CHAPTER 4

CPwE Solution Design—Manufacturing and Demilitarized Zones

Overview

This chapter provides an overview and basic design considerations for the Manufacturing and Demilitarized zones of the CPwE architecture. This solution guide offers basic design and implementation guidance for these zones, with which Industrial Automation and Control Systems (IACS) networking personnel could design and deploy a basic implementation. Often, these zones are where Enterprise IT networking resources or hybrid Plant-IT resources are involved in the design, implementation, and maintenance. For more complex deployments, Cisco and Rockwell Automation recommend that either external resources or Enterprise IT networking experts are used for the design, implementation, and maintenance.

Manufacturing Zone

The Manufacturing zone contains all IACS networks, devices, and controllers that are critical to controlling and monitoring plantwide operations. Hierarchically, the Manufacturing zone includes Site Manufacturing Operations and Control functions (Level 3) as well as multiple Cell/Area zones (Levels 0 to 2).

To preserve smooth plantwide operations and functioning of the IACS application and IACS network, this zone requires clear isolation and protection from the Enterprise zone via security devices within the Demilitarized zone (DMZ). This insulation not only enhances security segmentation between the Enterprise and Manufacturing zones, but may also represent an organization boundary where IT and manufacturing organizational responsibilities interface.

This approach permits the Manufacturing zone to function entirely on its own, irrespective 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 DMZ connectivity disruption. As a best practice, Cisco and Rockwell Automation recommend that all manufacturing assets required for the operation of the Manufacturing zone should remain there. Assets include FactoryTalk as well as applications and services such as Active Directory, DNS, and WINS.

Level 3, Site Manufacturing Operation and Control, has a dedicated Level 3 IACS network within the Manufacturing zone and contains the IACS software, such FactoryTalk. Cisco and Rockwell Automation recommend assigning a unique IP subnet and virtual LAN (VLAN) to this Level 3 IACS network.

The FactoryTalk application servers connect to a dedicated multilayer access switch, which aggregates into the Layer-3 distribution switches. To provide redundant default gateways to the Cell/Area zones, use the Cisco Catalyst 3750 StackWise Layer-3 distribution switches. If standalone distribution switches are used, use Gateway Load Balancing Protocol (GLBP) or Hot-Standby Routing Protocol (HSRP) between the distribution switches. Standalone distribution switches are not addressed in this version of the *Design and Implementation Guide (DIG)*. These protocols provide Layer 3 failover and load-balancing capabilities that are important to ensure communications between the Level 3 IACS network and the Cell/Area IACS network in the event of network disruption. FactoryTalk application-server redundancy is not addressed in CPwE 2.0.

An example of software applications that would be deployed within the Level 3 IACS network includes the following:

- FactoryTalk Services Platform
  - Directory
  - Activation
  - Security
  - Diagnostics
  - Audit
  - Live Data
  - Alarms and Events
- Application Servers
  - Factory Talk View SE
  - FactoryTalk AssetCentre
  - FactoryTalk Historian
  - FactoryTalk Transaction Manager
- Engineering Workstation
  - RSLogix™ 5000/500/5
  - RSNetWorx™

Key functions and features of the CPwE architecture for the Manufacturing zone include the following:

- Interconnecting the various Cell/Area IACS networks
- Interconnecting the Level 3 Site Manufacturing Systems
- Providing network management and security services to the Level 0 to 3 systems and devices
- Endpoint protection

The key Manufacturing zone design topics covered in this chapter include the following:

- Traffic flow
- Component selection
- Topology

- Routing
- High Availability and Network Resiliency
- IP addressing
- Security
- IACS Software, such as FactoryTalk, positioning within the Manufacturing zone

Multicast management is not included as a function in the Manufacturing zone. Although multicast traffic is routable traffic in many application types (video and voice), the most prevalent IACS network traffic applications ensure multicast traffic is contained to the Layer 2 network or subnet/VLAN with TTL=1. For the scope of this *DIG*, multicast traffic is constrained to the Cell/Area zone. IACS applications using routable multicast traffic or Precision Time Protocol (PTP) solutions are not yet common, but in the future will require design and implementation of multicast routing capabilities.

The Manufacturing zone is analogous to the core and distribution network hierarchy levels of the campus network architecture. This section refers to and includes many of the network recommendations from the following campus design guides:

- Overall Campus

  http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cOverall_design.html
- High Availability

  http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cHi_availability.html

Demilitarized Zone

In the design of the industrial Ethernet network, one of the critical elements is to ensure the separation between the IACS network and the enterprise network. In terms of the Purdue Reference Model, this is the separation between Levels 0 to 3 and Levels 4 to 5. This separation is necessary because real-time availability and security are the critical elements for the traffic in the IACS network. The impact of 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 schedule, 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 plant firewalls and a DMZ between the Manufacturing and Enterprise zones to securely manage the traffic flow between these networks.

It is a requirement to share data and services between the Manufacturing and Enterprise zones. Many of the benefits of a converged manufacturing and enterprise network rely on real-time communication and transfer of data between these zones. Without plant firewalls and a DMZ, this sharing is not possible while maintaining the security of the IACS network and its IACS systems. The plant firewall:

- Enforces authentication of users trying to access data or services
- Strictly controls traffic flow
- Performs stateful packet inspection and intrusion detection/protection
- Provides security and network management support
- Terminates VPN sessions with external or internal users

- Provides Web-portal services to offer proxies services, such as remote desktop, to specific servers in the Manufacturing zone

DMZ offers a network on which to place data and services to be shared between the Enterprise and Manufacturing zones. The DMZ enables the principle of not allowing direct communication between the Manufacturing and Enterprise zones, while meeting the requirement to share data and services. With the deployment of a DMZ and plant firewall, attacks and issues that arise in one zone cannot easily affect the other zone. In fact, by temporarily disabling the DMZ and plant firewall, an IACS or IT network administrator can protect a zone from being attacked until the situation is resolved in the other zone.

The DMZ network design covers the following:

- DMZ components
- DMZ topology
- Firewall design and implementation considerations

Key Requirements and Considerations

This section outlines the general requirements and considerations for the DMZ network and Manufacturing zone IACS networks. The requirements generally follow the requirements listed for the overall solution in Chapter 1, "Converged Plantwide Ethernet Overview." An additional consideration of application and service composition was also added to highlight the need to identify what key network, security and application services will be replicated or located in the various zones.

Industrial Characteristics

Most manufacturing facilities have environmentally controlled areas for certain types of applications and IT-related infrastructure. The Manufacturing and Demilitarized zone applications and systems typically reside in these environments. This suggests that the environmental requirements of the Cell/Area IACS network typically do not apply to the Manufacturing and Demilitarized zone network infrastructure. An exception exists where the distribution devices (Layer-3 switches or routers) or firewalls may potentially need to reside closer to the Cell/Area IACS networks and therefore meet certain levels of extended environmental tolerance.

Interconnectivity and Interoperability

A key requirement of the Manufacturing zone is to interconnect Cell/Area zones with each other and the systems, devices, and applications that make up the Manufacturing zone. This interconnectivity is achieved by applying routers or Layer-3 switches with an appropriate routing protocol.

As with the Layer-2 protocols discussed in the Cell/Area zone, there are a number of protocols used for routing, availability, and resiliency in the Manufacturing zone that have both proprietary and standard implementations. These are considered and recommendations are made for use in various scenarios.

This chapter includes the consideration and evaluation of the following standard features and functions in the Manufacturing zone:

- Routing protocols
- Router resiliency protocols
- EtherChannel or Link Aggregation Control Protocol (LACP) for link resiliency
- Quality-of-Service (QoS)

The DMZ is required to be the one and only connection point between the Manufacturing and Enterprise zone. The DMZ allows interconnectivity, but is designed to strictly control the types of traffic and traffic flow as well as apply a variety of security concepts.

Real-Time Communication, Determinism, and Performance

Manufacturing zone systems and applications do not have the real-time communications considerations that apply to devices within the Cell/Area zone. Network availability is critical, but the sensitivity of the devices to network performance (for example, latency and jitter) is significantly reduced because they tend to be standard IT servers and workstations relying upon TCP communications. Essentially, latency and jitter can vary more widely without significant disruption to the applications and devices in this zone.

IACS Implicit I/O and Explicit messaging network traffic may traverse between Cell/Area zones through the Manufacturing zone distribution switches. For this reason, it is important to apply similar QoS designs from the Cell/Area zone to the Manufacturing zone distribution switches.

By design, no IACS traffic should traverse the DMZ. Although the plant firewalls should process traffic in a timely manner, there is no specific need to carry QoS or implement other specific real-time functions for features to the plant firewalls or DMZ functions.

Figure 4-1 CPwE Overall Architecture

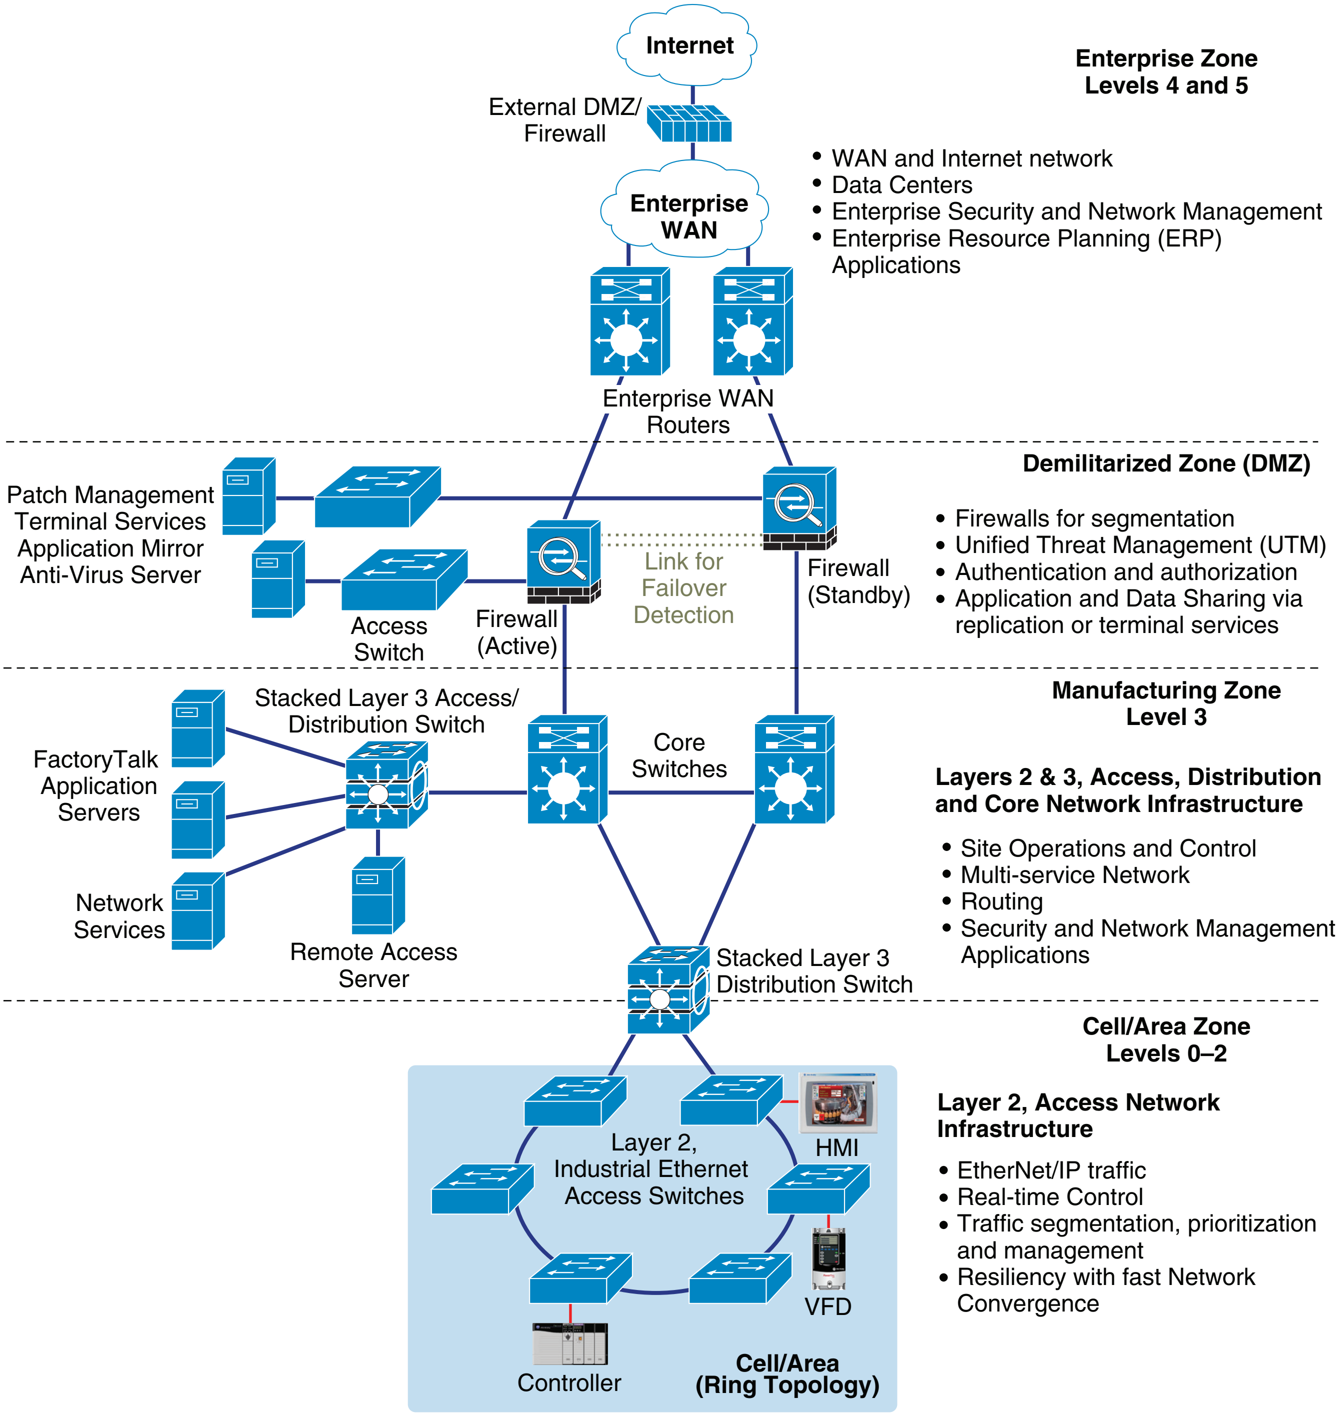

Availability

Availability of the network services is critical. Although the applications and services in the Manufacturing and Demilitarized zones may be more tolerant to network outages than the real-time communications in the Cell/Area zone, it is crucial that they stay available to maintain the operations in the Cell/Area zone. Without the services of the Manufacturing zone or DMZ, the IACS application may stop or be required to stop for compliance or regulatory reasons. Considerations discussed later in this chapter include the following:

- Equipment choice—Many aspects of the network infrastructure equipment impact the level of availability they will provide. In summary, these include the following:
  - Ease and speed of replacement features to reduce impact of a failure and reduce overall mean-time-to-repair (MTTR).
  - Support for network features and functions related to overall availability (e.g., resiliency protocols supported).
- Eliminate single points-of-failure in the network infrastructure, especially devices in critical roles (e.g., having redundant distribution and core switches).
- Multiple paths in the network uplink cabling from the Cell/Area zone access switches to the distribution switches, from the distribution switches to the core switches, and from the core switches to the plant firewalls.
- Resilient network protocols in place to meet application requirements.
- Applying a QoS approach to protect and prioritize key IACS network traffic.
- Segmentation to limit the impact of a failure or breach.

Security

The convergence of manufacturing and enterprise networks provides greater access to manufacturing data, which allows manufacturers to make more informed real-time business decisions. This business agility provides a competitive edge for manufacturers that embrace convergence. Convergence also calls for evolved security policies for IACS networks, which no longer remain isolated within a manufacturing area. Manufacturing computing and controller assets have become susceptible to the same security vulnerabilities as their enterprise counterparts. A security policy needs to protect manufacturing assets. This security policy needs to balance requirements such as 24x7 operations, low MTTR and high overall equipment effectiveness (OEE). Securing manufacturing 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 manufacturing assets. To address these obstacles, Cisco and Rockwell Automation recommend that manufacturers develop a manufacturing security policy, distinct from the enterprise security policy, based on the following considerations:

- Manufacturing operation requirements
- Enterprise security policy best practices
- Risk assessment
- A holistic security policy based on the defense-in-depth approach
- Industry security standards such as ISA-99

- Manufacturers' corporate standards
- Segmented Manufacturing IACS Network Security Framework
- A rigorous and well-documented patch management process

Manufacturing Security Policies

The key to a successful security strategy is understanding the potential problems that need to be solved, such as what to protect and how. Establishing a security policy focused on manufacturing needs provides a roadmap for applying security technologies and best practices to protect manufacturing assets, while avoiding unnecessary expenses and excessive restrictive access. Security services should not inhibit nor compromise the manufacturing operation.

As defined by ISA-99, a security policy "*enables an organization to follow a consistent program for maintaining an acceptable level of security.*" The security policy consists of both physical and electronic procedures that define and constrain behaviors by both personnel and components within the manufacturing system. A team consisting of IT, operations, and engineering professionals should work together to define manufacturing security needs. Security policy development starts with evaluating potential risks. Conducted by either an internal or external team, the risk assessment process identifies potential vulnerabilities and determines mitigation techniques through procedures and/or technology. For example, a procedure could restrict physical manufacturing systems access to authorized personnel. Technology mitigation techniques could involve change management software to authorize and authenticate user credentials.

Since security policies traditionally remained in the IT domain, IT has developed best practices to help identify and mitigate security vulnerabilities. Manufacturers can apply many of these policies and best practices to manufacturing as long as they account for differences between the needs of manufacturing applications and enterprise applications.

CPwE outlines general recommendations for deploying a holistic policy to help secure manufacturing assets. Many of the security requirements for the Cell/Area zone also apply to the Manufacturing zone. But, as the Manufacturing zone has some specific functions, those functions also need security considerations as well.

Security for the Manufacturing zone is covered in the "Manufacturing Zone IACS Network Design" section on page 4-10. Also, see security considerations including the "IACS Network Security Framework" section on page 6-13.

The DMZ and plant firewalls are important security considerations. Their key purpose is to securely provide interconnectivity to shared data and services between the Manufacturing and Enterprise zones.

Manageability

The systems and applications in the Manufacturing zone are typically administered and maintained by people with a focus on plant floor operations, not IT. Although more technologies that are standard will be applied to manage the network resources, they need to be easy to implement and use.

The DMZ and the plant firewalls typically require a level of security understanding that is rare in plant personnel. Therefore, the DMZ tends to be managed and supported by IT personnel. Tools are available and considered in the DMZ Network Design/Components.

Scalability

Plant floors come in a large variety of sizes. The Manufacturing zone IACS network in particular has to be flexible and robust to support this variety of sizes. To address plant scalability, Cisco and Rockwell Automation recommend the creation of multiple, smaller Cell/Area zones as building blocks, with interconnection and aggregation into the Manufacturing zone. The Manufacturing zone needs to scale up or down depending on those requirements.

The IACS network may include only a small number of devices (up to 50) to multiple 10,000s of devices. The IACS network solution architecture concepts and recommendations need to be applicable to that range, noting the considerations for various sizes. This version of the CPwE solution architecture focuses on basic concepts, tested in typical small-to-medium network installations.

The key features of the network design that enable scalability include the following:

- Topology
- Routing
- IP addressing

Scalability considerations for the DMZ usually involve volume of traffic handled and number of external synchronous users supported. These considerations are included in the DMZ composition. A DMZ typically does not have to support larger numbers of ports as the single point of interconnectivity.

Composition

For the DMZ, the key consideration is around which data and services need to be shared between the Manufacturing and Enterprise zones. Cisco and Rockwell Automation recommend that network developers carefully consider which applications, data, and services are considered part of the Manufacturing and Demilitarized zones.

The following are some of the key points to consider:

- How long can operations continue without these services?
- Must this service be configured specifically for the Manufacturing zone?
- How does the application and data need to interface with the Enterprise zone?
- What are the costs/complexities associated with either replicating or adding redundant services to the Manufacturing zone or DMZ that may also exist in the Enterprise zone?
- What are the security risks involved with placing the application or service into other zones and subsequent modification to the traffic flows?

Table 4-1 lists some of the key applications and services to consider.

Table 4-1 Key Applications and Services

Type	Critical	Optional
Manufacturing applications	· Historian · Asset management and security · Production floor visualization, monitoring and reporting · IACS application and network management and maintenance	· Manufacturing execution system · Batch Management
Network and security management	· Network management · Security management · Security monitoring, analysis, and response	
Common network-based Services	· Directory and domain services provide application security to Manufacturing zone applications · IP address allocation (for example, DHCP or BootP); if dynamic allocations services are used, this will be required · Dynamic Name Services—Although some IACS network devices do utilize dynamic names, most are applied with hard-code IP addresses. If dynamic names are used, a DNS service is required and is likely in addition to the DNS offered by IT services in the Enterprise zone. · Network Time Protocol (NTP) servers are required to coordinate clocks in various IACS applications, including to network infrastructure.	· Backup and restore—This function is commonly provided from the Enterprise zone, and for disaster recovery considerations, moving critical data off-site should be considered.

Manufacturing Zone IACS Network Design

This section outlines the following key requirements for an IACS network design.

Network Components

- Traffic flow—Flow of information between the various endpoints
- Network topology—Layout and orientation of the network equipment
- High availability and network resiliency
- IP addressing
- Routing
- Security

Manufacturing Zone Components

The Manufacturing zone consists of the following (see Figure 4-2):

- CPwE Level 3 IACS applications (such as FactoryTalk), workstations, and servers
- Depending on size and complexity, Layer-2 access switches to connect the CPwE Level 3 components
- Layer-3 switching and routing network infrastructure
- Network management applications
- Security management applications
- Endpoint security agent for endpoints with a common operating system (Linux and Microsoft)
- Remote Access Server (if being deployed) for remote personnel and partners to use to access Manufacturing zone applications

Figure 4-2 Manufacturing Zone Overview

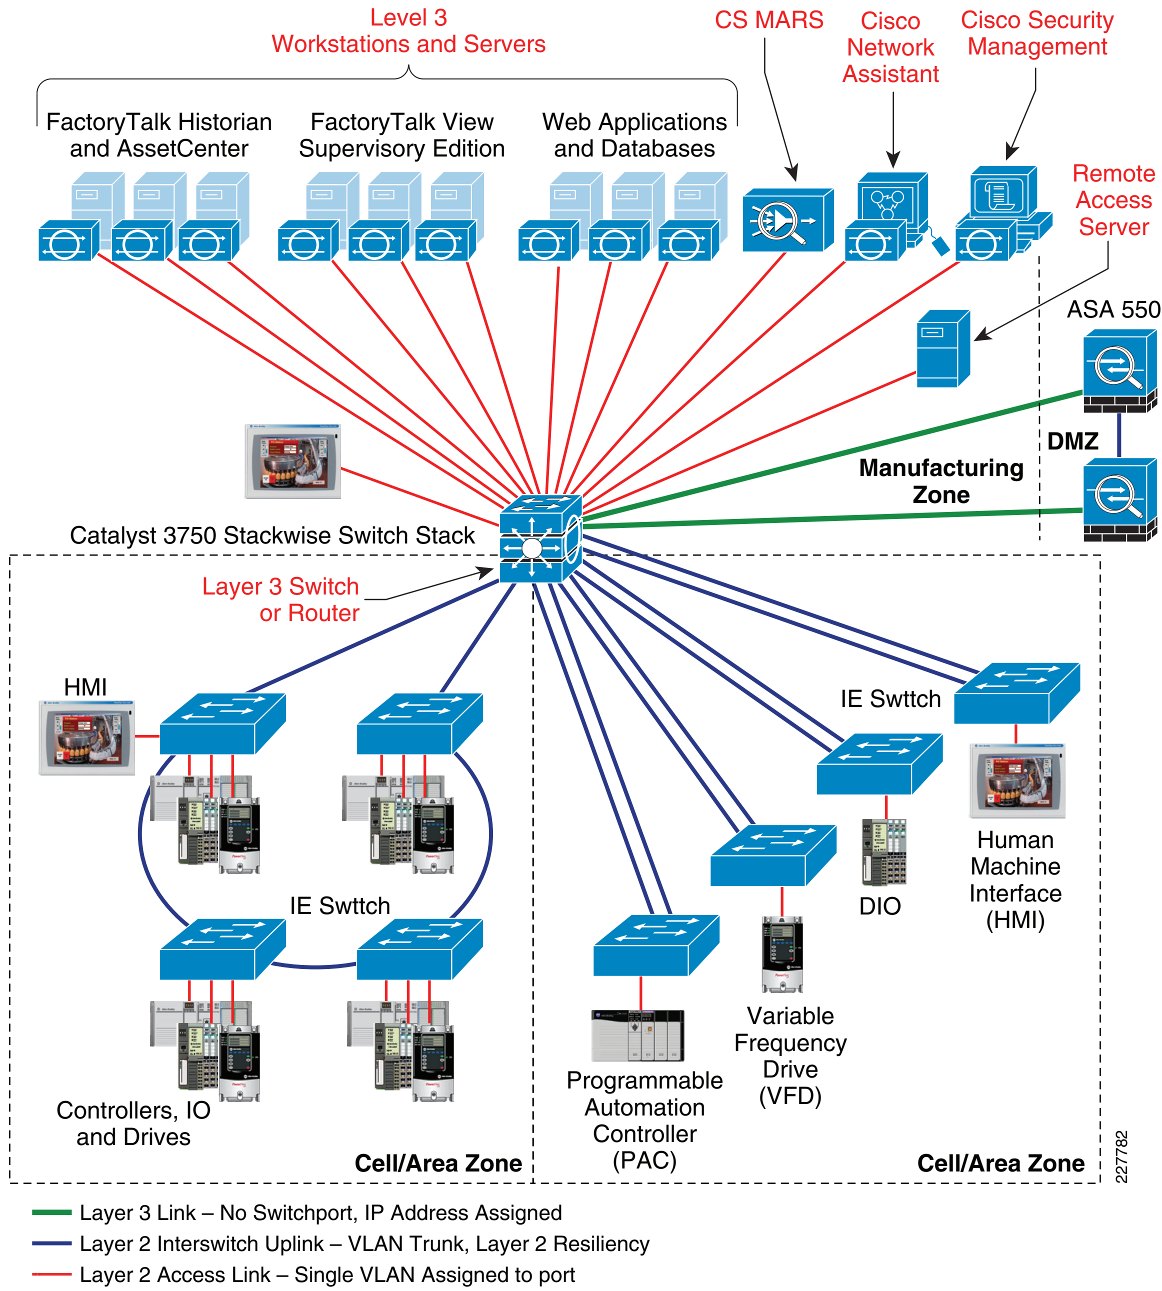
Level 3
Workstations and Servers
FactoryTalk Historian and AssetCenter
FactoryTalk View Supervisory Edition
Web Applications and Databases
CS MARS
Cisco Network Assistant
Cisco Security Management
Remote Access Server
ASA 5500
DMZ
Manufacturing Zone
Catalyst 3750 Stackwise Switch Stack
Layer 3 Switch or Router
HMI
IE Swttch
IE Swttch
Controllers, IO and Drives
Cell/Area Zone
Programmable Automation Controller (PAC)
Variable Frequency Drive (VFD)
DIO
Human Machine Interface (HMI)
Cell/Area Zone
227782
🟩 Layer 3 Link – No Switchport, IP Address Assigned
🟦 Layer 2 Interswitch Uplink – VLAN Trunk, Layer 2 Resiliency
🟥 Layer 2 Access Link – Single VLAN Assigned to port

This *DIG* does not provide guidance about the selection, design, or the implementation of the actual CPwE Level 3 Site Manufacturing Operations and Control equipment, workstations, servers, or the media used to connect the devices and switches.

The following are the Cisco components used in the Manufacturing zone:

- Optional Layer-2 access switches
- Layer-3 switching or routers
- Network management application
- Security management, analysis, and response applications

- Endpoint security for standard operating system workstations and servers (for example, Microsoft Windows and Linux)

Rockwell Automation application software examples that would be deployed within the Level 3 IACS network include the following:

- FactoryTalk Services platform such as Directory, Activation, Security, Diagnostics, Audit, Live Data Alarms, and Events
- FactoryTalk Application Servers such as View SE, AssetCentre, Historian, and Transaction Manager
- Engineering tools such as RSLogix 5000/500/5

The key considerations for the components in this zone are described in the following subsections.

Cost

Although cost is always a consideration in manufacturing facilities, the applications and devices in this zone tend not to be replicated as often as, for example, the Layer-2 switches found in Cell/Area zones. Therefore, there is no similar managed versus unmanaged question as in the Cell/Area zone; managed equipment is used by default.

Industrial Characteristics

As stated above, the industrial characteristics for this Manufacturing zone are less critical because it is assumed that controlled environments exists for the equipment.

It is recognized, however, that there is a need in some plant floor environments for the Layer-3 switching/routing functions to exist in an industrial packaging and to operate in the same conditions. That requirement is not addressed in this version of this *DIG*.

Performance and Real-Time Communications

Although not quite as critical as the Cell/Area zone, it is important for the Level 3 IACS network infrastructure to support real-time communications functions. The critical Explicit message IACS network traffic may traverse the Level 3 IACS network infrastructure. Note the following considerations:

- Bandwidth supported on Layer-3 switches and router ports (typically up to 1 Gbps) and any Layer-2 access ports (typically up to 100 Mbps) and uplink ports (typically up to 1 Gbps)
- VLAN trunking and inter-VLAN routing support
- QoS support is required, especially as critical IACS network traffic may traverse the Level 3 network infrastructure
- Load balancing protocols supported (for example, Gateway Load Balancing Protocol)
- Multicast routing protocols supported (this feature to be included in future version of the solution architecture) for CIP Sync and PTP applications

Availability

The network infrastructure availability is directly related to the overall availability of the IACS application. Thus, availability considerations are important and include the following:

- Availability options, including a routing resiliency (for example, hot-standby router or virtual router redundancy protocols), in-service upgradability, redundant components (for example, dual-processors, power, cooling), and other failover options (for example, stackable switch technology)
- Mean-time to break/fix ratings
- Storm control and rate-limiting to protect the network and other devices from out-of-control network communications
- Support for routing convergence protocols
- Support for Layer-2 resiliency protocols, such as EtherChannel/LACP or Flex Links, from Level-3 access switches to distribution switches

Manageability

Network and security management services and endpoint security are a part of this Manufacturing zone. These applications must be relatively easy to install, configure, and operate by plant floor personnel. Key considerations for this equipment includes the following:

- Intuitive Web-based interfaces via secure connections (for example, HTTPS)
- Ease of installation and upgradeability
- Ease of configuration and auto-detect functions to find and interface with appropriate network/security infrastructure
- Intuitive summarization of network and security status with easy-to-use drill-down features for problem solving
- Ability to develop templates for security/network management and to apply those throughout the Manufacturing zone
- Built-in knowledge repositories to assist plant and Control Engineers during problem resolution
- Ability to securely enable access to plant floor personnel and partners

In addition to the actual network and security management applications, there are also manageability considerations for the network infrastructure, especially the Layer-3 switches and routers. Basic management functions such as initial configuration, break/fix, and monitoring need to be relatively easy. Key considerations include the following:

- SNMP capable—Most network device vendors support management via the Simple Network Management Protocol (SNMP)v3. SNMPv3 is available on the crypto version of the Cisco IOS.
- Ease of installation, setup, and maintenance—The IACS network infrastructure should be easy to install, setup, and maintain with its key functions monitored and managed by IACS applications.
- Web-based, intuitive user interfaces.
- Application interfaces (for example, XML support) to interface with other applications.
- CIP support—The ability for the equipment or application to interface with the IACS network for basic management and monitoring functions greatly eases overall use and ongoing maintenance.

Security

The Manufacturing zone contains a number of security components including the security monitoring and analysis, security management, and endpoint security. Beyond these aspects, the key security considerations for each network component within the Manufacturing zone include the following:

- Access control lists (ACLs) allow users to configure security policies into a switch
- Support for VLANs
- Secure Shell (SSH) switch OS access
- SNMPv3 support for encryption of this important protocol used to manage and monitor the network infrastructure
- Port-based security to prevent access from unauthorized devices, including the following:
  - Limit the number of allowed MAC addresses on a physical port
  - Limit the allowed MAC address range on a switch port
  - MAC address notification—Notification via SNMP when any MAC-based port-security violations occur, for example more than one MAC address on an IACS end-device port
- Control-plane policing for switches and routers—Protection of the Layer-3 protocols used to support various network services
- Authentication and access servers to manage network and application security

Component Summary

For the purpose of testing, the products listed inTable 4-2 were part of the Manufacturing zone.

Table 4-2 Components

Role	Product/Platform	Comments
Distribution switch	Cisco Catalyst 3750 Series • Cisco Catalyst 3750G-24TS-24 Ethernet 10/100/1000 ports and four Small Form-Factor Pluggable (SFP) uplinks • Cisco Catalyst 3750G-24T-24 Ethernet 10/100/1000 ports • Cisco Catalyst 3750G-12S-12 Gigabit Ethernet SFP ports • Cisco Catalyst 3750G-24TS-1U-24 Ethernet 10/100/1000 ports and four SFP uplinks, 1-rack unit (RU) height • Cisco Catalyst 3750G-48TS-48 Ethernet 10/100/1000 ports and four SFP uplinks	Provide redundant distribution and core routing functions to Cell/Area and Manufacturing zone traffic
Core Switch	Catalyst 3750 Series (see above) Catalyst 4500 Series Catalyst 6500 Series:	Optional in medium-to-large operations to provide core networking functions

Table 4-2 Components (continued)

Role	Product/Platform	Comments
Security monitoring, analysis, and response	Cisco Security Monitoring, Analysis and Response Solution (CS-MARS)	Monitors security events from switches, routers, firewalls, and endpoint agents
Endpoint protection	CSA	Security protection for standard OS devices
Firewall configuration and management	Cisco Adaptive Security Device Manager	Firewall and intrusion protection services. Manages traffic flows between manufacturing, DMZ, and enterprise zones.
Endpoint security management	Cisco Security Manager	Manages endpoint security agent configuration
Network management	Cisco Network Assistant	Performs basic network management

Switching and Routing

The Cisco Catalyst 3750 switch (shown in Chapter 3, "CPwE Solution Design—Cell/Area Zone") was selected because it provides the best mix of features, performance, and cost for small-to-medium manufacturing facilities. Key considerations included the following:

- Lower cost base
- Already established in this role at a number of customer manufacturers
- Provides sufficient Layer-3 switching/routing features for most small-to-medium facilities
- Provides easy-to-configure resiliency and scalability with the StackWise connectivity to form a *virtual* switch
- Flexibility to grow with the manufacturing facility by adding additional stackable units
- In-service swappable and upgradeable components

For more information, refer to the Cisco Catalyst 3750 Series Switches at the following URL:

http://www.cisco.com/en/US/products/hw/switches/ps5023/index.html

Figure 4-3 shows the Cisco Catalyst 4500 switches.

Figure 4-3 Cisco Catalyst 4500 Series Switches

Figure 4-4 Cisco Catalyst 6500 Series Switches

An option that was strongly considered and is still believed to be a good option for medium-to-larger manufacturing facilities is the Cisco Catalyst 4500/6500, for the following reasons:

- Capacity or scalability is a concern; for example, when integrating a large number of Cell/Area IACS networks and CPwE Level 3 workstations and servers
- Need for a higher density of fiber ports
- Support for dual processors, cooling, and power
- Upgradeable processor and interfaces for longer-term viability
- Better failover features for availability; for example, in-service upgradeability

For more information, see the Cisco Catalyst 4500 Series Switches or Cisco Catalyst 6500 Series Switches at the following URL:
http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html

Figure 4-5 CS MARS

Security Monitoring, Analysis, and Response

The entry-level Cisco Security MARS (CS-MARS) appliance was selected. A wide variety of appliances is available that support increasing levels of events and network flow. CS-MARS is useful to simplify the security monitoring and response required to maintain a secure IACS network. CS-MARS provides the following capabilities:

- Identifies threats "learning" the topology, configuration, and behavior of the converged architecture with network behavior analysis and correlation technologies
- Makes precise recommendations for threat mitigation, including the ability to visualize the attack path and identify the source of the threat with detailed topological graphs that simplify security response at Layer 2 and Layer 3.

- Simplifies incident management and response with actionable E-mail incident notification, built-in case management, and the ability to configure firewall rules and intrusion prevention system (IPS) signatures through integration with Cisco Security Manager.

For more information, see the CS-MARS product overview at the following URL: http://www.cisco.com/en/US/partner/products/ps6241/index.html

For manufacturers interested in deploying CS-MARS in a number of manufacturing sites, global controller units are available, although this version of CPwE does not cover this case.

Endpoint Security

Cisco recommends the deployment of Cisco Security Agent (CSA) on the workstations and servers running common operating systems.

CSA security software provides the following:

- Threat protection for server and desktop systems
- Industry-leading defense against targeted attacks, spyware, rootkits, and day-zero attacks
- Proactive protection is offered against unknown, never-seen-before threats, brand new exploits, and variants trying to take advantage of recently announced vulnerabilities
- “Zero update” system integrity protection for critical servers that cannot be taken out-of-service to apply operating system or application-specific vulnerability patches. This greatly reduces the need for emergency patching of systems to respond to vulnerability announcements, minimizing patch-related downtime and plant man-hour expenses. Plants can patch on their own schedule, not in crisis mode, with a CSA deployment.
- Ability to integrate with CS-MARS and Cisco's intrusion detection and prevention solutions to mitigate and thwart complex attacks against IACS networks and devices.

Cisco recommends sufficient testing and “learning” is conducted with any CSA deployment. CSA is typically installed and operated in “learning” mode for a period of time to determine the base operational behaviors of the system. Once this phase is complete, CSA can be put in “restrictive” mode once policies have been established.

For more information, see the CSA product website at the following URL: http://www.cisco.com/en/US/partner/products/sw/secursw/ps5057/index.html

For information on CSA Management Center, refer to the following URL:

http://www.cisco.com/en/US/products/sw/secursw/ps5057/products_tech_note09186a0080769226.shtml

Network Management

The Cisco Network Assistant (CNA) is recommended to perform the network management functions for the Manufacturing zone. CNA supports up to 40 Cisco network devices, including the Stratix 8000, which meets the needs of the small-to-medium manufacturer. Key features include the following:

- No cost, downloadable at http://www.cisco.com/go/cna
- Configuration and maintenance of the network infrastructure devices via easy-to-use Web-based graphical user interface
- Inventory reports
- Event notification

- Task-based menu
- Software upgrades and operating system maintenance including IOS File management

For more information on CNA, refer to the following documents:

- CNA Overview— http://www.cisco.com/en/US/products/ps5931/index.html
- Getting started with CNA— http://www.cisco.com/en/US/partner/products/ps5931/products_getting_started_guide_book09186a00802b3c41.html

CiscoWorks is suggested as an option for more sophisticated and involved network management, such as the following:

- Multi-vendor network infrastructure must be supported (via SNMP)
- Cross-manufacturing site management is a current or future requirement
- More than 40 network devices at one site need to be managed
- CiscoWorks provides portfolio of network management. For more information, see the following URL: http://www.cisco.com/en/US/products/sw/netmgtsw/index.html

Security Management

Cisco and Rockwell Automation recommend the deployment of the Cisco Adaptive Security Device Manager to manage the firewalls in the DMZ, including the Adaptive Security Appliance (ASA). Key features include the following:

- Intuitive, easy-to-use web-based management interface to implement the DMZ, establish remote access and configure the firewalls
- Robust administration tools, real-time log viewer and monitoring dashboards that provide at-a-glance view of firewall appliance status and health
- Troubleshooting features such as packet trace and packet capture, providing administrators powerful debugging tools

For more information, see the following URL:
http://www.cisco.com/en/US/products/ps6121/index.html

Cisco recommends the deployment of CiscoWorks Management Center for CSA (when deployed) to manage the CSA and the endpoint security solution. Key features include the following:

- Centralized monitoring and management of CSA endpoint instances
- Role-based, web browser, intuitive user interface
- 20 preconfigured default policies
- Allows users to work in an IDS mode for learning and alerting (versus blocking)
- Allows for customizations to the policies and easy deployment to the agents

For more information, see the following URL:
http://www.cisco.com/en/US/products/sw/cscowork/ps5212/index.html

For network developers who are interested in more comprehensive security management solutions, Cisco and Rockwell Automation recommend considering the Cisco Security Manager, which incorporates the above applications. For more information, see the following URL: http://www.cisco.com/en/US/products/ps6498/index.html.

Traffic Flows

The traffic flows in the Level 3 IACS network resemble those of a decentralized client-server environment. Many of the CPwE Level 3 workstations, applications, and servers accomplish the following:

- Send detailed scheduling, execution, and IACS data to IACS controllers in the various Cell/Area zones
- Collect information from the Cell/Area IACS for historical and audit purposes
- Provide site-level operations management
- Perform application, network, and security administration and maintenance function for the overall Manufacturing zone, including the following:
  - Patch launch server
  - Remote access server
  - File server
  - Domain and Lightweight Directory Access Protocol (LDAP) services
  - Network and security management
- IACS reporting services (for example, cycle times, quality index, predictive maintenance) available to Manufacturing zone and via the DMZ to Enterprise zone users
- Provide data and services that will be shared through the DMZ to applications or users in the Enterprise zone

Traffic flows are outlined from the following two perspectives:

- IACS applications (for example, historian, asset management, IACS security, reporting)
- Network and security management

As with the Cell/Area zone, traffic from the Level 3 IACS network infrastructure protocols (for example, ARP and RPVST+) are not represented.

Figure 4-6 and Table 4-5 show the Level 3 IACS traffic flows to the Cell/Area IACS network.

Figure 4-6 Manufacturing Zone Traffic Flow—IACS Application

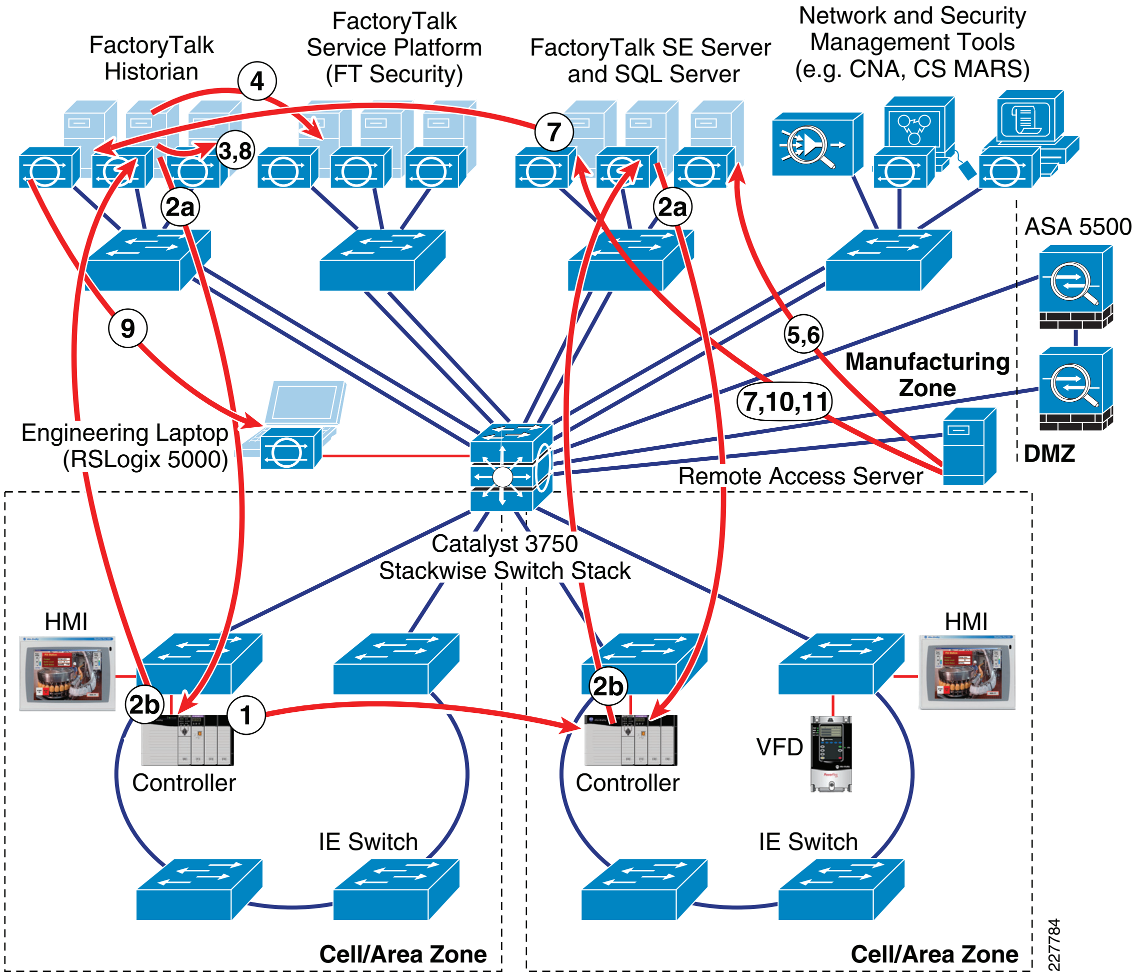
FactoryTalk Historian
FactoryTalk Service Platform (FT Security)
FactoryTalk SE Server and SQL Server
Network and Security Management Tools (e.g. CNA, CS MARS)
4
3,8
7
2a
2a
ASA 5500
9
5,6
Manufacturing Zone
7,10,11
Engineering Laptop (RSLogix 5000)
DMZ
Remote Access Server
Catalyst 3750 Stackwise Switch Stack
HMI
HMI
2b
2b
1
VFD
Controller
Controller
IE Switch
IE Switch
Cell/Area Zone
Cell/Area Zone
227784

Table 4-3 Manufacturing Zone Level 3 Traffic Flows

Ref. #	From	To	Description	Protocol	Type	Port(s)
1	Server	Cell/Area device	CIP diagnostic, configuration, information, uploads/downloads, and identification data. Example: a. FactoryTalk Historian or FactoryTalk View SE requests data b. Controller replies with data	EtherNet/IP	TCP/UDP	44818
2	Client/ server	Client/ server	FactoryTalk Transaction Manager	RPC	TCP	400–402

Table 4-3 Manufacturing Zone Level 3 Traffic Flows (continued)

Ref. #	From	To	Description	Protocol	Type	Port(s)
3	Client/ server	Client/ server	FactoryTalk Metrics—Production server	RPC	TCP	4120
			FactoryTalk Metrics—Server manager	RPC	TCP	4121
			FactoryTalk Metrics—Plant Metrics server	RPC	TCP	4122
			FactoryTalk Metrics—Task manager	RPC	TCP	4123
			FactoryTalk Metrics—Schedule server	RPC	TCP	4124
			FactoryTalk Metrics—Schedule CTP server	RPC	TCP	4125
4	Client/ server	Client/ server	FactoryTalk Service Platform support DCOM	Endpoint mapper	TCP	135
				DCOM	TCP	dynamic (1024-65535+)
5	Client/ server	Client/ server	FactoryTalk—Object RPC	rnaprpc	TCP	1330
			FactoryTalk—Service control	rnaserv	TCP	1331
			FactoryTalk—Server health	ranserverping	TCP	1332
			FactoryTalk—Directory server file transfer	rnadirft	TCP	3060
			FactoryTalk—Alarming server	rnaalarming	TCP	6543
			FactoryTalk—Event multiplexor		TCP	7600
			FactoryTalk—Event server		TCP	7700
			FactoryTalk—Directory server		TCP	7710
			FactoryTalk—License server		TCP	27000
6	Client/ server	Client/ server	FactoryTalk View SE—HMI server		TCP	7720
			FactoryTalk View SE—Server framework		TCP	7721
			FactoryTalk View SE—HMI Activation		TCP	7722
			FactoryTalk View SE—Historical data log reader		TCP	7723
7	Client/ server	Client/ server	FactoryTalk AssetCentre		TCP	1433
8	Client/ server	Client/ server	FactoryTalk AssetCentre	RPC	TCP	135
9	Server	Client-browser	FactoryTalk View SE and RSView 32	HTTP	TCP	80
10	Server	Client-browser	FactoryTalk Metrics—Reports and server manager	HTTP	TCP	8080 8081
11	Client	Mail server	FactoryTalk Metrics, FactoryTalk Transaction Manager, FactoryTalk View—Mail for event notification	SMTP	TCP	25

In summary, the traffic flow of the IACS application data depends on where the various clients and servers are placed within the framework (for example, DMZ or Manufacturing zone) to best support the required integration between the Enterprise and Manufacturing zones.

Topology Options Overview

The deciding factor in the design of the Manufacturing zone is the size and distribution of the IACS network. Large IACS networks need a more complex infrastructure to support the many Cell/Area zones. Small IACS networks can be simple with a single core/distribution switch for the entire Manufacturing zone. The following three different topology options have been developed based on the size of the IACS network:

- Small Manufacturing zone of up to 30-50 network infrastructure devices
- Medium Manufacturing zone of up to 200 network infrastructure devices
- Large Manufacturing zone of more than 200 network infrastructure devices

Small Manufacturing Zone Topology

The small Manufacturing zone topology includes a redundant pair of Layer-3 switches configured for redundancy (see Figure 4-7). All CPwE Level 3 IACS devices are connected directly to these switches. A set of Catalyst 3750 StackWise Layer-3 switches can support from 23 (two 12-port switches) to 468 ports (maximum 9 switches and maximum 48-port devices), so this configuration can support a small-to-medium plant. For the small Manufacturing zone topology, the Layer-3 switches provide inter-VLAN and inter-zone routing functions as well as Layer-2 connectivity to CPwE Level 3 workstations and servers.

The small Manufacturing zone topology essentially represents a collapsed core-distribution network routing services. This is representative for many small and medium manufacturing facilities.

Figure 4-7 Small Manufacturing Zone Topology

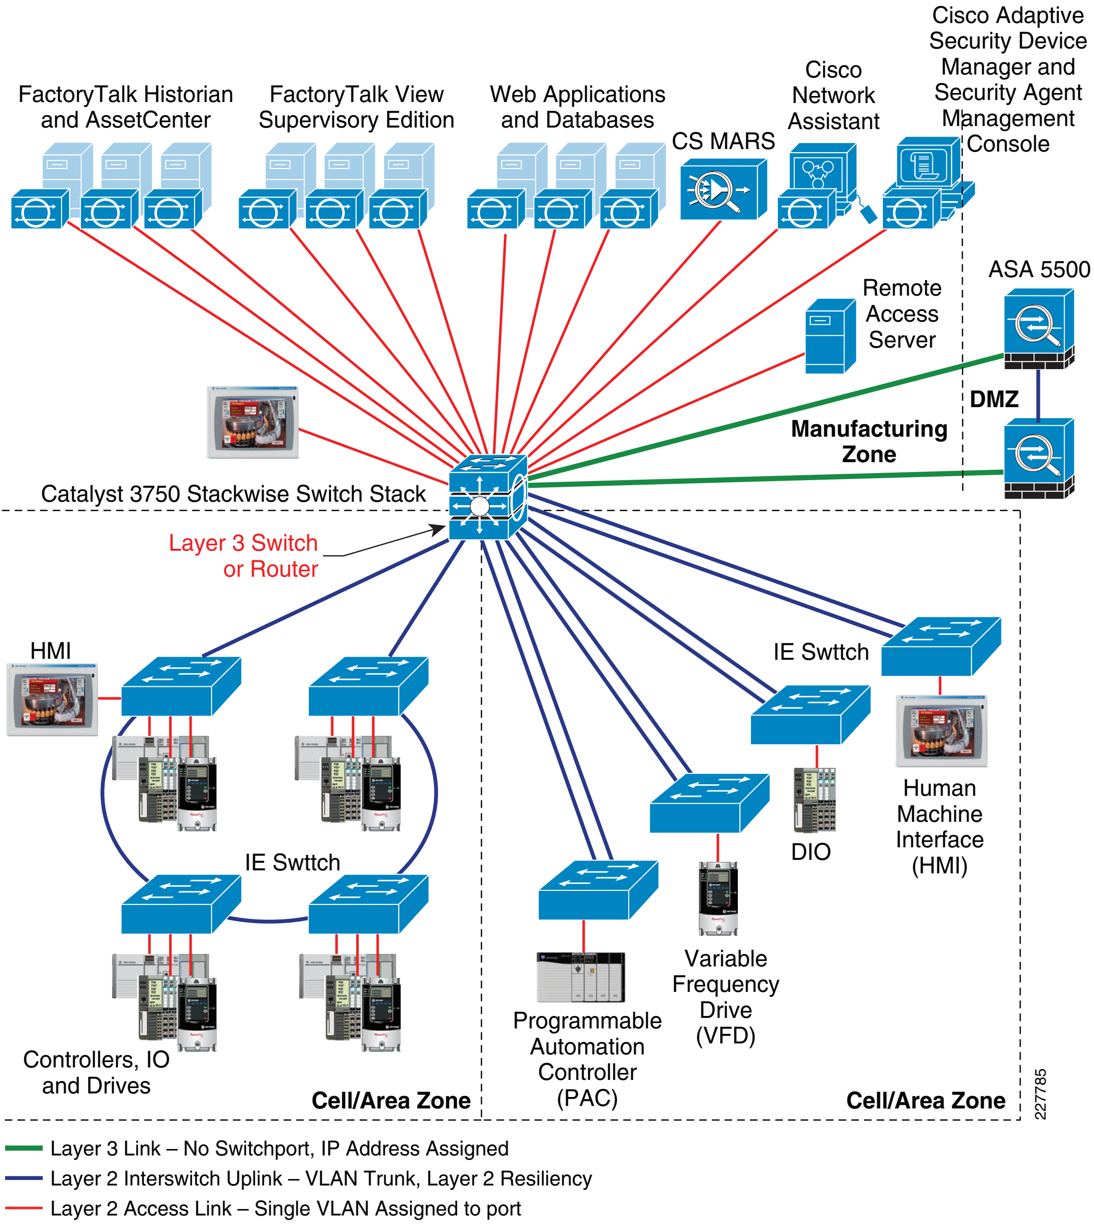
FactoryTalk Historian and AssetCenter
FactoryTalk View Supervisory Edition
Web Applications and Databases
CS MARS
Cisco Network Assistant
Cisco Adaptive Security Device Manager and Security Agent Management Console
ASA 5500
Remote Access Server
DMZ
Manufacturing Zone
Catalyst 3750 Stackwise Switch Stack
Layer 3 Switch or Router
HMI
IE Swttch
IE Swttch
DIO
Human Machine Interface (HMI)
Variable Frequency Drive (VFD)
Programmable Automation Controller (PAC)
Controllers, IO and Drives
Cell/Area Zone
Cell/Area Zone
227785
Layer 3 Link – No Switchport, IP Address Assigned
Layer 2 Interswitch Uplink – VLAN Trunk, Layer 2 Resiliency
Layer 2 Access Link – Single VLAN Assigned to port

Medium Manufacturing Zone Topology

The medium Manufacturing zone topology represents the separation of various network routing services and replication of these services to meet requirements in a larger manufacturing facility (see Figure 4-8). Although the small Manufacturing zone topology can support up to 200 network infrastructure nodes, there are situations even in this type of node count that may require a more segmented topology. The medium Manufacturing zone topology differs from the small Manufacturing zone topology as follows:

- Higher density, modular chassis-based switches
- Separate distribution switches enabling geographical separation

- Application of HSRP to provide routing resiliency between distribution switches

Figure 4-8 Medium Manufacturing Zone Topology

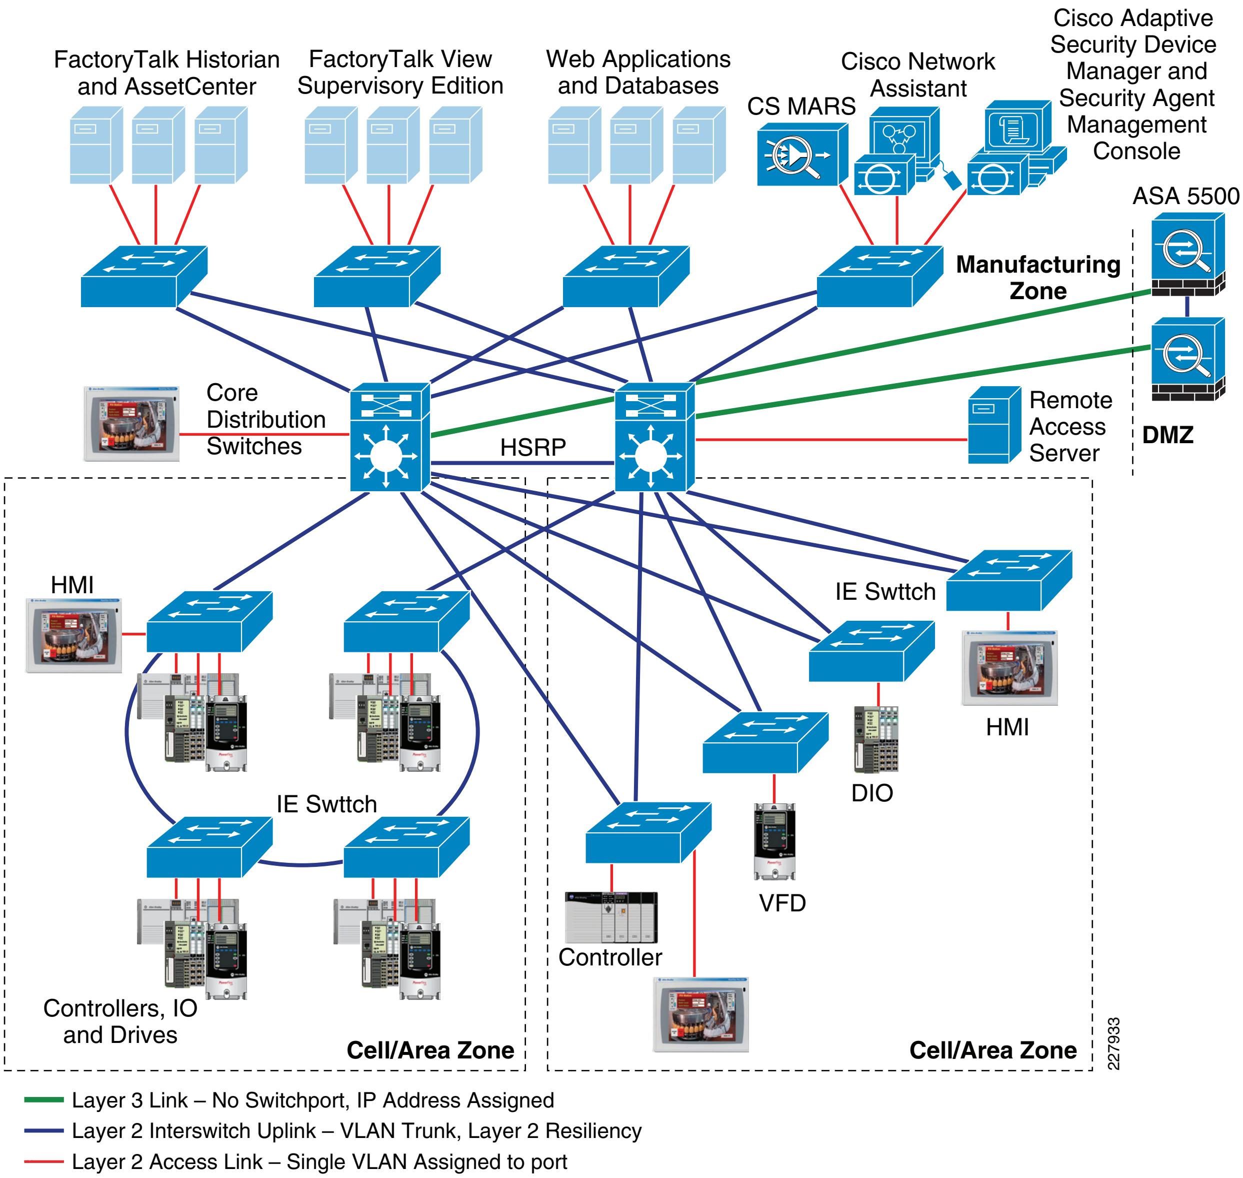
FactoryTalk Historian and AssetCenter
FactoryTalk View Supervisory Edition
Web Applications and Databases
CS MARS
Cisco Network Assistant
Cisco Adaptive Security Device Manager and Security Agent Management Console
ASA 5500
Manufacturing Zone
Core Distribution Switches
HSRP
Remote Access Server
DMZ
HMI
IE Swttch
IE Swttch
HMI
DIO
VFD
Controller
Controllers, IO and Drives
Cell/Area Zone
Cell/Area Zone
227933
Layer 3 Link – No Switchport, IP Address Assigned
Layer 2 Interswitch Uplink – VLAN Trunk, Layer 2 Resiliency
Layer 2 Access Link – Single VLAN Assigned to port

Large Manufacturing Zone Topology

The large Manufacturing zone topology represents the separation of various network routing services and replication of these services to meet requirements in a larger manufacturing facility (see Figure 4-9). Although the medium Manufacturing zone topology can more than 200 Ethernet nodes, there are situations even in this type of node count that may require a more segmented topology. The large Manufacturing zone topology differs from the medium Manufacturing zone topology as follows:

- Separate Layer-2 access switches to connect the CPwE Level 3 workstations and servers
- Additional stack of Layer-3 distribution switches for geographically distributed Cell/Area zones
- Additional pair of Layer-3 core routers to consolidate the Manufacturing zone traffic

Each of these enhancements can be implemented individually depending on the size and requirements of the IACS network. The following are some considerations for each of the scenarios:

- Separate Layer-2 access switch:
  - High-availability workstation or server environments may require redundant network connectivity to the workstations and servers. In these cases, Cisco and Rockwell Automation recommend having a separate Layer-2 access switch for the configuration of the relevant protocols. For more information, see "Server Farm" section on page 4-46.
- Adding a pair of Layer-3 distribution and core switches/routers:
  - Cell/Area zones in the plant are geographically distant from one another, where the wiring cost and complexity outweigh the cost and complexity of adding the additional pair.
  - Adding the additional pair for geographical reasons requires separate core and distribution switch/router pairs to manage the redundant interconnectivity between the DMZ, CPwE Level 3 workstations and servers, and other Cell/Area zones.

Figure 4-9 shows the resulting topology.

Figure 4-9 Large Manufacturing Zone Topology

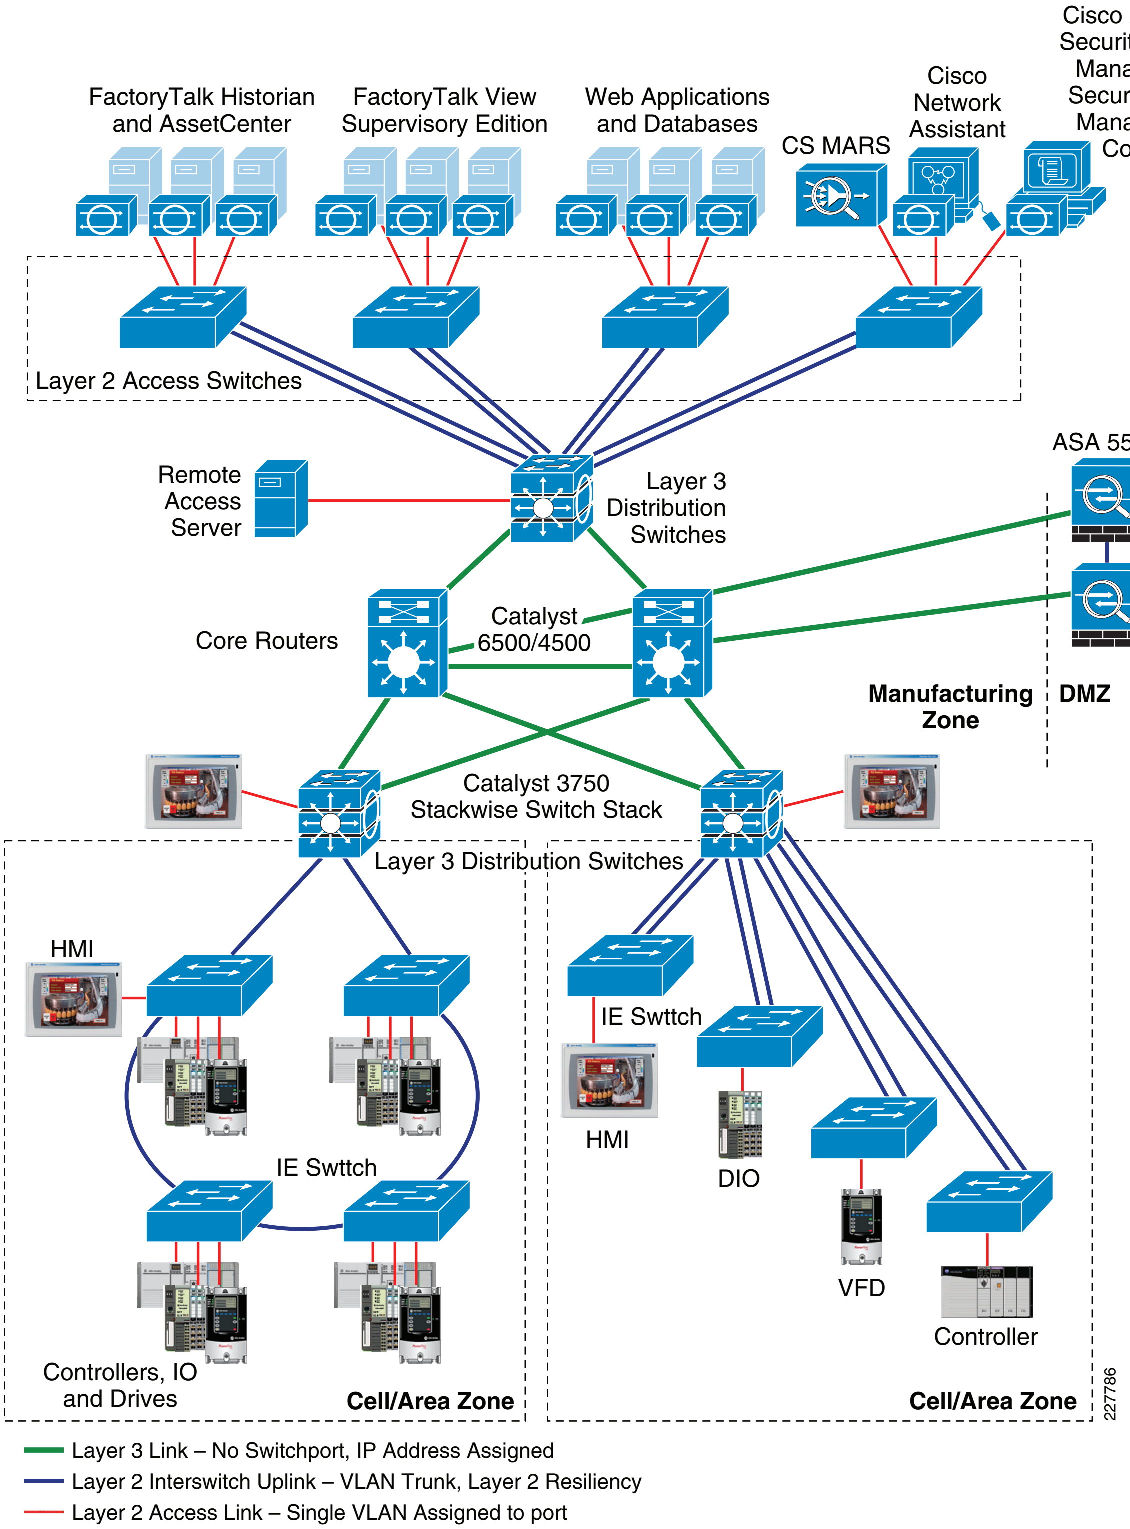
FactoryTalk Historian and AssetCenter
FactoryTalk View Supervisory Edition
Web Applications and Databases
CS MARS
Cisco Network Assistant
Cisco Adaptive Security Device Manager and Security Agent Management Console
Layer 2 Access Switches
Remote Access Server
Layer 3 Distribution Switches
ASA 5500
Core Routers
Catalyst 6500/4500
Manufacturing Zone
DMZ
Catalyst 3750 Stackwise Switch Stack
Layer 3 Distribution Switches
HMI
IE Swttch
IE Swttch
HMI
DIO
VFD
Controller
Controllers, IO and Drives
Cell/Area Zone
Cell/Area Zone
227786
🟩 Layer 3 Link – No Switchport, IP Address Assigned
🟦 Layer 2 Interswitch Uplink – VLAN Trunk, Layer 2 Resiliency
🟥 Layer 2 Access Link – Single VLAN Assigned to port

Manufacturing Zone Topology Summary

Cisco and Rockwell Automation do not have a specific recommendation between the small, medium, and large Manufacturing zone topology options presented. The IACS network requirements, in particular scalability, geographical dispersion, and availability requirements, determine which option to choose.

Note that the large Manufacturing zone topology option represents separating out the Level 3 access, distribution, and core networking functions into distinct equipment. In the small Manufacturing zone topology version, all three Level 3 switch functions are collapsed into the Layer-3 switch stack. It is also possible that only the access or core functions will be separated out, which produces more variations.

Routing

Routing is the process of finding a path to a destination host. Routers or Layer-3 switches forward packets from one network (i.e. sub-network or VLAN) to another based on network layer (IP or Layer 3) information. To do this, routers send each other information about the networks they know about by using various types of protocols—called routing protocols. Routers use this information to build a routing table that consists of the available networks, the cost associated with reaching the available networks, and the path to the next hop router.

For CPwE, routing begins at the Level 3 IACS network, in particular with distribution switches. The distribution switch (e.g., Catalyst 3750) is responsible for routing traffic between Cell/Area IACS networks (inter-VLAN) that it knows about, or into the core, to other routers, or the DMZ. No routing occurs in the Cell/Area IACS network itself. For more information on routing, refer to the following:

Internetwork Design Guidelines: Designing Large-scale IP Networks

I*nternetworking Technology Handbook: Routing Basics*

High-Availability Campus – Layer 3 Routing Protocols

Configuring IOS: IP Overview

For more information on routing basics, refer to the following URL: http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_chapter 09186a008075970b.html

Layer 3 Ports

In Chapter 3, “CPwE Solution Design—Cell/Area Zone,” two key types of Layer 2 ports are identified where the switch forwards incoming packets based on the Layer-2 MAC address:

- End-device or access ports with a specific VLAN assigned to it and other settings
- Uplink or trunk ports connecting switches that carry multiple VLANs in addition to other settings

The Layer-2 managed switch may use other fields in the processing of the packet (e.g., the Layer 3 DSCP field for QoS), but uses the Layer-2 MAC address to determine where to send the packet.

For the Manufacturing zone, a third-type port needs to be introduced, a Layer 3 or routed port. This is a port on which the Layer-3 switch or router will forward incoming traffic based on the Layer-3 IP address, in other words route the packet versus switch the packet. The next section reviews the routing protocols used by the Layer-3 switches or routers to build the routing table. This section simply identifies considerations for port configuration.

Layer 3 ports should be used between switches or routers when no VLANs need to span over that link. Layer 3 links with Layer 3 ports on either end are used:

- Between the distribution and core switches
- Between the plant firewall and core or collapsed core/distribution switches
- Key considerations for a Layer 3 port include the following:
  - Use the no switchport command to remove any Layer 2 switchport commands
  - Apply an IP address to the port

In addition to the physical ports, any switch/router that needs to route traffic from a VLAN, the VLAN definition will create an switch virtual interface (SVI) that is considered a Layer 3. The default gateway IP address for the subnet is typically assigned to this SVI, either directly or as part of the HSRP configuration for failover between two switches/routers. A Layer 3 link may also be applied to an EtherChannel port-channel, thus making the physical ports assigned to that port channel essentially Layer 3 connections.

Note that the assignment of IP addresses to the Layer 3 ports is considered in the "IP Addressing" section on page 4-38.

For an example of Layer 3 or routed ports configuration, refer to Appendix D, "Configurations."

Selection of a Routing Protocol

The correct routing protocol is selected based on the characteristics described in the following subsections.

Distance Vector versus Link-State Routing Protocols

Distance vector routing protocols (such as RIPv1, RIPv2, and IGRP) use more network bandwidth than link-state routing protocols, and generate more bandwidth overhead because of large periodic routing updates. Link-state routing protocols (OSPF, IS-IS) do not generate significant routing update overhead but use more CPU cycles and memory resources than distance vector protocols. Enhanced Interior Gateway Routing Protocol (EIGRP) is a hybrid routing protocol that has characteristics of both the distance vector and link-state routing protocols. EIGRP sends partial updates and maintains neighbor state information just as link-state routing protocols do. EIGRP does not send periodic routing updates as other distance vector routing protocols do.

Cisco and Rockwell Automation recommend using EIGRP or OSPF in IACS networks.

Classless versus Classful Routing Protocols

Routing protocols can be classified based on their support for variable-length subnet mask (VLSM) and Classless Inter-Domain Routing (CIDR). Classful routing protocols do not include the subnet mask in their updates, while classless routing protocols do. Because classful routing protocols do not advertise the subnet mask, the IP network subnet mask should be the same throughout the entire network, and should be contiguous for all practical purposes. For example, if you choose to use a classful routing protocol for a network 172.21.2.0 and the chosen mask is 255.255.255.0, all router interfaces using the network 172.21.2.0 should have the same subnet mask. The disadvantage of using classful routing protocols is that you cannot use the benefits of address summarization to reduce the routing table size, and you lose the flexibility of choosing a smaller or larger subnet using VLSM. RIPv1is an example of a classful routing protocol. RIPv2, OSPF, and EIGRP are classless routing protocols. It is very important that the manufacturing zone uses classless routing protocols to take advantage of VLSM and CIDR.

Convergence

Whenever a change in network topology occurs, every router that is part of the network is aware of this change (except if you use summarization). During this period, until convergence happens, all routers use the stale routing table for forwarding the IP packets. The convergence time for a routing protocol is the time required for the network topology to converge such that the router part of the network topology has a consistent view of the network and has the latest updated routing information for all the networks within the topology.

Link-state routing protocols (such as OSPF) and hybrid routing protocol (EIGRP) have a faster convergence as compared to distance vector protocols (such as RIPv1 and RIPv2). OSPF maintains a link database of all the networks in a topology. If a link goes down, the directly connected router sends a link-state advertisement (LSA) to its neighboring routers. This information propagates through the network topology. After receiving the LSA, each router recalculates its routing table to accommodate this topology change. In the case of EIGRP, Reliable Transport Protocol (RTP) is responsible for providing guaranteed delivery of EIGRP packets between neighboring routers. However, not all the EIGRP packets that neighbors exchange must be sent reliably. Some packets, such as hello packets, can be sent unreliably. More importantly, they can be multicast rather than having separate datagrams with essentially the same payload being discretely addressed and sent to individual routers. This helps an EIGRP network converge quickly, even when its links are of varying speeds.

Routing Metric

If a router has a multiple paths to the same destination, there should be some way for a router to pick a best path. This is done using a variable called a metric assigned to routes as a means of ranking the routes from best to worse or from least preferred to the most preferred. Various routing protocols use various metrics, such as the following:

- RIPv1 and RIPv2 use hop count as a metric and therefore are not capable of taking into account the speed of the links connecting two routers. This means that they treat two parallel paths of unequal speeds between two routers as if they were of the same speed, and send the same number of packets over each link instead of sending more over the faster link and fewer or no packets over the slower link. If you have such a scenario in the Manufacturing zone, it is highly recommended to use EIGRP or OSPF because these routing protocols take the speed of the link into consideration when calculating metric for the path to the destination.
- EIGRP uses a composite metric that is based on the combination of lowest bandwidth along the route and the total delay of the route.
- OSPF uses cost of the link as the metric that is calculated as the reference bandwidth (ref-bw) value divided by the bandwidth value, with the ref-bw value equal to 10^8 by default.

Scalability

As the network grows, a routing protocol should be capable of handling the addition of new networks. Link-state routing protocols such as OSPF and hybrid routing protocols such as EIGRP offer greater scalability when used in medium-to-large complex networks. Distance vector routing protocols such as RIPv1 and RIPv2 are not suitable for complex networks because of the length of time they take to converge. Although IS-IS is scalable, IS-IS is not commonly used in Enterprise networks due to the complexity to implement and the fact it does not use IP to communicate routing information. BGP is a protocol commonly found at the Internet edge of enterprise networks, and therefore not a relevant option for plant networks. Factors such as convergence time and support for VLSM and CIDR directly affect the scalability of the routing protocols.

Table 4-4 shows a comparison of routing protocols

Table 4-4 Routing Protocols Comparison

Name	Type	Proprietary	Function	Updates	Metric	VLSM	Summarization
RIP	Distance vector	No	Interior	30 sec	Hops	No	Auto
RIPv2	Distance vector	No	Interior	30 sec	Hops	Yes	Auto
IGRP	Distance vector	Yes	Interior	90 sec	Composite	No	Auto
EIGRP	Advanced Distance vector	Yes	Interior	Trig	Composite	Yes	Both
OSPF	Link-state	No	Interior	Trig	Cost	Yes	Manual
IS-IS	Link-state	No	Interior	Trig	Cost	Yes	Auto
BGP	Path vector	No	Exterior	Incr	N/A	Yes	Auto

In summary, the Manufacturing zone usually has multiple parallel or redundant paths for a destination and requires VLSM for discontinuous major networks. Cisco and Rockwell Automation recommend use of OSPF or EIGRP as the core routing protocol in the Manufacturing zone.

At the time of the writing of this DIG, test results show that EIGRP is better suited to a campus environment than OSPF. The ability of EIGRP to provide route filtering and summarization maps easily to the tiered hierarchical model, while the more rigid requirements of OSPF do not easily integrate to existing implementations and require more complex solutions.

The following are additional considerations when comparing EIGRP and OSPF:

- Within the campus environment, EIGRP provides for faster convergence and greater flexibility.
- EIGRP provides for multiple levels of route summarization and route filtering that map to the multiple tiers of the campus.
- OSPF implements throttles on Link-State Advertisement (LSA) generation and Shortest Path First (SPF) calculations that limit convergence times.
- When routes are summarized and filtered, only the distribution peers in an EIGRP network need to calculate new routes in the event of link or node failure.

Static or Dynamic Routing

The role of a dynamic routing protocol in a network is to automatically detect and adapt to changes to the network topology. The routing protocol decides the best path to reach a particular destination. If precise control of path selection is required, particularly when the path you need is different from the path of the routing protocol, use static routing. Static routing is hard to manage in medium-to-large network topologies, and therefore a dynamic routing protocol is preferred.

Applying the Routing Protocol

The following are the key recommendations to consider as routing is configured for all Manufacturing zone topologies:

- Enable IP directed broadcast. This feature is required to allow IACS software data servers, such as RSLinx Classic with RSWho functionality, to discover the IACS EtherNet/IP devices in the Cell/Area zones. Some Cisco security guidelines suggest disabling this feature, but when strong segmentation with a DMZ is in place, Cisco and Rockwell Automation recommend the feature be enabled within the Manufacturing zone.
- Specify the default gateway on all Cell/Area zone switches and IACS EtherNet/IP devices. This enables the end-devices and switches to send communications outside of the Cell/Area zone.
- Control peering across access layer links (passive interfaces). In most cases, core/distribution switches will not be interconnected via access layer switches, but in the cast that occurs, for example when a VLAN must span multiple distribution switches, steps should be taken to control Layer 3 peering. Unless you control Layer 3 peering in the hierarchical campus model, the distribution nodes establish Layer 3 peer relationships many times using the access nodes that they support, wasting memory and bandwidth.

For topologies where separate core and distribution are in place (medium-to-large):

- Use triangles when configuring core and distribution switch/router ports, See Figure 4-10.

  If a topology is built using triangles, with equal-cost paths to all redundant nodes, slower, timer-based convergence can be avoided. Instead of indirect neighbor or route loss detection using hellos and dead timers, you can rely on physical link loss to mark a path as unusable and reroute all traffic to the alternate equal-cost path. Figure 4-10 shows triangle and Figure 4-10 square network topologies. In a square network topology, depending on the location of the failure, the routing protocol may need to converge to identify a new path to the subnet/VLAN, slowing the convergence of the network.

Figure 4-10 Example of Triangle Network Topology

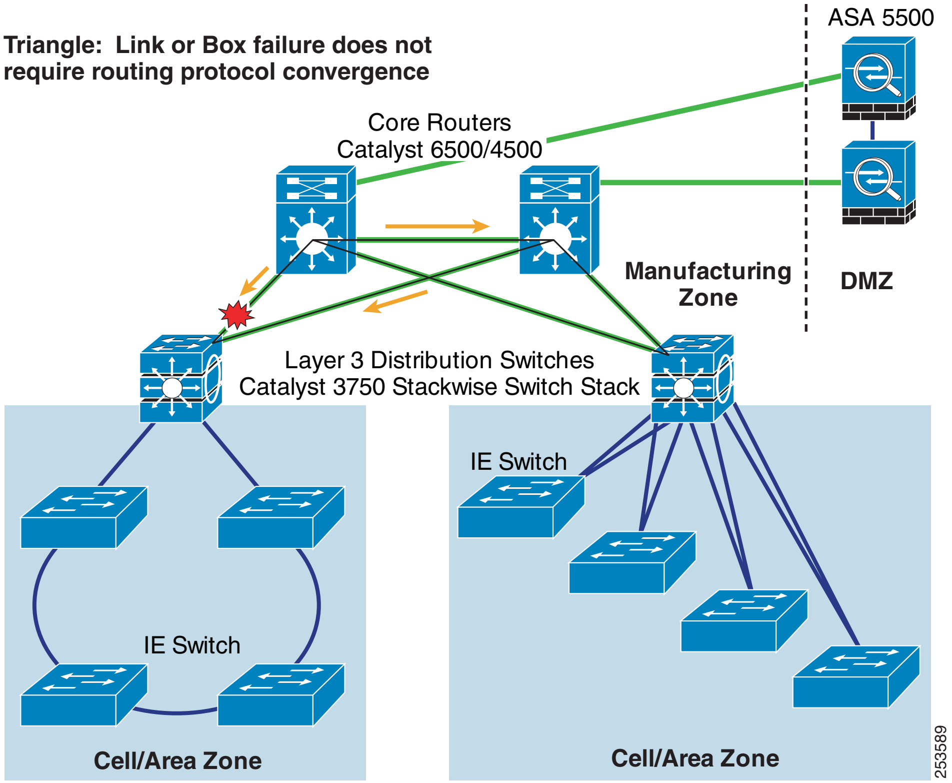
Triangle: Link or Box failure does not require routing protocol convergence
ASA 5500
Core Routers
Catalyst 6500/4500
Manufacturing Zone
DMZ
Layer 3 Distribution Switches
Catalyst 3750 Stackwise Switch Stack
IE Switch
IE Switch
Cell/Area Zone
Cell/Area Zone
253589

🟩 Layer 3 link–no switchport, IP address assigned

🟦 Layer 2 interswitch uplink–VLAN Trunk, Layer 2 resiliency

Figure 4-11 Example of Square Network Topologies

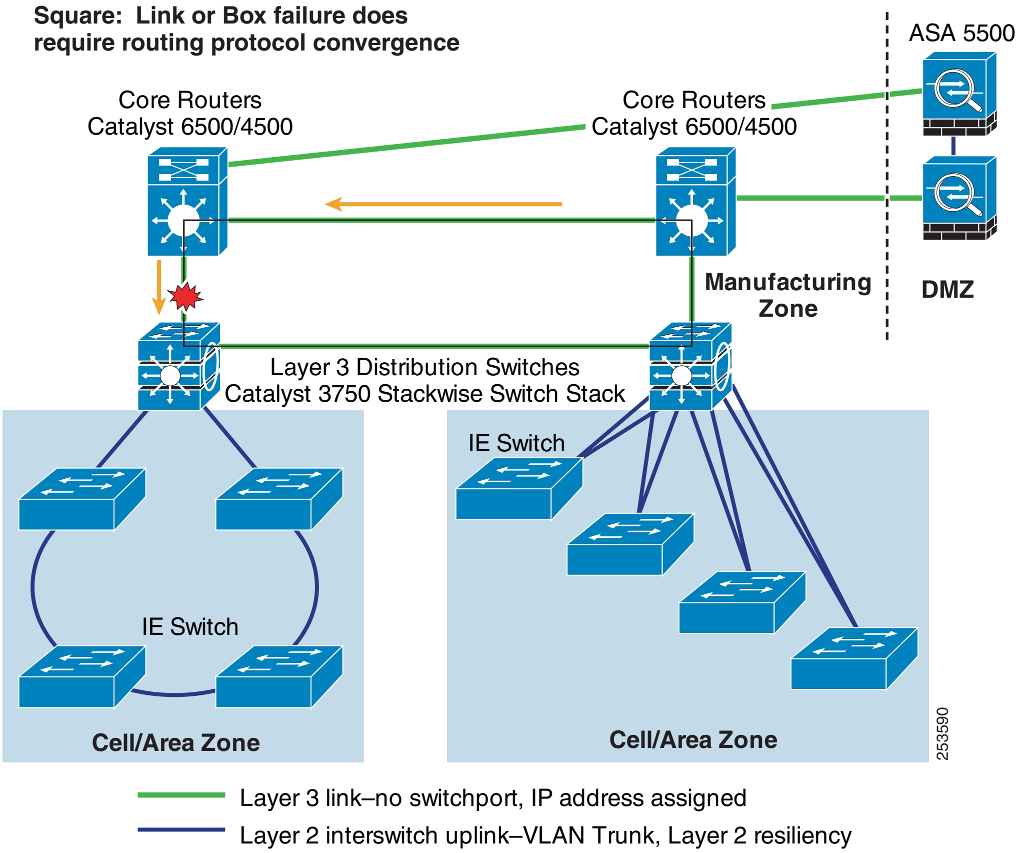
Square: Link or Box failure does require routing protocol convergence
ASA 5500
Core Routers Catalyst 6500/4500
Core Routers Catalyst 6500/4500
Manufacturing Zone
DMZ
Layer 3 Distribution Switches Catalyst 3750 Stackwise Switch Stack
IE Switch
IE Switch
Cell/Area Zone
Cell/Area Zone
253590
Layer 3 link–no switchport, IP address assigned
Layer 2 interswitch uplink–VLAN Trunk, Layer 2 resiliency

- Use equal-cost redundant connections from distribution to the core for fastest convergence and to avoid black holes.
- While it is tempting to reduce cost by reducing links between the distribution nodes to the core in a partial mesh design, the complexity and convergence tradeoffs related to this design are ultimately far more expensive.
- Summarization is required to facilitate optimum EIGRP or OSPF convergence. Summarize at the distribution.

  It is important to summarize routing information as it leaves the distribution nodes towards the core for both EIGRP and OSPF. When you force summarization at this layer of the network, bounds are implemented on EIGRP queries and OSPF LSA/SPF propagation, which optimizes both routing protocols for campus convergence.

This solution does not cover the option of using multiple "distribution" switches to support a VLAN or set of VLANS. In this case, a Layer 2 connection between the distribution switches is needed. For more information on. The key consideration in this mode include the following:

- Connect distribution nodes to facilitate summarization and maintain Layer 2 VLANs. The devices should have HSRP running to manage the routing. This way, the core has one port/direction for every VLAN/subnet. If summarization is being used, the distribution switches must be linked or routing black holes occur.

For more on configuration of Routing protocols, refer to Configuring IP Routing Protocols at the following URL: http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfodr.html

Logical Segmentation

Logical segmentation is important for the Level 3 IACS network, especially for the CPwE Level 3 workstations and servers. In the Cell/Area zone, it is important for endpoints that communicate Implicit Common Industrial Protocol (CIP) I/O traffic to be in the same VLAN for traffic flow and real-time communications reasons. In the Manufacturing zone, the key consideration for segmentation is security. Security policy may require that certain functions or roles have access to particular applications and services that reside in the Manufacturing zone. In addition, the IACS applications (CPwE Level 3) may need access only to a subset of Cell/Area zones. A well-designed segmentation design greatly improves the ability to apply and maintain a security policy.

The following key functional areas are good candidates for segmentation:

- IACS applications dedicated to particular functions on the plant floor (for example, a brewing control room)
- Remote access server(s)
- Security and network administration applications

As in the Cell/Area zone, a mixture of physical separation and VLANs is used to achieve segmentation.

In this context, there is one particular common practice that Cisco and Rockwell Automation *strongly discourages: dual-homing*. Dual-homing is the concept of having key Manufacturing zone workstations or servers installed with two network interfaces: one connected to the Manufacturing zone and the other directly to the Enterprise zone. Dual-homing facilitates the sharing of data and services between the two zones. This poses a significant security risk, because these workstations or servers typically are not secured as other devices can be, and are points of entry to the Manufacturing zone for malicious activity to target. The CPwE solution architecture identifies a DMZ with firewall functions to safely and securely share data and services between these zones.

Availability

Because the Cell/Area zone inter-connect functionality exists within the Manufacturing zone, the high availability of the routing infrastructure is critical to the optimal performance of the Manufacturing zone. This section describes design considerations for the following key manufacturing services:

- CPwE Level 3, Layer-2 connectivity
- Core routing and Layer-3 switching
- Network and systems management
- Endpoint security

Layer 2 Connectivity

The CPwE Level 3 workstations and servers are connected to LANs/VLANs. These VLANs also need to be designed with availability considerations. CPwE previously recommended that the redundant topology be applied; therefore, MSTP must be implemented in the Layer 2, Level 3 IACS network to prevent network loops and to recover after the loss of a connection.

Core Routing and Layer-3 Switching Resiliency

Key availability considerations in routing and switching can be divided into hardware and device level considerations and network level considerations.

Network Device Level Resiliency

Device level resiliency refers to techniques that protect against any failure of a device node so that it can continue processing traffic with no or minimum disruption. The techniques relevant to the control network environment are shown in Table 4-5.

Table 4-5 Network Device Level Resiliency Design

Feature	Description	Supported Platforms	Where to Apply in Industrial Ethernet Network
Redundant route processors (supervisors)	Active and standby supervisors operate in active and standby modes and provide a variety of redundancy mechanisms to handle failure scenarios. Requires redundant devices.	• Catalyst 6500 • Catalyst 4500 • Catalyst 3750—Virtual with StackWise	All, especially the Core routing function
StackWise	Uses stack interconnect cables to create a virtual switch fabric for stacks of the Catalyst 3750.	• Catalyst 3750 • N/A to Other Platforms	All, especially for distribution function
Redundant power supplies	Each system has dual power supplies so that the system operates normally upon failure of a power supply	• Catalyst 6500 • Catalyst 4500: Internal • Catalyst 3750: External • Stratix 8000: External • IE 3000: External	All
Redundant fans	Each fan tray has multiple fans	• Catalyst 6500 • Catalyst 4500	All
Line card online insert and removal (OIR)	New line cards can be added or removed without affecting the system or losing the configuration.	• Catalyst 6500 • Catalyst 4500	All
Control Plane Policing (CoPP)	Prevents malicious traffic from flooding the CPU to the point that the switch can no longer forward packets and perform functions. Achieved by configuring a QoS filter.	• Catalyst 6500 • Catalyst 4500	All
Nonstop Forwarding with Stateful Switchover (NSF with SSO)	Inter-chassis supervisor failover at Layers 2 through 4. Reduces the mean time to recovery (MTTR).	• Catalyst 6500 • Catalyst 4500	Whatever Layer 3 routing takes place
In-Service Software Upgrade (ISSU)	Ranges from full image upgrades to granular; selective software maintenance can be performed without service impact across all Cisco IOS-based products.	• Catalyst 6500 • Catalyst 4500	

Table 4-5 Network Device Level Resiliency Design (continued)

Feature	Description	Supported Platforms	Where to Apply in Industrial Ethernet Network
Automatic software upgrade for Catalyst 3750 StackWise	The Master 3750 transfers the same version of code to the remaining switches in the stack. The upgrade includes · Transfer the global configuration · Apply default configuration · Apply preconfigured configuration	Catalyst 3750	All
Generic Online Diagnostics (GOLD)	Online diagnostics to help ensure that a system that is booting up and a live system are healthy.	Catalyst 6500, 4500 and 3750: subset of GOLD	All
Configuration rollback	Capability to replace the current running configuration with any saved Cisco IOS configuration file	Catalyst 6500	

Network Level Resiliency

Network level resiliency refers to techniques that can route traffic around a failure point in the network. The techniques relevant to the IACS network environment are shown in Table 4-6.

Table 4-6 Network Level Resiliency Design

Feature	Description	Supported Platforms	Where to Apply in Industrial Ethernet Network
Link redundancy	Sends packets to their destinations over a backup link of a network device when its primary link fails because of link breakage, or failure of an interface or line card. Determined by the Layer 2 resiliency or a Layer 3 routing protocol.	All routers and switches	All
Hot Standby Router Protocol (HSRP)	Cisco proprietary redundancy protocol for establishing a fault-tolerant default gateway.	· Catalyst 6500 · Catalyst 4500 · Catalyst 3750—Virtual with StackWise	Between routers/switches performing similar functions (e.g. redundant core switches). It is not needed on a 3750-stack to manage the switches in the stack.
Incremental SPF Optimizations	Optimization of the OSPF algorithm to reduce computational load.	· Catalyst 6500 - Internal · Catalyst 4500—Internal · Catalyst 3750—External	All
IP dampening	Mechanism to suppress affects of excessive state changes (flapping).	· Catalyst 6500 · Catalyst 4500	All

In addition to these features, Cisco and Rockwell Automation recommend the following be applied at the core and distribution layers to improve resiliency:

- Use GLBP/HSRP millisecond timers.

Convergence around a link or node failure in the Layer-2/Layer-3 distribution boundary model depends on default gateway redundancy and failover. Millisecond timers can reliably be implemented to achieve sub-second (800 ms) convergence based on HSRP/GLBP failover.

- Tune GLBP/HSRP preempt delay to avoid black holes.

  Ensure that the distribution node has connectivity to the core before it preempts its HSRP/GLBP standby peer so that traffic is not dropped while connectivity to the core is established.

IP Addressing

IP addressing is an important concept for IACS networking. For all practical purposes, every IACS network device (server, endpoint, infrastructure device) on the industrial Ethernet network is assigned an IP address, and is thereby addressable within the Manufacturing zone.

This section addresses the following:

- Provides a brief background on IP addressing and some helpful links for more information
- Reviews best practices for IP address management in plant manufacturing environments
- Reviews best practices for IP address allocation in the Manufacturing zone

IP Addressing Background

As with most concepts described here, there are multiple versions of standards in use. That is also true for IP addressing. The two key IP addressing standards relevant in standard networking are IPv4 and IPv6. IPv4 defines an IP address as a 32-bit number and is currently the most prevalent IP addressing standard in use, not just in manufacturing but in general. IPv6 defines an IP address as a 128-bit number and is currently only adopted in a small percentage of devices and relatively unknown in IACS networks. Most IACS network devices available today do not support IPv6 addressing along with the relevant IACS protocols and standards (ODVA's CIP, EtherNet/IP included). Therefore, this *DIG* assumes IPv4 is the IP addressing standard in effect.

Refer to the following key documents about IP addressing:

- IP Addressing and Subnetting for New Users
- Internetworking Technology Handbook: IP
- Configuring IP Addressing

IP Address Management

IP address management is the process of allocating, recycling, and documenting IP addresses and subnets in a network. IP addressing standards define subnet size, subnet assignment, network device assignments, and dynamic address assignments within a subnet range. Recommended IP address management standards reduce the opportunity for overlapping or duplicate subnets, non-summarization in the network, duplicate IP address device assignments, wasted IP address space, and unnecessary complexity.

Developing an appropriate IP addressing schema for an IACS network with considerations for future expansion is critical. Changing IP addressing schemas are difficult, time consuming and involve IACS network downtime.

Address Space Planning

When planning address space, administrators must be able to forecast the IP address capacity requirements and future growth in every accessible subnet on the network. This is based on many factors such as number of end-devices, number of users working on the plant floor, number of IP addresses required for each application or each end-device, and so on. Even with plentiful availability of private address space, the cost associated with supporting and managing the IP addresses can be huge. With these constraints, it is highly recommended that administrators plan and accurately allocate the addressing space with future growth into consideration.

Hierarchical Addressing

Hierarchical addressing leads to efficient allocation of IP addresses. An optimized address plan is a result of good hierarchical addressing. A hierarchical address plan allows you to take advantage of all possible addresses because you can easily group them contiguously. With random address assignment, there is a high possibility of wasting groups of addresses because of addressing conflicts.

Another benefit of hierarchical addressing is a reduced number of routing table entries. The routing table should be kept as small as possible by using route summarization.

Summarization (also known as supernetting) allows aggregation of smaller subnets that reside on that network into a single route in the routing table. Route summarization is a way of having a single route represent multiple smaller routes, which can be very well accomplished when hierarchical addressing is used. By summarizing routes, you can keep the routing table entries small, which offers the following benefits:

- Efficient routing
- Reduced router memory requirements
- Reduced number of CPU cycles when recalculating a routing table or going through routing table entries to find a match
- Reduced bandwidth required because of fewer small routing updates
- Easier troubleshooting
- Fast convergence
- Increased network stability because detailed routes are hidden, and therefore impact to the network when the detailed routes fail is reduced

If address allocation is not done hierarchically, there is a high chance of duplicate IP addresses being assigned to end-devices. In addition, networks can be unreachable if route summarization is configured.

Hierarchical addressing helps in allocating address space optimally and is the key to maximizing address use in a routing-efficient manner.

Note Overlapping IP addresses should be avoided in the Manufacturing Cell/Area zone. If two devices have identical IP addresses, the ARP cache may contain the MAC (node) address of another device, and routing (forwarding) of IP packets to the correct destination may fail. CPwE recommends that IACS network devices should be hard-coded with a properly unique static IP address.

Note CPwE recommends that the traffic associated with any multicast address (224.0.0.0 through 239.255.255.255) used in the Manufacturing zone should not be permitted in the Enterprise zone; EtherNet/IP devices in the Manufacturing zone use an algorithm to choose a multicast address for their Implicit CIP I/O traffic. Therefore, to avoid conflict with multicast addresses in the Enterprise zone, use the DMZ to segment multicast traffic in the Manufacturing zone from multicast traffic in the Enterprise zone.

Centralized IP Addressing Inventory

Address space planning and assignment can be best achieved using a centralized approach and maintaining a central IP inventory repository or database. The centralized approach provides a complete view of the entire IP address allocation of various sites within an organization. This helps in reducing IP address allocation errors and also reduces duplicate IP address assignment to end-devices. The IT department is usually responsible for maintaining a centralized IP addressing schema and inventory for an enterprise.

Multicast IP Addresses

As stated earlier, IACS applications and in particular EtherNet/IP implementations generate multicast traffic. The IP standards have set aside a range of IP addresses for the exclusive use of multicast traffic (see the URL below). It is important to consider the multicast addresses used in the IACS network devices and to avoid address overlap with other applications that may use multicast and not allow the IACS multicast traffic to mix with enterprise traffic to ensure application consistency.

Internet Protocol IP Multicast Technology:

http://www.cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.html

For more information on the range of IP multicast addresses used with EtherNet/IP applications see (ODVA).

Private and Public IP Addresses

The IPv4 (RFC 1918 and 4193) standard also allots a range of IP address for private use. These addresses can be used within a private network but cannot be used on the Internet. Any company or individual can use these addresses within their internal networks.

The IACS network traffic is confined to the Manufacturing zone. Because of this, either private or public IP addresses can be used for the Manufacturing zone. Routable IP addresses are an extremely limited resource. Not all organizations have been assigned routable IP addresses. If an organization has public addresses, they typically have a relatively small number of addresses assigned to them. Using private IP addresses in the Manufacturing zone frees up the public addresses for other purposes. It is important that a unique summerizable block of IP addresses is assigned to the Manufacturing zone.

In the end, the use of private or public IP addresses is going to depend highly on the enterprise and the overall IP approach.

For more information and pros/cons on private IP addresses, see Address Allocation for Private Internets.

Subnetworks

IP networks are divided into smaller network called subnets. Each subnet represents a group of hosts on the network. Hosts on the same subnet communicate directly with each other over the Layer 2 network. Hosts on different subnets communicate with each other via their default gateways. Subnets divide the IP network into smaller, more manageable networks. In the CPwE architecture, each VLAN in the Manufacturing zone has a unique subnet assigned to it.

Network Address Translation (NAT)

NAT is another mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. NAT operates on a router connecting two networks together; one of these networks (designated as inside) is addressed with either private or duplicate public addresses that need to be converted into unique public addresses before packets are forwarded onto the other network (designated as outside). Devices on the inside assigned a unique IP address via NAT are also addressable from the outside.

Another reason NAT is attractive to IACS networking, besides the need to conserve IP addresses, is the opportunity to reuse IP address schemas for different parts of the IACS network. By reusing IP address schemas, manufacturers and their partners hope to reduce the test and implementation of cookie-cutter systems or lines on the plant floor by repeating the IP address schema for parts of the IACS network.

The disadvantages with NAT include the following:

- Additional overhead to develop and maintain NAT translation schemas
- NAT likely will leave a number of devices as un-addressable from outside of the NAT'd region, any that do not receive a NAT address
- NAT may impact applications that use embedded IP addresses in the data packet.
- Data servers for IACS application software may not properly support NAT

Cisco and Rockwell Automation do not recommend the use of NAT services in the Manufacturing zone.

- Cisco IOS NAT Overview: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_q_and_a_item09186a00800e523b.shtml
- NAT Frequently Asked Questions: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_q_and_a_item09186a00800e523b.shtml
- How NAT Works: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094831.shtml

Dynamic Name Services

Domain Name System (DNS) is used for translating names of network hosts into addresses. DNS is used as an abstraction layer between IP addresses, applications and services. DNS allows a host to be identified by its name instead of its IP address. If the IP address of the host is changed, DNS is updated to reflect the new address. The client application continues to communicate with the name and does not need to know that the address has changed.

Cisco and Rockwell Automation recommend referencing Cell/Area IACS network devices by their static IP address as opposed to their DNS name for simplicity and to avoid potential latency delays if the DNS server has issues. DNS resolution delays may not be acceptable depending on the IACS application. DNS may be used effectively for certain Level 3 IACS applications and systems. In this case, a DNS service must be established to perform the name translation.

Note that DNS is a network service commonly used throughout most enterprise networks. Due to the strong segmentation recommended by this CPwE solution, the DNS names found in the Enterprise zone do not need to be available in the Manufacturing zone; therefore replication of DNS services between Enterprise and Manufacturing zones may not be required.

Other IP Address Considerations

As IP address schemas for the IACS network are developed, here are some other considerations:

- If a Management VLAN is implemented, a unique subnet and VLAN should be applied to the switching infrastructure.
- Loopback Interfaces. Loopback addresses for system logging should be established for effective of system messages. If you have available IP address space or can allocate a range of IP subnets for the loopback interface, it will be easier for you to manage your network, utilizing some of the configuration techniques. A loopback interface requires its own IP subnet. But you can use a host mask (255.255.255.255 or /32) for this interface because no other device uses that subnet.
- If you are short on IP addresses, the loopback interface is probably not a viable option. In that case, consider “key” physical interfaces, such as backbone-oriented interfaces, for use as source IP addresses for SNMP traps or syslog messages because those interfaces should be up most of the time.
  - Source the traps from the Loopback0 interface using the **snmp-server trap-source Loopback0** configuration command. By doing this, you can control where traps are sourced from versus having multiple IP address sourcing the traps. By default, all traps are sourced from the outgoing physical interface's IP address. It is easier to track one IP address or hostname than multiple IP addresses or hostnames from a common host (Cisco device).

IP Address Allocation

This section covers the mechanisms used to deploy IP addresses to end-devices once an IP addressing schema has been established. The section looks at both static and dynamic methods and summarizes the recommendations.

Static IP Addressing

In the Manufacturing zone, the Level 3 workstations and servers are static. Additionally, it is recommended to statically configure Cell/Area IACS network devices.

The Level 3 servers send detailed scheduling, execution, and operator IACS data to controllers in the Manufacturing zone, and collect data from the controllers for historical data and audit purposes. Cisco and Rockwell Automation recommend manually assigning IP addresses to all the IACS networking devices including servers and Cisco networking equipment in the Manufacturing zone. For more information on IP addressing, see *IP Addressing and Subnetting for New Users* at the following URL:
http://www.cisco.com/en/US/customer/tech/tk365/technologies_tech_note09186a00800a67f5.shtml

Using Dynamic Host Configuration Protocol and DHCP Option 82

Dynamic Host Configuration Protocol (DHCP) is used in LAN environments to dynamically assign host IP addresses from a centralized server, which reduces the overhead of administrating IP addresses. DHCP also helps conserve limited IP address space because IP addresses no longer need to be permanently assigned to client devices; only those client devices that are connected to the network require an IP addresses. The DHCP relay agent information feature (option 82) enables the DHCP relay agent (Catalyst switch) to include information about itself and the attached client when forwarding DHCP requests from a DHCP client to a DHCP server. This extends the standard DHCP process by tagging the request with the information regarding the location of the requestor. See Figure 4-12.

Figure 4-12 DHCP Option 82 Operation

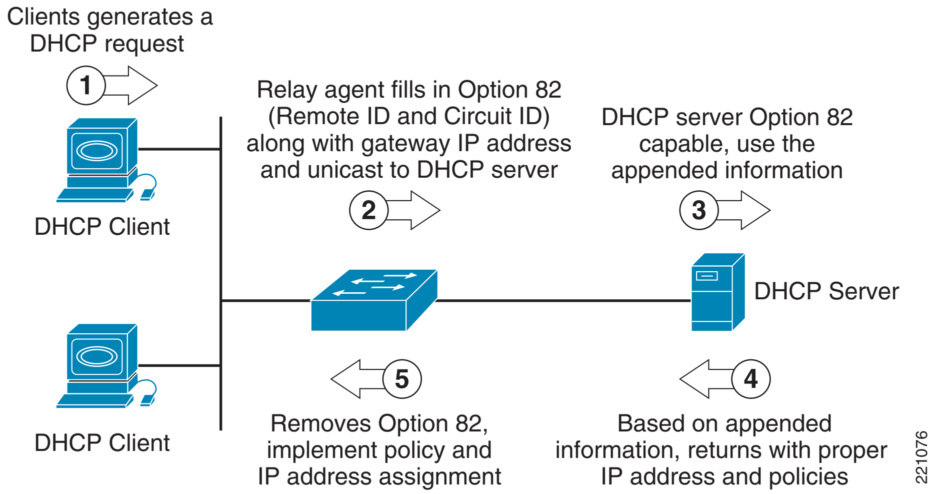
Clients generates a DHCP request
1
DHCP Client
DHCP Client
Relay agent fills in Option 82 (Remote ID and Circuit ID) along with gateway IP address and unicast to DHCP server
2
DHCP server Option 82 capable, use the appended information
3
DHCP Server
5
Removes Option 82, implement policy and IP address assignment
4
Based on appended information, returns with proper IP address and policies
221076

The following are key elements required to support the DHCP option 82 feature:

- Clients supporting DHCP
- Relay agents supporting option 82
- DHCP server supporting option 82

The relay agent information option is inserted by the DHCP relay agent when forwarding the client-initiated DHCP request packets to a DHCP server. The servers recognizing the relay agent information option may use the information to assign IP addresses and to implement policies such as restricting the number of IP addresses that can be assigned to a single circuit ID. The circuit ID in relay agent option 82 contains information identifying the port location on which the request is arriving.

For details on DHCP features, see the following URL:
http://www.cisco.com/en/US/products/ps7077/products_configuration_guide_chapter09186a008077a28b.html#wp1070843

The DHCP option 82 feature is supported only when DHCP snooping is globally-enabled and on the VLANs to which subscriber devices using this feature are assigned. DHCP and the DHCP option 82 feature have not been validated in the lab for CPwE version 2.0

At this time, Cisco and Rockwell Automation do not recommend dynamic IP address allocation in the CPwE architecture. Table 4-7 provides a summary of IP address allocation mechanisms used for this solution.

IP Address Summary

There are several ways to allocate IP addresses. Table 4-7 lists a number of variations and their advantages and disadvantages. Cisco and Rockwell Automation recommend that network developers use either a static IP addressing schema or DHCP persistence for the Manufacturing zone, especially for allocating IP addresses to IACS devices in the Cell/Area zone. DHCP persistence provides a dynamic IP allocation mechanism on a IES per-port basis. For more information on DHCP persistence, refer to Chapter 10, "DHCP Persistence in the Cell/Area Zone."

Table 4-7 Summary of IP Address Allocation mechanisms

Option	Description	Advantages	Disadvantages
Static	IACS network device hard coded with an IP Address through mechanical means such as a rotary switch	Simple to commission and replace	In large environments, can be burdensome to maintain Limited ranged of IP addresses and subnet Not all devices support
Static via BOOTP Configuration	Server, through manual intervention, assigns IACS network devices IP addresses Precursor to DHCP	Supported by every device	Requires technician to configure IP address/MAC address when a device is replaced Adds complexity and point of failure
DHCP	Server automatically assigns IP addresses from a pool (NOT RECOMMENDED for Cell/Area zone IACS devices)	Efficient use of IP address range Can reduce administration work load	More complex to implement and adds a point of failure Devices get different IP addresses when they reboot
DHCP Option 82	Server assigns consistent IP addresses from a pool (NOT RECOMMENDED)	Efficient use of IP Address range Can reduce administration work load	More complex to implement and adds a point of failure Mixed environments may not work
DHCP Persistence	Automatically assign IP address per physical switch port	Efficient use of IP Address range Eases commissioning and maintenance in large environments	Cisco/RA only Requires some maintenance and upkeep, on a per switch basis

Security Design

The security design for the Manufacturing zone network infrastructure builds upon the IACS network infrastructure security recommendations made for the Cell/Area zone, but includes concepts to protect the network features specific to the Manufacturing zone, especially routing and protecting the applications and servers specific to the Manufacturing zone (e.g., the FactoryTalk Integrated Production and Performance Suite). The key concepts applied in the Cell/Area zone that also apply to the Manufacturing zone include:

- Network Infrastructure Device Access, covering port security, password maintenance, administrative access
- Resiliency and survivability, covering redundancy and disabling un-necessary services
- Network Telemetry, covering network system message logging, SNMP, SSH, and network information to monitor
- Other zone security best practices, covering restricting broadcast domains, VLANs and protecting a variety of network protocols

This section will not go into detail on the design considerations as they are sufficiently covered in Chapter 3, "CPwE Solution Design—Cell/Area Zone," Section Security as well as Cisco's SAFE solution. The topics covered in this section include:

- Routing Infrastructure protection

  Note that more advanced protection mechanisms are also possible, but probably beyond the majority of IACS networks. These may be considered by advanced IACS network teams, larger installations or highly sensitive installations. Some of the advanced techniques include the following:

  – Establishing infrastructure ACLs (iACLs) to explicitly permit authorized control and management traffic bound to the IACS network infrastructure equipment
- Control Plane Policing and Protection to protect the network traffic and data used by the network infrastructure.
- Network Policy Enforcement that ensures traffic entering a network conforms to the network policy, including the IP address range and traffic types

All of these concepts and techniques are described in more detail in Cisco's SAFE and Network Security Baseline solutions at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html

Routing Infrastructure Protection

Routing is one of the key features for the Manufacturing zone IACS network. Therefore, steps to secure the routing protocol are critical to maintaining overall IACS security and availability. The routing protocol may be compromised via injection of improper routing data, denial-of-service or used to target attacks on other devices in the network. The key measures Cisco and Rockwell Automation recommend to protect the routing function include:

- Restrict routing protocol membership- Limit routing sessions to trusted peers, validate origin, and integrity of routing updates.
- Log status changes adjacency or neighbor sessions.

Restrict Routing Protocol Membership

The key considerations for an IACS network in this technique are to enable routing neighbor authentication and to statically configure a list of trusted neighbors. For most IACS networks, the routers are not going to be very dynamic therefore the effort involved will be minimal and ensures that dynamic mechanisms do not identify invalid routers. This should keep the routing function in a Manufacturing zone segmented by plant firewalls and a DMZ well protected. There are other considerations listed in the referred to documentation, but they are either for different types of networks or are already covered in other security recommendations.

For details on these techniques and implementation guidance, see Cisco's SAFE and Network Security Baseline solutions. For more on the authentication techniques for each protocol, see

- Configuring EIGRP Authentication—http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00807f5a63.shtml
- Configuring IS-IS Authentication—http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080093f36.shtml
- Configuring OSPF Authentication—http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080093f36.shtml

Log Status Changes

Logging of status changes that impact the routing protocols is important to maintaining security. If frequent or unexpected changes occur, they may indicate an ongoing attack or vulnerability. In most routing protocols, status change message logging is enabled by default. When enabled, every time a router session goes down, up, or experiences a reset, the router generates a log message. If syslog is enabled as recommended by this solution, the message is forwarded to the syslog server; otherwise is kept in the router's internal buffer.

For details on these techniques and implementation guidance, see Cisco's SAFE and Network Security Baseline solutions at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg

Server Farm

Types of Servers

The servers used in the Manufacturing zone can be classified into three categories.

- Servers that provide common network-based services such as the following:
  - DNS—Primarily used to resolve hostnames to IP addresses.
  - DHCP—Used by end-devices to obtain IP addresses and other parameters such as the default gateway, subnet mask, and IP addresses of DNS servers from a DHCP server. The DHCP server makes sure that all IP addresses are unique; that is, no IP address is assigned to a second end-device if a device already has that IP address.
  - Directory services—Set of applications that organizes and stores date about end users and network resources.

  - Network Time Protocol (NTP)—Synchronizes the time on a network of machines. NTP runs over UDP, using port 123 as both the source and destination. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. An NTP client makes a transaction with its server over its polling interval (64–1024 seconds,) which dynamically changes over time depending on the network conditions between the NTP server and the client. No more than one NTP transaction per minute is needed to synchronize two machines.
  - Precision Time Protocol (PTP) is not being addressed in this version of CPwE.

    For more information, see *Network Time Protocol: Best Practices White Paper* at the following URL:

    http://www.cisco.com/en/US/customer/tech/tk869/tk769/technologies_white_paper09186a0080117070.shtml

- Servers that provide security and network management services such as the following:
  - Cisco Security Monitoring, Analysis, and Response System (MARS)—Provides security monitoring for network security devices and host applications made by Cisco and other providers.
  - Greatly reduces false positives by providing an end-to-end view of the network
  - Defines the most effective mitigation responses by understanding the configuration and topology of your environment
  - Promotes awareness of environmental anomalies with network behavior analysis using NetFlow
  - Makes precise recommendations for threat removal, including the ability to visualize the attack path and identify the source of the threat with detailed topological graphs that simplify security response at Layer 2 and above

    For more information on CS-MARS, see the CS-MARS introduction at the following URL:

    http://www.cisco.com/en/US/customer/products/ps6241/tsd_products_support_series_home.html
  - Cisco Network Assistant (CNA)—PC-based network management application optimized for wired and wireless LANs for growing businesses that have 40 or fewer switches and routers. Using Cisco Smartports technology, Cisco Network Assistant simplifies configuration, management, troubleshooting, and ongoing optimization of Cisco networks. The application provides a centralized network view through a user-friendly GUI. The program allows network administrators to easily apply common services, generate inventory reports, synchronize passwords, and employ features across Cisco switches, routers, and access points.

    For more information, see the Cisco Network Assistant general information at the following URL: http://www.cisco.com/en/US/customer/products/ps5931/tsd_products_support_series_home.html
- CiscoWorks LAN Management Solution (LMS)—CiscoWorks LMS is a suite of powerful management tools that simplify the configuration, administration, monitoring, and troubleshooting of Cisco networks. It integrates these capabilities into a best-in-class solution for the following:
  - Improving the accuracy and efficiency of your operations staff

  - Increasing the overall availability of your network through proactive planning
  - Maximizing network security

  For more information, see CiscoWorks LMS at the following URL:

  http://www.cisco.com/en/US/customer/products/sw/cscowork/ps2425/tsd_products_support_series_home.html

- Servers that provide manufacturing services such as the following:
  - Data servers such as RSLinx Classic and Enterprise
  - FactoryTalk Servers such as Historian, View, AssetCentre, and Batch
  - FactoryTalk Services platform such as Directory, Security, Audit, Diagnostics, Live Data, Activation, and Alarms & Events
  - Engineering workstation such as RSLogix 5000/500/5

The recommendation is put the above three categories into three separate VLANs. If necessary, the manufacturing application servers can be further segregated based on their functionality.

Security Protection for Servers

The servers that provide network services, network management, or Level 3 Site Manufacturing Operations and Control should be provided with the following security protection:

- Reusable passwords—Users likely authenticate to their systems with username and passwords.
- Session security—Application crypto-any communication between a client to a server considered sensitive (based on the security policy) should be cryptographically protected with session-application crypto.
- OS/application hardening—Harden the OS and any application. Do not simply deploy every patch as it is released. Use some mechanism to do testing on updates before applying to IACS systems. Also, make sure to follow hardening guides for popular applications, such as Microsoft Internet Information Server (IIS) and Apache web server, used on the servers.
- Partitioning disk space—In the event of a problem, limit the ability of one rogue process to consume the entire disk space of the server. In Unix, for example, it is good practice to set aside separate partitions for the following components: /, /var, /home, /usr, and /tmp.
- Turning off unneeded services—If the host is a standard desktop, it probably does not need to run any services for other users such as FTP. If it is a server, the running services should be limited to those that are required to perform the job of the server. For example, this means running HTTP but not Telnet on a web server.
- Deploying the Cisco Security Agent (CSA)—The CSA protects critical servers by being a host-based IDS to help mitigate local attacks. See Endpoint Protection with Cisco Security Agent, page 5-33.

Endpoint Protection with Cisco Security Agent

No security strategy can be effective if the servers and desktop computers (endpoints) are not protected. Endpoint attacks typically run in stages: probe, penetrate, persist, propagate, and paralyze. Most endpoint security technologies provide early stage protection (and then only when a signature is known).

The Cisco Security Agent (CSA) proactively defends against damage to a host throughout all stages of an intrusion, and is specifically designed to protect against new attacks where there is no known signature. The CSA goes beyond conventional endpoint security solutions by identifying and preventing malicious behavior before it can occur, thereby removing potential known and unknown security risks that threaten enterprise networks and applications.

When an application attempts an operation, the agent checks the operation against the security policy of the application. The agent makes a real-time "allow" or "deny" decision on its continuation and determines whether that request should be logged. Because protection is based on blocking malicious behavior, the default policies stop both known and unknown attacks without needing updates. Correlation is performed both at the agent and the management center console. Correlation at the agent results in dramatically increased accuracy, identifying actual attacks or misuse without blocking legitimate activity. Correlation at the management center identifies global attacks such as network worms or distributed scans.

Server Farm Access Layer

Layer 2 Access Considerations

The Layer-2 access switch provides physical connectivity to the server farm. The applications residing on these servers for the Manufacturing zone are considered to be critical for the operation of the IACS. It is recommended that these servers be dual-homed to the Layer 2 access switches through NIC teaming.

The Layer-2 access switch is connected to the aggregation layer through an IEEE 802.1Q trunk. The first point of Layer 3 processing is at the distribution switch. There is no Layer-3 routing done in the access switch.

Spanning VLANs across Access Layer switches

For applications that require spanning VLANs across Layer 2 access switches, and Spanning Tree protocol (STP) is used as the resiliency protocol, take the following steps to make the best of this suboptimal situation:

- Use RPVST+ as the version of STP. When Spanning Tree convergence is required, RPVST+ is superior to PVST+ or STP.
- Provide a Layer 2 link between the two distribution switches to avoid unexpected traffic paths and multiple convergence events.

If you choose to load balance VLANs across uplinks, be sure to place the HSRP primary and the RPVST+ primary on the same distribution layer switch. The HSRP and RPVST+ root should be collocated on the same distribution switches to avoid using the inter-distribution link for transit.

For more information, see *Campus Network Multilayer Architecture and Design Guidelines* at the following URL:

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cOverall_design.html

Figure 4-13 shows an example of a Layer-2 access topology.

Figure 4-13 Layer 2 Access Topology

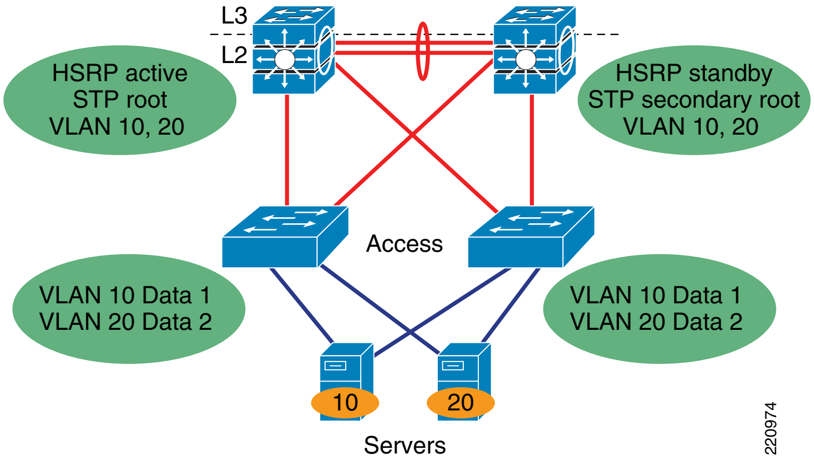

Layer-2 Adjacency Requirements

When Layer 2 adjacency exists between servers, the servers are in the same broadcast domain, and each server receives all the broadcast and multicast packets from another server. If two servers are in the same VLAN, they are Layer 2 adjacent. The requirement of Layer 2 adjacency is important for high availability clustering and NIC teaming.

NIC Teaming

Mission-critical business applications cannot tolerate downtime. To eliminate server and switch single point-of-failure, servers are dual-homed to two different access switches, and use NIC teaming drivers and software for failover mechanism. If one NIC card fails, the secondary NIC card assumes the IP address of the server and takes over operation without disruption.

NIC teaming features are provided by NIC vendors. NIC teaming comes with three options:

- Adapter Fault Tolerance (AFT)
- Switch Fault Tolerance (SFT)—One port is active and the other is standby, using one common IP address and MAC address.
- Adaptive Load Balancing (ALB) (a very popular NIC teaming solution)—One port receives and all ports transmit using one IP address and multiple MAC addresses.

Figure 4-14 shows examples of NIC teaming using SFT and ALB.

Figure 4-14 NIC Teaming

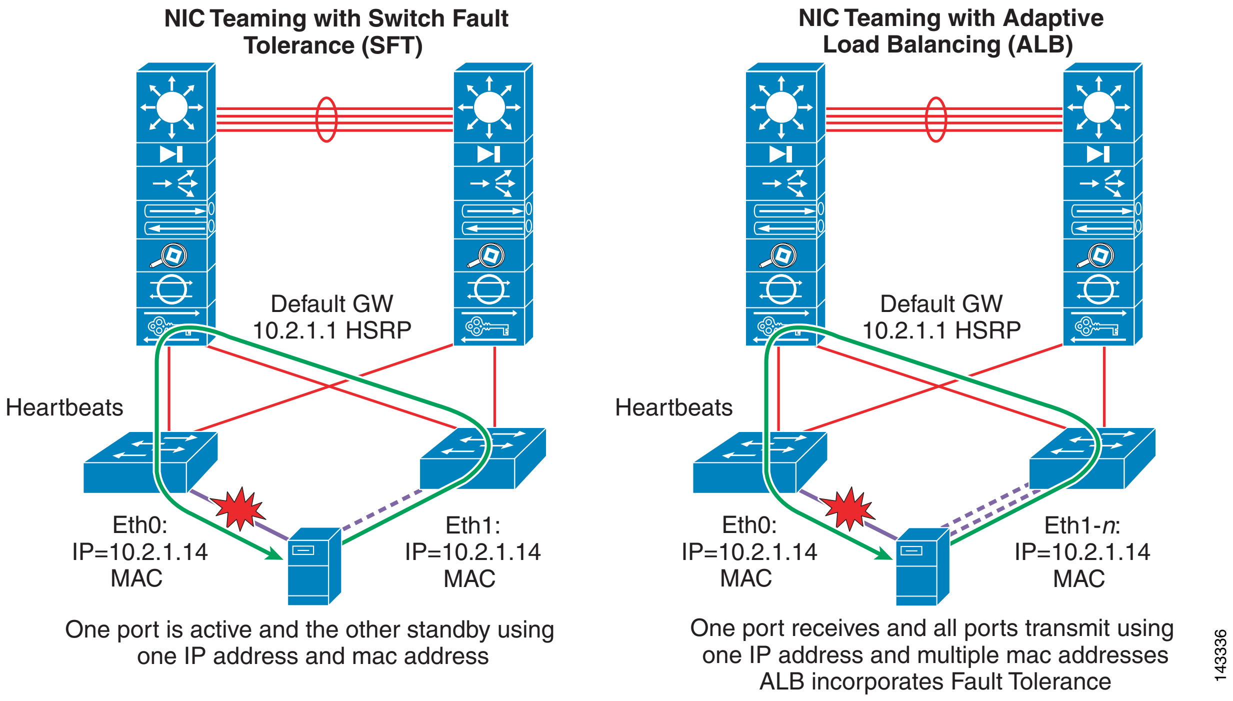
NIC Teaming with Switch Fault Tolerance (SFT)
Default GW
10.2.1.1 HSRP
Heartbeats
Eth0:
IP=10.2.1.14
MAC
Eth1:
IP=10.2.1.14
MAC
One port is active and the other standby using one IP address and mac address
NIC Teaming with Adaptive Load Balancing (ALB)
Default GW
10.2.1.1 HSRP
Heartbeats
Eth0:
IP=10.2.1.14
MAC
Eth1-*n*:
IP=10.2.1.14
MAC
One port receives and all ports transmit using one IP address and multiple mac addresses
ALB incorporates Fault Tolerance
143336

The main goal of NIC teaming is to use two or more Ethernet ports connected to two different access switches. The standby NIC port in the server configured for NIC teaming uses the same IP and MAC address of a failed primary server NIC, which results in the requirement of Layer 2 adjacency. An optional signaling protocol is also used between active and standby NIC ports. The protocol heartbeats are used to detect the NIC failure. The frequency of heartbeats is tunable to 1 to 3 seconds. These heartbeats are sent as a multicast or a broadcast packet and therefore require Layer 2 adjacency.

Security and Network Management

The security and network management services are in the Manufacturing zone for security and availability considerations; they are considered critical to the plant floor operations. Therefore, they can be used to attack a system, or with the service, the plant floor operations may be jeopardized. The most secure location for these services is behind the DMZ firewall in the Manufacturing zone.

These services are not typically critical to the operation of the plant floor. If they fail, services should be restored as soon as possible, but it is not likely that IACS will be directly impacted.

There are situations and environments where critical audit and control procedures may dictate that these systems be operational to maintain logs and audit trails of activity in the Manufacturing zone. In this case, these applications may then require a higher level of availability, which can be achieved in various ways.

Although this CPwE solution architecture does not provide specific implementation guidance, key considerations to increase availability include the following:

- All workstations or servers with security or network management applications should be backed up, and scheduled testing of the integrity of the backup should be performed.
- Redundant servers or workstations capable of continuing operations should be deployed.

- Redundant network connectivity on the servers running the applications adds a level of network resiliency.

Security Monitoring, Analysis, and Mitigation with CS-MARS

The Cisco Security Monitoring, Analysis, and Response System (CS-MARS) is an appliance-based, all-inclusive solution that allows network and security administrators to monitor, identify, isolate, and counter security threats. High-performance, scalable threat mitigation appliances fortify deployed network devices and security countermeasures by combining network intelligence with features such as Context Correlation, SureVector analysis, and AutoMitigate capability, empowering companies to readily identify, manage, and eliminate network attacks and maintain compliance.

Going beyond first- and second-generation security information management systems, CS-MARS more efficiently aggregates and reduces massive amounts of network and security data from popular network devices and security countermeasures. By gaining network intelligence, it effectively identifies network and application threats through sophisticated event correlation and threat validation. Verified attacks are visualized through an intuitive, detailed topology map to augment incident identification, investigation, and workflow. Upon attack discovery, the system allows the operator to prevent, contain, or stop an attack in real-time by pushing specific mitigation commands to network enforcement devices. The system supports manufacturer-centric rule creation, threat notification, incident investigation, and a host of security posture and trend reports.

CS-MARS appliances consist of standard Intel platforms with availability features accessible through a web-based user interface, hardened OS, embedded Oracle database, proprietary logic, and scalable architecture options.

FactoryTalk

At the beginning of this chapter, a guideline was established that all critical applications and services required for maintaining and operating the plant floor should be located within the Manufacturing zone. This includes the IACS applications that provide site-wide services and functions, such as FactoryTalk.

FactoryTalk consists of a services platform and modular IACS disciplines (hereafter referred to as applications) that tightly integrate with the Rockwell Automation Logix Control Platform, helping to deliver a seamless flow of valuable manufacturing data. The Rockwell Automation Integrated Architecture is comprised of FactoryTalk and Logix, together providing plantwide control and information. See Figure 4-15.

Figure 4-15 FactoryTalk

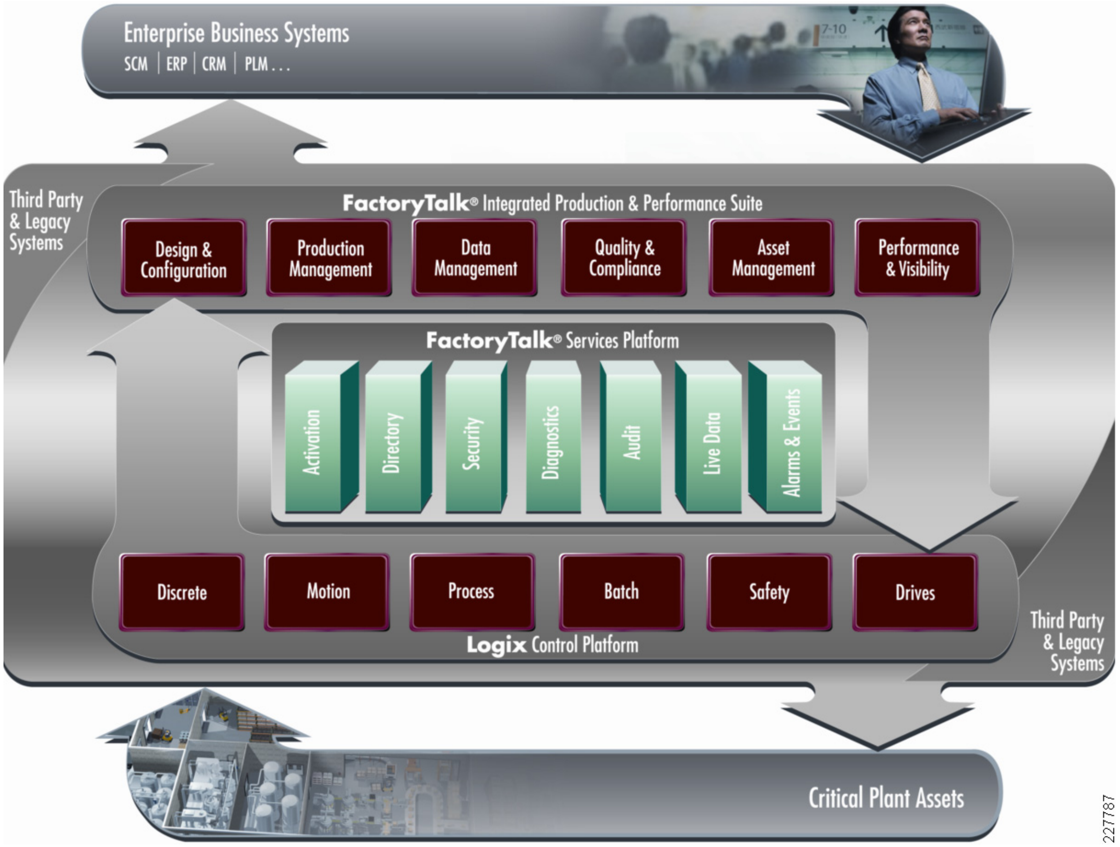

The modular system design of FactoryTalk supports incremental solution deployments to help maximize legacy technology investments, while improving the ability to incorporate new technologies.

The FactoryTalk Services platform is the foundation of the FactoryTalk applications. Comprised of a set of common software services that form a service-oriented architecture (SOA), FactoryTalk Services platform allows applications to be developed that share common definitions, administration, real-time data, and so on. The FactoryTalk Services platform is grouped by functionality.

- FactoryTalk Security allows centralized management of each user's rights and privileges based on their role and location.
- FactoryTalk Directory allows the sharing of common definitions such as users, tags, alarms or graphic displays.
- FactoryTalk Diagnostics and Audit provides common message formats, storage and viewing and also tracks the changes made in an application.
- FactoryTalk Live Data allows real-time communication between software applications as well as third-party OPC data servers.
- FactoryTalk alarms and events provide unified alarm definitions and common management between Logix programmable automation controllers (Logix PAC™) and software applications.

Figure 4-16 depicts the FactoryTalk application suite positioned within the CPwE architecture.

Figure 4-16 FactoryTalk Application

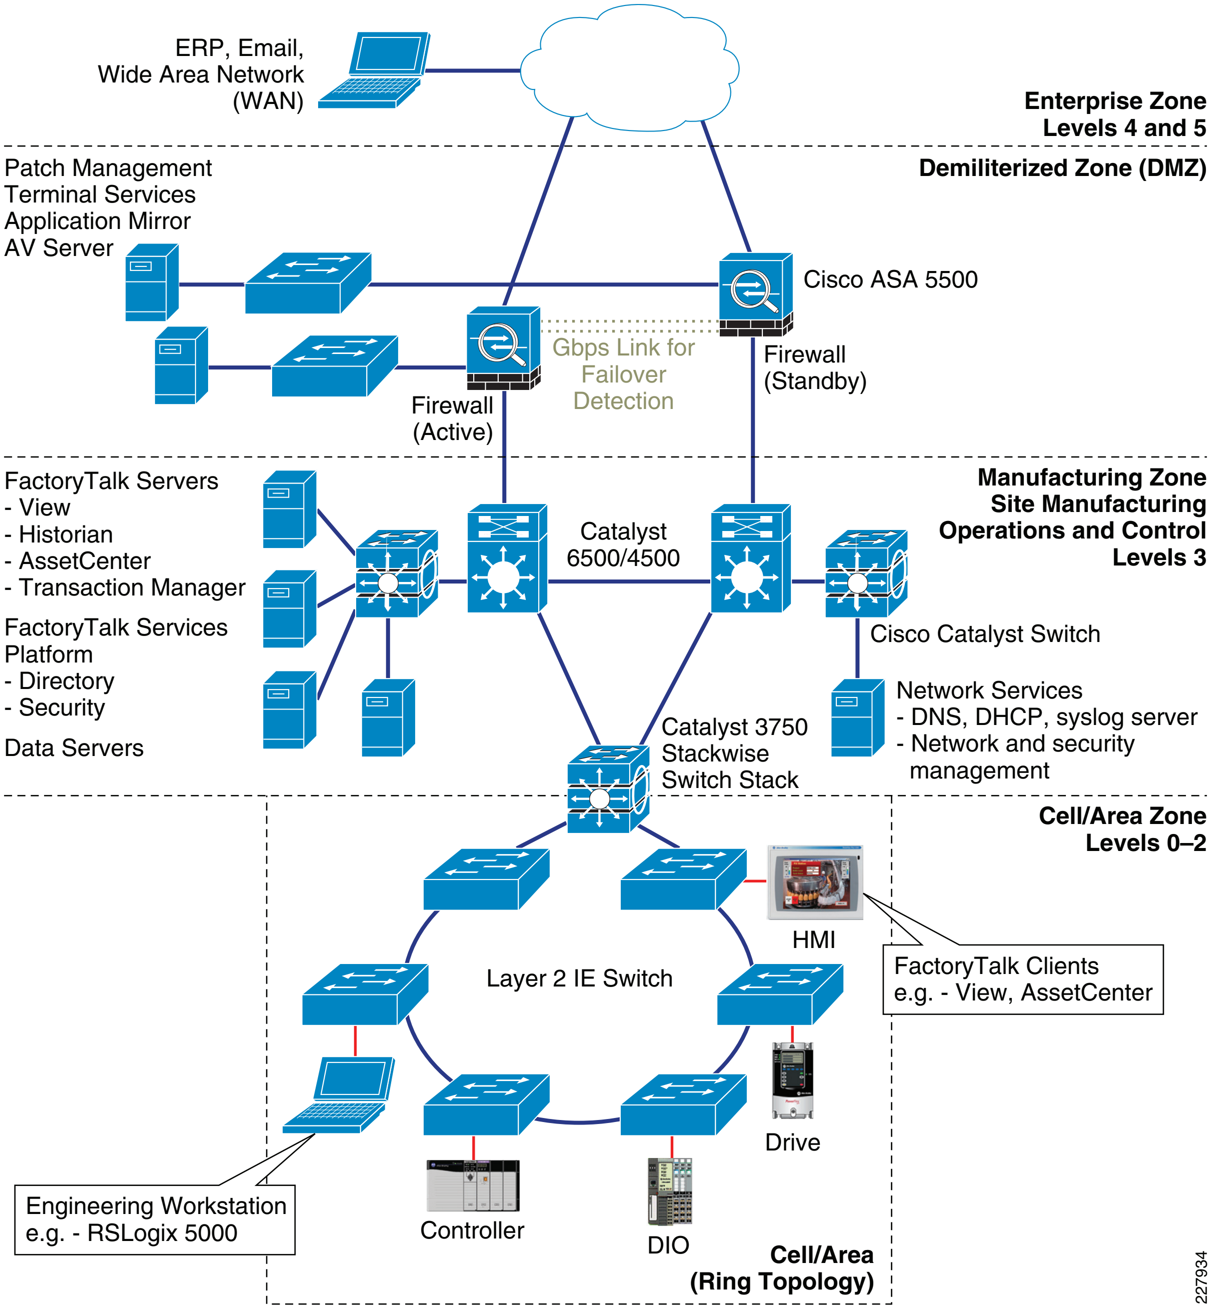

Demilitarized Zone Network Design

Given that the Enterprise zone and the Manufacturing zone have different requirements, priorities, policies, and implications of incidents, and that it's desirable that they be able to share data and access systems and applications, CPwE introduces a fourth zone to provide insulation. This is the Demilitarized zone. Systems and data that need to be accessed by both manufacturing and enterprise business systems reside in the DMZ, protecting information and accommodating the different security requirements of these major zones. As a best practice, all traffic should terminate in the DMZ, eliminating direct traffic flow between the Enterprise zone and the Manufacturing zone.

Although it is possible to deploy a firewall without a DMZ, Cisco and Rockwell Automation do not recommend this option because the required "holes" in the firewall to allow data to be shared weaken the defense-in-depth security stance.

Servers that users from both networks need to access are put in a separate Demilitarized zone (DMZ) network that is connected to the same firewall or separate firewalls. To provide more granular network access, the Cisco ASA provides authentication, authorization, and accounting (AAA) services by working in conjunction with the Cisco Secure Access Control Server (ACS). This provides a user database of which the Cisco ASA can inquire to identify and validate before permitting the transmission of traffic to the destination network.

In addition to controlling traffic access between the three zones, the Cisco ASA can optionally be installed with the Cisco Adaptive Inspection Prevention Security Services Module (AIP-SSM) to provide intrusion protection and detection services to prevent network attacks to those destinations to which the firewall function of the Cisco ASA permits network access.

Finally, all the servers placed in the DMZ need to be secured. Refer to the "Security Protection for Servers" section on page 4-48 for more information.

The DMZ network design covers the following:

- DMZ components
- DMZ topology
- Firewall design and implementation considerations

DMZ Components

The DMZ (see Figure 4-17) consists of the following:

- Plant firewall(s) to provide the strong segmentation between the zones
- DMZ switches to provide inter-connectivity between any servers and the firewall(s)
- Firewall and security management software to manage the firewall, although this may not reside in the DMZ network zone
- Servers to run any application mirrors or store data to be shared

Figure 4-17 DMZ Components

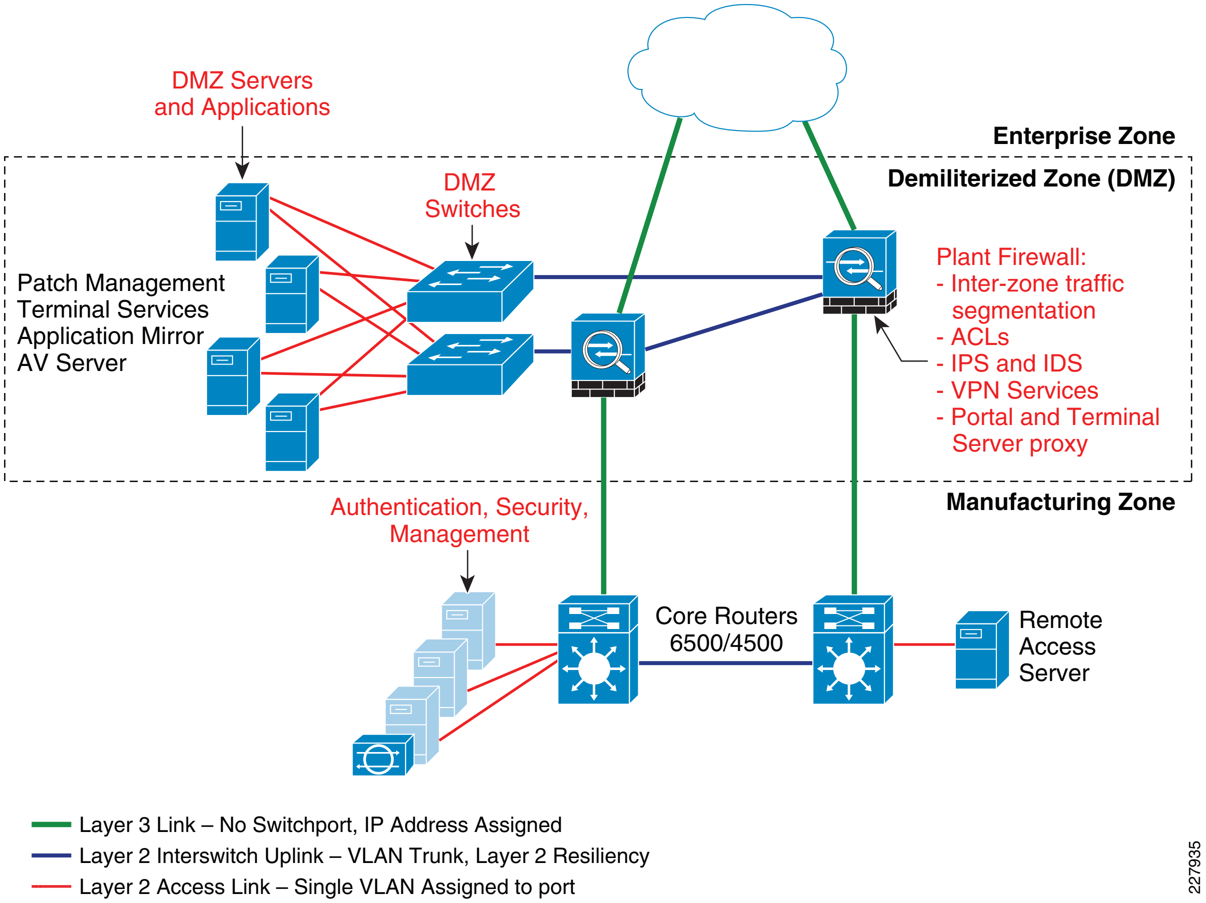
DMZ Servers and Applications
Enterprise Zone
Demiliterized Zone (DMZ)
DMZ Switches
Patch Management
Terminal Services
Application Mirror
AV Server
Plant Firewall:
- Inter-zone traffic segmentation
- ACLs
- IPS and IDS
- VPN Services
- Portal and Terminal Server proxy
Manufacturing Zone
Authentication, Security, Management
Core Routers 6500/4500
Remote Access Server
Layer 3 Link – No Switchport, IP Address Assigned
Layer 2 Interswitch Uplink – VLAN Trunk, Layer 2 Resiliency
Layer 2 Access Link – Single VLAN Assigned to port
227935

This section addresses the selection of the network components—namely the firewall, firewall and security management software, and DMZ switches. The selection and components involved with the DMZ applications and services must be determined by the plant personnel and other relevant groups.

The key considerations for the components in this zone are described in the following subsections.

Cost

Although cost is always a consideration in manufacturing facilities, the cost of the firewalls is dependent on the functionality and scalability required.

Industrial Characteristics

As stated above, the industrial characteristics for the DMZ are less critical because it is assumed the devices are in controlled environments.

It is recognized, however, that there is a need in some manufacturing environments for the firewall components to exist in industrial environments. This requirement is not addressed in this *DIG*.

Performance and Real-Time Communications

Traffic through the plant firewalls and DMZ is typically not critical in nature such as that of the Manufacturing and Cell/Area zones. Typical performance considerations for a firewall include the following:

- Firewall throughput measured in Mbps
- VPN throughput and concurrent sessions
- Number of Security Contexts (e.g., DMZ to Enterprise or DMZ to Manufacturing zone)[1]
- Clustering and load balancing, although not generally a requirement for plant firewalls

Information Convergence

Information convergence between manufacturing and business systems has provided manufacturers with greater business agility and opportunities for innovation. With these opportunities, come challenges. Manufacturing computing and IACS controller assets have become susceptible to the same security vulnerabilities as their enterprise counterparts due to this convergence. Securing manufacturing assets such as FactoryTalk 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. Mitigation techniques include the use of a defense-in-depth security approach that addresses internal and external security threats. This approach utilizes multiple layers of defense (physical and electronic) at separate manufacturing levels by applying policies and procedures that address different types of threats. For example, multiple layers of network security protect networked assets, data, end points, and multiple layers of physical security to protect high value assets. No single technology or methodology can fully secure an IACS.

Given the different requirements, priorities, policies, and implications of incidents between the Enterprise zone and the Manufacturing zone, and the desire to share data, a DMZ should be used as a mitigation technique to provide a buffer zone between the Manufacturing and Enterprise zones. The DMZ can allow data that needs to be accessed by manufacturing and business systems to be shared securely, protecting information and accommodating the different security requirements of these zones.

The methodology used to traverse information across the DMZ depends on the manufacturer's security policy, which determines the acceptable approach and risk.

One example is to temporarily transfer Manufacturing zone information into the DMZ, and then replicate this information up to the Enterprise zone. This can be either unidirectional or bidirectional. This example is shown in Figure 4-17. This example uses FactoryTalk Transaction Manager to provide two-way data exchange between tags, such as Logix programmable automation controller (PAC) or FactoryTalk View tags, and applications like an MSSQL server. These tags may contain key performance indicators (KPIs) or other important data that need to be integrated into an enterprise application.

In the example provided below, IACS data is collected and transferred to a business system in the Enterprise zone. The data is neither stored nor used in the Manufacturing zone, so DMZ connectivity disruption will not affect Manufacturing zone operations. A methodology should be deployed to buffer IACS data to and from the Enterprise zone in the event of DMZ connectivity disruption.

[1] Using multiple contexts and VPN technologies used for Secure Remote Access is not currently supported.

- The FactoryTalk Transaction Manager server (Level 3) uses the RSLinx Data Server to read/write tags to controllers in Level 1 using EtherNet/IP.
- This same FactoryTalk Transaction Manager server is configured to read/write its SQL data to and from an MSSQL server located in the DMZ.
- This MSSQL server replicates the data to and from the Enterprise zone MSSQL server.
- Business systems within the Enterprise zone only access the enterprise MSSQL server.

Figure 4-18 FactoryTalk Transaction Manager and MSSQL Server

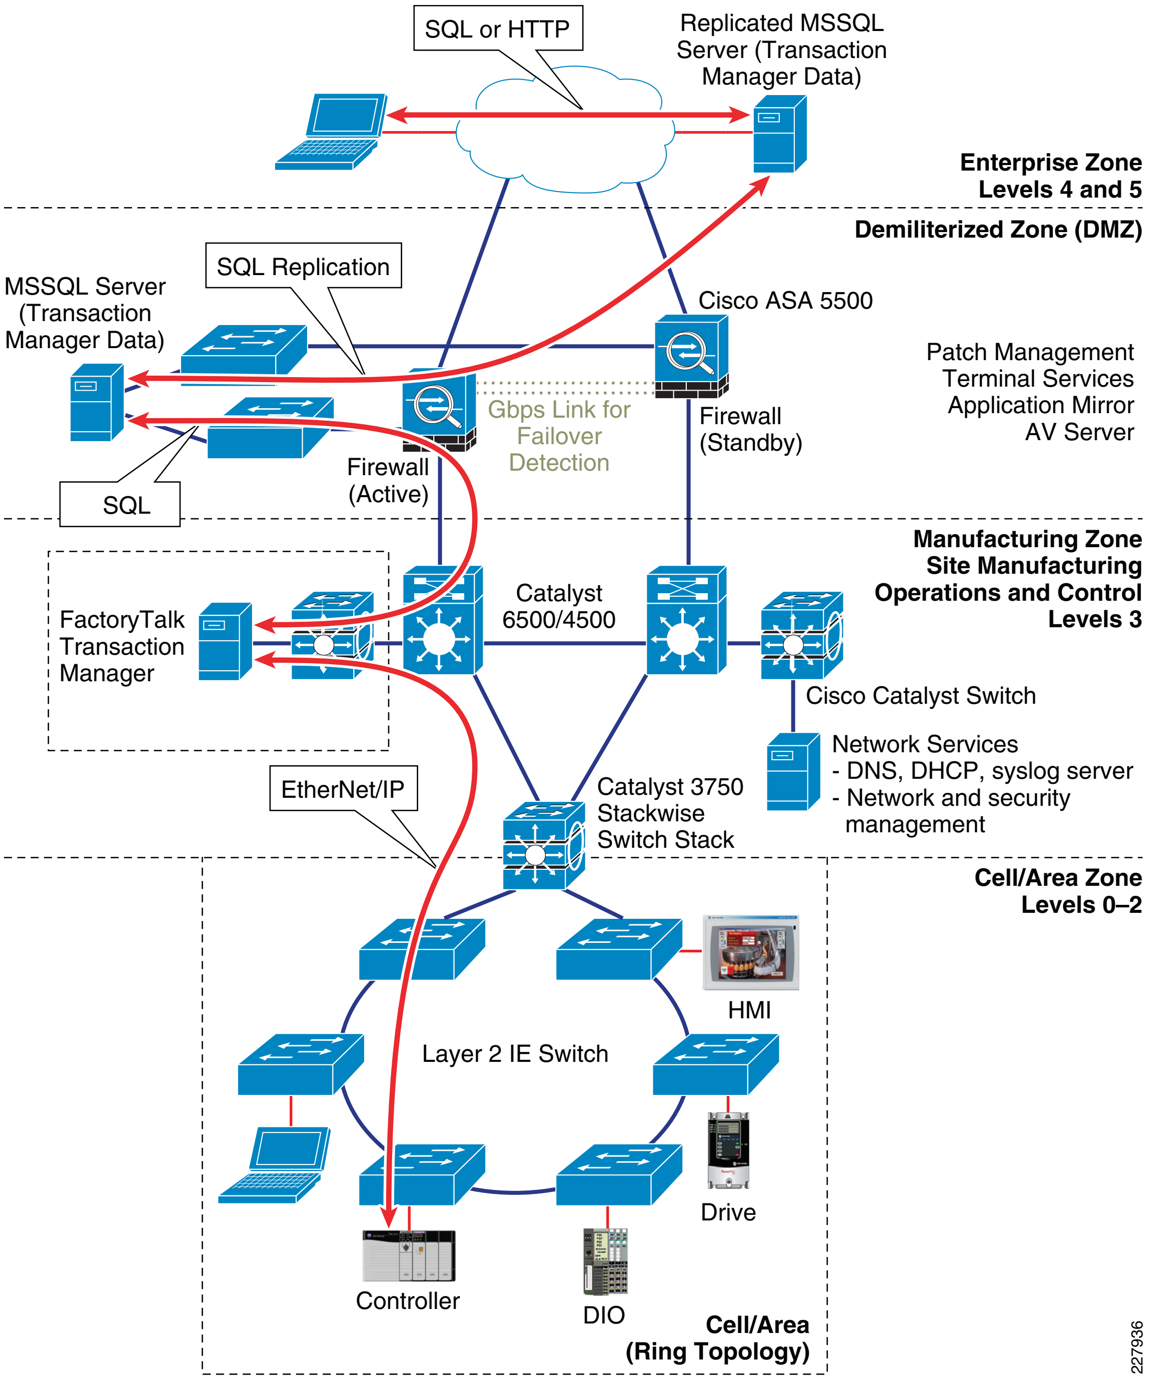

In addition to information convergence, applications such as web-based monitoring applications and personnel such as an engineer or partner may require remote access to manufacturing assets for the purpose of monitoring, management and configuration. This remote access to manufacturing assets can occur from either the enterprise or the Internet. This remote access is covered in Chapter 6, "IACS Network Security and the Demilitarized Zone."

The convergence of manufacturing and enterprise networks has provided greater access to manufacturing data, which has led to greater agility in making business decisions for manufacturers. The resulting agility has provided manufacturers who have embraced the convergence trend with a competitive edge.

Network convergence has also exposed IACS assets to security threats that were traditionally found in the enterprise. Securing IACS assets such as FactoryTalk requires a comprehensive security model based on a well-defined set of security policies, and the use of a defense-in-depth security approach that addresses internal and external security threats. This approach utilizes multiple layers of defense (physical and electronic) at separate IACS levels by applying policies and procedures that address different types of threats.

General recommendations for securing IACS assets include the following:

- Establish a DMZ between the Enterprise and Manufacturing zones.
- Keep FactoryTalk applications and Services Platform within the Manufacturing zone.
- Keep replicated services such as Active Directory within the Manufacturing zone.
- Use a team consisting of IT, operations and engineering professionals to define a Manufacturing zone security policy to address manufacturing needs:
  - DMZ information convergence—firewall and trust policies
  - Remote access for engineers and partners
- Use application data replication within the DMZ to converge Manufacturing and Enterprise zone information.

Availability

Although by definition the IACS application and IACS network should not rely on the DMZ infrastructure for availability, inter-connectivity to the Enterprise and functions like remote access depend on the DMZ. Thus, availability considerations are important and include the following:

- Support for redundant firewalls
- Mean-time to break/fix ratings

Manageability

Network and security management services are relevant to the DMZ, but these functions tend to be located in other network zones. These applications must be relatively easy to install, configure, and operate. They must relax the level of expertise and knowledge required by the plant floor personnel to maintain and operate the IACS network. Key considerations for this equipment include the following:

- Intuitive web-based interfaces via secure connections (for example, HTTPS)
- Ease of installation and upgradeability
- Ease of configuration and auto-detect functions to find and interface with appropriate network/security infrastructure
- Intuitive summarization of network and security status with easy-to-use drill-down features for problem solving
- Ability to develop templates for security/network management and to apply those throughout the Manufacturing zone

- Built-in knowledge repositories to assist plant floor personnel and Control Engineers during problem resolution
- Ability to securely enable access to plant floor personnel and partners
- Allow both IT and Manufacturing personnel to manage separate firewall "contexts" for segmentation of responsibilities[1]

In addition to the actual network and security management applications, there are also manageability considerations for the network infrastructure, especially the Layer-3 switches and routers. Basic management functions such as initial configuration, break/fix, and monitoring need to be relatively easy. Key considerations include the following:

- SNMP capable-Most network device vendors support management via the Simple Network Management Protocol (SNMP)v3.
- Ease of network infrastructure installation, setup, and maintenance.
- Warranty and support options for the expected lifetime of the equipment
- Web-based, intuitive user interfaces
- Application interfaces (for example, XML support) to interface with other applications

Security

The DMZ itself is a security key aspect of the defense-in-depth security stance. The plant firewall is critical to the security of the Manufacturing zone and the IACS systems that it contains. Therefore, the following features should be considered for the plant firewall appliances:

- Stateful packet inspection and access control support to manage traffic flows
- Content security
- Malware detection
- VPN technology
- Intrusion protection and detection services
- Modular policy support
- Authentication (AAA) enforcement and ability to setup local authentication
- Web portal, support for remote desktop and other proxy services
- NAT services
- Support for SSH, HTTPS, and SNMPv3

[1] Using multiple contexts and VPN technologies used for Secure Remote Access is not currently supported.

Component Summary

For the purpose of testing, the products listed in Table 4-8 were part of the DMZ.

Table 4-8 DMZ Components

Role	Product/Platform	Software Release	Comments
Plant Firewall	Cisco Adaptive Security Appliance 5500		Provide strong segmentation and security features
DMZ Switch	Catalyst 2960 or 3000 series		

Plant Firewall

For the solution testing, the ASA 5505, 5510, and 5520 devices were selected and tested. Note that ASA 5505 does not support IPS and IDS, but may be sufficient for small plants or Manufacturing zones. Any firewall appliance or module in the series should suffice for the plant firewall, depending on scalability requirements. See Figure 4-19.

Figure 4-19 ASA 5500 Firewall Appliances

Firewall Management

For the Firewall Management, the Adaptive Security Device Manager suffices for basic management. This user-friendly application enables quick configuration, monitoring, and helps troubleshoot Cisco firewall appliances and firewall service modules. Ideal for small or simple deployments, the Cisco Adaptive Security Device Manager provides the following:

- Setup wizards that help you configure and manage Cisco firewall devices without cumbersome command-line scripts
- Powerful real-time log viewer and monitoring dashboards that provide an at-a-glance view of firewall appliance status and health
- Handy troubleshooting features and powerful debugging tools such as packet trace and packet capture

DMZ Switches

For the DMZ switches, either the Catalyst 3560, 3750, or the 2960 Series switches will meet many application requirements. See Figure 4-20 and Figure 4-21.

Figure 4-20 Catalyst 3560 and 3750 Series Switches

Figure 4-21 Catalyst 2960 Series LAN switches

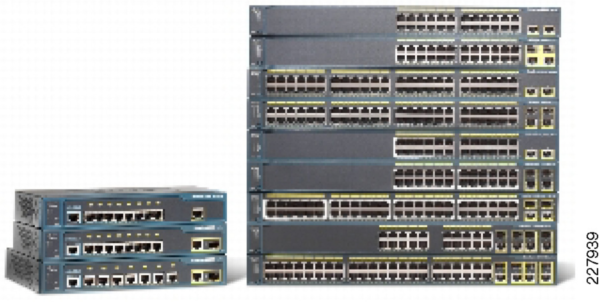

Topology Options

The topology for the DMZ highlights the following:

- Dual firewalls for high availability running in either active-active or active-standby mode, depending on the type of firewall procured and the speed of failover if a firewall fails.
- DMZ zone switched network

Figure 4-22 depicts the typical DMZ configuration.

Figure 4-22 Classic DMZ Topology

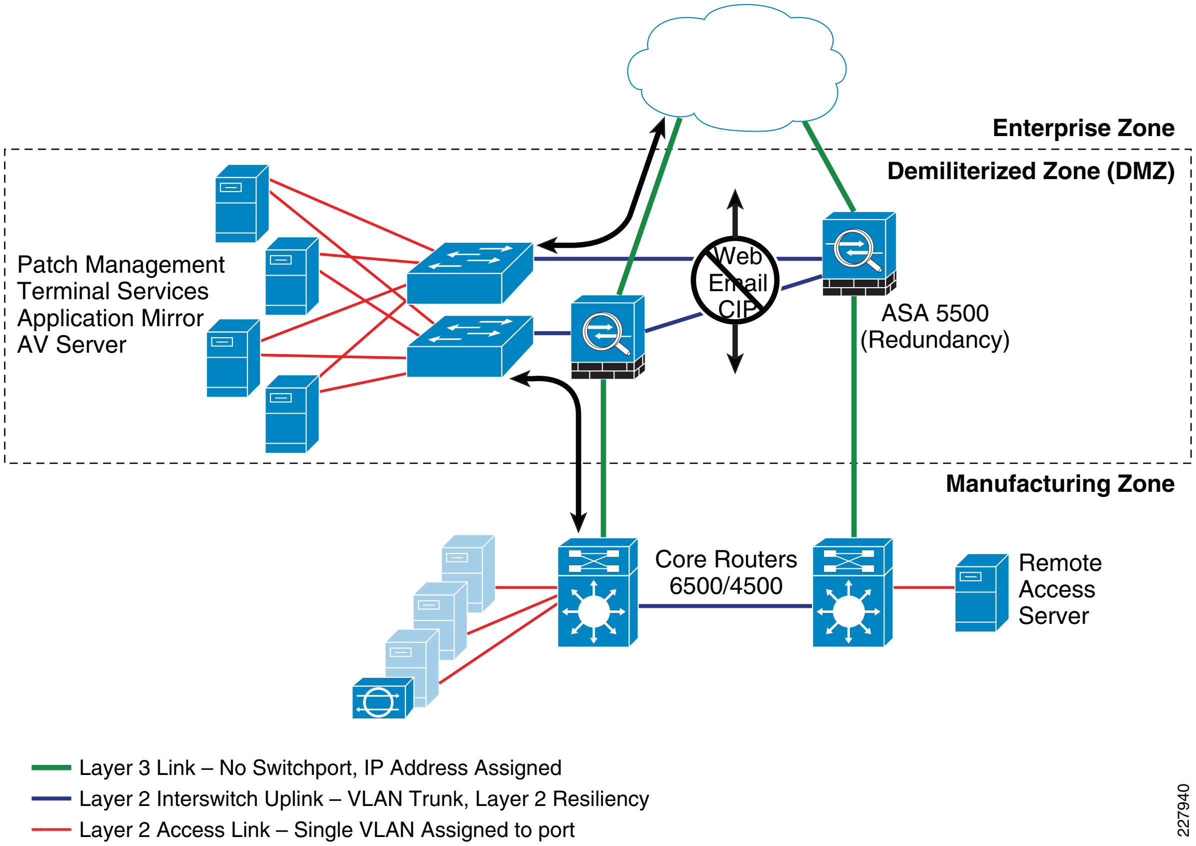

Figure 4-22 depicts a typical DMZ topology with dual firewalls for resiliency. Each firewall can support two or more firewall contexts to manage traffic between the Enterprise and DMZ as well as the DMZ and the Manufacturing zone. Separate "contexts" enables two organizations to separately manage and maintain firewall rules applied to the two traffic flows; one between the Enterprise zone and DMZ and two between the Manufacturing zone and DMZ. Multiple contexts and VPN technology used for secure remote access (see "Remote Access to the IACS Network" section on page 6-16) on the same firewall is not currently supported. If separating the operational responsibility and secure remote access is required, a second pair of firewalls is currently required.

The topology is scalable via more powerful versions of the ASA appliance.

Firewall Design and Implementation Considerations

Security Levels on the Cisco ASA Interfaces

The Cisco ASA uses the concept of assigning security levels to its interfaces. The higher the security level, the more secure an interface is. The security level is thus used to reflect the level of trust of this interface with respect to the level of trust of another interface on the Cisco ASA.The security level can be between 0 and 100. The most secure network is placed behind the interface with a security level of 100. The security level is assigned by using the security-level command.

In the *EttF 1.2 Design and Implementation Guide*, Cisco recommends creating three network zones in different security levels, as shown in Table 4-9.

Table 4-9 Network Security Levels

Network	Security Level	Interface
Enterprise network	0	G0/2
DMZ	50	G0/1
Manufacturing network	100	G0/0

Configuration Example

Refer to Figure 4-23 for the subsequent configuration example.

Figure 4-23 Security Levels on the Interfaces of the Cisco ASA 5500

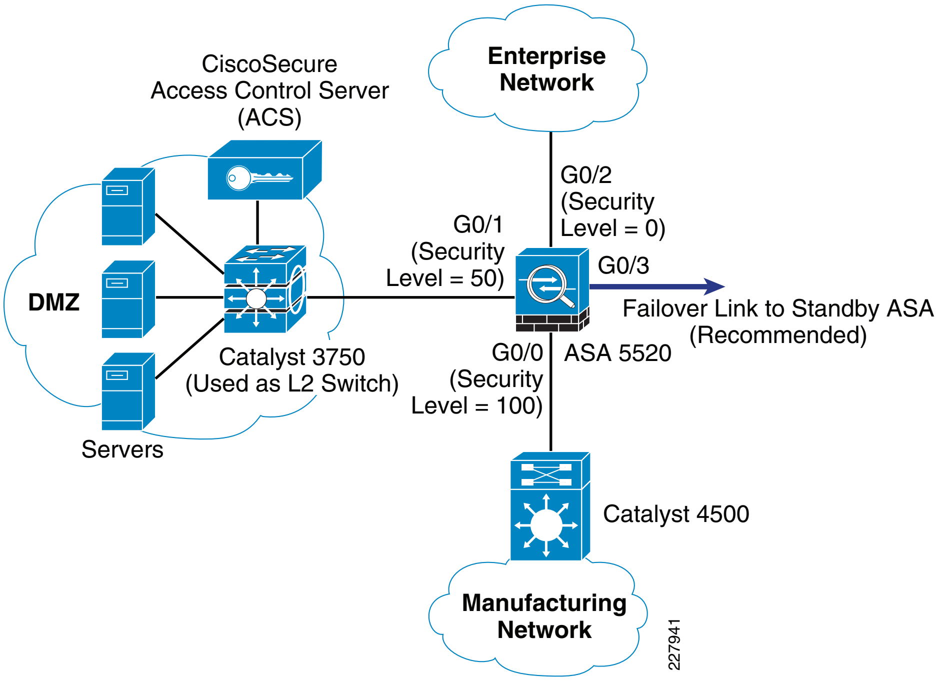

Based on the security level recommendations above, the following shows how to configure the levels on the interfaces of the Cisco ASA 5520 platform:

- GigabitEthernet 0/0 is the interface connected to the manufacturing IACS network. It is named inside. Because it is at security level 100, it has the highest security level.

```
interface GigabitEthernet0/0
nameif inside
security-level 100
ip address 10.18.1.1 255.255.255.0
```

- GigabitEthernet 0/1 is the interface connected to the manufacturing IACS network. It is named outside with security level set to 0.

```
interface GigabitEthernet0/1
nameif outside
security-level 0
ip address 10.13.2.1 255.255.255.248
```

- GigabitEthernet 0/2 is the interface connected to the DMZ. It is named DMZ with security level 50.

```
interface GigabitEthernet0/2
nameif dmz
security-level 50
ip address 10.19.2.9 255.255.255.248
```

The command name is used to assign a name to an interface. This interface name is used to set up any configuration feature associated to the given interface.

Note that the IP address configuration includes an optional parameter standby. It is used for configuring the standby Cisco ASA in the solution.

By default, the ASA 5500 implicitly permits traffic that enters the ASA via a high security level interface and leaves via a low security level interface, but the appliance implicitly denies traffic in the reverse direction. However, the CPwE solution recommends that traffic be denied going from the manufacturing IACS network (security level 100) to the enterprise network (security level 0). An ACL needs to be explicitly configured to meet this access policy.

Stateful Packet Filtering

The Cisco ASA in the DMZ between the manufacturing IACS network and enterprise network enables the definition of policies and rules that identify what traffic should be permitted in or out of an interface. It uses ACLs to drop unwanted or unknown traffic when it attempts to enter the trusted networks.

An ACL, starting with a keyword access-list, is a list of security rules and policies grouped together that allows or denies packets after looking at the packet headers and other attributes. Each permit or deny statement can classify packets by inspecting up to Layer 4 headers for a number of parameters:

- Layer 2 protocol information such as EtherTypes
- Layer 3 protocol information such as ICMP, TCP, or UDP
- Source and destination IP addresses
- Source and destination TCP or UDP ports

After an ACL has been properly configured, it can be applied to an interface to filter traffic with the keyword access-group. The Cisco ASA can filter packets in both the inbound and outbound direction on an interface. When an inbound ACL is applied to an interface, the security appliance inspects against the ACL parameters after receiving or before transmitting them. An incoming packet is screened in the following sequence:

Step 1 If this packet matches with an existing connection in the firewall connection table, it is allowed in. If it does not, go to Step 2.

Step 2 The firewall tries to match the packet against the ACLs sequentially from the top to the bottom. After the first matched ACL is identified, the packet is allowed in or dropped according to the action (permit or deny). If there is no match, go to Step 3.

Step 3 The security appliance drops all traffic that does not match any parameter defined in the ACL. There is an implicit deny at the end of all ACLs.

Note The interface ACL does not block packets destined for the IP addresses of the security appliance.

For the CPwE solution, general packet filtering recommendations are listed in Table 4-10 and shown in Figure 4-24.

Table 4-10 Packet Filtering Recommendations

Traffic Destination		Enterprise Network	DMZ	Manufacturing IACS Network
	Enterprise Network	N/A	Explicitly permitted by ACLs	Disallowed (explicitly denied by ACLs)
	DMZ	Explicitly permitted by ACLs	N/A	Explicitly permitted by ACLs
	Manufacturing IACS Network	Disallowed (implicitly denied by ACLs)	Explicitly permitted by ACLs	N/A

Figure 4-24 High-Level Packet Filtering Recommendations for the DMZ Between the Manufacturing IACS and Enterprise Networks

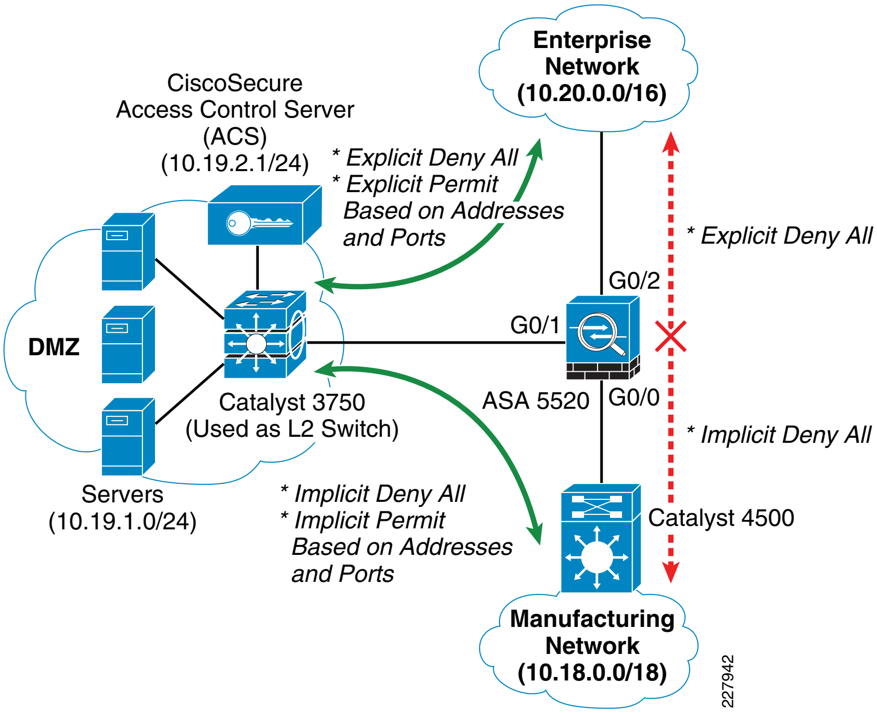
Enterprise Network (10.20.0.0/16)
CiscoSecure Access Control Server (ACS) (10.19.2.1/24)
* Explicit Deny All
* Explicit Permit Based on Addresses and Ports
* Explicit Deny All
G0/2
DMZ
G0/1
ASA 5520
G0/0
Catalyst 3750 (Used as L2 Switch)
* Implicit Deny All
Servers (10.19.1.0/24)
* Implicit Deny All
* Implicit Permit Based on Addresses and Ports
Catalyst 4500
Manufacturing Network (10.18.0.0/18)
227942

Configuration Example

Table 4-11 shows an example for ingress ACLs applied to the manufacturing IACS network-facing interface.

Table 4-11 Configuration Example for Ingress ACLs on the Manufacturing IACS Networking-Facing Interface

Applied To Interface	Traffic Direction	Permitted Traffic Types (Source to Destination)

Interface connected to the manufacturing IACS network (*inside*)	Inbound	HTTP (servers in the manufacturing IACS network to servers in DMZ)
		`access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq www`
		HTTPS (any in the manufacturing IACS network to servers in DMZ)
		`access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq https`
		Telnet (any in the manufacturing IACS network to host 10.19.1.10 in the DMZ)
		`access-list inside extended permit tcp 10.18.0.0 255.255.0.0 host 10.19.2.1 eq telnet`
		ICMP (any in the manufacturing IACS network to servers in the DMZ)
		`access-list inside extended permit icmp 10.18.0.0 255.255.0.0 10.19.2.0 255.255.255.0`
		Explicitly deny other traffic types to anywhere (i.e. DMZ and enterprise networks)
		`access-list inside deny 10.18.0.0 255.255.0.0`
		Apply the ACLs above to the ingress side of the manufacturing IACS network-facing interface
		`access-group inside in interface inside`

Table 4-12 shows an example for ingress ACLs applied to the enterprise network-facing interface.

Table 4-12 Configuration Example for Ingress ACLs on the Enterprise Networking-Facing Interface

Applied To Interface	Traffic Direction	Permitted Traffic Types (Source to Destination)

Interface connected to the enterprise network (*outside*)	Inbound	Telnet (any in the enterprise network to the DMZ [10.19.0.0/16])
		`access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq telnet`
		HTTP (any in the enterprise network to the DMZ [10.19.0.0/16])
		`access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq www`
		HTTPS (any in the enterprise network to the DMZ [10.19.0.0/16])
		`access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq https`
		Explicitly deny other traffic types to anywhere
		`access-list inside deny 10.20.0.0 255.255.0.0`
		Apply the ACLs above to the ingress side of the enterprise network-facing interface
		`access-group outside in interface inside`

Authenticating Firewall Sessions for User Access to Servers in the DMZ

When users in the manufacturing IACS network or enterprise network want to access servers in the DMZ, the best practice is to enable authentication on the Cisco ASA. This involves validating the users based on their identity and predetermined credentials, such as passwords. The Cisco ASA can be configured to maintain a local user database or to use an external server for authentication. To communicate with an external authentication server, the Cisco ASA supports various protocols such as RADIUS, TACACS+, RSA SecurID, Windows NT, Kerberos, and LDAP.

The following steps show how the Cisco ASA authenticates an HTTP session originated from the enterprise network before the Cisco ASA permits the session to access the web server in the DMZ:

Step 1 The user on the outside of the Cisco ASA attempts to create an HTTP connection to the web server behind the ASA in the DMZ.

Step 2 The Cisco ASA prompts the user for authentication.

Step 3 The Cisco ASA receives the authentication information (userid and password) from the user and sends an AUTH Request to the Cisco Secure ACS.

Step 4 The server authenticates the user and sends an AUTH Accept message to the Cisco ASA.

Step 5 The Cisco ASA allows the user to access the web server.

Note For more details of the Cisco ACS, see the following URL: http://www.cisco.com/en/US/products/sw/secursw/ps2086/products_configuration_guide_book09186a0080721d25.html

Configuration Example

The following example illustrates how to use firewall session authentication in a IACS network. Plant XYZ wants to define the following policies on the ASA to specify which source addresses have rights to access to a server at 10.18.1.2 in the DMZ:

- Any user in the enterprise network can access the server at 10.18.1.2. The permitted protocols are HTTP and HTTPS.
- Only users in the 10.17.0.0/16 subnets in the IACS network can access the server. The permitted protocols are Telnet, HTTP, and HTTPS.

The users residing in these legitimate addresses are required for authentication before reaching out to the server.

Step 1 Define an AAA server group named ETTF2 using TACACS+ as the protocol for authentication. This AAA server is at 10.19.2.11.

```
aaa-server ETTF2 protocol tacacs+
aaa-server ETTF2 host 10.19.2.11
key Cisco
```

Step 2 Add the Cisco ASA as an AAA client in the CiscoSecure ACS.

Step 3 Create an ACL named INSAUTH that requires authentication of HTTP and HTTPS traffic.

```
access-list INSAUTH extended permit tcp 10.17.0.0 255.0.0.0 host 10.18.1.2 eq telnet
access-list INSAUTH extended permit tcp 10.17.0.0 255.0.0.0 host 10.18.1.2 eq www
access-list INSAUTH extended permit tcp 10.17.0.0 255.0.0.0 host 10.18.1.2 eq 8080
```

Step 4 Define the AAA match command to match the source and destination addresses of the incoming Telnet, HTTP, and HTTPS traffic from the IACS network (inside) against the ACL group INSAUTH.

```
aaa authentication match INSAUTH inside ETTF2
```

Step 5 Create ACLs named OUTAUTH that require authentication of HTTP and HTTPS traffic.

```
access-list OUTAUTH extended permit tcp any host 10.18.1.2 eq www
access-list OUTAUTH extended permit tcp any host 10.18.1.2 eq 8080
```

Step 6 Define the AAA match command to match the source and destination addresses of the incoming HTTP and HTTPS traffic from the enterprise network (outside) against the ACL group OUTAUTH.

```
aaa authentication match OUTAUTH outside ETTF2
```

Step 7 Define the AAA match command to match the source and destination addresses of the incoming HTTP and HTTPS traffic.

If there is an ACL without authentication, the firewall session authentication can be customized in the following ways:

- Authentication exception based on users
- Authentication timeouts

- Customization of authentication prompts

Integrating the ASA 5500 Appliance with the Adaptive Inspection Prevention Security Services Module

The Cisco ASA supports the Adaptive Inspection Prevention Security Services Module (AIP-SSM) running the Cisco Intrusion Prevention System (CIPS) software. Although the Cisco ASA can also provide IPS support with the **ip audit** command if an AIP-SSM module is absent, it supports only a limited number of signatures compared to the module. Also, these built-in signatures are not upgradeable.

Note For details on how to upgrade the image or signatures of the module, see the following URL: http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_configuration_guide_chapter09186a00807517ba.html.

Note The Cisco ASA 5510 and 5520, which is the ASA model recommended for the CPwE, supports both the AIP-SSM10 and AIP-SSM20 modules.

Access to the AIP-SSM Module

An administrator can connect to the AIP-SSM module via the following:

- Telnet and SSH to the FastEthernet management interface port on the module
- Telnet and SSH to the FastEthernet management interface port on the ASA and then the **session** <*module-number*> command to the AIP-SSM module
- HTTPS to Adaptive Security Device Manager (ASDM) on the ASA

Note For the initialization and maintenance of the AIP-SSM module, see the ASA documentation at the following URL: http://www.cisco.com/en/US/products/ps6120/products_getting_started_guide_chapter09186a00806a8347.html.

Inline Versus Promiscuous Mode

The Cisco AIP-SSM supports both inline and promiscuous modes. In the inline mode, the module can be considered to be an intrusion protection system (IPS); in the promiscuous mode, it can be considered to be an intrusion detection system (IDS). Cisco and Rockwell Automation recommend using promiscuous mode unless sufficient testing is complete and the environment requires the additional protection and can support the additional operational impact of an IPS.

When configured as an inline IPS, the AIP-SSM module can drop malicious packets, generate alarms, or reset a connection, allowing the ASA to respond immediately to security threats and protect the network. Inline IPS configuration forces all traffic to be directed to the AIP-SSM. The ASA does not forward any traffic out to the network without the AIP-SSM first inspecting it.

Figure 4-25 shows the traffic flow when the Cisco ASA is configured in inline IPS mode.

Figure 4-25 Inline IPS Traffic Flow

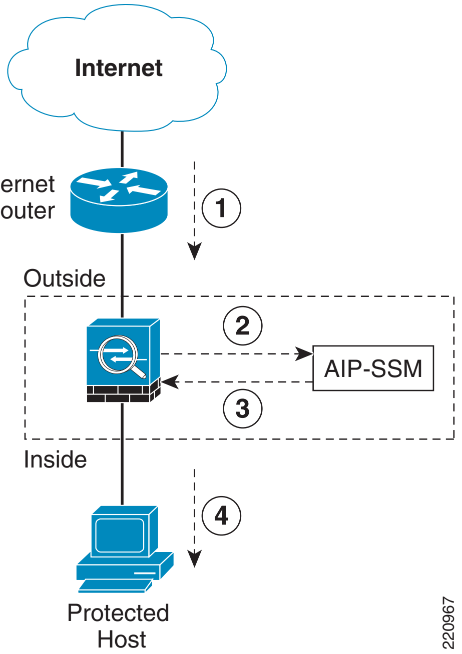

Figure 4-25 shows the following sequence of events:

Step 1 The Cisco ASA receives an IP packet from the Internet.

Step 2 Because the Cisco ASA is configured in inline IPS mode, it forwards the packet to the AIP-SSM for analysis.

Step 3 The AIP-SSM analyzes the packet and, if it determines that the packet is not malicious, forwards the packet back to the Cisco ASA.

Step 4 The Cisco ASA forwards the packet to its final destination (the protected host).

Note Inline IPS mode is the most secure configuration because every packet is inspected by the AIM-SSM. However, this may affect the overall throughput. The impact depends on the type of attack, signatures enabled on the system, and the amount of traffic passing through the application.

When the Cisco ASA is set up to use the AIP-SSM in promiscuous mode, the ASA sends a duplicate stream of traffic to the AIP-SSM. This mode has less impact on the overall throughput. Promiscuous mode is considered to be less secure than inline mode because the IPS module can only block traffic by forcing the ASA to shun the malicious traffic or send a TCP-RST (reset) message to terminate a TCP connection.

Note Promiscuous mode has less impact on performance because the AIP-SSM is not in the traffic path. A copy of the packet is sent to the AIM-SSM. If a packet is dropped, there is no effect on the ASA.

Figure 4-26 shows an example of how traffic flows when the AIP-SSM is configured in promiscuous mode.

Figure 4-26 Promiscuous Mode Traffic Flow

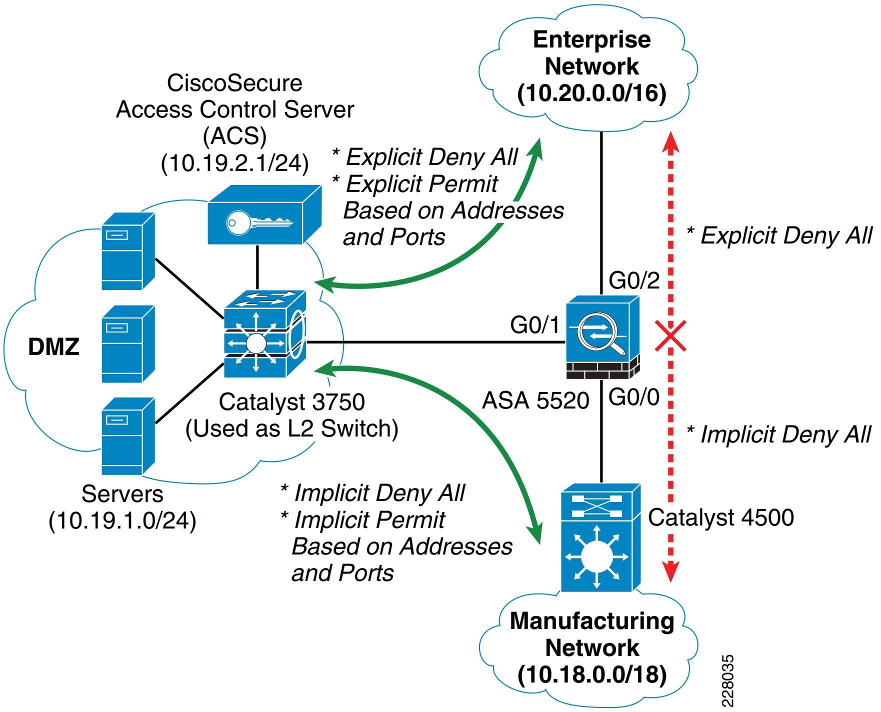
Enterprise Network (10.20.0.0/16)
CiscoSecure Access Control Server (ACS) (10.19.2.1/24)
* Explicit Deny All
* Explicit Permit Based on Addresses and Ports
* Explicit Deny All
G0/2
G0/1
DMZ
Catalyst 3750 (Used as L2 Switch)
ASA 5520
G0/0
* Implicit Deny All
Servers (10.19.1.0/24)
* Implicit Deny All
* Implicit Permit Based on Addresses and Ports
Catalyst 4500
Manufacturing Network (10.18.0.0/18)
228035

Figure 4-26 shows the following sequence of events:

Step 1 The Cisco ASA receives an IP packet from the Internet.

Step 2 Because the Cisco ASA is configured in promiscuous mode, the AIP-SSM silently snoops the packet.

Step 3 The ASA forwards the packet to its final destination (the protected host) if the packet conforms to security policies; that is, if it does not match any of the configured signatures.

Note If the ASA firewall policies deny any inbound packet at the interface, the packet is not inspected by the AIM-SSM. This applies to both inline and promiscuous IPS modes.