CHAPTER 3

CPwE Network Security Design Considerations

This chapter covers design considerations that must be considered by OT control system engineers and IT security architects when deploying CPwE Network Security solutions. These design considerations provide the rationale for choosing a particular option and are the basis for the deployment options described in Chapter 4, “Configuring the Infrastructure.”

- Traffic Flows in a Network
- Segmentation—High Level
- Segmentation Using Downloadable Access Control Lists (dACLs)
- Segmentation Using Layer 3 Access Control List
- Segmentation—TrustSec
- Enforcement Point
- Scalable Group Tag Exchange Protocol Considerations
- NetFlow Data Collection
- Stealthwatch Deployment Considerations
- Cisco ISE Deployment Considerations
- IPDT Considerations

Traffic Flows in a Network

Horizontal communication among peer-to-peer IACS devices in a network is called East-West communication. Figure 3-1 depicts East-West communication in a plant-wide architecture.

Figure 3-1 East-West Traffic Flow in Cell/Area Zone

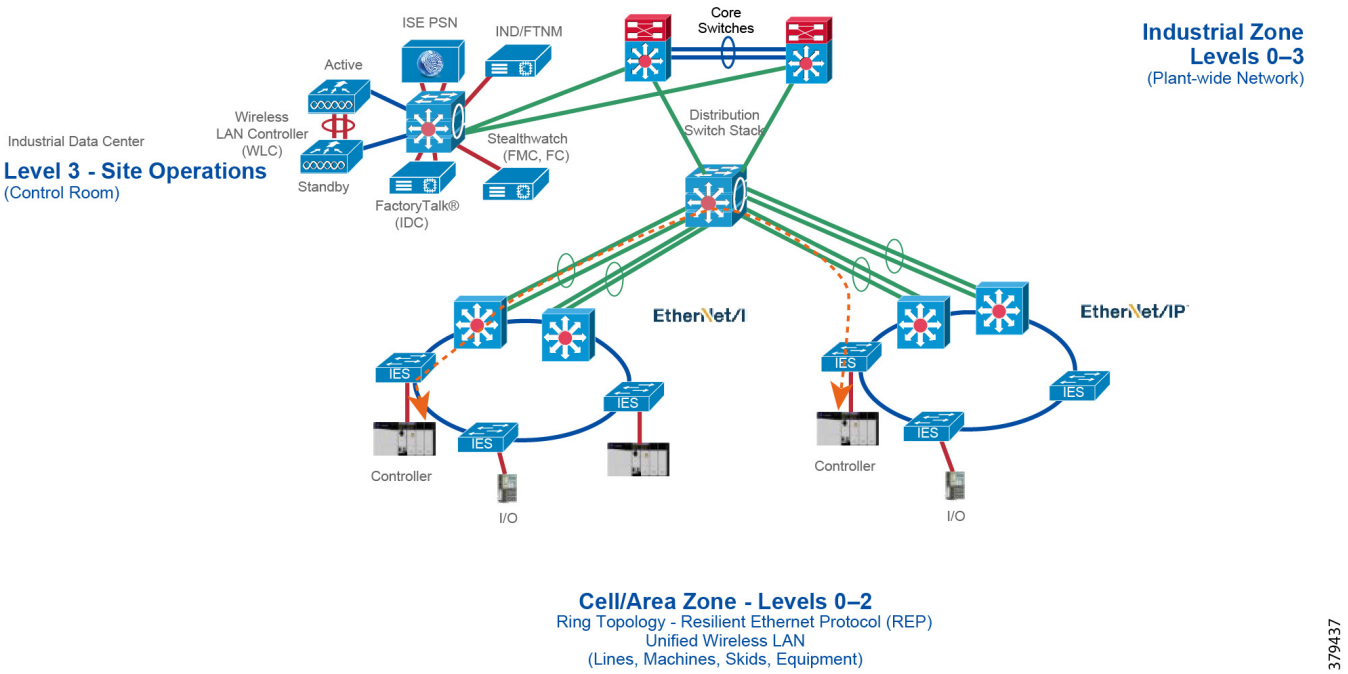

Allowing a server or any other device in Level-3 Site Operation, IDMZ, or Enterprise Zone to communicate with an IACS asset in the Cell/Area Zone is called North-South communication. In Figure 3-2, the Engineering Workstation (EWS) is accessing a controller in the Cell/Area Zone and this communication flow is defined as North-South communication.

Figure 3-2 North-South Communication in a Plant-wide Network

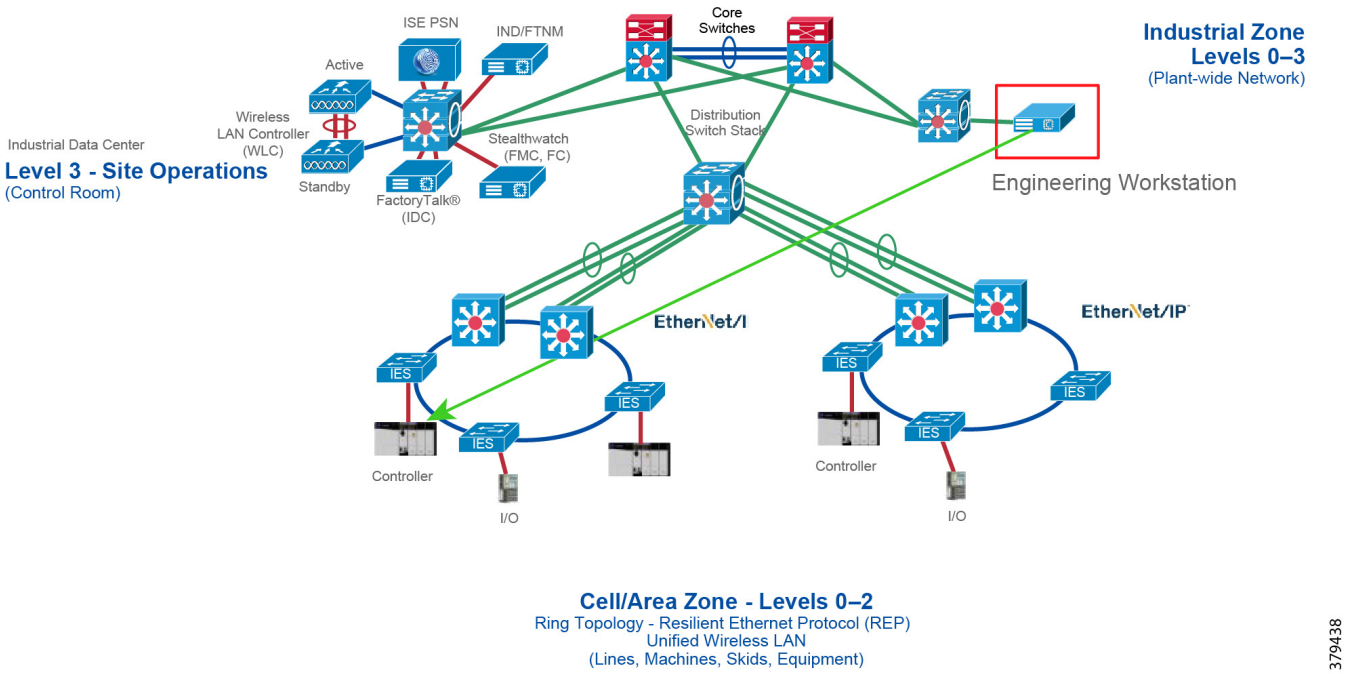

Segmentation—High Level

IT security architects in conjunction with a control system engineer should design an access policy that specifies the East-West and North-South communication flows that must be allowed in an IACS network. In an IACS network, having an open policy that allows every IACS asset to communicate with every IACS asset is convenient, but that approach increases the risk of cyber threat propagation. On the other hand, implementing a restrictive policy that does not allow any inter Cell/Area Zone communication is also counter-productive because certain IACS assets need to access other IACS assets that exist in different Cell/Area Zones. Since the exact requirements of a particular scenario are based on the current IACS application requirements, specifying a policy that would work for all the deployments is not possible. Hence in this CPwE Network Security CVD an access policy example is shown that can be customized for use in different environments.

Assumptions about the access policy for an IACS network:

- All the traffic within the Cell/Area Zone is implicitly permitted because it is assumed that a Cell/Area Zone is formed because a group of IACS assets need to communicate with each other, so no enforcement is applied to any IES in the Cell/Area Zone.
- All the traffic between any two different Cell/Area Zones will be enforced. As an example, in Figure 3-3 Controller_A in one Cell/Area Zone is allowed to access Controller_C in another Cell/Area Zone, but Controller_B is not allowed to access Controller_C.

Figure 3-3 Example of Enforcement in East-West Traffic Flow

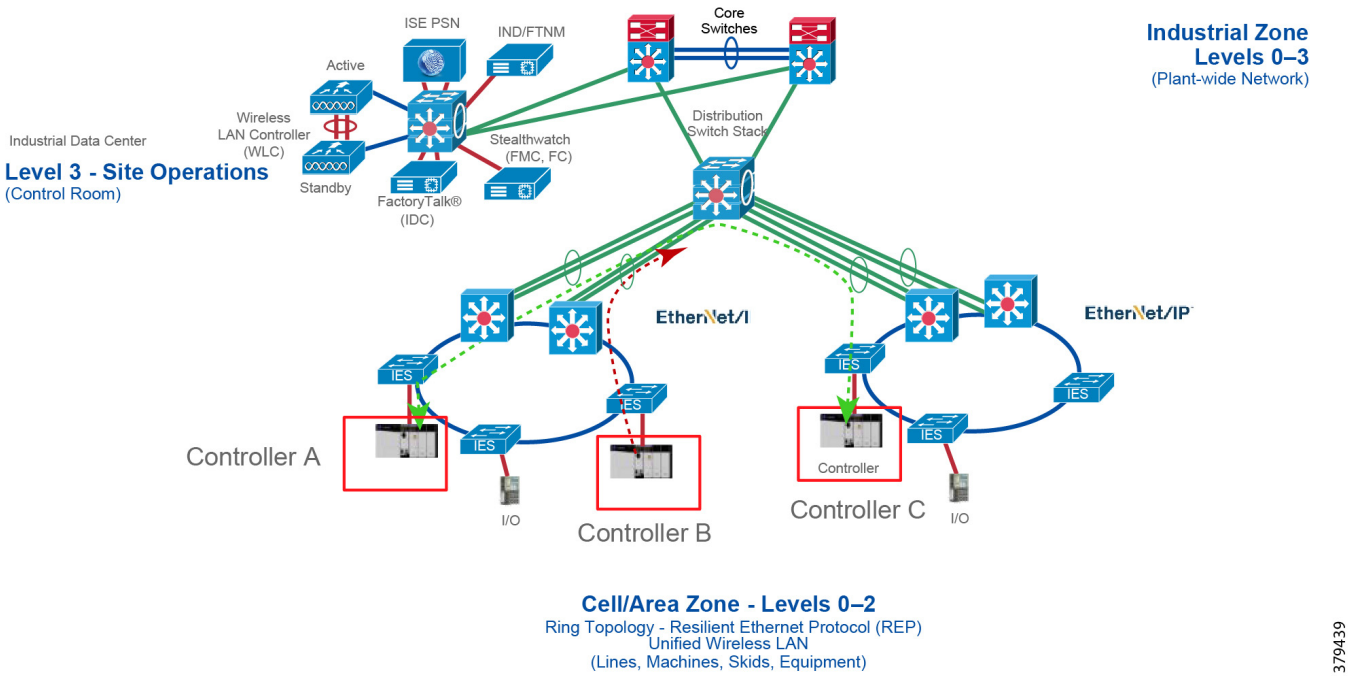

Segmentation Using Layer 3 Access Control List

When an IACS asset is not configured with MAC authentication bypass (MAB) and is unable to get a downloadable access control list (dACL) from ISE, use a static ACL on the distribution switch which is connecting different Cell/Area Zones. In Figure 3-4, ACL is applied on the distribution switch connecting the two Cell/Area Zones. In Figure 3-4, the ACL must allow communication between 10.20.25.10 and 10.20.8.10 so that Controller-A is able to establish communication with Controller-B.

Figure 3-4 Segmentation Using Layer 3 ACL

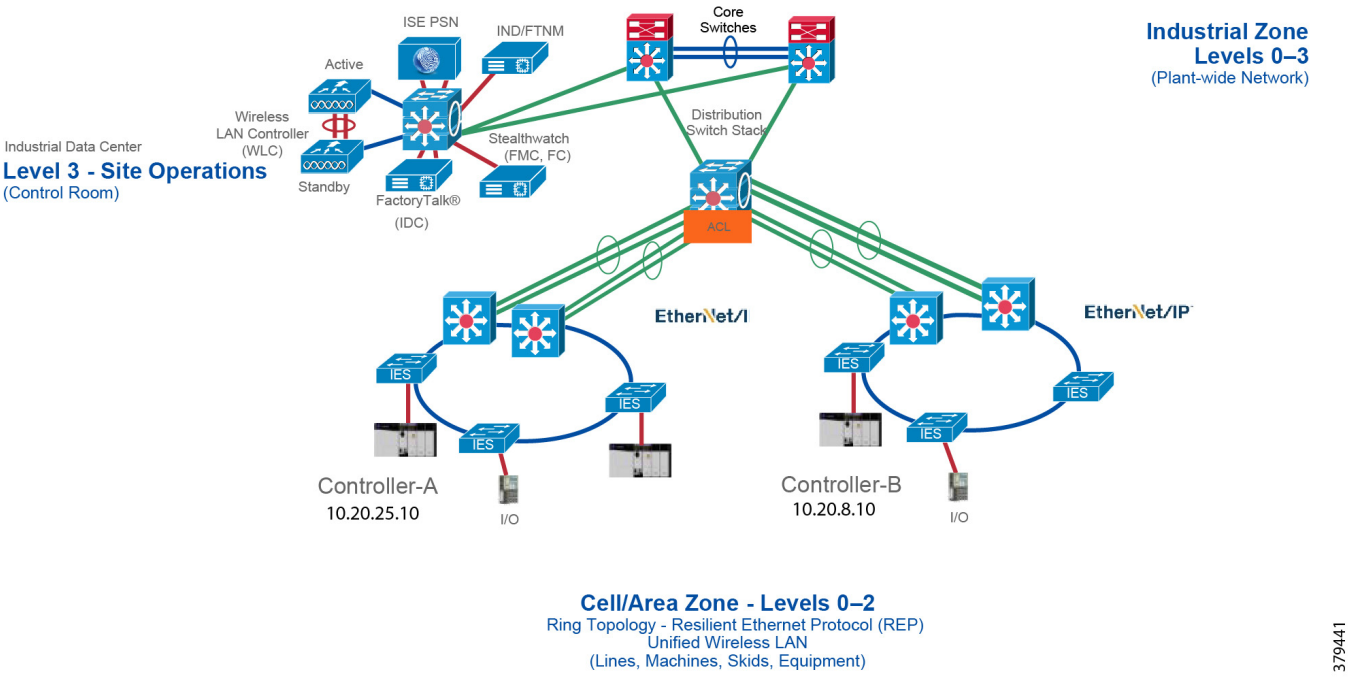

The above method has similar disadvantages in managing the ACL as the dACL. Whenever the controller IP address changes or is moved to a different location, then the ACL needs to be updated. The old entries need to be purged and the new entries added. This process can be burdensome and may lead to an IT security architect making mistakes.

Segmentation Using Downloadable Access Control Lists (dACLs)

As explained in Chapter 1, "CPwE Network Security Overview", Segmentation, segmentation is the practice of zoning the IACS network to create smaller domains of trust to help protect the IACS network from known and unknown risks in the network. This section describes the first approach to segmentation by using Downloadable Access Control Lists (dACLs). See Figure 3-5, which describes how a dACL is provisioned on a device when an IACS asset gets attached to the network. In Figure 3-5, there are two Cell/Area Zones connected via a distribution switch. There are two controllers: Controller-A (10.20.25.10) in Cell/Area Zone -1 and Controller-B (10.20.8.10) in Cell/Area Zone -2.

Figure 3-5 Segmentation Using dACL

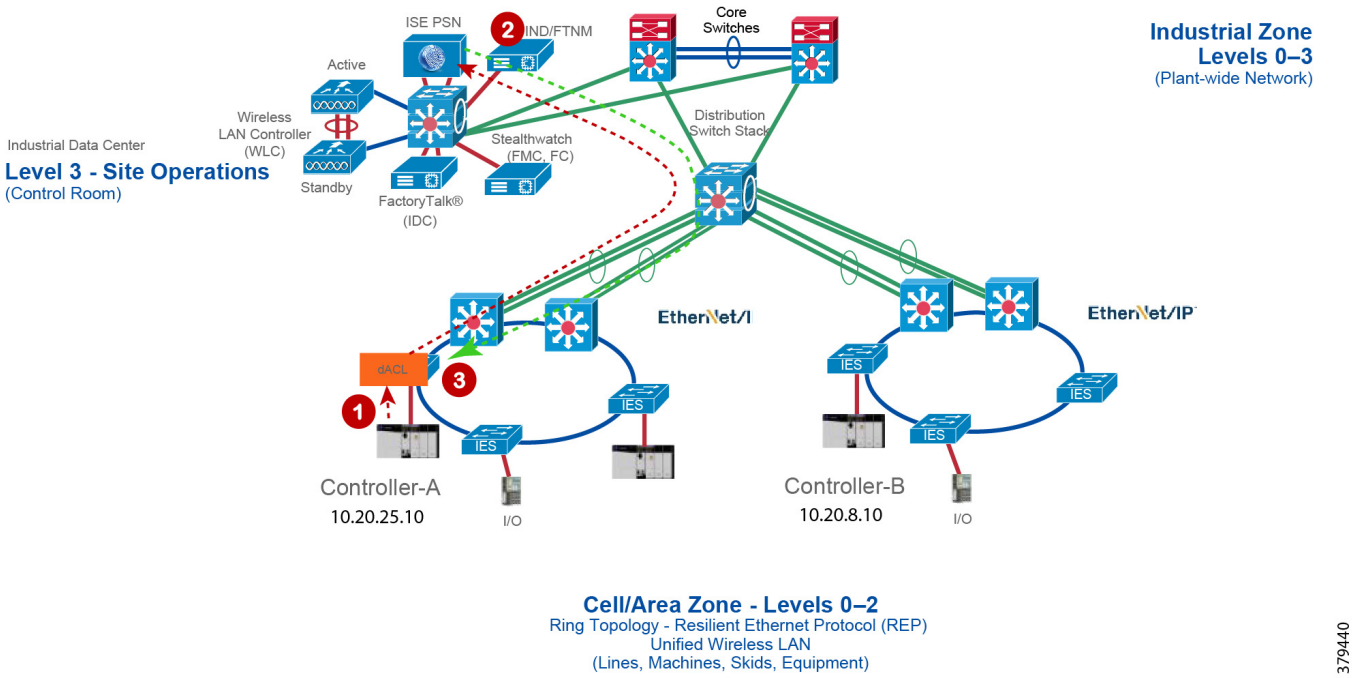

1. The Controller connects to an access port on the IES which in-turn sends an 802.1X MAB authentication request to the Cisco ISE.
2. The Cisco ISE, upon receiving the request, processes the request using the configured authentication and authorization policy and sends the authorization result as a dACL to the distribution layer switch.
3. The dACL installed on the IES to which Controller-A is attached, determines the destination IP addresses with which this Controller can communicate. If a control needs to be imposed, then add an entry in the dACL.

The dACL must have Access Control Entries (ACEs) specifying which IP address is allowed to communicate with which IP address. In Figure 3-5, if CONTROLLER-A with IP address of 10.20.25.10 is permitted to communicate with CONTROLLER-B with IP address of 10.20.8.10, then the ACE must have a permit statement with 10.20.25.10 to 10.20.8.10.

The above method works in controlling access to a Cell/Area Zone and also between the Cell/Area Zones. However, this method has the following disadvantages:

- Assume communication is allowed between CONTROLLER_A and CONTROLLER_B. If CONTROLLER_B moved to a new location with a different IP address, then the dACL needs to be updated.
- If a CONTROLLER_A is allowed to communicate with a particular server in the Industrial Zone and if the IP address of the server changes, then the dACL needs to be updated again.
- If there is a large dACL, then it could impact the performance of the distribution switch.

Segmentation—TrustSec

Cisco TrustSec technology assigns SGTs to IACS assets, networking devices, and users when they attach to a network. By using these tags, an IT security architect can define an access policy and enforce that policy on any networking device. Cisco TrustSec is defined in three phases: classification, propagation, and enforcement. When the users and IACS assets connect to a network, the network assigns them a specific SGT in a process called classification. Classification can be based on the results of the authentication and authorization policies and SGT is an end result of that process. For example, an IACS asset can be classified and assigned a specific tag if the IACS asset is a controller, I/O, HMI, or Windows workstation. Depending on the IACS asset type, a separate tag can be assigned to the IACS asset. Figure 3-6 shows how a controller is assigned an SGT value of 5. The process of SGT assignment is similar to how a dACL is pushed to the Cisco distribution switch when an IACS asset is attached to the IES. The only difference is that instead of a dACL, an SGT value is assigned. As shown in Figure 3-6, when Controller-A attaches to the IES, the IES goes through the 802.1X authentication and authorization with ISE and the result is a tag assignment to the IACS asset.

Apart from using Cisco TrustSec, customers can also use the methods described in the previous sections to segment the network, such as static ACLs and dACLs. However, the above two methods are difficult to manage, which can introduce errors during the deployment. Static ACLs need to be constantly managed, for example removing older entries and adding newer entries. Also, if the static ACL size becomes very large, then this can cause performance impact to the distribution switch. The second method using dACLs works well when the policy enforcement is done in the north-south communication flow—restricting communication from the Cell/Area Zone to any Zone above it. To restrict communication for the inter-Cell/Area Zone, dACL has the same limitation as static ACLs, namely the need to update the IP addresses whenever an IACS asset IP address changes.

Figure 3-6 Cisco TrustSec Device Classification

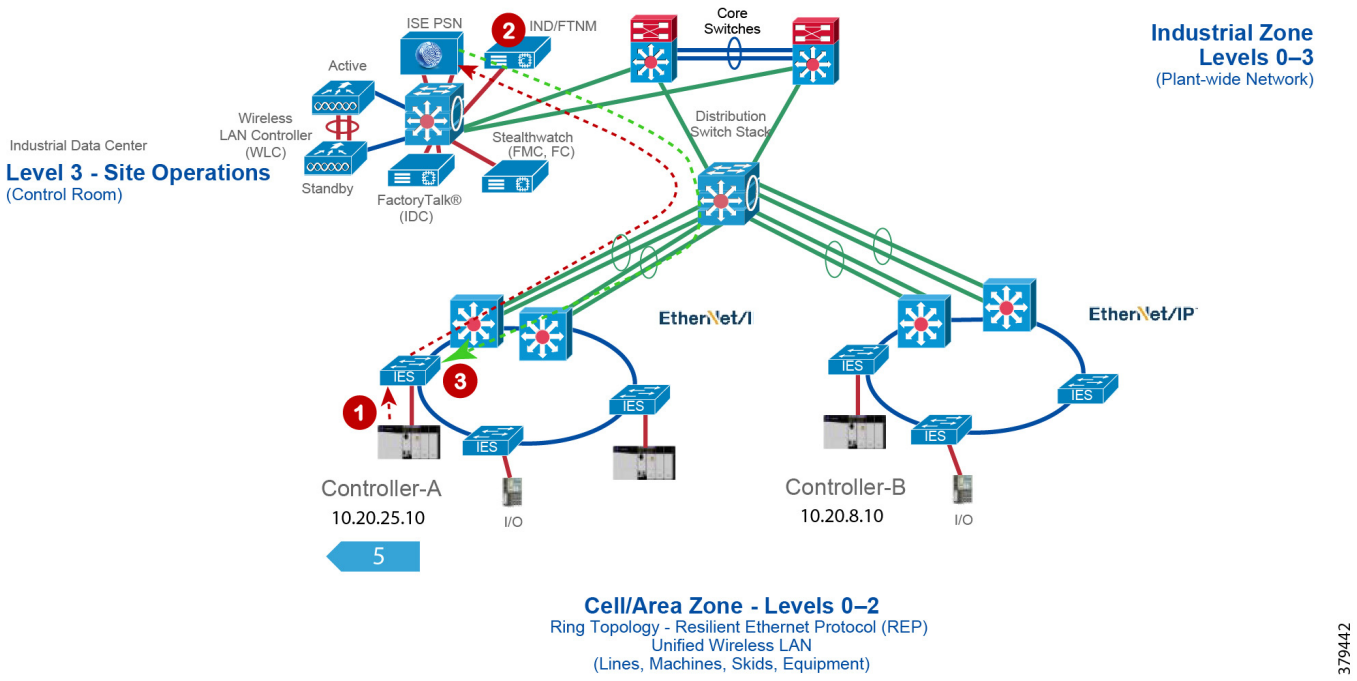

The next phase of TrustSec is propagation in which the SGT tag is made of the Ethernet frame and sent from one switch or router to another device. The SGT tag that is assigned to the IACS asset must propagate along with every packet generated by the IACS asset. Figure 3-7 shows how an SGT inserted frame is propagated in the network. In Figure 3-7, the Controller-A has IP address of 10.20.25.10 and is assigned an SGT value of 5. When an Ethernet frame is generated by Controller-A, the IES inserts the SGT value of 5 along with the IP address and sends it to the next switch. The incoming switch, if configured with SGT in-line tagging, propagates the same frame to the next switch and this information travels in hop-by-hop fashion to the destination.

Figure 3-7 Cisco TrustSec SGT Propagation

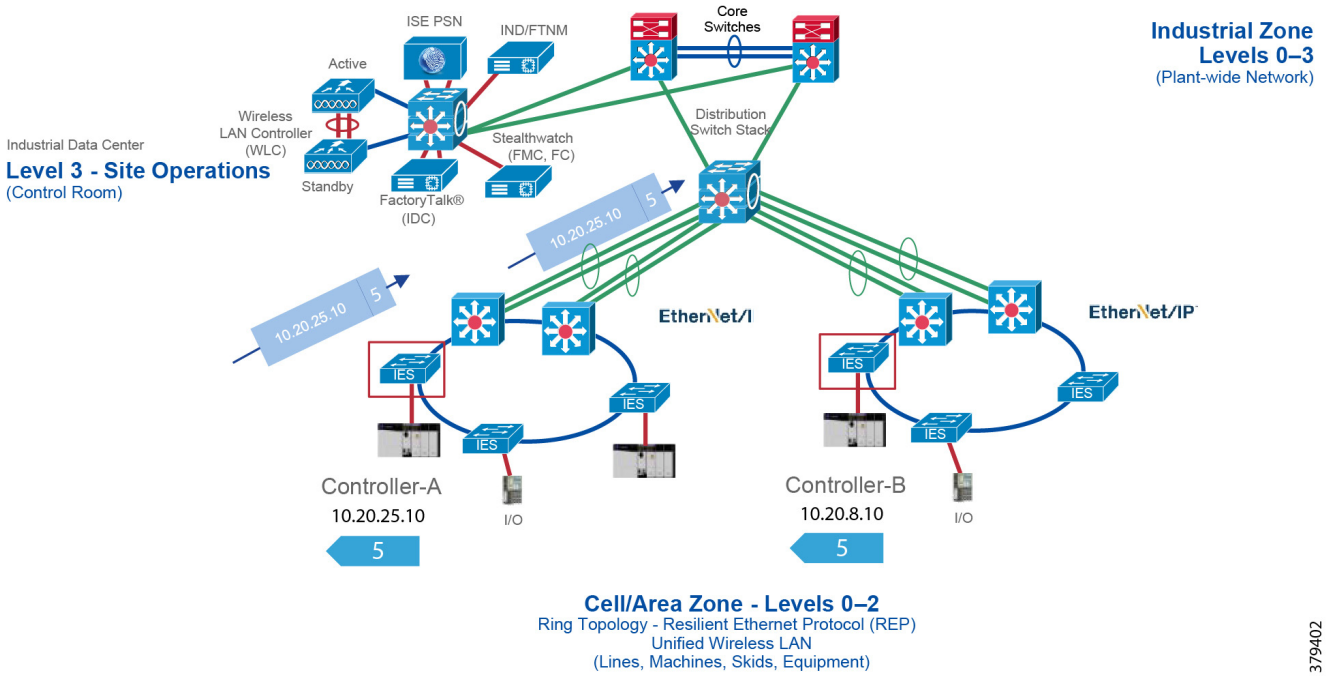

The previous phase describes a scenario for propagation using a method called in-line tagging. However, in certain network topologies there could be a situation where a switch in the path from the source to the destination does not support in-line tagging. When that scenario happens, the non-SGT capable switch would ignore the SGT in the frame and would send a normal Ethernet frame on the out-going interface. In other words, for in-line tagging feature to work, all the switches in the path must support this feature or the technology would not be applicable.

To circumvent that problem, Cisco TrustSec also supports a different mechanism to transport SGT frames over a path when a non-SGT capable IES (e.g., Allen-Bradley Stratix 5700 or Cisco IE 2000) is present by using an exchange protocol called SGT Exchange Protocol (SXP). SXP is used to securely share SGT- to-IP address mapping. Figure 3-8 shows how to exchange SGT binding over SXP tunnel. In Figure 3-8, Controller-A is establishing communication with Controller-B using an SGT tag value of 5. As you can see there is a non-SGT device in the path and this switch would ignore the SGT enabled frame coming from the distribution switch. For SGT information to be sent to Switch-B, an SXP tunnel is required between Switch-A and Switch-B. This tunnel would carry the binding information, which is 10.20.25.10 mapped to SGT 5.

Figure 3-8 Cisco TrustSec SGT Propagation Using SXP Tunnel

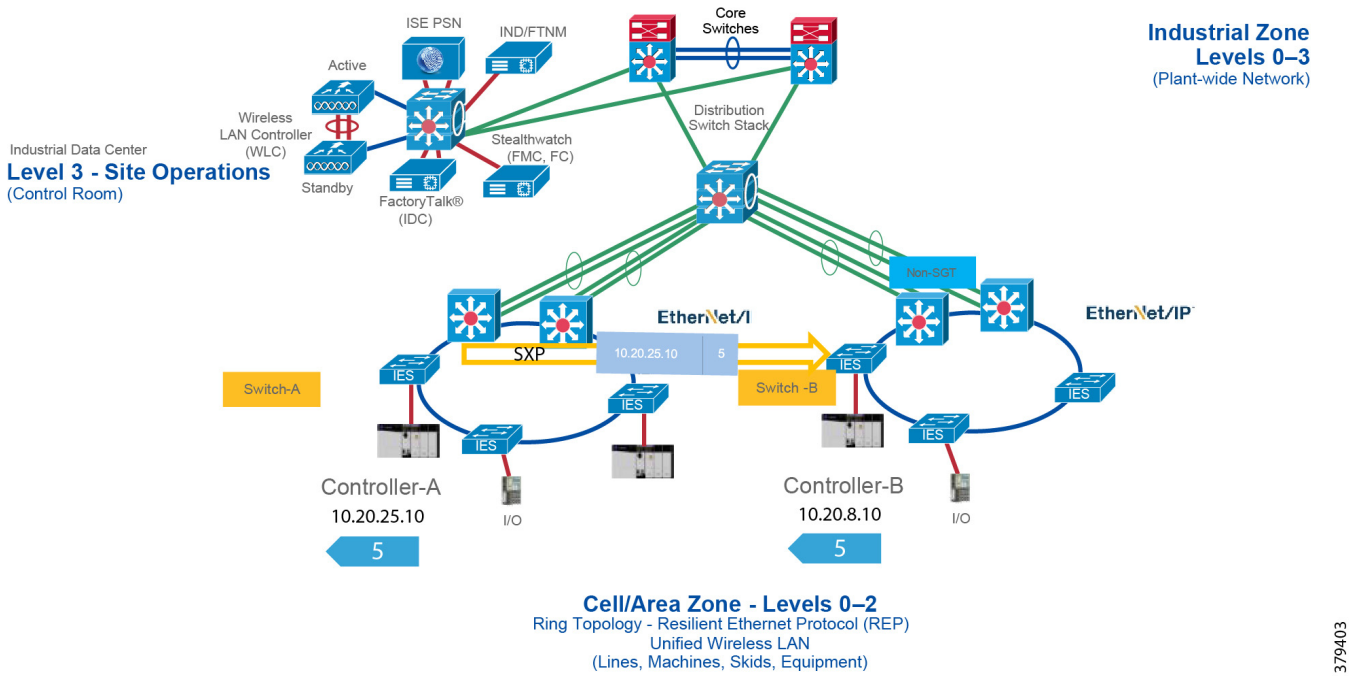

The next stage of Cisco TrustSec is policy enforcement. The enforcement device controls traffic based on the tag information. A TrustSec enforcement point can be a Cisco firewall, router, or switch. The enforcement device takes the source SGT and looks it up against the destination SGT to determine if the traffic should be allowed or denied. The advantage of Cisco TrustSec is that any switch, router, or firewall between the source and the destination can impose the policy, but the key requirement is that the enforcement point must be able to map the destination IP address to the tag value. This process is further explained in Figure 3-9. In this scenario Controller_A has been given SGT value of 5 and Controller_B, which is of similar device type, is also given an SGT value of 5. The IO device is given a different tag value because it is of a different device type. Now, in this scenario Controller_A is allowed to establish communication with Controller_C. However, the IO device is not allowed to establish communication with Controller_C. The access policy can be described in Table 3-1.

Table 3-1 Access Policy Example

	SGT 5	SGT 6
SGT 5	Yes	No
SGT 6	No	No

The next step is to determine where to apply the policy. As shown in Figure 3-9, the enforcement can be at switches A, B, C, or D. However, as previously indicated, for a switch to enforce a policy it must be able to derive the destination IP address to the tag value. For example, at point A there are two flows occurring: 1) 10.20.25.10, 5 ---- 10.20.8.10, 5 and 2) 10.20.25.12,6 --- 10.20.8.10,5. If the access policy at point A is imposed, then the switch would be only able to understand the source tag, but it has no knowledge of the destination IP address to tag mapping. Switch A would see that the destination IP address is 10.20.8.10, but it does not know that 10.20.8.10 is mapped to tag value of 5, which should be allowed. The same behavior

would be seen if the policy is applied at the point B. If the policy is applied at point C or D it would work because both C, which is Layer 2 adjacent to D, and Switch D, which is directly attached to Controller C, would be able to enforce the policy correctly because it would be able to derive the association between the destination IP and the associated SGT value.

Even though applying access policy at the point which is closest to the IACS asset is the preferred choice, in some situations a policy needs to be applied at a different point. Whenever away from the IACS asset, the knowledge of the mapping between the SGT value and the IP address is lost. To circumvent that problem establish SXP tunnels to the IES that has IACS assets attached to it. The details of using SXP for deriving the mapping information are covered below.

Figure 3-9 Access Policy Enforcement Example

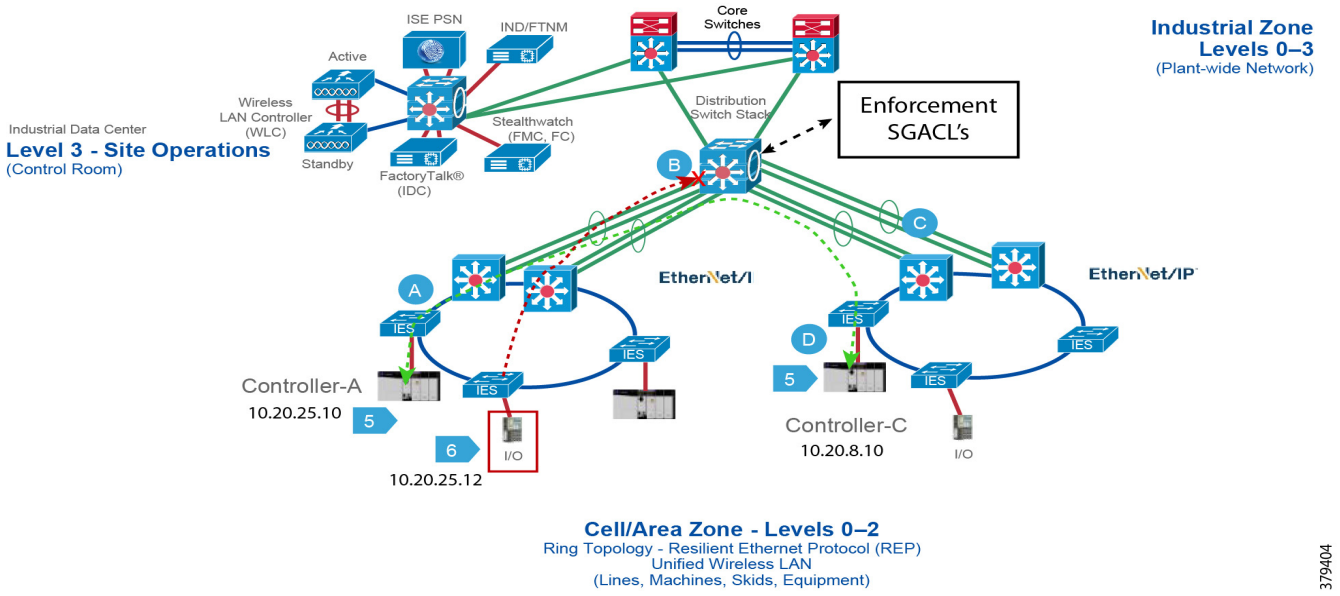

Enforcement Point

The IT security architect must next decide where in the design the access policy should be enforced. Policy enforcement occurs at the distribution switch and there are pros and cons associated with each design choice. For example, consider the case where the policy is enforced on an IES located in the Cell/Area Zone. As stated in the previous section, the basic assumption is that every IACS asset in the Cell/Area Zone must be able to access every other IACS asset. The second assumption is that policies are enforced on East-West communication going across the Cell/Area Zones. For example, two Cell/Area Zones, Cell/Area Zone-1 and Cell/Area Zone-2, and a PAC and I/O are both in the Cell/Area Zones. From a Cell/Area Zone-1 intra-zone policy perspective, every PAC and I/O in Cell/Area Zone-1 must be able to access one another. The inter-Cell/Area Zone security access policy is to block the communication between I/O in Cell/Area Zone-1 to PAC in Cell/Area Zone-2. This security access policy is shown in Table 3-2.

Table 3-2 Enforcement Point

	PAC-Cell/Area-1	I/O-Cell/Area-1	PAC-Cell/Area-2	I/O-Cell/Area-2
PAC-Cell/Area-1	Yes	Yes	No	No
I/O-Cell/Area-1	Yes	Yes	No	No
PAC-Cell/Area-2	No	No	Yes	Yes
I/O-Cell/Area-2	No	No	Yes	Yes

When designing a security policy using TrustSec, associate each IACS asset with a tag. If a PAC tag of 10 and I/O tag of 20 are assigned, designing the same matrix and restricting communication between 10 and 20, then two policy tables are needed: 1) Intra-Cell/Area Zone and 2) Inter-Cell/Area Zone.

Table 3-3 Intra-Cell/Area Zone Access Policy Enforcement

	SGT 10	SGT 20
PAC-Cell/Area-1 SGT 10	Yes	Yes
I/O-Cell/Area-1 SGT 20	Yes	Yes

Table 3-4 Inter-Cell/Area Zone Access Policy Enforcement

	SGT 10	SGT 20
PAC-Cell/Area-2 SGT 10	No	No
I/O-Cell/Area-2 SGT 20	No	No

As seen above, the Cell/Area Zone IES needs to have two tables implemented and that is not possible with the current design. The current TrustSec policy enforcement supports only a single matrix. To ensure both objectives are achieved, implement the security access policy on the distribution switch and do not have any enforcement on the Cell/Area Zone IES. By doing so, the Table 3-3 and Table 3-4 policy requirements have been met because when no policy is imposed on the Cell/Area Zone IES, then all the IACS assets on the Cell/Area Zone IES can communicate with each other. When Figure 3-9 is implemented on the distribution switch, then the inter-Cell/Area Zone or East-West communication can be restricted. Figure 3-10 shows the inter-Cell/Area Zone security access policy enforcement point. If the industrial security access policy requires intra-Cell/Area Zone access control, Cisco and Rockwell Automation recommend IACS application security such as CIP Security from ODVA, Inc.

Figure 3-10 Enforcement Point in CPwE Network Security

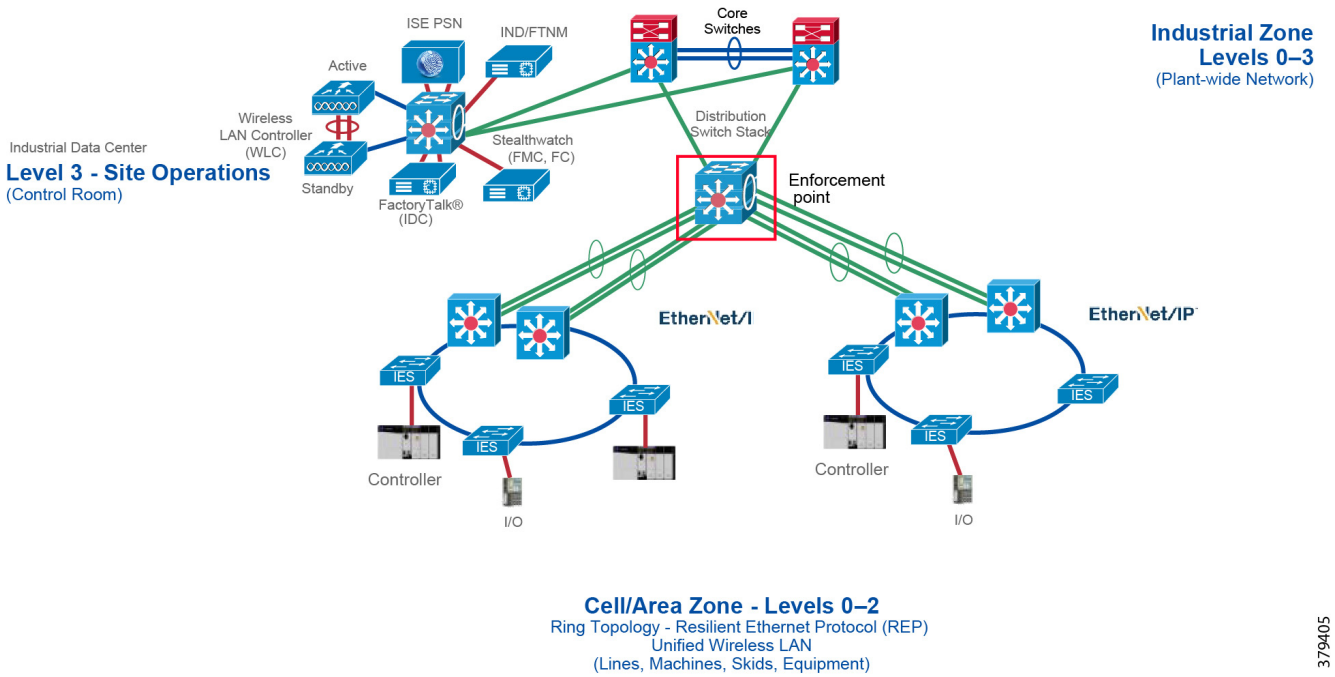

Scalable Group Tag Exchange Protocol Considerations

Scalable Group Tag Exchange Protocol (SXP) is used to propagate the SGTs across network devices that do not have hardware support for TrustSec. SXP is used to transport an endpoint's SGT along with the IP address from one SGT-aware network device to another. The data that SXP transports is called as IP-SGT mapping. The SGT to which an endpoint belongs can be assigned statically or dynamically and the SGT can be used as a classifier in network policies.

SXP uses TCP as its transport protocol to set up an SXP connection between two separate network devices. Each SXP connection has one peer designated as SXP speaker and the other peer as SXP listener. The peers can also be configured in a bi-directional mode where each of them acts as both speaker and listener. Connections can be initiated by either peer, but mapping information is always propagated from a speaker to a listener.

As shown in the previous section, the enforcement is moved to the distribution switch, so the distribution switch needs to derive the destination IP address to SGT. This is because the Ethernet frame has only the source SGT information and to enforce the policy the distribution switch needs to learn the SGT binding associated with the destination IP address. To help the distribution switch to derive the destination tag, SXP tunnels are needed from the access layer IES to the distribution.

In the current design, SXP tunnels are established from the access layer IES to the Cisco ISE and the distribution switch also has an SXP tunnel to the Cisco ISE. This way the IP-SGT binding information is sent to the Cisco ISE and the distribution switch learns the IP-SGT binding information from the Cisco ISE. Figure 3-11 depicts the design.

Figure 3-11 SXP Design in CPwE Network Security CVD

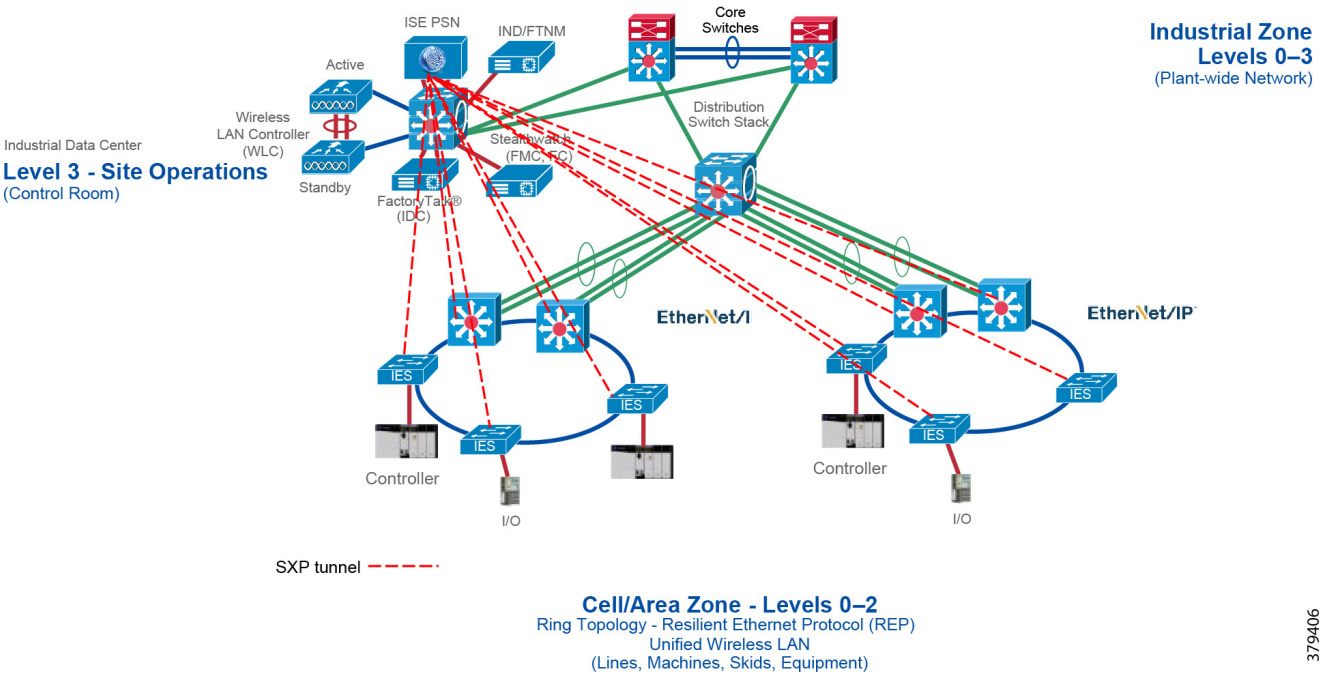

NetFlow Data Collection

A flow is a unidirectional connection between a source and a destination. To describe a full exchange between two devices, two independent unidirectional flows are needed. For example, when data is flowing between client and server, then there are two flows occurring: from client to server and from server to client. NetFlow is a protocol that creates flow records for the packets flowing through the switches and the routers in a network between the end devices and exports the flow records to a Flow Collector. The data collected by the Flow Collector is used for different applications to provide further analysis. Initially, NetFlow was used for providing traffic statistics in a network, but later it started gaining traction as a network security tool. In CPwE Network Security CVD, NetFlow is primarily used for providing security analysis, such as malware detection, network anomalies, and so on. There are many advantages in deploying NetFlow:

- NetFlow can be used for both ingress and egress packets.
- Each networking device in a network can be independently enabled with NetFlow.
- NetFlow does not require a separate management network to collect the traffic.

In a normal flow the 5-tuples information (source IP, destination IP, source port, destination port, and protocol) information is recorded as shown in Figure 3-12.

Figure 3-12 NetFlow Data Collection

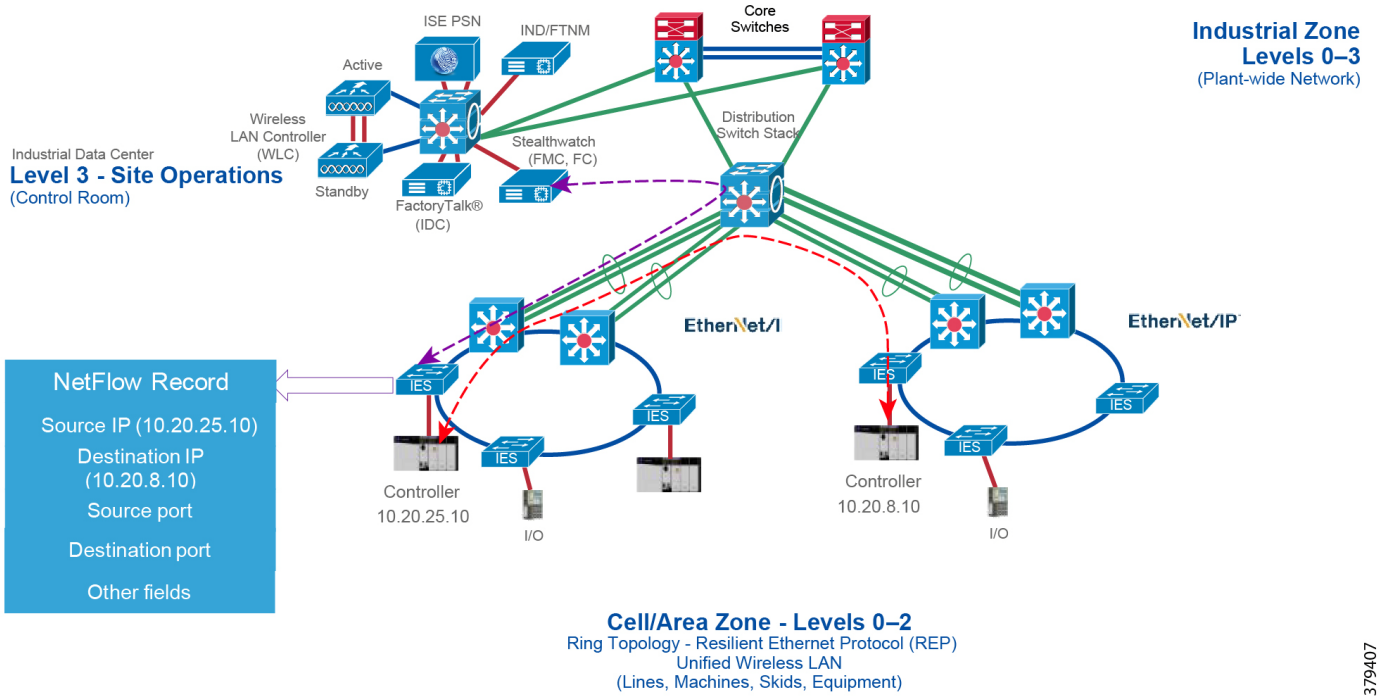

With the latest releases of NetFlow, the switch/router can gather additional information such as ToS, source MAC address, destination MAC address, interface input, interface output, and so on. For NMT and ISE integration, collecting the MAC address of the device is very critical. If NMT does not gather the MAC address, then the device is not imported into ISE. The following **show** command output describes the flow record information collected at the IES in the Cell/Area Zone:

```
flow record Cisco Stealthwatch_Record
 description NetFlow record format to send to Cisco Stealthwatch
 match datalink mac source address input
 match datalink mac destination address input
 match ipv4 tos
```

```
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 collect transport tcp flags
 collect interface input
 collect interface output
 collect counter bytes long
 collect counter packets long
 collect timestamp sys-uptime first
 collect timestamp sys-uptime last
!
```

Configuration of flow records can be done by using NMT, which is discussed in more detail in Chapter 4, "Configuring the Infrastructure." The next important consideration is on managing flows. As network traffic traverses IES and distribution switches, flows are continuously created and tracked. As the flows expire, they are exported from the NetFlow cache to the Stealthwatch flow collector. A flow is ready for export when it is inactive for a certain time (for example, no new packets are received for the flow) or if the flow is long lived (active) and lasts greater than the active timer (for example, long FTP download and the standard CIP I/O connections). There are timers to determine whether a flow is inactive or a flow is long lived.

After the flow timeout the NetFlow record information is sent to the flow collector and deleted on the switch. Since the NetFlow implementation is done mainly to detect security-based incidents rather than traffic analysis, Cisco and Rockwell Automation recommend leaving the default value in the switch, which is 30 seconds for inactive timeout and 60 seconds for active timeout.

The next consideration is on enabling NetFlow in the network. Since in this CPwE Network Security CVD NetFlow is enabled for security perspective, the recommendation is to enable NetFlow monitoring on all the IES and distribution switch interfaces in the CPwE network.

Stealthwatch Deployment Considerations

The Stealthwatch solution has two different components, both of which are installed on different systems:

- Flow Collectors
- Stealthwatch Management console

The Flow Collector collects the NetFlow data from the networking devices, analyzes the data gathered, creates a profile of normal network activity, and generates an alert for any behavior that falls outside of the normal profile. Based on the network flow traffic, there could be one or multiple Flow Collectors in a network. The Stealthwatch Management Console (SMC) provides a single point for the IT Security Architect to get a contextual view of the entire network traffic.

Note For example, if there is a single Flow Collector and a single SMC, then there are two virtual or hardware appliances and each appliance has its own IP address and its own device credentials.

The SMC client allows an IT Security Architect to access the SMC graphical user interface by using a web browser. SMC enables the following:

- Centralized management, configuration, and reporting for up to 25 Flow Collectors
- Graphical Charts for visualizing traffic
- Drill down analysis for troubleshooting
- Consolidated and customizable reports

- Trend analysis
- Performance monitoring
- Immediate notification of security breaches

Some important considerations when deploying a Stealthwatch solution include:

- Stealthwatch is available as both hardware (physical appliances) and virtual appliances. To install hardware and software appliances, refer to the Stealthwatch guide at: https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/system_installation_configuration/SW_7_0_Installation_and_Configuration_Guide_DV_1_0.pdf.
- The resources allocation for the Stealthwatch Flow Collector are dependent on the number of flows per second expected on the network and the number of exporters (networking devices that are enabled with NetFlow) and the number of hosts attached to the each networking device. The scalability requirements for the Flow Collector are available at: https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/system_installation_configuration/SW_7_0_Installation_and_Configuration_Guide_DV_1_0.pdf.
- The Data Storage requirements must be taken into consideration, which are again dependent on the number of flows in the network. The sizing table for Data Storage is available at: https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/system_installation_configuration/SW_7_0_Installation_and_Configuration_Guide_DV_1_0.pdf.
- A specific set of ports needs to be open to the Stealthwatch solution in both the inbound and outbound directions. For example, HTTPS needs to be open inbound so that a client can access the Stealthwatch solution for managing the appliances. Similarly, certain ports such as DNS, NTP, and external log sources should be open in the outbound direction so that the Stealthwatch solution can reach those services. For the complete list of ports that are recommended to be open, refere to: https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/system_installation_configuration/SW_7_0_Installation_and_Configuration_Guide_DV_1_0.pdf.

Stealthwatch Flow Collection

This section describes:

Launching the Stealthwatch Web Client

The Stealthwatch web client is a JAVA-based client that can be used to configure Stealthwatch and view flow data. To access the Stealthwatch web client, the IT Security Architect can use a web browser to establish a HTTPS connection to the Stealthwatch web client and then click the **Desktop Client** button, which is found in the upper-right corner of the screen as shown in Figure 3-13:

Figure 3-13 Desktop Client Button

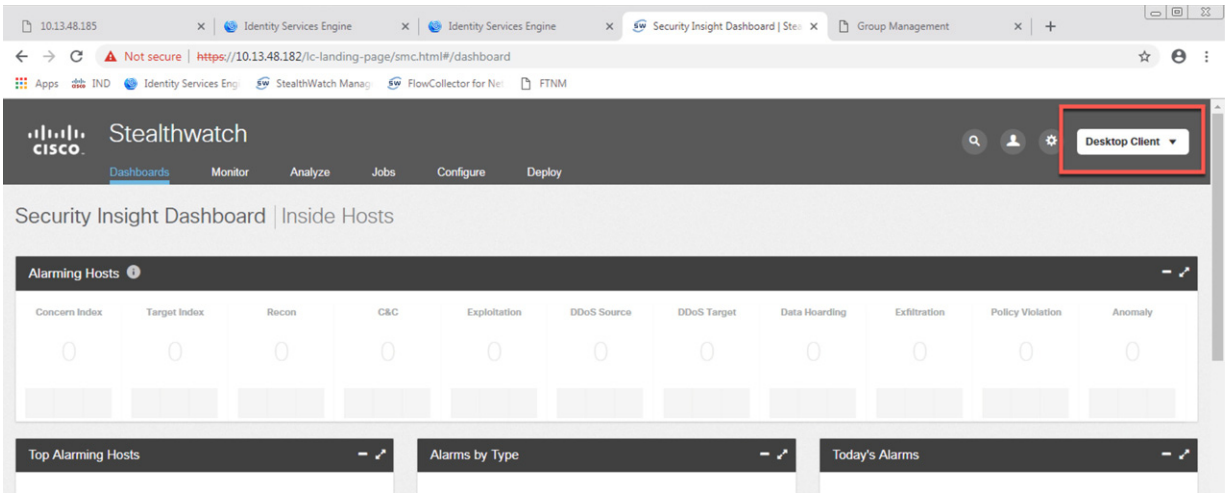

Clicking the **Desktop Client** button downloads the software on the computer where the HTTPS session was initiated. Opening the file displays Figure 3-14.

Figure 3-14 Software Installer

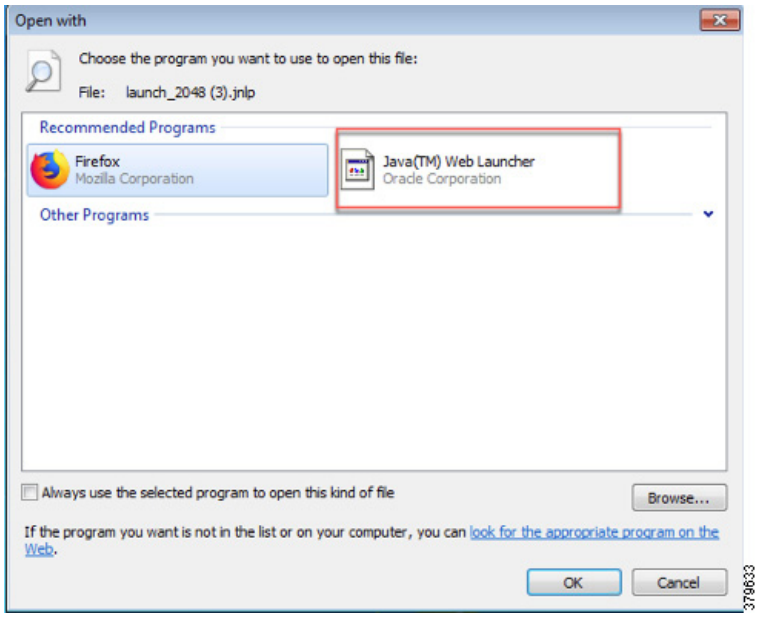

After a successful log in, the IT Security Architect sees the screen in Figure 3-15.

Figure 3-15 Stealthwatch Management Console

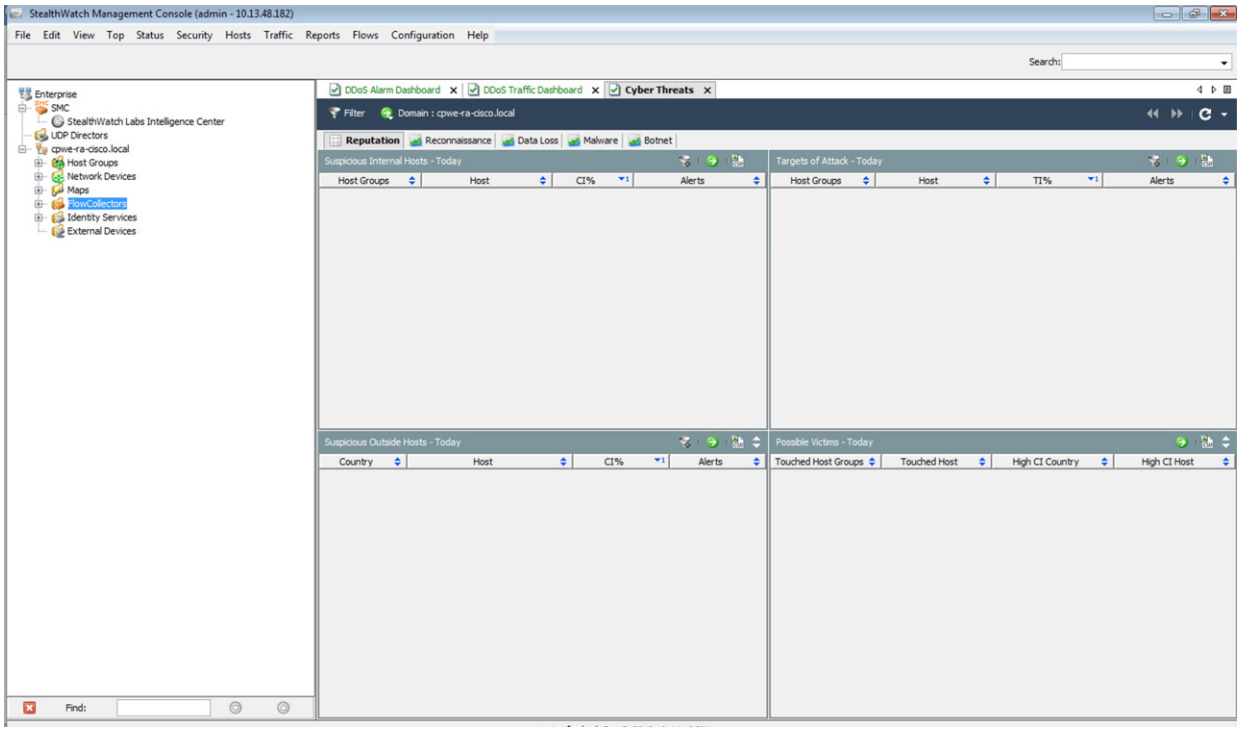

Configuring Flow Collector

Typically, the Flow Collectors are registered during the installation of the SMC. However, if a Flow Collector needs to be added to the SMC after the initial set up, then select the option **FlowCollector** -> **Add FlowCollector**.

Figure 3-16 Adding a Flow Collector

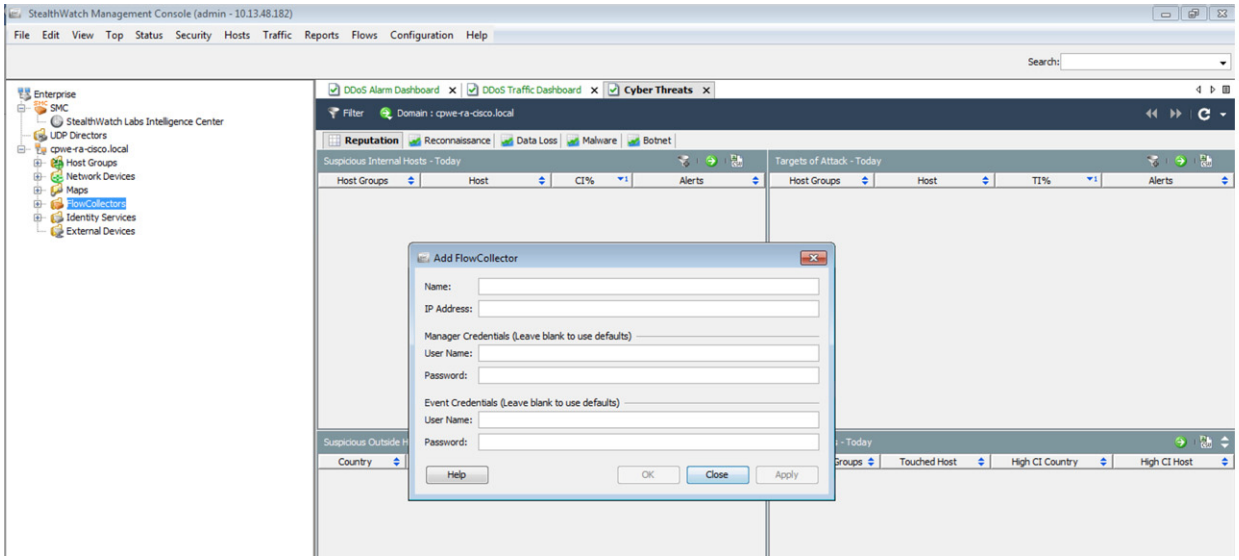

Configuring Flow Exporters

In Stealthwatch context, Flow Exporters are the networking devices that send the flows to the Stealthwatch Flow Collector. The configuration for the Flow Exporter is done by selecting **Exporter** -> **Configuration** -> **Add Exporter**, as shown in Figure 3-17.

Figure 3-17 Configure Flow Exporters

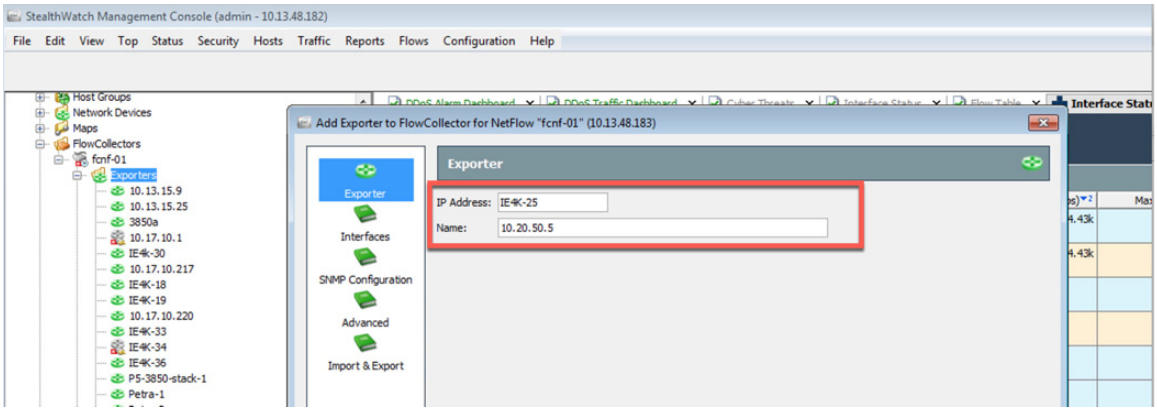

Once configured, the Flow Exporter should up in the tree under the Flow Collector, as shown in Figure 3-18.

Figure 3-18 Flow Exporters

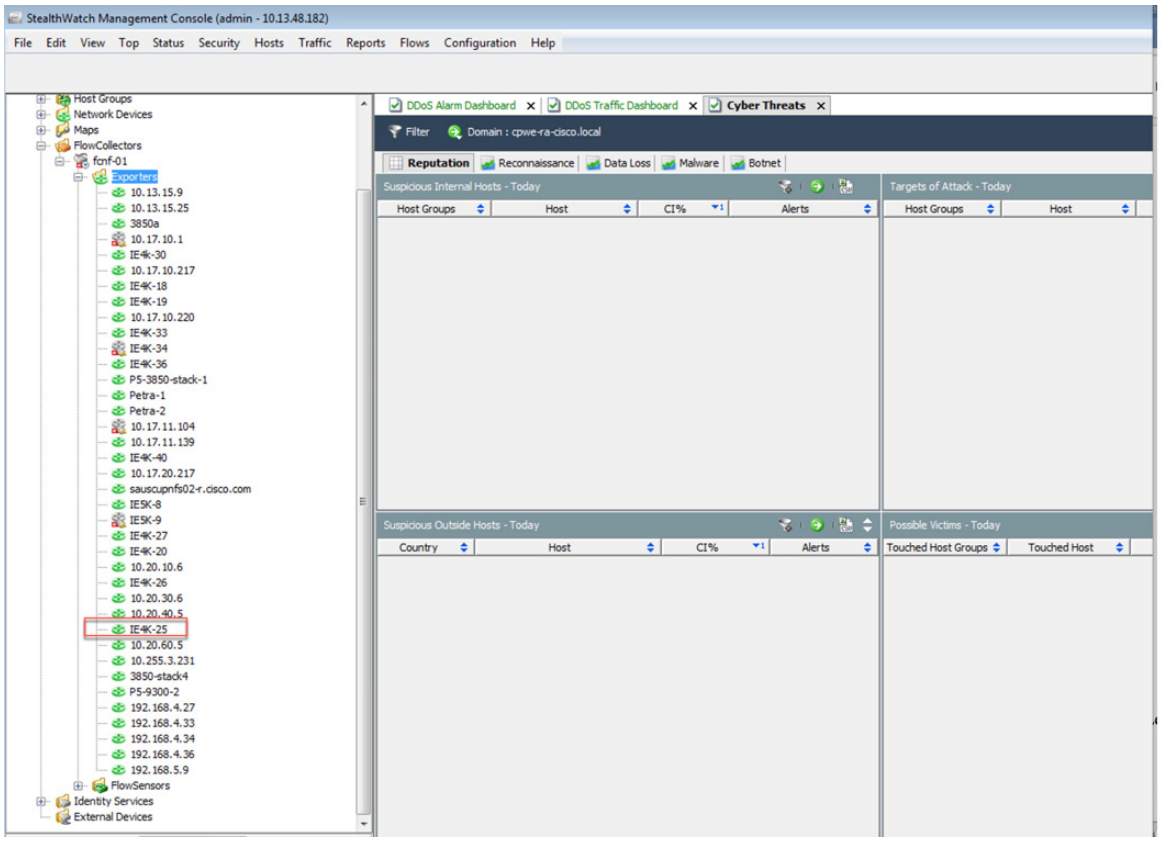

Configuring the Host Groups

A host group is a “container” of hosts or IP addresses that share attributes and policies. Host groups enable the establishment of different thresholds or the bypassing of alerts for certain behavior. Using host groups correctly in the Stealthwatch solution helps ensure that you are alerted correctly about events and that the information provided to you is relevant. The following are some of the different attributes that are typically grouped together:

- Shared functions
- Exhibits similar behavior
- Can be managed as a single object
- To which a single policy can be applied

For more information about Host Groups, refer to:
https://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/iot-threat-defense-mfg-design-implementation-guide.pdf

Viewing the Flows Generated by Flow Exporter

To view the flows generated by a Flow Exporter, select the **Flow Exporter** -> **Flows** -> **Flow Table**, as shown in Figure 3-19.

Figure 3-19 Flow Exporter Flows

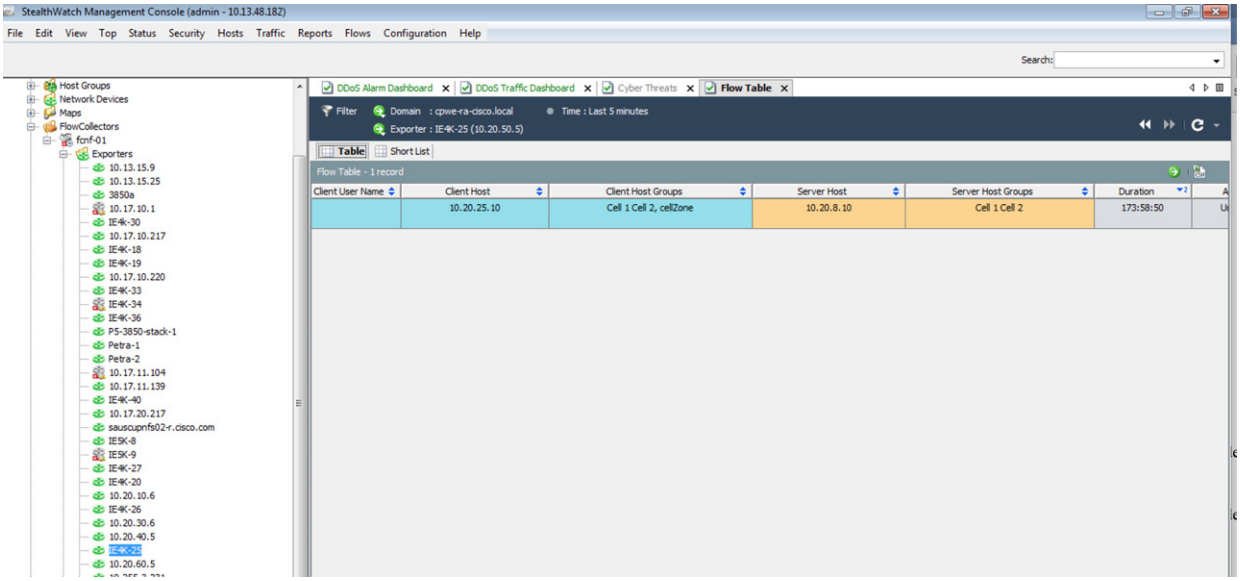

Cisco ISE Deployment Considerations

Deploying Cisco ISE in a large network requires an IT security architect to consider several factors such as scalability and high-availability of the Cisco ISE deployment. Cisco and Rockwell Automation have developed a design and implementation guide on Deploying Identity and Mobility Services within CPwE. This design guide covers in depth many factors related to deploying a large-scale Cisco ISE deployment

model. The design considerations listed in the CPwE Identity and Mobility Services CVD are very much related to the current CPwE Network Security CVD effort. Cisco and Rockwell Automation encourage the reader to read this DIG to develop a good understanding of large-scale solution deployments. Some of the key recommendations given in the design guide are shown here as a quick reference.

In the distributed installation, the Cisco ISE system is divided into three discrete nodes (personas)—Administration, Policy Service, and Monitoring—which are described as follows:

- Policy Administration Node (PAN) allows the Enterprise IT team to perform all administrative operations on the distributed Cisco ISE system. The PAN (located in the Enterprise Zone) handles all system configurations that are related to functionality such as authentication and authorization policies. A distributed Cisco ISE deployment can have one or a maximum of two nodes with the Administration persona that can take on the primary or secondary role for high availability.
- Policy Service Node (PSN) provides client authentication, authorization, provisioning, profiling, and posturing services. The PSN (located within the Industrial and the Enterprise Zone) evaluates the policies and provides network access to devices based on the result of the policy evaluation. At least one node in a distributed setup should assume the Policy Service persona and usually more than one PSN exists in a large distributed deployment.
- Monitoring Node (MnT) functions as the log collector and stores log messages and statistics from all the Administration and Policy Service Nodes in a network. The MnT (located in the Enterprise Zone) aggregates and correlates the data in meaningful reports for the Enterprise IT and operational technology (OT) personnel. A distributed Cisco ISE system can have at least one or a maximum of two nodes with the Monitoring persona that can take on primary or secondary roles for high availability.

For optimal performance and resiliency, Cisco and Rockwell Automation provide these recommendations for the CPwE Identity and Mobility Services architecture:

- Administration and Policy Service personas should be configured on different Cisco ISE nodes.
- Monitoring and Policy Service personas should not be enabled on the same Cisco ISE node. The Monitoring node should be dedicated solely to monitoring for optimum performance.
- A PSN should be placed in the Industrial Zone (Levels 0-3) to provide services for clients in the Industrial Zone. If the Enterprise and Industrial Zones become isolated, any existing clients will still be able to securely access the network. For best practices, see Appendix A, “References” for links to the CPwE IDMZ CVD DIG.
- A PSN should also be present in the Enterprise Zone to authenticate corporate mobile users who connect to the corporate network through the IDMZ in a secure data tunnel. This scenario is covered in detail later in the document.

Based on the recommendations above, a typical distributed Cisco ISE deployment in the CPwE architecture consists of the following nodes (hardware appliances or VM) as shown in Figure 3-20.

- One Primary Administration/Secondary Monitoring node
- One Secondary Administration/Primary Monitoring node
- One or several PSN in the Enterprise Zone
- One or several PSN in the Industrial Zone

Note The number of PSN in the Enterprise and Industrial Zones may depend on the company size, the number of active clients, redundancy requirements, and geographical distribution (for example, one PSN per each plant).

Figure 3-20 ISE Deployment Models

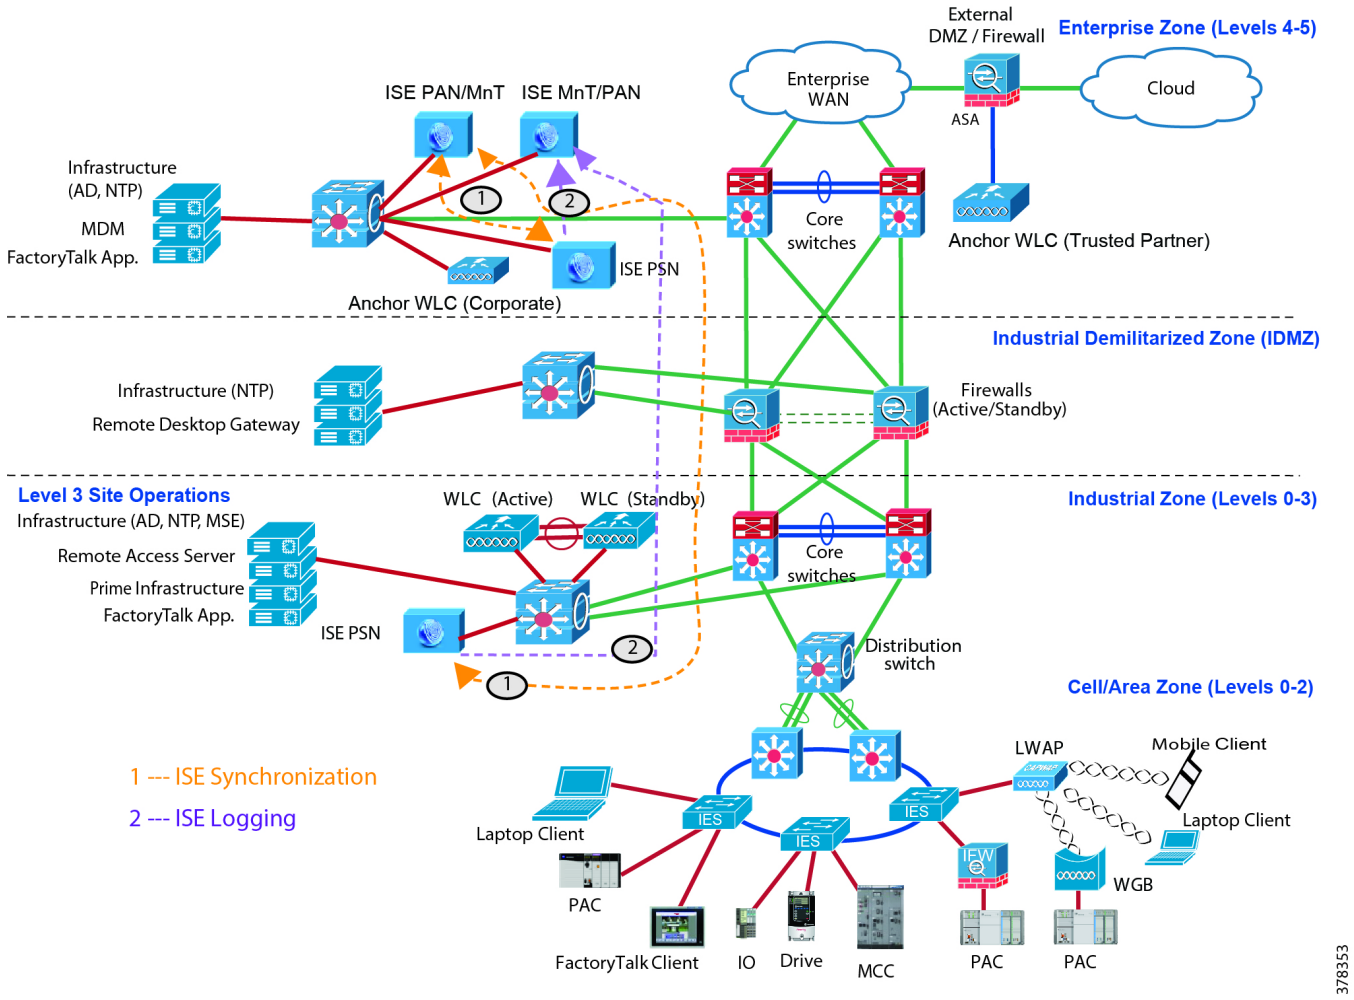

IPDT Considerations

IP Device Tracking (IPDT) is feature that allows an IES or any other switch or router to keep track of connected hosts attached to it. The IPDT feature must be enabled for several security features such as dot1x, MAB, Web-Auth, auth-proxy, and so on. The IPDT feature keeps mappings between IP addresses and mac-addresses. To do the tracking, the IES when enabled with IPDT feature sends an ARP probe with default interval of 30 seconds. The probes are implemented as per RFC5227 where the source IP address is set to 0.0.0.0. If the IPDT feature is enabled with a default source IP address of 0.0.0.0, then there could be conflict between the IES and an IACS asset that is also doing device tracking (the Duplicate IP address 0.0.0.0) problem is explained in: https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/8021x/116529-problemsolution-product-00.html).

To avoid this problem, the recommended option is to modify the standard IP address used with the IPDT feature prior to the implementation of IPDT. The following command can be used in an IES:

```
ip device tracking probe auto-source fallback 169.254.26.64 0.0.0.0 override
```

This command uses the source of the probe to SVI if present and falls back to 169.254.26.64, which is a link-local IP address. The rational for using a link-local IP address as a fallback is based on the assumption that any device attached to the switch does not have a link-local IP address. The link-local IP address is used only to route packets within a local segment and if a router receives a link-local IP address then it does not forward the packet. The IT security architect must verify if there is any link-local IP address present in the network before enabling the command.

Note IP Device Tracking (IPDT), which operates in accordance with RFC 5227, must be enabled on the IES to implement RADIUS downloadable ACL, NetFlow, and SGT. IPDT uses ARP probes to determine the IP addresses of hosts on different ports; this behavior may disrupt IACS assets devices and applications. IPDT should only be enabled in the following situations on IES ports with 802.1X authentication:

- Maintenance ports and/or designated non-IACS equipment ports
- IACS ports with MAC Authentication Bypass if DACL is required by the security policy, with proper IPDT workaround applied and tested with IACS assets devices and applications

By default, IPDT should not be enabled on ports connected to IACS assets devices and applications if DACL functionality is not required. Refer to the URL below for more details and IPDT workarounds: https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/5-0/Firewalls/DIG/CPwE-5-IFS-DIG.html