Summary: Network Time Protocol (NTP) is a widely consumed standard for distributing time across network devices. It is also a legacy system with known security vulnerabilities. The vulnerabilities can be mitigated through modern implementations of NTP that include internal and external redundancy, ensure fault tolerance, and reduce cyber threats. While Precision Time Protocol (PTP) is a preferred approach, this technical bulletin provides recommendations for leveraging NTP internally when architectures require it or devices are not PTP-enabled.

Across the power generation, transmission, and distribution systems, wide-area time distribution is utilized to keep devices synchronized, from grid sensors to operational technology (OT) controls. Today, timing is generally provided in the field by Global Navigation Satellite Systems (GNSS), and it is typically distributed via Network Time Protocol (NTP). NTP is commonly supported across commercial-off-the-shelf grid components, both power electronics and communications systems. Given the well-documented vulnerabilities associated with GNSS signals [1], the Center for Alternative Synchronization and Timing (CAST) was established to perform R&D on solutions that can augment the US Global Positioning System (GPS) to provide a resilient timing infrastructure supporting generation, transmission, and distribution operations. This includes demonstrating and evaluating the opportunities of Precision Time Protocol (PTP) (IEEE 1588 [2]) as an alternative to NTP for timing distribution, details of which can be found in other CAST documentation. PTP, however, is not as broadly integrated into the grid hardware ecosystem, so it is imperative to understand the value and risks of NTP for distributing time across a network. While NTP does have its vulnerabilities, there are architectural approaches that can mitigate these to ensure robust and secure synchronization.

With implementations as far back as 1980 [3], NTP was first introduced as an internet standard in 1988 with Internet Engineering Task Force (IETF) Request for Comment (RFC) 1059 [4]. Now on version 4 since 2010, NTP clients and servers synchronize to the Coordinated Universal Time (UTC) timescale used by national laboratories and disseminated by radio, satellite, and internet modem [4]. Currently, most organizations consume NTP from public sources via the internet. Organizations like the National Institute of Standards and Technology (NIST), Google, or the NTP Pool Project supply a timing service over the internet in the form of NTP packets. Typically, an organization configures its network infrastructure to allow incoming NTP traffic from these reputable sources through firewall rules. By permitting NTP packets from trusted entities, organizations establish a conduit for synchronizing their internal devices with an external time source. Within this framework, typically a designated server (internally) assumes the role of an NTP source, disseminating time to clients on the internal network, and establishes a bidirectional communication channel with the external provider that ensures close synchronization with UTC. NTP communication leverages port 123, which is typically open on firewalls and routers to facilitate synchronization.

Figure 1 below shows how NTP from a public source is allowed through an organizations firewall.

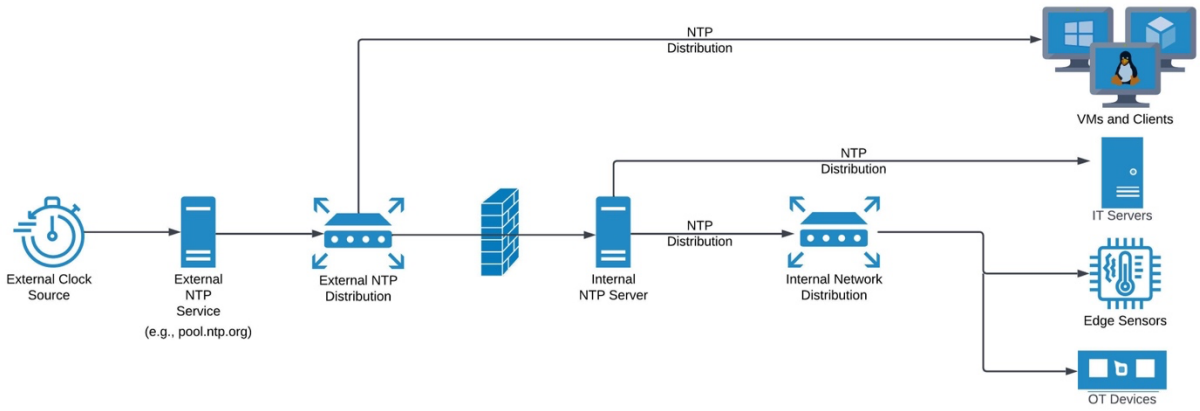

Figure 1: Standard NTP architecture with multiple diverse paths to NTP time, and persistently open firewall port for NTP server synchronization.

Utilizing NTP via a public service over the internet is a cheap and easy method to provide time and synchronization. Unfortunately, the practice of using public timing sources leaves organizations and their networks vulnerable to manipulation and attack. NTP over the internet can be intercepted and manipulated in various ways to cause harm to networks or shut them down all together. Particular NTP cyber vulnerabilities include:

- Man-in-the-Middle attack: an attacker intercepts and alters the NTP signals for malicious purposes, including deceiving systems into syncing with inaccurate time sources, obfuscating logs to mask attacker activities and causing potential system disruptions through inaccurate time synchronization. Man-in-the-Middle attack is form of active wiretapping attack in which the attacker intercepts and selectively modifies communicated data to masquerade as one or more of the entities involved in a communication association [7].

- Distributed Denial-of-Service attack (DDoS): an attacker floods a network with excessive traffic, overwhelming its capacity and disrupting normal operations. In 2014 the US-based company CloudFlare was hit with a massive DDoS attack using NTP. "CloudFlare chief executive Matthew Prince said the attack tipped 400Gbps, 100Gbps larger than the previous record DDoS attack which used DNS reflective amplification" [5].

- Masquerade attack: an attacker assumes the identity of a legitimate NTP source and then feeds incorrect NTP packets into a network. These attacks aim to disrupt operations and cripple services posing a significant risk to network integrity and operations. Masquerade attack is type of threat action whereby an unauthorized entity gains access to a system or performs a malicious act by illegitimately posing as an authorized entity [7].

NTP best practices can be found in RFC 8633: Network Time Protocol Best Practices [7]. These best practices specifically target version 4 of NTP, the default NTP version used in commercial equipment today. More information on NTP version 4 can be found in RFC5905 [6]. The recommendations are intended to help operators distribute time on their networks more accurately and securely, and they apply generally to a broad range of operations and networks. Some specific networks may have higher accuracy requirements that call for additional security techniques beyond what is documented. These best practices focus on general network security, time protocol-specific security, and NTP server/client configurations. Some of the best practices found in RFC 8633 are:

- Keeping NTP Up to Date: Users should keep up to date on any known attacks on their selected implementation and deploy updates containing security fixes as soon as it is practical. Updates can include firmware patches and bug fixes, and they should be done as soon is available or recommended by vendor. NTP users should select an implementation that is actively maintained.

- Using Enough Time Sources: An NTP implementation that is compliant with NTP version 4 takes the available public timing from providers such as NIST, Microsoft, and pool.ntp.org. The NTP sources are submited to sophisticated intersection, clustering, and combining algorithms to get the best estimate of the correct time. Multi-source NTP systems can also tolerate the loss of some time sources with marginal reduction in accuracy.

- Using a Diversity of Reference Clocks: Having a diversity of sources with independent implementations means that any one issue is less likely to cause a service interruption. This can include clocks from a more than one vendor (who may have different update/patch cycles), as well as both space and terrestrial reference sources. Having independent sources of truth ultimately helps a user more effectively identify issues with a time provider and maintain trust in system time.

- Monitoring: Operators should use remote monitoring capabilities for NTP to quickly identify servers that are out of sync and ensure correct functioning of the service. Operators should also monitor system logs for messages so that problems and abuse attempts can be quickly identified. For example, the CAST team uses the open-source time-series visualization platform Grafana to support monitoring, by visualizing raw clock data as well as analytics applied to these data. One could also implement higher level monitoring within a Network Operations Center (NOC), and using dedicated commercial NTP monitoring software.

Figure 2 shows how an internal authoritative time source can distribute secure NTP throughout a network.

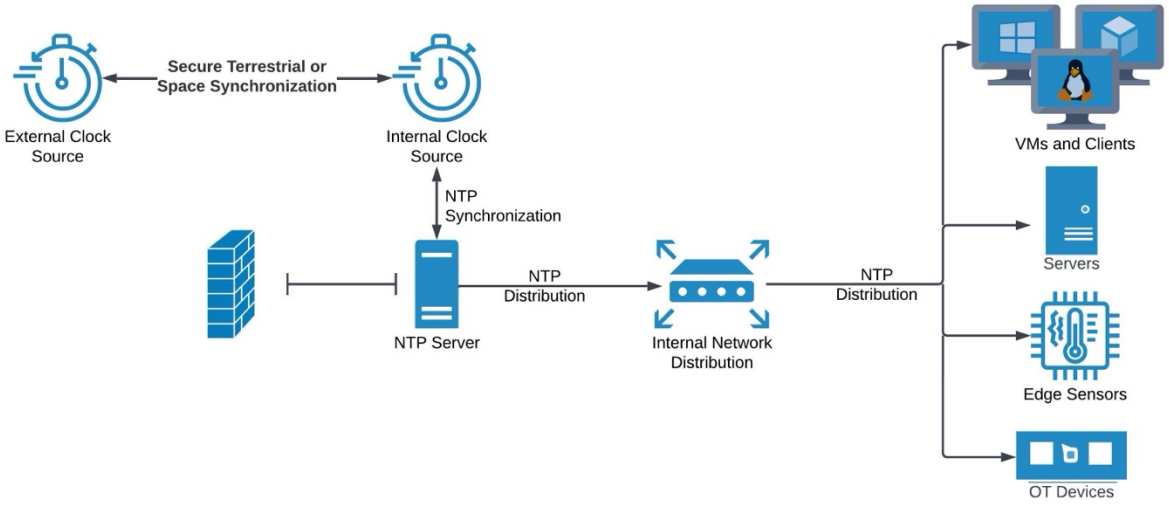

Figure 2: An internal clock source supporting NTP distribution reduces the attack surface and improves security for network time and synchronization.

An organization can bring timing and synchronization inside their network by introducing their own authoritative time source such as a Grand Master Clock (GMC). A Grand Master Clock typically uses GNSS as its time reference to UTC, instead of internet. While we have previously noted concerns around GNSS vulnerabilities, this is partially mitigated through two measures: (1) many clock vendor GMCs have built-in GNSS spoofing detection algorithms; and (2) GMCs can reliably deliver precision timing in a "holdover" state (meaning there is no external timing reference) for a number of weeks. (CAST is currently benchmarking multiple vendors for holdover stability, with reporting to be released in late 2024.)

PTP reduces the risk of internet cyber attacks by implementing a robust architecture that includes internal sources of time with a secure point-to-point synchronization. Doing so not only enables the provisioning of secure and highly precise time, but also closes the firewall pinhole required when relying on an external NTP server link/synchronization. Thus safeguarding critical infrastructure from potential outside cyber threats and ensuring continuing operation in the event of an emergency. PTP systems should include non-networked external references (e.g., GNSS) to prevent long-time scale drift, but these are not vulnerable to internet attacks.

In conclusion, while NTP has been crucial in synchronizing networks since its inception, it presents some potential vulnerabilities when using internet-based services as the primary source of time. Cyber vulnerabilities in NTP distribution can lead to severe cascading effects, from data loss to equipment damage and beyond. This potential for manipulation and attack highlights the need for organizations to fortify timing and synchronization by establishing internal authoritative timing sources on their networks. Yet, implementing these precise timing solutions can be a complicated endeavor. DOE's CAST is playing an important role in researching and benchmarking alternative timing and synchronization solutions to demonstrate their ability to support critical infrastructures like the power grid. Additionally, by equipping utilities with best practices and technical assistance, collectively we can bolster resilience and limit the potential impact of timing and synchronization disruptions from malicious actors.

References

[1] *Understanding Vulnerabilities of Positioning, Navigation, ...*, www.cisa.gov/sites/default/files/2023-04/fs_positioning-navigation-timing-vulnerabilities_508.pdf

[2] "IEEE SA - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems." IEEE Standards Association, standards.ieee.org/ieee/1588/4355/

[3] "History of NTP." NTPsec.org. https://docs.ntpsec.org/latest/history.html

[4] "Network Time Protocol (Version 1) Specification and Implementation." July 1988. https://www.rfc-editor.org/rfc/rfc1059.txt

[5] Pauli, Darren. "World's Largest Ddos Strikes Us, Europe." *iTnews*, 11 Feb. 2014, www.itnews.com.au/news/worlds-largest-ddos-strikes-us-europe-372033.

[6] Martin, Jim, et al. "RFC 5905: Network Time Protocol Version 4: Protocol and Algorithms Specification." IETF Datatracker, June 2010, datatracker.ietf.org/doc/html/rfc5905.

[7] Reilly, Denis, et al. "RFC 8633: Network Time Protocol Best Current Practices." *IETF Datatracker*, 2019, datatracker.ietf.org/doc/rfc8633/.

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.