Threat Modeling CIP Security™

David Smith
Cybersecurity Architect
Schneider Electric

Jack Visoky
Principal Engineer and Security Architect
Rockwell Automation

Joakim Wiberg
Head of Technology
HMS Networks

Presented at the ODVA
2022 Industry Conference & 21st Annual Meeting
March 9, 2022
San Diego, California, USA

Abstract

CIP Security™ brings a number of important cybersecurity protections to CIP™ and EtherNet/IP™ communication. However, CIP Security is not meant to defend against all possible threats, but rather stands as a part of a Defense-in-Depth approach to cybersecurity of industrial equipment. It is important for vendors and users to understand what types of protections CIP Security provides, as well as limitations of those protections and areas where other technologies might be able to boost overall defense. This paper provides a sample of some of the interesting and impactful threats where CIP Security provides protection, as well as areas where CIP Security is meant to fit into a layered approach to cyber protection. This paper is not meant to be a full Threat Model of CIP Security, but rather provides some illustrative examples around Threat Modeling and the Defense-in-Depth approach to security in which CIP Security plays a major role in protecting important plant assets.

Keywords

Cybersecurity, CIP Security, Industrial Ethernet, Threat Model, TLS, DTLS, OpenID Connect

Definition of terms

Term	Definition
ARP	Address Resolution Protocol: a routing protocol used to determine the path to a given IP address. Traditionally ARP doesn't not have any built in data assurances.
ARP Poison	A spoofing attack where IP packets are "tricked" into being sent to the wrong destination. This is a

	common technique used to launch a “person-in-the-middle” attack where packets are unknowingly routed to an attacker before being sent off to the intended destination, possibly in a modified form.
CA	Certificate Authority: A root and possibly intermediate certificate which can be used to sign new certificates. CAs are also responsible for signing revocation lists that denote certificates which are no longer trusted.
Compensating Controls	An outside mechanism that provides additional protection or information assurance properties for assets within a Threat Model. Examples include a given technology, a network configuration/topology, or a physical structure.
Defense in Depth	A security strategy where multiple security controls, or “layers” are applied to a system. The reasoning behind this approach is that although one control might be compromised, it will be substantially more difficult to compromise multiple protections. This is a common and well-accepted strategy in modern cybersecurity.
DHCP	Dynamic Host Control Protocol: a protocol that is mainly used to dynamically assign IP addresses.
DNS-SD	Domain Name Server Service Discovery: a protocol used to discover the location of services on the network.
DTLS	Datagram Transport Layer Security. TLS as applied to a datagram-based transport (e.g. UDP).
EST	Enrollment over Secure Transport. An IETF defined protocol for requesting a certificate for use in secure communication. This protocol is done over HTTPS and is used in the CIP Security Pull Model.
IDS	Intrusion Detection System: a product designed to detect cyber attacks. This can be run at network level or host level, or as a combination through a distributed system. Generally an IDS will create an alert when an attack is detected but not perform active measures while the attack is taking place.
IEEE 802.1X	A technology designed to authenticate an endpoint before it can generally communicate on a network. Prior to authentication only limited communication is permitted to allow for authentication.
IPS	Intrusion Protection System: Similar to an IDS but in this case actively prevents the attack from occurring. The mechanisms for attack prevention can vary and might include things like dynamically changing firewall configuration or modifying packet routing.
NTP	Network Time Protocol: a time synchronization protocol that is common in Internet and IT systems.

PKI	Public Key Infrastructure: A set of policies, tools, and keys that control the issuing, revocation, and management of secure identities. Generally the secure identities are implemented as digital certificates. A PKI usually includes one or more CAs and RAs.
PSK	Pre-Shared Key: A symmetric key used in a TLS or DTLS handshake for authentication and to derive a shared session key.
RA	Registration Authority: the policies and decision point regarding whether or not a certificate is to be granted. This might be automated through software or might involve the decision of a human, or some combination.
SAN	Subject Alternative Name: A field within a digital certificate that contains identifying information regarding the certificate owner. This might include an IP address, hostname, or other similar information.
TOFU	Trust On First Use: The idea that a device or system will trust the first client that attempts to configure it, but that client may narrow the trust by provisioning trust anchors or other security configuration.
Threat Model	The output of the Threat Modeling activity where protected resources are identified, trust boundaries determined, threats analyzed and mitigations identified. Threat Modeling is an iterative process in which the Threat Model is updated periodically when new information is discovered.
TLS	Transport Layer Security: A standard communication security technology that is commonly used in Internet and IT communication to protect data in transit.
Trust Boundary	A boundary identified where if data passes over, protections need to be applied to the data. A trust boundary is somewhat arbitrary in that it may be drawn at different places for various reasons, but it denotes a boundary over which some security controls are needed.

Introduction

CIP Security is a technology which provides robust cybersecurity protections for products and systems which use it. However, it is not enough to simply state this, users of this technology must be given an understanding of the specific protections provided by CIP Security and the threats for which it is meant to provide mitigation. Threat Modeling is a powerful technique used by security professionals to understand the threats present within a system as well as mitigations to those threats. This technique has been applied to CIP Security and will be published as an appendix to Volume 8 of the CIP Specification. In the creation of the Threat Model the well-known STRIDE technique was used. This technique instructs threats to be analyzed from the STRIDE acronym:

- Spoofing
- Tampering
- Repudiation
- Information disclosure
- Denial of service
- Elevation of privilege

The CIP Security Threat Model analyzes the threats on the system's protected resources via these six tenets and then analyzes mitigations provided by CIP Security. For each threat analyzed a mitigation is also described, with some mitigations involving just CIP Security Technology and others requiring additional countermeasures or protections.

A Threat Model relies on the idea of a "Trust Boundary". A trust boundary is a designated boundary over which information that crosses it needs additional protection (e.g. data encryption, data confidentiality, etc.). The designation of trust boundaries is somewhat arbitrary and depends on the goals of the system, but it is an important part of establishing how threats will be analyzed and mitigated. For each threat discussed with the CIP Security Threat Model, a trust boundary is denoted by a dashed red line.

Note that this paper does not contain all the Threat Modeling information in Volume 8, but rather a sample of it as well as some additional discussion about threats and mitigations. For a more detailed analysis of threats and mitigations provided by CIP Security please see Volume 8. The Threat Modeling done in Volume 8 is meant to be very generic and as such will not cover all the specific situations that might arise. It is meant to serve as guidance for a more specific and detailed Threat Model done by vendors and/or users.

CIP Security is meant to fit into a Defense-in-Depth architecture and as such is not expected to mitigate all threats on a system. Therefore, many threats and aspects of cybersecurity fall outside of the scope of CIP Security. Although not an exhaustive list, some of the areas that are outside of the scope are:

- Non-EtherNet/IP Communication
- Security of non-CIP endpoints (e.g. Certificate Authorities, NTP servers, DHCP servers, etc.)
- Network-based Denial of Service attacks (e.g. dropped packets, data storms, etc.). In general an attacker with local network access can drop packets or cause packet storms with sufficiently powerful hardware. A secure communication protocol at the transport or application layer will not protect against these IP-based attacks, as it does not prevent an attacker from access to that layer of the network stack.

Threats and Mitigations Sampling

The following section provides a sample of some of the threats and mitigations described within the CIP Security Threat Model. More details are available in Volume 8. These threats and mitigations were chosen because they serve as good examples of a given threat type or mitigation type. A brief discussion of how a particular threat and mitigation serves as an example is given after each description.

Threats Against Provisioning (EtherNet/IP Confidentiality Profile)

Background Info:

The EtherNet/IP Confidentiality Profile provides for two mechanisms for provisioning, the Push Model (where certificates and/or PSKs are "pushed" to the endpoint via EtherNet/IP) and the Pull Model (which allows certificates to be requested automatically via the EST protocol). For the Push Model a device simply waits to accept security configuration from any client which can connect to it. In the Pull Model a device will discover an EST server via DNS-SD and then request a certificate. However, the device has no information assurances for the DNS-SD exchange or the EST exchange.

Both of these mechanisms utilize a Trust On First Use (TOFU) mechanism (see RFC 7435 for a general discussion of TOFU). That is, a device in the Factory Default state will trust whatever configuration client is the first that connects to this. Note that vendors are free to further restrict this trust by vendor specific mechanisms, although the standard EtherNet/IP Confidentiality Profile provisioning is TOFU.

Being that the TOFU mechanism is utilized, there are no authenticity guarantees of the configuration software provided by the CIP Security protocol. This must be managed by securing the supply chain and/or by a vendor specific means.

Note however that in both the Push and Pull Model, the configuration software/EST server can possibly verify the device. If a device is shipped with a vendor-signed certificate, then the configuration client software or EST server can be pre-loaded with the root of trust for that vendor. The vendor-signed certificate is used for the initial TLS connection as the server certificate for the Push Model and as the client certificate in the Pull Model. This allows verification of device authenticity.

Trust Boundary Diagram(s)

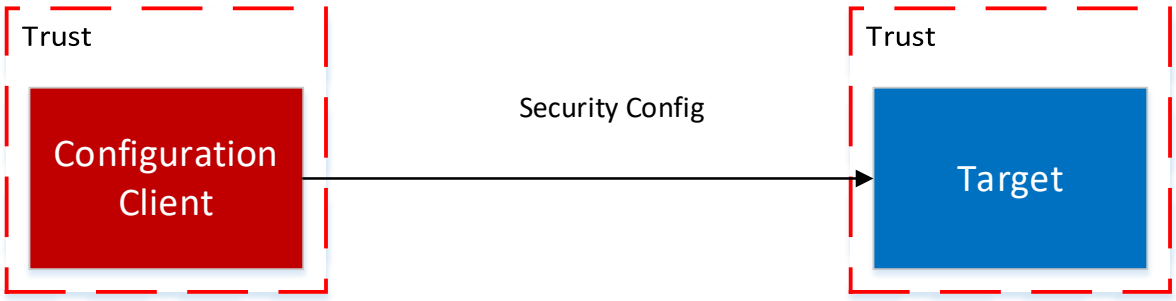

Figure 1: Trust boundary and data flow for CIP Security Push Model

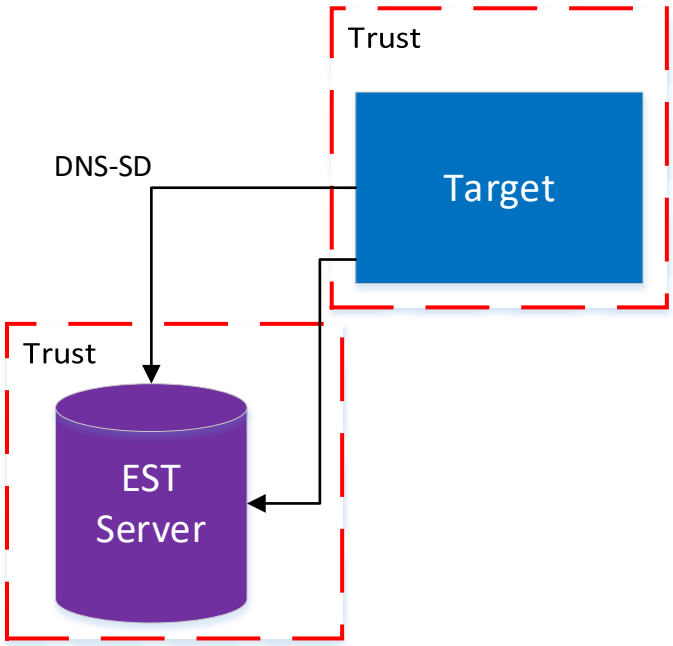

Figure 2: Trust boundary and data flow for CIP Security Pull Model

Threat – Spoofing

As mentioned, there is no trust pre-provisioned to the CIP Security endpoint before initial provisioning, as such there is no guarantee of authenticity for the configuration client. For the Pull Model, in general DNS-SD can be spoofed, as there are usually not any information assurances on the DNS communication.

Mitigation: This risk is accepted as a TOFU trust model is how CIP Security works by design. However, device authenticity can be provided by a vendor-signed certificate, therefore it is highly recommended for vendors to ship devices with a vendor certificate. Furthermore, vendors and users are free to include additional controls that go beyond a simple TOFU model if they deem this risk to be worth further mitigation. Examples of these types of controls could be pre-provisioning devices with roots of trust in manufacturing or using a compensating network control like 802.1X https://1.ieee802.org/security/802-1x/

Additional Discussion: This threat and mitigation is an example of a threat which in general is accepted due to industry and product requirements. Threats of this nature are not directly mitigated by CIP Security, but can work with other countermeasures to provide additional mitigations if desired by a user. In this particular example, something like 802.1X is provided as an additional countermeasure that can be deployed. However, this also serves as an example of a threat which a user needs to evaluate within their own unique environment to make an informed decision of whether or not additional countermeasures are required.

Threats Against Data in Transit (EtherNet/IP Confidentiality Profile)

Background Info:

Class 3 and Unconnected Messaging are used by EtherNet/IP endpoints for sending and receiving information in a structured, request-response manner. Without authentication of endpoints there is no guarantee that a connection is made with the correct Originator and Target. Furthermore, as the underlying transport for Class 3 and Unconnected Messaging is TCP/IP this messaging is subject to standard “Person-in-the-Middle” attacks. Similarly, Class 0 and 1 “implicit” messaging uses UDP and is subject to these same network level Person-in-the-Middle attacks.

Trust Boundary Diagram:

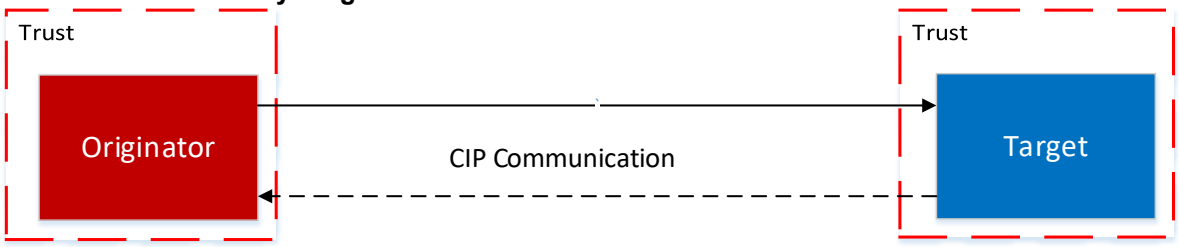

Figure 3: Trust boundary and data flow for general EtherNet/IP communication

Threat – Spoofing

A connection to a device from an unauthorized Originator represents a spoofing threat. Unauthorized Originators could affect configuration or I/O data by sending messages to the device. Like this, an Originator attempting to connect to a Target is susceptible to spoofing of the Target by an attacker.

Mitigation: Authenticators provided by (D)TLS such as certificates or PSKs provide for the authentication of both parties in the (D)TLS session. For certificates, CIP Security provides a configuration option via an attribute in the EtherNet/IP Security Object Instance that causes the Target to request and validate the client certificate. Mutual authentication during the (D)TLS handshake fully mitigates this risk. For systems in which the user determines that Originators do not need to be authenticated the option can be selected to only validate server certificates. Note that even with client and server authentication, there is no notion of Role-Based User Access

Control; for that the CIP Security User Authentication Profile is required. Also note that proper generation, storage, and protection of private keys is necessary for mitigation of this risk, this subject is specific to a given product and therefore outside the scope of the specification and Threat Model.

Threat – Tampering

Data may be tampered with through well-known or novel Person-in-the-Middle attacks such as ARP cache poisoning/ARP spoofing. This could result in a device receiving messages which are different from the intended message, and in some cases without the sender or receiver knowing of the change.

Mitigation: The information assurance properties of (D)TLS include data authenticity. This is realized by using an HMAC and/or an authenticated encryption algorithm. With either an HMAC and/or authenticated encryption in place through the (D)TLS cipher suite this risk is low and therefore generally mitigated.

Additional Discussion: This threat and mitigation is a good example of how CIP Security can provide a very strong mitigation. CIP Security, and the backing technologies of TLS and DTLS, were designed to specifically mitigate this type of threat, and therefore are well suited to this use. However, even with a case like this it is still important for a user to evaluate their unique system to ensure there aren't any extenuating circumstances that change their risk profile. However, most systems likely will have this type of threat sufficiently mitigated by CIP Security.

Threats Against Communication Redirection (EtherNet/IP Confidentiality Profile)

Background Info:

Even in a system in which CIP Security has been set up, an attacker may be able to affect packet routing through TCP/IP based attacks on network traffic (e.g. ARP spoofing). This could allow for an attacker to re-direct legitimate communication from the intended Target to a different one.

As an example, consider the case where a controller is sending data and commands to two drives, Drive A and Drive B, shown in figure 4. There is mutual trust between the controller and drives, and yet the commands sent to each drive are different. An attacker may attempt to re-direct traffic intended for Drive A to Drive B.

Trust Boundary Diagram:

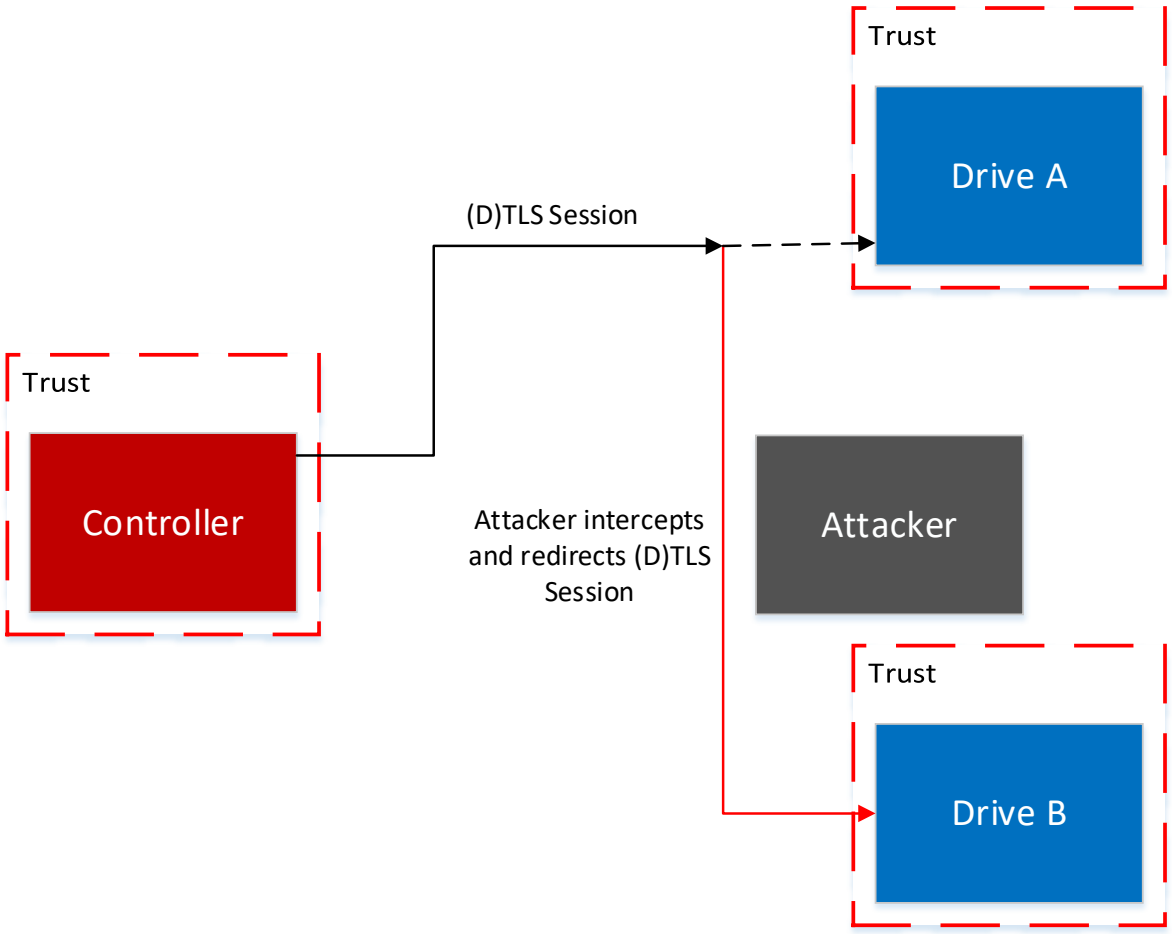

Figure 4: Trust boundary and data flow for redirection

Threat – Spoofing and Tampering

An attacker able to successfully re-direct traffic intended for one Target to another represents a spoofing threat in that the communication, although legitimate for the intended Target, can be considered spoofed on the new Target. Depending on the contents of the communication, tampering may also occur, as data could be modified in ways not intended by the legitimate user. Note that the attacker is not able to author any of the commands, but rather just redirect existing legitimate commands to an unintended Target. Furthermore there must be trust between the Originator and both Targets for this attack to be successful.

Note a special case of this attack would be a redirection of communication from the Originator back to the Originator. This was discussed extensively in “Selfie: reflections on TLS 1.3 with PSK” (https://eprint.iacr.org/2019/347).

Mitigation: Several mitigations exist for this vulnerability. Fundamentally this vulnerability is mitigated through the use of identifying information for the Target that can be trusted and verified by the Originator. One mechanism for this would be identifying information at the application layer (EtherNet/IP). Often times there may be route information, and/or product type and code information that would prevent this attack from occurring. However, this is not always the case, and as such a better mechanism for this would be to rely on identifying information within the cryptographic identity. Any identifying information within the certificate can be used, such as the Common Name or the Subject Alternative Name. CIP Security does provide mechanisms to set both of these, as well as to specifically check the Subject Alternative Name matches what is

expected. If the Originator verifies this information as part of the (D)TLS handshake then this vulnerability is mitigated. However, PSKs do not have any such identifying information, and as such, any usage of a PSK beyond two parties may be subject to this type of redirection attack. Note the PSK usage field prevents the special case of this attacker where an Originator's communication is reflected back to itself, as PSKs are only allowed to be used for Target or Originator functionality, but not both.

Additional Discussion: This is an example of a somewhat nuanced threat and why details of the configuration are important. Simply using CIP Security is not enough to provide a mitigation to this class of attacks, but rather setting specific configuration options is necessary. This drives the point that in Threat Modeling and mitigation analysis details are often very important.

Threats Against Proof of Authentication (User Authentication Profile)

Background:

After authentication occurs the Originator receives a signed Token that serves as proof of the authentication event. The Token provides proof to the Target that the authentication has occurred, as well as claims regarding the role/identity of the Originator. Several threats exist on this proof of authentication via the Token.

Trust Boundary Diagram:

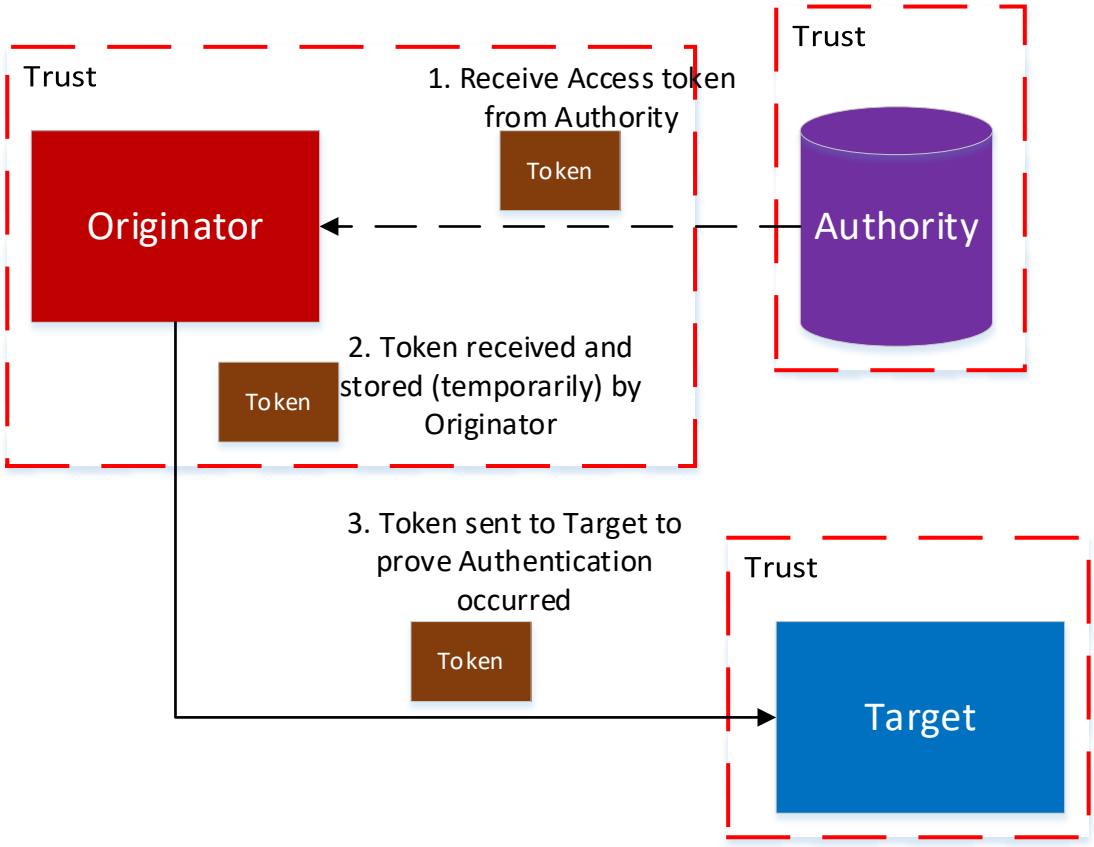

Figure 5: Trust boundary and data flow for proof of authentication via token

Threat – Spoofing

There are several spoofing threats for this data flow:

- An attacker might spoof the Originator to obtain the Token (interaction #1 in figure x)
- An attacker might spoof the Target to obtain the token (interaction #3 in figure x)
- An attacker might spoof the Token itself

Mitigation: (D)TLS provides endpoint authentication of both Originator and Target when Verify Client Certificate is enabled, providing assurances against the spoofing of the Originator or Target. Tokens are produced by the Authority with a unique digital signature generated by a private key the Authority controls, which provides mitigation against an attacker spoofing the token itself.

Threat – Tampering/Elevation of Privilege

Tampering and Elevation of Privilege threats are closely related. Tampering with the token can lead to a change of role or other claim information, granting the attacker a higher privilege. Similarly, a token could be replayed after expiration in order to obtain a level of privilege that is no longer granted

Mitigation: The digital signature of the token provides information assurance guarantees against tampering of the token.

Threat – Information Disclosure

An attacker might attempt to capture the token in order to impersonate the user; the Token is considered confidential as it can be used for impersonation. This might be done via standard network-based attacks, or via the spoofing threats discussed earlier in this section where the attacker spoofs either the Originator, the Target, or the Token itself.

Mitigation: While in transit over EtherNet/IP the confidentiality of the token is protected via a (D)TLS session with a confidentiality-based cipher suite. Note: the design of an Originator must ensure that the Token is not exposed outside of its trust boundary, although the internal structure of an Originator is outside the scope of this Threat Model and is vendor specific.

Additional Discussion: This threat is an example of a complex data flow in which a protected resource (the Token) is handled by various endpoints. Due to the complexity of this flow there are various threats, although mitigations are provided for each.

Threats Against Originator to Authority Authentication (User Authentication Profile)

Background

The Originator will send authenticators to the authority in order to prove its identity. Authenticators can include confidential information such as passwords. Note for external, non-CIP authorities this communication is not within the scope of the Threat Model. However, two CIP-based authenticators are supported: username/password and X.509 certificates. The communication of these CIP-based authenticators is within scope for this Threat Model and is subject to enumerated threats. Note that threats in this section all have the same mitigation, so only one mitigation is discussed.

Trust Boundary Diagram

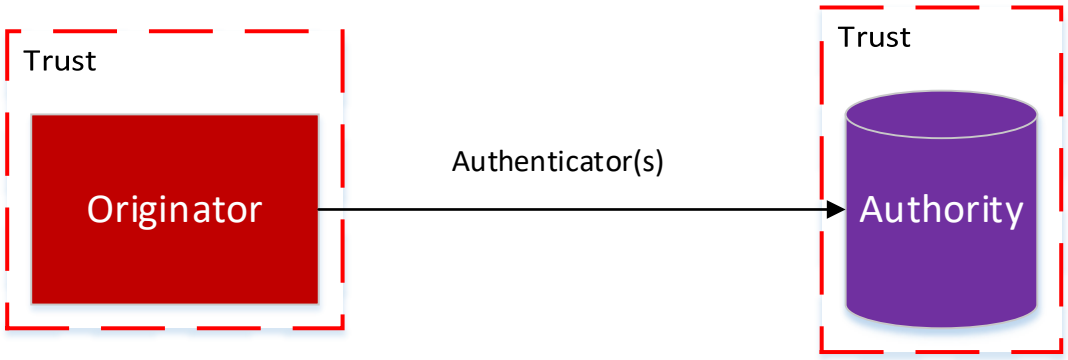

Figure 6: Trust boundary and data flow for authenticator exchange.

Threat – Information Disclosure

An attacker might capture confidential authenticators (e.g. passwords). This could be done either through spoofing the authority (described in C-3.4.2.1) or through passive network attacks where packets containing the authenticators are captured.

Mitigation A (D)TLS session between the Originator and Authority provides data authenticity and endpoint authentication. However, these assurances are only provided if bi-directional authentication is enabled (via the VerifyClientCertificate attribute of the EtherNet/IP Security Object) and if a cipher suite that utilizes confidentiality is used. Further note that non-CIP based Authorities may have other mitigations besides (D)TLS; in that case those Authorities must be evaluated against these threats.

Additional Discussion: This threat and mitigation is an example of a situation in which information assurances of one profile are used to protect resources of another profile. In this case the confidentiality assurances provided by TLS and DTLS from the EtherNet/IP Confidentiality Profile are used in mitigating risks against authenticators like passwords being exposed. This is one of the reasons which the CIP Security User Authentication Profile requires the EtherNet/IP Confidentiality Profile be supported.

Threats Against Discovery (User Authentication Profile)

Background

Before any User Authentication can occur, the Originator needs to discover the Target and the Authority. Threats exist against this process of discovery.

Trust Boundary Diagram

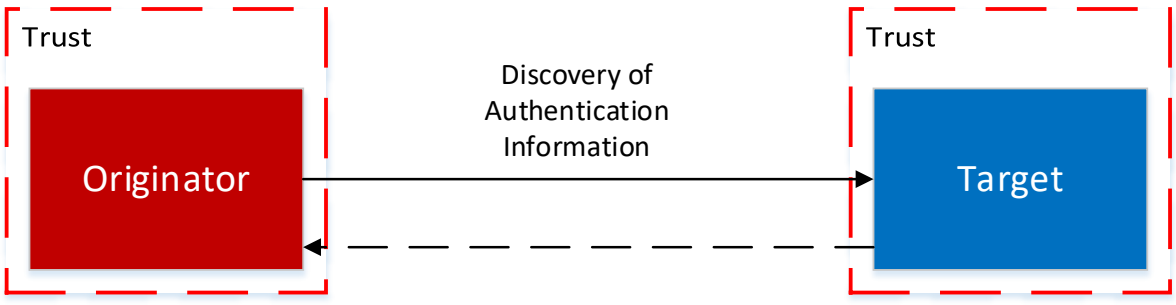

Figure 7: Trust boundary and data flow for discovery

Threat – Spoofing

An attacker might attempt to spoof the Target and send a malicious response to the Originator containing discovery information. This could in turn direct the Originator to a rogue authority, possibly leading to the leakage of authenticators.

Threat – Tampering

An attacker might attempt to tamper with the discovery information to direct the Originator to a rogue authority, possibly leading to the leakage of authenticators.

Mitigation: In the case of both the spoofing and tampering threat the mitigation is provided by the (D)TLS session over which the discovery is done. A (D)TLS session between the Originator and Target provides authentication of both the Originator and Target, if the VerifyClientCertificate option is set to true. The (D)TLS session also provides information assurance as to the authenticity of the data in transit, which in this case is the discovery information. Through the (D)TLS session a mitigation is provided against the tampering threat.

Additional Discussion: This threat and mitigation provides another example of information assurances from one profile being used to protect data in another. However, in this case the information assurance is not around data confidentiality, but rather the authenticity of data and authenticity of the endpoint. Again this is provided by the TLS session that is part of the EtherNet/IP Confidentiality Profile.

Best Practices

Threat Modeling CIP Security shows that CIP Security is not intended to provide a mitigation to all threats, but rather to fit into a defense-in-depth system. This section describes some of the supporting technologies for CIP Security, as well as some of the other defense-in-depth protections that can be applied to a CIP Security system. This section simply contains examples, there are other technologies that could also be used which are not listed.

Supporting Technologies

Public Key Infrastructure (PKI)

A PKI is a fundamental part of a system utilizing CIP Security. Although CIP Security can be used without a PKI (with PSKs or self-signed certificates), it is highly recommended to make use of a PKI for any but the simplest systems. A PKI allows for unique identity certificates to be issued and revoked, as well as for trust to be managed across the entire system. CIP Security endpoints can be configured for trusting multiple Certificate Authorities which can be managed by the PKI. Policies around how a certificate is granted are the domain of a Registration Authority (RA) within the PKI and are very important to constructing a secure PKI. The RA function is outside the scope of CIP Security, but directly impacts the security of an endpoint using CIP Security. Therefore, it is important to analyze and review the RA policies to ensure that proper authorization is required for a certificate to be granted, renewed, or revoked. There are many options for using commercially available software to set up a PKI, and for many users this will be a good option. IT departments often have a PKI already in place that could be used. However, in some cases the OT system will want to utilize a separate PKI that can be used to distribute certificates and trust anchors to CIP Security devices independent of IT trust. CIP Security can work with multiple CAs and multiple roots of trust, although a compromise of one or more CAs will likely have serious consequences to the effectiveness of CIP Security. Therefore it is very important to carefully consider the security of the PKI and CAs contained within.

OpenID Connect Identity Provider

For systems which use centralized authentication, an OpenID Connect Identity Provide is the technology chosen to work with CIP Security endpoints. There are many commercially available and open source OpenID Connect Identity Providers (see https://openid.net/developers/certified/ for examples). These Identity Providers issue the tokens which serve as proof of authentication, therefore it is very important that they are configured, used, and protected properly. Many OpenID Connect Identity Providers are available as a service, with the Identity Provider running in a cloud environment accessible over a secure Internet connection. In this case some of the protections are managed by the service provider, although it is still important for users to understand what types of protections are provided and how the Identity Providers are intended to be used. For an OpenID Connect Identity Provider running on-premise, more of the burden around configuration and protection will fall to the end user. Each particular environment will have nuanced needs which must be evaluated by the end user.

OpenID Connect Identity Providers generally support a wide range of authentication mechanisms. Many support various multi-factor authentication schemes, which may include the user of biometrics, smartcards, secure dongles, etc. It is important to understand the tradeoffs provided by various authentication schemes in terms of information assurance, ease-of-use, cost-to-deploy/maintain, etc. Systems with more advanced information assurance needs will likely want to use multi-factor authentication, although the exact details behind which scheme are important to work out through a Threat Model.

Given that the OpenID Connect Identity Provider issues the tokens which serve as proof of authentication in CIP Security it is a fundamental part of the security system. A compromise of the Identity Provider would very likely lead to a significant elevation of privilege, as tokens may be issued for an elevated role (e.g. Administrator). It's very important to ensure that tokens are only issued to properly authenticated parties, and that authentication is set up in such a way that it provides appropriate assurances of identity without causing undue burden to the users.

Secure Time Server (e.g. NTS)

System security depends on a synchronized time between endpoints in the system. Some components of the system are time sensitive, as such server and client must agree on time to ensure certificates are used within their window of validity, tokens have not expired, and events logged can be correlated. By default, expired tokens and certificates are rejected.

There are two well know abuse cases regarding time. The first is to manipulate time backwards to re-use previously expired credentials or certificates. This is a rare case as the attacker needs to compromise private keys in addition to manipulating time. The second abuse case is where the attack manipulates time to be outside of the validity window of all certificates and tokens causing a denial of service when they are rejected. Typically, time was manipulated by spoofing response messages to the NTP client with invalid time. As there was no method to authenticate the timestamp, the NTP client would adjust time towards the malicious time setting.

These abuse cases have been known for years, and NTPv3 included Message Authentication Code extensions to allow a client to authenticate the timestamp with a symmetric key. The key distribution mechanism was manual out of band management that posed non-trivial key management problems for implementors. NTPv4 attempted to solve the key distribution problem with the autokey protocol. Autokey included a key distribution and timestamp signing mechanism. However, autokey suffered from design flaws and was depreciated.

In 2020, the IETF released NTS4NTP (RFC 8915) to provide for uni-cast and multi-cast NTP environment an initial key agreement mechanism, and continuous update of a nonce used in MAC extension key. NTS4NTP allows automatic key distribution using a TLS channel and protocol handshake to give the NTP client the symmetric key and initial nonces used by the server. The client

then uses NTP with NTS extensions in the client to request a secured timestamp. The server responds with a timestamp and Authentication extension. The MAC allows the client to verify the authenticity of the timestamp to ensure the time received is accurate.

Mitigating Technologies

Firewall

A firewall is often one of the first tools that is brought to mind for cybersecurity protection. Firewalls range widely in their usage, features, and sophistication, but it is often true that a firewall can help provide additional protection for an industrial system. Firewalls often block certain types of traffic. They might block traffic based on the source or destination; for example preventing traffic from a class of IP addresses, or to a particular known-malicious domain name. More sophisticated rules might be applied dealing with certain types of traffic, for example blocking the insecure telnet protocol. More sophisticated tools can utilize techniques like deep packet inspection to apply more nuanced rules against certain types of traffic or network patterns.

Firewalls are often placed at network boundary locations as their function of blocking certain classes of traffic works well in the context of a network boundary. What constitutes a network boundary and what type of firewall to install there varies from system to system. They might be placed at the boundary between an internal network and the Internet, or between the IT network and OT network, or even between a cell or line within a plant and the rest of the plant network. Given the firewall's function at the network boundary, it typically does not interact directly with CIP Security devices and their communication, although in some cases it may. Since CIP Security is typically device to device, or computer to device, firewalls may not be blocking or inspecting CIP Security traffic. However, if they are, then it is of course important that the firewall be configured to allow CIP Security traffic. It is important to analyze the firewall configuration to ensure that it won't prevent legitimate CIP Security traffic from crossing a network boundary. Given that CIP Security uses well known ports and uses standard TLS and DTLS, writing rules for allowing CIP Security traffic should be quite achievable with most firewalls.

IDS/IPS

Another cybersecurity protection that can be applied is an Intrusion Detection System (IDS) or Intrusion Prevention System (IPS). IDS and IPS are often grouped together as they share many similar characteristics, with the main difference being whether or not the response to a threat is active and preventative as in the case of an IPS, or passive and more logging/alerting-based, as in the case of an IDS. Like a firewall, there is a wide range of sophistication of these tools; some are quite advanced and apply complex machine learning to determine if an attack is taking place, whereas others are quite straightforward in terms of simple packet matching for detection. One important consideration for deploying an IDS/IPS within a CIP Security system is that the CIP Security traffic may be encrypted, in which case an IDS/IPS that uses packet inspection will not be able to do any deep inspection of the CIP Security packets. That said, information can still be gleaned from the IP layer packets, such as source/destination, protocol, routing information, etc. However, if deep packet inspection is important to the security of the system then the user should consider using TLS and DTLS cipher suites which support authenticity only. Other than this consideration, IDS/IPS will likely work well with CIP Security devices, and can help bolster the defense-in-depth posture of the system.

Conclusion

Threat Modeling is an important activity for the security of any system involving protected resources and communication interfaces. This paper provides an overview of some of the Threat Modeling done for the CIP Security protocol within Volume 8 of the CIP Specification, as well as some additional technologies that are used by CIP Security or can be used to increase the defense-in-depth of a system using CIP Security. However, Threat Modeling and mitigation analysis is highly dependent on the particular details of a given system, therefore the information provided here is not meant to be a "one-size-fits-all" Threat

Model for systems and devices that use CIP Security. Rather, the intention is that this paper provides an introduction to the topic and activity of Threat Modeling with CIP Security, and serves as an aid to vendors and users who are using CIP Security and are creating a Threat Model for their system. Threat Models are not static but rather continuously updated as new information, including new attacks, become known. Therefore the information here is necessarily a snapshot in time and may need adjustments as time goes on. Note that the CIP Security Threat Model present in Volume 8 of the CIP Specification will be updated as conditions dictate. This paper and the information in Volume 8 show that CIP Security provides robust mitigation for a large class of cybersecurity threats, and that its usage is important in a system where CIP and EtherNet/IP communication are used.

References

1. STRIDE: https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=cs.20)?redirectedfrom=MSDN
2. Selfie: reflections on TLS 1.3 with PSK: https://eprint.iacr.org/2019/347
3. RFC 5246 (TLS 1.2) https://datatracker.ietf.org/doc/html/rfc5246
4. RFC 6347 (DTLS 1.2) https://datatracker.ietf.org/doc/html/rfc6347
5. OpenID Connect https://openid.net/developers/specs (version 1.0)
6. OpenID Connect Certificed Identity Providers https://openid.net/developers/certified/
7. RFC 8915 (NTS for NTP) https://datatracker.ietf.org/doc/html/rfc8915
8. IEEE 802.1X https://1.ieee802.org/security/802-1x/
9. RFC 7425 (Oppotrunistic Security: Some Protection Most of the Time) https://datatracker.ietf.org/doc/html/rfc7435
10. ODVA, Inc. The CIP Networks Library, Volume 8: CIP Security™, PUB00299