CHAPTER 2

System Design Considerations

This chapter includes the following major topics:

This chapter provides a high-level overview of the basic design considerations for the Industrial Demilitarized Zone (IDMZ) of the CPwE architecture. CPwE IDMZ offers basic design and implementation guidance for the IDMZ, which IACS networking personnel could use to design and deploy their architecture. Often, the IDMZ is where IT networking resources are involved in the design, implementation and maintenance. For more complex deployments, Cisco and Rockwell Automation recommend the involvement of either external resources or Enterprise IT networking specialists.

Note This chapter provides both general descriptions of product capabilities and specific design recommendations for the CPwE IDMZ architecture. Refer to Configuring the Infrastructure for more information about specific features and configuration steps that have been validated for the CPwE architecture.

CPwE IDMZ Overview

This section describes the concepts, objectives and main design principles of the IDMZ.

What is the IDMZ?

The Industrial Zone contains all IACS network and automation equipment that is critical to controlling and monitoring plant-wide operations. Hierarchically, the Industrial Zone includes Site Operations (Level 3) and multiple Cell/Area Zones (Levels 0 to 2).

To preserve smooth plant-wide operations and functioning of the IACS applications and network, the Industrial Zone requires clear segmentation and protection from the Enterprise Zone via security devices, replicated services and applications. The zone that separates the Enterprise Zone from the Industrial Zone is called the IDMZ. This insulation not only enhances security segmentation between the Enterprise and Industrial Zones, but may also represent an organizational boundary where IT and Operational Technologies (OT) responsibilities interface.

A Demilitarized Zone (DMZ) is sometimes referred to a perimeter network that exposes an organization's trusted external services and data to an untrusted network. Most of the time, the DMZ is understood as protecting a company's Enterprise assets from the Internet. A DMZ is a proven method to protect a trusted network like the Enterprise network from an untrusted network like the Internet.

In the context of securing the Industrial Zone from the Enterprise, the IDMZ is placed between a trusted network (the Industrial Zone) and an untrusted network (the Enterprise Zone). The IDMZ functions in the same manner as a traditional DMZ because it allows traffic between the zones to be terminated within the DMZ and to be inspected as it enters and exits the IDMZ.

The IDMZ is comprised of:

- Boundary or “edge” security appliances like firewall(s) that can inspect traffic as it enters and exits each security zone
- Appliances and servers that replicate services like web proxies, data proxies, file transfer proxies, application and operating system patch proxies, and application proxies

In the most basic terms, the IDMZ is a termination end point for traffic from the untrusted Enterprise network. The traffic from the Enterprise that is destined for the Industrial Zone is terminated on a server or application proxy within the IDMZ. The firewalls can inspect the traffic as it enters or exits the IDMZ. The firewall can be configured to allow remote access or file requests from certain users, but block “untrusted” users or devices from entering or exiting the IDMZ (see Figure 2-1).

Figure 2-1 IDMZ Concepts

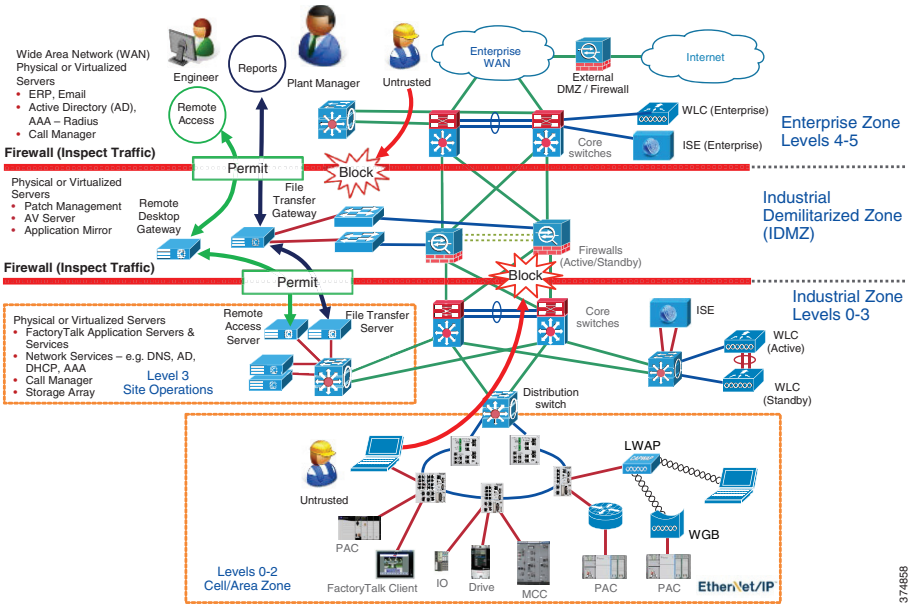

This approach permits the Industrial Zone to function entirely on its own, without taking account of the connectivity status to the higher levels. A methodology and procedure should be deployed to buffer IACS data to and from the Enterprise Zone in the event of IDMZ connectivity disruption. As a best practice, Cisco and Rockwell Automation recommend that all IACS assets required for the operation of the Industrial Zone should remain in the Industrial Zone.

This separation is necessary because real-time availability and security are the critical elements for the traffic in the IACS network. Downtime in an IACS network is much more costly than downtime of similar scale in an enterprise environment. The cost of capital, the loss of product and material, missed schedules and the wasted time of plant personnel drive this very concrete impact on revenue and efficiency. Therefore, Cisco and Rockwell Automation recommend the deployment of Industrial Zone firewalls and an IDMZ between the Industrial and Enterprise Zones to securely manage the traffic flow between these networks.

IDMZ Objectives

Data and services must be shared between the Industrial and Enterprise Zones. Many of the benefits of converged industrial and enterprise networks rely on real-time communication and transfer of data between these zones. Without Industrial Zone firewalls and an IDMZ, data cannot be shared while also maintaining the security of the IACS network and its IACS systems.

The Industrial Zone firewall:

- Enforces and strictly controls traffic from hosts or networks into and out of each security zone
- Performs stateful packet inspection
- Optionally can provide intrusion detection/prevention
- Provides security and network management support
- Terminates VPN sessions with external or internal users
- Provides web portal services such as proxy services
- Enables Remote Desktop connectivity services for secure remote access via Remote Desktop Connection client and ThinManager to servers in the Industrial Zone

IDMZ offers a network on which to place data and services to be shared between the Enterprise and Industrial Zones. The IDMZ doesn't allow direct communication between the Industrial and Enterprise Zones, but meets the data and service sharing requirement. With the deployment of an IDMZ and Industrial Zone firewall, attacks and issues that arise in one zone cannot easily affect the other zone. In fact, by temporarily disabling the IDMZ and the firewalls, an IACS or IT network administrator can help to protect a zone until the situation is resolved in the other zone.

The IDMZ network design covers the following:

- IDMZ components
- IDMZ topology
- Firewall design and implementation considerations
- IACS application interoperability

IDMZ Design Principles

To design an IDMZ, the first exercise is to fully understand:

- Which Enterprise systems need to interact with the Industrial Zone systems

- Which Industrial Zone systems need to interact with Enterprise systems
- Which Enterprise users must interact with Industrial Zone systems and the tasks they perform with these systems
- Which Industrial Zone users must interact with Enterprise systems and the tasks they perform
- How long the Enterprise systems can stay disconnected from the Industrial Zone before IDMZ connectivity is restored

After getting the answers to these questions, you will be able to define the services and data that need to be replicated or proxied within the IDMZ.

The IDMZ design process consists of gathering stakeholder requirements and designing a solution to meet the requirements. In order to do so, you will gather requirements from the people who design, operate, change and maintain these systems. Before designing an IDMZ, you must identify the assets within the Enterprise and Industrial Zones that are needed to support the IACS process.

- The **Reconnaissance Phase** is used to identify the assets or "types" of assets in the Industrial Zone used to support production and those that will feed data or reports to Enterprise systems or Enterprise users (see Figure 2-2). The Reconnaissance Phase is used to identify the systems that are located in the Industrial Zone that will interact with the Enterprise Zone and ultimately have to communicate through the IDMZ to do so. During the Reconnaissance Phase, it is also important to compile a list of asset owners so they can be interviewed and their requirement documented.

Figure 2-2 IDMZ—Network Reconnaissance (Design Pre-Work)

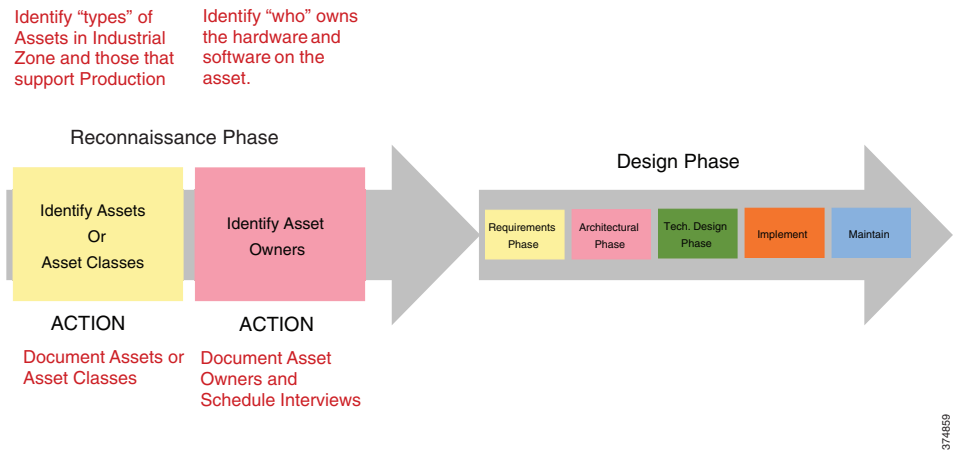

For example, it may be determined that access to a Human Machine Interface (HMI) server from the Enterprise engineering department is required. The asset owner, who would then be interviewed to gather their requirements, may state that gaining access to the HMI server to create faceplates or troubleshoot the system is required. All the system owners should be interviewed paying close attention to the tasks they perform within the system; those requirements are then the basis for design of the IDMZ solution (see Figure 2-3).

Figure 2-3 System Owners Interview Process

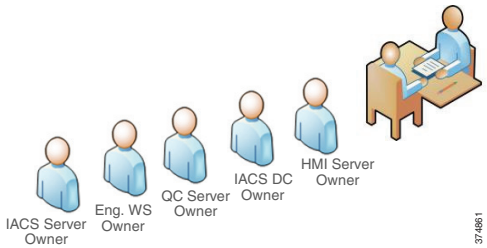

After the Reconnaissance Phase is completed, the IDMZ design phase can begin. An example of an IDMZ design cycle is listed below (see Figure 2-4). While this is not the only methodology available, it has been successfully used in the design and implementation of an IDMZ.

Figure 2-4 IDMZ Design Methodology

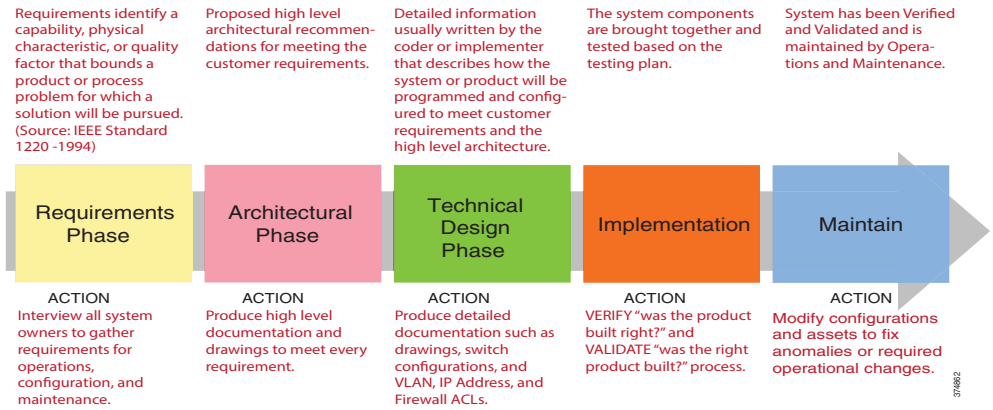

- The **Requirements Phase** is used to record all user and system requirements. The IEEE 1220 standard states that "requirements are a statement identifying a capability, physical characteristic or quality factor that bounds a product or process problem for which a solution will be pursued". All stakeholder and system requirements should be documented during this phase so technical solutions can be engineered to meet all the requirements. Requirements are derived from system users, designers, engineers and vendors.
- The **Architectural Phase** is used to propose a high-level solution to the stakeholders to see if it is acceptable prior to committing time to working on the technical solution.

  For example, let's suppose a requirement could be met with a solution that is available on a Linux operating system. Before moving forward, let's suppose the stakeholder is not familiar with nor can they support the Linux operating system. The Architectural Phase allows the technical team to propose a solution and gain consensus with the stakeholder before implementing the solution. In the last example, let's suppose the same solution is available in the Windows operating system and the stakeholder agreed to the solution. It is much better to gain consensus earlier when less technical implementation hours have been spent.
- The **Technical Design Phase** is when detailed information is written by the coder or the implementer that describes how the system or product will be programmed or configured to meet the customer's requirements. This reflects the proposed solutions in the Architectural Phase.
- The **Implementation Phase** is when the system components are brought together, tested, verified and validated per the testing plan.
- The **Maintain Phase** is when the operating systems are supported. Frequently configurations are modified to fix anomalies or to support operational changes.

IDMZ Security Policy

As mentioned previously, the IDMZ is meant to buffer and inspect traffic that is flowing between the Enterprise and Industrial Zones. When designing the IDMZ that help shape the security policies and design decisions, the following key points should be kept in mind:

- Eliminate direct traffic between the Enterprise and Industrial Zones. Every organization must assess the risk(s) if this rule is not followed. Exceptions are sometimes made if risk vs. reward metrics are accepted by the organization.
- Do not create firewall or security rules that allow IACS protocols through the Industrial Zone interface. IACS protocols are defined as those that are used by Distributed Control System (DCS) and Programmable Automation Controllers (PAC) vendors to communicate with controllers, I/O subsystems, human-machine interface (HMI) or computer systems that are used to program or monitor these types of equipment. An example of such a protocol is CIP.
- Where practical, use VLAN segmentation for the IDMZ assets. This policy will help to make it possible for the firewall to inspect traffic between the IDMZ hosts and make it more likely to catch a compromised IDMZ asset.
- Design the IDMZ to limit the number of inbound and outbound connections to help simplify the firewall and security rules. As a rule, IACS assets and their support systems should remain in the Industrial Zone as much as possible.
- Design the IDMZ with the ability to be disconnected from both the Enterprise and Industrial Zones. This could have a major impact on how the Industrial Zone or Enterprise Zone assets are deployed in order to support operations while the IDMZ is disconnected.
- Do not place permanent data stores in the IDMZ. The IDMZ is the buffer network between the Enterprise and Industrial Zone and is used for temporary data replication and services. If the IDMZ is compromised and an organization has placed valuable data stores in the IDMZ, it could affect the operation and compromise the critical data.

Cisco and Rockwell Automation recognize that each organization must determine their own risk tolerance as they design the IDMZ. The risk vs. financial investment will most likely have some impact on the technologies and architectures that are ultimately chosen for implementation. The best practices listed in this document are meant to provide solution examples that have been tested within the IDMZ.

CPwE IDMZ Security Policy Exceptions

As previously noted, the recommendation is to disallow direct communications between the Enterprise and Industrial Zones. Certain technologies, however, are not designed to be proxied through a demilitarized zone. Situations also exist where a customer makes a reasonable design decision that allows for more risk acceptance in order to trade for better performance or lowered cost of implementation and total life cycle cost.

The tested CPwE IDMZ architecture took security policy exceptions for the following systems and the rationale for each (see Table 2-1). These technologies are reviewed in more detail later in the document.

Note In some cases, an addition of an asset in the IDMZ may help to avoid direct communication through the IDMZ. However, these solutions have not been validated as part of CPwE IDMZ.

Table 2-1 IDMZ Security Policy Exceptions

Asset or Technology	Rationale for Exception	Can Additional Assets be Placed in IDMZ?
Domain Controller Replication	Transport and application Layer security; End-to-end encrypted communication; Total cost of ownership	Yes—Additional DC located in IDMZ that would synchronize with the Enterprise and Industrial Zone DC
Identity Services Engine policy Synchronization and Logging	Can use company-wide distributed ISE deployment with Policy Administration Node (PAN) in the Enterprise Zone; Total Cost of Ownership	Yes—ISE Policy Service Node (PSN) located in the IDMZ
NTP Time Synchronization	Better accuracy by direct connection of Industrial NTP server to Enterprise NTP; NTP traffic can be authenticated by servers.	Yes-IDMZ NTP server could synchronize time with the Industrial Zone NTP server.

IDMZ Data Flow Example

One of the results of developing an IDMZ security policy is a set of requirements for the network services and application data flow through the IDMZ. Figure 2-5 shows a high-level overview of what applications and protocols may have to be allowed through the IDMZ firewalls. As discussed in the previous section, certain network services may be allowed to communicate directly while IACS applications use IDMZ assets to exchange data.

Note The applications and services shown here have been validated as part of the CPwE IDMZ solution and discussed in more details later in the document. The requirements for a particular IACS network may differ depending on business needs, security policies and existing infrastructure.

Figure 2-5 IDMZ Data Flow Example

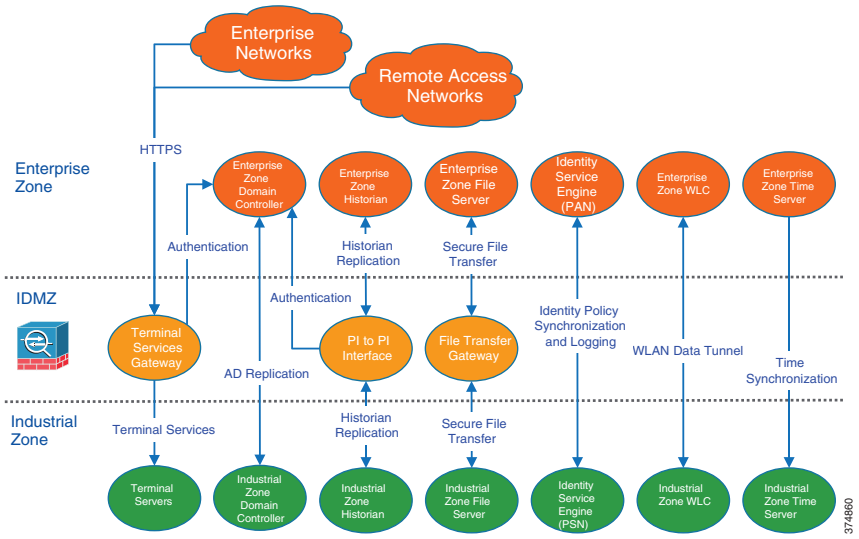

Industrial Zone Security Policy

The convergence of plant-wide and enterprise networks provides greater access to IACS data, which allows manufacturers to make more informed real-time business decisions. This business agility provides a competitive edge for manufacturers who embrace convergence. Convergence also calls for evolved security policies for IACS networks, which no longer remain isolated within a plant-wide or site-wide area. IACS computing and controller assets have become susceptible to the same security vulnerabilities as their enterprise counterparts. A security policy needs to protect IACS assets. This security policy needs to balance requirements such as 24x7 operations, low Mean Time to Repair (MTTR) and high overall equipment effectiveness (OEE).

Securing IACS assets requires a comprehensive security model based on a well-defined set of security policies. Policies should identify both security risks and potential mitigation techniques to address these risks. Manufacturers also face an unclear demarcation line of network ownership and cultural differences between deploying enterprise and IACS assets. To address these obstacles, Cisco and Rockwell Automation recommend that manufacturers develop a IACS security policy, distinct from the enterprise security policy, based on the following considerations:

- Plant-wide or site-wide operation requirements
- Enterprise security policy best practices
- Risk assessment results
- A holistic security policy based on the defense-in-depth approach
- Industry security standards such as ISA-99/IEC-62443
- Manufacturers' corporate standards
- Deploying Network Security within a Converged Plantwide Ethernet Architecture located at:
  - *Deploying Network Security within a CPwE Architecture*: https://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td019_-en-p.pdf
  - *Deploying Network Security within a Converged Plantwide Ethernet Architecture*: https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/5-1/Network_Security/WP/CPwE-5-1-NetworkSecurity-WP/CPwE-5-1-NetworkSecurity-WP.html
- A rigorous and well-documented patch management process

IACS Operations Security

This section outlines some recommended best Operations Security (OpSec) practices for the IACS and IDMZ.

- OpSec encompasses all networks, security systems, computer systems, applications and control systems that are required to support the production of a product or service and the duty to keep all of these systems running securely. While this is a broad topic, the intent of OpSec is to help confirm that people and services, such as computer applications, have the proper rights and privileges to the resources they require to do their job while preventing access to resources that they are not entitled to use.
- OpSec involves the use of technical controls like firewalls, access control lists, anti-virus, anti-malware, allow and deny listing technologies to name a few. It also involves using non-technical controls like policies and procedures to recommend or enforce behaviors with business assets or guiding business conduct.
- OpSec is often synonymous with protecting data that is considered proprietary or that contains intellectual property by means of technical and non-technical controls.

Defining Roles

Every person within the organization should be assigned a role so that consistent security credentials can be assigned to every organizational member. Role-based access control (RBAC) is prevalently used within companies and is leveraged throughout the plant-wide or site-wide systems. As a best practice, it is recommended that RBAC be used within the IACS environment to manage user authentication and authorization of all IACS assets.

While defining roles, one must also consider the concept of Least Privilege and Need to Know:

- **Least Privilege** means an individual should have enough permissions and rights to fulfill their role in the company and no more.
- **Need to Know** describes the restriction of access to the sensitive data. Under need-to-know restrictions, even if one has all the necessary official approvals (such as a security clearance) to access certain information, one would not be given access to such information unless one has a specific need to know; that is, access to the information must be necessary for the conduct of one's duties.

Separation of Duties

Separation of Duties (SoD) is a term used to describe when more than one person is required to complete a task and is used as a method for discovering or preventing fraud. SoD at many organizations is implemented within the enterprise context in a fuller and more robust manner, but often lacks granularity or roles and responsibilities within the IACS systems. For instance, one might see well-defined roles within enterprise descriptions but larger categories within IACS.

Table 2-2 shows an example of user groups in an organization. In this case, defining who can create accounts and give users their roles is an example of SoD.

Table 2-2 User Role Examples

Organizational Role	Core Responsibilities	Account Creation
IT Domain Controller Admin	Configure and maintain corporate domain controllers	Yes
Enterprise Database Admin	Create new database tables and SQL Queries Maintain database	No
Network Admin	Installs and maintains WAN/LAN equipment	Yes—Infrastructure only
Security Admin	Defines, configures and maintains security systems	No
Production Admins	Defines, configures and maintains Industrial Zone software assets that contain common enterprise software such as anti-virus and OS patches	Yes—IACS only
Engineers	Defines, configures and maintains Industrial Zone assets related directly to production systems	No
Maintenance	Maintains Industrial Zone assets related directly to production systems	No
Operators	Monitors production equipment to support the IACS process	No
Trusted Partner	A non-employee resource that is working for the company that needs access to certain assets	No

It is important to define all the roles within the organization and then define what each role is authorized to do with each system and application. Documenting these roles and responsibilities can help identify possible issues such as the conflict of a user being able to change their own security role or conflicting organizational reporting relationships such as the Security Administrator reporting to the Network Administrator which have conflicting business goals.

Data Classification

Data classification helps identify the value of the data to the organization so sensitive data can be organized and protected according to its sensitivity of theft, loss or unavailability. Data classified by the levels of confidentiality, integrity and availability attributes allows security administrators to determine the value to their organization and choose the appropriate controls necessary to protect the data.

Data classification has traditionally started at the Level 3 Site Operations without much regard to classifying the data at lower levels. Most security professionals are familiar with traditional controls to protect data while in motion and at rest in a traditional host. However, technology advances within modern Programmable Automation Controllers (PACs) like the ControlLogix® family make it possible to apply data classifications methods at the Cell/Area Zone level. Implementing FactoryTalk Security within the controller gives the ability to control **who** can do **what** to **which** controller and is also capable of restricting access to data. This type of functionality makes it possible to limit data tampering by allowing the PACs to participate in data security.

Network Redundancy and Availability

Networks are built out of numerous hardware and software components that may fail or that may be subject to attacks. Implementing redundant designs helps eliminate single points of failure (SPOF), which improves the availability of the network and makes it more resistant to attacks.

The CPwE architecture is built with many options for redundancy:

- Backup and redundant uplink interfaces
- Network hardware redundancy:
  - Redundant stackable distribution switches
  - Active/standby and active/active failover in the distribution and core layer
  - Firewall redundancy
- Topological redundancy: designs built with redundant paths at both network and data link layers

Typically, redundant network and control strategies are applied to operationally critical processes where a business determines that hardware failure or loss of visibility into the process cannot be tolerated. It is also recognized that non-critical processes that do not have this same high availability requirement exist and therefore the network and control architectures will not be designed in the same fashion as critical processes.

As a best practice, it is recommended that each business determine the types of processes within each Cell/Area Zone and classify the availability requirements. This type of classification exercise will determine the availability requirements within each Cell/Area Zone and drive the network requirements. Once this exercise is complete, one should design and test modular network architectures to support each availability requirement.

Network Infrastructure Hardening

This section reviews some of the best practices for securing IACS network infrastructure.

Note More information about network infrastructure hardening can be found in the following documents:

- *Cisco Guide to Harden Cisco IOS Devices*
  - http://www.cisco.com/c/en/us/support/docs/ip/access-lists/13608-21.html
- *Configuring Switch-Based Authentication*

– http://www.cisco.com/c/en/us/td/docs/switches/lan/cisco_ie2000/software/release/15_2_2_e/configuration/guide/scg-ie2000/swauthen.html

Disabling Unnecessary Services

Switches come out of the box with a list of services turned on that are considered appropriate for most network environments. Disabling these unnecessary services has two benefits: it helps preserve system resources and it helps to eliminate the potential of security exploits on the disabled services.

Cisco and Rockwell Automation recommend the following best practices:

- Global services should be disabled on all routers and switches unless explicitly needed. Note that some of the services are enabled by default (BOOTP, IP source routing, and PAD). Other global services such as finger, identification (identd), and TCP and UPD small servers are disabled by default.
- IP-directed broadcast should remain disabled on all Layer 3 IP interfaces except those required for access by RSLinx® data servers to browse for known or available IACS EtherNet/IP devices on a different subnet.
- Cisco Discovery Protocol (CDP) should be disabled on interfaces where the service may represent a risk; for example, on external interfaces such as those at the Internet edge and ports that connect end devices.
- Services on access ports such as MOP, IP redirects and Proxy ARP should be disabled unless required.
- The Stratix 5800 meets IEC 62443 4-2 Security Level 2 when properly configured to comply with the certification requirements. Table 2-3 provides which switch security features are required to be configured to meet the IEC 62443 4-2 certification requirements.

Table 2-3 Switch Security Features

Switch Security Feature	Required to Meet IEC 62443-2	Details
IOS Release is certified for IEC-62443 4-2	Yes	To verify if your IOS release is certified for IEC-62443 4-2, access product release notes from the Product Compatibility and Download Center at rok.auto/pcdc
Configure Certificate Authority (CA)	Yes	A CA provides a chain of trusts for devices in the network. This mechanism provides the ability for a user or process to trust the connection to one of these devices on the network by validating its identity. For more information, see the Security Configuration Guide at: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_pki/configuration/xe-17/sec-pki-xe-17-book/sec-pki-overview.html
Configure Authentication, Authorization, and Accounting (AAA)	Yes	AAA services provide flexible administrative control and accounting for network access. For more information, see the Security Configuration Guide at: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/17-4/configuration_guide/sec/b_174_sec_9300_cg.html
Disable Telnet	Yes	Telnet is disabled by default during Express Setup. Keep Telnet disabled to secure remote access to the switch, such as when you are using the command-line interface (CLI) to manage the switch from a computer.
Transport Layer Security (TLS) 1.2	Yes	TLS 1.2 is enabled by default during Express Setup. Keep this feature enabled to secure the exchange of data through encryption.
Configuration of Type 9 password hashing	Yes	Hashing makes password storage more secure by transforming a password into data that cannot be converted back to the original password. For more information, see the User Security Configuration Guide att: https://www.cisco.com/c/en/us/support/docs/ip/access-lists/13608-21.html

Applying Port Security

Access to the network starts with physically accessing ports on switches. A number of techniques to limit the ability to access the network exist.

First, network access cannot be achieved if the network devices are physically secure with limited access. Placing the industrial Ethernet switches in locked rooms or cabinets and installing port locks to prevent access to unused ports on a switch are all recommended best practices by Cisco and Rockwell Automation.

Further, industrial Ethernet switches themselves can be configured to secure their ports from unknown access or misuse. Switch port security limits the access to the network by unknown devices and limits the number of devices or media access control (MAC) addresses on any network port. Port security builds a list of secure MAC addresses in one of the two following ways, configurable on a per-interface basis:

- Dynamic learning of MAC addresses, which defines the maximum number of MAC addresses that will be learned and permitted on a port, is useful for dynamic environments, such as at the access edge
- Static configuration of MAC addresses, which defines the static MAC addresses permitted on a port, is useful for static environments such as a server farm or a DMZ

Note Although some implementers may consider static MAC address configurations per port for environments that need very high security, this method requires significant effort and network expertise to perform normal maintenance tasks such as replacing a failed device.

The Error Disable feature helps protect the switch and therefore the network from certain error conditions. For example, when the number of MAC addresses on a port is exceeded. When the error condition is discovered, the interface is put into the error disable state and does not pass traffic.

Cisco and Rockwell Automation recommend the following:

- Use dynamic learning to limit the number devices that can access a port. This allows, for example, only one MAC address to access an IACS network port on the industrial Ethernet switch.
- Apply the *errdisable recovery interval seconds* global configuration command to restore the port state. This command will periodically check to see if the error condition still exists on the interface. The interface will be enabled automatically when the error condition has cleared.
- Disable all unused ports on a switch and only enable them when required.

Securing Administrative Access

When considering the security of a network infrastructure, it is critical that the administrative access to network devices is protected. Cisco and Rockwell Automation recommend the following best practices:

- Set and protect local passwords:
  - Enable automatic password encryption
  - Define a strong local *enable* password using the *enable secret* global command
  - Configure local user accounts (individual usernames and passwords) on devices for administrative access, as opposed to a single password for every user
- For highly secure IACS networks, configure devices for Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System Plus (TACACS+) authentication against the centralized user database using a remote AAA server (for example, Cisco ISE). Use accounting features of the AAA to log access and configuration changes

Note Local user authentication should be used as a backup in case the remote AAA server is not available.

- Implement legal notification banners that are presented on all interactive sessions to confirm that users are notified of the security policy being enforced and to which they are subject
- Use Secure Shell (SSH) protocol when available rather than the unsecured Telnet. Use at a minimum 1024-bit modulus size. The SSH feature requires configuring AAA or local accounts on a device.
- If possible, use HTTPS for device management instead of clear-text HTTP
- If Simple Network Management Protocol (SNMP) is used for device management, use only SNMP v3 for read-write access. Configure the maximum security level using authentication and encrypted communication (authPriv). If SNMP v3 is not supported, use SNMP v1 or v2 for read-only access.
- Explicitly define the protocols allowed for incoming and outgoing sessions on the device. Restricting outgoing sessions prevents the system from being used as a staging host for other attacks.
- Configure access control lists (ACL) to restrict management traffic to the device. For example, an ACL can be configured to allow SNMP traffic only from the designated management servers.
- Set idle and session timeouts for remote access. Enable TCP keepalives to detect and close hung sessions.
- Protect switch configuration files and store them in a secure location. When sending the files externally (for example, to technical support), remove critical information such as user credentials, passwords and secret keys from the files (even if encrypted).

Note SSH, HTTPS and SNMP v3 require the cryptographic (K9) version of IOS on the Cisco IE Series and Allen-Bradley® Stratix industrial Ethernet switches. Some cryptographic features are subject to additional export and contract restrictions. For more information about the Cisco products, see *Export and Contract Compliance* at:

- http://www.cisco.com/web/about/doing_business/legal/global_export_trade/general_export/contract_compliance.html.

Contact your Rockwell Automation sales representative or distributor for details about Stratix products.

Computer Hardening

For computing assets within the Industrial Zone, implement IT best practices applied to Enterprise computers. Some best practices and general recommendations include the following:

- Secure physical access. Network equipment and servers should be in locked cabinets or rooms.
- Keep computers up-to-date on service packs and hot fixes, but disable automatic updates. Also, network developers should test patches before implementing them and schedule patching and regular network maintenance during operational downtime.
- Apply Microsoft updates that have been qualified by Rockwell Automation to computers running Rockwell Automation® software products. Before implementing qualified updates, verify them on a non-production system, or when the facility is non-active, to avoid unexpected results or side effects.
- Deploy and maintain anti-virus and antispyware software, but disable automatic updates and automatic scanning. Test definition updates before implementing them and schedule manually-initiated scanning during operational downtime since antispyware scanning can disrupt real-time operations. Automatic anti-virus and antispyware scanning has caused data loss and downtime at some IACS facilities.

- Prohibit direct Internet access. IACS assets should not have direct line of sight to the Internet. Any necessary communication (for example, firmware, patches and anti-virus updates) should be accomplished via the IDMZ proxy and application servers.
- Implement the best practice password policy such as enforcing password history, maximum password age, minimum password length and complex password requirements.
- Disable the guest account on clients and servers.
- Require that the built-in administrator account uses a complex password, has been renamed and has removed its default account description.
- Develop and deploy backup and disaster recovery policies and procedures. Test backups on a regular schedule.
- Implement a change management system to archive network, controller and computer assets (clients, servers and applications).
- Use Control+Alt+Delete, along with a unique user name and password to log in.
- Protect unnecessary or infrequently used USB ports, parallel and serial interfaces to prevent unauthorized hardware additions (modems, printers, USB devices, etc.).
- Uninstall the unused Microsoft Windows components, protocols and services not necessary to operate the plant-wide system.
- Install and run only legitimately purchased software.

Assessments and Baselining

Baselines are representative of a singular point in time in which a company can reference for future changes. Typically, the assessment process is used to obtain a baseline of:

- Computer systems
- Infrastructure components like firewalls, routers and switches
- Network traffic types, quantity and data flow diagrams
- Security control assessments
- Hardware, firmware and software inventories

Baselines are also used to define the minimum levels of security controls that should be implemented to adhere to an organization’s standards.

Cisco and Rockwell Automation recommend implementing consistent standards for assessment methodology by generating assessment and baselining policies. These activities should be performed periodically with a method to record the improvement or failure of the security program execution.

IDMZ Network Infrastructure Design

This section describes IDMZ CPwE design recommendations for the following network infrastructure components and protocols:

- Industrial Zone firewalls
- Industrial Zone core switches
- IDMZ server network
- Routing protocols between zones

Industrial Zone Firewalls

The industrial firewalls are an essential aspect of protecting the IACS network and applications. The combination of firewalls and an IDMZ zone concept are key aspects of the defense-in-depth approach for IACS network security. An Industrial Zone firewall provides the following functions:

- Implements an IDMZ where data and services between the Enterprise and Industrial Zones can be securely shared
- Establishes traffic patterns between the network zones via assigned security levels and access rules
- Provides stateful packet inspection of all traffic between the various zones
- Provides Intrusion Prevention Services (IPS) and Deep Packet Inspection (DPI) capabilities for inspecting application data between the zones designed to identify and potentially stop a variety of attacks
- Allows remote access to the IACS network for authenticated and authorized users

The following sections provide an overview of Cisco Firepower Threat Defense firewall platform and recommendations for deploying it as part of the CPwE solution.

Industrial Firewall Functionality

Cisco Firepower Threat Defense (FTD) brings distinctive threat-focused next-generation security services to the industrial network. It provides comprehensive protection from known and advanced threats, including protection against targeted and persistent malware attacks. Cisco Secure Firewall (Cisco Firepower Threat Defense) includes:

- Site-to-site and remote access VPN and advanced clustering provide highly secure, high-performance access and high availability to help achieve business continuity.
- Granular AVC supports more than 3,000 application-layer and risk-based controls that can launch tailored IPS threat detection policies to optimize security effectiveness.
- The industry-leading Cisco FTD provides highly effective threat prevention and full contextual awareness of users, infrastructure, applications, and content to detect multi-vector threats and automate defense response.
- Reputation- and category-based URL filtering offer comprehensive alerting and control over suspicious web traffic and enforce policies on hundreds of millions of URLs in more than 80 categories.
- Advanced Malware Protection (AMP) provides industry-leading breach detection effectiveness, a low total cost of ownership, and superior protection value that helps you discover, understand, and stop malware and emerging threats missed by other security layers

Firewall Resiliency

Configuring high availability, also called failover, requires two identical FTD devices connected to each other through a dedicated failover link and, optionally, a state link. FTD supports Active/Standby failover, where one unit is the active unit and passes traffic. The standby unit does not actively pass traffic, but synchronizes configuration and other state information from the active unit. When a failover occurs, the active unit fails over to the standby unit, which then becomes active.

The health of the active unit (hardware, interfaces, software, and environmental status) is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

Note More information about FTD resiliency features can be found in *Information About High Availability* at:

- https://www.cisco.com/c/en/us/td/docs/security/firepower/70/configuration/guide/fpmc-config-guide-v70/high_availability_for_firepower_threat_defense.html

---

Failover System Requirements

This section describes the hardware and software requirements for FTDs in a failover configuration.

The two units in a failover configuration must meet these hardware requirements:

- Be the same model
- Have the same number and types of interfaces
- Have the same modules installed (if any)
- Have the same amount of RAM installed

If you are using units with different flash memory sizes in your failover configuration, make sure the unit with the smaller flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger flash memory to the unit with the smaller flash memory will fail.

The two units in a failover configuration must meet these software requirements:

- Be in the same firewall mode (routed or transparent)
- Have the same software version
- Be in the same group or domain in Firepower Management Center (FMC)
- Have the same NTP configuration
- Be fully deployed in FMC with n uncommitted changed
- Not have DHCP or PPPoE configured on any interface
- (Firepower 4100/9300) Have the same offload mode, either both enabled or both disabled

Failover Link

The two FTD units in a failover pair constantly communicate over a failover link to determine the operating status of each unit. The following information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keepalives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

You can use any unused interface on the device as the failover link, however, you cannot specify an interface that is currently configured with a name. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface should only be used for the failover link (and optionally for the Stateful Failover link).

The failover link can be connected in one of the following two ways:

- Using a switch, with no other device on the same network segment as the failover interfaces of the FTD device.
- Using an Ethernet cable to connect the appliances directly, without the need for an external switch

Stateful Failover Link

To use Stateful Failover, firewalls must have a Stateful Failover link to pass all connection state information. Three options exist for selecting an interface for a Stateful Failover link:

- A dedicated Stateful Failover interface
- Sharing an interface with the failover link

Sharing a failover link is the best way to conserve interfaces. However, if you have a large configuration and a high traffic network, you must consider a dedicated interface for the state link and failover link.

Note By default, all information is sent in clear text over the failover and Stateful Failover links. For additional security, the failover communication can be encrypted using a failover key.

If you use the failover link as the Stateful Failover link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the Stateful Failover interface.

Avoiding Interrupted Failover and Data Links

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the FTD device can use the data interfaces to determine if a failover is required. Subsequently, the failover operation is suspended until the health of the failover link is restored.

See the following connection scenarios to design a resilient failover network.

If a single switch or a set of switches are used to connect both failover and data interfaces between two FTD devices, then when a switch or inter-switch-link is down, both FTD devices become active. Therefore, the two connection methods shown in Figure 2-6 and Figure 2-7 are NOT recommended.

Figure 2-6 Connecting with a Single Switch—Not Recommended

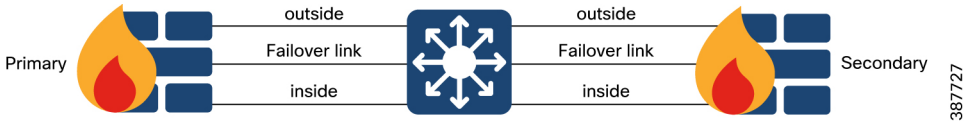

We recommend that failover links not use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in Figure 2-7.

Figure 2-7 Connecting with a Double Switch—Not Recommended

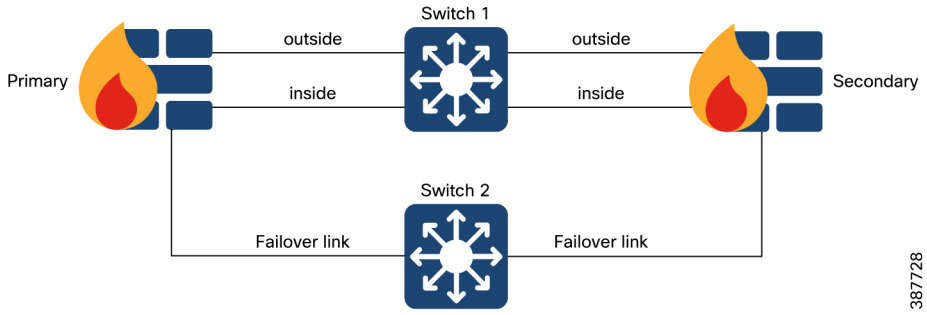

The most reliable failover configurations use a redundant interface on the failover link, as shown in Figure 2-8.

Figure 2-8 Connecting with Redundant Interfaces—Recommended

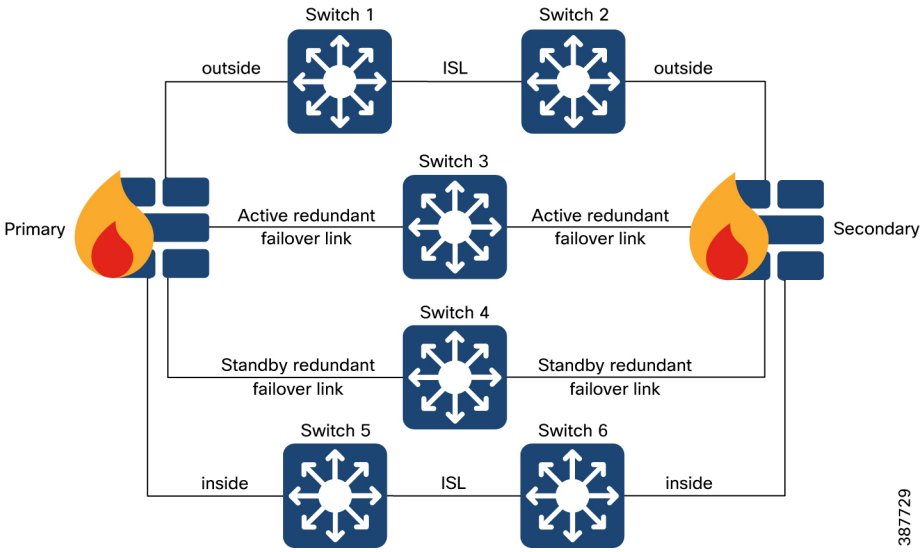

Firewall Policy Design

IDMZ firewall is positioned between the Industrial Zone and the Enterprise Zone which follows the IDMZ security policy as described previously. The firewall design is primarily based on what application traffic needs to be permitted or denied, and what hosts can originate application connections that are allowed through the firewall.

The recommended and usual practice is to implement a restrictive policy on a firewall:

- Deny any service unless it is expressly permitted
- Restrict who is allowed to communicate to the necessary minimum

Access control is a hierarchical policy-based feature that allows you to specify, inspect, and log (non-fast-pathed) network traffic.

Each managed device can be targeted by one access control policy. The network traffic data collected by the policy target devices can be used to filter and control that traffic based on:

- Simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- The latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- Realm, user, user group, or ISE attribute
- Custom Security Group Tag (SGT)
- Characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- Whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- Time and day (on supported devices)

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line defense.

Access Control Policy Default Action

A newly created access control policy directs its target devices to handle all traffic using its default action. In a simple access control policy, the default action specifies how target devices handle all traffic. In a more complex policy, the default action handles traffic that:

- Is not trusted by Intelligent Application Bypass
- Is not on a Security Intelligence Block list
- Is not blocked by SSL inspection (encrypted traffic only)
- Matches none of the rules in the policy (except Monitor rules, which match and log-but do not handle or inspect-traffic)

The access control policy default action can block or trust traffic without further inspection, or inspect traffic for intrusions and discovery data.

Deep Inspection Using File and Intrusion Policies

Deep inspection uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

- Intrusion policies govern the system's intrusion prevention capabilities.

  For complete information, see An Overview of Intrusion Detection and Prevention at: https://www.cisco.com/c/en/us/td/docs/security/firepower/70/configuration/guide/fpmc-config-guide-v70/overview_of_network_analysis_and_intrusion_policies.html

- File policies govern the system's file control and AMP for Networks capabilities.

  For complete information, see File Policies and Malware Protection at:

  https://www.cisco.com/c/en/us/td/docs/security/firepower/70/configuration/guide/fpmc-config-guide-v70/file_policies_and_advanced_malware_protection.html

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique pair of intrusion policy and variable set counts as one policy.

The creation and deployment of access control policies in FTD will be explored further in the document in relation to the use cases that are being secured.

Industrial Zone Core Network

The Industrial Zone core is the critical part of the plant network that is designed to be highly available and operate in an always-on mode. The core serves as the aggregator for all of the Cell/Area Zones and provides connectivity between end-devices, server-based applications and data storage. The Industrial Zone core connects via firewalls to the IDMZ.

The key design objectives for the core are:

- Provide the appropriate level of redundancy to allow for near immediate data-flow recovery in the event of any component (switch, supervisor, line card, or fiber) failure
- Permit the necessary hardware and software upgrade/change to be made without disrupting any network applications
- Avoid implementing any complex policy services in the core and have the minimal control plane configuration
- Do not have any directly attached user/server connections

In small-to-medium plants, it is possible to collapse the core into the two redundant distribution switches. However, for large plants, where a large number of Cell/Area Zones exist, this level of hierarchical segmentation is recommended.

Core Switch Architecture

The core switch architecture should meet the design requirements listed above to provide the required level of resiliency and performance. Large architectures normally use modular chassis-based core switches, such as Cisco Catalyst 4500/9300 or 9500 platforms.

Note For more information on the Industrial Zone design and topology options, refer to Deploying a Resilient Converged Plantwide Ethernet Architecture Design and Implementation Guide found at:

Rockwell Automation site:

- https://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td010_-en-p.pdf

Cisco site:

- https://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/4-0/Resiliency/DIG/CPwE_resil_CVD.html

Connection to Redundant Firewalls

The IDMZ CPwE architecture recommends configuring redundant industrial firewalls in the active/standby mode with EtherChannels to the core switches. When a Virtual Switching System (VSS) is used, the FTD interfaces within the same EtherChannel can be connected to separate switches in the VSS.

Note Separate EtherChannels should be created on the VSS switches for each FTD in an active/standby failover deployment (see Figure 2-9). A single EtherChannel on the VSS switch pair will not be established because of the separate FTD system IDs, and would not be desirable anyway because traffic should be sent only to the active FTD.

Figure 2-9 VSS and Active/Standby Firewalls

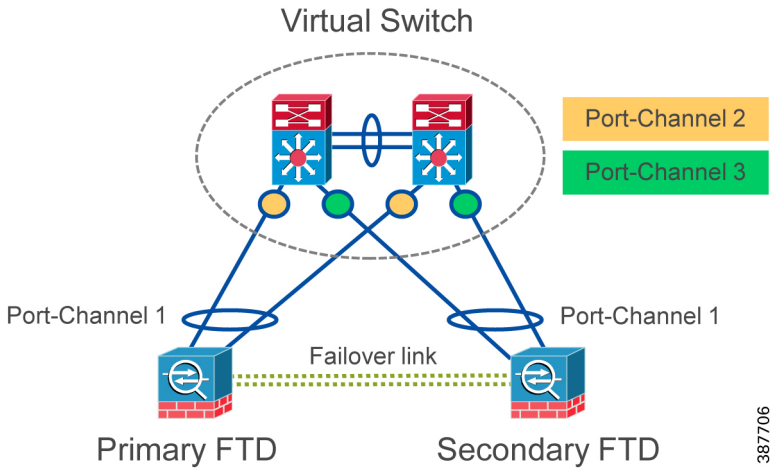

IDMZ Server Network

The IDMZ network hosts services that facilitate communication between the Enterprise and Industrial Zones, including RD Gateway for secure remote access via Remote Desktop Connection client and ThinManager, file transfer gateway, Historian connector (PI-to-PI Interface), and anti-virus and OS patch servers.

The design of the IDMZ server network will depend on the server farm size and IT management requirements and practices. Several scalable options exist:

- A redundant access/distribution switch pair (chassis-based, stack or stand-alone)
- A redundant distribution switch pair connecting multiple access switches in the redundant star topology
- Two or more switch blocks with separate distribution switches, for example to segregate servers into different management domains

The resiliency features should include:

- Redundant server connections to access or distribution switches
- Redundant connections between distribution switches and industrial firewalls

IDMZ VLAN Segmentation

The IDMZ server network should be designed to meet the availability requirements and also designed to support traffic inspection between the hosts. In the example (see Figure 2-10), every IDMZ host such as the RD Gateway or the Secure File Transfer server has been put onto its own network or VLAN.

Figure 2-10 VLAN Segmentation in IDMZ Network

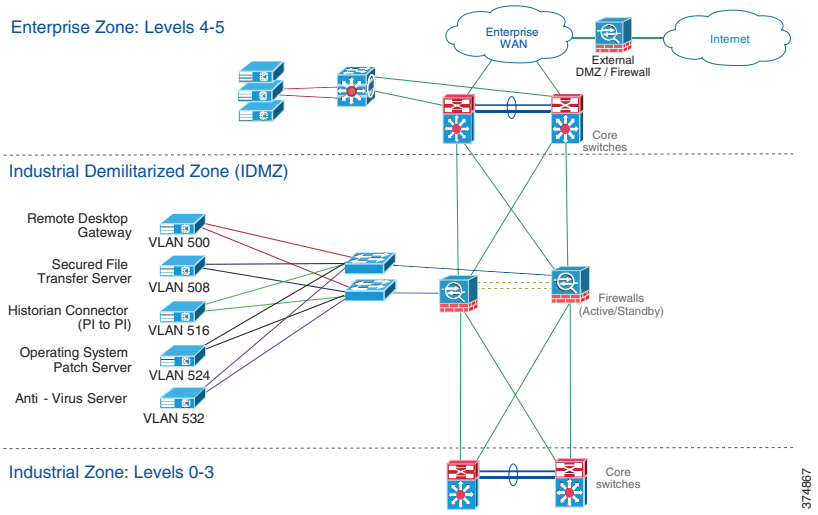

The IDMZ assets were placed on their own VLAN for strategic purposes. A “typical” piece of malware will compromise the asset and often attempt to either communicate outside the local network or it will attempt to infect other hosts on the same or other networks (see Figure 2-11).

If the compromised host attempts to communicate outside or within the IDMZ, a properly configured firewall will block the attempted pivot. If the firewall is configured to send alarm messages to a log or to a system monitor, then this incident can be investigated by the security team.

Figure 2-11 Compromised IDMZ Host

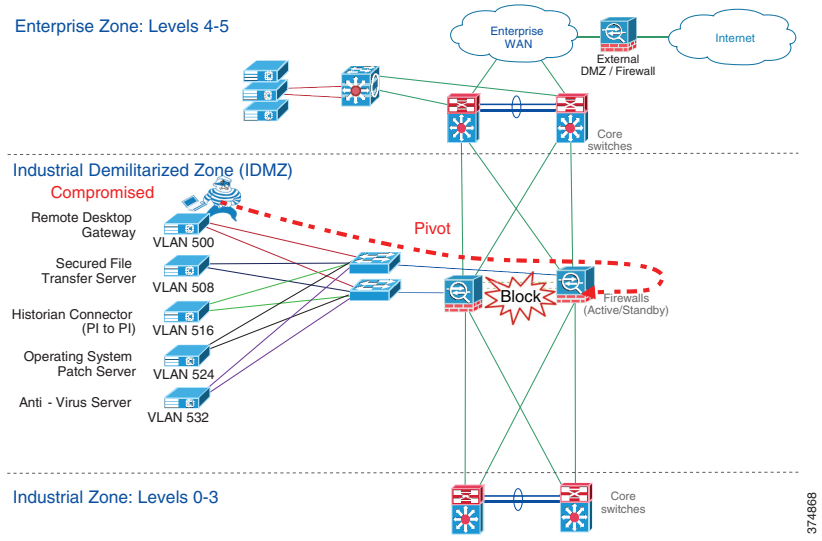

Routing Between Zones

In order to communicate between the Industrial and the Enterprise Zones, network infrastructure devices (Layer 3 switches, routers and firewalls) need to exchange IP subnet information via dynamic routing protocols or to have statically defined routes to destination IP subnets. This section provides recommendations and considerations for the routing protocol selection and design.

EIGRP Overview

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a Cisco proprietary routing protocol. EIGRP is an advanced distance vector routing protocol. The Diffusing Update ALgorithm (DUAL) is used to obtain a loop-free topology during the network convergence. All routers involved in a topology change are able to synchronize at the same time. Routers that are not affected by topology changes do not need to synchronize. The EIGRP convergence time rivals that of any other existing routing protocol.

Some of the many advantages of EIGRP are:

- Very low usage of network resources during normal operation
- Only routing table changes are propagated, and not the entire table, during the convergence
- Rapid convergence times for changes in the network topology

Note More information about EIGRP can be found in *Enhanced Interior Gateway Routing Protocol* at:

- http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html

EIGRP Design

This section contains EIGRP design options and considerations, such as scalability, routing policy and configuration complexity.

Single EIGRP Domain

A single EIGRP domain is the simplest configuration for the routing protocol. In this design, all routers in both the Enterprise Zone and Industrial Zones participate in a common routing protocol instance, which is defined by the Autonomous System (AS) number (see Figure 2-12). The IDMZ firewalls actively participate in the routing protocol and summarize routes between the Enterprise and Industrial Zones. The firewall advertises a single default route to the Industrial Zone routers. On the enterprise side, it advertises a summary route for the Industrial Zone networks.

Figure 2-12 Single EIGRP Domain

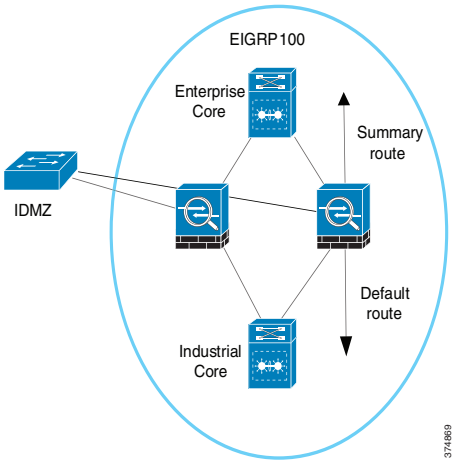

This design, which allows for end-to-end routing from the Industrial Zone to the Enterprise Zone, works best if a single administrative team is responsible for all network devices across the company. Since all routers are a part of a common routing domain, the risk that routing protocol instability in one zone could affect other zones exists. This solution fits best with small-to-medium sized networks.

Note A single EIGRP domain was used during the testing of this CPwE solution.

Multiple EIGRP Domains with Redistribution

In this design, the Enterprise Zone and each Industrial Zone are assigned a unique EIGRP domain. The routers in each zone use their own AS numbers while firewalls are configured for both AS. The IDMZ firewall acts as a boundary between the EIGRP process domains and redistributes routes between the processes (see Figure 2-13). In addition to redistribution, the firewalls also summarize routes advertising a single default route to the Industrial Zone. On the Enterprise side, it advertises summary routes for the Industrial Zone networks. The firewalls can also filter any routes that do not need to be advertised to either the Enterprise or Industrial Zones.

Figure 2-13 Multiple EIGRP Domains with Redistribution

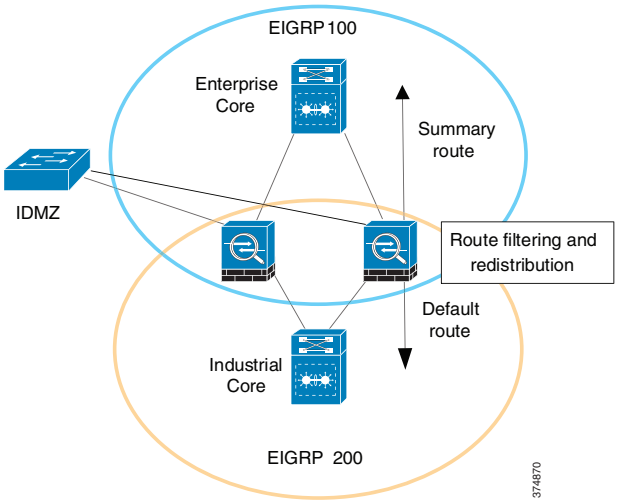

This design also allows for end-to-end routing between the Enterprise and Industrial Zones. However, this design divides the EIGRP routing process into smaller domains. This reduces the possibility of a routing protocol instability in one zone affecting another. This solution fits best with medium-to-large networks and easily accommodates environments with multiple Industrial Zones.

Multiple EIGRP Domains with Static Routes

In this case, similar to the previous design, the Enterprise and each Industrial Zone are assigned a unique EIGRP domain (AS number). The IDMZ firewalls act as a boundary between the EIGRP process domains; however, routes are not redistributed between zones.

Static routes must be configured on the routers connected to the firewalls and the firewalls themselves to forward traffic from the Enterprise Zone to the Industrial Zone and vice versa. The boundary router must also redistribute the static routes back into the routing protocol so the routes are reachable across the zone (see Figure 2-14).

Figure 2-14 Multiple EIGRP Domains with Static Routes

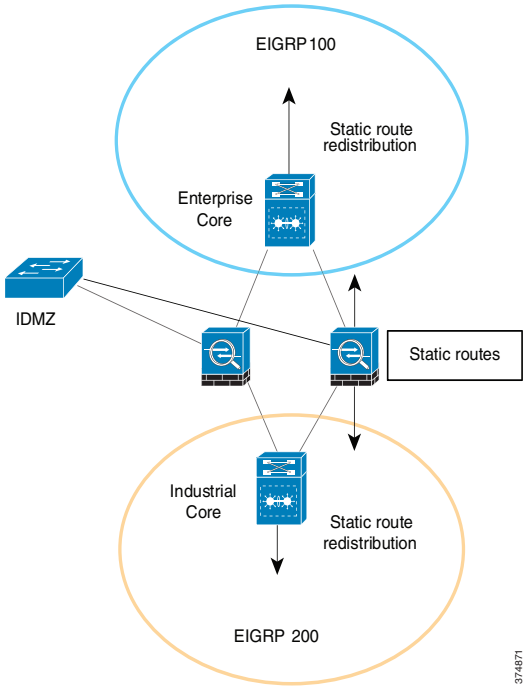

This design allows for end-to-end routing while completely isolating the routing processes in the Enterprise and Industrial Zones. This solution fits best when policy prevents running a routing protocol on the firewall or when the organizational structure has independent teams supporting routing and firewalls.

Protecting EIGRP

All routing protocols must be protected to prevent the distribution of faulty route information. The primary methods of protecting the integrity of the EIGRP routing table are:

- Passive interfaces
- Route authentication

By default, a router running the EIGRP routing protocol will attempt to establish a neighbor relationship with any routers on the local network. The risk that someone could introduce a rogue router and advertise false routes into the network exists. The *passive interface* command prevents the EIGRP process from establishing a neighbor relationship with any routers on the specified interface. This command is commonly used on LAN interfaces connecting to end devices.

The EIGRP protocol also supports route authentication. With route authentication, a shared key is configured on all routers. A router will only accept a route update from a neighbor that signed the update using a MD5 hash that includes the shared key.

OSPF Overview

The Open Shortest Path First (OSPF) routing protocol is an open standard protocol defined in IETF RFC 2328. The OSPF is an Interior Gateway Protocol used to distribute routing information within a single AS. OSPF uses Dijkstra's Shortest Path First algorithm in order to build and calculate the shortest path to all known destinations. OSPF is a link state protocol which means that each router must maintain a database of the state of each routed link in the network.

To reduce the overhead of the protocol, OSPF divides the network into multiple areas. Area 0 is the backbone of the network and all other areas must directly connect to the Area 0 through Area Border Routers (ABR). Dividing the routing protocol into multiple areas reduces the CPU and memory required to maintain the link state database.

Some of the many advantages of OSPF are:

- Open standard defined by the Internet Engineering Task Force (IETF)
- Multi-area design that reduces CPU and memory requirements on individual routers
- Link-state design for fast convergence

Note More information about OSPF can be found in *OSPF Design Guide* at:

- http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html

OSPF Design

This section contains OSPF design options and considerations, such as scalability, routing policy and configuration complexity.

Single OSPF Domain

In the single OSPF domain design, Area 0 is contained in the Enterprise Zone. Each Industrial Zone is a new area in the OSPF network. The IDMZ firewalls function as the ABRs between the corporate backbone (Area 0) and the Industrial Zone area.

The Industrial Zone area should be configured as a Totally Stubby Area to reduce the overhead of routing within the zone. Optionally, the Industrial Zone area can be configured as a Stub Area or a Not-So-Stubby Area (NSSA) depending on the needs of the application. See Figure 2-15.

Figure 2-15 Single OSPF Domain with Stub Areas

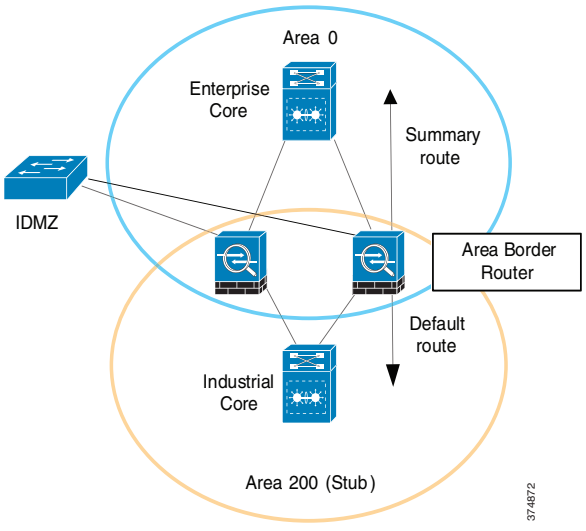

This design also allows for end-to-end routing between the Enterprise and Industrial Zones. OSPF naturally subdivides the routing protocol into areas. This solution fits best with medium-to-large networks and easily accommodates environments with multiple Industrial Zones.

Multiple OSPF Domains

This design treats the Enterprise and Industrial Zones as separate routing domains. Both zones have their own Area 0 backbone network. Because of this, the IDMZ firewalls do not run an instance of OSPF.

Static routes must be configured on the routers connected to the firewalls and the firewalls themselves to allow communications between the zones. The boundary routers must redistribute the static routes into the OSPF instance for the zone. See Figure 2-16.

Figure 2-16 Multiple OSPF Domains with Static Routes

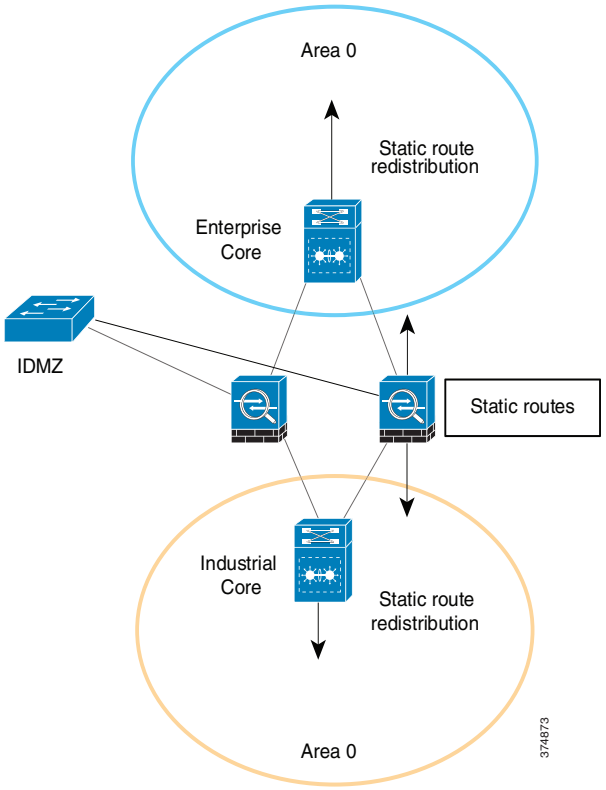

This design allows for end-to-end routing between the Enterprise and Industrial Zones while completely segregating the routing processes in the zones. This solution fits best when policy prevents running a routing protocol on the firewall or when the organizational structure has independent teams supporting routing and firewalls.

Securing OSPF

OSPF supports route authentication to reduce the chance of malicious routes being added to the routing protocol. When route authentication is enabled, the OSPF router signs its route update with a shared key. The neighboring router will only accept route updates that are signed with the correct key.

By default, OSPF sends the authentication information in clear text. It is important to use type 2 authentication which uses an MD5 hash for authentication.

Selecting Routing Design

Table 2-4 summarizes features of the different design options for EIGRP and OSPF protocols. The main criteria for selecting the appropriate design should be the network size, configuration complexity and IT policies. Often, the existing enterprise network design determines the routing configuration in the Industrial Zone.

Table 2-4 Routing Design Selection Criteria

Routing Design	Network Size			Network Policy			Complexity		
	Small	Medium	Large	Routing Protocol on Firewall	Single Administrative Domain	Subdivides Routing Processes	Low	Medium	High
Single EIGRP domain	X	X		X	X		X		
Multiple EIGRP domains (redistribution)		X	X	X		X		X	
Multiple EIGRP domains (static routes)		X	X			X			X
Single OSPF domain		X	X	X		X		X	
Multiple OSPF domains		X	X			X			X

IDMZ Design for Network Services

In a converged IACS network, Industrial and Enterprise Zones can share certain network services to reduce cost of deployment and support and to be able to use same IT management resources. From a security perspective, user authentication and authorization policies have to be managed and applied throughout the whole infrastructure.

These services use network protocols to enable data replication and configuration synchronization across the IDMZ. Design considerations for the following network services are reviewed in this section:

- Active Directory Services
- Certificate Services
- Network Time Protocol (NTP)
- Identity Services
- Multi-Factor Authentication (MFA)
- Licensing
- Windows Updates

Active Directory Services

Microsoft Active Directory (AD) services play an essential role in managing, authenticating and authorizing users and network assets in an enterprise. Companies need a central repository of information about people and their access rights that apply to both the Industrial and Enterprise Zones. AD services in the Industrial Zone should be designed to allow secure replication of information across the IDMZ while being able to operate independently if necessary.

The following sections describe AD and provide design recommendations for the CPwE IDMZ.

Active Directory Overview

Active Directory Domain Services (AD DS) provides a distributed database of information about network resources and application data. AD DS organize network elements, such as users, computers and other devices, into a hierarchical structure that includes the Active Directory forest, domains in the forest, and organizational units (OUs) in each domain. A server that is running AD DS is called a Domain Controller (DC).

- A **forest** acts as a security boundary for an organization and defines the scope of authority for administrators. By default, a forest contains a single domain, which is known as the forest root domain.
- A **domain** is a logical group of network objects (computers, users, devices) that share the same AD database. An AD domain supports a number of core functions including network-wide user identity, authentication, and trust relationships.

  Additional domains can be created in the forest to provide partitioning of AD DS data. Multiple domain structure can be used to control data replication and to scale globally over a network with limited bandwidth.
- **OUs** simplify the management of large numbers of objects by the delegation of full or limited authority to other users or groups. OUs are used more often than domains to provide structure and to simplify the implementation of policies and administration.

AD DS implements security with a logon authentication and access control to resources in the directory. Authorized network users and administrators can use a single network logon to access resources anywhere in the network. Policy-based administration allows simplifying management of even the most complex network.

Additional AD DS features include the following:

- A **schema** is the set of rules that defines the classes of objects and attributes that are contained in the directory, the name format and the constraints for these objects.
- A **global catalog** that contains information about every object in the directory. Users and administrators can use the global catalog to find directory information, regardless of which domain in the directory actually contains the data.
- A **replication service** that distributes directory data across a network. Any change to directory data is replicated to all DCs in the domain.
- **Operations master roles** (also known as Flexible Single Master Operations or FSMO) on designated DCs to perform specific tasks to confirm consistency and eliminate conflicting entries in the directory.
- **Active Directory Federation Services (AD FS)** can be deployed to manage access to protected resources for trusted partners including external third parties or other departments or subsidiaries in the same organization.

Note For information about AD DS, refer to *Active Directory Domain Services* at:

- https://technet.microsoft.com/en-us/windowsserver/dd448614

Active Directory Architecture in IDMZ

This section provides design recommendations for the AD DS in the CPwE IDMZ architecture.

AD DS Deployment Model

The tested and validated deployment of the AD DS in the CPwE architecture is based on the AD implementation in a single domain with multiple sites.

A single AD domain for the Enterprise and Industrial Zones allows maintaining a single identity and access policy repository for all employees in a company. This approach can bring many benefits for the CPwE architecture, for example, secure remote access to the Industrial Zone from the enterprise.

The first domain controller you install automatically creates the first site, known as the Default-First-Site-Name. After installing the first domain controller, all additional domain controllers are automatically added to the same site as the original domain controller. The Enterprise Zone and IDMZ can be part of the Default-First site.

To deploy the CPwE architecture topology, the addition of an Active Directory Domain Controller (AD DC) in the Industrial Zone is required. The Industrial Zone is placed in its own AD site. Establishing separate sites for the Industrial and Enterprise Zones provides the following benefits:

- Efficient use of bandwidth for replication in case of WAN connectivity
- Detailed control of replication behavior, for example schedule
- Industrial assets can authenticate to the local DC

Note AD DS should be installed in accordance with Microsoft best practices and deployment guidelines provided in *Deploy Active Directory Domain Services (AD DS) in Your Enterprise* at:

- https://technet.microsoft.com/en-us/library/hh472160.aspx

Active Directory Replication

The CPwE IDMZ architecture for AD implements bi-directional replication between the Enterprise DC and the Industrial Zone DC. An AD administrator should be able to create, delete and update accounts in the Industrial Zone and the changes will be replicated to the Enterprise Zone and vice versa.

Companies may also choose one-directional replication (Enterprise DC to Industrial DC only) due to security policies and management practices.

Site-to-site replication data can be compressed and sent on a schedule, depending on the available network bandwidth and requirements. The synchronous (scheduled) replication between sites is based on the Microsoft implementation of Distributed Computing Environment/Remote Procedure Calls (DCE/RPC) over TCP/IP.

Note For information about Active Directory replication, refer to the following resources:

- *How Active Directory Replication Works*
  - http://social.technet.microsoft.com/wiki/contents/articles/4592.how-active-directory-replication-works.aspx
- *Active Directory Replication Technologies*
  - https://technet.microsoft.com/en-us/library/cc776877%28v=ws.10%29.aspx

Firewall Design for AD Replication

Remote Procedure Call (RPC) dynamic port allocation is used by many server applications. RPC dynamic port allocation will instruct the RPC program to use a particular random port in the range configured for TCP and UDP, based on the implementation.

Some AD DS rely on Microsoft Distributed Component Object Model (DCOM) RPC for service replication. The default dynamic port range varies depending on the Windows platform (for example, 1025-5000 for Windows Server 2003 and 49152-65535 for Windows Server 2008), while the Cisco ASA used pinholing to limit the number of open ports across the firewall. Getting replication to function properly across security perimeters can be challenging. Three possible approaches exist:

- Open up the firewall to permit RPC's native dynamic behavior.
- Limit RPC’s use of TCP ports and open the firewall for a small range of ports.
- Encapsulate the DC-to-DC traffic inside IP Security Protocol (IPsec) and open the firewall for the IPSec only between the DCs.

Cisco FTD uses application detectors to identify the commonly used applications in your network. These application detectors can then be used in access control rules, to provide a granular method of handling network traffic cross multiple managed devices, enabling simple rule creation for AD replication.

Specific to DEC/RPCC, Cisco FTD provides a preprocessor which normalizes the packet data into formats that the intrusion rules engine can analyze, enabling traffic to not only be policed by the access control rule, but for potential exploits to be detected through the IPS engine. More information regarding the DEC/RPC preprocessor in Cisco FTD can be found at:

https://www.cisco.com/c/en/us/td/docs/security/firepower/70/configuration/guide/fpmc-config-guide-v70/application_layer_preprocessors.html#ID-2244-00000019

More details of the DEC/RPC operations are listed below:

- An RPC service configures itself in the registry with a universally unique identifier (UUID). UUIDs are well-known identifiers unique for each service and common across all platforms.
- When an RPC service starts, it obtains a free high port and registers that port with the UUID. Some services use random high ports; others try to use the same high ports all the time (if they are available). The port assignment is static for the lifetime of the service.
- Once the service restarts with a new process or network server reload, the port assignment changes. This makes it impossible to know in advance which port an RPC service will use. The DEC/RPC inspection monitors the communication between the Endpoint Mapper (EPM) on a server and a client on the well-known TCP port 135. The embedded server IP address and port number are received from the EPM response messages.

Note For detailed information, refer to *Active Directory and Active Directory Domain Services Port Requirements* at:

- https://docs.microsoft.com/en-us/troubleshoot/windows-server/networking/service-overview-and-network-port-requirements

Figure 2-17 illustrates the AD replication between the DCs in the Industrial and Enterprise Zones.

Figure 2-17 Domain Controller Bi-Directional Replication

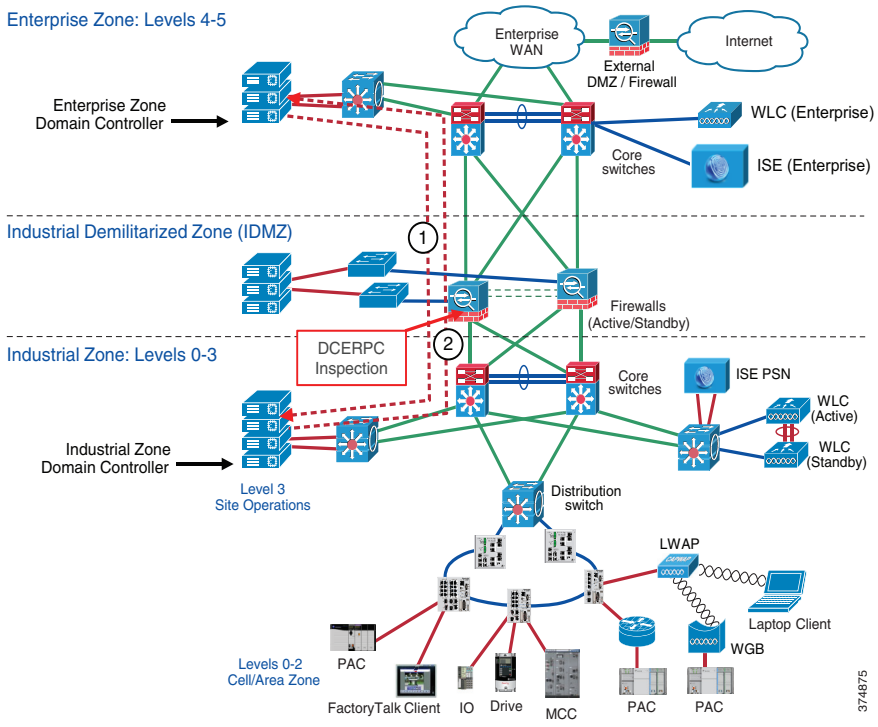

1. The Enterprise Domain Controller replicates any changes to the Industrial Zone Domain Controller using the RPC protocol. The firewall inspects the RPC traffic and dynamically opens necessary ports.
2. The Industrial Domain Controller replicates any changes to the Enterprise Zone Domain Controller. The firewall inspects the RPC traffic and dynamically opens necessary ports.

In addition to RPC, other ports may need to be opened between the DCs, depending on the implementation. Table 2-5 shows an example of protocols that may be required for AD replication. Note that this may not be a complete list depending on the AD configuration and the requirements.

Table 2-5 AD Replication Ports Example

Protocol / Service Name	TCP/UDP Port
SMB over IP	TCP 445
Kerberos	TCP/UDP 88
LDAP, LDAP SSL	TCP/UDP 389, 636
LDAP GC, LDAP GC SSL	TCP/UDP 3268, 3269
RPC	TCP/UDP 135

Authentication of IDMZ Resources

IDMZ hosts that belong to the AD domain have to authenticate to the Enterprise DC. Examples of such hosts include Terminal Services Gateway, anti-virus and Windows Update servers. To achieve this, the firewall access policy should allow certain protocols between the IDMZ and the Enterprise Zone. The policy should be restricted to specific IP addresses in the IDMZ that require authentication. The dynamic RPC inspection should also be included in the policy (see Active Directory Replication, page 2-30).

Table 2-6 shows an example of protocols that may be required for AD authentication. Note that this may not be a complete list depending on the AD configuration and the requirements.

Table 2-6 AD Authentication Ports Example

Protocol / Service Name	TCP/UDP Port
SMB over IP	TCP 445
Kerberos	TCP/UDP 88, 464
LDAP, LDAP SSL	TCP/UDP 389, 636
DNS	TCP/UDP 53
RPC	TCP/UDP 135

Note More information on AD port requirements can be found in *Active Directory and Active Directory Domain Services Port Requirements* at:

- https://docs.microsoft.com/en-us/troubleshoot/windows-server/networking/service-overview-and-network-port-requirements

Certificate Services

This section provides an overview of certificate services and public key infrastructure (PKI).

Note CPwE IDMZ does not cover PKI in depth nor does it recommend how to properly implement or manage PKI. For test purposes, firewalls and other devices using self-signed certificates as PKI management were beyond the scope of this CPwE DIG.

Certificate Services Overview

The Certificate Authority (CA) is a trusted entity that manages and issues security certificates and public keys that are used for secure communication in a public network. The CA is part of the PKI along with the Registration Authority (RA) who verifies the information provided by a requester of a digital certificate. If the information is verified as correct, the certificate authority can then issue a certificate.

PKI is a scalable architecture that includes software, hardware and procedures to facilitate the management of digital certificates. Certificate-based authentication methods can be required for:

- User network access, both wired and wireless
- Authentication of network devices, for example servers and wireless APs

Access Point (AP) Certificate Services can also be used to:

- Enroll users for certificates from the CA using the Web or the Certificates Microsoft Management Console (MMC) snap-in, or transparently through auto enrollment
- Use certificate templates to help simplify the choices a certificate requester has to make when requesting a certificate, depending upon the policy used by the CA
- Take advantage of the AD service for publishing trusted root certificates, publishing issued certificates, and publishing Certificate Revocation Lists (CRLs)
- Implement the ability to log on to a Windows operating system domain using a smart card

Note For more information about AD Certificate Services, refer to *Active Directory Certificate Services* at:

- https://technet.microsoft.com/en-us/windowsserver/dd448615.aspx

Certificate Authority Hierarchy

PKI supports a hierarchical structure with various CA roles in the network, depending on the scale of the system.

- Root CA is the most trusted CA in a CA hierarchy and is the first CA installed in the network. When a root CA remains online, it is used to issue certificates to the intermediate and subordinate CAs. Most times, the root CA remains offline to protect the private keys. The root CA rarely issues certificates directly to users, computers or services.
- Intermediate CAs are the next in hierarchy after the root CA. The intermediate CA issues certificates only to subordinate CAs.
- Subordinate CAs can be used to issue certificates to users and computers, or to issuing CAs.
- Issuing CA is used to issue certificates directly to users and computers.

AD Certificate Services can be deployed into Enterprise CA and stand-alone CA modes depend on the customer specific requirements:

- Enterprise CA is integrated with AD and use domain services for certificate management. Enterprise CAs are typically used for issuing user and computer certificates.
- Stand-alone CA is not dependent on AD and not part of a domain. A stand-alone mode is often used to implement a secure offline root CA.

Figure 2-18 shows the CA hierarchy and various deployment models depending on the system scale.

Figure 2-18 PKI Infrastructure Models Example

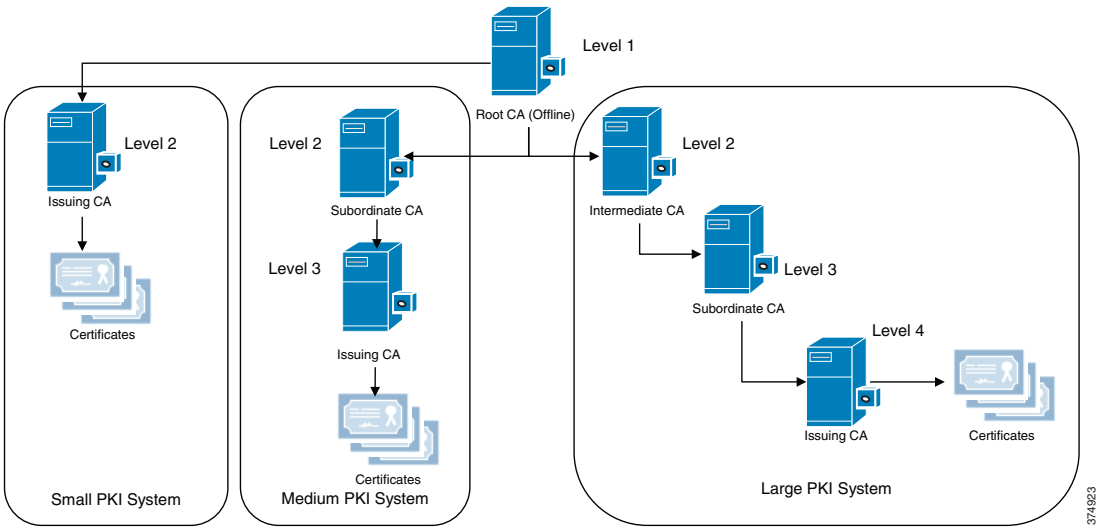

Certificate Services Architecture in IDMZ

Similar to AD DS service, the deployment of the Active Directory Certificate Services (AD CS) in the CPwE architecture is based on the AD implementation in a single domain. To provide a local CA for each Industrial Zone, the root CA should be configured in the Enterprise Zone, with a subordinate CA in the secured Industrial Zone.

Enterprise Zone root and subordinate CAs will have full-fledged functionalities to provide the following services:

- **Certification Authorities (CAs)**—Root, intermediate, subordinate and issuing CAs are used to issue certificates to users, computers, and services, and to manage certificate validity.
- **CA Web Enrollment**—Web enrollment allows users to connect to a CA by means of a Web browser in order to request certificates and retrieve CRLs.
- **Online Responder**—The Online Responder service accepts revocation status requests for specific certificates, evaluates the status of these certificates, and sends back a signed response containing the requested certificate status information.
- **Network Device Enrollment Service**—The Network Device Enrollment Service allows routers and other network devices that do not have domain accounts to obtain certificates.
- **Certificate Enrollment Web Service**—The Certificate Enrollment Web Service enables users and computers to perform certificate enrollment that uses the HTTPS protocol. Together with the Certificate Enrollment Policy Web Service, this enables policy-based certificate enrollment when the client computer is not a member of a domain or when a domain member is not connected to the domain.
- **Certificate Enrollment Policy Web Service**—The Certificate Enrollment Policy Web Service enables users and computers to obtain certificate enrollment policy information. Together with the Certificate Enrollment Web Service, this enables policy-based certificate enrollment when the client computer is not a member of a domain or when a domain member is not connected to the domain.

Subordinate CA behind the IDMZ firewall is responsible for issuing and validating client's **Certificate Signing Request (CSR)** and authentication requests inside the Industrial Zone. Issuing certificates to users or devices inside an Industrial Zone, instead of forwarding all requests to Enterprise Zone Root-CA, allows certificate services to operate in case of incidents when the enterprise CA is not available.

Multiple Subordinate CAs inside the Industrial Zone can also achieve plant-wide smooth operation during a failure of any single subordinate CA.

Network Time Protocol

Time synchronization is a critical requirement in most industrial systems. Network Time Protocol (NTP) is one of the most common protocols governing time transfer in computer networks. The following sections present recommendations and considerations for deploying NTP in the CPwE IDMZ architecture.

NTP Overview

Network Time Protocol (NTP) version 4 is an IETF standard defined in RFC 5905. NTP uses a hierarchy of clocks with each level referred to as a stratum. This hierarchy begins at stratum 0, which is the primary reference clock.

- The primary reference clock (stratum 0) synchronizes to Coordinated Universal Time (UTC) using a GPS, radio, or atomic clock.
- Stratum 1 clocks synchronize directly with the reference clock and are the first clocks connected to the network.

- Stratum 2 clocks will synchronize against multiple stratum 1 clocks. Stratum 3 clocks will synchronize with multiple stratum 2 clocks and so on.

An NTP-enabled device never synchronizes to a device that is not synchronized itself. Additionally, an NTP-enabled device compares the time reported by several NTP devices, and will not synchronize to a device whose time is significantly different than others.

In general, a lower stratum will have higher precision and accuracy than a higher stratum clock. However, the quality of the components used in the NTP servers has a large impact on the accuracy and precision of time. For example, the stratum 4 server that is part of a hierarchy with high quality clocks and well performing networks may have a higher precision and accuracy than a stratum 3 server that uses poor quality clocks and networks in the hierarchy.

No more than one NTP transaction per minute is necessary to achieve 1 millisecond synchronization on a high-speed LAN. For larger systems (wide-area networks), NTP can routinely achieve 10 millisecond synchronization. However, the level of synchronization is not guaranteed and can be affected by the infrastructure. Asymmetric routes and network congestion can cause errors of 100 ms or more.

Windows Time Service Overview

Microsoft Windows clients and servers use the Windows Time Service (W32Time) to synchronize time across the domain. By default, W32Time uses a combination of AD and NTP to propagate time throughout the domain hierarchy. While AD uses a multi-master model for directory updates, some updates must happen using a single master model or FSMO roles.

One of the key FSMO roles in an AD domain is the Primary Domain Controller (PDC) emulator. In the Windows Time Service model, the Domain Controller that holds the PDC emulator role acts as the master time source for the domain. The PDC emulator should synchronize its clock to at least two reliable NTP sources.

The other domain controllers in the domain synchronize their clocks to the PDC emulator. In addition, the Windows clients synchronize their clocks to the local domain controller.

Note It is important to understand that the W32Time service has limited accuracy and precision. It cannot reliably maintain time synchronization to more than a few seconds.

NTP Architecture in IDMZ

Various applications within the Industrial Zone use NTP for clock synchronization, for example:

- AD uses Kerberos protocol for authentication within the domain. Kerberos authentication uses timestamps to prevent replay attacks. By default, authentication request will fail if the client and server clocks differ by more than 5 minutes.
- Infrastructure devices such as routers, switches, and firewalls should synchronize their clocks via NTP. Many of these devices do not have onboard real-time clocks and will revert to a default date and time after a reboot. Devices such as these log critical event data to an internal or external syslog. Proper time stamps on these log entries are important for identifying and resolving faults in the device. Furthermore, synchronized clocks allow for system wide fault analysis involving multiple infrastructure devices.

Note NTP and especially W32Time are not appropriate for high precision applications such as CIP Motion™ and Sequence of Events (SOE) applications using FactoryTalk Alarms and Events. These applications must use IEEE 1588 Precision Time Protocol (PTP) sourced from a reliable reference clock.

NTP Server Choice

The Enterprise Zone should have two reliable NTP servers that serve time for the enterprise systems. The enterprise time servers should synchronize to privately owned reference clocks to provide the most accurate and precise time. GPS time servers, which are relatively inexpensive, are a good choice for enterprise reference clocks. These clocks can be backed up by public time servers available on the Internet. However, public Internet time servers may not be reliable enough to be the sole primary reference for systems where accurate and precise time is critical.

The Industrial Zone should also have two reliable NTP servers. In medium precision applications, the industrial time servers can sync directly with the enterprise servers. However, consider deploying reference clocks in the Industrial Zone if high precision timestamps are required for the application. A number of vendors sell reference clocks that function as a NTP stratum 1 clock as well as a PTP grandmaster clock for CIP Sync and CIP Motion applications.

NTP Synchronization through IDMZ

Because of the critical role that time synchronization plays in most networks, NTP is one of the few protocols that directly traverse from the Enterprise Zone to the Industrial Zone. The Industrial Zone NTP servers should be allowed to communicate directly with the Enterprise NTP servers on UDP port 123. NTP servers should use NTP authentication to validate the identity of the source clock during synchronization. In addition, the IPS/IDS in the firewall should inspect the integrity of NTP traffic passing through.

AD domain controllers also need to synchronize using the NTP protocol. The synchronization rules will vary depending on the domain structure. In the single domain model, the domain controllers need visibility to the PDC emulator. In a multi-domain model, the domain controllers need visibility to the PDC emulator or a domain controller in the parent domain.

https://docs.microsoft.com/en-us/windows-server/networking/windows-time-service/windows-time-service-top

Note For more information on the Windows Time Service, refer to *Windows Time Service Technical Reference* at:

- https://docs.microsoft.com/en-us/windows-server/networking/windows-time-service/windows-time-service-top

Figure 2-19 illustrates NTP data traversal across the IDMZ between NTP servers in different zones. In this example, the NTP server in the Enterprise Zone can be a corporate time source (stratum 1 or stratum 2) or a PDC emulator in the AD domain. The NTP server in the Industrial Zone can also be a domain controller that synchronizes to the PDC.

Note Depending on the requirements, the Industrial Zone may have its own reference clock for NTP synchronization, such as a GPS clock.

Figure 2-19 NTP Synchronization across IDMZ

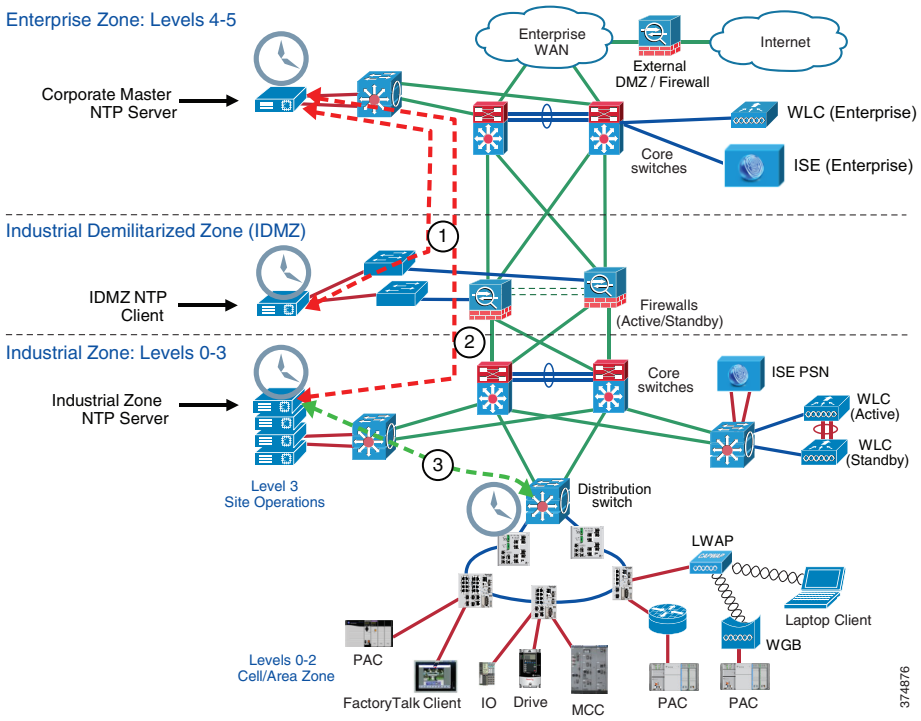

1. The NTP client in the IDMZ synchronizes time with the corporate Master NTP server in the Enterprise Zone.
2. The NTP server in the Industrial Zone synchronizes time with the corporate Master NTP server in the Enterprise Zone.
3. Industrial NTP clients synchronize their clocks with the Industrial NTP server.

Recommendations and considerations for the NTP deployment in the CPwE IDMZ architectures are summarized as follows:

- Deploy reference clocks (stratum 1) in the Enterprise Zone and Industrial Zone as needed.
- Use public NTP servers as a backup to private reference clocks.
- NTP servers should sync to at least two reliable clocks at a lower stratum.
- Synchronize the W32Time clock on the PDC Emulator to at least two reliable NTP servers.
- Be aware of limited accuracy and precision of the W32Time.
- Configure NTP authentication and inspect NTP traffic on the firewall.

Identity Services Engine

This section provides an overview of the distributed Cisco Identity Services Engine (ISE) architecture in the CPwE IDMZ.

Note For more information about ISE deployment in the CPwE, refer to the *Deploying Identity Services within a Converged Plantwide Ethernet Architecture Design and Implementation Guide* at:

- http://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/3-5-1/ISE/DIG/CPwE_ISE_CVD.html

ISE Overview

With the introduction of secure employee and contractor access, the use of the Cisco ISE as an identity and access control policy platform enables organizations to enforce compliance, enhance infrastructure security and streamline their service operations. The ISE architecture allows an organization to gather real-time contextual information from the network, users and devices to make proactive policy decisions by tying identity into various network elements including access switches and WLCs.

The ISE functions as the authentication and authorization server for the wired and wireless networks using RADIUS protocol. The ISE can use AD as an external identity database for resources such as users, machines, groups and attributes. Cisco ISE supports Microsoft AD Sites and Services when integrated with AD. ISE needs an identity certificate that is signed by a CA server so that it can be trusted by endpoints, gateways and servers.

Distributed ISE Architecture

In the distributed ISE architecture, multiple ISE nodes assume different roles (personas) in the network:

- **Policy Administration Node (PAN)** persona allows the Enterprise IT team to perform all administrative operations on the ISE system. The PAN handles all system-related configurations that are related to functionality such as authentication and authorization. An architecture can have one or a maximum of two PANs that can have the standalone, primary or secondary role.
- **Policy Service Node (PSN)** persona provides network access, plant personnel and guest access and client provisioning and profiling services. This persona evaluates the policies and provides network access to computers based on the result of the policy evaluation. More than one node can assume this persona and typically more than one PSN exists in a large distributed deployment.
- **Monitoring ‘n Troubleshooting (MnT) Node** persona functions as the log collector and stores log messages from all PANs and PSNs in a network. This persona provides advanced monitoring and troubleshooting tools that the Enterprise IT team can use to effectively manage a network and resources. A maximum of two MnTs can take on primary or secondary roles for high availability. At least one node in a distributed setup should assume the Monitoring persona. For optimum performance, an MnT persona should not be enabled on the same node as PSN or PAN and should be dedicated solely to monitoring.

ISE Architecture in IDMZ

Within the CPwE IDMZ architecture, the recommendation is to deploy the Cisco ISE platform as a distributed solution (see Figure 2-20).

- The corporate IT department maintains the management of the ISE platform via PAN in the Enterprise Zone. The MnT is also deployed in the Enterprise Zone.
- One or multiple PSNs are deployed in the Industrial Zone for identity services. The PAN synchronizes its policy configurations with PSNs.
- The IDMZ firewall is configured to allow ISE synchronization and logging traffic between the nodes (see ISE Configuration, page 3-14 for details).

Figure 2-20 Distributed ISE Architecture

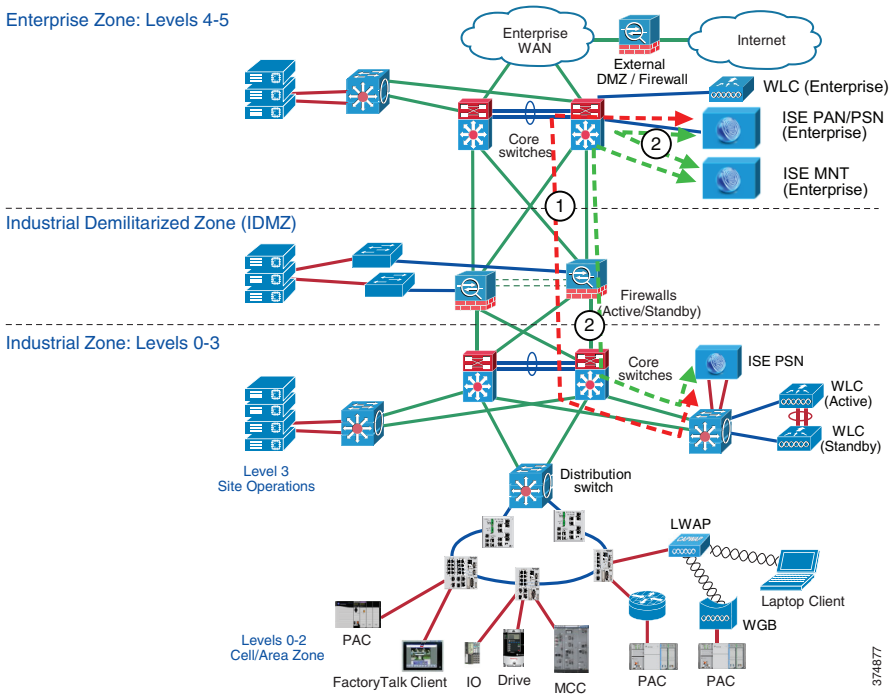

1. The Enterprise ISE PAN/PSN synchronizes its policy configurations with the Industrial ISE PSN.
2. The Enterprise and Industrial ISE PSNs send detailed logs to the Enterprise ISE Mn1T.

Multi-Factor Authentication

Multi-factor authentication (MFA) is a security process that requires users to respond to requests to verify their identities before they can access networks or other online applications. MFA may use knowledge, possession of physical objects, or geographic or network locations to confirm identity.

MFA Overview

Authentication based on usernames and passwords alone is unreliable and unwieldy, since users may have trouble storing, remembering, and managing them across multiple accounts, and many reuse passwords across services and create passwords that lack complexity. Passwords also offer weak security because of the ease of acquiring them through hacking, phishing, and malware.

MFA requires means of verification that unauthorized users won't have. Since passwords are insufficient for verifying identity, MFA requires multiple pieces of evidence to verify identity. The most common variant of MFA is two-factor authentication (2FA). The theory is that even if threat actors can impersonate a user with one piece of evidence, they will not be able to provide two or more.

Proper multi-factor authentication uses factors from at least two different categories. Using two from the same category does not fulfill the objective of MFA. Despite wide use of the password/security question combination, both factors are from the knowledge category--and don't qualify as MFA. A password and a temporary passcode qualify because the passcode is a possession factor, verifying ownership of a specific email account or mobile device.

Processes vary among the different MFA methods, but a typical 2FA transaction happens like this:

- The user logs in to the website or service with their username and password.
- The password is validated by an authentication server and, if correct, the user becomes eligible for the second factor.
- The authentication server sends a unique code to the user's second-factor method (such as a smartphone app).
- The user confirms their identity by providing the additional authentication for their second-factor method.

In adaptive authentication, authentication rules continuously adjust based on the following variables:

- By user or groups of users defined by role, responsibility, or department
- By authentication method: for example, to authenticate users via push notification but not SMS
- By application: to enforce more secure MFA methods--such as push notification or Universal 2nd Factor (U2F)--for high-risk applications and services
- By geographic location: to restrict access to company resources based on a user’s physical location, or to set conditional policies restricting use of certain authentication methods in some locations but not others
- By network information: to use network-in-use IP information as an authentication factor and to block authentication attempts from anonymous networks like Tor, proxies, and VPNs

Multi-Factor Authentication Architecture in IDMZ

Figure 2-21 Cisco Secure Access by Duo IDMZ Architecture

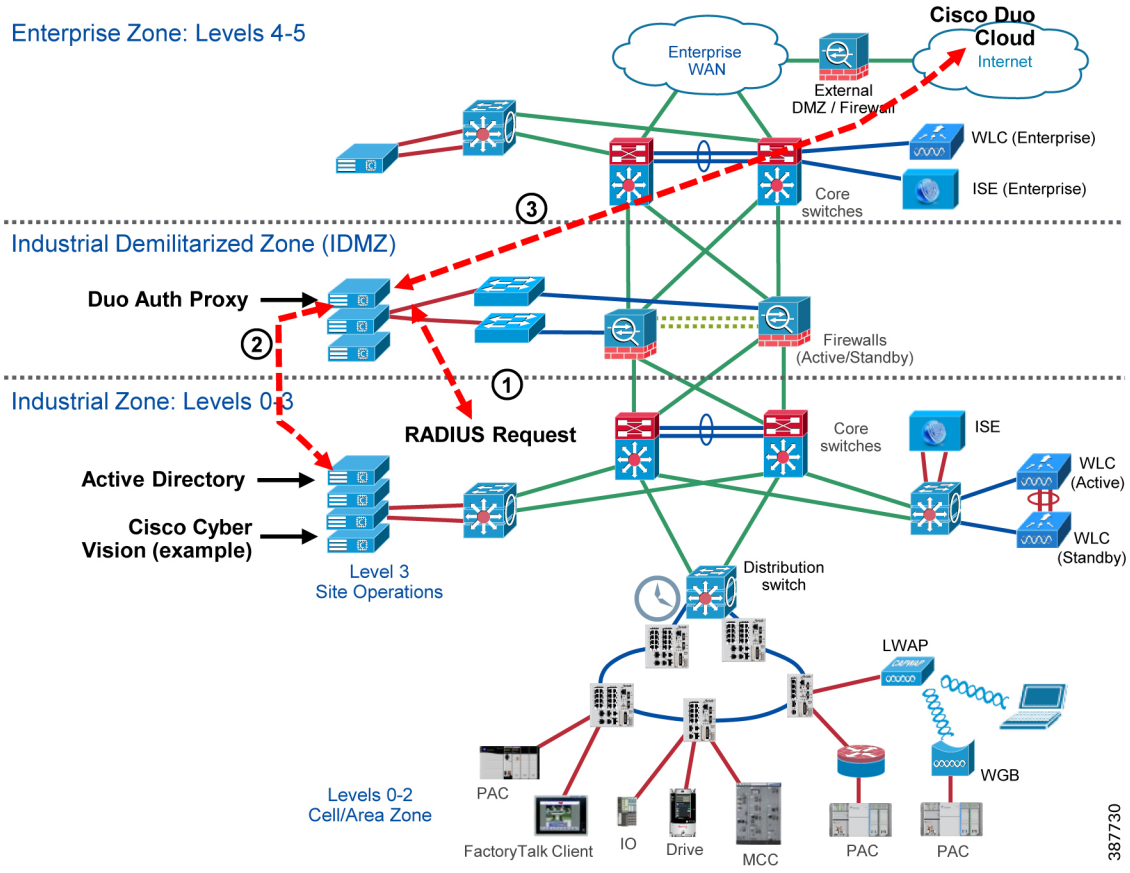

Multi-factor authentication from Cisco's Duo protects your applications by using a second source of validation, like a phone or token, to verify user identity before granting access. Many of Duo's application integrations do not require any local components. However, certain services do require a local Authentication Proxy service. This document makes use of the Duo Authentication Proxy for:

- Remote Access VPN
- Microsoft Remote Desktop Gateway
- Windows Logon and RDP

Duo's Authentication Proxy (sometimes referred to as the Authproxy) is a local service needed to properly configure certain Duo-protected applications. The Authentication Proxy can be installed on a physical or virtual host, on Windows or Linux machines. Once configured, Duo sends your users an automatic authentication request via Duo Push notification to a mobile device or phone call after successful primary login.

1. RADIUS Request is sent to Duo Authentication Proxy.
2. Duo Authentication proxy validates primary credentials with Active Directory.
3. Duo Authentication proxy, if credentials were valid, prompts the Duo cloud to send 2FA.
4. Duo Authentication proxy returns accept or deny response back to the application.

Licensing

Cisco Smart Licensing is a flexible software licensing method that simplifies the way you activate and manage licenses across your organization. With Smart Licensing, a pool of licenses is associated to a Cisco Smart Account. Like banking, new licenses are automatically deposited into your Smart Account, increasing your account balance of licenses (also known as entitlements). As licenses expire or are terminated, your inventory balance decreases.

Cisco Smart Software Manager On-Prem Overview

Cisco Smart Software Manager (SSM) On-Prem license server is a component of Cisco Smart Licensing. It works in conjunction with Cisco Smart Software Manager to intelligently manage customer product licenses, providing near-real-time visibility and reporting of the Cisco licenses that customers purchase and consume.

Cisco Smart Licensing requires products to be associated with Smart Accounts, which can be created on Cisco Software Central. A Smart Account is associated with a unique company ID and is like an online banking account containing Cisco entitlements and devices for that customer. From the Cisco Smart Software Manager, subaccounts (also called virtual accounts) can be created to represent various subdivisions or buying centers of the company.

Cisco SSM On-Prem is targeted for security-sensitive customers who are unable to manage their installed base with a direct Internet connection. For devices and applications in the industrial zone, the Cisco SSM On-Prem provides a mechanism to provide local control and management of their license usage, without the need for manual intervention.

Cisco Smart Software Manager On-Prem IDMZ Architecture

Figure 2-22 Cisco Smart Software Manager On-Prem IDMZ Architecture

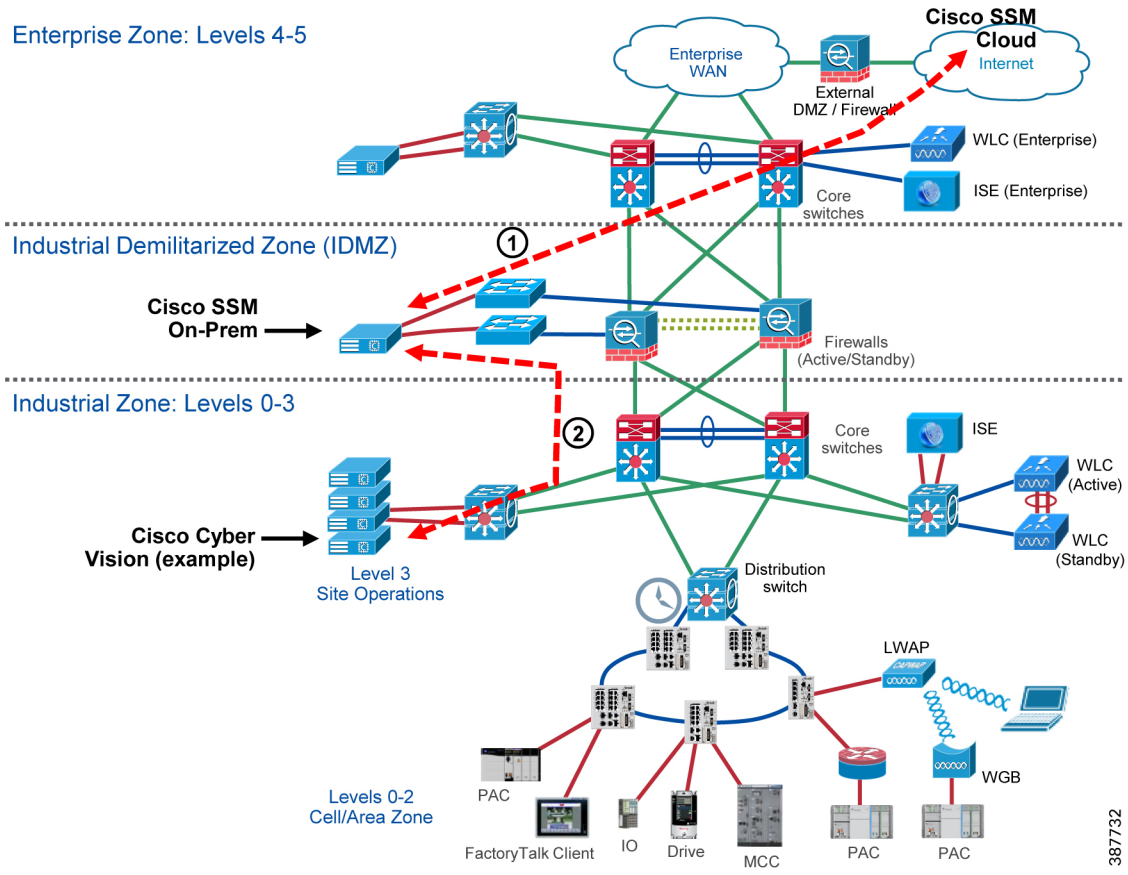

1. Cisco SSM On-Prem is installed in the IDMZ.
2. If the Cisco SSM is in networking mode, it will continually synchronize with Cisco Smart Licensing cloud.
3. Products in the Industrial Zone request license reservation from the Cisco SSM On-Prem server.
4. Cisco SSM On-Prem server returns valid licenses to products in the Industrial Zone.

Windows Updates

Update management is the process of controlling the deployment and maintenance of interim software releases into production environments. Windows Server Update Services (WSUS) enables information technology administrators to deploy the latest Microsoft product updates. You can use WSUS to fully manage the distribution of updates that are released through Microsoft Update to computers on your network.

Windows Server Update Services Server Role Description

A WSUS server provides features that you can use to manage and distribute updates through a management console. A WSUS server can also be the update source for other WSUS servers within the organization. The WSUS server that acts as an update source is called an upstream server. In a WSUS implementation, at least one WSUS server on your network must be able to connect to Microsoft Update to get available update information.

Windows Server Update Services IDMZ Architecture

For this design, the WSUS server role will be deployed in the IDMZ and connect directly to Microsoft Update for periodic update checks. All Windows devices in the Industrial Zone will connect to the WSUS for updates.

Note It is recommended that automatic updates are not enabled so that updates are not immediately installed. Updates should be installed during scheduled maintenance windows to reduce downtime.

Figure 2-23 Windows Server Update Services IDMZ Architecture

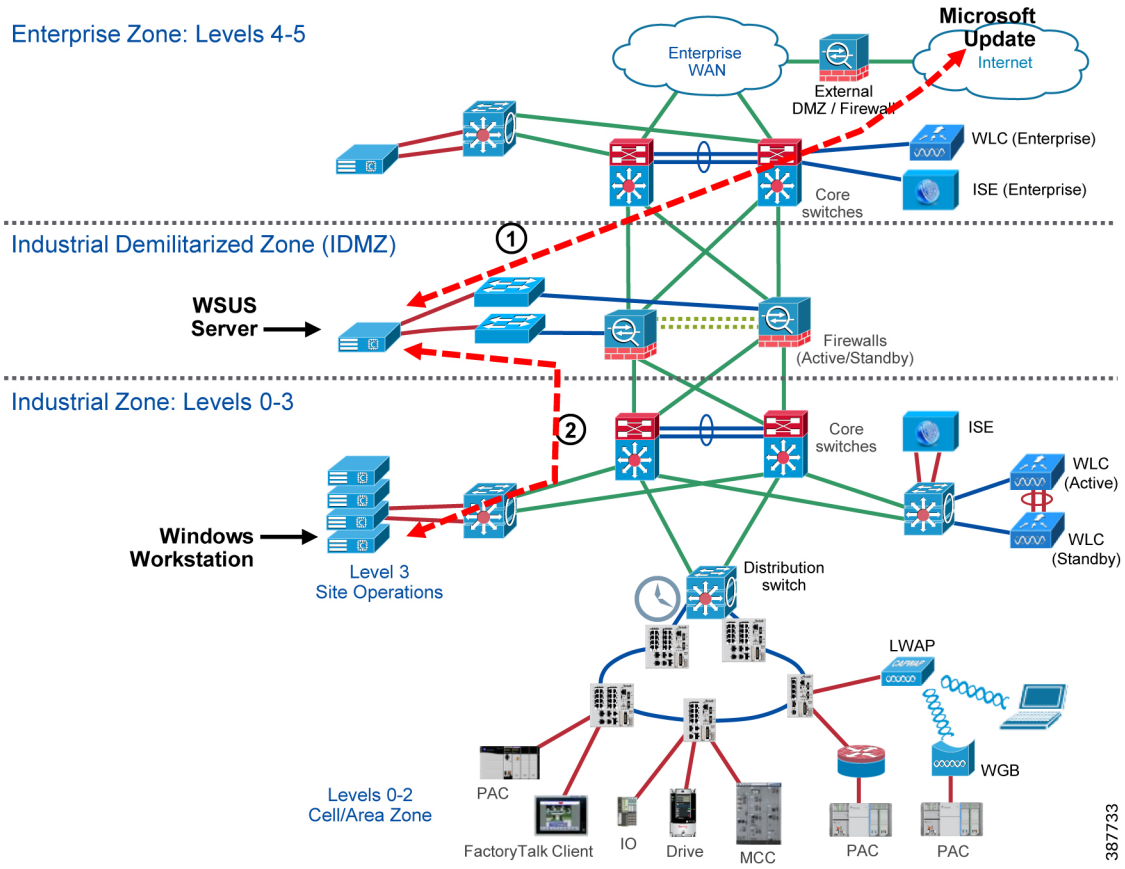

1. WSUS server role is installed on Windows Server in the IDMZ.
2. WSUS server is configured to use web proxy to pull Microsoft updates from the cloud.
3. Create computer groups to add client computers for update services.

4. Clients in the Industrial Zone will check the WSUS server for updates periodically. Updates can be installed locally on the client, but should not be automatically installed so changes can be approved by an administrator.

Data Transfer through the IDMZ

The key principle of the IDMZ is to meet the requirement to share necessary IACS data between the Industrial and Enterprise Zones while not allowing direct communication between the zones. This goal is achieved by placing gateways, application data mirrors, proxy servers and similar services in the IDMZ.

Two application examples in this section demonstrate how data can be transferred through the IDMZ in a secure way:

- FactoryTalk Historian data transfer
- Secure file transfer with Cisco Advanced Malware Protection

FactoryTalk Historian Data Transfer

Access to process and operational data is often a key requirement for setting up an industrial network. This data can, among other purposes, visualize process and production progress, identify improvement opportunities and assist in troubleshooting.

FactoryTalk Historian establishes a reliable foundation for capturing this data. The suite of software products can target a single machine (FactoryTalk Historian Machine Edition) or be a plant-wide system (FactoryTalk Historian Site Edition). It can also be extended across the global enterprise using the Rockwell Automation Global Enterprise Historian Strategy.

With FactoryTalk Historian Site Edition (SE), you can collect critical time-series data for various calculations, estimations, and statistical processes producing information to benefit a multitude of enterprise-wide processes and applications. An overview of the FactoryTalk Historian SE operation is provided below:

- At its core, FactoryTalk Historian Site Edition stores user defined data (tag + value pairs) into an archiving system on the FactoryTalk Historian SE server. A FactoryTalk VantagePoint client can access this data. These archives can also be queried through a specialized OLEDB connection (PI OLEDB) by means of a SQL query language like T-SQL.
- The data gets collected from PACs in the Industrial Zone through instances of RSLinx Enterprise and FactoryTalk® Linx running on separate servers. These servers have FactoryTalk® Live Data interfaces installed that relay the data collected by RSLinx to the FactoryTalk Historian SE server.
- It is best practice to install the FactoryTalk Historian SE server and two independent FactoryTalk® Live Data interfaces on separate physical or virtual server hardware for a more robust and redundant system.
- The FactoryTalk Historian SE server can be made part of a collective for even better redundancy.

Note For more information on FactoryTalk Historian, refer to:

- http://www.rockwellautomation.com/rockwellsoftware/products/factorytalk-historian.page?

Application Requirements

In the Industrial Zone, FactoryTalk® Historian SE server functions as data (archives) repository, central point for data queries, coordinator for FactoryTalk® Live Data interfaces and access point for system setup and maintenance. It can be installed as a single server or several servers bound into a collective. A recommended two separate servers should have a FactoryTalk® Live Data interface installed along with an install of RSLinx Enterprise.

A requirement for the Enterprise Zone is to have a custom enterprise level reporting at multiple clients, which includes Industrial Zone historical data and trending. A FactoryTalk Historian server can be installed in the Enterprise Zone to provide centralized data aggregation from site historians.

One of the main design goals of the CPwE is no direct traversal of data from the Enterprise Zone to the Industrial Zone or vice versa. In order to securely replicate production data to the Enterprise Zone, a PI-to-PI (Historian to Historian) replication service can be installed on a server in the IDMZ. The PI-to-PI Interface copies data between two instances of FactoryTalk Historian server (single server or a collective).

FactoryTalk Historian IDMZ Architecture

Figure 2-24 illustrates a recommended IDMZ architecture for the FactoryTalk Historian. Two Historian servers are installed, one in the Industrial Zone and one in the Enterprise Zone. A PI-to-PI Interface is installed in the IDMZ to copy data between the two instances of FactoryTalk Historian (either a single server or a collective).

A FactoryTalk VantagePoint server is also installed in the Enterprise Zone to collect data from the Historian in the Enterprise Zone.

Figure 2-24 FactoryTalk Historian Data Transfer

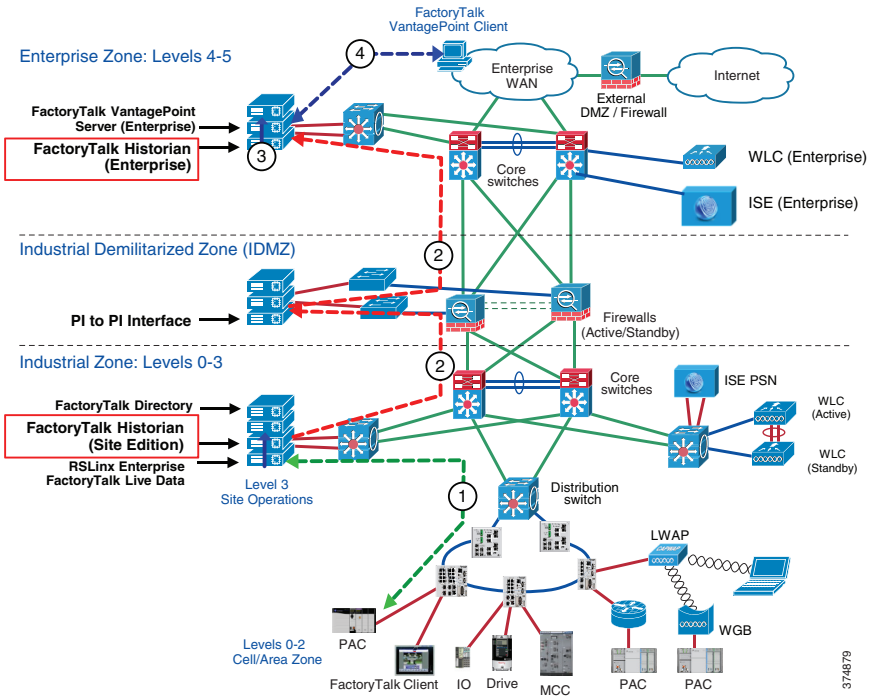

1. Controller data is sent to the FactoryTalk Historian SE data repository via RSLinx Enterprise and FactoryTalk® Live Data interfaces.

2. The PI-to-PI Interface pulls predefined data from the FactoryTalk Historian SE in the Industrial Zone and pushes the data to the FactoryTalk Historian in the Enterprise Zone.
3. FactoryTalk VantagePoint server in the Enterprise Zone gathers preconfigured data from the Enterprise Zone Historian to generate reports.
4. A FactoryTalk VantagePoint client requests a web report based on the data collected from the Enterprise Zone Historian data.

Note By installing a second FactoryTalk VantagePoint server in the Industrial Zone, data can be visualized in both the Enterprise and the Industrial Zones.

An overview of PI-to-PI Interface configuration and firewall rules for the Historian data transfer is provided in FactoryTalk Historian Data Transfer Configuration, page 3-21.

FactoryTalk VantagePoint Connectivity to Historian Server

The previous example described connectivity of an Enterprise VantagePoint server getting data from the Enterprise Zone Historian.

- It is recommended that the FactoryTalk VantagePoint clients do not cross security boundaries to obtain FactoryTalk Historian information or connect to a FactoryTalk VantagePoint server in another security zone. For instance, it is not recommended for an Enterprise FactoryTalk VantagePoint server or client to connect directly to the Industrial Zone FactoryTalk Historian (see Figure 2-25).
- If a client in the Enterprise zone wants access to an asset in the Industrial Zone, the client can then access the server via one of the remote access methods available (described later in this chapter).

Figure 2-25 FactoryTalk VantagePoint in Enterprise Zone

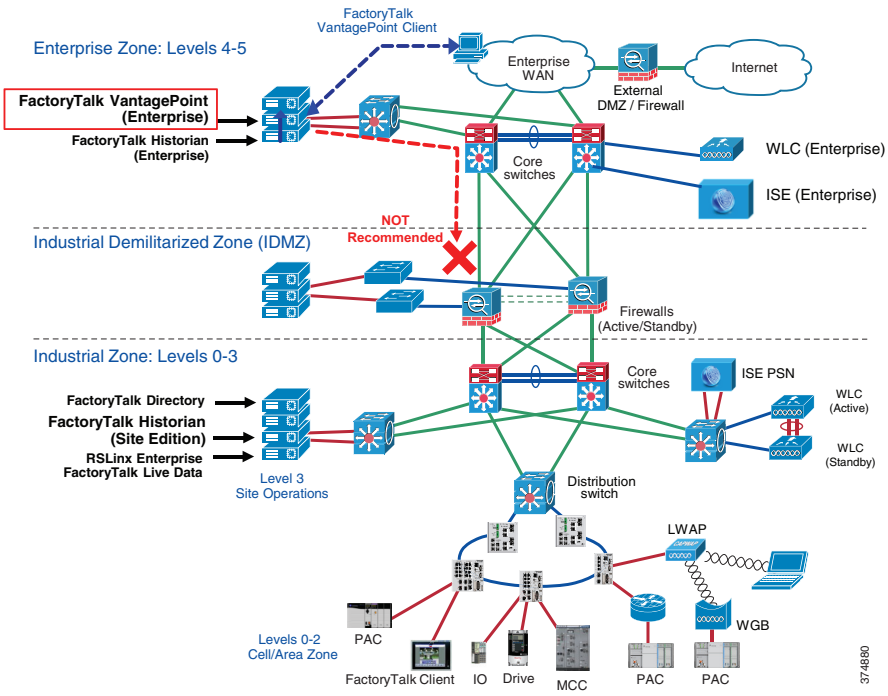

Secure File Transfer

Employees often need to transfer files between the Industrial and Enterprise Zones due to business requirements. Some examples of files that need to be passed between both zones are production reports, assembly line instructions, user manuals and software installation files. Traditional ways to move files, such as Windows file shares, email attachments, USB drives or third-party web-based solutions, can be unsecure or introduce significant risks to the industrial environment. Secure File Transfer (SFT) solutions provide a secure way to accomplish the task in compliance with the IDMZ design principles.

Overview of Managed File Transfer (MFT) Solutions

Several products on the market provide the solution for a managed secure file transfer between the Enterprise and the Industrial Zone. These solutions require the installation of a file transfer server gateway in the IDMZ with SFT servers located in the Enterprise and Industrial Zones.

This approach allows for secure file transfers between the zones using the IDMZ gateway as a proxy. An industrial SFT system should have the following characteristics:

- All incoming connections are accepted on a hardened SFT gateway in the IDMZ using a secure protocol, for instance Secure File Transfer Protocol (SFTP), FTP over SSL or HTTPS.
- Files can be inspected by the Firewall during transit using Cisco AMP.
- Users can be authenticated against the AD to confirm that only certain people can send or receive files. Depending on the security policy, different groups of authorized users can be created in the Industrial and Enterprise Zones.
- No inbound connections are allowed from the IDMZ to the Industrial Zone. Only authorized users can initiate file transfer connections.
- Files are not stored permanently in the IDMZ.
- The system can provide audit trail tracking and extensive reporting.

Secure File Transfer Architecture

Figure 2-26 provides an overview of the SFT architecture for the CPwE IDMZ.

Figure 2-26 Secure File Transfer

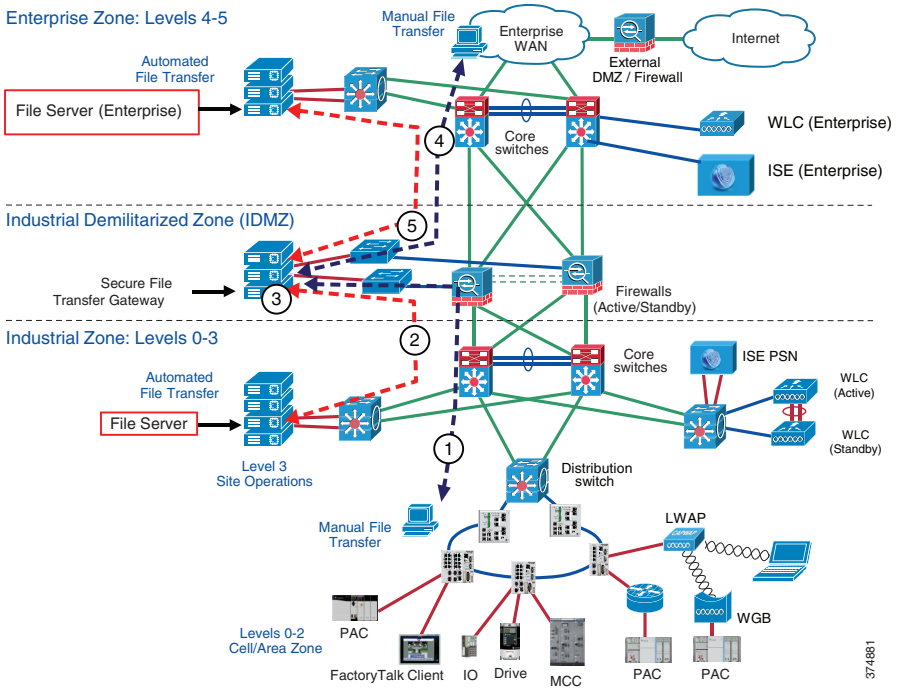

The steps below describe a manual or automated file transfer that initiated from the Industrial Zone.

1. A manual file transfer is initiated from the Industrial Zone. An industrial user connects to the Secured File Transfer Gateway in the IDMZ.
2. In case of an automated transfer, a file server in the Industrial Zone connects to the gateway.
3. The user is authenticated on the Secure File Transfer Gateway and the file is transferred, inspected and saved. The file transfer is done via a secure encrypted protocol such as SFTP or HTTPS.
4. The IDMZ firewall validates the file does not contain any known malware.
5. The enterprise user logs onto the Secure File Transfer Gateway and retrieves the file.
6. In case of an automated transfer, a file server in the Enterprise Zone retrieves the file.

If the file transfer is initiated from the Enterprise Zone, the process is reversed.

Data Brokering

Network and Security administrators use multiple tools to get their jobs done and the need for these tools is growing. From cloud to on-premises, big vendors to homegrown, these tools compete for precious access to a limited number of data feeds. As administrators feel the pressure to install, evaluate, and implement these tools, all while staying under budget, they are hindered by the constraints of the data consumers and exporters. Network and security administrators should not have to analyze their data using multiple tools built for specific protocols. Rather than burdening the customer's workflow to fit the needs of the data, the data should be groomed to fit the needs of the customers and their tools.

Cisco Telemetry Broker Overview

Cisco's Telemetry Broker has roots in the Secure Network Analytics UDP Director, which simply replicated UDP traffic to multiple destinations. The Cisco Telemetry Broker will build upon the successes of the UDP Director, while also creating a new market—the Telemetry Broker market. Cisco Telemetry Broker optimizes telemetry pipelines for the hybrid cloud. It vastly simplifies the consumption of telemetry data for customers' business critical tools by brokering hybrid cloud data, filtering unneeded data, and transforming data to a usable format.

Cisco Telemetry Broker provides several key functionalities that will address the growing concerns of our customers:

- **Brokering Data**—The ability to route and replicate telemetry data from a source location to multiple destination consumers and to quickly onboard new telemetry-based tools!
- **Filtering Data**—The ability to filter data that is being replicated to consumers for fine grain control over what consumers are able to see and analyze and save money by not sending data to expensive tools!
- **Transforming Data**—The ability to transform data protocols from the exporter to the consumer's protocol of choice and enable tools to consume multiple data formats!

Cisco Telemetry Broker IDMZ Architecture

Figure 2-27 Cisco Telemetry Broker IDMZ Architecture

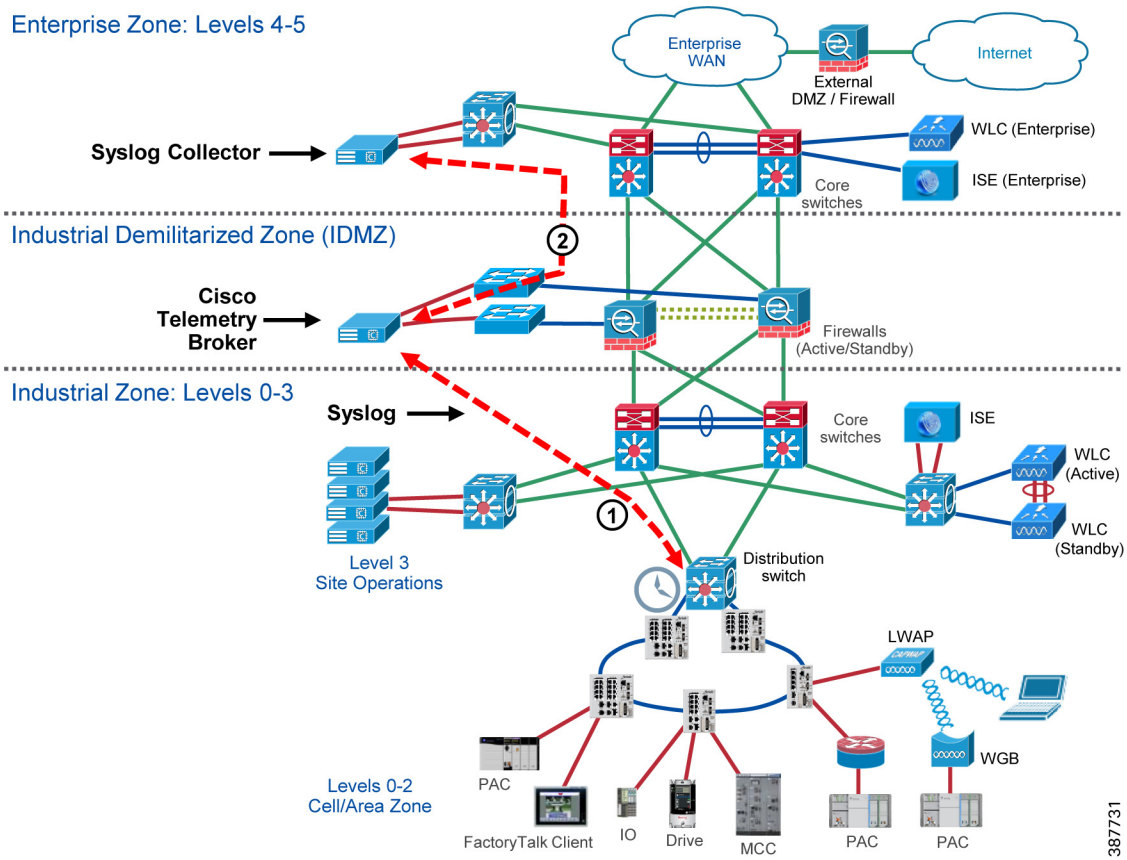

Cisco Telemetry Broker supports multi-node setups, where a single Cisco Telemetry Broker manager can manage multiple broker nodes. Broker nodes will exist in the IDMZ, as there role in the network is to broker data from the Industrial Zone. The manager node, if deployed to manage just the broker nodes for a single network, may also exist in the IDMZ. However, for deployment scenarios that span across multiple IDMZs, the manager node should exist in the Enterprise zone where it can be linked to all instances of the broker node across a distributed network architecture.

Remote Access Services

This section provides an overview of CPwE IDMZ remote access solutions and examples for FactoryTalk applications. Specific technologies include:

- SSL VPN access using the FTD platform
- Microsoft RD Gateway
- ThinManager via Microsoft RD Gateway
- Cisco Duo Authentication for Microsoft Remote Desktop Gateway

Remote Access Overview

Quick and effective response to issues on the plant floor often requires real-time access to information and status from IACS applications as well as the skills and knowledge to take corrective action or optimize the IACS process. Secure remote access to industrial assets, data and applications provides companies with the ability to apply the right skills and resources at the right time, independent of their physical location. Companies can use internal experts or the skills and resources of trusted partners and service providers, such as OEMs and system integrators, without needing someone onsite.

To deploy secure remote access, CPwE architecture includes a number of network services and technologies that are widely deployed in enterprise networks such as VPN, terminal services and web access portals.

Remote Access Design Principles

In the past, companies relied completely on onsite personnel to provide support for IACS applications, or used methods such as dial-up access and separate dedicated networks for remote support. These remote access methods have limited bandwidth and capabilities and are therefore limited to very basic monitoring and updating functionality. At the same time, they often circumvent perimeter security defenses and do not have the visibility and support of the IT organization. This creates the threat of "back doors" into the Industrial Zone and can represent a significant security risk. As manufacturers and partners want to provide more service and support remotely, and respond to issues in real time, these methods are no longer sufficient.

To truly leverage the full value of a converged enterprise, remote access needs to be scalable, regardless of location or company, and it needs to be done securely and in combination with the necessary tools to effectively communicate, diagnose problems and implement corrective actions. However, access needs to be limited to those individuals who are authorized to access systems, and their authorized actions need to be aligned to corporate and plant policies and procedures.

Several guiding principles should be maintained when allowing remote access to IACS data and resources:

- Use User Access and Authentication Policies and Procedures:
  - Access to resources and services should be monitored and logged.
  - Every user must be a known entity to the organization and use a unique account.

  - Users should be granted access to IACS data and resources based on the authorization policy on “as needed” basis.
  - Users should verify their identity using MFA.
  - Use of back-door solutions (such as modems, phone lines, and direct Internet access) may pose a significant risk and should be avoided.
  - Written policies should be implemented specifying under what conditions and who may be granted access into the secured Industrial Zone. Industrial personnel and trusted partners should sign a security agreement acknowledging their responsibilities.
- Control the Applications:
  - IACS protocols, such as CIP or FactoryTalk Live Data, should be contained to the Industrial Zone.
  - As a best practice, partners and remote engineers should use versions of IACS applications on controlled application servers in the Industrial Zone. By restricting remote users to applications running on a RAS, companies can enforce change management, version control and regulatory compliance of the applications being used.
  - This best practice prevents viruses or other compromises of the remote system from affecting the Industrial Zone applications and systems. The use of IACS applications on a remote user's computer introduces significant risk to the IACS and should be avoided.

- No Direct Traffic:
  - No direct traffic is permitted between the Enterprise Zone (including the Internet) and the Industrial Zone, with exception for certain highly controlled network services as outlined previously in this guide.
  - Remote access to devices on the IACS network should require connecting through the IDMZ firewall and logging into or at least proxying through a server.
- Only One Path In or Out:
  - The path from the IDMZ into the Industrial Zone should be the only path in or out. The path from the enterprise LAN into the IDMZ should be the only path connecting the two zones.

These guiding principles encapsulate the key concepts of strictly controlling the remote access of IACS applications rather than trusting that remote users are doing the right thing when accessing the IACS applications.

SSL VPN Access

The Firepower Management Center supports the following types of VPN connections:

- Remote Access VPNs on FTD devices.

  Remote access VPNs are secure, encrypted connections, or tunnels, between remote users and your company's private network. The connection consists of a VPN endpoint device, which is a workstation or mobile device with VPN client capabilities, and a VPN headend device, or secure gateway, at the edge of the corporate private network.

  FTD devices can be configured to support Remote Access VPNs over SSL or IPsec IKEv2 by the Firepower Management Center. Functioning as secure gateways in this capacity, they authenticate remote users, authorize access, and encrypt data to provide secure connections to your network. No other types of appliances, managed by the Firepower Management Center, support Remote Access VPN connections.

FTD secure gateways support the AnyConnect Secure Mobility Client full tunnel client. This client is required to provide secure SSL IPsec IKEv2 connections for remote users. This client gives remote users the benefits of a client without the need for network administrators to install and configure clients on remote computers since it can be deployed to the client platform upon connectivity. It is the only client supported on endpoint devices.

- Site-to-site VPNs on Firepower Threat Defense devices.

  A site-to-site VPN connects networks in different geographic locations. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers that comply with all relevant standards. These peers can have any mix of inside and outside IPv4 and IPv6 addresses. Site-to-site tunnels are built using the Internet Protocol Security (IPsec) protocol suite and IKEv1 or IKEv2. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

Client-based SSL VPN (Cisco AnyConnect)

The Cisco AnyConnect Secure Mobility Client is recommended for remote users who require full network connectivity. The Cisco AnyConnect client, which uses SSL, is designed for automated download and installation of the client software on the user's PC.

Other capabilities for the Cisco AnyConnect client include features that allow the client to reconnect if the tunnel goes down, to disable the tunnel if the client moves onto the trusted network or to bring up the tunnel if the client moves from a trusted to an untrusted network.

The Cisco AnyConnect VPN client provides secure SSL or IPsec (IKEv2) connections to the FTD for remote users with full VPN tunneling to corporate resources.

Without a previously installed client, remote users enter the IP address in their browser of an interface configured to accept SSL or IPsec-IKEv2 VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, remote users must enter the URL in the form https://address. After the user enters the URL, the browser connects to that interface and displays the login screen.

After a user logs in, if the secure gateway identifies the user as requiring the VPN client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure connection, and either remains or uninstalls itself (depending on the security appliance configuration) when the connection stops. In the case of a previously installed client, after login, the FTD security gateway examines the client version and upgrades it as necessary.

MFA—Cisco Secure Access by Duo

AAA servers enable managed devices acting as secure gateways to determine who a user is (authentication), what the user is permitted to do (authorization), and what the user did (accounting). Some examples of the AAA servers are RADIUS, LDAP/AD, TACACS+, and Kerberos. For Remote Access VPN on FTD devices, AD, LDAP, and RADIUS AAA servers are supported for authentication.

Duo MFA for Cisco FTD supports push, phone call, or passcode authentication for AnyConnect desktop and AnyConnect mobile client VPN connections that use SSL encryption. The Duo proxy server will receive incoming RADIUS requests from the FTD, contact your existing local LDAP/AD or RADIUS server to perform primary authentication if necessary, and then contact Duo's cloud service for secondary authentication.

Microsoft Remote Desktop (RD) Gateway

Remote Desktop (RD) Gateway, formerly Terminal Services Gateway, is an available option in the Remote Desktop Services server role included with Windows Server Operating Systems. A Windows Server with the RD Gateway role enabled allows authorized remote users and thin clients using ThinManager to connect to resources from an internal corporate or private network to assets in the Industrial Zone from any device that can run the Remote Desktop Connection (RDC) client.

RD Gateway uses the Remote Desktop Protocol (RDP) over HTTPS to establish a secure, encrypted connection between remote users and internal network resources. The remote desktop user via a thin client using ThinManager or Microsoft Remote Desktop Connection client will have access to the desktop and applications of the remote computer as if they are sitting locally and accessing the computers keyboard and mouse and viewing the local display.

Note For more information, refer to *Remote Desktop Services Overview* at:

- https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds-deploy-infrastructure

ThinManager

ThinManager by Rockwell Automation is a thin client and content delivery management platform that is purpose built around providing a safe and secure environment to mitigate risk in a connected industrial environment.

Note For more information, refer to ThinManager Product Profile at:

- https://thinmanager.com/profile/

Network segmentation is critical to maintaining a secure environment across multiple layers of a control network. A ThinManager server should be located inside the Industrial Zone and/or the Enterprise Zone, depending on the use case and business needs. This separates the Industrial Security Zone and the Enterprise Security Zone and does not permit any network traffic to traverse the zone without being redirected by the Remote Desktop Gateway. By placing ThinManager inside the Industrial Zone and/or the Enterprise Zone, no traffic is required to traverse directly across the IDMZ in order to deliver Remote Desktop Services content to requesting clients.

RD Gateway is a critical component of the deployment that will be required to deliver Remote Desktop Services content from the Industrial Security Zone to the Enterprise Security Zone or vice versa. This applies for Remote Desktop content that will be used on thin clients or mobile clients using WinTMC, aTMC, or iTMC.

RD Gateway Architecture

The RD Gateway is placed within the IDMZ and acts as the gateway to the Industrial Zone assets to enterprise users who wish to access these computers (see Figure 2-28).

Figure 2-28 Remote Desktop Gateway Architecture

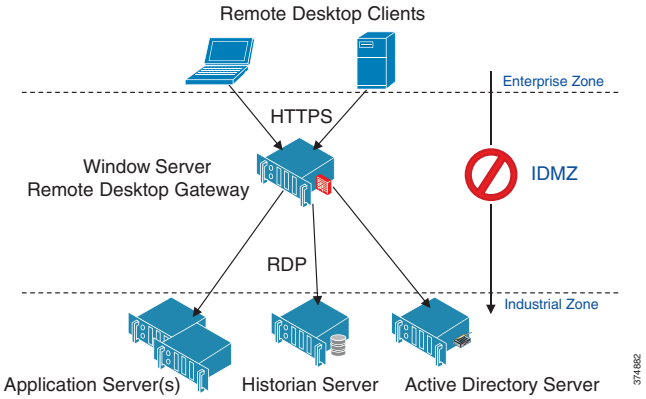

The RD Gateway uses two-factor authentication that verifies that a valid SSL certificate is being presented by the server and a valid user name and password is entered to authenticate the user's credentials.

The RD Gateway is designed to use HTTPS for the authentication process and initial connection establishment. Once the user is authenticated, the RD Gateway server connects to the requested Industrial Zone host via RDP. The firewall should be configured to allow HTTPS into the IDMZ from the Enterprise Zone and RDP from the remote desktop gateway to the Industrial Zone server(s).

RD Gateway Policies

The RD Gateway allows the administrator to configure **who** can connect to **what** through resource and connection policies.

- **RD Gateway Connection Authorization Policies (CAPs)** allow you to specify who can connect to the IDMZ RD Gateway server. The RD Gateway administrator can specify a user or user group that exists on the local RD Gateway server or in AD. The administrator can list specific conditions in each RD Gateway CAP, for example, you might require a group of users to use a smart card to connect through the RD Gateway.
- **RD Gateway Resource Authorization Policies (RAPs)** allow you to specify the Industrial Zone network resources that remote users can connect to through an RD Gateway server. When you create an RD Gateway RAP, you can use AD computer groups or single IP Addresses and associate it with the RD Gateway RAP.

Before CAPs and RAPs can be configured, the administrator should define user groups and computer groups for remote access and create security rules for those groups.

Microsoft Remote Desktop Gateway Configuration, page 3-32 has an example of defining a remote access rules and configuring the RD Gateway policies.

FactoryTalk Application Examples

Industrial Zone assets oftentimes require access to configure, maintain, and troubleshoot the process from outside the Industrial Zone. Security policies usually require that each user must be authenticated and their access to the Industrial Zone assets must be limited based on their credentials.

The following FactoryTalk applications can be accessible via remote access technologies:

- Studio 5000 Logix Designer®
- FactoryTalk AssetCentre

- FactoryTalk View Site Edition (SE)
- FactoryTalk ViewPoint
- FactoryTalk VantagePoint
- FactoryTalk Historian
- FactoryTalk Metrics

Each of these applications will have design and runtime programs that will need to be accessed by a remote user.

RD Gateway Access for FactoryTalk Applications

A solution that meets the application requirements listed above is to use a RD Gateway located in the IDMZ. Two variants of this solution are considered:

- Applications are installed and run on the terminal server in the Industrial Zone.
- Applications are installed and run on separate servers and accessed directly from the RD Gateway.

The choice of a solution depends on the existing deployment scheme, scale of operation, type of applications for remote access and whether a company chooses to implement more granular policies to restrict or control access.

RD Gateway Access to FactoryTalk Applications Installed on a Terminal Server

In this scenario, the required FactoryTalk design and runtime software is configured to run on the Terminal Server in the Industrial Zone. This scenario's workflow is described in Figure 2-29.

Figure 2-29 RD Gateway Access to FactoryTalk Applications Installed on Terminal Server

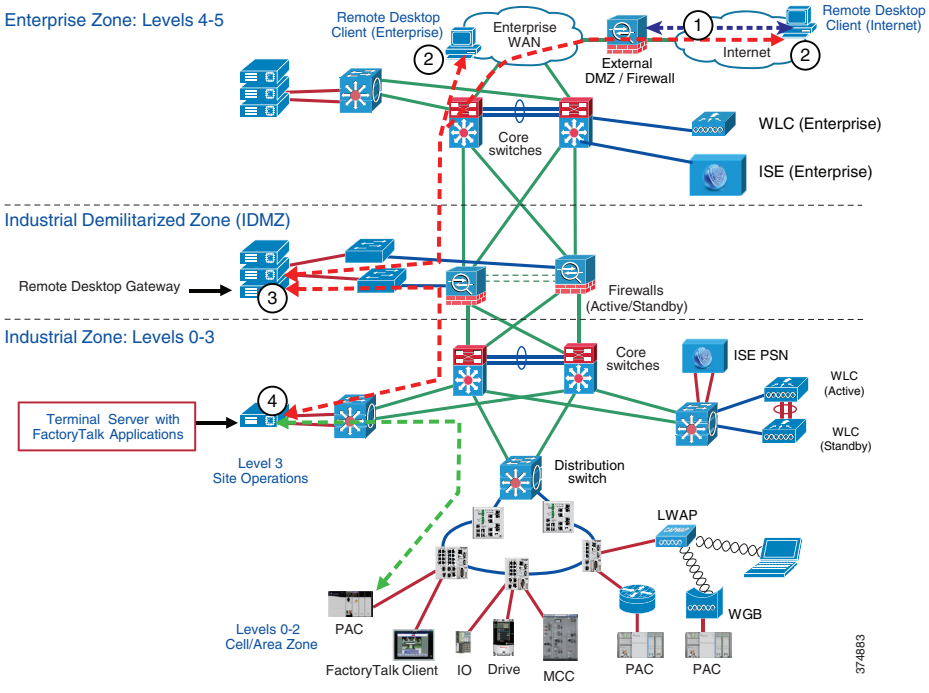

1. If the Remote Desktop client is outside the corporate network, a VPN session is established with the customer site.

2. The RDC application is launched from remote user's computer. The user enters Industrial Zone Remote Session Host's address (the Terminal Server running FactoryTalk applications) as the target desktop and starts the session.
3. The RD Gateway server in the IDMZ validates the SSL certificate and the username and password.
4. The Remote Session Host’s desktop is now presented to the remote desktop user.

Direct Access to FactoryTalk Applications via RD Gateway

The RD Gateway is capable of being configured to allow certain users such as production administrators or corporate engineers to have direct access to Industrial Zone assets for configuration, maintenance and troubleshooting purposes without going through a terminal server. A variation to the prior solution is to use a RD Gateway located in the IDMZ to access FactoryTalk and other Industrial Zone assets directly. This scenario's workflow is described in Figure 2-30.

Figure 2-30 Direct Access to FactoryTalk Applications via RD Gateway

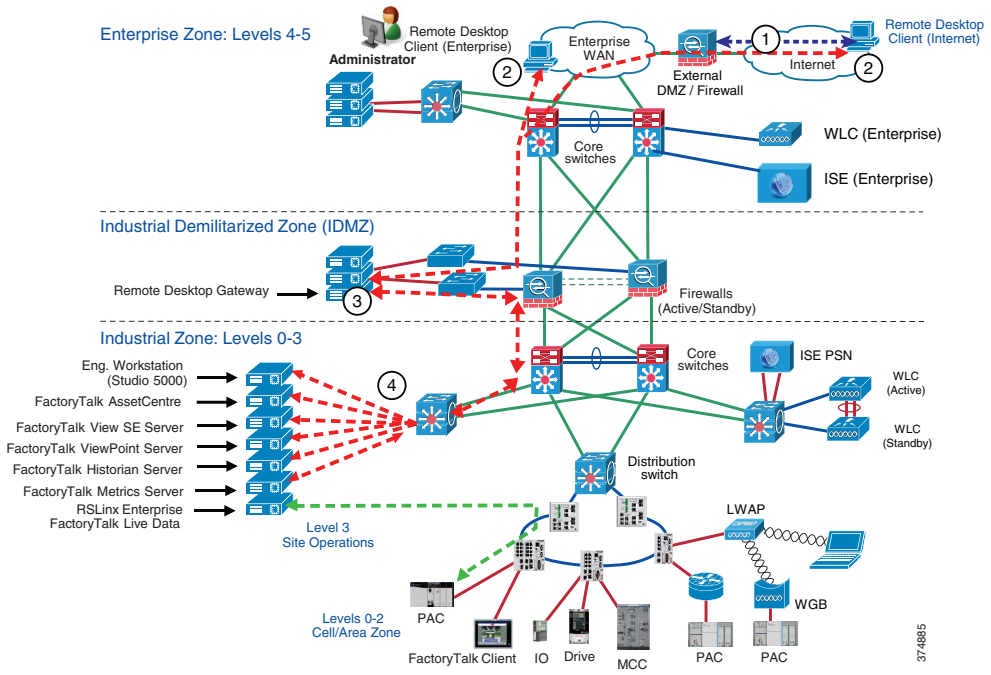

1. If the RD client is outside the corporate network, a VPN Session is established with the customer site.
2. Remote Desktop Connection application is launched from remote user's computer. The user enters Industrial Zone host's address as the target desktop and starts the session.
3. The RD Gateway server in the IDMZ validates the SSL certificate and the username and password.
4. The remote host’s desktop is now presented to the remote desktop user.

Examples of Non-Recommended Architectures

Previous examples described ways to access FactoryTalk applications remotely that comply with the security design principles for the IDMZ. Often companies try to deploy FactoryTalk applications across the IDMZ in a way that violates these principles, for convenience or cost saving purposes. This situation, which creates security risks, is strongly not recommended.

For example, Figure 2-31 shows an architecture where a FactoryTalk View SE client is installed in the Enterprise Zone and communicates to the FactoryTalk View server and FactoryTalk Directory server in the Industrial Zone. To enable this scenario, a wide range of ports needs to be open on the firewall, including DCOM dynamic port range.

This architecture does not align with IDMZ security design principles and is not recommended.

Figure 2-31 FactoryTalk View SE Client (Enterprise)—Not Recommended

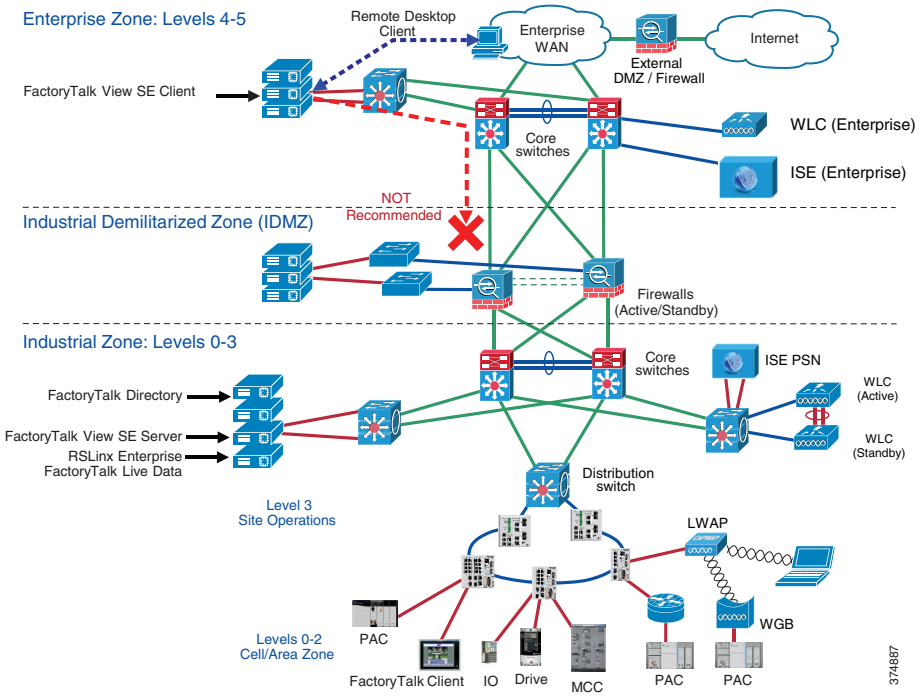

In another example (see Figure 2-32), FactoryTalk AssetCentre server and agents are placed in the Enterprise Zone. This scenario also requires multiple ports to be opened. In addition to that, agents will have direct access to the assets in the Industrial Zone (for example, PACs), which violates IDMZ security policy.

This architecture does not align with IDMZ security design principles and is not recommended.

Figure 2-32 FactoryTalk AssetCentre (Enterprise)—Not Recommended

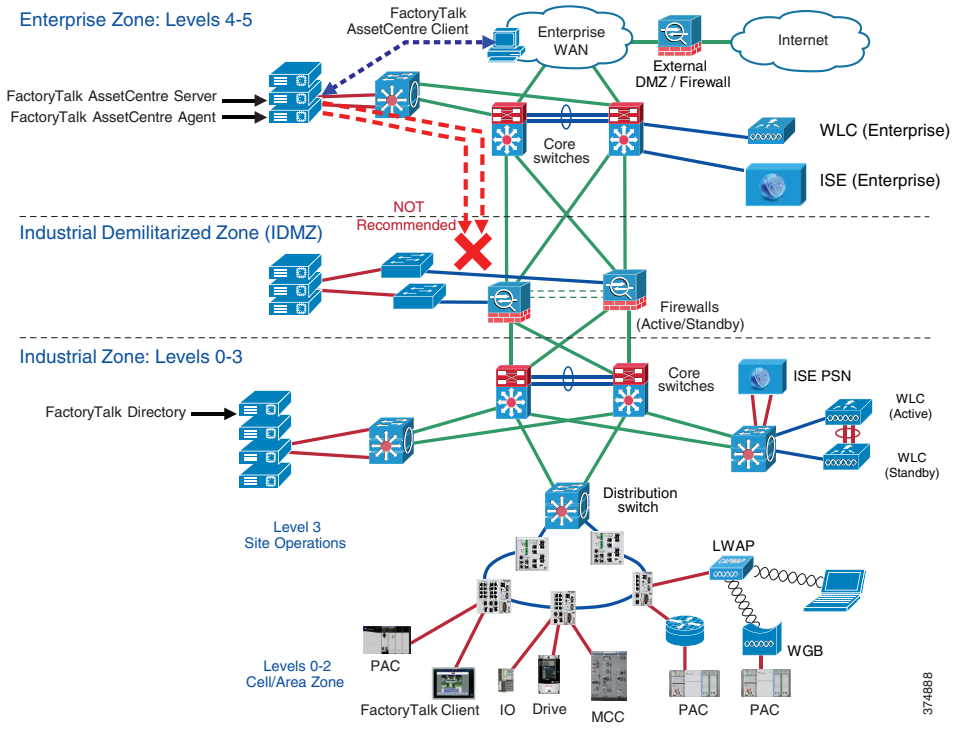

Application Security

Application security is the crucial part of the defense-in-depth security strategy. The components that provide application security may include:

- Application-level firewalls and proxies
- Application-level user authentication and authorization
- Application and OS hardening against threats such as code tampering, malware insertions, reverse-engineering and unauthorized use

The following sections review some of the application security methods:

- FactoryTalk Security
- Microsoft OS hardening

FactoryTalk Security

FactoryTalk Security is designed to provide a layer of application security. Its purpose is to protect against internal threats that are either malicious or accidental by limiting access to only those individuals who legitimately need access to specific automation assets.

FactoryTalk Security accomplishes this goal by allowing security administrators to define the answer to this question: “Who can carry out what actions upon which secured resources from where?”

- **Who** can use Rockwell Automation software products
  - ...to perform **what** specific actions

- ...on **which** Rockwell Automation hardware devices and other securable resources
- ...from **where** - that is, from which specific computers or workstations

How does FactoryTalk Security Protect the Application Layer?

When someone attempts to use a FactoryTalk-enabled software product to access a Rockwell Automation hardware device or other securable resource, FactoryTalk Security authenticates the person's identity and authorizes that person to access that resource and perform only allowed actions.

- **Authentication**—Verifies a user's identity and verifies that a request actually originated with that user.
- **Authorization**—Verifies a user's request to use a software product or to access a hardware device or secured resource against a set of previously defined access permissions.

FactoryTalk Security allows centralized administration of user accounts and access permissions. Security information, including user authentication and authorization, can be shared across all software products and hardware devices on a particular computer, throughout a plant or across an entire enterprise.

In the Windows domain environment, FactoryTalk Security accounts can be linked to the AD accounts and groups, which allows single identity for employees.

Note For further details about FactoryTalk Security, see the *FactoryTalk Security System Configuration Guide* at:

- http://literature.rockwellautomation.com/idc/groups/literature/documents/qs/ftsec-qs001_-en-e.pdf

An example of how to configure FactoryTalk Security is given in FactoryTalk Security Configuration, page 3-49.

Operating System Hardening

Software vulnerabilities and exploits have become an everyday part of life. Virtually every product has to deal with them and consequently users are faced with a stream of security updates. Security mitigation technologies are designed to make it impossible or more difficult for an attacker to exploit vulnerabilities in a given piece of software.

Rockwell Automation supports Microsoft AppLocker® as an OS hardening solution.

Note Full description and implementation guides to these solutions can be found on the Rockwell Automation knowledge base site, with a valid support center account:

- 546989—*Using Rockwell Automation Software Products with AppLocker*
  - https://rockwellautomation.custhelp.com/app/answers/answer_view/a_id/546989

Microsoft AppLocker Overview

In order to further harden the desktop environments, Rockwell Automation and Cisco recommends restricting administrator credentials. Normally, running as a standard, non-administrative user is recommended as it limits configuration changes that can be made in the desktop environment. However, running as a standard user does not prevent the installation or execution of unknown or unwanted applications in your organization.

To meet these challenges, Microsoft introduced a new feature in Windows 7 and Server 2008 R2 called AppLocker. AppLocker allows you to specify which users or groups can run particular applications in your organization based on unique identities of files. If you use AppLocker, you can allow or deny applications by creating rules to allow or deny applications from running.

Note

- It is strongly recommended to define allow lists rather than deny lists.
- For a more detailed explanation of Microsoft AppLocker, refer to the technical reference at:
  - https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-defender-application-control/applocker/applocker-technical-reference