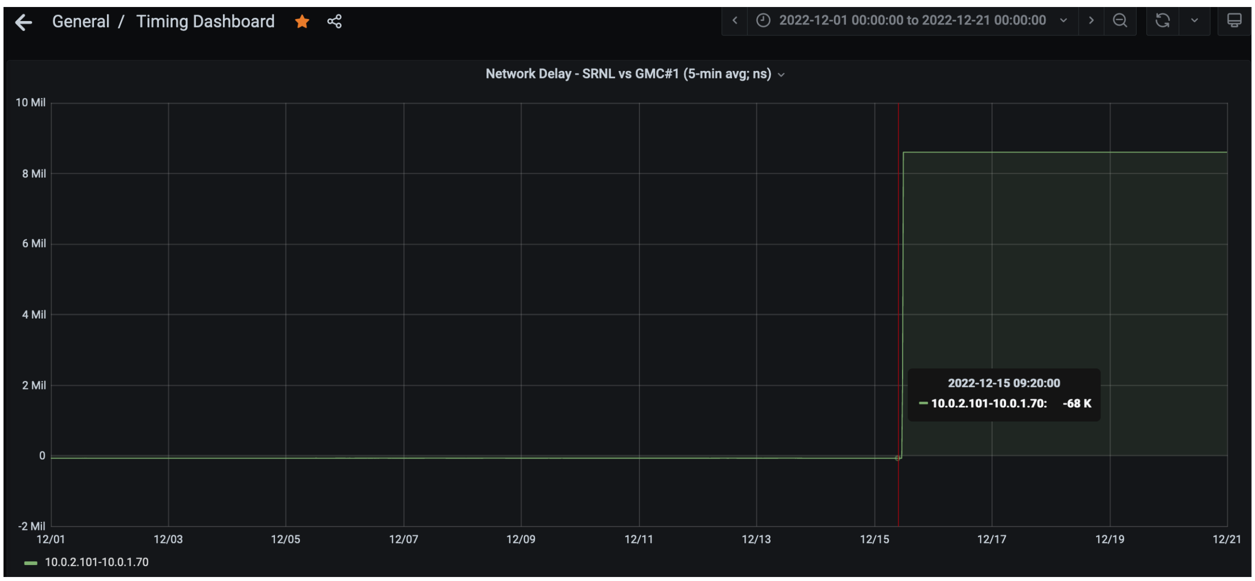

Figure 6. Network Latency - SRNL vs ORNL GMC (2022/12/01-2022/12/21)

Shortly after this, as depicted in figure 6, a large rise is seen. This is probably due to the automatic rerouting mentioned in the message received from ESnet support.

In figures 7 and 8, phase offset values and corresponding network latency values bouncing up and down, until just before 02/08 can be seen. According to the ESnet support message, quoted above, this was when the network issues were reported to have stabilized. Though, not for long!

Figure 7. Phase Offset - PTP Clock Time vs GNSS (2023/01/15-2023/02/24)

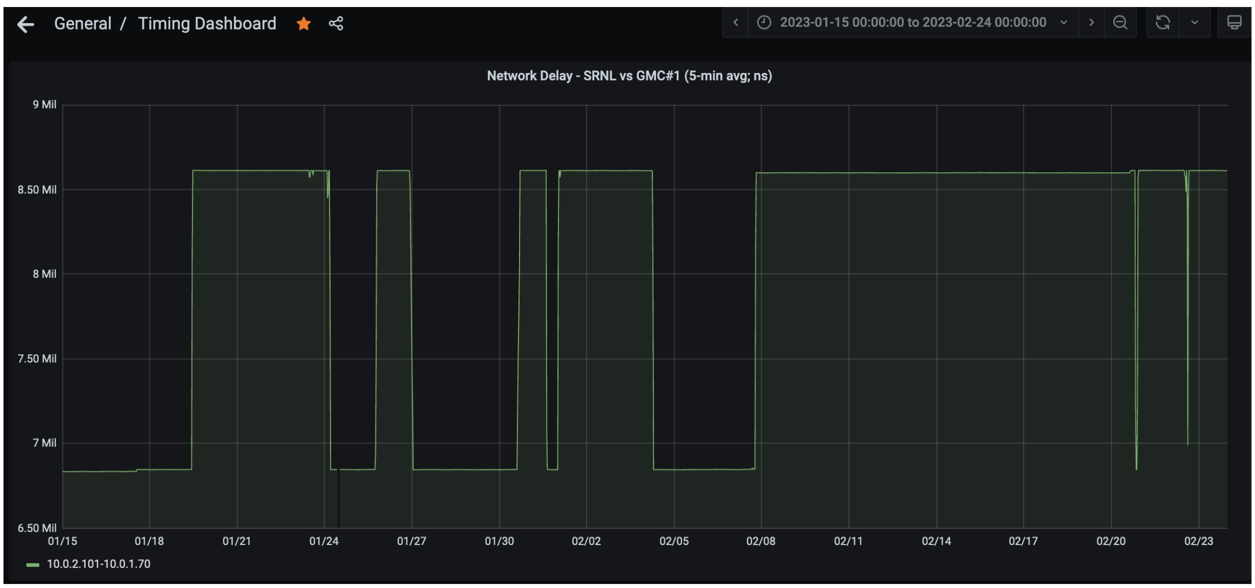

Figure 8. Network Latency - SRNL vs ORNL GMC (2023/01/15-2023/02/24)

7. Conclusions

On an Ethernet-type network, it would be almost impossible to see the time taken for a PTP message to travel from GMC to BC be the same as the time taken for a PTP message to travel from BC to GMC. According to the PTP protocol specification, the smaller the delay asymmetry values, the better the calculated offset and average one-way network delay, leading to better PTP accuracy.

It appears that on long network paths involving many network routers, network routes, etc. the accuracy of

Figure 9. Phase Offset - PTP Clock Time vs GNSS (2022/09/24-2023/03/21)

PTP is inconsistent and unpredictable–it very much depends on the state the network is in.

Figures 9 and 10 show that. The mean and standard deviation values are shown in table 1, rows 7 & 8, respectively. (Note: The blank sections in the graphs denote a period of no data collection due to outage of timing lab for hardware upgrades.)

From the all of the above, it can be seen that PTP accuracy that hovered around 250 ns dropped to around 2.3 μs (microseconds) due to network path issues.

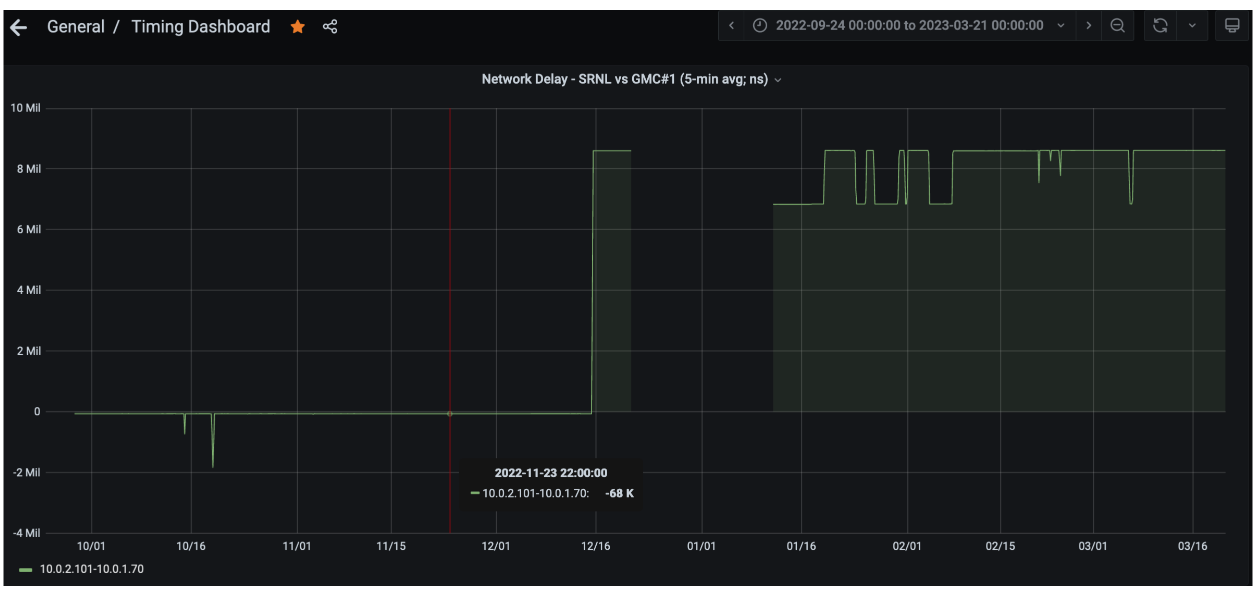

Figure 10. Network Latency - SRNL vs ORNL GMC (2022/09/24-2023/03/21)