



White Paper

Data Center Switching Solutions

INTRODUCTION

The Data Center is the consolidation point for provisioning multiple services that drive an Enterprise business. Ethernet switching technology is the foundation upon which many of these services are built. The requirement to serve the Data Center environment mandates that a solid design methodology be used to build the network. Moreover, support for key technologies must also be present to allow it to fulfill its role as the foundation for delivering multiple services.

This document will explore the key technologies that allow Ethernet switching platforms to successfully serve the needs to the Data Center. It will summarize the suggested topology and layers for a Data Center. Further detail will be provided to explain features that enhance the operational capabilities of the switching platforms at each layer of the Data Center topology.

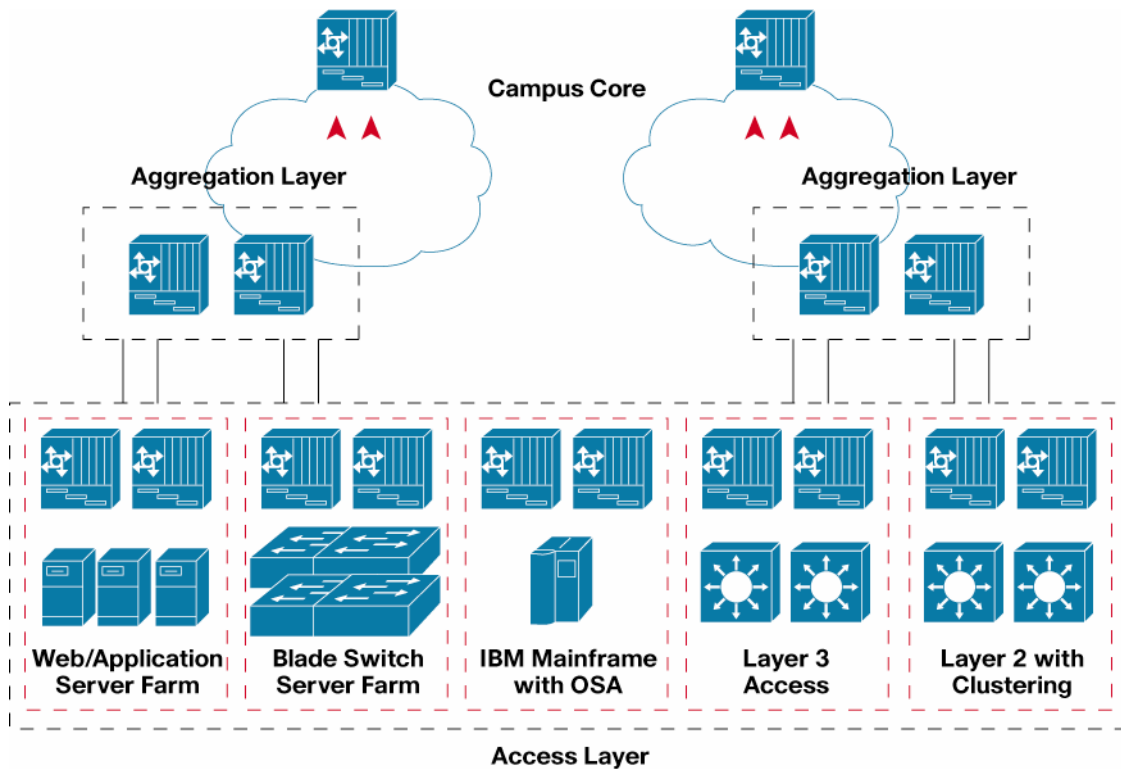
DESIGNING THE DATA CENTER LAYERS

The functional layers designed to serve the needs of the Data Center are typically built around three layers—the Core Layer, the Aggregation Layer and the Access Layer.

The Core Layer is central to the Data Center network and provides interconnection between the Aggregation Layers. Typically, the Core Layer utilizes high performance low latency switches providing high densities of 10GE. The use of 10GE to link up to the Aggregation Layer Switches is highly recommended. Switches at this layer operate exclusively as Layer 3 devices. The target Cisco device recommended to serve the needs of the Core Layer is the Catalyst 6500 switch.

The Aggregation Layer acts as a Services Layer for the Data Center. Services such as Load Balancing, SSL Optimization, Firewalling, etc are typically found at this layer. Multiple Access Layer switches will also use the Aggregation Layer as an interconnection point. The use of 10GE links to uplink into the Core Layer is a common practice. More of an emerging trend is the use of 10GE links to downlink into the Access Layer providing higher bandwidth and future proofing the network. The switch of choice for deployment in the Aggregation Layer is the Catalyst 6500 Switch.

Figure 1. Data Center Topology



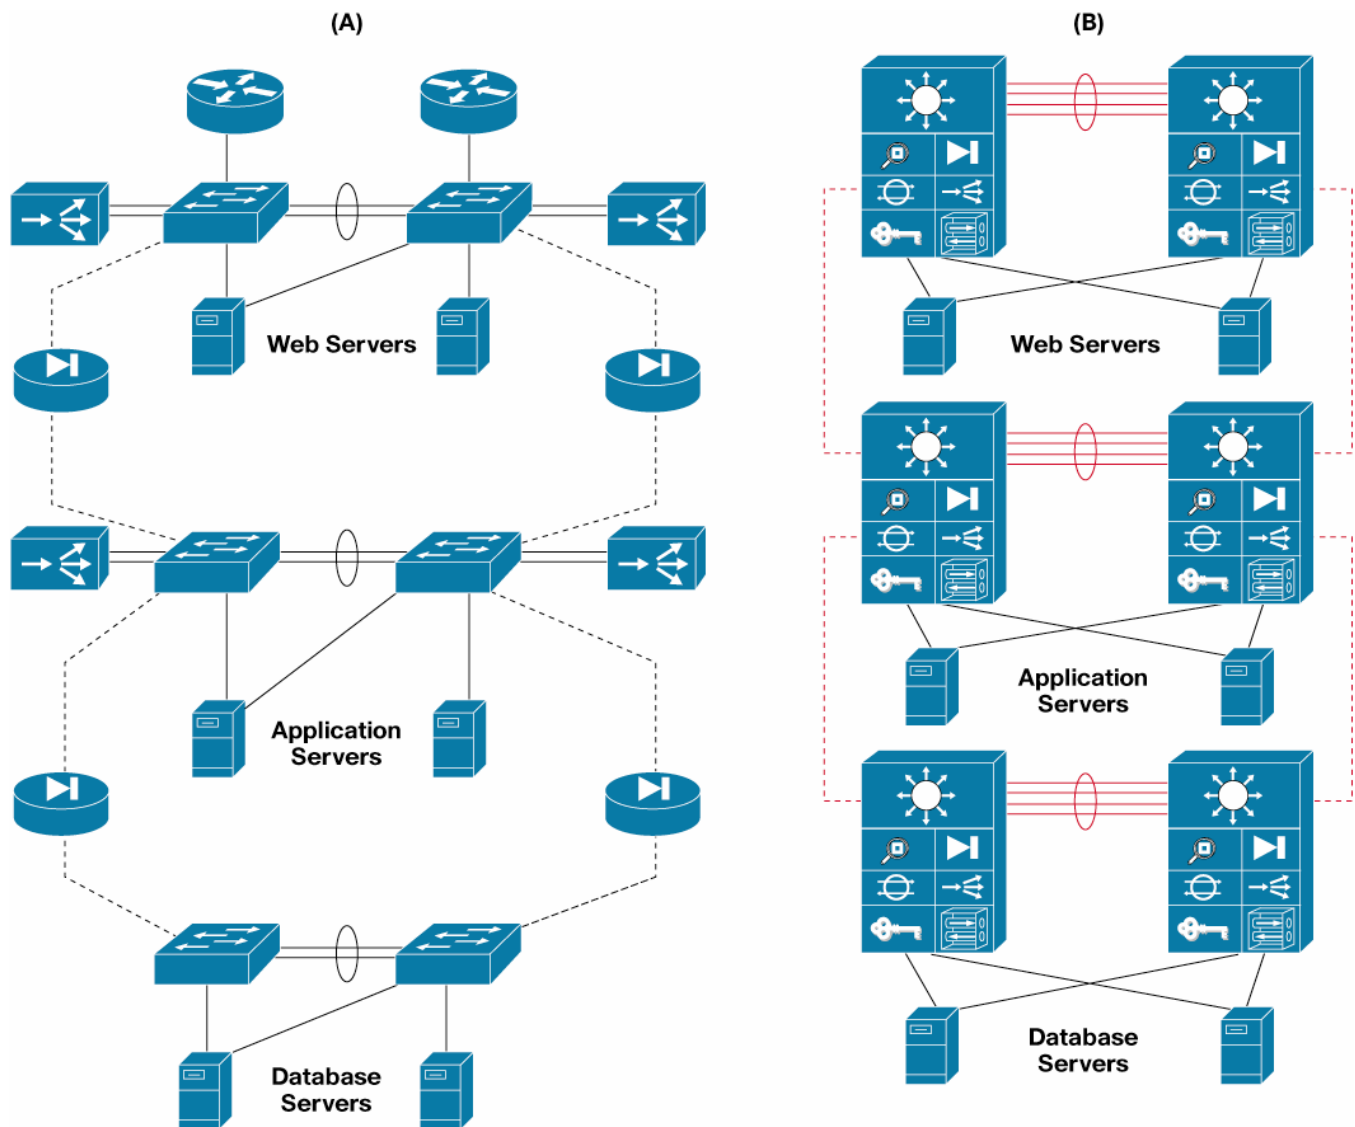
The Access Layer provides connectivity for the many servers that deliver application and web services to the business as well as the interconnections for a server cluster design. Access Layer switches can be configured for both Layer 2 and Layer 3 deployments. An access Layer switches may be required to support both single and dual homed servers. Flexibility to extend VLAN's between Access Layer switches is an important feature that should be catered for in the design. Use of multiple uplink options (to the Aggregation Layer) is also used to improve network resilience. These designs have been confirmed and ratified though internal Cisco testing.

Traditionally, the Catalyst 6500 has been deployed at most customers at the Access Layer. It can aggregate 100s of servers and has the scalability and resiliency to support hundreds of servers. In addition, two new Access layer switching topologies are emerging driven by the choice of servers—top-of-rack and blade switches. Cisco offers Catalyst 4948G as a top-of-rack switch and Catalyst Blade Switches as blade switches. The Catalyst 4948G can aggregate up to 48 servers, while the Catalyst Blade Switch usually aggregates 10-16 servers. Hence, the scalability, resiliency and manageability capabilities are designed to meet those needs.

Two standard design approaches (Multi-Tiered Design and Server Cluster Design) are often used when designing a Data Center and are based on tried and proven design methodologies tested in some of the largest Data Centers Cisco has had the opportunity to work with. These designs are engineered to provide improvements in performance, scalability, resiliency and flexibility.

The Multi-Tiered Design for the Data Center is the first of these design options and is geared to support a tiered approach to serving HTTP applications. This design approach happens to be the most common design model used for Data Centers today. Three tiers of servers are used, those being Web, Application and Database. These multi tiered server farms offer improvements in both resiliency and security.

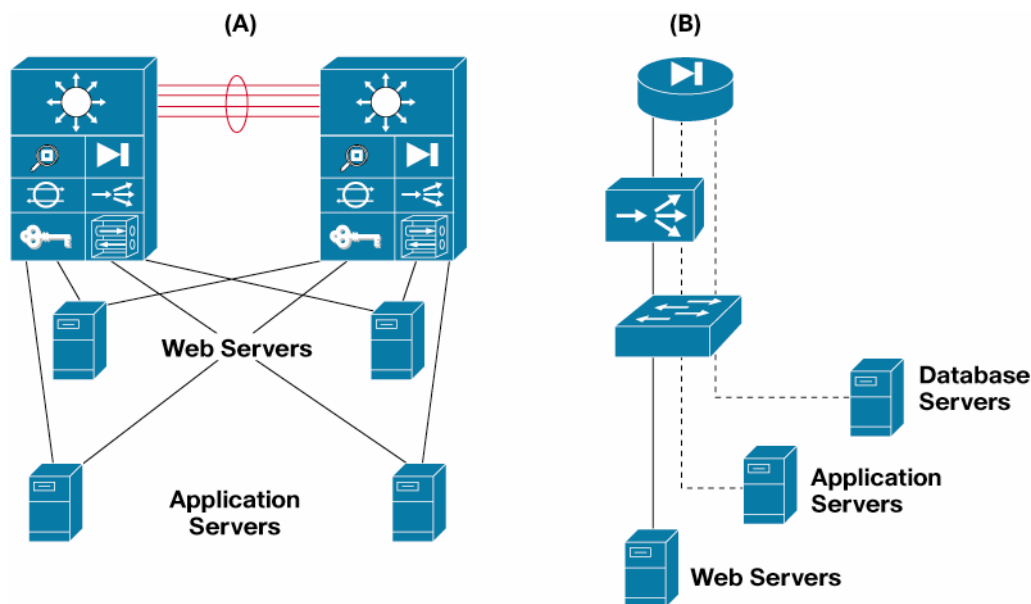
Figure 2. Multi-Tiered Design Model



Aggregating servers into separate functional layers further bolsters resiliency in this design model. A server can be taken out for maintenance, for example, while other servers in the same tier can continue to provide service. As can be seen in the diagram above, a three-tiered server approach is used to separate server functionality into a web layer, application layer and database layer. On the left is a view of the network utilizing stand-alone appliances while on the right is a view of the network achieving the same aim with an integrated services module approach. Added resiliency is achieved by load balancing traffic between the tiers.

Security enhancements are gained by utilizing firewalls between the tiers. VLAN's complement the implementation of firewalls by segregating the server farms into functional groups. The use of VLAN's lends itself to the switch architecture, which along with the firewalls and load balancers is VLAN aware and use VLAN technology to segregate server farms. Segmentation also improves performance when the choice of allowing each tier to use dedicated hardware is taken. Simplifying the complexity of managing the server farm is inherently built in through the use of VLAN's. The view of a server tier is depicted in the diagram below with the physical view on the left and the logical segmented view on the right.

Figure 3. Server Farm Segmentation with VLAN's

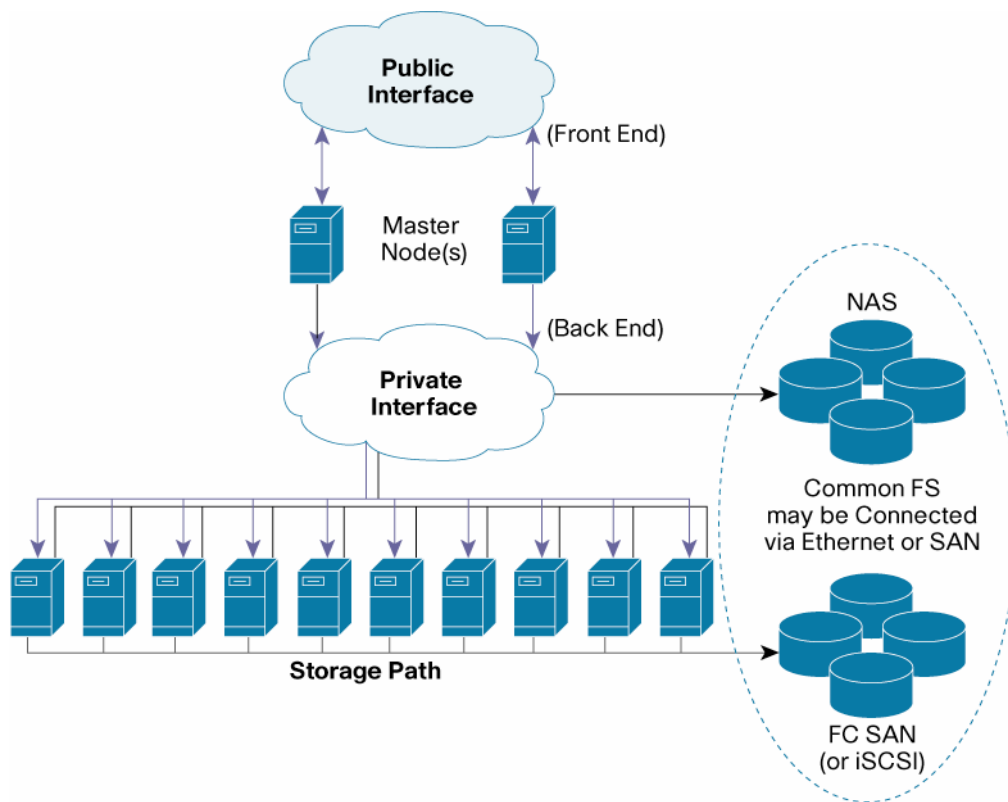


The Server Cluster Design model is the other major design model advocated for Data Centers. This model is engineered for high performance server clusters typically found in high performance computing centers. Server clusters have unique operational requirements and should be served by a slightly different design approach. High performance interconnects; low latency switching and large densities of high bandwidth interfaces are common inclusions in this design.

While this approach has typically been implemented in the past to meet specific scientific and military applications, more and more enterprise Data Centers are now adopting this model to meet the new application demands. Specifically, finance data centers requiring financial analysis, major film companies rendering animations for films, design modeling for manufacturing and search engines are all examples of enterprise-based customers looking to use this design model for their Data Center.

An example of the server cluster design can be seen in the following diagram.

Figure 4. Server Cluster Design



The elements of this design incorporate a master node, which is responsible for directing the other computing nodes that make up the cluster. These other computing nodes are collectively responsible for processing the application and associated data. A high-speed fabric typically links these nodes up to provide a low latency path between nodes. A storage path to a common file system rounds off the elements of this cluster providing a storage point for application data. High density (10GE) and low latency switching are but a few of the key drivers that the network must deliver to make this design successful. Other technologies discussed in the next section, will also play a critical role in the successful implementation of the server cluster design.

TECHNOLOGY DRIVERS FOR DATA CENTER SWITCHING

Data Center design is the starting point in driving towards a successful Data Center implementation. However, the switching technology that is proposed in the design should not only be about simply offering Ethernet connectivity. Features are key to a Switch successfully serving the needs of each layer of the Data Center, and more importantly, to the business needs served by the Data Center. Feature requirements can differ slightly depending on the target layer, the Data Center design model being implemented and where the switch is being deployed. The categories under which the required Data Center features can be grouped are Comprehensive Resilience, Architectural Scalability, and Operational Manageability. Examples of features that serve the requirements for each of these groups are explored in further detail in the following section.

Comprehensive Resilience (CR)

Business demands on application and data availability are always increasing, which in turn applies pressure on Data Center technology to maximize the uptime of network services. Data Center switches must look to implementing numerous High Availability features to maintain availability. The following section details some of the HA features that can be used to drive up network availability in the Data Center.

CR—Non-Stop Forwarding/Stateful Switchover (NSF/SSO)

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation, DC Access

Core and Aggregation Layers using dual Supervisors for improved resilience can take advantage of NSF/SSO to further bolster network availability. The concept behind NSF/SSO is to allow the redundant Supervisor to maintain state with the primary Supervisors forwarding tables. It does this such that in the event of a primary Supervisor failure, the redundant Supervisor can assume the primary role (stateful failover) and continue to forward data using the last known set of routes (non stop forwarding).

The key to this feature is the ability of the redundant Supervisor to de-couple the control and data plane during this failover. Upon switchover, the data plane will continue to forward packets while the control plane alerts its neighbors of the switchover by issuing an NSF graceful restart to “NSF aware” neighbors. This restart avoids the peers tearing down neighbor relationships reducing route flaps and unnecessary network instability that might occur from a full reconvergence. The control plane will rebuild the FIB using route information from its peers from which the data plane can then start using to forward packets.

It should also be noted that many servers are dual-homed to two separate switches, which would provide an alternative resiliency mechanism to the dual supervisor strategy.

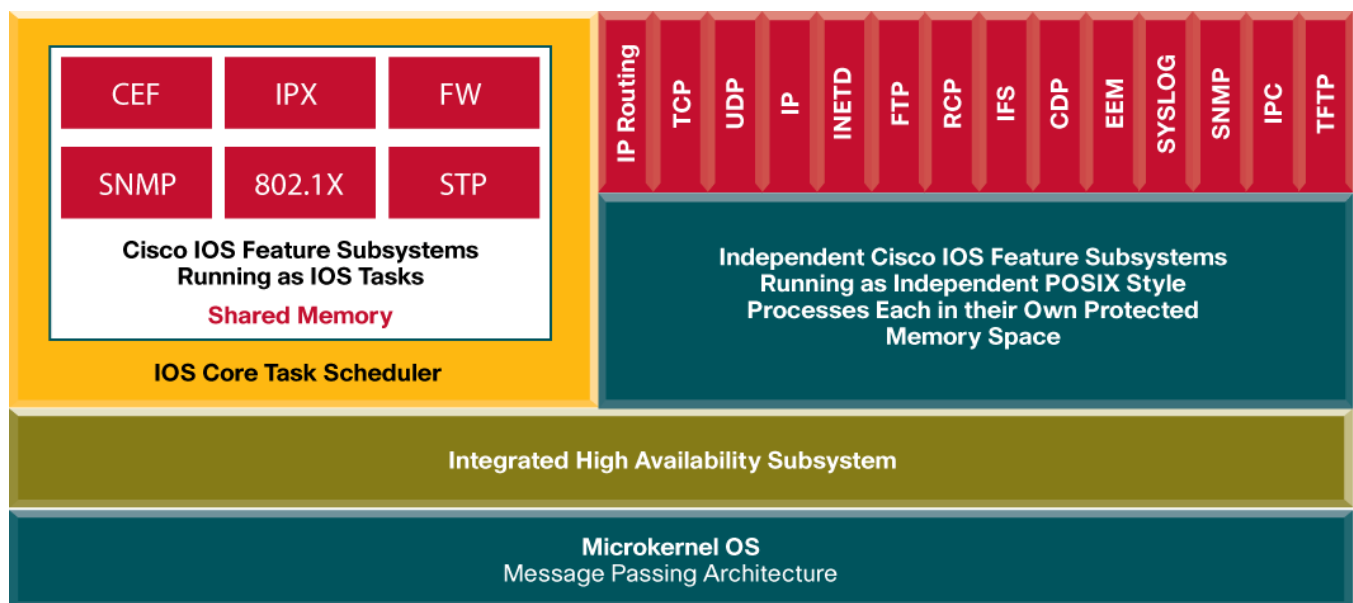
CR—IOS with Software Modularity

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

The introduction of IOS with Software Modularity in 12.2(18)SXF4 on the Catalyst 6500 adds more resiliency to IOS. IOS with Software Modularity adds subsystem ISSU (In Service Software Upgrades) allowing individual processes to be patched without the need to reboot the entire IOS and incur system downtime. Applying a patch to a running process will cause that individual process to be upgraded and restarted, allowing other processes to continue to operate with no operational impact.

Figure 5. IOS with Software Modularity Architecture



These processes can be manually restarted from the CLI or by the system (in the event of a processing failure), avoiding any outage that might normally have been incurred by other processes in prior versions of IOS.

In the first release of IOS with Software Modularity, more than 20 processes now run in their own protected memory space. As this software is developed in the future, more processes will be migrated into their own protected memory space, allowing IOS with Software Modularity to take Data Center switches to even better levels of availability.

CR—Multicast High Availability

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

With more multicast applications being deployed in Data Centers, so too is the reliance on HA features that enhance multicast availability. In a multicast deployment, the Rendezvous point (RP) plays a critical role in defining and building a path through the network over which multicast traffic can traverse. Loss of the RP can significantly impact the delivery of that multicast traffic affecting SLA's for those applications. To protect against the loss of the Multicast RP, IOS now supports Redundant Multicast RP's. A backup RP can now be configured to take over from a primary RP in the case of failure providing further resiliency for Multicast traffic.

Another important Multicast HA feature to consider is support for Multicast traffic during NSF/SSO switchover. In those switches using redundant Supervisors, now both Unicast and Multicast traffic flows can continue to be switched in the event of a Supervisor Failover.

CR—Route Dampening

Applicable Platforms: Catalyst 6500, Catalyst 4948G

Applicable Design Layers: DC Core, DC Aggregation, DC Access (if L3 to the Edge)

With the DC Core, DC Aggregation and possibly the DC Access Layers using Layer 3 protocols to maintain forwarding tables, so too the need to ensure that flapping routes do not impact network stability. Normally a link that moves between an UP and DOWN state will cause the Layer 3 routing protocol to re-converge resulting in potential packet drops. To protect against this scenario, the use of Route Dampening can be put into effect. Route Dampening effectively puts a flapping link into a temporary DOWN state until the link stabilizes. Once the link stabilizes, it is moved back to an online state. Moving the link into a temporary DOWN state eliminates network instability while the link flaps and removes the burden on the Layer 3 routing protocol from continually re-converging.

Currently, customers deploy the Catalyst Blade Switches in a layer 2 environment. Hence, this feature is not relevant in that architecture.

CR—Generic On-Line Diagnostics (GOLD)

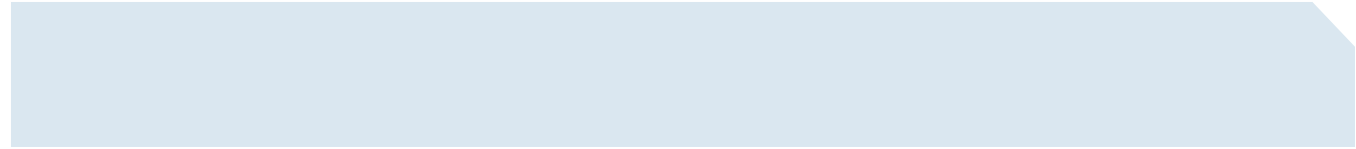
Applicable Platforms: Catalyst 6500, Catalyst 4500, Catalyst 4948G

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Diagnostics are performed on Catalyst switch hardware at boot up time. Many Data Centers run their switches for an indefinite period of time, so the inability to run boot-up diagnostics reduces the chance of detecting potential hardware issues that might otherwise arise. Generic Online Diagnostics (GOLD) is a feature that allows Diagnostics to be invoked at run time or scheduled at a pre-determined time while the switch is active. The use of GOLD can be used as a preemptive measure on the part of Network Administrators to ensure the hardware serving the needs of the Data Center is functioning correctly.

Architectural Scalability (AS)

As the enterprise business grows, so too must the Data Center adapt to meet the growing demands of the business. Implementing switching platforms that accommodate growth in connectivity and functionality is one key requirement of a Data Center Switch. However, scalability is much more than the ability to incrementally add more port density. Scalability needs to cover other areas of control and data plane scalability to truly meet the scalability demands of a Data Center switch.



The area of scalability should also address the migration of multiple services into the Data Center, which necessitates the tightening of security to secure hosts and applications from external (and internal) attacks. The use of Firewalls, Intrusion Prevention and Detection devices often spring to mind as the first line of defense. Hardening password policies, use of tokens to randomize passwords and implementation of AAA services also adds to the overall security strategy. The networking devices, however, also have a range of security features that should also be considered in parallel with those mentioned above.

The following section lists examples of some of the scalability features that should be considered for a Data Center switch.

AS—Hardware MAC Learning

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation, DC Access

To provide Layer 2 services, a Data Center switch will be required to learn MAC addresses of connected devices. The two methods of MAC learning are software-based learning and hardware-based learning. Of the two learning methods, hardware based MAC learning is preferred for a number of reasons. Firstly, performing MAC learning in hardware alleviates the switch control plane from that function, freeing up important Switch CPU processing capacity for other control plane bound tasks. Secondly, and perhaps more importantly, to secure the switch against a Layer 2 MAC attack (where multiple random MAC addresses are forwarded to the switch). This form of attack can have the intended effect of initiating a Denial of Service attack on the switch control plane, which if Hardware MAC learning is used, can effectively mitigate this attack.

Another important aspect is that as more blade servers (or blade switches) are added into the network, the CAM table is flushed for that VLAN leaving the switch to learn the MAC addresses for all devices in that VLAN. Hardware MAC learning avoids any potential impact on performance caused by an overloaded control plane that would otherwise have to relearn all MAC addresses in software. Consider another scenario where failover occurs to a redundant chassis and now the switch has to relearn all MAC addresses for dual homed hosts. Like the previous example, Hardware MAC learning offsets any performance hit that the control plane might otherwise incur through this learning process.

AS—Private VLAN

Applicable Platforms: Catalyst 6500, Catalyst 4948G

Applicable Design Layers: DC Access

In those situations where IP address space needs to be conserved, or where further subnetting of the IP address space needs to be avoided, Private VLAN's offers a way to minimize IP address space loss. The Private VLAN allows the creation of a VLAN within a VLAN. Within the primary VLAN are a series of sub VLAN's known as a community or isolated VLAN. All of the hosts in these "sub" VLANs share the address range from a common IP subnet. Communication between those "sub" VLAN's however is limited when compared to a normal VLAN. Hosts in the same community VLAN can communicate with one another but cannot communicate with hosts in other community or isolated VLAN's within the same Private VLAN without having an ACL defined. Hosts in an isolated VLAN cannot communicate with any other host in the Private VLAN (even hosts in the same isolated VLAN) without an ACL.

The use of Private VLAN's defers the need for subnetting of the IP address space allowing all hosts in the same Private VLAN to share the address space while securing communication between the "sub" VLAN's.

AS—Low Latency Switching

Applicable Platforms: Catalyst 6500, Catalyst 4948G, Catalyst Blade Switch

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Performance is a key driver behind the establishment of many Data Centers. Close analysis of the end-to-end transaction timeframes are performed to ensure any performance optimizations can be achieved. The Catalyst 6500, Catalyst 4948G and Catalyst Blade Switch provide low latency switching in the order of 4us to 6us for the Catalyst 4948, 10-12us for the Catalyst 6500 and 8-10us for the Catalyst Blade Switch . Minimizing switch latency is key to driving down overall transaction times.

AS—High Density 10 Gigabit Ethernet

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Figure 6. WS-X6708-10GE-3C



Driving higher levels of bandwidth, coupled with the emerging 10GE NIC technology for servers is requiring switching platforms to support more 10GE connections. The Catalyst 6500 platform recently introduced support for an 8-port 10GE module that increases 10GE port density to 64 x 10GE ports in a chassis. As with most high-density 10GE linecards provided by Switch vendors, this linecard is 2:1 oversubscribed. The key to handling oversubscription though, is to provide adequate congestion management tools to handle the oversubscription. In the case of this module, each 10GE port is equipped with 200Mb of per port buffering which is distributed between the receive queues and transmit queues. Coupled with the Weighted Random Early Discard (WRED) Congestion Management algorithm, the port is more than adequately ready to handle periods of oversubscription.

AS—Logical Port Density

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Virtualization is being embraced on many fronts, and adoption of server virtualization is one such example that is high on the agenda of many Data Centers. With tools like VMWare and Parallels, Data Centers can significantly reduce operational and environmental costs by reducing the physical count of servers while maintaining an adequate number of logical servers to meet the demands of their customers.

One of the scalability challenges that faces Data Centers implementing server virtualization is scaling support for logical ports on the switch. In the simplest sense, a logical port is a representation of a physical ports aggregate VLAN membership. So, for example, an access port is seen as a single logical port as it is part of a single VLAN. A trunk, however, is seen as multiple logical ports due to it supporting multiple VLAN's. To support a physical server running multiple logical server instances will require that switch port to support multiple logical ports for each physical switch port connection.

The number of logical ports that are supported in a switch is gated by factors such as the physical module hardware and the Spanning Tree algorithm used (i.e. MST, PVST+, RPVST+). Logical Port scalability is one area that the Catalyst 6500 excels in. In the DC Aggregation, up to 52,000 logical ports can be supported in a Catalyst 6500 switch, with up to 8700 logical ports per linecard. In the Access Layer, assuming 24 Access Layer switches with 7000 server connections, and 400 VLAN's with 10 virtual machines per physical server, the Catalyst 6500 could support up to 4640 logical ports per chassis with up to 800 logical ports per linecard.

AS—Hardware Unicast RPF (URPF)

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation, DC Access

The ability to protect the Switch from an IP Address spoofing attack can be addressed using Unicast Reverse Path Forwarding. The Catalyst 6500 provides support for this feature in hardware. Unicast Reverse Path Forwarding provides a means to inspect the forwarding tables when receiving a packet but using a reverse lookup. What this means is rather than inspecting the forwarding tables to determine where to send the packet, the forwarding table lookup looks at where the packet should have come from. For example, if a packet arrives on interface 3/1 and has a source IP address of 99.1.1.2, a lookup into the forwarding tables is performed. The tables are inspected to see which interface network 99.1.1.2 exists on. If the lookup returns that the network does indeed exist out that interface, the packet is forwarded. If, however, the lookup returns an interface that is not 3/1, then the system considers this a spoofed packet and will drop it.

Protecting against spoofing attacks allows a range of Denial of Service (DoS) attacks to be mitigated. More importantly, by performing this mitigating function in hardware eliminates the attack from causing a side effect, that of impacting the switch control plane, which in itself is a form of DoS attack.

AS—STP Root Guard

Applicable Platforms: Catalyst 6500, Catalyst 4948G

Applicable Design Layers: DC Aggregation

In cases where the Access Layer runs at Layer 2, Spanning Tree will invariably be used to protect the Data Center Switch network from loops. During the design of the Spanning Tree, a switch (typically at the Aggregation Layer) will be anointed with the task of serving as the Spanning Tree Root. This device is important in the Spanning Tree topology as it serves as the anchor for that STP instance originating STP BPDU's (Bridge Protocol Data Unit Traffic which serves to provide other switches with topology information) for that STP domain.

Should another device be inserted into the STP domain with a lower STP bridge priority or a lower ordered MAC Address (when having the same priority as the STP root), it could take over as the STP root changing the STP topology to a less efficient communication path. Even worse, it could be used by an attacker to redirect traffic via their device in order to snoop on user data.

STP Root Guard is a feature that protects the integrity of the STP root. Should a superior BPDU arrive at the STP Root (indicating that another device wants to take over as STP Root), then the Root Switch will place the link from which the BPDU arrived into a "Root Inconsistent" state effectively disabling that link. This action secures the Root Switch from being compromised.

Since the Catalyst Blade Switch is deployed at the Access Layer, this feature is not relevant to it.

AS—Security Services Modules

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation

The first line of defense in most Data Centers, as noted above, is to implement a range of devices such as Firewalls, Intrusion Detection and Intrusion Prevention. While these services can be addressed by implementing standalone appliances, the Catalyst 6500 offers the added advantage of integrated these services into the switch chassis. One of the primary benefits of this approach is the integrated services module typically provides a much higher level of performance than its equivalent standalone appliance equivalent. As the data load on the Data Center grows, so too the requirement for higher performance loads on these devices, which can be met by these modules.

The Firewall Module (FWSM) and Application Control Engine (ACE) module also offer Virtualization support. The use of these modules in a Data Center allows different user groups to virtualize an instance of that service module which, in turn, allows the creation of an operational configuration tuned to the specific requirements of that group. Environmental benefits come into play with virtualization avoiding the requirement to otherwise deploy multiple instances of the standalone appliance to meet the needs of multiple groups.

Cisco's Data Center architecture recommends services be applied at the aggregation layer. Hence, these modules are not relevant to the Access layer switches.

AS—Multicast Security

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation

With increasing deployments of Multicast based applications, so too the need to secure the network against unauthorized Multicast sources from distributing content. Recent enhancements in IOS now allow the Multicast Rendezvous Point (RP) to inspect PIM Register commands from Multicast sources in order to determine their authenticity. The RP can use an ACL check to verify the validity of the source and accept or reject that multicast source from establishing itself as a Multicast source. This mechanism further bolsters an aspect of security not often considered by many Network Administrators.

AS—Control Plane Policing

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Aggregation, DC Access

The Switch Control Plane drives many of the functional aspects of a switch, and as such is a critical resource that must be protected. Compromising the Switch Control Plane can have dire consequences for maintaining the on-going operational functionality of the switch. Control Plane Policing has been introduced on the Catalyst 6500 platform to protect it from inadvertent bursts of control plane traffic and also from more maliciously intended forms of attacks. A new interface called the "Control Plane interface" can have a rate limiting policy applied to it to limit the total amount of traffic destined to the Control Plane. By limiting the amount of traffic that can be forwarded to the Control Plane ensures that the operational viability of the switch is maintained.

Operational Manageability (OM)

When the Data Center is operational, the role of the administrator is to manage the devices to maximize uptime of available services. Understanding what is going on in the network is paramount to supporting capacity planning efforts as well as troubleshooting issue that might arise. The following section details some of the features that can be used to help better manage the Data Center network.

OM—Encapsulated Remote SPAN


Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Switched Port Analyzer (SPAN) is a hardware-enabled feature in the switch that allows traffic on a target port (or VLAN) to be mirrored to a SPAN port. SPAN currently limits the source port and the destination port to be in the same chassis. Encapsulated Remote SPAN (ERSPAN) is an enhancement that allows the source and destination ports to be in different chassis that are separated by a Layer 3 boundary. ERSPAN works by using the Supervisor hardware to encapsulate the SPAN traffic within a GRE tunnel. While Remote SPAN (RSPAN) offers a similar capability, it requires the source and destination chassis to be Layer 2 adjacent.

ERSPAN requires the use of a Supervisor 720 to be in the chassis where both the source and destination ports are located. Devices in between the source and destination chassis can be any device capable of forwarding IP packets.

Figure 7. Encapsulated Remote SPAN



ERSPAN offers Data Centers the benefit of being able to consolidate monitoring devices on one device and replicating traffic from a source port to that central device.

OM—Time Domain Reflectometry (TDR)

Applicable Platforms: Catalyst 6500, Catalyst 4948G

Applicable Design Layers: DC Access

TDR is a method that can be used to assist in the location and identification of copper cabling faults. It can be used for the detection of broken or shortened conductors, loose connectors, crimped or cut cables, sheath faults and much more. This feature is embedded in the hardware of many of the latest generations of Ethernet linecards. It offers administrators managing Data Center resources a fast and effective way to assist in the diagnosis of copper cabling problems that might arise.

OM—Power

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

The growth in Data Center use translates into additional network resources, which require network connectivity. The drive towards higher speed connections and greater densities of connected devices requires the network to accommodate greater power loads over time. A move to a modular based switch brings with it a range of power supply options that can be used to meet these growth demands.

While both AC and DC power supply options increase the choice for power deployments, so too the arrival of multiple input power supplies also add flexibility for network administrators and facilities managers. The latest round of Power Supplies now available for the Catalyst 6500 offer multiple inputs. The 6000W power supply, for example, supports two inputs. Using one input provides 3000W of power (@220V) and can be incrementally grown to support 6000W of power by simply activating the second input. This allows the Data Center to provision only the required number of power inputs to meet current needs, while with little effort, supporting incremental growth by simply activating the second input.

OM—Configuration and Capacity Planning Management

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

As services in the Data Center are added, so too will the configuration in the Networking Switches adapt to meet those needs. Changes to the configuration over a period of time can lead to configuration inaccuracies gradually creeping in. The Catalyst 6500 now adds support for a configuration check (command is “show diagnostic sanity”) that scans the configuration for errors reporting back any inconsistencies it finds. It will review SNMP, HA, Port Channeling, Trunk interfaces, UDLD, STP, IGMP, Duplex configurations and many more configuration elements. A report is presented back to the user highlighting any errors found.

Understanding the use of hardware resources in a switch has, up until now, also been a daunting task. This information is pertinent to the ongoing task of network capacity planning. Much of the capacity information exists, but is buried in amongst a multitude of show commands that are often tedious to correlate. A new command enhancement (command is “show platform hardware capacity”) is now available that lets an Administrator check the resource consumption of a variety of hardware resources. Resources such as QoS Policers, Switch Fabric load, Power consumption, VLAN's, PFC TCAM resources and much more can be viewed with this show command.

OM—Layer 2 Traceroute

Applicable Platforms: Catalyst 6500, Catalyst 4500, Catalyst 4948G, Catalyst Blade Switch

Applicable Design Layers: DC Aggregation, DC Access

The availability of a Layer 2 Traceroute command complements the existing layer 3 Traceroute tool that most Network Administrators are familiar with. In Data Centers that deploy Layer 2 at the edge, the use of Layer 2 Traceroute can be used to assist in the troubleshooting of Layer 2 connectivity issues. It works in a similar fashion to Layer 3 Traceroute providing hop-by-hop information of the nodes from a source to destination device through a Layer 2 connection.

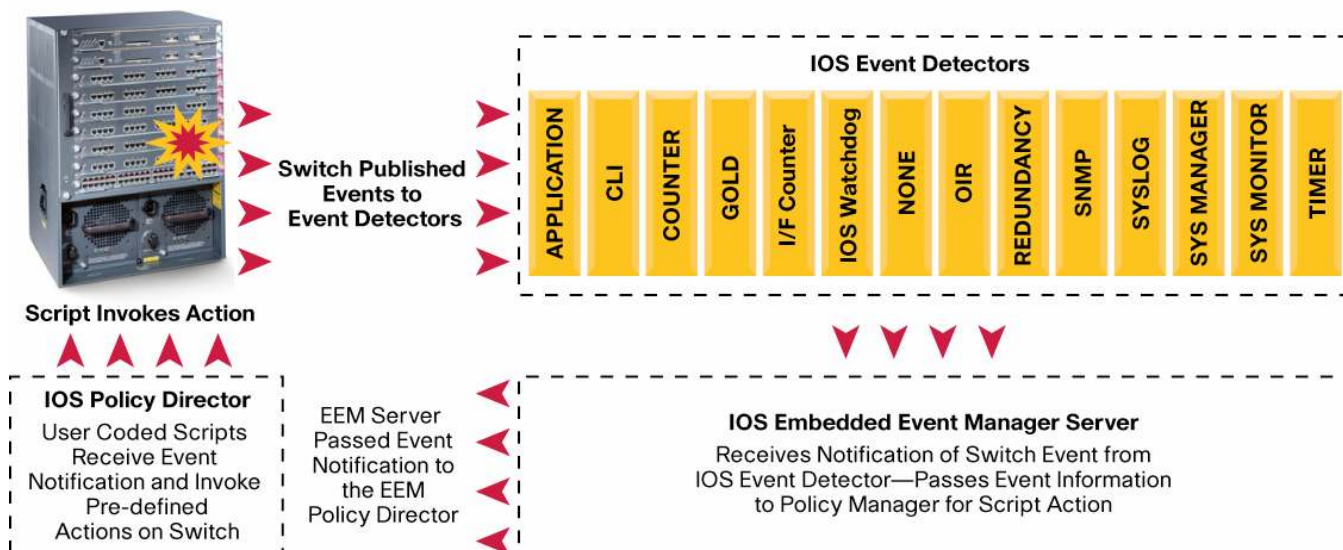
OM—Embedded Event Manager (EEM)

Applicable Platforms: Catalyst 6500

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Embedded Event Manager is a programmable subsystem now available on the Catalyst 6500 that allows a user script to invoke actions on the switch when a specific event occurs. The flexibility that EEM offers the Data Center is significant. With over 14 event detectors to detect events on the switch, the Administrator can implement scripts to manage a multitude of situations on a 24/7 basis. The event detectors can monitor SYSLOG messages, CLI commands, interface counters, insertion and removal of hardware, failovers, IOS processes, and much more. The scripts can invoke actions including generating emails and pagers, issuing CLI commands, restarting the switch, and generating custom SYSLOG and SNMP traps.

Figure 8. Embedded Event Manager



OM—Hardware Based Netflow

Applicable Platforms: Catalyst 6500, Catalyst 4948G

Applicable Design Layers: DC Core, DC Aggregation, DC Access

Better understanding traffic flows through the Data Center network for capacity planning or accounting purposes is a high priority for most Network Administrators. All modular Catalyst switch platforms offered by Cisco provide hardware-based collection of Netflow records. Recent extensions to Netflow support have been incorporated into the Catalyst 6500, allowing it to create Netflow v9 records. The implementation of this new record type allows the Catalyst 6500 to collect information on Multicast flows that transit the switch. The ability to complement Unicast flow record collection with Multicast flow record collection provides the Data Center with a better overall understanding of network traffic patterns. More importantly, it allows administrators to adjust key network parameters to meet service level agreements for key business users of the Data Center.

SUMMARY

The convergence of multiple services into the Data Center offers many operational benefits to the Enterprise business. By converging computational, storage, application and networking services, into the one operational center, however, places stringent demands on the Data Center design and the capabilities of the devices services these needs. Furthermore, performance, security and availability of application and data services are all key metrics that must also be met to provide a successful Data Center service.

The Modular Cisco Catalyst Switch family provides a solid foundation upon which to run today's Data Center. Its widespread deployment in customer networks today validates its ability to serve those needs. Moreover, its rich feature set provide the greatest flexibility for designing the Data Center to meet the specific needs of the Enterprise business it is serving.

As new server form-factors get deployed in the Data Center, Cisco has developed newer access layer switching solutions, like top-of-rack Catalyst 4948G and Catalyst Blade Switch, to meet the emerging networking needs. These switches offer compelling value-proposition based on the resiliency, scalability and manageability needs of the servers/applications they connect to the Data Center network.

OTHER REFERENCES

Cisco Data Center Infrastructure SRND 2.1

http://www.cisco.com/application/pdf/en/us/guest/netso/ns107/c649/ccmigration_09186a008073377d.pdf

Infrastructure Architecture SRND

http://www.cisco.com/application/pdf/en/us/guest/netso/ns304/c649/cdcont_0900aecd800e4d2e.pdf

Server Farm Security in the Business Ready Data Center

<http://www.cisco.com/warp/public/732/systems/docs/dcsrndbk.pdf>

Load Balancing Data Center Services SRND

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns656/net_design_guidance0900aecd800eb95a.pdf

Internet Edge Design Architecture SRND

http://www.cisco.com/application/pdf/en/us/guest/netso/ns304/c649/ccmigration_09186a008014ee4e.pdf



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

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.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuickStudy, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)