Protecting Time Sync Integrity in Industrial Systems: Detecting PTP Manipulation

Jack Visoky
Principal Engineer and Security Architect
Rockwell Automation, Inc.

Kevin Riederer
Embedded Software Architect
Rockwell Automation, Inc.

Todd Wiese
Principal Software Systems Architect
Rockwell Automation, Inc.

Presented at the ODVA
2026 Industry Conference & 24th Annual Meeting
October 21, 2026
Vilamoura, Portugal

Abstract

Accurate and reliable time synchronization has become a foundational capability for modern industrial automation systems. Applications such as motion control, coordinated drives, distributed I/O, expiration of credentials, and time correlated diagnostics increasingly depend on accurate time synchronization to maintain a common system time across controllers and devices. IEEE 1588 PTP is widely adopted in these applications due to its precision and scalability. However, IEEE 1588 PTP, like many other time synchronization technologies, was not originally designed with security as a primary consideration. As industrial systems become more connected and exposed, manipulation of system time—whether malicious or unintentional—represents a growing risk that can impact performance, reliability, and safety.

This session paper explores approaches for identifying and responding to security related topics for time synchronization technologies in industrial devices that utilize IEEE 1588 PTP and other common technologies. Since IEEE 1588 PTP is so important for CIP Motion, special attention is given to it. A proposed solution for securing IEEE 1588 PTP is given, which monitors time integrity at the CIP application layer, where deviations from expected behavior can be detected with minimal impacts to existing installations. The paper details that approach and examines tradeoffs of that solution against other potential solutions.

Keywords

CIP Security, CIP Sync, Clock Anomaly Detection, Distributed Clock Systems, Industrial Automation Systems, Network Based Time Attacks, Precision Time Protocol (PTP - IEEE 1588), Time Sync Diagnostic Clock, Time Manipulation Detection, Time Synchronization Integrity

 ©2026 ODVA, Inc.

Definition of terms

ARP:	Address Resolution Protocol
BTCA:	Best Time Transmitter Clock Algorithm
CIP:	Common Industrial Protocol
IEEE:	Institute of Electric and Electronic Engineers
IPSEC:	Internet Protocol Security
NTP:	Network Time Protocol
MITM:	Man-In-The-Middle Attack
MACSEC:	Media Access Control Security
PLC:	Programmable Logic Controller
PTP:	Precision Time Protocol

Introduction

Reliable time synchronization has become a foundational requirement in modern industrial automation systems, enabling coordinated motion control, distributed I/O, event sequencing, and time-correlated diagnostics. Technologies such as IEEE 1588 Precision Time Protocol (PTP) are widely adopted due to their ability to provide sub-microsecond synchronization across distributed devices. This is commonly used in applications such as CIP Motion or distributed I/O blocks. However, as industrial systems become increasingly interconnected, the integrity of synchronized time has emerged as a critical concern, particularly in applications where deterministic behavior and system safety depend on a trusted common time base. Protecting time synchronization integrity is also becoming more important as the EU Cyber Resilience Act (CRA) and other emerging regulations establish stronger cybersecurity requirements for products with digital elements.

Despite its precision and scalability, IEEE 1588 PTP was not originally designed with security as a primary consideration, leaving it vulnerable to various forms of manipulation. Attacks such as packet tampering, delay injection, and rogue grandmaster insertion can alter time distribution without necessarily disrupting immediate effects on time synchronization, making these threats difficult to detect using conventional mechanisms. As a result, systems may appear synchronized while operating on an incorrect time reference, introducing risks to system performance, reliability, and safety.

This paper presents an approach for protecting time synchronization integrity in industrial systems by introducing application-layer detection mechanisms that complement existing PTP infrastructure. Specifically, it proposes a dual-clock architecture in which an independent diagnostic clock operates alongside the PTP-derived system clock, enabling continuous validation of time behavior at the device level. By leveraging standard CIP communication mechanisms and CIP Security data exchange, this approach enables detection of time anomalies with minimal impact to existing deployments, providing a practical path toward improving resilience against both unintentional errors and malicious manipulation of system time.

Time Sync Basics

Time synchronization protocols are integral for coordination of devices in distributed systems. There are several methods that are used in practice to synchronize time across systems. Commonly used time synchronization protocols are Network Time Protocol (NTP) and PTP (IEEE 1588). GPS is another technology that provides high-precision time synchronization, but it is less commonly used in industrial systems, so will not be discussed further here.

 ©2026 ODVA, Inc.

Network Time Protocol

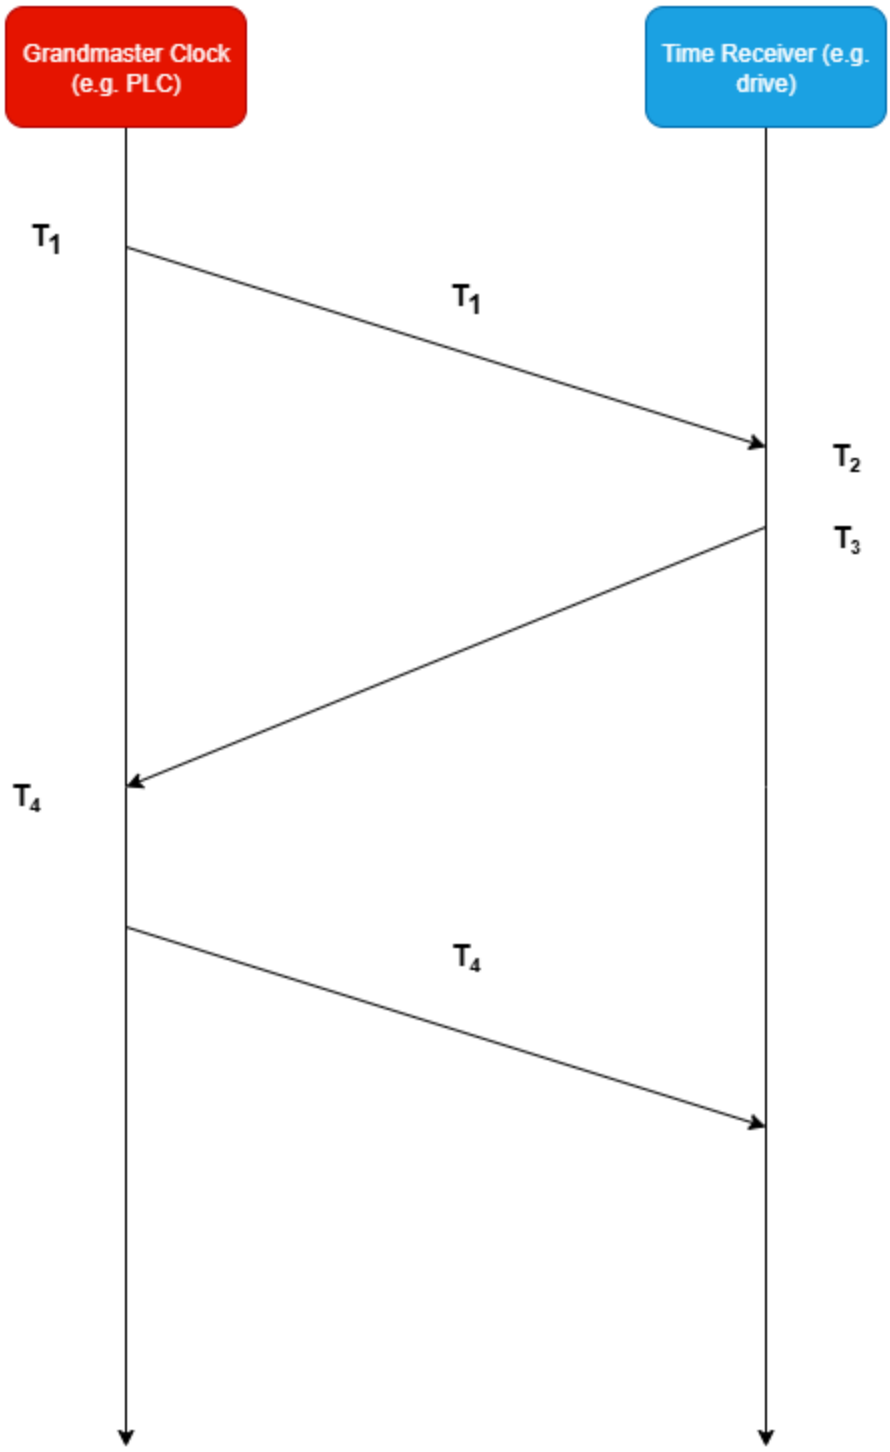
NTP synchronizes clocks across networked devices by having a client periodically exchanging timestamped messages with one or more time servers arranged in a hierarchical model. During each exchange, the client records the send and receive times and uses the server's timestamps to estimate network delay and clock offset, allowing it to adjust its local clock accordingly. By repeating this process and applying filtering to account for jitter and variable latency, NTP maintains relatively stable time synchronization across distributed systems, typically achieving millisecond-level accuracy on local networks. NTP is commonly used in applications where millisecond accuracy is sufficient, and communications are occurring over uncontrolled networks (e.g. internet). It does not rely on special hardware support to achieve this level of synchronization.

Precision Time Protocol

IEEE 1588 Precision Time Protocol (PTP) synchronizes clocks by organizing devices in a transmitter-receiver hierarchy, where a high-quality "grandmaster" clock distributes time to slave devices through periodic exchanges of precisely timestamped messages. Using these timestamps, the receiver computes both its clock offset relative to the transmitter and the network propagation delay, then continuously adjusts its local clock to align with the transmitter. This process is repeated regularly, often with hardware timestamping and compensation for intermediate switch delays, enabling sub-microsecond or even nanosecond-level synchronization in controlled networks. The diagram below illustrates how, with the sharing of four timestamps, the clock following can calculate its offset and network delay from a grandmaster clock.

 ©2026 ODVA, Inc.

The following equations show how the offset between grandmaster clock and time receiver are calculated:



$$T_{receiver} = T_{transmitter} + Offset$$

$$\mathbf{Offset = (T_2-T_1+T_4-T_3)/2}$$

 ©2026 ODVA, Inc.

Hardware Enhancements

To achieve sub-microsecond accuracy, PTP depends on several hardware enhancements. First, many ethernet NICs support hardware timestamping of PTP packets. By doing this in hardware, it reduces jitter that occurs as the packet is processed through the software stack. Lower jitter results in higher accuracy.

Additionally, PTP utilizes a frequency tunable clock, so allow synchronizing the frequency of the clock, rather than just the offset. Without this, the clocks can drift about during the time between PTP adjustments. This tunable clock functionality can be implemented in firmware, but hardware is often employed here to simplify the software application.

Finally, PTP supports ethernet switches that perform resident time corrections on PTP packets as they flow through the switch. This is referred to as a transparent clock. When a packet exits a switch node supporting transparent clock, it updates the packet with the amount of time spent in the switch. This field is used later to reduce the jitter in the clock follower's offset calculation.

Additional Features of PTP

PTP has additional features that make it both powerful and difficult to secure. For example, the configuration of a PTP clock domain is simple from a user perspective as it employs an algorithm where all potential clocks advertise their capabilities and the one with the most accurate (or otherwise highest priority) takes over and becomes the system-wide "grandmaster" clock. PTP also allows complex clock hierarchies, allowing it to scale to large systems. A system always has one "grandmaster" but can have other clocks in the system (boundary clocks) that synchronize to the "grandmaster" while downstream nodes synchronize to the boundary clock.

Potential Attacks

To evaluate the fitness of any given solution, it is important to first understand the nature of the attacks that are being protected against. In general, there are several categories of attacks to consider. These are discussed in the following section.

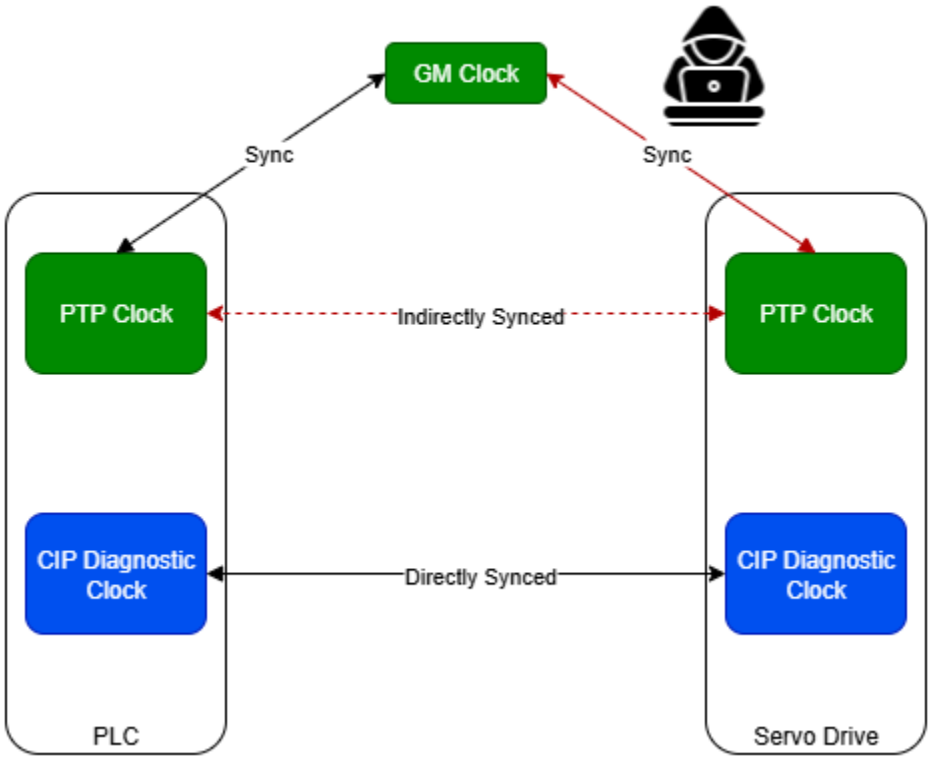
- **Clock Skew Attack:** A clock skew attack causes participating devices to disagree on the current time, resulting in a time offset between endpoints that should otherwise be synchronized. Potential consequences include loss of coordination between controllers and devices, incorrect sequencing of distributed actions, degraded motion synchronization, inaccurate event correlation, and difficulty diagnosing system behavior through timestamped data.

- **Clock Rate Attack:** A clock rate attack manipulates the rate at which synchronized clocks advance, causing multiple devices to share a common but incorrect notion of time. Because devices may remain synchronized to one another, these attacks can be especially difficult to detect. Potential consequences include gradual degradation of motion control performance, timing errors in scheduled I/O operations, incorrect expiration of time-dependent credentials or security functions, and accumulation of synchronization errors that affect system accuracy and reliability.

- **Denial-of-Service Attack:** A denial-of-service attack disrupts the exchange of PTP synchronization messages through packet loss, suppression, flooding, or other communication interference. As synchronization quality degrades, devices may enter a free-running state and begin drifting independently. Potential consequences include loss of synchronized motion, failure of coordinated control functions, increased system jitter, degraded network performance, and potential process downtime while synchronization is restored.

The following sections examine each attack category in greater detail and discuss representative attack techniques that can lead to these conditions.

Clock Skew Attack:

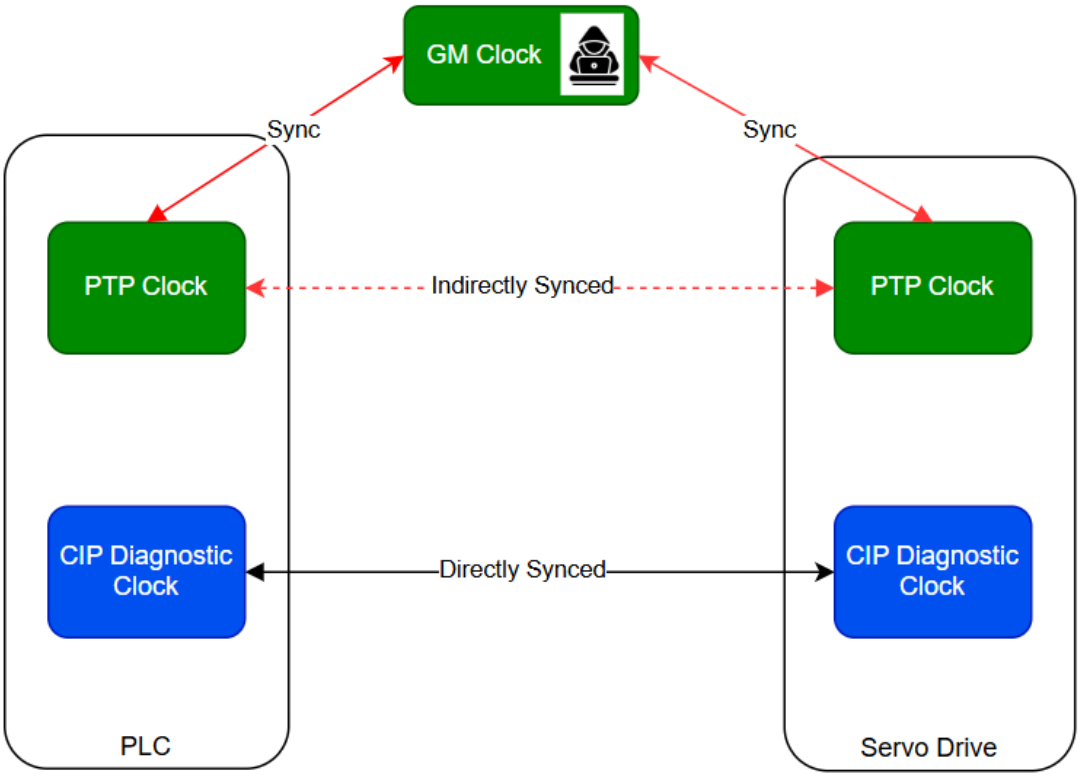


In this category of attack, a malicious actor is present on the local PTP network and is able to launch an attack such that the two application end nodes (e.g. PLC and servo drive) have a skew between their PTP times (e.g. they don't agree as to what time it is). Here are examples of attacks that can cause this condition:

- Packet Tampering via Man-in-the-Middle or Spoofing of Delay Req/Resp Packets – An attacker can intercept and modify Delay_req/Delay_resp PTP packets enroute between master clock and end node. This could be done via one of the many man-in-the-middle style attacks (e.g. ARP spoofing) to modify packets in flight, or by creating completely spoofed packets. By attacking these packets, which are specific to a given end node, an offset can be introduced between this node and the rest of the system.

- Delay Attack – Because PTP depends on a symmetric delay between sender and receiver, a man-in-the-middle attacker can intercept and delay packets going in one direction, thereby affecting the ability of the PTP client to accurately synchronize to the master clock. An attack like this need not even involve the packets directly but could be caused via changing the characteristics of the network to introduce extra unforeseen delays.

Clock Rate Attack:



Another category of attack occurs when an attacker manipulates the time on both ends of the application simultaneously. In this scenario, the clocks are synchronized, but they are ticking at the wrong rate – either faster or slower than reality. Here are examples of attacks that can cause this condition. Note the first two examples are the same as for the clock skew attack, but in this case the attacker is interested in changing the rate of the clock to be different from reality, not different between two individual nodes.

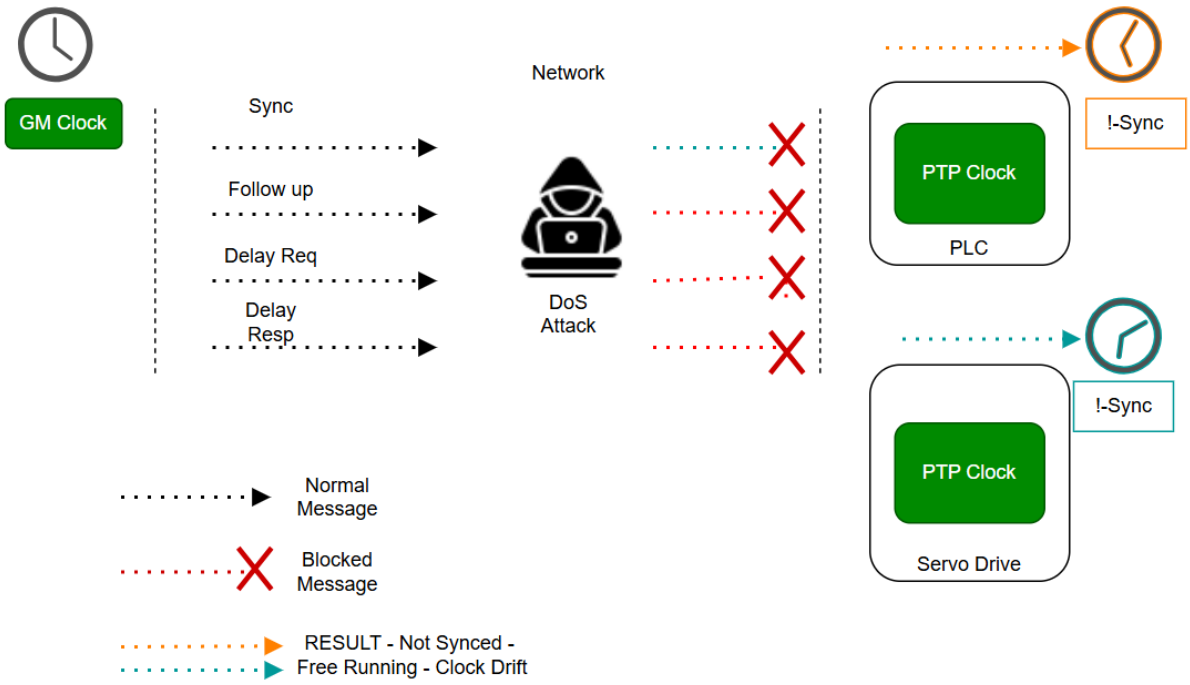
- Packet Tampering (man-in-the-middle) or Spoofing of Sync Packets – An attacker can intercept and modify PTP sync packets enroute between master clock and end node. This could be done via one of the many man-in-the-middle style attacks (e.g. ARP spoofing) to modify packets in flight, or by creating completely spoofed packets. By attacking these packets, which are global between many nodes, the attacked nodes will have similar times, but they can be wrong (either faster or slower than reality).

- Delay Attack – Because PTP depends on a symmetric delay between sender and receiver, a man-in-the-middle attacker can intercept and delay packets going in one direction, thereby affecting the ability of the PTP client to accurately synchronize to the master clock. An attack like this need not even involve the packets directly but could be caused via changing the characteristics of the network to introduce extra unforeseen delays.

- Rogue Grandmaster - In this type of attack, an attacker takes over the role of grandmaster in a PTP domain. It can do this because of the decentralized nature of the PTP grandmaster election process, which through the Best Time Transmitter Clock Algorithm (BTCA) elects the clock that claims to have the best attributes. An attacker can manipulate these attributes to raise its priority. Once an attacker has taken over as a rogue grandmaster, they can manipulate time, affecting all devices synchronizing to it.

- Time Source Degradation Attack – if an attacker can manipulate the underlying time source for a master, this will result in the clock rate of client devices being incorrect.

Denial of Service:



A third category of attack involves either intercepting and removing PTP packets, or a storm of PTP packets, leading to the PTP communication to not properly synchronize. This can lead to client clocks entering free-running mode and drifting out-of-sync.

Solution Goals and Selection Criteria

To identify an approach for protecting IEEE 1588 PTP time synchronization integrity in industrial automation systems, candidate solutions were evaluated against a common set of architectural and deployment criteria. The objective was to improve the ability of automation systems to detect loss of time integrity, including clock skew, clock rate manipulation, rogue grandmaster behavior, and related synchronization anomalies while preserving the deterministic clock synchronization performance characteristics required by motion control, scheduled I/O, and other time-sensitive applications. The evaluation also emphasized reuse of established security mechanisms, compatibility with existing CIP communication models, and the ability to support current deployed solutions without requiring wholesale replacement of system infrastructure.

A key conclusion of the evaluation was that deployment impact on existing installations must be treated as a first-order design consideration. Industrial automation systems frequently remain in service for many years, making large-scale replacement of controllers, drives, I/O modules, switches, or network infrastructure economically and operationally challenging. Consequently, solutions requiring pervasive infrastructure upgrades, extensive key-management deployment, or large modifications of existing IEEE 1588 implementations were considered less favorable. Preference was given to approaches that could coexist with existing CIP Sync deployments, enable incremental adoption, leverage existing CIP Security mechanisms, and avoid disruption to deployed brownfield systems that could have large downtime impacts. This emphasis on backward compatibility, staged rollout capability, and minimal infrastructure impact ultimately became one of the most heavily weighted selection criteria.

The resulting assessment balanced security effectiveness against deployment, implementation complexity, operational impact, and long-term interoperability. Candidate solutions were therefore evaluated not only on their ability to protect or detect manipulation of time synchronization, but also on how well they aligned with existing EtherNet/IP and CIP design patterns, their compatibility with heterogeneous industrial networks, their effect on installed systems, and their ability to be deployed incrementally as products and systems evolve.

Selection Criteria	Weighting	Rational
Impact on Existing Brownfield Systems / Backward Compatibility	High	Minimize disruption to deployed systems, avoid mandatory infrastructure replacement, and support staged adoption.
Ability to Detect Time Integrity Manipulation	High	Must detect clock skew, clock rate anomalies, rogue grandmaster conditions, and other synchronization integrity failures.
Leverage Existing CIP Security Mechanisms	High	Reuse established security capabilities rather than introduce entirely new security infrastructures.
Compatibility with Existing CIP Sync / EtherNet/IP Architectures	Medium	Preserve existing communication and connection models while enabling incremental enhancements.
Deployment and Operational Complexity	Medium	Reduce implementation effort, key management burden, and ecosystem-wide upgrade requirements.
Infrastructure Independence	Medium	Avoid dependence on specialized switches, transparent clocks, or network-wide upgrades.
Architectural Scalability and Future Alignment	Low	Important - but secondary - to ease of deployment and risk reduction for current industrial installations.

Potential Solutions

There are several security solutions suggested in the IEEE 1588 PTP specification, described in that specification as “prongs”. The specification does not mandate any of these solutions but rather suggests them as possible approaches, leaving the implementor to decide what approach would be most appropriate. Note the solutions suggested can be combined, and not every solution is likely to be appropriate for every use of PTP. As with any engineering decision, each approach comes with tradeoffs. A brief description is provided here. For more details, please see the IEEE 1588 PTP specification.

MACsec

MACsec is provided as a potential solution for securing PTP. In this approach each hop of the PTP network is secured via MACsec, providing data authenticity on each link of the PTP network. An advantage of this approach is that it works with transparent clocks, as they are part of the MACsec security system and can modify the packets as authenticated members. MACsec allows the transparent

clocks to fully participate in the security protections, whereas a more “end-to-end” security solution would not be well-suited to PTP packets being modified by transparent clocks as transparent clocks are typically implemented in network infrastructure like routers and switches. However, there are downsides to using MACsec. One constraint is that each hop (including transparent clocks) needs to implement MACsec. This becomes challenging when transparent clocks exist on networking infrastructure like switches and routers which now need to be part of the MACsec security model (note end-to-end security technologies typically are independent of the network infrastructure). Related to this is the key management challenges with MACsec; each one of the hops needs keys provisioned and managed, which increases the complexity of key management over a more end-to-end security technology like IPSec.

IPSec

IPSec provides similar information assurance guarantees to MACsec but is implemented at a higher layer and is not hop by hop but rather end-to-end. This alleviates some of the concerns with MACsec, especially around key management, but now transparent clocks are not able to modify the PTP packets since the transparent clock is not part of the IPSec tunnel. Therefore, a big constraint of IPSec is the difficulty in using transparent clocks. This can be alleviated with a dedicated IPSec tunnel to the transparent clocks, although now the network infrastructure is part of the security model, bringing the complexity disadvantage described in the MACsec discussion. In either case, key management is still an issue with IPSec, although the problem may be lesser than with MACsec if the boundary clock problem is ignored due to the smaller number of keys to manage.

Authentication TLV

The Authentication TLV option allows for the PTP application date to be cryptographically authenticated via keys for this specific purpose. This can be done through a shared group key that allows each receiver to authenticate the data it receives in line with the processing, or it could be done through an intermediary that denotes the transmitters authenticated for the receivers. In either case it would require specific changes to the PTP protocol and data processing, rather than IPSec or MACsec where the security exists at a different layer from the PTP data. This has the advantage of being specific for PTP, but the disadvantage is that changes are needed to the PTP data processing. Furthermore, this is a new security mechanism made specifically for PTP, rather than one which exists in more general-purpose libraries and implementations. This approach would still have similar key management concerns as IPSec and MACsec, although the degree of the key management problem is dependent on how the packets are processed. If each receiver authenticates the packets themselves then each needs a key for this, which is similar to the MACsec approach. If the intermediary is used then it is possible to reduce the number of keys needed, similar to the IPSec approach, depending on how many nodes participate in the authentication.

Hardware Changes for Cryptographic Protections

For IPSec, MACsec and Authentication TLV, there is a potential issue with PTP implemented directly in hardware. In these cases, it is generally infeasible to add the cryptographic support required to the existing hardware; therefore a user would need to install new products that support the chosen technology. This also raises the concern of technology compatibility; these three solutions do not interoperate seamlessly, and none are mandated by the IEEE specification. Therefore, a user would need to ensure that all of the PTP nodes support at least one of these technologies to leverage them to secure the PTP system. Also notable is that for both IPSec and MACsec, the IEEE spec mentions these may be implemented in a device for other reasons outside of PTP security, and if so, they could be leveraged for securing PTP as well. However, this is generally not a good assumption for EtherNet/IP devices, as CIP Security relies on TLS and DTLS rather than IPSec or MACsec. Although TLS/DTLS do not prevent the implementation of IPSec or MACsec, it does make it less likely that a device would implement a similar security technology like this when TLS/DTLS is supported.

Redundancy and Monitoring

The IEEE 1588 PTP specification also describes security mechanisms for PTP around redundancy and monitoring. In this case the PTP traffic itself does not have any direct security mechanisms such as cryptographic protection, rather the protection is done outside of the PTP data. Redundancy involves including more than one channel that would need to be compromised in order to compromise the PTP data. Ideally the channels do not have the same compromise paths, that is, a separate and independent effort is needed to compromise each channel. Monitoring involves using some kind of method to monitor the PTP data for irregularities that are likely to be evidence of tampering; for example, statistical methods looking for large, unexpected variations in timing. Both of these methods have an advantage in that they do not require the PTP data itself to be modified, therefore there is not likely to be a need to modify the hardware for these solutions. However, the downside is that these methods could be more difficult to implement in such a way as to robustly foil attacks, since they do not rely on standard information protections like cryptography or secure protocols. Even so, these remain interesting options for securing PTP that can provide significant benefit if done carefully.

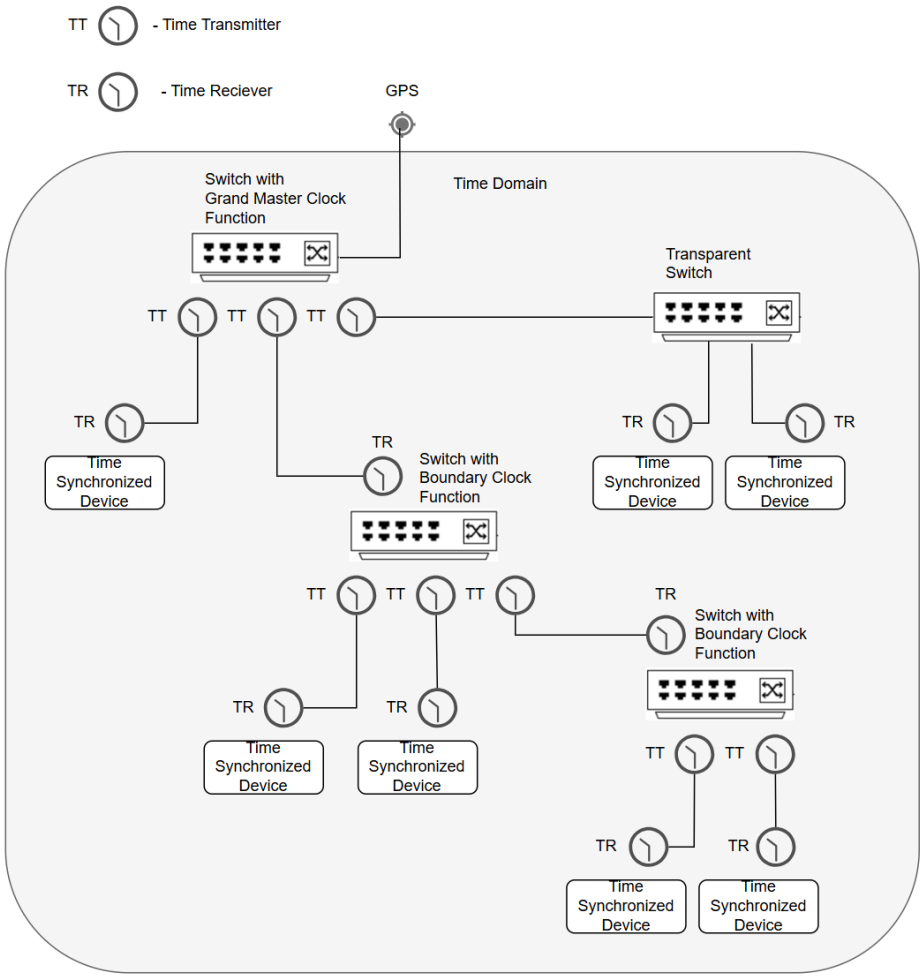
Proposed Solution

CIP Sync Present State:

CIP Sync provides time synchronization across industrial automation systems by leveraging the IEEE 1588 Precision Time Protocol (PTP). Through PTP, devices are synchronized to a common system time with sub-microsecond accuracy, enabling deterministic behavior for applications such as motion control, scheduled I/O, and event sequencing.

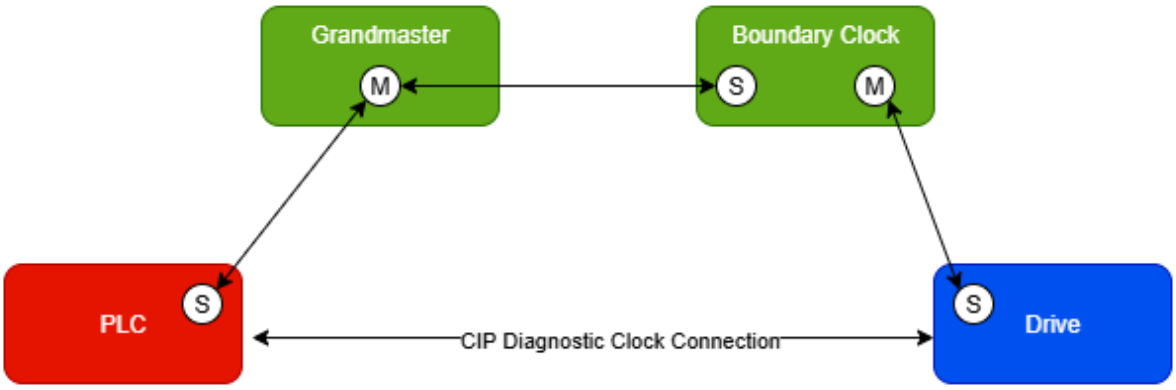
CIP Sync systems are composed of a hierarchy of clocks, including ordinary clocks, boundary clocks, transparent clocks, and time receivers. These roles define how devices participate in time distribution and how system time propagates throughout the network and time domain.

Hierarchy of Clocks:



Proposed Solution:

The proposed solution enhances existing CIP Sync time synchronization systems by introducing a secondary, independent diagnostic clock that validates the integrity of IEEE 1588 PTP-based system time. Rather than replacing PTP, which remains the primary mechanism for distributing precise time across the network, this approach establishes a parallel time reference operating on the application layer between controllers and devices. The diagnostic clock leverages standard CIP I/O communications to exchange timestamp information, enabling each device to compare the independently derived diagnostic time against the PTP-derived system time. This dual-clock architecture provides a mechanism to detect time integrity anomalies that would otherwise remain undetected within standard PTP operation.



The architecture is designed to integrate seamlessly into existing CIP deployments through IO connections, minimizing disruption to established network infrastructure and device implementations. PTP continues to provide high-accuracy synchronization across distributed devices, while the diagnostic clock introduces an additional layer of validation using secure, point-to-point communication protected by CIP Security. By embedding diagnostic timestamp data within Class 0/1 I/O connections, the solution enables continuous - runtime verification of synchronized behavior without requiring modifications to intermediate network infrastructure such as switches or transparent clocks. This approach ensures compatibility with both existing and future systems while enabling incremental adoption based on a device's implemented capabilities.

Through continuous comparison of the time sources, the system can detect conditions such as clock skew, clock rate anomalies, and grandmaster transitions that may result from network disturbances, configuration changes, or malicious manipulation. When time deviations exceed defined thresholds, devices can generate diagnostic events or faults, enabling higher-level equipment (e.g. PLC's) to take appropriate action based on system requirements. This detection-based strategy provides a method for improving the trust of time synchronization in industrial automation systems, aligning with emerging cybersecurity expectations by preserving the performance and architectural benefits of CIP Sync.

The proposed architecture also provides a mechanism for detecting certain rogue grandmaster scenarios. Because the diagnostic clock is established independently through secure CIP controller-to-device communications, it is not directly dependent upon the grandmaster selection process used by IEEE 1588 PTP. If a rogue grandmaster assumes control of the PTP domain and introduces a significant change in system time or clock rate, the PTP-derived system clock will diverge from the independently maintained diagnostic clock. This divergence can be detected through the same clock offset and rate comparison mechanisms used for other time integrity anomalies, allowing devices to identify unexpected grandmaster transitions or suspicious changes in system time and generate appropriate diagnostic indications. This approach does not prevent rogue grandmaster election but provides a means of detecting when a change in grandmaster behavior results in loss of time synchronization integrity.

It should be noted that detection depends on the diagnostic clock maintaining an independent and trusted time reference; scenarios in which both time sources are compromised in a coordinated manner may not be distinguishable from normal operation.

Capability Discovery:

To enable scalable deployment of secure time synchronization, devices may provide a standardized mechanism to advertise their time synchronization and diagnostic capabilities. A proposed spec enhancement will introduce the use of the CIP Sync Application Profile as a means for devices to declare their support for time synchronization capabilities and features. This mechanism allows devices to populate the Application Profiles attribute (25) within the Identity Object, allowing engineering tools and controllers to query and determine device capabilities.

The CIP Sync Application Profile proposes definition of a set of feature indicators that describe how a device participates in time synchronization, including supported clock roles (e.g., ordinary clock, boundary clock, or time receiver) as well as transport and implementation capabilities. Additionally, this profile extends beyond traditional role identification by including feature bits that identify functions, such as support for the Time Sync Diagnostic Clock capability. This addition provides a consistent method for determining whether a device can participate in the proposed dual-clock validation architecture, ensuring that systems can detect and configure devices based on their supported features.

Through this capability discovery approach, systems can support incremental adoption of time synchronization diagnostics while still maintaining backward compatibility with existing devices. Controllers and configuration tools then enable diagnostic clock functionality on devices that advertise support, while continuing to interoperate with legacy devices that don't. This design enables a migration path toward support for secure time synchronization architectures, without introducing disruption to established CIP Sync deployments.

Time Sync Diagnostic Object:

The proposed solution introduces a new Time Sync Diagnostic Object within a device CIP object model to provide a standardized mechanism for validating the integrity of IEEE 1588 PTP-based time synchronization. This object defines an interface to the application-layer diagnostic clock, operating independently of the existing Time Sync Object but enabling devices to detect discrepancies between the PTP-derived system time and the secondary clock, reference time maintained through secure controller-to-device communication. The object is designed to integrate into existing CIP Sync systems with minimal disruption and leverage standard Class 0/1 I/O connections and CIP Security to ensure secure exchange of diagnostic data and configuration parameters.

Rather than integrating with the current Time Sync Object, the Time Sync Diagnostic Object is implemented as a separate, multi-instance object, allowing devices to maintain independent diagnostic relationships for each I/O connection. This design enables per-connection validation of time synchronization behavior, which can be important in systems with multiple controllers or where multiple time domains may be present. By providing a consistent object model for diagnostic clock operation, configuration, and fault reporting, the solution ensures interoperability across devices and supports adoption within the CIP ecosystem.

Functional Role:

The Time Sync Diagnostic Object establishes an interface for the parallel diagnostic clock that operates independently from the IEEE 1588 PTP subsystem, enabling validation of time synchronization at the application endpoints. The diagnostic capability is constructed through timestamp exchanges embedded within the cyclic I/O communications between the controller (originator) and the device (target). By maintaining this independent time reference, the device can locally compare the diagnostic clock against the PTP-derived system clock to determine if the synchronization remains within acceptable bounds. This approach allows real-time detection of synchronization anomalies without requiring changes to the underlying PTP infrastructure. Because the diagnostic clock operates over secure CIP communication channels, it provides a trusted reference for validating PTP behavior and detecting conditions such as skew, delay asymmetry, or manipulation of the time distribution mechanism, ensuring that time integrity is evaluated where it directly impacts an application.

Object Characteristics:

The Time Sync Diagnostic Object is defined as a multi-instance supported object implemented on the target device, where each instance corresponds to a specific controller-to-device connection (and device represented synchronized clock). This allows devices to maintain independent diagnostic clocks for multiple concurrent connections, enabling fine-grained monitoring of synchronization behavior across application contexts. This design is needed because the diagnostic clock is not a global system function, but rather an independent connection-specific validation mechanism.

In addition, the object enforces secure configuration practices by requiring that any modification of settable attributes be performed through CIP Security mechanisms. This ensures that diagnostic parameters, such as thresholds and operational settings, cannot be altered without appropriate integrity protection. By combining per-connection instantiation with secure configuration requirements, the object provides the ability to manage time synchronization diagnostics across industrial devices.

Core Diagnostic Functions:

The Time Sync Diagnostic Object defines the set of attributes that enable devices to monitor, evaluate, and report the status of time synchronization. These attributes include mechanisms to report the status of diagnostic operation, measure the offset between the diagnostic clock and the originator reference, define acceptable clock skew thresholds, and report active fault conditions. In addition, the object maintains a record of fault history and provides a representation of the current diagnostic clock state, enabling both real-time monitoring and post-event analysis.

Through the capabilities, devices can continuously evaluate the relationship between the PTP-derived system time and the independently maintained diagnostic clock. This provides detection of abnormal timing behavior and supports adaptive thresholding approaches, where the acceptable deviation range may be automatically determined based on observed system conditions. The inclusion of diagnostic state and historical information further provides visibility into time synchronization performance, supporting troubleshooting and system validation activities.

Fault Detection Mechanisms:

The Time Sync Diagnostic capability and object also provides an interface to standardized fault detection mechanisms to identify conditions that may indicate loss of time synchronization integrity. These include detection of excessive clock skew between endpoints, abnormal clock rate adjustments, and transitions in the PTP grandmaster clock which may cause a “step in time” across the system. Each of these conditions represents a potential risk to deterministic system behavior and is evaluated against defined or dynamically calculated thresholds within the device.

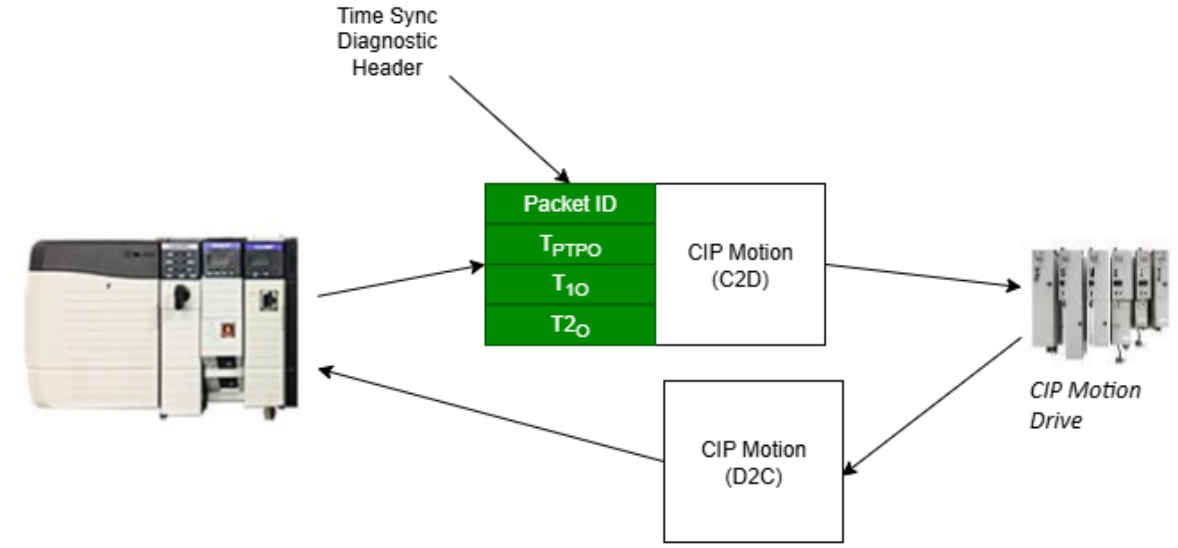
Detected faults are reported through a structured set of fault indicators and are typically latched until cleared, ensuring that transient or intermittent issues are not lost. These fault conditions may be propagated to higher-level application logic or control systems, enabling coordinated system responses to alarms, controlled shutdowns, or continued operation with notifications depending on the application requirements. By standardizing fault detection and reporting, the Time Sync Diagnostic Object enables consistent handling of time integrity events across CIP-compliant devices and applications.

System-Level Operation:

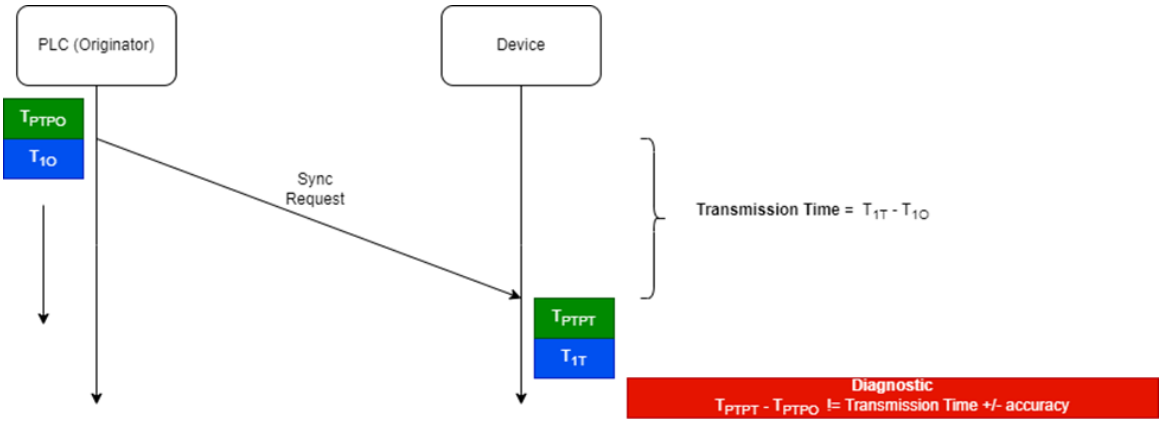
In a CIP Sync system implementing the proposed solution, IEEE 1588 PTP continues to serve as the primary mechanism for distributing system time across all devices, maintaining sub-microsecond synchronization through the established clock hierarchy. The Time Sync Diagnostic capability operates in parallel by introducing a point-to-point diagnostic clock between the controller (originator) and each target

device. This diagnostic clock is established and maintained over standard Class 0/1 I/O communications, allowing each device to independently validate the integrity of the PTP-derived system-time without altering the underlying PTP infrastructure.

The diagnostic clock is realized through the inclusion of Time Sync Diagnostic headers embedded within the cyclic I/O connection. These headers carry timestamp information captured at both the transmitting and receiving endpoints for each communication cycle. Specifically, it includes originator transmit timestamps, target receive timestamps, return path timestamps, and corresponding PTP timestamps. These values allow both endpoints, particularly the target device to compute clock offset using a symmetric time exchange model comparable to existing PTP message processing.



Using this data, the target device constructs and continuously maintains an independent diagnostic clock synchronized directly to the controller reference clock. The device then performs real-time comparison between the diagnostic clock and the local PTP-derived system clock. By evaluating differences in offset, delay, and rate of change between the two time sources, the system can detect discrepancies such as delays, clock skew, or abnormal clock rate adjustments. When these deviations exceed configured or dynamically calculated thresholds, the device flags a diagnostic condition, indicating potential issues such as network disturbances or malicious manipulation of PTP timing.



This system-level approach ensures that time synchronization integrity is validated at the application endpoints where it directly impacts control behavior, rather than relying solely on the correctness of the

distributed PTP infrastructure. By embedding diagnostic validation into standard, I/O communication messaging and leveraging a redundant time reference, the architecture provides continuous monitoring of synchronization health while maintaining compatibility with existing CIP Sync systems and enabling incremental deployment across heterogeneous devices currently being used in systems.

Time Sync Accuracy/Threshold detection:

One key element of the solution is the threshold for determining when a diagnostic fault has occurred. This threshold should be chosen carefully to balance the competing goals of detecting potential attacks, while also minimizing the possibility of false alarms.

The threshold value is dictated by how tightly the diagnostic clock can be synchronized. The tighter synchronization of this clock, the lower the threshold for detection can be. This synchronization accuracy can vary significantly depending on the system. Systems with less traffic and fewer switch hops will have tighter time synchronization than larger systems with more hops. This is also true of PTP as well, however because PTP can utilize PTP Transparent Clock functionality built into many standard ethernet switches to improve accuracy, its performance degrades slower than the diagnostic clock.

Instead of utilizing a “one size fits all” approach where we set the threshold to an expected worst-case performance, we will instead utilize a dynamic threshold. This allows the threshold to vary between systems depending on their network conditions. Our leading approach to this is to have a baseline phase early on during the CIP connection creation. This phase would determine a suitable threshold based on diagnostic clock performance measurements on the system and set the threshold accordingly. Once this threshold determination phase has completed, the user can read an attribute to determine what level of time protection is being provided by the time diagnostic.

Note that this is an area of active development, so the final algorithm has not been finalized. It is also expected to be product-specific, providing flexibility to the implementors.

The table below summarizes how the proposed diagnostics protect time:

Diagnostic	Description	Attacks Mitigated
Clock Skew Fault	Detects when the offset between the diagnostic clock and the PTP-derived system clock exceeds the configured clock-skew threshold.	Clock Skew Attacks - Asymmetric tampering or spoofing of time data, Rogue Grandmaster
Clock Rate Fault	Detects when the PTP subsystem adjusts the tunable clock beyond the rate limits specified by IEEE 1588.	Clock Rate Attacks – symmetric tampering/spoofing of time data, Rogue Grandmaster
Grandmaster Change	Detects when Grandmaster has changed. This is an indicator of a possible Rogue Grandmaster	Rogue Grandmaster

The following table provides a threat mapping and description of what effects and protection will be realized:

Threat	Description	Attack Effect	Protection
Packet Manipulation	Modification of network packets between the time transmitter and time receiver	Miscalculation of device offset leading to time skew	Clock Skew Diagnostic
Spoofing	Creation of new packets by a third party attributed to either the time transmitter or time receiver	Miscalculation of device offset leading to time skew	Clock Skew Diagnostic
Replay Attack	Repeat transmission of packets sent by the time transmitter or time receiver	Miscalculation of device offset leading to time skew	Clock Skew Diagnostic
Packet Interception and Removal	Packet Interception and Removal	DoS – times drift apart	Clock Skew Diagnostic
Packet Delay Manipulation	Manipulating the time to arrive for packets on the network	Miscalculation of device offset, time skew	Clock Skew Diagnostic
Denial of Service Against Time Protocol	Specific attacks that consume or destroy resources allocated to the time protocol	Miscalculation of device offset leading to time skew or clocks drift apart	Clock Skew Diagnostic
Grandmaster Time Transmitter Source Attack	Altering the primary source of time information in the network	Clock rate attack	Clock Rate Diagnostic

Solution Evaluation

Using the criteria defined, the solution proposed in this paper can be evaluated against the other solutions discussed in the IEEE specification.

Selection Criteria	Weighting	Authentication TLV	MACsec	IPSec	Proposed Solution
Impact on Existing Brownfield Systems / Backward Compatibility	High	Bad	Bad	Bad	Good
Ability to Detect Time Integrity Manipulation	High	Good	Good	Good	Medium
Leverage Existing CIP Security Mechanisms	High	Bad	Bad	Bad	Good
Compatibility with Existing CIP Sync / EtherNet/IP Architectures	Medium	Good	Bad	Bad	Good
Deployment and Operational Complexity	Medium	Bad	Bad	Medium	Good
Infrastructure Independence	Medium	Bad	Bad	Medum	Good
Architectural Scalability and Future Alignment	Low	Good	Good	Good	Good

 ©2026 ODVA, Inc.

Note that the Redundancy and Monitoring solution discussed in the IEEE PTP specification is not independently evaluated since the solution proposed in this paper is a specific implementation of these techniques.

The proposed solution scores very well on all criteria, with the only one not rated as good being the ability to detect time integrity manipulation. Although it certainly can do this, some of the proof requires tuning and implementation testing, whereas methods relying on existing cryptographic techniques would have a higher inherent assurance of being able to detect any manipulation. That said, the cryptography-based solutions score poorly on many of the other criteria due to their impact on existing PTP implementations, including hardware, as well as specific architecture considerations for these technologies. As such, the proposed solutions is deemed to be quite promising for protection of PTP within a CIP communication network. That said, the other techniques could still be valid depending on the needs and constraints of a given system.

Conclusion

Protecting the integrity of synchronized time is becoming increasingly important as industrial automation systems rely more heavily on precise timing for motion control, distributed I/O and other precision applications. While IEEE 1588 PTP provides the accuracy and scalability required by these applications, its limited security mechanisms create opportunities for malicious actors to exploit.

The CIP layer time diagnostic clock approach proposed in this paper provides a practical and incremental method for improving the security of PTP without replacing existing PTP network infrastructure. This approach leverages standard CIP communications and CIP Security protections, to establish a redundant diagnostic clock which continually monitors the PTP subsystem for potentially malicious attacks.

Although threshold selection and product-specific implementation details require further refinement, the proposed approach offers a practical path toward stronger time synchronization in both existing and future industrial automation systems.

References

[1] ODVA, The CIP Networks Library, Volume 1: Common Industrial Protocol, Ann Arbor: ODVA, Inc., 2001-2026.
[2] ODVA, The CIP Networks Library, Volume 2: EtherNet/IP Adaptation of CIP, Ann Arbor: ODVA, Inc., 1999-2026.
[3] ODVA, The CIP Networks Library, Volume 8, CIP Security, Ann Arbor: ODVA, Inc., 2015-2026.
[4] ODVA, The CIP Networks Library, Volume 9, CIP Motion, Ann Arbor: ODVA, Inc., 2016-2026.
[5] IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, IEEE Std 1588-2019.

****
The ideas, opinions, and recommendations expressed herein are intended to describe concepts of the author(s) for the possible use of ODVA technologies and do not reflect the ideas, opinions, and recommendation of ODVA per se. Because ODVA technologies may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, the reader and those responsible for specifying ODVA networks must determine for themselves the suitability and the suitability of ideas, opinions, and recommendations expressed herein for intended use. Copyright ©2026 ODVA, Inc. All rights reserved. For permission to reproduce excerpts of this material, with appropriate attribution to the author(s), please contact ODVA on: TEL +1 734-975-8840 FAX +1 734-922-0027 EMAIL odva@odva.org WEB www.odva.org. CIP, Common Industrial Protocol, CIP Energy, CIP Motion, CIP Safety, CIP Sync, CIP Security, CompoNet, ControlNet, DeviceNet, and EtherNet/IP are trademarks of ODVA, Inc. All other trademarks are property of their respective owners.