



Ethernet-to-the-Factory 1.1 Design and Implementation Guide

Cisco Validated Design I

July, 2007

This design and implementation guide represents a collaborative development effort from Cisco Systems and Rockwell Automation. It is built on, and adds to, design guidelines from the Cisco Ethernet-to-the-Factory solution and the Rockwell Automation Integrated Architecture™.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Rockwell Automation Global Headquarters
1201 South Second Street
Milwaukee, WI 53204-2496
Tel: 414 382-2000
www.rockwellautomation.com

Customer Order Number:
Text Part Number: OL-14268-01

Cisco Validated Design

The Cisco Validated Design Program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/validateddesigns.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCVP, the Cisco Logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

The following are trademarks or registered trademarks of Rockwell Automation, Inc: Integrated Architecture, RSLogix, FactoryTalk, PowerFlex, ControlLogix, and PanelView Plus. Trademarks not belonging to Rockwell Automation are the property of their respective companies. (0612R)



CONTENTS

Preface xi

Document Organization xi

CHAPTER 1

Ethernet-to-the-Factory Solution Overview 1-1

Executive Summary 1-1

Introduction 1-2

Cisco EttF 1.1 Solution—Description and Justification 1-2

Target Customer 1-3

Plant Managers and Control Engineers 1-5

Manufacturing IT 1-6

Applications and Services Supported by the Cisco EttF Solution 1-6

Cisco EttF Solution Benefits 1-7

Cisco EttF Solution Features 1-8

Compatibility 1-8

Real-Time Communication, Determinism, and Performance 1-9

Logical Segmentation 1-10

Physicality and Topology 1-10

Availability 1-14

Security 1-15

Manageability 1-16

Scalability 1-17

Scope of the Cisco EttF Solution 1-17

Industrial Automation and Control Background 1-18

History of Industrial Automation and Control Networks 1-18

Industrial Automation and Control System Components 1-19

Physical Layer 1-19

Networking Equipment 1-19

Industrial Automation and Control Devices 1-21

Industrial Computing 1-22

Industrial Automation and Control System Communication Protocols 1-22

Communication Model 1-22

Industrial Automation and Control Protocol Overview 1-23

Common Industrial Protocol Overview 1-26

Solution Architecture 2-1

- Overview 2-1
 - Industrial Automation and Control Reference Model 2-1
 - Safety Zone 2-2
 - Cell/Area Zone 2-2
 - Manufacturing Zone 2-4
 - Enterprise Zone 2-5
 - Ethernet-to-the-Factory Framework 2-6
 - Campus Network Reference Model 2-9
- Cell/Area Zone 2-10
 - Overview 2-10
 - Components 2-11
 - Unmanaged versus Managed Switches 2-13
 - Physicality and Environmental Considerations 2-13
 - Real-time Communications 2-13
 - Availability 2-14
 - Flexibility 2-14
 - Manageability 2-14
 - Security 2-15
 - Component Summary 2-15
 - Traffic Flows 2-17
 - Topology Options Overview 2-19
 - Cell/Area Network—Trunk-Drop (Daisy Chain) Topology 2-20
 - Cell/Area Network—Ring Topology 2-21
 - Cell/Area Network—Star Topology 2-22
 - Cell/Area Topology Comparison 2-23
 - Network Design Overview 2-25
 - Logical Segmentation 2-25
 - Availability 2-28
 - Multicast Management 2-29
 - Quality of Service 2-30
 - Security 2-31
- Manufacturing Zone 2-31
 - Overview 2-31
 - Components 2-33
 - Cost 2-35
 - Physicality and Environmental Considerations 2-35
 - Performance and Real-time Communications 2-35
 - Availability 2-35

Manageability	2-36
Security	2-36
Component Summary	2-37
Traffic Flows—Manufacturing Zone	2-41
Topology Options Overview	2-44
Small Manufacturing Zone Topology	2-44
Medium Manufacturing Zone Topology	2-45
Manufacturing Zone Topology Summary	2-47
Network Design Overview	2-47
Logical Segmentation	2-47
Availability	2-48
Routing	2-51
Manageability	2-51
Demilitarized Zone	2-51
Components	2-52
Topology Options Overview	2-53
Network Design Overview	2-54
Software Versions	2-55

CHAPTER 3**Basic Network Design 3-1**

Overview	3-1
Assumptions	3-1
IP Addressing	3-1
Static IP Addressing	3-1
Using Dynamic Host Configuration Protocol and DHCP Option 82	3-2
IP Addressing General Best Practices	3-3
IP Address Management	3-3
Address Space Planning	3-3
Hierarchical Addressing	3-3
Centralized IP Addressing Inventory	3-4
Routing Protocols	3-5
Selection of a Routing Protocol	3-5
Distance Vector versus Link-State Routing Protocols	3-5
Classless versus Classful Routing Protocols	3-5
Convergence	3-5
Routing Metric	3-6
Scalability	3-6
Static or Dynamic Routing	3-7
Server Farm	3-7

- Types of Servers 3-7
- Server Farm Access Layer 3-9
 - Access Layer Considerations 3-9
 - Layer 2 Access Model 3-9
 - Spanning VLANs across Access Layer switches 3-9
 - Layer 2 Adjacency Requirements 3-10
 - NIC Teaming 3-10

CHAPTER 4

Implementation of the Cell/Area Zone 4-1

- Cell/Area Zone Network Device Provisioning 4-2
- Virtual LAN Segmentation 4-3
 - VLAN Overview 4-3
 - VLAN Details 4-4
 - VLANs In the Cell/Area Zone 4-5
 - VLAN Highlights of Ring Topology 4-5
 - VLAN Recommendations 4-6
 - VLAN Benefits for EttF 4-6
- Spanning Tree Protocol Design 4-7
 - STP Overview 4-7
 - STP Configurable Parameters 4-7
 - More on STP Redundancy 4-8
 - STP Topology for EttF 4-10
 - STP Considerations for the Ring 4-11
 - Control Device Placement 4-11
 - Trunk Ports or Access Ports 4-11
 - Sample Trunk Configuration 4-12
 - VLAN 1 Minimization 4-12
 - Location of the Root Bridge 4-13
 - PortFast on Access Ports 4-13
 - PortFast Operational Overview 4-13
 - STP Limitations 4-14
 - RSTP+ Convergence Process 4-14
- Multicast Design 4-15
 - EtherNet/IP Multicast Traffic Patterns 4-15
 - IGMP Snooping 4-17
 - IGMP Querier and EtherNET/IP Traffic 4-19
 - IGMP Configurations 4-20
- Switch Troubleshooting Toolkit 4-21

CHAPTER 5**Implementation of Security 5-1**

Overview	5-1
Network Device Hardening	5-3
Router	5-4
Basic Hardening Settings	5-4
Authentication Settings	5-5
Management Access	5-8
Layer 2 Security Design	5-11
Precautions for the Use of VLAN 1	5-12
Trust Level of Switch Ports	5-12
Spanning Tree Protocol Security	5-13
MAC Flooding Attack	5-16
VLAN Hopping	5-16
ARP Spoofing Attack	5-17
DHCP Attacks	5-18
Security Design for the Manufacturing Zone	5-19
Security Design for the Catalyst 3750 Series Switch That Aggregates Cell/Area Zone Networks	5-19
Security Design for the Catalyst 4500 Series Switch for the Core of the Control Network	5-19
Security Design for the Catalyst 3750 Series Switch in the Server Farm	5-20
Security Protection for Servers	5-21
Security Design for the Demilitarized Zone	5-21
Security Levels on the Cisco ASA Interfaces	5-22
Configuration Example	5-22
Stateful Packet Filtering	5-23
Configuration Example	5-26
Authenticating Firewall Sessions for User Access to Servers in the DMZ	5-28
Configuration Example	5-29
Integrating the ASA 5500 Appliance with the Adaptive Inspection Prevention Security Services Module	5-30
Access to the AIP-SSM Module	5-30
Inline Versus Promiscuous Mode	5-30
Endpoint Protection with Cisco Security Agent	5-33
Security Monitoring, Analysis, and Mitigation with CS-MARS	5-33

CHAPTER 6**Implementation of High Availability 6-1**

Benefits of an HA Design	6-3
Best Practices and HA Modeling	6-4
HA Design in the Cell/Area Zone	6-5

HA Design and Implementation in the Manufacturing Zone	6-6
First Hop Redundancy	6-8
NSF/SSO	6-8
Summary of Features in the Manufacturing Zone	6-9
HA Design and Implementation in the DMZ	6-10
Cisco ASA Redundancy Design	6-10
Primary/Secondary Status and Active/Standby Status	6-11
Determination of the Active Unit	6-11
Failover Triggers	6-11
Configuration Synchronization	6-11
Passage of State Information to the Standby Unit	6-12
Active/Standby Failover Configuration	6-13
Selecting the Failover Link	6-13
Assigning Failover IP Addresses	6-13
Setting Failover Key (Optional)	6-14
Designating the Primary Cisco ASA	6-14
Enabling Stateful Failover (Optional)	6-14
Enabling Failover Globally	6-15
Configuring Failover on the Secondary Cisco ASA	6-15

CHAPTER 7

Implementation of Network Management 7-1

Cisco Network Assistant	7-1
CNA Security Considerations	7-2
Cisco Adaptive Security Device Manager	7-2

APPENDIX A

Characterization of the EttF Cell/Area Zone Design A-1

STP Testing	A-1
STP Test Methodology	A-1
STP Test Topology	A-2
STP Test Scenarios	A-3
Test Suite 1—Bidirectional Traffic (Tx1 <-> Tx2)	A-3
Test Suite 2—Bidirectional Traffic (Tx3 <-> Tx4)	A-6
Test Tools	A-8
STP Test Results	A-8
Suite 1 Test Cases	A-8
Suite 2 Test Cases	A-10
Sample Trend Line for Link Failure Between Adjacent Switches	A-12
Sample Trend Line for Link Failure To Root Bridge	A-13
16-Switch Ring—STP Testing	A-13

Test Suite 1—Bidirectional Traffic from (Tx1 <-> Tx2) **A-14**

Test Suite 2—Bidirectional Traffic (Tx3 <-> Tx4) **A-14**

Redundant Star Topology—STP Testing **A-15**

Latency/Jitter Testing **A-17**

IGMP Testing **A-18**

IGMP Snooping Test Methodology **A-18**

IGMP Snooping Test Topology **A-18**

IGMP Snooping Test Results **A-19**

APPENDIX B

Configuration of the EttF Cell/Area Zone **B-1**

Layer 2 Configuration **B-1**

Layer 3 Configuration **B-4**

APPENDIX C

Configuration of the EttF Demilitarized Zone **C-1**

Security Configuration **C-1**

ASA Configuration **C-1**

IPS Configuration **C-3**

APPENDIX D

EttF High Availability Testing **D-1**

HA Test Methodology **D-1**

HA Test Topology **D-1**

HA Test Scenarios **D-3**

Test Suite 1—HA in the Cell/Area Zone (Tx1 † Tx2) **D-3**

Test Suite 2—HA in the Manufacturing Zone (Tx1 † Tx2) **D-4**

Test Suite 3—HA in the DMZ (Tx1 † Tx3) **D-6**

Test Tools **D-7**



Preface

This document describes the Cisco Ethernet-to-the-Factory (EttF) Architecture, which offers value inside industrial operations when part of a larger integrated, industrial automation architecture.

Document Organization

This document contains the following chapters and appendices.

Chapter or Appendix	Description
Chapter 1, “Ethernet-to-the-Factory Solution Overview.”	Provides an overview of the Cisco Ethernet-to-the-Factory solution.
Chapter 2, “Solution Architecture.”	Provides an overview of the EttF solution architecture, as a means to describe the various systems, components, and their relation to each other to give context to the networking function and technical requirements.
Chapter 3, “Basic Network Design.”	Provides guidelines and best practices for IP addressing, and the selection of routing protocols based on the manufacturing zone topology and server farm access layer design.
Chapter 4, “Implementation of the Cell/Area Zone .”	Outlines recommendations, best practices, configurations, and caveats for implementing a cell/area zone architecture in an EttF environment.
Chapter 5, “Implementation of Security.”	Describes the security components of the EttF solution that address the major security concerns of defending against threat, establishing trust boundaries and verifying identity, and securing business communications.
Chapter 6, “Implementation of High Availability.”	Describes design considerations and best practices for high availability in the cell/area zone, manufacturing zone, and the DMZ, as well as testing results.
Chapter 7, “Implementation of Network Management.”	Describes the use of network management tools such as Cisco Network Assistant and Cisco Adaptive Security Device Manager.

Chapter or Appendix	Description
Appendix A “Characterization of the EttF Cell/Area Zone Design.”	Outlines the validation methodology and the corresponding results of the testing.
Appendix B “Configuration of the EttF Cell/Area Zone.”	Provides sample configurations for the cell/area zone.
Appendix C “Configuration of the EttF Demilitarized Zone.”	Provides sample configurations for the DMZ.
Appendix D “EttF High Availability Testing.”	Outlines the validation methodology and the corresponding results of the high availability testing.



CHAPTER 1

Ethernet-to-the-Factory Solution Overview

Executive Summary

This design and implementation guide represents a collaborative development effort from Cisco Systems and Rockwell Automation. It is built on, and adds to, design guidelines from the Cisco Ethernet-to-the-Factory solution and the Rockwell Automation Integrated Architecture™.

Faced with internal pressures to cut costs and external demands for better products and services, manufacturers are realizing the benefits of a converged network, such as the following:

- Greater visibility
- Better data integration
- Shorter lead times
- Increased turnaround
- Reduced costs
- Simplified management

The key targets are industrial automation and control systems, which benefit greatly from the transition to modern networking technologies from the factory-optimized networks typically in use today. New services and streamlined efficiency result when the information contained within these automation and control systems is available and shared throughout the larger enterprise. Access to existing production information is presently gated by disparate, proprietary, and closed systems. Manufacturers and their industrial suppliers are discovering that standard communication and uniform networking of industrial systems is the key to optimized services, greater visibility, and lower total cost of ownership (TCO). They are starting to embrace standard information technology, particularly Ethernet and IP, for industrial automation and control environments.

Although most manufacturers recognize that Ethernet and the IP protocol suite will be the de-facto networking standard in manufacturing environments in the near future, only a few have fully adopted standards-based Ethernet network architectures for industrial automation. Much of this resistance can be attributed to the aversion to disrupting existing systems, the accounting realities of fully-depreciated assets, and the general ebb and flow of production investment cycles. Resistance to migration also comes from the market being serviced by small niche vendors with narrowly-designed products or limited support capabilities. As bigger players start to enter the market and create an industry-wide industrial networking standards organization, the market is poised to explode.

Cisco and Rockwell Automation believe standard networking technology offers value inside industrial operations when the technology is part of larger integrated, industrial automation architectures. Cisco calls this the Ethernet-to-the-Factory (EttF) Architecture. Rockwell Automation offers their Integrated Architecture solution.

The purpose of this architecture is to accelerate the convergence of standard networking technologies with the industrial automation and control environment. This solution architecture and relevant design and implementation guidelines will give customers, partners, and the marketplace the confidence and background necessary to employ EttF. This solution architecture must be tailored to support automation and control systems. By adopting the solution architecture, the manufacturing process will have to operate at higher levels of performance, efficiency, and uptime as under the previous solutions. At the same time, it must also safely and securely integrate these systems into the broader manufacturing environment; only at this point will all the benefits be available to the manufacturing enterprise.

Introduction

Cisco EttF 1.1 Solution—Description and Justification

The industrial manufacturing environment of today is very similar to the IBM legacy mainframe environments of the mid 1990s. Although these legacy industrial systems are functional, they are costly to maintain, difficult to connect, and slow to evolve. With their factory floor-optimized protocols, specific operating requirements, and separate staffs, manufacturers are also struggling to evolve. Whether their industrial automation and control systems are discrete, process, batch, or hybrid, manufacturers need their systems to interact in real-time with the other enterprise applications, supply chain partners, and end customers. To accomplish this, manufacturers are bringing their industrial automation systems online. When doing this, manufacturers encounter a number of challenges, such as the following:

- **Production reliability**—As manufacturing operations become globally integrated, manufacturers are challenged to provide consistent access to data while making the manufacturing environment programmable and flexible. Security, availability, and asset use are critically important to manufacturing companies because industrial automation and control equipment is mission-critical, and efficiency is important to remain competitive.
- **Cost**—Legacy industrial automation and control systems, although often fully depreciated in existing manufacturing environments, can be difficult to bring online and can require significant investment.
- **Product design integration**—Data silos and closed systems hinder the ability to reduce time to market for new products.
- **Service integration**—In an effort to provide differentiated service, manufacturers are struggling to create closed loop systems to capture and incorporate data from the field.
- **Data interaction and management**—Incorporating real-time factory productivity and operational data into manufacturing execution systems (MES), customer relationship management (CRM), supply chain management (SCM), and other enterprise resource planning (ERP) systems is an increasingly complex data translation exercise.
- **Partner connections**—With an aging workforce and production complexity, manufacturers are trying to find ways to leverage relationships with industrial automation and control vendors.

These challenges are pushing manufacturers to adopt standard Ethernet and IP technologies throughout the manufacturing environment. By moving to standard technologies, manufacturers can:

- **Realize significant cost savings**—Standard Ethernet and IP technology has greater market penetration and thus is more likely than proprietary technology to have shorter and more consistent cost reduction life cycles.

- Provide better maintenance—As access to skilled production staff becomes difficult, legacy industrial automation and control technology is becoming more complex to maintain than standard Ethernet and IP networking technology.
- Enhance their flexibility—Standard Ethernet and IP technology allows for rapid production gains, new functionality, and evolving capabilities in the manufacturing environment and beyond.
- Increase efficiency—Standard Ethernet and IP technology eases integration with business systems.

Manufacturing organizations and their production operations want to use the newer standard networking technologies in industrial automation and control networks, but there has been little guidance from industrial networking or automation suppliers to date. In addition, although the automation and control industry as a whole has embraced standard networking over legacy, proprietary networking, some industrial automation and control vendors continue to suggest that standard Ethernet and IP technology is not good enough for manufacturing environments. The principle argument has been that deterministic and time-sensitive manufacturing environments require more than what Ethernet and IP technologies can deliver. Others question the inherent determinism, reliability, and resiliency of Ethernet and IP technologies. Some have even asserted that standard networking technology in production environments makes manufacturers more susceptible to security risks. Although there is some basis for these concerns, there is little substantive data to support these claims. Modern, full-duplex, switched Ethernet networks offer real-time performance, including latency, jitter, and (non) packet loss capabilities, that equals or surpasses the older fieldbus networks they replace. In addition, these modern networks have mature and tested technologies to safely secure the network and the systems they interconnect beyond what is available for the older fieldbus networks.

EttF is an architecture that provides standards-based network services to the applications, devices, and equipment found in modern industrial automation and control systems, and integrates them into the wider enterprise network. The Cisco EttF 1.1 solution gives design and implementation guidance to achieve the real-time communication requirements needed for determinism as well as the reliability and resiliency required by the industrial and automation control systems. By bringing the Cisco EttF Architecture to market, Cisco can enable manufacturing customers to meet all the challenges of a fully-integrated industrial automation system. The Cisco EttF Architecture also enhances the status of Cisco as a trusted business partner, not only for manufacturing customers but also for industrial automation partners.

Target Customer

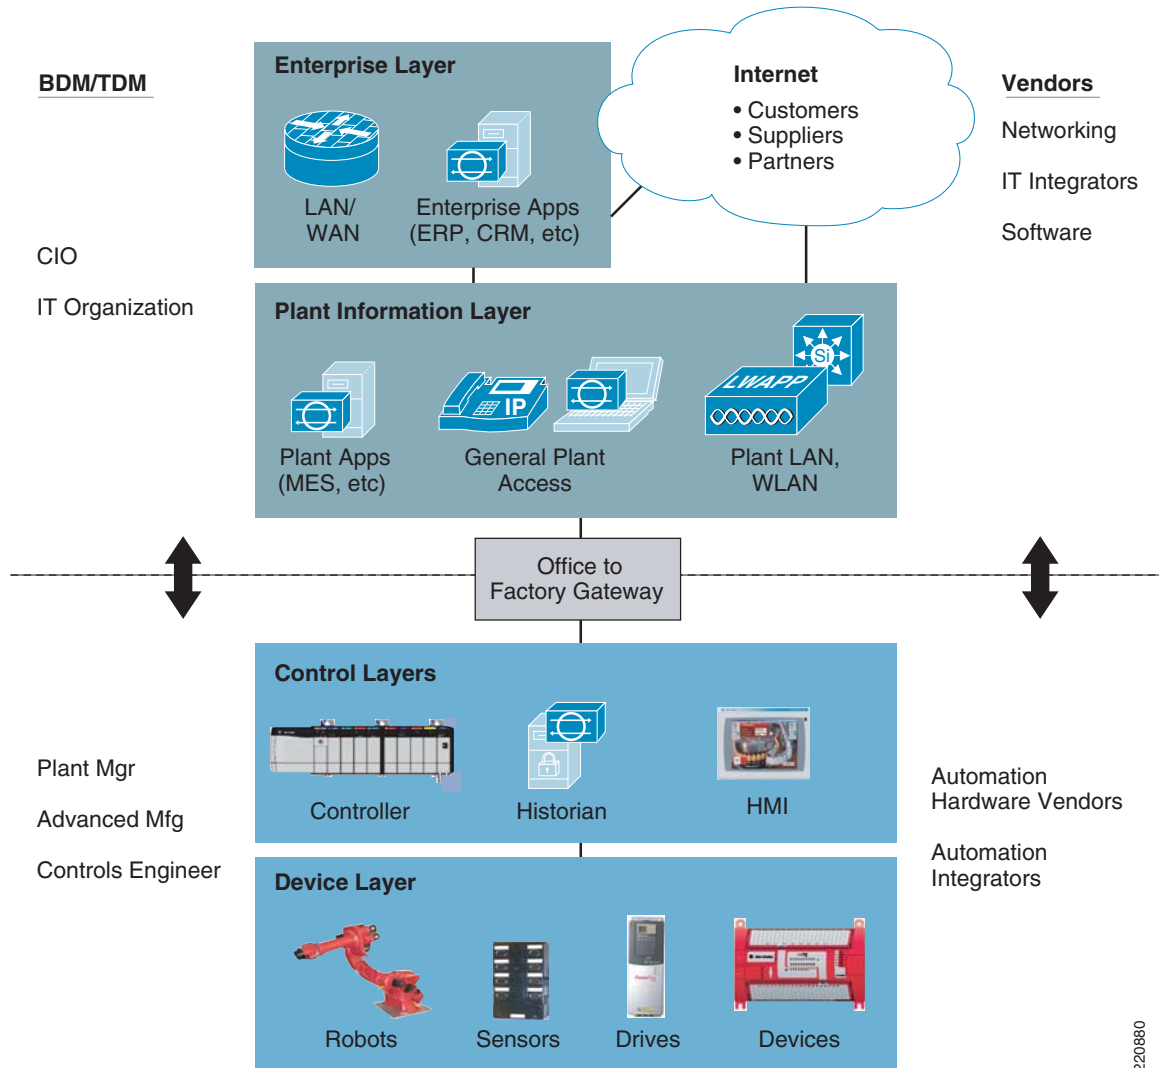
The Cisco EttF solution is targeted at manufacturing customers seeking to integrate or upgrade their industrial automation and control networks to standard networking technologies. These customers want to do the following:

- Lower the TCO of their current industrial automation and control network approach
- Integrate the industrial automation and control systems with the wider enterprise
- Take advantage of the networking innovations provided by using standards-based technologies

Decisions impacting industrial automation and control networks of the factory floor are typically driven by plant managers and control engineers, rather than the IT department. Additionally, they rely on a completely different vendor and support supply chain for their industrial automation and control systems than typically found in the IT department. This is driven by the different requirements of a factory floor. That being said, the IT departments of manufacturing customers are increasingly engaging with plant managers and control engineers to leverage the knowledge and expertise in standard networking technologies.

The Cisco EttF solution recognizes and targets the IT department *and* the plant managers and control engineers. Each camp has different perspectives and requirements for a successful EttF implementation (see Figure 1-1).

Figure 1-1 Business/Technical Decision Makers—IT versus Automation and Control



For the IT department, it is critical to understand the various factory floor requirements and operating environment, and to implement an appropriate solution. For the plant managers and control engineers, a deeper knowledge of the capabilities and functioning of standard networking technologies is required. The Cisco EttF solution includes a large number of references to “basic” networking concepts to recognize the need to raise the level of knowledge and expertise of business and technical decision makers.

To increase its value and impact, the Cisco EttF solution will be brought to market with a key industrial automation and control systems partner, Rockwell Automation. This will allow the Cisco EttF solution to benefit not only from the deep expertise in industrial automation and control systems found in this partner, but also to more effectively target the various business and technical decision makers.

220880

To summarize, the industrial automation and control systems toward which the Cisco EttF solution is targeted see various business and technical decision makers introduced into the decision process. These decision makers often have differing business objectives and concerns that must be recognized. These decision makers rely on different vendors and integrators for solutions and their implementation. In addition, the typical decision makers (IT) stay involved, but may need awareness levels raised concerning the differences and challenges posed by the manufacturing environment.

Plant Managers and Control Engineers

As mentioned, plant managers and control engineers are key decision makers for the Cisco EttF solution.

Plant managers are business decision makers for this solution and are responsible for achieving production targets by ensuring plant reliability, uptime, and energy efficiency. Their performance is typically measured by plant profitability, throughput, quality, and return on assets. Technology decisions are made related to reliability, risk-free operation, environment fit, and company-wide standards. Plant managers usually depend on vendors for support based on track record and industry knowledge.

Control engineers are technical decision makers for this solution, and are responsible for the design, implementation, and operations of the industrial automation and control systems that operate the production facility. They are responsible for the automation equipment that supports the basic manufacturing process. They have a direct share of the responsibility of the quality and consistency of the end product, and often report to the plant manager.

For both these decision makers, the key business drivers include the following:

- **Reliability**—The solution must support the operational availability of the production facility.
- **Cost**—Capital comes at a premium, and additional costs (or costlier components) must add clear value that is understood by the plant manager.
- **Ease of integration**—Not just with enterprise applications, but ease of integrating remote or vendor expertise in a secure manner.
- **Flexibility**—The ability to rely on common off-the-shelf (COTS) equipment, provided by a number of vendors and supported from a common expertise (often found in the IT department).

Key concerns include the following:

- **Performance**—Ability of the network infrastructure to meet the real-time communications requirements of the industrial automation and control systems
- **Availability**—Both the ability to limit the impact on operations of upgrading or maintaining the Cisco EttF solution, and the reliability of the supported base network infrastructure features to handle outages with minimal impact
- **Manageability**—Ease of configuring, maintaining, and fixing the Cisco EttF solution
- **Compatibility**—How the network infrastructure supports various types of industrial communications (see [Industrial Automation and Control System Communication Protocols, page 1-22](#)) and the devices, controllers, human-machine interfaces (HMIs), and applications already in use

Both plant managers and control engineers typically rely on vendors with strong knowledge and track records in industrial automation and control. These vendors have varying degrees of capability and knowledge in deploying standards-based networking technologies and the relevant technical issues. By going to market with this solution jointly with a key vendor, the objective is to bring the relevant partners, channels, and integrators up to speed on the availability and capabilities of industrial Ethernet in general and specifically the Cisco EttF solution.

Manufacturing IT

Although IT managers are typically the business and technical decision makers for network infrastructure, they have not typically been involved with network infrastructure for industrial automation and control systems for a wide variety of reasons. They are often seen by the plant managers and control engineers as an obstacle to be avoided, rather than a partner to be relied on for skills, expertise, and services. They are usually making decisions to focus on standardized solutions, to re-use whenever possible, and to reduce cost. There is often a cultural gap between IT and the manufacturing world. However, because IT managers often have the deepest knowledge and expertise in standard networking technologies within the enterprise, their involvement is often required for a truly successful implementation of industrial Ethernet. To help overcome the cultural gap, the Cisco EttF solution does the following:

- Raises IT awareness of the particular challenges and requirements for industrial automation and control systems
- Outlines a solution and relevant design and implementation guidance that allows both to focus on a mutually-acceptable solution
- Pulls IT into the environment to deliver expertise and services based on their strength in standards-based networking technologies

Applications and Services Supported by the Cisco EttF Solution

The Cisco EttF solution primarily supports industrial automation and control systems and their integration into the overall enterprise network. Industrial automation and control systems consist of the following:

- Automation devices, such as robots, sensors, actuator, and drives
- Human-machine interfaces (HMIs) that provide visual status reports and control of the automated manufacturing process
- Controllers such as programmable automation controllers (PACs) and the distributed control system (DCS)
- Higher level plant systems, including the manufacturing execution system (MES) and historians

This version of the architecture focuses on the above items that support Ethernet/IP, which is driven by the Common Industrial Protocol (CIP) (see [Industrial Automation and Control System Communication Protocols, page 1-22](#)) and in particular are tested with Rockwell Automation devices, controllers, and applications.

The key networking services that are supported in this version of the EttF architecture include the following:

- Local area networking (typically defined as OSI Layers 1 and 2) to all the above items, including topology, port configuration, subnet and VLAN configuration, network protocols for spanning tree, and quality of service (QoS)
- Routing (typically defined as Layer 3) for all the above items, as well as to other areas of an enterprise network
- Design and implementation recommendations for network technical considerations such as topology, resiliency, and redundancy (including Spanning Tree Protocol), and handling of multicast traffic (including Internet Group Management Protocol configuration)
- IP address allocation, assigning, and related services (for example, DHCP, BootP, and DNS)
- Basic network management

- Network security for the industrial automation and control systems including demilitarized zone (DMZ), firewall, intrusion protection, endpoint security, and security monitoring, analysis, and response

These will be applied to small (up to 50 Ethernet nodes) to medium (up to 200 Ethernet nodes) environments.

Cisco EttF Solution Benefits

The value proposition for the Cisco EttF solution is as follows:

- Enables and simplifies integration of industrial automation and control systems with enterprise networks to improve the flow and integration of production information into business systems.
- Enables remote access for production engineers, partners, and industrial automation and control equipment vendors for diagnostics and maintenance. Increases efficiency and response time and enables industrial automation and control vendors to provide services to customers that may have limited subject matter expert (SME) resources.
- Reduces operating and capital costs by using open standards to eliminate the need to support multiple protocols in industrial automation and control networks and to provide manufacturing companies more options when purchasing automation equipment.
- Integrates more quickly advances in networking technology that come from working with standards-based technologies (for example, voice, video, and security).

Integrating advanced technologies and working with leading industrial automation and control vendors such as Rockwell Automation allows Cisco to have a unique value proposition relative to the rest of the industry by providing benefits beyond those associated with integration and use of open standards, including the following:

- Combining two areas of expertise: the networking expertise of Cisco with the industrial automation and control expertise of Rockwell Automation for the benefit of the customer
- Providing integrated security specifically configured for industrial automation and control networks to protect vital manufacturing assets, limit access to production equipment, and help address issues such as patch management
- Providing a foundation for deploying additional advanced technologies such as voice, video, and wireless on the converged network at the control level as the technology matures and the business requires
- Simplifying deployment and helping to bridge the gap that often exists between IT and industrial automation and control networks by integrating and validating architectures with leading partners in the industrial automation and control market that ensure compliance with relevant industry standards.

The above capabilities depend on the deployment of technologies based on standard Ethernet and IP, and help demonstrate the value of open standards to differentiate Cisco and its partners from other “standards-based” Ethernet and non-standard solutions on the market.

Cisco EttF Solution Features

Industrial automation and control network environments have evolved over the years, driven by a number of key design features. These features are not specific to industrial Ethernet, but to networking for industrial automation and control systems in general. In the move towards industrial Ethernet, many of these design features still apply, although the importance sometimes shifts. For example, with Ethernet and IP-based industrial networks, security is a pressing issue, particularly if there are no access restrictions between industrial automation and control systems and the larger business system. This section defines the following six key features that the industry expects as best practices:

- Compatibility
- Real-time communication, determinism, and performance
- Logical segmentation
- Physicality and topology
- Availability
- Security
- Manageability
- Scalability

This document provides details on why and how to take advantage of these benefits. The industry, and especially manufacturing participants such as plant managers, control engineers, and their partners and vendors, are looking for simple guidelines and recommendations. Each chapter in this document highlights key recommendations and steps to follow when designing and implementing an industrial Ethernet solution.

Compatibility

By definition, industrial Ethernet (IE) protocols should operate on standard networking technologies and infrastructure if they so specify. However, standard networking technologies have a wide range of service and configuration options that need to be considered to effectively support the industrial automation and control application. As well, various IE protocols specify various networking features that then must be available to operate at required performance levels.

EttF must show compatibility with the IE protocols and communication models of the applications that run on it. This typically means supporting the types of traffic they generate, such as TCP and UDP (multicast and unicast), as well as any features and functions they expect of the network, such as quality of service (QoS). A large number of types of traffic may exist in an industrial Ethernet network, including automation and control protocols such as CIP, Modbus/TCP or OPC, as well as common protocols such as web browsing (HTTP), file transfer (FTP), and many others. The Cisco EttF solution outlines how to design and implement compatible network architectures.

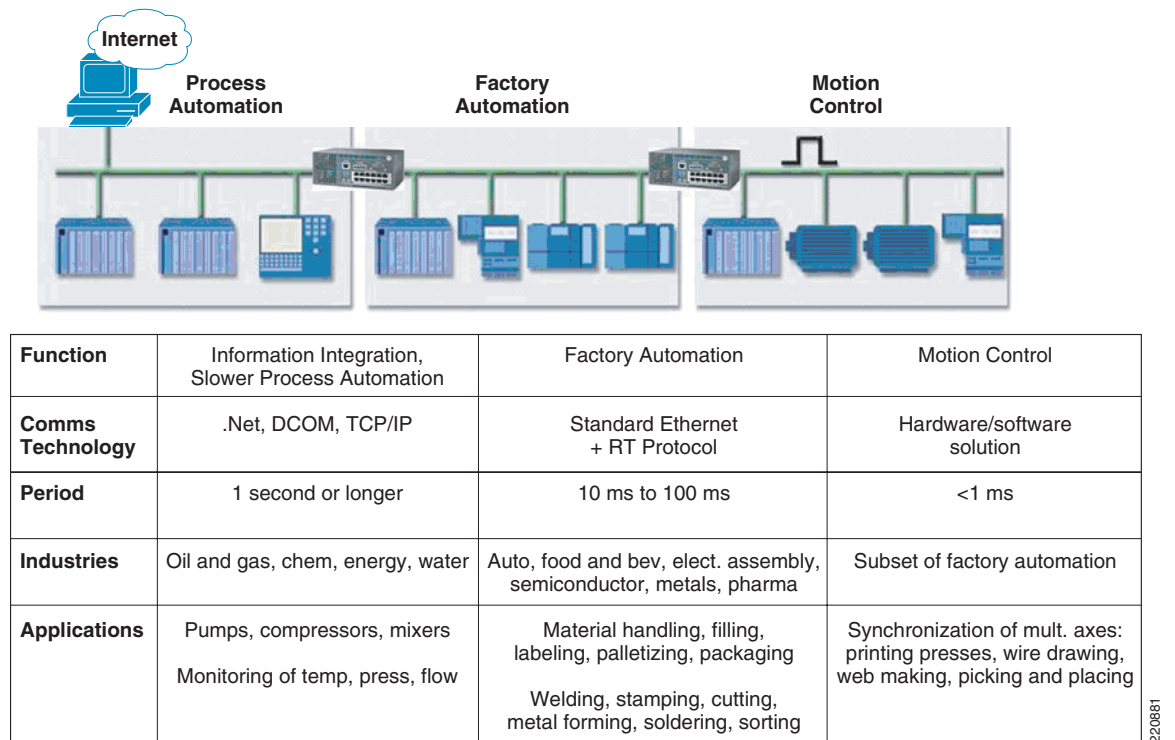
[Industrial Automation and Control System Communication Protocols, page 1-22](#) lists the relevant general industrial protocols and the corresponding industrial Ethernet versions. This solution architecture focuses on the Common Industrial Protocol (CIP). Other network protocols are considered (see the sub-sections on traffic flows in [Cell/Area Zone, page 2-10](#) and [Manufacturing Zone, page 2-31](#)).

Real-Time Communication, Determinism, and Performance

Industrial automation and control systems differ from their IT counterparts in their need to support real-time communications, which means communicating messages with minimal latency (time delay between message sent and message received) and jitter (the variance of the latency). Real-time communications help the industrial automation and control systems become more deterministic. Although the network plays a role in determinism, a number of other factors, such as end-device latency and response time, are also involved. Therefore, the capabilities of standards-based networks to support challenging real-time communications are described in this document.

Industrial automation and control networks have various real-time communications requirements based on the type of application, as shown in Figure 1-2.

Figure 1-2 Real-Time Applications (Source: ARC Research, 2006)



The Cisco EttF solution provides design and implementation guidance to help customers achieve the real-time communications requirements of their industrial automation and control systems.

Key considerations in achieving real-time communications include the following:

- Number of switches and routers and amount of traffic in the Layer 2 network, which affects latency and jitter.
- Ratio of LAN ports to uplink ports based on traffic loads and patterns. Typically, this means using 10/100 Mbps to devices and 10/100/1000 Mbps for uplinks.
- Use of Internet Group Management Protocol (IGMP) to manage the efficient delivery of multicast traffic.
- Use of QoS parameters to meet the real-time requirements of various traffic flows.
- Reduction of excessive Rapid Spanning Tree Protocol (RSTP) traffic to avoid bandwidth flooding of the Layer 2 domain. To accomplish this, use PortFast, Backbone Fast, and so on.

220861

Logical Segmentation

Logical segmentation integrates logically (or physically) isolated networks of the production facility with the enterprise and external networks to safely and securely share data, services, and access from the industrial automation and control systems. Logical segmentation is critical for industrial Ethernet because it helps ensure that availability, determinism, performance, manageability, and security requirements are maintained. Logical segmentation means allowing required communication between devices while preventing extraneous traffic from interfering with critical communications between devices on the industrial automation and control network. Logical segmentation is required because industrial Ethernet network architectures may generate traffic that is not readily compatible with general enterprise traffic, and vice versa. For example, multicast traffic in a manufacturing environment may use multicast addresses that overlap with those in the enterprise zone, or traffic in either zone may set QoS markings that create issues in the other zone. The fundamental tenet of logical segmentation is that the manufacturing traffic be separate from the enterprise traffic.

Insulation can be achieved via numerous mechanisms. The Cisco EttF solution provides design and implementation guidelines on the key considerations and mechanisms that can be applied, including the following:

- Using an additional physical or logical DMZ to segregate the manufacturing control network from the corporate IT network, especially to do the following:
 - Halt the mixing of incompatible traffic
 - Create clear administrative boundaries to manage organizational control and configuration differences between the manufacturing and enterprise zones
 - Safely and securely share data and services between the zones
- Using hierarchically-tiered switches inside the industrial automation and control network to further segment manufacturing functional areas. (See the following subsections for related parameters in this situation.)
- Limiting the number of devices per Layer 2 domain in industrial automation and control networks to devices that must talk to each other, to maintain more control over performance characteristics and easily develop a more granular security model.
- Using virtual LANs (VLANs) to create logical structures around Layer 2 domains.
- Using routers/Layer 3 switches to interconnect VLANs.
- Controlling broadcast, multicast, or unicast storms with port-level rate controls where appropriate.

Physicality and Topology

Another key differentiator of industrial automation and control systems is the environment in which they must perform, which depends on the environment in which the manufacturing process is occurring. Physical constraints in the manufacturing industry are significant. The networking systems need to recognize challenges in spatial and environmental conditions. End devices, such as controllers, drives, and HMIs, located in harsh environments such as the production floor often need to meet environment specifications such as IEC529 (ingress protection) or National Electrical Manufacturers Association (NEMA) specifications. The end device and control network layers may be located in physically disparate locations (up to miles away), and in non-controlled or even harsh conditions in terms of temperature, humidity, vibration, noise, explosiveness, electronic interference, and so on. These requirements are met by the switch and are not a specific focus of the Cisco EttF solution. Additionally, the physical media infrastructure is also driven by the physical requirements of the environment, with special consideration given to the potential for high noise, but is not currently a specific focus of this solution.

The physical layout of the manufacturing facility or the automation equipment also impacts the network topology for automation networks. Unlike traditional IT networks, which are largely redundant star topology networks, industrial automation and control networks have significant physical limitations that drive the use of topologies such as bus and ring. In manufacturing plants with long production lines, or equipment with long runs and interconnected operations (such as a printing press, or similar types of equipment), it is often not feasible or cost-effective to use a redundant star topology. In manufacturing environments, the costs of cabling are significantly higher than typical office conditions to meet the harsh physical requirements. The redundant star topology offers the best resiliency, convergence, and overall performance, but the additional cabling complexity and constraints of a redundant star limit its applicability.

In addition, many industrial automation and control applications do not need significant bandwidth at this time, and are therefore not significantly impacted by the potential bandwidth limitations of a ring architecture. In many cases, the industrial automation and control network is a combination of topologies, with large rings connecting multiple star-based manufacturing cells.

Cost considerations also drive the architectural and technology directions of many manufacturing companies. Given the physical layout of a manufacturing plant and industrial automation and control equipment, it is often significantly cheaper to implement a ring topology than a redundant star topology.

Based on these considerations, the design guidelines provide information regarding the trade-offs between the various topologies to help customers, partners, and account teams to make appropriate design decisions. Because of their significant use in manufacturing, bus topologies are discussed, as well as the associated trade-offs between bus, ring, and redundant star architectures (such as availability, and so on).

For a summary of the advantages and disadvantages of each topology, see [Cell/Area Topology Comparison](#), page 2-23.

Figure 1-3 shows a star topology.

Figure 1-3 Star Topology

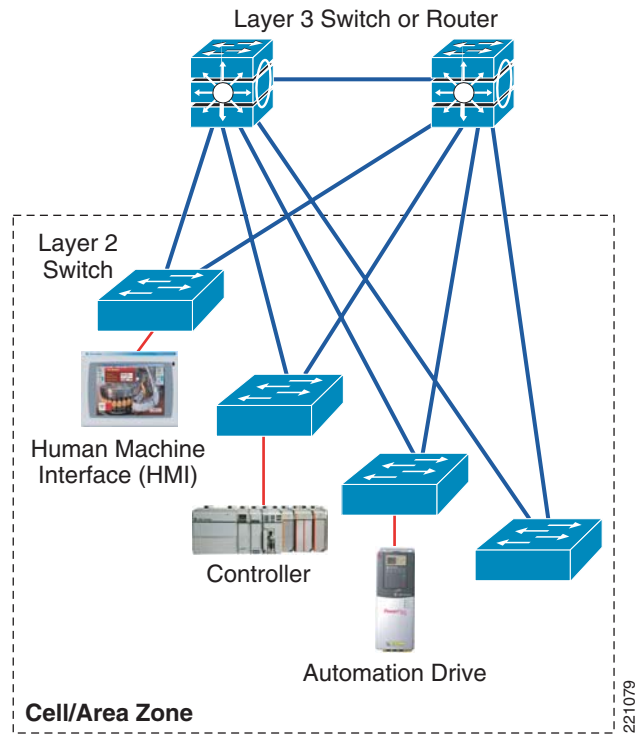


Figure 1-4 shows a ring topology.

Figure 1-4 Ring Topology

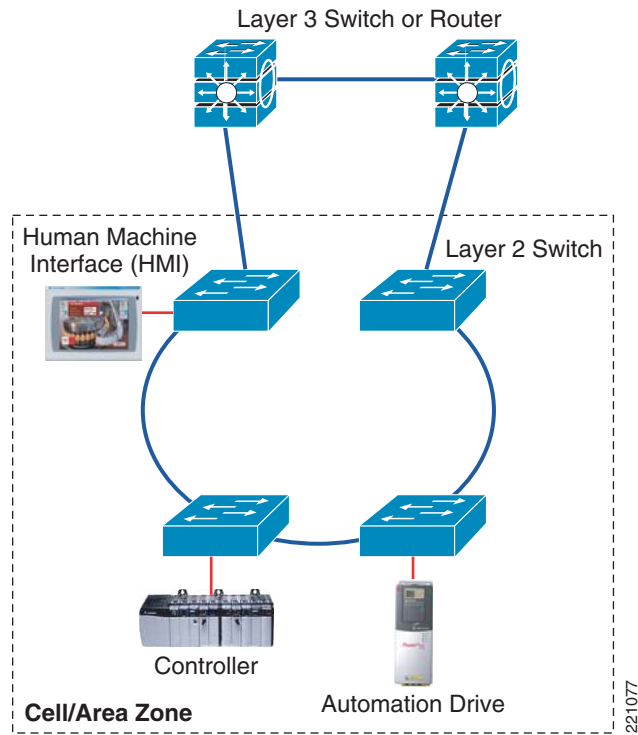
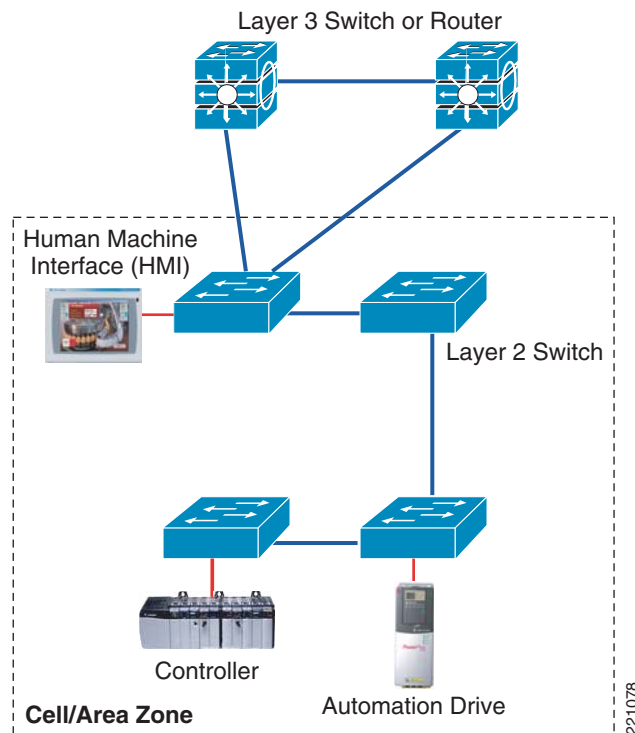


Figure 1-5 shows a bus topology.

Figure 1-5 Bus Topology



The Cisco EttF solution design and implementation guidelines include the following key considerations:

- Choose a topology that meets the performance, cost, and spatial requirements of the automation and control application.
- The layout of plant operations, conduit/wiring paths, cost, and desired level of availability determine whether the network topology follows a tree, ring, star, or trunk and drop topology, or a hybrid.
- Use ruggedized/hardened network devices in the factory environment where needed, but consider using non-industrial routers, switches, and firewalls where possible to reduce cost.
- The number of automation devices and grace ports for programming/troubleshooting and 10 percent spare for future expansion determines the type and size of switch needed at various levels.
- Hierarchically-layered switches may be required to address density, distance, or communication path challenges.

Availability

Availability of the industrial automation and control systems has a direct correlation to the operational efficiency of a production facility. Because the network is a key aspect of the overall system, these requirements translate directly to the network.

Note that limitations in the network technology may also limit the application of high availability features. For example, the lack of the ability of the network to converge quickly enough and the cost associated with redundant wiring have often led to non-redundant topologies being implemented in

industrial networking environments. The Cisco EttF solution outlines the capabilities so as to let customers and integrators make decisions on the level of network availability needed for the overall system.

High availability considerations are identified in each aspect of the Cisco EttF solution. Key considerations include the following:

- Creating alternative data communication paths, regardless of physical layout. Risk profile, opportunity cost, culture, and other variables determine how much and to what level redundant paths are required.
- Eliminating single points of failure with critical operations, including such items as dual power supplies, alternate routes for redundant media, redundant industrial automation and control network infrastructure: routers, switches, firewalls, and so on.
- Using advanced network resiliency and convergence techniques to improve availability, such as EtherChannel/trunks, 802.1w Rapid Spanning Tree Protocol (RSTP), Hot Standby Routing Protocol (HSRP), and so on.
- Although redundant star topology offers the best convergence capabilities, consider alternative ring recovery techniques when configured in a ring topology.
- Using routing configurations and protocols (for example, OSPF) to achieve high availability.

Security

IP-based networking facilitates interconnection of the industrial automation control system with the enterprise LAN. Many industries have implemented enterprise applications for more efficient production, as well as Internet business applications to communicate more efficiently with their suppliers, customers, and business partners. Internet-based enterprise resource planning (ERP) and supply chain management (SCM) systems simplify connections both to other organizations and to internal business processes. These connections can enable greater efficiencies in processes and manufacturing. In large manufacturing or utility operations, small percentage increases in efficiency can translate into significant cost savings.

However, connecting the industrial automation and control network to the enterprise network brings all the security risks of the Internet and enterprise network to the industrial automation and control system. Mitigating these risks is more difficult and more critical than in the enterprise network because of the higher requirement for availability in an industrial automation and control system and the sensitivity of these systems to various disruptions. Of the three security properties of confidentiality, integrity, and availability, control systems are primarily concerned with availability and integrity. Many of the applications that industrial automation and control networks support cannot be stopped or interrupted without serious physical or loss of productivity with measurable financial damage. On the other hand, in enterprise networks that are the primary design consideration for the Internet Protocol (IP) suite, confidentiality and integrity are the primary design considerations. For example, it is preferable for an e-commerce server to be temporarily unavailable rather than for it to lose transactions or divulge credit card numbers. Consequently, the network architectures, firewall configurations, intrusion detection configurations, and other aspects of a security deployment require tuning and customization to properly support industrial automation and control systems. The industrial automation and control systems industry has been struggling for several years to determine how to build secure, reliable control systems based on IP.

Although standards bodies such as ISA SP99 are still debating security design axioms, there is at least an approximate consensus on what a secure industrial automation and control architecture should provide. This includes an industrial automation and control network that is highly available and

redundant, has fast convergence, thus being more deterministic and therefore more suitable for real-time control, and is secure against both outside and inside threats. The specific security principles of the EttF architecture are as follows:

- Control data flows between different levels (ACLs, firewall rules, and so on).
- Prevent direct communication between industrial automation and control systems and enterprise systems.
- Restrict real-time production data to the industrial automation and control network.
- Restrict enterprise access to the mirror version or copies of production data to the DMZ.
- Authenticate and authorize user access based on the level within the industrial automation and control network and the role (read/read-write/local/remote/vendor/partner).
- Control rogue access inside the industrial automation and control network (port level MAC address controls, administratively shutdown unused ports, and so on).
- Control which devices can be plugged into the switch (for example, port security, DHCP snooping).
- Detect and mitigate malicious traffic originating from infected devices that are plugged into the industrial automation and control network.
- Detect and mitigate malicious traffic originating from the corporate IT network.
- Secure connectivity for remote access to automation devices.
- Use DMZ design options based on costs and levels of security and redundancy required.
- Limit rogue network communication activity from impacting networking devices (set root bridge, SNMP capabilities, and so on).
- Regarding data and services in the DMZ, connection initiation should originate from either the manufacturing or enterprise zone and terminate in the DMZ. Connections originating from the DMZ should be exceptions.
- Document and define policy and risk appropriate for the environment.

The above are provided as principles, with the understanding that customers may choose to make exceptions.

Manageability

Manageability is a key consideration for industrial automation and control systems. The network must be manageable by a group of people that are not trained or experts in standard network technologies or administration of this technology.

Key manageability concerns include the following:

- Configuring switches using the command-line interface (CLI), element management system (one GUI configures one switch), solution management system (one GUI configures multiple switches), or by downloading pre-defined templates
- Leveraging existing SNMP-based management systems when and where they make sense
- Using other network devices such as routers and security appliances that use similar functionality for configuration
- Using SmartPort templates for easy port configuration based on application types
- Assigning consistent IP addresses to devices because IP addresses are often coded into the logic and configurations of various industrial automation and control devices, rather than using dynamic IP address translation techniques such as dynamic naming services or DNS

- Considering various easy replacement options for network infrastructure elements
- Using systems that offer notification of industrial automation and control network events (for example, if an Ethernet link goes up or down), and the means to diagnose and debug problems with the network infrastructure
- Staging software upgrades for network devices
- Allowing for patch management of Windows-based automation devices
- Standardizing hardware and software elements wherever possible
- Driving the integration of basic network administration into the existing applications based on various industrial automation and control network protocols

Scalability

Industrial automation and control systems, once installed, tend not to grow, but rather are replaced or have additional lines, systems, or functions. Industrial automation and control systems come in a wide range of sizes, from the small OEM solutions to the extremely large factory complexes (for example, an automotive plant). The industrial automation and control system may include only a small number of devices (up to 50) to multiple 10,000s of devices. The solution architecture concepts and recommendations need to be applicable to that range, noting the considerations for various sizes.

This version of the solution architecture focuses on basic concepts, tested in typical small-to-medium network installations. The test environment and testing for this version of the solution architecture was not focused on full-range scalability and performance testing, but rather defining and testing core concepts that are applicable to a full range of factory floor sizes. The basic concepts in this guide are applicable to the range of industrial automation and control systems.

Key scalability considerations include the following:

- Network infrastructure sizing and performance constraints
- Network infrastructure tiering to meet spatial, size, and performance criteria
- Link aggregation to achieve higher bandwidth requirement
- IP addressing schema and allocation mechanism
- Maintenance and management considerations as manual tasks have greater impact in large environments

Scope of the Cisco EttF Solution

This phase of the Cisco EttF solution is meant to introduce a basic network architecture based on standards-based technologies to provide services to industrial automation and control systems. The first phase is a starter kit for customers, partners, and vendors seeking to implement a basic EttF solution.

Key aspects of this phase include the following:

- The Cisco EttF 1.1 solution focuses on wired solutions for the industrial automation and control systems.
- The Cisco EttF 1.1 solution is designed for small (less than 50 Ethernet endpoints or nodes) to medium (less than 200 Ethernet nodes) manufacturing environments.
- The Cisco EttF 1.1 solution introduces key technical considerations such as the following:
 - Topology
 - Real-time communications

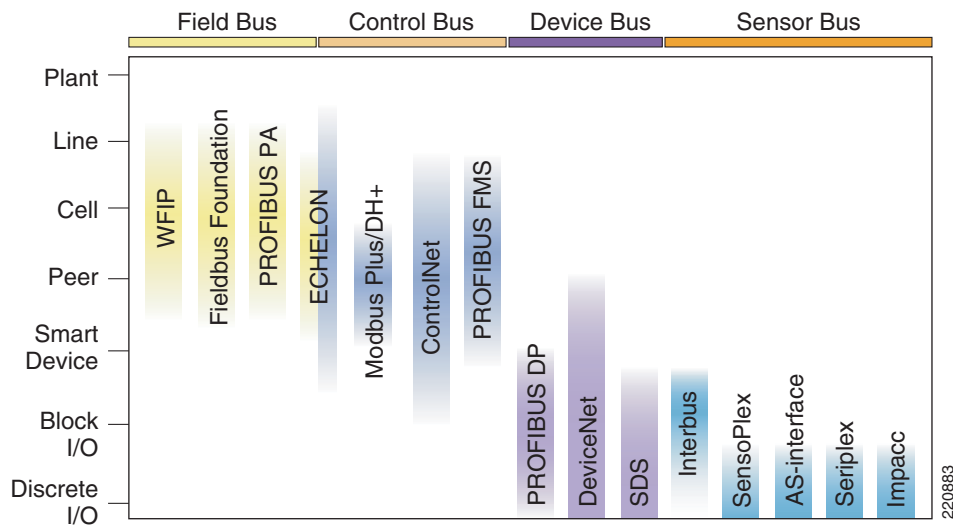
- OSI Layers 2 and 3 configuration including basic routing protocols
 - Insulation and segmentation including VLANs and DMZ design
 - Multicast traffic handling including IGMP protocol
 - Quality of service (QoS)
 - Redundancy and resiliency (including application of the standard RSTP)
 - IP address allocation, assigning, and related services (for example, DHCP and DNS)
 - Basic network management
 - Network security for the automation and control systems including DMZ, firewall, intrusion protection, endpoint security, and security monitoring, analysis, and response
- Design and implementation is based on Ethernet/IP (driven by CIP) based automation and control systems.

Industrial Automation and Control Background

History of Industrial Automation and Control Networks

From the beginning, manufacturing environments have relied on numerous technologies to enable communication at the plant, cell, or I/O level. Typically, the technologies deployed were purpose-built and vendor-specific. [Figure 1-6](#) provides a list of some of the types of protocols used in manufacturing environments.

Figure 1-6 Control Protocols Overview (Source: David Humphries, ARC)



The industrial automation and control industry as a whole has been moving away from the purpose-built and vendor-specific communication protocols for reasons that include the following:

- Difficulty of finding and training people who can debug a specific communication network technology
- Difficulty of extracting data for production reporting with older fieldbuses

- Expense of using vendor-specific technology to tie industrial automation and control systems together
- End user frustration in procuring industrial automation and control systems because of the confusion related to various fieldbus technologies
- Complexity of integrating various technologies into the overall industrial automation and control system

Ethernet and the IP protocol suite are now the ultimate solution to the multiple standalone industrial automation and control protocols. Ethernet and the IP protocol suite are standard technologies that provide a robust, cost-effective, easy-to-implement, and easy-to-troubleshoot mechanism for transmitting industrial automation data. Industrial networks based on standard Ethernet and IP technologies define the physical and transport layer for moving data. However, these technologies do not replace fieldbus communication standards per se. For example, fieldbus communication standards still define the data and its meaning and determine how messaging occurs. Each technology has its purpose, depending on the protocol and the data that is in the device.

Industrial Automation and Control System Components

Physical Layer

Many of the purpose-built and vendor-specific industrial technologies have specific physical media requirements that often require unique cables (such as co-axial) and specialized termination (such as serial connectors). These various physical layer specifications result in a forklift physical media upgrade when migrating from one system to another. In comparison, industrial Ethernet uses standard Ethernet wiring; either twisted pair cables, or multimode or single mode fiber. The connectors for these various types of Ethernet wiring are also standardized with RJ45 connectors for copper cables, and SC or ST connectors for fiber optic cables. The benefit of Ethernet is that after the Ethernet physical plant is installed, it can be used to connect hardware and software from various vendors as long as they support standard Ethernet.

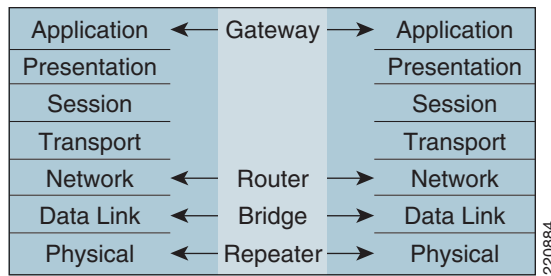
Operating speeds across conventional Ethernet wiring have options for 10 Mbps, 100 Mbps, and 1 Gbps. 10 Gbps is also available but is mainly being deployed in enterprise-wide backbone networks. Most industrial automation and control installations are using 10 Mbps or 100 Mbps Ethernet; although in some cases, Gigabit Ethernet is appearing in industrial system backbones. Based on performance requirements, a specific Ethernet physical layer may be needed to improve the reliability of industrial automation and control applications, such as high-noise applications. In extreme cases, this may even require sealed connectors. These operating constraints may limit the type of cabling and speed options, but there are typically physical Ethernet options.

The physical layout and communication requirements of a manufacturing environment dictate how various Ethernet-based resources are physically connected. Typical Ethernet environments have full duplex connection via a redundant star topology. Other options are possible such as ring, trunk and drop, and daisy chain. Specific operating constraints when using Ethernet in these other models are discussed in [Chapter 4, “Implementation of the Cell Zone .”](#)

Networking Equipment

As the industrial automation and control industry adopts standard Ethernet and IP technologies, it receives the added benefit of access to a wide range of standard networking equipment. The type of device required depends on many factors, the first being what type of communication protocol is in use. As [Figure 1-7](#) shows, various types of devices work at different layers of the OSI model.

Figure 1-7 OSI Model



In the beginning, many early factory floor Ethernet networks used simple, cheap repeaters (also known as hubs) to connect industrial automation and control systems together. In many cases, these were the same Ethernet hubs that were handling front-office workstations. As a multi-port broadcast device, a hub does the following:

“Creates one big collision domain, with all traffic shared. As more network nodes are added or traffic increases, every node in the collision domain has a greater chance of slowing communication or having a collision. Additionally, because industrial automation and control networks are not configured to differentiate between the relative importance of Ethernet packets, it is possible for non-essential traffic on the network (perhaps people backing up their computers to the network server or printing a large document across the network) to slow or collide with essential traffic (such as inter-PLC communication, or HMI polling.)”

(Source: http://www.industrialnetworking.com/support/general_faqs_info/hub_vs_switches.shtml)

The next advancement in industrial network design was to use switches; a type of multi-port Layer 2 bridge. Switches can divide networks into VLANs that segment devices into logical work groups. Ethernet switches also typically have a fast internal backbone, which helps eliminate collisions among data packets. Switches separate collision domains and map Ethernet nodes based on address and port. When an industrial automation and control device is directly connected to a non-blocking switch, potential collisions are eliminated. This occurs because full-duplex Ethernet devices can both send and receive packets of Ethernet data at the same time. This increases the level of determinism of Ethernet, assuring that packets arrive with much greater certainty, and that each port has more bandwidth available at any time.

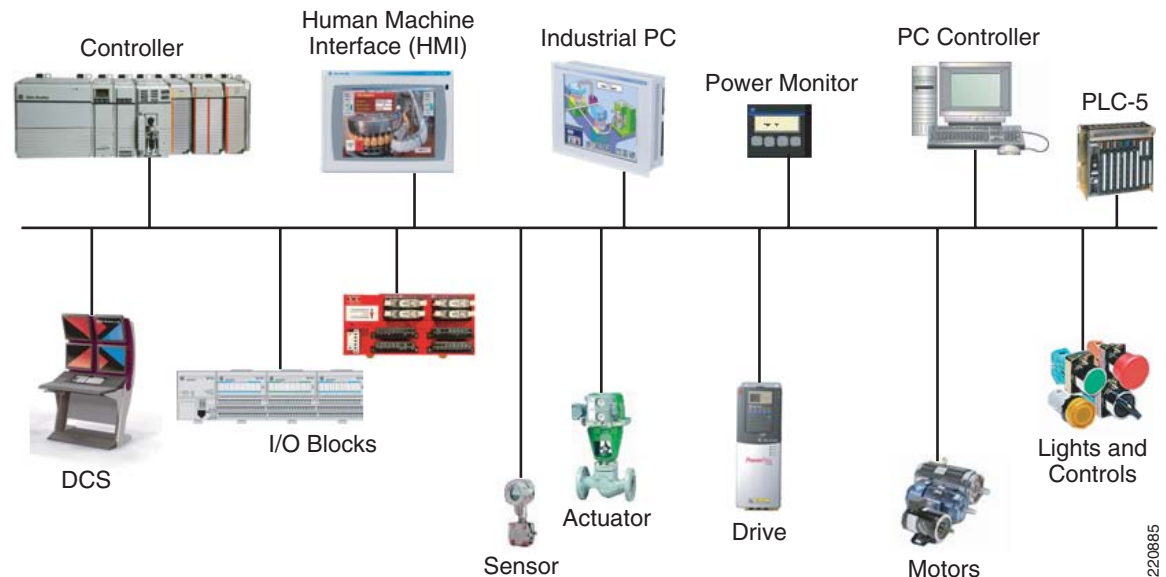
Adding some intelligence to the switch improves traffic management capabilities, meaning that the switch can provide more granular quality of service for industrial automation and control networks. One example is the management of multicast traffic to communicate critical I/O data applied in most implementations of EtherNet/IP. Management of the multicast (rather than treating it as broadcasts) significantly reduces the number of messages that end devices and network infrastructure must process, leading to better network and device performance. As another example, by assigning a priority to time-sensitive data, intelligent Ethernet switches can elevate that traffic above lower-priority data. This ensures that high-priority traffic always traverses the network, even if the network becomes congested. Switches can also classify, reclassify, police, mark, and even drop incoming data packets as application priorities require. The use of managed versus unmanaged switches is a key consideration facing those implementing industrial automation and control networks today. Both Cisco and Rockwell Automation highly recommend the use of managed switches. For further details on managed versus unmanaged switches, see [Network Design Overview, page 2-25](#).

In some cases, Layer 3 switches, or routers, are used in manufacturing environments. This occurs primarily when different VLANs or different IP domains need to be connected. Regardless of the specific layer being connected, switches provide industrial automation and control networks with many of the safeguards that were provided by distinct, proprietary-control networks. The specifics of how a Layer 2 switch is used compared to a Layer 3 switch and how QoS can be implemented is addressed in [Cell/Area Zone, page 2-10](#).

Industrial Automation and Control Devices

Numerous types of devices are used in industrial automation and control systems. Some are small, simple, single function sensors or I/O blocks, while others are complex, programmable automation controllers (PACs). The breadth and depth of available devices is driven primarily by industrial automation and control vendors and their partners and suppliers. Figure 1-8 shows some of the various types of devices used in the manufacturing environment.

Figure 1-8 Industrial Devices



Older lower-level industrial automation and control devices tend to be simple, low-cost devices that use a specific industrial automation and control protocol and are capable of only low data rates and volumes. More advanced industrial automation and control devices have internal logic optimized for I/O control with the ability to support higher data rates and volumes. Many of these newer industrial automation and control devices now come standard with more communication options including Ethernet and IP. For example, controllers now come with options of 512 K to 100+ MB of memory, integrated serial (integrated RS-232-C, RS-422 or RS-485 ports for SCADA, ASCII, or peer-to-peer communication), modular and scalable EtherNet/IP, and ControlNet and/or DeviceNet communications.

The trend with most industrial automation and control devices is to add more functionality and capabilities at all levels. This is occurring because of the continual evolution in the microelectronics industry and access to lower cost components. The low cost of microcontrollers is already making it easy for design engineers to include Ethernet and IP in a growing number of products that tie into industrial systems. As with many electronic technologies, after a few high-end products incorporate a feature or function, it rapidly becomes a common attribute on many of the emerging new products.

Even so, there is and will continue to be a place for simple, low cost, and lower capability devices in industrial automation and control systems. When Ethernet and IP represents too much of a cost and capability increase for the end device itself, these devices will continue to communicate via simple, non-Ethernet I/O networks; for example, a distributed I/O device used as an Ethernet gateway connecting a number of simple devices, such as a pushbutton, to a controller.

Industrial Computing

Computing technology has been used for years in purpose-built and vendor-specific manufacturing environments. Just as with IT, the technology has migrated from mainframes and mini-computers with dumb terminals to standalone, dedicated computing platforms. With the cost of computing highly commoditized, the trend now is to put computing anywhere in the industrial automation and control network using high performance CPUs. By using fanless and diskless PCs with features such as capacitive touchscreens, class 1 division 2 certification, and mission-critical solid-state drives, computing platforms are now suitable for any harsh industrial or embedded application.

From an operating system perspective, most industrial automation and control vendors have moved away from legacy or custom-built operating systems to common off-the-shelf operating systems based on Microsoft or Unix derivatives (including Linux) for many applications. The benefit of this development is a simpler and faster application programming environment both for vendors as well as end users. This migration has coincided with the overall general trend in the software industry towards Internet browser-based technology. This gives automation vendors the ability to embed web interfaces directly into industrial automation and control devices.

The downside of all these developments is a significant amount of system complexity related to security and patch management. The specific application requirement of industrial automation and control systems is discussed in [Chapter 2, “Solution Architecture.”](#)

Industrial Automation and Control System Communication Protocols

Communication Model

The communication messaging model in manufacturing environments has only loose ties to traditional client-server or peer-to-peer IT models. Unlike the typical IT environment, standards-based Ethernet and IP industrial automation and control communications have different patterns, loads, and frequencies required by the manufacturing process they support. Standards-based industrial automation and control communications are also driven by status polling between devices, cyclic data transfer, or change of state message patterns. The various requirements of the layers previously discussed have led key industrial automation and control providers to define a variety of communication models, including OSI layer 1–7 networking protocols.

These communication models have both strong commonalities and differences. In common, they differentiate the control or I/O traffic between devices and the PACs (EttF levels 0–1) and administration traffic within the upper layer applications down to the PAC (EttF levels 1–3). This differentiation is made to meet the stringent requirements at these lower levels (see [Automation and Control Reference Model, page 2-1](#)). However, the models can differ greatly at the control or I/O level. One example is the producer-consumer model applied in the Open Device Vendor Association (ODVA) Common Industrial Protocol (CIP). This model describes how devices “produce” data to be “consumed” by other devices; in particular, the PACs that take action on their data and control their behavior. These models are incorporated into the industrial automation and control protocols described below. They are important because they impact or shape the network traffic that is produced by the applications that use them.

CIP, for example, defines two distinct message types: *explicit* messages and *implicit* messages. For explicit messages, the action is explicitly implied in the message; for example, read the value of a variable. Explicit is a request/response, client/server-like protocol typically used for “information” and administrative messaging over TCP. For implicit messages, the communicating parties inherently know how to parse the message content because of contextual knowledge. Implicit is a unicast or multicast, producer/consumer protocol designed for “Input/Output” or “control” data transfer over UDP. Explicit messages are information messages used for additional device configuration and diagnostics of features

of the industrial automation and control device. Explicit messages are highly variable in both size and frequency based on configuration and application. On the other hand, real-time control or implicit messaging is sent at specified intervals, and although the size can vary, it is consistent after the configuration is set and is generally smaller than explicit messages. Implicit messages contain control data that must be interpreted very quickly by the receiving device, which demands network and end-device performance that is different than other traffic. With implicit traffic, the communication model is UDP-based (either unicast or multicast), to minimize processing resources and time on the end device. The communication model in manufacturing environments can be unicast, multicast, or broadcast driven by master/slave, multi-master, or peer-to-peer relationships.

These differing communication models and the protocols into which they are embedded drive various configuration considerations for the networks that support the automation and control systems. For example, the CIP use of multicast traffic generates different network configuration considerations. However, these differences are focused on the lower levels of the network where the networking requirements are the most significantly different than standard IT networks; for example, levels 0–1 (see [Chapter 2, “Solution Architecture,”](#) for definitions of industrial automation and control levels). Often, the various communication models apply standard Ethernet and the IP protocol suite for the traffic at higher levels of the industrial automation and control network, but use varying degrees of the standard Ethernet and IP protocol suite, if at all, for the control-level traffic.

Industrial Automation and Control Protocol Overview

Most Ethernet and IP-based industrial automation and control protocols have a common core. This includes the data transmission technology (Layer 1), the bus access method (CSMA/CD, Layer 2), the Internet Protocol (IP, Layer 3), the TCP and UDP protocols (Layer 4), the Hypertext Transfer Protocol (HTTP), the File Transfer Protocol (FTP), and the Simple Network Management Protocol (SNMP). All these are established in the IT industry and are being implemented to varying degrees, unchanged in industrial automation and control environments.

The goal of an Ethernet and IP-based industrial automation and control network is to ensure that the control protocol of choice, assuming it is based on standard Ethernet and IP, is supported to meet the operating constraints of the industrial automation and control systems.

Table 1-1 shows a list of some industrial automation and control protocols that support or partially support standard networking.

Table 1-1 Control Network Protocols

Fieldbus Protocol	Ethernet Implementation	Leading Vendors	Standards Body	Application
DeviceNet, ControlNet	EtherNet/IP (EIP)	Rockwell Automation, Schneider (EIP), Omron, Eaton	ODVA	Industrial automation process control
PROFIBUS DP, PA, and so on	PROFINET CBA, I/O, IRT, and so on	Seimens	PROFIBUS Foundation	Industrial automation process control
Modbus	Ethernet Modbus TCP	Schnieder	Modbus.org	Industrial automation process control
Foundation Fieldbus	Foundation Fieldbus High-Speed Ethernet	Emerson, Honeywell, ABB	Fieldbus Foundation	Process control
CAN/Lightbus	EtherCAT	Beckhoff	EtherCAT Technology Group	Motion control
CAN/CAN-Bus	ETHERNET Powerlink	Bernecker, + Rainer	ETHERNET Powerlink Standardization Group	Motion control
Sercos Interface	Sercos III	Bosch Rexroth	SERCOS International	Motion control

However, there are some differences in the application protocols for real-time communication as well as the object and engineering models for system configuration. These differences lead to different considerations and deployments of industrial automation and control networks. Of these protocols, EttF Architecture Phase 1 is focused on exploring only the ODVA implementation of CIP on the Ethernet and the IP protocol suite referred to as EtherNet/IP.

In addition to the approach taken to integrate with Ethernet (physical and data layers) and the IP protocol suite, these application protocols have also identified various messaging frameworks that dictate the type of traffic and traffic patterns found in the industrial automation and control network.

Table 1-2 briefly describes some of the key characteristics of the various protocols.

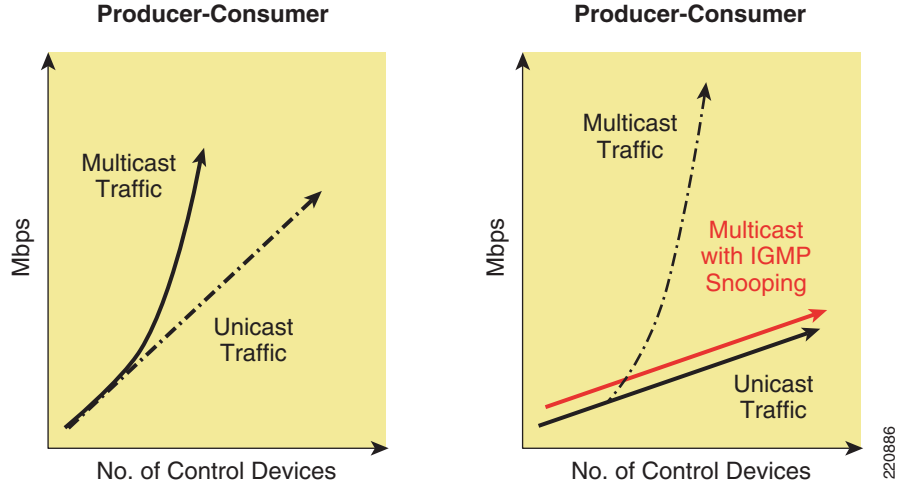
Table 1-2 *Various Features of Different Industrial Ethernet Protocols*

IE Protocol	Encapsulated Telegram	TCP/IP UDP/IP	Port Usage	Profile/Object Support
EtherNet/IP	Common Industrial Protocol (CIP)	TCP/IP explicit UDP/IP implicit	44818 2222	Legacy
Modbus/TCP	Modbus	TCP/IP	502	Legacy
PROFINET CBA PROFINET I/O and IRT	Profibus Plus	TCP/IP Special Data link	Dynamic	ORPC
OPC (OLE for Process Control)	DCOM/XML	TCP/IP	Dynamic	DCOM / XML
MMS TCP/IP	MMS	TCP		MMS
.NET for Manufacturing	COM	TCP/IP	80	DCOM/ XML
Foundation Fieldbus HSE	H1	UDP/TCP Optimized	Dynamic	Legacy plus
iDA	N/A	UDP/IP	Dynamic	XML
AADS-net	N/A	UDP/IP	Dynamic	Possible

The various protocols and their application of the Ethernet/TCP/IP stack drive particular considerations in the configuration of the network. Using CIP and the “producer-consumer” model as an example, the control-level devices use UDP unicast and/or multicast to send information out on the network. Although the choice to use multicast or unicast is the choice of the device vendor, multicast is the prevalent mode of communication of I/O data in CIP implementation in Ethernet/IP.

The ability to control multicast traffic in the control levels of the network is a very important aspect of the network devices. Figure 1-9 shows how without multicast control features, the bandwidth requirements in an industrial automation and control network application increase exponentially (versus a linear increase) with the increase in the number of devices. This is just an example of the type of network design and implementation considerations specific to industrial automation and control protocols.

Figure 1-9 Producer-Consumer Network Impact

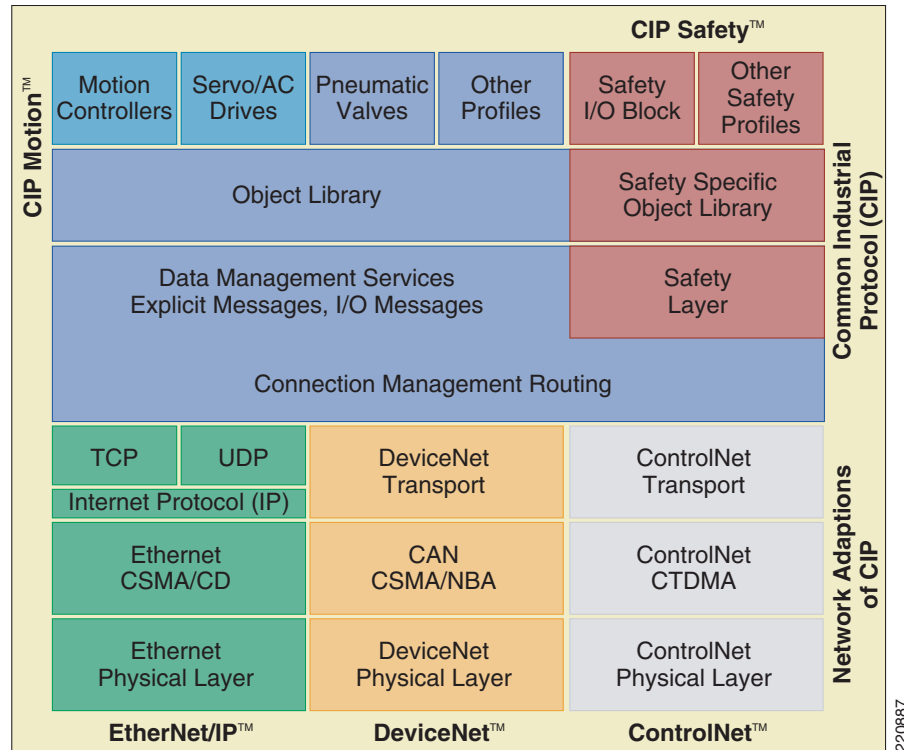


In summary, a wide number of protocols are in operation in industrial automation and control networks. Design and implementation guidelines need to consider the various protocols and their underlying communication models. This initial version covers EtherNet/IP and the CIP protocol along with the producer-consumer communication model. Over time, this architecture and the subsequent deliverables will take into account the various communication relationships, protocols, and Ethernet/TCP/IP implementations when designing, implementing, and operating an industrial automation and control network.

Common Industrial Protocol Overview

CIP is a messaging protocol that defines how various industrial automation and control devices, systems, and applications come together to form an industrial automation and control system, as shown in [Figure 1-10](#). CIP is an application-layer protocol (OSI Layers 5–7). EtherNet/IP extends the application of Ethernet TCP/IP to the factory floor for CIP-based applications.

Figure 1-10 Common Industrial Protocol (Source: ODVA)



CIP is a connection-based protocol and offers two main types of messaging: explicit and implicit. The protocol specifies a set of objects and services used to develop industrial automation and control systems. CIP is implemented on three network layers: DeviceNet, ControlNet, and EtherNet/IP. This document is concerned only with EtherNet/IP.

For more information on CIP and the various network layers, see the ODVA website at the following URL: <http://www.odva.org>; and ControlNet International at the following URL: <http://www.controlnet.org>.

The important aspects of the CIP implementation of EtherNet/IP are the various types of messaging that are used and how they are implemented in standard Ethernet TCP/IP.

Table 1-3 provides a brief overview of the CIP messaging types and their key networking characteristics.

Table 1-3 CIP Communication Overview

CIP mode	CIP message type	Description	Response time requirements	Layer 4 type	Packet Size (Bytes) ¹	Port ²
Unconnected	Unconnected	Basically used to open a CIP connection with another device. This mode is only temporarily used.	Seconds	TCP	~500	44818
Connected	Explicit	Non-time-critical information data. For example, between a controller and a manufacturing historian application.	100s of milliseconds to seconds	TCP	~500	44818
	Implicit or I/O	Time-critical control information usually passed on regular intervals in a “producer-consumer” multicast communication model. For example, between a controller and a drive (PAC to device) or between controllers (PAC-to-PAC).	< Millisecond to 10s of milliseconds	UDP multicast and unicast	100 - 200	2222

1. These are typical numbers, although depending on the application and customer can be different

2. These are registered ports for EtherNet/IP, although non-registered ports may be used in EtherNet/IP.

Other key technical considerations for EtherNet/IP implementations include the following:

- The producer-consumer model specifies that “producers” of I/O data communicate via UDP unicasts or multicasts. The consumers (for example, controllers) typically respond with UDP unicast messages. Rockwell Automation and the ODVA therefore recommend the application of IGMP to manage the multicast traffic flow.
- Multicast traffic in current installations is stamped with a time-to-live (TTL) value of 1, rendering the multicast packets un-routable. This limitation forces all nodes that produce and consume information from one another to exist in the same subnetwork/VLAN. The capability to change/increase this value has been outlined in the most recent version of Volume 2 (EtherNet/IP Adaptation of CIP) of the CIP Specification release 1.3, which was published in December, 2006. For the purpose of this solution architecture, it is assumed CIP-based multicast traffic is not routable.
- By the current EtherNet/IP standard, a multicast group is created for each Ethernet adapter that “produces” information, and for each “produced” tag (shared piece of data) established by a controller. EtherNet/IP specifies an algorithm to establish the multicast address and the commands to join and leave multicast groups. Current EtherNet/IP multicasting is based on IGMP version 2, although there are devices (producers) that may still be based on IGMP version 1. IGMP version 1 devices should not be an issue or be a consideration in the Cisco EttF solution.
- Depending on the device producer, options may be enabled to configure whether the traffic generated by the PAC for each “produced” tag is unicast and multicast. This allows more flexibility in cell/area design and a means to manage the number of multicast groups within a cell/area.
- No CIP-EtherNet/IP QoS guidelines have been developed, so devices and applications typically do not mark either MAC-layer class-of-service (CoS) or IP-layer Differentiated Services Code Point (DSCP) fields. The ODVA has included a placeholder in the specification for QoS and work is currently ongoing to develop an approach.



CHAPTER 2

Solution Architecture

Overview

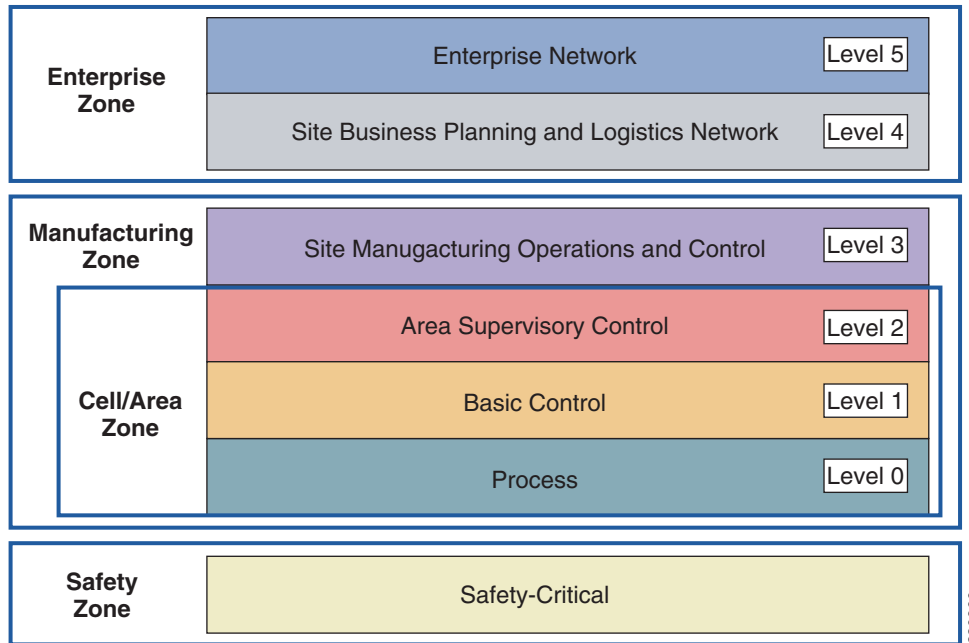
This chapter provides an overview of the Ethernet-to-the-Factory (EttF) solution architecture, as a means to describe the various systems, components, and their relation to each other to give context to the networking function and technical requirements. EttF is an architecture that provides network and security services to the devices, equipment, and applications found in industrial automation and control systems and integrates them into the wider enterprise network. As noted before, the networking requirements of a production facility often differ from a typical IT network. This solution architecture overview provides the background and description of an industrial automation and control network model and highlights the differences between the EttF architecture and the IT network.

Reuse is an objective of any architecture, and this is the case with the EttF solution architecture. Industrial automation and control systems are deployed in a large variety of industries, such as automotive, pharmaceuticals, consumer goods, pulp and paper, oil and gas, and energy. Industrial automation and control systems are also deployed in a wide variety of types of manufacturing, such as batch, discrete, process, and hybrid manufacturing. Size of deployments include small (less than 50 devices), medium (less than 200 devices), and large (from 200 up to 10,000s of devices). This architecture is meant to be a model/structure to be used in all these types of manufacturing environments, but clearly it must be tailored to the industry, type of manufacturing, size, and eventually the customer.

Industrial Automation and Control Reference Model

To understand the security and network systems requirements of an industrial automation and control systems in a production facility, this guide uses a framework to describe the basic functions and composition of a manufacturing system. The Purdue Model for Control Hierarchy (reference ISBN 1-55617-265-6) is a common and well-understood model in the industry that segments devices and equipment into hierarchical functions. It has been incorporated into many other models and standards in the industry. Based on this segmentation of the production floor technology, the Instrumentation, Systems, and Automation Society (ISA) SP 99 Committee for Manufacturing and Control Systems Security has identified the levels and framework shown in [Figure 2-1](#). Each zone and the related levels are then subsequently described in detail.

Figure 2-1 Six Level Plant Architecture



This model identifies “levels” of operations and defines each level. In this document, “levels” generally refer to this concept of levels of operations. The OSI model is also commonly referred to when discussing network architectures. The OSI model refers to “layers” of network communication functions. In this document unless specified, “layers” refer to layers of the OSI model.

Safety Zone

Safety is considered the highest priority function in industrial automation and control systems. Historically, safety sub-systems have been hard-wired. More recently, these systems have been implemented in totally dedicated infrastructure to ensure that the industrial automation and control equipment does not pose a threat to people or the environment. In some industries, these sub-systems have not shared any resources (power, network, and so on) with the rest of the industrial automation and control system with which they work. These sub-systems do have specific protocols and networking technologies, but because of the reliability and the impact of failure, this will probably be the final frontier for Ethernet and IP technologies. More recently, however, there have been enhancements to industrial automation and control networks such as the Open DeviceNet Vendors Association (ODVA) Common Industrial Protocol (CIP) safety solution where the safety protocol runs on the same network infrastructure as the standard protocol. CIP safety systems on DeviceNet and EtherNet/IP have been successfully developed and installed.

Cell/Area Zone

The cell/area zone is a functional area within a production facility. In an automotive plant, it may be a body shop or a sub-assembly process. In a food and beverage facility, it may be the batch mixing area. It may be as small as a single controller and its associated devices on an assembly line, or multiple controllers on several tanks. Each production facility defines the cell/area zone demarcation differently and to varying degrees of granularity. For the purposes of this architecture, a cell/area zone is a set of devices, controllers, and so on, that are involved in the real-time control of a functional aspect of the

manufacturing process. To control the functional process, they are all in real-time communication with each other. Most production facilities have multiple cell/area networks. This zone has essentially three levels of activity occurring, as described in the following sections.

Level 0—Process

Level 0 consists of a wide variety of sensors and devices. These devices perform the basic functions of the industrial automation and control system, such as driving a motor, measuring variables, setting an output, and performing key functions such as painting, welding, bending, and so on. These functions can be very simple (temperature gauge) to highly complex (a moving robot). See [Industrial Automation and Control Devices, page 1-21](#) for a more detailed explanation.

These devices take direction and communicate status with the control devices in the next level of the model. In addition, other devices or users may need to directly access devices to perform maintenance or resolve problems on the devices.

Level 0 devices usually do the following:

- Drive the real-time, deterministic communication requirements
- Measure the machine variables and control process outputs based on time
- Exist in challenging physical environments that drive topology constraints
- Vary according to the size of the network from a small (10s) to a large (1000s) number of devices
- Once implemented, are not replaced all together until the production line is overhauled or replaced, which is typically five or more years

Because historically these requirements have not been met by the Ethernet and TCP/IP technologies, a wide variety of proprietary network protocols has arisen. These protocols often cover Layers 1–7 of the OSI model. Ethernet and TCP/IP are being integrated into their frameworks, but with differing approaches. See [Industrial Automation and Control System Communication Protocols, page 1-22](#) for an overview of these protocols.

Control engineers such as electrical, process, and so on, and *not* the IT departments, typically design and implement these devices and the networks that support them.

Level 1—Basic Control

Level 1 consists of interfaces to the Level 0 devices (I/O, linking devices, bridges, and so on) and controllers. In manufacturing, this is typically a process-level controller (PLC), also known as a programmable logic controller, rather than a distributed control system (DCS) found in distributed manufacturing. For the purposes of this solution architecture, this document uses the terms *controller* or *programmable automation controller* (PAC), which refer to the general range of controllers used across manufacturing, both process and discrete.

Most PACs run proprietary operating systems that are programmed and configured from workstations or other advanced control systems. PACs are basically very simple, modular computers that consist of some or all of the following:

- A controller that computes all the data and executes programs loaded onto it
- I/O or network modules that communicate with devices, human-machine interfaces (HMIs), or advanced control systems
- Power modules that deliver power to the rest of the PAC and potentially other devices

PACs are the brains of the industrial automation and control system, making the basic decisions based on feedback from the devices found at Level 0. PACs act alone or in conjunction with other PACs to manage the devices and thereby the manufacturing process. PACs are programmed via a workstation,

and configured and managed via an external device referred to as an HMI, which is considered a Level 2 device. PACs also communicate with information and production control systems (historian, asset manager, manufacturing execution system, production scheduler, and so on) in Levels 2 and 3. PACs provide status and data about the actual process being controlled as well as take input for execution (for example, fulfill an order).

Thus, PACs produce network traffic in three directions:

- Downward with the devices in Level 0 that they control and manage
- Peer-to-peer with other PACs to manage the automation and control for a cell/area or production line
- Upward with HMIs and information and production control systems

The PAC performs a hub function in the manufacturing control area. The PAC translates high-level parameters (for example, recipes) into executable orders and manages those parameters throughout the manufacturing process. They also consolidate the I/O traffic from devices and pass data on to the upper-level plant floor functions. In some implementations, the PAC is also the physical network hub for the control network as the only device where connections to Level 0 devices, HMIs, and advanced control systems exist.

PACs must also meet the requirements being driven by the Level 0 devices, as described above.

Level 2—Area Control

Production facilities are usually laid out in areas or cells where a particular aspect of the manufacturing process occurs. In an automotive plant, this might be a body shop, paint shop, or general assembly. In a process solution, it might be a batch mixing area. Level 2 represents the systems and functions associated with the runtime supervision and operation of an area of a production facility. These include the following:

- Operator interfaces or HMIs
- Alarms or alerting systems
- Process historian batch management systems
- Control room workstations

Depending on the size or structure of a facility, these functions may exist at the site level (level 3). These systems communicate with the PACs in Level 1 and interface or share data with the site or enterprise (Level 4/5) systems and applications. These systems are primarily based on standard computing equipment and operating systems (Unix-based or Microsoft Windows). For this reason, these systems are more likely to communicate with standard networking protocols (Ethernet and TCP/IP).

Additionally, because these systems tend to be more aligned with standard IT technologies, they may also be implemented and supported by IT-skilled personnel, although typically they are implemented by the manufacturing organization. These people may or may not belong organizationally to IT.

Manufacturing Zone

The manufacturing zone comprises the cell/area networks and site-level activities. It typically contains multiple cell/area zones. The manufacturing zone is important because *all* the systems, devices, and controllers critical to monitoring and controlling the factory floor operations are in this zone. To preserve smooth plant operations and functioning of the systems and network, this zone requires clear logical segmentation and protection from the above levels of plant/enterprise operations. Beyond the cell/area networks, there is one additional level of activity that comprises the manufacturing zone.

Level 3—Site Level

Level 3, the site level, represents the highest level of industrial automation and control systems. The systems and applications that exist at this level manage site-wide industrial automation and control functions. Levels 0 through 3 are considered critical to site operations. The systems and functions that exist at this level include the following:

- Production reporting (for example, cycle times, quality index, predictive maintenance)
- Plant historian
- Detailed production scheduling
- Site-level operations management
- Asset and material management
- Control room workstations
- Patch launch server
- File server
- Other domain, AD, terminal server
- Staging area
- Administration and control applications (for example, domain servers, patch distribution, terminal services)

These systems may communicate with the PACs in Level 1, function as a staging area for changes into the production zone, and share data with the enterprise (Levels 4/5) systems and applications. These systems are primarily based on standard computing equipment and operating systems (Unix-based or Microsoft Windows). For this reason, these systems are more likely to communicate with standard networking protocols (Ethernet and TCP/IP).

Additionally, because these systems tend to be more aligned with standard IT technologies, they may also be implemented and supported by IT-skilled personnel. These people may or may not belong organizationally to IT.

Enterprise Zone

Level 4—Site Business Planning and Logistics

Level 4 is where functions and systems exist that need standard access to services provided by the enterprise network. This level is viewed as an extension of the enterprise network. The basic business administration tasks are performed here and rely on standard IT services. These functions and systems include wired and wireless access to enterprise network services such as the following:

- Internet access
- E-mail
- Non-critical production systems such as manufacturing execution systems and overall plant reporting (for example, inventory, performance, and so on)
- Enterprise applications such as SAP, Oracle, and so on

Although important, these services are not viewed as critical to the industrial automation and control system and thus the factory floor operations. Because of the more open nature of the systems and applications within the enterprise network, this level is often viewed as a source of threats and disruptions to the industrial automation and control network.

Users and systems in Level 4 often require summarized data and information from the lower levels of the industrial automation and control network. The network traffic and patterns here are typical of a branch or campus network found in general enterprises.

This level is typically under the management and control of the IT organization.

Level 5—Enterprise

Level 5 is where the centralized IT systems and functions exist. Enterprise resource management, business-to-business, and business-to-customer services are typically located here. Often the external partner or guest access systems exist here, although it is not uncommon to find them in lower levels of the framework to gain flexibility that may be difficult to achieve at the enterprise level.

The industrial automation and control systems must integrate with the enterprise applications to exchange production and resource data. Direct access to the industrial automation and control systems is typically not required, with the exception of partner access. Access to data and the industrial automation and control network must be managed and controlled to maintain the availability and stability.

The services, systems, and applications at this level are directly managed and operated by the IT organization.

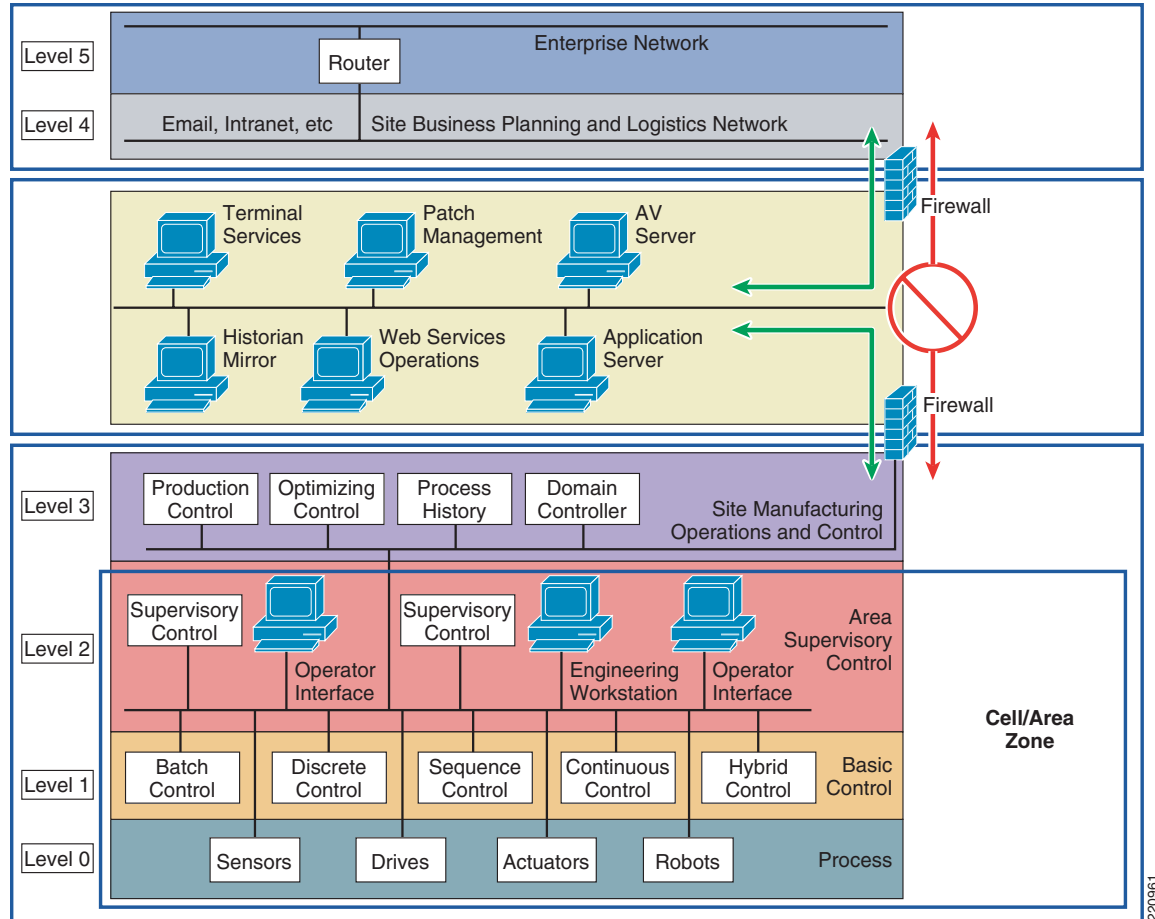
Ethernet-to-the-Factory Framework

The Purdue Model and ISA SP99 have identified levels of operations and key zones for industrial automation and controls systems. In addition to the levels and zones, Cisco includes an additional demilitarized zone (DMZ) between the enterprise and manufacturing zones. The purpose of the DMZ is to provide a buffer zone where data and services can be shared between the enterprise and manufacturing zones. The introduction of the DMZ is critical in maintaining availability, addressing security vulnerabilities, and abiding by regulatory compliance mandates. In addition, the DMZ allows for segmentation of organizational control; for example, between the IT organization and production. This segmentation allows different policies to be applied and contained. For example, the production organization may apply security policies that are different from the IT organization, and apply them to the manufacturing zone. The DMZ is where the policies and organizational control can be divided.

These levels and zones form the base framework around which the network infrastructure and services are designed for the EttF solution (see [Figure 2-2](#).)

The following sections contain a more detailed description of each zone, including the DMZ and their related functions and components.

Figure 2-2 Ethernet-to-the-Factory Framework



220961

Table 2-1 provides a short summary of each level.

Table 2-1 *Purdue Model for Control Hierarchy*

Level	Name	Description
Enterprise Zone		
Level 5	Enterprise	Corporate level applications (for example, ERP, CRM, document management) and services (for example, Internet access, VPN entry point) exist in this layer.
Level 4	Site business planning and logistics	Production facility IT services exist in this layer and may include scheduling systems, material flow applications, manufacturing execution systems (MES), and local IT services (for example, phone, printing, security/monitoring).
Demilitarized Zone		
	DMZ	Provides a buffer zone where services and data can be shared between the manufacturing and enterprise zones. In addition, the DMZ allows for easy segmentation of organizational control. Cisco recommends that the DMZ be designed so that no traffic traverses the DMZ. All traffic should originate/terminate in the DMZ.
Manufacturing Zone		
Level 3	Manufacturing operations and control	Includes the functions involved in managing the work flows to produce the desired end products. Examples include dispatching production, detailed production scheduling, reliability assurance, site-wide control optimization, security management, network management, and potentially other required IT services (for example, DHCP, LDAP, DNS, file servers, and so on).
Cell/Area Zone		
Level 2	Area supervisory control	Control room, controller status, control network/application administration, and other control-related applications (for example, supervisory control, historian).
Level 1	Basic control	Multiple controllers, dedicated HMIs, and other applications may talk to each other to run a part or whole production line.
Level 0	Process	Where devices (for example, sensors, actuators) and machines (for example, drives, motors, robots) talk to controller or multiple controllers for redundancy.
Safety Zone		
	Safety-critical	Devices, sensors, and other equipment used to manage the safety functions of industrial automation and control systems.

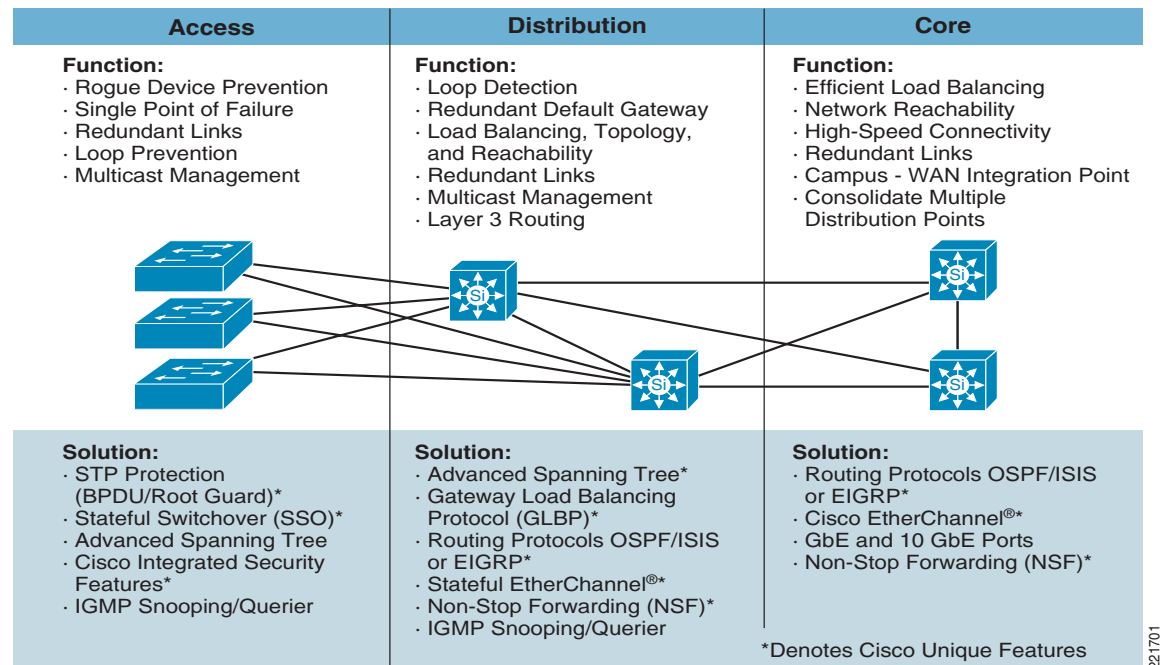
Campus Network Reference Model

The EttF framework reflects the basic functions of a production facility. This is the key model for this solution architecture. However, as identified earlier, the goal of this architecture is to integrate the knowledge and expertise from both a manufacturing perspective as well as an IT perspective where expertise in standard networking technologies exists. An important and relevant model for network architectures is the Cisco Enterprise Campus Network. The enterprise campus solution architecture incorporates key networking concepts and models. The EttF solution architecture comprises many of the concepts and models of the enterprise campus solution architecture, although it does not incorporate the entire scope of that solution because not all concepts are applicable to the production facility.

This section briefly introduces the campus network and some of the key concepts of its solution architecture. The Cisco Enterprise Campus Network combines a high-availability core infrastructure of intelligent switching and routing with an overlay of productivity-enhancing technologies, including IP communications, mobility, and advanced security. This document refers to the campus network documentation and the concept of core, distribution, and access. The campus network solution architecture documentation includes design and implementation guides, presentations, and other valuable material relevant to an industrial Ethernet network. This material is meant to enhance this solution architecture. Note that not *all* aspects of the campus network are reflected in this solution architecture; for example, wireless mobility and unified communications.

Figure 2-2 shows the three-layer model of access, distribution, and core, which is a very common and well-known network approach in the IT world.

Figure 2-3 Campus Network—Core High Availability Network



22.1701

The access layer provides the first level of access to the network. Layer 2 (OSI model) switching, security, and QoS reside at this layer. The distribution layer aggregates the access layer switches and provides security and network policy enforcement. Layer 3 protocols are used at this layer to provide load balancing, fast convergence, and scalability. The core is the backbone of the network. This layer is designed to be fast converging, highly reliable, and stable. This layer aggregates the distribution

switches and often integrates connectivity to the DMZ in this solution architecture, or other locations or the enterprise WAN when applied in other contexts. Also designed with Layer 3 protocols, the core provides load balancing, fast convergence, and scalability.

This three-layer design provides high availability with redundant hardware, redundant software features, redundant network connections/paths, and automatic procedures for reconfiguring network paths when failures occur. The highly available campus network architecture recommends no single points of failure on critical links or systems, and failover should be transparent to users and the applications operating on the network.

In addition to the high availability switching and routing network, the enterprise campus architecture incorporates the following three core networking functions:

- Network security based on the Cisco Self-Defending Network
- IP-based communications
- Mobility and wireless LAN services

For more information on the enterprise campus network, see the following URL:

http://www.cisco.com/en/US/netsol/ns340/ns394/ns431/networking_solutions_packages_list.html.

Cell/Area Zone

Overview

The cell/area network is the major building block of the EtoF architecture. This is the network that connects devices, controllers, and any other devices that need to communicate in real-time (I/O communication) with these devices. This section outlines the key technical considerations that need to be addressed in the test environment and design guidelines. In addition to the specific technical considerations for a cell/area network, the design guidelines also provide guidance on the design of a cell/area zone. The approach generally recommended is to design the cell/area by functional role of the devices and systems, rather than a design by device type, which is a common consideration. The design guidelines must indicate how customers can arrive at a cell/area network design schema that meets the performance requirements of the industrial automation and control systems.

An important distinction in this solution architecture is that the cell/area zone is considered a Layer 2 or LAN network. Layer 3 or IP routing features and functions are considered as part of the manufacturing zone (also including level 3). Even so, a Layer 3 switch or router is an important component of the cell/area zone because it performs critical Layer 2 network roles, such as Spanning Tree Protocol (STP) root bridge and Internet Group Management Protocol (IGMP) querier.

A cell/area zone network is very different from a typical IT access layer network (the term that IT may use for this network). The difference is related to the following:

- Environment in which it operates—Production facilities operate in conditions required by the manufacturing process taking place. This process can lead to extended temperature, humidity, invasive materials, shock, vibration, and multiple types of noise. The equipment operating in these environments must be designed and tested for these conditions.
- Industrial automation and control devices and applications—Relatively “dumb” to sophisticated Level 1 devices talking to a level 2 controllers and level 3 workstations and HMIs.
- Industrial automation and control systems also are very demanding of the cell/area network. Level 0 devices can be very simple devices with limited software and processing capabilities, which makes them susceptible to network-related disruptions or extraneous communication. In addition, a very

quickly changing manufacturing process (for example, a paper mill), or complex automation (for example, multi-axis robot) demand very high levels of determinism in the industrial automation and control system. These then require real-time communication from the network infrastructure.

The combination of the demanding environmental conditions of the manufacturing process and the industrial automation and control systems drive particular requirements of the cell/area network design. In summary, key design considerations are as follows:

- **Environment**—The conditions of the factory floor must be taken into consideration, because the equipment must be able to perform in these conditions. The network topology must be shaped to fit appropriately into the factory floor environment.
- **Real-time communications and network performance**—A cell/area network must be designed to meet the latency and jitter requirements of the industrial automation and control system it supports. This can impact the size of the LAN, the number of hops, the VLAN configuration, and a number of other network parameters.
- **Availability**—The availability of the cell/area network is directly attributable to the uptime of the manufacturing process it supports. The network must also be able to recover from network impacting events (for example, connection break) faster than the cycle time of the industrial automation control system to avoid the system automatically shutting down. Availability impacts the network design, topology, and even the type of equipment used.
- **Manageability**—The factory floor is usually not supported in the same manner as an IT network. First, the factory floor maintenance personnel tend not to have the networking expertise to perform anything beyond the most basic tasks. The setup and maintenance of network equipment and configuration must be simplified to meet the expertise level of the production floor maintenance personnel.
- **Security**—The factory floor tends to be protected physically from attack or exposure, but is the most sensitive area in that the devices are highly susceptible to network-borne attacks (for example, denial of service). Other sections of this architecture provide various forms of insulation for the cell/area zone, but certain precautions in the cell/area zone as well can significantly improve security.
- **Unmanaged versus managed**—The network infrastructure may not represent a large proportion of the factory floor (implementation or maintenance), but the same cost reduction mentality is applied as to other aspects of the production facility. In fact, because of a lack of understanding of the qualities of a managed, intelligent network, the additional costs they represent may lead customers to choose less intelligent solutions based purely on cost considerations; only later do they determine that the cheaper, unmanaged infrastructure cannot scale, perform, integrate, or be as easily maintained as an intelligent, managed network.

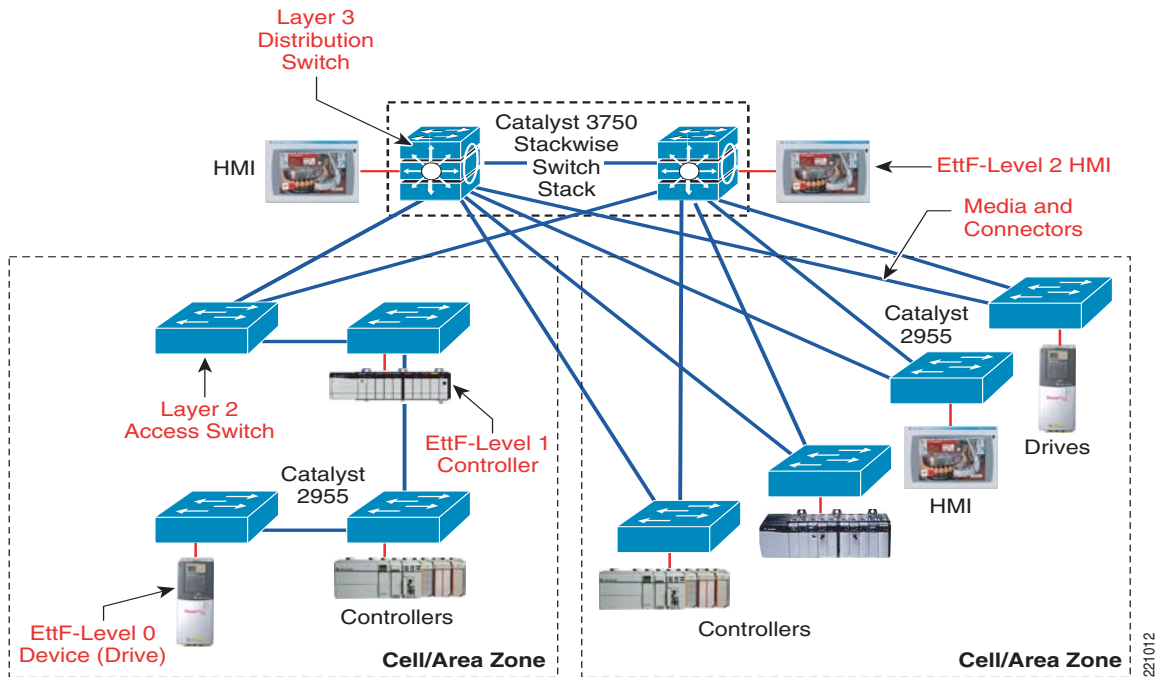
All these factors directly impact the components, topology, traffic flow, and network design, each of which is explored in the following sections.

Components

A cell/area zone comprises the following (see [Figure 2-4](#)):

- EttF Levels 0, 1, and 2 components; for example, devices, controllers, and HMIs
- Layer 2 access switches
- Layer 3 distribution switches or routers
- Media to connect all of the above

Figure 2-4 Cell/Area Components



This document does not provide guidance about the selection or implementation of the actual industrial automation and control equipment or the media used to connect the devices and switches. The equipment included in the test lab that was used to validate the overall solution is listed.

The key considerations customers go through when choosing network infrastructure include the following:

- Cost—For example, should I use expensive managed or cheap, unmanaged switches?
- Environment—Does the switch meet the environmental conditions in which the equipment must operate?
- Availability—How critical is the process being supported by the cell/area network to overall production? What level of operation is the cell/area network expected to operate? What is the cost of downtime?
- Flexibility—What variations of power, number of ports, type of media connections, mounting, and so on, does the switch support to meet the variety of situations in the production environment?
- Manageability—Can the device be easily maintained? What support and warranty options are available? Often, industrial automation and control systems can be operational for 5+ years, even into decades.
- Security—What security capabilities does the switch provide?
- Support—What type of support is available? What are the warranty options available?

Unmanaged versus Managed Switches

There is a significant distinction in the network infrastructure between intelligent, managed switches and unmanaged switches. Unmanaged switches require minimal or no configuration, but they do not support advanced features such as multicast management, port mirroring, security, diagnostics, or quality of service (QoS).

This design recommends the use of industrialized, managed, intelligent switches in all parts of the network infrastructure. Although unmanaged switches may initially meet the objectives of small, un-integrated networks, customers will eventually want to integrate and manage those networks, and without the capabilities of a managed, intelligent switch, will not be able to do so. [Table 2-2](#) shows the advantages and disadvantages of managed and unmanaged switches.

Table 2-2 *Managed and Unmanaged Switch Comparison*

	Advantages	Disadvantages
Managed switches	<ul style="list-style-type: none"> • Ability to manage multicast traffic • Provide diagnostics data • Provide security options • Provide other advanced features (see next section) 	<ul style="list-style-type: none"> • More expensive • Require some level support and configuration to start up and replace
Unmanaged switches	<ul style="list-style-type: none"> • Inexpensive • Simple to set up • “No config” replacement 	<ul style="list-style-type: none"> • No management capabilities • No security • No diagnostic information provided • Difficult to troubleshoot

Physicality and Environmental Considerations

Critical to cell/area levels are the environmental conditions in which the network infrastructure operates. Important considerations include the following:

- Extended temperature ranges supported
- Humidity tolerance
- Shock resistance
- Noise immunity
- Ingress protection or IP ratings defining the level of protection from physical intrusion

Often, network equipment may reside encased in a cabinet on the factory floor, which may reduce some of the environmental considerations, but this is not always the case.

Cisco works with external suppliers that supply casings to meet various needs.

Real-time Communications

A switch plays a key role in real-time communications. Key considerations for a switch performance include the following:

- Bandwidth supported on both access ports (typically up to 100 Mbps) and uplink ports (typically up to 1 Gbps).

- VLAN support; segmenting traffic flows is a key capability to achieve performance. VLANs play a key role in achieving this kind of segmentation.
- QoS support at both the Ethernet/CoS and IP/ToS layers.
- Multicast management features (for example, IGMP snooping).

Availability

The switch impacts overall availability of the industrial automation and control system because the switch is often a single point of failure if devices are connected only to a single switch. Thus, availability considerations are important and include the following:

- Passive cooling or no moving parts (for example, fans)
- 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 convergence protocols, such as STP and Rapid STP (RSTP)

Flexibility

The flexibility of the industrial Ethernet network is also a consideration. To efficiently support an industrial automation and control system, the network infrastructure should come in variations that include the following:

- Multiple port configurations
- Connections supported, such as fiber, small form-factor pluggables (SFPs), and copper/RJ45
- Power support—AC/DC in wide varieties as well as potential for redundant power
- Mounting support

Manageability

The manageability of the network infrastructure is also important. The switch is typically maintained by factory floor operations personnel who may have minimal network expertise. 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.
- Smartport configurations—Smartports allow pre-defined port configurations to be used that ease configuration of a switch.
- Ease of installation, setup, and maintenance. The network infrastructure should be easy to install, set up, and maintain with key basic functions supported by the key applications production floor user applications; that is, the industrial automation and control applications.
- Warranty and support.
- CIP support—The ability for the switch to interface with the industrial automation and control systems for some basic functions greatly eases ongoing maintenance.

Security

The Layer 2 access switch can play an important role in security as a port of entry to the manufacturing and cell/area zones. Some key considerations include the following:

- Access control lists (ACLs) to configure security policies into a switch
- Secure Shell (SSH) switch OS access
- SNMPv3 support for encryption of this important protocol for managing and monitoring the network infrastructure
- MAC address notification
- Port Security via MAC address identification

Component Summary


Table 2-3 lists the EttF testing lab component selections for the cell/area networks.

Table 2-3 Cell/Area Network Components

Role	Product/Platform	Software Release	Comments
Layer 2 access switch	Cisco Catalyst 2955 T-12, 12 10/100 ports, and two fixed 10/100/1000BASE-T uplink ports Catalyst 2960 for non-industrial environments	12.1(22)EA6	Connects EttF-Level 0-2 devices to the network The only industrial Ethernet switch Cisco currently offers For more details, see http://www.cisco.com/go/2955
Layer 3 distribution switch	<ul style="list-style-type: none"> • 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 	12.2(35)SE1	Provides inter-connection to cell/area zones. In cell/area VLANs, performs some LAN roles; for example, in STP root bridge and IGMP querier. The price/performance and port density of this switch has already made it a dominant choice for this role in existing EttF implementations.

The Cisco Catalyst 2955 switch is selected because it is the only industrial switch currently in the Cisco portfolio. (See [Figure 2-5](#).)

Figure 2-5 Cisco Catalyst 2955



If environmental requirements allow commercial grade switches, the key alternative to the Catalyst 2955 is the Catalyst 2960 (for details, see <http://www.cisco.com/en/US/products/ps6406/index.html>).

[Figure 2-6](#) shows the Cisco Catalyst 3750.

Figure 2-6 Cisco Catalyst 3750



The Catalyst 3750 Layer 3 switch was chosen rather than the Catalyst 4500 switch for the following considerations:

- Lower cost base
- StackWise feature allows scalability and redundancy
- Already established in this role at a number of customer accounts

Cisco recommends the Catalyst 4500 switch in the following situations:

- Capacity or scalability is a concern; for example, when integrating a large number of cell/area networks
- Upgradeable processor and interfaces for longer-term viability
- Better failover features for availability; for example, in-service upgradeability

The components in consideration for this phase of the solution architecture are connected via single connections to the network infrastructure. This is common for the industrial automation and control systems applying the CIP protocol. Although controllers may and often do have more than one Ethernet connection, they are typically not working in a high-availability configuration where one card assumes the IP address of the other in the event of failure. Dual-connected for high availability cell/area devices are not considered in this solution architecture at this point in time.

Traffic Flows

Traffic flow in a cell/area network is largely determined by the design and implementation of the industrial automation and control systems. These systems produce very different traffic patterns than the client-server and Internet-based applications in the IT domain. For example, 80–90 percent of the cell/area traffic is local as compared to a typical IT LAN in which perhaps less than 10 percent of the traffic is local. This is primarily driven by the cyclical I/O data being communicated on very short intervals (milliseconds) from devices to controllers and workstations/HMIs all on the same LAN or VLAN.

Cisco recommends that customers understand the traffic flows that a network should support so as to design the network to meet these various requirements, especially in relation to performance and real-time communications. Customers have choices (for example, network segmentation) that can change the network traffic flows and thereby significantly impact network performance and the real-time nature of the network.

Key considerations when designing traffic flows include the following:

- Current EtherNet/IP implementations have traditionally had the limitation that multicast traffic cannot be routed because the time-to-live field in the IP packet is set to 1. Although the recently released CIP EtherNet/IP specifications (CIP Specifications version 1.3, Volume 2 EtherNet/IP Adaptation of CIP, December 2006) call for this limit to be removed, these design and implementation guides are based on the limitation because the routing of multicast traffic requires a more complex set of protocols and considerations to be applied.
- The use of multicast for implicit traffic is a vendor choice, and is the prevalent choice. The most recent version of the Rockwell Controller program application (RSLogix) allows customers to choose multicast rather than unicast delivery for certain types of data. These design and implementation guidelines are based on the mode where all producer-generated I/O is multicast. Devices and controllers that talk to each other need to be in the same cell/area.



Note However, Cisco recommends that customers consider and apply this feature to maintain the size of the cell/area networks. This suggests choosing to use unicast delivery for PAC-to-PAC data/tags shared between two PACs where using multicast forces the cell/area network to be larger.

- Cisco recommends that customers perform an exercise to estimate the amount of CIP/IO traffic that will be generated on the network. Based on this, an estimate of the network bandwidth required to ensure the cell/area network size is sufficiently small to avoid oversubscribing connections or overwhelming the network infrastructure, and to maintain the real-time communication requirements.
- There is also traffic generated by the various network protocols (ARP, SNMP, RSTP, IGMP, and so on) that is not shown here. Properly configured, this is a minimal amount of the overall traffic. In an IT network, this is referred to as “control” traffic.

Figure 2-7 shows various cell/area zone traffic flows.

Figure 2-7 Cell/Area Zone Traffic Flows

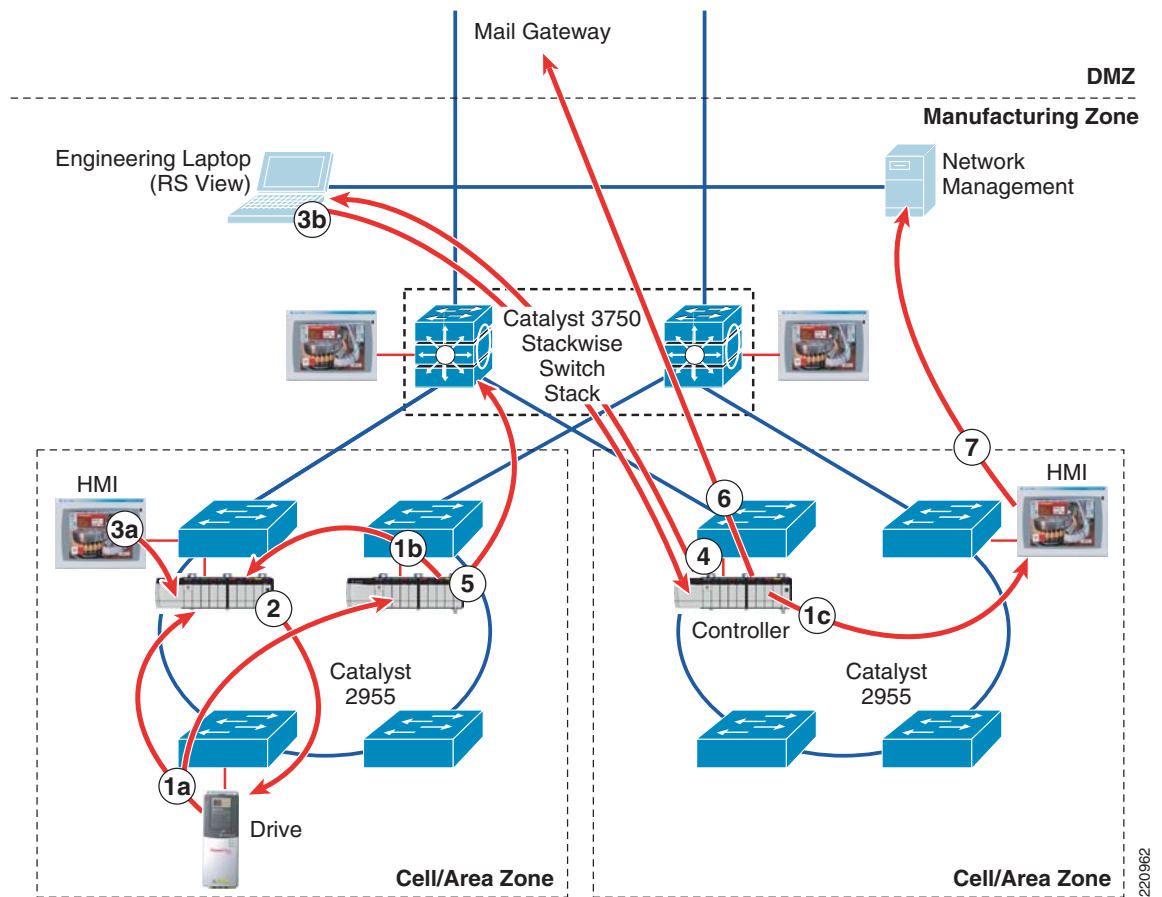


Table 2-4 describes the traffic flows shown in Figure 2-7.

Table 2-4 Cell/Area Zone Traffic Flows

Ref. #	From	To	Description	Protocol	Type	Port
1 a,b,c	Producer (for example, device)	Consumer (for example, PAC)	A producer (for example, device, PAC, or HMI) communicates data via CIP Implicit/IO (UDP multicast) traffic to multiple consumers, also known as input. <ul style="list-style-type: none"> a—Represents device to controller IO b—Represents PAC–PAC IO c—Represents controller reporting real-time status to HMI 	EIP	UDP	2222
2	Consumer	Producer	Consumer (for example, PAC or HMI) responds with output data or a heartbeat via CIP IO (UDP unicast) traffic to the producer.	EIP	UDP	2222

Table 2-4 Cell/Area Zone Traffic Flows (continued)

3a, b	Device	Device	CIP diagnostic, configuration, information, uploads/downloads, and identification data. For example, an HMI wants to open a CIP-connection with a controller. The CIP-connection request is communicated via TCP. Not shown, but the PAC responds with a TCP message. <ul style="list-style-type: none"> a—HMI opens a CIP connection b—Engineering workstation downloads a program 	EIP	TCP/UDP	44818
4	Device	Workstation/laptop	Most EtherNet/IP devices can provide diagnostic and monitoring information via web browsers (HTTP)	HTTP	TCP	80
5	Device	DHCP/BootP server	Devices or clients at startup for IP address allocation	DHCP/BootP	UDP	67-88
6	Controller	Mail server	Mail messages as warnings or for informational status	SMTP	TCP	25
7	Device	Network manager	All network infrastructure (for example, switches and routers) and many Ethernet devices can send SNMP messages	SNMP	UDP	161

Topology Options Overview

The cell/area network is where the various physical topologies are required to meet physical constraints of the factory floor. The network infrastructure (cabling, switches, and so on) must fit into the layout of the manufacturing process. A large variety of network topologies must be considered. This document considers the redundant star, ring, and trunk-drop.



Note

This document provides no specific design and implementation guidance for other topologies that may be supported as well, such as tree and hub-and-spoke.

Topology starts with considering how devices are connected to the network. In many industrial automation and control systems, the devices themselves support only single network connections, and therefore are connected via only a single connection to a single access switch. Where availability is critical and the devices support multiple connections, they should be connected to multiple switches to avoid single points of failure. In those cases, the network infrastructure should be configured in such a way to support the resiliency/redundancy of the overall manufacturing process.

Key considerations include the following:

- Physical layout—The layout of the manufacturing environment is a key driver of topology design. For example, a long conveyor belt system does not easily lend itself to a redundant star configuration, but rather a trunk-drop or ring.

- **Availability**—*Cisco recommends using resilient network topologies (for example, redundant star and ring) over non-redundant topologies.* These allow the network to continue to function after an event such as connection loss or switch failure. Although some of these events may still lead to downtime of the industrial automation and control systems, a resilient network topology may reduce that chance and should improve the recovery time.
- **Real-time communications**—Latency and jitter are impacted by a large variety of factors, but primary by the amount of traffic and number of hops a packet must make to reach its destination. The amount of traffic in a Layer 2 network is driven by various factors, but the number of nodes is important. Key guidelines include the following:
 - Amount of latency introduced per Layer 2 hop.
 - Bandwidth should not consistently exceed 50 percent of the interface capacity on any switch.
 - CPU should not consistently exceed 50–70 percent utilization. Above this level, the chances increase significantly that the switch may not properly process control packets and start behaving abnormally.

The key connectivity considerations made for the test environment include the following:

- Devices are connected to a switch via a single network connection or an IP-enabled I/O block or linking device if they do not support Ethernet natively. Note that most devices (including PACs) have limited or no failover capabilities and therefore cannot effectively use redundant network connections.
- Redundant network connections were not considered for this phase. Redundant connections may be used in certain industries and applications; mostly process-related industries applied to critical infrastructure.
- Switches may be arranged in a star, redundant star, trunk-drop, or ring network

Part of the validation phase is to generate guidelines for the size of a cell/area network and the limits of determinism that can be achieved as the cell/area network increases. The cell/area network in the test environment contains up to 15 switches, in the configurations shown in the following sections.


Cell/Area Network—Trunk-Drop (Daisy Chain) Topology

A trunk-drop topology is where switches are connected to each other to form a chain of switches. Key characteristics include the following:

- The connection between the Layer 3 switch and the first Layer 2 switch is a natural bottleneck and more susceptible to oversubscription, which can degrade network performance.
- Simple, easy-to-implement configuration.
- Minimal amount of cabling required to establish a network.
- No resiliency to loss of a connection.
- High level of flexibility for factory floor layout.

Figure 2-8 shows the trunk-drop (daisy chain) topology for the cell/area network.

Figure 2-8 Cell/Area Network—Trunk-Drop Topology



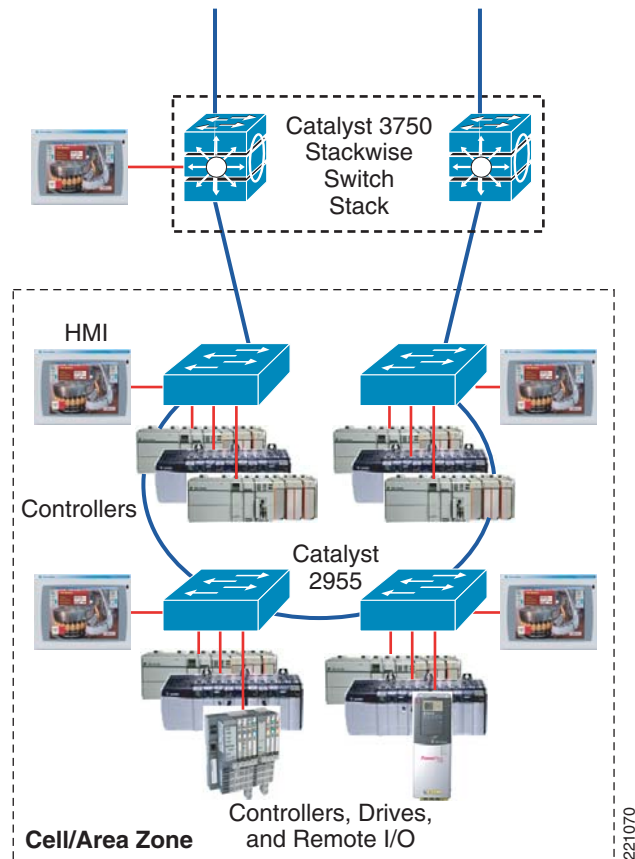
Cell/Area Network—Ring Topology

A ring topology is similar to a trunk-drop topology except that the last switch in the chain is connected to the Layer 3 switch, which forms a network ring. In a ring, if a connection is lost, each switch maintains connectivity to the other switches. Key considerations of the ring topology include the following:

- Additional cable connection to close the loop.
- Minimal level of network resiliency in that the network can recover from the loss of a single connection.
- More difficult to implement because it requires additional protocol implementation and Rapid Spanning Tree.
- High level of flexibility for the factory floor layout.
- Although better than the trunk-drop, the top of the ring (connections to the Layer 3 switches) can become a bottleneck and is susceptible to oversubscription, which can degrade network performance.

Figure 2-9 shows the ring topology for the cell/area network.

Figure 2-9 Cell/Area Network—Ring Topology



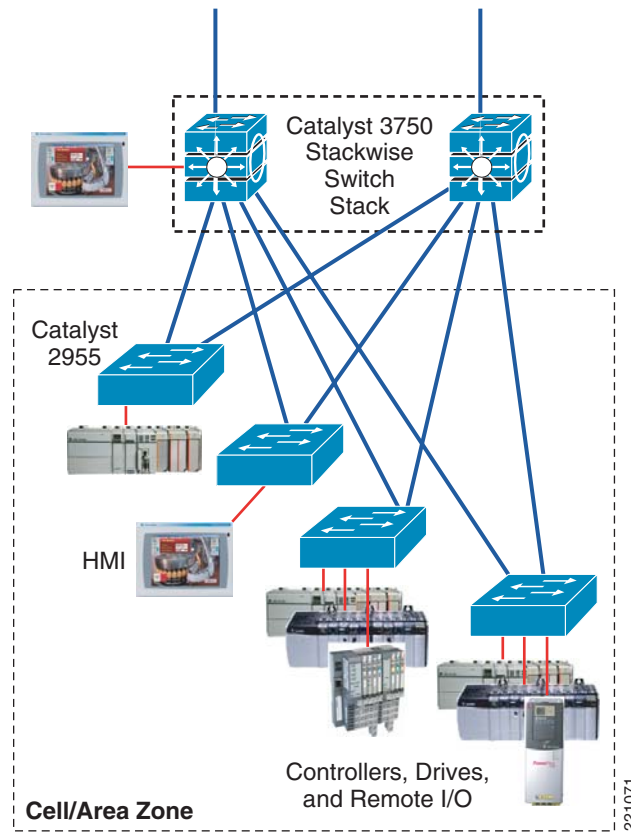
Cell/Area Network—Star Topology

A redundant star topology is essentially where every Layer 2 access switch has dual connections to a Layer 3 distribution switch. Devices are connected to the Layer 2 switches. This topology has the following advantages:

- Always only two hops from another Layer 2 switch
- No natural bottlenecks in the Layer 2 network because each switch has dual connections to the Layer 3 devices
- Layer 2 network is maintained even if multiple connections are lost
- Most complex cabling infrastructure required to establish dual connectivity of each switch to the Layer 3 switch
- More complex RSTP implementation

Figure 2-10 shows the star topology for the cell/area network.

Figure 2-10 Cell/Area Network—Star Topology



Cell/Area Topology Comparison

Cisco recommends that customers plan, design, and implement network topologies based on the redundant star configuration. This topology provides maximum network performance and availability. A redundant star provides protection against multiple connection failures and the quickest recovery in the case of such a failure.

However, customers implement other topologies because of production floor limitations and the complexity of the redundant star. Therefore, [Table 2-5](#) provides design and implementation guidance for the various topologies.

Table 2-5 *Cell/Area Topology—Advantages and Disadvantages*

Type	Advantages	Disadvantages
Redundant star	<ul style="list-style-type: none"> • Resiliency from multiple connection failures • Faster convergence to connection loss • Consistent number of hops (typically two in a flat design) provides predictable and consistent performance and real-time characteristics • Fewer bottlenecks in the design reduces chances of segment over-subscription 	<ul style="list-style-type: none"> • Additional wiring (and relevant costs) required to connect Layer 2 access switches directly to a Layer 3 distribution switch • Additional configuration complexity (for example, Spanning Tree with multiple blocks)
Ring	<ul style="list-style-type: none"> • Resiliency from loss of one network connection • Less cabling complexity in certain production floor layouts • Multiple paths reduces potential for oversubscription and bottlenecks 	<ul style="list-style-type: none"> • Additional configuration complexity (for example, Spanning Tree with a single block) • Longer convergence times • Variable number of hops makes designing predictable performance more complex
Trunk-drop	<ul style="list-style-type: none"> • Easy to design, configure, and implement • Least amount of cabling (and associated cost) 	<ul style="list-style-type: none"> • Loss of network service in case of connection failure (no resiliency) • Creates bottlenecks on the links closest to Layer 3 device, and varying number of hops make it more difficult to produce reliable performance.

Table 2-6 provides information to help decide which topology is appropriate based on key customer concerns.

Table 2-6 *Cell/Area Topology—Advantages and Disadvantages*

Key Concerns	Topology to Use
<ul style="list-style-type: none"> • Highly available network with minimal convergence • High performance network with minimal bottlenecks • Straightforward network design 	Redundant star
<ul style="list-style-type: none"> • Cabling complexity is a major concern • Highly available network is important • Cost is important 	Ring
<ul style="list-style-type: none"> • Cost and simplicity over availability and performance 	Trunk-drop

Network Design Overview

The sections above outline the following key requirements for a network design:

- Endpoints connected to a network
- Flow of information between the various endpoints
- Topology of the network (where everything is located)

This section outlines the key technical considerations in designing a cell/area network, including the following:

- Logical segmentation
- Availability
- Multicast management
- Traffic management via quality of service (QoS)
- Security

For a basic overview of networking concepts, see the following URL:

http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_book09186a00807594e5.html.

Logical Segmentation

Logical segmentation is the process of outlining which endpoints need to be in the same LAN. Segmentation is a key consideration for a cell/area network. Segmentation is important to help manage the real-time communication properties of the network, and yet support the requirements as defined by the network traffic flows. Security is also an important consideration in making segmentation decisions. A security policy may call for limiting access of factory floor personnel (such as a vendor or contractor) to certain areas of the production floor (such as a functional area). Segmenting these areas into distinct VLANs greatly assists in the application of these types of security considerations.

In fact, there are competing objectives. On one hand, by the assumptions made for this design guide, all Level 0–2 devices that need to communicate multicast I/O between each other must be in the same LAN. On the other hand, the smaller the VLAN, the easier it is to manage and maintain real-time communications. Real-time communications are harder to maintain as the number of switches, devices, and the amount of network traffic increase in a LAN.

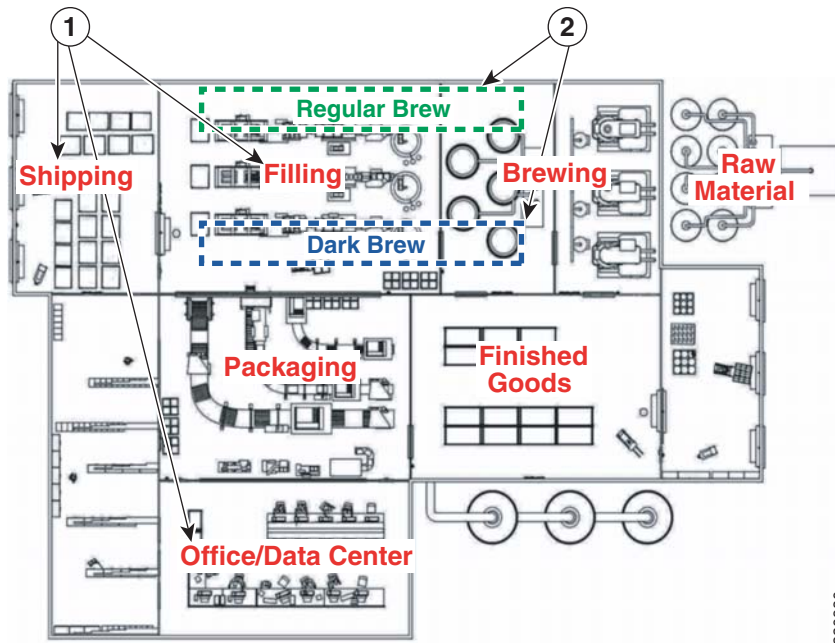


Note

Cisco therefore recommends that customers strive to design smaller LANs or VLANs, while recognizing that the traffic patterns of industrial automation and control systems may make this difficult.

There are many approaches to segmenting a network. Production facility networks can be divided by functional sections of the factory floor (see #1 in [Figure 2-11](#)), product lines (see #2 in [Figure 2-11](#)), and traffic type (for example, I/O, PAC-to-PAC, and explicit traffic). To achieve the goal of minimizing VLAN sizes, a mixture of all three may be used.

Figure 2-11 Sample Factory Floor—Brewing and Bottling




Segmentation can be achieved though via the following two key mechanisms in the cell/area network:

- Physical—Use of separate cabling and Layer 2 access switches to achieve segmentation
- VLAN (802.1Q)—Use of the VLAN protocol to achieve a VLAN that can be implemented on the same physical infrastructure

Physical segmentation is a very common approach in current Ethernet implementations, but is applied to an extreme. For example, a common approach in current Ethernet deployments is to physically separate I/O traffic from explicit traffic and not to connect the I/O traffic to any interconnected Layer 3 distribution switch. In these cases, a PAC has separate network connections to each network, and the only means to communicate between the two networks is over the backbone of the PAC. The I/O network is therefore reachable only via the PAC backplane that processes only CIP traffic. (See [Figure 2-12.](#))

Figure 2-12 CIP Gateway Configuration Example



The effects of this include the following:

- Devices on subnet #2 are not accessible via non-CIP protocols (such as SNMP or HTTP), limiting overall interconnectivity.
- PAC was not designed to router/switch traffic, and may introduce significant delays when used in this manner.

- Network-based services (such as security, management, IP address allocation, and so on) must either be replicated in each network or are not available.
- Increased costs occur because the available network resources in subnet #1 (for example, open ports) are not available in subnet #2.

Although physical segmentation dedicates network resources to these various traffic types and helps increase the level of certainty that the traffic receives sufficient network resources, Cisco recommends that these networks be at least connected to Layer 3 switches so as to enable interconnectivity via other methods than the PAC. Additionally, Cisco recommends that customers consider other ways (for example, application of QoS) to ensure that critical network traffic (such as implicit I/O) receives appropriate network performance.

**Note**

Cisco recommends the use of VLANs in addition to any physical segmentation, and that all cell/area LANs be connected to Layer 3 distribution switches to maintain connectivity.

VLANs offer the following features:

- Broadcast control—Just as switches isolate collision domains for attached hosts and forward only appropriate traffic out a particular port, VLANs refine this concept further and provide complete isolation between VLANs. A VLAN is a bridging domain, and all broadcast and multicast traffic is contained within it.
- Security—VLANs provide security in two ways:
 - High-security users can be grouped into a VLAN, possibly on the same physical segment, and no users outside of that VLAN can communicate with them.
 - Because VLANs are logical groups that behave like physically separate entities, inter-VLAN communication is achieved through a router. When inter-VLAN communication occurs through a router, all the security and filtering functionality that routers traditionally provide can be used because routers are able to look at Layer 3 information. In the case of non-routable protocols, there can be no inter-VLAN communication. All communication must occur within the same VLAN.
- Performance—The logical grouping of users allows, for example, an engineer making intensive use of a networked CAD/CAM station or testing a multicast application to be assigned to a VLAN that contains just that engineer and the servers he or she needs. The work of this engineer does not affect the rest of the engineering group, which results in improved performance for the engineer (by being on a dedicated LAN) and improved performance for the rest of the engineering group (whose communications are not slowed down by the single engineer using the network).
- Network management—The logical grouping of users, divorced from their physical or geographic locations, allows easier network management. It is no longer necessary to pull cables to move a user from one network to another. Adds, moves, and changes are achieved by configuring a port into the appropriate VLAN. Expensive, time-consuming recabling to extend connectivity in a switched LAN environment is no longer necessary because network management can be used to logically assign a user from one VLAN to another.

For more background information on VLANs, see the following:

- VLANs and VLAN trunking—
http://www.cisco.com/en/US/partner/tech/tk389/tk689/tsd_technology_support_protocol_home.html
- LAN switching and VLANs—
http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_chapter09186a008075983b.html

- Internetwork design guide—LAN switching—
http://www.cisco.com/en/US/partner/tech/tk1330/technologies_design_guide_chapter09186a008066670e.html

Any end device to be connected to multiple VLANs typically requires multiple network interface cards (NICs) available to the device. For example, controllers can have multiple NIC cards installed because of their modularity, and therefore have direct Layer 2 access to multiple VLANs. This may also be a consideration in the segmentation of the network.

Availability

Depending on the topology selected, various availability options can be designed into the network. If a topology is chosen with resiliency (for example, redundant star or ring), some form of network protocol is required to eliminate loops in the network. Loops are created when Layer 2 network devices are connected with multiple paths to reach the same destination. If left unmanaged, loops can cause serious network issues by creating broadcasts storms (messages continuously passed around the network) that eventually disrupt network service. Both standard and proprietary protocols have been developed to manage the loops and to react to connection losses by maintaining a consistent network.

The protocols identify (either manually or automatically) one or more connections to be virtually turned off to eliminate loops. When a connection is lost, the protocols must recognize the disruption and re-activate a closed connection to restore network viability. The speed at which a network protocol recognizes the disruption, opens a closed connection to restore network interconnectivity, and resumes normal network services is called the convergence time. During the convergence time, some portion of the traffic is dropped by the network because inter-connectivity does not exist. If the convergence time is longer than the cycle time in the industrial automation and control system, the systems on the affected portion of the network may begin to stop operating and bring parts of the factory floor to a halt. Thus, production and control engineers may decide that the additional cost and complexity of a resilient network may not provide sufficient value because they cannot recover in sufficient time to avoid disruption.



Note

Although network convergence may not be fast enough to ensure industrial automation and control system uptime, Cisco recommends the use of resilient network topologies because they allow the manufacturing operations to continue when industrial automation and control systems are re-started without waiting on lost connections to be restored.

As mentioned, there are standard and proprietary protocols to manage resilient network topologies. The standard protocols are based on STP, which implements the 802.1D IEEE algorithm by exchanging Bridge Protocol Data Unit (BPDU) messages with other switches to detect loops, and then removes the loop by shutting down selected bridge interfaces. This algorithm guarantees that there is only one active path between two network devices. RSTP, based on IEEE 802.1w, is an evolution of the STP 802.1D standard and provides for faster spanning tree convergence after a topology change. The standard also includes features equivalent to Cisco PortFast, UplinkFast, and BackboneFast for faster network reconvergence.

For resilient network technologies and customers who want to either implement multi-vendor environments or to rely on standard technologies, Cisco recommends using RSTP in the network.

For all-Cisco environments, Cisco recommends the use of Per Vlan Spanning Tree Plus (PVST+).

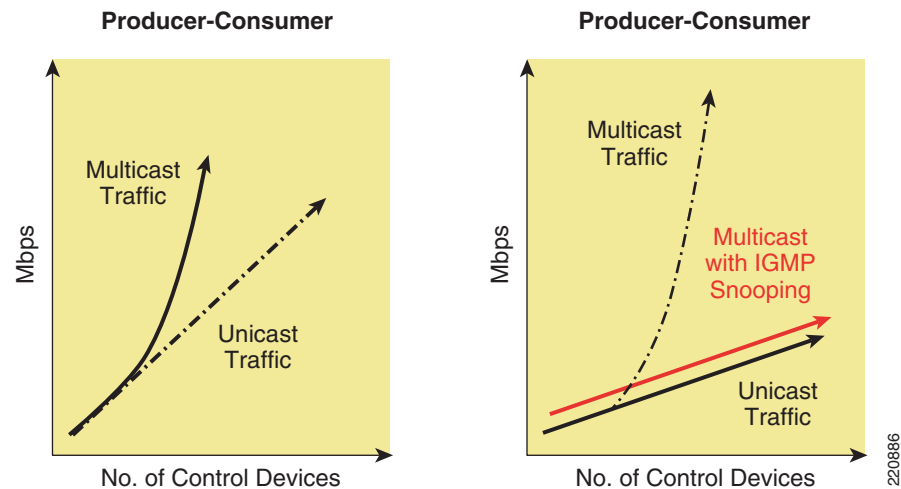
For more information on Spanning Tree Protocol and related technologies, see the Spanning Tree Protocol Introduction at the following URL:

http://www.cisco.com/en/US/partner/tech/tk389/tk621/tsd_technology_support_protocol_home.html

Multicast Management

Multicast traffic is an important consideration of a cell/area network because it is used by some of the key industrial automation and control communication protocols, such as CIP. Unmanaged, multicast traffic is treated by the network infrastructure as a Layer 2 broadcast; every endpoint on the network receives the message. The impact increases exponentially as more multicast producing endpoints are added to the LAN. Internet Group Management Protocol (IGMP) is the standard method to manage multicast traffic. IGMP enables the network infrastructure to understand which endpoints are interested in which multicast data, and thus to forward the messages only to those endpoints that want them. This reduces the amount of traffic the network and endpoints must handle. (See [Figure 2-13.](#))

Figure 2-13 IGMP Impact of Network Traffic



Note

Cisco recommends that the network infrastructure be configured to manage multicast traffic. Ethernet switches should be configured to perform IGMP snooping. When IGMP snooping is enabled, the switch listens to IGMP traffic and develops a table that lists the multicast groups and the end devices. Thus, when a multicast packet is received, the switch forwards it only to end devices that want it. In addition, the Layer 3 distribution switch where the LAN is connected should be configured to perform the IGMP Querier function.

Although the number of multicast addresses in a VLAN or subnet is not typically an issue, it is a consideration under certain scenarios. EtherNet/IP devices can support up to 32 multicast addresses. Typically, however, an EtherNet/IP device uses a single multicast address. PACs can potentially use more for doing peer communications, but that may also be alleviated by choosing unicast messaging (an option in recent RA firmware updates). This is important because the network infrastructure has limits on the number of multicast addresses that can be supported. For example, the Catalyst 2955 can handle up to 256 multicast address. Unless steps are taken to use more multicast address than normal, these limits do not come into play. It is theoretically possible to configure a VLAN with nine or more EtherNet/IP devices to overrun the number of multicast addresses that the Catalyst 2955 switches can handle. This can be avoided using standard EtherNet/IP configuration guidelines. The impact of overrunning the switch multicast address space is that the multicast messages are treated as broadcast, introducing more traffic than necessary.

Although in this version of the solution architecture, multicast packets do not mix with enterprise or IT traffic via DMZ segmentation, if they did, there is the distinct potential of redundant use of multicast group addresses that would lead to disruptions in both the industrial automation and control system and

the relevant IT application. For this and many other reasons, this solution architecture recommends a demilitarized zone (DMZ) between the manufacturing and enterprise zone to ensure that industrial automation and control multicast traffic and IT-driven multicast traffic remain separate.

For information on IP multicasting, visit Cisco Technology Support at the following URLs:

- http://www.cisco.com/en/US/partner/tech/tk828/technologies_white_paper09186a0080092942.shtml
- http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_chapter09186a00807598c3.html
- http://www.cisco.com/en/US/partner/tech/tk828/tsd_technology_support_protocol_home.html

For more information on configuring and implementing IGMP, see Chapter 4, “Implementation of the Cell/Area Zone .”

Quality of Service

Quality of service (QoS) refers to control mechanisms that can provide various priorities to different users or data flows, or guarantee a certain level of performance to a data flow in accordance with requests from the application program. QoS guarantees are important if the network performance is critical, especially for real-time industrial automation and control systems.



Note

Cisco recommends the application of QoS to the critical CIP-based implicit I/O as well as the explicit traffic generated by industrial automation and control systems.

As seen in [Traffic Flows, page 2-17](#), non-CIP traffic (such as Internet HTTP) is very likely on any industrial automation and control network. The industrial automation and control devices can support various traffic types natively, and certain functions (for example, monitoring) are implemented using common protocols. As well, EttF level 3 workstations and servers in the manufacturing zone produce traffic of various types that may mix with the cell/area traffic. In addition, in future versions of this architecture, Cisco believes that manufacturers will want to take advantage of the openness that standard networks provide to introduce other services into the manufacturing zone without sacrificing the performance required by the industrial automation and control systems. QoS will be a key mechanism to achieve this goal.

Unfortunately, at the time of this writing, the ODVA is still working on a set of QoS specifications for EtherNet/IP traffic. Cisco therefore has decided not to add specific design and implementation guidelines at this time. Design and implementation guidance based on the ODVA QoS specifications will be produced in later phases of this solution architecture. Until such time as these specifications are designed and tested, Cisco recommends that detailed network design and analysis be completed with the introduction of QoS in an industrial Ethernet network.

For more information on QoS, see the following URLs:

- Quality of Service—
http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_chapter09186a0080759886.html
- Implementing QoS—
http://www.cisco.com/en/US/partner/tech/tk543/tk757/technologies_white_paper09186a008017f93b.shtml

- Cisco IOS QoS—
http://www.cisco.com/en/US/partner/tech/tk389/tk813/technologies_white_paper0900aecd802b68b1.shtml
- Understanding QoS—
http://www.cisco.com/en/US/partner/tech/tk543/tk762/technologies_white_paper09186a00800b0828.shtml#intro

Security

An overall security approach is presented in [Chapter 5, “Implementation of Security.”](#) The security design considerations for the cell/area network include the following:

- Port configuration (including MAC filtering, limited ACL configurations, QoS Trust)
- Infrastructure protection (hardening from a security perspective) of the network infrastructure to prevent unauthorized access
- Layer 2 security—Networking services can be disrupted through attacks on the protocols and standards to which they adhere. Layer 2 security protects the network services from attacks on the key protocols, including the following:
 - Quality of service (QoS)
 - Address Resolution Protocol (ARP)
 - Rapid Spanning Tree Protocol (RSTP)
 - Dynamic Host Configuration Protocol (DHCP)
 - MAC address flooding
- Monitoring of network infrastructure administration data (syslogs and SNMP)
- VLANs for isolation

Manufacturing Zone

Overview

The primary function of the manufacturing zone is to separate the services and applications that are critical to operating the factory floor from the rest of the enterprise, most importantly the industrial automation and control systems. The manufacturing network represents all the critical industrial automation and control systems required to operate the factory floor. These are the systems and applications insulated from the enterprise zone by the DMZ. This insulation not only provides security to both the enterprise and manufacturing zones, but may also represent an organization boundary where IT and manufacturing organizational responsibilities interface. By definition, the cell/area networks described above are part of the manufacturing zone. In fact, most manufacturing zones contain many cell/area zones. Key functions and features of the EttF architecture for the manufacturing network include the following:

- Interconnecting the various cell/area networks
- Integrating the level 3 site manufacturing systems
- Providing network management and security services to the Level 0–3 systems and devices
- Endpoint protection

Note that Cisco recommends that customers carefully consider which applications and services are considered part of the manufacturing zone. Key points to consider include the following:

- 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 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 2-7 lists some of the key applications and services to consider.

Table 2-7 Key Applications and Services

Type	Critical	Optional
Manufacturing applications	<ul style="list-style-type: none"> • Historian • Asset management and security • Production floor monitoring and reporting • Industrial automation and control system management and maintenance 	<ul style="list-style-type: none"> • Manufacturing execution system
Network and security management	<ul style="list-style-type: none"> • Network management • Security management • Security monitoring, analysis, and response 	
Common network-based Services	<ul style="list-style-type: none"> • 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—Most industrial automation and control systems do not use dynamic names and use instead hard-coded IP addresses. If dynamic names are used, a DNS service is required and is likely separate from the IT services. • Network Time Protocol (NTP) servers are required to coordinate clocks in various industrial automation and control systems, including to network infrastructure. 	<ul style="list-style-type: none"> • 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.

For the purpose of this solution architecture, Cisco placed the applications and systems listed above in the manufacturing zone.

After the set of applications and services for the manufacturing zone are determined, the design of the network can begin. The following are key considerations of the manufacturing zone that are applied in the following sections:

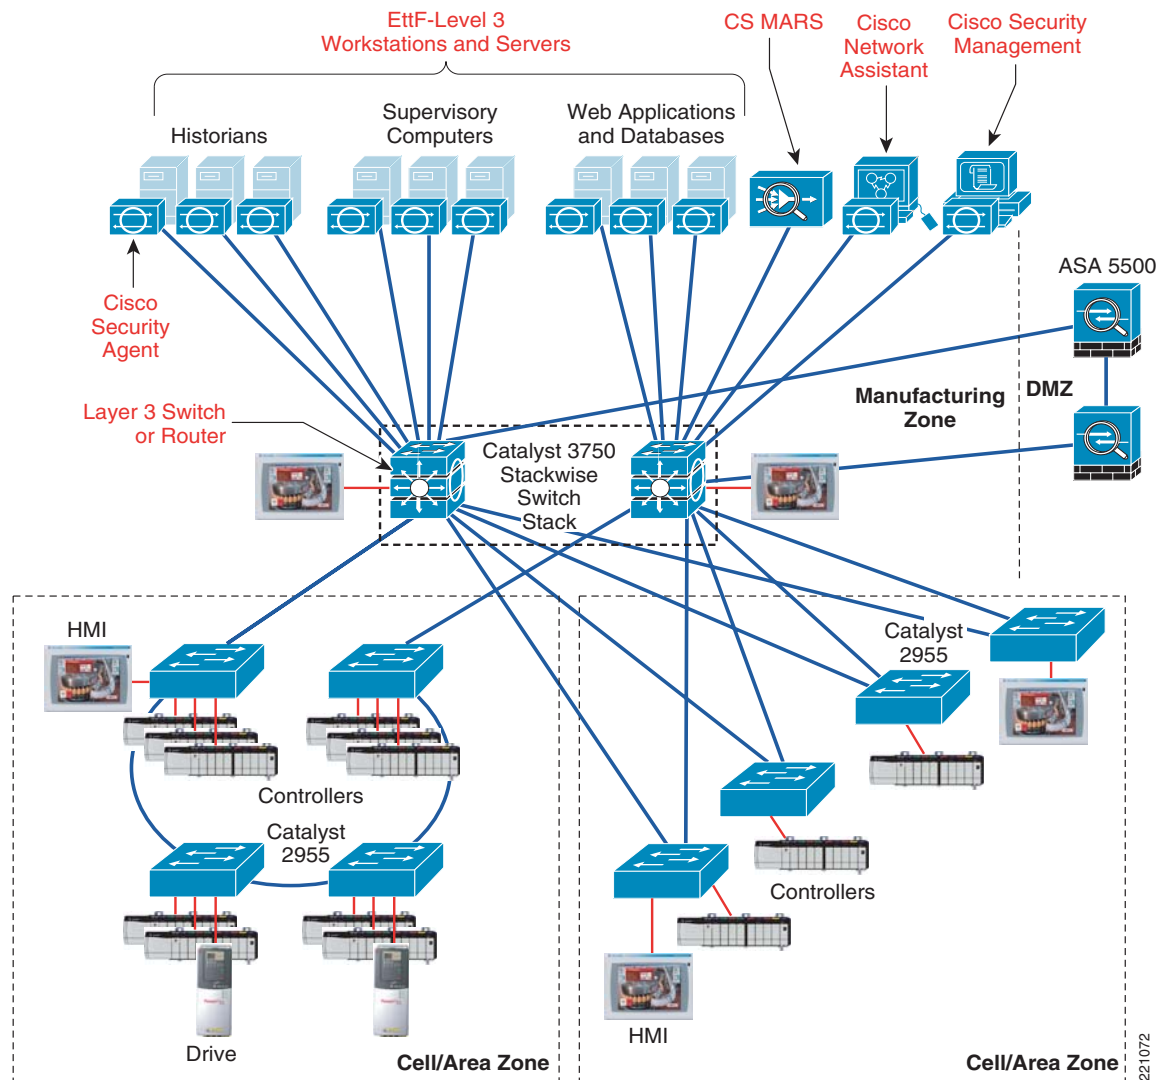
- **Environmental**—Most production facilities have controlled areas for certain types of applications and IT-related infrastructure. The EttF level 3 applications and systems of the manufacturing zone typically reside in these environments. This suggests that the environmental requirements of the cell/area network typically do not apply to the manufacturing zone network infrastructure. An exception exists where the distribution devices (Layer 3 switches or routers) may potentially need to reside closer to the cell/area networks and therefore meet certain levels of extended environmental tolerance.
- **Real-time communications**—Because the systems and applications are more similar to IT, there are not the same real-time communications considerations of the manufacturing zone as in 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. Essentially, latency and jitter can vary more widely without significant disruption to the applications and devices in this zone.
- **Availability**—Availability of the network services is critical. Although the applications and services in the manufacturing zone may be more tolerant than more variable real-time communications, it is crucial that they stay available to maintain the operations in the cell/area zone. Without the services of the manufacturing zone, the production floor would soon stop.
- **Manageability**—The systems and applications in the manufacturing zone are typically administered and maintained by people with a focus on production floor operations, not IT or network. Although more standard technologies will be applied to manage the network resources, they need to be easy to implement and use.
- **Security**—As more of the applications and systems are based on standard computing resources (workstations, servers, and common operating systems), they will be susceptible to the wider range of security risks for this equipment. Additionally, security functions will be applied in this zone to provide service to cell/area zones and other EttF level 3 components and applications.

Components

The manufacturing zone consists of the following (see [Figure 2-14](#)):

- EttF level 3 applications, servers, and workstations
- Depending on size and complexity, Layer 2 access switches to connect the EttF 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)

Figure 2-14 Manufacturing Zone Overview



This document does not provide guidance about the selection design or the implementation of the actual EttF level 3 industrial automation and control equipment and equipment, or the media used to connect the devices and switches. The equipment included in the test lab that was used to validate the overall solution is listed.

Cisco included components for the following in this 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)

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

Cost

Although cost is always a consideration in production facilities, the applications and devices in this zone will tend not to be replicated as often as, for example, the Layer 2 switches found in cell/area zones. Therefore, there is not the similar managed versus unmanaged question as in the cell/area zone; managed equipment is used by default.

Physicality and Environmental Considerations

As stated above, the environmental considerations for this zone are less critical because it is assumed that controlled environments will exist for the equipment.

It is recognized, however, that there is a need in some production floor environments for the Layer 3 switching/routing functions to exist in an industrial packaging and to operate in the same conditions. Cisco does not currently offer these capabilities in an industrial packaging.

Performance and Real-time Communications

Although not quite as critical as the cell/area zone, it is critical for the network infrastructure to support real-time communications functions. The critical implicit and explicit traffic may traverse the manufacturing zone 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 support
- QoS support
- Load balancing protocols supported (for example, Gateway Load Balancing Protocol)
- Multicast routing protocols supported

Endpoint security has also been included in this solution architecture. As such, consider the overall network and system performance the application has on the system on which it is running.

Regarding Cisco security management, analysis, and response (MARS), a key performance factor is the number of events and amount of network traffic that a device can support.

Regarding network management, the key performance criteria are the number and types of network infrastructure devices that can be supported.

Availability

The network infrastructure availability is directly related to overall availability of the industrial automation and control system. Thus, availability considerations are important and include the following:

- Availability options available (for example, Hot Standby Router Protocol) and 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

Manageability

Network and security management services and endpoint security are a part of this zone. These applications must be relatively easy to install, configure, and operate. They must relax the level of expertise and knowledge required by the production floor personnel to maintain and operate the production floor 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 production and control engineers during problem resolution
- Ability to securely allow access to appropriate operational or maintenance personnel

In addition to the actual network and security management applications, there also are 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 installation, setup and maintenance—The network infrastructure should be easy to install, set up, and maintain with key basic functions supported by the key applications production floor user applications; that is, the industrial automation and control applications.
- 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.
- CIP support—The ability for the equipment or application to interface with the industrial automation and control systems for basic management and monitoring functions greatly eases overall use and ongoing maintenance.

Security

The Layer 2 access switch can play an important role in security as a port of entry to the manufacturing and cell/area zones. Some key considerations include the following:

- Access control lists (ACLs) allow users to configure security policies into a switch.
- Secure Shell (SSH) switch OS access.
- SNMPv3 support for encryption of this important protocol used to manage and monitor the network infrastructure.

Component Summary

For the purpose of testing, the products listed in [Table 2-8](#) were part of the manufacturing zone.

Table 2-8 **Components**

Role	Product/Platform	Software Release	Comments
Distribution switch	Cisco Catalyst 3750 Series <ul style="list-style-type: none"> • 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 	12.2(35)SE1	Provide redundant distribution and core routing functions to cell/area and manufacturing zone traffic
Core router	Catalyst 3750 Series (see above) Catalyst Series 4500: Chassis: <ul style="list-style-type: none"> • Catalyst 4503, 4506 (without Supervisor Engine redundancy capability) • Catalyst 4507R or Catalyst 4510R (with Supervisor Engine redundancy capability) Supervisor: Supervisor IV or up Line cards (The selection depends on the uplink type of the C2955 model of choice): <ul style="list-style-type: none"> • Cisco Catalyst 4500 48-port 1000BASE-LX (SFP) • Cisco Catalyst 4500 24-port 10/100/1000 Module (RJ-45) • Cisco Catalyst 4500 48-port 10/100/1000 Module (RJ-45) • Cisco Catalyst 4500 Enhanced 48-port 10/100/1000 Module (RJ-45) 	12.2(31)SGA	Optional in medium-to-large operations to provide core networking functions
Security monitoring, analysis, and response	Cisco Security MARS 20R	4.2.3	Monitors security events from switches, routers, firewalls, and endpoint agents
Endpoint protection	CSA	5.1	Security protection for standard OS devices

Table 2-8 Components (continued)

Firewall configuration and management	Cisco Adaptive Security Device Manager	5.2	Firewall and intrusion protection services. Manages traffic flows between manufacturing, DMZ, and enterprise zones.
Endpoint security management	Cisco Security Agent Management Console	5.1	Manages endpoint security agent configuration
Network management	Cisco Network Assistant	5	Performs basic network management

Switching and Routing

The Catalyst 3750 switch was selected because it provides the best mix of features, performance, and cost for small-to-medium production facilities. Key considerations included the following:

- Lower cost base
- Already established in this role at a number of customer accounts
- 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 production facility by adding additional stackable units

For more information, see the following URL:

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

Figure 2-15 shows the Cisco Catalyst 4500 switches.

Figure 2-15 Cisco Catalyst 4500 Switches



An option that was strongly considered and that is believed to be a good option for larger production facilities is the Catalyst 4500, for the following reasons:

- Capacity or scalability is a concern; for example, when integrating a large number of cell/area networks and EttF level 3 workstations and servers
- Upgradeable processor and interfaces for longer-term viability
- Better failover features for availability; for example, in-service upgradeability

For more information, see the following URL:

<http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html>

For large and extremely large manufacturing environments, the Catalyst 6500 should also be considered, but was not considered for this version of the architecture. For more information, see the following URL: <http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html>

Figure 2-16 shows the CS-MARS.

Figure 2-16 CS MARS



Security Monitoring, Analysis, and Response

The entry-level Cisco Security MARS appliance was selected. A wide variety of appliances is available that support increasing levels of events and network flow. For more information, see the CS MARS product overview at the following URL:

<http://www.cisco.com/en/US/partner/products/ps6241/index.html>.

A wide variety of devices that support increasing levels of events are available.

Additionally, for customers interested in deploying in a number of manufacturing sites, global controller units are also available, although this version of the solution architecture does not cover this case.

Endpoint Security

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

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>.

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, 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 management
- Troubleshooting advice (Catalyst Express 500 Series)
- Inventory reports
- Event notification
- Network security settings (Catalyst Express 500 Series)
- Task-based menu
- File management
- Drag-and-drop Cisco IOS Software upgrades

For more information on CNA, see the following URLs:

- 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 is a portfolio of network management. For more information, see the following URL: <http://www.cisco.com/en/US/products/sw/netmgts/index.html>.

Security Management

Cisco recommends the deployment of the Cisco Adaptive Security Device Manager to manage the firewalls in the DMZ. Key features include the following:

- Intuitive, easy-to-use web-based management interface
- Intelligent wizards
- Robust administration tools
- Versatile monitoring services

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 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 pre-configured 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 customers that are interested in more comprehensive security management solutions, Cisco recommends considering 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—Manufacturing Zone

The traffic flows in a manufacturing zone resemble those of a de-centralized client-server environment. Many of the EttF level 3 workstations, applications, and servers do the following:

- Send detailed scheduling, execution, and control data to controllers in the various cell/area zones
- Collect information from the cell/area networks 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
 - Terminal services
 - File server
 - Domain and Lightweight Directory Access Protocol (LDAP) services
 - Network and security management
- Production reporting services (for example, cycle times, quality index, predictive maintenance) available to manufacturing zone and via the DMZ and 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 two perspectives:

- Industrial automation and control applications (for example, historian, asset management, automation and control security, reporting)
- Network and security management

As with the cell/area zone, traffic from the network infrastructure protocols (for example, ARP, STP, and so on) is not represented.

[Figure 2-17](#) and [Table 2-9](#) show the manufacturing zone traffic flow.

Figure 2-17 Manufacturing Zone Traffic Flow—Industrial Automation and Control Application

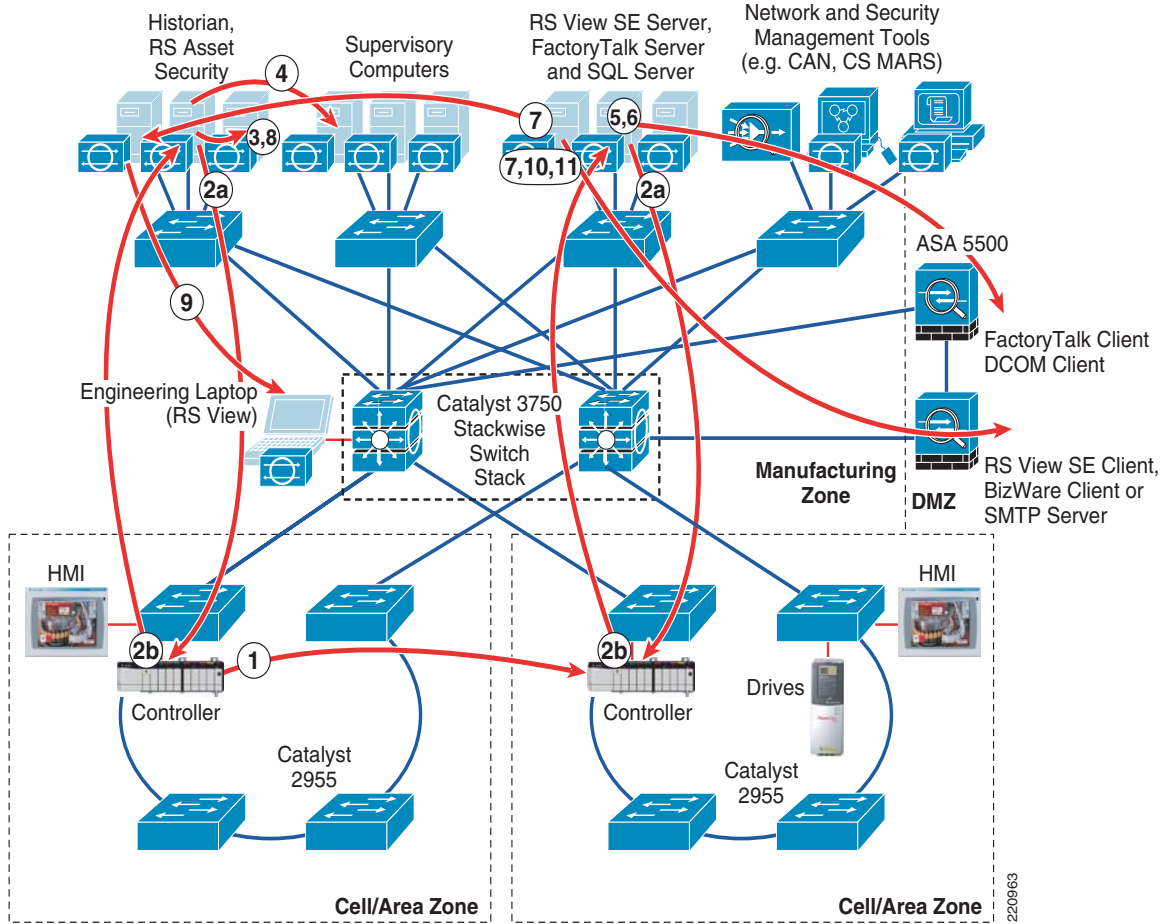


Table 2-9 Manufacturing Zone Traffic Flows

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	EIP	TCP/UDP	44818
2	Client/server	Client/server	FactoryTalk Transaction Manager	RPC	TCP	400–402
3	Client/server	Client/server	FactoryTalk Metrics—Production server	RPC	TCP	4120
			FactoryTalk Metrics—Server manager	RPC	TCP	4121
			FactoryTalk Metrics—PlantMetrics 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

Table 2-9 Manufacturing Zone Traffic Flows (continued)

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	ranserve rping	TCP	1332
			FactoryTalk—Directory server file transfer	rnadirft	TCP	3060
			FactoryTalk—Alarming server	rnaalarm ing	TCP	6543
			FactoryTalk—Event multiplexor		TCP	7600
			FactoryTalk—Event server		TCP	7700
			FactoryTalk—Directory server		TCP	7710
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 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 industrial automation and control 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

As the manufacturing zone is assumed to be installed in a controlled environment and not on the factory floor, the topology options are not going to be driven by the physical constraints of the factory floor, but rather the size of the environment to be supported. The various options range from a single pair of Layer 3 switches to a set of Layer 2 access switches and two (or more) sets of Layer 3 switches/routers for the large manufacturers by introducing separate access, distribution, and core networking services. Topology options include the following:

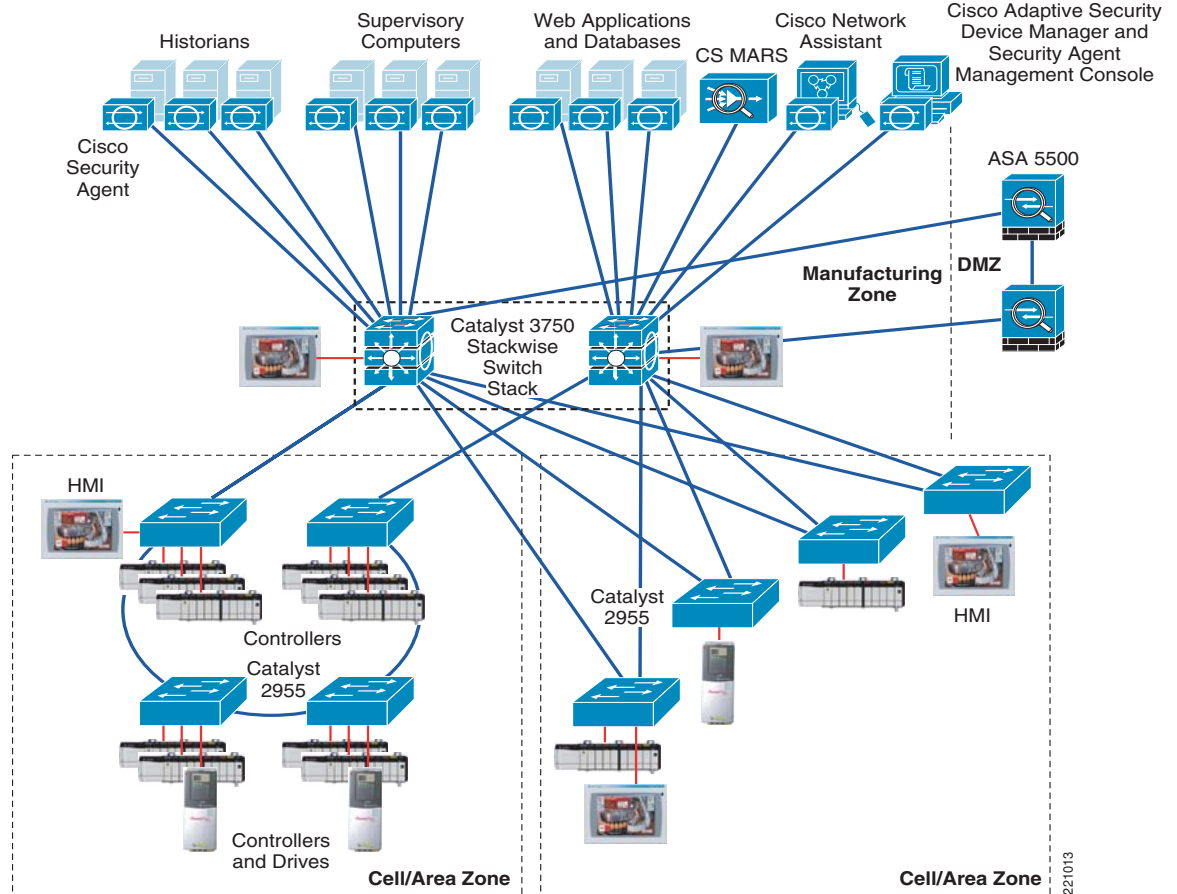
- Small manufacturing zone of up to 50 nodes
- Medium manufacturing zone of up to 200 nodes
- Large manufacturing zone of more than 200 nodes

Small Manufacturing Zone Topology

The small manufacturing zone includes a redundant pair of Layer 3 switches configured for redundancy (see [Figure 2-18](#)). All EttF level 3 devices are connected directly to the switches. A set of stacked 3750 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 manufacturer. In this version, the Layer 3 switches provide inter-VLAN and zonal routing functions as well as Layer 2 connectivity to EttF level 3 workstations and servers.

The small manufacturing zone essentially represents a collapsed core-distribution network routing services. This should suffice for small and many medium manufacturing facilities.

Figure 2-18 Small Manufacturing Zone Topology



Medium Manufacturing Zone Topology

The medium topology represents the separation of various network routing services and replication of these services to meet requirements in a larger production facility (see Figure 2-19). Although the small topology can easily support up to 200 Ethernet nodes, there are situations even in this type of node count that may require a more segmented topology. This topology differs from the small manufacturing topology as follows:

- Separate Layer 2 access switches to connect the EtoF level 3 workstations and servers
- Additional pair of Layer 3 distribution switches for geographically distributed cell/area zones
- Additional pair of Layer 3 core routers to consolidate the traffic

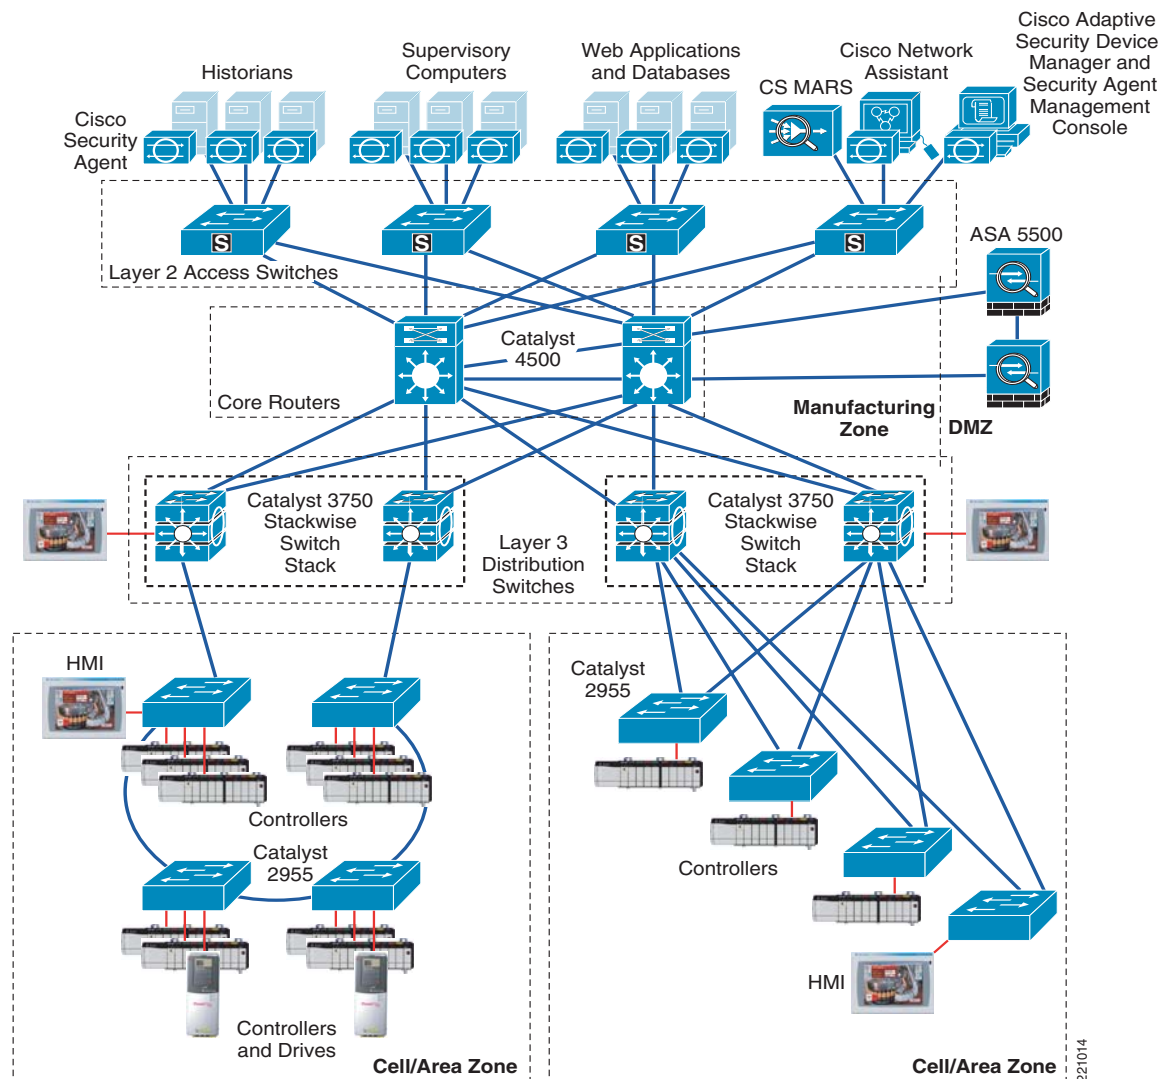
Not all three of these topology enhancements need be implemented; each one could be added. Following are some considerations for each of the scenarios:

- 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 recommends having a separate Layer 2 access switch for the configuration of the relevant protocols. For more information, see [Server Farm Access Layer](#), page 3-9.
- 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, EttF level 3 workstations and servers, and other cell/areas.

Figure 2-19 shows the resulting topology.

Figure 2-19 Medium Manufacturing Zone Topology



Manufacturing Zone Topology Summary

Cisco does not have a specific recommendation between the small and medium options presented. The customer requirements, in particular scalability, geographical dispersion, and availability requirements, determine which option to choose.

Note that the medium option represents separating out the access, distribution, and core networking functions into distinct equipment. In the small version, all three are supported by the pair of Layer 3 switches. It is also possible that only the access or core functions will be separated out, which produces more variations.

Network Design Overview

The sections above have outlined the following key requirements for a network design:

- Endpoints connected to a network
- Flow of information between the various endpoints
- Topology of the network (where everything is located)

The next sections outline the key technical considerations in designing a cell/area network, which include the following:

- Logical segmentation
- Availability
- Routing
- Traffic management via QoS
- Security
- Manageability

Logical Segmentation

Logical segmentation is important at this level, especially for the EtfF level 3 workstations and servers. In the cell/area zone, it is important for endpoints that communicate implicit IO 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 industrial automation and control applications (EtfF 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:

- Industrial automation and control systems dedicated to particular functions in the factory floor (for example, a brewing control room)
- 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 *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 Cisco solution architecture identifies a DMZ with firewall functions to safely and securely share data and services between these zones.

Availability

Because the cell/area inter-connect functionality exists in this 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:

- EttF level 3-Layer 2 connectivity
- Core routing and Layer 3 switching
- Network and systems management
- Endpoint security

Layer 2 Connectivity

The EttF level 3 workstations and servers are connected to LANs/VLANs. These VLANs also need to be designed with availability considerations. Cisco previously recommended that the redundant topology be applied; therefore, RSTP must be implemented in the Layer 2 networks 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.

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 2-10](#).

Table 2-10 *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.	<ul style="list-style-type: none"> • Catalyst 4500 • Catalyst 3750—Virtual with StackWise 	All
StackWise	Uses stack interconnect cables to create a virtual switch fabric for stacks of the Catalyst 3750.	<ul style="list-style-type: none"> • Catalyst 3750 • N/A to Other Platforms 	All
Redundant power supplies	Each system has dual power supplies so that the system operates normally upon failure of a power supply	<ul style="list-style-type: none"> • Catalyst 4500: Internal • Catalyst 3750: External • Catalyst 2955: External 	All
Redundant fans	Each fan tray has multiple fans	Catalyst 4500	All

Table 2-10 **Device Level Resiliency Design (continued)**

Line card online insert and removal (OIR)	New line cards can be added or removed without affecting the system or losing the configuration.	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 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 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 4500	
IOS Software Modularity	Minimizes unplanned time, If an error occurs in a modular process, the system determines the best recovery action. The recovery options include: <ul style="list-style-type: none"> Restart a modular process Switchover to standby supervisor Remove the system from the network Allows a modular process to be patched and restarted without any downtime.	Catalyst 6500	All
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 <ul style="list-style-type: none"> 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 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 control network environment are shown in [Table 2-11](#).

Table 2-11 Network Level Resiliency Design

Feature	Description	Supported Platforms	Where to Apply in Industrial Ethernet Network
Link redundancy—dual homing	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 L2 STP or a L3 routing protocol.	All routers and switches	All
Hot Standby Router Protocol (HSRP)	Cisco proprietary redundancy protocol for establishing a fault-tolerant default gateway.	<ul style="list-style-type: none"> Catalyst 4500 Catalyst 3750—Virtual with StackWise 	Wherever Layer 3 routing takes place
Incremental SPF Optimizations	Optimization of the OSPF algorithm to reduce computational load.	<ul style="list-style-type: none"> Catalyst 4500—Internal Catalyst 3750—External 	All
IP dampening	Mechanism to suppress affects of excessive state changes (flapping).	Catalyst 4500	All

Security and Network Management

The security and network management services are in the manufacturing zone for security considerations; they require access to critical network features. Therefore, they can be used to attack a system. The most secure location for these services is behind the 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 production 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 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 add a level of network resiliency.

Endpoint Security

CSA can be designed to operate in a managed or unmanaged mode. This solution architecture recommends a managed mode so as to better manage the security stance of the protected endpoints in the manufacturing zone. In the managed mode, CSA-protected endpoints communicate with the CSA Management Console (MC). If the CSA MC fails, CSA continues to operate (as if in unmanaged mode) until the service is re-established.

Routing

Routing is a key feature of the manufacturing zone. For more information on routing basics, see the following URL:

http://www.cisco.com/en/US/tech/tk1330/tsd_technology_support_technical_reference_chapter09186a008075970b.html

Manageability

The following are key considerations for the availability of the network and security management functions:

- Scripting (security issues, troubleshooting, complex setup and recovery)
- Network sniffer and protocol analyzer; for example, Wireshark (formerly known as Ethereal) for basic network troubleshooting

Demilitarized Zone

The enterprise zone and the manufacturing zone have different requirements, priorities, policies, and implications of incidents, but they should be able to share data and access systems and applications. The EttF solution architecture introduces the DMZ into the EttF architecture to provide logical segmentation between the enterprise and manufacturing 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 and operational 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.

The DMZ is a key aspect of the overall security approach for the EttF solution architecture. The DMZ is a strong form of logical segmentation between the manufacturing zones (and all that encompasses) and the enterprise zone (and via that the Internet). It is in the DMZ that firewall and intrusion protection solutions are introduced. These are applied to all traffic entering or exiting the DMZ, either to or from the manufacturing or enterprise zones. In addition, the concept of not allowing traffic to traverse the DMZ also adds the capability of setting clear access and authorization principles. For example, enterprise users can be forced to authenticate against manufacturing-specific security services to ensure that they have right to the services and data made available in the DMZ.

The DMZ is also a demarcation where different operational and security policies can be applied to meet objectives from various perspectives. For example, the DMZ demarks where QoS settings change. In addition, the DMZ also demarks where critical I/O traffic from the manufacturing zone stops and is not mixed with enterprise traffic. The DMZ can be used to apply different operational settings (for example, authorizations, configurations, monitoring, and so on) to allow different network operational models to exist between the manufacturing environment and the IT-managed enterprise. As a last resort, the DMZ is also a point where access can easily be shut off if issues or threats arise in a zone that threatens operations in other zones.

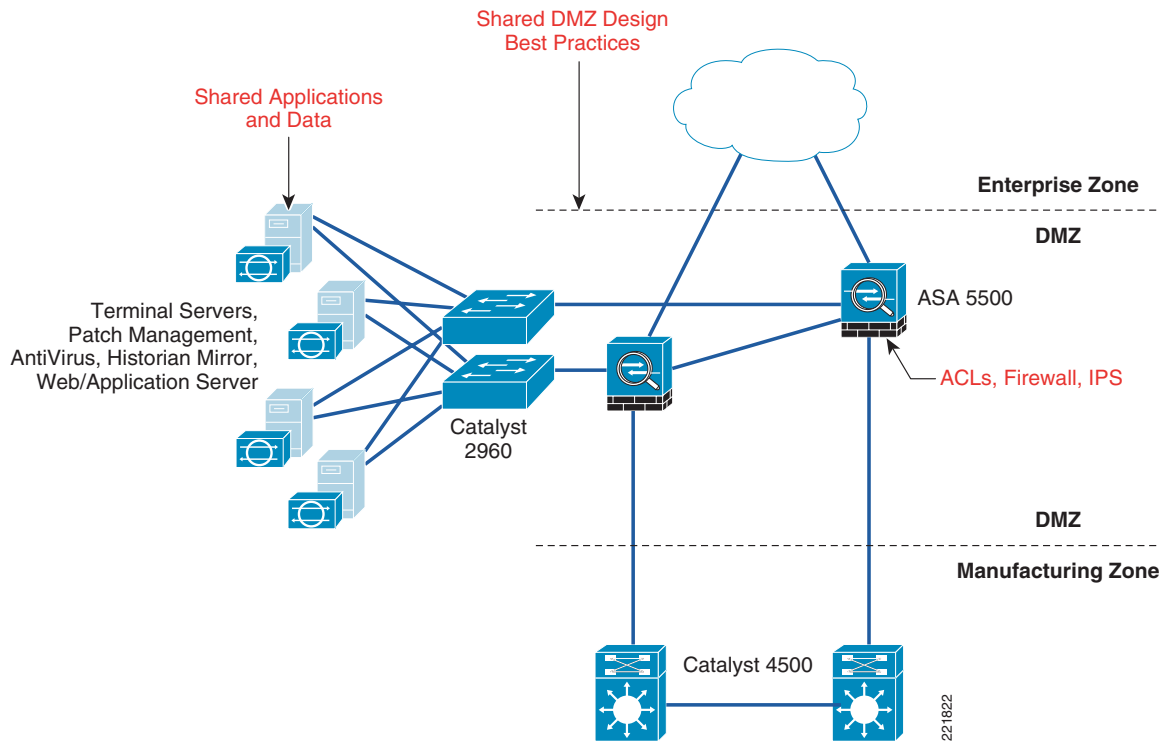
The DMZ is not a new concept, but is prevalent in enterprise networks where applications and data are shared with other enterprises or made available to the Internet. The concept and application of the DMZ between the enterprise and manufacturing zones is very similar to the DMZs applied at the Internet and enterprise interface. Those DMZs apply strong traffic control, in-depth packet inspection, and enforced authorization and authentication for privileged access.

The DMZ provides the following:

- Threat control and containment
- Area to safely and securely share applications and data
- Demarcation for organizational and responsibility reasons

Figure 2-20 shows the DMZ and related features.

Figure 2-20 DMZ and Related Features



Components

As with the network infrastructure in the manufacturing zone, the DMZ components are assumed to be housed in environmentally-controlled locations, where the stringent factory floor physical constraints do not come into play. All DMZ components are not assumed to be “industrialized” or “ruggedized”.

The key components of the DMZ are the firewalls. The firewalls act as routers/switches between the manufacturing and enterprise zone, and provide in-depth, stateful packet inspection for firewall and intrusion detection.

Figure 2-21 shows the Cisco ASA 5520 switch.

Figure 2-21 Cisco ASA 5520



In addition, the DMZ also contains switching and routing domains where the servers that share data and applications exist. Cisco recommends deployment of the Catalyst 2960. For more details regarding the Catalyst 2960 Series Switches, see the following URL:

<http://www.cisco.com/en/US/products/ps6406/index.html>.


Topology Options Overview

Only one topology was considered for this version of this guide. This topology includes the following:

- Single pair of firewalls, one operating in stand-by mode for high availability. Each firewall has a connection to the enterprise zone, manufacturing zone, and access switches of the DMZ.
- A pair of switches to provide Layer 2 access connectivity between the servers in the DMZ and the firewalls.

Figure 2-22 shows the DMZ topology.

Figure 2-22 DMZ Topology



A topology that was not tested as part of the development of this guide, but may be applicable, is dual-paired firewalls. The purpose of two pairs of firewalls is to highly segment operational control between the firewall configuration to the enterprise zone and the firewall configuration to the manufacturing zone.

Network Design Overview

Key design considerations for the DMZ include the following:

- Services and data that need to be shared

Careful thought and implementation must be applied to the services and data that need to be shared via the DMZ. The services and data in the DMZ that are intended to be shared between the manufacturing and enterprise zones must exist in the DMZ. For data, replication mechanisms need to be considered and applied if, for example, historical manufacturing data from the manufacturing zone must be made available to enterprise applications. Typical applications for the DMZ include the following:

- Patch management servers to transfer patches and upgrades destined for the manufacturing zone.
- Proxy or terminal servers to provide controlled and secured access to manufacturing zone applications.
- File servers as points to store, and server data or files for users or applications in the manufacturing or enterprise zones.
- In some cases, customers may choose to deploy “client” or “view” versions of their applications in the DMZ to allow access/view into manufacturing.

The applications and data to be shared directly impacts the following key services.

- Access and authentication to DMZ resources

As with any set of shared applications and data, designed and implemented access and authentication services are required for the DMZ. In particular, because the zone is shared, some specific DMZ version of access and authentication services may be required to maintain operational management objectives. The use of VLANs in the DMZ as well should be applied to segment access to applications and data in the DMZ. For example, a DMZ VLAN could be established for a particular set of partners/vendors who should have access to only a subset of applications/data in the DMZ (also translating to subset of resources in the manufacturing zone).

- Firewall configuration

Which applications and their representative ports must be allowed in which directions should be specified. The default firewall configuration should “not allow” traffic unless explicitly identified. In this case, analysis is required of what type of traffic is required to support access to DMZ data and services. In addition, careful design, testing, and implementation of the intrusion protection is required; in particular, choosing between inline and promiscuous modes (see [Chapter 5, “Implementation of Security.”](#))

- Availability

By definition, the DMZ servers and applications should not be “critical” to the operations of the manufacturing zone. Therefore, other availability priorities may apply. In particular, the DMZ design and implementation may need to take into consideration that the DMZ may be used to quickly and effectively separate the manufacturing zone from the enterprise zone if either is compromised and threatens the other zone.

- Operational management

The DMZ is a kind of demarcation between organizational and operational activities. Typically, IT sets the operational priorities and conditions for the enterprise zone. Although IT may be involved with operations in the manufacturing zone, the manufacturing organization will most likely also be involved and have different operational priorities. For example, the manufacturing organization may want to control who from the enterprise zone has access to data and applications in the DMZ (and further into the manufacturing zone). As mentioned before, it is also at the DMZ where various network configuration concepts are segmented, such as different QoS settings for the traffic to/from the enterprise zone than traffic to/from the manufacturing zone.

For more details on the implementation of the DMZ including firewalls, intrusion protection, and access/authorization, see [Chapter 5, “Implementation of Security.”](#)

Software Versions

[Table 2-12](#) provides list of the software versions related to the Cisco components that are supported by this version of the solution architecture.

Table 2-12 **Software Versions**

Product / Platform	Zone	Software Release
Catalyst 2955	Cell/area	12.1(22)EA6
Catalyst 3750G	Cell/area, manufacturing	12.2(35)SE1
Catalyst 4500	Manufacturing	12.2(31)SGA
Cisco ASA 5500	DMZ	7.2
Cisco Network Assistant	Manufacturing	5
CS-MARS	Manufacturing	4.2.3
Cisco Security Agent	Manufacturing	5.1
Cisco Adaptive Security Device Manager	Manufacturing	5.2
Cisco Security Agent Management Console	Manufacturing	5.1



CHAPTER 3

Basic Network Design

Overview

The main function of the manufacturing zone is to isolate critical services and applications that are important for the proper functioning of the production floor control systems from the enterprise network (or zone). This separation is usually achieved by a demilitarized zone (DMZ). The focus of this chapter is only on the manufacturing zone. This chapter provides some guidelines and best practices for IP addressing, and the selection of routing protocols based on the manufacturing zone topology and server farm access layer design. When designing the manufacturing zone network, Cisco recommends that future growth within the manufacturing zone should be taken into consideration for IP address allocation, dynamic routing, and building server farms.

Assumptions

This chapter has the following starting assumptions:

- Systems engineers and network engineers have IP addressing, subnetting, and basic routing knowledge.
- Systems engineers and network engineers have a basic understanding of how Cisco routers and switches work.

IP Addressing

An IP address is 32 bits in length and is divided into two parts. The first part covers the network portion of the address and the second part covers the host portion of the address. The host portion can be further partitioned (optionally) into a subnet and host address. A subnet address allows a network address to be divided into smaller networks.

Static IP Addressing

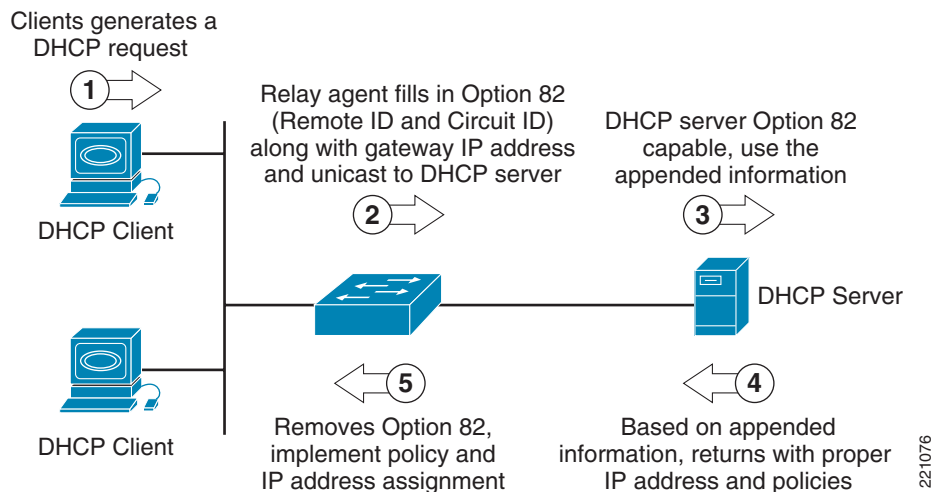
In the manufacturing zone, the level 3 workstations and servers are static. Additionally, it is recommended to statically configure level 2 and level 1 control devices. These servers send detailed scheduling, execution, and control data to controllers in the manufacturing zone, and collect data from the controllers for historical data and audit purposes. Cisco recommends manually assigning IP addresses to all the 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. In addition, Cisco recommends referencing devices by their IP address as opposed to their DNS name, to avoid potential latency delays if the DNS server goes down or has performance issues. DNS resolution delays are unacceptable at the control level.

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 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 basically extends the standard DHCP process by tagging the request with the information regarding the location of the requestor. (See Figure 3-1.)

Figure 3-1 DHCP Option 82 Operation



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

**Note**

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.

**Note**

DHCP and the DHCP option 82 feature have not been validated in the lab for EttF version 1.1. At this time, Cisco recommends considering only DHCP with option 82 for the application servers at level 3.

IP Addressing General Best Practices

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.

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 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. Because the control traffic is primarily confined to the cell/area zone itself, and never crosses the Internet, Cisco recommends using a private, non-Internet routable address scheme such as 10.x.y.z, where x is a particular site, y is a function, and z is the host address. These are guidelines that can be adjusted to meet the specific needs of a manufacturing operation. For more information on private IP addresses, see RFC 1918 at the following URL: <http://www.ietf.org/rfc/rfc1918.txt>.

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 all the host and device individual IP addresses that reside on that network into a single route. Route summarization is a way of having single IP address represent a collection of IP addresses, 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. Cisco recommends that automation systems in manufacturing should be hard-coded with a properly unique static IP address.

**Note**

Cisco 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 allowed in the enterprise zone because the EtherNet/IP devices in the manufacturing zone use an algorithm to choose a multicast address for their implicit traffic. Therefore, to avoid conflict with multicast addresses in the enterprise zone, multicast traffic in the manufacturing zone should not be mixed with 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.

Routing Protocols

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 EttF 1.1, routing begins at the manufacturing zone, or distribution layer. The Catalyst 3750 is responsible for routing traffic between cells (inter-VLAN), or into the core, or DMZ. No routing occurs in the cell/area zone itself.

Selection of a Routing Protocol

The correct routing protocol can be selected based on the characteristics described in the following sections.

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.

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 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 also lose the flexibility of choosing a smaller or larger subnet using VLSM. RIPv1 is 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 re-calculates 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:

- RIP uses hop count.
- 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.
- 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.

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. Factors such as convergence time and support for VLSM and CIDR directly impact the scalability of the routing protocols.

Table 3-1 shows a comparison of routing protocols.

Table 3-1 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

Table 3-1 Routing Protocols Comparison (continued)

Name	Type	Proprietary	Function	Updates	Metric	VLSM	Summarization
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 also requires VLSM for discontinuous major networks. The recommendation is to use OSPF or EIGRP as the core routing protocol in the manufacturing zone. For more information, see the Cisco IP routing information page at the following URL:

http://www.cisco.com/en/US/tech/tk365/tsd_technology_support_protocol_home.html

Static or Dynamic Routing

The role of a dynamic routing protocol in a network is to automatically detect and adapt changes to the network topology. The routing protocol basically 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 dynamic routing protocols should be used.

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. IP address pool management is done by the server.
 - Directory services—Set of applications that organizes and stores data 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, which in turn runs over IP. 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.



Note 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

- Security and network management servers
 - 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



Note 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—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.



Note 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

**Note**

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

- Manufacturing application servers—Consists of the following:
 - Historian
 - RS Asset Security Server
 - Supervisory computers
 - RSView SE Servers
 - RSLogic Server
 - Factory TalkServer
 - SQL Server

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.

Server Farm Access Layer

Access Layer Considerations

The access layer provides physical connectivity to the server farm. The applications residing on these servers for the manufacturing zone are considered to be business-critical and therefore necessary to be dual-homed to the access layer switches.

Layer 2 Access Model

In the Layer 2 access model, the access switch is connected to the aggregation layer through an IEEE 802.1Q trunk. The first point of Layer 3 processing is at the aggregation switch. There is no Layer 3 routing done in the access switch. The layer model provides significant flexibility by supporting VLAN instances through the entire set of access layer switches that are connected to the same aggregation layer. This allows new servers to be racked in anywhere and yet still reside in the particular subnet (VLAN) in which all other applications-related servers reside.

Spanning VLANs across Access Layer switches

If your applications require spanning VLANs across access layer switches and using STP as an integral part of your convergence plan, take the following steps to make the best of this suboptimal situation:


- Use Rapid PVST+ as the version of STP. When spanning tree convergence is required, Rapid PVST+ is superior to PVST+ or plain 802.1d.
- Provide an L2 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 STP primary on the same distribution layer switch. The HSRP and Rapid PVST+ root should be co-located 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/application/pdf/en/us/guest/netso/ns656/c649/cdccont_0900aecd804ab67d.pdf

Figure 3-2 shows an example of a 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. There are certain features such as private VLANs that allow groups of Layer 2 adjacent servers to be isolated from each other but still be in the same subnet. The requirement of Layer 2 adjacency is important for high availability clustering and NIC teaming.

NIC Teaming

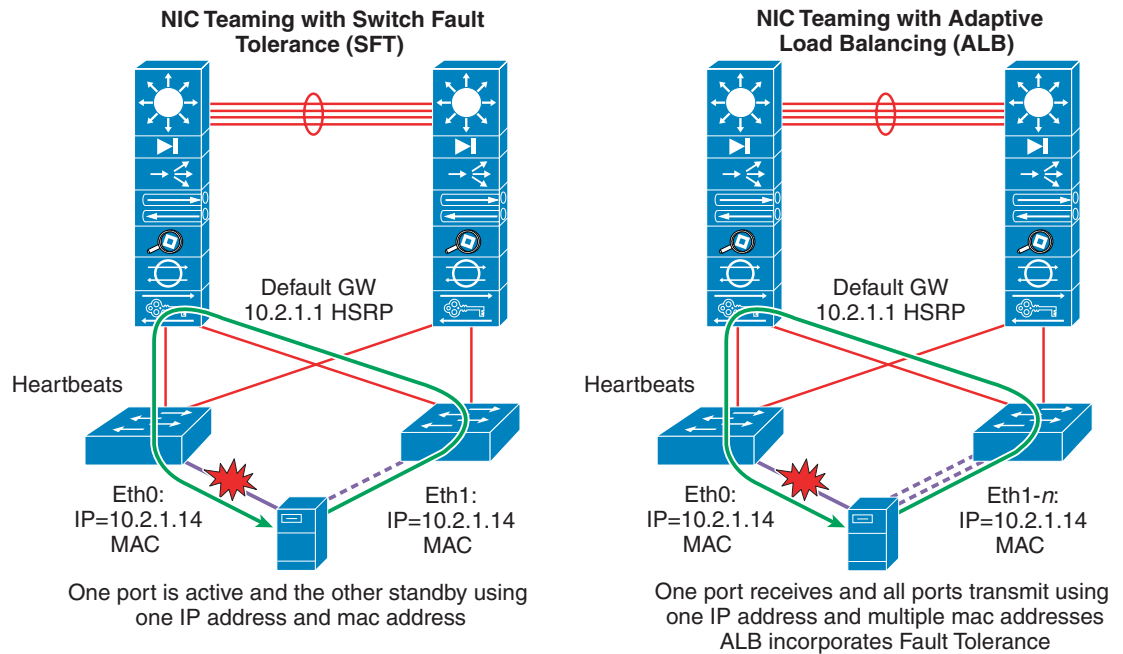
Mission-critical business applications cannot tolerate downtime. To eliminate server and switch single point of failure, servers are dual-homed to two different access switches, and use NIC teaming drivers and software for failover mechanism. If one NIC card fails, the secondary NIC card assumes the 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 3-3 shows examples of NIC teaming using SFT and ALB.

Figure 3-3 NIC Teaming



143396

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 1–3 seconds. These heartbeats are sent as a multicast or a broadcast packet and therefore require Layer 2 adjacency.



CHAPTER 4

Implementation of the Cell/Area Zone

This chapter outlines recommendations, best practices, and configurations for implementing a cell/area zone architecture in an EttF environment. The cell/area zone is where actual end nodes connect into the network, so careful planning must be done to achieve the optimal design from both the network and device perspective.

As mentioned in earlier chapters, EtherNet/IP is the enabling standard at this layer. Ethernet networks have been successfully used on the factory floor for the past 15 years, mainly in non-time-critical applications. Ethernet technology (more accurately IEEE standard 802.1 and IEEE standard 802.3 technologies) has evolved from a 10 Mbps, half-duplex, bus/tree topology into a 100 Mbps and 1 Gbps, full-duplex, switch/router-based hierarchical star topology. This evolution has created an opportunity for using Ethernet in industrial networks that support time-critical applications. EtherNet/IP is a communication system suitable for use in industrial environments and time-critical applications. EtherNet/IP uses standard Ethernet and TCP/IP technologies and an open Application Layer protocol called Control and Information Protocol (CIP). CIP is also used in ControlNet and DeviceNet networks. In EtherNet/IP networks, exchange of time-critical data is based on the producer/consumer model where a transmitting device (host or end node) produces data on the network, and many receiving devices can consume this data simultaneously. Implementation of the producer/consumer data exchange is based on the IP multicast service mapped over the Ethernet multicast service. EtherNet/IP-supported functions include the following:

- Time-critical data exchange
- Human-machine interface (HMI)
- Device configuration and programming
- Remote access to web pages embedded in EtherNet/IP devices
- Device and network diagnostics

The configuration details outlined below (for example, VLAN numbers, hostnames, port numbers, and so on) are merely examples and should be adjusted accordingly to a particular factory environment.

Cell/Area Zone Network Device Provisioning

Networking devices in the cell/area zone include Cisco Catalyst 2955s and the downlinks on the Catalyst 3750. The recommended Cisco IOS Software version at the time of this writing is C2955-12.1.22-EA9 (Crypto Image) and C3750-12.2.25-SEB4 (Advanced Crypto Image). The images can be downloaded from the following URL: <http://www.cisco.com/kobayashi/sw-center/index.shtml>.

Beginning with the Catalyst 2955s, the startup process is as follows:

-
- Step 1** Load the Cisco IOS image on all devices.
- Step 2** Configure all uplink 1 GE ports (gi0/1–2) as trunk ports carrying only one VLAN:
- ```
interface GigabitEthernet0/1
switchport trunk native vlan 20
switchport trunk allowed vlan 20
switchport mode trunk
end
```
- Step 3** Configure all FastEthernet (fa0/1–fa0/12) interfaces as switchport access ports for the particular VLAN carried on the uplink trunk:
- ```
!
interface FastEthernet0/1
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpduguard enable
ip dhcp snooping limit rate 10
end
```
- Step 4** On FastEthernet ports connected to an end device, manually configure speed and duplex settings to those supported by the end device:
- ```
cell-c2955-12#conf t
Enter configuration commands, one per line. End with CNTL/Z.
cell-c2955-12(config)#int f0/1
cell-c2955-12(config-if)#speed 100
cell-c2955-12(config-if)#duplex full
```
- The end device must also be configured to match the settings from above.
- Step 5** Enable broadcast suppression filters on all Gigabit Ethernet uplinks to help prevent broadcast floods in case of a misconfiguration or a rogue device:
- ```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# storm-control broadcast level 20
```
- Step 6** Shut down all interfaces that are not in use.
-

For the Catalyst 3750, the startup process is as follows:

-
- Step 1** Load the correct Cisco IOS image on both devices in the stack,
- Step 2** Configure downlink GE ports (gi1/0/14 and gi2/0/14) as trunk ports carrying the one VLAN configured for the Catalyst 2955s. Note that each of these is on a different switch in the stack.

```
!
interface GigabitEthernet1/0/14
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 999
 switchport trunk allowed vlan 20
 switchport mode trunk

interface GigabitEthernet2/0/14
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 999
 switchport trunk allowed vlan 20
 switchport mode trunk
```

- Step 3** Configure the L3 SVIs where the VLANs will terminate. This will also be the IGMP querier interface needed for IGMP snooping.

```
!
interface Vlan20
 ip address 10.17.20.1 255.255.255.0
 ip pim sparse-dense-mode
end
!

interface Vlan30
 ip address 10.17.30.1 255.255.255.0
 ip pim sparse-dense-mode
end
```

- Step 4** Enable broadcast suppression filters on all Gigabit Ethernet uplinks to help prevent broadcast floods in case of a misconfiguration or rogue device:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet1/0/14 , gigabitethernet2/0/14
Switch(config-if)# storm-control broadcast level 20
```

- Step 5** Shut down all interfaces that are not in use.
-

Virtual LAN Segmentation

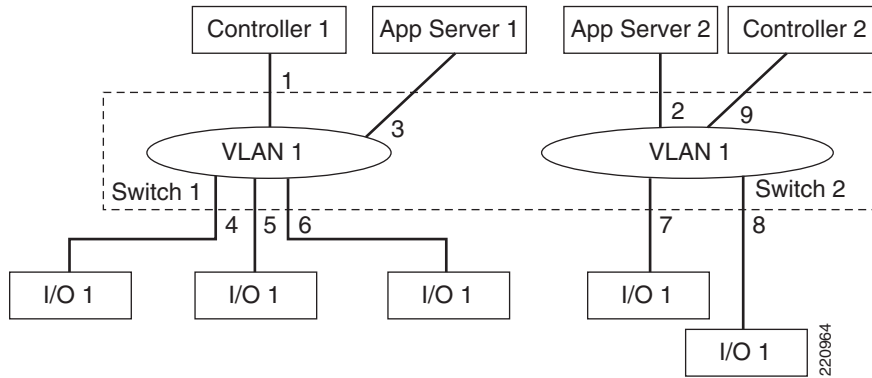
VLAN Overview

A virtual LAN (VLAN) is a switched network segmented on a functional, application, or organizational basis as opposed to a physical or geographical basis. Switches filter destination MAC addresses and forward VLAN frames only to ports that serve the VLAN to which the traffic belongs. A VLAN consists of several end systems, either hosts or network equipment (such as switches and routers), all of which are members of a single logical broadcast domain. A VLAN no longer has physical proximity constraints

for the broadcast domain. This VLAN is supported on various pieces of network equipment (for example, LAN switches) that support VLAN trunking protocols between them. Each VLAN supports a separate spanning tree (IEEE 802.1d).

A VLAN can span multiple switches such that in the topology shown in Figure 4-1, PAC 1 is controlling I/Os 1, 2, and 3 on Switch 1; and PAC 2 is controlling I/Os 4 and 5 on Switch 2 on the same VLAN. In this case, all devices on this VLAN are on the same broadcast domain.

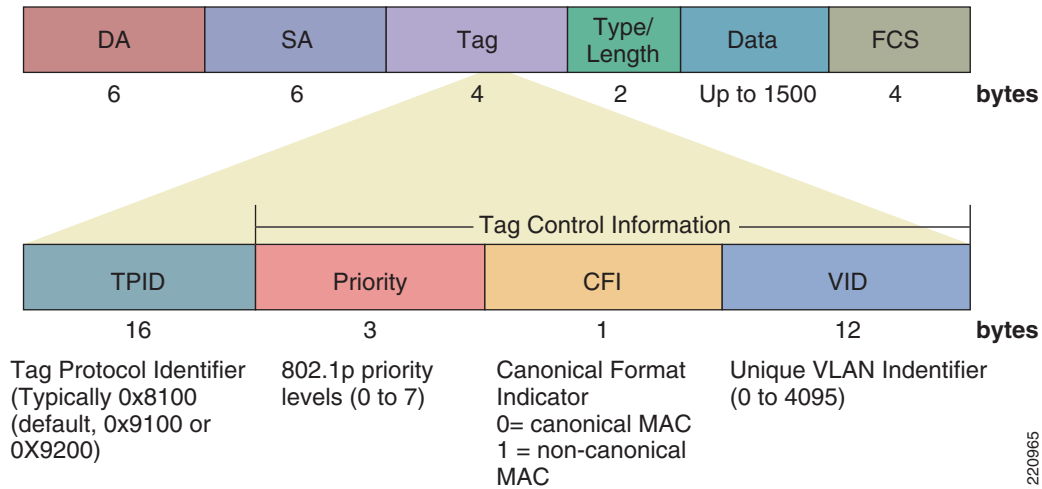
Figure 4-1 VLAN Spanning Multiple Switches



VLAN Details

A VLAN is created by inserting a four-byte VLAN header into the basic Ethernet frame between the source address and length/type fields, as shown in Figure 4-2.


Figure 4-2 VLAN Spanning Multiple Switches



VLANs In the Cell/Area Zone

Cell/area zone devices on the factory floor should include only traffic (application, consumer/producer) that is relevant to running that particular cell. For this reason, EttF 1.1 recommends logically segmenting traffic with the use of VLANs. As shown in Figure 4-3, only one VLAN is recommended for all data traffic relevant to that particular cell/area zone. Because 80–90 percent of traffic is local to one cell, this is the optimal design. Furthermore, producer/consumer traffic is technically confined to a single VLAN because of a feature in Rockwell Automation device firmware. The Catalyst 3750 aggregates all VLANs in the cell/area zone and terminates them with L3 switched virtual interfaces (SVIs).

Figure 4-3 VLANs in the Cell/Area Zone



220959

VLAN Highlights of Ring Topology

Following are VLAN highlights of the ring topology:

- All downward-facing FastEthernet (100 Mbps) ports connected to devices are configured as access ports for a single VLAN.
- All uplinks (GigEthernet, 1000 Mbps) ports are connected as dot1q trunks carrying only the VLAN defined above.
- At the top of the ring, the Catalyst 3750 terminates all VLANs in the ring below with L3 SVIs configured.
- The Catalyst 3750 provides inter-VLAN routing functionality.

VLAN Recommendations

The following are VLAN recommendations for EttF phase 1.1:

- Use one VLAN per ring topology for all manufacturing traffic per cell/area zone.
- If network traffic in one ring is consistently above 2500 packets per second (pps), consider dividing communicating PAC, I/O, and HMI groups into other VLANs.
- If non-manufacturing traffic (PC, and so on) must exist in the ring, it should be on a separate VLAN.
- Remove VLAN 1 from trunk ports and assign a new native VLAN. (See [Spanning Tree Protocol Design, page 4-7.](#))
- Configure VTP Mode as “transparent” to avoid operational error because very few VLANs are used.

VLAN Benefits for EttF

In a flat, bridged network, all broadcast packets generated by any device in the network are sent to and received by all other network nodes. The ambient level of broadcasts generated by the higher layer protocols in the network, known as *broadcast radiation*, typically restricts the total number of nodes that the network can support. In extreme cases, the effects of broadcast radiation can be so severe that an end station spends all its CPU power on processing broadcasts.

VLANs have been designed to address the following problems inherent in a flat, bridged network:

- Scalability issues of a flat network topology
- Simplification of network management by facilitating network reconfigurations

VLANs offer the following features:

- Broadcast control—Just as switches isolate collision domains for attached hosts and forward only appropriate traffic out a particular port, VLANs refine this concept further and provide complete isolation between VLANs. A VLAN is a bridging domain, and all broadcast and multicast traffic is contained within it.
- Security—VLANs provide security in the following two ways:
 - High-security users can be grouped into a VLAN, possibly on the same physical segment, and no users outside of that VLAN can communicate with them.
 - Because VLANs are logical groups that behave like physically separate entities, inter-VLAN communication is achieved through a router. When inter-VLAN communication occurs through a router, all the security and filtering functionality that routers traditionally provide can be used because routers are able to look at Layer 3 information. In the case of nonroutable protocols, there can be no inter-VLAN communication. All communication must occur within the same VLAN.
- Performance—The logical grouping of users allows, for example, a control engineer making intensive use of a networked PAC to be assigned to a VLAN that contains just that engineer and the I/O devices he or she needs. The work of the engineer does not affect the rest of the engineering group, which results in improved performance for the engineer (by being on a dedicated LAN) and improved performance for the rest of the engineering group (whose communications are not slowed down by the engineer using the network).
- Network management—The logical grouping of users, divorced from their physical or geographic locations, allows easier network management. It is no longer necessary to pull cables to move a user from one network to another. Adds, moves, and changes are achieved by configuring a port into the

appropriate VLAN. Expensive, time-consuming recabling to extend connectivity in a switched LAN environment is no longer necessary because network management can be used to logically assign a user from one VLAN to another.

Spanning Tree Protocol Design

STP Overview

Spanning Tree Protocol (STP) is a Layer 2 protocol designed to run on bridges and switches. Its main purpose is to ensure that loops are avoided when there are redundant paths by deterministically blocking appropriate interfaces. If a link failure occurs in such a network, the STP of choice is responsible for establishing a new path for data traffic. The IEEE specification is 802.1D, which has evolved to include the following:

- Common Spanning Tree
- Per VLAN Spanning Tree (PVST)
- Per VLAN Spanning Tree Plus (PVST+, a Cisco proprietary superset of 802.1D)
- Classic STP (802.1D)
- Multiple Instance Spanning Tree (MISTP/802.1S)
- Rapid Spanning Tree (RSTP/802.1W)

Cisco has a recommended Spanning Tree toolkit that includes the following:

- PortFast—Lets the access port bypass the listening and learning phases
- UplinkFast—Provides 3–5 second convergence after link failure
- BackboneFast—Cuts convergence time by MaxAge for indirect failure
- Loop Guard—Prevents the alternate or root port from being elected unless Bridge Protocol Data Units (BPDUs) are present
- Root Guard—Prevents external switches from becoming the root
- BPDU Guard—Disables a PortFast-enabled port if a BPDU is received
- BPDU Filter—Prevents sending or receiving BPDUs on PortFast-enabled ports

For more information on Spanning Tree, see the following URL:

<http://www.cisco.com/warp/public/473/146.html>

EttF version 1.1 recommends only RSTP, IEEE 802.1w, which includes the features from the Cisco Spanning Tree toolkit.

STP Configurable Parameters

With RSTP, IEEE 802.1w, Cisco does not recommend making many changes to the default STP settings. Only the bridge priority should be changed unless you have a valid reason for making other changes.

STP parameters include the following:

- Bridge priority—A configurable value to be used as portion of the bridge identifier. This is the first consideration of STP when root bridge determination is taking place.

The default value of the bridge priority is 32768. In root bridge calculation, the bridge with the lowest value is declared the root bridge. If two or more bridges are involved in a tie, the bridge address (MAC) is used as the final determining factor.

- Hello time—The time interval between the transmission of configuration BPDUs by a bridge that is attempting to become the root or is the root.

The root bridge generates BPDU packets every *HelloTime* seconds, which according to the IEEE standards should be two seconds (2 sec). Each port of the switch has a timer associated with the BPDU information and receiving the BPDUs refreshes this timer.

- MaxAge—Maximum age of received protocol information before it is discarded.

The information associated with a port is considered to be stale if the timer reaches *MaxAge*. The default MaxAge is twenty seconds (20 sec). When a switch stops receiving BPDUs from its root port and the MaxAge expires, the switch looks for a new root port, from the pool of blocking ports. If no blocking port is available, it claims to be the root itself on the designated ports.

- Forward delay—Time spent by a port in the listening state and the learning state before moving to the learning or forwarding state, respectively. It is also the value used for the aging time of dynamic entries in the filtering database, while received configuration messages indicate a topology change.

The default value of the *Forward Delay* is fifteen seconds (15 sec).

- Diameter—Maximum number of bridges between any two points of attachment of end stations. Although this is not configurable directly, it can be manipulated by changing the max age variable from above.

Network diameter can have a profound effect on the operation of STP and the network as a whole because the latency of BPDUs increases as the network grows in diameter. The default value for *Diameter* is seven (7).

- Port cost(s)—Contribution of the path through this port, when the port is the root port, to the total cost of the path to the root for this bridge.

The cost of each port between a given bridge and the root bridge contributes to the overall path cost. Some of the default values used are as follows.


- Gig E (1000 Mbps)—4
- OC-12 (622 Mbps)—6
- OC-3 (155 Mbps)—14
- FastE (100 Mbps)—19
- Ethernet (10 Mbps)—100

More on STP Redundancy

A Layer 2 loop is defined as the existence of two paths between any two Layer 2 devices within a single network. These loops can create many different problems within a network. The problems usually manifest themselves in the form of a “storm” incident. This is the continuous propagation of one or more packets within the network, such as a broadcast storm. When a Layer 2 switch receives a broadcast, it is sent to each of its ports including ports connected to other switches. If a loop exists in the network, it is possible for a switch to process the broadcast endlessly.

In [Figure 4-4](#), Host A sends a broadcast to Switch 1.

Figure 4-4 Layer 2 Loop

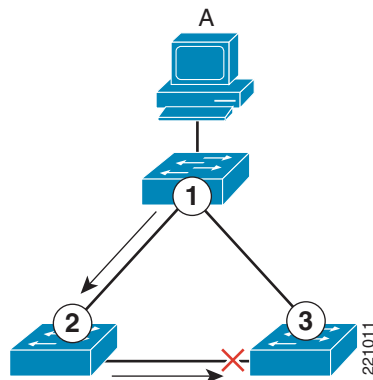


If a loop exists, as shown in [Figure 4-4](#), the broadcast is forwarded to Switch 2, Switch 3, and back to Switch 1. Upon arrival in Switch 1, the broadcast is then resent to Switch 2. This is repeated until a break in the cycle is experienced. The loop is also experienced in the opposite direction and on any other loops present in the network. This loop eventually diminishes network performance. The use of STP prevents Layer 2 loops.

The implementation of STP allows the switches to communicate through BPDUs to form a loop-free topology at Layer 2. During this negotiation process, the switches use the algorithm to decide on the final state of all ports: blocking or forwarding. Blocking strategic ports prevents Layer 2 loops in the network.

In [Figure 4-5](#), through the implementation of STP, Switch 3 is in the blocking state on its port facing Switch 2.

Figure 4-5 Use of STP to Block Loops




This effectively breaks the Layer 2 loop shown in [Figure 4-4](#). Switch 3 still receives the broadcast; however, it comes from Switch 1. Furthermore, Switch 3 does not forward the broadcast received from Switch 1 to Switch 2 because the port is blocking.

STP Topology for EttF

For EttF version 1.1, the STP topology consists of n number of devices in a ring configuration. The purpose of a ring architecture is to achieve a level of redundancy at the lowest possible cabling cost. Furthermore, ring designs are popular in control network environments and are well-aligned with customer expectations. The ring devices are sitting at the access layer in a Cisco three-tier architecture with the distribution layer devices responsible for L2 (downlink) and L3 (uplink) functionality.

For EttF version 1.1, there are essentially 1 to many (1 to n) cell/area zones, each constituting a separate L2 STP domain. Only one VLAN is carried per cell/area zone. EttF device traffic flows both intra-cell (within a cell) and inter-cell (across cells). All multicast producer/consumer traffic is confined to within a cell/area zone because of the TTL=1 limitation on multicast traffic. To achieve a more deterministic design, Cisco recommends that a root bridge be chosen rather than letting the network choose one automatically. In Figure 4-6, Device CZ-C3750-1 is configured to be the root bridge.

Figure 4-6 STP Topology for EttF



The root bridge can be viewed as the top of the network hierarchy with which every other switch in the network must communicate. The root bridge concept is crucial in determining a loop-free topology at Layer 2. In this determination, each device selects the best path possible towards the root bridge. The result of this is the spanning tree topology. Alternate paths towards the root bridge that would result in Layer 2 loops are in a blocked state.

STP Considerations for the Ring

As a general rule, RPVST+ should be deployed because of its faster convergence and ease of use. Enter the following global command to set the spanning-tree in Rapid-PVST mode:

```
spanning-tree mode rapid-pvst
```

The following sections describe further STP considerations and best practices for a ring topology.


Control Device Placement

Knowing which link is blocking is important in designing where to place devices; for example, a PAC talking to an I/O device. If the control engineer knows that PAC-A will be communicating with I/O-B for the majority of the time, it is advisable to connect them either on the same switch or on adjacent switches that are not STP blocked. This ensures the minimum convergence time if a network failure occurs. This also saves bandwidth and reduces latency because the traffic does not have to traverse the entire ring.

Trunk Ports or Access Ports

On the Cisco 2955, all uplink 1 GE ports should be configured as trunk ports carrying only one VLAN. All FastEthernet (fa0/1–fa0/12) interfaces should be configured as access ports for the particular VLAN carried on the uplink trunk (see Figure 4-7).

Figure 4-7 Trunk Port Configuration



Sample Trunk Configuration

Following is a sample trunk configuration:

```
interface GigabitEthernet0/1
switchport trunk native vlan 20
switchport trunk allowed vlan 20
switchport mode trunk
end
Sample Access Port Configuration:
!
interface FastEthernet0/1
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
end
```

VLAN 1 Minimization

VLAN 1 has a special significance in Catalyst networks. When trunking, the Catalyst Supervisor Engine always uses the default VLAN, VLAN 1, to tag a number of control and management protocols such as Cisco Discovery Protocol (CDP), VLAN Trunking Protocol (VTP), and Port Aggregation Protocol (PAgP); as well as management protocols such as Simple Network Management Protocol (SNMP), Telnet, Secure Shell (SSH), and Syslog. All switch ports are configured by default to be members of VLAN 1, and all trunks carry VLAN 1 by default. When the VLAN is used in this way, it is referred to as the native VLAN. The default switch configuration sets VLAN 1 as the default native VLAN on the Catalyst trunk ports. You can leave VLAN 1 as the native VLAN, but keep in mind that any switches that run Cisco IOS Software in your network set all interfaces that are configured as Layer 2 switch ports to access ports in VLAN 1 by default. Most likely, a switch somewhere in the network uses VLAN 1 as a VLAN for user traffic.

The main concern with the use of VLAN 1 is that, in general, the Supervisor Engine should not be constantly interrupted by a lot of the broadcast and multicast traffic that end stations generate (user data). Multicast applications in particular tend to send a lot of data between servers and clients. The Supervisor Engine does not need to see this data. If the resources or buffers of the Supervisor Engine are fully occupied as the Supervisor Engine listens to unnecessary traffic, the Supervisor Engine can fail to see management packets that can cause a bridging loop.

VLAN 1 tags and handles most of the control plane traffic. VLAN 1 is enabled on all trunks by default. With larger campus networks, you need to be careful of the diameter of the VLAN 1 STP domain. Instability in one part of the network can affect VLAN 1 and can influence control plane stability and STP stability for all other VLANs. To limit the VLAN 1 transmission of user data and operation of STP on an interface, Cisco recommends doing the following:

- Clear VLAN 1 from the trunk to avoid control plane traffic being part of STP on those links, and allow only VLANs that have data traffic flowing through them:

```
Switch(config)#interface type slot/port
Switch(config-if)#switchport trunk allowed vlan 20
```

- Change the native VLAN from 1 to an arbitrary VLAN that is not in use:

```
Switch(config)#interface type slot/port
Switch(config-if)#switchport trunk native vlan 999
```


To verify, issue the following:

```
Switch#show int trunk

Port      Mode      Encapsulation  Status      Native vlan
Gi0/1     on        802.1q         trunking    999
Gi0/2     on        802.1q         trunking    999

Port      Vlans allowed on trunk
Gi0/1     20
Gi0/2     20

Port      Vlans allowed and active in management domain
Gi0/1     20
Gi0/2     20

Port      Vlans in spanning tree forwarding state and not pruned
Gi0/1     20
Gi0/2     20
```

You should see that the link is successfully trunking and carrying only VLAN 20, and that the native VLAN is something other than 1.

Location of the Root Bridge

The STP root bridge should be forced by setting this bridge to have the lowest priority. In the Etf 1.1 design, Cisco recommends that the Catalyst 3750 be elected as the root bridge because logically it sits at the top of the ring. This makes troubleshooting easier if there is a need to look at the STP state of the ring devices.

```
CZ-C3750-1(config)# spanning-tree vlan 1-1024 priority 8096
```

PortFast on Access Ports

You can use PortFast to bypass normal spanning tree operation on access ports. PortFast speeds up connectivity between end stations and the services to which end stations need to connect after link initialization. The Microsoft DHCP implementation needs to see the access port in forwarding mode immediately after the link state goes up to request and receive an IP address. Some protocols, such as Internetwork Packet Exchange (IPX)/Sequenced Packet Exchange (SPX), need to see the access port in forwarding mode immediately after the link state goes up to avoid get nearest server (GNS) problems.

PortFast Operational Overview

PortFast skips the normal listening, learning, and forwarding states of STP. This feature moves a port directly from blocking to forwarding mode after the link is seen as up. If this feature is not enabled, STP discards all user data until it decides that the port is ready to be moved to forwarding mode. This process can take up to twice the ForwardDelay time, which is 30 seconds by default.

PortFast mode prevents the generation of an STP topology change notification (TCN) each time a port state changes from learning to forwarding. TCNs are normal. However, a wave of TCNs that hits the root bridge can extend the convergence time unnecessarily. A wave of TCNs often occurs in the morning when people turn on their PCs.

The following are PortFast recommendations for EttF 1.1:

- Set STP PortFast to “on” for all enabled host ports connected to either a PAC, I/O device, or HMI.
- Explicitly set STP PortFast to “off” for switch-switch links and ports that are not in use.

STP Limitations


As mentioned earlier, EttF 1.1 implements only RPVST+ because it can achieve faster convergence times than PVST or PVST+ by relying on an active bridge-to-bridge handshake mechanism rather than depending on network-wide timers specified by the root bridge. With standard STP, convergence times are expected in the range of 30–60 seconds, depending on network conditions and timer settings. This is unacceptable in a control zone environment. However, with RPVST+, these numbers can be reduced to sub-second levels depending on conditions (type of failure, number of MAC addresses, traffic load). Also, the number of devices in the ring determine how fast STP converges. For a complete list of testing results, see [Appendix A, “Characterization of the EttF Cell/Area Zone Design.”](#)

RSTP+ Convergence Process

Figure 4-8 shows the sequence of STP events in the RSTP+ convergence process.

Figure 4-8 RSTP+ Convergence

CZ-C3750-1



The sequence of STP events that take place given a link failure between Switch A and Switch B is as follows:

1. As soon as the failure occurs, Switch B loses its root port and claims to be the new root bridge.
2. This new root bridge information circulates down in the ring in a clockwise manner, from B to C, C to D, D to E and E to F.
3. When bridge F receives this inferior BPDU (it contains worse information than the one emanating from switch A, the “real” root), it in turn “replies” back to its upstream neighbors (E-D-C-B) to let them know it can still reach root switch A.
4. With RSTP+, there is no MaxAge parameter to depend on before bridge F can react to the reception of the inferior BPDU of E.
5. As soon as bridge F receives that inferior BPDU, it immediately transitions port 2/49 to forwarding from blocking and informs switch E via a “proposal” mechanism that it wants to become its designated bridge.
6. The same mechanism then takes place very rapidly between each pair of switches all the way towards bridge B.
7. Bridge F also initiates a topology change notification when opening up port 2/49 to flush stale learning table information in the other bridges of the ring.
8. The STP network is now converged.

Thus, L2 convergence is defined as the completion of all STP state changes and the completion of updating the MAC addresses in the CAM table as measured by data traffic.

Multicast Design

EtherNet/IP Multicast Traffic Patterns

Although traditional multicast services (usually video feeds) tend to scale with the number of streams, the EtherNet/IP model is implemented as a many-to-many model and scales differently because of a built-in feedback mechanism. In the specification, devices generate data for consumption by other devices. The devices that generate the data are called producers, and the devices receiving the information are called consumers. The data exchange model is therefore referred to as the producer-consumer model. Multicast is more efficient over unicast in that in many cases, multiple consumers want the same information from a particular producer. However, because every consumer of traffic needs to respond with a heartbeat, a significant load of unicast packets is generated on the network.

Most EtherNet/IP devices generate very little data. However, networks with a large number of nodes can generate a large aggregate amount of multicast traffic. If a method to control this is not deployed, this aggregate traffic can swamp some or all of the end devices in the network. This is aggravated by the fact that there is likely unicast traffic (FTP, HTTP, and so on) going to the device. This is even more critical if there are no other mechanisms in place (such as QoS) to prioritize real-time traffic.

In general, end devices can be overwhelmed by the following:

- The port speed is overrun
- More packets are received than the network interface controller can handle
- More packets are received than the host processor can process

If the aggregate data exceeds the port speed (that is, 20 Mbps going to a 10 Mbps configured port), traffic is dropped because of congestion.


If a device on the network can process only 900 packets per second (pps), the network design must ensure that this particular end device does not see more than 900 pps. The following list is an example of a network that produces more than 900 pps. In this example, there are 12 producers. It is assumed that each producer is also a consumer. Each producer is generating only 100 pps, but because all consumers see all producer traffic, they actually see 1200 pps, and because they also are acknowledging at least one stream, they are also sending 100 pps of unicast traffic.

- Number of PACs—12
- RPI—10 Msecs
- Packets of multicast—100 pps
- Size of multicast packets—600 bytes
- Bandwidth of each multicast stream—0.5 Mbps
- Packets of unicast—100 pps
- Size of unicast packets—100 bytes
- Bandwidth of each unicast stream (hello echoes)—0.096 Mbps
- Backplane traffic—7.1 Mbps
- Traffic seen on each port (multicast + unicast)—6.0 Mbps
- Number of packets received—1300 pps

IGMP snooping ensures that only the multicast traffic requested by the particular end device is received on its inbound interface. With IGMP snooping, each end device processes only 200 pps instead of the aggregate 1300.

Figure 4-9 shows an example of multicast with the producer-consumer model.

Figure 4-9 Multicast on Producer-Consumer Model



220969

IGMP Snooping

Layer 2 switches can use IGMP snooping to constrain the flooding of multicast traffic by dynamically configuring the multicast traffic to be forwarded to only those access interfaces associated with devices requesting the multicast group. As the name implies, IGMP snooping requires the LAN switch to snoop on the Internet Group Management Protocol (IGMP) transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry (see [Figure 4-10](#)); when it receives an IGMP leave group message from a host, it removes the host port from the table entry (see [Figure 4-11](#)). It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

Figure 4-10 IGMPv1 Join Process





Figure 4-11 IGMPv2 Leave Process



The multicast router sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch creates one entry per VLAN in the Layer 2 forwarding table for each MAC group from which it receives an IGMP join request. It sets a flag for each port that will receive the particular group.

Multicast group membership lists can consist of both user-defined and IGMP snooping-learned settings. If an STP TCN, a port group, or a VLAN ID change occurs, the IGMP snooping-learned multicast groups from this port on the VLAN are deleted, and must be relearned on the next IGMP query message.

When a host connected to the switch wants to join an IP multicast group, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from the router, it forwards the query to all ports in the VLAN. Hosts wanting to join the multicast group respond by sending a solicited report (join) message to the switch. The switch creates a multicast forwarding table entry for the group if one is not already present. It also adds the interface where the join message was received to the forwarding table entry. The host associated with that interface receives multicast traffic for that multicast group.

When hosts want to leave a multicast group, they can either silently leave by not responding to an IGMP query message, or they can send an IGMPv2 leave message. When the switch receives a leave message from a host, it sends out a MAC-based general query to determine whether any other devices connected


to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

I/O or produce tag traffic does not pass through a router. By design, the time-to-live (TTL) parameter is configured in the Rockwell Automation firmware for a value of 1. This value is decremented by any router and then discarded. A value of 1 was selected to avoid attempts to implement I/O control (high-speed) through a slow network device (router) or through a slow network. I/O or produce tag traffic uses multicast messaging.

IGMP Querier and EtherNET/IP Traffic

The purpose of the IGMP snooping with querier is to generate periodic query messages. Consumers of particular multicast groups should respond with an IGMP report message stating that they still want to receive the data stream. If the client is finished with the stream, it does not issue the report message. The Ethernet switch, using IGMP snooping, should prune off that port, which has the same effect as the client issuing an unsolicited IGMP leave message. In the EttF architecture, IGMP snooping is turned on by default in all the switches in the cell/area zone (Catalyst 2955) and the Catalyst 3750. The Catalyst 3750 acts as the IGMP querier. A querier must be configured for each switched virtual interface (SVI) associated with each VLAN. The querier knob can be turned on by configuring **ip pim sparse-dense-mode** in the SVI-Vlan 20 (see Figure 4-12).

Figure 4-12 EttF Architecture



Another option is to configure the global IP IGMP querier address from the CLI as follows: **ip igmp snooping querier address ip_address**. For more information on configuring IGMP snooping with querier, see the following URL:
http://www.cisco.com/en/US/partner/products/hw/switches/ps5532/products_configuration_guide_chapter09186a008081bb8c.html#wp1130762.

IGMP Configurations

Configure IGMP querier on the Catalyst 3750 multicast router (**ip pim sparse-dense-mode**) as follows:

```
C3750-1#
Building configuration...

Current configuration: 87 bytes
!
interface Vlan20
 ip address 10.17.20.1 255.255.255.0
 ip pim sparse-dense-mode
end
```

The **show ip igmp int vlan <>** command verifies the IGMP querying router and the multicast designated router:

```
C3750-1#show ip igmp int vlan 20
Vlan20 is up, line protocol is up
 Internet address is 10.17.20.1/24
 IGMP is enabled on interface
 Current IGMP host version is 2
 Current IGMP router version is 2
 IGMP query interval is 60 seconds
 IGMP querier timeout is 120 seconds
 IGMP max query response time is 10 seconds
 Last member query count is 2
 Last member query response interval is 1000 ms
 Inbound IGMP access group is not set
 IGMP activity: 10 joins, 5 leaves
 Multicast routing is enabled on interface
 Multicast TTL threshold is 0
 Multicast designated router (DR) is 10.17.20.1 (this system)
 IGMP querying router is 10.17.20.1 (this system)
 No multicast groups joined by this system
C3750-1#
```

The multicast groups attached to the Catalyst 3750 (SVI 20) are as follows (**show ip igmp groups vlan 20**):

```
C3750-1#show ip igmp groups vlan 20
IGMP Connected Group Membership
Group Address      Interface      Uptime      Expires      Last Reporter
239.192.21.96     Vlan20        1w0d        00:02:09     10.17.20.164
239.192.21.160    1w1d         00:02:11    10.17.20.164
239.192.24.160    Vlan20        6d01h       00:02:13     10.17.20.164
239.192.24.161    Vlan20        6d01h       00:02:13     10.17.20.164
239.192.24.192    Vlan20        6d02h       00:02:11     10.17.20.164
C3750-1#
```


The output from the **show ip igmp snooping vlan <>** command verifies IGMP snooping as well as the multicast router learning mode. If IGMP snooping with querier is not enabled on the Catalyst 3750, the bottom portion of the output with the multicast router mode is not present:

```
cell-c2955-12#show ip igmp snooping vlan 20
Global IGMP snooping configuration:
-----
IGMP snooping                : Enabled
IGMPv3 snooping (minimal): Enabled
Report suppression           : Enabled
TCN solicit query            : Disabled
TCN flood query count        : 2
Last member query interval: 1000

Vlan 20:
-----
IGMP snooping                : Enabled
Immediate leave               : Disabled
Multicast router learning mode : pim-dvmrp
Source only learning age timer : 10
Last member query interval    : 1000
CGMP interoperability mode    : IGMP_ONLY
cell-c2955-12#
```

The output from the Catalyst 2955 in the cell/area zone (**show ip igmp snooping querier**) verifies the IGMP querier in the Catalyst 3750:

```
cell-c2955-12#show ip igmp snooping querier
Vlan      IP Address      IGMP Version      Port
-----
20        10.17.20.1      v2                 Gi0/1
cell-c2955-12#
```

Switch Troubleshooting Toolkit

The following are some general troubleshooting techniques on the cell/area zone switches:

1. Use Port Mirroring (SPAN) when troubleshooting and you need to characterize the traffic flow.

```
Switch(config)# monitor session 1 source interface fastethernet0/1
Switch(config)# monitor session 1 destination interface fastethernet0/2 <- A Sniffer
would be connected here
```

2. Use **show ip traffic** to get a quick snapshot of traffic statistics to check whether there are any skewed data points; for example, excessive broadcasts indicate a broadcast storm.

Sample output is as follows:

```
IP statistics:
Rcvd: 98 total, 98 local destination
      0 format errors, 0 checksum errors, 0 bad hop count
      0 unknown protocol, 0 not a gateway
      0 security failures, 0 bad options
Frag:0 reassembled, 0 timeouts, 0 too big
      0 fragmented, 0 couldn't fragment
Bcast:38 received, 52 sent
Sent: 44 generated, 0 forwarded
      0 encapsulation failed, 0 no route
ICMP statistics:
Rcvd: 0 format errors, 0 checksum errors, 0 redirects, 0 unreachable
      0 echo, 0 echo reply, 0 mask requests, 0 mask replies, 0 quench
```

```

    0 parameter, 0 timestamp, 0 info request, 0 other
Sent: 0 redirects, 3 unreachable, 0 echo, 0 echo reply
    0 mask requests, 0 mask replies, 0 quench, 0 timestamp
    0 info reply, 0 time exceeded, 0 parameter problem
UDP statistics:
Rcvd: 56 total, 0 checksum errors, 55 no port
Sent: 18 total, 0 forwarded broadcasts
TCP statistics:
    Rcvd: 0 total, 0 checksum errors, 0 no port
Sent: 0 total
EGP statistics:
Rcvd: 0 total, 0 format errors, 0 checksum errors, 0 no listener
    Sent: 0 total
IGRP statistics:
Rcvd: 73 total, 0 checksum errors
    Sent: 26 total
HELLO statistics:
Rcvd: 0 total, 0 checksum errors
Sent: 0 total
ARP statistics:
Rcvd: 20 requests, 17 replies, 0 reverse, 0 other
Sent: 0 requests, 9 replies (0 proxy), 0 reverse
Probe statistics:
Rcvd: 6 address requests, 0 address replies
    0 proxy name requests, 0 other
Sent: 0 address requests, 4 address replies (0 proxy)
    0 proxy name replies

```

3. To troubleshoot link problems (port flapping, errdisable, and so on), see the guide at the following URL:
http://www.cisco.com/en/US/partner/products/hw/switches/ps700/products_tech_note09186a008015bfd6.shtml
4. Verify that CPU usage is below 50 percent average utilization. High CPU can indicate a broadcast storm, multicast data traffic flooding (for example, IGMP snooping configured incorrectly), or an attack.

```

cell-c2955-9#show proc cpu | exc 0.00
CPU utilization for five seconds: 1%/0%; one minute: 0%; five minutes: 0%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min  TTY Process
 22   6918340     3448871    2005   0.40%  0.39%  0.40%  0 Calhoun Statisti
 68         60         43       1395   0.90%  0.09%  0.02%  0 Exec

```



CHAPTER 5

Implementation of Security

Overview

The number of skilled hackers has multiplied, and a variety of sophisticated hacking tools are freely available on the Internet. These tools exploit the way the network is designed to work, and are simple enough for even a novice to use. This combination has dramatically increased the risk to networks.

Some of the more dangerous types of attacks include the following:

- **Packet sniffer**—Software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a particular collision domain. Sniffers are used legitimately in networks today to aid in troubleshooting and traffic analysis. However, because several network applications send data in clear text (Telnet, File Transfer Protocol [FTP], Simple Message Transfer Protocol [SMTP], Post Office Protocol [POP3], and so on), a packet sniffer can provide meaningful and often sensitive information, such as usernames and passwords. One serious problem with acquiring usernames and passwords is that users often reuse their login names and passwords across multiple applications and systems.
- **IP spoofing**—A hacker inside or outside a network impersonates the conversations of a trusted computer. The hacker uses either an IP address that is within the range of trusted IP addresses for a network, or an authorized external IP address that is trusted and to which access is provided to specified resources on a network. IP spoofing attacks are often a launch point for other attacks. The classic example is to launch a denial-of-service (DoS) attack using spoofed source addresses to hide the identity of the hacker.
- **Distributed denial-of-service (DDoS) attacks**—Multiple compromised systems flood the bandwidth or resources of a targeted system, usually one or more web servers. Although the attack does not flood the entire network with traffic, it overwhelms the specific device and takes it out of service. These systems are compromised by attackers using a variety of methods. Malware can carry DDoS attack mechanisms; one of the more well-known examples of this was MyDoom. Its DoS mechanism was triggered on a specific date and time. This type of DDoS attack involved hard-coding the target IP address before release of the malware. No further interaction was necessary to launch the attack.
- **Network reconnaissance**—Learning information about a target network by using publicly available information and applications. When hackers attempt to penetrate a particular network, they often need to learn as much information as possible about the network before launching attacks. This can take the form of DNS queries, ping sweeps, and port scans. DNS queries can reveal such information as who owns a particular domain and what addresses have been assigned to that domain. Ping sweeps of the addresses revealed by the DNS queries can present a picture of the live hosts in a particular environment. After such a list is generated, port scanning tools can cycle through all well-known ports to provide a complete list of all services running on the hosts discovered by the

ping sweep. Finally, the hackers can examine the characteristics of the applications that are running on the hosts. This scenario can lead to specific information that is useful when the hacker attempts to compromise that service.

- **Unauthorized access**—Although unauthorized access attacks are not a specific type of attack, they refer to most attacks executed in networks today. A brute-force attack on a Telnet login requires the Telnet prompt on a system. On connection to the Telnet port, a message might indicate “authorization required to use this resource.” If the hacker continues to attempt access, their actions become “unauthorized.” These kinds of attacks can be initiated on both the outside and inside of a network.
- **Virus and Trojan horse applications**—The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses refer to malicious software that is attached to another program to execute a particular unwanted function on a user workstation. An example of a virus is a program that is attached to `command.com` (the primary interpreter for Windows systems), which deletes certain files and infects any other versions of `command.com` that it can find. A Trojan horse is different only in that the entire application is written to look like something else, when in fact it is an attack tool. An example of a Trojan horse is a software application that runs a simple game on the user workstation. While the user is occupied with the game, the Trojan horse mails a copy of itself to every other user in the address book of the user. Other users then get the game and play it, thus spreading the Trojan horse.
- **Password attacks**—Hackers can implement password attacks using several methods, including brute-force attacks, Trojan horse programs, IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually refer to repeated attempts to identify a user account or password. These repeated attempts are called brute-force attacks. Often, a brute-force attack is performed using a program that runs across the network and attempts to log into a shared resource, such as a server. When hackers successfully gain access to resources, they have the same rights as the users whose accounts have been compromised to gain access to those resources. If the compromised accounts have sufficient privileges, the hackers can create back doors for future access without concern for any status and password changes to the compromised user accounts.

The goal of the comprehensive model provided here is to prevent attacks by keeping the outsiders out and the insiders honest. Specific goals include the following:

- Prevent external hackers from getting access to the network
- Allow only authorized users into the network
- Prevent those inside the network from executing deliberate or inadvertent attacks
- Provide various levels of access for various types of users

To be truly effective, the security policy must do this in a way that is transparent to the users and easy to administer, and that does not disrupt the operations of the plant floor.

To accomplish all this, the solution needs to provide the following:

- Network-wide security that is fully embedded into the network infrastructure
- Protection, prevention, and self-protection
- Control over who has network access and what they can do

The following security components of the EttF 1.1 solution address the major security concerns of defending against threat, establishing trust boundaries and verifying identity, and securing business communications:

- Device hardening
- Threat defense—Guard the network against malicious as well as unintentional attack. Threat defense can be further broken down into the following goals:
 - Defending the edge—Using Cisco Adaptive Security Appliance (ASA) integrated firewalls and intrusion detection systems (IDS) to fortify the network edge against intrusion and attack.
 - Protecting the interior—Enabling Cisco IOS security features on routers and switches to protect the network against emerging internal attacks.
 - Guarding the endpoints—Using the Cisco Security Agent (CSA) to proactively defend against infection and damage to hosts, such as human-machine interfaces (HMIs), servers, and PCs.
 - Trust and identity—Controlling who has access from the enterprise network to the plant floor network. This control is provided by CiscoSecure Access Control Server (ACS).
- Secure communications—Protecting the confidentiality of internal and external data communication.

Network Device Hardening

Device hardening refers to changing the default posture of a system out of the box to make it more secure. These network devices include, among others, routers, switches, firewalls, and network-based intrusion detection system (NIDS). The default security of these devices can differ, which changes the amount of work required to harden a particular device.

An important characteristic of all these devices is the availability of a console port. The console port has privileged access to these devices because it generally implies physical access to the device (though this could be a modem). The console port defaults to having initial authentication that is weak or nonexistent and is able to send a break signal to the device upon boot. This is used to reset most of these types of devices or to recover from a lost password.

Because of the capabilities of a console port, it is important to control physical access to networking devices whenever possible.



Note

This section on network devices assumes that the devices are not running on general-purpose operating systems. If they are, be sure to run the host operating system-hardening as well as the network device-hardening steps.

From a configuration perspective, the methods for hardening a router or switch are very similar.

[Table 5-1](#) summarizes the device hardening techniques needed for the platforms supported by the EttF 1.1 solution. The detailed configuration is presented in the following sections.

Table 5-1 **Device Hardening Techniques**

	Catalyst 2955	Catalyst 3750	Catalyst 4500
Disable unneeded services—DNS lookup	Yes	Yes	Yes
Disable unneeded services—Small services	Yes	Yes	Yes
Disable unneeded services—BootP server	N/A	Yes	Yes

Table 5-1 Device Hardening Techniques (continued)

Disable unneeded services—Source routing and directed broadcast	N/A	Yes	Yes
Disable unneeded services—Proxy ARP	N/A	Yes	Yes
Disable unneeded services—ICMP redirects	N/A	Yes	Yes
Password encryption	Yes	Yes	Yes
Authentication settings—Enable secret	Yes	Yes	Yes
Authentication settings—Login banner	Yes	Yes	Yes
Authentication settings—Line access	Yes	Yes	Yes
Authentication settings—Set up usernames	Yes	Yes	Yes
Authentication settings—Secure Shell (SSH)	Yes (supported only by crypto image)	Yes (supported only by crypto image)	Yes (supported only by crypto image)
Management access—HTTP server	Yes	Yes	Yes
Management access—NTP	Yes	Yes	Yes
Management access—ACL Options	Yes	Yes	Yes

Router

Router hardening has recently gained attention because attacks have increasingly targeted routed infrastructure. This section outlines steps to take when hardening a router; configuration examples are for Cisco IOS devices. For more information about router hardening, see the following URLs:

- Infrastructure Protection on Cisco IOS Software-Based Platforms: http://www.cisco.com/application/pdf/en/us/guest/products/ps1838/c1244/cdcont_0900aecd804ac831.pdf
- Improving Security on Cisco Routers: <http://www.cisco.com/warp/public/707/21.html>
- Building Bastion Routers Using Cisco IOS: <http://www.phrack.com/phrack/55/P55-10>
- NSA Router Security Configuration Guide (290 pages!): <http://www.nsa.gov/snac/cisco/>

Basic Hardening Settings

The following hardening steps are useful on almost every router you deploy in a network. These steps include disabling unneeded services and ensuring that passwords are encrypted whenever possible.

Disable Unneeded Services

Turn off DNS lookups for the router with the following command:

```
Router(config)#no ip domain-lookup
```

Although not strictly security-related, this is the first command to type on a fresh router before doing any other configuration (assuming, of course, you do not need domain resolution for a feature you plan to use). Otherwise, be careful to avoid input errors. Typing the command **enable** instead of **enable** results in a long timeout while the router tries to find host “enable” and communicate with it.

Disable small services such as echo, chargen, and discard, as well as the finger service. After Cisco IOS Release 11.3, these services are disabled by default, but it never hurts to have these commands as part of the script you use to harden a device. These small services should almost always be turned off because they have no legitimate use.

```
Router(config)#no service tcp-small-servers
Router(config)#no service udp-small-servers
Router(config)#no service finger
```

Disable the BootP server with the following command if you are not using it on your network (most do not):

```
Router(config)#no ip bootp server
```

Disable source routing and directed broadcast. These should be off by default on reasonably current routers, but verify this with the following commands:

```
Router(config-if)#no ip directed-broadcast
Router(config)#no ip source-route
```

You can disable Proxy ARP in most situations, assuming your devices are routing aware:

```
Router(config-if)#no ip Proxy-arp
```

ICMP redirects should be sent only to end systems that have multiple outbound routes from which to choose. In situations in which IP redirects are unnecessary, disable them with the following command:

```
Router(config-if)#no ip redirects
```

Password Encryption

The following command enables a simple Vigenere cipher, which encrypts most passwords on a router that would otherwise be shown as clear text in the configuration:

```
Router(config)#service password-encryption
```

This cipher, as implemented on Cisco routers, is very weak and can easily be broken. It is enabled primarily to prevent a casual observer from noting your passwords. For example, you might not want a coworker observing your work to learn the password for your router after you type **wr t**.

Authentication Settings

This section outlines authentication-related settings, including the use of **enable secret**, login banners, line access, usernames stored locally or through AAA servers, and device access by SSH.

Enable Secret

Enable strong MD5-hashed passwords for router enable mode. The following password should be used instead of the basic **enable password** encrypted by using **service password-encryption**. It is much more secure, though it has the same susceptibility to dictionary attacks as any hashed password. Choosing strong passwords mitigates dictionary attacks.

```
Router(config)#enable secret password
```

Login Banner

Enable a warning banner to be presented to users when they connect to the device. This sort of banner can aid in prosecution in some jurisdictions and should generally at least include a statement saying that unauthorized access is prohibited. Be sure not to disclose any information that would be useful to the attacker such as platform type, software version, owner, location, and so on.

```
Router(config)#banner motd ^
Enter TEXT message. End with the character '^'.
Enter your warning banner message here.
^
```

Line Access

On a standard Cisco router, there are three primary ways to log on:

- VTY line (**line vty 0 4**, though some routers go to 15)
- Console port (**line con 0**)
- Auxiliary port (**line aux 0**)

Fresh out of the box, only the console and aux ports can be used to access the device. Generally, only the console port is needed and not the aux port. To set up the console port, enter the following commands:

```
Router(config)#line con 0
Router(config-line)#exec-timeout 5 0
Router(config-line)#password password
Router(config-line)#login
```

These commands enable login with a local password and time out the connection after 5 minutes and 0 seconds of inactivity.

To disable the aux port, type the following commands:

```
Router(config)#line aux 0
Router(config-line)#no exec
```

Turning off exec prevents logon to the device. Additional commands such as **transport input none** or **exec-timeout 0 1** are not going to make you more secure. Controlling VTY access is separate and requires the following commands:

```
Router(config)#line vty 0 3
Router(config-line)#exec-timout 5 0
Router(config-line)#password password
Router(config-line)#login
Router(config-line)#transport input protocol
```

Typically, a router has 5 VTY lines. The preceding four commands set up access in a very similar fashion to the console port. Replace *protocol* with your method of access, preferably SSH.



Note

SSH is supported only by the IOS crypto images of the respective Catalyst switching platforms.

The following eight lines reserve the last VTY port for a specific IP address. This is useful if someone is attempting to deny service to the login process on the router (which can be done without the password). You can use the access class settings referenced here for lines 0 to 3 as well. If you do, open the access control list (ACL0) to allow a wider range of IP addresses to access (for instance, your entire management subnet).


```

Router(config)#line vty 4
Router(config-line)#exec-timeout 5 0
Router(config-line)#password password
Router(config-line)#login
Router(config-line)#transport input protocol
Router(config-line)#access-class 99 in
Router(config)#access-list 99 permit host adminIP
Router(config)#access-list 99 deny any logs

```

Setting Up Usernames

If you do not have access to TACACS+ or RADIUS, local usernames can be configured on a system as follows:

```

Router(config)#username username password password
Router(config)#line vty 0 4
Router(config-line)#login local

```

The preceding commands set up a local username and password and then configure the VTY lines to use a local database.

To configure TACACS+ access to a system, you must first enable the AAA system:

```

Router(config)#aaa new-model

```

You must then define the TACACS+ host and password:

```

Router(config)#tacacs-server host ipaddr
Router(config)#tacacs-server key password

```

After setting up the host, you must define the authentication methods. The following uses TACACS+ as the default authentication but also defines the authentication method **no-tacacs**, which can be used for the console port. Using AAA for the console port is not recommended because if the network is down, you are not able to log on to the box.

```

Router(config)#aaa authentication login default group tacacs+
Router(config)#aaa authentication login no-tacacs line

```

The line parameters can then be modified based on which method you want to use to authenticate:

```

Router(config)#line vty 0 4
Router(config-line)#login authentication default
Router(config)#line con 0
Router(config-line)#login authentication no-tacacs

```

So far, these authentication, authorization, and accounting (AAA) commands have dealt only with authentication. Assume, for example, that you wanted to have a detailed log of every command typed on a router as well as when an administrator logged in or out. The following commands enable TACACS+ accounting for these events:

```

! Enable login and logout tracking for router administrators
Router(config)#aaa accounting exec default start-stop group tacacs+
! Enable command logging for exec level 1 commands (basic telnet)
Router(config)# aaa accounting commands 1 default start-stop group tacacs+
! Enable command logging for exec level 15 commands (enable mode)
Router(config)# aaa accounting commands 15 default start-stop group tacacs+

```

AAA can be very complicated, with many options. For more information about configuring AAA on Cisco devices, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/fsaaa/index/htm.

Secure Shell (SSH)

Use SSH instead of Telnet whenever possible. To configure it, you must first define a hostname and domain name, and generate keys:

```
Router(config)#hostname hostname
Router(config)#ip domain-name yourdomain.com
Router(config)#crypto key generate rsa
```

From here, you can refer to the **transport input** command in [Line Access, page 5-6](#). To set up the VTY lines to accept only SSH, enter the following command:

```
Router(config)#line vty 0 4
Router(config)#transport input ssh
```

There are a few other options with respect to SSH configuration. For more information, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/fothersf/scfssh.htm.

Management Access

This section outlines basic settings for hardening management access, including security settings for the HTTP server, Simple Network Management Protocol (SNMP), Cisco Discovery Protocol (CDP), syslog, Network Time Protocol (NTP), and various ACL logging options.

HTTP Server

If not in use, disable the HTTP server for router management with the following command:

```
Router(config)#no ip http server
```

The embedded web server in routers has had vulnerabilities in the past, so unless you have a specific need for the HTTP functionality (such as a specific management application), it is best to disable it. If you need access to the HTTP server, use the **http access-class** command as shown:

```
Router(config)#ip http access-class 10
Router(config)#access-list 10 permit host http-mgmt-ip
Router(config)#access-list 10 deny any log
```

You should also require HTTP authentication with the following command:

```
Router(config)#ip http authentication ?
enable Use enable passwords
local Use local username and passwords
tacacs Use tacacs to authorize user
```

TACACS+ is preferred; otherwise, a local username and password can be used. Try to avoid using the enable password.

SNMP

SNMP is widely used as a network management protocol. Unfortunately, it is UDP-based (port 161) and, until version 3, had no real security options. Earlier versions of SNMP use a community string for authentication, and it is sent in the clear with the rest of the SNMP datagram. Although version 3 offers more security, most network management applications use SNMP version 1 or version 2c.

In EttF 1.1, you need to enable SNMP if CS-MAR is implemented. If you do not plan to deploy CS-MARS or to manage a device with SNMP, you should disable it:

```
Router(config)#no snmp-server
```

If you must use SNMP v1 or v2c, consider using read-only as opposed to read-write. Much of the damage an attacker can cause with SNMP is prevented if you remove the ability to write changes. In either case, the community string should be set and managed like the root password on any system (change it regularly, and so on). At the bare minimum, an ACL should be defined that allows only your SNMP devices to query the management agents on the network device, as follows:

```
Router(config)#snmp-server community password ro 98
Router(config)#snmp-server community password rw 98
Router(config)#access-list 98 permit host snmp-server-ip
Router(config)#access-list 98 deny any log
```

If you are using SNMP v3 or want more information on the rest of the SNMP configuration, see the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_c/fcfrpt3/fc014.htm.

CDP

CDP is a proprietary Cisco protocol that provides a mechanism for Cisco devices to exchange information. The following two commands show how to globally disable CDP or, alternately, to disable it only on a specific interface:

```
Router(config)#no cdp run
Router(config-if)#no cdp enable
```



Note

One exception is that CDP should remain enabled when a device supports CDP and sends its capabilities information over the protocol to the switch to which the device is attached. This enables the switch to set up the proper configuration on this access port for this endpoint. An example of such a device is the Cisco IP Phone. This is not relevant to the current EttF architecture.

Syslog

Using syslog on a router is one of the easiest ways to troubleshoot your network. Syslog servers are free (besides the hardware), and the messages generated by syslog are usually easy to understand. If you are using any kind of ACLs on a router, you need syslog; even if you are not, it is a very good idea. Enabling syslog is easy. Just enter one or more logging hosts and make sure timestamps are enabled:

```
Router(config)#service timestamps log datetime localtime msec show-timezone
Router(config)#logging syslog-ip-addr
```

Sometimes viewing messages locally on the router can be useful. Besides viewing messages as they are generated on the console, you can optionally have them buffered to router memory. You do not need a larger buffer here because these are simple text messages; even 512 KB saves lots of messages. Be sure you do not use up a significant portion of your device memory, or you might affect packet forwarding. (That is, if you have 8 MB of memory on your router, do not set the buffer size to 6 MB.) Enter the following command to enable this functionality:

```
Router(config)#logging buffered buffersize
```

You can use the **logging trap** command to set the level of logging information you receive; there is no rule for where to set this, except that the highest level of logging is almost always too much information and the lowest level does not provide enough information. Try a few different levels on your own device to determine the amount of information that makes sense in your environment. Syslog has a number of

additional options. For more information, see the following URL:

http://www.cosco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_c/fcfprt3/fcf013.htm#1001168.

NTP

Without proper timestamps, router syslog messages are nearly useless in troubleshooting. Your networking devices can be synchronized to the same clock with NTP. Configuring NTP on a router is a simple matter of locally configuring the time zone and then pointing the router to the NTP server. In the following example, NTP authentication is enabled, and an ACL restricting NTP access to the configured NTP server is applied:

```
Router(config)#clock timezone PST -8
Router(config)#clock summer-time PDT recurring
Router(config)#ntp authenticate
Router(config)#ntp authentication-key 1 md5 password
Router(config)#ntp trusted-key 1
Router(config)#ntp access-group peer 96
Router(config)#ntp server ntp-svr-ip key 1
Router(config)#access-list 96 permit host ntp-svr-ip
Router(config)#access-list 96 deny any log
```

Although there are several free NTP services on the Internet, it is not advisable to use them for security reasons. If your time source is corrupted, your log data is useless. Instead, consider setting up a local time source that connects to a reliable, known atomic clock to maintain accurate time. NTP can be disabled on interfaces that do not expect to receive valid NTP information. Use the following command:

```
Router(config-if)#ntp disable
```

More information on NTP is available at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/ffun_c/fcfprt3/fcf012.htm#1001170.

ACL Options

By default, the last line in an ACL is an implicit deny all. Matches to this list are not logged, however. If you want to enable logging, a manual entry should be added to the ACL denying all traffic and informing the ACL to log the violation. It is possible to log permits as well, but this tends just to fill up a syslog server. To drop all traffic and log violations in a standard IP ACL, use the following command:

```
Router(config)#access-list 1 deny any log
```

For an extended IP ACL, use this command:

```
Router(config)#access-list 101 deny ip any any log
```

In addition to the basic log keyword, log input is usually available for extended ACLs. Log input adds the source interface and MAC address to the usual IP address and port number message associated with the ACL entry.



Note

After hardening a router, it is a good idea to scan it with your favorite port scanner. This ensures that you are not running any services you thought you turned off.

Layer 2 Security Design

Unlike hubs, switches are able to regulate the flow of data between their ports by creating almost “instant” networks that contain only the two end devices communicating with each other at that moment in time. Data frames are sent by end systems, and their source and destination addresses are not changed throughout the switched domain. Switches maintain content-addressable memory (CAM) lookup tables to track the source addresses located on the switch ports. These lookup tables are populated by an address-learning process on the switch. If the destination address of a frame is not known or if the frame received by the switch is destined for a broadcast or multicast address, the switch forwards the frame out all ports, except for the port that the frame entered to the switch. With their ability to isolate traffic and create the “instant” networks, switches can be used to divide a physical network into multiple logical or virtual LANs (VLANs) through the use of Layer 2 traffic segmentation. In general, Layer 2 of the OSI reference model is subject to network attacks in unique ways that include the following:

- Vulnerability of the use of VLAN 1
- Spanning tree attack
- MAC flooding attack
- VLAN hopping
- 802.1Q tagging attack
- ARP attacks
- MAC spoofing attack
- DHCP starvation attack
- Rogue DHCP server attack

In EttF 1.1, the implementation of Layer 2 security protection is needed on all switches (that is, Catalyst 3750 and Catalyst 2955) in the following network areas:

- Cell/area zone
- Server farm in the manufacturing zone
- Server farm in the DMZ

Equally important is that all switches need hardening. [Network Device Hardening, page 5-3](#) discussed how to harden network devices. This section focuses on Layer 2 security for these devices.



Note

For details on how to protect the Catalyst 3750 against L3 security threats from the manufacturing zone perspective, see [Security Design for the Manufacturing Zone, page 5-19](#).

[Table 5-2](#) summarizes the L2 vulnerabilities for which these platforms can provide protection.

Table 5-2 Layer 2 Security Threats and Switch Protection

L2 Attacks/Vulnerabilities	Catalyst 3750	Catalyst 2955
Vulnerability of the use of VLAN 1	Yes	Yes
Trust level of switch ports	Yes	Yes
Spanning tree attack	Yes	Yes
MAC flooding attack (Port Security)	Yes	Yes
Broadcast/multicast storm control	Yes	Maybe

Table 5-2 Layer 2 Security Threats and Switch Protection

VLAN hopping	Yes	Yes
ARP attack	Yes	Lack of Dynamic ARP Inspection feature support
MAC spoofing attack	Yes	Yes
DHCP starvation attack	Yes	Yes
Rogue DHCP server attack	Yes	Yes

Precautions for the Use of VLAN 1

The reason VLAN 1 became a special VLAN is that L2 devices needed to have a default VLAN to assign to their ports, including their management port(s). In addition to that, many L2 protocols such as Cisco Discovery Protocol (CDP), Port Aggregation Protocol (PAgP), and VLAN Trunking Protocol (VTP) needed to be sent on a specific VLAN on trunk links. For all these purposes VLAN 1 was chosen.

As a consequence, VLAN 1 may sometimes end up unwisely spanning the entire network if not appropriately pruned and, if its diameter is large enough, the risk of instability can increase significantly. In addition, the practice of using a potentially omnipresent VLAN for management purposes puts trusted devices to higher risk of security attacks from untrusted devices that by misconfiguration or pure accident gain access to VLAN 1 and try to exploit this unexpected security hole.

To redeem VLAN 1 from its bad reputation, a simple common sense security principle can be used: as a general security rule, the network administrator should prune any VLAN, and in particular VLAN 1, from all the ports where that VLAN is not strictly needed.

Therefore, with regard to VLAN 1, the above rule simply translates into the following recommendations:

- Do not use VLAN 1 for inband management traffic; preferably pick a different, specially-dedicated VLAN that keeps management traffic separate from industrial Ethernet and other user data traffic.
- Prune VLAN 1 from all the trunks and from all the access ports that do not require it (including not connected and shutdown ports).

As a general design rule, it is desirable to prune unnecessary traffic from particular VLANs. For example, it is often desirable to apply VLAN ACLs and/or IP filters to the traffic carried in the management VLAN to prevent all Telnet connections and to allow only SSH sessions. Alternatively, it may be desirable to apply QoS ACLs to rate limit the maximum amount of ping traffic allowed.

If VLANs other than VLAN 1 or the management VLAN represent a security concern, automatic or manual pruning should be applied as well. In particular, configuring VTP in transparent or off mode is commonly considered as the most effective method:

```
Switch(config)#vtp mode transparent
```



Note

The discussion of having more than one VLAN applies only to the server farm in the manufacturing zone in EttF 1.1. The cell/area zone in this solution phase is assumed to have only one VLAN.

Trust Level of Switch Ports

After proper handling of VLAN 1 has been decided on and implemented, the next logical step is to consider other equally important best practices commonly used in secure environments. The general security principle applied here is to connect untrusted devices to untrusted ports, trusted devices to trusted ports, and to disable all the remaining ports.

The recommendations are as follows:

- If a port on a Catalyst switch in the cell ring is connected to a “foreign” device, such as a drive, HMI, I/O, PAC, or historian, make sure to disable CDP, DTP, PAgP, UDLD, and any other unnecessary protocol, and to enable switch port mode access, PortFast, and BPDU Guard on it, as in the following example:

```
Switch(conf)#vtp mode transparent
Switch(conf)#interface type/slot port
Switch(config-if)#switchport access vlan vlan number
Switch(config-if)#switchport mode access
Switch(config-if)#no cdp enable
Switch(config-if)#spanning-tree portfast
Switch(config-if)#spanning-tree bpduguard enable
Switch(config-if)#spanning-tree bpdufilter enable
```

- Enable Root Guard on the Catalyst 3750 interfaces to which the cell ring is connected. This prevents a directly or indirectly connected STP-capable device from affecting the Catalyst 3750 being the root bridge:

```
Switch(config)#interface type/slot port
Switch(config-if)# spanning-tree guard root
```

- Configure the VTP domains appropriately or turn off VTP altogether if you want to limit or prevent possible undesirable protocol interactions with regard to network-wide VLAN configuration. This precaution can limit or prevent the risk of an administrator error propagating to the entire network, and the risk of a new switch with a higher VTP revision overwriting by accident the VLAN configuration of the entire domain.

```
Switch(conf)#vtp mode transparent
```

- By default, all the LAN ports on all the Catalyst switches are configured as “untrusted”. This prevents attached devices from manipulating QoS values inappropriately. In the EttF 1.1 design, only the trunk ports are recommended to be set as “trusted” if QoS is ever implemented in the network. All the access ports (for example, those on the Catalyst 2955) should remain “untrusted.”

```
Switch(conf)#interface type/slot port
Switch(config-if)#no mls qos trust
```

- Disable unused ports and put them in an unused VLAN. By not granting connectivity or by placing a device into a VLAN not in use, unauthorized access can be prevented by fundamental physical and logical barriers.

```
Switch(conf)#interface type/slot port
Switch(config-if)#shutdown
```

Spanning Tree Protocol Security

STP is a useful protocol, but it does not implement any authentication and encryption to protect the exchange of Bridge Protocol Data Units (BPDUs). Because of the lack of authentication, anyone can speak to an STP-enabled device. An attacker could very easily inject fraudulent BPDUs, triggering a topology recalculation. A forced change to the STP topology could lead to a DoS condition, or leave the attacker as a man-in-the-middle. In addition, because BPDUs are not encrypted, it is fairly simple to intercept BPDUs in transit, revealing important topology information.

Catalyst 3750 and 2955 Series switches support a set of features that help protect bridged networks using the Spanning Tree Protocol. The following are the recommended best practices:

- Disable VLAN auto-negotiated trunking on user ports
- Disable unused ports and put them into an unused VLAN (as explained in the previous section)

- Use Per-VLAN Spanning Tree (PVST)
- Implement Port Security (as explained in a subsequent section)
- Configure BPDU Guard
- Configure STP Root Guard

Disabling Auto-negotiated Trunking

By default, all Ethernet ports on Catalyst switches are set to auto-negotiated trunking mode, which allows switches to automatically negotiate ISL and 802.1Q trunks. The negotiation is managed by Dynamic Trunking Protocol (DTP). Setting a port to auto-negotiated trunking mode makes the port willing to convert the link into a trunk link, and the port becomes a trunk port if the neighboring port is set as a trunk or configured in desirable mode.

Although the auto-negotiation of trunks facilitates the deployment of switches, somebody can take advantage of this feature and easily set up an illegitimate trunk. For this reason, auto-negotiation trunking should be disabled on all ports connecting to end users.

To disable auto-negotiated trunking, use the **switchport mode access** command. Setting the port mode to **access** makes the port a nontrunking, nontagged single VLAN Layer 2 interface. The following example shows how to set a port as nontrunking, nontagged single-VLAN Layer-2:

```
Switch(config)# interface type slot/port
Switch(config-if)# switchport mode access vlan 10
Switch(config-if)#
```

BPDU Guard

BPDU Guard is a feature that prevents a host port from participating in spanning tree. Under normal circumstances, Layer 2 access ports connected to a single workstation or server should not participate in spanning tree. When enabled on a port, BPDU Guard shuts down the port as soon as a BPDU is received in that port. In this way, BPDU Guard helps prevent unauthorized access and the illegal injection of forged BPDUs.

BPDU Guard requires STP PortFast to be configured on the port first. STP PortFast causes a Layer 2 LAN port configured as an access port to enter the forwarding state immediately, bypassing the listening and learning states. PortFast can be used on Layer 2 access ports connected to a single workstation or server to allow those devices to connect to the network immediately, instead of waiting for STP to converge.

BPDU can be configured per port or globally. When configured globally, BPDU Guard is effective only on ports in the operational PortFast state.

To enable BPDU Guard on an interface, use the **spanning-tree bpduguard** command. Make sure to first enable PortFast on the port.

```
Switch(config)# interface type/slot port
Switch(config-if)# spanning-tree portfast
Switch(config-if)# spanning-tree bpduguard enable
```

BPDU Guard can be globally enabled on systems running Cisco IOS by using the **spanning-tree portfast bpduguard default** command. When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast state:


```
Switch(config)# spanning-tree portfast bpduguard
```


STP Root Guard

STP Root Guard is a feature that enforces the placement of the root bridge. STP Root Guard is a feature that is enabled on selected ports to prevent surrounding switches from becoming the root switch. The Root Guard feature forces a port to become a designated port so that no switch on the other end of the link can become a root switch. If a port configured for Root Guard receives a superior BPDU, the port immediately goes into a root-inconsistent (blocked) state. In this way, STP Root Guard blocks other devices trying to become the root bridge by sending superior BPDUs.

Figure 5-1 illustrates the placement of the STP Root Guards in the ring topology.

Figure 5-1 Placement of STP Guards



Note

Do not enable Loop Guard and Root Guard on a port at the same time. Root Guard forces a port to always be designated as the root port. Loop Guard is effective only if the port is a root port or an alternate port.

To enable STP Root Guard on an interface, use the **spanning-tree guard root** command. Make sure to first enable PortFast on the port. The following example shows how to enable STP Root Guard on an interface:

```
Switch(config)# interface type/slot port
Switch(config-if)# spanning-tree guard root
```

MAC Flooding Attack

All switches have a finite hardware learning table to store the source addresses of all received packets; when this table becomes full, the traffic that is directed to addresses that can no longer be learned is permanently flooded. Packet flooding, however, is constrained within the VLAN of origin; therefore, no VLAN hopping is permitted.

One corner case behavior can be exploited by a malicious user that wants to turn the switch to which the user is connected into a dumb pseudo-hub and sniff all the flooded traffic. On non-intelligent switches, this problem arises because the L2 identity of a sender is not checked; therefore, the sender is allowed to impersonate an unlimited number of devices simply by counterfeiting packets. Cisco switches support a variety of features whose only goal is to identify and control the identities of connected devices. The security principle on which they are based is very simple: authentication and accountability are critical for all untrusted devices, including PACs, I/Os, drives, and human-machine interfaces (HMIs) attached to a switch in the cell network.

Port Security can be used to constrain the connectivity of a device. With Port Security, preventing any MAC flooding attack becomes as simple as limiting the number of MAC addresses that can be used by a single port; the identification of the traffic of a device is thereby directly tied to its port of origin.

In ETT 1.1, Cisco recommends enabling Port Security on an access port (not a trunk port that is used to form the L2 backbone network) and to set the maximum number of secure addresses to 1. The violation mode is the default; no static secure MAC addresses are configured.

```
Switch(config)# interface type/slot port
Switch(config-if)# switchport mode access vlan vlan number
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 1
Switch(config-if)# switchport port-security violation restrict
Switch(config-if)# end
```

It is a security violation when one of the following situations occurs:

- The maximum number of secure MAC addresses have been added to the address table, and a station whose MAC address is not in the address table attempts to access the interface.
- An address learned or configured on one secure interface is seen on another secure interface in the same VLAN.

Packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. Specifically, an SNMP trap is sent, a syslog message is logged, and the violation counter increments.

More information about the feature can be found at the following URL:

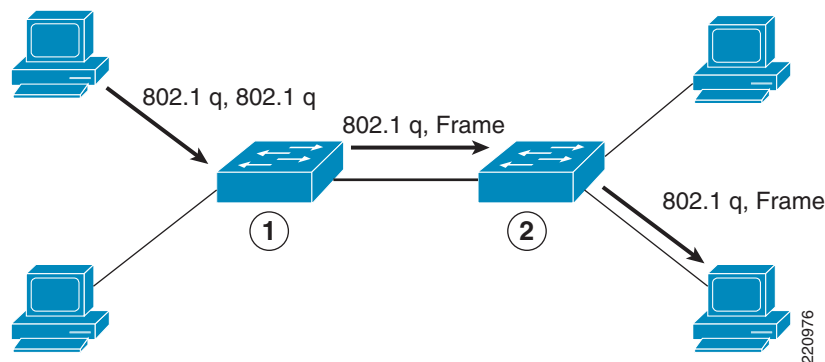
http://www.cisco.com/en/US/products/hw/switches/ps628/products_configuration_guide_chapter09186a0080648dd5.html.

VLAN Hopping

Tagging attacks are malicious schemes that allow a user on a VLAN to get unauthorized access to another VLAN. For example, if a switch port is configured as DTP auto and receives a fake DTP packet, it might become a trunk port and it might start accepting traffic destined for any VLAN. Therefore, a malicious user can start communicating with other VLANs through that compromised port.

Another version of this network attack is called double tagging, and involves tagging the transmitted frames with two 802.1q headers to forward the frames to the wrong VLAN (see [Figure 5-2](#)).

Figure 5-2 VLAN Hopping with Double-Encapsulated 802.1q Traffic



The first switch to encounter the double-tagged frame (1) strips the first tag off the frame and forwards the frame. The result is that the frame is forwarded with the inner 802.1q tag out all the switch ports (2), including trunk ports configured with the native VLAN of the network attacker. The second switch then forwards the packet to the destination based on the VLAN identifier in the second 802.1q header.

VLAN hopping attack can be prevented by setting DTP to “off” on all non-trusted ports:

- If you do not intend to trunk across those links, use the **switchport mode access interface** configuration command to disable trunking.

```
Switch(config)# interface type/slot port
Switch(config-if)# switchport mode access
```

- To enable trunking to a device that does not support DTP, use the **switchport mode trunk** and **switchport nonegotiate** interface configuration commands to cause the interface to become a trunk but to not generate DTP frames.

```
Switch(config)# interface type/slot port
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport nonegotiate
```

Sometimes, even when simply receiving regular packets, a switch port may behave like a full-fledged trunk port (for example, accepting packets for VLANs different from the native), even if it is not supposed to do so. This is commonly referred to as “VLAN leaking”. Fortunately, the Catalyst switches have been designed in their hardware and software to always enforce proper traffic classification and isolation on all their ports.

ARP Spoofing Attack

Address Resolution Protocol (ARP) is used to map IP addressing to MAC addresses in a LAN segment where hosts of the same subnet reside. Normally, a host broadcasts an ARP request to find the MAC address of another host with a particular IP address, and an ARP response comes back from the host whose address matches the request. The requesting host then caches this ARP response. Within the ARP protocol, another provision is made for hosts to perform unsolicited ARP replies. The unsolicited ARP replies are called gratuitous ARPs (GARPs). GARPs can be exploited maliciously by an attacker to spoof the identity of an IP address on a LAN segment. Typically, this is used to spoof the identity between two hosts or all traffic to and from a default gateway in a man-in-the-middle attack.

By crafting an ARP reply, a network attacker can make their system appear to be the destination host sought by the sender. The ARP reply causes the sender to store the MAC address of the system of the network attacker in the ARP cache. This MAC address is also stored by the switch in its CAM table. In

this way, the network attacker has inserted the MAC address of their system into both the CAM table of the switch and the ARP cache of the sender. This allows the network attacker to intercept frames destined for the host being spoofed.

The use of DHCP snooping along with Dynamic ARP Inspection (DAI) mitigates various ARP-based network exploits. These Catalyst features validate ARP packets in a network and permit the interception, logging, and discarding of ARP packets with invalid MAC address to IP address bindings.

DHCP snooping provides security by filtering trusted DHCP messages and then using these messages to build and maintain a DHCP snooping binding table. DHCP snooping considers DHCP messages originating from any user-facing port that is not a DHCP server port or an uplink to a DHCP server as untrusted. From a DHCP snooping perspective, these untrusted, user-facing ports should not send DHCP server-type responses such as DHCP Offer, DHCP Ack, or DHCP Nak.

The DHCP snooping binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information corresponding to the local untrusted interfaces of a switch; it does not contain information regarding hosts interconnected with a trusted interface. An untrusted interface is an interface configured to receive messages from outside the network or firewall. A trusted interface is an interface that is configured to receive only messages from within the network. The DHCP snooping binding table can contain both dynamic as well as static MAC address to IP address bindings.

DAI determines the validity of an ARP packet based on the valid MAC address to IP address bindings stored in a DHCP snooping database. Additionally, DAI can validate ARP packets based on user-configurable ACLs. This allows for the inspection of ARP packets for hosts using statically configured IP addresses. DAI allows for the use of per-port access control lists (PACLs) and VLAN access control lists (VACLs) to limit ARP packets for specific IP addresses to specific MAC addresses.

```
Switch(config)# ip dhcp snooping
Switch(config)# ip dhcp snooping vlan vlan_id
Switch(config)# ip arp inspection vlan vlan_id
Switch(config)# ip arp inspection validates src-mac dst-mac ip
Switch(config)# interface type slot/port
Switch(config-if)# ip dhcp snooping trust
Switch(config-if)# ip dhcp snooping limit rate rate
Switch(config-if)# ip arp inspection trust
```

DHCP Attacks

There are two common types of DHCP attacks: DHCP starvation attack and rogue DHCP server attack.

A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses. This is easily achieved with attack tools such as Gobbler. If enough requests are sent, the network attacker can exhaust the address space available to the DHCP servers for a period of time. The attack can be mitigated by configuring Port Security on the Catalyst switch as described in [MAC Flooding Attack, page 5-16](#).

In a rogue DHCP server attack, the attacker sets up a rogue DHCP server on their system and responds to new DHCP requests from clients on the network. The network attacker can provide clients with addresses and other network information. Because DHCP responses typically include default gateway and DNS server information, the network attacker can supply their own system as the default gateway and DNS server, resulting in a man-in-the-middle attack.

Use the following commands to mitigate these attacks:

```
Switch(config)# ip dhcp snooping
Switch(config)# ip dhcp snooping vlan vlan number
Switch(config)# ip dhcp snooping information option
```

Security Design for the Manufacturing Zone

Because the security design strategy of the manufacturing zone is identical to that of the enterprise campus network, this section simply provides description of the required best practices. References are provided for their detailed implementation.

Security Design for the Catalyst 3750 Series Switch That Aggregates Cell/Area Zone Networks

Note the following:

- Device hardening (see [Network Device Hardening, page 5-3](#))
- Layer 2 security for L2 ports (see [Layer 2 Security Design, page 5-11](#))
- Ingress/egress filtering—RFC 1918 and RFC 2827 filtering should be implemented to protect against spoofed denial-of-service (DoS) attacks (http://www.cisco.com/en/US/tech/tk59/technologies_white_paper09186a0080174a5b.shtml).
- Routing protocol authentication—This is to prevent an attacker from sharing incorrect routing information between a rogue router and a valid one. The intent of the attack is to trick the router into not only sending data to the incorrect destination but also possibly to put it out of service. The recommended method is to check the integrity of routing updates by authentication using MD5-HMAC. See the following URLs:
 - 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

Security Design for the Catalyst 4500 Series Switch for the Core of the Control Network

The Catalyst 4500 is assumed to provide only L3 routing. Thus, only device hardening and L3 security best practices are needed. Note the following:

- Device hardening (see [Network Device Hardening, page 5-3](#))
- Ingress/egress filtering—RFC 1918 and RFC 2827 filtering should be implemented to protect against spoofed DoS attacks (http://www.cisco.com/en/US/tech/tk59/technologies_white_paper09186a0080174a5b.shtml).
- Router with ACL—The Catalyst 4500 should be configured to block traffic flows at L3/L4 based on your trust domains and security policies.

- Routing protocol authentication—This is to prevent an attacker from sharing incorrect routing information between a rogue router and a valid one. The intent of the attack is to trick the router to not only send data to the incorrect destination but to also possibly put it out of service. The recommended method is to check the integrity of routing updates by authentication using MD5-HMAC.
 - 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
- Control Plane Policing (CoPP)—This feature protects the CPU from unnecessary or DoS traffic by giving priority to important control plane and management traffic. The idea is to protect most of the CPU bound traffic and ensure routing stability, reachability, and packet delivery. Most importantly, CoPP is often used to protect the CPU from the DoS attack. For more information, see the following URL:
http://www.cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008062ce7b.html.

**Note**

The following URL provides more in-depth discussion about infrastructure protection for the Catalyst 4500:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a0080747eb2.pdf

Security Design for the Catalyst 3750 Series Switch in the Server Farm

This Catalyst 3750 switch provides VLAN separation for different servers for network services, network management, and site manufacturing operations and control. Because the Catalyst 4500 performs routing for these VLANs, this Catalyst 3750 provides only Layer 2 switching. Thus, its security protection simply includes device hardening and L2 security (see [Network Device Hardening, page 5-3](#) and [Layer 2 Security Design, page 5-11](#)).

**Note**

The following URLs provide more in-depth discussion about infrastructure protection for the Catalyst 3750:
http://www.cisco.com/application/pdf/en/us/guest/products/ps1838/c1244/cdcont_0900aec804ac831.pdf

Security Protection for Servers

The servers that provide network services, network management, or site manufacturing operations and control should be provided at least with the following security protection:

- Reusable passwords—Users likely authenticate to their systems with username and passwords.
- Session-application crypto—Any communication between a client to a server considered sensitive (based on your 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 production 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, you do not want 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](#).

Security Design for the Demilitarized Zone

In the design of the industrial Ethernet network, one of the critical elements is to ensure the separation between the control network and enterprise network. In terms of the Purdue Reference Model, this is the separation between levels 1–3 and levels 4–5. This separation is necessary because real-time availability and security are the critical elements for the traffic in the control network. You do not want enterprise traffic that has very different traffic characteristics to enter the control network and cause any disruption to the ongoing operations. Acting as a firewall, the Cisco ASA5500 provides this separation of the two networks.

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. To provide more granular network access, the Cisco ASA provides authentication, authorization, and accounting (AAA) services by working in conjunction with the CiscoSecure 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 networks, the Cisco ASA can optionally be installed with the Cisco Adaptive Inspection Prevention Security Services Module (AIP-SSM) to provide intrusion detection or intrusion protection 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. See [Security Protection for Servers, page 5-21](#).

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 1.1 solution, Cisco recommends creating three networks in different security levels, as shown in [Table 5-3](#).

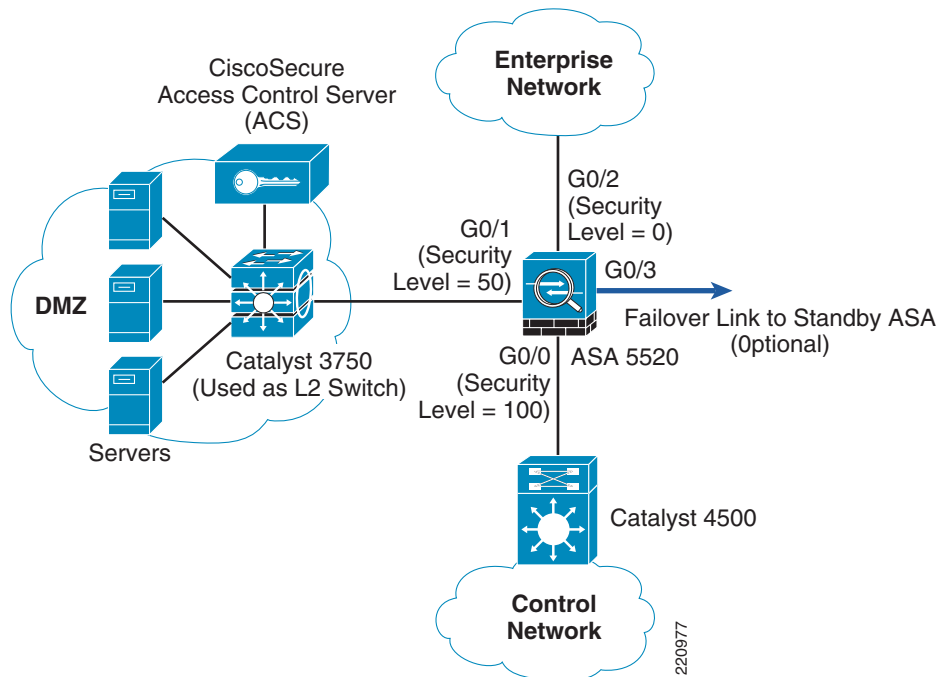
Table 5-3 Network Security Levels

Network	Security Level	Interface (see Figure 5-3)
Enterprise network	0	G0/2
DMZ	50	G0/1
Control network	100	G0/0

Configuration Example

Refer to [Figure 5-3](#) for the subsequent configuration example.

Figure 5-3 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 control 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 control 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 **nameif** 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 E2F 1.1 solution recommends that traffic be denied going from the control 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 control network and enterprise network enables the operator to define 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:

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.
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.
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**

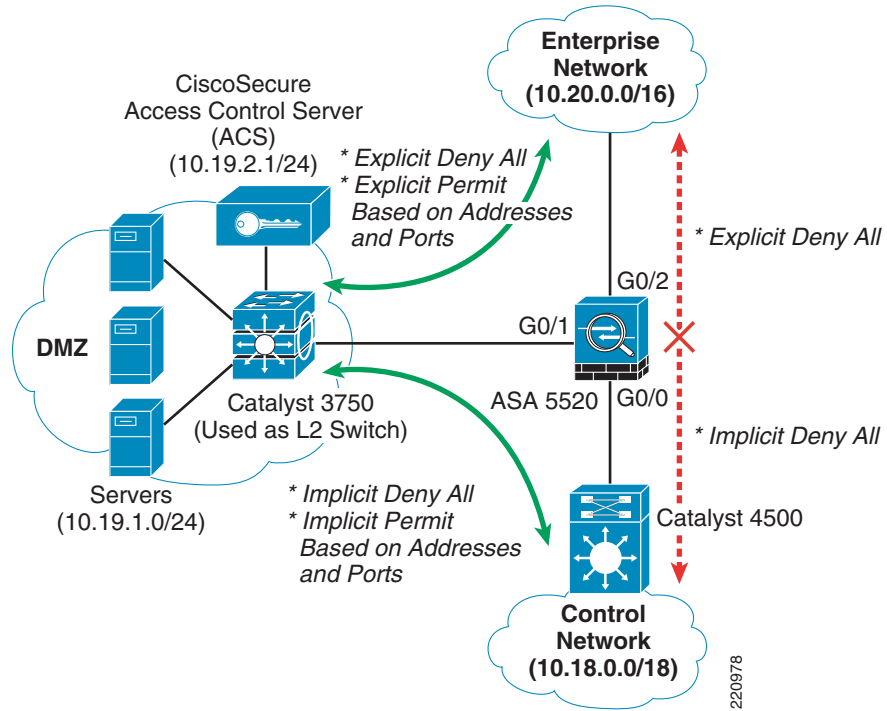
The interface ACL does not block packets destined for the IP addresses of the security appliance.

For the EttF 1.1 solution, general packet filtering recommendations are listed in [Table 5-4](#) and shown in [Figure 5-4](#).

Table 5-4 Packet Filtering Recommendations

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

Figure 5-4 High-Level Packet Filtering Recommendations for the DMZ between the Control and Enterprise Networks



Configuration Example

See [Table 5-5](#) for an example for ingress ACLs applied to the control network-facing interface.

Table 5-5 Configuration Example for Ingress ACLs on the Control Networking-Facing Interface

Applied To Interface	Traffic Direction	Permitted Traffic Types (Source to Destination)
Interface connected to the control network (<i>inside</i>)	Inbound	<ul style="list-style-type: none"> HTTP (servers in the control network to servers in DMZ) <pre>access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq www</pre> HTTPS (any in the control network to servers in DMZ) <pre>access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq https</pre> Telnet (any in the control network to host 10.19.1.10 in the DMZ) <pre>access-list inside extended permit tcp 10.18.0.0 255.255.0.0 host 10.19.2.1 eq telnet</pre> ICMP (any in the control network to servers in the DMZ) <pre>access-list inside extended permit icmp 10.18.0.0 255.255.0.0 10.19.2.0 255.255.255.0</pre> Explicitly deny other traffic types to anywhere (i.e. DMZ and enterprise networks) <pre>access-list inside deny 10.18.0.0 255.255.0.0</pre> Apply the ACLs above to the ingress side of the control network-facing interface <pre>access-group inside in interface inside</pre>

See [Table 5-6](#) for an example for ingress ACLs applied to the DMZ-facing interface.

Table 5-6 Configuration Example for Ingress ACLs on the DMZ -Facing Interface

Applied To Interface	Traffic Direction	Permitted Traffic Types (Source to Destination)
Interface connected to the DMZ (<i>dmz</i>)	Inbound	<ul style="list-style-type: none"> • Telnet (servers in the DMZ to the control and enterprise networks) <pre>access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq telnet</pre> • HTTP (servers in the DMZ to the control and enterprise networks) <pre>access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq www</pre> • HTTPS (servers in the DMZ to the control and enterprise networks) <pre>access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq https</pre> • ICMP (servers in the DMZ to the control and enterprise networks) <pre>access-list dmz extended permit icmp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 access-list dmz extended permit icmp 10.19.1.0 255.255.255.0 10.20.0.0 255.255.0.0</pre> • Explicitly deny other traffic types to anywhere <pre>access-list inside deny 10.19.0.0 255.255.0.0</pre> • Apply the ACLs above to the ingress side of the DMZ-facing interface <pre>access-group dmz in interface inside</pre>

See [Table 5-7](#) for the example for ingress ACLs applied to the enterprise network-facing interface.

Table 5-7 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 (<i>outside</i>)	Inbound	<ul style="list-style-type: none"> Telnet (any in the enterprise network to the DMZ [10.19.0.0/16]) <pre>access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq telnet</pre> HTTP (any in the enterprise network to the DMZ [10.19.0.0/16]) <pre>access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq www</pre> HTTPS (any in the enterprise network to the DMZ [10.19.0.0/16]) <pre>access-list outside extended permit tcp 10.20.0.0 255.255.0.0 10.19.1.0 255.255.255.0 eq https</pre> Explicitly deny other traffic types to anywhere <pre>access-list inside deny 10.20.0.0 255.255.0.0</pre> Apply the ACLs above to the ingress side of the enterprise network-facing interface <pre>access-group outside in interface inside</pre>

Authenticating Firewall Sessions for User Access to Servers in the DMZ

When users in the control 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:

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.
2. The Cisco ASA prompts the user for authentication.
3. The Cisco ASA receives the authentication information (userid and password) from the user and sends an AUTH Request to the CiscoSecure ACS.

4. The server authenticates the user and sends an AUTH Accept message to the Cisco ASA.
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 plant floor network. Factory 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 control floor 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 control 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](http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_configuration_guide_chapter09186a00807517ba.html).



### Note

The Cisco ASA 5520, which is the ASA model recommended for the EtfF 1.1 design, 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](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).



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 5-5 shows the traffic flow when the Cisco ASA is configured in inline IPS mode.

**Figure 5-5** Inline IPS Traffic Flow

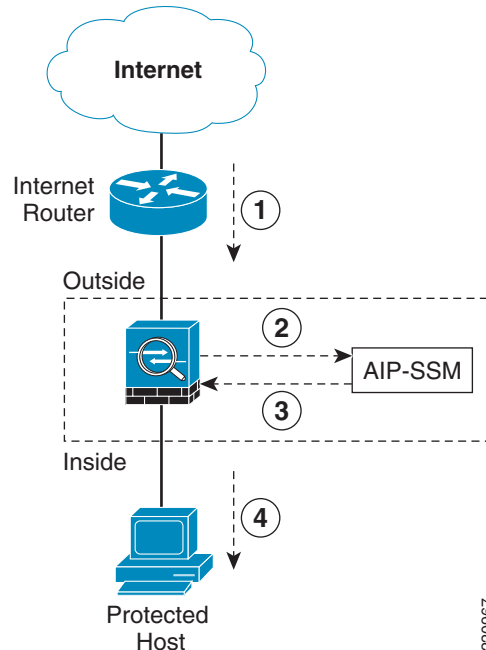


Figure 5-5 shows the following sequence of events:

1. The Cisco ASA receives an IP packet from the Internet.
2. Because the Cisco ASA is configured in inline IPS mode, it forwards the packet to the AIP-SSM for analysis.
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.
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 5-6 shows an example of how traffic flows when the AIP-SSM is configured in promiscuous mode.

**Figure 5-6 Promiscuous Mode Traffic Flow**

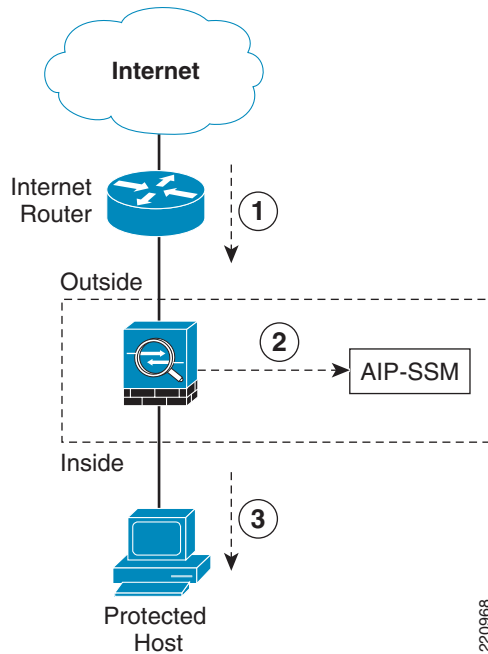


Figure 5-6 shows the following sequence of events:

1. The Cisco ASA receives an IP packet from the Internet.
2. Because the Cisco ASA is configured in promiscuous mode, the AIP-SSM silently snoops the packet.
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 AIP-SSM. This applies to both inline and promiscuous IPS modes.

# 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.

## 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 ContextCorrelation, 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 customer-centric rule creation, threat notification, incident investigation, and a host of security posture and trend reports.

The entire solution is cost-effectively delivered in an appliance platform that affords low adoption costs and flexible use. 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 with various performance characteristics and price points to address a broad range of customer sizes and deployment scenarios.





## CHAPTER 6

# Implementation of High Availability

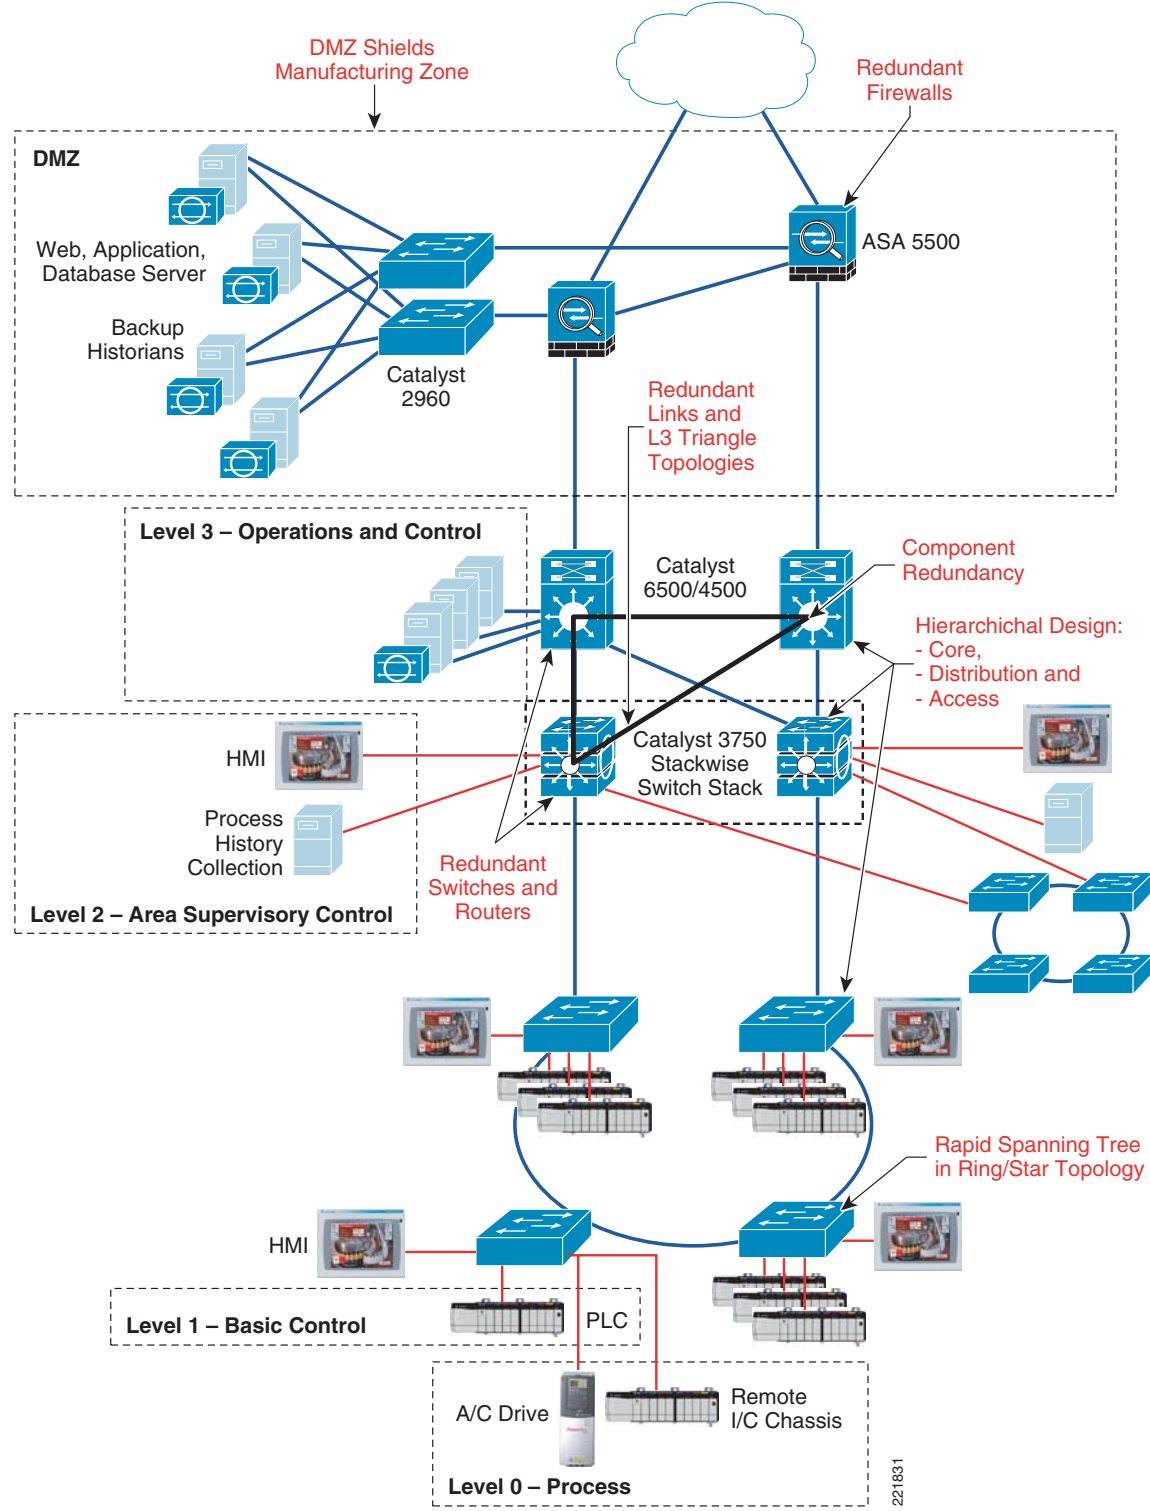
---

High availability (HA) is not a standalone feature, but instead an approach to implementing a variety of interrelated features as tools to ensure business resilience and maintain end-to-end availability for services, users, devices, and applications. High availability should be incorporated at many layers. With a sound design, network stability is easy to achieve, troubleshooting is easier, and human error is reduced. Key aspects of an HA network include the following:

- Hierarchical network design based on both the EttF logical architecture and the enterprise campus core-distribution-access model
- Network and component redundancy that includes both redundant network systems, links, and systems with redundant components
- Foundation services that apply the network software features to maintain network availability when links, components, or other failures occur

Figure 6-1 shows the key high availability features that Cisco recommends for the EttF solution architecture.

Figure 6-1 Recommended High Availability Features



221831

Cisco HA is technology delivered in Cisco IOS Software that enables network-wide resilience to increase network availability. Network applications must cross different network layers from the access, distribution, core, and DMZ. High availability in a manufacturing environment consists of both network resiliency and system resiliency, which when combined result in transparent fault detection and recovery to the user community. An unscheduled network failure that is not resolved can result in termination, interruption, or violation of service-level agreements (SLAs) for manufacturing business-critical applications.

This chapter includes the following topics:

- Overall benefits of an HA design
- Best practices and HA modeling
- Design considerations and best practices for HA in the cell/area zone
- Design considerations and best practices for HA in the manufacturing zone
- Design considerations and best practices for HA in the DMZ

For an overview of general HA topics, see *Designing a Campus Network for High Availability* at the following URL:  
[http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdcont\\_0900aecd801a8a2d.pdf](http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdcont_0900aecd801a8a2d.pdf).

## Benefits of an HA Design

Unscheduled downtime has many costs. On a manufacturing shop floor, downtime can result in revenue losses that directly affect the bottom line. Table 6-1 shows availability percentages and downtime per year.

**Table 6-1**      **Availability and Downtime**

| Availability | Defects Per Million | Downtime Per Year   |
|--------------|---------------------|---------------------|
| 99.9000%     | 1000                | 8 hours, 46 minutes |
| 99.9500%     | 500                 | 4 hours, 23 minutes |
| 99.9900%     | 100                 | 53 minutes          |
| 99.9990%     | 10                  | 5 minutes           |
| 99.9999%     | 1                   | 30 seconds          |

A highly available network design improves SLA support, reduces unplanned downtime, and increases operational efficiencies. High availability is a function of the application as well as the end-to-end network connectivity between a factory floor node and a specific service. These services can be device communication within the cell/area zone, or application traffic to or from the manufacturing zone and DMZ. With an effective HA design, a more deterministic network is available, which influences overall network availability. Key considerations are the mean time between failures (MTBF) and mean time to repair (MTTR).

## Best Practices and HA Modeling

For the network to be deterministic, the design must be as simple and highly structured as possible. This is achieved by implementing a network hierarchy. Recovery mechanisms must be considered as part of the design process. Recovery timing is determined in part by the nature of the failure; for example, total device failure, direct link failure, indirect link failure, and so on. Several key components and design concepts are examined in the following sections at each network layer.

Before implementing an HA solution, it is wise to use a modeling exercise, for both the network infrastructure and the network connections, to validate the architecture/design, justify costs, and analyze design tradeoffs.

Device modeling concerns include each field-replaceable unit as well as critical components such as supervisors, power supplies, and line cards. Device modeling includes the following steps:

1. Describing the system and traffic path
2. Building a system block diagram
3. Describing failure scenarios
4. Running the model

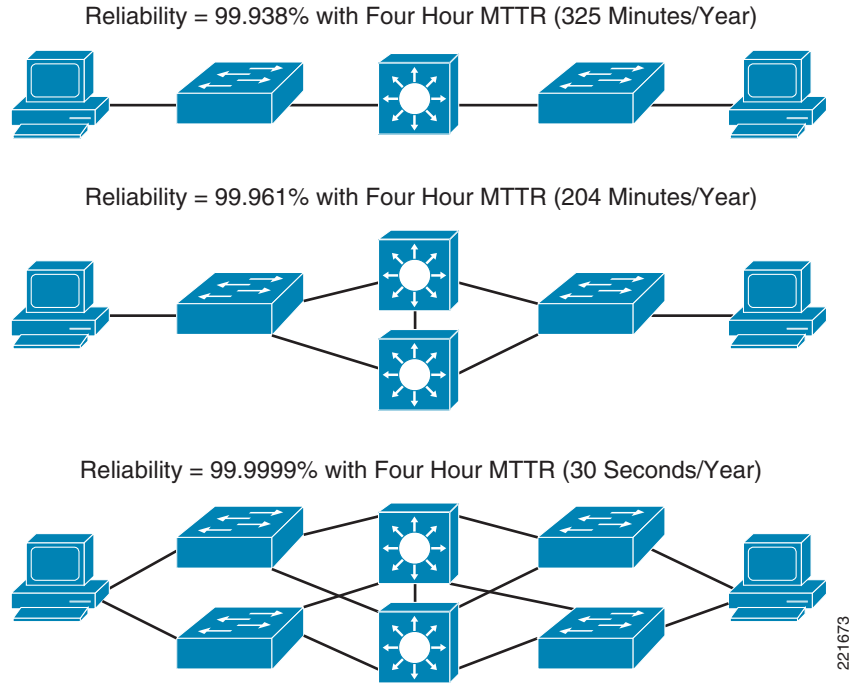
Network modeling is concerned with network resiliency, which involves redundant links and alternate paths. Network modeling includes the following steps:

1. Describing the system and traffic path
2. Building a system block diagram
3. Describing failure scenarios
4. Running the device model from above
5. Running the network model

See [Figure 6-2](#) for an example of network reliability models with an MTTR of four hours.



**Figure 6-2 Network Reliability Models**



## HA Design in the Cell/Area Zone

The cell/area zone devices, Cisco Catalyst 2955 Series switches, rely mainly on a redundant network design (star or ring) to achieve high availability. These are the Layer 2-only access devices in the Cisco core-distribution-access model, and care must be taken to achieve some resiliency at this level because this is the first network connection point from the end-node perspective. (See [Table 6-2.](#))

**Table 6-2 Summary of Features in the Cell/Area Zone**

| Feature                      | Description                                                                                                                                                                                                                                                                          |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Redundant power supplies     | Each Cisco Catalyst 2955 can have an external power supply installed.                                                                                                                                                                                                                |
| Redundant paths              | Depends on topology (either star or ring).                                                                                                                                                                                                                                           |
| StackWise                    | Uses stack interconnect cables to create a virtual switch fabric for stacks of the Catalyst 3750.                                                                                                                                                                                    |
| Broadcast storm control      | Storm control uses rising and falling thresholds to block and then restore the forwarding of broadcast, unicast, or multicast packets. You can also set the switch to shut down the port when the rising threshold is reached. This is good for STP misconfiguration or a bad cable. |
| Rapid Spanning Tree Protocol | Layer 2 protocol to prevent loops in the network. This mode of spanning tree is preferred for the best convergence times if a failure occurs.                                                                                                                                        |

# HA Design and Implementation in the Manufacturing Zone

Depending on the topology choice, either two access layer switches (ring) or all access layer switches (star), have uplinks terminating on a pair of redundant Layer 3 switches, which act as an aggregation point. This aggregation layer is referred to as the distribution layer. The distribution layer is the first place where routing and Layer 3 switching occur in the multilayer model. Important features provided by this layer include the following:

- Default gateway redundancy
- Intelligent best-path selection towards other modules of the network
- Wire-speed Layer 3 switching

For EttF 1.1, two stacked Cisco Catalyst 3750s are chosen to fulfil this role (see [Figure D-1](#)). With StackWise technology, switch redundancy is achieved, but logically only one switch is being managed and configured.

The uplinks on the Catalyst 3750s from the distribution layer to the core layer are cross-stack Layer 3 EtherChannels, which provide yet another level of redundancy and resiliency in the event of a link failure (see [Figure D-1](#)). EtherChannels also provide additional bandwidth by aggregating up to eight interfaces into one logical interface. The core layer, with dual Cisco Catalyst 4507R switches, achieve the maximum device redundancy with redundant supervisors and power supplies. From a link and alternate path perspective, meshed connections down to the Cisco Catalyst 3750 are recommended.

Following is a sample Catalyst 3750 cross-stack EtherChannel configuration:

```
!
interface GigabitEthernet1/0/25
 no switchport
 no ip address
 channel-group 1 mode active
end

CZ-C3750-1#show run int g2/0/25
Building configuration...

Current configuration : 98 bytes
!
interface GigabitEthernet2/0/25
 no switchport
 no ip address
 channel-group 1 mode active
end

CZ-C3750-1#show run int p1
Building configuration...

Current configuration : 108 bytes
!
interface Port-channel1
 description CZ-C4500-1
 no switchport
 ip address 10.18.3.100 255.255.255.0
end

CZ-C3750-1#show ether summ
Flags: D - down P - in port-channel
 I - stand-alone s - suspended
 H - Hot-standby (LACP only)
 R - Layer3 S - Layer2
 U - in use f - failed to allocate aggregator
```

u - unsuitable for bundling  
w - waiting to be aggregated  
d - default port

Number of channel-groups in use: 4  
Number of aggregators: 4

| Group | Port-channel | Protocol | Ports                   |
|-------|--------------|----------|-------------------------|
| 1     | Po1(RU)      | LACP     | Gi1/0/25(P) Gi2/0/25(P) |
| 2     | Po2(SD)      | -        |                         |
| 3     | Po3(RU)      | LACP     | Gi1/0/27(P) Gi2/0/27(P) |
| 10    | Po10(SD)     | LACP     | Gi1/0/1(D) Gi2/0/1(D)   |

Sample 4500 EtherChannel Configuration

```
interface GigabitEthernet4/9
no switchport
no ip address
logging event link-status
channel-group 1 mode active
end
```

CZ-C4500-1#show run int g4/10  
Building configuration...

```
Current configuration : 123 bytes
!
interface GigabitEthernet4/10
no switchport
no ip address
logging event link-status
channel-group 1 mode active
end
```

CZ-C4500-1#show run int p1  
Building configuration...

```
Current configuration : 96 bytes
!
interface Port-channel1
ip address 10.18.3.101 255.255.255.0
logging event link-status
end
```

CZ-C4500-1#show ether summ  
Flags: D - down P - in port-channel  
I - stand-alone s - suspended  
R - Layer3 S - Layer2  
U - in use f - failed to allocate aggregator  
u - unsuitable for bundling  
w - waiting to be aggregated  
d - default port

Number of channel-groups in use: 4  
Number of aggregators: 4

| Group | Port-channel | Protocol | Ports              |
|-------|--------------|----------|--------------------|
| 1     | Po1(RU)      | LACP     | Gi4/9(P) Gi4/10(P) |

## First Hop Redundancy

For EttF 1.1, first hop redundancy is the default gateway that end nodes have configured in case they need to communicate across subnets. For cell/area-level devices (PACs, VFDs, I/O), the default gateway is configured on the SVI terminating the VLAN(s) carried in the cell/area zone. Two Catalyst 3750s are stacked together to form one logical router, so in the case of a single router failure, the second stacked Catalyst 3750 takes over as the default gateway. Note, however, that the default behavior is for the newly active stack master to assign a new stack MAC address. This can be a problem for end nodes that do not support Gratuitous Address Resolution Protocol (GARP), which sends a message to hosts to clear their ARP cache and to assign a new IP/MAC binding. As such, the **stack-mac persistent timer 0** command should be used to ensure that the original master MAC address remains the stack MAC address after a failure. This makes it transparent to endpoints, so that they do not have to learn a new IP/MAC pair. The following is a sample session:

```
CZ-C3750-1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
CZ-C3750-1(config)#
CZ-C3750-1(config)#stack-mac persistent timer ?
 <0-0> Enter 0 to continue using current stack-mac after master switchover
 <1-60> Interval in minutes before using the new master's mac-address
 <cr>

CZ-C3750-1(config)#stack-mac persistent timer 0
```

For application servers, which are connected to a separate Layer 2 switch that is connected to the dual Catalyst 4507R switches in the manufacturing zone (see [Figure D-1](#)), Cisco recommends using Hot Standby Routing Protocol (HSRP) as the method of providing first hop redundancy. A router is elected as the active router and is responsible for answering ARP requests for the virtual IP address. HSRP routers discover each other via hello packets, which are multicast packets sent to the 224.0.0.2 “all-routers” address (knowledge of this address may be important for troubleshooting purposes). Remember that HSRP is completely independent of the routing protocol in use on the router. For more information about HSRP, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs009.htm>

## NSF/SSO

From a Layer 3 perspective, the default behavior of a router after a failover event is to tear down routing protocol neighbor relationships, re-establish the neighbor, relearn routing information, and re-populate the Cisco Express Forward table. This can take minutes, depending on the size of the routing table. In a manufacturing environment, this is unacceptable, so Nonstop Forwarding with Stateful Switchover (NSF/SSO) is recommended to mitigate that risk. For EttF 1.1, this applies only to the Catalyst 4507s and the Catalyst 3750s.

NSF works in conjunction with SSO to ensure Layer 3 integrity following a switchover. This allows a router experiencing the failure of an active supervisor to continue forwarding data packets along known routes while the routing protocol information is recovered and validated. This forwarding can continue to occur even though peering arrangements with neighbor routers have been lost on the restarting router. NSF relies on the separation of the control plane and the data plane during supervisor switchover. The data plane continues to forward packets based on pre-switchover Cisco Express Forwarding information. The control plane implements graceful restart routing protocol extensions to signal a supervisor restart to NSF-aware neighbor routers, reform its neighbor adjacencies, and rebuild its routing protocol database following a switchover. An NSF-capable router implements the NSF functionality and continues to forward data packets after a supervisor failure. An NSF-aware router understands the NSF

graceful restart mechanisms: it does not tear down its neighbor relationships with the NSF-capable restarting router, and can help a neighboring NSF-capable router restart, thus avoiding unnecessary route flaps and network instability. An NSF-capable router is also NSF-aware.

To configure SSO on the Catalyst 4507s, perform the following configuration:

```
CZ-C4500-1#conf t
CZ-C4500-1(config)#redund
CZ-C4500-1(config-red)#mode ?
 rpr Route Processor Redundancy
 sso Stateful Switchover

CZ-C4500-1(config-red)#mode sso
CZ-C4500-1(config-red)#
```

To configure NSF on the Catalyst 4507s, add the **nsf** keyword when configuring a dynamic routing protocol (note that OSPF is the selected IGP):

```
router ospf 100
router-id 1.1.1.1
log-adjacency-changes
nsf
```

## Summary of Features in the Manufacturing Zone

Table 6-3 provides a summary of features in the manufacturing zone.

**Table 6-3** Summary of Features in the Manufacturing Zone

| Feature                                                    | Description                                                                                                                                     | Where To Apply in I.E Network                                                                                                             |
|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Redundant paths                                            | Meshed connections for multiple paths.                                                                                                          | <ul style="list-style-type: none"> <li>Catalyst 4500</li> <li>Catalyst 3750</li> </ul>                                                    |
| 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. | <ul style="list-style-type: none"> <li>Catalyst 4500</li> <li>Catalyst 3750—Virtual with StackWise</li> </ul>                             |
| StackWise                                                  | Uses stack interconnect cables to create a virtual switch fabric for stacks of the Catalyst 3750.                                               | <ul style="list-style-type: none"> <li>Catalyst 3750</li> <li>N/A to other platforms</li> </ul>                                           |
| Redundant power supplies                                   | Each system has dual power supplies so that the system operates normally upon failure of a power supply.                                        | <ul style="list-style-type: none"> <li>Catalyst 4507R—Internal</li> <li>Catalyst 3750—External</li> <li>Catalyst 2955—External</li> </ul> |
| Redundant fans                                             | Each fan tray has multiple fans.                                                                                                                | Catalyst 4507R                                                                                                                            |
| Line card online insert and removal (OIR)                  | New line cards can be added or removed without affecting the system or losing the configuration.                                                | Catalyst 4507R                                                                                                                            |
| 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 4507R                                                                                                                            |

**Table 6-3 Summary of Features in the Manufacturing Zone (continued)**

|                                                        |                                                                                                                                                                                                                                                                                                      |                                                                                            |
|--------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Nonstop Forwarding Awareness                           | Processes NSF messages from restarting neighbor and does not tear down neighbor relationship.                                                                                                                                                                                                        | Catalyst 3750                                                                              |
| In-Service Software Upgrade (ISSU)                     | Ranges from full image upgrades to granular, selective software maintenance are able to be performed without service impact across all Cisco IOS-based products.                                                                                                                                     | Catalyst 4507R                                                                             |
| 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: <ul style="list-style-type: none"> <li>• Transferring the global configuration</li> <li>• Applying default configuration</li> <li>• Applying preconfigured configuration</li> </ul> | Catalyst 3750                                                                              |
| Generic Online Diagnostics (GOLD)                      | Online diagnostics to help ensure that a system booting up and a live system are healthy.                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li>• Catalyst 4507R and 3750—Subset of GOLD</li> </ul> |
| EtherChannel                                           | Link aggregation for bandwidth and redundancy; both PAgP and LACP are supported.                                                                                                                                                                                                                     | Catalyst 4507R and 3750                                                                    |

## HA Design and Implementation in the DMZ

### Cisco ASA Redundancy Design

Cisco ASA supports the following two types of failover:

- Active/standby
- Active/active

Active/standby is the supported design for EtrF 1.1, in which the active ASA is responsible for passing traffic. Active/standby failover allows a standby security appliance to take over the functionality of a failed unit. When the active unit fails, it changes to the standby state and the standby unit changes to the active state. The unit that becomes active assumes the IP addresses (or, for transparent firewall, the management IP address) and MAC addresses of the failed unit and begins passing traffic. The unit that is now in standby state takes over the standby IP addresses and MAC addresses. Because network devices see no change in the MAC-to-IP address pairing, no ARP entries change or time out anywhere on the network.

## Primary/Secondary Status and Active/Standby Status

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby; that is, which IP addresses to use and which unit actively passes traffic. However, the following differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units startup at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit is active, and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

## Determination of the Active Unit

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, the primary unit becomes the active unit and the secondary unit becomes the standby unit.

## Failover Triggers

The unit can failover if one of the following events occurs:

- An administrator manually switches over from active to standby.
- The standby Cisco ASA stops receiving keepalive packets on the failover command interface.
- The command interface link goes down.
- The link state of an interface goes down.
- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- The **no failover active** command is entered on the active unit, or the **failover active** command is entered on the standby unit.

## Configuration Synchronization

Configuration synchronization occurs when one or both devices in the failover pair boot. Configurations are always synchronized from the active unit to the standby unit. When the standby unit completes its initial startup, it clears its running configuration (except for the failover commands needed to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

The following commands are replicated to the standby unit:

- all configuration commands except for the **mode**, **firewall**, and **failover lan unit** commands
- **copy running-config startup-config**

- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands are not replicated to the standby unit:

- all forms of the **copy** command except for **copy running-config startup-config**
- all forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **mode**
- **show**

## Passage of State Information to the Standby Unit

The standby ASA monitors the status of the active ASA by sending keepalive messages over a dedicated Gigabit Ethernet failover connection between the two (also known as LAN-based failover). The active also does the reverse to the standby ASA. This failover design is stateful, which means that the active ASA maintains the connection table and replicates it to the standby ASA. Whenever there is a change in the table, the active ASA sends a stateful update to the standby.

The state information passed to the standby unit includes the following:

- NAT translation table
- TCP connection states
- UDP connection states
- ARP table
- Layer 2 bridge table (when running in transparent firewall mode)
- HTTP connection states (if HTTP replication is enabled)
- ISAKMP and IPsec SA table
- GTP PDP connection database

The information that is not passed to the standby unit when stateful failover is enabled includes the following:

- HTTP connection table (unless HTTP replication is enabled)
- User authentication (uauth) table
- Routing tables
- State information for security service modules
- DHCP server address leases
- L2TP over IPsec state information



## Active/Standby Failover Configuration



### Note

Before configuring the failover, verify that the second Cisco ASA is turned off. Also verify that the activation key on the Cisco ASA supports failover and both Cisco ASAs are using the same mode (single or multiple).

### Selecting the Failover Link

The first step is to decide which interface will be used to send failover control messages. The failover control link interface is defined by using the **failover lan interface** command followed by the interface name. The following example shows that the Cisco ASA is using GigabitEthernet0/2 as the failover control interface. In this example, the LAN interface is given a name of *FOCtrlIntf*. However, you can specify any name for this interface.

```
Chicago(config)# failover lan interface FOCtrlIntf GigabitEthernet0/2
```



### Note

If an interface already has the **nameif** statement configured, the security Cisco ASA displays an error stating that the interface is already in use. For example:

```
Chicago(config)# failover lan interface FOCtrlIntf GigabitEthernet0/2
interface already in use
```

To fix this issue, issue the **no nameif** command under that interface.

After the **failover lan interface** command is configured, the Cisco ASA adds a description under the failover interface configuration. It also clears any configuration parameters on that interface, as follows:

```
Chicago# show running | begin interface GigabitEthernet0/2
interface GigabitEthernet0/2
 description LAN failover Interface
```

### Assigning Failover IP Addresses

After selecting which failover control interface the Cisco ASA is going to use, the next step is to configure the physical interfaces for the system and the standby IP addresses. The active Cisco ASA uses the system IP addresses, while the standby Cisco ASA uses the standby IP addresses. The following example shows that the Chicago Cisco ASA is using 209.165.200.225 and 192.168.10.1 as the system IP addresses, and 209.165.200.226 and 192.168.10.2 as the failover IP addresses on the outside and inside interfaces, respectively.

```
Chicago(config)# interface GigabitEthernet0/0
Chicago(config-if)# nameif outside
Chicago(config-if)# security-level 0
Chicago(config-if)# ip address 209.165.200.225 255.255.255.224 standby 209.165.200.226
Chicago(config-if)# exit
Chicago(config)# interface GigabitEthernet0/1
Chicago(config-if)# nameif inside
Chicago(config-if)# security-level 100
Chicago(config-if)# ip address 192.168.10.1 255.255.255.0 standby 192.168.10.2
```

For two security Cisco ASAs to communicate, the designated failover control interface should be configured with an IP address as well. Following is the complete command syntax to configure an IP address on the failover control interface:

```
failover interface ip interface_name ip_address mask standby ip_address
```

*interface\_name* is the designated interface used for failover. The first IP address is the interface IP address used by the active Cisco ASA, and the second IP address is the IP address used by the standby Cisco ASA. The active unit uses its address to synchronize the running configuration with the standby. In the following example, the active Cisco ASA is assigned a 10.10.10.1 IP address along with a standby IP address of 10.10.10.2.

```
Chicago# configure terminal
Chicago(config)# failover interface ip FOctrlIntf 10.10.10.1 255.255.255.252 standby
10.10.10.2
```

## Setting Failover Key (Optional)

To secure the failover control messages sent between the Cisco ASAs, an administrator can optionally specify a shared secret key. The shared secret key encrypts and authenticates the failover messages if they are susceptible to unauthorized users. The following example shows how to configure a failover shared secret key of *cisco123*.

```
Chicago# configure terminal
Chicago(config)# failover key cisco123
```

The failover key uses DES or AES, depending on the installed license. It also uses MD5 as the hash to authenticate the message. Therefore, it is important that both Cisco ASAs use the same cipher license key.



### Note

---

If a failover key is not used, the active Cisco ASA sends all information in clear text, including the UDP/TCP states, the user credentials, and the VPN-related information.

---

## Designating the Primary Cisco ASA

The two security Cisco ASAs send failover control messages through a network cable that has identical ends. Unlike a Cisco PIX firewall, in which the failover cable decides which firewall becomes primary, it is impossible to designate a Cisco ASA as primary based on the cable. To resolve the problem of which device should act as primary or secondary, you must designate the primary and secondary status through software configuration by using the **failover lan unit** command. In the following example, FO1 is designated as the primary failover Cisco ASA.

```
Chicago# configure terminal
Chicago(config)# failover lan unit primary
```

## Enabling Stateful Failover (Optional)

As discussed earlier, the stateful failover feature in the Cisco ASA replicates the state and translation tables from the active unit to the standby unit. In the event of a failure, the standby unit takes over the connections, and data flows are not interrupted. The stateful failover requires a network interface to replicate the states. Cisco ASA can use either a dedicated or the failover control interface to replicate the updates. A stateful link interface is defined by using the **failover link** command followed by the

name of the interface. In the following example, the primary Cisco ASA is using GigabitEthernet0/3 as the stateful interface. The interface IP is 10.10.10.5 and the standby IP address is 10.10.10.6. The administrator uses *statefullink* as the interface name.

```
Chicago(config)# failover link statefullink GigabitEthernet0/3
Chicago(config)# failover interface ip statefullink 10.10.10.5 255.255.255.252 standby
10.10.10.6
```

**Note**

Like the **failover lan interface** command, the Cisco ASA adds a description under the stateful link interface and clears any configuration on that interface.

For stateful failover, you can use the failover LAN interface if the stateful updates do not oversubscribe the interface bandwidth. Set up a different interface for stateful failover if you are concerned about possibly oversubscribing the failover control interface. If the security Cisco ASA uses the same interface for both control and stateful messages, you have to connect the security Cisco ASA through a switch. Crossover cable is not supported.

The stateful failover does not replicate HTTP-based connections. HTTP connections usually have a short lifetime and therefore are not replicated by default. Additionally, they add considerable load on the security Cisco ASA if the amount of HTTP traffic is large in comparison to other traffic.

If the HTTP connections need to be replicated to the standby Cisco ASA, use the **failover replication http** command, as shown by the following example:

```
Chicago(config)# failover replication http
```

## Enabling Failover Globally

The last step in configuring failover on the primary Cisco ASA is to enable failover globally. The following example shows how to enable failover in the Chicago FO1 Cisco ASA:

```
Chicago(config)# failover
```

## Configuring Failover on the Secondary Cisco ASA

In the Cisco failover feature, there is no need to manually configure the secondary Cisco ASA. Instead, you just need to configure some basic information about failover. After that, the primary/active Cisco ASA starts synchronizing its configuration. The bootstrap configuration includes the following five configuration parameters:

- Failover designation
- Failover link interface
- Failover interface IP address
- Failover shared key
- Failover enable

The following example shows the bootstrap configuration of the secondary Cisco ASA needed in LAN-based failover:

```
failover lan unit secondary
failover lan interface FOctrlIntf GigabitEthernet0/2
failover key cisco123
failover interface ip FOctrlIntf 10.10.10.1 255.255.255.252 standby 10.10.10.2 failover
```

**Note**

---

After failover is enabled on both Cisco ASAs, their running configuration is identical except for the **failover lan unit** command.

---

For a detailed explanation of configuring the ASA redundancy, see the following URL:  
[http://www.cisco.com/en/US/products/ps6120/products\\_configuration\\_guide\\_chapter09186a008063b31a.html#wp1058096](http://www.cisco.com/en/US/products/ps6120/products_configuration_guide_chapter09186a008063b31a.html#wp1058096)



## CHAPTER 7

# Implementation of Network Management

---

## Cisco Network Assistant

Cisco Network Assistant (CNA) is a PC-based network management application optimized for wired and wireless LANs. It simplifies configuration, deployment, and ongoing management of networks. CNA is available at no cost and can be downloaded from <http://www.cisco.com/go/cna>. CNA gives a centralized network view through a user-friendly GUI for configuration, monitoring, troubleshooting, and maintenance. It allows network administrators to easily apply common services, generate inventory reports, synchronize passwords, upgrade software, and employ features across Cisco switches, routers, and access points.

CNA uses the concept of the community to group the switches and routers in the network. In each community, candidates are network devices that have IP addresses but are not part of a community. Members are network devices that are currently part of a community. To join a community, a candidate must meet the following requirements:

- The candidate has an IP address.
- Cisco Discovery Protocol (CDP) version 2 is enabled (the default) if you want the device to be autodiscovered.
- The candidate has HTTP (or HTTPS) enabled.

CNA uses CDP to automatically discover all the available devices in the network. Beginning with the IP address for a starting device and the port numbers for HTTP (or HTTPS) protocols, CNA uses CDP to compile a list of community candidates that neighbor the starting device. CNA can discover candidate and member devices across multiple networks and VLANs as long as they have valid IP addresses. By default, CNA in community mode discovers up to four hops away. If CNA fails to discover a device, you can add it manually through the IP management IP address.

The latest CNA release, version 5.1, can be installed on the following:

- Windows XP with Service Pack 1 or later
- Windows 2003 with Service Pack 1 or later
- Windows 2000 with Service Pack 3 or later

CNA supports networks with 40 or fewer switches and routers. Within each community, CNA supports up to 20 devices, including up to 4 Catalyst 4500 series switches (modular), 16 Catalyst 2900/3500 switches, 2 routers, and 2 PIX firewalls. The network components in EtoF 1.1 that CNA 5.1 supports include those in the cell/area and manufacturing zones: Catalyst 2955, Catalyst 3750, and Catalyst 4500.

The server running CNA is recommended to be placed in the manufacturing zone (level 3).

The following three steps are needed to start using CNA:

1. Install the CNA application on the PC.
2. Configure HTTP or HTTPS on all the switches and routers that CNA will contact.
3. Launch CNA to connect to the switches and routers.

The detailed instructions of these procedures are available at the following URL:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/configuration/guide/netasist.html#wp1080326>

## CNA Security Considerations

Note the following security considerations:

- Use HTTPS instead of HTTP for CNA access to routers and switches to ensure that the information exchanged between the web browser and CNA is encrypted.
- Because CDP is recommended to be disabled for security reasons, the operator can add routers and switches to the CNA device list manually.

## Cisco Adaptive Security Device Manager

Cisco Adaptive Security Device Manager (ASDM) delivers security management and monitoring through an intuitive, easy-to-use web-based management interface. The ASDM accelerates security appliance deployment with intelligent wizards, robust administration tools, and versatile monitoring services that complement the Cisco ASA 5500 Series Adaptive Security Devices. The ASDM provides the following features:

- Packet tracer utility—Verifies the impact of real traffic flows on the entire system configuration. This utility sketches animated results as each policy is rigorously tested and provides direct links to correct failed tests for exploration-free policy tuning.
- Profile-based management for all application inspection and control capabilities—Uses preconfigured low, medium, and high security profiles for each of the application inspection engines for rapid deployment in any security environment. The ASDM enables the granular customization of any of the security profiles to cater to the needs of advanced applications. It provides easy integration of user-defined regular expressions into existing security policies to allow rapid threat mitigation against new and upcoming application attacks.
- High-availability and scalability wizard—Simplifies the deployment of active/active and active/standby high availability. It helps ensure comprehensive connectivity testing and error verification for smooth and accurate deployment.
- Integrated security policy and access control table—Enhances the policy configuration and management experience by providing a stream-lined, in-depth perspective into all the access rules, AAA, and security policies of the system. The ASDM facilitates rapid troubleshooting through a new rule query option that enables administrators to quickly search for network elements and the policies employing them. It enables the rapid editing of all network and service object groups via a new object group selector panel.
- Easy troubleshooting—Integrates syslog references to provide brief explanations and recommended actions for each message for isolating and resolving security issues quickly. The ASDM enables the parsing of syslog messages for customizable views based on time, date, syslog IDs, and IP

addresses. It also provides traceroute support for network connectivity testing and verification. An ASDM Assistance Guide provides task-oriented methods to configure features such as AAA, logging filters, and SSL VPN clients.







## APPENDIX **A**

# Characterization of the EttF Cell/Area Zone Design

---

All factory floor devices are connected at the cell/area zone layer of EttF 1.1. From a network design perspective, a solid spanning tree and multicast design are critical for reliability and to meet predefined service-level agreements (SLAs). As mentioned in earlier sections, Cisco recommends that Rapid Spanning Tree Protocol (RSTP 802.1w) and IGMP snooping with querier be deployed at this layer. This appendix outlines the validation methodology and the corresponding results of the testing.

## STP Testing

### STP Test Methodology

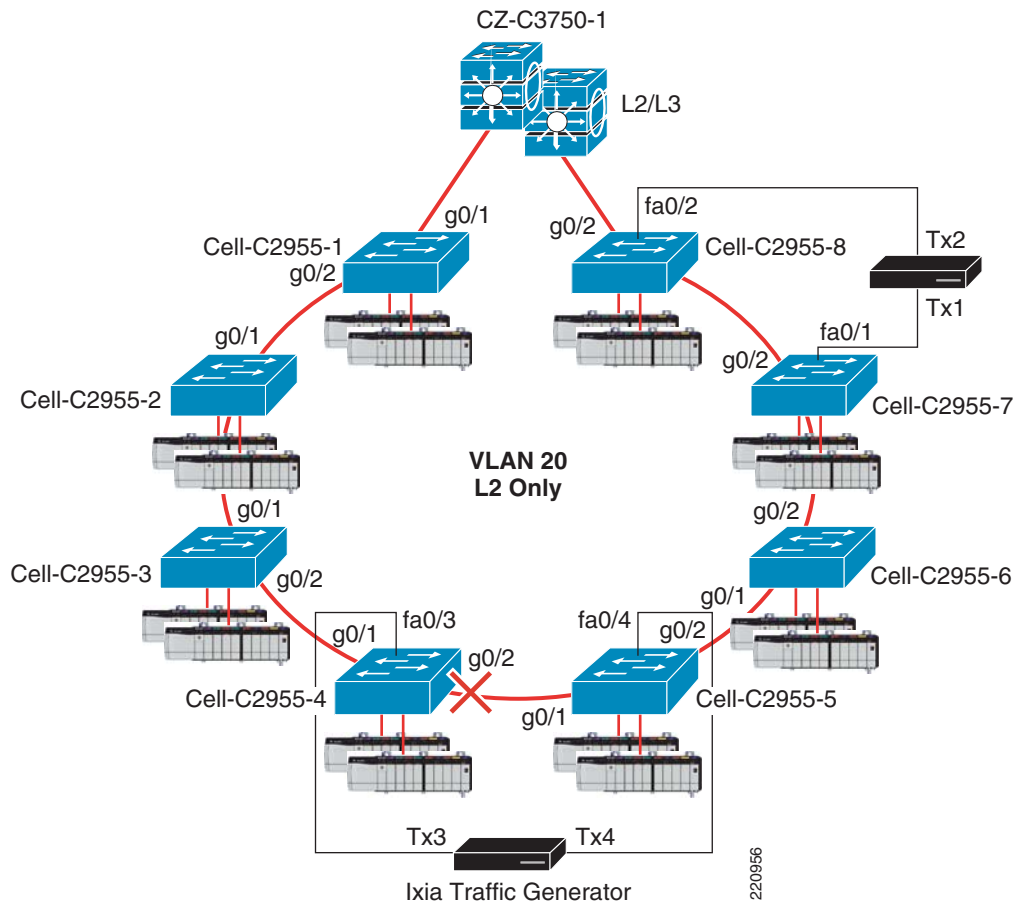
Eight Cisco Catalyst 2955 switches, each running Native IOS 12.1(22)EA9, plus one Catalyst 3750 running Native IOS 12.2(25)SEB4, are linked together in back-to-back fashion via 802.1q Gigabit Ethernet trunks to form a ring topology (See [Figure A-1](#)).

All nine devices are running RSTP 802.1w carrying only one VLAN (vlan 20). This topology creates a Layer 2 loop, and therefore one port must be blocked by STP. By configuring switch CZ-3750-1 with the lowest spanning tree priority, it is elected as the root bridge for this VLAN. The STP parameters on all the other switches are left at their default settings. Accordingly, switch Cell-2955-4 is now the furthest away from the root bridge in terms of path cost and must select a port to block (GigabitEthernet 0/2). The root port on devices Cell-2955-1 through Cell-2955-4 is GigabitEthernet 0/1, and the root port on devices Cell-2955-5 through Cell-2955-8 is GigabitEthernet 0/2. A traffic generator (Ixia 400 Tf) is attached to switches Cell-2955-7, Cell-2955-8, Cell-2955-4, and Cell-2955-5 via ports Tx1, Tx2, Tx3, Tx4 respectively. The traffic generator is used to measure the convergence time in various failure scenarios.

## STP Test Topology

Figure A-1 shows the test topology.

**Figure A-1 Test Topology**



Multiple bidirectional traffic streams are configured on Ixia between Tx1 <-> Tx2 and Tx3 <-> Tx4. Each of these pairs is referred to as a traffic suite with their own test cases. Each stream for each test suite is designed to source 1, 50, 100, and 200 MAC addresses destined to 1, 50, 100, and 200 MAC addresses. Thus, twice the number of configured source MAC addresses are traversing the ring for each test case. Each packet is a mix of 500 and 64 bytes in length sent in a continuous fashion until the STP has re-converged and the network has stabilized. At steady state for Suite 1 (Tx1 <-> Tx2), the traffic flow is as follows:

1. Packet egresses Tx1 of Ixia and ingresses port fa0/1 of Cell-2955-7
2. Packet then egresses gi0/2 on Cell-2955-7 and ingresses port g0/1 on Cell-C2955-8
3. Packet then egresses fa0/2 on Cell-C2955-8 and finally ingresses Tx2 of Ixia

This flow is reversed for streams going in the opposite direction (Tx2 to Tx1).

At steady state for Suite 2 (Tx3 <-> Tx4), the traffic flow is as follows:

1. Packet egresses Tx3 of Ixia and ingresses port fa0/3 of Cell-C2955-4

2. Because STP is blocking port gi0/2, the packet egresses gi0/1 of Cell-C2955-4 and traverses the entire ring in the clockwise direction until it reaches Cell-C2955-5
3. Packet ingresses gi0/2 on Cell-C2955-5 and egresses fa0/4
4. Packet finally ingresses Tx4 of Ixia

This flow is reversed for streams going in the opposite direction (Tx4 to Tx3).

## STP Test Scenarios

Two test suites are explored, each with a different traffic flow. Within each test suite, multiple failure scenarios are introduced to simulate various STP changes in the ring topology. With each failure, convergence time is measured using the following formula:

$$[(Tx - Rx) / \text{packet rate}] * 1000$$

Where:

Tx = Packets transmitted


Rx = Packets received

PPS = 10,000 pps

### Test Suite 1—Bidirectional Traffic (Tx1 <-> Tx2)


At steady state, traffic is flowing bidirectionally point-to-point between Cell-C2955-7 and Cell-C2955-8 (essentially a best-case scenario). (See [Figure A-2](#).)

Figure A-2 Test Suite 1—Bidirectional Traffic Flow



However, after simulating a failure between these two switches, the traffic must then traverse the entire ring (becoming the worst-case scenario) to reach its destination. (See [Figure A-3](#).)

Figure A-3 Test Suite 1—Worse-Case Scenario



The eight failure scenarios in Suite 1 are as follows:

- Failure 1—Software shut link between Cell-C2955-7 and Cell-C2955-8
- Failure 2—Software unshut link between Cell-C2955-7 and Cell-C2955-8
- Failure 3—Physically remove link between Cell-C2955-7 and Cell-C2955-8
- Failure 4—Physically re-insert link between Cell-C2955-7 and Cell-C2955-8
- Failure 5—Root bridge down
- Failure 6—Root bridge up
- Failure 7—Stack master down on CZ-C3750
- Failure 8—Stack master re-established on CZ-C3750


For each failure scenario, the following is measured and verified:

- Convergence time
- Verify Rockwell Automation (RA) equipment functioning properly after disruption
- Measure CPU and memory on cell devices that are sending and receiving Ixia traffic

## Test Suite 2—Bidirectional Traffic (Tx3 <-> Tx4)


At steady state, traffic is flowing bidirectionally between Cell-C2955-4 and Cell-C2955-5 across the entire ring. (See [Figure A-4](#).)

**Figure A-4** Test Suite 2—Bidirectional Traffic Flow



Because STP is blocking g0/2 on Cell-C2955-4, this constitutes the worse-case scenario before any failures are introduced. However, after a network failure and a subsequent STP topology change, traffic flows to its directly-connected neighbor and becomes the best-case scenario. (See [Figure A-5](#).)

Figure A-5 Test Suite 2—Best-Case Scenario



The eight failure scenarios in Suite 2 are as follows:

- Failure 1—Software shut link between Cell-C2955-1 and CZ-C3750
- Failure 2—Software unshut link between Cell-C2955-1 and CZ-C3750
- Failure 3—Physically remove link between Cell-C2955-1 and CZ-C3750
- Failure 4—Physically re-insert link between Cell-C2955-1 and CZ-C3750
- Failure 5—Root bridge down
- Failure 6—Root bridge up
- Failure 7—Stack master down on CZ-C3750
- Failure 8—Stack master re-established on CZ-C3750

For each failure scenario, the following is measured and verified:

- Convergence time
- Verify RA equipment functioning properly after disruption
- Measure CPU and memory on cell devices that are sending and receiving Ixia traffic

## Test Tools

The following equipment is needed for performing these tests.

- 16 Cisco Catalyst C2955T-12 industrial switches
- 2 Cisco Catalyst WS-C3750G-24PS (stacked)
- 1 Ixia traffic generator
- 1 set of RA equipment

## STP Test Results

### Suite 1 Test Cases

**Table A-1** Test Case 1—Software Shut Link between Cell-C2955-7 and Cell-C2955-8

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 444.92          | 446.60          | 1105.3        | 1107.4        | 1436.7         | 1488           | 2015           | 2017.1         |
| 2     | 408.744         | 410.49          | 1362.5        | 1364.5        | 1264           | 1266.1         | 1750.5         | 1747.5         |
| 3     | 319.11          | 320.81          | 1080.6        | 1082.6        | 1064           | 1066           | 1027.4         | 1029.4         |
| Avg   | 390.9252242     | 392.6380775     | 1182.8        | 1184.833333   | 1254.9         | 1273.366667    | 1597.633333    | 1598           |

**Table A-2** Test Case 2—Software Unshut Link between Cell-C2955-7 and Cell-C2955-8

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 229.29          | 232.18          | 790.9         | 802.7         | 1990.6         | 2005.4         | 2005.4         | 1486.9         |
| 2     | 171.02          | 174.62          | 992.3         | 1002.9        | 1993.2         | 2004.8         | 2004.8         | 1584.6         |
| 3     | 314.75          | 317.79          | 993.6         | 1003.5        | 1987.9         | 2005.9         | 2005.9         | 1981.7         |
| Avg   | 238.35          | 241.53          | 925.60        | 936.37        | 1990.57        | 2005.37        | 2005.37        | 1817.73        |

**Table A-3** Test Case 3—Physically Remove Link between Cell-C2955-7 and Cell-C2955-8

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 511.44          | 510.21          | 2071.7        | 2071.7        | 1738.3         | 1738.3         | 2140.6         | 2140.6         |
| 2     | 465.48          | 462.35          | 2197.5        | 2197.5        | 1970.6         | 1970.6         | 2505.1         | 2505.1         |
| 3     | 441.89          | 439.76          | 1952.6        | 1952.6        | 2040.6         | 2110.2         | 1803           | 1803           |
| Avg   | 472.94          | 470.77          | 2073.93       | 2073.93       | 1916.5         | 1939.7         | 2149.57        | 2149.57        |



**Table A-4** Test Case 4—Physically Re-insert Link between Cell-C2955-7 and Cell-C2955-8

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 174.017         | 175.616         | 1810.6        | 1822.6        | 1989.9         | 2001           | 1979.3         | 1989.8         |
| 2     | 474.074         | 473.718         | 1989.7        | 2002.3        | 1863           | 1879.1         | 1986.3         | 2003.1         |
| 3     | 473.340         | 473.079         | 1992.1        | 2004.1        | 1968.3         | 2001.1         | 1891.7         | 1916.8         |
| Avg   | 407.143         | 407.471         | 1930.8        | 1943          | 1940.4         | 1960.40        | 1952.43        | 1969.9         |

**Table A-5** Test Case 5—Root Bridge Down

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 0               | 0               | 0             | 0             | 0              | 0              | 0              | 0              |
| 2     | 0               | 0               | 0             | 0             | 0              | 0              | 0              | 0              |
| 3     | 0               | 0               | 0             | 0             | 0              | 0              | 0              | 0              |
| Avg   | 0               | 0               | 0             | 0             | 0              | 0              | 0              | 0              |

**Table A-6** Test Case 6—Root Bridge Up

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 173.03          | 174.24          | 809.5         | 795.7         | 898.2          | 879.5          | 923.6          | 899.9          |
| 2     | 57.83           | 55.93           | 863.3         | 848.7         | 733.1          | 715.2          | 912.5          | 882.7          |
| 3     | 82.74           | 80.09           | 743           | 726.9         | 857.5          | 836.8          | 853.9          | 838.3          |
| Avg   | 104.53          | 103.42          | 805.27        | 790.43        | 829.6          | 810.50         | 896.67         | 873.63         |

**Table A-7** Test Case 7—Stack Master Down on CZ-C3750

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 8.3             | 13.40           | 16.1          | 25.6          | 14.1           | 28.6           | 28.7           | 31.2           |
| 2     | 9.8             | 15.49           | 16.3          | 24.3          | 14.8           | 26.6           | 22.7           | 28.2           |
| 3     | 8.6             | 15.40           | 16.7          | 23.3          | 11.1           | 25.6           | 25.7           | 30.2           |
| Avg   | 8.90            | 14.76           | 16.37         | 24.40         | 13.33          | 26.93          | 25.70          | 29.87          |

**Table A-8 Test Case 8—Stack Master Re-Established on CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 153.10          | 145.19          | 982.5         | 968.6         | 919.2          | 898            | 883.2          | 865.3          |
| 2     | 155.10          | 148.19          | 891.8         | 859.6         | 949.2          | 924.5          | 984.2          | 978.2          |
| 3     | 161.39          | 161.39          | 851.8         | 895.6         | 986.9          | 974.7          | 927.3          | 915.4          |
| Avg   | 156.53          | 151.59          | 908.70        | 907.93        | 951.77         | 932.40         | 931.57         | 919.63         |

## Suite 2 Test Cases

**Table A-9 Test Case 1—Software Shut Link between Cell-C2955-1 and CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 29.60           | 48.00           | 43            | 50.1          | 51.5           | 56.5           | 73.3           | 95.3           |
| 2     | 32.60           | 43.49           | 39.1          | 46.1          | 47.2           | 59.8           | 68.3           | 90.7           |
| 3     | 39.70           | 54.79           | 45.1          | 47.2          | 45             | 57.4           | 66.2           | 88.3           |
| Avg   | 38.20           | 48.06           | 42.4          | 47.8          | 47.9           | 57.9           | 69.26          | 91.43          |

**Table A-10 Test Case 2—Software Unshut Link between Cell-C2955-1 and CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 1275.3          | 1289.054069     | 1944.6        | 1956.2        | 1923.7         | 1934.2         | 1916.8         | 1942.6         |
| 2     | 1361.1          | 1377.855617     | 1954          | 1965.7        | 1981.6         | 1987.5         | 1999.9         | 1964.8         |
| 3     | 1469.7          | 1475.2          | 1940.1        | 1951.3        | 1927.7         | 1933.8         | 1909.1         | 1930.7         |
| Avg   | 1368.7          | 1380.703229     | 1946.233333   | 1957.733333   | 1944.333333    | 1951.833333    | 1941.933333    | 1946.033333    |

**Table A-11 Test Case 3—Physically Remove Link between Cell-C2955-1 and CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 537.4           | 537.5           | 474.8         | 480           | 570.4          | 575            | 492.3          | 512.3          |
| 2     | 511.6           | 511.7           | 526.3         | 526.1         | 519.6          | 529.7          | 567.2          | 587.4          |
| 3     | 423.6           | 423.7           | 497.6         | 497.2         | 492.1          | 510.2          | 571.2          | 576.6          |
| Avg   | 490.8666667     | 490.9666667     | 499.5666667   | 501.1         | 527.3666667    | 538.3          | 543.5666667    | 558.7666667    |

**Table A-12 Test Case 4—Physically Re-insert Link between Cell-C2955-1 and CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 1964.9          | 1970.8          | 1983.2        | 1991.3        | 1924.5         | 1938.2         | 1928.5         | 1946.4         |
| 2     | 1971.1          | 1976.9          | 1965.2        | 1974.5        | 1925.3         | 1939.3         | 1883.1         | 1894.5         |
| 3     | 1980.9          | 1986.5          | 1950.3        | 1957.6        | 1982.5         | 1962.2         | 1950.3         | 1992.1         |
| Avg   | 1972.3          | 1978.066667     | 1966.233333   | 1974.466667   | 1944.1         | 1946.566667    | 1920.633333    | 1944.333333    |

**Table A-13 Test Case 5—Root Bridge Down**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 466.4           | 466.6           | 539.1         | 544.3         | 580.7          | 591            | 859.5          | 862.3          |
| 2     | 477             | 477.2           | 529.1         | 534.3         | 797.3          | 797.8          | 1034.9         | 822.6          |
| 3     | 424.4           | 424.9           | 556.6         | 556           | 790.3          | 788.4          | 767.5          | 787.9          |
| Avg   | 455.9333333     | 456.2333333     | 541.6         | 544.8666667   | 722.7666667    | 725.7333333    | 887.3          | 824.2666667    |

**Table A-14 Test Case 6—Root Bridge Up**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 2280.2          | 2283.1          | 2040.4        | 2192.7        | 1969.2         | 1954.6         | 1931.6         | 2183           |
| 2     | 1476            | 1477.8          | 1973.7        | 1965.7        | 1935.9         | 1918.8         | 2125           | 2112.3         |
| 3     | 1979.9          | 1974.8          | 2111.5        | 2099.8        | 2357.8         | 2364.4         | 2013.2         | 1990.6         |
| Avg   | 1912.033333     | 1911.9          | 2041.866667   | 2086.066667   | 2087.633333    | 2079.266667    | 2023.266667    | 2095.3         |

**Table A-15 Test Case 7—Stack Master Down on CZ-C3750**

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 1429.6          | 1429.9          | 1740.8        | 1740.9        | 1927.3         | 1927.4         | 1937.4         | 1932           |
| 2     | 1650.1          | 1650.4          | 1748.6        | 1759.3        | 1607.2         | 1610.6         | 1646.5         | 1664.2         |
| 3     | 1139.9          | 1140.5          | 1718.8        | 1719.1        | 1417.3         | 1412.6         | 1346.5         | 1346.7         |
| Avg   | 1406.533333     | 1406.933333     | 1736.066667   | 1739.766667   | 1650.6         | 1650.2         | 1643.466667    | 1647.633333    |


Table A-16 Test Case 8—Stack Master Re-Established on CZ-C3750

| Run # | Baseline A to B | Baseline B to A | 50 MAC A to B | 50 MAC B to A | 100 MAC A to B | 100 MAC B to A | 200 MAC A to B | 200 MAC B to A |
|-------|-----------------|-----------------|---------------|---------------|----------------|----------------|----------------|----------------|
| 1     | 1981.8          | 1977.2          | 1948.3        | 1939.9        | 1928.4         | 1923.2         | 1925.1         | 1904           |
| 2     | 1978            | 1983.2          | 1968.5        | 1975.6        | 1935           | 1935.5         | 1904.1         | 1930.5         |
| 3     | 1996.1          | 1991.6          | 1952.1        | 1938.7        | 1953.2         | 1926.9         | 1990.8         | 1982.7         |
| Avg   | 1985.3          | 1984            | 1956.3        | 1951.4        | 1938.866667    | 1928.533333    | 1940           | 1939.066667    |

## Sample Trend Line for Link Failure Between Adjacent Switches

Figure A-6 shows the trend line for link failure between the C2955-7 and C2955-8 switches.

Figure A-6 Link Failure—C2955-7 to C2955-8



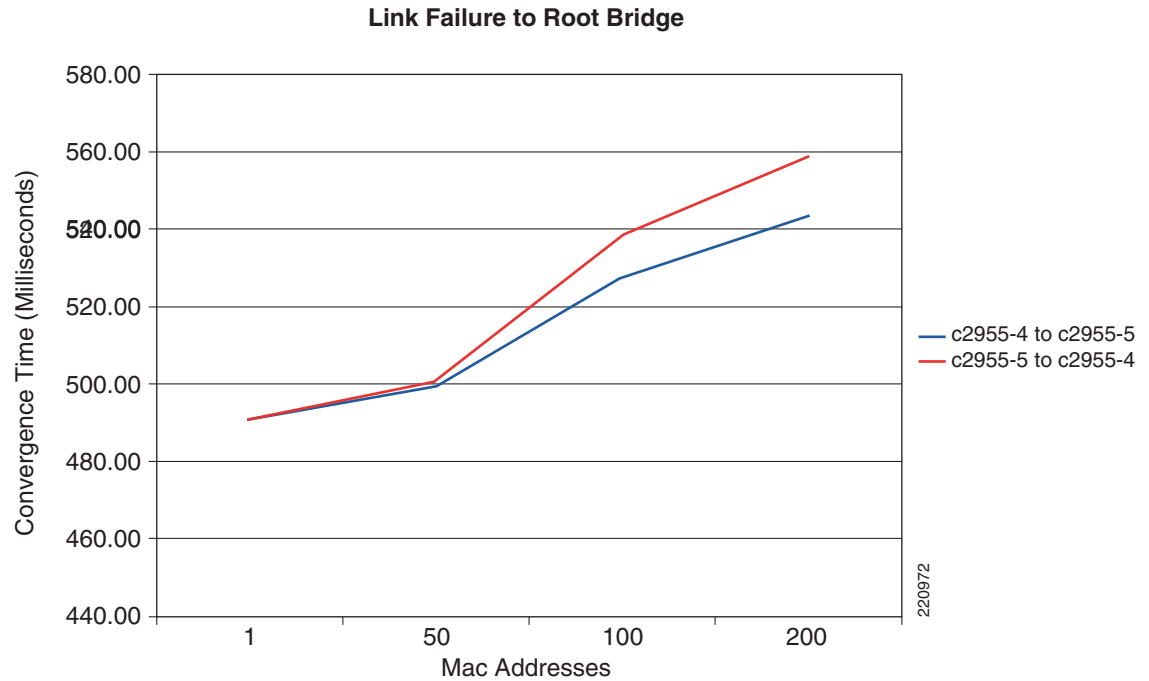
Key findings were as follows:

- ~.5 seconds with one MAC address, and without CIP and producer/consumer traffic
- ~2 seconds with between 50–200 MAC addresses with producer/consumer traffic
- Traffic load (CIP, producer/consumer) increases convergence time
- Number of MAC addresses not a large influence on convergence time

## Sample Trend Line for Link Failure To Root Bridge

Figure A-7 shows the trend line for link failure to the root bridge.

**Figure A-7** Link Failure to Root Bridge



Key findings were as follows:

- Convergence time is ~490–590 milliseconds, with a slight upward trend depending on number of MAC addresses.

## 16-Switch Ring—STP Testing

To verify some scaling parameters, testing was performed with double the amount of 2955s in the cell/area zone (16), and a spot check of certain test cases was performed to compare against the 8-switch ring STP tests. The same traffic flow and test methodology exists from the 8-switch tests except that there are more L2 hops from source to destination. Following are the tests and the corresponding results.

## Test Suite 1—Bidirectional Traffic from (Tx1 <-> Tx2)

**Table A-17** Test Case 1—Physically Remove Link between Cell-C2955-7 and Cell-C2955-8

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 1522.6           | 1522.5           | 1907.7            | 1902              | 3131.4            | 3131.4            |
| 2     | 1363.8           | 1363.7           | 2299.1            | 2289.2            | 2872.7            | 2858.7            |
| 3     | 1843             | 1834.4           | 1934.2            | 1923              | 2273.8            | 2273.7            |
| Avg   | 1576.466667      | 1573.533333      | 2047              | 2038.066667       | 2759.3            | 2754.6            |

**Table A-18** Test Case 2—Physically Reinsert Link between Cell-C2955-7 and Cell-C2955-8

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 954.6            | 968.9            | 1993.5            | 2005.5            | 1715.4            | 1743.9            |
| 2     | 1135.2           | 1149.3           | 1779.1            | 1795.2            | 1987.4            | 2004.4            |
| 3     | 1993.3           | 2003.8           | 1789.5            | 1805.9            | 1977.6            | 2003.2            |
| Avg   | 1361.033333      | 1374             | 1854.033333       | 1868.866667       | 1893.466667       | 1917.166667       |

Key findings were as follows:

- Very similar results as 8 switch tests
  - Slightly longer convergence times with 200\*2 MAC Addresses from ~2.1 seconds to ~2.7 seconds upon link failure
- Other tests not performed because of the similarity of the worst case scenario test from above



**Note** No baseline measurements were gathered (background traffic was always running).

## Test Suite 2—Bidirectional Traffic (Tx3 <-> Tx4)

**Table A-19** Test Case 1—Physically Remove Link between Cell-C2955-1 and CZ-C3750

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 585.8            | 586.2            | 653.1             | 645.1             | 615.7             | 598.8             |
| 2     | 606              | 601.1            | 508.3             | 500.5             | 673.8             | 654.2             |
| 3     | 581.2            | 576.6            | 573.6             | 564.4             | 648.1             | 628.4             |
| Avg   | 591              | 587.9666667      | 578.3333333       | 570               | 645.8666667       | 627.1333333       |

**Table A-20** Test Case 2—Physically Reinsert Link between Cell-C2955-1 and CZ-C3750

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 1809.3           | 1818.7           | 1758.9            | 1777.3            | 1787.5            | 1775.5            |
| 2     | 1820.8           | 1872.5           | 1783.9            | 1786.3            | 1707.8            | 1680.8            |
| 3     | 1817.2           | 1827.4           | 1815.4            | 1824.2            | 1795              | 1820.2            |
| Avg   | 1815.766667      | 1839.533333      | 1786.066667       | 1795.933333       | 1763.433333       | 1758.833333       |

Key findings were as follows:

- Very similar results as 8-switch tests
  - Slightly longer convergence time on link failure test (all MAC addresses), which was expected
- Other tests not performed because of the similarity of the worst-case scenario test from above



**Note** No baseline measurements were gathered (background traffic was always running).

## Redundant Star Topology—STP Testing

Although the majority of the testing was done with the ring topology, the redundant star topology was also tested for comparison purposes. (See [Cell/Area Network—Star Topology, page 2-22](#).) Ixia connections were used between two adjacent 2955 switches. The following test cases were performed.

**Table A-21** Test Case 1—Shut Non-Blocking Link on C2955-12

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 24.5             | 43.9             | 35.1              | 61.9              | 45.5              | 88.2              |
| 2     | 23.9             | 50.3             | 29.8              | 65.8              | 44.9              | 85.2              |
| 3     | 23               | 46.5             | 31.8              | 59.4              | 47.2              | 96.5              |
| Avg   | 23.8             | 46.9             | 32.23333333       | 62.36666667       | 45.86666667       | 89.96666667       |

**Table A-22** Test Case 2—Shut Non-Blocking Link on C2955-12

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 1999.7           | 2015             | 2017.9            | 2038.4            | 1986.8            | 2016.5            |
| 2     | 2000.5           | 2012.9           | 1997.7            | 2015              | 1992.6            | 2017.1            |
| 3     | 1998.8           | 2015.7           | 1991.3            | 2011.9            | 1994.8            | 2015.6            |
| Avg   | 1999.666667      | 2014.533333      | 2002.3            | 2021.766667       | 1991.4            | 2016.4            |

**Table A-23 Test Case 3—Physically Remove Non-Blocking Link on C2955-12 (Wire Cut Simulation)**

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 555.8            | 598.6            | 449.9             | 478.7             | 523.5             | 563.2             |
| 2     | 489.3            | 510.7            | 523.5             | 548.1             | 479.3             | 522.8             |
| 3     | 382              | 404.7            | 492.7             | 518.9             | 509.1             | 543.6             |
| Avg   | 475.7            | 504.6666667      | 488.7             | 515.2333333       | 503.9666667       | 543.2             |

**Table A-24 Test Case 4—Physically Re-Insert Non-Blocking Link on C2955-12 (Wire Cut Simulation)**

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 2000.7           | 2016.3           | 1993.5            | 2013.5            | 1990.6            | 2012.9            |
| 2     | 2004             | 2017.7           | 1996              | 2014.8            | 1992.8            | 2014              |
| 3     | 1994.9           | 2011.1           | 2001.2            | 2018.2            | 1993.9            | 2016.6            |
| Avg   | 1999.8666667     | 2015.0333333     | 1996.9            | 2015.5            | 1992.4333333      | 2014.5            |

**Table A-25 Test Case 5—Fail Root Bridge (Slot 1 on 3750)**

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 591              | 591.3            | 576               | 501.9             | 566.3             | 520               |
| 2     | 554.3            | 588.9            | 584.1             | 545.4             | 483.9             | 570.1             |
| 3     | 562.6            | 603              | 532.4             | 569.1             | 567               | 595               |
| Avg   | 569.3            | 594.4            | 564.1666667       | 538.8             | 539.0666667       | 561.7             |

**Table A-26 Test Case 5—Re-establish Root Bridge (Slot 1 on 3750)**

| Run # | 50 MAC<br>A to B | 50 MAC<br>B to A | 100 MAC<br>A to B | 100 MAC<br>B to A | 200 MAC<br>A to B | 200 MAC<br>B to A |
|-------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| 1     | 131.2            | 197.7            | 225.7             | 240.1             | 323.3             | 400.1             |
| 2     | 118.5            | 193.9            | 210               | 308.6             | 298.2             | 380               |
| 3     | 123.2            | 196.7            | 200.2             | 337.3             | 311               | 281.7             |
| Avg   | 124.3            | 196.1            | 211.9666667       | 295.3333333       | 310.8333333       | 353.9333333       |



Key findings were as follows:

- Best convergence times compared to ring-8 or ring-16, averaging 500ms consistently
- More consistent numbers
- Worse case for traffic flow is only two L2 hops away

## Latency/Jitter Testing

To characterize different network topologies under steady state, latency and jitter measurements were captured. Unlike spanning tree convergence testing, no failures were introduced. These tests assume that the network is functioning normally with typical control device traffic running in the background. Simulated source/destination patterns are worse-case scenarios (in the ring topologies) with traffic traversing the entire ring. This is done by having knowledge of the STP-blocked port before the testing begins.

The following test cases were completed for latency/jitter measurements:

- 8-switch ring—Bidirectional traffic 2955-5 <> 2955-4
- 16-switch ring—Bidirectional traffic 2955-8 <> 2955-9
- Hub/spoke—Bidirectional traffic 2955-12 <> 2955-13

The following results are in microseconds ( $\mu$ ):

**Table A-27 Latency/Jitter Test Results**

| Use Case | Latency<br>Tx3 >Tx4 | Jitter<br>Tx3 >Tx4 | Latency<br>Tx4 >Tx3 | Jitter<br>Tx4 >Tx3 |
|----------|---------------------|--------------------|---------------------|--------------------|
| 1        | 43.068              | 30.94              | 42.822              | 33.8               |
| 2        | 65.902              | 36.92              | 65.606              | 36.94              |
| 3        | 27.022              | 36.34              | 25.447              | 34.7               |

Key findings were as follows:

- Consistent with results from disruptive tests from above
  - Hub/spoke has the best latency, followed by 8-ring and 16-ring.
  - Jitter was consistent across all tests.

# IGMP Testing

## IGMP Snooping Test Methodology

The same eight devices that were used for the STP tests were also used to verify IGMP snooping. The testing exercised various combinations of IGMP snooping with and without a querier on various switches in the network, with active producer-consumer traffic running between PACs and a variable frequency drive. However, as per the Cisco recommended deployment, IGMP snooping works properly only in the presence of a querier. (For more information, see the following URL:

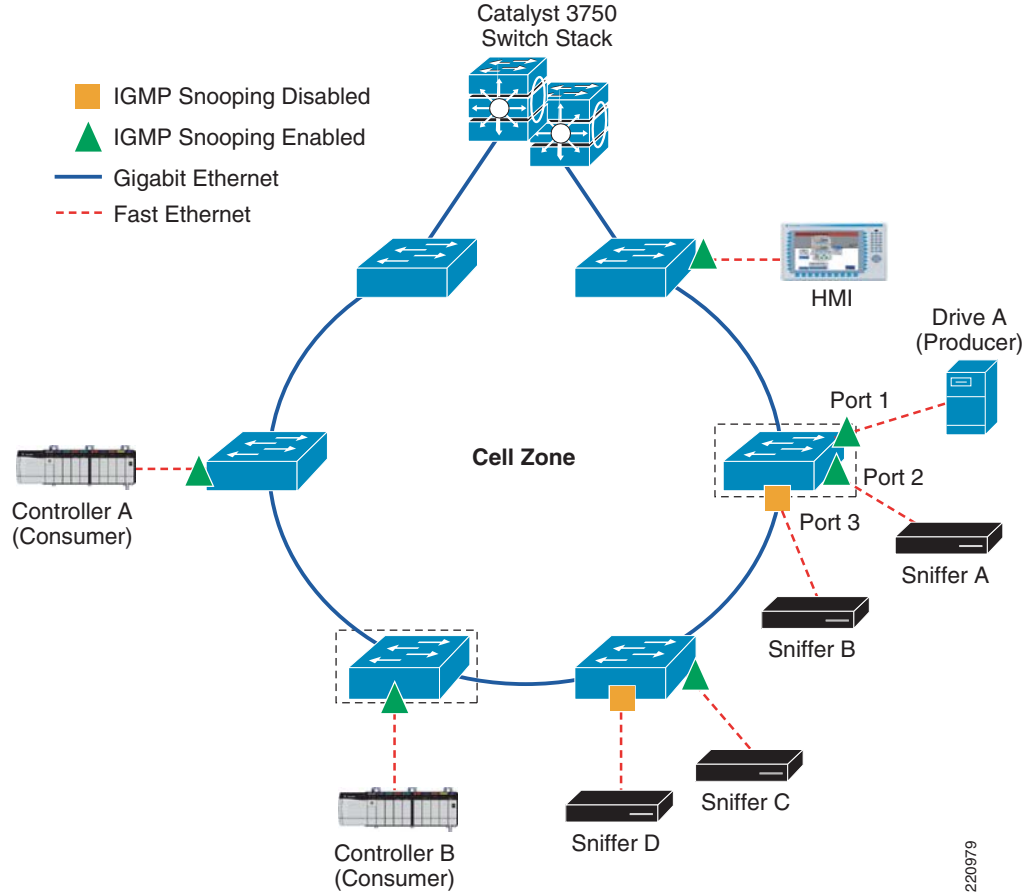
[http://www.cisco.com/en/US/products/hw/switches/ps708/products\\_tech\\_note09186a008059a9df.shtml](http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008059a9df.shtml)).

Thus, the only relevant test results are IGMP snooping with querier, and no IGMP snooping with or without querier enabled. A protocol analyzer (<http://www.wireshark.org>) was used to verify the presence of multicast data traffic.

## IGMP Snooping Test Topology

Figure A-8 shows the IGMP snooping test topology.

Figure A-8 IGMP Snooping Test Topology



In this topology, Controller A (Receiver/Consumer) is consuming traffic from Drive A (Producer/Source) across the ring topology. On the same switch that is producing multicast data traffic, a sniffer (Sniffer A) is connected on a different port that has IGMP snooping enabled to verify that this traffic is *not* seen on this port. The same verification is done on another port with IGMP snooping disabled. Finally, the test was repeated on a different switch in the ring with Sniffer C and Sniffer D, respectively.

## IGMP Snooping Test Results

Table A-28 shows the IGMP snooping test results. Note the following:

- Yes = Receiving multicast data traffic (destined to 239.x.x.x)
- No = Not receiving multicast data traffic (destined to 239.x.x.x)

Table A-28 IGMP Snooping Test Results

| Querier/Snooping Enabled | IGMP Client | Non-IGMP Client Same Switch | Non-IGMP Client Different Switch |
|--------------------------|-------------|-----------------------------|----------------------------------|
| OFF/OFF                  | Yes         | Yes                         | Yes                              |
| ON/ON                    | Yes         | No                          | No                               |





## APPENDIX **B**

# Configuration of the EttF Cell/Area Zone

---

## Layer 2 Configuration

Following is a sample configuration of one of the Layer 2 devices in the ring topology:

```
Current configuration : 3447 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname cell-c2955-9
!
enable password factory0
!
ip subnet-zero
!
!
!
spanning-tree mode rapid-pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 1-9,11-19,21-1024 hello-time 1
!
!
!
alarm profile defaultPort
!
alarm facility temperature primary relay major
alarm facility temperature primary syslog
alarm facility temperature primary notifies
!
!
interface FastEthernet0/1
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/2
description temp_L2_isolation_link
switchport access vlan 20
switchport trunk native vlan 20
```

```

switchport trunk allowed vlan 20
switchport mode trunk
shutdown
!
interface FastEthernet0/3
switchport access vlan 20
switchport mode access
shutdown
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/4
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/5
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/6
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/7
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/8
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/9
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/10
switchport access vlan 20
switchport mode access

```

```
spanning-tree portfast
spanning-tree bpduguard enable
spanning-tree bpdufilter enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/11
switchport access vlan 20
switchport mode access
spanning-tree portfast
spanning-tree bpdufilter enable
spanning-tree bpduguard enable
ip dhcp snooping limit rate 10
!
interface FastEthernet0/12
switchport access vlan 100
switchport mode access
spanning-tree portfast
!
interface GigabitEthernet0/1
switchport trunk native vlan 20
switchport trunk allowed vlan 20
switchport mode trunk
!
interface GigabitEthernet0/2
switchport trunk native vlan 20
switchport trunk allowed vlan 20
switchport mode trunk
!
interface Vlan1
no ip address
no ip route-cache
shutdown
!
interface Vlan2
no ip address
no ip route-cache
shutdown
!
interface Vlan20
ip address 10.17.20.50 255.255.255.0
no ip route-cache
!
ip default-gateway 10.17.20.1
ip http server
!
line con 0
line vty 0 4
password factory0
login
line vty 5 15
login
!
!
!
monitor session 1 source interface Fa0/12
end
```

## Layer 3 Configuration

Following is a sample configuration of the distribution/aggregation switch:

```

Current configuration : 10758 bytes
!
version 12.2
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname CZ-C3750-1
!
enable password factory0
!
username root privilege 15 password 0 factory0
aaa new-model
aaa authentication login default local
aaa authorization exec default local
aaa authorization network default local
!
aaa session-id common
clock timezone pst -8
switch 1 provision ws-c3750g-24ps
switch 2 provision ws-c3750g-24ps
vtp mode transparent
ip subnet-zero
no ip source-route
ip routing
ip cef load-sharing algorithm universal F9C26989
no ip domain-lookup
ip domain-name cisco.com
!
ip dhcp snooping vlan 10
ip dhcp snooping
ip multicast-routing distributed
ip ssh time-out 60
ip ssh authentication-retries 2
ip scp server enable
!
mls qos
!
crypto pki trustpoint TP-self-signed-1835000704
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-1835000704
 revocation-check none
 rsakeypair TP-self-signed-1835000704
!
!
crypto ca certificate chain TP-self-signed-1835000704
certificate self-signed 01
 30820290 308201F9 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
 54312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
 69666963 6174652D 31383335 30303037 30343121 301F0609 2A864886 F70D0109
 02161244 4D5A2D53 572D412E 63697363 6F2E636F 6D301E17 0D393330 33303130
 30303931 305A170D 32303031 30313030 30303030 5A305431 2F302D06 03550403
 1326494F 532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3138
 33353030 30373034 3121301F 06092A86 4886F70D 01090216 12444D5A 2D53572D
 412E6369 73636F2E 636F6D30 819F300D 06092A86 4886F70D 01010105 0003818D
 00308189 02818100 B86E69EB 3AD7C959 9F2CD10B BDFAB40D 6BF1DD24 06AB79E6
 4A27520F 5896ACE0 B9BE5788 A63AD836 2FD31A48 5C646E3D 2E1E19FE 2858CB63
 DB826F7E 09149DBD C5AE578E C859059A C6A4727F CD1BDB06 C24632C3 E7D7A082

```



```

C00FCAD9 F84166F5 8D1E5202 742398FF D55D5323 1AAA7050 9880BE4C 08C363E3
2E46C259 6BF053E5 02030100 01A37230 70300F06 03551D13 0101FF04 05300301
01FF301D 0603551D 11041630 14821244 4D5A2D53 572D412E 63697363 6F2E636F
6D301F06 03551D23 04183016 80140689 AC22B76B 6ED2E37D 87E03F3E 0ED65D3F
C313301D 0603551D 0E041604 140689AC 22B76B6E D2E37D87 E03F3E0E D65D3FC3
13300D06 092A8648 86F70D01 01040500 03818100 73C19D50 C99E2764 95C874E7
84B1302F 5A0DDD98 E197BBEE 494B4C34 F1A30F05 55E1773D 957D3F05 69DAF284
648E4AB9 62F3716A 612AEE09 A35D122D B67644C4 84836AD5 DB17AFE2 CDC9781A
8A54FBD0 CAF9763D E32C4C8E 07D4BB89 8699E62E 9CABE244 FE93A53C FF48CF4F
C50EF6E1 4D522967 6C3020A5 9D80D5FF 66E6C1AD
quit
!
!
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree extend system-id
spanning-tree vlan 1-1024 priority 4096
!
vlan internal allocation policy ascending
!
vlan 2-100,200
!
vlan 250
name management
!
class-map match-all cip-priority-class
match access-group name cip-priority
class-map match-all cip-consumer-class
match access-group name cip-consumer
class-map match-all cip-producer-class
match access-group name cip-producer
!
!
policy-map cip-policy
class cip-producer-class
set ip precedence 4
class cip-consumer-class
set ip precedence 4
policy-map cip-egress-policy
!
!
!
interface Port-channel1
description CZ-C4500-1
no switchport
ip address 10.18.3.100 255.255.255.0
!
interface Port-channel2
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 10
switchport mode trunk
!
interface Port-channel3
description CZ-C4500-2
no switchport
ip address 10.18.4.100 255.255.255.0
!
interface GigabitEthernet1/0/1
switchport mode access
spanning-tree portfast
!
interface GigabitEthernet1/0/2
no switchport

```

```
ip address 172.28.212.12 255.255.255.0
!
interface GigabitEthernet1/0/3
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/4
 switchport access vlan 10
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/5
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/6
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/7
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/8
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/9
 no switchport
 no ip address
 channel-group 1 mode active
 spanning-tree portfast
!
interface GigabitEthernet1/0/10
 no switchport
 no ip address
 channel-group 1 mode active
 spanning-tree portfast
!
interface GigabitEthernet1/0/11
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/12
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet1/0/13
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport trunk allowed vlan 10
 switchport mode trunk
 uddl port aggressive
 spanning-tree guard root
 spanning-tree vlan 10 port-priority 0
!
interface GigabitEthernet1/0/14
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 20
 switchport trunk allowed vlan 20
 switchport mode trunk
 uddl port aggressive
!
interface GigabitEthernet1/0/15
```

```
 switchport access vlan 30
 switchport mode access
 !
interface GigabitEthernet1/0/16
 switchport access vlan 3
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/17
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/18
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/19
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/20
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/21
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast
 !
interface GigabitEthernet1/0/22
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
 !
interface GigabitEthernet1/0/23
 switchport access vlan 30
 switchport mode access
 !
interface GigabitEthernet1/0/24
 switchport access vlan 30
 switchport mode access
 !
interface GigabitEthernet1/0/25
 no switchport
 no ip address
 channel-group 1 mode active
 !
interface GigabitEthernet1/0/26
 no switchport
 no ip address
 channel-group 1 mode active
 !
interface GigabitEthernet1/0/27
 no switchport
 no ip address
 channel-group 3 mode active
 !
interface GigabitEthernet1/0/28
 no switchport
 no ip address
```

```
channel-group 3 mode active
!
interface GigabitEthernet2/0/1
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/2
 no switchport
 no ip address
 spanning-tree portfast
!
interface GigabitEthernet2/0/3
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/4
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/5
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/6
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/7
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/8
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/9
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/10
 switchport access vlan 2
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/11
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10
 switchport mode trunk
 channel-group 2 mode active
 spanning-tree portfast
!
interface GigabitEthernet2/0/12
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10
 switchport mode trunk
 channel-group 2 mode active
```

```
 spanning-tree portfast
!
interface GigabitEthernet2/0/13
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 10
 switchport trunk allowed vlan 10
 switchport mode trunk
 uddld port aggressive
 spanning-tree guard root
 spanning-tree vlan 10 port-priority 16
!
interface GigabitEthernet2/0/14
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 20
 switchport trunk allowed vlan 20
 switchport mode trunk
 uddld port aggressive
!
interface GigabitEthernet2/0/15
 switchport access vlan 30
 switchport mode access
!
interface GigabitEthernet2/0/16
 switchport access vlan 3
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/17
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/18
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/19
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/20
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/21
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast
!
interface GigabitEthernet2/0/22
 switchport access vlan 250
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/23
 switchport access vlan 200
 switchport mode access
 spanning-tree portfast
!
interface GigabitEthernet2/0/24
 switchport access vlan 200
```

```

switchport mode access
spanning-tree portfast
!
interface GigabitEthernet2/0/25
!
interface GigabitEthernet2/0/26
!
interface GigabitEthernet2/0/27
!
interface GigabitEthernet2/0/28
!
interface Vlan1
no ip address
no ip mroute-cache
!
interface Vlan2
no ip address
!
interface Vlan10
ip address 10.17.10.1 255.255.255.0
ip pim sparse-dense-mode
!
interface Vlan20
ip address 10.17.20.1 255.255.255.0
ip pim sparse-dense-mode
!
interface Vlan30
ip address 10.17.30.1 255.255.255.0
ip pim sparse-dense-mode
!
interface Vlan250
ip address 172.16.250.3 255.255.255.0
!
router rip
version 2
redistribute connected metric 1
network 10.0.0.0
!
ip default-gateway 172.28.212.1
ip classless
ip route 0.0.0.0 0.0.0.0 172.28.212.1
ip http server
ip http port 2222
ip http authentication local
ip http secure-server
!
!
ip access-list extended cip-consumer
permit udp any eq 2222 any
ip access-list extended cip-priority
permit ip any any tos max-throughput
ip access-list extended cip-producer
permit udp any any eq 2222
!
logging source-interface Vlan10
logging 10.18.2.201
snmp-server community public RO
snmp-server community private RW
snmp-server community marstring RO
snmp-server host 10.18.2.201 marstring
radius-server source-ports 1645-1646
!
control-plane
!

```

```
!
line con 0
 exec-timeout 0 0
line vty 0 4
 rotary 1
 transport input ssh
line vty 5 15
 password factory0
 rotary 1
!
!
end
```







# APPENDIX **C**

## Configuration of the EttF Demilitarized Zone

---

### Security Configuration

#### ASA Configuration

```
ASA Version 7.2(2)
!
hostname DMZ-ASA-1
domain-name cisco.com
enable password 7w22FjI5eWall1BPD encrypted
names
!
interface GigabitEthernet0/0
nameif inside
security-level 100
ip address 10.18.1.1 255.255.255.0 standby 10.18.1.3
!
interface GigabitEthernet0/1
nameif outside
security-level 0
ip address 10.13.2.1 255.255.255.248 standby 10.13.2.3
!
interface GigabitEthernet0/2
nameif DMZ
security-level 50
ip address 10.19.2.9 255.255.255.248 standby 10.19.2.10
!
interface GigabitEthernet0/3
description LAN/STATE Failover Interface
!
interface Management0/0
nameif management
security-level 100
ip address 172.28.212.31 255.255.255.0
management-only
!
passwd 2KFQnbNIdI.2KYOU encrypted
boot system disk0:/asa722-k8.bin
ftp mode passive
dns server-group DefaultDNS
domain-name cisco.com
access-list outside extended permit tcp any any eq telnet
access-list outside extended permit tcp any any eq www
access-list outside extended permit icmp any any
access-list INSAUTH extended permit tcp any host 10.19.2.5 eq telnet
```

```

access-list INSAUTH extended permit tcp any host 10.19.2.5 eq www
access-list INSAUTH extended permit tcp any host 10.19.2.5 eq 8080
access-list dmz extended permit tcp any any eq telnet
access-list dmz extended permit tcp any any eq www
access-list dmz extended permit icmp any any
access-list DMZ_authentication extended permit tcp any any eq telnet
access-list inside extended permit tcp any any eq www
access-list inside extended permit tcp any any eq https
access-list inside extended permit icmp any any
access-list inside extended permit tcp any host 10.19.2.1 eq telnet
access-list OUTAUTH extended permit tcp any host 10.18.1.2 eq telnet
access-list OUTAUTH extended permit tcp any host 10.18.1.2 eq www
access-list ips-acl extended permit ip any any
access-list ips-acl extended permit icmp any any
pager lines 24
logging enable
logging buffered debugging
logging trap debugging
logging host management 172.28.212.22
mtu inside 1500
mtu outside 1500
mtu DMZ 1500
mtu management 1500
ip verify reverse-path interface inside
ip verify reverse-path interface outside
ip verify reverse-path interface DMZ
ip verify reverse-path interface management
failover
failover lan unit primary
failover lan interface failover GigabitEthernet0/3
failover replication http
failover link failover GigabitEthernet0/3
failover interface ip failover 10.18.2.33 255.255.255.248 standby 10.18.2.34
icmp unreachable rate-limit 1 burst-size 1
asdm image disk0:/asdm-521.bin
asdm history enable
arp timeout 14400
access-group inside in interface inside
access-group outside in interface outside
access-group dmz in interface DMZ
route inside 10.17.0.0 255.255.0.0 10.18.1.5 1
route inside 10.18.0.0 255.255.0.0 10.18.1.5 1
route DMZ 10.19.0.0 255.255.0.0 10.19.2.1 1
route management 171.70.0.0 255.255.0.0 172.28.212.1 1
route management 172.0.0.0 255.0.0.0 172.28.212.1 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout uauth 0:05:00 absolute
aaa-server ETTF2 protocol tacacs+
aaa-server ETTF2 (DMZ) host 10.19.2.11
key cisco
username root password /bieFEvWpEclHwvP encrypted privilege 15
aaa authentication match OUTAUTH outside ETTF2
aaa authentication ssh console LOCAL
aaa authentication serial console LOCAL
aaa authentication http console LOCAL
aaa authentication match INSAUTH inside ETTF2
aaa authentication match DMZ_authentication DMZ ETTF2
http server enable
http 0.0.0.0 0.0.0.0 management
snmp-server host management 172.28.212.22 community marstring
no snmp-server location

```

```

no snmp-server contact
snmp-server community marstring
snmp-server enable traps snmp authentication linkup linkdown coldstart
virtual telnet 10.18.1.254
telnet 0.0.0.0 0.0.0.0 DMZ
telnet timeout 1440
ssh scopy enable
ssh 10.18.0.0 255.255.0.0 inside
ssh 10.17.0.0 255.255.0.0 inside
ssh 10.19.0.0 255.255.0.0 DMZ
ssh 0.0.0.0 0.0.0.0 management
ssh timeout 10
console timeout 0
!
class-map ips-class
match access-list ips-acl
class-map inspection_default
match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
parameters
message-length maximum 512
policy-map ips-policy
class ips-class
ips inline fail-close
policy-map global_policy
class inspection_default
inspect dns preset_dns_map
inspect ftp
inspect h323 h225
inspect h323 ras
inspect netbios
inspect rsh
inspect rtsp
inspect skinny
inspect esmtp
inspect sqlnet
inspect sunrpc
inspect tftp
inspect sip
inspect xdmcp
!
service-policy global_policy global
service-policy ips-policy interface inside
service-policy ips-policy interface outside
service-policy ips-policy interface DMZ
webvpn
csd image disk0:/securedesktop-asa-3.1.1.29-k9.pkg
csd enable
prompt hostname context
Cryptochecksum:dd189225023b09b212fb39b73974edad
: end

```

## IPS Configuration

```

! -----
! Current configuration last modified Thu Mar 29 23:03:06 2007
! -----
! Version 6.0(1)
! Host:

```

```

! Realm Keys key1.0
! Signature Definition:
! Signature Update S263.0 2006-12-18
! Virus Update V1.2 2005-11-24
! -----
service interface
exit
! -----
service authentication
exit
! -----
service event-action-rules rules0
exit
! -----
service host
network-settings
host-ip 172.28.212.32/24,172.28.212.1
host-name dmz-ssm-1
telnet-option disabled
access-list 0.0.0.0/0
login-banner-text You are logging on to AIP-SSM of DMZ-ASA-1
exit
time-zone-settings
offset 0
standard-time-zone-name UTC
exit
exit
! -----
service logger
exit
! -----
service network-access
exit
! -----
service notification
trap-destinations 172.28.212.22
trap-community-name marstring
trap-port 162
exit
enable-notifications true
enable-set-get true
read-only-community marstring
read-write-community marstring
trap-community-name marstring
exit
! -----
service signature-definition sig0
signatures 2000 0
status
enabled false
exit
exit
signatures 2004 0
alert-severity high
engine atomic-ip
event-action produce-alert|produce-verbose-alert
exit
alert-frequency
summary-mode fire-all
summary-key AxBx
exit
exit
status
enabled true

```

```
exit
exit
exit
! -----
service ssh-known-hosts
exit
! -----
service trusted-certificates
exit
! -----
service web-server
exit
! -----
service anomaly-detection ad0
exit
! -----
service external-product-interface
exit
! -----
service analysis-engine
exit
```





## APPENDIX **D**

# EttF High Availability Testing

---

High availability is critical to maximize network and system uptime, thereby meeting predefined SLAs. This section outlines the validation methodology and the corresponding results of the testing.

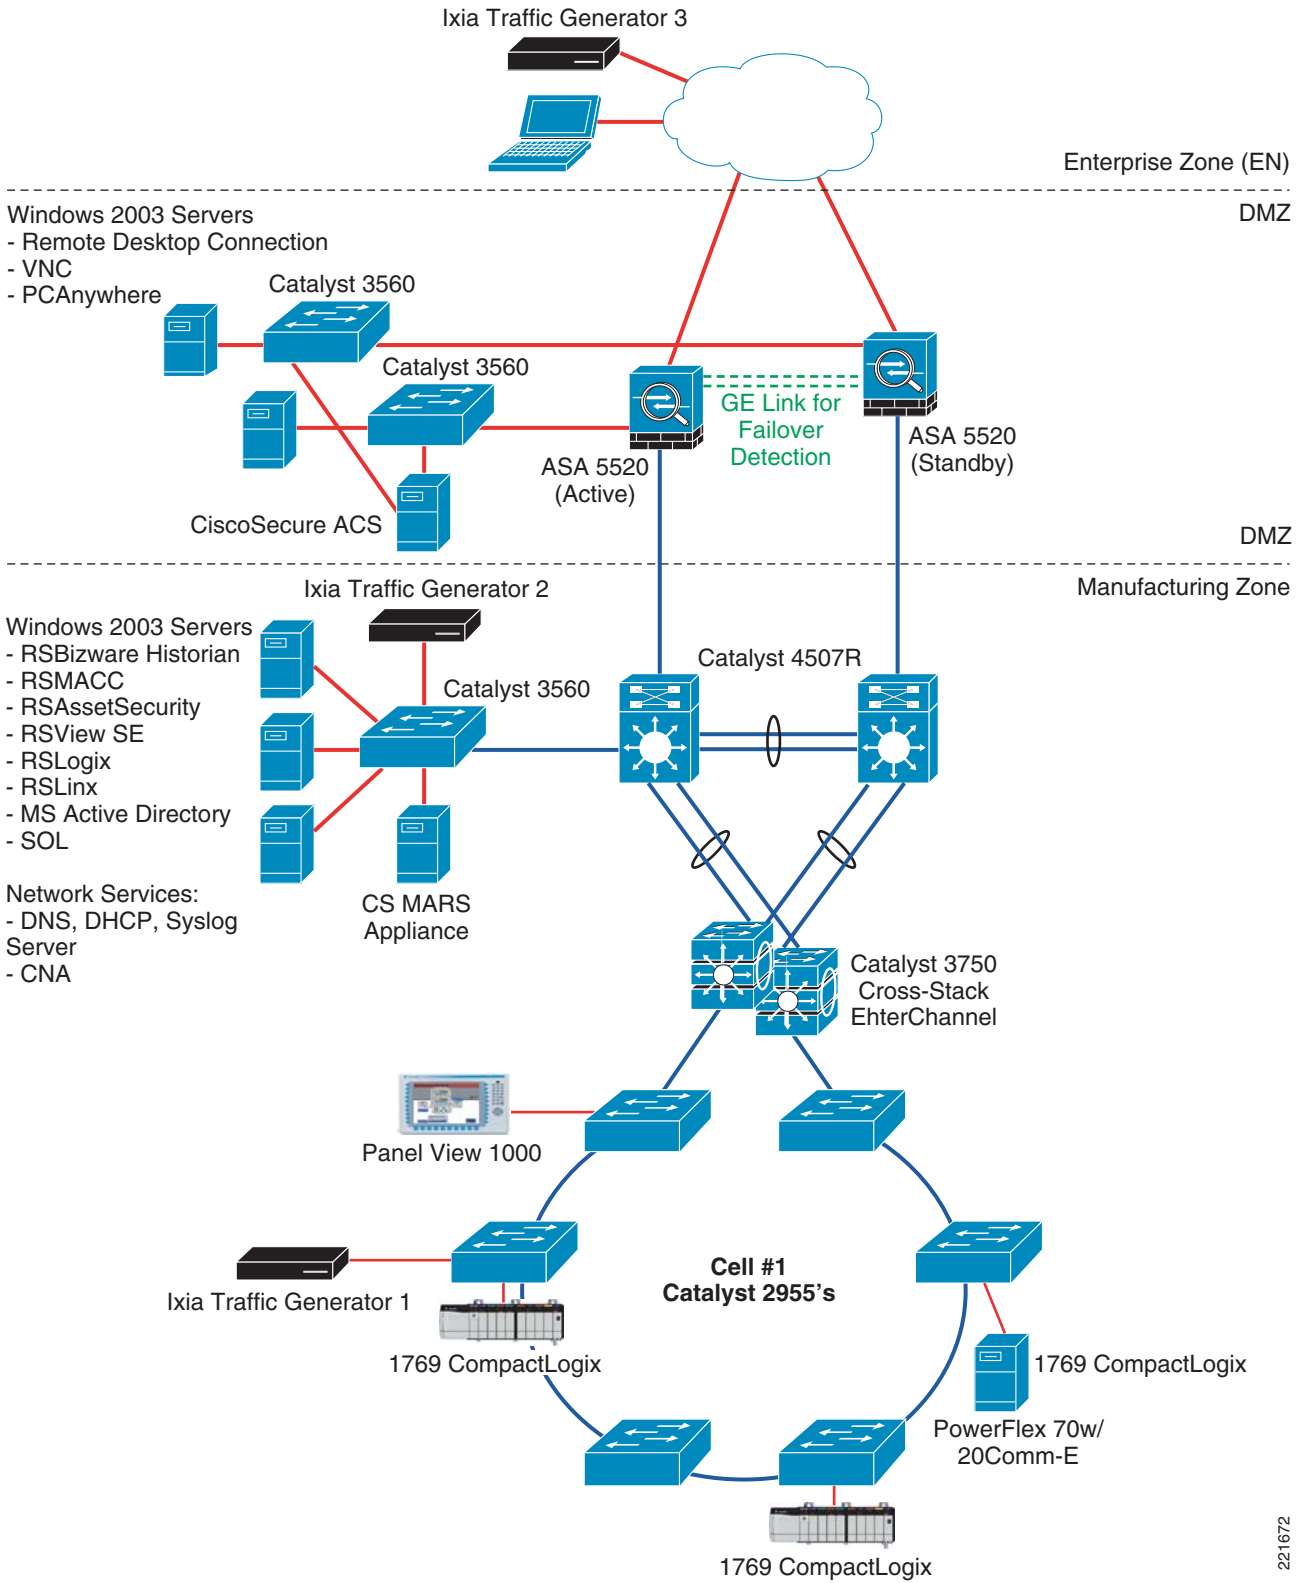
## HA Test Methodology

Each layer in the EttF solution is verified for HA functionality and recovery times. Simulated Layer 3 traffic is traversing end-to-end from the cell/area zone through the manufacturing zone, and from the cell/area zone through the DMZ to the outside world. Various failures are triggered at each layer, and the convergence time is measured to characterize and quantify the impact on availability.

## HA Test Topology

[Figure D-1](#) shows the test topology.

Figure D-1 HA Test Topology



221672



Layer 3 traffic is flowing from Tx1 → Tx2 and from Tx1 → Tx3. Tx2 and Tx3 are injecting 1000 simulated OSPF routes into the network and Tx1 is sending to all the simulated routes. The idea is to simulate TCP-based traffic as if originating from an application server such as Historian.

## HA Test Scenarios

Three test suites are explored to characterize different failure/recovery times at different layers in the EttF design. Various disruptions are initiated at the cell/area zone, manufacturing zone, and DMZ levels. With each failure, convergence time is measured using the following formula:

$$[(Tx - Rx) / \text{packet rate}] * 1000$$

Where:

Tx = Packets transmitted

Rx = Packets received

PPS = 10,000 pps

### Test Suite 1—HA in the Cell/Area Zone (Tx1 → Tx2)

- Use Case 1—Fail master stack with **stack-mac persistent** enabled
- Use Case 2—Fail slave stack with **stack-mac persistent** enabled
- Use Case 3—Fail master stack with HSRP subsecond timers configured
- Use Case 4—Fail slave stack with HSRP subsecond timers configured



#### Note

HSRP with only one logical router was selected so that the end nodes would have a VIP and Virtual MAC that remains persistent. It turns out that the better solution is to enable the **stack-mac persistent 0** feature as indicated by the results.

### Suite 1 Test Results

The Suite 1 test results (all in milliseconds) are provided in the following tables.

**Table D-1** Test Case 1—Fail Master Stack with *stack-mac persistent Enabled*

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 1988.2     |
| 2     | 1998.6     |
| 3     | 2160.3     |
| Avg   | 2049.033   |

**Table D-2 Test Case 2—Fail Slave Stack with stack-mac persistent Enabled**

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 1201.4     |
| 2     | 1091.4     |
| 3     | 1183.6     |
| Avg   | 1158.8     |

**Table D-3 Use Case 3—Fail Master Stack with HSRP Subsecond Timers Configured**

| Run # | Tx1 -> Tx2  |
|-------|-------------|
| 1     | 14994.5     |
| 2     | 13726.1     |
| 3     | 14106.9     |
| Avg   | 14275.83333 |

**Table D-4 Use Case 4—Fail Slave Stack with HSRP Subsecond Timers Configured**

| Run # | Tx1 -> Tx2  |
|-------|-------------|
| 1     | 2136.9      |
| 2     | 2126.9      |
| 3     | 2086.1      |
| Avg   | 2116.633333 |

## Test Suite 2—HA in the Manufacturing Zone (Tx1 → Tx2)

- Use Case 1—Fail physical link in EtherChannel to 4500-1
- Use Case 2—Fail physical link in EtherChannel to 4500-2
- Use Case 3—Supervisor failover on 4500-1
- Use Case 4—Supervisor failover on 4500-2

## Suite 2 Test Results

The Suite 2 test results (all in milliseconds) are provided in the following tables.

**Table D-5**      **Test Case 1—Fail Physical Link in EtherChannel to 4500-1**

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 1.4        |
| 2     | 1.4        |
| 3     | 1.6        |
| Avg   | 1.466667   |

**Table D-6**      **Test Case 2—Fail Physical Link in EtherChannel to 4500-2**

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 1.8        |
| 2     | 1.4        |
| 3     | 1.4        |
| Avg   | 1.533333   |

**Table D-7**      **Use Case 3—Supervisor Failover on 4500-1**

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 16.3       |
| 2     | 16.1       |
| 3     | 16         |
| Avg   | 16.13333   |

**Table D-8**      **Test Case 4—Supervisor Failover on 4500-2**

| Run # | Tx1 -> Tx2 |
|-------|------------|
| 1     | 15.9       |
| 2     | 18.2       |
| 3     | 16.1       |
| Avg   | 16.73333   |

## Test Suite 3—HA in the DMZ (Tx1 → Tx3)

Test suite 3 removes the control network-facing interface on the active ASA, as follows:

1. Active ASA link failure on control network-facing interface
2. Active ASA link failure on DMZ-facing interface
3. Standby ASA link failure on control network-facing interface
4. Standby ASA link failure on DMZ-facing interface
5. Reload of active ASA
6. Reload of standby ASA
7. Run “failover active” on the active ASA
8. Run “failover active” on the standby ASA
9. 4500-1 switchover to standby supervisor (HSRP active)
10. 4500-1 chassis failure (HSRP active)
11. 4500-2 switchover to standby supervisor
12. 4500-2 chassis failure

### Suite 3 Test Results

**Table D-9 Suite 3 Test Results**

| Use Case | Tx1 → Tx3 |
|----------|-----------|
| 1        | 5610.9    |
| 2        | 6646.3    |
| 3        | 0         |
| 4        | 0         |
| 5        | 7189.9    |
| 6        | 0         |
| 7        | 0         |
| 8        | 134.2     |
| 9        | 270.8     |
| 10       | 25759.6   |
| 11       | 0         |
| 12       | 0         |

## Test Tools

The following equipment is needed for performing these tests:

- 16 Cisco Catalyst C2955T-12 industrial switches
- 2 Cisco Catalyst WS-C3750G-24PS (stacked)
- 2 fully-redundant Cisco Catalyst 4507R switches with Supervisor IV
- 1 Ixia traffic generator
- Various Rockwell Automation (RA) equipment

