ODVA®

TECHNOLOGY OVERVIEW SERIES

OVERVIEW OF CIP SECURITY™

CIP Security™

AT A GLANCE

With IT/OT convergence being driven by IIoT and Industry 4.0, ODVA saw the need to enhance the defensive capability of devices connected to EtherNet/IP and other CIP Networks. This added approach is an important level of defense in a defense-in-depth architecture. The ultimate goal is to allow vendors to build interoperable EtherNet/IP devices that can defend themselves, the communications between them, and communications with third parties.

This approach is being realized through CIP Security™, ODVA's enhancement to *The EtherNet/IP Specification* for cybersecurity.

Introduction

Industrial automation networks were originally developed as a means to simplify the wiring of remote I/O devices and save wiring cost. Over time, this connectivity evolved to allow remote diagnostics and configuration of these devices. The Common Industrial Protocol (CIP™) is a peer-to-peer object-oriented protocol that provides connections between industrial devices (sensors, actuators) and higher-level devices (controllers). CIP has two primary purposes:

- Transport of control-oriented data associated with I/O devices
- Transport of other information that is related to the system being controlled, such as configuration parameters and diagnostics.

These networks were considered secure because they were physically isolated from other networks, they were constrained to geographies that could be secured by physical means (locked doors, etc.) and they could be monitored for unauthorized access. Over time, these once-isolated networks began getting connected with enterprise systems for the purpose of exchanging information to improve productivity, make better use of assets, energy savings and improved decision making. The value of this connectivity is obvious but it comes with certain security risks. These threats include: theft of intellectual property, tampering with plant systems, disruption of plant operations, and possibly damage to equipment.

In order to address these security issues, adoption of a defense-in-depth security architecture has been recommended for many years (see figure below). This architecture is based on the idea that multiple layers of security would be more resilient to attack. The expectation is that any one layer could be compromised at some point in time while the automation devices at the innermost layer would remain secure.

Figure 1: Defense-in-Depth Security

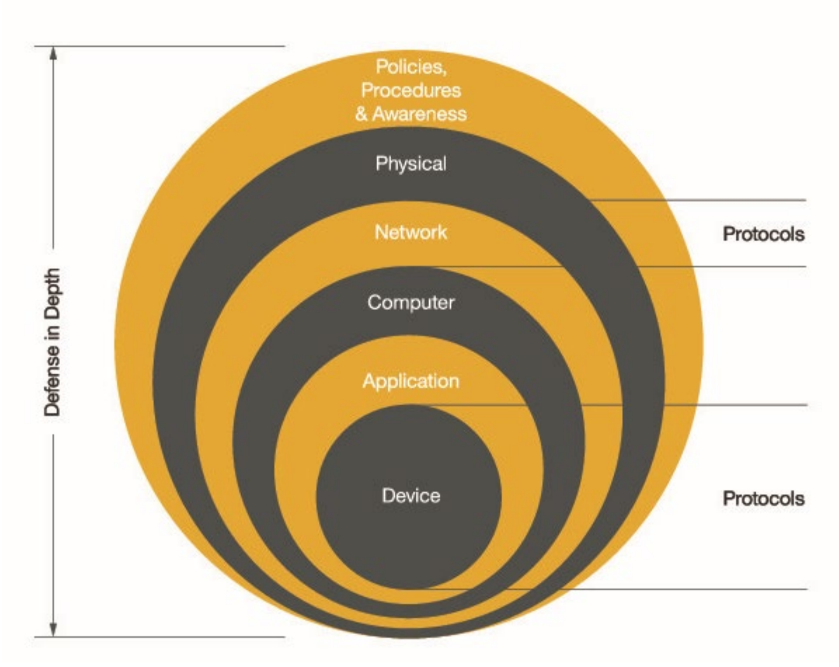

The goal of CIP Security is to improve the defensive capability of the CIP-connected device – a critical level of defense – in a defense-in-depth architecture. The ultimate goal of CIP Security is to build CIP devices that are able to defend themselves.

A fully self-defending CIP device would be able to:

- Reject data that has been altered (integrity)
- Reject messages sent by untrusted people or untrusted devices (authenticity)
- Reject messages that request actions that are not allowed (authorization)

CIP Security makes the following basic assumptions (the first three presuppositions are principles also found in "Zero Trust"):

- The network connected to the device should generally be considered untrusted
- All entities – both people and devices -- that attach to the network are considered untrusted until they can be authenticated
- Network access to a device should not be allowed until authorized by the device
- Physical access to a device will be limited to only trusted individuals (this is not covered by this specification)

Security Threats and Attack Vectors

It is important to understand the security threats and attack vectors to which a CIP device may be subjected, in order to mitigate those threats.

STRIDE is a system developed by Microsoft for thinking about and modeling security threats. It provides a mnemonic for security threats in six categories. The threat categories are:

- **S**poofing of user identity
- **T**ampering
- **R**epudiation
- **I**nformation disclosure (privacy breach or data leak)
- **D**enial of Service (DoS)
- **E**levation of privilege

The STRIDE name comes from the initials of the six threat categories listed. It was initially proposed for threat modeling, but is now used more broadly. The Microsoft-developed STRIDE [1] model is a tool that can be used to evaluate security threats.

The following table lists the different STRIDE threat types and security properties that apply to each.

Table 1: STRIDE

Threat Type	Threat Description	Security Property
Spoofing identity	An example of identity spoofing is illegally accessing and then using another user's authentication information, such as username and password.	Authentication
Tampering with data	Data tampering involves the malicious modification of data. Examples include unauthorized changes made to persistent data, such as that held in a database, and the alteration of data as it flows between two computers over an open network, such as the Internet.	Integrity
Repudiation	Repudiation threats are associated with users or devices who deny performing an action without other parties having any way to prove otherwise. Nonrepudiation refers to the ability of a system to counter repudiation threats. For example, a user who purchases an item might have to sign for the item upon receipt. The vendor can then use the signed receipt as evidence that the user did receive the package.	Non-repudiation
Information disclosure	Information disclosure threats involve the exposure of information to individuals who are not supposed to have access to it—for example, the ability of users to read a file that they were not granted access to, or the ability of an intruder to read data in transit between two computers.	Confidentiality
Denial of service	Denial of service (DoS) attacks deny service to valid users—for example, by making a Web server temporarily unavailable or unusable. You must protect against certain types of DoS threats simply to improve system availability and reliability.	Availability
Elevation of privilege	In this type of threat, an unprivileged user gains privileged access and thereby has sufficient access to compromise or destroy the entire system. Elevation of privilege threats include those situations in which an attacker has effectively penetrated all system defenses and become part of the trusted system itself, a dangerous situation indeed.	Authorization

Given the general description of STRIDE threat types in Table 1, the following table presents the threats that may apply to CIP based devices:

Table 2: Threat Description in CIP Data Flow Mapped to STRIDE

Threat Type	Threat Description in CIP Data Flow	Security Property
Spoofing identity	**Unauthorized session:** An attacker is able to establish a CIP connection to a target device and send arbitrary CIP packets. **Session hijacking:** An attacker is able to hijack an existing CIP connection and send arbitrary CIP packets. **Message replay:** An attacker is able to capture valid CIP packets and replay them at a later time. **Rogue server:** An attacker is able to spoof the identity of a valid server and accept messages from an unknowing client. **Notes** The source of the malicious messages could be the attacker's device connected to the network at a point of attachment (e.g., switch port), or could be a compromised device already on the network.	Authentication
Tampering with data	**Message alteration:** An attacker is able to intercept and alter or drop CIP packets in a man-in-the-middle (MITM) attack.	Integrity
Repudiation	**Log alteration:** An attacker is able to tamper with a local audit log, crash dump file or diagnostic file on a device. The user would have no ability to assure that the file was originally created by a specific device. This is a major concern in regulated industries where validated audit records are common place.	Non-repudiation
Information disclosure	**Message eavesdropping:** An attacker is able to capture CIP messages between two end points and see their contents.	Confidentiality
Denial of service	A number of threats listed above could result in denial-of-service, by virtue of the sending malicious messages to the CIP end point: 1. Unauthorized session 2. Session hijacking 3. Message alteration 4. Message replay	Availability
Elevation of privilege	**Unauthorized change:** An attacker with permissions of "get only" access to a CIP object somehow elevates the permissions to include both "get and set" access. Since legacy CIP does not support user authentication, every user and attacker has the highest access privilege that the object is designed to support. This is the problem that adding user authentication and device authorization solves.	Authorization

When using STRIDE, the items in the threat-mitigation table below represent possible techniques that can be employed to mitigate the threats shown in Table 2:

Table 3: Possible Techniques to Mitigate Threats

Threat Type	Threat Description in CIP Data Flow
Spoofing identity	Appropriate authentication Protect secret data Don't store secrets
Tampering with data	Appropriate authorization Hashes MACs Digital signatures Tamper resistant protocols
Repudiation	MACs Digital signatures Timestamps Audit trails
Information disclosure	Authorization Privacy-enhanced protocols Encryption of Data Protection of Secrets Not storing Unnecessary Secrets
Denial of service	Appropriate authentication Appropriate authorization Filtering Throttling Quality of service
Elevation of privilege	Run with least privilege

Additional threat modelling approaches to minimize the abilities for attackers to be successful include DREAD and Attack Tree. The DREAD threat model includes damage, reproducibility, exploitability, affected users, and discoverability. Attack Trees or Threat Trees are a logical approach to better understand how an intrusion could happen in a detailed, logical, and step by step manner.

Regardless of the threat modelling approach employed, it’s important to conduct these analyses on a regular basis to take into account changes in manufacturing footprints and newly discovered threats. Another valuable approach is to think of security as a state of mind within an organization. The best policies, procedures, and systems can be overcome, but a host of vigilant employees can be much harder to defeat. Security is an invaluable investment that should be treated with the same care as safety when designing, updating, and operating an industrial facility.

CIP Security Threat Model

CIP Security offers a sample threat model as an appendix to the EtherNet/IP specification to allow vendors and users to better understand potential security threats, and specifically how CIP Security can either prevent or mitigate the effects of cyber-attacks against the system. This is particularly helpful when seeking certification for a product or system that implements CIP Security.

A threat model is an essential part of designing a system with information security assurances. Available threat models such as STRIDE, DREAD, or Attack Tree analyze CIP Security in a general sense and are not intended to apply to a specific product implementation. These threat models can be useful to product vendors and end users who must analyze the information assurances and threat mitigations of a given product or system. Threats are dispositioned as accepted or mitigated within these threat models. However, these dispositions are provided as a starting point for a product or system specific threat model; an analysis of a given product or system may conclude on different dispositions due to unique risks and attributes. General threat models are meant to provide guidance, but not meant to be the final and complete document for any given product or system that uses CIP Security.

The CIP Security threat model lists critical assets, details the scope of the model, and what the trust boundaries are. Specific threats are outlined, such as threats against discovery, spoofing, tampering, and information disclosure, denial of service, threats against provisioning, threats against data in transit, threats against configuration data, elevation of privilege, threats related to redirection of communication, and threats against cryptography. Visit odva.org to obtain the latest version of The EtherNet/IP Specification including CIP Security.

Figure 2: CIP Security Threat Model Overview

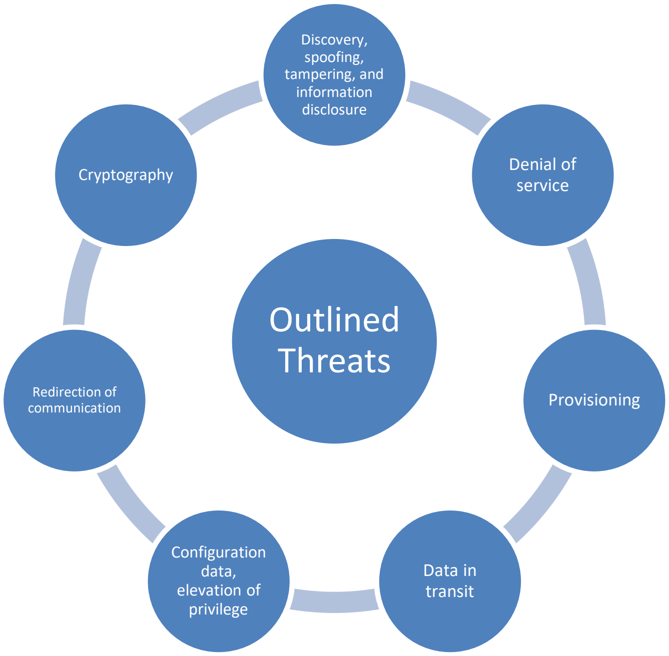
Outlined Threats
Discovery, spoofing, tampering, and information disclosure
Denial of service
Provisioning
Data in transit
Configuration data, elevation of privilege
Redirection of communication
Cryptography

CIP Security Approach

CIP Security specifies security-related requirements and capabilities for CIP devices. CIP Security comprises Volume 8 of *The EtherNet/IP Specification* and includes material that is network-independent as well as material that is CIP network-specific (e.g., EtherNet/IP).

The specification at present is focused on EtherNet/IP, as EtherNet/IP-connected devices represent the largest risk due to enterprise network connectivity. The specification at present defines the mechanisms, common behaviors, and requirements to provide a secure transport for EtherNet/IP communications. Additional CIP Security material will be added to the specification over time to address additional security properties.

It is not required that all CIP Security enabled devices provide support for all CIP Security properties, however, it is very important for customers of CIP Security enabled products to easily determine the security properties that are supported by the products they are purchasing. In order to simplify the ability for a customer to identify which products support a specific set of security features, a set of Security Profiles are available as shown in the table below.

CIP Security provides device authentication, a trust domain (both broad across a group of devices and narrow by user and role), device identity (including user), device integrity, data confidentiality, user authentication, policy enforcement (authorization), and fixed user authentication. This is accomplished through five separate security profiles that provide flexibility for vendors in adding security features to their device depending on the intended application(s) and use case(s). A security profile is a set of well-defined capabilities to facilitate device interoperability and end-user selection of devices with the appropriate security capability.

The first security profile is the EtherNet/IP Confidentiality Profile, which provides secure communications between EtherNet/IP endpoints to assure endpoint authentication, data confidentiality, and data authenticity. The second profile is the CIP User Authentication Profile, which provides Authentication at a user level for CIP communications. This is used as a basis for Authorization and Role Based Access Control. CIP Security's ability to authenticate via the device or through a central server allows for simplicity in smaller, simple systems and efficiency in large, complicated installations. In the future, CIP Security may make use of a CIP authorization profile that will enhance CIP to provide additional security properties such as general, flexible authorization where access policy can be based on any attribute of the user and/or system. The third profile is the Resource-Constrained CIP Security Profile, which provides a lightweight version of the protections afforded by the first two CIP Security profiles specifically for highly resource-constrained devices. Access policy information is included to allow a more capable device, such as a gateway, to be used as a proxy for user authentication and authorization of a resource constrained device. Implementation of CIP Security for resource-constrained devices requires only DTLS (Datagram Transport Layer Security) support instead of DTLS and TLS (Transport Layer Security), as it is used only with low-overhead UDP communication.

The fourth security profile is the Pull Model Profile that enables ease of use for device replacement and commissioning, using EST and DNS-SD technologies. Certificates involve a private key that is stored on a device. When a device fails it needs a brand-new certificate. The Pull Model allows a device to automatically discover and request a certificate using DNS-SD for discovery and EST for certificate request. The automatic discovery and request/grant of a certificate allows automatic device replacement to proceed even when security is being used. The fifth security profile is the Device-Based Firewall Profile, which provides a simple mechanism to filter traffic based on IP Address/port/protocol. The Device-Based Firewall works much like the “IP Tables” program that has been present in Linux/Unix for many years, and is implemented via a new object (called the Ingress Egress Object).

Table 4: Supported Security Profiles

Security Profile	General Description
EtherNet/IP Confidentiality Profile	Provides secure communications between EtherNet/IP endpoints to assure data confidentiality. Includes the EtherNet/IP Integrity profile as a subset
CIP User Authentication Profile	Provides Authentication at a user level for CIP communications. This is used as a basis for Authorization and Role Based Access Control.
Resource-Constrained CIP Security Profile	Provides a lightweight version of the protections afforded by the first two CIP Security profiles specifically for highly resource-constrained devices.
Pull Model Profile	Secure automatic device replacement/commissioning workflows using EST and DNS-SD technologies.
Device-Based Firewall Profile	Simple traffic filter (firewall) similar to the Linux IP Tables.

Each of the Security Profiles shown in Table 4 is targeted at providing security properties to mitigate the threats described previously as follows:

Table 5: Supported Security Properties

Security Properties	EtherNet/IP Confidentiality Profile	CIP User Authentication Profile	Resource-Constrained CIP Security Profile	Pull Model Profile	Device-Based Firewall Profile
Device Authentication	X		X		X
Trust Domain	Broad – group of devices	Narrow – Users/Roles	Broad; option to be narrow via gateway or proxy		
Device Identity	X	X (Identity of User)	X (via PSK)	X	
Data Integrity	X		X		
Data Confidentiality	X		X		
User Authentication		X	Via gateway or proxy		
Change Detection (Audit)					
Policy Enforcement (Authorization)		Fixed	Via gateway or proxy		

The development of the various CIP Security Profiles follows a number of key guidelines:

- The EtherNet/IP Security Profiles provide a secure transport mechanism for EtherNet/IP, with relatively little change to the CIP application layer.
- The CIP User Authentication and Resource-Constrained CIP Security Profile enhance CIP to provide additional security properties such as user authentication, and potentially extending CIP Security to support other non-EtherNet/IP networks.

CIP Security mechanisms in general should have the following attributes:

- Utilize proven-in-use, open security standards wherever possible
- Provide security options and/or scalable properties compatible with different risk profiles and device capabilities (e.g., apply encryption for confidentiality if required)
- Maximize compatibility with existing network infrastructure (switches, routers, firewalls, etc.)
- Require no custom cryptography to maximize security and minimize any possible import and export restrictions
- Implementations should be available as both commercial and open-source supporting many different OS platforms (embedded, PC, Linux, etc.) where possible
- Devices that support CIP Security must still be able to interoperate with devices that do not support CIP Security, on the same network. It should be a matter of end user configuration to allow or disallow such a mix of devices on the network. When mixing devices with secure and non-secure communications, it is the end user's responsibility to manage the device and network configuration appropriately. The user may need to provide additional controls such as firewalls or physical security means.
- Implementations should be compatible with other IP based security protocols such as IPSec or SSL-based VPN CIP Security should be capable of running over VPN connections to address remote access applications.

Figure 2 shows the relationship between the existing network protocols with no security (CIP, EtherNet/IP and IP) and those that support security enhancement delivered as part of CIP Security:

Figure 2: Security and Standard Network Relationship

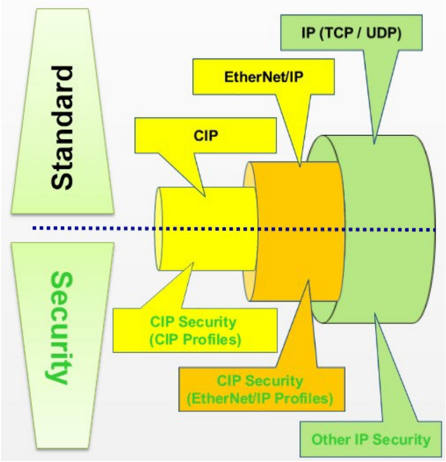

As Figure 2 illustrates, the mechanisms defined for the CIP Profiles build upon the EtherNet/IP Profiles, and make use of the secure transport for EtherNet/IP traffic.

Security Technologies

CIP Security makes extensive use of proven-in-use open security technologies such as:

- X.509v3 Digital Certificates used to provide cryptographically secure identities to users and devices
- Pre-Shared Key (PSK) for both client and server functionality. This is best used with simpler installations and simpler devices.
- TLS (Transport Layer Security) and DTLS (Datagram Transport Layer Security) cryptographic protocols used to provide secure transport of EtherNet/IP traffic
- Hashes or HMAC (keyed-Hash Message Authentication Code) as a cryptographic method of providing data integrity and message authentication to EtherNet/IP traffic
- AES symmetric encryption algorithm designed to be efficient with both hardware and software
- Encryption, that is continually updated based on the latest standards, as a means of encoding messages or information in such a way as to prevent reading or viewing of EtherNet/IP data by unauthorized parties
- OAuth 2.0 and OpenID Connect for cryptographically protected token-based user authentication and JSON Web Tokens (JWT) as proof of authentication, usernames and passwords

Secure EtherNet/IP transport provides the following security attributes:

- Authentication of the endpoints — ensuring that the target and originator are both trusted entities. End point authentication is accomplished using X.509 certificates or pre-shared keys.
- Message integrity and authentication — ensuring that the message was sent by the trusted endpoint and was not modified in transit. Message integrity and authentication is accomplished via TLS message authentication code (HMAC).
- Message encryption — optional capability to encrypt the communications, provided by the encryption algorithm that is negotiated via the TLS handshake.

Guide to the Specifications

CIP Security specifies security-related requirements and capabilities for CIP devices and includes material that is CIP network-specific (e.g., EtherNet/IP) in addition to material that is network-independent.

In its present form, the specifications for CIP Security include the following material:

- Chapter 1: Introduction to CIP Security
  The introduction duplicates information found in this technical overview.
- Chapter 2: CIP Security
  CIP security requirements and behaviors that are independent of the particular CIP network. Currently empty, this chapter is expected to include information on CIP-level authentication and authorization.
- Chapter 3: EtherNet/IP Security
  Requirements and behavior specific to EtherNet/IP. Primary material is the mechanism for secure transport over EtherNet/IP using TLS and DTLS.
- Chapter 4: Commissioning and Configuration
  Requirements and behavior related to device security commissioning and configuration.
- Chapter 5: Object Library
  CIP Objects related to security.
- Chapter 6: Certificate Management
  Requirements and behavior related to X.509 certificate usage in devices.
- Chapter 7: EDS Files
  EDS file content specific to security capabilities.
- Chapter 8: Security Profiles
  Explicit definition of requirements and recommendations that define each of the security profiles.
- Appendix: Threat Model
  Allows vendors and users to better understand how CIP Security can either prevent or mitigate the effects of cyber-attacks against the system.

References

[1] https://www.owasp.org/index.php/Application_Threat_Modeling