CISCO SYSTEMS

Intelligent **Ethernet** Factory Network Design

Factory automation networks, whether small or large, need an intelligent network infrastructure for availability and scalability, IP Multicast, quality of service (QoS), and network security. The network foundation hosting these technologies for a factory network needs to be robust, scalable, secure, and manageable, and must provide high availability. This network foundation is designed to run a convergence of commercial and industrial applications over a common IP infrastructure, with due consideration for QoS, bandwidth, latency, and high performance in automation applications.

This paper is an implementation guide for deploying an intelligent Ethernet factory network infrastructure. It looks beyond a single-box approach, and focuses on designing best practices. It also makes an effort to highlight topology, suggest best practice recommendations, show actual device configuration, and explain the features and protocols required to build a converged voice, video, and data network. It will also address product recommendation and suggest an upgrade path to accommodate future growth.

This document concentrates on the access layer (switches on the factory floor), and provides an overview on many topics. Please refer to the Cisco Systems® Software Configuration Guide and Command Reference for particular switches and routers, or for specific functions or commands.

Network Design Best Practices

1. Build a hierarchical network model
2. Design the factory backbone with Layer 3 protocols
3. Limit size of Layer 2 domains to factory cells
4. Provide interdomain (Layer 3 routing) at the distribution layer
5. Provide redundant paths for all access layer (Layer 2) switches

A hierarchical network design distributes networking functions at each layer through a layered organization. Modular designs are made out of building blocks. Modules and can be added or removed without redesigning the network. A modular design is also easier to grow and troubleshoot. The multilayer design Cisco® uses is an example of a modular hierarchical design model. The key elements of the structured hierarchy are the core, distribution, and access layers in a network.

With the advent of hardware-accelerated Layer 3 switching offering intelligent network services and routing at Layer 2 switching rates, there is no reason to extend a Layer 2 domain across the factory. Cisco recommends designing the core and distribution layers with Layer 3 protocols that provide Layer 2 handoff to industrial Ethernet switches on the factory floor

(access layer). Layer 3 has many advantages derived from routing protocols such as load balancing, optimal path selection, summary addressing, and deterministic traffic pattern in case of failure. A Layer 3 boundary also limits the size of the spanning-tree and broadcast domain, and provides a Layer 2 loop-free topology so that all links actively forward traffic.
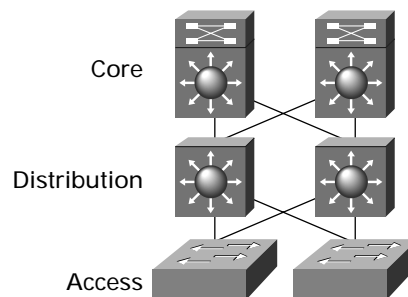
Routing protocols also offer load balancing, path redundancy, and fast failover without the complexity of spanning-tree. A good multilayer design in factory networks uses a balance between both Layer 2 and Layer 3 technology without adding the complexity of either.

In a structured design, one IP subnet maps to a single VLAN, which is carried to the factory floor work cell switch. A good IP addressing scheme can take advantage of Layer 3 features to exchange summarized routing information, rather than learning the path to every host in the whole network. Summarization is critical to the scalability of routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP). Each Layer 3 switch also provides multicast routing between the work cell VLANs. An example of subnet and VLAN numbering scheme is discussed later.

Physical Topologies

Factory network topologies generally fall into one of two designs: tree and ring. Figure 1 depicts a dual-homed tree network topology. In this example, each access layer switch has two uplinks. Each uplink goes to a separate distribution layer switch. Using Rapid Spanning-Tree Protocol (RSTP) and Multi-Instance Spanning Tree Protocol (MSTP), these two uplinks provide redundancy and load balancing. RSTP and MSTP will be discussed in the next section. The tree topology provides for a stable network design, but is geographically limited to the distances that each uplink can support. For example, if the uplinks are copper-based, the distance limitation is 100 meters. By using single-mode fiber (100BASE-FX), uplinks can extend to two kilometers, and multimode fiber (100BASE-LX) can extend to 15 kilometers. In order to cover large factory floors, longer cable runs made need to be created.
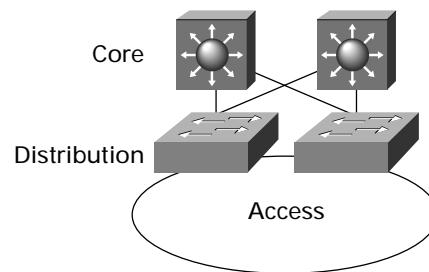
Figure 1
Dual-Homed Tree Network Design

In ring topologies, each access layer switch is placed on a ring so that each uplink connects to another access layer switch (Figure 2). Each access layer switch still has primary and redundant uplinks. However, the primary and redundant paths go through another access layer switch. Should a link between two switches get broken, or if a switch on the ring goes down, traffic can still run around the ring in the other direction. Switches are generally distributed among several rings. The number of switches per ring depends on the physical media being used for the uplinks, the amount of area to be covered, and the amount of traffic to be allowed on the ring. Typically, seven to ten switches are placed on each ring.

Usually, two switches at the top of the ring provide the distribution layer. Generally, these switches support Layer 3 routing functions. The primary advantages to the ring topology is that it saves on distribution-layer switch ports and allows greater areas to be covered with less cabling. However, two failures in a particular ring can isolate switches between those failures and prevent traffic from flowing up to the distribution layer. Also, depending on the number of switches (and devices connected to those switches), the ring may become congested with multicast traffic.

Figure 2
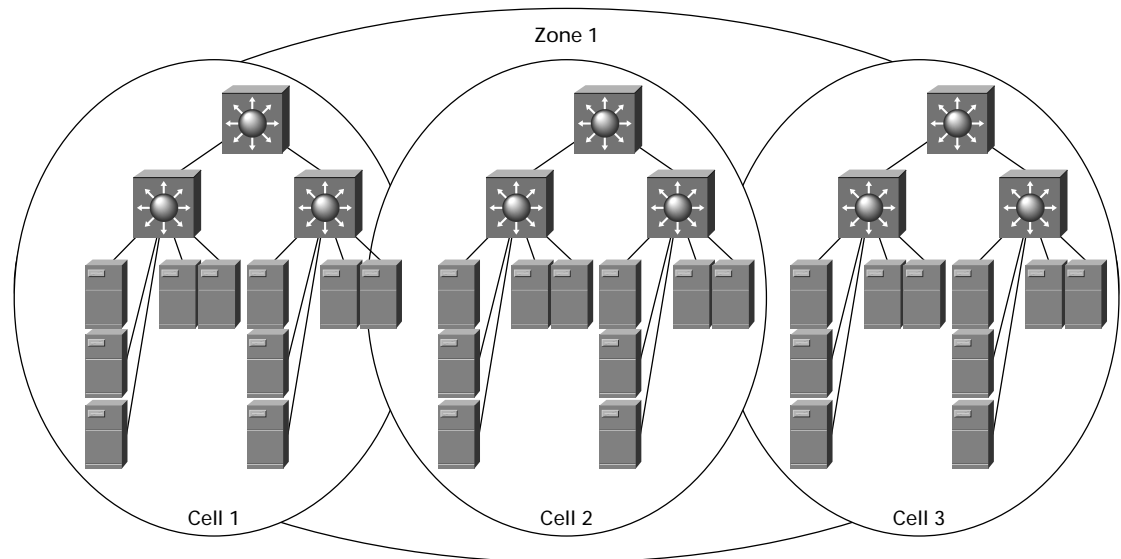Ring Topology

Availability in Factory Network

In factory networks, it is easy to draw similarities between work cells and VLANs. Generally, each work cell comprises a single VLAN. A work cell may also be limited to a single switch. However, there may be cases where there are more than one work cell per switch or a work cell uses more than one switch. Layer 2 trunking allows for the VLAN information to travel from the access layer to the distribution layer (Figure 3).

Factory-floor VLANs should be isolated on the access-layer switches and terminated on the first Layer 3 device. This will limit the size of the Layer 2 broadcast domain. Make sure to prune unnecessary VLANs from each trunk on the access layer switches. IP routing protocols support intelligent forwarding and route determination based on tunable metrics to achieve fast convergence and load balancing in Layer 3 switched environments. Cisco switches provide MSTP (IEEE 802.1s) and RSTP (IEEE 802.1w) advanced spanning-tree functions at Layer 2.

An example might be a set of robotic welders. An assembly line may contain 20–30 welders grouped into work functions. Since the control equipment needs to talk to all the welders, all welders and control equipment may be assigned to the same VLAN. However, in some cases the welders may need to be managed into subgroups (frame, door panel, hood and trunk, for example), where each of these subgroups is assigned its own VLAN.
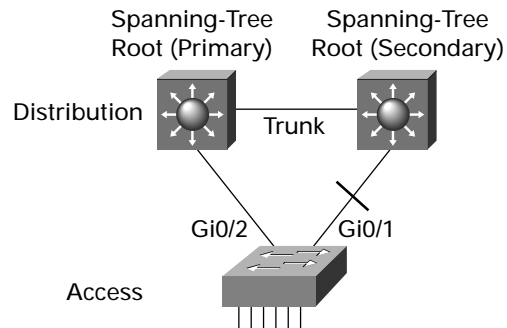
Figure 3
Work Zone Diagram

The recommended design access-layer switches are only configured with the VLANs they need to avoid unnecessary flooding. Trunking provides an optimal way to carry two or more VLANs on a single uplink to the distribution layer switch. If the switch supports Dynamic Trunking Protocol (DTP) the recommended trunk settings are *desirable* in the distribution layer and *auto* at the access layer. Auto makes the port willing to convert the link to a trunk link. The port becomes a trunk port if the neighboring port is set to *on* or *desirable* mode.

Network availability is a function of redundancy. Redundancy is achieved through three main components—external redundancy, (redundant device), internal redundancy (redundant internal sub-systems like power supply), and redundant data paths. Redundancy comes with a percentage of the cost, but increases availability by a bigger percentage of the cost. For example, failure of an access device should not render an entire network useless. Various high-availability protocols and features need to be optimized in factory networks to achieve end-to-end network availability.

Spanning-Tree 802.1s and 802.1w shows a simple redundant link configuration where the access layer switch is dual homed to distribution switches. The IEEE extensions 802.1s and 802.1w to the Spanning-Tree Protocol provide quicker recovery time and uplink trunk load balancing. RSTP (IEEE 802.1w) can provide subsecond convergence times in networks with up to 75 VLANs. By providing redundant paths with rapid failover, the access layer switch can minimize packet loss during a link outage. MSTP (IEEE 802.1s) allows the administrator to assign each VLAN a primary path (possibly different from other VLANs). This way if a particular switch has multiple uplinks, then each uplink can be a primary for one VLAN and the redundant uplink for another VLAN.

Figure 4
Simple Spanning-Tree Diagram

Below are commands for setting up and configuring the Cisco Catalyst® 2955 for MSTP:

```
2955#configure terminal
2955(config)#spanning-tree mode mst
2955(config)#spanning-tree mst config
2955(config-mst)#name cellA123
2955(config-mst)#revision 1
2955(config-mst)#instance 1 vlan 1-5
2955(config-mst)#instance 2 vlan 6-10
2955(config-mst)#interface gig 0/1
2955(config-if)#spanning-tree mst 1 port-priority 1
2955(config-if)#spanning-tree mst 2 port-priority 2
2955(config-if)#int gig0/2
2955(config-if)#spanning-tree mst 2 port-priority 1
2955(config-if)#spanning-tree mst 1 port-priority 2
2955(config-if)#
```

Spanning-Tree PortFast, BPDU Guard, and BPDU Filtering

Spanning-Tree PortFast is a recommended configuration on access ports because it transitions the port directly into forwarding mode after linkup, rather than going through spanning-tree transition states that delay the link bringup by forward delay time (default 15 seconds) at each transition.

```
2955#configure terminal
2955(config)#interface FastEthernet 0/4
2955(config-if)# spanning-tree portfast
```

In a valid configuration, PortFast-enabled ports are connected to edge devices and should not receive Bridge Protocol Data Unit (BPDU) packets. Receiving a BPDU packet on a PortFast-enabled port indicates connection of an unauthorized device to the edge port. BPDU Guard protects the network by disabling the port if a BPDU packet is received. BPDU Guard (when enabled globally) applies to all PortFast-enabled interfaces.

```
2955#spanning-tree portfast bpduguard
```

BPDU Filtering prevents ports that are in a PortFast operational state from sending or receiving BPDUs. The ports still send a few BPDUs at linkup before the switch begins to filter outbound BPDUs. BPDU Filtering should be globally enabled on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a PortFast-enabled port, the port loses its PortFast-operational status, and BPDU Filtering is disabled.

```
2955(config)#spanning-tree portfast bpdufilter default
```

BPDU Guard and BPDU Filtering can also be applied at the interface level. As shown in the example below, the interface configuration overrides the global configuration.

```
2955#configure terminal
2955(config)#interface gigabitEthernet 0/4
2955(config-if)#spanning-tree bpduguard enable
2955(config-if)#spanning-tree bpdufilter enable
```

Spanning-Tree Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports in the event of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network. When the switch is operating in Per-VLAN Spanning-Tree (PVST) mode, loop guard prevents alternate and root ports from becoming designated ports, and the spanning tree does not send BPDUs on root or alternate ports. When the switch is operating in MSTP mode, BPDUs are not sent on nonboundary ports only if the port is blocked by loop guard in all MSTP instances. On a boundary port, loop guard blocks the port in all MSTP instances.

```
2955#configure terminal
2955(config)#interface gigabitEthernet 0/4
2955(config-if)#spanning-tree guard loop
2955(config)#interface gigabitEthernet 0/5
2955(config-if)#spanning-tree guard loop
2955(config-if)#exit
```

Link Aggregation

Cisco Fast EtherChannel® and Gigabit EtherChannel provide bandwidth scalability within the factory by bundling up to eight Fast Ethernet or Gigabit Ethernet ports together to form a channel. Cisco EtherChannel is an effective way of doubling bandwidth by aggregating multiple links, and also provides link redundancy.

```
2955#configure terminal
2955(config)#interface range gig0/1 -2
2955(config-if-range)#switchport mode access
2955(config-if-range)#channel-group 1 mode desirable
Creating a port-channel interface Port-channel1
```

QoS

QoS is a requirement in factory networks for control applications that are sensitive to the delay, loss, and jitter found in data networks. QoS is essentially a buffer-management technique that helps to ensure that high-priority traffic is placed in a separate queue to prevent packet loss, even in times of congestion. By using access control lists (ACLs) to identify control traffic, the switch or router can generate a policer that will mark this traffic as high-priority, placing it into the high-priority queue. An example of commands for use with Rockwell Automation equipment follows:

```
access-list 101 permit udp any eq 2222 any
access-list 101 permit tcp any eq 2222 any
access-list 102 permit udp any eq 44818 any
access-list 102 permit tcp any eq 44818 any
class-map match-all CIP-IMPLICT
  match access-group 101
class-map match-all CIP-Other
policy-map ciptraffic
   class CIP-IMPLICIT
     set ip dscp 40
   class CIP-Other
     set ip dscp 32
```

Unicast, Multicast, and Broadcast Storm Controls

Basically, devices in factory networks do not generate enough traffic to flood the network. However, to help ensure that rogue devices do not flood the network, storm controls should be placed on all access ports. The administrator needs to set these values carefully, as to not impede the normal traffic flow, while protecting the network from unnecesary traffic. Typical automation equipment generates no more than 1–3 Mbps per port. Since storm controls are specified as the percentage of the port speed, this would equate to one to three percent. Below is an example for applying storm controls for an access port:

```
2955(config)#interface fast0/7
2955(config-if)#storm-control action trap
2955(config-if)#storm-control broadcast level 0.1
2955(config-if)#storm-control unicast level 3.0
2955(config-if)#storm-control multicast level 3.0
2955(config-if)#
```

IP Multicast and IGMP Snooping

IP Multicast using Interior Gateway Management Protocol (IGMP) snooping is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to many clients without network or CPU processing overhead. Since the producer-consumer model for data exchange is based on IP Multicast, IGMP snooping is required to keep unwanted traffic flows to a minimum. To enable this feature, use the following commands:

```
2955(config)#ip igmp snooping vlan 2
2955(config)#ip igmp snooping vlan 2 mrouter interface gig 0/1
2955(config)#
```

In addition to IGMP snooping, multicast management can be aided by making sure VLAN pruning is done on the trunk ports. Since multicast traffic needs to flood to other devices across trunk ports, it can be blocked on uplinks that do not extend the particular VLAN that is carrying that traffic. It is important to make sure that IP Multicast routing is enabled on all Layer 3 devices.

Security

As networks enable more applications and become available to more users, they also become prone to greater security threats. As factory networks are connected to the rest of the corporate network, malicious applications or users can cause outages on the factory floor, leading to loss of revenue. Security is an integral part of all factory designs that cannot be overlooked.

Security ACLs

Most factory floor managers do not want to open up the network to general traffic. In this case, ACLs can be used to limit which devices can talk to which other devices and by which protocol. By using an ACL similar to the one mentioned above for QoS, traffic patterns can be classified and explictly permitted or denied on individual ports.

Security ACLs can also limit access to a particular port or switch, based on the MAC address. By using a particular MAC address and a mask value, it is possible to create a filter that would allow only a specific vendor's Programmable Logic Controller (PLC) to be connected a particular port, regardless of the unique MAC address on a single device. Since all vendors are given a certain range of MAC addresses, the first three bytes are the vendor code and can be used in this ACL scenario.

Port Security

Port security allows network administrators to limit port access to particular MAC addresses. It can also be used to limit the number of devices that access a particular port.

802.1x

The function of the IEEE 802.1x standard is similar to the port security feature. In 802.1x, each client device must authenticate via a central server before the access to the port is opened to normal traffic. By using the Extensible Authentication Protocol (EAP), each client device sends the device code and password. These are passed through the access-layer switch onto the authentication server. If the client is authenticated, a message will travel back to the switch telling it to open that port. In most cases, 802.1x is used for public access ports. It is hoped that future factory automation devices include support for 802.1x in their protocol stacks.

MAC Address Notification

In any network (factory or otherwise), the administrator should know when devices join or leave the network. MAC address notification sends a trap message to the network administration console, with the particular MAC address and the port and switch that it is connecting to or disconnecting from. This can aid in intrusion detection or fault management.

Private VLAN Edge (Protected Ports)

Some applications require that no traffic be forwarded between access ports on the same switch so that one device does not see the traffic generated by another device. In such an environment, the use of protected ports helps to ensure that there is no exchange of nnicast, broadcast, or multicast traffic between these access ports on the switch. A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.

Manageability

A factory network needs to use advanced features to ensure that high-priority control data takes precedence over non-real-time traffic. High-availability and fault-tolerance features help prevent the network from costly downtime. Access features like QoS and multicast provide advanced services to the network. Internet edge, security, and voice features on data networks are becoming mainstream in today's factories. The factory network infrastructure provides a foundation for all services and features to work together. Some of the important minimum features required for building a factory network are discussed below.

DHCP Option 82

The Dynamic Host Configuration Protocol (DHCP) has become a standard in factory networks for automating the configuration of end stations and devices like PLCs. DHCP can be used by the `ip helper-address` command in a Layer 3 switched network, where DHCP servers might be on a different VLAN from the client. Option 82 adds the ability to tag the DHCP request packets with the switch and port number that the requester is connected to. For example, this allows for each exchange of damaged equipment since the new device will replace the broken one in the same location—regardless of the MAC address of the original or replacement device, the same IP address will get assigned.

Cisco IE 2100

The Cisco IE 2100 is a configuration-management tool for centralizing router and switch configurations. When used with DHCP Option 82, replacement network devices can get their IP addresses and their configurations from this device. Therefore, the maintenance person does not need to look up and download the configurations during the replacement cycle; it is handled automatically.

Cisco Cluster Management Suite

The Cisco Cluster Management Suite (CMS) is a built-in, Web-based management application offered in fixed-configuration Cisco Catalyst switches ranging from the Cisco Catalyst 1900 to the 3550 (Table 1). Cisco CMS Software is embedded in each switch. Regardless of location, media, or device types supported, up to 16 devices can be managed from a single IP address. Cisco CMS provides configuration and management assistance via a standard Web browser and is designed for small networks.

CiscoWorks

CiscoWorks is a collection of management tools that reside on a central management server. It includes LAN Manager (providing network configuration and monitoring), QoS Policy Manager (QPM) (providing a consistent QoS standard across the network of routers and switches), and other useful tools. CiscoWorks integrates with Cisco CMS to provide device-level configuration and monitoring. CiscoWorks is based on Simple Network Management Protocol (SNMP) and can be integrated with other management tools such as HP Openview.

Table 1 Cisco Catalyst Switch Feature Comparison

	Cisco Catalyst 3550	Cisco Catalyst 2950	Cisco Catalyst 2955
Multilayer services	Layers 2–3 services	Layer 2 services	Layer 2 services
Backplane	Up to 24 Gbps	Up to 13.6 Gbps	6.4 Gbps
Forwarding rate	Up to 17 Mpps	Up to 10.1 Mpps	6.6 Mpps
Rack units	1–1.5	1	N/A (DIN mount)
Redundant power supply	External	External	External
Inline power	Yes	No	No
10/100-Mbps copper ports	24, 48	12, 24, 48	12
10/100/1000 copper ports	2, 10	2	2
100-Mbps fiber ports	2,10, 24	2	2
Gigabit interface converter (GBIC) ports	2, 10	2	0
Class of service to differentiated services code point (CoS->DSCP) remarking	Yes	Yes	Yes
Rate-limiting policers	Yes	Yes	Yes
Weighted round robin (WRR)	Yes	Yes	Yes

Table 1 Cisco Catalyst Switch Feature Comparison

	Cisco Catalyst 3550	Cisco Catalyst 2950	Cisco Catalyst 2955
Weighted Random Early Detection (WRED)	Gigabit ports only	No	No
Strict priority queue	Yes	Yes	Yes
Strict priority scheduling	No	Yes	Yes
Hardware ACLs	Yes	Yes	Yes
Routes in hardware	12,000	N/A	N/A
Routed interfaces, Switched Virtual Interface (SVI)	64-12T/G 32-24/48	N/A	N/A
EIGRP/IGRP/OSPF/Routing Information Protocol (RIP)/Border Gateway Protocol (BGP)	Yes	N/A	N/A
Multicast routing and IGMP snooping	Yes/Yes	No/Yes	No/Yes
VLANs/VLAN IDs	1000/4096	250/4096	250/4096
MAC addresses	12,000	8000	8000

CISCO SYSTEMS