HALMSTAD UNIVERSITY

Securing Enclave
Lecture 6

Urban Bilstrup

Urban.Bilstrup@hh.se

Perimeter Defense

- Once an enclave is identified, it must be mapped to the network so that clear electronic perimeters can be defined.
- It is a necessary process that should be performed for any industrial network regardless of regulatory concerns, as an enclave can only be secured if there are defined and control entry points.
- In many cases the demarcation of the enclave will be very clear; for example, there may be a single network connection between a control center's supervisory LAN and the control system network.
- In some instances, multiple connections might exist; all network connections into or out of an enclave comprise that enclave's electronic perimeter.

Perimeter Defense

- In some instances a single enclave may consist of multiple, geographically or otherwise separated groups, the enclave is still considered to be a single enclave.
- If there are any network connections between separated locations, they should be held to the same controls as the rest of the enclaves.
- There should be no communications across those links that do not originate and terminate within the enclave, and **if outside communication is required, it must occur through defined and secure access points.**
- A common method **of interconnecting distributed enclaves** is the use of a dedicated **VPN or other encrypted gateway**, while for extremely critical enclaves, a dedicated network connection or fiber cable may be used so that **physical separation** is maintained.

Perimeter Defense

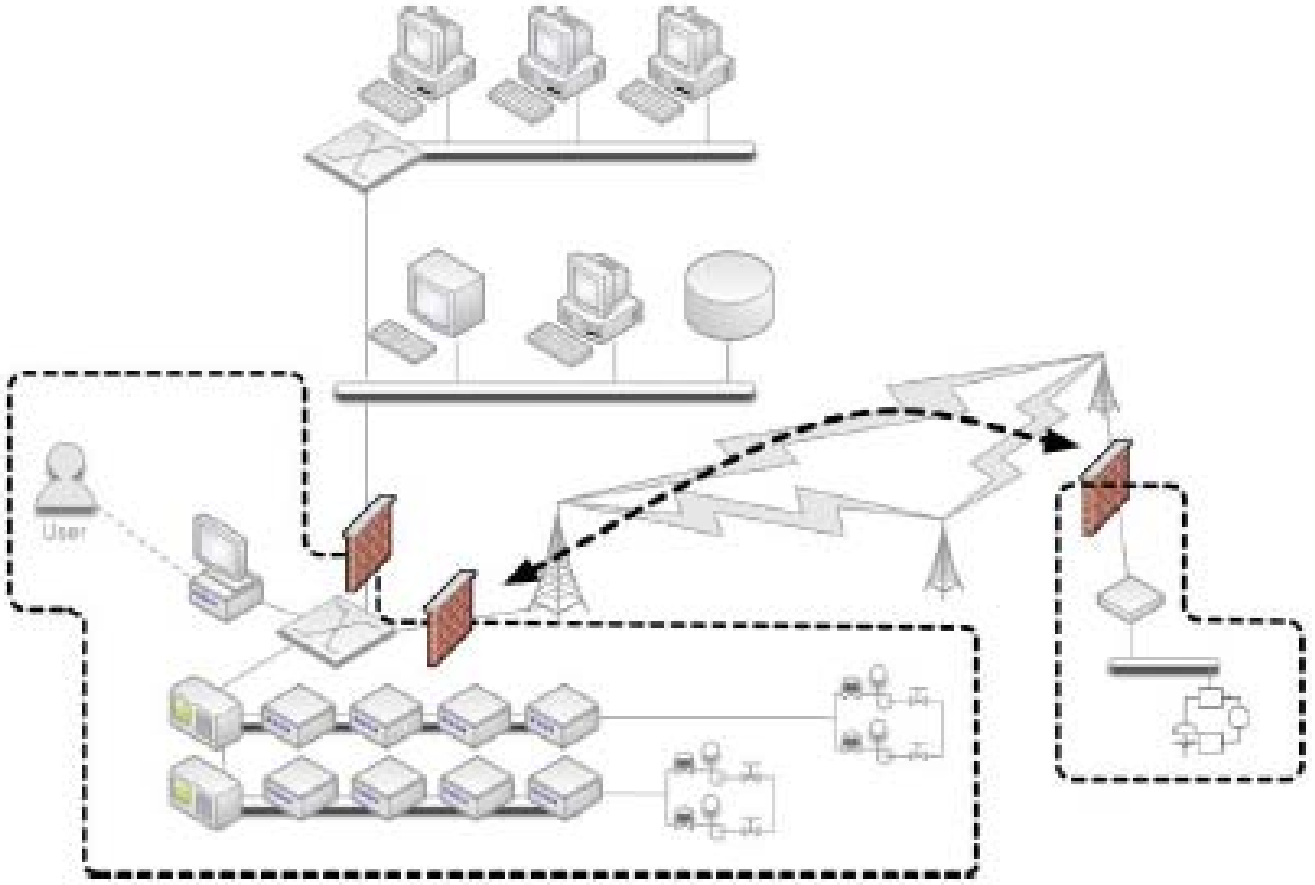

FIGURE 7.12

A Geographically Split Enclave.

Perimeter Defense

- The goal is that each enclave be isolated as strictly as possible, with as few connections as possible between that enclave and any other directly adjacent (or surrounding) enclave.
- By providing a single access point in and out of an enclave, that point can be secured using a perimeter security device such as a firewall or IPS.
- In the event of a single enclave that is split (geographically or by another enclave), inter-enclave communication can still be allowed: in this case through the use of perimeter firewalls, which effectively enforce a point-to-point route between the split enclaves (this path should also be encrypted).

Perimeter Defense

To establish an Electronic Security Perimeter (ESP) and effectively secure inbound and outbound traffic, two things must occur:

1. All inbound and outbound traffic must be forced through one or more known network connections that can be monitored and controlled.
2. One or more security devices must be placed in-line at each of these connections.

For each enclave, appropriate security devices should be selected and implemented using the recommendations below. **Typically, the criticality of the enclave dictates the degree of security that is required.**

Perimeter Defense

Table 7.1 Perimeter Security Requirements by Criticality

Criticality	Required Security	Recommended Enhancements
4 (highest)	NRC CFR 73.54: Unidirectional Perimeter, NERC CIP 005: Firewall or IDS or IPS	Application layer monitoring, Firewall, IDS and IPS
3	NRC CFR 73.54: Unidirectional Perimeter, NERC CIP 005: Firewall or IDS or IPS	Application layer monitoring, Firewall, IDS and IPS
2	NERC CIP 005: Firewall or IDS or IPS	Firewall and IDS and IPS
1	NERC CIP 005: Firewall or IDS or IPS	Firewall and IPS
0 (lowest)	NERC CIP 005: Firewall or IDS or IPS	Firewall and IPS

Perimeter Defense

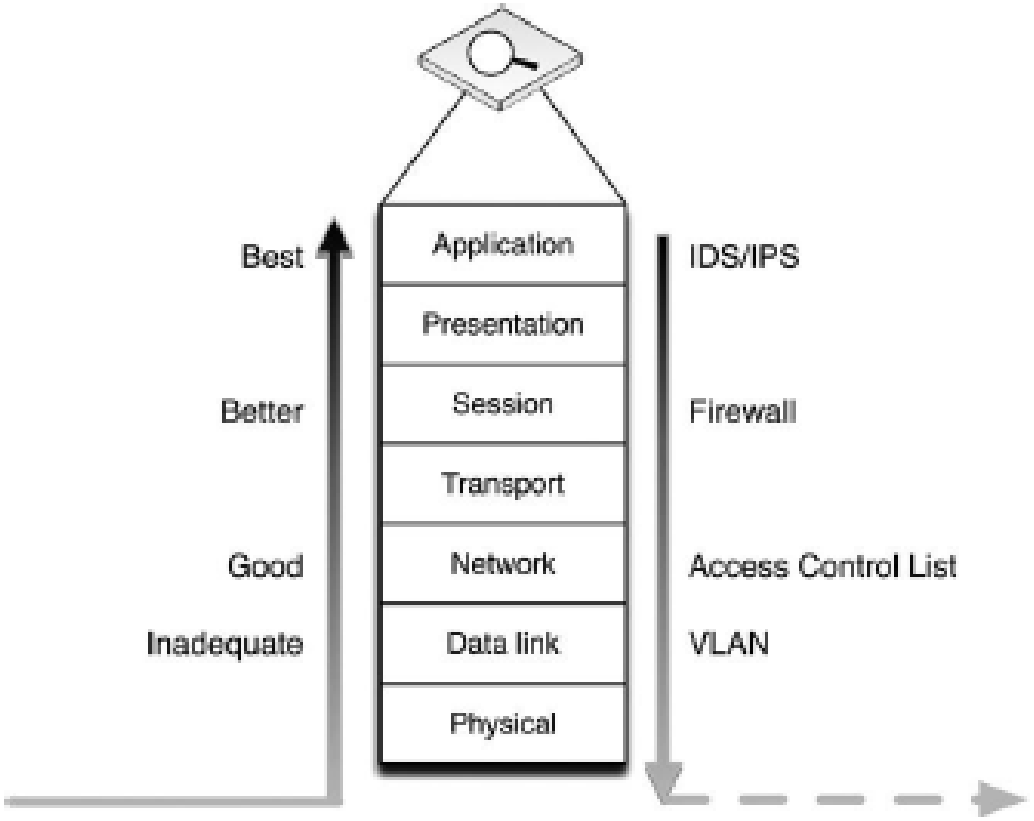

FIGURE 7.14

Relative Capabilities of Common Security Devices.

Perimeter Defense

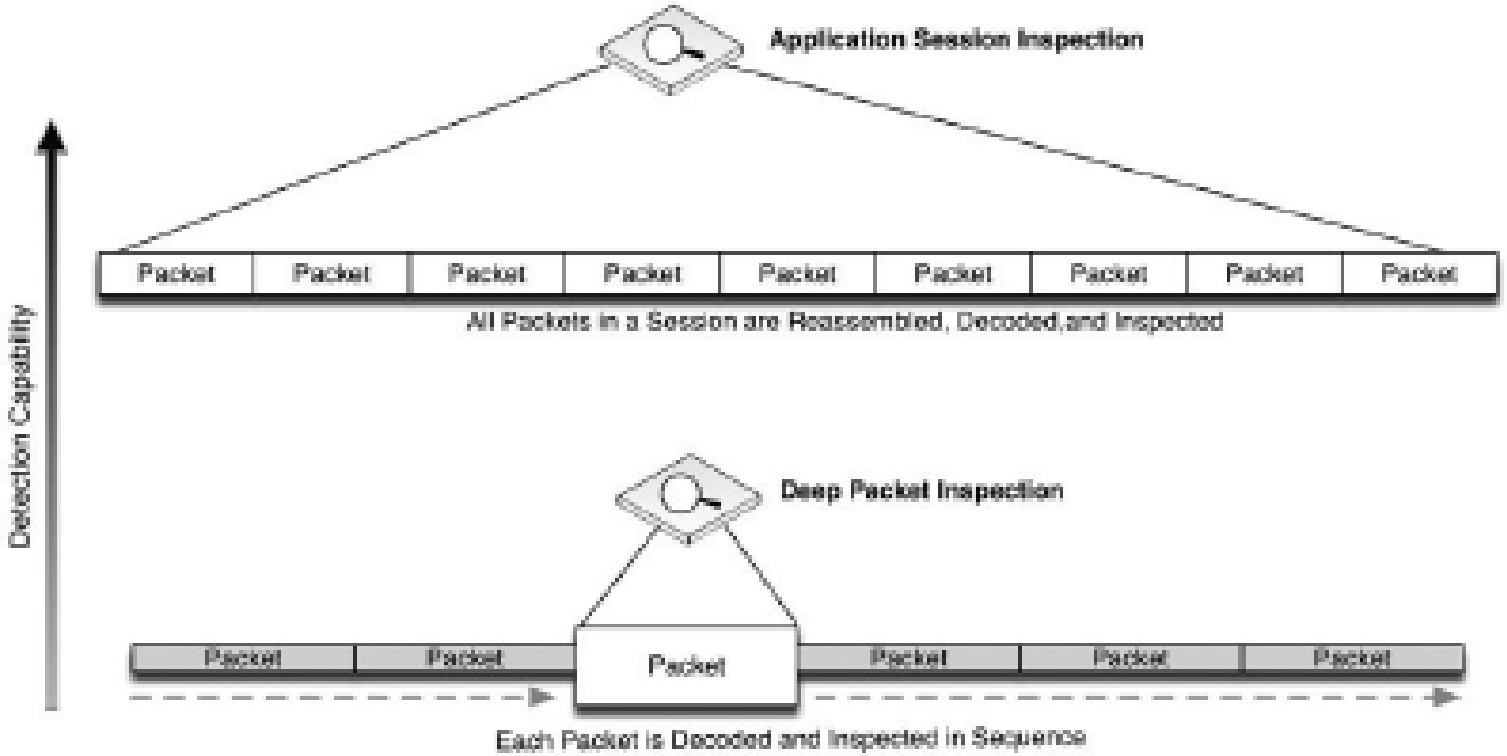

FIGURE 7.15

Application Session Inspection vs. Deep Packet Inspection.

Firewall Configuration

- Firewalls control communication using a defined configuration policy, typically consisting of **Accept** (allow) and **Drop** (deny) statements.
- Most firewalls will enforce a configuration in sequence, such that starting with a broadly defined policy, such as Deny All, which will drop all inbound traffic by default.
- Which then be overruled by subsequent, more focused rules.
- The following firewall policy would only allow a single IP address to communicate outside of the firewall on port 80 (HTTP).
  - Deny All
  - Allow 10.0.0.2 to Any Port 80

Firewall Configuration

- Determining what rules should be configured is typically easier in an industrial network because the nature of an industrial network is such that there is no need to accommodate the full diversity of applications and services typically found in an enterprise network.

- This is especially true when configuring a specific firewall against a specific enclave: the enclave will by its nature be limited in scope, resulting in concise firewall policies.

Firewall Configuration

- The method of properly configuring an enclave firewall is as follows:

1. Begin with bidirectional Deny All rules.

2. Configure specific exceptions, using the defined variables $ControlSystem_Enclave01_Devices and $ControlSystem_Enclave01_PortsServices.

3. Verify that all Allow rules are explicitly defined (i.e., no All rules).

- Guidelines National Infrastructure Security Coordinator Center (NISCC)

Firewall Configuration

Table 7.2 NISCC Firewall Configuration Guidelines with Enclave Variables[a]

NISCC Recommendations	Example Rule Using Enclave Variables	Notes
Start with universal exclusion as a default policy	`Deny All / Permit None`	Firewalls should explicitly deny all traffic inbound and outbound as the default policy.
Ports and services between the control system environment and an external network should be enabled and permissions granted on a specific case by case basis	`Allow 10.2.2.120 port 162 to 192.168.1.15 port 162` `#Allow SNMP traps from router ip 10.2.2.120 to network management station ip 192.168.1.15, authorized by John Doe on April 1 2005`	Comments used within the firewall configuration file can be used to document special cases, permissions, and other details.
All "permit" rules should be both IP address and TCP/UDP port specific, and stateful if appropriate, and shall restrict traffic to specific IP address or range of addresses	N/A	This guideline can be enforced by using `$ControlSystem_Enclave01_Devices` and `$ControlSystem_Enclave01_PortsServices` to define rules.
All traffic on the SCADA and DCS network(s) are typically based only on routable IP protocols, either TCP/IP or UDP/IP; thus, any non-IP protocol should be dropped	N/A	By using `$ControlSystem_Enclave01_PortsServices` within all defined rules, only protocols explicitly allowed within that enclave will be accepted by the firewall, and all others will be dropped by the overarching `Deny All` rule.
Prevent traffic from transiting directly from the Process Control / SCADA network to the enterprise network; all traffic should terminate in the DMZ	`Deny [Not $Neighboring Enclave1, Not $Neighboring Enclave2] to $ControlSystem_Enclave01_Devices` `Deny $ControlSystem_Enclave01_Devices to [Not $Neighboring Enclave1, Not $Neighboring Enclave2]`	By configuring a rule on each enclave that explicitly denies all traffic to and from any enclave that is NOT a neighboring enclave will prevent any transitive traffic. All traffic will need to be terminated and reestablished using a device local to that enclave.

Any protocol allowed between the DCS and the SCADA DMZ is explicitly NOT allowed between SCADA DMZ and enterprise networks (and vice versa)	At the demarcation between the enterprise network and SCADA DMZ:
	`Deny $ControlSystem_Enclave01_PortsServices to $EnterpriseNetwork_Enclave01_Devices`
	At the demarcation between the DCS and SCADA DMZ:
	`Deny $EnterpriseNetwork_Enclave01_PortsServices to $ControlSystem_Enclave01_ Devices`
		These rules enforce the concept of "disjointing" protocols, and further prevents transitive communication from occurring across an enclave.

Firewall Configuration

Allow outbound packets from the PCN or DMZ only if those packets have a correct source IP address assigned to the PCN or DMZ devices	N/A	Explicitly defined `Deny All` rules combined with explicitly defined known-good IP addresses using `$ControlSystem_Enclave01_Devices` ensures that all outbound packets are from a correct source IP. Firewalls may also be able to detect spoofed IP addresses. In addition, network activity monitoring using a Network Behavior Anomaly Detection (NBAD), Security Information and Event Management (SIEM), or Log Management solution may be able to detect instances of a known-good IP address originating from an unexpected device based on MAC Address or some other identifying factor (see Chapter 9, "Monitoring Enclaves")
Control network devices should not be allowed to access the Internet	At the Internet firewall: `Deny [$ControlSystem_Enclave01_ Devices, $ControlSystem_Enclave02_ Devices, $ControlSystem_Enclave03_ Devices, $ControlSystem_Enclave04_ Devices]`	Because all devices in all enclaves have been identified and mapped into variables, these devices can be explicitly denied at the Internet firewall.
Control system networks shall not be directly connected to the Internet, even if protected via a firewall	N/A	Using the enclave approach, no control system should be directly connected to the Internet (see "Establishing Enclaves").
All firewall management traffic be: 1. Either via a separate, secured management network (e.g., out of band) or over an encrypted network with two-factor authentication 2. Restricted by IP address to specific management stations	N/A	This recommendation supports the establishment of a Firewall Management enclave using the methods described earlier under "Establishing Enclaves." By placing all firewall management interfaces and management stations in an enclave, which is isolated from the rest of the network, the traffic can be kept separate and secured.

[a]National Infrastructure Security Coordination Center, NISCC Good Practice Guide on Firewall Deployment for SCADA and Process Control Networks. British Columbia Institute of Technology (BCIT). February 15, 2005.

Intrusion detection and prevention configuration

- IDS and IPS devices inspect network packets for signs of malicious code or exploits.
- Intrusion Detection refers to passive inspection.
- An IDS examines packets and compares them against a set of detection signatures, and issues an alert when there is a match.
- Intrusion Prevention refers to active inspection, where traffic is matched against IDS rules, but where specific actions can be taken in addition to alerting.

Intrusion detection and prevention configuration

- IDS actions can include Alert (generate a custom message and log the packet), Log (log the packet), and Pass (ignore the packet), while IPS actions can also include Drop (drop the packet and log it), Reject (drop the packet and initiate a TCP reset to kill the session), and sDrop (drop the packet, but do not log it).
- In addition, both IDS and IPS rules can use the Activate and Dynamic actions, the former of which activates another rule, and the latter of which remains idle until activated by an Activate rule.
-

Intrusion detection and prevention configuration

- Both IDS and IPS devices can be deployed either out-of-line using a network span or tap port or in-line using two network interfaces, although an IPS can only actively block traffic if it is deployed in-line
- An enabled collection of IDS/IPS detection signatures is referred to as an IDS/ IPS policy, and this policy will dictate what types of threats may be detected by the device, as well as the degree and scope of events that will be generated.
- While active blocking of malicious traffic is important, the IDS/IPS events that are generated can also be analyzed to provide other important indicators—including network behavior, larger threat incidents.

Intrusion detection and prevention configuration

- Signatures generally follow a format similar to a firewall rule, where there is an identified source and destination address and/or port, as well as an action.
- In addition, IDS/IPS signatures may match against specific contents of a packet, looking for patterns within the packet that indicate a known exploit (i.e., a “signature”). Common IDS/IPS signature syntax follows the de facto standards defined by Snort, an open-source IDS project owned by SourceFire.

Intrusion detection and prevention configuration

- An example signature is written as follows:

[Action] [Protocol] [Source Address] [Source Port] [Direction Indicator] [Destination Address] [Destination Port] [Rule Options]

- which when written in correct syntax looks like

drop tcp 10.2.2.1 80 -> 192.168.1.1 80 (flags: <optional snort flags>; msg: “<message text>”; content: <this is what the rule is looking for>; reference: <reference to external threat source>;)

Intrusion detection and prevention configuration

- To highlight the difference between a firewall rule and an IDS/IPS signature, consider the following example:

  drop tcp 10.2.2.1 80 -> any any

- Without any rule options, the previous rule is essentially the same as the firewall rule

  Deny 10.2.2.1port 80

- which would block all traffic originating from 10.2.2.1 on port 80, effectively preventing that user from accessing the web (via HTTP port 80).

Intrusion detection and prevention configuration

- However, the ability to match packet contents within the rule options enables an IDS/IPS device to control traffic at a much more granular level, such as:

drop tcp 10.2.2.1 80 -> any any (msg: “drop http POST”; content: “POST”;)

- This rule functions differently, only dropping traffic from the source address in question if the HTTP traffic contains a POST request (used by many web forms or applications attempting to upload a file to a web server over HTTP).

Intrusion detection and prevention configuration

- Determining the exact IDS/IPS policy to be enforced is the first step in correctly configuring the device.
- The enclave variables defined earlier under “Establishing Enclaves” are valuable tools that can be used to write succinct and highly relevant signatures.
- Unlike a firewall which starts with a simple Deny All rule, an IDS/IPS should be deployed “large”—with many active signatures—and then pruned back to the specific requirements of the enclave.

Intrusion detection and prevention configuration

A method of properly configuring an IDS/IPS is as follows:

1. Begin with a more robust signature set, with many active rules.

2. If a protocol or service is not allowed in the enclave, replace any specific detection signatures associated with that protocol or service with a broader rule that will block all traffic from that protocol or service (i.e., drop unauthorized ports and services).

3. If a protocol or service is allowed in the enclave, keep all detection signatures associated with that protocol or service active.

3a. For all active signatures, assess the appropriate action, using Table 7.3.

4. Keep all IDS signatures current and up to date.

Table 7.3 Determining Appropriate IDS/IPS Actions

Allowed Port or Service?	Source	Destination	Criticality of Service	Severity of Event	Recommended Action	Note
No	Any	Any	Any	Any	`Reject`	Any communication not explicitly allowed within the enclave should be `Rejected` to disrupt unauthorized sessions and deter an attack.
Yes	Inside Enclave	Inside Enclave	High	Any	`Alert`	Active blocking or rejection of traffic that originates and terminates within an enclave could impact operations. For example, a false positive could result in legitimate control system traffic being blocked or rejected.
Yes	Inside Enclave	Inside Enclave	Low	Any	`Alert` or `Pass`	For noncritical services, logging is recommended but not necessary (`Alert` actions will provide valuable event and packet information that could assist in later incident investigations).
Yes	Outside Enclave	Inside Enclave	High	Low (events from obfuscated detection signatures or informational events)	`Alert`	Many detection signatures are broad to detect a wider range of potential threat activity. These signatures should `Alert` only to prevent unintentional interruption of control system operations.

Yes	Outside Enclave	Inside Enclave	High	High (explicit malware or exploit detected by a precisely tuned signature)	`Block, Alert`	If inbound traffic to a critical system or asset contains known malicious payload, the traffic should be blocked to prevent outside cyber incidents or sabotage.
Yes	Inside Enclave	Outside Enclave (explicitly allowed destination address)	Any	Any	`Alert`	This traffic is most likely legitimate. However, alerting and logging the event will provide valuable event and packet information that could assist in later incident investigations.
Yes	Inside Enclave	Outside Enclave (unknown destination address)	Any	Any	`Block` or `Reset`	This traffic is most likely illegitimate. Generated alerts should be addressed quickly: if the event is a false positive, necessary traffic could be unintentionally blocked; if the event is a threat, it could indicate that the enclave has been breached.

Recommended IDS/IPS Rules

Basic recommendations for IDS/IPS configuration include active block rules to

1. Prevent any undefined traffic from crossing enclave boundaries (where the disruption of the communication will not impact the reliability of a legitimate service).

2. Prevent any defined traffic containing malware or exploitation code from crossing enclave boundaries.

3. Detect and log suspicious or abnormal activity within an enclave.

4. Log normal or legitimate activity within an enclave, which may be useful for compliance.

Recommended IDS/IPS Rules

The greater the extent of functional isolation and separation into defined enclaves, the more concise and effective the IDS/IPS policy will be.

Some basic IDS and IPS rules suitable for use in enclave perimeters include the following:

- Block any industrial network protocol packets that are the wrong size or length.
- Block any network traffic that is detected inbound to or outbound from any enclave where that is not expected or allowed.

Recommended IDS/IPS Rules

- Block any industrial network protocol packets that are detected in any enclave where that protocol is not expected or allowed.
- Alert any authentication attempts, in order to log both successful and failed logins.
- Alert any industrial network port scans.

Recommended IDS/IPS Rules

- Alert any industrial network protocol function codes of interest, such as:
  - “Write” functions, including codes that write files or that clear, erase, or reset diagnostic counters.
  - “System” functions, including codes that stop or restart a device.
  - “System” functions that disable alerting or alarming.
  - “Read” functions that request sensitive information.
  - “Alarm” or “Exception” codes and messages.