Cisco Systems | Corporate Alliance

hp invent

Integrating the Cisco Gigabit Ethernet Switch Module (CGESM) for HP BladeSystem p-Class into the Cisco Data Center Network Architecture

Design Guide

"Design practices for deploying Cisco Gigabit Ethernet Switch Modules (CGESMs) for the HP BladeSystem p-Class into the Cisco Data Center Network Architecture"

Rev 3
March 15, 2005

Contents

This document provides best design practices for deploying the Cisco Gigabit Ethernet Switch Modules (CGESM) for the HP BladeSystem p-Class enclosures within the Cisco Data Center Networking Architecture. This document describes the internals of the blade enclosure and CGESM and explores different methods of deployment. It includes the following sections:

- The Cisco Gigabit Ethernet Switch Module (CGESM)
- CGESM Features
- Design Goals
- Design and Implementation Details

The Cisco Gigabit Ethernet Switching Module (CGESM)

This section briefly describes the CGESM and explains how the blade servers within the HP BlaseSystem are physically connected to the switching module.

The CGESM provides enhanced Layer 2 services (known as L2+ or Intelligent Ethernet) to the HP BladeSystem p-Class. The CGESM extends the capabilities of a Layer 2 Ethernet switch to include Cisco proprietary protocols, ACLs, and QoS based on Layer 3 information. With SNMP, CLI, or HTTP management options available and a robust set of IOS switching features, the CGESM 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 (PVST+), 802.1w, 802.1s, BDPU Guard, Loop Guard, PortFast, UplinkFast and UniDirectional Link Detection (UDLD)
- Advanced management protocols, including Cisco Discovery Protocol(CDP), 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 HP BladeSystem p-Class enclosure consists of eight server bays and two network-interconnect bays. Figure 1 depicts the BladeSystem p-Class architecture using two CGESM housed in the network interconnect bays and eight BL20P G3 servers.

Figure 1 BladeSystem p-Class Switch Architecture

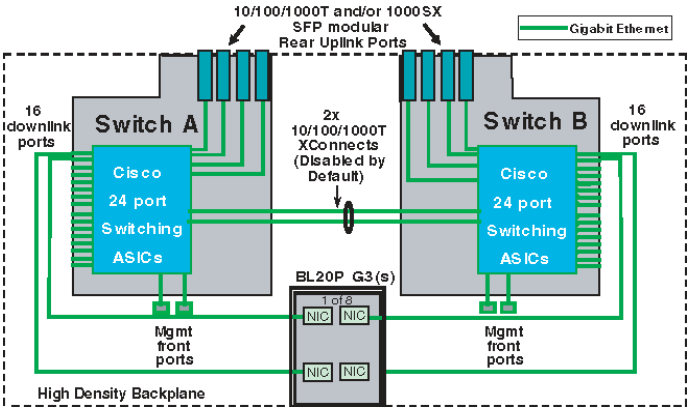

The HP BladeSystem p-Class backplane provides power and network connectivity to the blades. The interconnect bays house a pair of CGESMs, which provide a highly available and multi-homed environment where each server blade is Gigabit attached to each CGESM. Figure 2 illustrates how the HP BladeSystem p-Class logically provides Ethernet connectivity.

Figure 2: Blade Enclosure Ethernet Connectivity

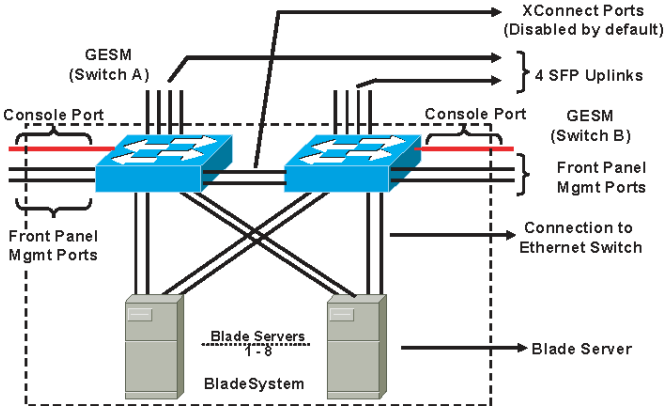

Note Figure 2 is based on the use of the HP BladeSystem p-Class using BL20P G3 servers. The remainder of this document will use the BL20P G3 server for all figures.

In the illustration of Figure 2 above, two CGESMs within the blade enclosure connect the blade server

modules to external network devices such as aggregation layer switches. Each Ethernet switch provides six external Ethernet ports for connecting the blade enclosure to the external network. Four SFP ports provide 1000 Base-SX and 10/100/1000 Base-T links on the rear of the enclosure and two 10/100/1000 Base-T ports provide connectivity on the front panel. All six of these ports can be grouped to support the 802.3ad link aggregation protocol. In Figure 2 above, each blade server is connected to the backplane via the available Gigabit Ethernet network interface cards (NICs). The number of NICs on each blade server varies. Table 1 provides more detail on the connectivity options available with each HP blade server and the maximum number of blade servers a single enclosure can support.

Table 1 Blade Server Options

Blade Server	**Maximum Number of Server Blades per Enclosure**	**NICs Available**
BL20P G2	8	3 10/100/1000T NICs 1 dedicated iLO interface
BL20P G3	8	4 10/100/1000T NICs 1 dedicated iLO interface
BL30P	16	2 10/100/1000T NICs 1 dedicated iLO interface
BL40P	2	5 x 10/100/1000T NICs 2 Slots for SAN Connectivity 1 dedicated iLO interface

Note In Table 1 iLO refers to the Integrated Lights-Out interface. It supports the iLO management subsystem that resides on each server blade. For more information on the iLO system refer to the "Management" section of this document.

In Figures 1 and 2 above, two NICs on each blade server connect to CGESM A and CGESM B. The blade servers connect to the CGESM switches over the HP BladeSystem p-Class backplane. There are sixteen 10/100/1000 internal ports on each CGESM dedicated to the blade server modules.

The Figures 1 and 2 also depicts two internal 10/100/1000 ports interconnecting the two blade enclosure switches over the backplane. These ports are disabled by default, but are configurable to carry all traffic types between the two blade enclosure switches. These ports support trunking and can be configured as channels.

Note In Figure 1 above, if the BL20p G1 and G2 servers were used, each server would dedicate one NIC connected to the CGESM B side to the iLO port. This NIC is only capable of 100 MB. The enclosure with an enhanced backplane provides blade server connectivity to an embedded iLO module on the backplane. See "Management" section for more details.

The HP BladeSystem p-Class enclosure with enhanced backplane consists of eight server bays and two network-interconnect bays. The HP BladeSystem p-Class sleeve option allows the enclosure to support 16 Proliant BL30p servers or two BL30p servers per bay. Figure 3 illustrates the Gigabit Ethernet connectivity of an enhanced backplane enclosure and 16 BL30p servers.

Figure 3 HP BladeSystem p-Class with 16 servers

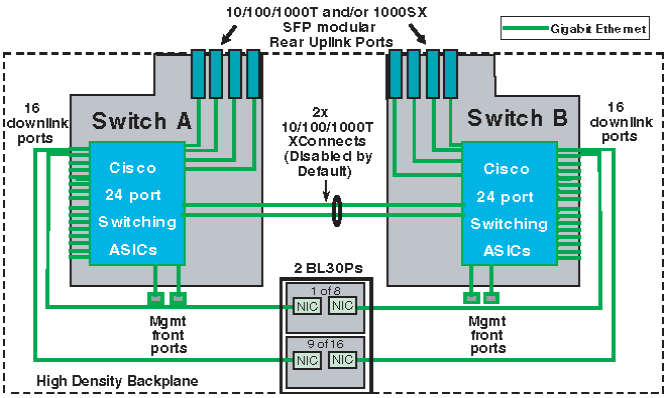

For more information about the HP BladeSystem p-Class, refer to the following URL: http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/documentation.html

CGESM Features

This section highlights information about the protocols and features provided by the CGESM that help integrate the HP BladeSystem p-Class enclosure into the Cisco Data Center Network Architecture. This section includes the following topics:

Spanning Tree

Traffic Monitoring

Link Aggregation Protocols

Spanning Tree

The CGESM supports different versions of the Spanning Tree Protocol (STP) and associated features, including the following:

Rapid Spanning Tree (RSTP) based on 802.1w

Multiple Spanning Tree (MST) based on 802.1s with 802.1w

Per VLAN Spanning Tree Plus (PVST+)

Rapid Per VLAN Spanning Tree Plus (RPVST+)

Loop Guard

Unidirectional Link Detection (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 key 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 There is no need to enable BackboneFast or UplinkFast.

In terms of convergence, STP algorithms based on 802.1w are much faster than the traditional STP 802.1d algorithms. The proposal agreement mechanism allows the CGESM to decide new port roles by exchanging proposals with its neighbors.

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

This seems to indicate that STP convergence time could be as long as 6 seconds. However, because 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 below *1 second* rather than 6 seconds.

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 misconfiguration of a device connected to a 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 CGESM supports Per-VLAN Spanning Tree (PVST) and a maximum of 128 spanning tree instances. RPVST+ is a combination of Cisco PVST Plus (PVST+) and Rapid Spanning Tree Protocol. RPVST+ provides the flexibility of one spanning tree instance per VLAN and fast convergence benefits of 802.1w. Multiple Spanning Tree (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 are 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.

Note The 802.1w protocol is enabled by default when running spanning tree in RPVST+ or MST mode on the CGESM. CGESM 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 EtherChannels. To allow for more granular STP calculations, enable the use of a 32-bit value instead of the default 16-bit value. The *longer* path cost better reflects changes in the speed of channels and allows STP to optimize the network in the presence of loops.

Note The CGESM supports IEEE 802.1t, which allows for spanning tree calculations based on a 32-bits path cost value instead of the default 16 bits. For more information about the standards supported by the CGESM, refer to the "Cisco Gigabit Ethernet Switch Module (CGESM) Overview" document.

For more information regarding spanning tree and Layer 2 design in the data center, refer to the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/cdccont_0900aecd800e4d2e.pdf

Traffic Monitoring

The CGESM supports 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 for 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.

Note RSPAN does not require a dedicated reflector port to mirror traffic from either a source port or source VLAN.

Link Aggregation Protocols

Fast EtherChannel (FEC) and Gigabit EtherChannel (GEC) are logically bundled physical interfaces that 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 CGESM supports 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.

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

Blade Enclosures in the Data Center Architecture

For details about this architecture, refer to the document at the following URL: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/cdccont_0900aecd800e4d2e.pdf

Data Center Network Components

The terms front-end or 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 4 depicts the front-end network and the services available at each layer.

Figure 4: Data Center Front-end Network

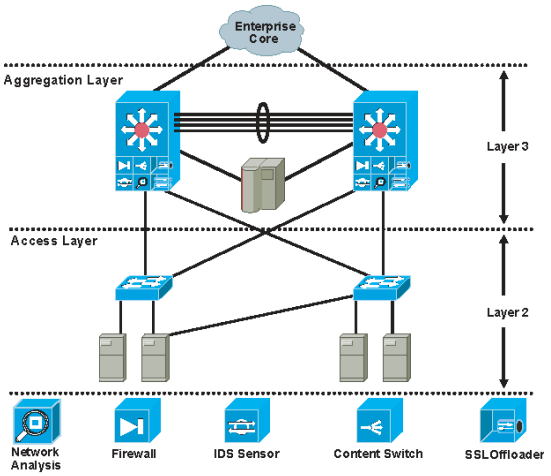

Aggregation Layer

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 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 server to 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) offloaders
- 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 farm 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 ensure a stable Layer 2 domain. A predictable access layer allows the spanning tree to converge and recover quickly during failover and fallback scenarios.

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 the 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. STP 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 Routing 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 impact 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 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 provide an even greater level of availability for the server and the applications it supports.

Using the HP BladeSystem p-Class enclosure in the Data Center Architecture

The HP BladeSystem p-Class enclosure supports a maximum of two internal CGESM enhanced Layer-2 switches. Each switch provides sixteen internal Gigabit Ethernet ports to support the blade servers within the enclosure. The HP BladeSystem p-Class supports up to eight blade servers (see Figure 5), each having multiple ProLiant NC series NICs (see Table 1). Note that the enclosure with enhanced backplane and supports up to sixteen blade servers (see Figure 6). Figure 5 and 6 illustrate the physical layout of the chassis. The two interconnect bays house the CGESM switches that are connected via two 10/100/1000 cross connects on the backplane. Each switch also has separate dual-backplane connections to the individual server blade bays. This indicates that each server blade is dual-homed to the two internal switches.

Figure 5: HP BladeSystem p-Class Enclosure with Proliant BL20p G3 Servers

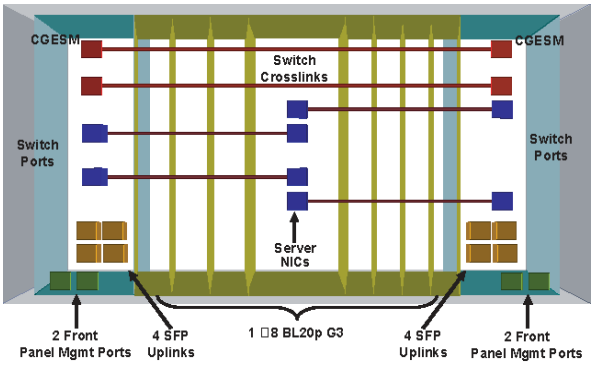

Figure 6: HP BladeSystem p-Class Enclosure with Proliant BL30p Servers

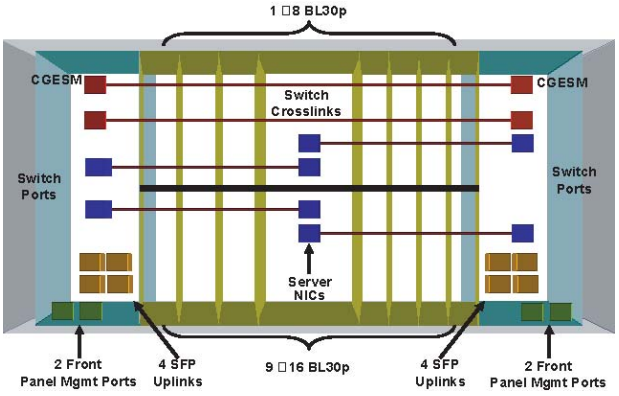

Each CGESM has four external SFP ports supporting 1000Base-SX and 10/100/1000Base-T on the rear of the enclosure and two external 10/100/1000BaseT ports on the front panel. These six ports provide connectivity to the data center or other external network. For more information, refer to the HP Blade Systems at the following URL:

http://h71028.www7.hp.com/enterprise/cache/80632-0-0-0-121.aspx#Servers

Design Goals

This section describes the design goals when deploying blade servers and the functionality supported by the CGESM in data centers. It includes the following topics:

High Availability

Scalability

Management

High Availability

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 around the clock. 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 NIC cards 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. An HP BladeSystem p-Class, using the CGESM, has a number of features and characteristics that contribute to a reliable, highly available network.

High Availability for the Blade Enclosure Switching Infrastructure

High availability between the HP BladeSystem p-Class CGESMs and the aggregation layer switches requires link redundancy. Each CGESM in the HP BladeSystem p-Class uses four SFP uplinks for connectivity to the external network which allows for redundant paths using two links each for more redundancy. Redundant paths implemented between the HP BladeSystem p-Class and each aggregation layer switch when each path uses two links provides a highly resilient design. However, this introduces the possibility of Layer 2 loops; therefore, a mechanism is required to manage the physical topology. The implementation of the Rapid Spanning Tree Protocol (RSTP) ensures a fast converging, predictable Layer-2 domain between the aggregation layer and access switches (the CGESMs) when redundant paths are present.

The recommended design is a triangle topology, 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.

As illustrated in

Figure 7, each CGESM switch has two direct port-channel connections to the Layer-2/Layer-3 aggregation layer where the primary STP root switch resides.

Figure 7 Blade Enclosure Redundant Topology

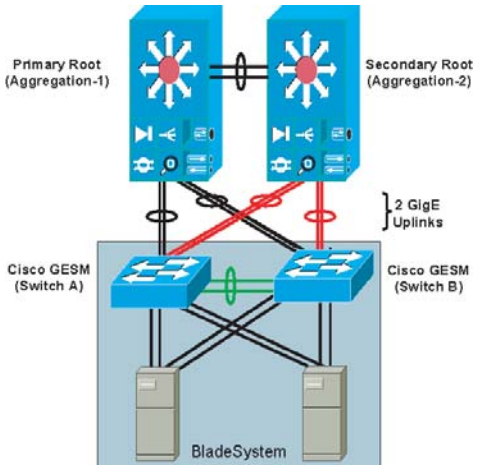

These channels support the publicly available subnets in the data center and traffic between servers. The server-to-server traffic that uses these uplinks is logically segmented through VLANs and may take advantage of network services available in the aggregation layer. There is also a port channel defined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

It is recommended to use RPVST+ 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 enclosure 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

The HP BladeSystem p-Class enclosure provides high availability to blade servers by multi-homing each server to the CGESMs. The two CGESMs housed in the interconnect bays are connected to the blade server over

the backplane. Four backplane Gigabit Ethernet connections are available to every blade-server slot.

Multi-homing the server blades allows the use of network adapter (NIC) teaming driver, which provides another high availability mechanism to failover and load balance at the server level. The ProLiant NC series NICs support three modes of teaming:

Network Fault Tolerance (NFT)

Transmit Load Balancing (TLB)

Switch Assisted Load Balancing (SLB)

NFT team 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 team 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 team 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 CGESM supports the IEEE 802.3ad standard and Gigabit port channels. For more information about the ProLiant NIC teaming features, refer to the following URL:
http://h18000.www1.hp.com/products/servers/networking/whitepapers.html

Scalability

The ability of the data center to adapt to increased demands without compromising its availability is a key 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

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.

As shown in

Figure 7, each CGESM requires four Gigabit Ethernet ports from the aggregation layer switches. The number of physical ports that an aggregation-layer switch can support equals the number of ports per-slot times the number of available slots.

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.

Table 2 lists the number of blade systems supported by a single line card with varying port counts. This table is based on the recommended topology depicted in Figure 4, where each blade system is dual-homed to the aggregation layer over two Gigabit Ethernet port channels.

Table 2: Blade System Support per Aggregate Switch Line Card

Type of Line Card	**Uplinks per CGESM**	**Total Uplinks /Blade System Enclosure (Two CGESM/Enclosure)**	**Blade Systems per Line Card**
8-port Gigabit Ethernet	2	4	2
16-port Gigabit Ethernet	2	4	4
48-port Gigabit Ethernet	2	4	12

Table 2 highlights the finite number of BladeSystems supported by a single aggregate switch line card. This table implies that the aggregate layer must provide linecard density for a scalable BladeSystem environment.

Slot Count

The data center infrastructure must be flexible enough to allow growth both in 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 sixteen 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 9 depicts a service switch deployment.

Figure 8 Data Center Scaling with Service Switches

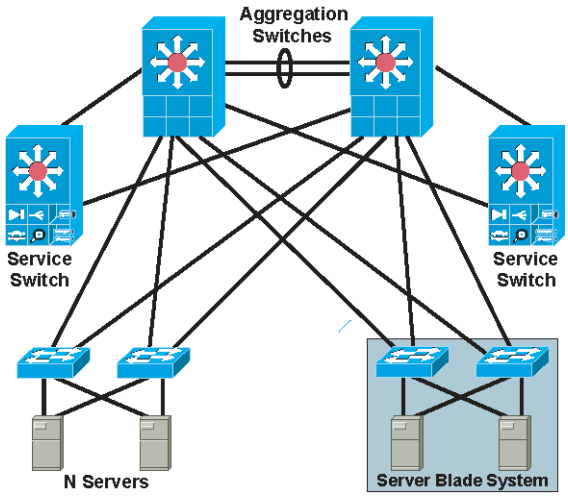

The data center core is a mechanism to replicate and horizontally scale the data center environment. In the recommended design the aggregation and access layer is regarded as a module that can be duplicated to extend the enterprise. Each data center module provides its own network services locally in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

Figure 9 depicts the data center core design. The aggregation switches for each data center module are Layer 3 attached to the core. In addition, the aggregation switches house the service modules required to support the server farms.

Figure 9: Data Center Core Design

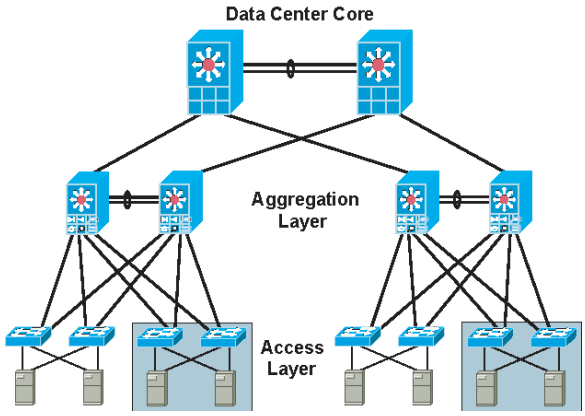

Management

The CGESM 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 CGESM. The remote management of the blade servers within the HP BladeSystem p-Class enclosure is critical to an efficient and scalable data center. The iLO connectivity options provided via the enclosure to the blade servers is also discussed.

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 BladeSystems. Out-of-band management isolates the management and data traffic and provides a more secure environment.

Figure 10 illustrates the interfaces available for connectivity. It is recommended to use the front panel ports for connectivity to the management domain.

The CGESM allows only one switched virtual interface (SVI) to be active. By default, the SVI is created as VLAN 1 and it is disabled in an administratively down state. Cisco recommends that a VLAN other than VLAN 1 be used as the management VLAN. Therefore, it is important to create an SVI with another VLAN and allow this VLAN on the external front panel ports.

For best practices in selecting the management VLAN, refer to the following URL:

http://www.cisco.com/en/US/partner/products/hw/switches/ps700/products_white_paper09186a00801b49a4.shtml

Figure 10: Blade Enclosure

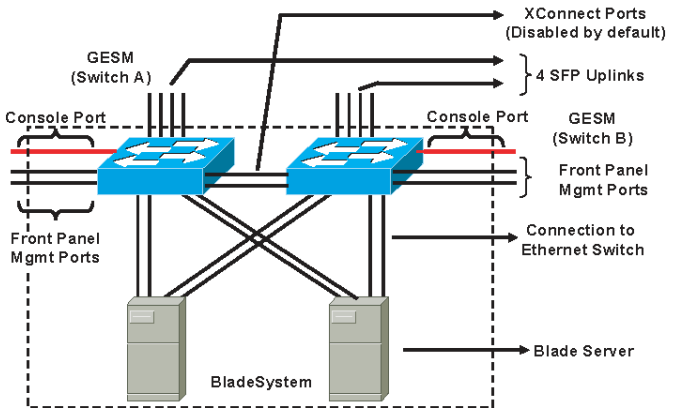

In-Band Management

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 CGESMs for management. It is recommended to use a VLAN other than VLAN 1 for management.

Serial/Console Port

The front panel of the CGESM has a single RJ-45 serial port that can be used to manage the switch through the command-line interface (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.

Management Options

The CGESM switch is manageable through the following methods:

HTTP-based device manager GUI
SNMP-based management applications
Cisco IOS software CLI

The embedded device manager on the CGESM provides a GUI interface 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. CiscoWorks2000 and HP OpenView are two such management applications. SNMP versions 1, 2, and 3 are available on the switch (Cisco IOS software crypto image).

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 manger, refer to the online help on the switch CLI.

For more information about Cisco MIBs, refer to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

For more information about the management options for the HP Blade System, refer to the following URL: http://h18004.www1.hp.com/products/blades/components/management.html

HP BladeSystem p-Class iLO Connectivity

The iLO provides remote management capabilities and is standard with BL p-Class server blades. Remote power, console, and diagnostics are just a few of the advanced functions iLO provides. Table 1 indicated that each of the blade servers supports a dedicated iLO NIC. The HP BladeSystem p-Class enclosure provides two methods to access this management interface:

Through the CGESM located in interconnect bay B

Through an enhanced BladeSystem enclosure

An HP BladeSystem p-Class enclosure without the enhanced backplane features provides connectivity to the dedicated iLO NIC through the CGESM located in interconnect bay B. The iLO NIC autonegotiates to 100 Mbps and uses one of the CGESM ports assigned to that server bay. Figure 11 illustrates the iLO interface connection in the absence of an enhanced backplane.

Figure 11 HP BladeSystem p-Class iLO Connectivity

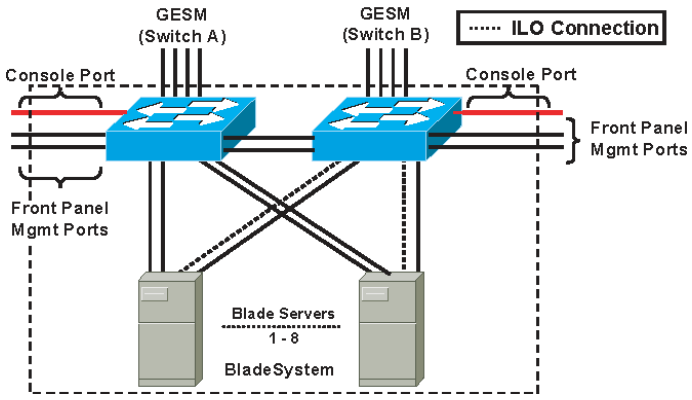

The enhanced backplane of the HP BladeSystem p-Class enclosure allows each server to use a dedicated Ethernet port for iLO connectivity. As shown in Figure 12, the iLO connection is independent of the CGESM. The blade server management module located on the rear of the chassis provides access to each of the iLO interfaces through a single RJ-45 cable.

Figure 12: HP BladeSystem pClass with Enhanced Backplane iLO Connectivity

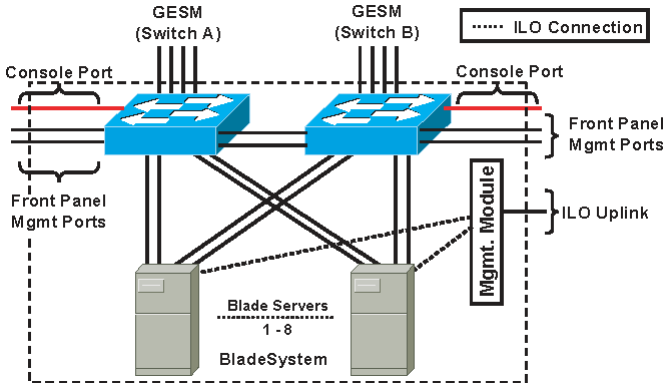

Note The Proliant BL30p server blade requires the use of an enhanced backplane.

Design and Implementation Details

Network Management Recommendations

Network Topologies Using CGESM

Configuration Details

Network Management Recommendations

An out-of-band (OOB) network is recommended for managing the CGESM switch. 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 CGESM switch has two external Gigabit Ethernet ports located at the front of the chassis that may be used to support network monitoring devices and network management traffic. Using secure protocols, such as SSH or HTTPS maintains the integrity of communications between the switch and the management station. The console port positioned at the front of the CGESM is another option for connectivity to the OOB network.

Network Topologies using the CGESM

The following physical topologies are discussed in this section:

Recommended Topology – Classic "V" shaped topology with STP

Alternate Topology – Square topology with STP

These network designs emphasizes 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 as the aggregation layer platform. This Layer-2/Layer-3 switching platform supports slot density and integrated network services required by data centers deploying blade systems. An HP BladeSystem p-Class server blade enclosure with two CGESMs composes the Layer-2 access layer.

Recommended Topology

Figure 13 depicts a blade system deployment in the data center using the classic triangle topology. There is no single point of failure in this deployment model. The CGESMs are dual-homed to the aggregation layer, providing link redundancy. STP manages the physical loops created by the uplinks between the aggregation and access switches, assuring a predictable and fast converging topology. In Figure 13 the black links are in spanning tree forwarding state and the red links are in blocking state. The green links represent the internal cross connects that are disabled by default.

Figure 13: Recommended Topology HP BladeSystem p-Class with CGESMs

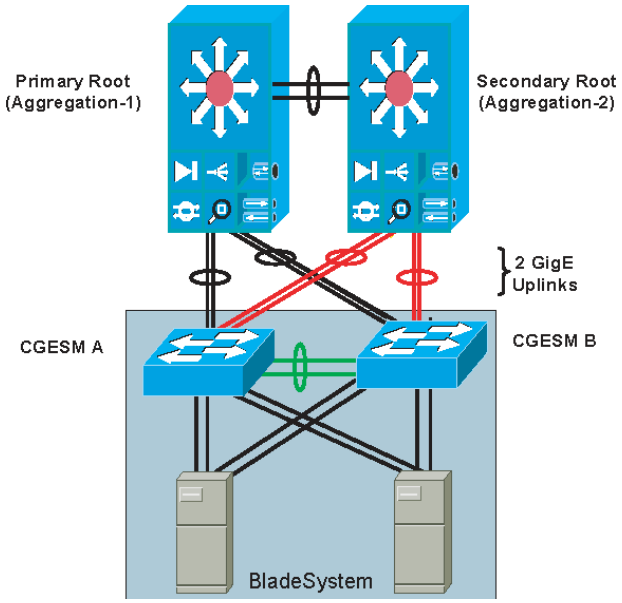

RPVST+ fulfills the high availability requirements of this design and is the recommended mode of spanning tree operation. RPVST+ provides fast convergence (less than 1 second) in device or uplink failure scenarios. In addition, RPVST+ offers enhanced Layer 2 features for the access layer with integrated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

The connection between the two internal blade switches in Figure 13 supports local traffic limited to the HP BladeSystem; for example, clustering applications, or management traffic such as remotely mirrored (RSPAN) traffic. This connection does not carry a publicly accessible subnet (for example, a VLAN that exists on the uplinks to the aggregation switches). If this were the case, another set of interfaces would have to be accounted for in the STP calculations. Therefore, to create a less complex STP domain, these cross-connect interfaces are removed from the equation by clearing the public VLANs from the links.

The HP BladeSystem p-Class server blade NICs support the logical separation of VLANs via trunking. This allows each NIC to accommodate the public and the private VLANs on the CGESMs. In addition, servers such as the BL20P G3 series are dual-homed to each of the two CGESMs in the HP BladeSystem enclosure (see Figure 13). This structural design allows for the physical

separation of public and private VLANs between two NICs homed to the same CGESM.

A series of network convergence tests were performed to verify and characterize the high availability features of the recommended design. These tests consisted of passing traffic between an external client device and the blade servers while monitoring packet loss. The following test cases were used:

- Uplink failure and recovery between Switch-A and the primary root
- Uplink failure and recovery between Switch-B and the primary root
- Switch-A failure and recovery
- Switch-B failure and recovery
- Primary root switch failure and recovery
- Secondary root switch failure and recovery

These tests revealed the intricacies of fast convergence in the data center and the necessity for a holistic approach to high availability.

Test cases that did not involve the failure of the active HSRP aggregation switch resulted in an average failover time of about 1 second. Failing the active HSRP device requires convergence at Layer 3 and resulted in a recovery time that reflected the settings of the HSRP timers.

It is possible to tune the HSRP timers for sub-second convergence However, when multiple HSRP devices are involved the recovery time is typically in the 5-second range.

In this topology, two Gigabit Ethernet links compose the port channel uplinks between the access and aggregation layers. This configuration allows a single link to fail without triggering STP convergence.

Note The default gateway for the servers is the HSRP address of the Layer 3 aggregation switches. Failover times may be affected if the default gateway of the server is located on another device, such as a load-balancer or firewall.

The recommended topology provides a high level of availability to the blade servers except in one failure scenario. If both the uplinks to each of the aggregation switches from a single CGESM are unavailable, the server NICs homed to that CGESM are not notified. The blade servers are unaware of the disconnection between the access layer switches (CGESMs) and the aggregation layer switches and continue to forward traffic.

To address this breakdown in network connectivity, use one of the following methods:

- Use the NIC teaming features of the ProLiant blade servers
- Deploy load balancing in front of the blade server farm

In addition, the NIC teaming features of the ProLiant 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 interface or a small configuration file, depending on the operating system, to construct the team. HP also offers network-aware teaming software to verify and detect network routes. For more information about these features, refer to the ProLiant Essential Intelligent Network Pack at the following URL:

http://h18004.www1.hp.com/products/servers/proliantessentials/inp/index.html

A load balancer, by monitoring the health of a server farm, can bypass the network failure by redirecting traffic to available servers. This helps ensure fulfillment of end user requests despite the network failure.

The recommended network topology illustrated in Figure 13 allows for traffic monitoring either locally or remotely using Switched Port Analyzer (SPAN). Local SPAN supports monitoring of network traffic within one switch, while remote SPAN (RSPAN) allows the destination of mirrored traffic to be another switch within the data center. The source of mirrored traffic for a SPAN or RSPAN session may be one or more ports or VLANs.

Local SPAN is readily supported by the CGESM over one of the two external Gigabit Ethernet ports located on the front panel of the switch. This RJ-45 connection is an ideal location to attach intrusion detection or other network analysis device.

RSPAN requires a VLAN to carry the mirrored traffic to the remote destination switch. In the recommended topology, the secondary aggregation switch is the RSPAN destination, where an analysis device, such as the integrated Network Analysis Module (NAM), resides.

Figure 14 illustrates the traffic path of the RSPAN VLAN. The RSPAN VLAN uses the uplink between the blade switch and the secondary aggregation switch. This uplink is blocking under normal conditions for regular VLANs. As a result, bandwidth utilization is only a concern when the uplink is forwarding and sharing the path with production traffic.

Figure 14: RSPAN Traffic Path

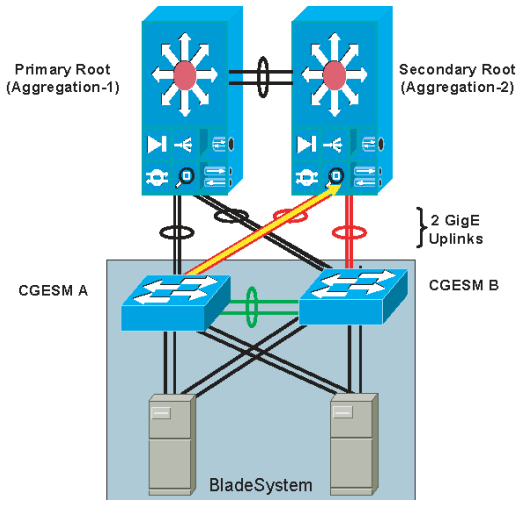

Configuring the Aggregate Switches

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 and CGESM switches

Step 6: Trunking the 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 CGESM Switches

Complete the following steps on the CGESM switches:

Step 1: VLAN configuration

Step 2: RPVST+ configuration

Step 3: Configuration of port channels between the CGESM and aggregate switches

Step 4: Trunking port channels between the CGESM and aggregate switches

Step 5: Configuration of server ports on the CGESM

Note The "Configuration Details" section describes each of these steps.

Additional Aggregation Switch Configuration

The following recommendations help integrate the CGESM switches 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 switch 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 CGESM 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 switches 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.

Alternate Topology

Figure 15 illustrates an alternative topology that relies on STP, specifically RPVST+, to account for redundant paths in the network. The two aggregate switches connect to each other via a port-channel supporting the server farm VLANs. The four external uplinks of each CGESMs are channeled and connected to one of the two aggregate

switches. The internal connections between the two CGESMs complete the loop.

Figure 15: Alternate Topology HP BladeSystem p-Class with CGESMs

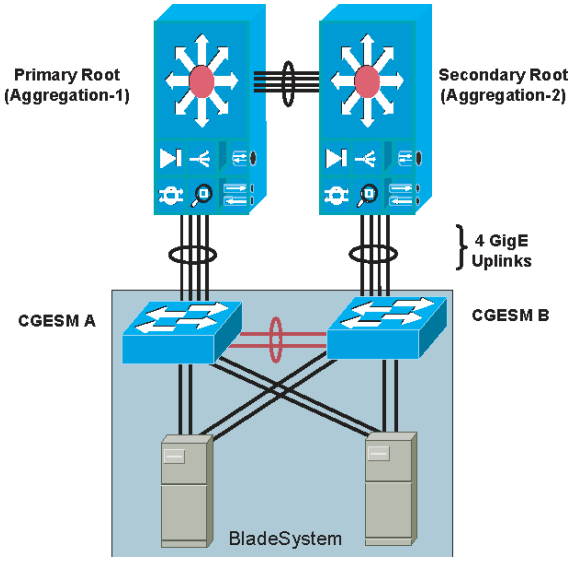

This design uses the links between the two CGESMs as a redundant path for blade server traffic. In Figure 15, the black links are in spanning tree forwarding state and the red links are in blocking state. These links are in blocking state by default. The use of a longer path cost value provides for a more granular calculation of the topology based on the available link bandwidth (see the "CGESM 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.

Note This design uses a lower bandwidth path when an uplink failure occurs on either CGESM A or CGESM B. To increase the bandwidth of the redundant path between the CGESMs, consider using the external ports of CGESM A and B in the EtherChannel.

The following convergence tests were conducted against this alternate 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.

The design in Figure 15 supports traffic monitoring via SPAN and/or RSPAN. For example, a network analysis device connected to the external ports on the front of the CGESM may capture locally mirrored traffic. Alternatively, RSPAN traffic may be carried on the CGESM uplinks if bandwidth utilization is not a concern. For the steps to configure traffic monitoring, refer to the "Configuration Details" section.

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 and CGESM switches

Step 6: Trunking the 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 CGESM Switches

Complete the following steps on the CGESM switches:

Step 1: VLAN configuration

Step 2: RPVST+ configuration

Step 3: Configuration of port channels between the CGESM and aggregate switches

Step 4: Trunking port channels between the CGESM and aggregate switches

Step 5: Configuration of server ports on the CGESM

Note The "Configuration Details" section describes each of these steps.

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+

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 as follows:

```
(config)# vtp domain <domain name>
(config)# vtp mode transparent
```

Configure the server farm VLANs as follows:

```
(config)# vlan 60
(config-vlan)# name bladeservers
(config-vlan)# state active
```

RPVST+ Configuration

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

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

Configure the path cost to use 32 bits in the STP calculations:

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

Configure the root switch as follows:

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

Configure the secondary root switch as follows:

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

Inter-Switch Link Configuration

The topologies discussed in this guide require connectivity between the switches. The following three types of inter-switch connections exist:

Aggregate-1 to Aggregate-2

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

HP BladeSystem Switch-A to Switch-B

Each of these connections are Layer 2 EtherChannels 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 EtherChannels 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 CGESM Switch-A as follows:

```
(config)# interface GigabitEthernet12/1
(config-if)# description <<** Connected to Switch-A **>>
(config-if)# channel-protocol lacp
(config-if)# channel-group 1 mode active
(config)# interface GigabitEthernet11/1
(config-if)# description <<** Connected to Switch-A **>>
(config-if)# channel-protocol lacp
(config-if)# channel-group 1 mode active
```

Configure the passive LACP members on CGESM Switch-A as follows:

```
(config) # interface GigabitEthernet0/19
(config-if)# description <<** Connected to Aggregation-1 **>>
(config-if)# channel-group 1 mode on
(config) # interface GigabitEthernet0/20
(config-if)# description <<** Connected to Aggregation-1 **>>
(config-if)# channel-group 1 mode on
```

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
```

Note The autonegotiation of a trunk requires that the ports be in the same VTP domain and be able to pass 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
```

Server Port Configuration

A blade server is assigned a specific port on the blade switch. This is pre-determined by the physical slot the blade server occupies in the chassis. Table 3 correlates the server and switch ports.

Table 3: Correlation of Server and Switch Ports

IOS CLI Identifier	Actual Port Location in 8-Slot Server Chassis	Actual Port Location in 16-Slot Server Chassis
GigabitEthernet 0/1	Server Slot 1	Server Slot 1
GigabitEthernet 0/2	Server Slot 1	Server Slot 2
GigabitEthernet 0/3	Server Slot 2	Server Slot 3
GigabitEthernet 0/4	Server Slot 2	Server Slot 4
GigabitEthernet 0/5	Server Slot 3	Server Slot 5
GigabitEthernet 0/6	Server Slot 3	Server Slot 6
GigabitEthernet 0/7	Server Slot 4	Server Slot 7
GigabitEthernet 0/8	Server Slot 4	Server Slot 8
GigabitEthernet 0/9	Server Slot 5	Server Slot 9
GigabitEthernet 0/10	Server Slot 5	Server Slot 10
GigabitEthernet 0/11	Server Slot 6	Server Slot 11
GigabitEthernet 0/12	Server Slot 6	Server Slot 12
GigabitEthernet 0/13	Server Slot 7	Server Slot 13
GigabitEthernet 0/14	Server Slot 7	Server Slot 14
GigabitEthernet 0/15	Server Slot 8	Server Slot 15
GigabitEthernet 0/16	Server Slot 8	Server Slot 16
GigabitEthernet 0/17	Cross Connect Port 1	Cross Connect Port 1
GigabitEthernet 0/18	Cross Connect Port 2	Cross Connect Port 2
GigabitEthernet 0/19	SFP/Uplink Port 1	SFP/Uplink Port 1
GigabitEthernet 0/20	SFP/Uplink Port 2	SFP/Uplink Port 2
GigabitEthernet 0/21	SFP/Uplink Port 3	SFP/Uplink Port 3
GigabitEthernet 0/22	SFP/Uplink Port 4	SFP/Uplink Port 4
GigabitEthernet 0/23	RJ45/Front Panel Port 1	RJ45/Front Panel Port 1
GigabitEthernet 0/24	RJ45/Front Panel Port 2	RJ45/Front Panel Port 2

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 (i.e., 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 STP topology by preventing the blade server from receiving BPDUs. A port disabled via 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 (i.e. the use of a virtual MAC address) must be considered when configuring Port Security.

```
interface GigabitEthernet0/1
 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 protect the gateway from being a single point of failure and create a highly available date 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 VLAN Interfaces (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 CGESM 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 <session ID> destination remote vlan
<remote vlan ID>
```

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
www.cisco.com/go/ibm

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185 USA
http://www.hp.com