



CHAPTER 3

Basic Network Design

Overview

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

Assumptions

This chapter has the following starting assumptions:

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

IP Addressing

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

Static IP Addressing

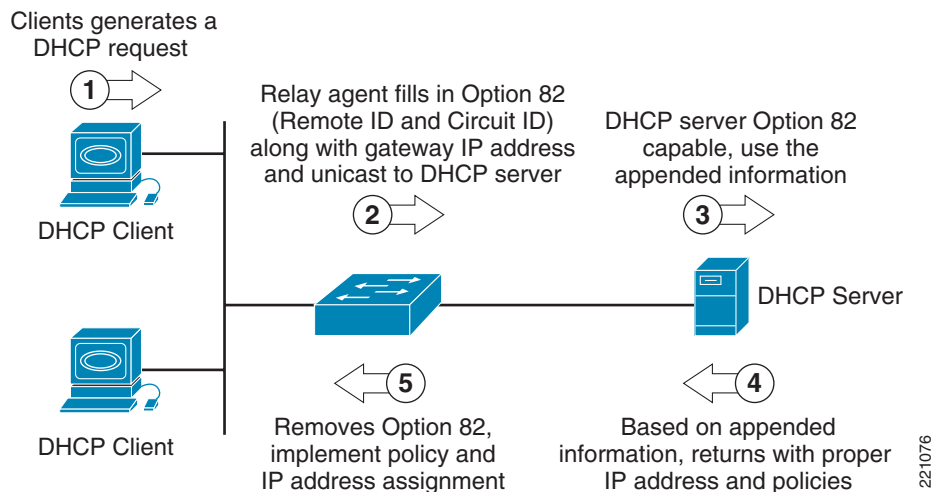
In the manufacturing zone, the level 3 workstations and servers are static. Additionally, it is recommended to statically configure level 2 and level 1 control devices. These servers send detailed scheduling, execution, and control data to controllers in the manufacturing zone, and collect data from the controllers for historical data and audit purposes. Cisco recommends manually assigning IP addresses to all the devices including servers and Cisco networking equipment in the manufacturing

zone. For more information on IP addressing, see *IP Addressing and Subnetting for New Users* at the following URL:
http://www.cisco.com/en/US/customer/tech/tk365/technologies_tech_note09186a00800a67f5.shtml. In addition, Cisco recommends referencing devices by their IP address as opposed to their DNS name, to avoid potential latency delays if the DNS server goes down or has performance issues. DNS resolution delays are unacceptable at the control level.

Using Dynamic Host Configuration Protocol and DHCP Option 82

Dynamic Host Configuration Protocol (DHCP) is used in LAN environments to dynamically assign host IP addresses from a centralized server, which reduces the overhead of administrating IP addresses. DHCP also helps conserve limited IP address space because IP addresses no longer need to be permanently assigned to client devices; only those client devices that are connected to the network require IP addresses. The DHCP relay agent information feature (option 82) enables the DHCP relay agent (Catalyst switch) to include information about itself and the attached client when forwarding DHCP requests from a DHCP client to a DHCP server. This basically extends the standard DHCP process by tagging the request with the information regarding the location of the requestor. (See [Figure 3-1](#).)

Figure 3-1 DHCP Option 82 Operation



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

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

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

For details on DHCP features, see the following URL:

http://www.cisco.com/en/US/products/ps7077/products_configuration_guide_chapter09186a008077a28b.html#wp1070843

**Note**

The DHCP option 82 feature is supported only when DHCP snooping is globally enabled and on the VLANs to which subscriber devices using this feature are assigned.

**Note**

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

IP Addressing General Best Practices

IP Address Management

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

Address Space Planning

When planning address space, administrators must be able to forecast the IP address capacity requirements and future growth in every accessible subnet on the network. This is based on many factors such as number of end devices, number of users working on the floor, number of IP addresses required for each application or each end device, and so on. Even with plentiful availability of private address space, the cost associated with supporting and managing the IP addresses can be huge. With these constraints, it is highly recommended that administrators plan and accurately allocate the addressing space with future growth into consideration. Because the control traffic is primarily confined to the cell/area zone itself, and never crosses the Internet, Cisco recommends using a private, non-Internet routable address scheme such as 10.x.y.z, where x is a particular site, y is a function, and z is the host address. These are guidelines that can be adjusted to meet the specific needs of a manufacturing operation. For more information on private IP addresses, see RFC 1918 at the following URL: <http://www.ietf.org/rfc/rfc1918.txt>.

Hierarchical Addressing

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

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

Summarization (also known as supernetting) allows aggregation of all the host and device individual IP addresses that reside on that network into a single route. Route summarization is a way of having single IP address represent a collection of IP addresses, which can be very well accomplished when hierarchical addressing is used. By summarizing routes, you can keep the routing table entries small, which offers the following benefits:

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

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

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

**Note**

Overlapping IP addresses should be avoided in the manufacturing cell/area zone. If two devices have identical IP addresses, the ARP cache may contain the MAC (node) address of another device, and routing (forwarding) of IP packets to the correct destination may fail. Cisco recommends that automation systems in manufacturing should be hard-coded with a properly unique static IP address.

**Note**

Cisco recommends that the traffic associated with any multicast address (224.0.0.0 through 239.255.255.255) used in the manufacturing zone should not be allowed in the enterprise zone because the EtherNet/IP devices in the manufacturing zone use an algorithm to choose a multicast address for their implicit traffic. Therefore, to avoid conflict with multicast addresses in the enterprise zone, multicast traffic in the manufacturing zone should not be mixed with multicast traffic in the enterprise zone.

Centralized IP Addressing Inventory

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

Routing Protocols

Routers send each other information about the networks they know about by using various types of protocols, called routing protocols. Routers use this information to build a routing table that consists of the available networks, the cost associated with reaching the available networks, and the path to the next hop router. For EtfF 1.1, routing begins at the manufacturing zone, or distribution layer. The Catalyst 3750 is responsible for routing traffic between cells (inter-VLAN), or into the core, or DMZ. No routing occurs in the cell/area zone itself.

Selection of a Routing Protocol

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

Distance Vector versus Link-State Routing Protocols

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

Classless versus Classful Routing Protocols

Routing protocols can be classified based on their support for variable-length subnet mask (VLSM) and Classless Inter-Domain Routing (CIDR). Classful routing protocols do not include the subnet mask in their updates while classless routing protocols do. Because classful routing protocols do not advertise the subnet mask, the IP network subnet mask should be same throughout the entire network, and should be contiguous for all practical purposes. For example, if you choose to use a classful routing protocol for a network 172.21.2.0 and the chosen mask is 255.255.255.0, all router interfaces using the network 172.21.2.0 should have the same subnet mask. The disadvantage of using classful routing protocols is that you cannot use the benefits of address summarization to reduce the routing table size, and you also lose the flexibility of choosing a smaller or larger subnet using VLSM. RIPv1 is an example of a classful routing protocol. RIPv2, OSPF, and EIGRP are classless routing protocols. It is very important that the manufacturing zone uses classless routing protocols to take advantage of VLSM and CIDR.

Convergence

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

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

Routing Metric

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

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

Scalability

As the network grows, a routing protocol should be capable of handling the addition of new networks. Link-state routing protocols such as OSPF and hybrid routing protocols such as EIGRP offer greater scalability when used in medium-to-large complex networks. Distance vector routing protocols such as RIPv1 and RIPv2 are not suitable for complex networks because of the length of time they take to converge. Factors such as convergence time and support for VLSM and CIDR directly impact the scalability of the routing protocols.

Table 3-1 shows a comparison of routing protocols.

Table 3-1 Routing Protocols Comparison

Name	Type	Proprietary	Function	Updates	Metric	VLSM	Summarization
RIP	Distance vector	No	Interior	30 sec	Hops	No	Auto
RIPv2	Distance vector	No	Interior	30 sec	Hops	Yes	Auto

Table 3-1 Routing Protocols Comparison (continued)

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

In summary, the manufacturing zone usually has multiple parallel or redundant paths for a destination and also requires VLSM for discontinuous major networks. The recommendation is to use OSPF or EIGRP as the core routing protocol in the manufacturing zone. For more information, see the Cisco IP routing information page at the following URL:

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

Static or Dynamic Routing

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

Server Farm

Types of Servers

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

- Servers that provide common network-based services such as the following:
 - DNS—Primarily used to resolve hostnames to IP addresses.
 - DHCP—Used by end devices to obtain IP addresses and other parameters such as the default gateway, subnet mask, and IP addresses of DNS servers from a DHCP server. The DHCP server makes sure that all IP addresses are unique; that is, no IP address is assigned to a second end device if a device already has that IP address. IP address pool management is done by the server.
 - Directory services—Set of applications that organizes and stores data about end users and network resources.
 - Network Time Protocol (NTP)—Synchronizes the time on a network of machines. NTP runs over UDP, using port 123 as both the source and destination, which in turn runs over IP. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. An

NTP client makes a transaction with its server over its polling interval (64–1024 seconds,) which dynamically changes over time depending on the network conditions between the NTP server and the client. No more than one NTP transaction per minute is needed to synchronize two machines.



Note For more information, see *Network Time Protocol: Best Practices White Paper* at the following URL:
http://www.cisco.com/en/US/customer/tech/tk869/tk769/technologies_white_paper09186a080117070.shtml

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



Note For more information on CS-MARS, see the CS-MARS introduction at the following URL:
http://www.cisco.com/en/US/customer/products/ps6241/tsd_products_support_series_home.html

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



Note For more information, see the Cisco Network Assistant general information at the following URL:
http://www.cisco.com/en/US/customer/products/ps5931/tsd_products_support_series_home.html

- CiscoWorks LAN Management Solution (LMS)—CiscoWorks LMS is a suite of powerful management tools that simplify the configuration, administration, monitoring, and troubleshooting of Cisco networks. It integrates these capabilities into a best-in-class solution for the following:
 - Improving the accuracy and efficiency of your operations staff
 - Increasing the overall availability of your network through proactive planning
 - Maximizing network security

**Note**

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

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

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

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

Server Farm Access Layer

Access Layer Considerations

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

Layer 2 Access Model

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


Spanning VLANs across Access Layer switches

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

- Use Rapid PVST+ as the version of STP. When spanning tree convergence is required, Rapid PVST+ is superior to PVST+ or plain 802.1d.
- Provide an L2 link between the two distribution switches to avoid unexpected traffic paths and multiple convergence events.
- If you choose to load balance VLANs across uplinks, be sure to place the HSRP primary and the STP primary on the same distribution layer switch. The HSRP and Rapid PVST+ root should be co-located on the same distribution switches to avoid using the inter-distribution link for transit.

For more information, see *Campus Network Multilayer Architecture and Design Guidelines* at the following URL:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns656/c649/cdcont_0900aecd804ab67d.pdf

Figure 3-2 shows an example of a Layer 2 access topology.



Layer 2 Adjacency Requirements

When Layer 2 adjacency exists between servers, the servers are in the same broadcast domain, and each server receives all the broadcast and multicast packets from another server. If two servers are in the same VLAN, they are Layer 2 adjacent. There are certain features such as private VLANs that allow groups of Layer 2 adjacent servers to be isolated from each other but still be in the same subnet. The requirement of Layer 2 adjacency is important for high availability clustering and NIC teaming.

NIC Teaming


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

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

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

Figure 3-3 shows examples of NIC teaming using SFT and ALB.

Figure 3-3 NIC Teaming



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

