

# Web and Application Hosting Services

## A Cisco End-to-End Design Guide

### Overview

The world of business today offers many exciting opportunities for companies to compete in innovative new ways. As the fundamental paradigm of commerce shifts to take advantage of opportunities offered by the Internet, companies that do not follow suit could be left behind. Never before has there been a medium like the Internet providing such widespread access to target markets at relatively minimal cost.

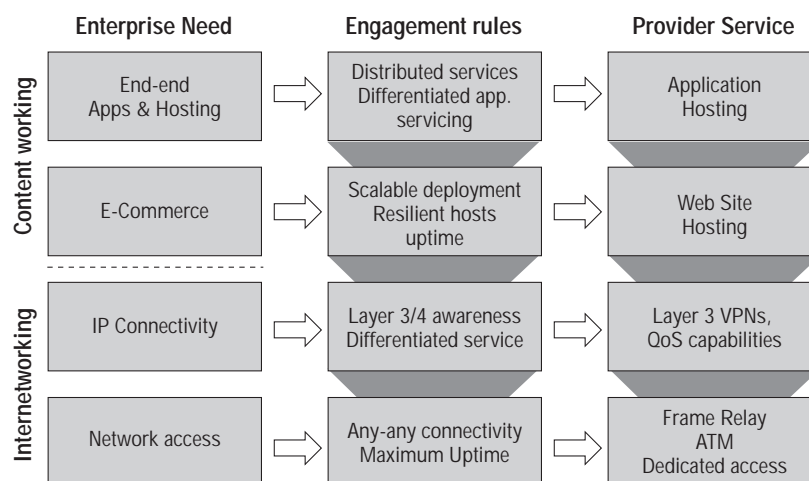
While today's companies restructure to embrace an e-commerce business model, many are forced to focus on their core competencies as a company because of increasing competitive pressure. Such focus, as in the past, has opened many new doors for service providers to offer outsourced intelligent services, thereby off-loading the burden from the enterprise customer.

With the continual growth of enterprise outsourcing opportunities, the intelligence of the network needed to provide these services has grown as well. Figure 1 shows the major milestones of service provider offerings relative to the ongoing requirements of the enterprise customer. Currently

one of the major opportunities offering high-revenue potential to the service provider is application and Web hosting services. In order for enterprise customers to engage in an e-commerce model, they must not only have a presence on the Internet, but they must also possess a highly scalable and resilient solution. Prospective customers of the e-commerce model will enter via a newer "Web-based" door. This door must remain open at all times and be able to handle expected traffic demands.

Today, Cisco and its partners are able to offer the service provider all integrated elements required to build a solid Web hosting service. Building on its reputation to provide highly scalable and resilient Layer 1 to 3 infrastructure, Cisco continues to add services of higher intelligence to its proven design model. Adding Layer 4 to 7 services in the same scalable and resilient manner to the infrastructure allows Cisco to provide a superior one-vendor, end-to-end solution. Such offerings from Cisco and its partners allow the service provider to minimize costs associated with growth, management, and maintenance related to building a Web hosting environment.

Figure 1 Evolving Service Provider Offerings



In order to build a successful Web hosting service, three key characteristics of the service must be considered, namely scalability, high availability, and manageability. Scalability must be evaluated on more than just the ability to increase device performance and link bandwidth—it must also be considered on a geographic basis. The ability to create multiple service data centers and dynamically migrate, host, and synchronize distinct Web content across such centers in a manageable fashion becomes a key requirement. Not only does a geographic distribution of the service make sense from a scalability aspect, but it also lends well to providing high availability and optimal resource usage. The ability to replicate content across geographically dispersed data centers allows for access to multiple copies of identical content to maximize availability. In addition, users within various global regions can access content which is situated in closer proximity making the browsing experience more pleasurable.

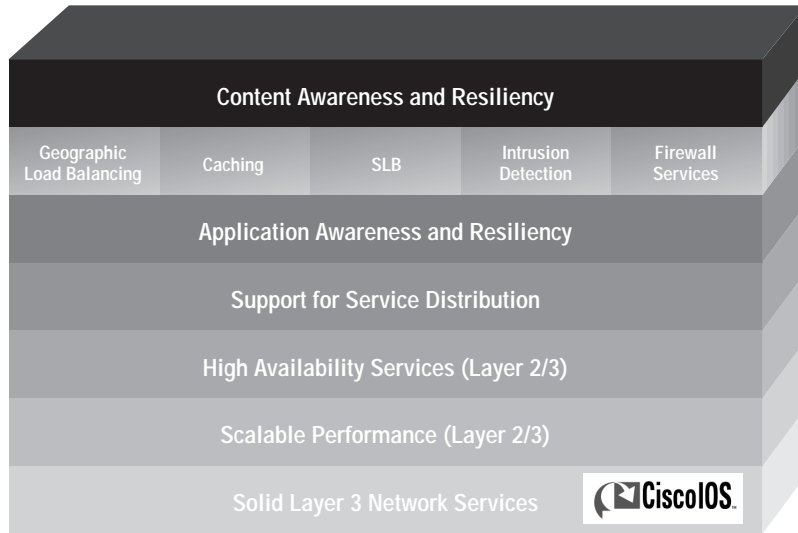
High availability is present in many layers of the Web hosting service all with the intent of ensuring maximum accessibility to Web content. In its most basic form, high availability within a data center ensures Layer 2 protocols such as Spanning Tree Protocol (STP) are kept in check and do not create lengthy outages during a link failure. On a more global basis, high availability amongst Layer 3 protocols ensures fast rerouting is possible during core or data center IP transport failure. However, in a Web hosting environment, high availability is also critical at higher layers of the service. The purpose of a Web hosting environment is to service content requests from a large numbers of Web clients in a resilient manner. Included in this service description is not to just simply service the request, but service the request in a way that will offer the most pleasurable experience for the Web client. Measures to ensure such serviceability include returning content to the client that is accurate, recent, and in closest proximity to the original request. Pitfalls of a Web hosting service can include uneven loading of content requests amongst eligible servers, requests forwarded to inactive servers, and the return of expired content from a server or a Web cache facility. In order to bypass all potential Web hosting pitfalls, a highly sophisticated suite of protocols and services are required within the Web hosting environment that are deployed in a tightly integrated fashion.

Manageability is always an important aspect of any large-scale network service. Probably most important is the ability to retrieve valuable information from the network for

planning and billing purposes. The ability to collect statistics including Netflow, Remote Monitoring (RMON), and Simple Network Management Protocol (SNMP) data at key points within a Cisco infrastructure allow the service provider many options for detailed service billing and analysis. Equally important is the ability to configure devices in a scalable manner. Cisco IOS® software technology presents a well known user interface which has been replicated across all relative devices within a Cisco Web hosting solution. A growing concern for service providers within a data-center environment is the ongoing task of managing IP address deployments. Many service providers are faced with a complex task of ensuring efficient allocation of address space to Web hosting clients. The task of managing multiple variable-sized subnets is a large task in itself. With its ability to create a “private virtual LAN (VLAN),” has a solution for IP address management by shifting the paradigm of IP address deployment within a multiclient Web hosting environment.

This paper presents the Cisco end-to-end solution for a solid Web hosting service. Many facets of Cisco technology are integrated to provide the overall solution (see Figure 2). In addition, individual technologies offered by Cisco are explored in the context of their integration into the overall Web hosting service design. A typical Web hosting service consists of numerous data centers interconnected through a resilient core IP service. This paper focuses on the higher-layer services required to build a highly successful Web hosting service. Please refer to the Cisco Web site at <http://www.cisco.com> for more information on designing large scale IP networks using Cisco technology.

Figure 2 Elements of a Cisco Web Hosting Service

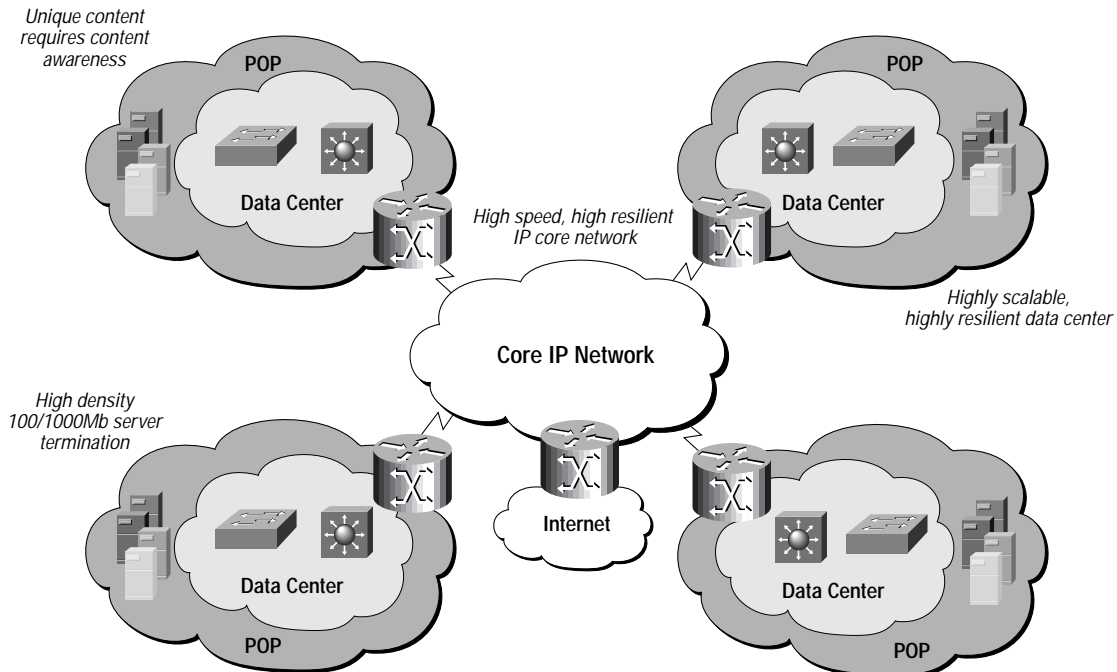


### The Web Hosting Service Design

As shown in Figure 3, a typical Web hosting service design consists of several geographically distributed data centers. Web hosting data centers are interconnected using a high-speed, highly resilient IP core service. For scalability, resiliency, and usability reasons, Web content may be replicated and dispersed throughout the entire service bounds. Several advanced networking services are required to ensure the appropriate content is always accessible and

accurate. Although many point-product vendors offer partial solutions to this service requirement, Cisco and its partners have developed an integrated suite of advanced network services, thus allowing such dynamic behavior. Using an end-to-end Cisco Web hosting solution, several advanced protocols and platforms work in conjunction to provide the desired service characteristic commonly referred to as “content-awareness” or “content-working.”

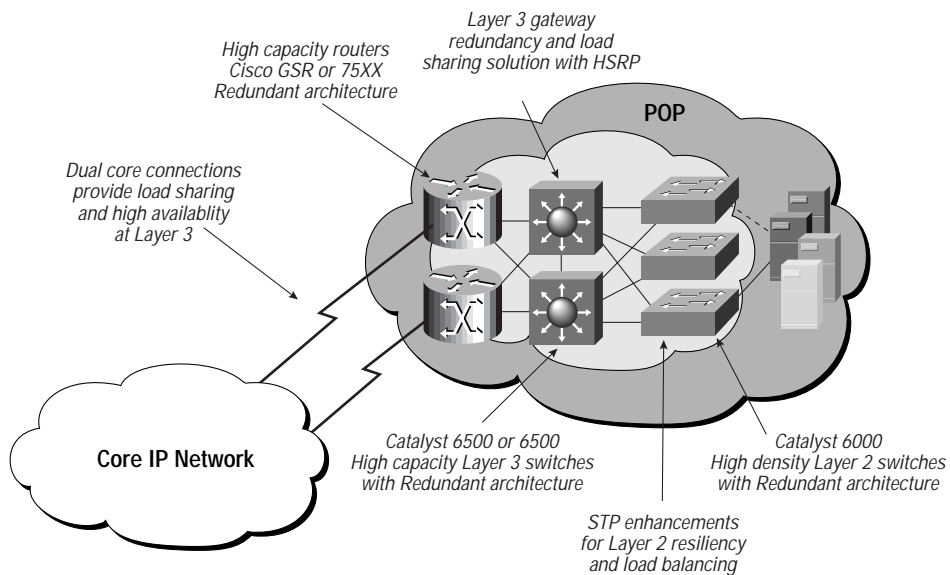
Figure 3 Components of a Distributed Web Hosting Service



A closer look at a data center shows that many advanced services are required to create a content-aware service. In addition to content awareness, the design must provide for high availability at all layers. In order to build a solid Web hosting service, one must start with a solid Layer 1 to 3 infrastructure that displays the same desired scalability and high availability of the overall service. Figure 4 highlights a data-center conceptual design in which resiliency is optimally used in a modular fashion. Redundancy at Layer 2 and Layer 3 is provided to ensure minimal downtime is experienced in

the occurrence of a link or device failure. An advantage to the Cisco redundant design is the ability to load share traffic amongst the redundant components, thereby using all procured resources. Load sharing is provided through the Cisco Hot Standby Router Protocol (HSRP) at Layer 3 and the Cisco per-VLAN spanning tree (PVST) at Layer 2.

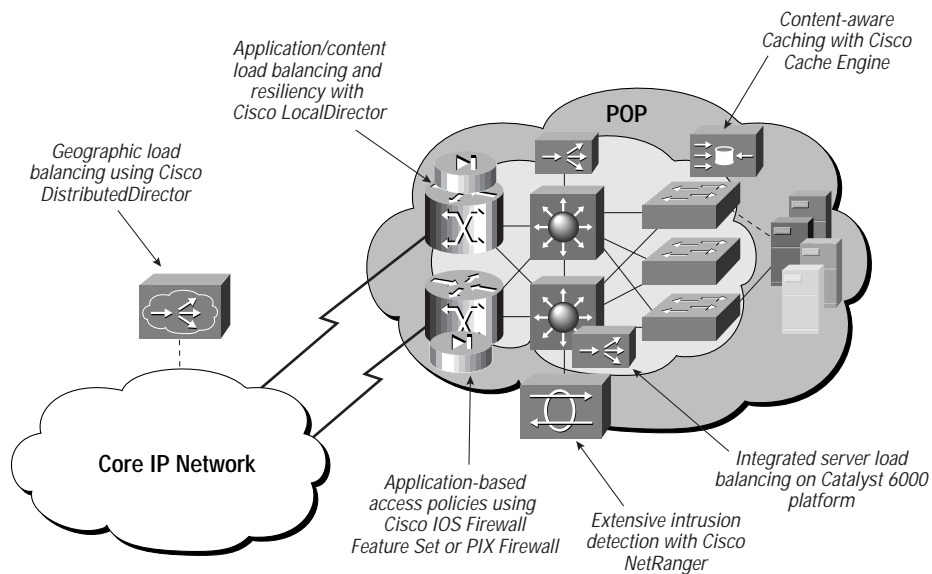
Figure 4 Resiliency Within a Web Hosting Center



As previously mentioned, the foundation to a solid Web hosting environment is a Layer 1 to 3 infrastructure that provides scalability and high availability. Building upon this infrastructure, we can now start to incorporate the advanced protocols and devices required to build the desired Web hosting service. Figure 5 shows the required advanced services as an overlay to the previous Layer 1-3 infrastructure. Many of the Cisco Internet appliance products and technologies are integrated to provide the desired solution. With the evolution of the Cisco Layer 3 switch products such as the Catalyst® 6000 series, many of the appliance technologies are being incorporated within the product to provide a one-box solution.

Just as there exists high availability and scalability within a Cisco Layer 1 to 3 infrastructure, the same scalability and high availability exist for the advanced Cisco services. Each Internet appliance technology possesses the ability to be configured in a highly fault-tolerant arrangement while also providing for scalable performance. In addition, several feedback mechanisms exist between components of the design to ensure utmost availability and optimal resource utilization. A more in depth look is given to the individual Internet appliance technologies and their scalability and availability characteristics later in this white paper.

Figure 5 Overlay of Advanced Services within Data Center



## Advanced Design Components

As seen in Figure 5, several products and technologies exist within a Cisco advanced Web hosting data center to provide utmost scalability and availability. Many products and technologies work together through feedback protocols to ensure optimal resource utilization. This section of the white paper looks at the various products and technologies that integrate into the Cisco Web hosting solution. Focusing on content scalability and high availability, the Cisco solution helps to scale Web content and applications from a single server to multiple servers, and finally to multiple sites.

### How Internet Applications Work

Internet applications can be loosely defined as the applications people utilize over a large-scale IP network. The large scale IP network could comprise the global Internet or perhaps a company intranet. Common applications and their TCP/User Datagram Protocol (UDP) port numbers utilized over IP networks include:

Web requests (Hypertext Transfer Protocol/Hypertext Transfer Protocol Secure [HTTP / HTTPS])—TCP port 80/443

- File Transfer Protocol (FTP)—TCP port 21
- News (NNTP)—TCP port 119
- Streaming audio/video—UDP port range
- Instant messaging (IRC)—UDP port 194
- E-mail retrieval (point of presence [POP])—TCP port 109/110
- E-mail Transmission (Simple Mail Transfer Protocol [SMTP])—TCP port 25

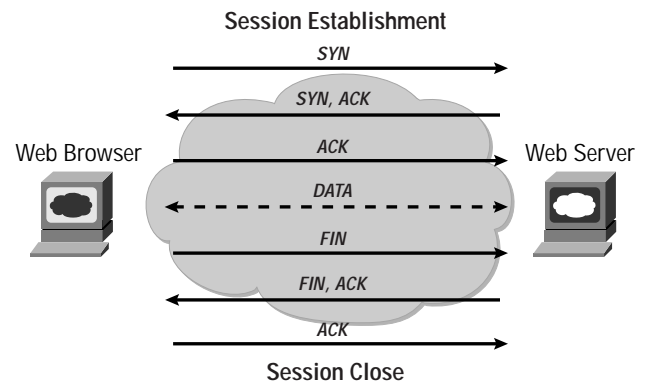
Of most importance to a Web hosting service are TCP ports 80 and 443 associated with nonsecure and secure Web content distribution, respectively. As any flow represented by these port numbers originates from a Web client, special attention must be given to these requests.

Users of Internet applications typically do not remember specific IP addresses or the port number for a particular application. More commonly the user simply remembers and easily recognizes its URL. From this point, the easily remembered URL is translated into a related IP address and port number yielding the location of the desired content or application. Hence, one of the keys to providing truly distributed Internet content and applications is to utilize flexible name-to-address binding services (Domain Name System [DNS]). A flexible DNS service ensures that clients are

connected to servers that provide the desired content or application with the best possible service. As an example, the best service might be the server closest to the user or perhaps the one that is currently least loaded. This requirement is key, ensuring that Web client enjoy their browsing experience and will likely become repeat customers.

When the IP address and well-known TCP port number are known for the server hosting a client's request, a TCP connection is made from the client to the server. Being connection-oriented, TCP follows a sequence of handshakes to establish a session between client and server. Figure 6 shows the stages of session creation and session termination between a client and a server. As seen in the figure, a series of three handshakes must take place before any Web content flows from server to client.

Figure 6 Stages of a TCP Connection



The key to providing a highly scalable and resilient Web hosting service is to ensure this connection is reliably formed for every client request. In most cases, however, the client does not require knowledge of the location of the information source. What is mandatory is for the client to receive the desired content from a source that is reliable and provides adequate service performance. As discussed earlier, it is desirable for many reasons to replicate and distribute identical content on a geographic basis. We can, therefore, steer clients to the most appropriate source of content by steering the establishment of the TCP connection.

To summarize, scaling a Web hosting service requires solutions focused on the following two areas:

- Replication and distribution of Web content—making the content globally available
- Content request to server binding—steering connection establishment to appropriate server

These two problems are solved using advanced network services that:

- Dynamically connect the client to the most appropriate content location based on requested URL
- Dynamically connect the client to the optimal server at determined content location
- Distributing Content from one to many servers


The scaling of the Web content from a single server to multiple servers can be a daunting task. Many of the associated challenges include:

- Content synchronization—ensures that all servers possess most recent content

- Content accessibility—ensures that content is accessible regardless of individual server status
- Request load distribution—ensures that content requests are distributed evenly within a data center

The Web hosting server farm is a collection of Web servers which are used together to scale the performance and capacity of a Web site. Web farms may be centrally located or globally distributed offering access to the same content from various locations. Figure 7 shows a typical configuration of a Web server in a Web hosting environment today is that it will have an interface facing the Internet (front) and an interface facing an internal network (back) that is used for functions such as backup, content replication, server management, and database access.

Figure 7 Server Configuration within Web Hosting Environment



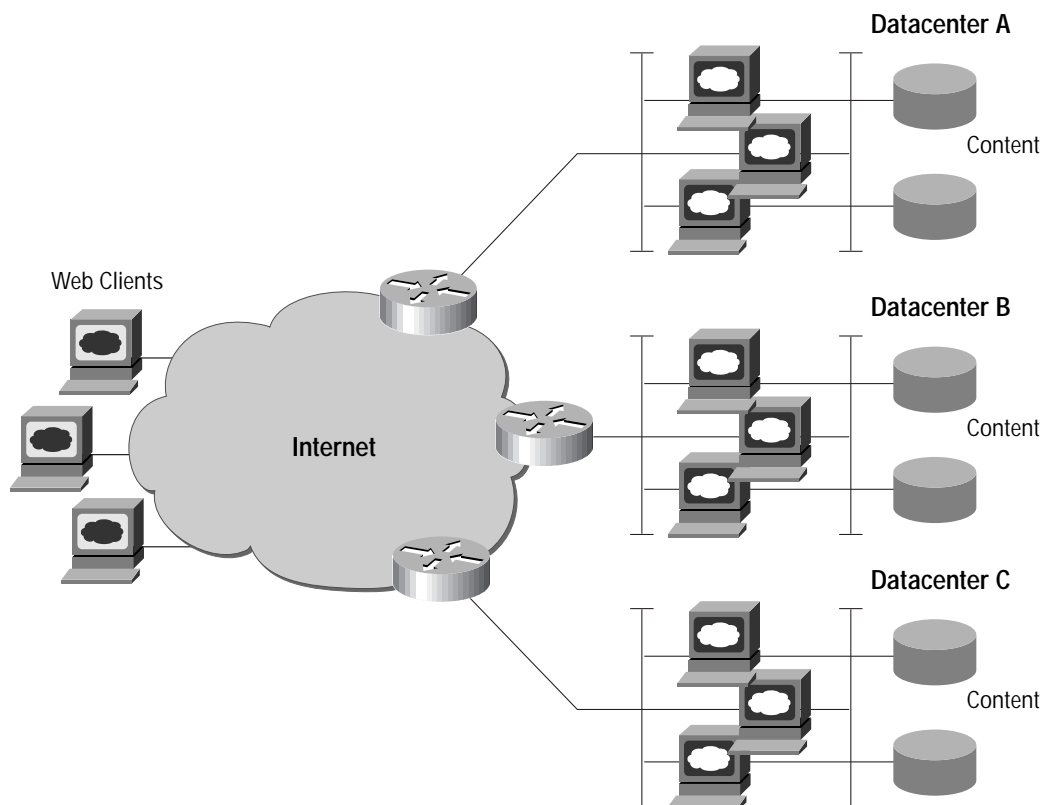
The challenge is how to allow a collective group of Web servers to handle many more page requests, while maintaining the appearance of a single entity to the outside world. The area of technology used to describe this feature is server load balancing (SLB). There are related techniques from the host operating system world that are generally termed clustering. In clustering, many systems participate in the delivery of an application while appearing as one system through the use of a single IP address. This scenario scales Web performance within a single cluster to a certain point, but any desire to geographically disperse parts of the content delivery mechanism are best handled using a form of server load balancing. A detailed look at the Cisco solution for server load balancing is presented later in the paper.

#### Distributing Content from One to Many Locations

Once Web content can be reliably scaled within a Web hosting data center, it then becomes an issue of how to deploy a regional/global solution. Regardless of where clients are located, they must always gain reliable access to desired Web content. The challenge is to scale a single physical site to multiple sites, each being geographically dispersed and all delivering the same content, as shown in Figure 8.

The solution to this problem is to have devices within the network capable of dynamically binding Web server names to appropriate IP addresses. The decision of which address to bind is based upon the location of the requesting client and response times of the content servers. These devices require knowledge of IP routing tables within the network to accurately determine the proximity of the client to the closest server based on routing metrics. In many cases, response times of content servers are also a deterministic factor in the decision of where to forward the client request. If a particular client is in Europe and the European servers are heavily congested, it may be better to resolve the client's request by forwarding it to another geographic region with more available resources. Although this solution may seem odd at first, giving consideration to time zones and their associated regular busy periods, under-utilized servers in different time zones may be able to handle excess load more efficiently. A detailed look at the Cisco solution for geographic load balancing is presented later in the paper.

Figure 8 Geographic Distribution of Web Content



### Server Load Balancing—The Cisco Approach


The crux of the Cisco server load-balancing techniques is to represent a series of content servers arranged in a redundant configuration as one common server to the client community. A common feature of SLB techniques is the mapping of a virtual IP address and port numbers to the real or virtual IP addresses of the redundant server arrangement. To provide load balancing, successive client connection requests are distributed across all available servers that contains required content. SLB can be achieved using two fundamental techniques:

- Dispatch mode
- Directed mode

Within each of these techniques, a variety of algorithms and associated metrics are used to distribute connections between the redundant content servers. These load-balancing algorithms include:

- Least utilized
- Least active connections
- Lowest response time

Figure 9 Server Load Balancing—Directed Mode



#### SLB—Directed Mode

In directed mode, the virtual server can be assigned an IP address that is not known to any of the real servers. The SLB service translates the address in packets exchanged between clients and real servers, translating the virtual server IP address to a real server address via Network Address Translation (NAT). In the example in Figure 9, the virtual IP address of three servers is 10.10.10.1 while the real addresses of the servers are 10.10.20.1, 2 and 3. The server load balancer in this situation performs NAT on inbound connections from clients and changes the address and Media Access Control (MAC) addresses on the frame to the real MAC and IP address of a physical server. When the server sends traffic back to the client, the SLB mechanism will perform invert the original NAT operation. Directed mode can work in four variations:

- *Directed mode—normal*

Within normal directed mode, the servers do not need to be in the same subnet as the load balancer. Packets destined to or from the servers must traverse through the load balancer as per topology restrictions. The source and destination MAC addresses along with destination IP addresses are rewritten. For traffic destined to the servers, the destination IP address is rewritten to the servers real IP address of the server. Traffic originating from the servers will have its source IP address rewritten to that of the virtual address. Load-balancing decisions are made based on the first packet.

- *Directed mode—source NAT*

Within source-NAT directed mode, the servers again do not need to be in the same subnet as the load balancer. Packets originating from the servers, however, are directed back to the load balancer. This scenario is accomplished by performing a NAT of the source IP address of the original

client request to that of the load balancer. The load balancer, therefore, by default sends its return traffic back to the load balancer, thereby eliminating the topology restrictions of normal mode. Traffic returning from the load balancer to the clients experiences a NAT of the source and destination MAC addresses and source IP address of the load balancer, so the client community sees one virtual server. Load-balancing decisions are made based on the first packet.

- *Directed mode—TCP proxy*

Within TCP-proxy directed mode, the servers again do not need to be in the same subnet as the load balancer. Packets originating from the servers are directed back to the load balancer by topology restrictions. Connections from clients are terminated at the load balancer and new connections are established or existing connections may be reused from the load balancer to real servers. Load-balancing decisions are made by inspecting the content request (URL-based load balancing). On return traffic from the balancer to the clients, the source and destination MAC addresses are rewritten, and the TCP sequence numbers in both directions are adjusted after the connection has been unproxied. Also, a rewrite of the source IP address is performed to that of the virtual server IP address.

- *Directed mode—TCP proxy—source NAT*

TCP-proxy source-NAT directed mode operates the same as TCP-proxy directed mode with the exception that the client's IP address is translated to the address of the load balancer for requests originating from clients. This scenario allows for the topology restriction to be lifted as in the case of Source-NAT directed mode.

## Dispatch Mode

In dispatch mode, the virtual server address is known to the real servers and the SLB process simply redirects packets to the real servers at the MAC layer. In this mode, the real servers must be Layer 2 adjacent to the load balancer, and they need to have knowledge of the virtual IP address. The example in Figure 10 shows three servers, each with two IP addresses configured on it. One address is its normal address and the other is the virtual IP address for the server cluster. This feature is referred to by different names under different operating systems. In Windows NT, it is referred as secondary addresses, and in UNIX and Linux, it is called IP aliasing. When SLB is running in dispatch mode, it will rewrite only the MAC-layer headers of the frame because the virtual IP address is configured upon each server. This setup does result in some loss of flexibility because destination port mapping for service scalability cannot be accomplished.

Directed mode can work in four variations:

- *Dispatch mode—normal*

Within normal dispatch mode, the servers must be in the same subnet as the load balancer. Packets originating from the servers are directed back to the load balancer as per topology restrictions. Only the source and destination MAC addresses are rewritten. Load-balancing decisions are based on the first packet.

- *Dispatch mode—source NAT*

Within source-NAT dispatch mode, the servers again must be in the same subnet as the load balancer. Packets originating from the servers, however, are directed back to the load balancer. This scenario is accomplished by performing a NAT of the source IP address of the original client request to that of the load balancer. The load balancer, therefore, by default sends its return traffic back to the load balancer, thereby eliminating the topology restrictions of normal mode. The source and destination MAC addresses and source IP addresses are rewritten for client traffic directed to the servers. The destination IP address is rewritten by the load balancer for traffic destined back to the clients. Load-balancing decisions are based on the first packet.


- *Dispatch mode—TCP proxy—normal*

Within TCP-proxy normal dispatch mode, the servers again have to be in the same subnet as the load balancer. Packets originating from the servers are directed back to the load balancer as per topology restrictions. Connections from clients are terminated at the load balancer and new connections are established (or existing connections may be reused) from the load balancer to the real servers. Load-balancing decisions are made by inspecting the packet content (URL load balancing). On return traffic from the balancer to the clients, the source and destination MAC addresses are rewritten, and the TCP sequence numbers in both directions are adjusted after the connection has been unproxied.

- *Dispatch mode—TCP proxy—source NAT*

TCP-proxy source-NAT dispatch mode operates the same as TCP-proxy dispatch mode with the exception that the client's IP address is translated to the address of the load balancer for requests originating from clients. This setup allows for the topology restriction to be lifted, as in the case of source-NAT dispatch mode.

Figure 10 Cisco LocalDirector—Redundant Configuration



### Server Load Balancing from Cisco

Cisco provides several approaches to SLB services, depending on specific performance and resiliency requirements. An initial look at the various SLB products from Cisco is presented below. Following the product overview, a scalable and resilient Web hosting service design is presented, integrating advanced Cisco Web hosting technologies.

#### LocalDirector


Cisco LocalDirector is a network appliance that provides server load balancing by portraying a bridge in the path of packets destined to servers. The servers can be either directly adjacent to the LocalDirector or connected across routers. LocalDirector supports load balancing in many ways using different algorithms to determine effective connection load distribution. The selection criteria it can use to balance server connections include:

- *Least connections*—Servers with least number of open connections are favored.

- *Weighted*—Servers with higher assigned weights will receive a proportional numbers of connections.
- *Fastest*—Servers with the fastest response time receive a higher percentage of connections.
- *Round robin*—Servers are allocated connections in an even distribution.
- *Loaded*—Servers are assigned connections based on a combined metric of weigh and round robin.
- *Limit*—Servers are assigned limits as to the total number of simultaneous connections they can serve.

Cisco LocalDirector can also be configured in a redundant configuration through the use of two devices connected via a serial cable. The LocalDirector high-availability solution allows for stateful fail-over of active sessions if one of the LocalDirector devices is removed from service. Figure 11 highlights a configuration using redundant LocalDirector devices.


Figure 11 Cisco LocalDirector—Redundant Configuration



LocalDirector Acceleration with Catalyst 6000 Family  
 LocalDirector acceleration works through the use of a Catalyst 6000 family switch, a policy feature card (PFC), and an external LocalDirector. This form of SLB acceleration is supported only in dispatch mode. All TCP handshake packets, namely TCP-SYN, TCP-FIN, and TCP-RST are forwarded to the LocalDirector; UDP packets are also forwarded to the LocalDirector. It is the responsibility of the PFC with the Catalyst 6000 family switch to identify and forward these control packets. The result of this arrangement is that the Catalyst 6000 family switch handles a majority of server farm traffic once the initial load balancing decisions have been made through the LocalDirector. Because the Catalyst 6000 family switch is capable of much higher forwarding rates than the LocalDirector itself, the SLB process is able to scale to much higher forwarding rates and handle higher numbers of servers.

In the example shown in Figure 12, three servers are providing the same Web content. A sample configuration of these devices is shown in Appendix A. The advantages with this model of SLB functionality come from the integrated features within the Catalyst platform. In addition to SLB functionality, the Catalyst 6000 platform offers features including security access control lists (ACLs), quality of service (QoS), and accounting capabilities.

Figure 12 Cisco LocalDirector Acceleration



## Multinode Load Balancing

Multinode Load Balancing (MNLB) is a hardware and software combination SLB solution that is an implementation of Cisco ContentFlow Architecture. MNLB is a solution for providing extremely enhanced scalability of Internet appliance functionality such as the LocalDirector. The function provided by MNLB extends itself across multiple Cisco IOS software-based router and switch platforms.

Multinode Load Balancing consists of the following the components:

- *Service manager (SM)*

The service manager function is the heart of MNLB; it is an implementation of ContentFlow architecture's flow management agent. The SM(s) controls services applied to candidate flows within the network based on an installed policy. However, contrary to a traditional 'in-line' Internet appliance, the decision is based on the first packet of the flow. After applying the decision to a flow, the SM is no longer involved until a flow is terminated. This architecture allows the SM to be involved only as a control center and not as the prime forwarding path for all flows. The resultant action that is applied to a flow by the SM is called an *affinity*.

- *Forwarding agents (FA)*

The forwarding agent function is the workhorse of MNLB; it is an implementation of ContentFlow architecture's Flow Delivery Agent. The FA(s) is placed in line with the traffic flow and is responsible for forwarding traffic at very high packet-per-second (pps) rates. The FA does not independently dictate action to be applied to various flows; it captures and forwards flow packets that are considered "interesting" to the SM. As an example, the initial TCP-SYN flow initiator packet is classified as an interesting packet and is to be forwarded to the SM for service application. After the SM dictates the action to be applied to the flow, the resultant affinity is cached by the FA. Caching of the SM affinity allows the FA to forward successive packets of the flow at very high pps rates based on information in the affinity cache. The FA must be educated by the SM on which traffic it is to consider as being "interesting." The classification of interesting traffic can be granular to flag network-specific, host-specific, or

even protocol-specific traffic. The SM instructs the FA on which traffic to consider interesting and the resultant is placed in a *wildcard cache* within the FA.

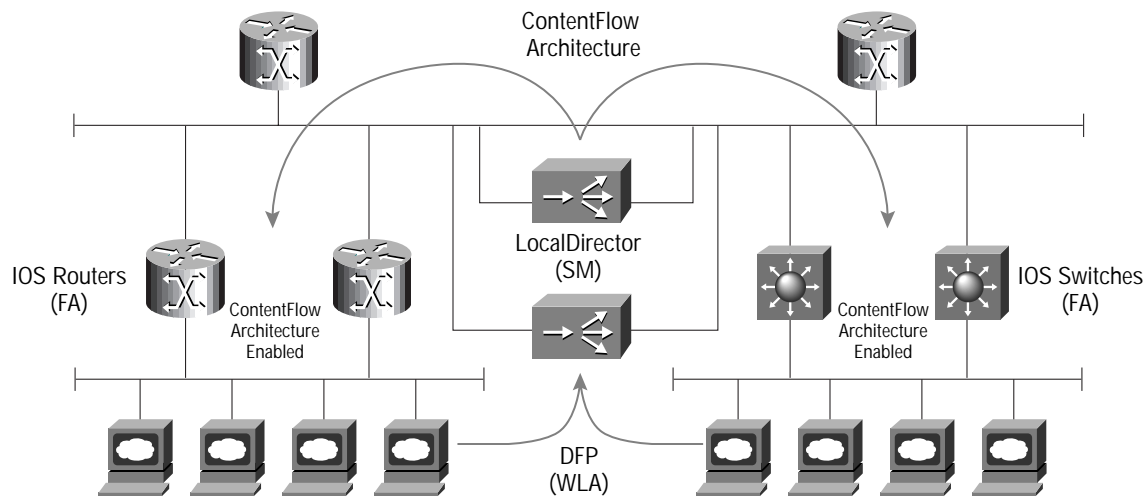
- *Workload agent (WLA)*

The workload agent is the major contributor to extensive scalability of the MNLB architecture; it comprises a software agent that resides on servers responsible for delivering content and applications. The workload agent provides real-time feedback as to the status of an application server. The status report includes statistics relative to the health of the server and its application. A report may include characteristics such as disk usage, memory usage, and CPU usage. This information can then be fed back into the system through the SM to affect policy that is applied to successive flows.

The MNLB system implements elements of the ContentFlow architecture. The role of the SM or ContentFlow architecture FMA is played by the LocalDirector. The component that directs interesting traffic to the SM, namely the FA or ContentFlow architecture FDA is served by multiple Cisco IOS software-based switches and routers, including the Cisco 7200/7500 series routers. Finally, the WLA capability is provided through the Dynamic Feedback Protocol (DFP), which provides a mechanism for servers to provide feedback information that the SM within the MNLB system can use to dynamically adjust its service application. Figure 13 shows the components of the MNLB service.

The MNLB system offers great advantages for high availability applications. The architecture offers high availability options in several ways while maintaining scalability. First, each SM has the ability to work in partnership with several FAs. In addition, each FA can receive instruction from multiple SMs. In this scenario, the FA will favor the instruction of one SM over another until there is an SM failure. As shown in Figure 13, each of the two SMs build affinities for the four FAs. Although each FA is acting only on affinities created by one of the SMs, should an SM fail, FAs will be able to revert to the backup SM.

Figure 13 Cisco Multinode Load Balancing Service Implementation



#### Geographic Load Balancing—The Cisco Approach

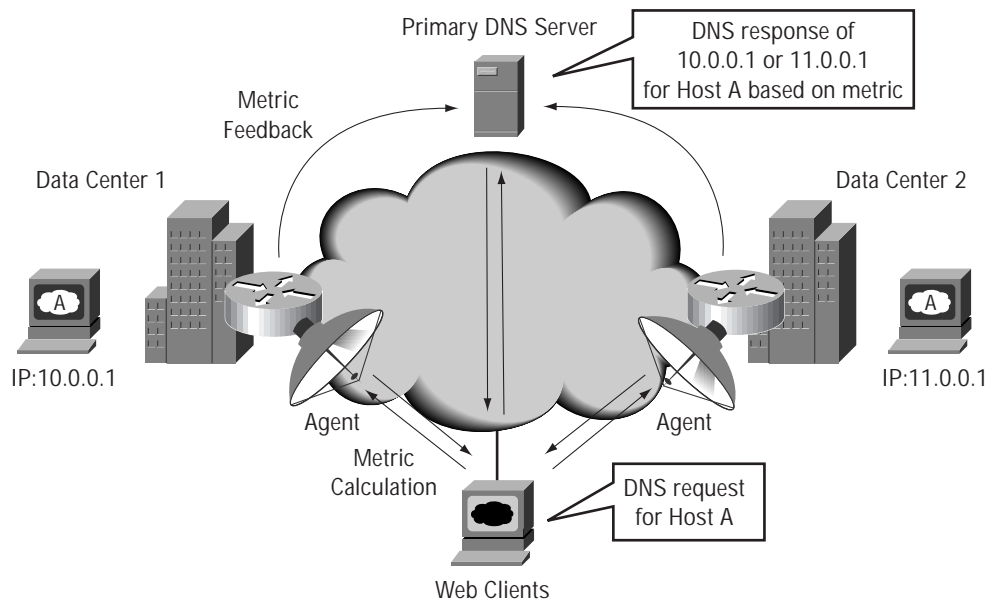
Cisco SLB tools provides the required functionality within a Web hosting solution to ensure high availability and scalability of content and applications within a data center. However, service providers willing to accept the challenge of providing a solid Web hosting service will quickly scale beyond the limits of a single data center. Target markets for Web hosting services within the Internet community are globally distributed worldwide thereby mandating a distributed Web hosting environment. It is not only important to build presence in international regions, one must ensure costly network resources are used sparingly such as transoceanic bandwidth. In order to optimize resource utilization and maximize service performance from a client perspective, the service model favors distributing Web content to remote data centers and serving requests locally. In addition, high availability requirements dictate that each data center shall provide backup content access should the local Web content become unavailable.

Similar challenges exist within the distributed Web hosting environment as seen within the SLB environment. These challenges, now potentially on a global scale, include:

- *Content Synchronization*—ensure all distributed servers possess most recent content
- *Content Accessibility*—ensure content is accessible regardless of individual distributed server status
- *Request Load Distribution*—ensure content requests are distributed favoring local access

As resiliency and scalability within an individual data center are covered by the SLB offerings, the challenge is reduced to simply ensuring content requests from clients are directed to the most appropriate data center. As within the SLB solution whereby an element, namely the server load balancer, was central to all requests entering a particular data center, geographic load balancing also requires a central control point for all content requests destined to any data center. The central point for a geographic load balancing function resides with DNS. All Web clients must contact a DNS server at some point prior to requesting content from a server within a data center. Based on the fact that geographically replicated content resides on servers with unique IP addresses, unique DNS responses can be provided to queries for the same Web server based on a series of metrics. Metrics are dynamically calculated and updated by a distributed set of agents throughout the network. Based on metric calculations between the series of agents and the requesting client, an appropriate decision can be made as to the direction a client's content request should be forwarded. Figure 14 shows the relationship between the various components of the Cisco global load balancing solution.

Figure 14 Cisco Geographic Load Balancing Solution



### Geographic Load Balancing by Cisco

In order to provide geographic load balancing functionality, Cisco has created a system comprising of an Internet appliance and a series of distributed agents named DistributedDirector. Following the previous discussed model, the Cisco solution provides dynamic DNS services based on a wide range of metrics. Cisco Distributed Director is an integrated piece of an overall powerful Web hosting solution

#### DistributedDirector

Cisco DistributedDirector is an IOS-based software product designed to run on the Cisco 2501/02 and 4700M router platforms. The agents that participate in the DistributedDirector system are protocols embedded within Cisco IOS-based routers. In order for a Cisco router to participate as a DistributedDirector agent, it must be running IOS software version 11.3(2)T or later. Agents are typically configured within data center border routers facing the Internet as these represent useful endpoints for measurement. DistributedDirector communicates to its agents through the Director Response Protocol (DRP). DRP is used for two main purposes:

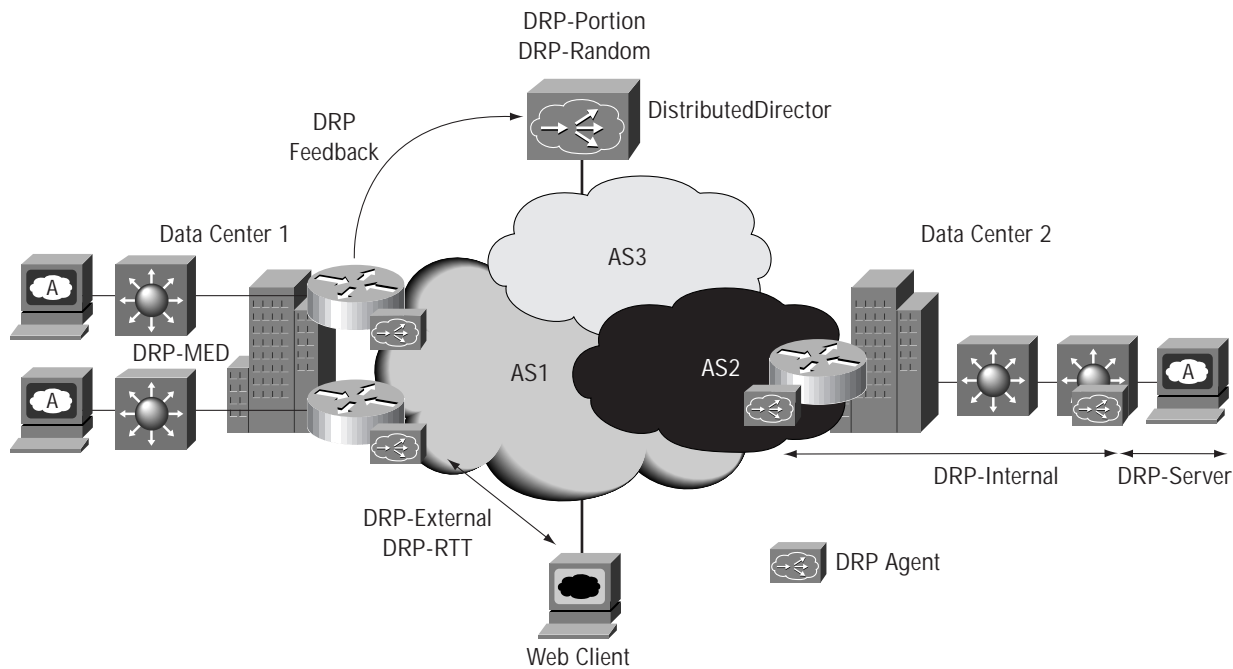
1. Query DRP server agents in the network for BGP and IGP routing table metrics between the data centers and clients to determine client-to-server topological proximities
2. Query DRP server agents in the network for client-to-server link latency metrics

DistributedDirector can use one or both metrics above to decide where to forward a particular content request. Several overall metrics can be used in combination that fit into the two above categories. The metrics that can be used by DistributedDirector include the following:

1. DRP-External—calculate BGP distances (AS hop counts) between DRP agents and querying client
2. DRP-Internal—calculate IGP distance between DRP agents and their closest BGP edge border router
3. DRP-Server—calculate IGP distance between DRP agents and their supported server farm
4. DRP-MED—calculation factors BGP MED value where two DRP agents are in same AS
5. DRP-RTT—calculate round trip times (TCP probe) between DRP agents and querying client
6. DRP-Portion—calculation based on assigned weight. Connection frequency proportional to weight
7. DRP-Random—calculation is random. DRP agent request is not necessary
8. Administrative Cost—assign statistical preference of one server over another—maintenance application

Figure 15 shows the relationships of all metrics used by DistributedDirector within a DistributedDirector system. As you can see, there are a variety of metric combinations that can be used to determine which server will be accessed by the client.

Figure 15 Cisco DistributedDirector System



#### Content Caching—The Cisco Approach


Typical Web caching solutions involve situation a series of caching devices in close proximity to a specific user community. For example, an ISP may install a series of cache devices at the border points between their network and upstream ISPs in order to minimize uplink bandwidth utilization. Caching solutions offer several benefits to the service provider that opts to install cache services. These benefits include:

- Maintenance of high percentage localized traffic to minimize upstream bandwidth usage
- Quicker response for client community accessing cached content
- Greater service scalability due to high percentage of localized traffic
- Totally transparent cache solutions from the client perspective

However, in a typical Web hosting environment, it is not feasible to locate content caches in close proximity to the user community. Service providers providing Web hosting services typically tie into multiple Internet access points with no guarantee of close proximity to a large concentration of the Web host user community. From this standpoint, it is impossible for a Web hosting service provider to situate caches close to the user community.

Content caching does however have a powerful application within a Web hosting environment. Web hosting providers can deliver accelerated hosting services to their customers by front-ending Web server farms with cache engine clusters. In this application, content requests are redirected to a cache engine cluster instead of directly forwarding them to the server farm. When the cache cluster fulfills these requests, it off-loads traffic from the server farm thereby minimizing content download latency and increasing server farm capacity. Therefore, once a particular piece of content is requested by a client, it is cached so that successive requests are not directed repeatedly to a server. Figure 16 shows the cache application within a hosting environment. Within this environment, the cache engine cluster is only able to cache the content that is available on the local servers. This arrangement is referred to as Reverse Proxy Caching function.

Figure 16 Web Cache Cluster Application

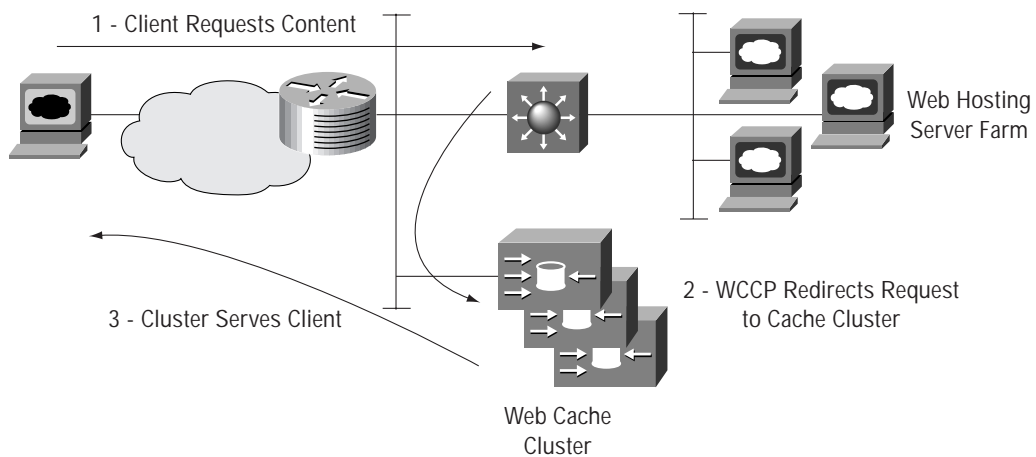


The key to deploying cache services is to ensure the resultant function will be transparent to the user community. It is not feasible to expect a user population to change their endstation configuration such as define a proxy server in order to utilize the cache devices. In order to deliver transparent service, Cisco has developed the Web Cache Communication Protocol (WCCP). WCCP is a protocol that runs between Cisco IOS router and switch platforms and Cisco Cache Engine products. The underlying function of WCCP is to redirect HTTP requests for content from a user community to an appropriate Web cache engine or cluster so that the request may be filled from cached information. Cisco

IOS-based devices equipped with WCCP and transparently redirect Web content requests to a cache engine transparent to the user.

The operation of transparent caching is shown in Figure 17. When a client requests Web content, the request is immediately redirected towards the cache engine cluster. If the content exists within the cluster, it is served to the client. However, if the content is not present within the cluster, the cluster itself will retrieve the content from the server, cache the content, and provide the content to the requesting client.

Figure 17 Web Cache Cluster Transaction



While WCCP Version 1 provided the basic functionality to build a transparent caching solution, WCCP Version 2 adds tremendous functionality in relation to high availability.

WCCP provides content management within a cache cluster. Each member of a cache cluster is responsible for managing a unique “bucket” of content. The “bucket” scheme allows a predetermined configuration of cache engine mapping to cacheable content thereby optimizing response time. The cluster is self-healing in the sense that if a cache engine malfunctions, the other cache engines will assume the “buckets” and continue to provide full cache service. If an entire cache cluster is rendered inactive, the WCCP-enabled routers will detect this condition and cease rerouting Web requests to the cache cluster. Two routers can share a cache cluster thereby allowing for fail-over functionality. Using caching and HSRP, if the primary HSRP gateway router fails, secondary HSRP gateway router can resume service with the cache cluster.

WCCP also provides additional functionality for added flexibility and resiliency. Enhanced functionality includes the following:

- *Overload Bypass*—if a cache cluster is unable to keep up with content demand, it can signal the WCCP-enabled router to cease redirecting requests to the cache cluster until traffic subsides.
- *Dynamic Client Bypass*—Special provisions exist within WCCP to dynamically allow certain requests that are not cacheable to bypass the cache cluster. Web authentication pages are an example of content that must be accessed directly by the user and is allowed to bypass the cache cluster through WCCP.

#### Content Caching by Cisco

Cisco provides many options for building a Web cache solution. A combination of WCCP agents and Cisco cache engines are chosen based on performance requirements. WCCP agents within IOS-based routers and switches are supported in the following IOS releases:

- WCCP v1
  - 11.1(14)CA and beyond (CA only)
  - 11.2(10)P and beyond (P only)
  - WCCP v1 is NOT supported in 11.1/11.2 mainline nor 11.3 any release
  - 12.0 any release
- WCCP v2
  - 12.0(3)T and beyond (T only) (WCCP v2 only)
- No release currently supports both WCCP v1 and v2

Most IOS-based routers and switches support a WCCP version with the exception of the Catalyst 8500 Series and the Cisco GSR router.

#### Cisco Cache Engine 500 Series

Cisco provides two main platforms for cache engine products. The Cisco Cache Engine 500 series comprises of the Cache Engine 505 and the Cache Engine 550. While the two models offer similar functionality, they differ in performance characteristics. The three main criteria for choosing a cache engine are based on the following:

- Transactions per second (TPS)
- Number of simultaneous connections
- Disk storage requirements

The Cache Engine 505 supports approximately 75 TPS, up to 500 simultaneous connections, an internal 9 GB disk storage and an external storage connector. The Cache Engine 550 supports approximately 200 TPS, up to 3000 simultaneous connections, an internal 18 GB disk storage and can be configured with external storage as necessary. Clustering of cache engines must be considered to maximize cache availability and allow for larger storage scaling.

### IP Address Management and Distribution—The Cisco Approach

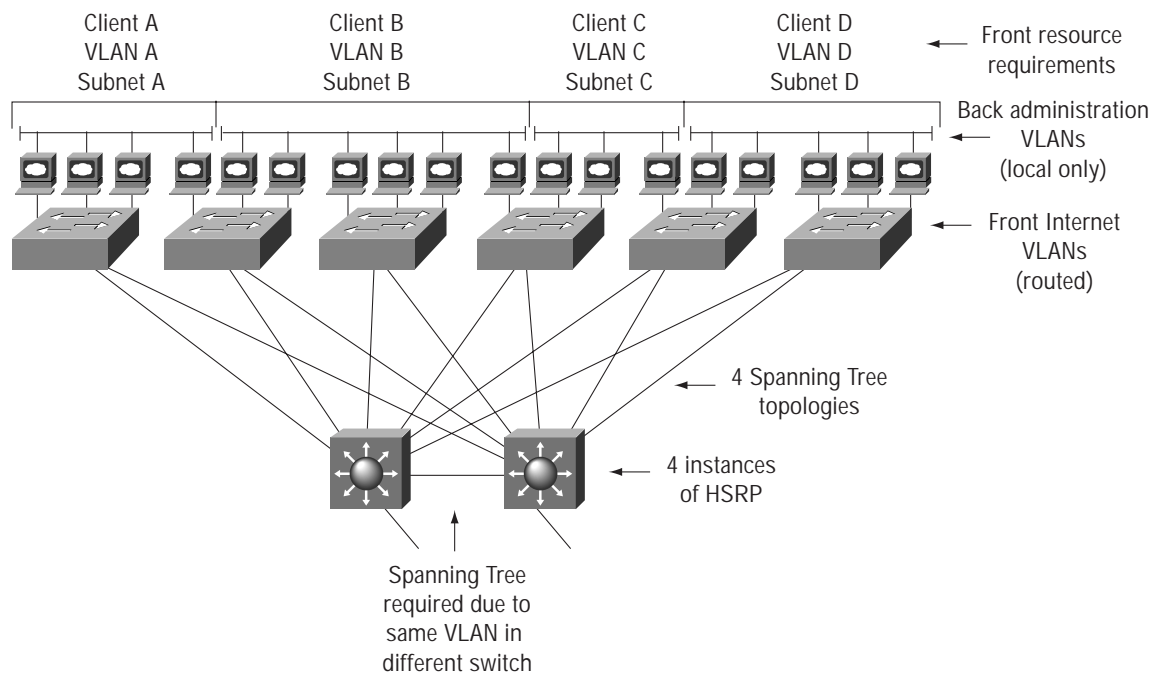
One of the biggest challenges for the Web hosting service provider today is the tedious management and distribution of IP address space. Due to the nature of a Web hosting service, servers must reside within Internet-routable address space. Routable IP address space can be an extremely limited resource for the service provider. Address space must be allocated efficiently to allow for maximum service coverage and adequate justification when applying for additional address space through ARIN (American Registry for Internet Numbers). Although Network Address Translation (NAT) is a potential solution to limited routable IP address availability, typically NAT performance rates do not yield wire-speed results as required.

A Web hosting environment is representative of a multi-client server farm. Each client hosts their Web services on a series of servers within a common data center. As such,

security is of utmost concern between servers belonging to different clients. A common way to ensure security between Web hosting clients is to deploy their servers in a unique VLAN and associated IP subnet. Using the VLAN mechanism, each client is segregated from other customers at Layer 2 preventing any malicious activity and any ethernet snooping of information. However, this model of assigning unique VLANs and IP subnets poses great scalability limitations. These limitations include the following:

- **VLAN limitation**—LAN switches have an inherent VLAN numbering limitation
- **STP complexity**—With each VLAN, an associated spanning tree topology must be managed
- **IP address inefficiencies**—IP subnets must be deployed based on address boundaries
- **Routing limitations**—With each subnet is an associated redundant gateway configuration HSRP

Figure 18 Multiclient Server Farm—Traditional Deployment



All of the above limiting factors lead to eventual exhaustion of network resources. Figure 18 shows a typical representative multiclient data center design. In this design, each client has two VLANs assigned to them. The 'front' VLAN is used for access to the Internet and the "back" VLAN is used for administrative purposes. The "back" VLANs are local in that they are only used for communication between servers and commonly built using separate hardware. As such, the "back" VLAN remains local and can be addressed using private IP address space (RFC1918) and are of little concern. What is of great concern is the usage of registered IP address space and spanning VLANs on the "front" side of the servers. While this is an example, it is easy to see that the inefficiencies grow rapidly as one considers a traditional data center might house thousands of customers. Table 1 show the IP address calculations required and the levels of wasted address in this example. These numbers do not take into consideration the complexity of managing switch port assignments, router configuration statements and VLAN trunk configurations. While this model may work for smaller multiclient data centers, it will likely not scale to larger Web hosting environments.

Table 1 IP address Allocation Plan - Traditional

	Data Center Client			
	A	B	C	D
Total Servers	4	6	3	5
Total Broadcasts Req.	2	2	2	2
Total Gateways Req.	2	2	2	2
Total Hosts Req.	8	10	7	9
Subnet Size Req.	/29	/28	/29	/28
<b>Wasted Addresses</b>	0	6	1	7

When one considers the purpose of the "front" network, it is used strictly to provide connection from the servers themselves to the Internet. The "back" channel is used for administrative tasks including backup, network management, and content replication. Therefore, server NICs connecting to the "front" network only require access to the default gateways and not necessarily to other servers. This understanding of the purpose of the "front" network is the foundation of the Cisco solution to addressing large multi-client Web hosting environments.


#### IP Address Management and Distribution by Cisco

What is required in most cases in terms of connectivity within the "front" network of a multiclient data center is connectivity to the Internet. However this must be provided in a secure implementation such that servers from one client cannot interfere with servers from another client at Layer 2 or above. The "back" network is typically used to facilitate administrative tasks between servers of the same client. Therefore in most cases servers belonging to the same client may not need to communicate with one another on the "front" network. Corner cases where this communication is required in the "front" network will be considered shortly. The Cisco IP address management solution involves a new VLAN mechanism whereby servers are only able to communicate with default gateways on the "front" network while residing in the same VLAN. This new VLAN feature is called a private VLAN.

#### Cisco Private VLANs

A private VLAN is a Layer 2 mechanism whereby two types of access ports of varying security levels are created within the same Layer 2 domain. Ports where servers are to be attached are called "private ports." A private port is restricted at Layer 2 and above to only being able to send/detect traffic to/from promiscuous ports. A "promiscuous port" does not have the restrictions of a private port and typically is connected to a router or Layer 3 switch interface. Simply put, within a private VLAN, traffic received on private ports will only travel to the promiscuous ports. Traffic received on promiscuous ports will travel to all ports both promiscuous and private. Figure 19 shows the relationships between these two types of ports within a private VLAN.

Figure 19 A Private VLAN



The application of a private VLAN is most effective within the “front” network of a multi-client data center. Since the “front” network need only provide connectivity between servers and their default gateways, a private VLAN provides this secure connectivity without the need for multiple VLANs and subnets. In this model, all stations, regardless of which client to which they belong, are connected to a private VLAN. All stations therefore will have connectivity to the default gateway(s) without having access to any other servers within the private VLAN.

A private VLAN has several design characteristics that must be considered. Private VLANs are currently only supported on the Catalyst 6000 Family of switches. Private VLANs and normal VLANs can exist on the same Layer 2 switch. Multiple promiscuous ports are permissible within the same private VLAN. Multiple promiscuous ports allow for multiple default gateways in an HSRP arrangement. The mechanisms of the private VLAN that provide for its implementation and security are enabled through ASIC configuration and do not inhibit the performance of the switch.

If we consider the example from Figure 18 as implemented with a single private VLAN, the numbers begin to look quite favorable. Figure 20 along with Table 2 show the positive effects of using a private VLAN implementation. When building a private VLAN implementation, the subnet size becomes irrelevant. The private VLAN implementation consists of only one IP subnet and hence all IP addresses are usable by any customer while still maintaining security. The addressing paradigm shifts from assigning IP subnets to individual customers to assigning a series of individual IP addresses based on the number of servers the client possesses. The private VLAN model eliminates the waste and complexity caused by highly fragmented IP address deployments. It was virtually impossible in the old model to allocate a client the ideal subnet size to accommodate all future growth. This lead to the likelihood of the clients receiving fragments of IP address space for their servers. Therefore allocating potentially discontinuous pools of IP addresses within the same network to clients in the private VLAN model does not pose any additional difficulty.

Figure 20 Multiclient Server Farm Using Private VLANs

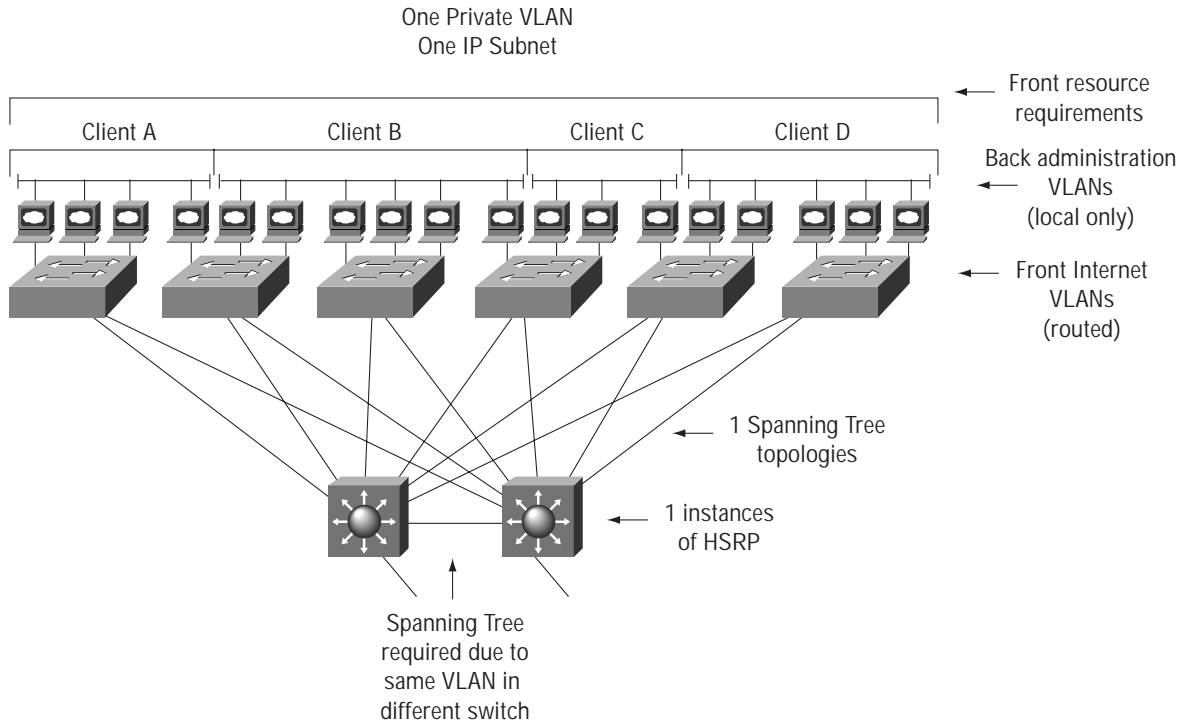


Table 2 IP address Allocation Plan—Private VLAN

	Data Center Client			
	A	B	C	D
Total Servers	4	6	3	5
Total Broadcasts Req.		2		
Total Gateways Req.		2		
Total Hosts Req.		22		
Subnet Size Req.		Irrelevant		
Usable Addresses		10		

Several other applications exist for a private VLAN within a data center. Some other examples are listed below:

- **Common backup VLAN**—server NICs within the “back” network from multiple clients can reside in a private VLAN and not have communication with one another. The “private” ports are those with servers attached. The “promiscuous” ports are used to connect backup servers to the VLAN allowing connectivity from the backup server to all connected client servers.
- **Common network management VLAN**—similar to the common backup network example, the “private” ports are connected to the servers and the “promiscuous” ports are connected to network management stations allowing access from the management consoles to all servers in a secure manner.


Cisco Private VLAN Communities

Within a data center, there may be situations whereby communication is required between servers in the “front” network belonging to a particular client. Examples of this requirement are listed below:

- *Fault tolerant server NICs*—are arrangements whereby a server has two NICs that provide resiliency for the server. The two NICs send a ‘heartbeat’ between one another to ensure operational state.
- *Server clustering*—is an arrangement whereby multiple servers participate in each hosting a portion of an application. Typically there exists a great deal of intra-cluster communication that takes place within the same VLAN.
- *Content replication*—may occur if the data center server deployment does not include a dedicated “back” network. In this case, all management and content replication must happen between servers on the “front” network.

Private VLAN Communities is an enhancement to Cisco private VLANs whereby a subgroup of servers within a private VLAN are able to communicate with one another while remaining in the same private VLAN and IP subnet. It is a concept of subgroups in which smaller Layer 2 broadcast domains are created within the private VLAN. While a private VLAN community does allow for flexible IP addressing space while still possessing the ability to create smaller “sub-VLANs,” one should not rely simply on communities. Each private VLAN community within a private VLAN consumes an additional system VLAN thereby yielding a community limit.

Figure 21 A Private VLAN Community



### The Total Cisco Solution—Piecing it Together

In order to build a solid reliable Web hosting service, many key technologies must be integrated together into a manageable design. Cisco offers all required components and integration capabilities to build a scalable Web hosting service. Building upon the Cisco proven reliability and expertise in building solid routed and switched networks, integrating Cisco Web hosting services becomes a relatively simple task. Many product and technology choices exist in all areas of the design based on scalability requirements.

Figure 22 shows the Cisco proven multilayer data center design capable of scaling to 10,000+ servers. This design consists of many Layer 2 and Layer 3 high availability features mentioned previously to guarantee a highly-resilient underlying infrastructure. In order to add the required Web hosting functionality, one simply integrates the previously described services into the design. Due to the modular nature of this design, one can scale an individual data center by

building additional “distribution” blocks and tying them into the data center core. To build additional data centers, one simply replicates the individual data center design.

Figure 23 shows the modular data center design with integrated Web hosting services. The three main services, namely SLB, Geographic Load Balancing, and Web Caching have been integrated into the design. High availability and scalability are present in these new services through a variety of mechanisms. Within the SLB function, high availability and scalability take form in multiple load balancers with stateful failover capability. Multiple DRP agents exist within the design to provide resiliency for the Geographic Load Balancing function. To gain resiliency and scalability with Web caching, one can install multiple cache clusters which can be accessed through WCCP. Redundant WCCP agents within the distribution Layer 3 switches provide resiliency and cache accessibility should a distribution switch be removed from service.

Figure 22 Cisco Multilayer Data Center Design

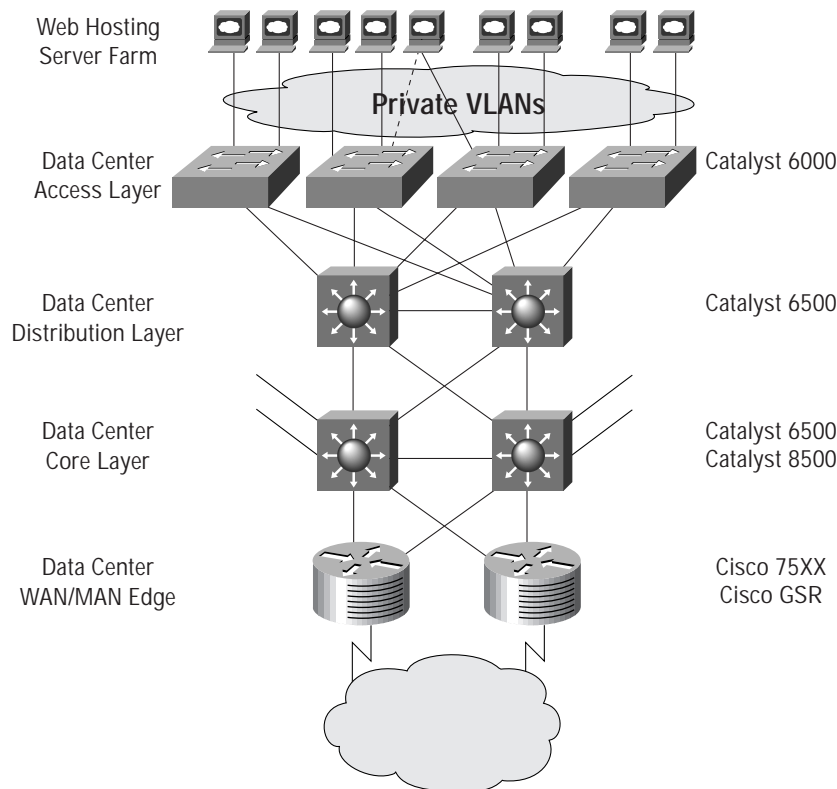
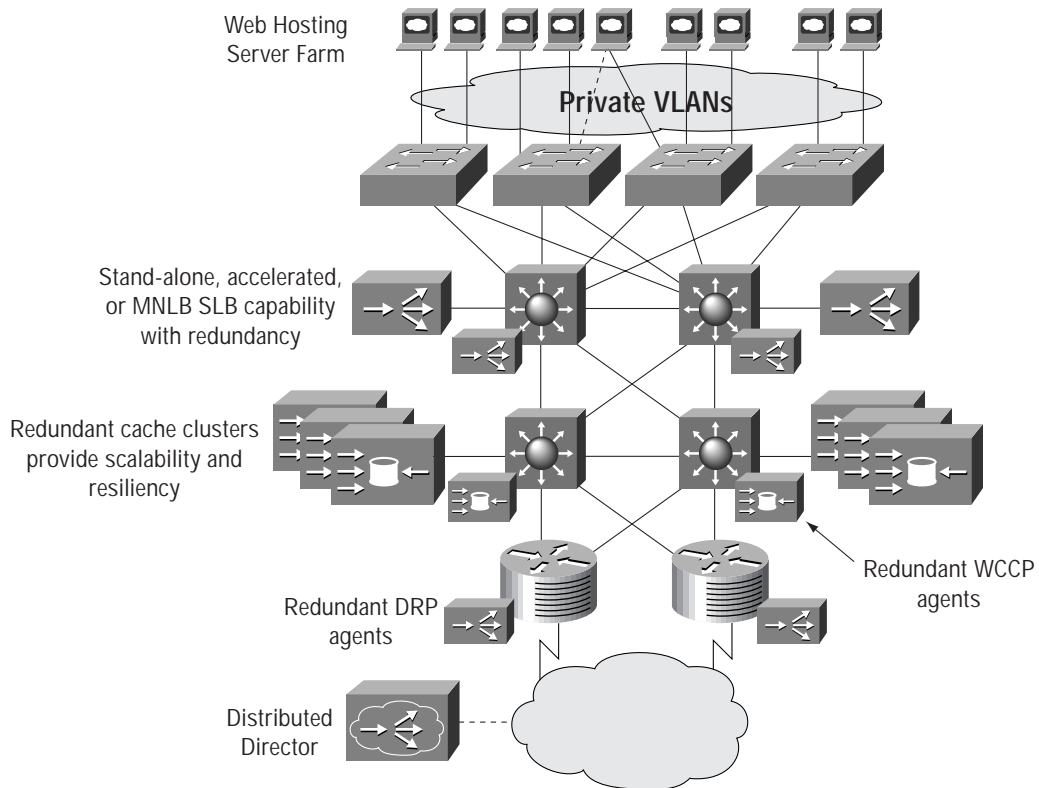


Figure 23 Cisco Web Hosting Integrated Service Design



#### Feedback—A Key to Scalability

One of the key advantages to the Cisco Web hosting solution is the dynamic nature of the service management. Every component of Cisco Web hosting service is engaged in some form of dynamic feedback, ultimately affecting the overall resource allocation of the service. As it is not practical to insist on service management staff to continually monitor every performance variable of the network on an ongoing basis in great detail, Cisco has constructed many automatic mechanisms within its Web hosting services to allocate resources automatically. Although many have been mentioned throughout this document, this section serves as a summary of the major feedback mechanisms. It is through these mechanisms that one can guarantee the most optimal dynamic allocation of costly network resources. The feedback systems within this design are built to allow for a very responsive system able to track network environmental changes and react proactively and expediently.

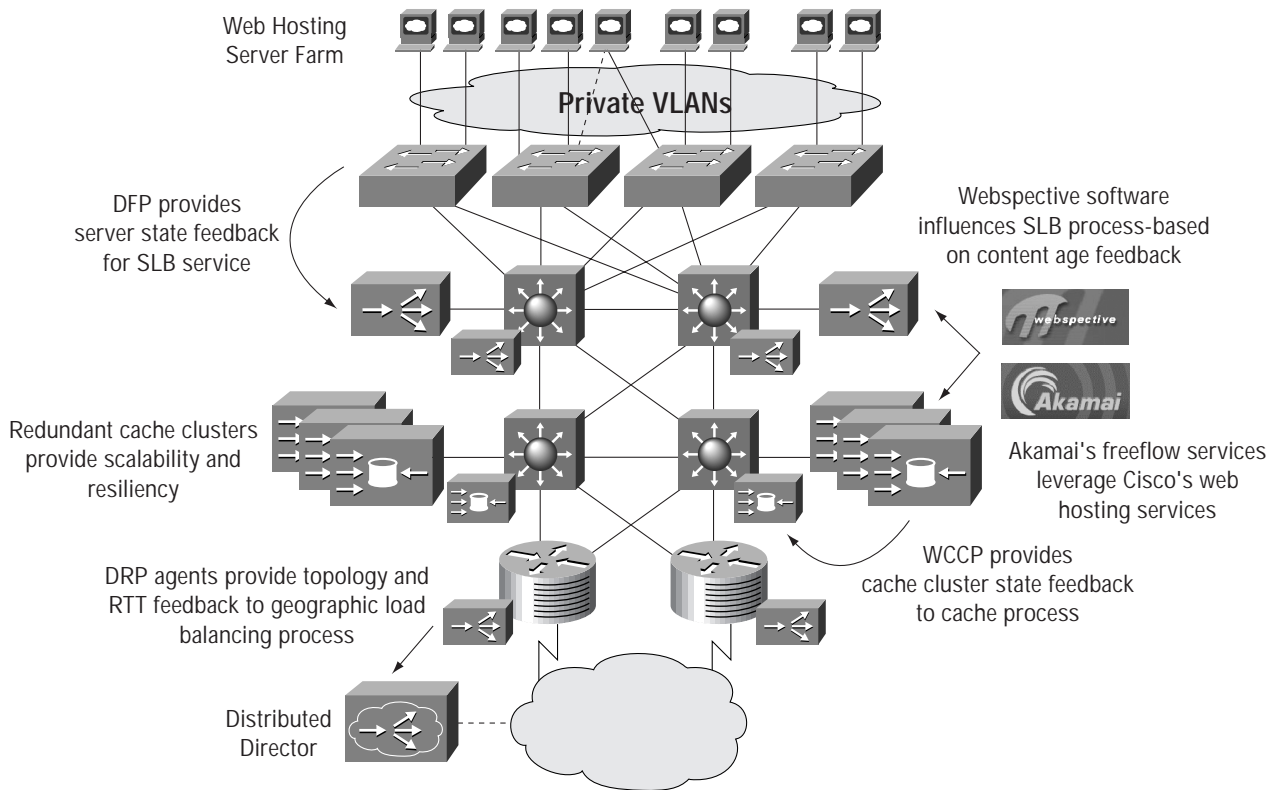
The following service feedback mechanisms are present within the Cisco Web hosting service. Figure 24 is a graphic representation these service feedback mechanisms. In addition, many of the Cisco partners are continuing to write enhanced Web hosting service “enhancers” which heavily leverage the Cisco feedback mechanism.

1. Dynamic Feedback Protocol (DFP)—facilitates the feedback of server state information to the LocalDirector SLB mechanism. Servers can be increased or decreased in priority for connection servicing based on the status of server environmental variables including CPU, memory, and disk usage.
2. LocalDirector User Interface (LUI)—is a graphical Web-based interface featuring application health probing for Web servers. A series of application probes periodically gather application availability information from Web servers. The resultant data can be dynamically fed into the SLB algorithm of LocalDirector to immediately affect load sharing decisions. LUI ships with every LocalDirector.

3. Director Response Protocol (DRP)—provides network topology and service response time feedback to the DistributedDirector. DRP agents reside within IOS-based router and switch platforms. DRP provides the information to DistributedDirector yielding the appropriate decision on where to send incoming content requests. As information updates are forwarded via DRP to the DistributedDirector, geographic load balancing policies are dynamically updated.
4. Web Cache Communication Protocol (WCCP)—provides a feedback mechanism from the Cisco Cache Engine products to WCCP-enabled Cisco router and switch platforms. As state changes within a cache cluster due to added/removed cache engines, request congestion, or other availability factors, the cache cluster can use WCCP to notify supporting routers and switches so they may alter their cache practices if necessary.
5. Akamai Freeflow Technology – is a partnership between Cisco Systems and Akamai Technologies targeted at integrating Akamai’s powerful content distribution

- network services with Cisco Web-hosting services. Together, Akamai’s Freeflow software will have direct feedback into Cisco Web hosting mechanisms.
6. Webspective Technology – is a partnership between Cisco Systems and Webspective Software targeted at integrating Webspective’s content distribution and verification software with Cisco Web hosting services. Webspective software is a content management and distribution system whereby it can deliver content to distributed Web servers and cache engines in a timely fashion. Together with Cisco, Webspective software will be able to integrate with the Cisco Cache Engine and LocalDirector services to ensure not only that caches possess current content but that the SLB function performed by Cisco services can make SLB decisions based on the age of content in addition to other metrics.

Figure 24 Feedback—A Key to Service Scalability



## Conclusion

Although Web hosting services is not a completely new concept, the scale at which many service providers are forecasting does warrant special attention. In order to build a scalable and highly resilient Web hosting service, one must start with a solid infrastructure. A last report, approximately 80 percent of the Internet's traffic travels through a Cisco infrastructure. Time-tested Layer 1-3 scalability, resiliency, and stability have allowed Cisco master the challenge of building solid data networks. Building upon this base, a suite of advanced networking services, specifically tailored to Web hosting environments must be seamlessly integrated. Many of Cisco Web hosting technologies front many of the large content providers of the Internet today. While Cisco has a comprehensive suite of the necessary advanced services, they also possess the technology to integrate them to provide a complete end-to-end service. To ensure a Web hosting service operates at its level of optimal resource utilization, multiple feedback mechanisms within the network are required. Cisco has continued to develop upon its highly successful IOS technology to include agents and service managers capable of integrating into the entire feedback process.

Today Cisco possesses all the necessary technology and expertise to build a complete robust end-to-end Web hosting service.



### Corporate Headquarters

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

### European Headquarters

Cisco Systems Europe s.a.r.l.  
Parc Evolic, Batiment L1/L2  
16 Avenue du Quebec  
Villebon, BP 706  
91961 Courtaboeuf Cedex  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 69 18 61 00  
Fax: 33 1 69 28 83 26

### Americas

**Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-7660  
Fax: 408 527-0883

### Asia Headquarters

Nihon Cisco Systems K.K.  
Fuji Building, 9th Floor  
3-2-3 Marunouchi  
Chiyoda-ku, Tokyo 100  
Japan  
<http://www.cisco.com>  
Tel: 81 3 5219 6250  
Fax: 81 3 5219 6001

**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the**

**Cisco Connection Online Web site at <http://www.cisco.com/offices>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE  
Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia  
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore  
Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela