White Paper

Honeywell

Honeywell Process Solutions

Experion Network Best Practices

Author/Editor: Jay Gustin

Document ID: WP-07-02-ENG (formerly ENBP-WP)
Original Issue Date: April 2004
Revised: June 2007
Version: 3.2

Table of Contents

1. Introduction

Scope

This document is intended to provide “best practices” advice for planning the installation of Experion FTE networks, and connecting them into the plant IT network.

Users

Intended users of this document include:

- Honeywell System Consultants, Technical Assistance Center, Project Services
- Honeywell clients

Definitions

Definitions	
ACE	Advanced Control Environment- An Experion node used for high-level control
ACL	Access Control List- A Cisco command for filtering traffic
CDA	Control Data Access- The Experion data access layer
DC	Domain Controller
DHEB	Data Hiway Ethernet Bridge
DSA	Distributed System Architecture- The Experion method of sharing data.
ESVT	Experion Server TPS
ES-T	Experion Station TPS
ACE	Experion application node
FIM	Fieldbus Interface Module
FTE	Fault Tolerant Ethernet- the control network of Experion PKS
FTEB	FTE Bridge- FTE interface for C200 controllers and FIMs
GBIC	Gigabit Interface Converter module for Cisco switches
HSRP	Hot Standby Router Protocol
IP	Internet Protocol
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol- a client-server protocol for accessing a directory service
MAC	Media Access Controller
NAT	Network Address Translation
NTP	Network Time Protocol
PHD	Process History Database- Honeywell’s data historian solution
SFP	Small form factor plug-in interface module
TCP	Transport Control Protocol
Uplink	Any interface that connects switches to switches or switches to routers

2. FTE Network Infrastructure

Overview

An FTE network is comprised of a variety of node types and networking components. This section describes the considerations and requirements for connecting and configuring these elements to provide a system that has significant security and reliability improvements over a simple Ethernet network.

Plant Network Levels

A plant network has four layers or levels. The following table briefly describes these levels. Level numbers are used to simplify the description of the node location within the network hierarchy. The FTE network of an Experion PKS system includes levels 1 and 2. Sections 3 through 6 of this document provide further details on these levels, including specific network requirements.

FTE Communities

An FTE community is a group of nodes that have fault tolerant communication coverage using FTE test messages. These nodes are all members of the same broadcast domain. Nodes that are either single attached, or are dual attached but do not run FTE, may also be members of the FTE community. Honeywell does not recommend multiple FTE communities in the same broadcast domain. The FTE node number limits may seem to inhibit large systems using FTE. This is not the case, however, as FTE communities may be interconnected using a router. The individual FTE communities should be designed to include those nodes that have critical intercommunication requirements. Data can be shared between routed FTE communities via Distributed Server Architecture (DSA). Using this technique, a very large system can be constructed of FTE nodes with a wide geographical distribution.

Best Practice Architecture

The drawings shown in this paper represent the Honeywell recommended best practice for a large installation. While variations of the architecture are possible, this topology represents the highest level of security and reliability. The emphasis is on isolating critical areas of function with layers of switches such that local peer-peer control is most important, peer to external peer is next most important, controller to server/station is next most important and server to station, ACE and other L2 nodes is next most important. Communication from L2 to L3 is generally less critical and more restriction can be placed on this path.

Level	Node Descriptions
Level 4	Plant Level Applications
Level 3	Advanced Control and Advance Applications (Non-Critical Control Applications)
Level 2	Supervisory Control, Operator HMI (HMI, and Supervisory Controllers)
Level 1	Real Time Control (controllers and IO)

3. Level 1

Description

Level 1 nodes are the heart of the control system. This network segment contains controllers, FTEB-based I/O, and Series A or Series C FIM nodes.

Level 1 Best Practices

The best practice is to put Level 1 nodes on a separate switch or Control Firewall pair. This allows critical peer-to-peer traffic that cannot tolerate a communication delay of longer than 250 ms following an FTE cable fault. It also gives controllers a level of isolation from other nodes during catastrophic failure or network disturbance. The user should arrange for the most critical elements of control to be connected to this switch. Because Level 1 nodes include controller nodes, the critical control traffic must have adequate bandwidth. The following sections describe how to accomplish this.

Control Firewall Best Practice

Experion R300 introduced the Control Firewall. This is an appliance that offers a level of protection of the embedded controller nodes against unwanted traffic from Level 2 and above. All Series C nodes including C300, FIM, and FTEB-based 1756 I/O, must be connected to the internal interfaces of a Control Firewall. The uplink of the Control Firewall is attached to a Level 2 Cisco switch using an interface configured for 100Mbps full duplex and port fast. An interface configured as an uplink can be used with the “spanning-tree portfast” attribute added to the configuration of the interface, or one configured as a normal Level 2 node connection can be used. Control Firewalls cannot be cascaded.

The Control Firewall has the following features:

- Allows only CDA connected traffic through by using TCP port filtering
- Limits broadcasts to ARP and Bootp and limits the rate
- Limits the rate of connection to mitigate SYN flood attacks
- Limits multicast to FTE messages
- Allows NTP time sync packets, but limits the rate
- Prioritizes internal packets over external packets
- No user configuration required

NOTE: Computer nodes running a Windows operating system with file sharing enabled must not be attached to the inside of the Control Firewall. NetBIOS messages will be blocked from entering the CF9, and the internal node will become the master browser as it can’t see any other nodes in the system. This will have the effect of preventing file sharing for Level 2 nodes.

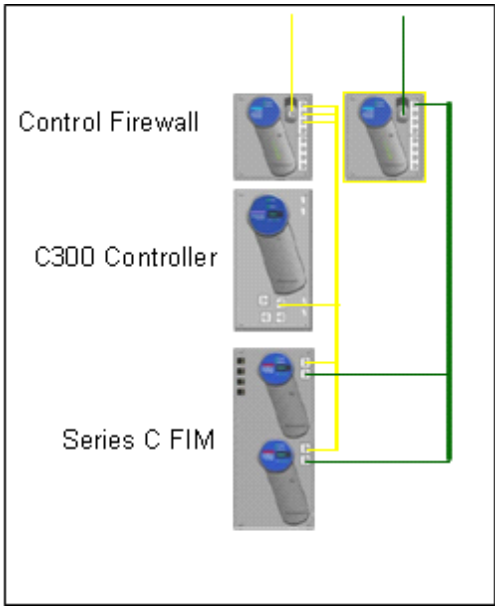

Series C Level 1 LAN Cluster

- Citizenship
  - Controller (C300)
  - Series C Fieldbus Interface Module
  - Control Firewall
- Purpose
  - Peer-peer Control
- Level 1 Control Firewalls
  - Provide point to point connectivity
  - Cyber Security
  - Prioritization of inside over outside packets

C200 With FTEB Best Practice

Installations with C200 controllers connected to FTE with the FTEB must be connected to a Cisco switch with a Level 1 configuration installed. Several configuration settings in the Honeywell scripts enable protection for the Level 1 traffic.

First, the TCP ports that are used for critical control and display traffic will be fixed and well known. Reception of a packet with those TCP port values informs the Cisco switches that this packet must be given priority. The output queue in the switches is configured to ensure traffic priority as follows:

- Control traffic is sent to the highest priority output queue.
- Display traffic is sent to the second level priority output queue.
- Any remaining traffic is sent to the lowest priority output queue.

Second, the uplink interface on the Cisco Level 1 switch is configured to limit the amount of broadcast and multicast traffic. Broadcast or multicast traffic levels that exceed the limit will be cut off, but other traffic will not be affected.

The use of a separate IP address range for Level 1 nodes is no longer being recommended as an overall best practice due to the difficulty of configuration. This scheme is still recommended for those installations where Level 1 address reuse is required. There is a discussion of this in the section on IP addressing:

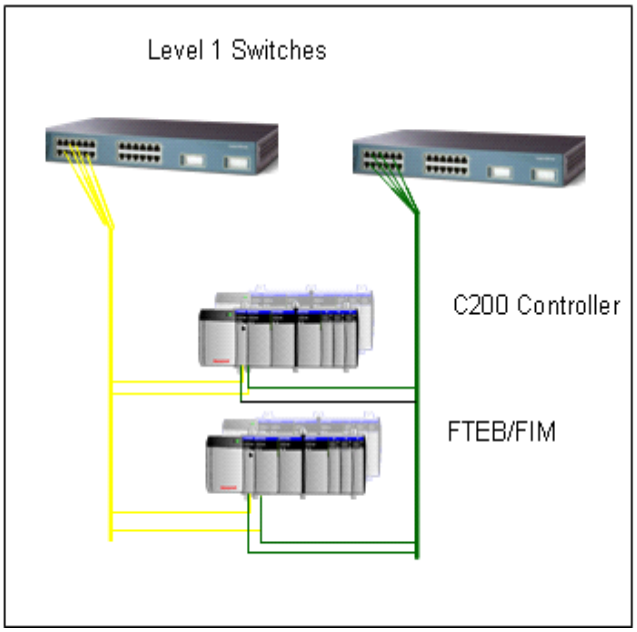

Series A Level 1 LAN Cluster

- Citizenship
  - Controller (C200)
  - Fieldbus Interface Module
  - Cisco Switches
- Purpose
  - Peer-to-peer Control
- Level 1 Switches
  - Provide point to point connectivity for FTE devices in cabinet
  - High Reliability Configuration
    - Always redundant
    - Pre-configure CDA traffic in high priority switch queue
    - Pre-configure view traffic in second highest priority queue
    - Pre-configure other traffic in low priority switch queue

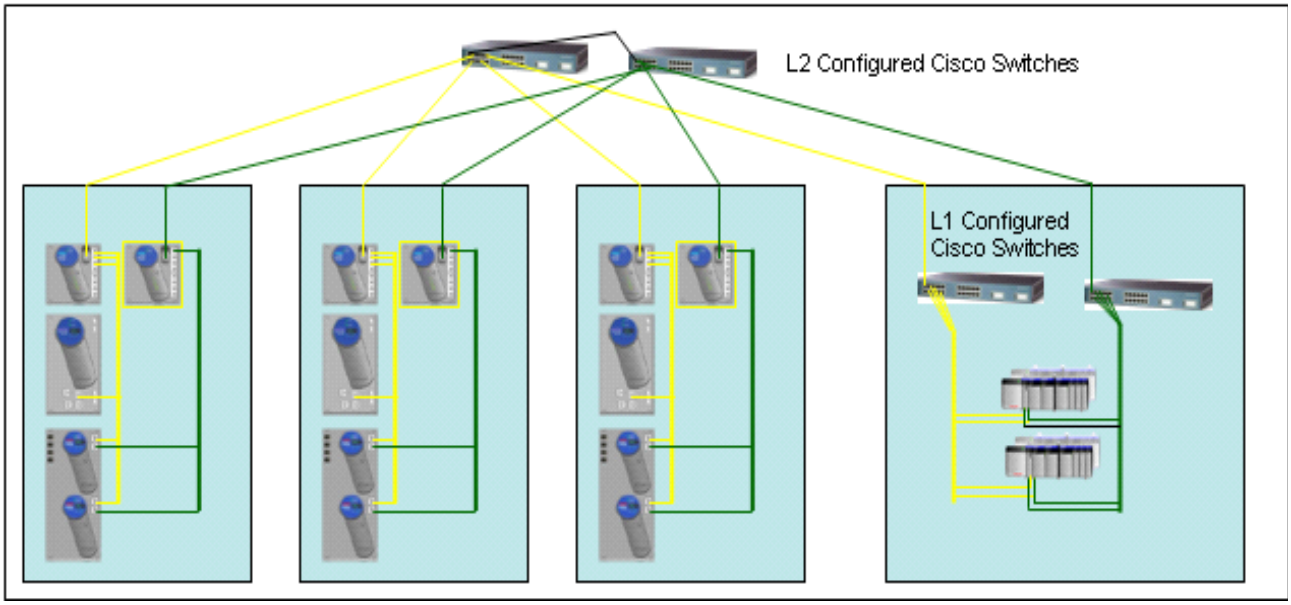

Connectivity Between Level 1 LAN Clusters

- Cisco Switches
  - Connect Level 1 clusters
  - High Reliability Configuration
    - Pre-configured bandwidth limits for broadcast, multicast storm suppression:
      - High traffic conditions will trigger Cisco switch to disable offending ports
      - Automatic port enabling when traffic profile returns to normal
    - Dual Cisco switch faults impact inter-cabinet traffic only

(alternatively, one pair of switches with the L1|L2 split configuration might be used)

Connection of Level 1 Nodes that Intercommunicate

The best practice is to connect Level 1 nodes that intercommunicate to the same switch pair, so that they will have the shortest communication path and the lower cable fault detection time. If intercommunicating Level 1 nodes cannot be contained on a single switch due to size of the installation or geographic dispersion, then their communications may go through the Level 2 switches. Level 2 switches are configured to have the same quality of service approach as the Level 1 Cisco switches. The same TCP ports are given the prioritization scheme described for Level 1. The control traffic entering from a Level 1 switch will be tagged with the highest priority at the ingress. The output queue to the destination Level 1 node will send the control traffic before any other traffic. Communications redundancy is provided for this peer to peer traffic by always having two “pipes” for peer to peer and using FTE to provide four possible paths. In addition, the Level 2 switches are configured to have storm protection on the interfaces where Windows operating system nodes will reside. This storm protection will prevent broadcast or multicast storms caused by a node that is infected and using a denial-of-service attack. If a node reaches a limit of 20% of the connection bandwidth being used for broadcast or multicast, then the interface is cut off until the traffic level falls below 18%. Normal FTE traffic of broadcast and multicast is well below 2% for each.

4. Level 2

Description

Level 2 nodes are the primary server, view and advanced control nodes for the process control system. Examples of Level 2 nodes include servers, stations, ACE nodes, and PHD nodes. These nodes are essential for operation of the process, but not as critical to control as the Level 1 nodes.

Level 2 Best Practices

The Cisco switches in Level 2 are configured to provide the security and reliability described in the Level 1 to Level 2 discussion. The nodes that reside on this level are more susceptible to attacks by viruses or software glitches because of the open nature of the operating system and the customized software that is running on these nodes. Thus, protection for broadcast and multicast storms on the interfaces to these open nodes is configured in the Cisco switches. Also the display traffic as with the control traffic is given a higher priority so the traffic for view to the process will take precedence over other traffic on the switch. This is especially important if there is a “bad actor” on the LAN that is generating high traffic. The higher priority control and view traffic will get to the destination first.

An important best practice is to avoid connecting a computer node to multiple networks. Connection of a server, for example, to two networks (“dual-homed”) turns that node into a router, which is a poor practice. Instead, the Experion network structure provides for the use of routers to join Level 2 nodes to Level 3 networks or to other Level 2 networks. A built-for-purpose router must be used in order to provide security and reliability through the use of Access List filtering.

There are exceptions where a third NIC card can be used for private connection to a single device that uses Ethernet. One example is the Honeywell DHEB for bridging to the Data Hiway.

There are nodes other than the Experion server, console and application nodes that can connect to Level 2 switches. Some of these devices have dual Ethernet connections. FTE is compatible with dual Ethernet nodes; they will not have the FTE protection, but no interference will occur and both types of nodes can intercommunicate.

Level 2 node examples:

- Safety Manager
- Terminal servers
- Matrikon servers
- PLCs

Single attached nodes such as terminal servers or subsystem devices also can attach to Level 2. If there are a large number of single attached nodes, then a separate switch can be used to aggregate these nodes. This switch will be counted as a level for spanning tree purposes, so it must not be connected to a FTE switch that is at the third level. This switch must not be connected to Level 1 switches. It can be connected to either the Yellow or Green side. The Yellow side is preferred.

Embedded operating systems may not have enough processing power to handle the volume of multicast and broadcast traffic generated by FTE test messages and Address Resolution Protocol (ARP) packets. This type of node must either be connected at Level 3 or protected with Access List filtering on a separate switch on Level 2. The recommended switch is one of the qualified Experion switches. Honeywell Network Services can be consulted for proper configuration of this switch.

FTE networks require a single crossover cable at the top of the hierarchy. In large systems it is recommended that a gigabit connection be used for this crossover connection. In the case of multiple faults, the backbone traffic will pass through this connection so the highest bandwidth will be available for this traffic. A determination of the necessity for greater than 100 Mbps for this crossover can be made by adding the total of the average bandwidth of all of the cluster servers. If this is higher than 20 Mbps, then a 1Gbps connection is recommended.

The best practice for the crossover cable is to use only one per FTE community. The placement of this crossover can be between any of the Level 2 Yellow and Green switches, as long as the rule of 3 levels of switches is preserved. The crossover must not be connected to Level 1 switches.

Implementing Level 2 Best Practices

To increase reliability and security, Level 2 nodes must be divided into two IP address ranges. Using two ranges simplifies the use of access lists for filtering as described below.

- Servers need access to nodes on other subnets as well as access to certain nodes on Level 3 and possibly Level 4. Communication to other nodes may include Distributed Server Access (DSA), as well as Engineering access to load control schemes and high-level control.
- Other nodes on Level 2 do not need to be accessed by nodes on Level 3 and should be protected from such access.

To accomplish node access control, filtering is done in either the router, or the switch interface that connects to the router. Filtering, which is implemented by creating specific access lists for the Cisco equipment, must accomplish the following:

- Allow servers to have complete two-way communication with other nodes on all levels of the network.
- Allow non-server nodes to communicate with Domain Controllers for authentication and name service.
- Allow Level 2 nodes to initiate communication with Domain Controllers on Level 3.

Communication between Level 2 nodes and Domain Controllers on Level 3 can be accomplished by adding access lists that enable established communications to return TCP packets from the Level 3 nodes to the initiating Level 2 nodes. In addition, communications can occur using UDP packets for Kerberos and LDAP. The filter must allow specific UDP port numbers used for these packets. See the section on “Example Cisco Router Configuration Statements” for examples of access lists to be used for filtering.

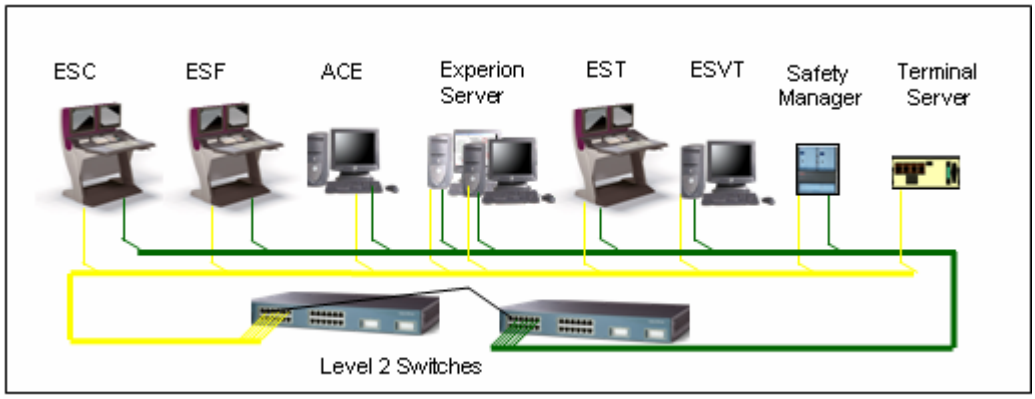

Level 2 LAN

- Citizenship
  - PKS Server
  - Server Stations
  - Direct Stations
  - ACE
  - EST
  - ESVT
  - Subsystem Interfaces
  - Cisco Switches
- L2 Cisco Switches
  - Point-to-point connectivity for L2 Devices
  - Pre-configured bandwidth limits for broadcast, multicast, storm suppression:
    - disables ports with high traffic conditions
    - enables ports when traffic profile returns to normal
  - Preconfigure CDA traffic in high-priority switch queue (ACE-ACE, ACE-Controller)
  - Preconfigure non-CDA traffic in low-priority switch queue

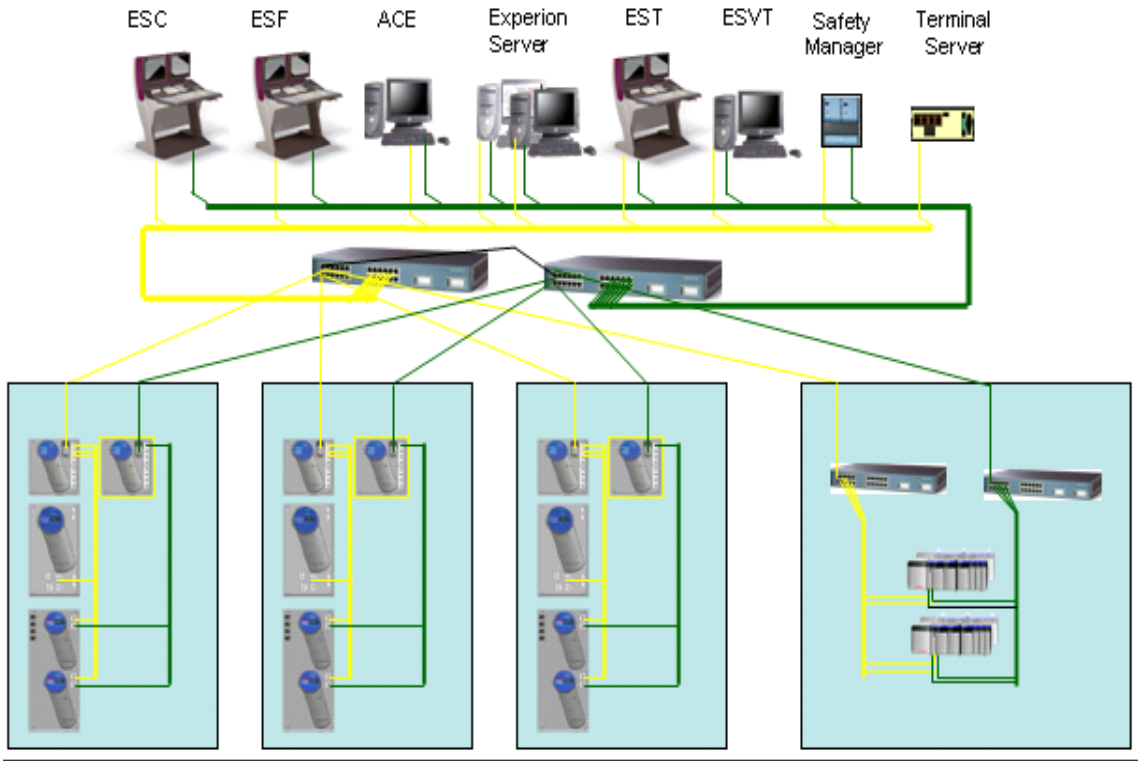

L2 to L1 Connectivity – Complete FTE Community

- L1 Control Firewall
  - Blocks traffic not needed for control.
  - Higher level of protection for peer-peer nodes on same Control Firewall.
  - Prioritizes internal traffic over external.
- L1 Cisco Switches
  - Prioritize ingress traffic; Non-CDA in low priority queue.
  - Ensures L2 – L1 supervisory traffic cannot disrupt L1 control
- L2 Cisco Switches
  - Provide L1-L2 connectivity
  - Broadcast, multicast storm suppression
  - Preconfigure CDA traffic in high-priority switch queue (e.g., ACE-ACE, ACE-Cx, ACE-FIM, Server-Cx, Server-FIM)
  - Preconfigure non-CDA traffic in low-priority switch queue

(alternatively, one pair of switches with the L1|L2 split configuration might be used)

5. Level 3

Description

In Level 3, all of the subnets on the plant-wide network, including FTE communities, are tied together. Additionally, the Level 3 router may be connected to Level 4.

Level 3 Best Practices

In order to accomplish control strategies from one FTE subnet to another FTE subnet, complete access between servers on each subnet must be allowed.

Implementing Level 3 Best Practices

The following list summarizes the networking configuration requirements for Level 3 of the FTE network:

- Provide access between FTE community subnets by grouping servers into an IP address range that can be separated from the other Level 2 nodes through use of a subnet mask, as discussed in Section 8.
- The use of unicast for DSA keep alive messages is the recommended best practice. Multicast is less desirable, but if it is used, enable IP multicast routing for the DSA multicast address, which is 225.7.4.103, and create an access-list filter to allow only this multicast address to pass to the FTE subnets. Redirection Manager may also use multicast addresses as described in the paragraph Redirection Manager below.
- Configure each FTE subnet to be in a separate VLAN, which protects the FTE community from unintended access by other nodes on the router.
- Connect only Switch A (Yellow tree) to the router. If multiple connections are desired see the paragraph Multiple connections from L2 to L3 Best Practice below. The router interface connected to FTE must be a routed interface. The “no switchport” configuration statement must be attached to the interface.
- Configure access list filters for the FTE communities that:
  - Permit complete access only to the server IP range, and
  - Allow established access to the remainder of the FTE subnet.
  - Deny all other access to the FTE subnet.
- If not using SFP/GBIC connections, configure the FTE switch’s router interfaces for 100-Mbps Full Duplex.

  NOTE: The router must be connected to a Level 2 switch interface that is configured as an uplink port, or to a SFP/GBIC-based interface.

Redirection Manager

Honeywell’s Redirection Manager (RDM) can use the FTE multicast test message multicast from the servers to keep track of when the primary OPC server goes off line. It is recommended to use the multicast only when the OPC client is in the same FTE community as the servers. When the OPC client resides in Level 3, or when the client is in another FTE community, then a mechanism using ICMP must be selected. In this case, ICMP must be allowed between L3 nodes and subnets.

Multiple Connections from L2 to L3 Best Practice

Dual connections between the FTE backbone switches and Level 3 may be desired. The best practice in this case is to use two routers that are running the Hot Standby Router Protocol (HSRP). HSRP will provide a redundant level of protection in both connection and equipment for the Level 3 router. The Level 3 nodes can connect redundantly to both routers using dual Ethernet, FTE or can be single attached to the primary router. The HSRP algorithm will protect against Level 2 cable failures when the Level 3 node is single attached. The configuration of the router is not possible with a

standardized configuration file. It is recommended that Honeywell Network Services group be contacted for router configuration consultation.

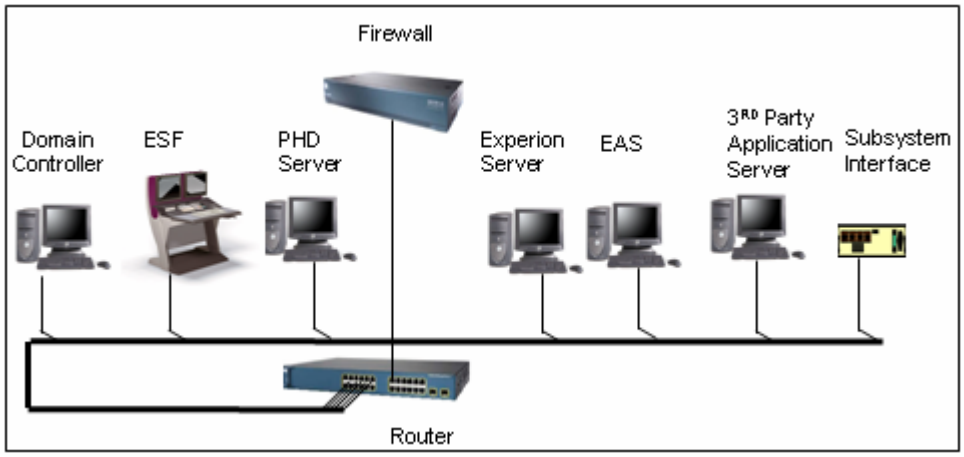

Level 3 LAN

Citizenship

- Plant Historians
- Applications
- Advanced Control
- Advanced Alarming
- Router / Switch
- Secure Gateway to L4
- Domain Controllers
- Subsystem Devices
- DSA Connected PKS Servers
- Server Stations (Monitoring)
- Engineering Stations

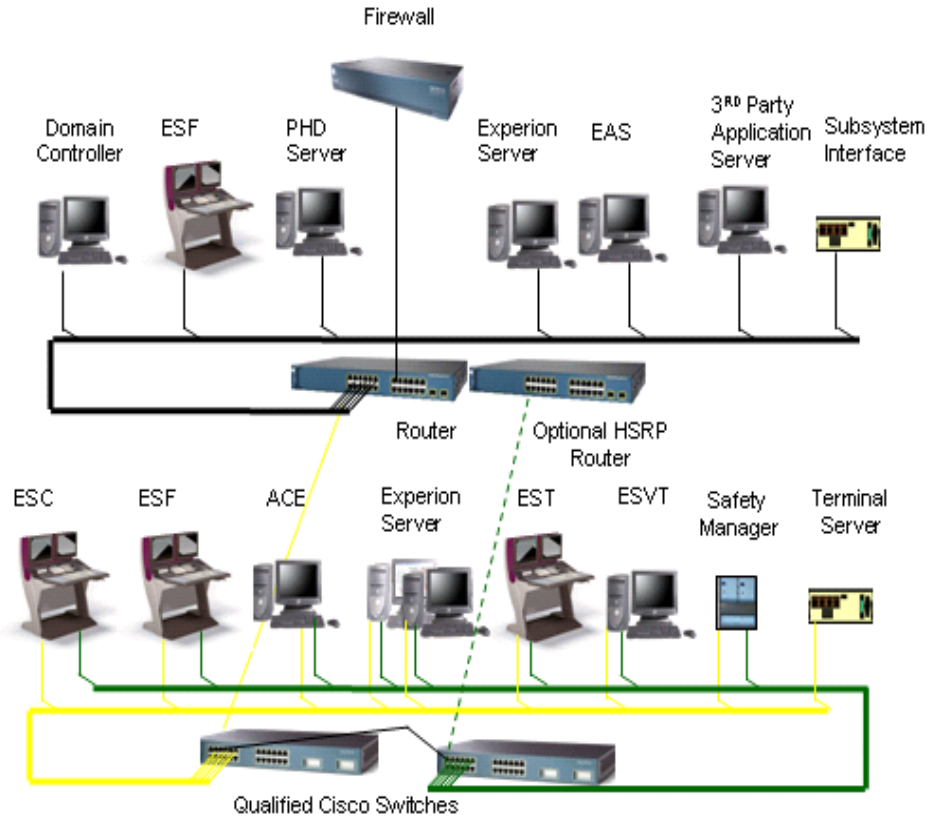

L3 to L2 Connectivity – Routing

- Cisco 3560 or 3750-- recommended router between L3 and L2
  - Security Filter to permit communications to and from specific nodes (may be implemented in Cisco PIX Firewall)

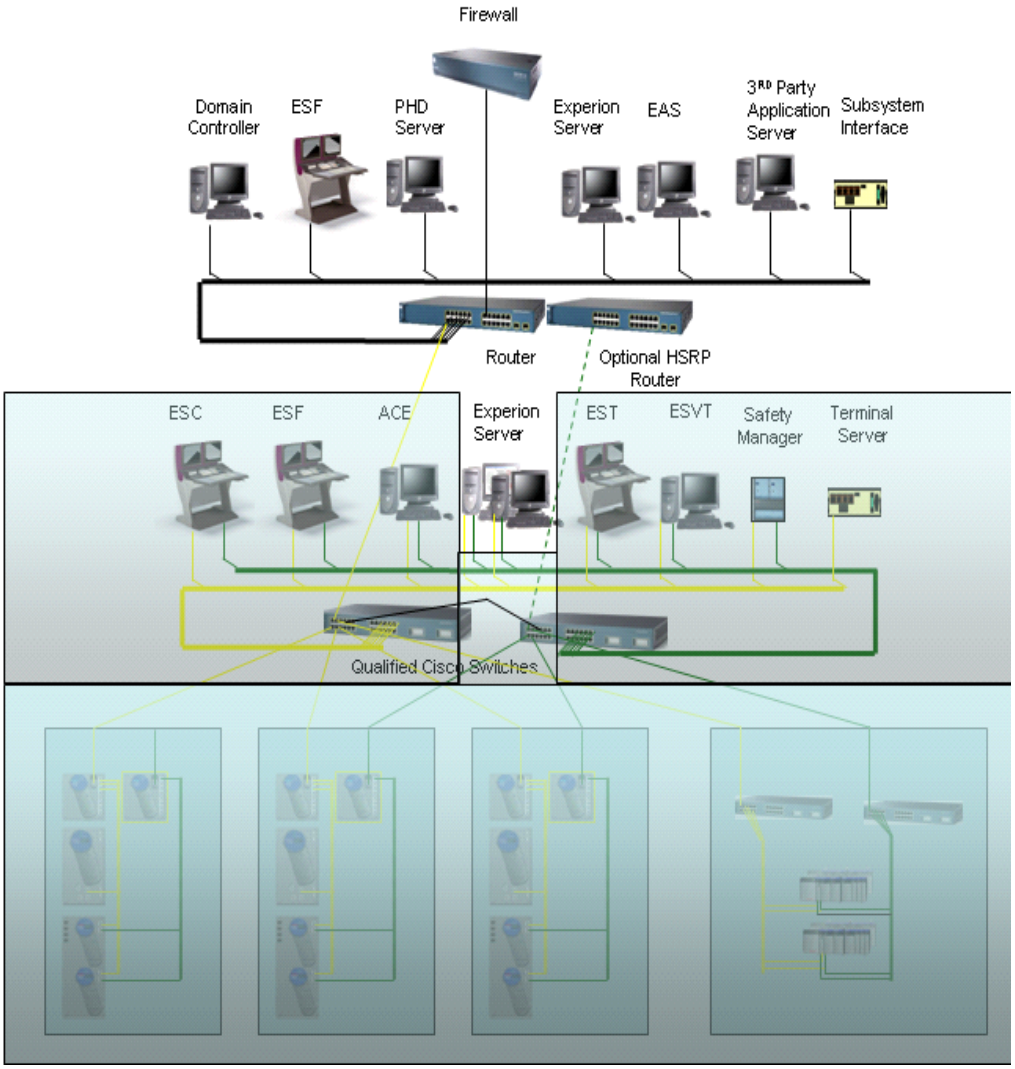

View of L2 from L3 with Routing and Filter

- Level 3 Router/Switch (Cisco 3560, 3750 or equivalent)
    - Provides connectivity for L3 devices and L2 networks
    - Has customer-defined route between L3 and L2
        - Routes between Enterprise IPs on L3 to Private L2
    - Implements Access List Filtering
        - Domain Controller / Management (L3 DCs and L2 Nodes requiring authentication)

6. Level 4

Description

Level 4 is not part of the control network. Communication on this level may not be as secure as that on Level 1, Level 2 or Level 3.

Level 4 Best Practices

Because Level 4 is a different security and networking environment, Honeywell strongly recommends that Level 3 and Level 4 be separated by a firewall.

Implementing Level 4 Best Practices

Requirements for a firewall between Level 4 and Level 1, 2 and 3:

- If there is a need to use DSA or any other form of communication with Level 2 that requires Microsoft RPC or DCOM APIs, then the firewall must not use Network Address Translation. See Section 8 for detailed information on IP address range selection, including the use of NAT.
- The firewall should limit communication to only those nodes on Level 4 that require access to nodes on Level 3 or Level 2.
- Level 1 nodes must not be allowed to communicate with nodes on Level 3 or on Level 4.
- Level 1 nodes may only communicate with Level 2 nodes on the same subnet.

Router

The router-to-firewall connection should be a single point of connectivity. This will enable higher security and improved management. A major advantage is the user just needs to pull a single cable to make an “air gap” between Level 3 and Level 4. The connection to the firewall isolates Enterprise LAN Broadcast and Multicast traffic while enabling connectivity between the PCN and Enterprise LAN.

Firewall

The firewall implements a restrictive security policy for traffic between Level 4 and Level 3. The firewall should deny all access to the PCN unless it is explicitly permitted. A best practice is to use IP address source and destination filtering. Only specific nodes on the enterprise network are permitted to communicate with specific nodes on the PCN. Permitted traffic must be limited to Server – Server traffic only (e.g., Experion Server or PHD). TCP Port Filtering is recommended to stop denial-of-service attacks to well-known ports

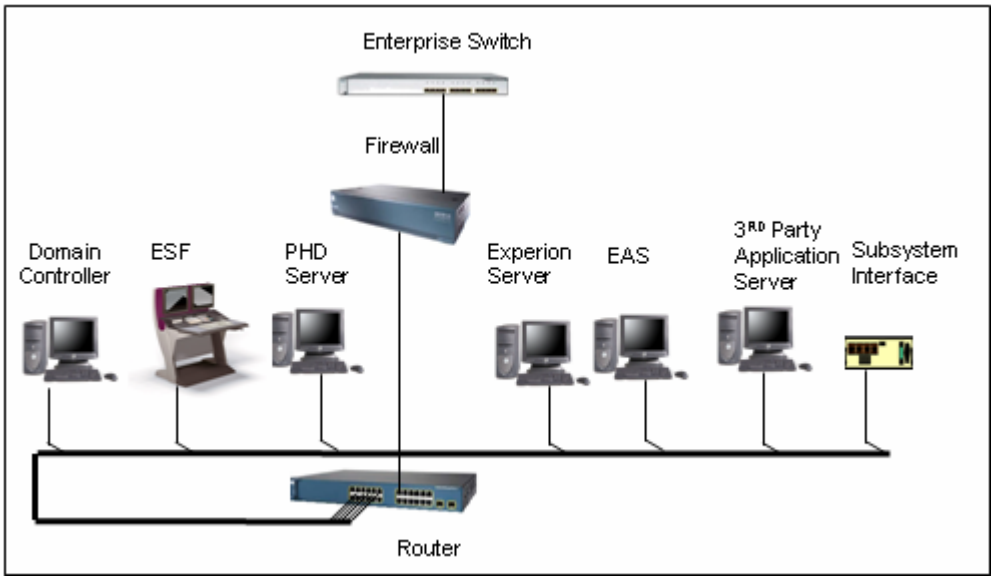

Process Control Network to Business Network

DMZ

Systems that require L4 nodes to access data on L3 or possibly L2 are recommended to use a “DMZ” or Level 3.5. Further, only nodes on L3.5 are allowed access from L4. These nodes are also accessible from L3 and L2 if necessary. Data for enterprise servers can be obtained by having an Experion server in L3.5 with DSA access up to L4 and down to L3. Terminal servers and virus update file servers can also be placed in the DMZ. The DMZ can either be a third leg on the firewall, or a separate network between L4 and L2 with a firewall between both L3.5 and L4 and L3.5 and L3. Further discussion of this best practice can be found in the Honeywell Security Planning Guide document.

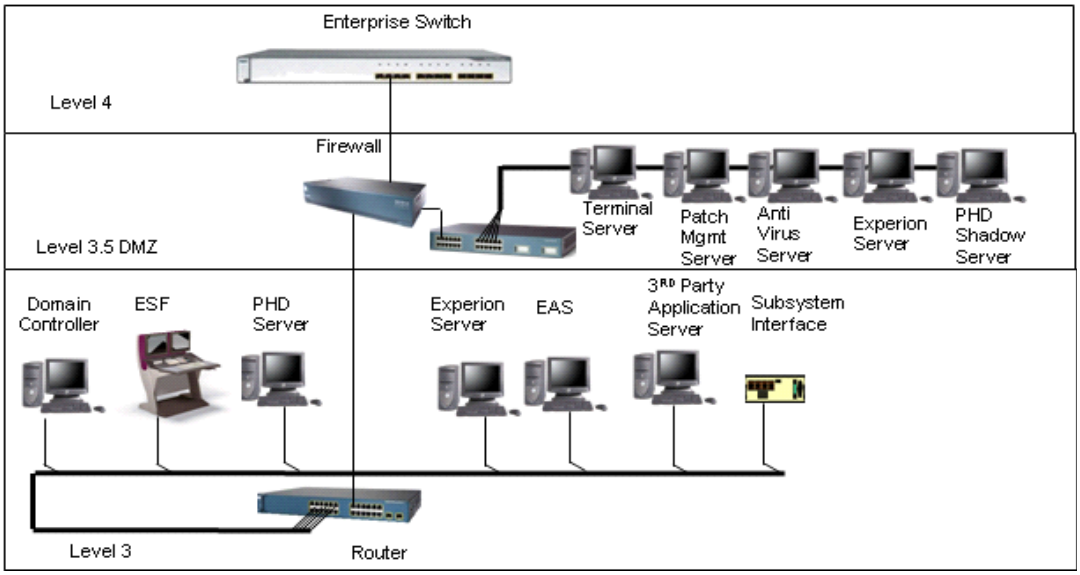

L3 to L4 connection with DMZ

7. Variations on Best Practice

Remote Locations

It may be necessary due to geographic limitations to make certain changes to the best practice architecture. It may be necessary to add one L2 console station node at a satellite control area for a roving operator to have a view to the process, or in case of a catastrophic break in the communications paths to the control room. In this case, it is acceptable to put the L2 station directly on the L1-configured switches.

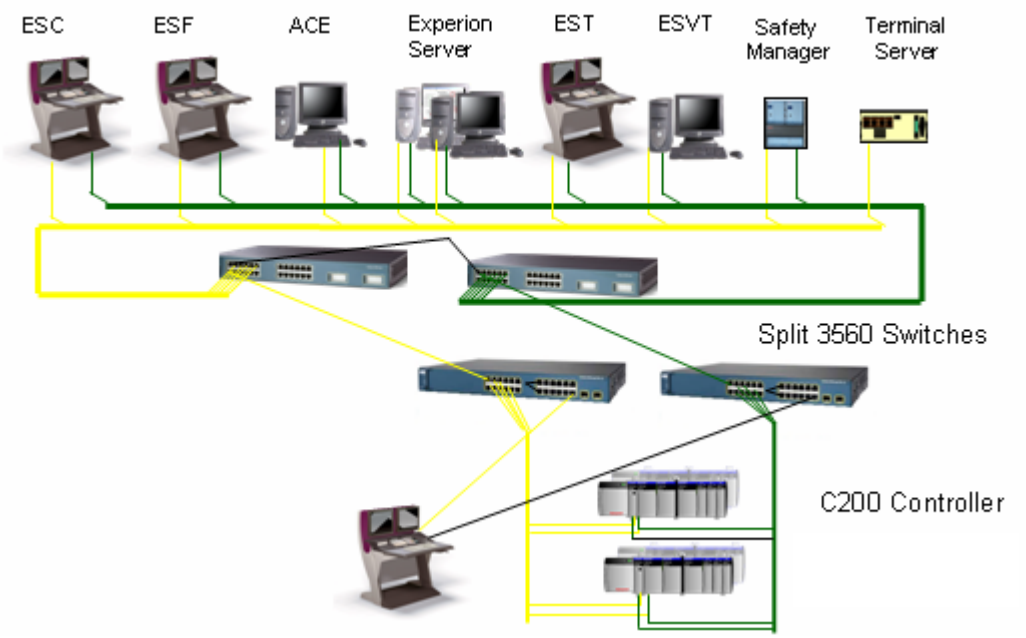

System with Console on Split L1|L2 Switches

But if it is necessary to have multiple L2 nodes at the remote location, then it is recommended that separate switches be used for the L1 controllers with uplinks to the switches where the servers and stations reside. The flow of data should be from L1 switches to local L2 switches then to the top-level switch pair at the central location. Or, a pair of switches with a split L1|L2 configuration could be used, where one section of a switch has L1 configuration and the other section has L2 configuration. Switches that can support L1|L2 split configuration are listed in FTE Specification EP03-500-300 (and later).

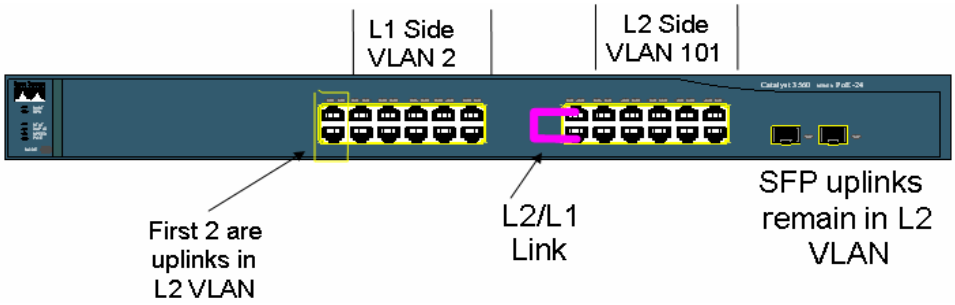

Split Switch Configuration

- Switch is split in two pieces- one for L2, one for L1
- A New VLAN is created for the L1 side, L2 uses the FTE community VLAN
- A cross-level cable connects the two VLANs and L2 to L1. It must be a crossed cable.
- Spanning tree is configured to prevent blocking between sides
- Filtering on the input to the L1 side passes all CDA TCP ports and all established traffic, all UDP and NTP.
- Multicast policing @ 2 Mbps and broadcast storm limits at 1 Mbps are configured

Small Experion Systems with FTE

The Experion system is expandable from very small systems with only a few nodes to very large multi-cluster and multi-FTE community installations. For small systems where all the FTE units are co-located, the best practice topology can be less restrictive to save cost. In this case, all units can be on the same switches. The split switch configuration file would again be used for this installation. When the installation requires multiple layers of switches or is geographically spread, then the Honeywell best practices should be followed.

Small system with single layer of switches.

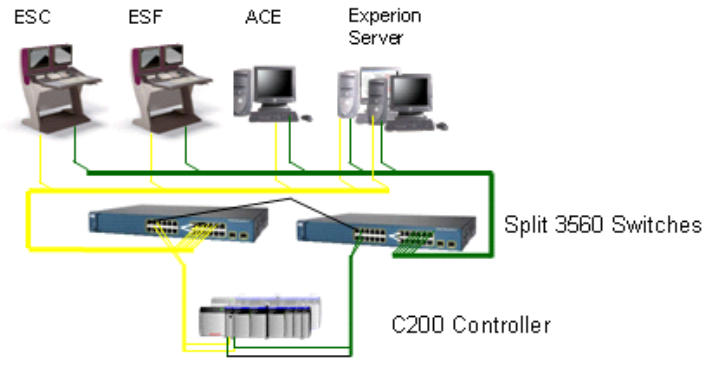

8. DVM Best Practices

The Digital Video Manager is capable of consuming a great deal of bandwidth depending on the configuration. For this reason, the recommendation for DVM is as follows:

- Create a separate L3 subnet for the cameras and DVM server
- Utilize separate display nodes in this subnet for heavy traffic DVM displays
- Limit the traffic in Station and Server nodes to less than 20% of the bandwidth
- Baseline CPU utilization for required DVM displays
- Always use unicast for DVM. Multicast will trip storm limits in the Cisco switches

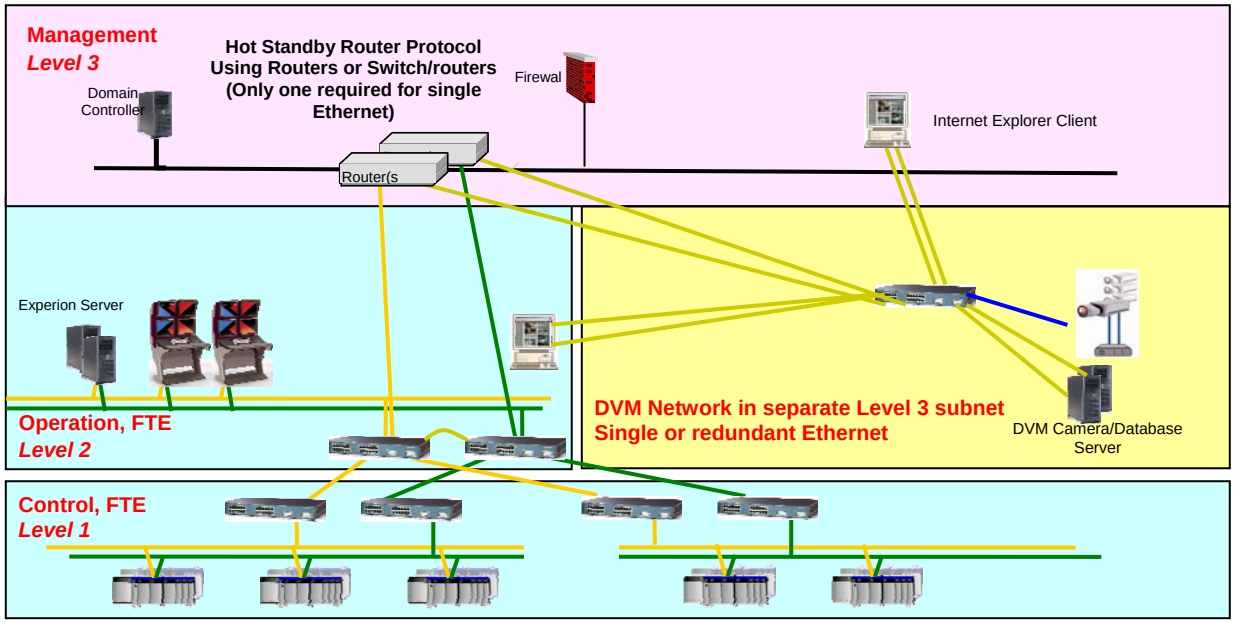

DVM Network

9. IP Addressing

With Experion controller nodes in the FTE community, communication reliability and security must be carefully planned.

The best form of security is “air gap”; that is, no connection between the control LAN and any other users in the plant. Unfortunately for security, most installations must have some form of communication between the control LAN and the plant LAN, so we must pay careful attention to IP Address management. Honeywell has developed several recommendations for IP address range selection to increase the security when connecting the Control LAN to outside communications networks. Another goal is to simplify selection of IP addresses for FTE networks. The examples in this paper all use a range of 10.n.n.n.

The types of networks are:

- Completely isolated FTE community
- Multiple FTE communities isolated from Level 4 networks
- FTE communities connected to Level 4 with private IP addresses
- FTE communities connected to Level 4 with corporate IP addresses

Completely Isolated FTE community

Even if there is complete isolation of the control LAN from the IT LAN, IP address ranges and rules should follow the best practices of the multiple isolated or DSA-connected communities described below. If the network expands so that a router is needed at a later day, the IP addresses will already conform to the Honeywell best practices for connected networks.

Multiple FTE Communities Isolated from Level 4 Networks

Plant-wide networks may contain several FTE communities connected by routers. If this network arrangement is isolated from the IT LAN, then Honeywell recommends that private IP addresses be used.

For ease of configuration, a simple address range of 10.CN.X.Y can be used for IP address distribution. CN stands for FTE community number. Multiple FTE communities can be connected together with a router. For example, the first FTE community subnet could be 10.1.x.y; the second could be 10.2.x.y, etc.

FTE Communities Connected to Level 4 with Private IP Addresses

For a plant-wide network that has a Level 3 network that connects multiple FTE communities and other plant Ethernet based nodes, Honeywell recommends using private IP addresses with Network Address Translation (NAT) for communication with Level 4. The NAT can be accomplished with a firewall: Honeywell recommends dedicated firewall equipment from Cisco. A Windows-based computer with firewall software is not recommended.

The private address distribution is similar to the previous scheme where the FTE communities are 10.1.X.Y, 10.2.X.Y, etc. X stands for the range of addresses where the two types of nodes exist. The servers must be in a separate range from other L2 nodes. An example would be 10.1.0.Y for server nodes, 10.1.1.Y for station nodes, and 10.1.2.Y for any other nodes such as ACE, PHD and third party IP-based nodes. Y stands for any address between 1 and 255. If the FTE community is connected to a router, the router interface IP address should be in the range where the servers are configured. In the above example, the router interface IP address would be 10.1.0.1.

Level 1 nodes should be in the address space above the other nodes on L2 and outside of the range of the subnet mask of the router interface, but within the subnet mask of the nodes that need to communicate. Thus, using the previous examples, the L1 addresses would appear in the range 10.0.4.Y. Nodes on L3 must not be able to communicate with the L1 nodes. The nodes will have the following subnet masks:

- L2 Servers and console stations with communication to L1 nodes: 255.255.248.0
- L2 nodes with no communication with L1 nodes: 255.255.252.0
- L1 controller nodes: 255.255.248.0
- L3 router interface to L2 255.255.252.0

IP Addressing Level 4 Connected Networks with Corporate IP Addresses

When it is necessary to use IP addresses from the corporate allocation, the L2/L3 addresses must be unique and compatible with L4 addresses, and NAT cannot be used. To minimize the number of corporate IP addresses used, the above addressing scheme cannot be used. Honeywell recommends a method that conserves addresses but is more difficult to configure, which is to obtain a subnet size that will cover all of the L2 nodes. The server range is contained in the lower addresses and the other L2 nodes would start on a power of 2 boundary. This is necessary so that the ACL filter used in the router to limit full access to the server nodes can be configured with a subnet mask defining the server range.

The following is an example of a FTE community subnet containing:

- 5 servers
- 10 stations
- 2 ACE
- 10 terminal servers
- 10 Controllers with FTEB

A range of addresses is obtained from the corporate range, which for this example is 164.1.0.0 with enough addresses for 126 nodes, the subnet default gateway and the subnet broadcast address. The address distribution would be:

164.1.0.1	The routed interface IP address with subnet mask of 255.255.255.192, enough for 62 usable nodes, the subnet mask and the subnet broadcast address.
164.1.0.2-15	Server nodes (5 servers 2 addresses each starting at address 2 rounded up to power of 2). The subnet mask is 255.255.255.128 to cover both L2 and L1 nodes
164.1.0.16-63	Stations, ACE terminal servers plus some spares. The subnet mask is 255.255.255.128, to cover the L2 and L1 nodes
164.1.0.64-127	FTEB (controller addresses must be outside of the subnet mask of the router interface). The subnet mask is 255.255.255.128, to cover the L1 and L2 range
164.1.0.64-127	The router interface to the FTE community blocks all access from L3 by the subnet mask of 255.255.255.192.

10. IP Address Reuse

L1 devices have the potential to consume many thousands of IP addresses in a corporate IP address space. To conserve Corporate IP addresses, an address reuse scheme is recommended by Honeywell. Only systems that have a need for address reuse should employ this IP addressing scheme. Systems that do not have this requirement must use one of the IP addressing schemes discussed above.

One range of addresses for L1-only should be requested from the corporate pool. This range can be reused in other FTE communities that are separated by a router. This range must be large enough to accommodate all of the L1 nodes on this subnet, both now and in the future. If a subnet is later added with a larger number of L1 nodes than the range obtained originally, then a new range must be requested. Existing L1 nodes would not need to have their addresses changed.

For L2 nodes that must communicate with L1 nodes in the reusable address space, a “route add” command must be configured in each such L2 node. A new service has been added for automatic insertion of the static route. This service is loaded with Experion Servers, direct consoles and ACEs. The service runs on node startup and queries the server for the address range and subnet mask of the controllers. If the address of the node running the service is not in the range of the controllers, then the static route to the controller will be added to the Yellow interface. The service will test every 10 minutes for changes in the server data base and to be sure the static route is still connected to the Yellow interface. Any errors or problems will be notified in the application event log.

For nodes prior to R300, a static route must be added by hand or by a batch file that runs at node startup. Nodes that do not communicate with the L1 nodes do not need the “route add”. The following example has the L2 address range of 164.1.0.0-164.1.7.255 and the L1 address is 164.0.0.0 – 164.0.2.255. The command for an L2 node would be:

Route ADD 164.0.0.0 MASK 255.255.252.0 164.1.3.10 –p

- 164.0.0.0 is the base address of the L1 subnet programmed in Control Builder
- 255.255.252.0 allows 1024 L1 FTE nodes
- 164.1.1.10 is the Yellow interface IP address of the node being configured with the route add.
- -p makes it persistent across reboots.

The L1 nodes will receive the address range of the L1 nodes and the L2 nodes. The L1 nodes will then calculate and add a static route to their IP stack to enable communication with L2. For releases prior to R300, in order for L1 nodes to communicate with L2, the L2 address range must be a subset of the L1 range so that a subnet mask will allow the L1-L2 connection. For the above example, if the L2 address range is 164.1.0.0 – 164.1.7.255, then the L1 range in the Route Add example would start at 164.0.0.1. A subnet mask of 255.0.0.0 can be set in L1 nodes via Control Builder and communications will be open to the L2 addresses. The range can be larger than the actual L2 address range because communications will not go outside of the FTE community subnet.

11. TPS Upgrade Best Practice

Existing TPS systems have the ability to add Experion capabilities with the ESVT, ES-T, and ACE. TPS nodes that are currently connected to an Ethernet Plant Control Network (PCN) can be connected to the FTE network in one of 3 ways.

- The PCN is a stand alone network, that is, it has only control system nodes connected to the switch(es). In this case, the top of the PCN network can be connected to the top of the FTE switch tree. The yellow switch is recommended.
- The PCN is part of a plant wide network. In this case, the FTE network must be connected to the L3 network through the existing router with the required filtering described in this document on the interface that connects to the FTE network. If the plant wide network is a single network, meaning there is no router, or the existing router does not have the required filtering capability, then the FTE network must connect to L3 through a firewall with the same required filtering.
- A conversion of the PCN to FTE. In this case, qualified FTE switches must replace existing PCN switches.

12. Example Cisco Router Configuration Statements

In order to configure the FTE community filtering requirements in Cisco routers the following configuration commands are used. Cisco uses an Access Control List (ACL) to describe what should pass or not pass through an interface.

Below is an example of a set of ACLs used to accomplish the filtering:

access-list 101 permit tcp 10.0.0.0 0.0.0.255 any established	Established connections are allowed in the whole FTE community subnet The range of addresses in this FTE community is 10.0.0.2-255.
access-list 101 permit udp host 225.7.4.103 any access-list 101 permit udp any host 225.7.4.103	The DSA multicast address, 225.7.4.103 is allowed to pass in both directions.
access-list 101 permit ip 10.0.0.0 0.0.0.240 any access-list 101 permit ip any 10.0.0.0 0.0.0.240	The server range is 10.0.0.2-15.
access-list 101 permit udp any any eq domain	Access to a domain controller TCP port is allowed.
access-list 101 permit udp any any eq 88	Access to a Kerberos server is allowed
access-list 101 permit udp any any eq 389	Access to a LDAP server is allowed
There is an assumed “deny all” at the end of the list. This means that any other address range is denied access.	

Router interfaces connected to FTE communities MUST NOT have VLANs associated with them. The following is a typical interface configuration

interface FastEthernet0/3	This example has a connection to a 3560 interface in the third fastethernet port.
No switchport	This configuration statement will create a routed port for the FTE community
duplex full speed 100	The speed and duplex if the interface is fixed to avoid problems with autosensing.
ip address 10.0.0.1 255.255.255.0	The FTE community’s default gateway address is 10.0.0.1. The subnet mask of 255.255.255.0 will allow traffic in this range to pass to the ACL filters
ip access-group 101 out	Access-group 101 uses the ACLs described above in access-list 101
no ip proxy-arp	Proxy arp must be disallowed to avoid possible issues
ip pim dense-mode	PIM dense-mode is needed for the DSA multicasts to be routed.

13. Switch Configuration Files

Overview

After installation, a Cisco switch pair must be configured for FTE using the switch's command line interface and the correct switch startup configuration file. Switch configuration files, which are copied to the hard drive when the FTE Driver package is installed, are used to configure the various switch and port options as listed in Table 2-2. Additionally, the configuration files contain Quality of Service parameters that are attached to the ports.

Cisco Switch and Port Options	
Option Types	**Available Options**
Switch options	NE-SW224S (C2960-24TC-L)**
	NE-SW248S (C2960-48TC-L)**
	NE-SW324S (C3560-24TS-S)**
	NE-SW312S (C3750G-12S-S)
	NE-SW512C (C2955C-12)
	(switches supported but no longer offered)
	NE-SW224G (C2950G-24-EI)
	NE-SW248G (C2950G-48-EI)
	NE-SW312G (C3550-12G)
	NE-SW324F (C3550-24-FX-SMI)
Port configuration options	Number of uplink ports
	Number of full duplex auto speed FTEB ports
	Number of full duplex 100 Mbps ports
	Whether the switch ports have VLAN101 configured.
	L1 only, L2 only, or **L1|L2 split

Configuration Order for Switch Ports

The chosen configuration file defines the switch options and how each switch port is configured. Uplink ports are configured first, FTE Bridge ports are configured second, and Full Duplex 100 Mbps ports are configured third. The following table summarizes the switch port configuration settings. Complete descriptions of the switch configuration files can be found in the FTE Overview and Implementation Guide found in the Experion Knowledge Base.

Configuration Order	Port Type	Spanning Tree	Status	Duplex	Speed
1st	Uplink ports	Uplink Fast	Enable	Full	100 Mbps
2nd	FTE Bridge ports	Fast	Enable	Full	Auto
3rd	FTE	Fast	Enable	Full	100 Mbps

More Information

For more information on any of Honeywell's Products, Services, or Solutions, visit our website www.honeywell.com/ps, or contact your Honeywell account manager.

Automation & Control Solutions
Process Solutions
Honeywell
2500 W. Union Hills Dr.
Phoenix, AZ 85027
Tel: +1-602-313-6665 or 877-466-3993
www.honeywell.com/ps

WP-07-02-ENG
June 2007

Honeywell