# Integrating the Cisco Catalyst Ethernet Switch Module 3110 (CBS3110) Family for the IBM BladeCenter into the Cisco Data Center Network Architecture

Design Guide

# Contents

# Introduction

This guide provides best design practices for deploying the Cisco® Catalyst® Ethernet Switch Module 3110 (CBS3110) for the IBM BladeCenter enclosure family within the Cisco Data Center Networking Architecture. This guide describes the internal components of the blade server enclosure and Cisco CBS3110 and explores different methods of deployment. It includes the following sections:

- The IBM BladeCenter Enclosure Overview
- The Cisco Catalyst Ethernet Switch Module 3110 Product Overview
- Design Goals
- Cisco Catalyst Ethernet Switch Module 3110 Features
- Network Topology Options
- Design and Implementation Details

### IBM BladeCenter Enclosure Overview

There are several models of IBM BladeCenter Enclosures:

BladeCenter (also now known as BladeCenterE)

BladeCenterT (for Telco applications)

BladeCenterH (newer enclosure supporting high speed I/O slots and servers)

BladeCenterHT (new Telco design based on BladeCenterH)

BladeCenterS (new low end SMB enclosure, announced Fall 2007)

With exception of BladeCenterS, the CBS3110 works in switch bays 1 through 4 in all the remaining enclosures. The BladeCenterS enclosure will be supported by the CBS3012 exclusively. For this document, however, only the BladeCenterH enclosure will be discussed. For more information on the other enclosures contact IBM or look at the IBM BladeCenter website.

The IBM BladeCenter H is IBM's next generation blade based solution, and offers a vast array of networking and server options to end users. Up to 14 servers can be installed into the front of the 9U chassis, where the servers also share a common Media tray. The rear provides access to up to 10 I/O module bays and up to 2 Advanced Management modules. Figure 1 below shows the front, with its 14 server slots, shared media tray to the right, and power supplies above and below the server slots, as well as a snapshot of the rear showing the many I/O module bays, blowers and Advanced management Module bays.

**Figure 1.** Front and Rear View of the IBM BladeCenterH Enclosure

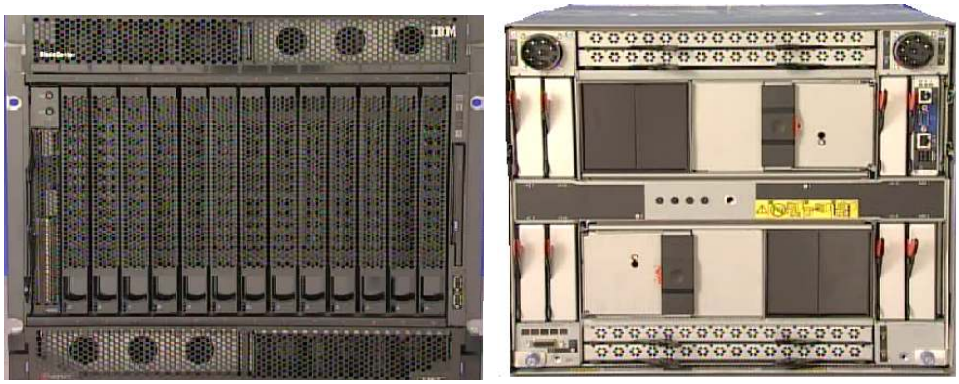

**Figure 2.** Rear view showing I/O module bays and numbering

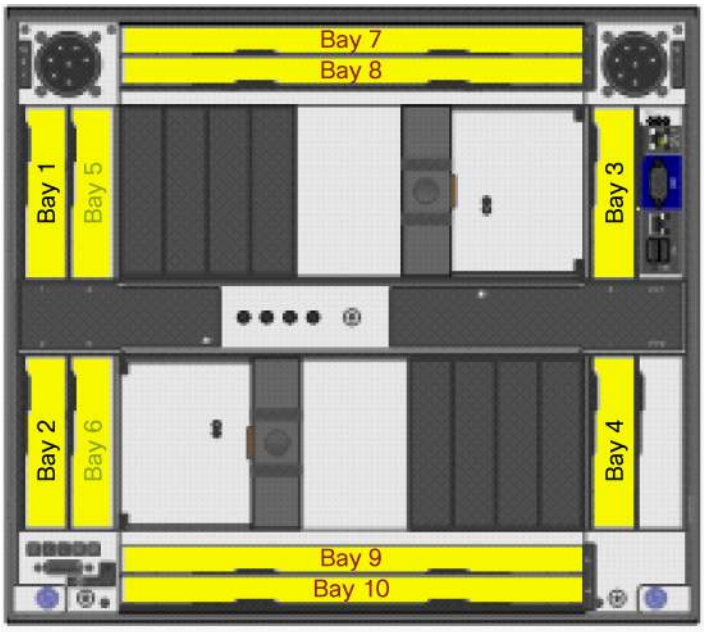


Figure 2 shows the rear of the BladeCenter H and provides bay numbering for the combination of Legacy I/O module bays, bridge I/O module bays and high speed I/O module bays.

- Bays 1 and 2 are dedicated Ethernet bays and could be switch or pass-thru
  - All servers have a connection to these bays
- Bays 3 and 4 can be any legacy module to servers or bridge module to bays 7-10
- Bays 5 and 6 can be bridge modules to bays 7-10 (no direct connection to servers)
- Bays 7-10 are high speed bays (4x IB, 10G Ethernet)
  - Can also be used for Legacy modules with a Multi Switch Interface Module (MSIM) carrier installed in bay 7/8 and/or 9/10

While bays 1 and 2 are dedicated to Ethernet, all other bays can be used by various I/O modules (Fibre Channel, Ethernet, InfiniBand or pass-thru), subject to the daughter cards installed in the servers. In most cases, if one server uses a specific type of daughter card connecting to a set of I/O module bays, then the other servers are restricted to the same type of daughter card. In other words, you can not install an Ethernet daughter card on one server to connect to bays 3 and 4, and then an InfiniBand daughter card in another server to also connect to bays 3 and 4.

The IBM BladeCenterH midplane provides, among other things, the network connectivity between the servers in the front and the I/O modules in the rear. As noted, the CBS3110 products can go into the low speed bays 1 through 4, providing a highly available and multi-homed environment where each server blade is Gigabit attached to each Cisco Catalyst Ethernet Switch Module 3110.

Note that when an MSIM is placed into the top or bottom set of high speed slots, two more slots become available for Ethernet module use. At the time of this writing, it is not known if the CBS3110 will be supported in the MSIM. Please contact your Cisco or IBM sales representative for more up to date information.

The enclosure also contains at least one and up to two Advanced Management Modules (aMM). The aMM provides chassis configuration and management. Each switch module has Ethernet connections to the pair of aMMs. Certain types of server access can be obtained via the aMM's, including concurrent Keyboard/Video/Mouse (cKVM) for all 14 servers (assuming the servers have the appropriate cKVM daughter card). The IBM enclosure also supports Serial over LAN (SoL) for those servers that do not contain video interfaces. The SoL traffic and cKVM traffic is carried between the servers and the aMM via the switch in I/O bay 1 in most cases.

# Cisco Catalyst Ethernet Switch Module 3110 Details

**Figure 3.** CBS3110 Product Family


Two products form the CBS3110 family: CBS3110G and CBS3110X. The CBS3110G provides four copper RJ45 based Gigabit Ethernet uplink ports. The CBS3110X has a single 10GE X2 based uplink port. This port supports SR, LRM, CX4 and TwinGig modules. Future support for LR and LX4 modules will be provided post FCS.

Each Ethernet switch provides 14 internal server connections. Each blade server is connected to the midplane using the available Gigabit Ethernet network interface cards (NICs). The switch also has a Fa0 interface that is accessed via the aMM. When the CBS3110 is in VBS mode, the aMM does not manage the switch and the Fa0 interface is not available for use. Therefore, initial configuration on the Master switch should be done before connecting it to other member switches. The IBM aMM will show the VLAN1 IP address as the management address for the virtual switch on every slot that contains a member switch.

### Server Port Configuration

A blade server is assigned a specific port on the blade switch. This is predetermined by the physical slot the blade server occupies in the enclosure. Table 1 correlates the server and switch ports.

**Table 1.** Correlation of Server and Switch Ports

| Cisco IOS Software CLI Identifier | Port Location in the Enclosure |
|---|---|
| **GigabitEthernet x/0/1** | Server Slot 1 |
| **GigabitEthernet x/0/2** | Server Slot 2 |
| **GigabitEthernet x/0/3** | Server Slot 3 |
| **GigabitEthernet x/0/4** | Server Slot 4 |
| **GigabitEthernet x/0/5** | Server Slot 5 |
| **GigabitEthernet x/0/6** | Server Slot 6 |
| **GigabitEthernet x/0/7** | Server Slot 7 |
| **GigabitEthernet x/0/8** | Server Slot 8 |
| **GigabitEthernet x/0/9** | Server Slot 9 |
| **GigabitEthernet x/0/10** | Server Slot 10 |
| **GigabitEthernet x/0/11** | Server Slot 11 |
| **GigabitEthernet x/0/12** | Server Slot 12 |
| **GigabitEthernet x/0/13** | Server Slot 13 |
| **GigabitEthernet x/0/14** | Server Slot 14 |
| **GigabitEthernet x/0/15** | Copper Uplink 1 (CBS3110G only) |
| **GigabitEthernet x/0/16** | Copper Uplink 2 (CBS3110G only) |
| **GigabitEthernet x/0/17** | Copper Uplink 3 (CBS3110G only) |
| **GigabitEthernet x/0/18** | Copper Uplink 4 (CBS3110G only) |
| **TenGigabitEthernet x/0/1** | 10 GE Uplink Port (CBS3110X only) |
| **Fa0** | Management Interface |

For interface names above the “x” indicates the member number. Interface FastEthernet0 (fa0) is connected internally to the IBM Advanced Management Module (aMM). Only the Fa0 interface on the master switch will be active when the switch is operating in Virtual Blade Switch (VBS) mode.

# Cisco Catalyst Ethernet Switch Module 3110 Overview

The Cisco Catalyst Ethernet Switch Module 3110 provides enhanced Layer 2 and Layer 3 services to Blade Servers. The CBS3110 ships standard with the IPBase software license. This image proves all the layer 2+ features, Data Center customers are used to in the other Catalyst products. The CBS3110 family of switches now supports basic routing functions: RIP and Static Routing, and EIGRP stub. The Cisco CBS3110 enhances basic Layer 2/3 switching by including Cisco proprietary protocols, access control lists (ACLs), and quality of service (QoS) based on Layer 3 information. With Simple Network Management Protocol (SNMP), command-line interface (CLI), or HTTP management options available and a robust set of Cisco IOS® Software switching features, the Cisco Catalyst Ethernet Switch Module 3110 naturally integrates into the data center environment. The following features highlight this capacity:

- Loop protection and rapid convergence with support for Per VLAN Spanning Tree Plus (PVST+), 802.1w, 802.1s, BPDU Guard, Loop Guard, PortFast, UplinkFast, and Unidirectional Link Detection (UDLD)
- Advanced management protocols, including Cisco Discovery Protocol, VLAN Trunking Protocol (VTP), and Dynamic Trunking Protocol (DTP)
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for link load balancing and high availability
- Support for authentication services, including RADIUS and TACACS+ client support
- Support for protection mechanisms, such as limiting the number of MAC addresses allowed or shutting down the port in response to security violations

The CBS3110 brings one additional feature not found in other Ethernet Blade switches: The ability to form a Virtual Blade Switch. By combining up to 9 Blade switches together via high speed interconnect cables, the CBS3110 can emulate a redundant Top of the Rack (ToR) solution

A Virtual Blade Switch offers the following features not present in other Blade Switches:

- A single point of management for the entire Rack of up to 9 Ethernet Switches
  - One IP address for Management
  - Single upgrade for the entire rack of VBS enabled switches. The user only upgrades the master switch and the master upgrades the members
  - Single configuration file for the entire rack. The VBS Ring has one master switch which provides the management functions. In the case of failure of the master switch, each member switch is capable of taking over the master functions. Additional switches can be added to the configuration without taking down the existing members. If a member switch, fails, the high speed ring will loop back around the failed device
  - Single STP instance for L2 networks and a single router for L3 networks
- Consolidation of uplink port – in a VBS environment, there is no longer a requirement to uplink each individual switch into the aggregation layer

Figure 4 shows a diagram of a Virtual Blade Switch configuration. In this configuration, two member switches have a 10GE Ethernet connection to the aggregation. These two cables can be EtherChannelled together to form a 20GE pipe. The raw capacity of the ring is 64 Gbps. Actual throughput depends on traffic patterns and traffic type. This will be discussed later in this paper.

**Figure 4.** Virtual Blade Switch Configuration

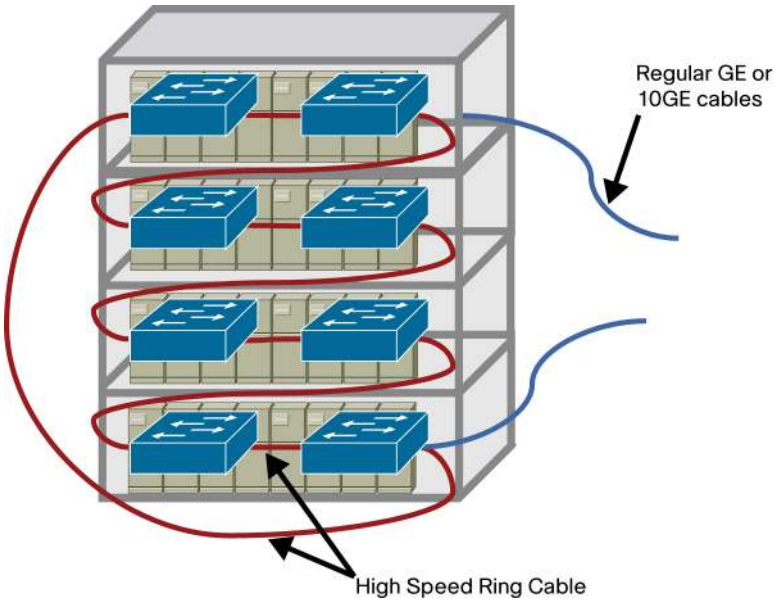

# Data Center Network Architecture

The architecture of the data center infrastructure must address the requirements necessary to create a highly available, scalable, and secure network. This section describes the basic architecture necessary to meet these goals.

It is a synopsis of the Cisco Data Center Network Architecture and includes the following topics:

- Data Center Network Components
- Aggregation Layer
- Access Layer
- High Availability
- BladeSystem in the Data Center Architecture

For details about this architecture, refer to the guide at: http://www.cisco.com/en/US/solutions/ns340/ns517/ns224/ns304/net_design_guidance0900aecd800e4d2e.pdf.

### Data Center Network Components

The terms front-end network and back-end network define the devices that comprise the infrastructure of the data center and their general role. The front-end network is the IP routing and switching environment. It provides client-to-server, server-to-server, and server-to-storage network connectivity. The back-end network supports the SAN fabric and connectivity between servers and other storage devices such as storage arrays and tape drives.

The front-end network contains two distinct functional areas called the aggregation and access layers. Figure 5 depicts the network and the services available at each layer.

**Figure 5.** Data Center Front-End Network

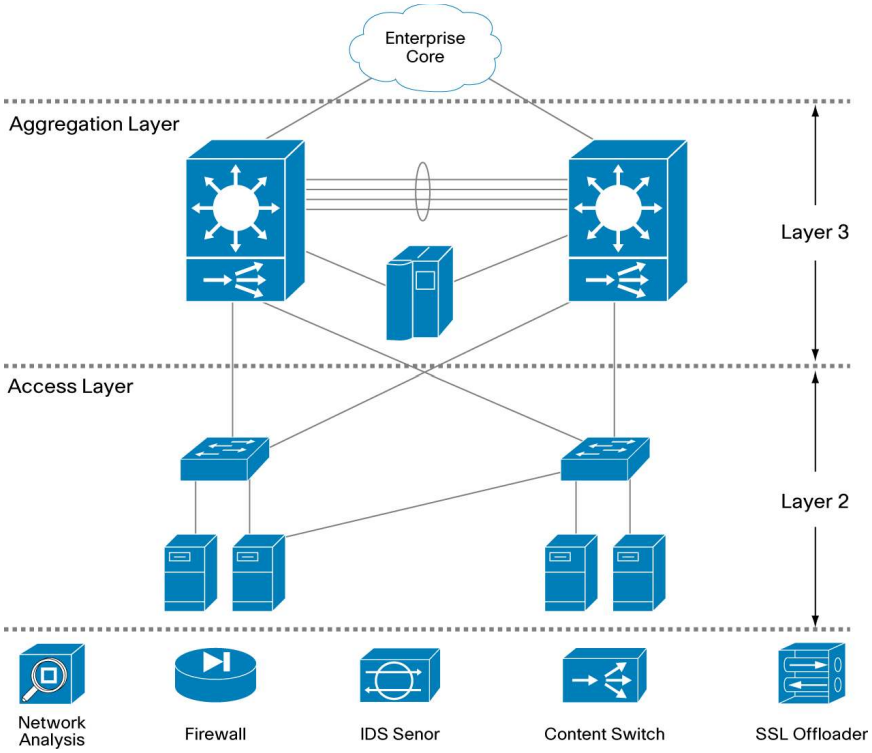


## Aggregation Layer

The aggregation layer is a point of convergence for network traffic that provides connectivity between server farms and the rest of the enterprise. The aggregation layer supports Layer 2 and Layer 3 functionality and presents an ideal location for deploying centralized application, security, and management services. These data center services are shared across the access layer server farms and provide an efficient, scalable, predictable, and deterministic behavior common to server farm needs.

The aggregation layer provides a comprehensive set of features for the data center. The features are supported by the following devices:

- Multilayer aggregation switches
- Load-balancing devices
- Firewalls
- Intrusion detection systems
- Content engines
- Secure Sockets Layer (SSL) off loaders
- Network analysis devices

### Access Layer

The primary role of the access layer is to provide the server farms with port density. In addition, it must be a flexible, efficient, and predictable environment supporting client-to-server and server-to-server traffic. A Layer 2 domain meets these requirements by providing the following:

- Adjacency between servers and service devices
- A deterministic, fast converging, loop-free topology

Layer 2 adjacency, in the server farms, allows for the deployment of servers or clusters that require the exchange of information done at Layer 2 only. It also readily supports access to network services in the aggregation layer such as load balancers and firewalls. This enables an efficient use of shared, centralized network services by the server farms. In contrast, if services are deployed at each access switch, the benefit of those services is limited to the servers directly attached to the switch. It is easier to insert new servers into the access layer when the aggregation layer is responsible for data center services, and the Layer 2 environment provides the flexibility to scale the number of ports. This is another benefit provided in a Layer 2 access layer.

The access layer must provide a deterministic environment to help ensure a stable Layer 2 domain. A predictable access layer allows the spanning tree to converge and recover quickly during failover and fallback scenarios.

### Design Goals

This section describes the design goals when deploying blade servers and the functionality supported by the Cisco Catalyst Ethernet Switch Module 3110 in data centers. It includes the following topics:

- High Availability
- Scalability
- Management

### High Availability

High availability in the data center is a goal that must be achieved systematically. A highly available environment is attainable by addressing each layer of the data center and each of the devices that comprise that particular data center layer. Network and software features help achieve high availability, as well as physical redundancy of links and devices.

The aggregation and access layers use redundant devices and links to help ensure there is no single point of failure. The Layer 2 and/or Layer 3 features supported by these switches also create a highly available infrastructure. Spanning Tree Protocol support on both the aggregation and access switches creates a deterministic topology that converges quickly. Logical redundancy or fault tolerance may be achieved with Layer 3 technologies such as Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP). These protocols allow the gateways for servers or clients to be virtualized across the physical routing devices in the network. This virtualization mitigates the effect of a routing device failure on the availability of data center services. Load-balancing services deployed in the aggregation layer allow the network to monitor server health and application availability. These devices and features combined produce a more resilient application environment.

Dual-homing a server in relation to separate access layer switches is another method to achieve a higher level of availability in the data center. NIC teaming removes the possibility of a single NIC failure isolating the server. It requires the server to have two separate NICs that support teaming software. Typically, teaming software detects failures over an external network probe between members of the team or by monitoring the local status of each NIC in the team. The combination of dual-homed servers and a network load balancer provides an even greater level of availability for the server and the applications it supports.

Data centers are the repository of critical business applications that support the continual operation of an enterprise. These applications must be accessible throughout the working day during peak times, and some on a 24-hour basis. The infrastructure of the data center, network devices, and servers must address these diverse requirements. The network infrastructure provides device and link redundancy combined with a deterministic topology design to achieve application availability requirements. Servers are typically configured with multiple NICs and dual-homed to the access layer switches to provide backup connectivity to the business application.

High availability is an important design consideration in the data center. The Cisco Catalyst Ethernet Switch Module 3110 has a number of features and characteristics that contribute to a reliable, highly available network.

### High Availability for the Blade Server Switching Infrastructure

High availability between the Cisco Catalyst Ethernet Switch Module 3110s in the blade server enclosure and the aggregation layer switches requires link redundancy. Each Cisco Catalyst Ethernet Switch Module 3110 offers multiple ports for uplink connectivity to the external network, which allows for redundant paths using two links each for more redundancy. However, this introduces the possibility of Layer 2 loops; therefore, a mechanism is required to manage the physical topology. The implementation of Rapid Spanning Tree Protocol (RSTP) helps ensure a fast converging, predictable Layer 2 domain between the aggregation layer and access switches (the Cisco Catalyst Ethernet Switch Module 3110s) when redundant paths are present. For customers who want to implement the Access Layer without Spanning Tree, the CBS3110 supports FlexLinks. FlexLinks associates a “back up” interface with each forwarding interface. Therefore, the customer can maintain a redundant topology without the use of STP.

The recommended design is a triangle topology (as shown in Figure 6 earlier), which delivers a highly available environment through redundant links and a spanning tree. It allows for multiple switch or link failures without compromising the availability of the data center applications.

The access layer uplink EtherChannels support the publicly available subnets in the data center and traffic between servers. The server-to-server traffic that uses these uplinks is logically segmented through VLANs and may use network services available in the aggregation layer. There is also a port channel defined between the two blade enclosure switches. This path provides intra-enclosure connectivity between the servers for VLANs defined locally on the blade enclosure switches. Clustering applications that require Layer 2 communication may use this traffic path, as well as mirrored traffic. Each of these port channels is composed of two Gigabit Ethernet ports or two 10GE ports.

RPVST+ is recommended as the method for controlling the Layer 2 domain because of its predictable behavior and fast convergence. A meshed topology combined with RPVST+ allows only one active link from each blade switch to the root of the spanning tree domain. This design creates a highly available server farm through controlled traffic paths and the rapid convergence of the spanning tree. The details of the recommended design are discussed in a later section.

### High Availability for the Blade Servers

Blade enclosures provide high availability to blade servers by multihoming each server to the Cisco Catalyst Ethernet Switch Module 3110s. The two Cisco Catalyst Ethernet Switch Module 3110s housed in the interconnect bays are connected to the blade server over the backplane. Two LAN on Motherboard (LOM) Gigabit Ethernet connections are available to every blade-server slot. Additional interfaces can be installed by using the optional Mezzanine slots and additional Ethernet Switches.

Multihoming the server blades allows the use of a NIC teaming driver, which provides another high-availability mechanism to fail over and load balance at the server level. Three modes of teaming are supported:

Network Fault Tolerance (NFT)

Transmit Load Balancing (TLB)

Switch Assisted Load Balancing (SLB)

NFT teaming creates a virtual interface by grouping the blade server network adapters into a team. One adapter is the primary active interface, and all other adapters are in a standby state. The virtual adapter uses a single MAC address and a single Layer 3 address. NFT provides adapter fault tolerance by monitoring the state of each team member network connection. The standby NICs only become active if the primary NIC loses connectivity to the network.

TLB teaming supports adapter fault tolerance (NFT) and adds more functionality in the server for load balancing egress (transmit) traffic across the team. Note that a TLB team uses only one NIC to receive traffic. The load-balancing algorithm is based on either the destination MAC or IP address. This teaming method provides better use of the bandwidth available for egress traffic in the network than NFT.

SLB teaming extends the functionality of TLB by allowing the team to receive load-balanced traffic from the network. This requires that the switch can load balance the traffic across the ports connected to the server NIC team. The Cisco Catalyst Ethernet Switch Module 3110supports the IEEE 802.3ad standard and Gigabit port channels. The CBS3110 now supports SLB. Servers can now operate in Active/Active configurations. This means that each server team can provide 2 Gigabit of Ethernet Connectivity to the Switching fabric. Failover mechanisms are automatically built into the LACP protocol.

For more information on NIC teaming, please visit:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5070766&brandind=5000020;

http://www.linux-foundation.org/en/Net:Bonding#Bonding_Driver_Options.

### Scalability

The capability of the data center to adapt to increased demands without compromising its availability is a crucial design consideration. The aggregation layer infrastructure and the services it provides must accommodate future growth in the number of servers or subnets it supports.

When deploying blade servers in the data center there are two primary factors to consider:

- Number of physical ports in the aggregation and access layers
- Number of slots in the aggregation layer switches

### Aggregation Layer Switch Physical Port Count

The introduction of blade systems into the data center requires greater port density at the aggregation layer. Blade systems, deployed with internal switches, provide their own access layer. The cabling and maximum number of servers per enclosure is predetermined. Scaling the aggregation-layer ports to accommodate the blade system uplinks is an area that requires attention.

It is important to remember that aggregation switches provide data center services such as load balancing, security, and network analysis that may require dedicated ports for appliances or slots for integrated services. This directly affects the number of ports available for access-layer connectivity.

### Aggregation Layer Switch Slot Count

The data center infrastructure must be flexible enough to allow growth in both server capacity and service performance. Connecting a blade system directly into the aggregation layer places more significance on the number of slots available to accommodate blade system uplinks and integrated services.

Traditionally, the access layer provides the port density necessary to allow the physical growth of server farms. Modular access-layer switches offer connectivity to densely packed server farms over a few uplinks. The aggregation-layer switches support a limited number of uplinks from the access layer. With this model, the number of servers supported per uplink is high.

Blade systems use more aggregation-layer resources per server than this traditional deployment model. Each uplink from a blade enclosure provides connectivity to a maximum of 14 servers. The aggregation layer must be flexible enough to manage the increased demand for ports and slots in this blade server system environment.

To scale the server farm, use an aggregation-layer switch that provides an ample number of slots for line cards and/or service module expansion.

In addition, consider using the following two options (which are not mutually exclusive):

- Deploying service switches in the aggregation layer (as depicted in Figure 8)
- Using a data center core to accommodate multiple aggregation-layer modules

Service switches are deployed in the aggregation layer to host integrated data center services such as load balancing, intrusion detection, and network analysis. Relocating these services to a separate switch frees ports and slots in the aggregation-layer switches. This design allows the aggregation switches to commit more slots and, ultimately, more ports to the Layer 2 connectivity of the server farms.

Figure 6 depicts a service switch deployment. The aggregation switches for each data center modules are Layer 3 attached to the core. In addition, the aggregation switches house the service modules required to support the server farms. The aggregation layer provides the connectivity layer for the various access layer switches. In a typical Data Center network, this layer will have connections from Blade and non-Blade switches, providing connectivity for all the servers in the network

**Figure 6.** Data Center Scaling with Service Switches

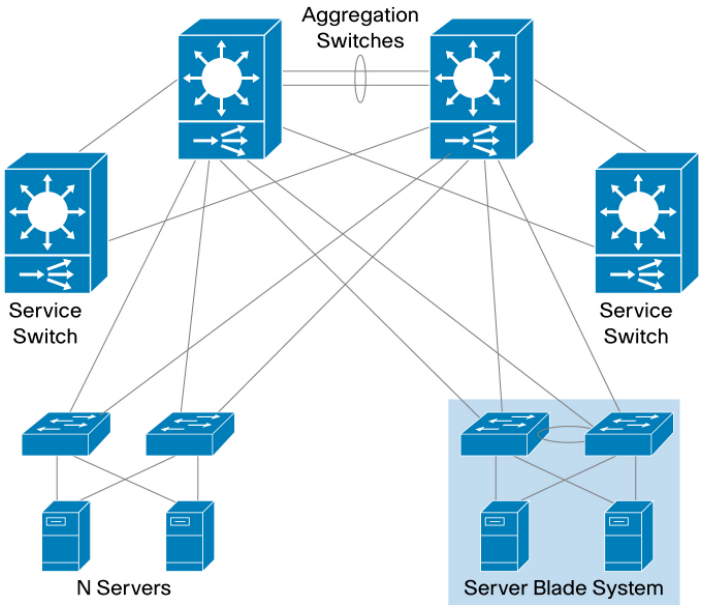


The data center core is a mechanism to replicate and horizontally scale the data center environment. In the recommended design the aggregation and access layers are regarded as a module that can be duplicated to extend the enterprise. Each data center module provides its own network services locally in the aggregation switches. This approach allows the network administrator to determine the limits of each data center module and replicate as necessary.

### Management

The Cisco Catalyst Ethernet Switch Module 3110 is accessible for management and configuration by any of the following traffic paths:

- Out-of-band management
- In-band management
- Serial console port

These traffic paths provide three different management options for network administration and support different user and application interfaces to the Cisco Catalyst Ethernet Switch Module 3110. The remote management of the blade servers within the blade enclosure is critical to an efficient and scalable data center. The various Blade Server vendors offer connectivity options provided using the enclosure to the blade servers are also discussed. See the particular Appendix for your enclosure type.

### Out-of-Band Management

Out-of-band management is the practice of dedicating an interface on the managed device for carrying management traffic. It is also the recommended management method for blade systems. Out-of-band management isolates the management and data traffic and provides a more secure environment.

The Cisco Catalyst Ethernet Switch Module 3110 contains an additional Fast Ethernet port, which connects to the Advanced Management Module (aMM), providing out-of-band management. The user may also use this path to access the CLI functions of the switch, transfer SNMP information, and upload software images and configuration files. This path is totally independent of the switch fabric. This Fast Ethernet port (Fa0) is assigned an IP address via the aMM when in standalone mode. When the CBS3110 are stacked together the Fa0 interface is disable. In-Band Management

In-band management uses logical isolation to separate management traffic from data traffic. VLANs segregate the two traffic types that are sharing the bandwidth of the uplink ports. This practice is common where applications running on the servers must be managed along with the network infrastructure devices.

In-band management traffic uses the uplink trunk ports located on the rear of the Cisco Catalyst Ethernet Switch Module 3110s for management. It is recommended that you do not up the user ports in the same VLAN as the management VLAN.

The Cisco Catalyst Ethernet Switch Module 3110supports multiple switched virtual interfaces (SVIs) to be active at the same time; however, the Cisco Catalyst Ethernet Switch Module 3110 does not perform any routing functions between SVIs. By default, the SVI is created as VLAN 1 and enabled during the setup phase of the installation. The VLAN is often referred to as the “management VLAN.” Cisco Systems® recommends that the user change the management VLAN to something other than VLAN 1. Therefore, it is important to create an SVI with another VLAN and allow this VLAN on the external front panel ports. In addition, you can manage the switch using the Fa0 port using the Onboard Administrator on the rear of the enclosure.

For best practices in selecting the management VLAN, please visit: Cisco.com Login Page.

#### Serial Console Port

The front panel of the Cisco Catalyst Ethernet Switch Module 3110 has a serial port that can be used to manage the switch through the CLI. The CLI can be accessed by connecting directly to the console port with the serial port of a workstation or remotely by using terminal servers and IP connectivity protocols such as telnet. Even when the switches are connected via the high speed ring ports, all console ports are active. However, traffic is routed to the CPU on the Master switch. Therefore, every console port acts the same. If the user is connecting the switches to a console server, then the user should hook up at least two console cables per Virtual Blade Switch, to insure connectivity in case of switch outage.

**Management Options**

The Cisco Catalyst Ethernet Switch Module 3110switch is manageable through the following methods:

- Cisco IOS Software CLI (via console, telnet, SSH)
- HTTP-based device manager GUI
- Cisco Device Manager (CNA, version 5.3 and greater)
- SNMP-based management applications

The embedded device manager on the Cisco Catalyst Ethernet Switch Module 3110 provides a GUI to configure and monitor the switch through a Web browser. This requires using either in-band or out-of-band management and enabling the HTTP/HTTPS server on the switch. The HTTP server and SSL are enabled by default.

SNMP-compatible management utilities are supported through a comprehensive set of MIB extensions and through four remote monitoring (RMON) groups. CiscoWorks 2000, IBM Director and HP OpenView are two such management applications. SNMP Versions 1, 2, and 3 are available on the switch (Cisco IOS Software crypto image is required for SNMP V3 and SSH).

The CLI delivers the standard Cisco IOS Software interface over telnet or the console port. The use of SSH for CLI access is recommended.

**Note:** For more information about the embedded device manager, refer to the online help on the switch CLI.

# Design and Implementation Details

This section includes the following topics:

- Network Management Recommendations
- Network Topologies Using Cisco Catalyst Ethernet Switch Module 3110
- Configuration Details

### Network Management Recommendations

An out-of-band (OOB) network is recommended for managing the Cisco Catalyst Ethernet Switch Module 3110. OOB management provides an isolated environment for monitoring and configuring the switch. Isolation is achieved by deploying a physically separate management network or by logically separating the traffic with management VLANs.

The Cisco Catalyst Ethernet Switch Module 3110 has four external Gigabit Ethernet ports; any of them may be used to support network monitoring devices and network management traffic. By using secure protocols, such as SSH or HTTPS, the network maintains the integrity of communications between the switch and the management station. The console port positioned at the front of the Cisco Catalyst Ethernet Switch Module 3110 is another option for connectivity to the OOB network.

If the CBS3110s are connected together via the high speed ring connectors, then placement of the management monitoring devices may impact traffic flow. If the user is sniffing local ports, then traffic does not traverse the ring. However, even if the user is using SPAN to sniff traffic from ports not on the local member switch, that traffic must travel on the high speed ring. This may reduce the ring capacity and cause or add to an oversubscription environment.

### Cisco Catalyst Ethernet Switch Module 3110 Features

This section highlights information about the protocols and features provided by the Cisco Catalyst Ethernet Switch Module 3110 that help integrate blade server enclosures into the Cisco Data Center Network Architecture. This section includes the following topics:

- VBS Ring Architecture
- Spanning Tree
- FlexLink
- Downstream EtherChannel to servers
- Upstream Cross Switch EtherChannel
- Traffic Monitoring
- Link Aggregation Protocols

### VBS Ring Architecture

Each CBS3110 switch has two connectors that allow it to become a member of a Virtual Blade Switch (VBS). The switches are connected on a Ring. The ring is made up of two counter rotating uni-directional traffic lanes. Each lane has a raw capacity of 16 Gbps. A member on the ring can transmit and receive in both directions. If the ring is composed of four or more switches, the ring capacity has a raw data rate of 64 Gbps due to spatial reuse (multiple senders at the same time for each ring).

The ring access is done on a token passing and credit assignment basis. When each member needs to transmit data, it waits for a token to arrive form either direction. Once it receives the token, it is allowed to transmit a block of traffic. If it exceeds its credits, it must wait for the token to return and then transfer again. Each time the member request credits, it may get them from either direction.

The maximum number of physical switches per VBS is 9. Due to the fact that most customers use them in pairs, the normal limit is 8. This works out to approximately one VBS per rack, assuming you have four enclosures per rack. A single VBS then services up to 56 servers. A VBS also eliminates up to 112 Ethernet cables.

The primary difference between VBS and standalone operation is from the server side. In a standalone configuration, the CBS3110 operates much like the CBS3012. Each server sees two separate switches. In the VBS configuration, the server only sees one switch for both NICs. Therefore, in the VBS configuration, the server can enable Active-Active NIC teaming.

A VBS seen from the upstream aggregation switch is no different than a standalone switch. Therefore the decision to select the network topology for connecting to the aggregation switches is the same for both. The options for network topologies will be discussed later.

Additional option is when the left switches in the rack are placed into one VBS configuration and all the right switches are placed in another. In this case, the upstream interfaces from each VBS follow the above model, but the downstream interfaces facing the servers now look like those in the standalone model. The user may choose this configuration, if completely isolated network fabrics are required. The user can not use Active-Active NIC teaming in this configuration.

### Ring Design and Capacity

As mentioned before, each member switch has two ring connections. Each connection is made up of two 16 Gbps counter rotating uni-directional rings. Each member switch can transmit and receive 32 Gbps.  When the VBS Ring is composed of four or more member switches, spatial reuse increases the ring capacity to 64 Gbps.

Much like Ethernet, the user available traffic on the ring is less. Whereas, Ethernet has a 20 byte Inter Packet Gap (IPG) overhead, the ring does not. However, packets on the ring have an additional 24 byte ring header. Therefore, the capacity of the ring is impacted by packet size. Each unicast packet sent on the ring is acknowledged by the receiving switch with a 16 Byte ACK packet. Ring Capacity for unicast packets is +/- 48 Gbps depending on packet size.

For Multicast and Broadcast packets, packets must travel the entire ring, and therefore reduce this capacity further. A ring made up of entirely Multicast packets may get as little as 26 Gbps of capacity. Since normal traffic is a mix of multiple packet types, the ring capacity will vary between 26 and 48 Gbps, and may exceed 48 Gbps in some rare cases.

### Spanning Tree

The Cisco Catalyst Ethernet Switch Module 3110 supports different versions of the Spanning Tree Protocol and associated features, including the following:

- Rapid Spanning Tree (RSTP), based on 802.1w
- Multiple Spanning Tree (MST), based on 802.1s (and includes 802.1w support)
- Per VLAN Spanning Tree Plus (PVST+)

Rapid Per VLAN Spanning Tree Plus (RPVST+)

Loop Guard

UDLD

BPDU Guard

PortFast

UplinkFast (Cisco proprietary enhancement for 802.1d deployments)

BackboneFast (Cisco proprietary enhancement for 802.1d deployments)

The 802.1w protocol is the standard for rapid spanning tree convergence, while 802.1s is the standard for multiple spanning tree instances. Support for these protocols is essential in a server farm environment for allowing rapid Layer 2 convergence after a failure occurs in the primary path. The primary benefits of 802.1w include the following:

The spanning tree topology converges quickly after a switch or link failure.

Convergence is accelerated by a handshake, known as the proposal agreement mechanism.

**Note:** PortFast, BackboneFast, or UplinkFast apply only to PVST+ based networks.

In terms of convergence, Spanning Tree Protocol algorithms based on 802.1w are much faster than the traditional Spanning Tree Protocol 802.1d algorithms. The proposal agreement mechanism allows the Cisco Catalyst Ethernet Switch Module 3110 to decide new port roles by exchanging proposals with its neighbors.

With 802.1w, as with other versions of the Spanning Tree Protocol, bridge protocol data units (BPDUs) are by default sent every 2 seconds (called the hello time). If three BPDUs are missed, Spanning Tree Protocol recalculates the topology, which takes less than 1 second for 802.1w.

Since the data center is made of point-to-point links, the only failures are physical failures of the networking devices or links. 802.1w is able to actively confirm that a port can safely transition to forwarding without relying on any timer configuration. This means that the actual convergence time is less than 1 second.

A scenario where BPDUs are lost may be caused by unidirectional links, which can cause Layer 2 loops. To prevent this problem, you can use Loop Guard and UDLD. Loop Guard prevents a port from forwarding as a result of missed BPDUs, which might cause a Layer 2 loop that could bring down the network.

UDLD allows devices to monitor the physical configuration of fiber optic or copper Ethernet cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected port and generates an alert. BPDU Guard prevents a port from being active in a spanning tree topology as a result of an attack or a misconfigured device connected to the switch port. The port that sees unexpected BPDUs is automatically disabled and must then be manually enabled. This gives the network administrator full control over port and switch behavior.

The Cisco Catalyst Ethernet Switch Module 3110supports Per VLAN Spanning Tree (PVST) and a maximum of 128 spanning tree instances. RPVST+ is a combination of Cisco PVST Plus (PVST+) and RSTP. RPVST+ provides the flexibility of one spanning tree instance per VLAN and the fast convergence benefits of 802.1w. MST allows the switch to map several VLANs to one spanning tree instance, reducing the total number of spanning tree topologies the switch processor must manage. A maximum of 16 MST instances is supported. In addition, MST uses 802.1w for rapid

convergence. MST and RPVST+ create a more predictable and resilient spanning tree topology, while providing downward compatibility for integration with devices that use 802.1d and PVST+ protocols.

Figure 7 illustrates an example for Spanning Tree Protocol when using two switches in the crossover configuration. Each blade switch is dual homed to each aggregation switch via a 2-port EtherChannel®. In this figure the blocked links are indicated in red. In this example, only four of the eight uplinks from each blade switch are in use. The network designer can make those EtherChannel uplinks more robust (up to four ports each), or use them to connect other devices such as intrusion detection systems (IDS) or standalone servers.

**Figure 7.** Spanning Tree Example with the Cisco Catalyst Ethernet Switch Module 3110s

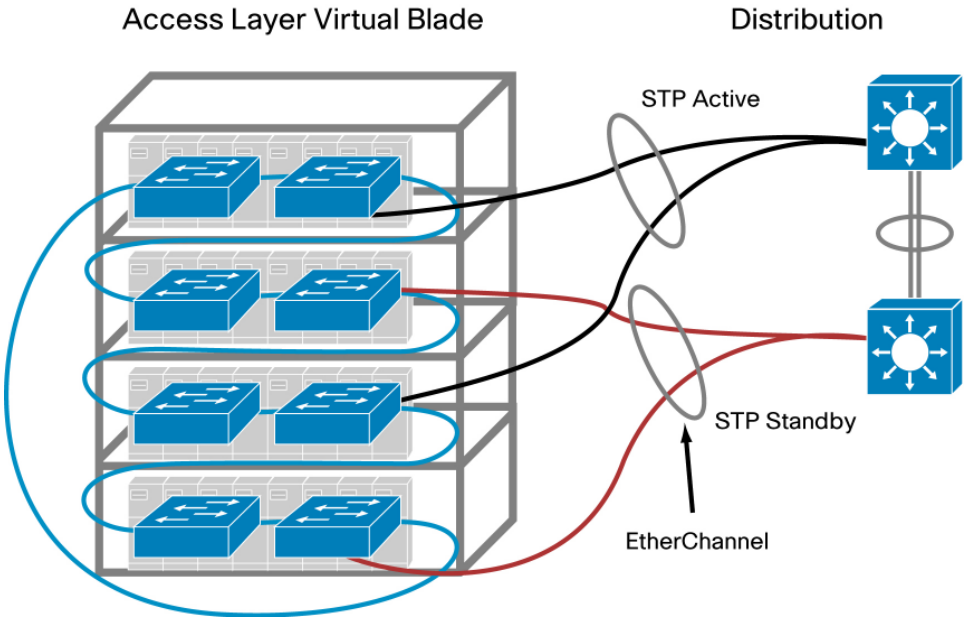


The 802.1w protocol is enabled by default when running spanning tree in RPVST+ or MST mode on the Cisco Catalyst Ethernet Switch Module 3110. The Cisco Catalyst Ethernet Switch Module 3110 enables PVST+ for VLAN 1 by default.

Spanning tree uses the path cost value to determine the shortest distance to the root bridge. The port path cost value represents the media speed of the link and is configurable on a per-interface basis, including Cisco EtherChannel interfaces. The longer path cost better reflects changes in the speed of channels and allows Spanning Tree Protocol to optimize the network in the presence of loops.

The Cisco Catalyst Ethernet Switch Module 3110supports IEEE 802.1t, which allows for spanning tree calculations based on a 32-bit path cost value instead of the default 16 bits. For more information about the standards supported by the Cisco Catalyst Ethernet Switch Module 3110, refer to the “Cisco Catalyst Ethernet Switch Module 3110 Overview” document.

For more information regarding spanning tree and Layer 2 design in the data center, refer to: http://www.cisco.com/en/US/solutions/ns340/ns517/ns224/ns304/net_design_guidance0900aecd800e4d2e.pdf.

### FlexLinks

An alternative to STP is FlexLinks. FlexLinks gives you the ability to make active-standby pairs and provide redundancy without needing to use STP. FlexLinks can be implemented on a VLAN by VLAN basis much like STP. For the example in Figure 5, the top EtherChannel can be active for one VLAN and be standby for another. This allows both uplinks to be active at the same time and share the network load.

### Traffic Monitoring

The Cisco Catalyst Ethernet Switch Module 3110supports the following traffic monitoring features, which are useful for monitoring blade enclosure traffic in data center environments:

Switched Port Analyzer (SPAN)

Remote SPAN (RSPAN)

SPAN mirrors traffic transmitted or received on source ports or source VLANs to another local switch port. This traffic can be analyzed by connecting a switch or RMON probe to the destination port of the mirrored traffic. Only traffic that enters or leaves source ports or source VLANs can be monitored using SPAN.

RSPAN enables remote monitoring of multiple switches across your network. The traffic for each RSPAN session is carried over a user-specified VLAN that is dedicated to that RSPAN session for all participating switches. The SPAN traffic from the source ports or source VLANs is copied to the RSPAN VLAN. This mirrored traffic is then forwarded over trunk ports to any destination session that is monitoring the RSPAN VLAN.

### Link Aggregation Protocols

Fast EtherChannel interfaces and Gigabit EtherChannel interfaces are logically bundled and provide link redundancy and scalable bandwidth between network devices. The Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) help automatically create these channels by exchanging packets between Ethernet interfaces and negotiating a logical connection. PAgP is a Cisco proprietary protocol that can be run only on Cisco switches or on switches manufactured by vendors that are licensed to support PAgP. LACP is a standard protocol that allows Cisco switches to manage Ethernet channels between any switches that conform to the 802.3ad protocol. Because the Cisco Catalyst Ethernet Switch Module 3110supports both protocols, you can use either 802.3ad or PAgP to form port channels between Cisco switches.

When using either of these protocols, a switch learns the identity of partners capable of supporting either PAgP or LACP and identifies the capabilities of each interface. The switch dynamically groups similarly configured interfaces into a single logical link, called a channel or aggregate port. The interface grouping is based on hardware, administrative, and port parameter attributes. For example, PAgP groups interfaces with the same speed, duplex mode, native VLAN, VLAN range, trunking status, and trunking type. After grouping the links into a port channel, PAgP adds the group to the spanning tree as a single switch port.

## Network Topologies Using the Cisco Catalyst Ethernet Switch Module 3110

The network designs emphasize high availability in the data center by eliminating any single point of failure and by providing deterministic traffic patterns and predictable behavior during times of network convergence. The configuration example included uses a pair of Cisco Catalyst 6513 Switches as the aggregation-layer platform. This Layer 2/Layer 3 switching platform supports the slot density and integrated network services required by data centers deploying blade systems. A generic Blade Server cabinet and at least two Cisco Catalyst Ethernet Switch Module 3110s composes the Layer 2 access layer.

### Cross Over Design (AKA “V” or Triangle configuration)

Typical deployment in the data center uses the classic triangle topology. There is no single point of failure in this deployment model. It does not matter if the Catalyst Ethernet Switch Module 3110s are operating in the standalone environment or in VBS mode. At least two switches are dual-homed to the aggregation layer, providing link redundancy. Spanning Tree Protocol manages the physical loops created by the uplinks between the aggregation and access switches, assuring a predictable and fast converging topology.

RPVST+ fulfills the high-availability requirements of this design and is the recommended mode of spanning tree operation. RPVST+ provides fast convergence (less than 1 second) in device or uplink failure scenarios. In addition, RPVST+ offers enhanced Layer 2 features for the access layer with integrated capabilities equivalent to PortFast, UplinkFast, and BackboneFast.

If the user wants to eliminate STP from the network, the Cross Over Design can be implemented using FlexLinks.

The connection between blade switches (in VBS mode) supports local traffic limited to the local rack: for example, clustering applications or management traffic such as remotely mirrored (RSPAN) traffic. Local traffic does not need to go up to 6513s and back. It can stay within the local ring. Of course, if the CBS3110s switches are running in standalone mode, they can not take advantage of this.

The server NICs support the logical separation of VLANs by trunking. This allows each NIC to accommodate the public and the private VLANs on the Cisco Catalyst Ethernet Switch Module 3110s. In addition, servers are dual-homed to each of the two Cisco Catalyst Ethernet Switch Module 3110s in the enclosure. This structural design can also allow for the physical separation of public and private VLANs between two NICs homed to the same Cisco Catalyst Ethernet Switch Module 3110.

In VBS mode, the cross-over topology provides a high level of availability to the blade servers. If all the uplinks from any blade switch in the ring to any of the aggregation switches are unavailable, the server NICs homed to that Cisco Catalyst Ethernet Switch Module 3110 will route traffic onto the ring to another member switch which has an upstream connection. The blade servers are unaware of the disconnection between the access-layer switches (Cisco Catalyst Ethernet Switch Module 3110s) and the aggregation-layer switches and continue to forward traffic.

In the standalone mode, this is not the case. Since there is no ring connection, server NICs could end up black holing traffic to a switch that has no uplinks. To address this breakdown in network connectivity, use the following methods:

- Use the NIC teaming features of the blade servers

Deploy Layer 2 trunk failover feature in the Cisco Catalyst Ethernet Switch Module 3110s

In addition, the NIC teaming features of the blade servers provide redundancy at the network adapter level. Stagger the preferred primary NICs between the two Cisco switches in the enclosure to increase server availability. Assigning the primary NIC is a straightforward process. The NIC teaming software provides a GUI or a small configuration file, depending on the operating system, to construct the team.

**Non-Loop Design (AKA "U" or Square configuration)**

When the CBS3110 switches are used in standalone environment, not only can the network topology follow the Cross-Over design, but also, a Non-Loop design can be deployed. The Non-Loop topology follows the same design guide lines as for the CBS30x0 products. Basically all the uplinks from one blade switch run to one Distribution switch and all the uplinks from the other run to the second Distribution switch. Here it is important to turn on Layer 2 trunk failover as discussed above. In this design, STP is not needed, and FlexLinks does not apply since there is no alternate path for either switch to forward traffic.

One disadvantage is the fact, that in the standalone operation, the servers can not form Active-Active NIC teams since each NIC is attached to a different switch.

One additional configuration is when all the left switches are placed into one VBS configuration and all the right switches are in another. I this case, the upstream interfaces from each VBS follow the above model. However, the downstream side facing the servers looks like the standalone model. The user may choose this configuration, if he requires two completely isolated network fabrics.

## Multiple Pairs of Switches per Blade Enclosure

Nothing in this design guide limits the user to a pair of Blade switches per enclosure. Some users may require more than two NICs per Blade. VMWare typically wants four or six interfaces. By place two or three pairs of switches per enclosure, the server capacity can be increased. However, the number of enclosures grouped together by the VBS is reduced. At no time, can the limit of 9 switches per VBS be exceeded. Also, increasing the number of NICs per server may increase the traffic on the Ring. In the next section, the throughput of the ring will be discussed.

# Configuration Steps

### Configuring the Aggregate Switches

Complete the following steps on the aggregate switches:

Step 1. VLAN configuration

Step 2. RPVST+ configuration

Step 3. Primary and secondary root configuration

Step 4. Configuration of port channels between aggregate switches

Step 5. Configuration of port channels between aggregate switches and Cisco Catalyst Ethernet Switch Module 3110s

Step 6. Trunking of port channels between aggregate switches

Step 7. Configuration of default gateway for each VLAN

**Note:** The “Configuration Details” section describes each of these steps.

### Configuring the Cisco Catalyst Ethernet Switch Module 3110s

Complete the following steps on the Cisco Catalyst Ethernet Switch Module 3110s:

Step 1. VLAN configuration

Step 2. RPVST+ configuration

Step 3. Configuration of port channels between the Cisco Catalyst Ethernet Switch Module 3110 and aggregate switches

Step 4. Trunking of port channels between the Cisco Catalyst Ethernet Switch Module 3110 and aggregate switches

Step 5. Configuration of server ports on the Cisco Catalyst Ethernet Switch Module 3110

### Additional Aggregation Switch Configuration

The following recommendations help integrate the Cisco Catalyst Ethernet Switch Module 3110s into the data center.

Step 1. Enable Root Guard on the aggregate switches’ links connected to the switches in the blade enclosure. The spanning tree topology is calculated and one of the primary parameters involved in this equation is the location of the root switch. Determining the position of the root switch in the network allows the network administrator to create an optimized forwarding path for traffic. Root Guard is a feature designed to control the location of the root switch. The aggregation switches should employ the spanning-tree guard root command on the port channel interfaces connected to the blade switches.

Step 2. Allow only those VLANs that are necessary on the port channel between the aggregate and the blade switches. Use the switchport trunk allowed vlan vlanID command to configure the port channel interfaces of the aggregate switch to allow only those VLANs indicated with the vlanID option.

### Additional Cisco Catalyst Ethernet Switch Module 3110 Configuration

Step 1. Enable BPDU Guard on the internal server ports of the switch. Use the spanning-tree bpduguard enable command to shut down a port that receives a BPDU when it should not be participating in the spanning tree.

Step 2. Allow only those VLANs that are necessary on the port channels between the aggregate and the blade switches. Use the switchport trunk allowed vlan vlanID command to configure the port channel interfaces of the switch to allow only those VLANs indicated with the vlanID option.

This design uses the links between the two Cisco Catalyst Ethernet Switch Module 3110s as a redundant path for blade server traffic. The use of a longer path cost value provides for a more granular calculation of the topology based on the available link bandwidth (see the “Cisco Catalyst Ethernet Switch Module 3110 Feature” section). This feature is enabled with the spanning-tree pathcost method long CLI command. RPVST+ should be used in this network design for its fast convergence and predictable behavior.

The following convergence tests were conducted against this alternative topology:

Uplink failure and recovery between Switch-A and the primary root

Uplink failure and recovery between Switch-B and the secondary root

Failure and recovery of Switch-A and Switch B

Failure and recovery of the primary and secondary root switches

These tests yielded results similar to the recommended topology. Layer 2 convergence occurs in approximately 1 second. As stated previously, recovery at Layer 3 is dependent on the HSRP settings of the aggregate switches (see “Recommended Topology” section). In our testbed, the failure of the active HSRP device typically increased the convergence time to 5 seconds.

This design supports traffic monitoring using SPAN and/or RSPAN. For example, a network analysis device connected to the external ports on the front of the Cisco Catalyst Ethernet Switch Module 3110 may capture locally mirrored traffic. Alternatively, RSPAN traffic may be carried on the Cisco Catalyst Ethernet Switch Module 3110 uplinks if bandwidth utilization is not a concern

### Configuration Details

This section describes the configuration steps required for implementing the topologies discussed in this guide. The configuration for the following are discussed:

- VLAN
- RPVST+ or FlexLink
- Inter-Switch Link
- Server Port
- Server Default Gateway
- RSPAN

### VLAN Configuration

To configure the VLANs on the switches, complete the following tasks:

Set the VLAN trunking-protocol administrative domain name and mode and create the server farm VLANs as follows:

```
(config)# vtp domain <domain name>
(config)# vtp mode transparent
(config)# vlan 60
(config-vlan)# name bladeservers
(config-vlan)# state active
```

### RPVST+ Configuration

Configure Spanning Tree Protocol to manage the physical loops in the topology. It is recommended to use RPVST+ for its fast convergence characteristics. Set the Spanning Tree Protocol mode on each aggregation switch as follows:

```
(config)# spanning-tree mode rapid-pvst
```

Configure the path cost to use 32 bits in the Spanning Tree Protocol calculations:

```
(config)# spanning-tree pathcost method long
```

Configure the primary and secondary root switches as follows:

```
(config)# spanning-tree vlan <vlan range> root primary | secondary
```

### FlexLinks Configuration

Alternative to RPVST+, is FlexLinks. To enable FlexLinks, use the following commands:

```
(config)# interface tengigethernet1/0/1
(config-int)# switchport backup interface tengigethernet2/0/2
```

By doing this, the interface on member switch two will become the back up to the 10GE interface on member switch one. There are additional settings, including ones for the MAC Address Move function, including in the Software Configuration Guide for the particular switch model you are using. Please refer to those instructions for more details.

### Inter-Switch Link Configuration

The topologies discussed in this guide require connectivity between the switches. The following two types of interswitch connections exist:

- Aggregate-1 to Aggregate-2
- Aggregate-1 or Aggregate-2 to Blade Switch-A or Blade Switch-B

Each of these connections maybe a Layer 2 EtherChannel connection consisting of multiple physical interfaces bound together as a channel group or port channel. These point-to-point links between the switches should carry more than one VLAN; therefore, each is a trunk.

### Port Channel Configuration

Link Aggregate Control Protocol (LACP) is the IEEE standard for creating and managing EtherChannel connections between switches. Each aggregate switch uses this feature to create a port channel across the line cards. The use of multiple line cards within a single switch reduces the possibility of the point-to-point port channel becoming a single point of failure in the network.

Configure the active LACP members on Aggregate-1 to Cisco Catalyst Ethernet Switch Module 3110switch-A as follows:

```
(config)# interface TenGigabitEthernet1/0/1
(config-if)# description <<** Connected to Switch-A **>>
(config-if)# channel-protocol lacp
(config-if)# channel-group 1 mode active
(config)# interface TenGigabitEthernet2/0/1
(config-if)# description <<** Connected to Switch-A **>>
(config-if)# channel-protocol lacp
(config-if)# channel-group 1 mode active
```

**Trunking Configuration**

Use the following guidelines when configuring trunks:

- Allow only those that are necessary on the trunk
- Use 802.1q trunking
- Tag all VLANs over a trunk from the aggregation switches

Configure trunks using the standard encapsulation method 802.1q as follows:

```
(config-if)# switchport trunk encapsulation dot1q
```

Define the VLANs permitted on a trunk as follows:

```
(config-if)# switchport trunk allowed vlan <VLAN IDs>
```

Modify the VLANs allowed on a trunk using one of the following commands:

```
(config-if)# switchport trunk allowed vlan add <VLAN IDs>
(config-if)# switchport trunk allowed vlan remove <VLAN IDs>
```

Define a port as a trunk port as follows:

```
(config-if)# switchport mode trunk
```

The autonegotiation of a trunk requires that the ports be in the same VLAN Trunking Protocol (VTP) domain and be able to pass Dynamic Trunking Protocol (DTP) frames.

To secure and enforce a spanning tree topology, configure the root guard feature on the aggregate switch interfaces that connect to the blade switches. The following is an example of the interface configuration between the aggregate and blade switch with root guard enabled:

```
(config)# interface GigabitEthernet12/13
(config-if)# description <text>(config-if)# no ip address
(config-if)# switchport
(config-if)# switchport trunk encapsulation dot1q
(config-if)# switchport trunk native vlan <vlan id>
(config-if)# switchport trunk allowed vlan <vlan id>
(config-if)# switchport mode trunk
(config-if)# spanning-tree guard root
(config-if)# channel-protocol lacp
(config-if)# channel-group <group id> mode active
```

The server ports on the blade switch support a single VLAN access and trunk configuration modes. The operational mode chosen should support the server's NIC configuration (that is, a trunking NIC is attached to a trunking switch port). Enable PortFast for the edge devices.

The BPDU Guard feature disables a port that receives a BPDU. This feature protects the Spanning Tree Protocol topology by preventing the blade server from receiving BPDUs. A port disabled with the BPDU Guard feature must be recovered by an administrator manually. Enable the BPDU Guard feature on all server ports that should not be receiving BPDUs.

Port Security limits the number of MAC addresses permitted to access the blade switch port. Configure the maximum number of MAC addresses expected on the port.

**Note:** The NIC teaming driver configuration (that is, the use of a virtual MAC address) must be considered when configuring Port Security.

```
interface GigabitEthernet1/0/1
 description <<** BladeServer-1 **>>
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10, 60
 switchport mode trunk
 switchport port-security aging time 20
 switchport port-security maximum 1 vlan 10, 60
 no cdp enable
 spanning-tree portfast trunk
 spanning-tree bpduguard enable
end
```

#### Server Default Gateway Configuration

The default gateway for a server is a Layer 3 device located in the aggregation layer of the data center. This device may be a firewall, a load balancer, or a router. Using protocols such as HSRP protects the gateway from being a single point of failure and creates a highly available data center network. HSRP allows the two aggregate switches to act as a single virtual router by sharing a common MAC and IP address between them. Define a switched virtual interface (SVI) on each aggregate switch and use the HSRP address as the default gateway of the server farm.

Configure Aggregation-1 as the active HSRP router. The priority command helps to select this router as the active router because it has a greater value.

```
interface Vlan10
 description <<** BladeServerFarm – Active **>>
 ip address 10.10.10.2 255.255.255.0
 no ip redirects
 no ip proxy-arp
 arp timeout 200
 standby 1 ip 10.10.10.1
 standby 1 timers 1 3
 standby 1 priority 51
 standby 1 preempt delay minimum 60
 standby 1 authentication <password>
end
```

Configure Aggregation-2 as the standby HSRP router as follows:

```
interface Vlan10
 description <<** BladeServerFarm – Standby **>>
 ip address 10.10.10.3 255.255.255.0
 no ip redirects
 no ip proxy-arp
 arp timeout 200
 standby 1 ip 10.10.10.1
 standby 1 timers 1 3
 standby 1 priority 50
 standby 1 preempt delay minimum 60
 standby 1 authentication <password>
end
```

### RSPAN Configuration

RSPAN allows for remote traffic monitoring in the data center. Define source and destination sessions to mirror interesting traffic to a remote VLAN captured by network analysis tools. Configure a VLAN for RSPAN on the Cisco Catalyst Ethernet Switch Module 3110 and the aggregate switch as follows:

```
(config)# vlan <vlanID>
(config-vlan)# name <vlan name>
(config-vlan)# remote-span
```

Create a source session as follows. This is the interface or VLAN that contains interesting traffic.

```
(config) # monitor session <session id> source vlan <VLAN IDs>
```

Configure the RSPAN VLAN as the target for the mirrored traffic as follows:

```
(config) # monitor session <ID> destination remote vlan <remote vlan
ID>
```

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.



Printed in USA C07-455528-00 05/08