Building Automation System over IP (BAS/IP) Design and Implementation Guide

Cisco Validated Design

15 August 2008 v8.1

This design and implementation guide represents a collaborative development effort from Cisco Systems and Johnson Controls. It is built on, and adds to, design guidelines from the Cisco Connected Real Estate program and the Johnson Controls Network and Information Technology Considerations Technical Bulletin.

Cisco
Connected Real Estate Practice
170 West Tasman Drive
San Jose, CA 95134-1706
http://www.cisco.com

Johnson Controls
Building Efficiency
507 E. Michigan Street
Milwaukee, WI 53202
www.johnsoncontrols.com

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.

Table of Contents

Chapter 1 Solution Overview

Executive Summary

This Design and Implementation Guide (“DIG”) represents a collaborative effort between Cisco Systems and Johnson Controls in support of the JCI-Cisco Strategic Alliance. The intent of this document is to provide guidance to the respective sales and technical organizations of each party in the Alliance for the design, implementation and operation of secure, scaleable and repeatable Building Automation System (“BAS”) networks running on Cisco IP networks.

The Building Automation Systems market is accelerating quickly towards a converged model where CTOs and CIOs of organizations are working with their respective counterparts in the Facilities and Real Estate departments of their organizations to deploy integrated networks. There has been resistance in the market to deploy BAS systems over production IT networks due to concerns about security, access and availability of critical BAS functionality during network outages.

Traditional BAS systems have historically used direct-digital control communication protocols over RS-485 low voltage control networks on dedicated wiring at the field-bus and device level. Until recently, the protocols that support these communications have not provided any security in terms of data transmission to the extent that is required or expected of a typical IT network.

With the right approach to network security and provisioning of a BAS system using the BACnet protocol, it is possible to support the deployment and operation of BAS system over an IP network. This Design and Implementation Guide is a foundational document intended to support a wide spectrum of BAS applications in secure network environments. Subsequent versions of this document will include use cases for varying building types and applications.

This DIG represents the **Foundational Infrastructure** phase of a comprehensive Connected Real Estate (“CRE”) framework. The foundational elements comprising this infrastructure include:

- Advanced Routing and Switching
- Network Security
- Active Directory with Core Users
- LAN/WAN Firewalls / Security Provisioning
- VLAN Segmentation
- Network Management
- VPN Access Control
- Device Level Security
- BAS – Building Automation Systems, including but not limited to, control systems for HVAC, Lighting, Fire Alarm Monitoring, Elevator Controls and Energy-Monitoring Systems

Other CRE Solutions that will be deployed on top of this foundation will include such applications such as IP Telephony, Unified Communications, Data Center Management, Physical Security, Digital Signage and many more specialized applications. These applications will be defined as subsequent use cases beyond the scope of this DIG. The next phase beyond this DIG will require definition and testing of the appropriate use cases that will be prioritized as part of the CRE Solutions Roadmap and development strategy.

Document Organization

This document contains the following chapters and appendices:

Chapter or Appendix	Description
Chapter 1 – "Solution Overview"	Provides an Executive Summary, Building Automation Systems – Industry Background and Context, Description and Justification for BAS over IP, Target Market Opportunities and Solution Benefits
Chapter 2 – "Solution Architecture"	This chapter provides an overview of the Johnson Controls Metasys Facility Management Networking (FMN) 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"	The focus of this chapter is on basic Cisco network design principles and the networking of those IP enabled devices in each of the subsystems in the Facilities Management Network.
Chapter 4 – "Implementation of Security"	This chapter describes the implementation of the Secure Architecture for Intelligent Facility Applications known as SAIFA v1.0.
Appendix A – "Reference Architecture Diagrams"	Visio diagrams illustrating Johnson Controls systems running on a Cisco IP Network
Appendix B – "Glossary and Acronym List"	A table of most, if not all, of the terms and acronyms used in this document.

Building Automation Systems - Industry Background and Context

It could be said that the Building Automation Systems (BAS) Industry has been in existence since Warren Johnson patented the first temperature control system in 1895 [1]. The industry has evolved in many ways since then, and we are entering into a new era of building intelligence, machine-to-machine communications and expanded functionality in ways that we have not yet imagined.

The latest trend that is dramatically impacting our industry is that of controlling and monitoring building automation controls over IP networks. This trend has accelerated in the past 3-5 years with the availability and proliferation of IP-based control systems and adoption of web services over IP networks.

The adoption and prevalence of this industry trend is explained in detail in a report prepared by Frost & Sullivan entitled "*Impact Analysis of IP Protocols on Building Automation"* [2]

" IP-based building automation systems are gaining momentum, thanks to the various contributing factors ranging from internet penetration to cheaper computing devices and platforms. Information technology is a powerful tool, and enterprises could effectively exploit the existing infrastructure to integrate building systems into them, enabling remote access, management and distributed control.

Technologists designed the next generation IP technology, called IPv6, in such a way that there is an IP address available for virtually every grain of sand on earth. Obviously, the number of devices, applications and services based on IP technologies are growing exponentially, andit is imperative to have sufficient IP addresses to cater to the same.

Pertaining to the application type and requirements, building owners and enterprises can exploit the flexibility of IP technologies to realize interoperability and convergence. It is to be noted that there are many advantages in opting IP technologies, poor/careless planning of the network infrastructure would lead to disruption in business and damage to property. The network must be customized to provide a perfect balance between the capital and the security of the network. It has been found that in most cases the users contribute to security problems, and hence their knowledge and perspective of network security must be enhanced for the benefit of the building."

This Design & Implementation Guide ("DIG") is specifically focused on the Johnson Controls Metasys system program and related technologies from Cisco Systems. In this DIG, we are going to concentrate on those technologies and services that are available today.

While there are many exciting developments on the horizon, the scope of this DIG will be limited to generally available hardware and software systems that are in current release, are supported by Johnson Controls and Cisco Systems and can be easily obtained by our mutual customers using established channels and methods of procurement, installation and support.

[1] The first complete Automatic Temperature Control System [economical to install and operate, long-lasting, and extremely effective in maintaining a constant temperature] was patented in 1895 by Warren S. Johnson. (per ASME - American Society of Mechanical Engineers)

[2] *IMPACT ANALYSIS OF IP PROTOCOLS ON BUILDING AUTOMATION – Frost & Sullivan – Report DA09* - © 2007 Frost & Sullivan

There are several standards that describe how Building Automation Systems are to be designed and implemented. The most predominant standards that affect this DIG are the following:

Construction Specifications Institute - Over the last forty years, *MasterFormat*™ has become the leading standard for organizing nonresidential construction specifications, and is now almost omnipresent in the AEC industry, thanks to the many applications for which it has been utilized. In 2001, the Construction Specifications Institute (CSI), along with sister organization Construction Specifications Canada (CSC), charged the *MasterFormat* Expansion Task Team with examining whether there was a need to revise and possibly expand the 1995 edition of *MasterFormat* to accommodate changes that have taken place in the industry since that version was published.

MasterFormat 2004 Edition: Numbers and Titles is a master list of numbers and subject titles for organizing information about construction work results, requirements, products, and activities into a standard sequence. Construction projects use many different delivery methods, products, and installation methods. Successful completion of projects requires effective communication among the people involved. Information retrieval is nearly impossible without a standard filing system familiar to each user. *MasterFormat Numbers and Titles* facilitate standard filing and retrieval schemes throughout the construction industry. *MasterFormat Numbers and Titles* are suitable for use in project manuals, for organizing cost data, reference keynotes on drawings, for filing product information and other technical data, for identifying drawing objects and for presenting construction market data.

Each *MasterFormat* number and title defines a “section,” arranged in “levels” depending on their breadth of coverage. The broadest collections of related construction products and activities are level one titles, otherwise known as “divisions.” Each division in the *MasterFormat 2004 Edition: Numbers and Titles* is made up of level two, level three, and occasionally level four numbers and titles assigned by *MasterFormat*, each of which delineate a gradually more detailed area of work results to be specified.

The sections and levels most applicable to this DIG are the sections that pertain to the following sections in the Facility Services Subgroup:

23 – Heating Ventilating and Air Conditioning: HVAC subjects relocated from Division 15 in *MasterFormat 1995 Edition*.
25 – Integrated Automation: Expanded integrated automation subjects relocated from Division 13 in *MasterFormat 1995 Edition*.
26 – Electrical: Electrical and lighting subjects relocated from Division 16 in *MasterFormat 1995 Edition*.
27 – Communications: Expanded communications subjects relocated from Division 16 in *MasterFormat 1995 Edition*.

The second major standard that impacts this DIG is **ISO Standard 16484.**

ISO 16484-2:2004 specifies the requirements for the hardware to perform the tasks within a building automation and control system (BAS). It provides the terms, definitions and abbreviations for the understanding of ISO 16484-2 and ISO 16484-3. ISO 16484-2:2004 relates only to physical items/devices, i.e. devices for management functions, operator stations and other human system interface devices; controllers, automation stations and application specific controllers; field devices and their interfaces; cabling and interconnection of devices; engineering and commissioning tools.

ISO 16484-2:2004 shows a generic system model to which all different types of BACS and their interconnections (BACS network) can fit. A graphical concept of the BACS network in terms of LAN topology will be provided in ISO 16484-5.

A copy of this standard can be downloaded from the ISO website:

http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=29682

An excellent view and presentation of this standard within the context of this DIG is provided by Steve Tom, PE, PhD, Director of Technical Information, Automated Logic Corporation in an article published by Automated Buildings.com. Here is the link to the article:

(http://www.automatedbuildings.com/news/dec04/articles/alc/stom.htm)

An excerpt of this article is provided below:

“Within the BAS industry, Web services are already being used world-wide to import HVAC after-hours use and utility meter readings into accounting systems and automatically generate tenant bills. They are also being used to automate commissioning tests, calculate and compare energy use by similar facilities, and to create “virtual thermostats” that give users control over their office environments. Test programs are integrating building automation systems with utility systems, implementing control options based upon real-time utility pricing and implementing energy curtailment during emergencies. Universities and other large complexes are experimenting with using Web services to create interactive web pages, integrating utility consumption, maintenance management, cost accounting, record drawings, and other facility systems into a “facilities portal,” a single user interface that can be used to access all of these systems. See Figure 1-1 below. Projects under consideration include using weather forecasts to optimize ice storage systems, boiler start-ups, and morning pre-cooling. Universities are exploring the possibility of using their central classroom scheduling computer to automatically schedule HVAC, lighting, and other classroom services.

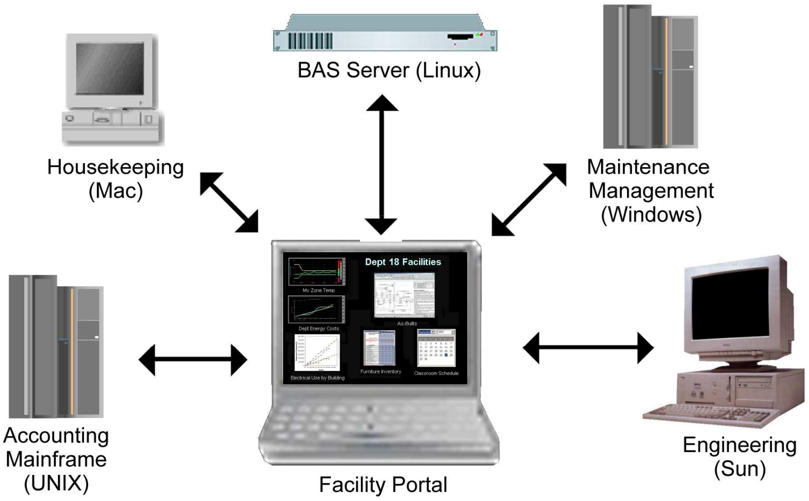

Figure 1-1- Integrating information from multiple systems into a Facility Portal.

If BAS vendors are already providing Web services, where does ASHRAE fit in? ASHRAE is establishing a standard means of using Web services to integrate facility data from multiple sources. The IT world established standards for the mechanism of Web services, but these standards say nothing about the actual data being exchanged. (This is analogous to the way the telecommunications industry establishes standards for telephone systems but does not specify what languages or conversations the system can carry.) Vendors can claim support for Web services while making as little or as much data available as they wish. They can also use whatever data structure they please and can make it very easy or very difficult to locate data in their system.

Even if every vendor tried to create useful Web services interfaces to their system, chances are no two interfaces would be alike and connecting two dissimilar systems would require hours and hours of custom programming. Some of the more farsighted members of our industry foresaw these problems, and three years ago they used the ASHRAE website to call for a standard information model. ASHRAE answered the call, and began gathering input from facility engineers, equipment manufacturers, government agencies, and universities. They used this information to develop a standard for using Web services in building automation systems. This standard covers the types of data to be exchanged, the path used to locate the data, and attributes of commonly used data objects such as analog inputs or binary outputs. The services required to read or write values are defined, as well as services needed to obtain information about the available data or to return error messages if a service fails. The standard covers arrays as well as scalar data, making it particularly useful for handling trend logs.

Because this standard is designed for use with Building Automation Systems, it was developed by the technical committee that is in charge of standards for Building Automation Control networks, i.e. the BACnet committee. Once approved, it will become an addendum to the BACnet standard, which means it will also become an ANSI, CE, and ISO standard. Naturally the standard is compatible with the BACnet protocol, but it is not limited to BACnet. Indeed, one of its most useful applications may be to serve as a standard for exchanging data between building automation systems using different protocols. Web services could be an ideal way to make a “top end” connection between systems running BACnet, LonWorks, MODBUS, or any proprietary protocol. Engineers would not have to learn the details of each individual protocol to program the connections, they would only have to understand the Web services standard. A Web services connection would also avoid the problems with incompatible baud rates, wire types, proprietary communication chips, and all the other issues that can come into play when a gateway is used to connect dissimilar protocols. (See Figure 2)

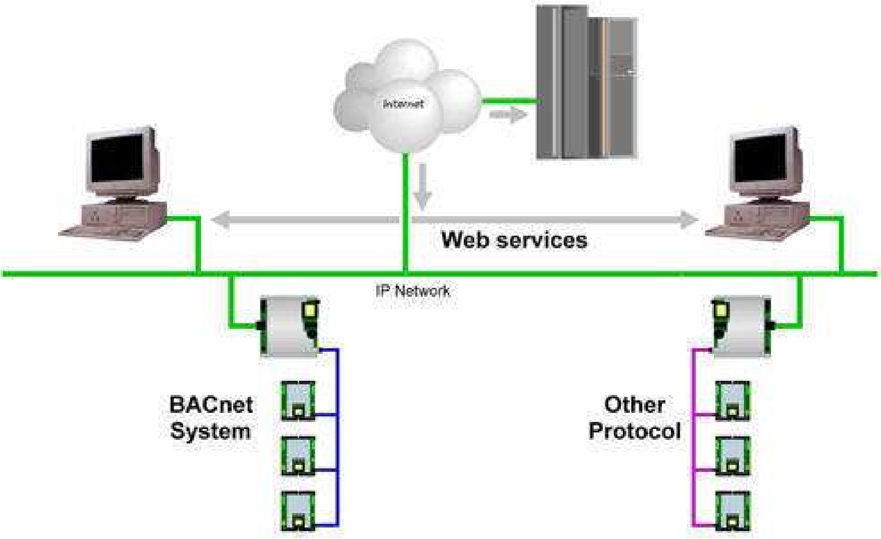

Figure 1-2 - Web services used to integrate BAS running dissimilar protocols, and to connect to a mainframe computer over the Internet.

Since Web services have quickly become the standard for B2B communications, it’s only natural to wonder if they will then replace BACnet, LonWorks, and other protocols within the BAS. That’s not likely, for several reasons. To begin with, no one has developed a set of Web services that covers all the functions needed by a BAS. Broadcasts, alarms, time synchronization, backup and restore – there are a host of BAS functions that simply are not covered in the proposed Web service standard. Certainly such a standard could be developed, but it would in essence become one more BAS protocol fighting for acceptance in the marketplace. It would not be a protocol that was well suited for a BAS because Web services require more “overhead” than most BAS controllers can provide. By definition, Web services use XML to communicate over an IP network.

Similarly, XML is a very “verbose” way to package data. It’s designed to be human understandable, flexible, and self-documented. These characteristics also mean it needs to be processed by a powerful computer and transmitted over a high-speed network. This is beyond the capabilities of the price-sensitive controllers typically used for small HVAC equipment like VAV boxes. This may be a temporary limitation, as inexpensive microprocessors gain power and speed with each passing year, but since existing protocols like BACnet are already developed, are a more efficient way of integrating controllers, and are open for use by any equipment manufacturer there is very little incentive to switch these controllers to Web services.

When you try to integrate with systems outside the BAS, such as the local utility system, the situation changes dramatically. To begin with, the systems you are trying to integrate with are not using BACnet, LonWorks, or any other building protocol, and the people who manage these systems have no interest in providing a special connection for a BAS. They would much rather provide a general-purpose interface that can be used by any external computer system. Their system is already running on a high-end computer connected to a high-speed IP network, which is exactly the situation Web services were created for. The computers and the networks have the “horsepower” to handle Web services. There will probably be certain amount of custom linking, if not custom programming, required to make the connections, but the self-documenting characteristics of XML simplify the programmer’s task. Chances are the programmer is already familiar with Web services from previous B2B integrations, which further simplifies the job. (A customer in Texas who was contracting for a custom interface between their BAS and a billing system found the contractor cut their price in half when they learned the BAS supported Web services.) The addition of a new ASHRAE standard to the Web services world promises even greater simplification, using IT technology and the foundation of BACnet to take building automation to the next level.

References:

1. BSR/ASHRAE Addendum c to ANSI/ASHRAE Standard 135-2004 Public Review Draft, American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE), www.ASHRAE.org

2. Information Model: The Key to Integration. Craton, Eric and Robin, Dave, AutomatedBuildings.com, Jan 02

There are many other references available in the market for understanding the industry background and context for this DIG. More information on this topic is available on these websites:

BACNET Website: http://www.bacnet.org/

OBIX Website: http://www.obix.org/

CABA Website: http://www.caba.org/index.html

Description and Justification for BAS over IP ("BAS/IP")

An excellent description of a BAS/IP system is provided in a report prepared by Frost & Sullivan entitled "*Impact Analysis of IP Protocols on Building Automation*"[3]. This specific description pertains to the Johnson Controls Metasys product, the system described below in subsequent chapters of this DIG.

"Networking technologies have come a long way with numerous enhancements and standards over the past two decades, and have become a part of our socioeconomic well-being. The present day building owners look for a comprehensive solution that would enable them to live with the already installed legacy systems, in a more sophisticated way by converging with the omnipresent information technology systems. Johnson Controls Inc., based out of Milwaukee, Wisconsin, seamlessly enables the integration of building automation systems with the information technology systems, thanks to its XML-based technology. With a higher degree of stress on enabling mobility, Johnson Control's Metasys building management system accounts for flexibility and scalability.

Built around Web-based technologies, IPs and standards, building infrastructure managers can use their Web browser-enabled PDAs, PCs or laptops to access, monitor, and control building assets. Periodic alerts and critical information about various events can be automatically delivered to mobile phones, tablet PCs, and the like, with ease of navigation, coordination, and control over the information received. Old buildings that have installed legacy building automation controls from Johnson Controls or other vendors, can seamlessly embrace wireless and Web-based technologies using the Metasys building management system, saving the building owners from unnecessary capital expenses. Metasys enables integration of open building automation protocols such as BACnet over IP, BACnet over MS/TP, N2 and LonWorks' LonTalk, facilitating a truly heterogeneous network of building systems.

With IPs as the network communications medium, services such as XML, SOAP, SNMP, and dynamic host configuration protocol (DHCP) are facilitated by Metasys servers and network control engines (NCEs) via standard Web browsers. As mentioned before, since Metasys insists on mobility and being untethered (wireless technology), expenditures incurred due to extensive cabling of large buildings are considerably contained. Metasys' wireless features are scalable and flexible, and the wireless technology can be put to the best use depending on the application requirements. "As we move forward we see that the IP-based communication is required to enable other functionalities such as the wireless, and the benefits of the wireless in terms of mobility, flexibility far outweigh the extra cost involved in using IP communications," says Terry Hoffmann, director of marketing, Building Management Systems, Johnson Controls Inc.

Apart from providing encoding protocols and enabling security, structured query language (SQL) database is supported to facilitate data storage and retrieval. IP network connectivity, Metasys software user interface and network supervisory capabilities are the features of the Metasys' network control engine (NCE), which enables direct digital control Impact Analysis of IP Protocols on Building Automation capabilities of its field equipment controllers. Specifically designed for integrating central plants and large air handlers, NCE series controllers are a lucrative solution. Buildings and enterprises with already installed IT and IP network infrastructure can seamlessly integrate with the NCE, facilitating communications over the Intranet, Internet protected by firewall, wide area network (WAN), and the like.

Without the need for separate software, the NCE can be accessed, monitored, and controlled via a standard Web browser and a network connection. Another important aspect is that as the system can be connected over the Internet, users can access the NCE via digital subscriber line DSL/cable or a normal dial-up connection, providing the flexibility of remote management. One can configure, archive data, monitor, manage, and control through the Web browser, from anywhere in the world."

Commercial real estate owner/developers are finding that convergence provides other opportunities to use the integrated communications network as well. Audio-visual, IP television, enhanced cellular coverage, wireless point of sale, building management systems, security access and surveillance, help point, and car parking systems are just a few of the deployments possible over a converged network.

[3] *IMPACT ANALYSIS OF IP PROTOCOLS ON BUILDING AUTOMATION – Frost & Sullivan – Report DA09* - © 2007 Frost & Sullivan

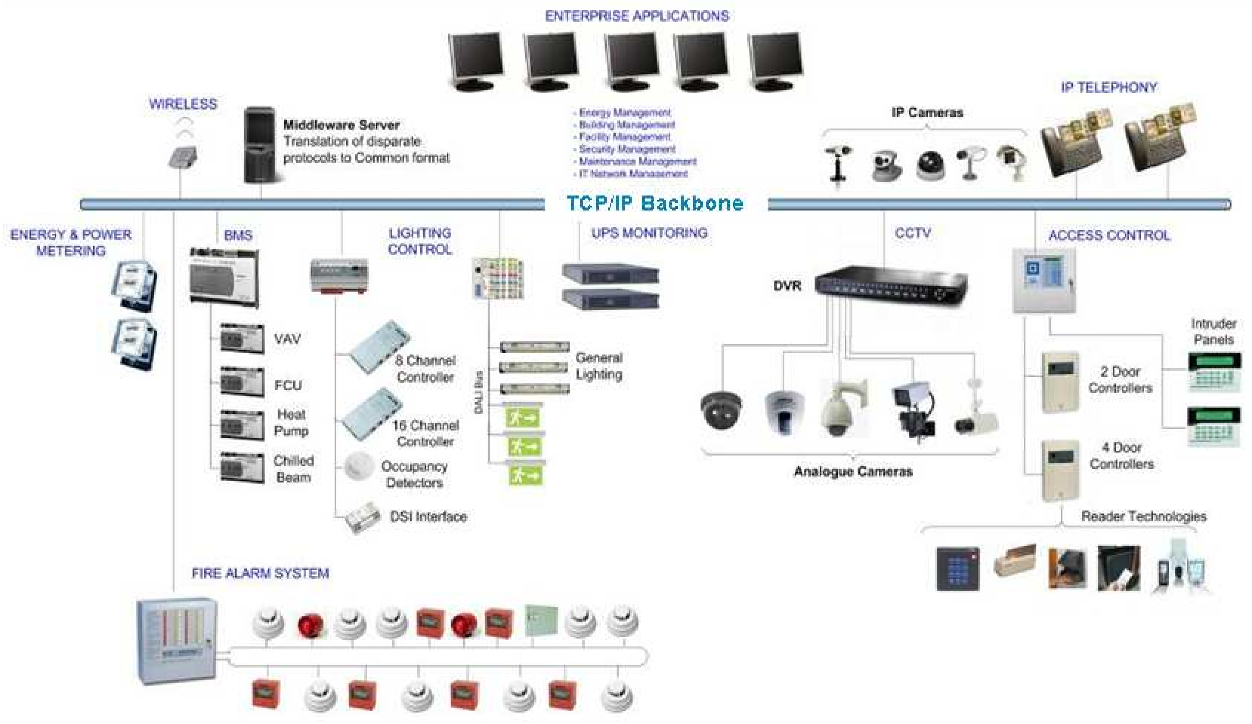

Figure 1-3 **– Typical Integrated Communications Network Diagram**

Some of the key definitions related to an Integrated Communications Network specification include:

Communications network
A network used to connect addressable field control devices such as outstations and unitary controllers.

Convergence
The integration of data, voice, and video solutions onto a single Internet Protocol (IP) based network.

Building Management System (BMS)
A BMS has at least one central database server to either include or support at least one permanent operator workstation connected directly or via a communications network to integrate a number of building services electronic systems into a common user interface.

Internet Protocol Closed Circuit TeleVision (IP CCTV)
Closed circuit television using the Ethernet communication network.

Local Area Network (LAN)
The physical communication network cabled throughout the site.

Virtual Local Area Network (VLAN)
A virtual communication network connecting devices through a virtual private network.

Virtual Private Network (VPN)
The virtual segmentation of particular service network traffic in the communication network within and off site.

Voice over Internet Protocol (VoIP)
Voice communication across the Ethernet network.

Wireless Fidelity (WiFi)
The wireless Ethernet communication network.

Domain Name System (DNS)
A system for converting host names and domain names into IP addresses on the Internet or on local networks that use the TCP/IP protocol.

Dynamic Host Configuration Protocol (DHCP)
Software that automatically assigns temporary IP addresses to client stations logging onto an IP network. It eliminates having to manually assign permanent "static" IP addresses. DHCP software runs in servers and routers.

Service Oriented Network Architecture (SONA)
A Service Oriented Network Architecture SONA is the framework for enterprises to connect network services to applications delivering business solutions.

Simple Network Management Protocol (SNMP)
A widely used network monitoring and control protocol.

Simple Mail Transfer Protocol (SMTP)
The standard e-mail protocol on the Internet and part of the TCP/IP protocol suite, as defined by IETF RFC 2821.

The business case for justifying a BAS/IP system is becoming increasingly less difficult to justify.

Trends in the marketplace show us that this traditional industry is changing dramatically.

Research done by i&I from the UK shows that in the next 12 months, more IP enabled building controls devices will be produced than the typical proprietary systems. The way is paved for us to play in this market, and Cisco intends to further drive and accelerate this trend forward. See Figure 1-4.

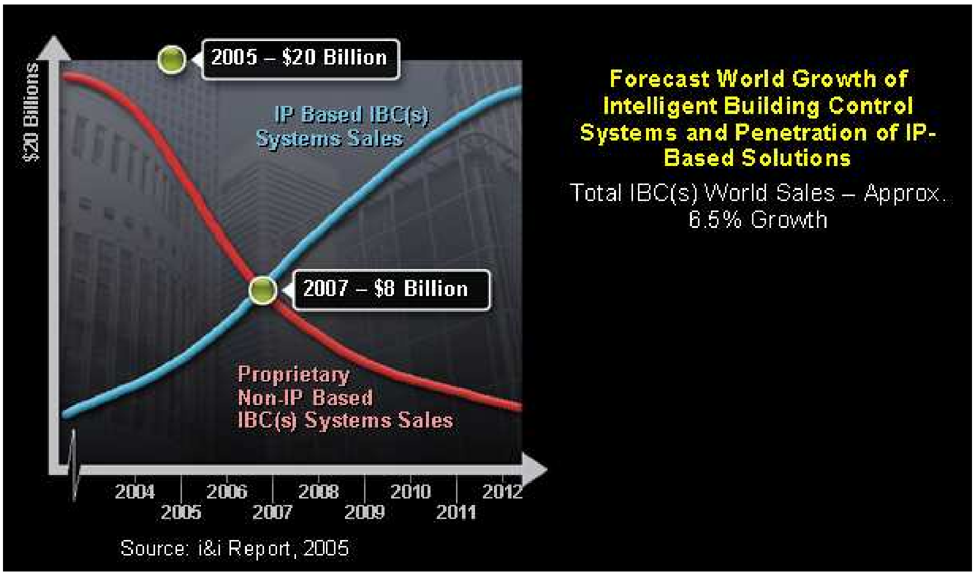

Source: i&i Report, 2005

Figure 1-4 - Industry Convergence Toward IP

Target Market Opportunities

In a report prepared by the ARC Advisory Group in 2005, forecasting through 2009, the outlook for growth of BAS over IP has never looked stronger.[4]

An excerpt from this 134 page report analyzing the market opportunity is provided below.

"..As the Building Automation Systems (BAS) market continues to redefine itself, suppliers are being forced into a period of transition. BAS hardware is becoming commoditized to the point that value-added software and information management solutions are now the focus of attention. Whereas many of the leading BAS suppliers' traditional focus was on designing hardware solutions to control HVAC equipment, they are now increasingly being asked to provide integrated solutions capable of not only controlling, but optimizing, all aspects of building automation including HVAC, lighting control, security & access control, and fire alarm & safety.

The adoption of Internet communication standards and Web Services in the BAS market is further extending the concept of smart buildings by including intelligent analysis of all building data. In sharp contrast to traditional BAS solutions, the new requirements for BAS solutions include providing facilities managers the tools to perform the same sophisticated business intelligence analysis typically reserved for business applications. Recognizing the emerging need for increasing business intelligence, more BAS suppliers are focusing on providing these capabilities.

The goal is to develop intelligent BAS solutions capable of providing facilities managers the ability to base operational decisions on real-time performance data and finally uncover hidden costs, and opportunities to save money, though comprehensive facilities management. Increasing demand from facilities managers, and the need for BAS suppliers to broaden their core strengths, is creating an environment rife for consolidation. Leading BAS suppliers are on a ruthless market consolidation drive, while tier two BAS suppliers are focusing on protecting their niche markets. As competition reaches new heights, many suppliers are seeking cross domain expertise to remain competitive. Mergers and acquisitions as a strategy for growth and expansion are in fashion in the current environment.

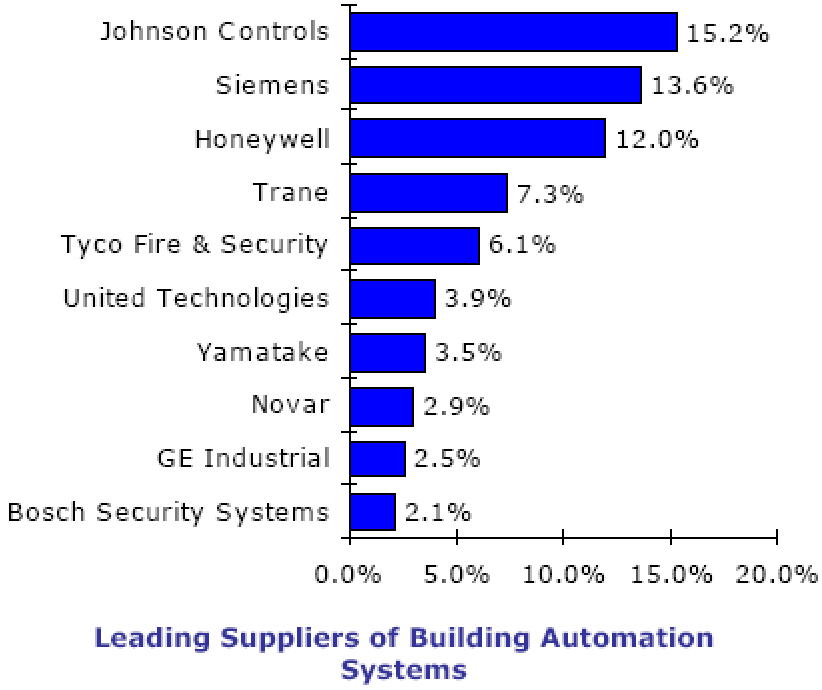

Figure 1-5 - Leading Suppliers of Building Automation Systems

[4] Building Automation Systems Worldwide Outlook - Market Analysis and Forecast through 2009, Copyright © 2005 ARC Advisory Group

Since the 2002 edition of this report, Johnson Controls has managed to overtake Siemens as the leading supplier of BAS worldwide. Two key factors contributing to Johnson Controls' success in the BAS market include the company's strong service organization and its ability to provide a complete solution for all building automation needs. Developing a strong local footprint to service customers around the globe has been a strong focus for Johnson Controls in recent years. Having created a global force of nearly 12,000 technicians, mechanics, and general maintenance personnel, Johnson Controls' can boast of having the leading BAS services organization worldwide. Johnson Controls' strong local presence and service capabilities are key factors in the company's recent success, as more facilities managers focus on choosing a BAS supplier capable of supporting them on a global basis. Although Siemens has dropped behind Johnson Controls for the top spot in the global BAS market, the company remains a strong player, maintaining its leadership position in many subcategories including Fire & Safety, Healthcare, and EMEA. Siemens Building Technologies (SBT) is capable of providing comprehensive BAS solutions including HVAC, energy management, fire detection, access control, video surveillance, and alarm systems. SBT is also capable of providing BAS solutions tailored to numerous vertical markets including Life Sciences, sports stadiums, hospitals, high-tech corporations, and hotels. SBT believes its in-house expertise of the business processes, unique to the various vertical markets, is critical to the company's success in the BAS market."

The worldwide BAS market continues to grow at a steady rate as corporations in developed countries continue investing in BAS solutions to help strategically manage existing building assets, and companies in developing countries continue constructing new state-of-the-art commercial and industrial buildings. For companies in developed regions, strategic management of existing building assets is one of the best ways to increase productivity, with little to virtually no operational upsets, while meeting the company's goals and objectives. On the flip side, corporations in Asia continue expanding operations to meet growing domestic demand, which require new construction projects, many of which are incorporating state-of-the-art BAS solutions."

To translate this opportunity into the Total Addressable Market ("TAM") for Cisco, we need to understand how the market opportunity for BAS relates to Cisco hardware and software.

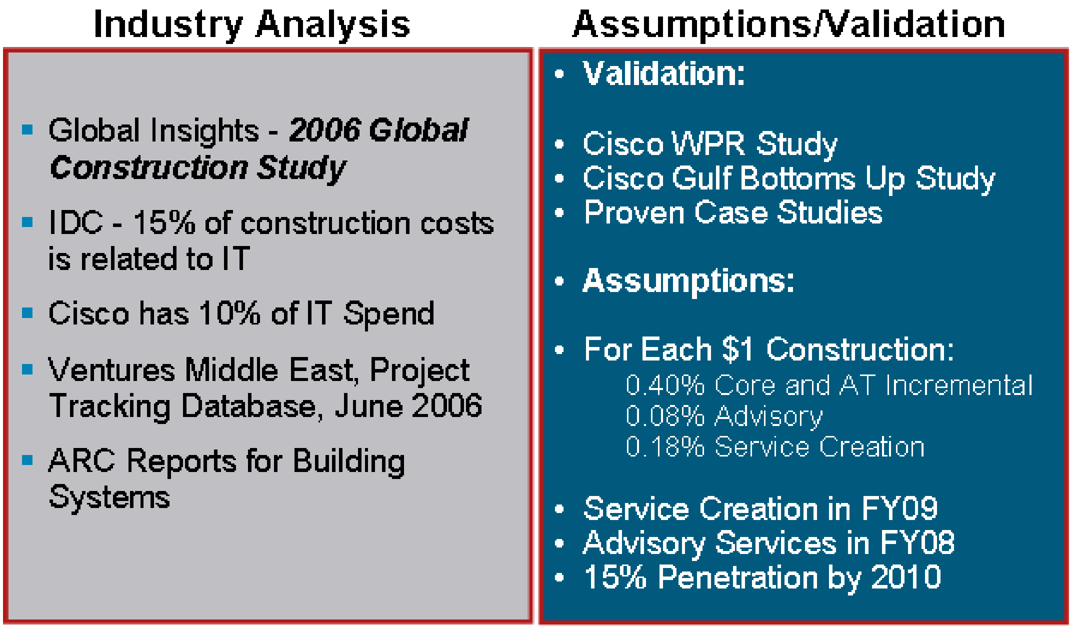

Industry Analysis

- Global Insights - ***2006 Global Construction Study***
- IDC - 15% of construction costs is related to IT
- Cisco has 10% of IT Spend
- Ventures Middle East, Project Tracking Database, June 2006
- ARC Reports for Building Systems

Assumptions/Validation

- **Validation:**
- Cisco WPR Study
- Cisco Gulf Bottoms Up Study
- Proven Case Studies
- **Assumptions:**
- For Each $1 Construction:
  0.40% Core and AT Incremental
  0.08% Advisory
  0.18% Service Creation
- Service Creation in FY09
- Advisory Services in FY08
- 15% Penetration by 2010

Figure 1-6 – Market Sizing – Industry Analysis

Even the most conservative projections indicate a sizable market opportunity for Cisco and Johnson Controls working together.

Applications and Services Supported by BAS/IP

The implications for applications in a BAS/IP environment are numerous. On the surface, applications in the Heating, Ventilation, and Air Conditioning (HVAC), Energy Management, Lighting, Access Control, Fire Alarm, and Environmental Monitoring are a few opportunity areas. Imagine if a fire occurs in a tenant's space. The fire system could immediately signal the air flow system to close the dampers, immediately restricting air flow. At the same time, it could signal the building access system to release all door locks. Elevators could be instructed to return to the nearest floor, open and cease operating, and video cameras could be instructed to begin recording at specific locations. At the same time, IP phone calls could be automatically generated to the fire department as well as to the tenants, faculty, students, and staff. This could all occur within seconds, helping to save lives and limit property destruction.

Optimized Energy Consumption is another BAS/IP application getting a lot of attention due to increased market rates for power throughout the world. With nearly 50% of all energy consumption occurring through the operation of buildings, BAS/IP is a mechanism to provide intelligent monitoring and management with energy utilization trending. The ability for a building to monitor and self-regulate energy consumption has enormous potential. As BAS/IP is integrated with access control and lighting through the use of sensors, controllers, and IP enabled devices, it is now possible to 'know' when systems are in demand and turn them off or reduce their usage when there is no reason to have them operational. This behavior approach to resource utilization optimizes energy consumption and results in significant operating expense (OPEX) reductions.

BAS/IP can provide personalized comfort control in office or hotel spaces via IP phone touch screens for convenience and savings. By monitoring sensors and providing personal control over small spaces, end-users can customize their environment to suit their needs providing convenience and productivity enhancements. The BAS/IP controls are made available over the same IP network that computers, telephones, a video devices use saving CAPEx and providing OPEX benefits.

Solution Benefits

BAS/IP provides us the opportunity to dramatically lower costs, improve services, and drive productivity increases on almost every level. BAS/IP solution benefits include the ability to reduce both CAPEX and OPEX costs. CAPEX reductions include construction of few mechanical and electrical installations, such as lighting, cooling, heating, fire alarm, telephone, and/or cable. Some buildings can have up to 15 separate systems. Accounting for a sizable part of this cost is that each system requires its own proprietary, separate network of wires and cabling combined with proprietary protocols for control and communications. OPEX reductions include better engineering staff utilization (higher value work by eliminating need to physically monitor and maintain separate systems), detailed monitoring and reporting of utility usage, optimized energy resource utilization and potential to negotiate energy rates based on usage trends. Many 'smart buildings' that deploy BAS/IP solutions report 30-60% reductions in operating costs associated with this type of integrated system. For large, multi-building real estate developments, this can add up to hundreds of thousands of dollars in annual savings.

Detailed Energy Cost Tracking – BAS/IP allows sensors and controllers to relay usage information in real time and have that information saved for future reporting and trending. Facility managers have the ability to understand their usage patterns, use this data to negotiate service levels and rate agreements with their utility providers and proactively manage adverse events when they occurred. In one situation, the operations manager was able to produce reports showing excessive power surge spikes, which resulted in failed equipment, and was able to pursue compensation from his utility provider for causing the failures. In another situation, consumption data was reconciled against utility provider usage billing which resulted in $600,000 in reduced charges. Having access to granular information provides the ability to understand and act on more accurate data allowing better business decisions to be made.

Improved Productivity For Engineering Staff – By managing by exception as opposed to actively managing all systems simultaneously, building engineers can have access to conditions and status at any time, receive alerts to exception situations, and modify system behavior as needed before users notice and call in support requests.

Sustainability and Green Benefits – Increases in green house gas emissions and global warming is seen as a potential threat to the ecological system of our planet. Governments and organizations are constantly under pressure to issue norms that could curb the energy waste and decrease the leakage of green house gases in to the atmosphere. By enabling BAS/IP systems to monitor and report equipment power status, sending alerts when equipment has been left on, the BAS/IP system can power down unused machines. Upgrading and installing Variable Frequency Drives (VFDs) on condenser and chilled water pumps provides not only significant Kilowatt-hour (Kwh) savings but a reduced carbon footprint for CO2 emissions. BAS/IP reporting can also provide a usage basis for expanding a companies Green initiatives into Green Power Purchase agreements. More companies are investing in Green Purchase Agreements to meet their social, environmental, and sustainability goals.

Health & Safety - Central to BAS/IP applications are monitoring and reporting from many types of sensors including carbon dioxide and other gases, temperature (in case of fires), and humidity (in case of flooding) to assure health and safety. BAS/IP systems provide the capture and routing mechanisms for alerting to other IP-based systems. Imagine your office phone with a broadcast alert of a fire alarm in an adjacent building or CO2 alert in the break room. With the convergence of IP-based devices, real time alerting of health and safety communications is now possible.

Solution Features

BAS/IP allows customers to effectively exploit the existing infrastructure to integrate building systems into that network, enabling remote access, management and distributed control. BAS/IP integration allows building owners to minimize life/safety situations in a building. In addition,
Owners and facility operators have a built in incentive to implement BAS/IP to enable insurance cost reductions due better management of life/safety systems.

The below list summarizes many of the expected features of a modern BAS/IP solution.

- UL-864 program certified to govern fire and smoke operations in commercial buildings
- UUKL program certified - The smoke abatement certification, UUKL, is an adjunct function of the fire system that automatically or manually purges the fire and directs smoke safely out the building by exhausting smoke from affected areas while simultaneously shutting down adjacent dampers and therefore the oxygen supplies.
- Sufficient integration to allow automatic evacuation systemic operation to be activated as part of the Fire/Smoke Control application or may be activated for other reasons such as terrorist threats
- Integrated HVAC systems which determine the earliest possible time it can shut down heating/cooling yet still control the set points to meet the requisite parameters
- Integrated Lighting which determines when lights can be extinguished as soon as they are no longer needed
- IP Telephone interface that allows occupancy sensors in meeting rooms and any late workers to override the normal HVAC and lighting schedules simply by dialing into the system and specifying their locale
- Energy integration to modulate or shutoff large equipment temporarily without affecting environment comfort. FMS system will constantly monitor real-time energy usage and automatically turn unneeded equipment off (or reduce the control set point)

- Access, monitoring, and control via standard web browser for Network Control Engine (NCE) devices. Ability to configure, archive data, monitor, manage, and control through the Web browser, from anywhere in the world
- Ability for building managers to use their Web browser-enabled PDAs, PCs or laptops to access, monitor, and control building assets
- Ability for Periodic alerts and critical information about various events can be automatically delivered to mobile phones, tablet PCs, and the like, with ease of navigation, coordination, and control over the information received
- System enabled integration of open building automation protocols such as BACnet over IP, BACnet over MS/TP, N2 and LonWorks' LonTalk, facilitating a truly heterogeneous network of building systems
- Use of advanced web services such as XML, SOAP, SNMP, and dynamic host configuration protocol (DHCP) to allow integration to servers and network control engines (NCEs) via standard Web browsers.
- Use of Web services used to automate commissioning tests, calculate and compare energy use by similar facilities, and to create "virtual thermostats" that give users control over their office environments.
- Reliance on mobility and being untethered (wireless technology) so expenditures incurred due to extensive cabling of large buildings are considerably contained. Wireless features are scalable and flexible and should be leveraged appropriately.
- Use of Web services to create interactive web pages, integrating utility consumption, maintenance management, cost accounting, record drawings, and other facility systems into a "facilities portal," a single user interface that can be used to access all of these systems.
- A network control engine with intuitive software user interface with network supervisory capabilities to enable direct digital control Impact Analysis on Building Automation capabilities of field equipment controllers.
- The BAS/IP is specifically designed for integrating central plants and large air handlersAn enterprise campus architecture incorporating a Network security design based on the Cisco Self-Defending Network, an IP-based communications schema, Mobility and wireless LAN services
- BAS/IP security to prevent packet sniffing, IP spoofing, Distributed Denial of service, Network Reconnaissance, unauthorized access, virus and Trojan horse applications and password attacks
- Ability to assign and track multiple levels of access for various types of users
- Support for multiple media types including Ethernet (802.3 and IP), EIA-485, Arcnet, LON and RS-232 and ZigBee wireless mesh
- Sensors, actuators, area controllers, zone controllers, and building controllers all utilize the BACnet protocol. The BACnet (Building Automation Control Network) is an ISO world-wide Standard protocol designed to maximize interoperability across many products, systems and vendors in commercial buildings.
- Compliance with *MasterFormat*™, the leading standard for organizing nonresidential construction specifications, for numbers and subject titles for organizing information about construction work results, requirements, products, and activities into a standard sequence
- Compliance with ISO 16484-2:2004 specifies the requirements for the hardware to perform the tasks within a building automation and control system (BACS) which provides the terms, definitions and abbreviations for the understanding devices for management functions, operator stations and other human system interface devices.
- Compliance with ASHRAE, a standard means of using Web services to integrate facility data from multiple sources using XML to communicate over an IP network.
- Compliance with Construction Specifications Institute (CSI), and sister organization Construction Specifications Canada (CSC)

Chapter 2 Solution Architecture

Overview

This chapter provides an overview of the Facility Management Networking (FMN) 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. FMN is an architecture that provides network and security services to the devices, equipment, and applications found in Facility Management Systems (FMS) as integrated into the wider enterprise network. The networking requirements of a real-time mission critical facility management system often differ from a typical IT network. This solution architecture overview provides the background and description of a facility management network model and highlights the differences between the FMN architecture and the IT network.

Reuse is an objective of any architecture, which is the case with the FMN solution architecture. Facility Management Systems are deployed in a large variety of commercial facilities, such as universities; hospitals; government facilities; K-12; pharmaceutical manufacturing facilities; and single-tenant or multi-tenant office buildings. Facility Management systems are deployed in a wide variety of commercial building topologies, including single buildings, multi-building single site environments such as university campuses and widely dispersed multi-building multi-site environments such as franchise operations. These buildings range in size from 100K sqft structures (5 story office buildings), to 1M sqft skyscrapers (110 story Shanghai World Financial Center) to complex government facilities (Pentagon). The FMN architecture is meant to be the model to be used in all these types of environments, but clearly it must be tailored to the building class, building tenant and vertical market being served.

The following sections describe the FMS system from the lowest layer to the highest layers in the hierarchy. Each section describes the basic functionality of the layer, its networking model, power requirements and a brief description of the communication requirements. The entire section references the block diagram noted in Figure 2.1a. This figure notes that there are 5 major subsystems comprised in an FMS. These subsystems all have layered solutions starting at the sensor layer and moving upward in complexity to the enterprise. While these five subsystems are common to most facilities, they are by no means the exhaustive list - a chemical facility may require a complete fume hood management system; a manufacturing facility may require interfacing to the PLC subsystem; or a multi-tenant facility might require a comprehensive power management subsystem. The objective in the overall design of the JCI Metasys system is to integrate all common functions into the system yet allow maximum flexibility to modify these systems and add other systems as dictated by the job by the JCI field engineers.

Facility Management Reference Model

FMS Topological Introduction

To understand the IT security and network systems requirements of a facility management system in a commercial building, this guide uses a framework to describe the basic functions and composition of the system. A FMS is a horizontally layered system of sensors and controllers. Additionally, an FMS may also be divided vertically across alike, but different building subsystems as noted in Figure 2-1a.

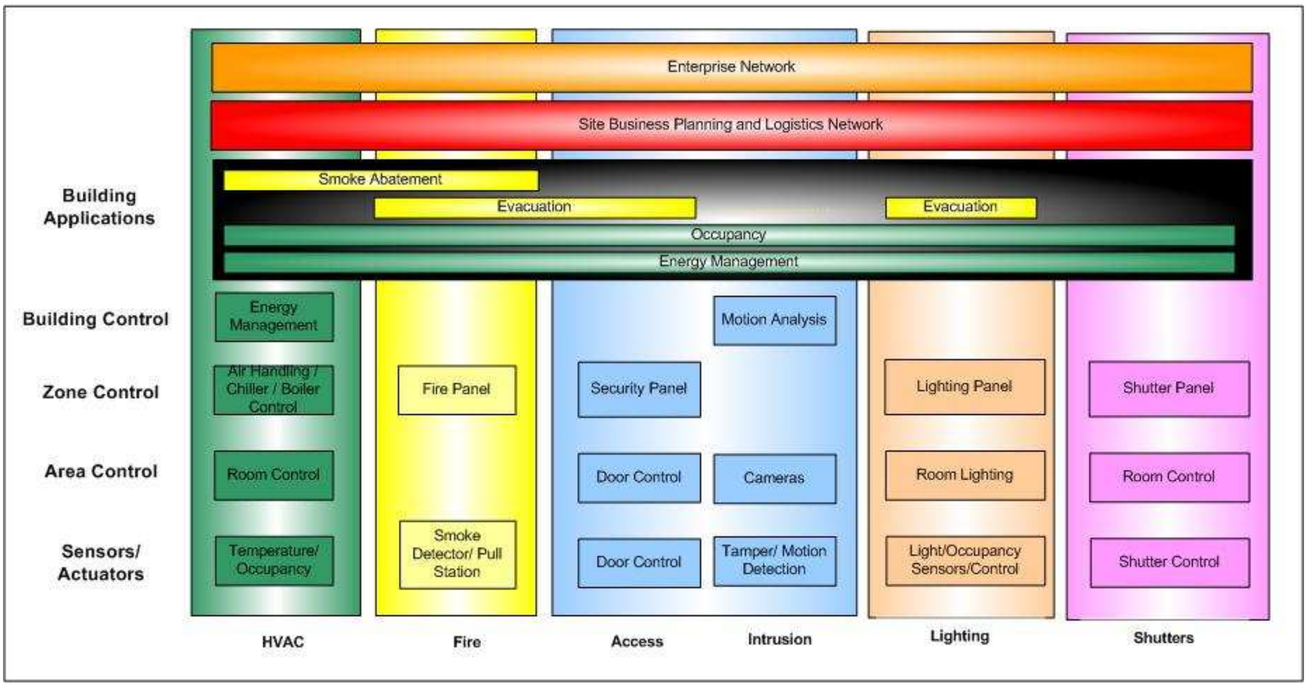

Figure 2-1a FMS Functional Domains

Much of the makeup of a FMS is optional, other than the sensors and actuators layers, all upper layers have standalone functionality. These devices can optionally be tethered together to form a more synergistically robust system. The customer can decide how much of this vertical 'silo' should be integrated to perform the needed application within the facility. This approach also provides excellent fault tolerance since each node is designed to operate in an independent mode if the higher layers are unavailable.

As depicted in Figure 2.1a, Heating, Ventilation and Air Conditioning (HVAC); Fire; Security and Lighting are components that can be tethered together into a cohesive set of all encompassing applications tailored to the customer's whim. Shutter control is an emerging application domain prevalent in the European market. These major subsystems are connected logically through application software called Building Applications. This horizontal stack follows the vertical stack design in that each silo is optional. The customer can integrate all the subsystems at once or add them as the facility or budgeting dictates.

Communication Media

The FMS is tied together via three network technologies; EIA-485, Ethernet and ZigBee.

The sensors, actuators, area controllers, zone controllers, and building controllers are connected via EIA-485 3-wire twisted pair serial media operating nominally at 38400 to 76800 baud. This allows runs to 5000 ft without a repeater. With the maximum of three repeaters, a single communication trunk can serpentine 15000 ft. Figure 2.1b defines the devices and protocols of the FMN wired network.

The HVAC, Fire, Access, Intrusion and Lighting subsystems are integrated using LAN based Ethernet technology. These enterprise devices connect to standard Cat-5e through workgroup switches. WLAN communications can replace the Ethernet connection if the application can operate within the WLAN performance characteristics. Currently all building controllers support only a RJ-45 connection. WLAN connections require an external wireless bridge. Multi-building sites can also connect onto the facility intranet if the intranet performance matches the application requirements.

The sensors, area controllers and zone controllers can optionally integrate onto the HVAC silo via ZigBee mesh at 2.4gHz. See Figure 2.2c. These devices can be a mixed wired and wireless set as required by the application parameters. ZigBee technology may also be used on other silos as the technology matures.

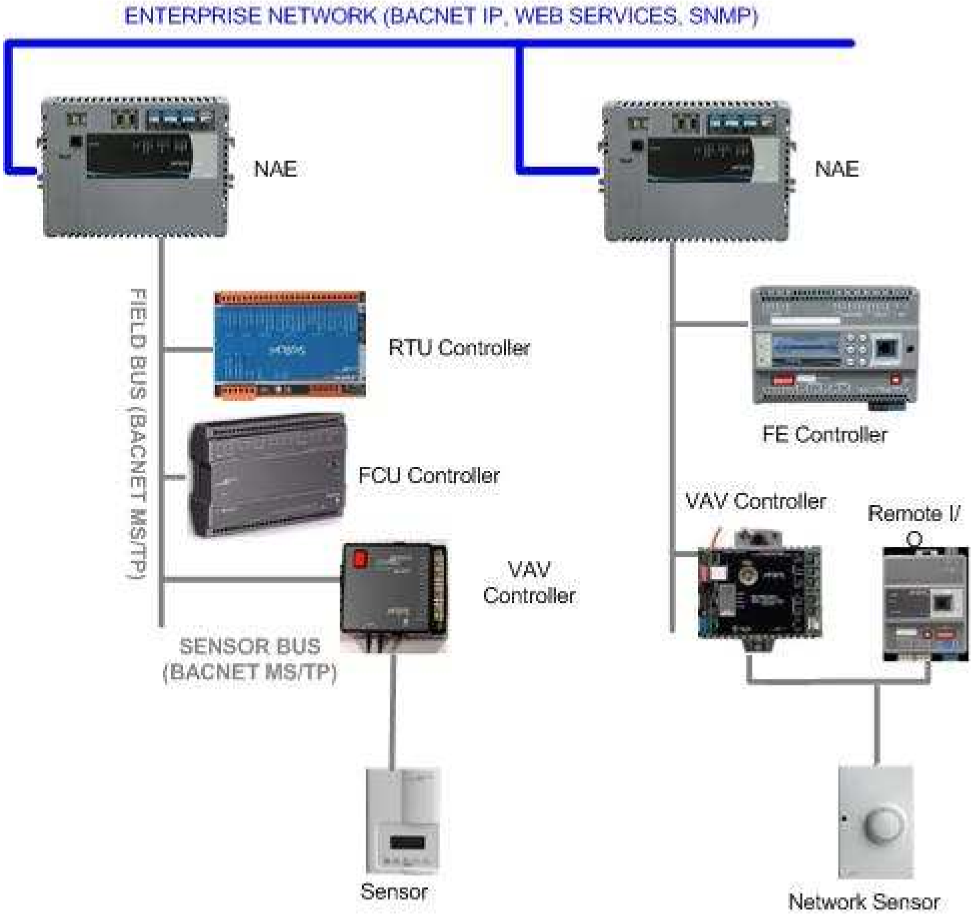

Figure 2-1b Network Names and Wired Protocols

Controller/Sensor/Actuator Communication Protocol

The sensors, actuators, area controllers, zone controllers, and building controllers all utilize the BACnet protocol. BACnet (Building Automation Control Network) is an ISO world-wide Standard protocol designed to maximize interoperability across many products, systems and vendors in commercial buildings. BACnet was conceived in 1987 and released in 1995 for the HVAC industry. Since that time Fire, Security and Lighting functionality has been added.

BACnet supports five media types including Ethernet (802.3 and IP), EIA-485, Arcnet, LON and RS-232. BACnet soon will add ZigBee wireless mesh to its media types supported.

BACnet MS/TP is merely an alternate BACnet data link for EIA-485 networks. MS/TP is a token passing protocol (implemented in software) allowing master/slave and peer-to-peer communication simultaneously. Devices must designate themselves as slaves or masters on the network. Slave devices may only access the network when solicited by a master device. Masters may communicate to any node on the network whenever it holds the token.

BACnet supports all expected network services including functions such as device and object discovery; unicast and broadcast messaging; full routing; flow control and fragmentation, and security policies.

BACnet addressing differs depending on the data link implemented. BACnet/IP currently supports IPv4 addressing. IPv6 is in discussion within the committee. BACnet MS/TP has a 1-octet MAC address allowing for a maximum of 254 devices per network segment. (Address 255 is reserved for broadcast designation). Table 2.1a describes the network parameters in tabular form.

Network Name	Media Type	Communication Rate	Protocols Supported	MAC Addressability	JCI Devices Supported
Sensor Bus	EIA-485	9.6- 76.8 kbps	BACnet MS/TP	8-bit	1- 16
Field Bus	EIA-485	38.4 – 76.8 kbps	BACnet MS/TP	8-bit	1 - 100
Enterprise Network	Cat-5e	10/100 mbps	BACnet IP Web Services SNMP	IPv4	thousands

Table 2.1a Network Parameters

Enterprise Communication Protocol

Multiple protocols are supported at the enterprise level of the FMS since this layer supports not only the embedded control operation but also the user interface and end-user enterprise applications.

Peer-to-peer Controller Communication

Building Controllers, often termed Supervisory Controllers, orchestrate the overall FMS system operations. Control and data access functions implemented at this level utilize BACnet IP. BACnet IP provides the complete building object model and requisite services across all the FMS silos. Since BACnet is deployed on the lower layers of the system, utilizing it to control operations at the highest layer of the system is prudent. BACnet IP implements UDP/IP with its own transport layer. It is designed to operate efficiently and transparently on all IP networks. It typically utilizes UDP port address xBAC0 (47808_{10})

Enterprise Communication

While BACnet is the control protocol of choice; it is out of scope for most enterprise applications. Web Services and SNMP has been added to the enterprise layer to assist in integration with end-user applications and Network Management Systems respectively. The enterprise level also supports most ancillary IT protocols such as SMTP, SNTP, DHCP and DNS.

Sensors/Actuators

As Figure 2a indicates an FMS may be composed of many functional silos that are interoperably woven together via Building Applications. Each silo has an array of sensors that monitor the environment and actuators that affect the environment as determined by the upper layers of the FMS topology,

The sensors typically are the leaves of the network tree structure providing environmental data into the system. The actuators are the sensors counterparts modifying the characteristics of the system based on the input sensor data and the applications deployed. Traditionally, sensors were wired devices deployed on proprietary networks. The proprietary nature of the protocols reduced interoperability options across silos.

In 1995, the BACnet protocol was released by ASHRAE that defined interoperable objects and services within the HVAC silo. BACnet has grown to be an international standard now including extensions for Fire, Access, Intrusion and Lighting functions.

Sensor and actuator performance is dictated by the class of device. Table 2.1b and 2.1c defines typical performance characteristics for various sensors and actuators respectively. Figures 2.1c and 2.1d summarize the basic functional characteristics of the sensors and actuators.

Sensor Type	Expected Response Time	Security Policy
Space Temperature	10 minutes	Heartbeat
Duct Temperature	1 minute	Heartbeat
Smoke Detection	10 seconds	Supervised
Occupancy	1 minute	Heartbeat
Door Access	1 second	Supervised
Static Pressure	100 milliseconds	Heartbeat

Table 2.1b Sensor Expected Performance Characteristics

Sensor Type	Expected Response Time	Performance Assurance
Air Flow Damper	30 seconds	Sensor Feedback
Evacuation	60 seconds	Supervised
Admittance	1 second	Supervised
Lighting	100 milliseconds	Optical Sensing
Smoke Control Dampers	10 seconds	Supervised
Smoke Abatement	60 seconds	Supervised

Table 2.1c Actuator Expected Performance Characteristics

Sensor Characteristics

- Setup
    - MAC Address set via onboard switches
- Operation
    - Associate to controller(s) of interest
    - Periodical (or upon event) sense the environment, encode the information and forward to the requesting controller(s)
- Reporting
    - Report to controller erroneous events such as unreliable sensor reading, obfuscation, or low battery
- User Interface
    - Display local information (optional)

Figure 2.1c Sensor Characteristics

Actuator Characteristics

- Setup
    - MAC Address set via onboard switches
- Operation
    - Execute command actions as received by controller
    - Prioritize command actions as needed to meet the customer requirements
    - Confirm actuation completed as directed.
    - Maintain requested setpoint via closed loop control
- Reporting
    - Report failed actuations to controller for further analysis
- User Interface
    - n/a

Figure 2.1d Actuator Characteristics

In 2005, JCI introduced wireless sensing. These devices sense space temperatures (as do their wired counterparts) and forward temperature information wirelessly to its room controller. Wireless communication reduced installation cost by easing sensor installation. These devices deployed an 802.15.4 star architecture. In 2007 a mesh technology sensor expanded the coverage area for reporting temperature data by transmitting the temperature data across the mesh. Since the sensors monitor the environment and inject status data onto the network, many times these devices can be deployed using battery power. This is not true for their actuator counterparts. Actuators change the environment by modulating dampers, opening and closing doors and the like. The very nature of these devices most often deems battery power insufficient to perform the task. Since actuators for the most part require line power, the installation cost reduction to communicate wireless is thwarted. JCI has no immediate plans to build wireless actuators.

Fire sensing and response is considered the highest priority function in Facility Management systems. Security systems rank second followed by HVAC and Lighting applications. Historically, fire and safety sub-systems have been hard-wired or have been implemented on totally dedicated infrastructure to ensure that the fire and security systems are not affected by the HVAC and lighting sub-systems. Market and customer pressure however, is changing this approach since customers want application interaction across these systems with the HVAC and lighting sub-systems.

Sensor/Actuator Functionality

Sensors are normally fixed function devices deployed on an 8-bit microcontroller running in 32K to 128K memory space. MAC Addressing is set via local dip switches. Some sensors may employ a user interface for example to adjust the temperature setpoint, extend the occupancy or set other local parameters. Most sensors though are self-contained fixed-function devices with no user interface.

Sensor/Actuator Addressing

Sensors reside on the Sensor bus. The sensor bus incorporates the same 8-bit MAC address as the field bus. On this network the area controller will take a MAC address of 0 by convention. The sensor bus carries a limit of 16 devices. These devices may be slave or master devices.

Sensor/Actuator Emergency Power

Most sensors and actuators do not support any emergency power capability. Sensors and actuators are designed to latch their last known value in case of a power interruption. All parametric data is persisted. Sensors often are battery powered devices which eliminates any need for emergency power. Low battery conditions are resolved since each update includes a battery status indication. This allows months of operation in a low power state before the batteries must be replaced. The controllers have fail-soft functionality. If communication is lost to one or more of the sensors, the controller will continue to operate in a depleted mode until the sensor is replaced. Actuators many times are configured with spring returns that will dictate its default position in case of a power failure.

Sensor/Actuator Communications

Sensors are simple devices having a limited protocol repertoire. Sensors are traditionally 3-wire twisted pair devices on an EIA-485 multi-drop communication network utilizing some of the primitive functionality of BACnet MS/TP. Bus length can run to 5000 feet without a repeater upwards to 15000 feet with the maximum of 3 repeaters. Most sensor buses run no longer than 50 feet since the sensors tend to reside in close proximity to the controller.

Wireless temperature sensors utilizing the ZigBee mesh protocol have been recently introduced. These sensors are less expensive to install and provide better sensing capability since the sensor can be placed in the optimal sensing location. Cost savings are further realized as buildings are retrofitted due to remodeling or new tenant requirements.

To date, JCI has investigated deploying sensors as IP devices on the network. However, the cost of running a star (switch-client) topology and adding the required backhaul infrastructure is not currently economical. Investigation of utilizing 802.11 wireless infrastructure was also considered for sensor inputs. However, the radio cost and expected battery life limitations cannot warrant this technology. As radio costs decrease and the WLAN protocols are enhanced, JCI may consider the technology in the future.

Sensor/Actuator Network Security

Sensors typically are affixed on a wall and could be pilfering or vandalism candidates; however they are not hacked devices. Most actuators sit above the drop ceiling or in locked down equipment rooms. There is currently no market requirement for any network security for HVAC or Lighting sensors or actuators.

Area Control

An area describes a small physical locale (300 – 500 ft^2) within a building, typically a room. As noted in Figure 2.1a the HVAC, Security and Lighting functions within a building address area or room level applications. Area controls are fed by sensor inputs that monitor the environmental conditions within the room. Common sensors found in many rooms that feed the area controllers include temperature, occupancy, lighting load, solar load and relative humidity. Sensors found in specialized rooms (such as chemistry labs) might include air flow, pressure, CO_2 and CO particle sensors. Room actuation includes temperature setpoint, lights and blinds/curtains.

The controllers deployed within a room are most often standalone devices that can provide the necessary functionality without further assistance by the higher layers of the system. However when these devices are connected to the higher system layers, these controllers can provide manual override, time series and event data to the higher layers for further analysis. Likewise, the enterprise level can then override the local control from a centralized location. When connected to the higher layers, the controllers deploy a fail-soft algorithm that reverts to local control if the higher order communication is lost.

Room temperature controllers are soft real-time devices implementing typically 60 second control loops. Environmental data is provided to the controller by its sensors in either a polled or event driven fashion. The controller then analyzes the data and modulates the actuators accordingly to meet the application requirements. Actuators are modulated each minute to maintain proper temperature, airflow and humidity.

Door control requires much higher performance. Persons entering a facility will expect a latency of no more than 500msec between swiping the access card and entry approval. Camera pan-tilt-zoom commands need to execute with less than 250msec latency.

Room lighting control also requires real-time performance. Room lights themselves need to have near instantaneous response to a light switch activation. The lighting operator will expect to see some change in the scene within 500msec after a complex lighting command has been executed. A list of area controller characteristics is defined in Figure 2.2.

Area Controller Characteristics

- Setup
  - MAC Address set via onboard switches
  - Programmed via Q&A queries to define application required
  - Archive application in non-volatile memory
- Operation
  - Monitor all required and optional sensor inputs for timely updates.
  - (re)calculate real-time control algorithms periodically as inputs change
  - Issue actuator directives as required to perform the required application
  - Monitor all required and optional actuators and sensors for error conditions.
- Reporting
  - Forward sensor data (e.g. outdoor air temp) to other system nodes at their request
  - Report sensor/actuator failures to higher layers
  - Report application alarms (i.e. failure to meet desired goals) to higher layers
  - Maintain statistics such as total runtime, communication errors, and operational longevity for system diagnostics.
- User Interface
  - Report current control state and sensor data on demand (optional)

Figure 2.2 Area Controller Characteristics

Area Controller Functionality

Some Area controllers are fixed function, but most are configurable embedded devices equipped with 64K to 256K flash memory. A proprietary configuration tool is used to define the application functionality in a Q&A format. The JCI Branch tech answers questions posed from the configuration tool. The configuration tool will develop the resulting controller download file based on the Q&A session. The tool will also produce the control flow drawings, the BOM and define the requisite input sensors and output actuation requirements.

The 8-bit MAC address is set either via an onboard switch bank or via the configuration tool. Some area controllers may employ a user interface for example to adjust the temperature setpoint, extend the occupancy or set other local parameters; however, most area controllers are self-contained devices with no user interface.

Area Addressing

Area controllers reside on the MS/TP Field Bus. An 8-bit MAC address is set either via an onboard switch bank or the configuration tool. The 8-bit address provides an addressable range of 254 devices, since address 255 is reserved for broadcast designation. A given field bus will carry nominally 25 devices. This bus may extend to over 100 devices depending on the application. All devices on the field bus are masters allowing peer-to-peer communication and hosting of temporary configuration devices.

Area controllers also support a local MS/TP Sensor Bus. The sensor bus incorporates the same 8-bit MAC address as the field bus. On this network the area controller will take a MAC address of 0 by convention. The sensor bus carries a limit of 16 devices. These devices may be slave or master devices.

Area Emergency Power

Area controllers provide actuation function such as moving a damper, opening the door or turning on the lights. Since this function requires a high power energy source, area controllers are driven from a continuous 24 VAC supply. This lower voltage is sufficient to drive most actuation requirements, yet is considered low voltage allowing lower cost installation procedures.

Most area control will cease during a power outage. The design of the controllers and actuators defaults to a ‘safe’ state during a power outage. In the case of a Variable Air Volume (VAV) room controller, the damper will mechanically be limited to a minimum air flow. Proper state and operation will automatically return when power is restored. Area controllers and actuators in mission critical applications such as Operating Rooms, or Clean Rooms require UPS support to assure continued operation during power outages.

Area Controller Communications

Area controllers need to communicate to higher order (zone) controllers as well as its subordinate sensors and actuators. The communication network is implemented with 3-wire twisted pair media on an EIA-485 multi-drop BACnet network. Because the media and protocol are consistent from the sensor to the controller to supervisory controller these devices could all reside on the same network. However, the area controller most often deploys two physical EIA-485 buses; one for the sensors (Sensor Bus) and one for the controllers (Field Bus) as shown in Figure 2.2a.

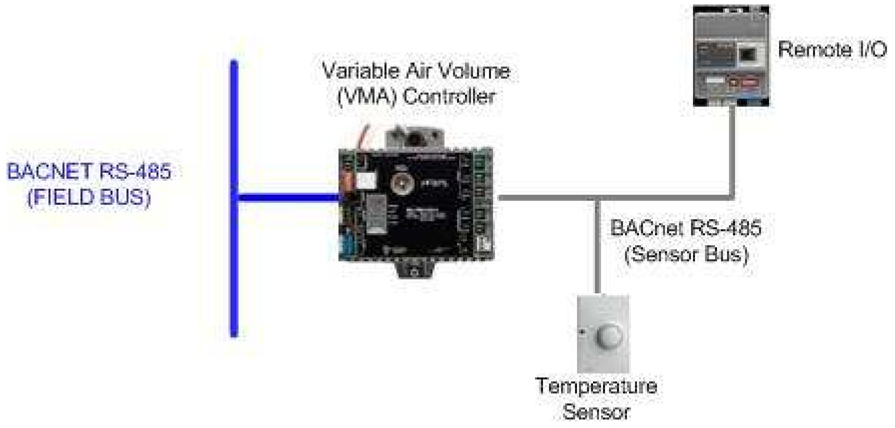

Figure 2-2a Sensor and Field Wired Bus Communication

Wireless mesh communication of sensors and area controllers was deployed in 2007 allowing for less expensive installation and retrofit costs. Here the sensors and controllers all reside on the same ZigBee PAN. The architecture allows intermixing wired and wireless sensors and controllers. This allows the field application engineer to decide the best tradeoff for the application. Figure 2.2b depicts wireless communication; Figure 2.2c a wired/wireless solution.

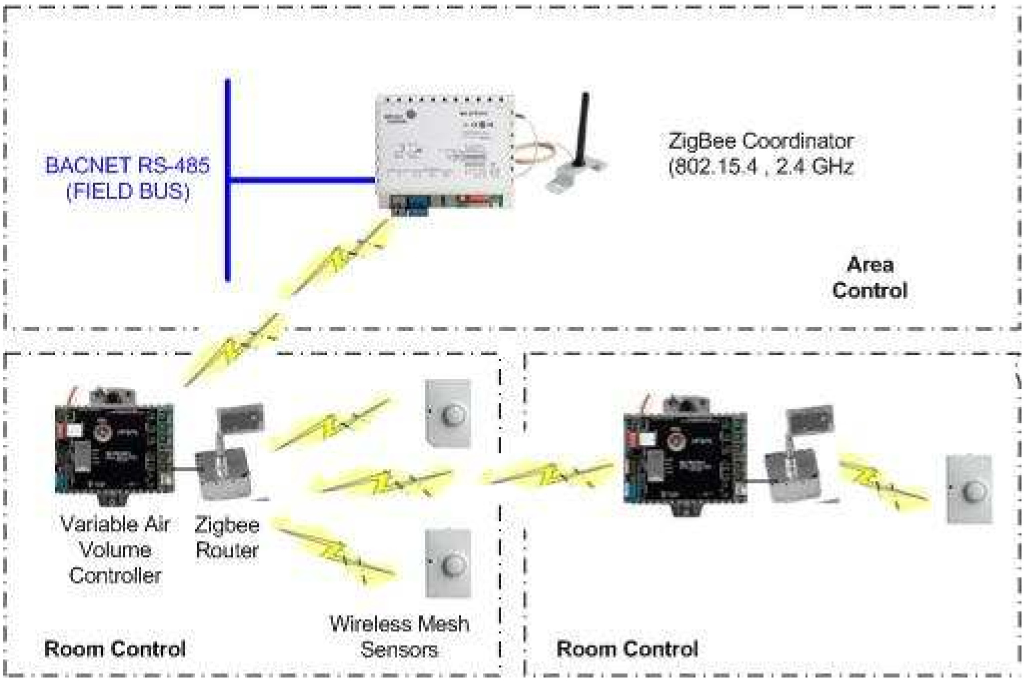

Figure 2-2b Sensor and Field Wireless Bus Communication

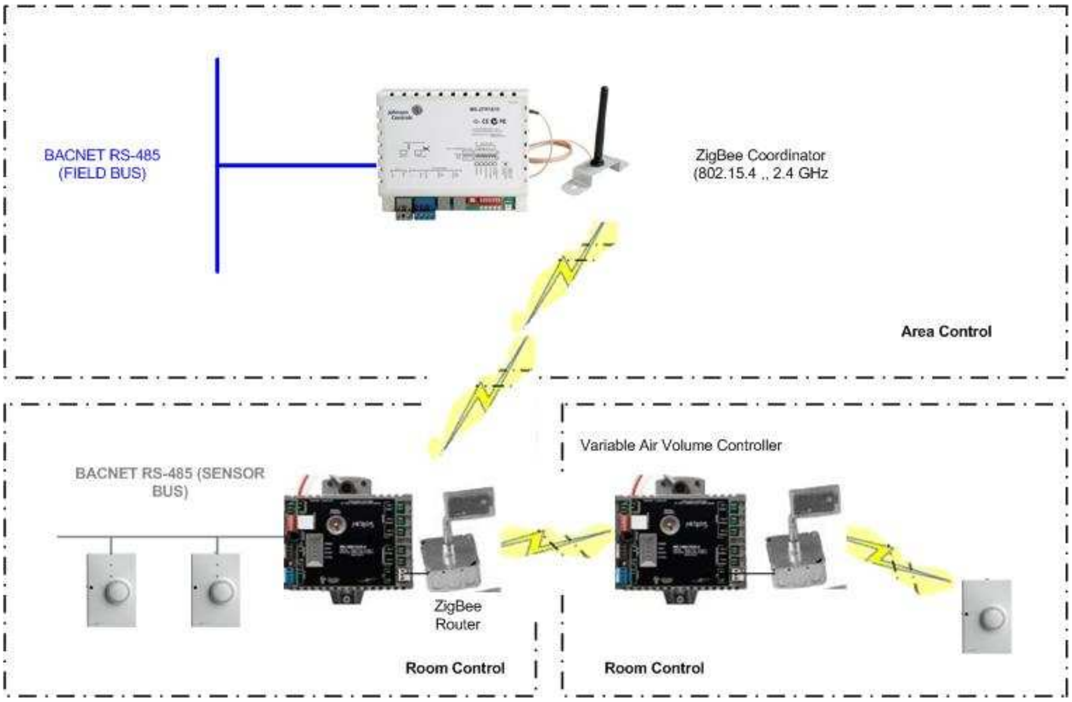

Figure 2-2c Sensor and Field Wired and Wireless Bus Communication

JCI has investigated deploying area controllers as IP devices on the network. However, the cost of running a star (switch-client) topology on a controller network is economically restrictive at this time. Some controllers have been redesigned to also reside on the enterprise (Ethernet) network. This is currently only economically viable for a small set of the area controllers.

JCI has developed IP segment extenders that allow a segment of the EIA-485 trunk to tunnel through UDP/IP to a remote location (Figure 2.2d). This product has proven effective in WAN based applications such as School Districts. Here, the higher level controllers can reside on a server farm in a centralized locale. The area controllers can then be deployed in each of the remote locations (e.g. schools). The IP network then acts as a wide-area transport allowing the devices to be connected into a logical LAN although actually deployed as a WAN. JCI deployed this extender device as both an Ethernet and a WLAN extender. To date, most applications have opted for the Ethernet appliance.

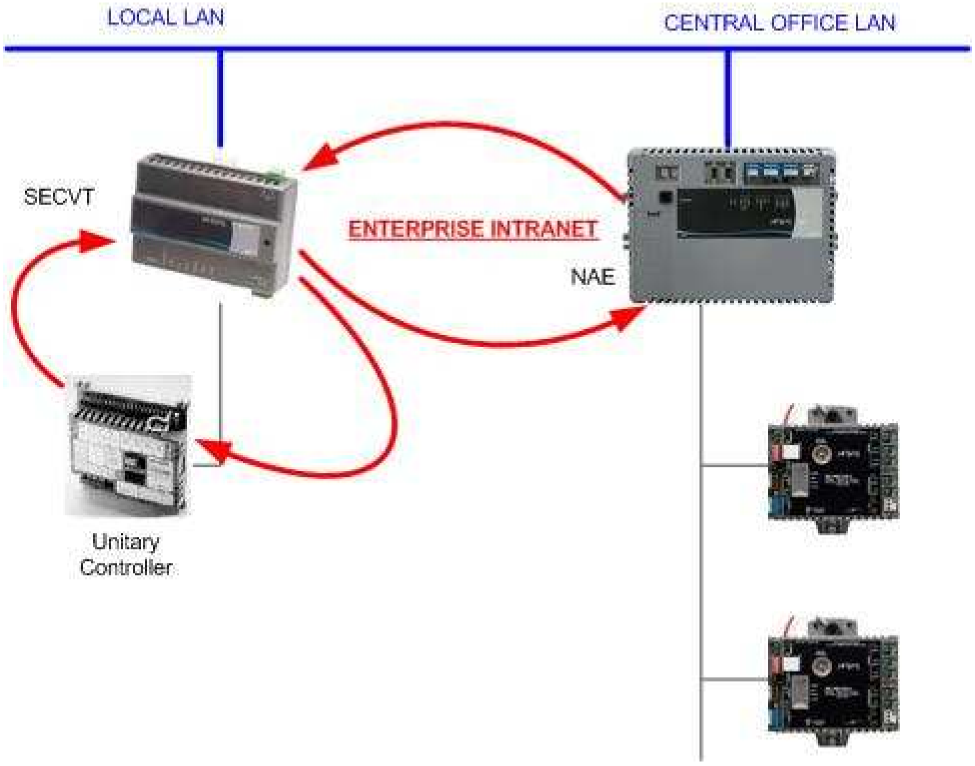

Figure 2-2d Tunneling EIA-485 Control Over Enterprise WAN

Area Controller Network Security

Most area controllers sit above the drop ceiling or in locked down equipment rooms. There is currently no market requirement for any network security for wired or wireless area controllers. The ZigBee mesh controllers adopted in 2007 support AES-128 encryption which may be deployed when the requirement surfaces.

Zone Control

Zone Control supports a similar set of characteristics as the Area Control albeit to an extended space. A zone is normally a logical grouping or functional division of a commercial building. A zone may also coincidentally map to a physical locale such as a floor. Table 2.2 describes some examples of zones for the various functional domains within a commercial building.

Functional Domain	Zone
HVAC	Air Handler – the area served by a single fan system; typically a floor or adjacent floors in a building.
Lighting	A bank of lights that all operate consistently
Fire	An area of a facility that will all operate consistently for example fed by the same fan system; covered by the same set smoke detectors or follows the same pressurization and annunciation rules. The zone may also be a functional grouping when a certain area is governed by a set of fire dampers.
Security	A subset of the building operating in a similar fashion for example a logical collection of lockable doors.

Table 2.3 Example of Commercial Zones

Zone Control may have direct sensor inputs (smoke detectors for fire), controller inputs (room controllers for air-handlers in HVAC) or both (door controllers and tamper sensors for security). Like area/room controllers, zone controllers are standalone devices that operate independently or may be attached to the larger network for more synergistic control.

Zone controllers may have some onboard sensor inputs and also provide direct actuation; however, zone controllers will also direct the actions of its underlings via commands as well as respond to environmental changes reported by its underlings. For example, an Air Handler controller might directly sample the duct pressure, the supply air temperature and return air temperature. However, it may also send commands to other networked devices querying the outdoor air temperature and relative humidity. Similarly, a fire panel may have all the smoked detectors directly wired; yet send commands to other adjacent fire panels to request their status if a fire condition arises. A list of zone controller characteristics is defined in Figure 2.3.

Zone Controller Characteristics

- Setup
  - MAC Address set via onboard switches
  - Programmed via Q&A queries to define application required
  - Archive application in non-volatile memory
- Operation
  - Monitor all required and optional sensor inputs to timely updates. Report failures to higher layers.
  - (re)calculate real-time control algorithms periodically as inputs change
  - Issue actuator directives as required to perform the required application
  - Issue commands to area controllers and actuators to maintain proper control per the application defined.
  - Monitor all required and optional sensors and actuators for error conditions. Report failures to higher layers.
- Reporting
  - Forward sensor data (e.g. outdoor air temp) to other system nodes at their request
  - Report sensor/actuator failures to higher layers
  - Report application alarms (i.e. failure to meet desired goals) to higher layers
  - Maintain statistics such as total runtime, communication errors, and operational longevity for system diagnostics.
- User Interface
  - Report current control state and sensor data on demand (optional)

Figure 2.3 Zone Controller Characteristics

Zone Controller Functionality

Zone controllers must meet a diverse set of application scenarios. They are typically completely field programmed. Most zone controllers are programmed via a Q&A session as done with the Area controllers. Zone controllers are memory based 16-bit devices with upwards to 1mb of ROM and 256kb of RAM. The JCI Branch tech answers questions posed from the configuration tool. The configuration tool will develop the resulting controller download file based on the Q&A session. The tool will also produce the control flow drawings, the BOM and define the requisite input sensors and output actuation requirements. Off-box data references to other global information (e.g. Outdoor Air) will be dynamically discovered via the BACnet protocol at system boot time. Some zone controllers may employ a user interface for example to adjust the temperature setpoint, extend the occupancy or set other local parameters; however, most zone controllers are self-contained devices with no user interface.

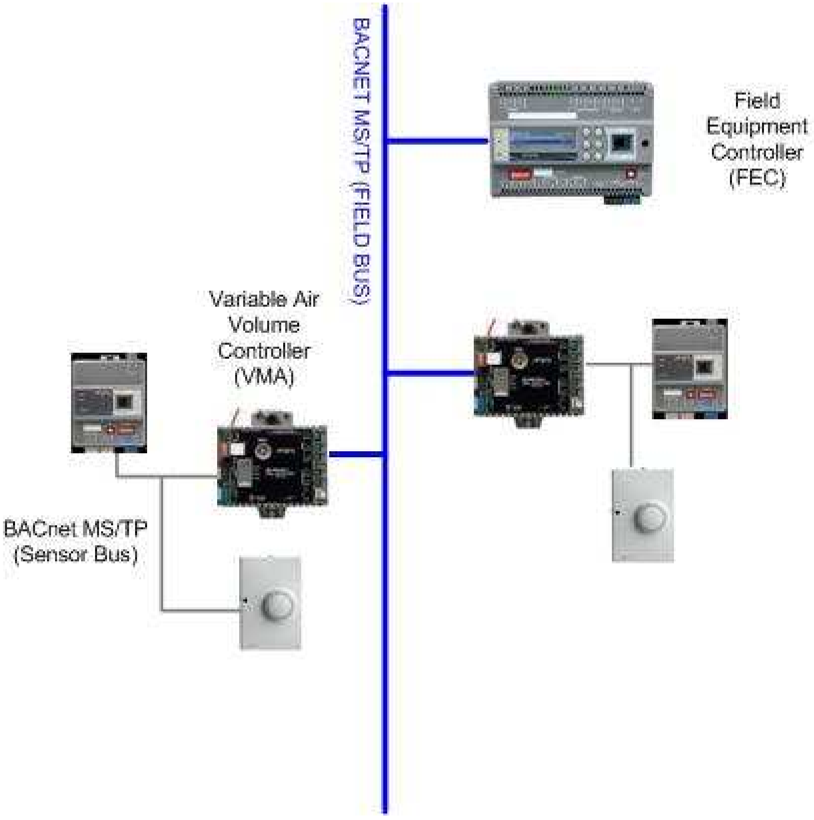
BACNET MS/TP (FIELD BUS)
Field Equipment Controller (FEC)
Variable Air Volume Controller (VMA)
BACnet MS/TP (Sensor Bus)

Figure 2-4 Zone Controller (FEC), Area Controllers (VAVs) and Sensors

Zone Addressing

Zone controllers reside on the MS/TP Field Bus. An 8-bit MAC address is set either via an onboard switch bank or the configuration tool. The 8-bit address provides an addressable range of 254 devices, since address 255 is reserved for broadcast designation.   A given field bus will carry nominally 25 devices. This bus may extended to over 100 devices depending on the application. All devices on the field bus are masters and hence peer-to-peer operation is supported.

Some zone controllers also support a local MS/TP Sensor Bus. The sensor bus incorporates the same 8-bit MAC address as the field bus. On this network the zone controller will take a MAC address of 0 by convention.   The sensor bus carries a limit of 16 devices. These devices may be slave or master devices.

Zone Emergency Power

Zone controllers are line voltage or 24VAC devices. These devices will typically cease operation in case of a power outage. Lighting, Security and Fire zone controllers are often tied to the emergency power systems to continue operation in a curtailed mode during a power outage. HVAC control most often ceases operation except in mission critical applications such as manufacturing control, white rooms and hospital operating rooms.

Zone Controller Communications

Zone controllers need to communicate to higher order building controllers as well as its subordinate area controllers, sensors and actuators. The communication network is implemented with 3-wire twisted pair media on an EIA-485 multi-drop BACnet network. Because the media and protocol are consistent from the sensor to the controller to supervisory controller these devices could all reside on the same network. However, the area controller most often deploys two physical EIA-485 buses; one for the sensors and one for the controllers as shown in Figure 2.2a.

Zone Controller IP Communications

JCI has investigated deploying zone controllers as IP devices on the network. However, the cost of running a star (switch-client) topology on a controller network is economically restrictive at this time. Some zone controllers have been redesigned to also reside on the enterprise (Ethernet) network. This is currently only economically viable for a small set of the area controllers.

Zone Controller Network Security

Most zone controllers sit above the drop ceiling or in locked down equipment rooms. There is currently no market requirement for any network security for wired or wireless zone controllers. The wireless mesh zone controllers adopted in 2007 support AES-128 encryption which will be deployed when the requirement surfaces.

Building Control

Building Control (aka Supervisory Control) provides the overall orchestration of the system. While the sensor, area and zone controllers provide real-time narrow focused applications; the Building Controllers provide broad systemic functionality. The building controllers provide the view ports into the embedded real-time systems for the operator, integrators and enterprise applications. Building controllers will cache and archive important real-time data from the controllers and act as an agent to the Building Servers layer for long-term data archival and retrieval. Building Controllers receive event information for the lower layers and forward the information to all needed devices and systems.

Building Controller Characteristics

- Setup
    - BACnet MAC Address set via onboard switches
    - IP Address – Static or DHCP settable
    - Scans all sensor and controllers to define its database
    - Full programming language for application customization.
- Operation
    - Monitors all subsystems expected behavior.
    - Overrides local control as needed to provide systemic operation
    - Implements system applications such as Electrical Demand Limiting
    - Cooperates with other silos to add complete integration support for HVAC, Fire, Security and Lighting applications
    - Integrates various controller protocols (N2, BACnet, LON) into a single

object model

  - Interoperates with 3rd party control devices
  - Provides all IT ‘friendly’ client capabilities including DNS, DHCP, SMTP and SNMP support
  - Provides rigorous IT security policies
- Reporting
  - Receives event and alarm indications from lower layers
  - Directs alarms notifications to requested users and processes. Caches alarms until acknowledged by users
  - Caches time series data for underlings. Uploads time series data to server farms
  - Provides modem interface for dial up connections
- User Interface
  - Provides Web Server support to users
  - Provides SNMP Server (get, trap) to Network Management Systems
  - Provides alternate indications to uses via pagers, printers
- Data Access
  - Provides BACnet data access for reading/writing data
  - Provides SNMP Server (get, trap) to Network Management Systems
  - Supports public web service interfaces

Figure 2.3 Building Controller Characteristics

Building Controller Functionality

Building Controllers are completely field programmable devices that are designed to integrate all system control operations. These devices also contain the user interface support for access by facility operations.

The HVAC Building Controllers are designated Network Automation Engines or NAEs. These devices come in multiple sizes ranging from embedded Win CE running 128mb flash memory to models handling dozens of controllers to Windows XP Server class models supporting thousands of controllers. These models also support various numbers and types of communications trunks including BACnet MS/TP, BACnet IP, LON, N1 and N2. N1 and N2 are legacy JCI trunks.

The Fire subsystem application is standalone in many cases dictated by the fire codes. However, the NAE most often monitors the Fire subsystem as a secondary reporting device. Here the smoke detectors, pull boxes, strobes and evacuation subsystems are integrated into the NAE for viewing and monitoring by building operations. By regulation, the HVAC system cannot affect changes to the fire system.

However, the fire subsystem may be further integrated into the NAE in cases where the HVAC system operates in concert with the Fire subsystem to provide a smoke abatement application. This application is further explained in the Building Application section following.

The Security subsystem will also standalone from a control point-of view. As noted above, local door controllers will support building entry algorithms. Cameras may be controlled from a centralized location. An optional centralized video server may be deployed to allow remote wireless viewing of cameras. This server may also support motion alerts on unexpected changes in the camera’s view. The Security system can also be tied into the HVAC system to facilitate the experience of someone entering a facility. This application will also be explained further in the Building Application section.

Lighting applications are most often localized to a room or area. Lighting manufacturers do not deploy server level devices to control the entire facility. Rather, they provide application ‘hooks’ into the lighting panels that allow the FMS to monitor and override the local lighting algorithms.

The NAE provides all the overall monitoring and control of these silos. As expected, the application requirements are job specific and hence require significant local effort to meet the solution expected by the customer. The JCI Branch network is well positioned to provide all needed tailoring of the system to meet the specific job requirements. Figure 2.4a shows the complete HVAC hierarchy. Figure 2.4b expands this to the entire FMS.

The Building Controller is completely field programmable and can be extended to provide other capabilities beyond those described. These include elevator control; fume hood monitoring and control; PLC monitoring and maintenance management. Due to the generic protocol interfaces employed and its field programmability, the Building Controller can be configured to interface directly to most any commercial building device employing a network connection. The JCI protocols have been available to any vendor for integration into its products since 1992. To date, the Metasys system has been tested for compatibility with over 500 3rd party products.

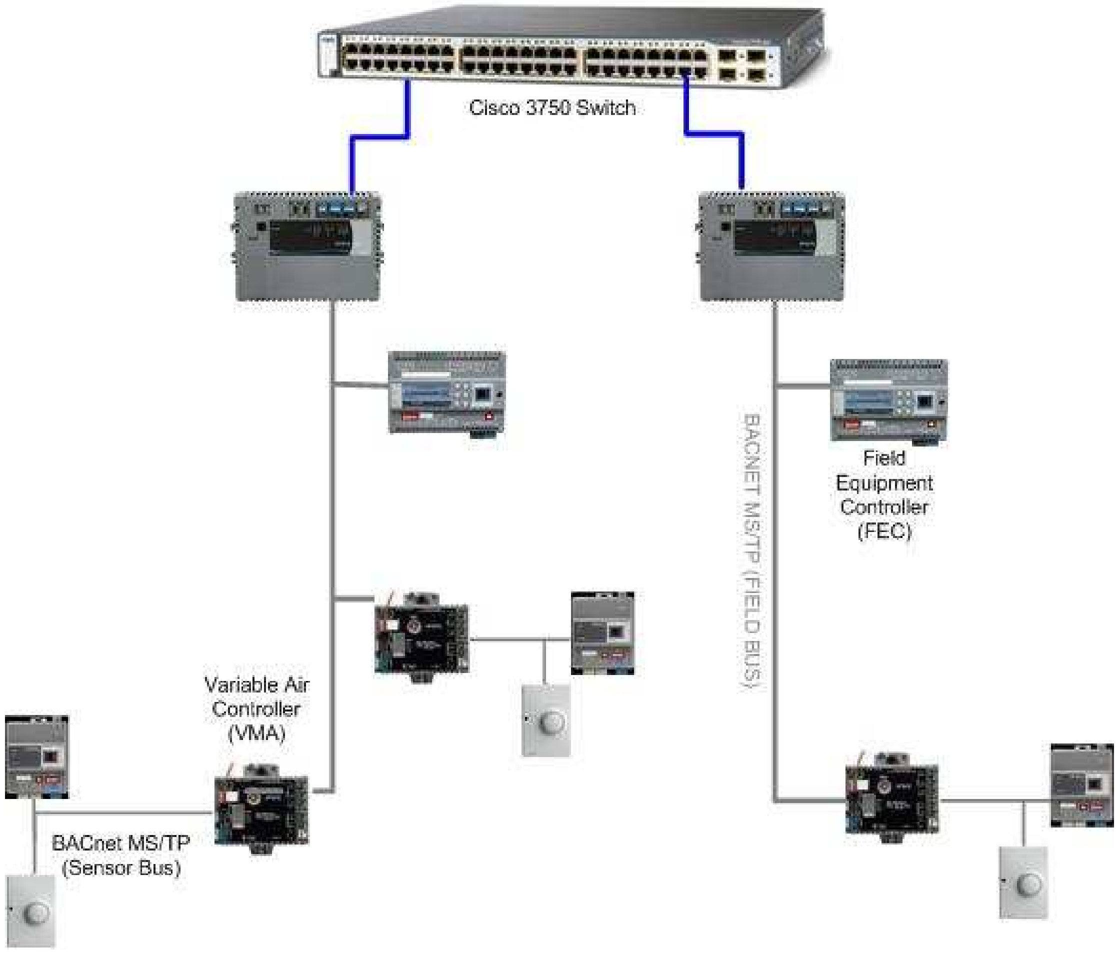
Cisco 3750 Switch
BACNET MS/TP (FIELD BUS)
Field Equipment Controller (FEC)
Variable Air Controller (VMA)
BACnet MS/TP (Sensor Bus)

Figure 2-4a NAEs connecting JCI Controllers to the Enterprise through a Cisco Workgroup Switch

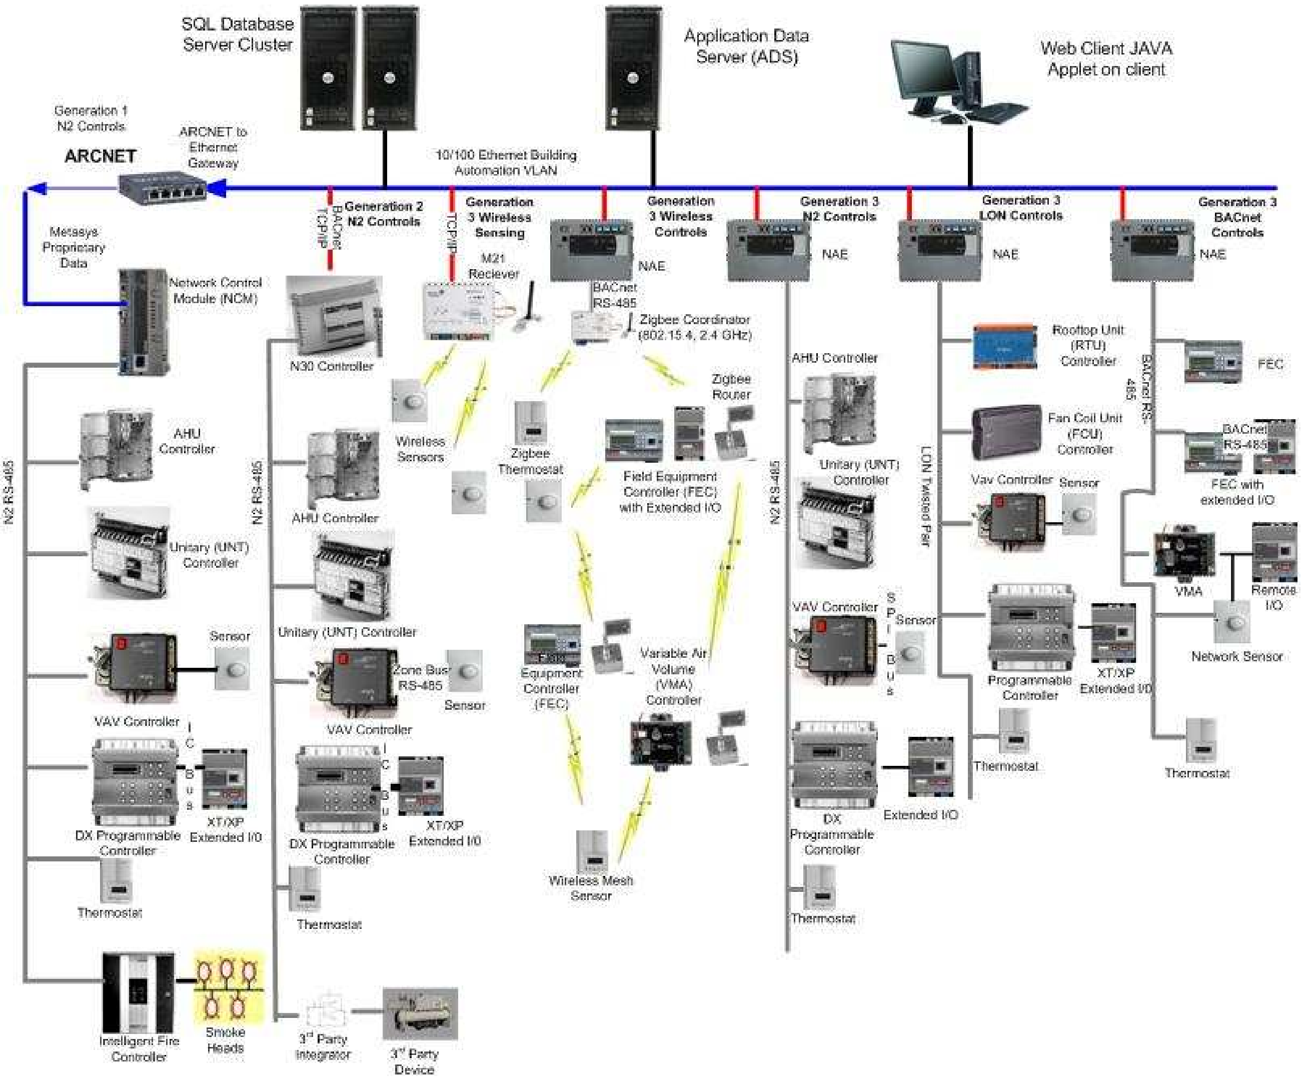

Figure 2-4b Inclusive set of HVAC Controller Options

Building Controller Addressing

The NAEs are integrating devices that morph many diverse systems into a single logical model. In this regard, the NAE will require multiple addresses consistent with the technologies involved. Table 2.4a lists the potential addresses and address assignments required in a single NAE

Media	Address
IP	IPv4. This address may be assigned via DHCP or set as a static address
N1	IPv4. This address may be assigned via DHCP or set as a static address
N2	N/A. The N2 is a Master/Slave protocol. The NAE is the bus master and needs no MAC.
MS/TP	8-bit (0:255) The NAE by convention will be assigned MAC Address 0
LON	7-bit (0..127) The NAE by convention will be assigned MAC Address 0

Table 2.4a MAC Address Assignments in an NAE

The Metasys network can expand horizontally with no stated limit. Small Metasys systems will incorporate 1 to 5 NAEs; large systems will support over 100 NAEs on the network. As noted above each NAE can support 200 controllers, each controller can support up to 16 sensors.

Building Controller Emergency Power

HVAC Building Controllers (NAEs) most often do not require emergency power. If an HVAC system is supporting the Fire system in the smoke abatement functions, the NAEs will be required to be on a UPS. The NAEs are designed to query all subservient devices upon reboot and regain its quiescent state once power is reapplied.

Building Controller Communications

As stated in the overview, the Building Controller supports various protocols as needed by the applications (see Table 2.4b).

Protocol	Media	Application
AT Command Set	EIA-232	Dial-up Communication
n/a		Printing
N2	EIA-485	JCI Legacy Control
BACnet MSTP		JCI Controllers and Sensors 3rd Party Integration
LON		JCI Controllers and Sensors
Private Web Services/SOAP	TCP/IP	JCI Internal Secure Communications
Public Web Service/SOAP		3rd party FMS Access
SNMP	UDP/IP	NMS Access
BACnet IP		NAE to NAE Control 3rd Party Integrations

Table 2.4b Protocol and Application Cross Reference

A USB port on the NAE can optionally be defined to support a printer or modem. These devices are most often utilized in a small office environment. Larger systems typically will deploy enterprise printers or remote access servers to fill these rolls.

The NAE supports three control protocols, N2, LON and BACnet. N2 is a JCI legacy proprietary protocol developed in the 1980's. This protocol was opened in 1991. JCI has promoted this protocol as a standard and has integrated upwards to 300 commercial building products since that time with the objective of providing the customer a highly synergistic cross-vendor system. Due to this extensive partner network, JCI will support the N2 protocol for many years.

LON is a standard open protocol developed by Echelon. It provides sensor and controller connectivity. Its applications are well defined allowing customers to find 'pin compatible' products across vendors. LON devices, however, require single-sourced communication chip, the neuron, for node-to-node communication. LON also has a limited set of configuration tools available to configure the resident applications. LON devices are most prominent at the sensor and room controller layers.

BACnet is a software only protocol supporting 5 media types; a MAC, network and transport layers. BACnet received ISO status in 2002. The latest generation Metasys systems support BACnet sensors and controllers. The software architecture will convert all LON and N2 device communication into the BACnet object model at the lowest application layers of the system. Higher order applications will then act on LON and N2 devices in a consistent manner with native BACnet devices.

NAE-to-NAE communications on the Ethernet also utilizes BACnet. Allowing BACnet on the Ethernet network allows JCI devices to cooperate with other vendor BACnet devices to interoperate in a client and/or server mode.

Secure communication required at the Ethernet level is implemented using private web services. This communication typically takes the form of a dialog between the NAEs and the enterprise servers. Some of these web services have been documented and opened as public web services. This allows applications to access real-time environmental data from the FMS. Archived data is also readily accessible via the optional SQL server albeit not for real-time data.

The N1, BACnet/IP and SNMP protocols all run over UDP/IP. N1 is a JCI legacy proprietary protocol that ran on the Network Control Module (NCM). The NCM was the predecessor supervisory controller to the NAE. The NCM and NAE performed very similar control and orchestration functions of the subservient controllers. The NCM, however did not have the hardware capabilities of running an embedded Windows OS or a web server. The N1 protocol was opened in 1991 over Arcnet. This protocol was revamped in 1996 to run on Ethernet IP networks. The N1 uses the datagram services of UDP and adds guaranteed delivery and flow control as its transport layer. Coincident with the Ethernet effort, JCI released an Arcnet to Ethernet gating device as a migration path for existing Arcnet based customers. Over the past decade most Metasys customers have replaced their Arcnet based supervisory devices with Ethernet devices. While we occasionally still see an Arcnet customer needing to migrate, these cases have become increasingly rare.

The N1 protocol was also superseded with BACnet/IP in 2001 as part of JCI's support of the industry standard. An N1 to BACnet/IP gating device, the Network Integration Engine (NIE), was developed as a migration strategy to allow existing customers to reap the benefits of using a standards based protocol.

BACnet/IP is the predominant control protocol on the Ethernet. Like the N1, BACnet/IP runs UDP/IP with network and transport layer functionality added. Running inter-NAE communications using BACnet completes the end-to-end BACnet implementation across all physical and functional layers of all JCI sensors, actuators and controllers. BACnet/IP also allows easy 3rd party interaction with any vendors supporting the BACnet stack.

As strong as BACnet is on providing a solid and stable control application model; it is weak in terms of an API interface to the enterprise users and applications. The NAE therefore morphs BACnet data into public web services and SNMP for easy interface to the enterprise application suite. The public web services, while JCI defined are open and accessible on the JCI corporate web sites. BACnet has recently augmented its protocol with web service definition. These web services align to the existing JCI web services. While there is currently no plan to convert the JCI web services to the BACnet set, this effort may be executed on future customer requests.

The NAE supports SNMP 'gets' and 'traps' interfaces. Currently, the NAE will not allow 'sets' via SNMP as a FMS application security measure. JCI has a registered MIB supporting the SNMP interface. The interface allows NMSs (e.g. HP Openview) to access most interesting real-time building data (gets); and FMS alarm and event information (traps).

Building Applications

The Building Application layer is a software layer that binds the various system silos into a cohesive systemic application. This discussion in not meant to be inclusive. Rather it is meant to show how these diverse systems can be coordinated to provide innovated synergistic applications for the customer. These applications are rooted at the development centers but are highly customized on a per job basis by the JCI field organization.

Fire and Smoke Abatement

Most local codes now require commercial buildings to incorporate comprehensive fire and life/safety systems into a building. It is well documented that loss of life in a building is mainly caused by smoke inhalation rather than the fire itself. UL has a fire certification program (UL-864) that governs fire and smoke operations in commercial buildings. This program requires very rigorous interactive testing with UL for certification. In addition to the obvious need to minimize life/safety situations in a building, facility operators are highly encouraged to implement these systems due to insurance cost reductions.

The UL fire and smoke systems operate in either a manual or automatic mode. The manual mode provides critical fire and smoke information on a display to be controlled by a Fire Marshal. The automatic mode is a preprogrammed set of events that control the fire automatically. In practice, the fire system will be set to automatic mode and operate accordingly until the Fire Marshall arrives. At that point the system is normally overridden to manual mode so that the Fire Marshall can control operations from the command center as deemed necessary.

UL-864 is comprised of the fire system operations (UOJZ) and smoke control (UUKL). UOJZ certification allows all fire and smoke operations, events and alarms to be controlled from a Fire Workstation. Local Fire panels can only be accessed and commanded from this workstation. Operator authentication and command authorization are required for all operations. Alarms can only be acknowledged from this(these) device(s). One and only one Fire Workstation can ever govern a given area at a time to assure that destructive control operations cannot inadvertently occur by two operators controlling a space simultaneously.

The smoke abatement certification, UUKL, is an adjunct function of the fire system that automatically or manually purges the fire and directs smoke safely out the building by exhausting smoke from exit passageways and refuge areas by judicially adjusting pressures and dampers in the affected areas. Furthermore, it will actually assist in putting out the fire by starving the fire of oxygen in the affected area while simultaneously routing smoke out the building in the adjacent areas.

While the smoke abatement operation could be the province of the fire system alone, economics dictates that the fire system off-load the smoke abatement operation to the HVAC system. In practice, the fire system will receive the initial fire indication by one or more of its smoke detectors. It will then inform the HVAC system of the physical locale of the fire. The HVAC system will then take charge of the smoke abatement operation by automatically adjusting the air handlers and dampers. The HVAC system must incorporate a comprehensive prioritization scheme throughout its system. This prioritization scheme must allow all smoke operations to take control precedence over all other control operations including manual operator control. All affected devices must support a supervision policy that assures that all operations requested were executed properly. The system must automatically return to normal operation once the smoke situation has abated.

Many buildings will also trigger the evacuation application (see below) coincidentally with a smoke control situation. The evacuation application will assist building inhabitants to safely leave a building. Elevator control policies may restrict inhabitants from calling for the elevators while simultaneously posting the elevators to the ground floor by use of the fire personnel.

Evacuation

Evacuation is another systemic operation that may be activated as part of the Fire/Smoke Control application, or may be activated for other reasons such as terrorist threats. Evacuation requirements most often will activate subsystems of the Fire, Security and Lighting silos. The Fire system normally supports the intercom subsystem in the facility. The intercom system will then trigger the recorded voice evacuation instructions. This may be in concert with the fire system audio indications if a fire situation is active or standalone. The lighting subsystem will be activated to turn on the lights and evacuation paths to aid in the evacuation. The security system will coincidentally open all doors to allow a smooth safe egress from the building. If the building also supports elevator control, the elevators will operates as directed by a preprogrammed evacuation policy.

Occupancy/shutdown

A major energy saving technique in commercial buildings is to automatically commence HVAC and lighting operations prior to building occupancy. Conversely, building shutdown allows the systematic reduction in HVAC and lighting operations as the building becomes unoccupied.

The HVAC system is usually charged with defining occupied and unoccupied times. The Fire and Security operations are always operable and lighting is most often subservient to HVAC. These times are typically programmed into the system by facility operations; however, it could be learned adaptively by the security's access control system. The target occupancy time drives the HVAC subsystem to turn on all ventilation equipment at an optimal time so that each space is ready for occupancy at the prescribed time. These algorithms will be adaptive over time but also include systemic instrumentation such as outdoor air and relative humidity to turn on the equipment at the last possible moment yet still meet the target environmental needs just before occupancy. The lighting systems will also be turned on just prior to occupancy.

Conversely, the HVAC systems will also determine the earliest possible time it can shut down heating/cooling yet still control the setpoints to meet the requisite parameters. Lighting again gets off easier since the lights can be extinguished as soon as they are not needed. Building owners may use the lighting systems to pace the janitorial service providers by defining a strict timetable that the lights will be on in a given area. Here, the janitorial service providers will need to keep in step to complete their work prior to the lights being turned off.

The Metasys system also includes a telephone interface that allows any late workers to override the normal HVAC and lighting schedules simply by dialing into the system and specifying their locale. The lights and fan system will continue to operate for a few extra hours in the immediate vicinity. The same applies to occupancy sensors in meeting rooms. Either by automatic sensing or a simple push of the occupied switch, the HVAC and lighting schedules will extend the normal schedule for the meeting room.

Energy Management

The occupancy/shutdown applications noted above optimize runtime of large equipment. This in itself is a major component of energy savings. However, even during occupancy large equipment can be modulated or shutoff temporarily without affecting environment comfort. This suite of applications run in the HVAC domain, however the HVAC silo will interact with the lighting system to reduce the lighting load to help in the overall reduction of energy.

The load rolling and demand limiting applications allow for the sequencing of equipment to reduce the overall energy profile or to shave off peak energy demands in the facility. The FMS system will constantly monitor real-time energy usage and automatically turn unneeded equipment off (or reduce the control setpoint) to stave off peaking the facility’s electrical profile. Demand peaks set by commercial facilities are frowned upon heavily by utilities and are often accompanied by huge energy charge increases for upwards to 1 year.

Recently real-time pricing has furthered the ability to save energy. This allows a facility to proactively either use or curtail energy based on the price/KWH of the energy. Again, the HVAC subsystem takes the lead in this application. It can either poll the price structure from the Utility off the Internet, or the current pricing will be forwarded to the facility by the Utility. The HVAC subsystem can then automatically defer unneeded operation or temporarily reduce the cooling or lighting load as the cost warrants.

Chapter 3 Basic Network Design

Overview

The main function of the Facilities Management Network is to logically group critical building application services that are important for the proper functioning of the building automation systems and then isolate these application services from the enterprise or corporate network. These application services are responsible for the monitoring and control of subsystems (Fire, Lighting, Elevator, Security, and HVAC) over a common converged IP network. The number of devices that require IP connectivity in each of the above subsystems can range from a single device (Fire) to hundreds of devices (IP based video systems). The focus of this chapter is on basic Cisco network design principles and the networking of those IP enabled devices in each of the subsystems in the Facilities Management Network.

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.

Network Design Concepts

When Layer 2 VLAN switching technology was first introduced, it gained widespread popularity by achieving the ever-growing demand for high bandwidth aggregation and high-speed packet forwarding rates in Enterprise campus backbone networks. The Layer 3 switching devices such as routers were considered a bottleneck. The benefits of Layer 2 switches, for obvious reasons, evolved campus backbones over time as high-speed L2 networks and pushed the routers to the edge of the Enterprise campus network. These designs are often referred to or characterized as “flat” networks and they are most often based on the campus-wide VLAN model where a set of VLANs span the entirety of the network. This type of architecture favored the “departmental segmentation approach” where, for example, all finance traders needed to exist on the same broadcast domain to avoid crossing “slow” routers, or where old legacy applications dictated a Layer 2 network. Because these departments or applications could exist anywhere within the network, VLANs had to span the entire network.

The subsequent development of Layer 3 (and higher) switching provides the advantages of routing, with the added performance boost from packet forwarding handled by specialized hardware. The majority of campus networks now leverage this technology. Layer 3 switching in the Distribution Layer and backbone of the campus network (and the access layer to a lesser extent) allows segmentation of the campus into smaller, more manageable pieces. The benefit of this approach eliminates the need for campus-wide VLANs, allowing for the design and implementation of a far more scalable architecture. This approach is commonly referred to as the ‘multilayer’ approach and combines Layer 2 switching with Layer 3 switching to achieve robust, highly available campus networks.

Definitions

High Availability

High availability is a function of the application as well as the end-to-end network connectivity between a client workstation and a specific service. Deterministic network design is the major factor influencing network availability. The Mean Time Between Failures (MTBF) of individual components also needs consideration.

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) and by the timer-values of the protocols that are used in the network. Several key components and design concepts are examined here.

Hierarchy

This is a very general concept from which many features of the network can be derived. Hierarchy is a characterization of the traffic flows in a network. It implies that flows increase as they pass through points of aggregation (nodes) and tend to follow a specific direction or pattern. This is a direct consequence of client-server type of applications. Likewise, the network topology and equipment dimensioning will reflect the traffic flow hierarchy. This concept allows us to distribute the functions of each piece of network equipment in an optimal way through a layered hierarchy. Equipment within the same level of hierarchy will have similar properties and behave in a predictable way. With the help of such a classification, we can derive rules of thumb concerning the bandwidth required on each link or the backplane capacity needed on the network components. On the other hand, hierarchy imposes the way we use the network, where we place servers, how many users are within a single VLAN, where we put multicast sources etc. Many high level protocols (OSPF, PIM, etc) are hierarchical in nature and therefore are more easily implemented on a hierarchical network. Hierarchy is the base for many other network features; it leads to Scalability, Modularity and Predictability among others.

Scalability

This allows a network to grow considerably without making drastic changes or needing any redesign. It is a product of Hierarchy and Modularity.

Modularity

Modularity means that the network is made up of distinct building blocks, each having a precise set of features and behaviors. Its main advantage is when making changes in the network. Blocks can be added and removed without redesigning the network each time. Addressing is made much easier too. Modularity also means isolation; blocks are separated and interact through specific pathways thereby easing control and security. They are independent from each other, changes in one block does not affect other blocks.

Predictability

Once the network is built, many future decisions will need to be made whereby a precise knowledge of the traffic behaviors is required, for example when implementing Quality of Service (QoS) or when deciding how to provide for backup (fail over) scenarios. Hence, the network must be built such that traffic flows are easily depicted, delays are predictable within reasonable bounds, and fail over paths easily identifiable.

Fault-Tolerance

This aspect is hidden in the very definition of the term 'network' which generally implies a certain degree of meshing. An intelligent network relies on this property to provide redundant routes from one node to the other implying the ability to work around failures. If we have hierarchy, we don't need to provide redundancy between all points, as the network does not need to be a full mesh. Instead we'll be able to locate critical nodes where redundancy is important. Along with this, come features like fast-convergence, determinism etc.

Policy Domain

Access policy is usually defined on the routers or Layer 3 switches in the campus intranet. A convenient way to define policy is with ACLs that apply to an IP subnet. Thus a group of servers with similar access policy can be conveniently grouped together in the same IP subnet and the same VLAN. Other services such as DHCP are defined on an IP subnet basis.

IP Subnet

An IP subnet also maps to the Layer 2 switched domain; therefore, the IP subnet is the logical Layer 3 equivalent of the VLAN at Layer 2 (that is, one VLAN equals one subnet). The IP network address is defined at the Layer 3 switch where the Layer 2 switch domain terminates. By implementing a sensible IP addressing scheme, one offers Layer 3 switches the possibility to exchange summarized routing information, rather than learning the path to every host in the whole network. Summarization is key to the scalability of routing protocols.

In an ideal, highly structured design, one IP subnet maps to a single VLAN, which maps to a single switch in a wiring closet. This design model is somewhat restrictive, but pays huge dividends in simplicity and ease of troubleshooting.

Campus Design Solutions

Figure 3-1 below is an example of a hierarchical network design. It distributes networking functions at each level through layered organization. Modular designs are made out of building blocks. Modules can be added or removed without redesigning the network. A modular design is also easier to grow and troubleshoot. Cisco's Multi-Layer design is an example of modular hierarchical design model. The key elements of the structured hierarchy are the Core, Distribution and Access Layer in a network.

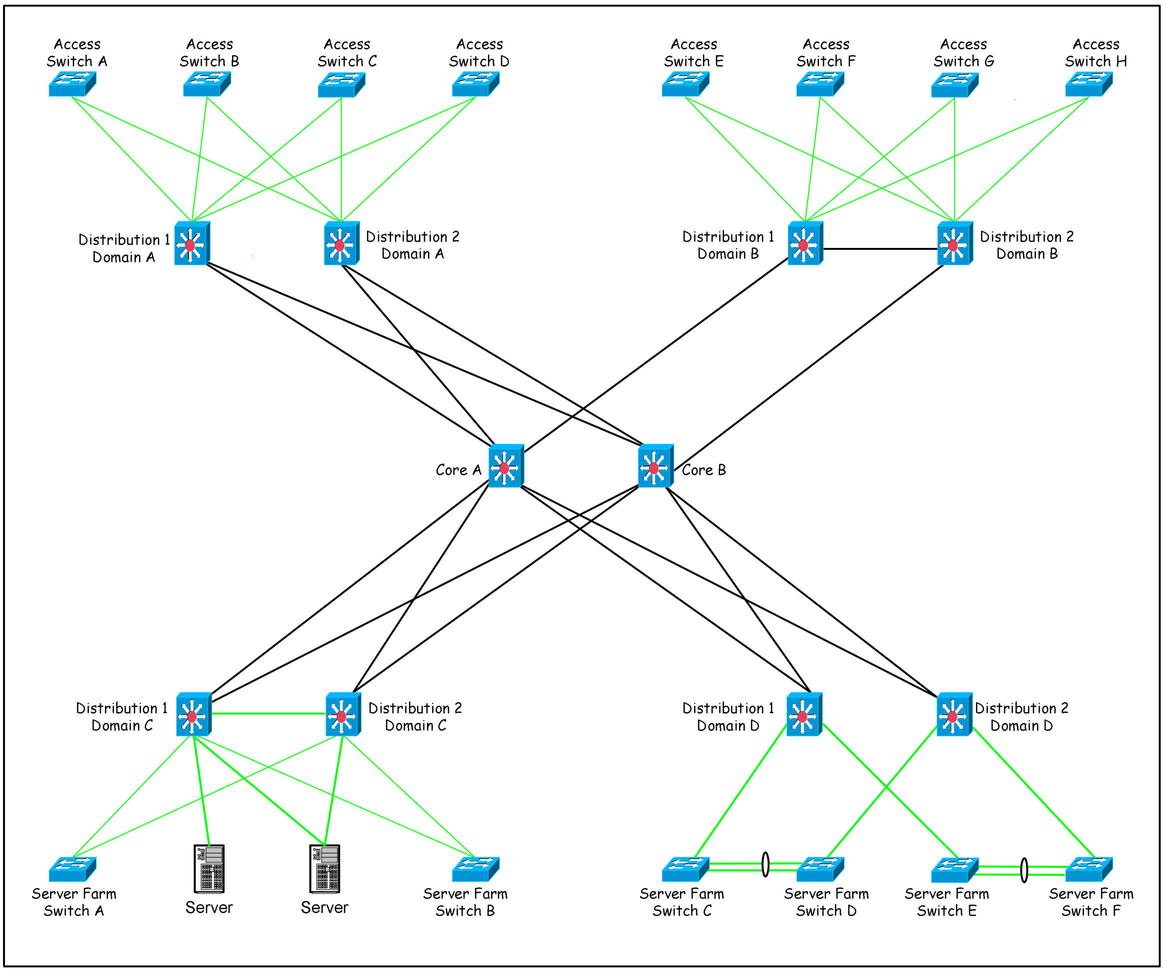

Figure 3-1 - The Multilayer Model

The Core Layer

This is the backbone of the network which aggregates the Distribution Layer switches and plays the primary role of connecting other network building blocks. The core provides redundant, Layer-3 paths for traffic traversing the network. The core switches should be as fast as possible and run as few complicated services as possible to maintain maximum efficiency and reliability for the network.

The Distribution Layer

The distribution layer is used for aggregating multiple access (or closet) switches, and sends their traffic to the core. The distribution layer is typically the demarcation point between the Layer-2 and Layer-3 domain in the campus. As such, equipment connected to the access layer depends on the distribution layer to act as a default gateway or to provide the necessary routing. The Distribution switches do not inject routing updates into the Access Layer under normal circumstances. Although the distribution consists of only two L3 switches, there are several paradigms for implementing and interconnecting them to the other Layers.

As represented in Figure 3-1, some distribution switch pairs have a link interconnecting the two switches together. This can be a source of confusion when designing networks; When to use the link? When is it required? What is its function? Should it be Layer 2 or Layer 3?

Uses of the Layer-3 Distribution Link

While a P-to-P Layer-3 routed link between switches can provide additional routing redundancy in certain topologies, it is required in some cases for the following reasons:

- HSRP tracking is not reliable in the event of specific failed conditions (i.e. the MSFC fails but the Layer-2 port remains active).
- Maintaining continuity for route summarization. This is especially true when implementing the VLAN in a box model and route summarization at the distribution layer.
- Accommodating single attached resources to the distribution (i.e. servers, WAN routers, etc.). The Layer-3 link ensures connectivity in a failed condition.
- Serving as a backup path for asymmetric routed flows within the access layer of the same distribution triggered by a failed condition.

Uses of the Layer-2 Distribution Link

The distribution layer, or in some cases the “collapsed Core-Distribution”, should be connected via a Layer-2 link in the following circumstances:

- Implementing building wide VLANs. The specific reasons for this are discussed in the next sections.
- Accommodating servers or hosts that are directly attached to the distribution.
- Providing connectivity for network components that utilize the U-shaped design paradigm. This might include routers and/or firewalls in DMZ networks.

The Access Layer

This Layer is typically made up of pure L2 switches each with uplinks to two L3 Distribution switches. There are several different design concepts for which to connect the access layer to the distribution. The four most common scenarios are discussed in some detail below. Note that all of these methods provide some form of redundancy and/or load balancing.

VLAN in a Box

The first model is used in Domains A and B of Figure 3-1. Note that the link between distribution switches in Domain B is an L3 routed link. This access model may also be referred to as the “V-shaped” design or even more commonly as the standard model - meaning that it should always be preferred and is considered to be a Cisco best practice design. The reason for this is that it does not rely on the spanning tree protocol for redundancy or convergence and therefore provides the fastest and most reliable fail over. It should be pointed out that while there are no L2 loops in this topology Cisco advocates that STP be enabled to help protect against any hardware, cabling or configuration mishaps.

In this model, distribution switches exchange HSRP hellos through the access layer. This is possible only because there is no layer 2 loop in the topology, therefore no blocking ports in the spanning tree. Any routing protocols in use with this model can typically run in passive mode on each access VLAN to prevent it from ever becoming a transit network. Configuring passive-interfaces on user facing VLANs also has the added benefit of reducing overhead associated with sending and receiving routing updates, processing hello messages and other maintenance or control plane traffic. The active HSRP router can also be configured to track the uplink(s) towards the core so that a switchover can be triggered in the event of a failure. This design does allow for multiple VLANs per closet as long as the VLANs do not span multiple access switches. A detailed representation of the standard HSRP model is depicted in Figure 3-2, below.

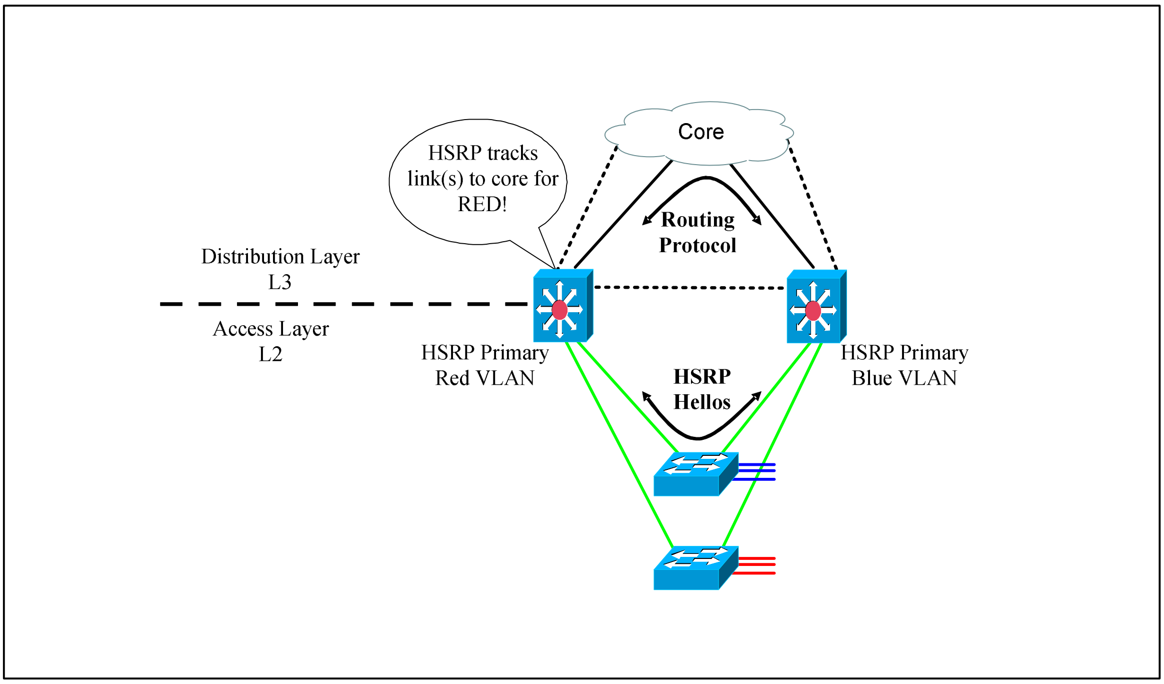

Figure 3-2: VLAN in a Box - The Standard Model

Note that the focus of this diagram is on the distribution and access layers. The minimum connectivity required from distribution to core is represented by the two solid black lines. The dotted black lines represent the additional connectivity required to complete either the U-shape or the dual home design. From the diagram above it is easy to see that a host sending unicast traffic to any other host on a different subnet will send the packets to its default gateway for proper forwarding, i.e. the HSRP primary router for that VLAN. However since the core is using a routing protocol instead of HSRP the path back to that VLAN, and more specifically to the MAC address of that originating host, may or may not be through the same set of components and links, that is the path back to VLAN RED may be through the opposite distribution switch. The conclusion is that using HSRP in conjunction with a hierarchical redundant network topology (i.e. equal cost L3 paths) lends itself to asymmetric routing.

VLAN in 2-Box

There are several reasons why restricting a VLAN or set of VLANs to a single access switch may not meet the requirements of certain networks, to include:

- Dual homing server NICs that require presence on the same L3 subnet
- Supporting applications that rely on L2 mechanisms for redundancy and failover (i.e. keepalives or heartbeat protocols)
- Sheer port density on a floor or within a specific user domain
- Cabling and infrastructure limitations

All of these reasons can pose significant challenges when designing campus networks and may very well prevent a network administrator from implementing the standard HSRP model throughout their environment. The VLAN in 2-Box model is a viable design option and may be preferred to the Building Wide VLAN model for those administrators who have made the effort, painfully so in some cases, to reduce their reliance on the spanning tree protocol. The VLAN in 2-Box model, also referred to as the U-shape or Horseshoe design, is depicted in Figure 3-3 below.

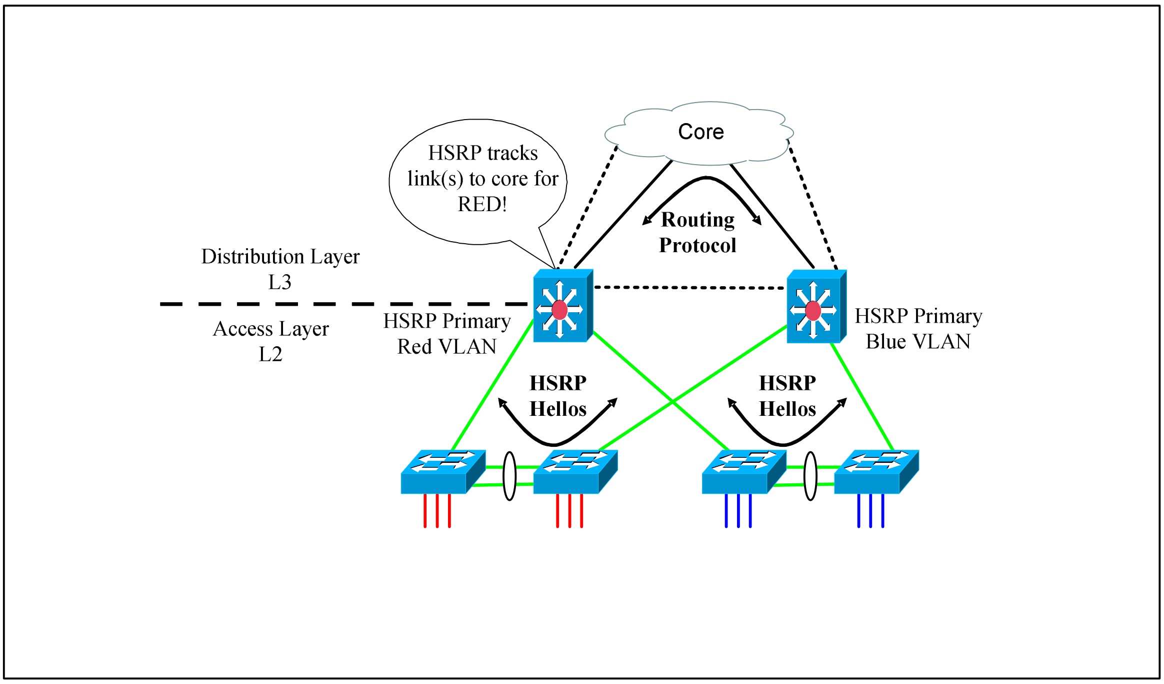

Figure 3-3: VLAN in 2-Box

Notice that there is still no L2 loop in this design and therefore it functionally resembles the VLAN in a Box model. Specifically, distribution switches still exchange HSRP hellos through the access layer, routing protocols operate in passive-mode facing user subnets, HSRP tracking can still be used to enhance failover and the two access switches can support multiple user VLANs. However, since each access switch has only a single connection to one distribution switch there is an opportunity for a discontiguous subnet upon a link failure between the two access switches. Etherchannel is used to mitigate the risk associated with this failure scenario. Optimally, the connections between the access switches are striped across different modules to achieve the necessary hardware and link redundancy. Note again that since HSRP is used in the access layer and equal cost paths are available back to those user subnets from the core, the opportunity for asymmetric routing still exists.

Building Wide VLANs

Considerations should always be made for simplifying network topologies and design wherever possible. This includes challenging the requirements put forth by specific applications and their administrators dictating that L2 services be extended throughout parts or all of a campus. Many times the advantages of L2 network topologies to that of implementing L3 services are merely perception and can be overcome by using good design practices. Still there are many networks that do not or can not employ the design paradigms detailed in the previous two sections. Network topologies that allow VLANs to span multiple access switches rely heavily on spanning tree for redundancy and convergence. This model is commonly known as the Building Wide VLAN model or simply the spanning tree model. A spanning tree topology can be well controlled and optimized by adhering to best practice configuration guidelines in conjunction with implementing features such as portfast, uplinkfast and VLAN pruning. Those topics are discussed later in this chapter.

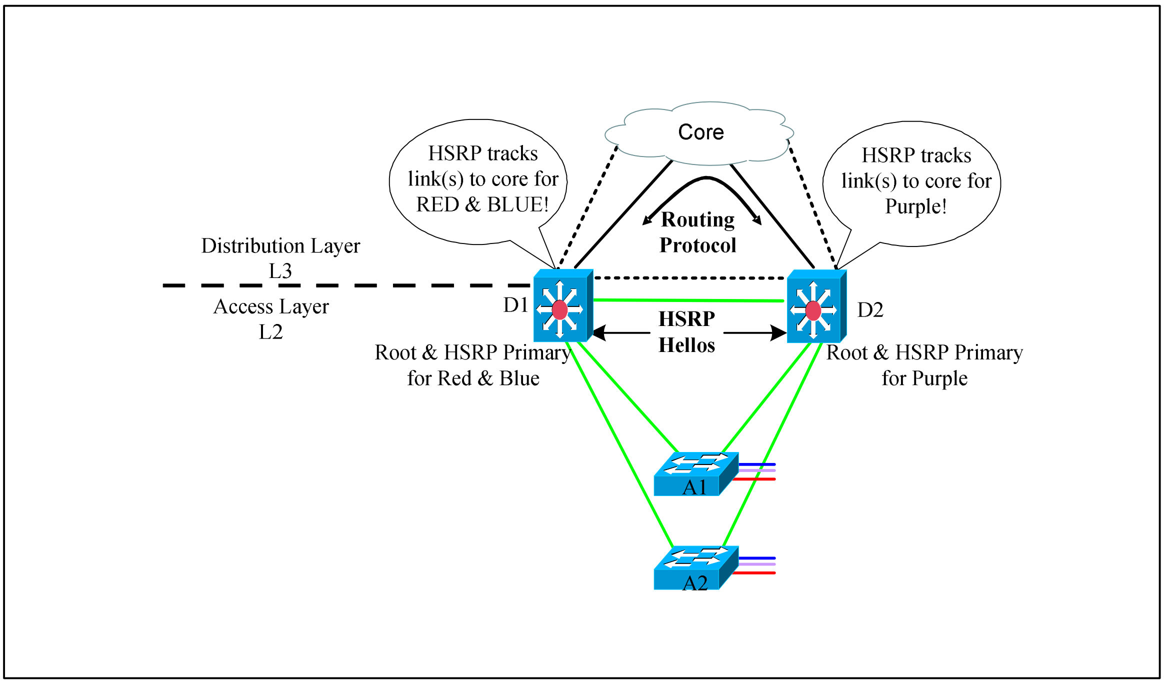
Core
HSRP tracks link(s) to core for RED & BLUE!
HSRP tracks link(s) to core for Purple!
Routing Protocol
Distribution Layer L3
Access Layer L2
D1
D2
HSRP Hellos
Root & HSRP Primary for Red & Blue
Root & HSRP Primary for Purple
A1
A2

Figure 3-4: Building Wide VLAN Model

Considering the topology in Figure 3-4, one can see two differences to that of the VLAN in a box design; a) an additional L2 trunk link is added between the two distribution switches and b) access VLANs are used on more than one switch. As part of this design it is recommended that the two distribution switches be configured as the spanning tree root or backup root bridge for all VLANs in this domain. One can also rotate root and backup root switch placement between the two distributions, i.e. odd and even VLANs. This allows the outbound traffic to spread across both distribution switches, although classifying this as load balancing is imperfect. Additionally, the HSRP primary router should match the root bridge for its respective VLAN to provide an optimum L2 path to the default gateway.

While the additional L2 link connecting the distribution switches is not absolutely required it does have several advantages in this design. First, it provides a more predictable topology which simplifies troubleshooting. Assuming that D1 and D2 are configured per the diagram, the L2 blocked port for any given VLAN is always on the access switch facing the backup root switch, assuming normal operation. This is not the case without that L2 link in place since one port from an access switch towards the backup root will have to forward packets. This is true for each VLAN and becomes more difficult as the number of access switches and VLANs increase. Having the L2 distribution link in place also has the added benefit of allowing control plane traffic to travel directly from distribution to distribution (HSRP hellos, routing protocol updates, STP BPDU's, etc.) as opposed to flowing through the access switch. However, it should be noted that this is perfectly viable in the event of a failure.

The biggest difference of the spanning tree model to that of the VLAN in a Box or VLAN in 2-Box models with respect to unknown unicast flooding and asymmetric routing is the number of ports that terminate at the distribution for each VLAN; the spanning tree model has more than one where the other two models have precisely one.

Layer 3 in the Access

Another option for consideration is to extend the routing domain right to the access layer. Using a routed access layer configuration presents equal cost L3 paths at the access towards the campus infrastructure. It too supports multiple VLANs in each access switch but requires, as does the VLAN in a Box model, that the set of VLANs be unique to that switch. This design may provide several advantages over that of the solutions discussed in the previous sections, to include:

- Decreased convergence times; re-routing around a link failure is based purely on L3 protocols.
- Load balancing across uplinks is based on L3 routing algorithms (CEF/OSPF/EIGRP) and not based on HSRP default gateway or Spanning Tree Root Bridge placement.
- Provides better isolation of broadcast domains further mitigating the possibility of L2 issues proliferating beyond the problem link or access switch.

These advantages aside, this design paradigm is not widely deployed. The primary reason for this is cost. Implementing an L3 access layer also imposes additional IP subnetting requirements to support the uplinks to the distribution. While the advantages noted above are genuine, one has to question the benefit of deploying L3 engines in the access layer to that of the VLAN in a Box model. After all, the prerequisites for implementing both of these models are similar amongst the two. Specifically, both models require that VLANs are unique to the access switch. The VLAN in a Box design is tried and proven without the additional routing hardware in the access layer. Still the L3 access design offers the fastest convergence and the highest level of isolation from Spanning Tree related issues.

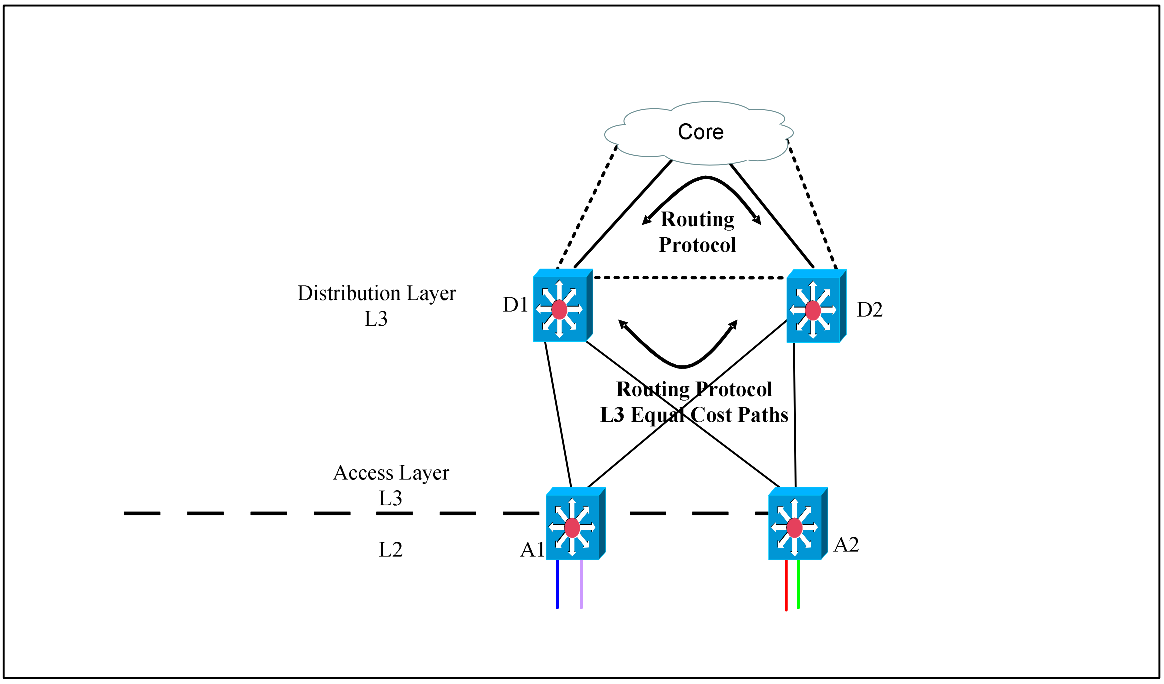

Figure 3-5: Layer 3 Access Model

The Server Farm

A Server farm is implemented as a high-capacity building block attached to the campus backbone, and is treated as its own Distribution block. A server farm is an aggregation point for much of the traffic from the whole campus. As such, it makes sense to design the server farm with lower over subscription ratios to that of a user building block; this includes switching backplane capacity and available bandwidth (i.e. trunks, uplinks and host ports).

Campus Network Design Considerations

Failure Domain

A Layer 2 switched domain is considered to be a failure domain because a misconfigured or malfunctioning workstation can introduce errors that impact or disable the entire domain. A jabbering network interface card (NIC) might flood the entire domain with broadcasts or undesirable frames at a very high rate. A protocol malfunction (for example, spanning-tree error or misconfiguration) can inhibit a large part of the network. Problems of this nature can be very difficult to localize.

The scope of a failure domain should therefore be reduced as much as possible. The best way to achieve this is by restricting its scope to a single Layer 2 switch in one wiring closet. In other words, only one or a few unique VLANs should exist per wiring closet switch. Typically, there is one VLAN for user data traffic, one VLAN for voice and video over IP, and possibly another one reserved for switch management. The keyword "unique" means that these VLANs should not span multiple Access Layer switches (Access Layer switches are defined in subsequent sections).

To implement this type of architecture, the deployment of VLAN trunking should be tightly controlled, and only the necessary VLANs are allowed on any given trunk. If no voice or video VLAN is necessary, then ideally only one VLAN (IP subnet) should exist in a single wiring-closet switch. This eliminates the requirement for trunking on the Gigabit uplinks from each wiring-closet switch and allows direct connection to routed interfaces on the Layer 3 switches.

Broadcast Domain

MAC-layer broadcast, multicast, and unknown unicast packets flood throughout the Layer 2 switched domain. Implementing network segmentation by utilizing Layer 3 switching in a structured design will help to reduce the scope of broadcast domains. In addition, intelligent, protocol-aware features of Layer 3 switches further contain broadcast packets such as DHCP by converting them into directed unicast packets as appropriate. Flooding of multicast traffic can be constrained to a set of interested ports by using IGMP snooping or the Cisco Group Membership Protocol (CGMP).

BBMD – BACnet Broadcast Manager Device

Broadcast domain management in a BACnet IP network is accomplished through the capabilities of a single device called a BACnet Broadcast Management Device (BBMD). One of the Cisco campus network design models states that an FMS network be composed of more than one IP subnet so that large campus-like installations can be deployed simply while maintaining high availability and resiliency. Many of BACnet's capabilities, such as dynamic name binding and unsolicitied change-of-value notification, stem from the use of broadcast messages, so there needs to be a way to support them in a hierarchical, distributed network.

Very simply, BBMDs receive broadcast messages on one subnet and forward them to another subnet. Each BACnet IP network comprised of two or more subnets shall have one device per subnet configured as a BBMD. Each BBMD shall possess a table called a Broadcast Distribution Table (BDT) which shall be the same in every BBMD in a given BACnet IP network. If the BBMD has also been designated to register foreign devices, it shall also possess a Foreign Device Table (FDT).

There are two ways that a BBMD may distribute broadcast messages to remote IP subnets. The first is to use IP "directed broadcasts" (also called "one-hop" distribution). This involves sending the message using a BACnet IP address in which the network portion of the address contains the subnet of the destination IP subnet and the host portion of the address contains all 1's. While this method of distribution is efficient, it requires that the IP router serving the destination subnet be configured to support the passage of such directed broadcasts.

Since not all IP routers are configured to pass directed broadcasts, a BBMD may be configured to send a directed message to the BBMD on the remote subnet ("two-hop" distribution) which then transmits it using the BACnet IP broadcast address. Since the use of one-hop distribution requires an IP router configuration that may or may not be possible, while the two-hop method is always available, the choice of which method to use in any given case is a local matter.

Spanning-Tree Protocol

IEEE 802.1d Spanning-Tree Protocol (STP) is used to prevent Layer 2 loops in the network. If loops are present in the Layer 2 design, then redundant links are put in "blocked" state and do not forward traffic. Well-designed campus networks rely as little as possible on STP to provide load balancing and link resiliency. Instead, the use of loop-free Layer 3 -based topologies is favored as much as possible so that all links actively carry traffic. STP is still left enabled to ensure that mis-patched cables etc. do not introduce loops that can't be recovered from.

With Layer 2 topologies that have loops inherently contained, the default STP convergence times are between 30 and 50 seconds minimally. Avoiding Layer 2 loops is especially important in the mission-critical parts of the network such as the campus backbone. To prevent STP reconfiguration events in the campus backbone, all links between core and distribution switches should be point-to-point routed links with only one unique VLAN defined per link. These links should not be VLAN trunks. Using Layer 3 links also constrains the broadcast and failure domains, as explained previously. Where possible the Interior Gateway Protocol should handle load balancing, redundancy, and fault recovery.

Cisco supports the following 802.1D IEEE specifications:

- 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

Virtual LAN

A VLAN is essentially an extended Layer 2 switched domain; that is, a broadcast domain that extends as far as the VLAN reaches. If several VLANs coexist across a set of Layer 2 switches, each individual VLAN has the same characteristics of a failure domain, broadcast domain, and spanning-tree domain as described previously. Therefore, although VLANs can be used to segment the campus network logically, deploying pervasive VLANs throughout the campus introduces complexity and reduces the deterministic behavior of the network. Avoiding loops and restricting a set of unique VLANs to a single Layer 2 switch in one wiring closet minimizes the complexity.

One of the goals of VLAN technology is to take advantage of high-speed Layer 2 switching. With the advent of high-performance Layer 3 (and beyond) switching in hardware, the use of VLANs is no longer related to performance. VLANs are best used for implementing policy. A VLAN can be used to logically associate a workgroup with a common access policy as defined by access control lists (ACLs). Similarly, VLANs can be used within a server farms to associate a group of servers with a common access policy as defined by ACLs.

Each of the subsystem networks in an FMS should include only traffic that is relevant to running that particular subsystem. For this reason, the recommendation is to logically segment traffic with the use of VLANs. As shown in Figure 3-6 below, one VLAN is used for all data traffic relevant to that particular subsystem. The Cisco Catalyst switch in the access layer aggregates all the VLANs in the zone area and terminates those VLANs at the core or distribution layer, if it exists. Because 80–90 percent of traffic is local to one zone, this is the optimal design for this scenario. Depending on the size of the network and the specific application requirements, additional VLANs per subsystem may be recommended to comply with the various Cisco campus network design models discussed earlier in this chapter.

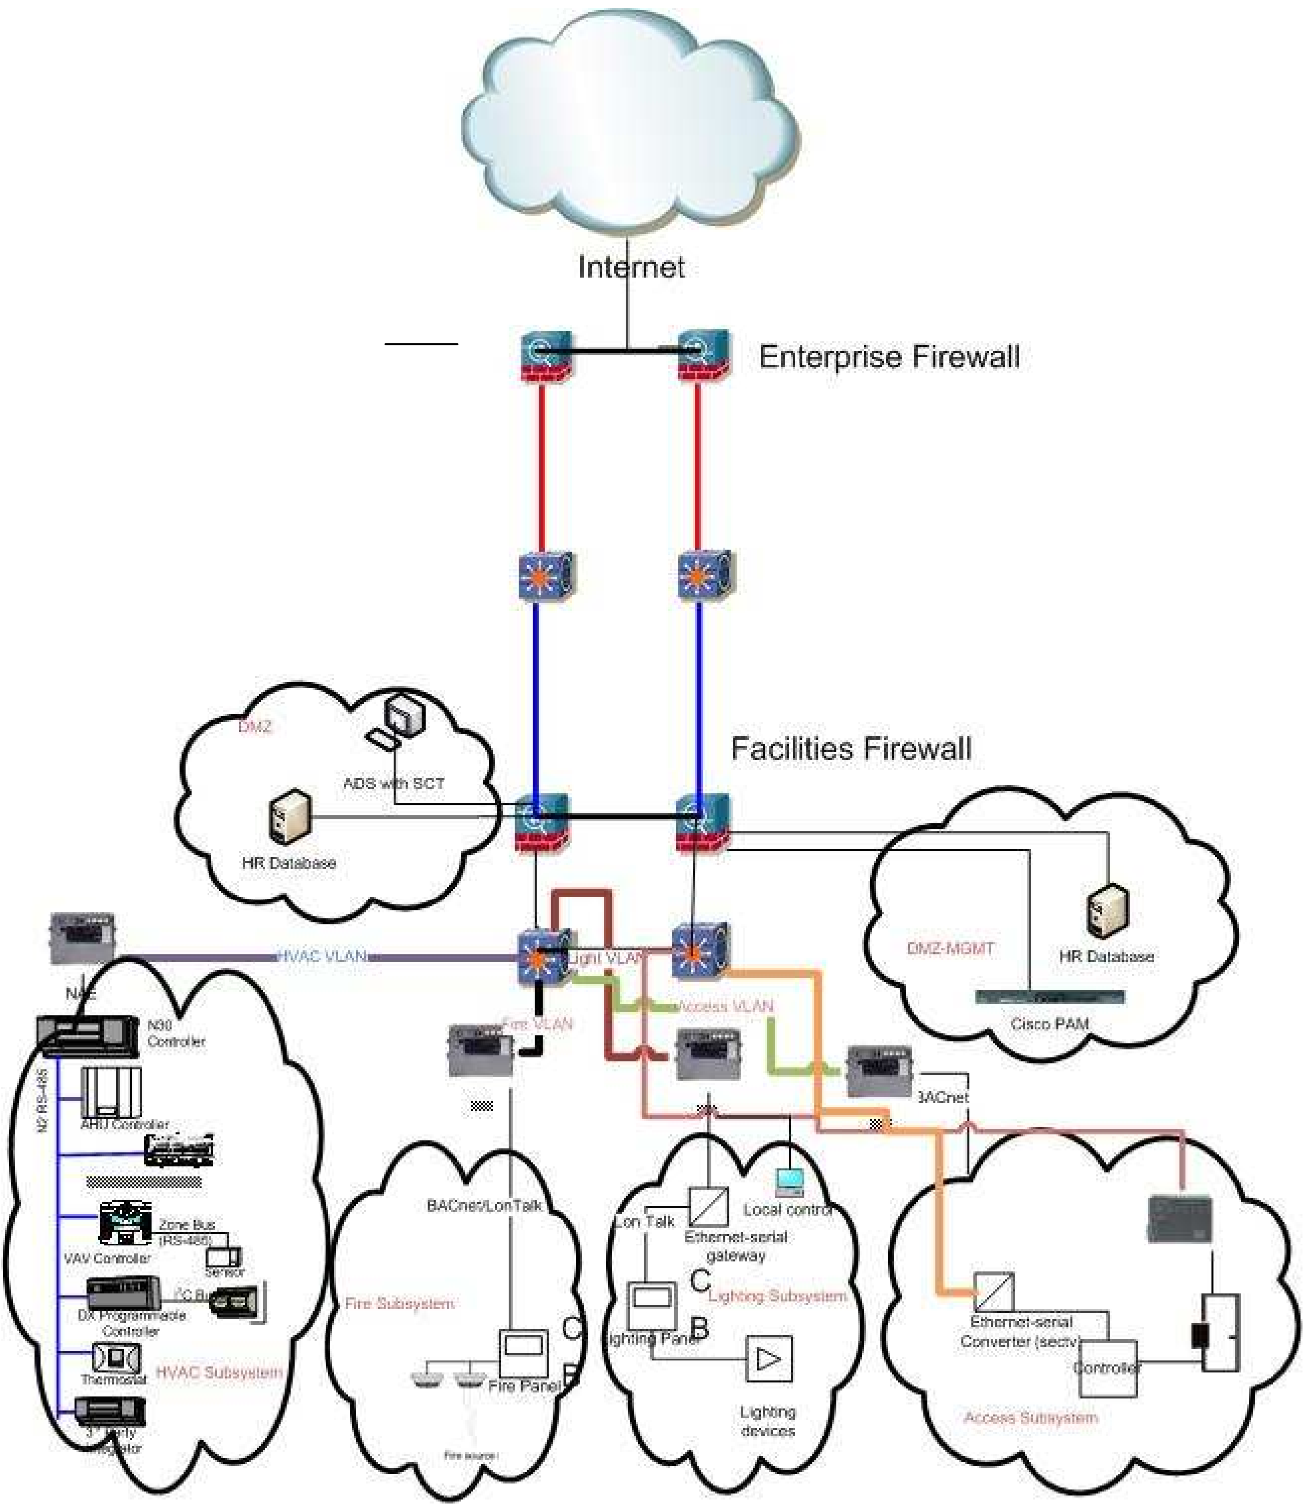

Figure 3-6 - VLAN Segmentation for BAS Subsystems

Asymmetric Traffic Flows

As eluded earlier in this document, using HSRP first hop redundancy in conjunction with equal cost path routing (ECPR) lends itself to asymmetric traffic flows within the network. While asymmetric routing is not necessarily undesirable, in fact using HSRP and ECPR is a best practice recommendation, it may present some symptoms for which network administrators must be aware.

LAN switches use CAM tables to direct traffic to specific ports based on the VLAN number and the destination MAC address of the frame. When there is no entry corresponding to the frame's destination MAC address in the incoming VLAN, the (unicast) frame will be sent to all forwarding ports within the respective VLAN. This is known as an unknown unicast flood. While limited flooding is part of the normal switching process, excessive flooding can cause adverse performance effects on the network, to include: saturating lower bandwidth links, consuming host or network resources (CPU, memory, etc.), or even limiting a stations ability to transmit data (especially true of half duplex connections).

As discussed previously, a host sending unicast traffic will always send packets to its default gateway (HSRP primary). Since most networks do employ the use of equal cost L3 paths for redundancy, the path back to that host may be through the alternate distribution switch (the HSRP secondary). By default, a switch’s CAM table will age out in 5 minutes. Since a router’s ARP table will not age out for four hours (the default configuration) it is easy to determine that after 5 minutes the “non-primary” distribution router will have an ARP entry for a given MAC address but the switch in which it resides will have no CAM entry in its table for the destination. Keep in mind that CAM tables are built solely on source MAC address. The result presents itself as an unknown unicast frame to the switch which will be flooded to all ports in that VLAN.

While there are several variations of this problem, the above describes the most basic behavior. As with many problems there are always different means for resolution - and network administrators may argue their point of preference. For the purposes of this document only the most common method for resolving this issue is provided. Lowering the routers ARP timeout such that it is equal to that of the switch’s CAM aging timer (5 minutes by default) is an effective and simple solution. The implications of doing so is well tested and widely deployed.

IP Addressing of Devices in the Subsystems

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 many of the subsystems, the level 3 workstations and NAE servers (which are used for control as described below) are static. These NAE servers send scheduling, execution, and control data to controllers in the FMS, 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 FMN. For more information on IP addressing, see *IP Addressing and Subnetting for New Users* at the following URL: http://www.cisco.com/en/US/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 3-7)

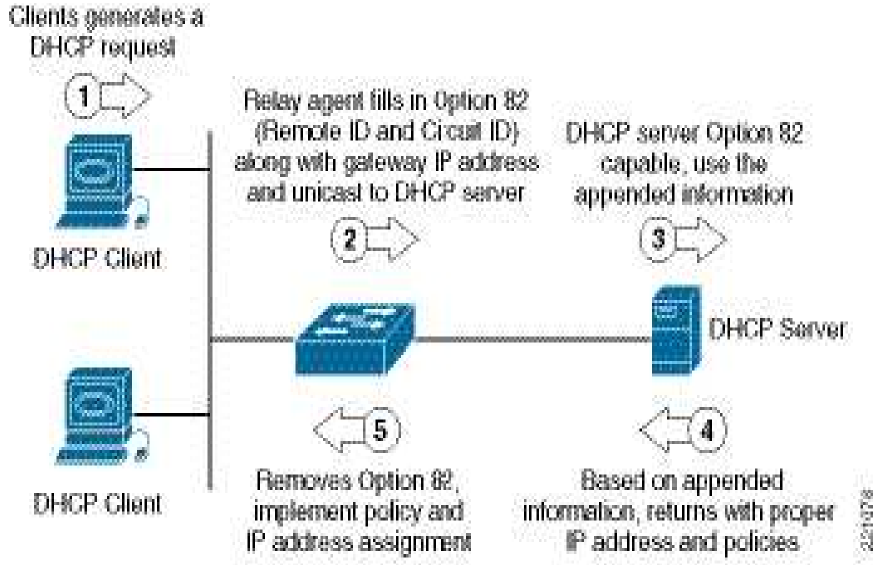
Clients generates a DHCP request
1
DHCP Client
DHCP Client
Relay agent fills in Option 82 (Remote ID and Circuit ID) along with gateway IP address and unicast to DHCP server
2
DHCP server Option 82 capable, use the appended information
3
DHCP Server
5
Removes Option 82, implement policy and IP address assignment
4
Based on appended information, returns with proper IP address and policies
221078

Figure 3-7 - 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. In subsystems where DHCP is required (such as video and voice networks where there are large number of devices), Cisco recommends DHCP option 82 for finer control over IP address assignment.

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 FMN 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.

For the building automation traffic that is primarily confined to the FMS 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 facilities operation. For more information on private IP addresses, see RFC 1918 at the following URL: http://www.ietf.org/rfc/rfc1918.txt.

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 FMS, routing begins at the subsystems. The Cisco Catalyst switches in the distribution or core layer, depending on the campus design model, are responsible for routing traffic between subsystems (inter-VLANs), into the core, or through the DMZ. No routing occurs in a particular subsystem itself unless the campus network design model deployed dictates otherwise.

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 service area 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 service area, 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 below 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
IGRP	Distance vector	Yes	Interior	90 sec	Composite	No	Auto
EIGRP	Advance d 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, a service area 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 service area.

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 Considerations

Types of Servers

The servers used in the service area 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 date about end users and network resources.
  - Network Time Protocol (NTP)—Synchronizes the time on a network of machines. NTP runs over UDP, using port 123 as both the source and destination, 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/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/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/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/products/sw/cscowork/ps2425/tsd_products_support_series_home.html

- Application servers—Consists of the following:
  - NAE/NIE
  - ADS/ADX

Chapter 4 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.

The following attacks are considered within the scope of this document and are addressed herein as the security provisioning known as the **Secure Architecture for Intelligent Facility Applications** (“**SAIFA v1.0**”)

- 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 (Telnet, File Transfer Protocol [FTP], Simple Message Transfer Protocol [SMTP], Post Office Protocol [POP3], and so on), and building control applications (BACnet, LonTalk and so on) send data in a binary encoded format. A packet sniffer can provide meaningful and often sensitive information, such as usernames and passwords.
- 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. BACnet used UDP port and has no authentication or integrity built into the protocol so it is trivial for an attacker to spoof BACnet messages.
- 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 can overwhelm a specific critical device (such as a NAE) 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.
- 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 IP based building control devices. This scenario can lead to specific information that is useful when the hacker attempts to compromise that service.
- Unauthorized access—Unauthorized access refers to an user being able to access a system (run applications, run specific commands, send uninteded packets and so on) that he should not or need not have access to. Although unauthorized access attacks are not a specific type of attack, they refer to most attacks executed in networks today.
- 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. Viruses typically scan for vulnerable hosts sending a flood of packets inside the network, which might cause un-intended consequences on the building automation network.

- 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 SAIFA 1.0 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 4-1 summarizes the device hardening techniques needed for the platforms supported by the SAIFA 1.0 solution. The detailed configuration is presented in the following sections.

Table 4-1 - 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
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/cdccont_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 **enadle** instead of **enable** results in a long timeout while the router tries to find host “enadle” 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. 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-mgmnt-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 SAIFA 1.0 solution, 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/fcfprt3/fcf014.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
```

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 SAIFA 1.0, the implementation of Layer 2 security protection is needed on all switches (both access and distribution) in the following network areas:

- Building automation subsystems
- Server farm
- Server farm in the DMZ

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 FMS 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
```

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 bpdufilter enable
Switch(config-if)#spanning-tree bpduguard 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
- 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.

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
```

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 VLLAN.

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 DHCPOffer, DHCPAck, or DHCPNak.

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 **Error! Reference source not found.** 4-14.

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 Building Automation Subsystem

Because the security design strategy of any of the sub system 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 Building Subsystem Networks and the Server Farm

Note the following:

- Device hardening (see Device Hardening 4-3)
- Layer 2 security for L2 ports (see Layer 2 Security Design 4-11)
- 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 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 4-35

Security Design for the Segmentation of Facilities Network: Facilities Firewall

In the design of the FMS network, one of the critical elements is to ensure the separation between the facilities management network and enterprise network. This separation is necessary because real-time availability and security are the critical elements for the traffic in the FMN. You do not want enterprise traffic that has very different traffic characteristics to enter the facilities network and cause any disruption to the ongoing operations. Acting as a firewall, the Cisco ASA5500 provides this separation of the two networks. We call this the facilities firewall and this is different and in addition to the corporate firewall which separates enterprise or corporate from the Internet.

The facilities firewall should be placed in the “inside interface” of the corporate firewall. However under certain circumstances if partners need access to facilities network, it can be placed inside the DMZ of the corporate firewall for a short period of time. Figure 4-1 below represents the two ways, notice the difference in the security levels of the corporate firewall in the two placements.

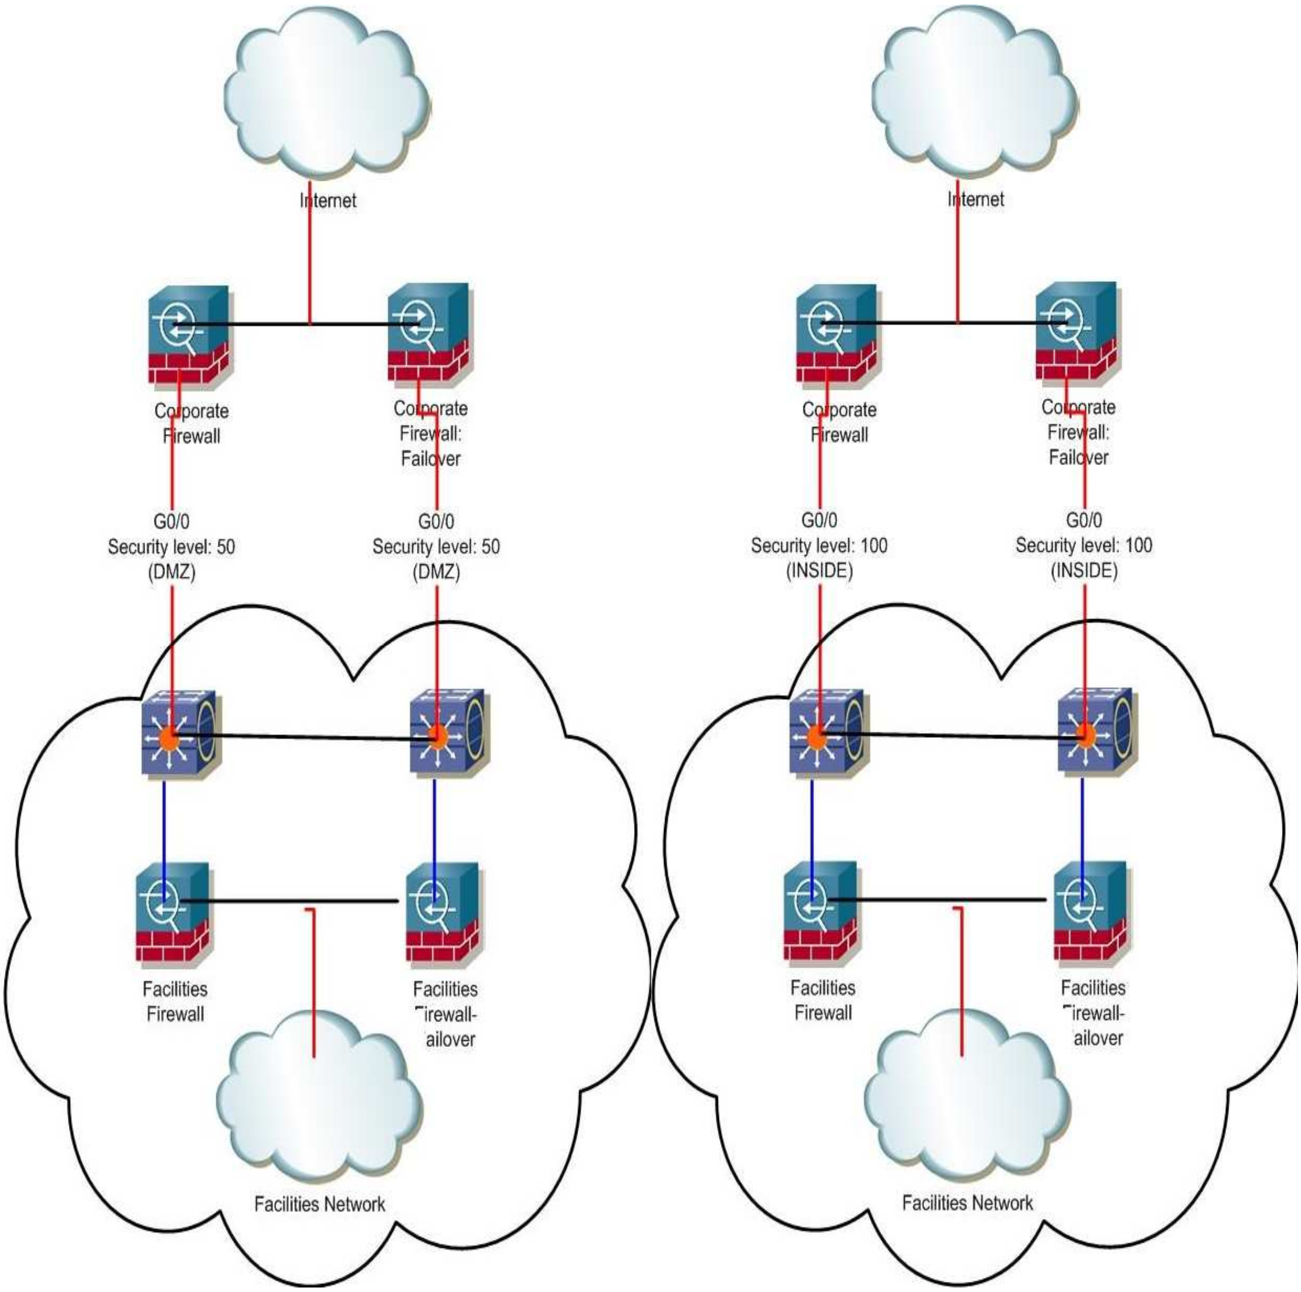

Figure 4-1 - Facilities Firewalls

Security Design for the Demilitarized Zone

IServers 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 4-18 .

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 DIG, Cisco recommends creating three networks in different security levels, as shown in Table 4-2.

Table 4-2 - Network Security Levels

Network Security Levels

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

Note on Security Levels

Each interface must have a security level in the range 0 to 100 (from lowest to highest). For example, you should assign your most secure network, such as the inside business network, to level 100. The outside network connected to the Internet can be level 0. Other networks can be in between. Also interfaces can be assign to the same security level if required.

The level controls the following behavior:

•Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an access list to the interface.

For same security interfaces, there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

•Inspection engines—Some application inspection engines are dependent on the security level. For same security interfaces, inspection engines apply to traffic in either direction.

•NAT control—When you enable NAT control, you must configure NAT for hosts on a higher security interface (inside) when they access hosts on a lower security interface (outside).

Without NAT control, or for same security interfaces, you can choose to use NAT between any interface, or you can choose not to use NAT. Keep in mind that configuring NAT for an outside

Configuration Example

Refer to Figure 4-2 below for the subsequent configuration example.

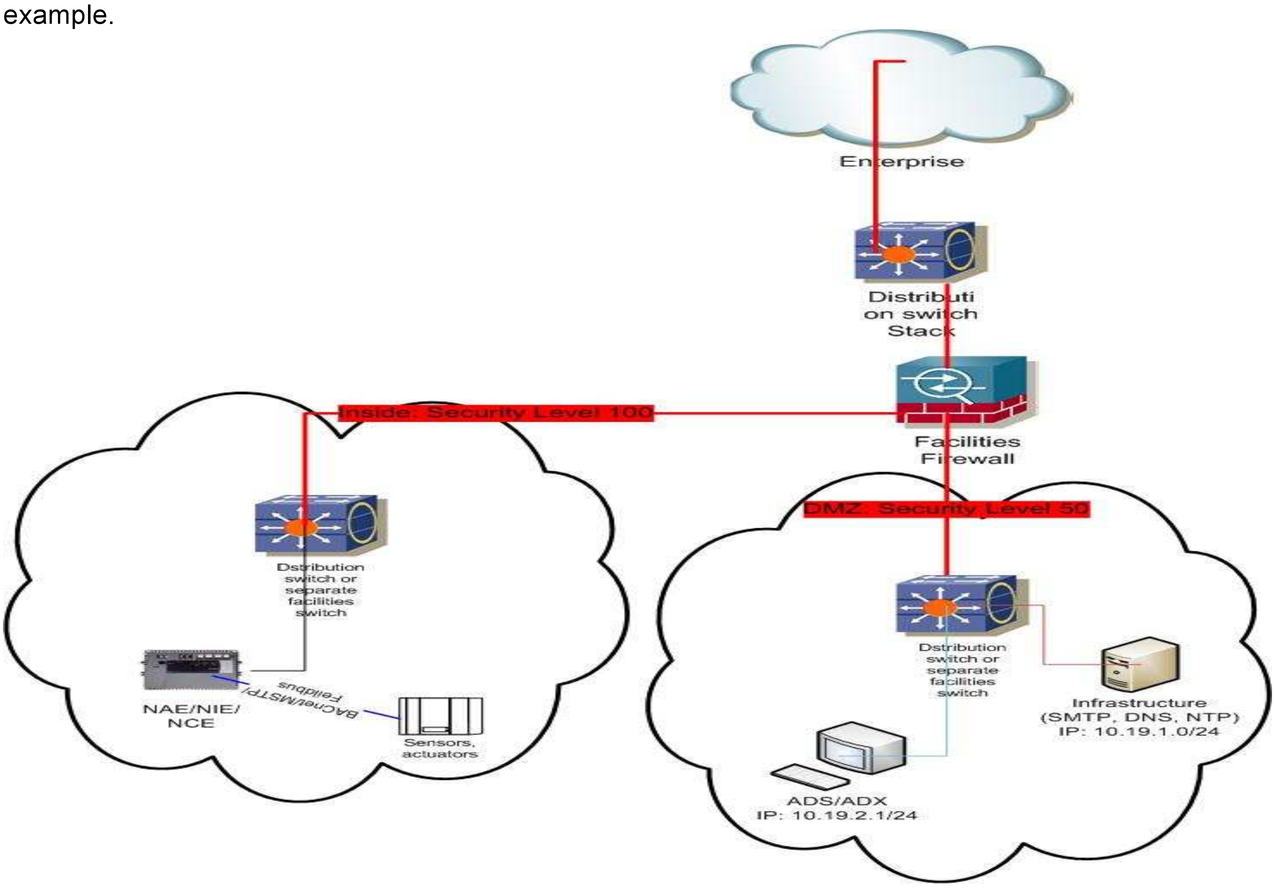

Enterprise
Distribution switch Stack
Facilities Firewall
Inside: Security Level 100
DMZ: Security Level 50
Dstribution switch or separate facilities switch
NAE/NIE/ NCE
BACnet/MSTP/ Feildbus
Sensors, actuators
Dstribution switch or separate facilities switch
Infrastructure (SMTP, DNS, NTP) IP: 10.19.1.0/24
ADS/ADX IP: 10.19.2.1/24

Figure 4-2 – Network Security Configuration Example

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 x.x.x.x 255.255.255.0
  ```

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

  ```
  interface GigabitEthernet0/1
  nameif outside
  security-level 0
  ip address x.x.x.x 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 x.x.x.x 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 DIG recommends that traffic be denied going from the facility 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 facility 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 SAIFA 1.0 solution, general packet filtering recommendations are listed below in Table 4-3.

Table 4-3 - Packet Filtering Recommendations

		Traffic Source		
		Enterprise Network	**DMZ**	**Facilities Management 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
	Facilities Management Network	Disallowed (implicitly denied by ACLs)	Explicitly permitted by ACLs	N/A

High-Level Packet Filtering Recommendations for the DMZ between the Facilities and Enterprise Networks

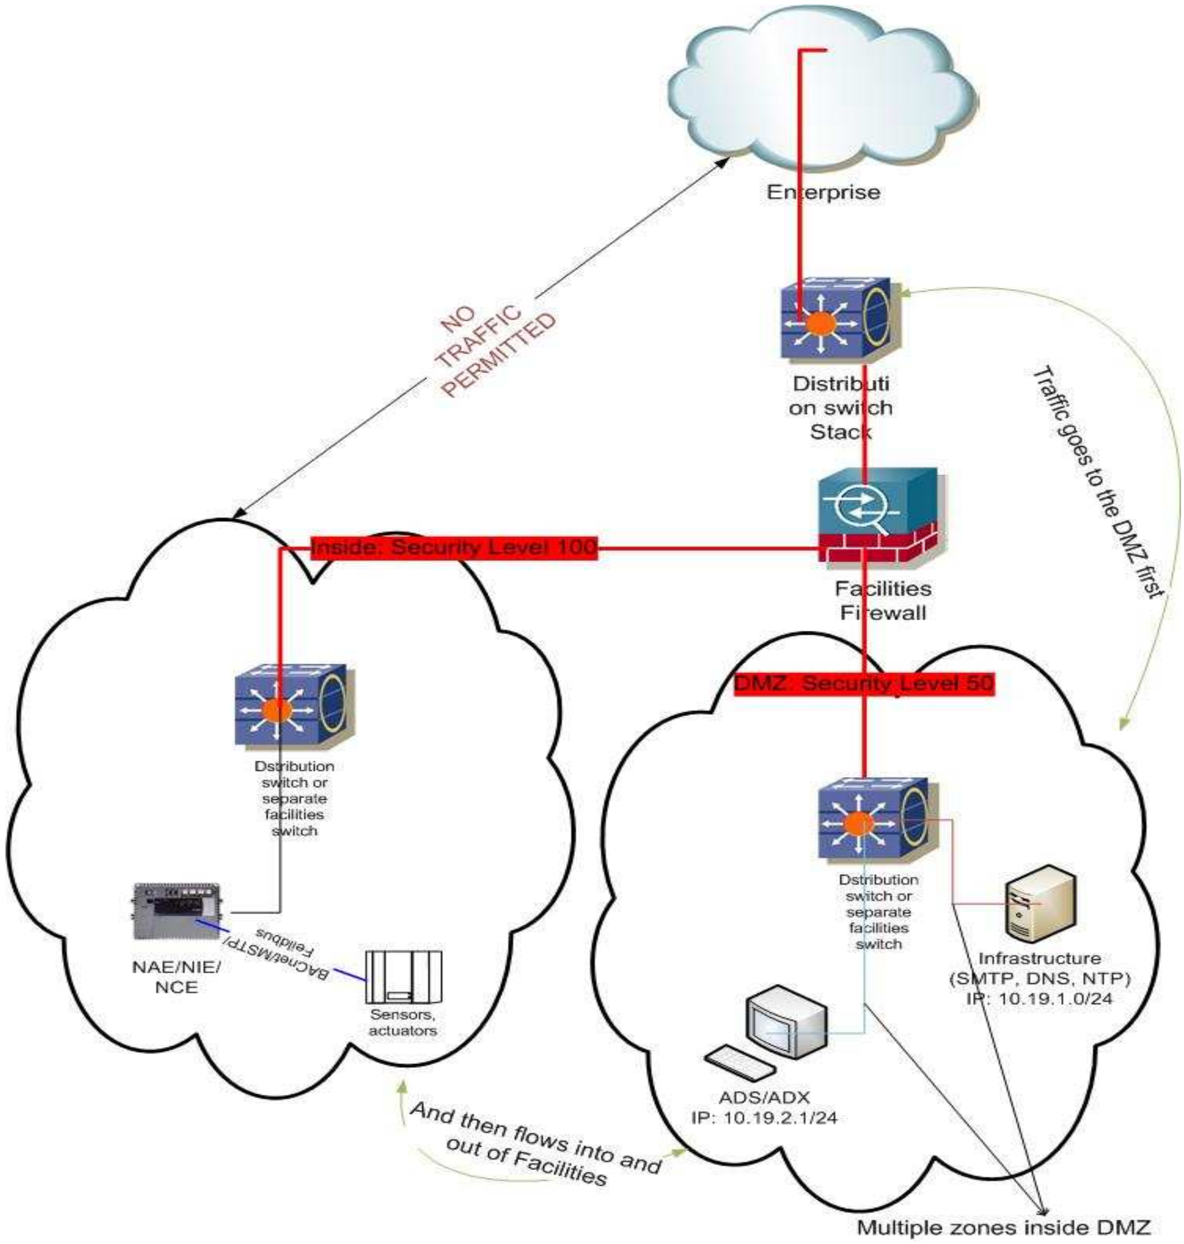
Enterprise
NO TRAFFIC PERMITTED
Distribution switch Stack
Traffic goes to the DMZ first
Facilities Firewall
Inside: Security Level 100
DMZ: Security Level 50
Dstribution switch or separate facilities switch
NAE/NIE/NCE
BACnet/MSTP/ Fieldbus
Sensors, actuators
Dstribution switch or separate facilities switch
Infrastructure (SMTP, DNS, NTP) IP: 10.19.1.0/24
ADS/ADX IP: 10.19.2.1/24
And then flows into and out of Facilities
Multiple zones inside DMZ

Figure 4-3 - High-Level Packet Filtering Recommendations for the DMZ between the Facilities and Enterprise Networks

Configuration Example

See Table 4-4 below for an example for ingress ACLs applied to the facilities network-facing interface.

Table 4-4 – Configuration Example for Ingress ACLs on the Facilities Networking-Facing Interface

Applied To Interface	Traffic Direction	Permitted Traffic Types (Source to Destination)
Interface connected to the facilities network (*inside*)	Inbound	• HTTP (servers in the facilities network to servers in DMZ such as ADS/ADX) access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq www • HTTPS (any in the facilities network to servers in DMZ) access-list inside extended permit tcp 10.18.0.0 255.255.0.0 10.19.0.0 255.255.255.0 eq https • Telnet (any in the facilities network to host 10.19.1.10 in the DMZ) access-list inside extended permit tcp 10.18.0.0 255.255.0.0 host 10.19.2.1 eq telnet • ICMP (any in the facilities network to servers in the DMZ) access-list inside extended permit icmp 10.18.0.0 255.255.0.0 10.19.2.0 255.255.255.0 • Explicitly deny other traffic types to anywhere (i.e. DMZ and enterprise networks) access-list inside deny 10.18.0.0 255.255.0.0 • Apply the ACLs above to the ingress side of the FMN-facing interface access-group inside in interface inside

See Table 4-5 below for an example for ingress ACLs applied to the DMZ-facing interface.

Table 4-5 - 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 (*dmz*)	Inbound	• Telnet (servers in the DMZ to the control and enterprise networks) access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq telnet • HTTP (servers in the DMZ to the control and enterprise networks) access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq www • HTTPS (servers in the DMZ to the control and enterprise networks) access-list dmz extended permit tcp 10.19.1.0 255.255.255.0 10.18.0.0 255.255.0.0 eq https • ICMP (servers in the DMZ to the control and enterprise networks) 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 • Explicitly deny other traffic types to anywhere access-list inside deny 10.19.0.0 255.255.0.0 • Apply the ACLs above to the ingress side of the DMZ-facing interface access-group dmz in interface inside

See Table 4-6 below for the example for ingress ACLs applied to the enterprise network-facing interface.

Table 4-6 - Configuration Example for Ingress ACLs on the Enterprise Networking-Facing Interface

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

Modular Policy Framework:

Overview :

MPF provides a consistent and flexible way to configure security appliance features. For example, you can use MPF to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications.

MPF supports these features:

- TCP normalization, TCP and UDP connection limits and timeouts, and TCP sequence number randomization
- CSC
- Application inspection
- IPS
- QoS input policing
- QoS output policing
- QoS priority queue

The configuration of the MPF consists of four tasks:

1. Identify the Layer 3 and 4 traffic to which you want to apply actions. Refer to Identifying Traffic Using a Layer 3/4 Class Map for more information.
2. (Application inspection only) Define special actions for application inspection traffic. Refer to Configuring Special Actions for Application Inspections for more information.
3. Apply actions to the Layer 3 and 4 traffic. Refer to Defining Actions Using a Layer 3/4 Policy Map for more information.
4. Activate the actions on an interface. Refer to Applying a Layer 3/4 Policy to an Interface Using a Service Policy for more information.

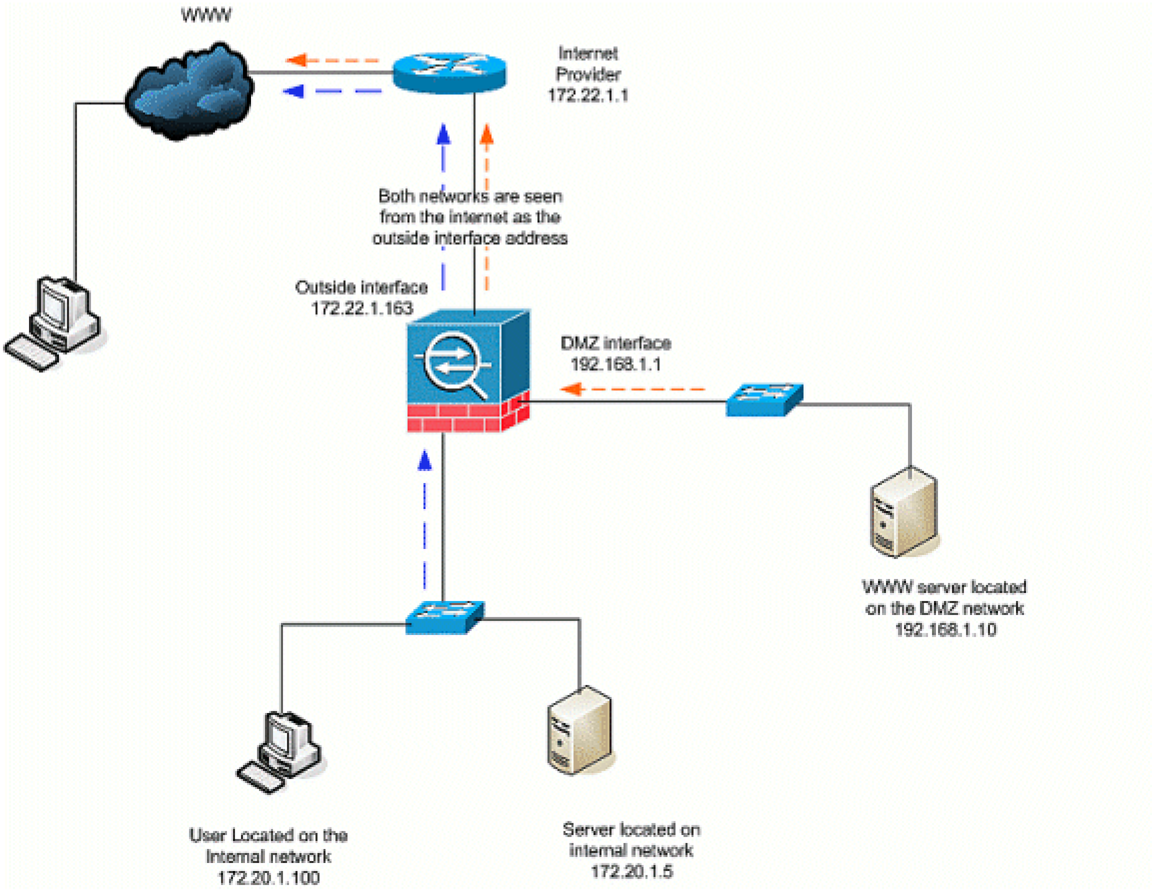
WWW
Internet Provider 172.22.1.1
Both networks are seen from the internet as the outside interface address
Outside interface 172.22.1.163
DMZ interface 192.168.1.1
WWW server located on the DMZ network 192.168.1.10
User Located on the Internal network 172.20.1.100
Server located on internal network 172.20.1.5

Figure 4-4 – Configuration Example

Regular Expression Overview:

A regular expression matches text strings either literally as an exact string, or with metacharacters, so you can match multiple variants of a text string. You can use a regular expression to match the content of certain application traffic; for example, you can match a URL string inside an HTTP packet.

In order to create a regular expression, use the regex command, which can be used for various features that require text matching. For example, you can configure special actions for application inspection with Modular Policy Framework with an inspection policy map (see the policy map type inspect command). In the inspection policy map, you can identify the traffic you want to act upon if you create an inspection class map that contains one or more match commands, or you can use match commands directly in the inspection policy map. Some match commands let you identify text in a packet with a regular expression; for example, you can match URL strings inside HTTP packets. You can group regular expressions in a regular expression class map (see the class-map type regex command).

Sample configuration, Controlling access using Regular Expressions:

```
regex urllist1 ".*\.([Ee][Xx][Ee]|[Cc][Oo][Mm]|[Bb][Aa][Tt]) HTTP/1.[01]"
regex urllist2 ".*\.([Pp][Ii][Ff]|[Vv][Bb][Ss]|[Ww][Ss][Hh]) HTTP/1.[01]"
regex urllist3 ".*\.([Dd][Oo][Cc]|[Xx][Ll][Ss]|[Pp][Pp][Tt]) HTTP/1.[01]"
regex urllist4 ".*\.([Zz][Ii][Pp]|[Tt][Aa][Rr]|[Tt][Gg][Zz]) HTTP/1.[01]"
regex contenttype "Content-Type"
regex applicationheader "application/.*"
boot system disk0:/asa802-k8.bin
ftp mode passive
dns server-group DefaultDNS
 domain-name default.domain.invalid
```

```
access-list inside_mpc extended permit tcp any any eq www
access-list inside_mpc extended permit tcp any any eq https
access-list inside_mpc extended permit tcp any any eq 8080
!
class-map type regex match-any DomainBlockList
 match regex domainlist1
 match regex domainlist2
 match regex domainlist3


class-map type inspect http match-all BlockDomainsClass
 match request header host regex class DomainBlockList


class-map type regex match-any URLBlockList
 match regex urllist1
 match regex urllist2
 match regex urllist3
 match regex urllist4


class-map inspection_default
 match default-inspection-traffic

class-map type inspect http match-all AppHeaderClass
 match response header regex contenttype regex applicationheader



class-map httptraffic
 match access-list inside_mpc




class-map type inspect http match-all BlockURLsClass
 match request uri regex class URLBlockList

policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
policy-map type inspect http http_inspection_policy
 parameters
  protocol-violation action drop-connection
 class AppHeaderClass
  drop-connection log
 match request method connect
  drop-connection log
 class BlockDomainsClass
  reset log
 class BlockURLsClass
  reset log

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

policy-map inside-policy
 class httptraffic
  inspect http http_inspection_policy

!
service-policy global_policy global
service-policy inside-policy interface inside
```

Authenticating Firewall Sessions for User Access to Servers in the DMZ

When users in the facilities 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

6. MSEA authentication required in form of username and password.

Configuration Example

The following example illustrates how to use firewall session authentication in a plant floor network. User 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 facilities network can access the server. The permitted protocols are Telnet, HTTP, and HTTPS.

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

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

```
aaa-server SAIFA protocol tacacs+
aaa-server SAIFA 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 facilities network (*inside*) against the ACL group *INSAUTH*.

```
aaa authentication match INSAUTH inside SAIFA
```

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 SAIFA
```

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

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

- Authentication exception based on users
- Authentication timeouts
- Customization of authentication prompts

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

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

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

Note The Cisco ASA 5520, which is the ASA model recommended for the SAIFA 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.

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 4-5 shows the traffic flow when the Cisco ASA is configured in inline IPS mode.

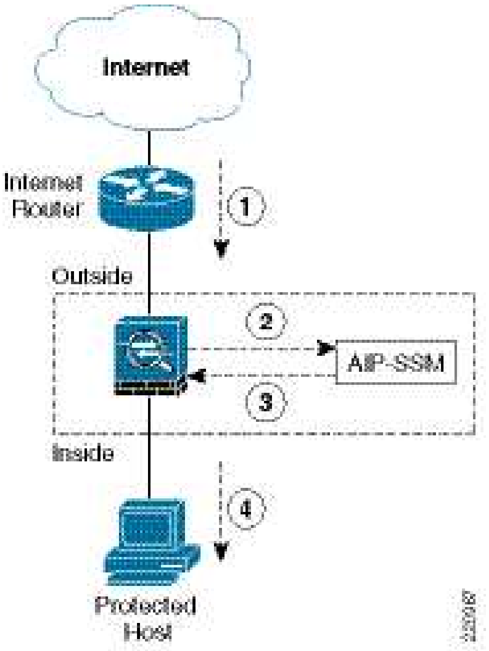

Figure 4-5 - Inline IPS Traffic Flow

Figure 4-5 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 4-6 shows an example of how traffic flows when the AIP-SSM is configured in promiscuous mode.

Figure 4-6 - Promiscuous Mode Traffic Flow

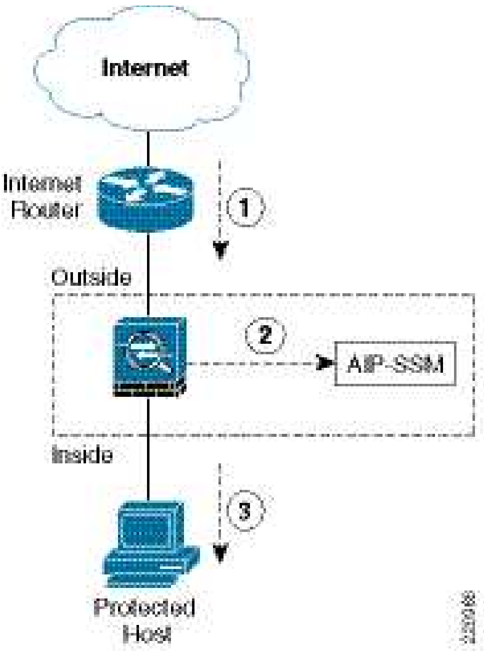

Figure 4-6 shows the following sequence of events:

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 AIM-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.

About This Document

History

Version No.	Issue Date	Author	Comments
1	6 March 2008	J.Martocci	Initial Draft
2	27 March 2008	Dave Clute	Established overall document structure, content and organization
3	1 May 2008	Dave Clute	Updated Chapter 1 – Solution Overview
4	17 June 2008	Chris Pirics Venkat Pothamsetty	Produced content for Chapter 3 Produced content for Chapters 4 & 5
5	1 July 2008	Dave Clute	Combined all content into updated version – format and nomenclature changes
5.3	15 July 2008	Dave Clute	Merged changes from Chapters 3-5, deleted old Chapter 4, updated chapter numbering and appendices
6	21 July 2008	Dave Clute	Added content for Chapter 1 – Executive Summary, Updated Table and Figure numbering for Chapter 4, Added Appendices A and B
7	30 July 2008	Dave Clute	Incorporated review comments and suggested revisions
8	15 August 2008	Dave Clute	Included updated diagrams and revisions from v7

Review

Reviewer's Details	Version No.	Date
Jerry Martocci / Ted Humpals	V6.0	25 July 2008
Dave Newgard	V6.0	26 July 2008
Nick Chong	V6.0	21 July 2008
Vikash Sharma	V6.0	25 July 2008
Venkat Pothamsetty	V7.0	5 August 2008

This document will be kept under revision control.

Appendix A – Reference Architecture Diagrams

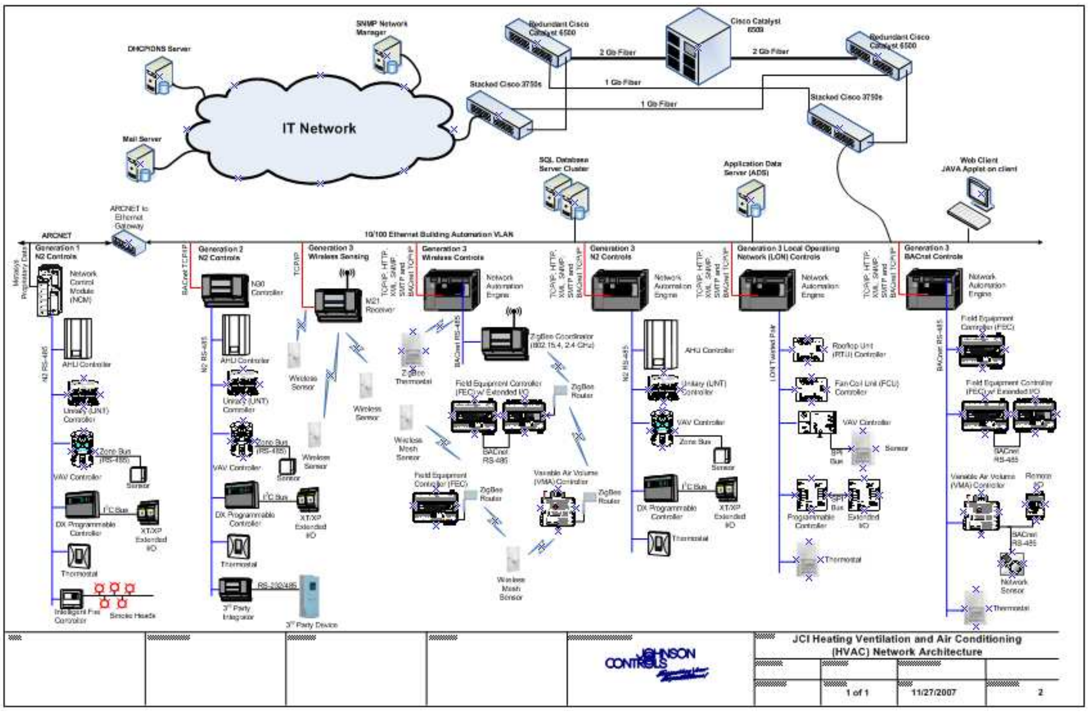
SNMP Network Manager
DHCP/DNS Server
Mail Server
IT Network
Redundant Cisco Catalyst 6500
Cisco Catalyst 6509
Redundant Cisco Catalyst 6500
2 Gb Fiber
2 Gb Fiber
1 Gb Fiber
1 Gb Fiber
Stacked Cisco 3750s
Stacked Cisco 3750s
SQL Database Server Cluster
Application Data Server (ADS)
Web Client JAVA Applet on client
ARCNET to Ethernet Gateway
ARCNET
10/100 Ethernet Building Automation VLAN
Metasys Proprietary Data
Generation 1 N2 Controls
Network Control Module (NCM)
N2 RS-485
AHU Controller
Unitary (UNT) Controller
Zone Bus (RS-485)
VAV Controller
Sensor
I²C Bus
DX Programmable Controller
XT/XP Extended I/O
Thermostat
Intelligent Fire Controller
Smoke Heads
BACnet TCP/IP
Generation 2 N2 Controls
N30 Controller
N2 RS-485
AHU Controller
Unitary (UNT) Controller
Zone Bus (RS-485)
VAV Controller
Sensor
I²C Bus
DX Programmable Controller
XT/XP Extended I/O
Thermostat
RS-232/485
3rd Party Integrator
3rd Party Device
TCP/IP
Generation 3 Wireless Sensing
M21 Receiver
Wireless Sensor
Wireless Sensor
Wireless Sensor
TCP/IP, HTTP, XML, SNMP, SMTP and BACnet TCP/IP
Generation 3 Wireless Controls
Network Automation Engine
BACnet RS-485
ZigBee Coordinator (802.15.4, 2.4 GHz)
ZigBee Thermostat
Field Equipment Controller (FEC) w/ Extended I/O
ZigBee Router
Wireless Mesh Sensor
BACnet RS-485
Field Equipment Controller (FEC)
ZigBee Router
Variable Air Volume (VMA) Controller
ZigBee Router
Wireless Mesh Sensor
TCP/IP, HTTP, XML, SNMP, SMTP and BACnet TCP/IP
Generation 3 N2 Controls
Network Automation Engine
N2 RS-485
AHU Controller
Unitary (UNT) Controller
VAV Controller
Zone Bus
Sensor
I²C Bus
DX Programmable Controller
XT/XP Extended I/O
Thermostat
TCP/IP, HTTP, XML, SNMP, SMTP and BACnet TCP/IP
Generation 3 Local Operating Network (LON) Controls
Network Automation Engine
LON Twisted Pair
Rooftop Unit (RTU) Controller
Fan Coil Unit (FCU) Controller
VAV Controller
SPI Bus
Sensor
SPI Bus
Programmable Controller
Extended I/O
Thermostat
TCP/IP, HTTP, XML, SNMP, SMTP and BACnet TCP/IP
Generation 3 BACnet Controls
Network Automation Engine
BACnet RS-485
Field Equipment Controller (FEC)
Field Equipment Controller (FEC) w/ Extended I/O
BACnet RS-485
Variable Air Volume (VMA) Controller
Remote I/O
BACnet RS-485
Network Sensor
Thermostat
JOHNSON CONTROLS
JCI Heating Ventilation and Air Conditioning (HVAC) Network Architecture
1 of 1
11/27/2007
2

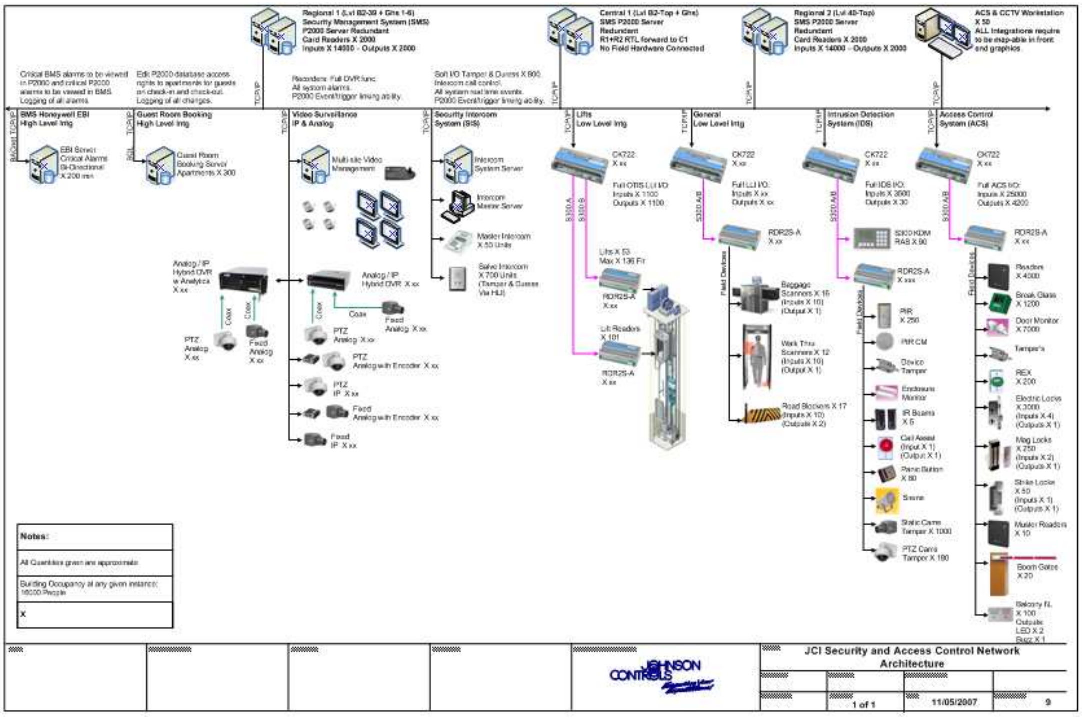
Regional 1 (Lvl B2-39 + Gns 1-6) Security Management System (SMS) P2000 Server Redundant Card Readers X 2000 Inputs X 14000 - Outputs X 2000
Central 1 (Lvl B2-Top + Gns) SMS P2000 Server Redundant R1+R2 RTL forward to C1 No Field Hardware Connected
Regional 2 (Lvl 40-Top) SMS P2000 Server Redundant Card Readers X 2000 Inputs X 14000 – Outputs X 2000
ACS & CCTV Workstation X 50 ALL Integrations require to be map-able in front end graphics
Critical BMS alarms to be viewed in P2000 and critical P2000 alarms to be viewed in BMS. Logging of all alarms.
Edit P2000 database access rights to apartments for guests on check-in and check-out. Logging of all changes.
Recorders: Full DVR func. All system alarms. P2000 Event/trigger timing ability.
Soft I/O Tamper & Duress X 800. Intercom call control. All system real time events. P2000 Event/trigger timing ability.
TCP/IP
BACnet TCP/IP
BMS Honeywell EBI High Level Intg
EBI Server Critical Alarms Bi-Directional X 200 min
TCP/IP
SQL
Guest Room Booking High Level Intg
Guest Room Booking Server Apartments X 300
TCP/IP
Video Surveillance IP & Analog
Multi-site Video Management
Analog / IP Hybrid DVR w Analytics X xx
Analog / IP Hybrid DVR X xx
Coax
Coax
Coax
Coax
PTZ Analog X xx
Fixed Analog X xx
PTZ Analog X xx
Fixed Analog X xx
PTZ Analog with Encoder X xx
PTZ IP X xx
Fixed Analog with Encoder X xx
Fixed IP X xx
TCP/IP
Security Intercom System (SIS)
Intercom System Server
Intercom Master Server
Master Intercom X 50 Units
Salve Intercom X 700 Units (Tamper & Duress Via HLI)
TCP/IP
Lifts Low Level Intg
CK722 X xx
Full OTIS LLI I/O: Inputs X 1100 Outputs X 1100
S350-A
S350-B
Lifts X 53 Max X 136 Flr
RDR2S-A X xx
Lift Readers X 101
RDR2S-A X xx
TCP/IP
General Low Level Intg
CK722 X xx
Full LLI I/O: Inputs X xx Outputs X xx
S350-A/B
RDR2S-A X xx
Field Devices
Baggage Scanners X 16 (Inputs X 10) (Output X 1)
Walk Thru Scanners X 12 (Inputs X 10) (Output X 1)
Road Blockers X 17 (Inputs X 10) (Outputs X 2)
TCP/IP
Intrusion Detection System (IDS)
CK722 X xx
Full IDS I/O: Inputs X 3500 Outputs X 30
S350-A/B
S300 KDM RAS X 90
RDR2S-A X xxx
Field Devices
PIR X 250
PIR CM
Device Tamper
Enclosure Monitor
IR Beams X 5
Call Assist (Input X 1) (Output X 1)
Panic Button X 80
Sirens
Static Cams Tamper X 1000
PTZ Cams Tamper X 180
TCP/IP
Access Control System (ACS)
CK722 X xx
Full ACS I/O: Inputs X 25000 Outputs X 4200
S350-A/B
RDR2S-A X xx
Field Devices
Readers X 4000
Break Glass X 1200
Door Monitor X 7000
Tamper's
REX X 200
Electric Locks X 3000 (Inputs X 4) (Outputs X 1)
Mag Locks X 250 (Inputs X 2) (Outputs X 1)
Strike Locks X 50 (Inputs X 1) (Outputs X 1)
Muster Readers X 10
Boom Gates X 20
Balcony FL X 100 Outputs: LED X 2 Buzz X 1
Notes:
All Quantities given are approximate
Building Occupancy at any given instance: 16000 People
X
JOHNSON CONTROLS
JCI Security and Access Control Network Architecture
1 of 1
11/05/2007
9

JCI Architecture Overview Rev1.vsd – Visio 2003 File dated 27 November 2007

Obtain current version from J. Martocci – Lead Staff Engineer – Johnson Controls

Appendix B – Glossary & Acronym List

A table of most, if not all, of the terms and acronyms used in this document follows:

AAA	Authentication, Authority and Accounting
ACS	Access Control Server
ACL	Access Control List
ADS	Application Data Server
AES	Advanced Encryption Standard
AGA	Advanced Graphics Application
AI	Analog Input Object
AO	Analog Output Object
APDU	Application Protocol Data Units
APIPA	Automatic Private Internet Protocol Addressing
ASA	Adaptive security Appliance
ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers
AT	Advanced technologies: covers IP Telephony, security, storage and wireless
B2B	Business to Business
BACnet®	Building Automation Control Network
BAS	Building Automation System
BAS/IP	Building Automation System over Internet Protocol
BBMD	BACnet Broadcast Management Device
BDT	Broadcast Distribution Table
BIBBS	BACnet® Interoperability Building Blocks
BMS	Building Management System
BOM	Bill of Materials
BPDU	Bridge Protocol Data Units
BU	Business Unit: a Cisco product manuafacturing organization
CAM	Content Addressable Memory
CAPEX	Capital Expense
Case	Case is used to track calls in the CARE database
CCIE	Cisco Certified Internetwork Engineer
CCT	Controller Configuration Tool
CPE	Customer Premises Equipment
CSA	Cisco Security Agent

CSC	Construction Specifications Canada
CSI	Construction Specification Institute
DCN	Data Communications Network
DDA	Destination Delivery Agent
DDOS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DIG	Design & Implementation Guide
DLLR	Demand Limiting and Load Rolling
DMZ	Demilitarized Zone
DP	Dew Point
DSL	Digital Subscriber Line
DNS	Domain Name System
DTS	Device Time Server
DX	Extended Digital Controller
EA	Enterprise Architecture
ECPR	Equal Cost Path Routing
EFT	Early Field Test: a robust release for testing the final product in a customer environment
EMEA	Cisco theater of operation that includes Europe, the Middle East and Africa
ENDP	Enthalpy - Dew Point
ENRH	Enthalpy - Relative Humidity
EOI	End of Interval Pulse (used with DLLR)
EOL	End of Life
EPROM	Erasable Programmable Read-Only Memory
ETHERNET	The most prevalent networking standard defined by the IEEE 802.3
FC Bus	Field Controller Bus
FDT	Foreign Device Table
FEC	Field Equipment Controller
FES	Field Engineering Services
FMN	Facility Management Networking
FMS	Facility Management System
FQIR	Fully Qualified Item Reference
GA	Generally Available
GTM (G-T-M)	Go to Market
HSRP	Hot Standby Routing Protocol
HTML	Hypertext Markup Language

HTTP	Hypertext Transfer Protocol
HVAC	Heating, Ventilating, and Air Conditioni
IDS	Intrusion Detection Service
IEEE	Institute of Electrical and Electronics Engineers
IFC	Intelligent Fire Controller
ILC	Intelligent Lighting Controller
IO	Input / Output
IOM	Input/Output Module
IP	Internet Protocol
IPv6	Internet Protocol Version 6
ISO	International Standards Organization
IT	Information Technology
JNLP	Java Network Launching Protocol
LAN	Local Area Network
LCT	Logic Connector Tool
LON	Local Operating Network
LONMARK®	A standards organization founded by LONWORKS network users
MAC's	Moves Adds & Changes
MADD	Mixed Air Dual Duct application
MASD	Mixed Air Single Duct application
MECVT	Master-Slave/Token-Passing (MS/TP) to Ethernet Converter
MIST	Multiple Instance Spanning Tree
MPF	Modular Policy Framework
MSA	Master Service Agreement
MSTP	Master-Slave/Token-Passing
MTBF	Mean Time Between Failure
MSFC	Multilayer Switch Feature Card
MVE	Metasys for Validated Environments Extended Architecture
NAE	Network Automation Engine
NAT	Network Address Translation
NCE	Network Control Engine
NIC	Network Interface Card
NIE	Network Integration Engine
NOC	Network Operations Center
NRFU	Network Ready for Use
NTP	Network Time Protocol

OBJECT	Self-contained functional items in the Metasys system that contain processes to manage building automation system components
OPEX	Operating Expense
PACL	Port Access Control Lists
PDIO / PDIOO / PPDIOO	(Prepare,) Planning, Design, Implementation, Operations (and Optimization). Cisco's methodology to define the continuous lifecycle of services required by the end customer: Prepare – the evaluation and qualification period for the customer to solve a particular business need prior to making a decision on a product or solution. Plan – the assessment and needs analysis period to further define the architecture and applications requirements to meet the customer solution Design – from the output of the planning phase detail how the application design and hardware infrastructure come together to define the architected system. Implement – deliver project planning, staging, installation, and configuration of the system; execution of test plans; training to operations staff; and end-user training Operate – on-going operation of infrastructures: monitoring of applications and equipment, monitoring of Help Desk and problem escalations; on premise resources and tools; remote resources and tools; Change Control and provisioning processes; escalation policies and procedures; Service Level Agreements (SLA's) tied to deliverables Optimize – the continuous process of PDIO to provide fine tuning of the existing deployment to get best performance level; the initiation of a new PDIO phase based on a tactical or strategic change in the customer business
PID	Project Initiation Document
PMO	Program/Project Management Office
POE	Power Over Ethernet
PPP	Point-to-Point Protocol
PRD	Program Requirements Document: Mandatory element of the NPI Process
Product Life Cycle	Sequence of stages in the marketing of a product that begins with commercialization and ends with removal from the market
Project	Distinct set of short term activities performed to achieve the results stated in an initiative
QOS	Quality of Service
R&S (R/S)	Routing & Switching
ROI	Return on Investment
RTU	Rooftop Unit
SAB	Sensor/Actuator Bus
SAIFA	Secure Architecture for Intelligent Facility Applications
SCT	System Configuration Tool

SECVT	Serial to Ethernet Converter
SKU	Stock Keeping Unit: a number associated with a product for inventory purposes.
SLA	Service Level Agreement
SMARTnet	Core support package, which provides software updates and upgrades, 24-hr global telephone support from expert technicians, advanced hardware replacement, and registered access to online tools and technical assistance available on the Cisco.com site
SNMP	Simple Network Management Protocol
SNTP	Simple Network Time Protocol
SONA	Service Oriented Network Architecture
STP	Spanning Tree Protocol
TAC	Technical Assistance Center, the TS function that is responsible for the successful resolution of all customer support incidents. The TAC provides 24x7 telephone support to customers worldwide.
TCO	Total Cost of Ownership
TCP/ IP	Transmission Control Protocol/Internet Protocol
UI	User Interface
VACL	VACL - VLAN Access Control Lists
VAV	Variable Air Volume
VAVDD	Variable Air Volume (VAV) Dual Duct application
VAVSD	Variable Air Volume (VAV) Single Duct application
VFD	Variable Frequency Drive
VLAN	Virtual Local Area Network
VoIP	Voice Over Internet Protocol
VPN	Virtual Private Network
VSD	Variable Speed Drive
WAN	Wide Area Network
WAP	Wireless Access Point
WIFI	Wireless Fidelity
WLAN	Wireless Local Area Network
XIL	External Interface File
XML	Extensible Markup Language