

# Cisco i-commerce Solution

## A Cisco End-to-End Design Guide

- UK New-Ventures Team

### Overview

While today's companies restructure to embrace an iCommerce business model, many are faced with the technological and operational challenges posed by the required infrastructure. In order for a company to engage in an iCommerce model, one must not only have an internet presence, but one must also possess a highly scalable and resilient solution. Prospective customers of the iCommerce offering 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 customers all the integrated elements required to build a solid iCommerce solution. Building on Cisco's reputation to provide highly scalable and resilient layer 1-3 infrastructure, Cisco continues to add services of higher intelligence to its proven design model. Adding layer 4-7 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 a company to minimize costs associated with growth, management, and maintenance of an iCommerce network infrastructure.

In order to build a successful iCommerce implementation, one must consider three key characteristics of the service, namely scalability, high availability, and manageability. When evaluating scalability, one must consider more than simply the ability to increase device performance and link bandwidth. A point in time may exist where the company must also consider scalability on a geographic basis. 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.

Looking at the characteristics of an iCommerce solution, important scaling parameters must be considered. In a typical enterprise server farm environment, one is mainly concerned with throughput capabilities of the network infrastructure. Throughput typically takes the form of bits-per-second and packets-per-second capabilities within network devices and data links. However, in designing an iCommerce solution, one must also give consideration to the connection rate and state of simultaneous customer transactions. The consideration for transaction volume scaling within the network infrastructure itself represents a shift in design practice. The iCommerce network gains intelligence through its awareness of transaction state; a necessity for many of the associated technology services. Figure 1 shows the relationship of the three performance characteristics of an iCommerce

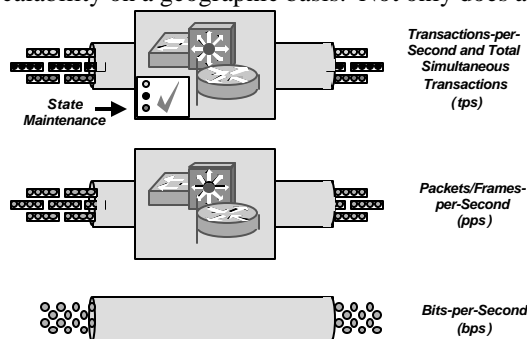


Figure 1. Scaling the Network Infrastructure

network design.

High availability must be maximized at all layers of an iCommerce network design ensuring maximum accessibility to web content and transaction systems. In its most basic form, high availability design ensures layer-2 protocols such as Spanning Tree (STP) are kept in check and do not create lengthy outages during a link failure within a layer-2 domain. On a more global basis, high availability amongst layer-3 protocols ensures fast rerouting is possible during IP transport failure. However, in an iCommerce environment, high availability is also critical at higher layers of the service. The purpose of an iCommerce network is to service transaction requests from a large numbers of web clients in a resilient manner. The service description does not simply ensure the request is served, but the request must be served in a way that will offer the most pleasurable experience for the iCommerce customer. Pitfalls of a web hosting service can include uneven loading of content requests amongst eligible servers, requests forwarded to inactive servers, and requests forwarded to servers with partially inactive pages and broken links. In order to bypass all potential iCommerce pitfalls, a highly sophisticated suite of protocols and services are required within the iCommerce environment that are deployed in a tightly integrated fashion.

Manageability is always an important aspect of any iCommerce network. The ability to collect statistics including Netflow®, RMON, and SNMP data at key points within a Cisco infrastructure allows the customer many options for detailed network analysis and appropriate capacity planning. Equally important is the ability to configure devices in a common manner. Cisco's IOS® software technology presents a well known user interface which has been replicated across all relative devices within a Cisco iCommerce solution.

The purpose of this paper is to present Cisco's end-to-end solution for a solid iCommerce design. Many facets of Cisco technology are integrated to provide the overall solution. Individual technologies offered by Cisco will be explored in the context of their integration into the overall iCommerce service design. Please refer to Cisco's website at <http://www.cisco.com> for more information on designing large scale IP networks using Cisco technology.

## The iCommerce Design

### Transaction Details

Prior to designing for an iCommerce network solution, it is best to fully understand the mechanics of a customer transaction from a network perspective. A typical customer transaction consists of a mix of IP connections that can occur within various parts of the design. Several connection profiles are evident during a transaction process and are identified by their IP protocol characteristics. Each connection stage of a typical customer transaction is described below.

The first connection of the transaction is shown in figure 3a. Upon the customer entering the URL of the iCommerce site in their web browser (eg. [www.forsale.com](http://www.forsale.com)), a resolution is made by a DNS server responsible for the site domain (eg. [www.forsale.com](http://www.forsale.com)). The resolution returns the IP address of the web server housing the requested URL.




Figure 3a. Connection #1:DNS Request

The second connection of the transaction is shown in figure 3b. Once the IP address of the web server is known, typically several connections are made to the server to download the main web page. Successive similar connections will be made as the customer browses the site while shopping. Two different behaviors can occur here depending on the capabilities of the client browser and associated web server. The usual behavior for a browser to follow when downloading a web page is to open a new TCP connection for each separate element of the web page including graphics. As iCommerce websites tend to embed many graphics and scripts into their web pages, this can result in high numbers of concurrent TCP connections per user and overall. An enhancement is offered as an option in HTTP/1.0 and as default for HTTP/1.1 to create a 'persistent' web server connection. Using a persistent HTTP connection, several elements of a webpage including graphics can be downloaded through one HTTP-TCP session. The client browser and associated web server negotiate how many persistent connections are allowed and how long to keep such a session open before closing it. In addition, by increasing the value of keep-alive timers settings for persistent connections on the server, data from multiple webpages can be also be downloaded in a common HTTP-TCP connection. A balance must be maintained between total number of persistent connections allowed vs. the persistent connection expiry timer in order for relatively few users to hog resources of the web server. In designing an iCommerce solution, one should maximize the use of persistent HTTP connections.




Figure 3c. Connection #3:Secure Transaction

The third connection of the transaction is typically invoked when a customer either accesses their own personal account or actually enters the 'checkout' portion of the website whereby a customer actually makes the purchase. At this point a secure encrypted connection is established between the client and the iCommerce web server.




Figure 3d. Connection #4:'Back-end' Processing

The secure connection is formed using the Secure Socket Layer (SSL) protocol and is identified by the 'https://' designation on the client web browser. Any successive data transfer or new connections between the customer and the iCommerce server will be established and transmitted using the SSL protocol designated by TCP port 443. Figure 3c shows the SSL connections established during this part of the transaction.

The fourth set of connections actually consists of a series of connections that can occur at any time throughout the transaction. These connections establish access to 'back-end' transaction and database servers. As the user proceeds through an iCommerce transaction, several databases are typically consulted for such information as product specifications and pricing, order history, and product inventory. In addition, specialized application servers may exist for transaction processing, inventory management, and credit card authorization. Finally 'hooks' into ERP applications may exist to automatically track financials, supplier management, and other logistics. Although these

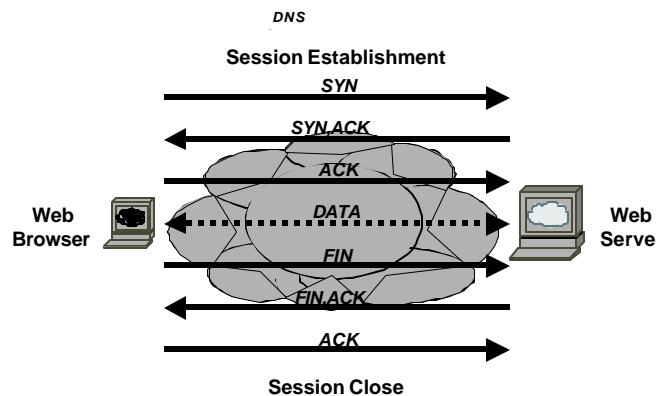


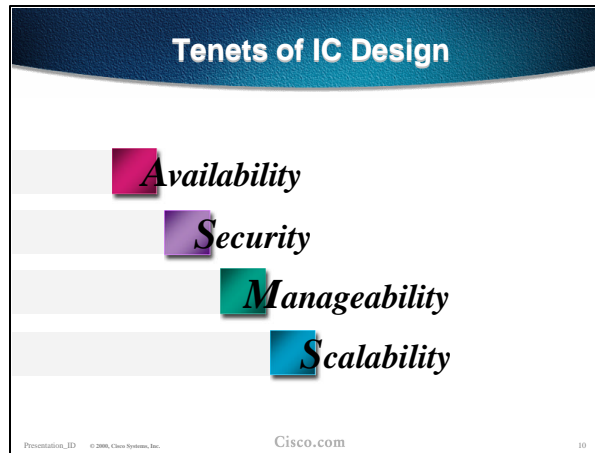
Figure 4. Stages of a TCP Connection

connections can use a series of TCP/UDP ports and can occur throughout the transaction, they occur behind the web servers on the 'back-end'. Figure 3d shows the relevant connections.

Each TCP connection between customer and server follows a particular pattern. TCP, being connection-oriented, follows a sequence of handshakes to establish a session between client and server. Figure 4 shows the stages of session creation and session termination between a client and a server. As seen in figure 4, a series of three handshakes must take place before any web content will flow from server to client. It is this handshake pattern that allows devices within the iCommerce design to track TCP connection state. Connection state awareness is an important characteristic of many advanced network services that will be discussed below.

Let's continue by looking at the typical requirements of icommerce companies.

## 4 Tenets of Cisco's i-commerce Architecture



There are four requirements of i-Commerce network design that Cisco always try to deliver the customer. They are below:

### Scalability

E-Commerce Networks are like “Mini-Service Provider” networks. They might start out with a requirement for a single T-1 but in 4 months with the help of a marketing blitz a T-3 might not be capable of supporting their peak hours. The networks we design need to be EXTREMELY modular and not lock the customer in to a single architecture or rely to heavily on a “one box does it all” mentality. If we can deliver on this in the architecture then as our clients network demands increase we can simply and effectively upgrade the network.

### Manageability

In the i-Commerce network, especially the remote Co-located network, Manageability is critical. If your network is a 2 hour drive away and you have a failure of any component the ability to our-of-band diagnose, troubleshoot, and repair and outage without being on-site is paramount.

### Security

This is the tenet that is probably the most important in customer confidence. Fortunately, to date, there has not been a MAJOR hack of an i-Commerce site. Credit card numbers have not been posted on a web page or medical patient records have not been distributed through email. Let's keep it that way. Security is more than a firewall. True i-Commerce security requires a combination of firewalls, IOS Extended ACLs, intrusion detection, active audit, and a bunch of server based technologies like PGP encryption of sensitive files, SSL v3, and SSH host management. It also requires the creation, adoption, and enforcement of a comprehensive security policy from the client

### **High Availability**

This is the tenet most commonly sacrificed because of cost. Building a resilient network is not an inexpensive proposition. Certainly we have marketing programs and special pricing to make it easier to swallow; however, it still requires a considerable investment from the client. But then again, a 4 hour

outage of a publicly held E-Commerce company cost them \$2 billion dollars in market capitalization. The worth of resilience can simply be measured by the question “What does downtime cost?” This might not be as tangible as lost revenue; but more about reduced customer confidence in the content provider’s ability to deliver on service and product.

Providing a redundant self-healing network is much more difficult than just throwing around some additional power supplies and adding redundant Supervisor Engines. There are three types of resilience we can build into a network.

- Network Level

Network Resilience is what we built into the E-Commerce architecture. Not only is each component redundant and robust, but the network is capable of self-healing. In an architecture designed with this level of redundancy in mind a component or multiple components can be powered down and the network continues to function.

- System Level

System Resiliency involved building a single component of the network as reliably as possible. Dual Power Supplies, Dual Processors, and warm sparing are all mechanisms of providing this.

- Maintenance Contract/Sparing

Cold sparing of parts locally or in close proximity, as well as quick response time maintenance contracts can mitigate a potentially huge problem in a resilient network but if an E-Commerce network is down this long it is too late already.

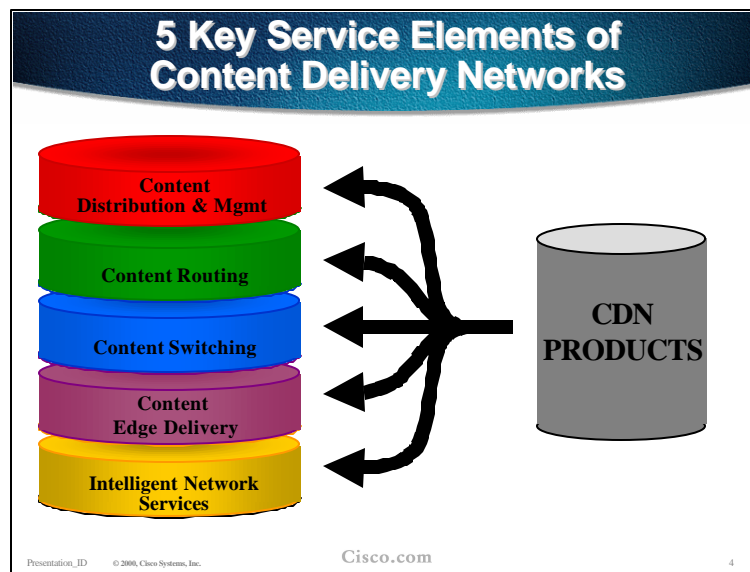
When a company is ready to debut and launch their business keeping it up reliably is paramount. If while still in its infancy a company suffers a long term outage the burgeoning business might be crippled and lose critical customer base and more important customer confidence.

## ***iCommerce Network Services***

Now that the framework of an iCommerce transaction is understood, the next step is to understand the functions and components incorporated into an iCommerce design. Several connections are commonly involved in the process of a client executing a transaction within an iCommerce site as previously shown. Each of these connections will pass through a series of services to provide high availability and scalability within the iCommerce network design.

An iCommerce design consists of a multilayer model of advanced network services. At each layer, services are invoked to ensure proper and highly resilient handling of the customer transaction. While many point-product vendors offer partial solutions to this service requirement, Cisco along with its partners have developed an integrated suite of advanced network services allowing such resilient behavior. Using an end-to-end Cisco iCommerce solution, several advanced protocols and platforms work in conjunction to provide the desired service characteristic commonly referred to as 'content-awareness' or 'content-networking'. As each advanced network service plays a critical role in the overall iCommerce solution, they will be discussed in detail.

## **Cisco Content Delivery Network**



A content delivery network is a globally coordinated network of devices designed to accelerate the delivery of information over Internet infrastructure. By leveraging services in the core IP network and layer 4-7 content-aware capabilities, companies are able to accelerate and improve the use of rich content, such as broadband streaming media, improve network performance and eliminate the stream of rich media on the infrastructure.

Because of the architecture of the Internet, many bottlenecks can exist at various locations that slow the delivery of information requested by a user. Since the Internet is really a series of interconnected networks, data typically traverses several networks before reaching the user. Whether bottlenecks occur at the last mile, the content source, or the "middle mile" between the two, the Internet lacks any kind of management system that can find the optimal route for content to travel. In addition, these performance problems can be

heightened during periods of heavy demand for specific content. Known as the "flash crowd" problem, this can often overwhelm the capacity of a Web site or communications links, leading to delays and site crashes.

Content delivery networks bypass potential sources of congestion by distributing the load across a collection of content engines located close to the viewing audience. Rich Web and multimedia content is replicated to the content engines and users are routed to an optimal content engine. The results are high-impact Web applications that translate into more powerful communications, premium revenue, and loyal audiences.

A complete end-to-end solution for building content delivery networks, includes:

- Content distribution and management
- Content routing
- Content switching
- Content edge delivery
- Intelligent network services

Until now, customers have had to combine products from various vendors or subscribe to a content delivery service to deliver rich media. Now, they are able to build their own content delivery networks—whether that solution includes caching, distributed applications, or live and on-demand content delivery—with end-to-end Cisco systems, including all five elements for content networking.

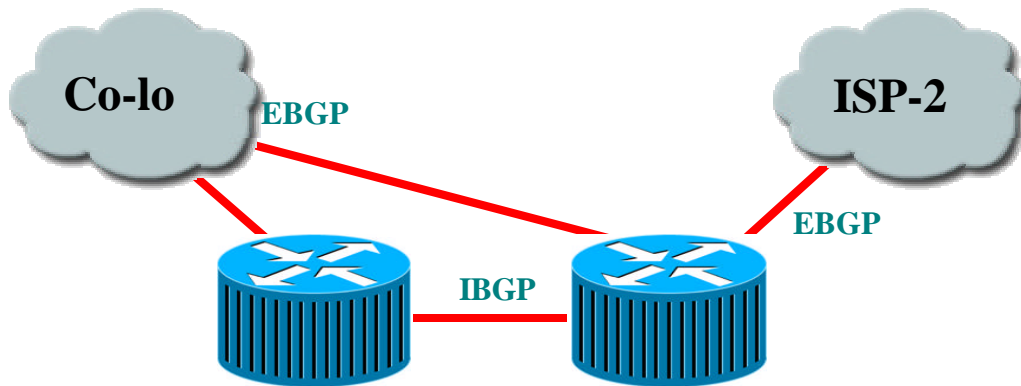
Cisco offers the industry's only complete solution for enterprises and service providers that want to transform their Internet or private IP networks into rich content delivery networks. Three key product components within this solution are:

- Content distribution managers, the administrative tool that manages the network policies and device settings, along with capturing usage and billing data
- Content engines, the "smart nodes" that reside at the edges of the Internet, either at the service provider's points of presence (POPs) or on the enterprise local-area network (LAN), and deliver content to the end user as requested. Content Engines provide caching, on-demand, and live delivery of media.
- Content routers, which route Web page "clicks" to an appropriate content engine based on load, geography, and other critical factors.
- Content switches, which provide server load balancing and content delivery, acceleration, and intelligence

CDNs enable companies to add a number of new services to their networks to improve performance and generate new revenues. Some of these services include:

- Caching of broad Internet content. Caches can be used in a number of applications to reduce bandwidth usage, improve Web site performance and reduce costs
- Accelerating static Web sites for Internet audiences
- Delivering rich streaming media to intranet, extranet, and Internet audiences
- Hosting of live broadband streaming media broadcasts
- Delivering high-impact e-business applications, such as e-learning, communications, and e-commerce applications

## Routed Front-End



Fronting any iCommerce network is a module consisting of ISP-edge routers. Now most people ask “Why do I need a Router in a collocation site?” Simply put, no network is infallible. There are planned and unplanned outages in the BEST networks in the world. The function of these routers is to propagate the ‘front-end’ IP network prefix into the internet so customers can reach the site. In addition, the ISP-edge routers participate in routing protocols, namely BGP, to route return traffic through the appropriate ISP in the case of a dual-homed iCommerce network. At the lowest end, an iCommerce network can be connected to one ISP or co-location provider however producing a large single point of failure. To increase resiliency, dual-connections can be established to a single provider to eliminate hardware failure risks. For ultimate protection, dual-homing to two ISPs or a co-location provider and an external ISP provides protection against a complete ISP failure. In the third case, full BGP routes are exchanged with both providers to ensure return traffic is sent to the ISP with the shortest path to the end customer. Using multiple connections to one or multiple providers offers a high availability solution while allowing for load balancing. There are 3 ISP connectivity options for the iCommerce network. In option 1 and option 2, default network statement(s) are all that are required from the iCommerce network perspective since either path lead to the same ISP. Only when connecting to two different ISPs is there an advantage to exchanging full BGP routes because return traffic can take the optimal return path through the optimal ISP.

Additionally, an edge router gives us the ability to provide security by blocking all ports except TCP-80 (HTTP) and TCP-443 (SSL). We can take this a step further with Firewall Feature Set on the routers. In fact, in 12.0.5T code and higher we support 59 Atomic (single packet) attacks of intrusion detection within IOS! We can use the shunning feature on a NetRanger to block multiple-packet attacks in conjunction with the edge routers by dynamically creating ACLs on the outbound, collocation edge, interface of the router. We can even use the router as a DRP agent in conjunction with Distributed Director to scale web sites in a geographically disparate topology. We will cover this more in a bit.

Finally, the front edge router can rate-limit nonessential traffic that exceeds prespecified thresholds in order to mitigate against (D)DoS attacks. This rate limiting can drop most undesired traffic when it exceeds a prespecified amount of the available bandwidth. The key is to correctly flag traffic as undesired. Common forms of DDoS attacks are ICMP floods, TCP SYN floods, or UDP floods. In an i-commerce environment, this type of traffic is fairly easy to categorize. Only when limiting a TCP SYN attack on port 80 (http) does an administrator run the risk of locking out legitimate users during an attack. Even then, it is better to temporarily lock out new legitimate users and retain routing and management connections, than to have the router overrun and lose all connectivity.

More sophisticated attacks use port 80 traffic with the ACK bit set so that the traffic appears to be legitimate Web transactions. It is unlikely that an administrator could properly categorize such an attack because acknowledged TCP communications are exactly the sort that you want to allow into your network.

One approach to limiting this sort of attack is to follow the guidelines outlined in RFC 1918 and RFC 2827. RFC 1918 specifies the networks that are reserved for private use and should never be seen across the public Internet. For inbound traffic on the edge router that is connected to the Internet, you could employ RFC 1918 and 2827 filtering to prevent unauthorized traffic from reaching the corporate network. When implemented at the ISP, this filtering prevents DDoS attack packets that use these addresses as sources from traversing the WAN link, potentially saving bandwidth during the attack. While this strategy does not directly prevent DDoS attacks, it does prevent such attacks from masking their source, which makes traceback to the attacking networks much easier.

Furthermore, because of the enormous security threat that they create, the router is configured to drop most fragmented packets that should not generally be seen for standard traffic types on the Internet. Any legitimate traffic lost because of this filtering is considered acceptable when compared to the risk of allowing such traffic.

### Security for the edge routers

Routers control access from every network to every network. They advertise networks and filter who can use them, and they are potentially a hacker's best friend. Router security is a critical element in any security deployment. By their nature, routers provide access and, therefore, you should secure them to reduce the likelihood that they are directly compromised. You can refer to other documents that have been written about router security. These documents provide more detail on the following subjects:

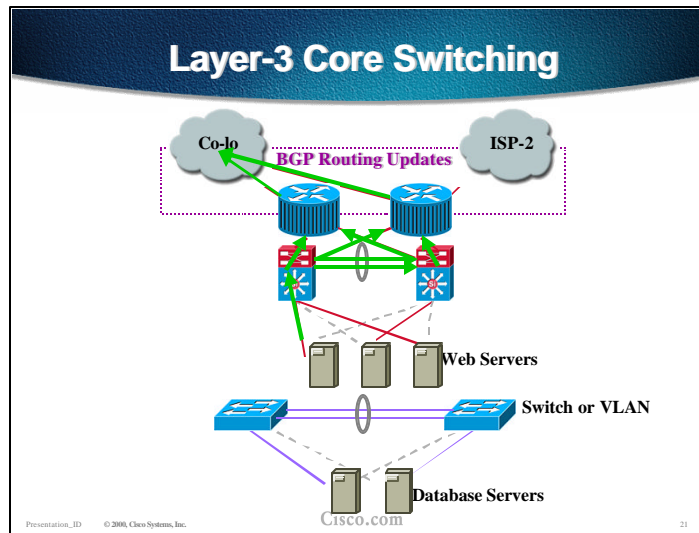
- Locking down telnet access to a router
- Locking down Simple Network Management Protocol (SNMP) access to a router
- Controlling access to a router through the use of Terminal Access Controller Access Control System Plus (TACACS+)
- Turning off unneeded services
- Logging at appropriate levels
- Authentication of routing updates

The most current document on router security is available at the following URL:  
<http://www.cisco.com/warp/customer/707/21.html>

## Summary - Edge Routers

- **BGP4 for load balancing and failover to secondary ISP**
- **QoS when allowing multiple services through ingress circuit**
- **Rate Limiting (CAR) - DOS attacks**
- **Security with firewall feature set**
- **Shunning of denial of service attacks when coupled with Cisco Secure IDS**

## Core Switching



The core switches will offer connectivity for the servers and most other devices in the co-location cage. Adding L3 to the core switching solution is imperative in terms of resilience. In conjunction with some of the current and future developments in server load balancing as a native service on the MSFC it makes sense to deploy L3 switching wherever cost effective for the client. For high availability, the servers should be dual-homed to two core switches. The two core switches would have an uplink to both core routers. The following are some additional capabilities that Cisco can offer in its Layer-3 switching offerings:

### Hot Standby Routing Protocol

HSRP provides protection against a default gateway failure. In a infrastructure with a L3 core switch, the L3 address of the core switch will be the default gateway for the web servers connected to the switch. There are some HSRP features that allow us to truly scale the traffic. First off the “Standby Track” feature allows us to monitor the state of the upstream FE interface. If it goes down then the virtual address will move to the second core switch so that multiple hops do not have to be incurred. Cisco also has support for M-HSRP or Multiple Host Standby Router Protocol. Since Web Farms are statically assigned rather than DHCP we can have some of the hosts have one default gateway, and thus point to one virtual address that exists on one core switch. Meanwhile the other half of the hosts have a default gateway that points to the virtual address that exists on the other core switch. This allows us to utilize both core switches and have them back each other up. Used in conjunction with failover NIC architectures this can become a powerful solution.

### Multimodule EtherChannel

This is quite simply EtherChannel that spans line cards. No longer are we constrained to using contiguous ports on a line card. This technology allows us to use Port 1 and 2 on the Supervisor Engine in a bundle with ports 7 and 8 from an eight port GE module as a single, logical pipe. This buys us the ability to avoid a spanning tree outage in the event of a line card failure. So we could literally reset a blade on the 6500, lose 50% of our available bandwidth between switches and NOT suffer a five to forty-five second systemic outage because of a Spanning Tree recalculation.

## E-SPAN and R-SPAN

The Catalyst 6500 features can now support multiple SPAN sessions simultaneously. This allows us to simultaneously run Cisco Secure IDS, a Switchprobe, and a Sniffer. We can support up to 4 sessions simultaneously so we still have one port available for connecting to a for web server accounting and performance verification. R-SPAN is the really exciting feature though. R-SPAN, available in CatOS 5.3.1CSX allows us to remotely monitor a port. So across the four of eight port GEC connection we can mirror a FE port. One Sniffer can be used to debug the entire network.

Determining what switching platform to deploy at an early stage development is oftentimes a tough one. For scalability reasons the Catalyst 6500 platform is non-pareil. However, the client might only have 4-10 web servers and a handful of database, application, and management servers. A switch commonly deployed here is a Catalyst 4000. The Layer-3 switching blade that comes with 32 10/100 ports and L3 services the makes Catalyst 4000 a viable platform, plus the flexibility of the 4K allows us to redeploy it when L3 services, server load balancing, and port density requirements start demanding the Catalyst 6500 family.

## Private VLANs

A Private VLAN is a layer 2 network structure not uncommon to but rather an extension of the common VLAN. Within a Private VLAN exists three separate port designations. Each port designation has its own unique set of rules which regulate a connected endpoint's ability to communicate with other endpoints connected to ports within the Private VLAN. The three port designations are promiscuous, isolated, and community.

An endpoint connected to a promiscuous port has the ability to communicate with any endpoint within the Private VLAN. Multiple promiscuous ports may be defined within a single Private VLAN. Within the previously mentioned hosting environment, the layer 3 switch default gateways are commonly connected to promiscuous ports.

Isolated ports are typically used for those endpoints that only require access to a limited number of endpoints. An endpoint connected to an isolated port will only possess the ability communicate with those endpoints connected to promiscuous ports. Endpoints connected to adjacent isolated ports cannot communicate. Within a web hosting environment, isolated ports are reserved for hosts only requiring access to default gateways or load balancers.

### Security for the core switch layer


Like routers, switches (both Layer 2 and Layer 3) have their own set of security considerations. Unlike routers, not as much public information is available about the security risks in switches and what can be done to mitigate those risks. You should take the following precautions:

- Ports without any need to trunk, should have any trunk settings set to off, as opposed to auto. This prevents a host from becoming a trunk port and receiving all traffic that would normally reside on a trunk port.
- Make sure that trunk ports use a virtual LAN (VLAN) number not used anywhere else in the switch. This prevents packets tagged with the same VLAN as the trunk port from reaching another VLAN without crossing a Layer 3 device. For more information, refer to the following URL:  
<http://www.sans.org/newlook/resources/IDFAQ/vlan.htm>
- Set all unused ports on a switch to a VLAN that has no Layer 3 connectivity. Better yet, disable any port that is not needed. This prevents hackers from plugging in to unused ports and communicating with the rest of the network.

- Avoid using VLANs as the sole method of securing access between two subnets. The capability for human error, combined with understanding that VLANs and VLAN tagging protocols were not designed with security in mind, makes their use in sensitive environments inadvisable. When VLANs are needed in security deployments, be sure to pay close attention to the configurations and guidelines mentioned above.


Within an existing VLAN, private VLANs provide some added security to specific network applications. As mentioned previously, Private VLANs work by limiting which ports within a VLAN can communicate with other ports in the same VLAN. Isolated ports within a VLAN can communicate only with promiscuous ports. Community ports can communicate only with other members of the same community and promiscuous ports. Promiscuous ports can communicate with any port. This is an effective way to mitigate the effects of a single compromised host. Consider a standard public services segment with a Web, FTP, and Domain Name System (DNS) server. If the DNS server is compromised, a hacker can pursue the other two hosts without passing back through the firewall. If private VLANs are deployed, once one system is compromised, it cannot communicate with the other systems. The only targets a hacker can pursue are hosts on the other side of the firewall.

### Layer 3 Core Switching Options




**Catalyst  
2948G-L3**

**Fixed Switch**



**Catalyst  
4000-L3**

**Mid-Level Chassis  
Switch**



**Catalyst  
6500**

**Flexible Chassis Switch**

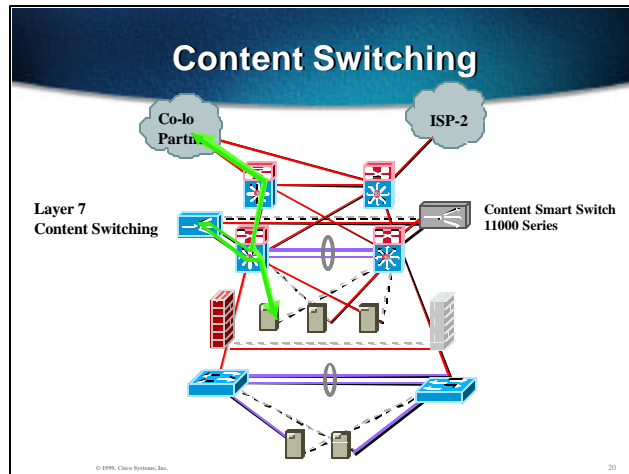
Presentation\_ID © 2000 Cisco Systems, Inc. Cisco.com 24

### Summary - Features on Core Switching

- Multimodule EtherChannel on 6500 so we avoid systemic Spanning Tree outages
- Integrated L3 switching enables "full-mesh" design architecture
- Multiple SPAN sessions allows-
  - Intrusion Detection Systems
  - Sniffer analysis
  - RMON probes
- Private VLANs allow-
  - Conservation of IP address
  - Additional Security

Presentation\_ID © 2000 Cisco Systems, Inc. Cisco.com 25


## Content Switching




A content switch, also known as a load balancer, is the device that will decide which servers are best suited to serve up content for a specific content request. Unlike traditional Layer 3/4 switch vendors that have only recently started to enhance their products with basic limited cookie and Universal Resource Locator (URL) filtering software capabilities, Cisco Web NS software, part of the Cisco 11000 series content switches, was designed from the ground up to provide “full-inspection” URL and cookie switching for highly optimized Web commerce sites.

### Content Switches - Layer 7

- Scalable product family
  - CSS 11050 - 8 FE, 1 GBE
  - CSS 11150 - 12-16 FE, 2 GBE
  - CSS 11800 - 64 FE, 32 GBE
  - Up to 11 billion hits/day
- Smart
  - Full content switching using content, application, network and user info
  - Flash Crowd support
  - Intelligent Cache Engine Bypass
- Fast
  - Up to 130,000 HTTP transactions per second



Content Switch  
11000 Family



© 2000 Cisco Systems, Inc. Cisco.com

At the heart of Cisco CSS 11000 series content services switches, first introduced in April 1998, is specialized Cisco Web NS software that is designed to allow customers to optimize the use of the Web for e-commerce and content delivery. Only the Cisco content switching architecture was purposely built to enable the construction of a new Web network services layer—a virtual overlay on top of the existing Internet infrastructure. This new layer seamlessly connects a Web browser with the best of a multiplicity of servers to handle each request for specific content, based on its full URL, at a given moment in time—ensuring a consistently positive Web-site user experience.

Running on the entire family of Cisco CSS 11000 series switch platforms, Cisco Web NS software delivers the industry’s most comprehensive and highest-performance URL—and cookie-based intelligent switching. All business-critical Web sites benefit from the rich feature set in Cisco Web NS that enables Web-site

security, content and transaction assurance, scalability, and content delivery. In addition, Cisco Web NS gives Web hosts and service providers the ability to offer high-value services to their customers—from shared to managed Web hosting to content delivery and distribution services—maximizing margins and customer satisfaction.

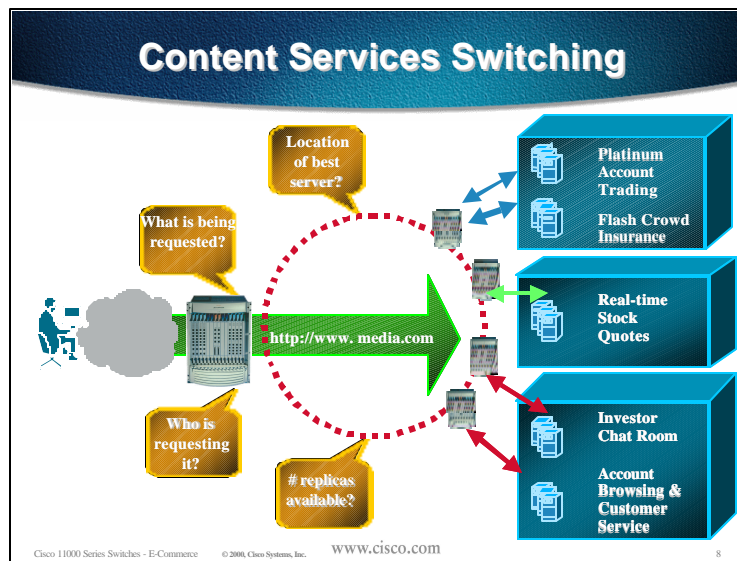
Cisco Web NS software helps businesses put their Web site to work by solving specific problems that impede e-business. Today's Web "business killers" include increased traffic and sudden "flash crowds" that swamp Web sites. For example, denial-of-service (DoS) attacks and lost shopping carts can lead to lost revenue, and slow downloads and cryptic errors often cause customers to become frustrated and leave Web sites. In addition, "content-blind" traffic routing can yield poor server utilization, and site outages can cause a Web business' brand and valuation to suffer. With Cisco Web NS, customers can eliminate these Web "business killers" to ensure that their e-commerce and content-delivery businesses are successful.

## Cisco Web NS

With the industry's most comprehensive URL- and cookie-based switching, Cisco Web NS software lets network managers tailor customer or content-specific service agreements, offer premium services for preferred customers, and deploy content-delivery services for streaming audio and video, distance learning, and Internet audio and video broadcasting. Support for "sticky" connections based on IP address, Secure Sockets Layer (SSL) session ID, and cookies ensures reliability and security for e-commerce transactions. The unique Cisco content-replication technology enables dynamic expansion of site capacity in response to sudden "flash crowds" for "hot" content, or seasonal peaks in traffic that can overwhelm servers.

Running Cisco Web NS software, with patented content switching technology, Cisco CSS 11000 services switches know:

Who the customer is, based on full visibility to the user cookie located anywhere within the HTTP header  
What information or transaction the customer is requesting  
Where best to service the customer from anywhere within a globally distributed Web infrastructure, based on comprehensive and current information on network, application, and server conditions



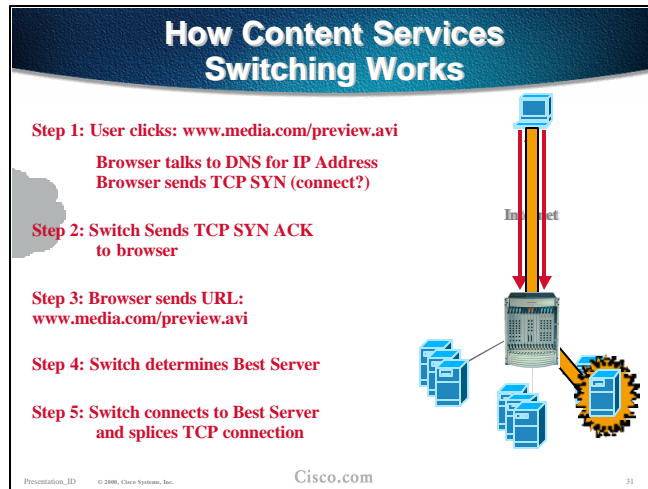
With this innovative software, Cisco CSS 11000 series switches can access information deep in the TCP and Hypertext Transfer Protocol (HTTP) headers, including the complete URL and "mobile" cookies that change location within the header between requests. This information is used to enable advanced load balancing, routing of requests, security (DoS and access control), priority access for important customers, and sticky connections. Cisco CSS switches were designed for name-based switching and are the only switches to use the entire URL and cookie to select the best site and server for the customer's inquiry or purchase at any given moment.

## Cisco Web NS Technologies

Cisco Web NS is based on a foundation of five core technologies, which together enable a set of core services for content delivery and distribution, e-commerce, as well as Web and application hosting. This section explores the key technology components—user and content policies, content delivery, dynamic capacity scaling, security, and resource assurance—that make up Cisco Web NS.

### User and Content Policies

Within Cisco Web NS are three "core" technologies that enable a variety of strategic services based on the individual user requesting content or services, and the specific content the user is seeking:



- *Delayed binding*—Current load-balancing solutions, including Layer 4 switches, route incoming packets based on the destination network address. The content request is routed to a server without an understanding of what content is being requested. Unlike traditional switching solutions, Cisco Web NS selects the site and server for a specific user request only after receiving the specific HTTP request. Called delayed binding, this process involves the switch intercepting the initial connection, but delaying the actual connection to the destination server. Instead, the switch sends the acknowledgment to the client that triggers the browser to transmit the HTTP request header, which contains the URL and the cookie.
- *URL and cookie switching with “full inspection”*—Cisco Web NS uses the information in the HTTP request to route incoming Web requests flows to the best site and servers. By looking deep into the HTTP payload down to the URL and cookie, a Web switch “knows” what content is being requested, and who is requesting it. With this knowledge, the switch can select the best site or server for that request based on access control, server response time, application availability, and the level of priority required for specific content or users.
- *Sticky connections*—In any authenticated Web application, it is necessary to provide a persistent connection between a browser (the user) and the Web or database server to which it is connected. Examples of these applications include shopping baskets, financial transactions, and some forms of interactive gaming. Because HTTP does not carry any state information for these applications, it is important for the browser to be mapped to the same server for each HTTP request until a user’s transaction is complete. This ensures that the user is not load balanced in midsession to a different server and forced to log in again. Cisco CSS 11000 series switches support sticky connections based not only on IP address/TCP port and SSL session ID, but also on the user’s cookie embedded in the HTTP header.

#### Content Delivery

In addition to routing Web requests to the best server, Cisco Web NS also optimizes the delivery of the requested content back to the user. This is accomplished through three methods:

- A purpose-built architecture that switches traffic at wire speed after the flow is established
- The ability to monitor both the control and data channels for complex Web flows such as port-based File Transfer Protocol (FTP) and streaming media

- The ability to route a request to the site and ensure the response is delivered over the shortest Internet path
- Purpose-built hardware architecture—Powered by Cisco Web NS, Cisco CSS 11000 series switches use a balance of centralized and distributed resources for the two key parts of the job: flow setup and flow forwarding. The only Web switch built from the ground up for Web traffic, Cisco CSS 11000 series switches combine substantial centralized symmetric multiprocessing and memory resources needed for flow setup and “rules processing” with distributed application-specific integrated circuit (ASIC)-based processors for wire-speed forwarding of Web flows, all in a single platform.
- Bidirectional load balancing—Conventional load balancers and Layer 4 switches perform simple TCP session distribution among Web servers. Cisco CSS 11000 Series Switches optimize complex server-initiated flows (inbound and outbound) by using layer 5-7 load distribution metrics. This offers greater control over Web traffic and back-end servers for complex flows (for example, streaming media).
- Network Address Translation (NAT) peering—channels on different ports. The Cisco CSS 11000 series switch sets up the initial request, selecting the best site and server for the content being requested, and then maintains security and prioritization policies over both the control and data channel for the duration of the flow.

#### **Best-Path Distributed Content Delivery**

Cisco CSS 11000 series switches use NAT peering to direct requests to the best site with the requested content based on URL or file type, geographic proximity, and server/network loads, avoiding the limitations of DNS-based site selection and the overhead HTTP redirect. If a Cisco CSS 11000 series switch receives a request for content and the cache or server hosting that content is down or busy, the switch uses information communicated among its peers (that is, other Cisco switches) at other points of presence (POPs) to locate the best site and server. When the best destination site is determined, the request is directed to that site and the source IP address is translated to the “new” IP address where the content resides. Cisco NAT peering preserves the original source IP address from the client so that the new site can send the response directly back to the client. NAT peering acts as “triangulation protocol,” allowing the response to be delivered directly to the user over the shortest Internet path.

#### **Dynamic Capacity Scaling**

To be competitive, today’s Web businesses must be able to provide consistently high levels of service, regardless of rapidly increasing daily traffic, seasonal peaks in usage, and sudden traffic surges for “hot content.” Cisco Web NS features “core” technology that enables dynamic scaling of site capacity without the need to “overbuild” server and network infrastructure for peak traffic loads or manual replication of popular content.

- Smart Content Replication with RPC Bypass—Some of the biggest challenges for a Web site include unpredictable traffic and flash crowds caused by suddenly “hot” content, such as the Starr Report. Conventional load balancers distribute traffic by using information in the TCP header, including IP addresses and TCP port numbers, to direct TCP sessions to a particular server. Load balancers generally support NAT, so in the process of setting up a TCP session they can remap IP addresses and TCP ports to direct packets to the target server. But they have no concept of what content is being requested, and as a result, require all content to be replicated between load-balanced servers. In addition, they cannot explicitly track content requests or detect, anticipate, or replicate hot content in response to flash crowds.

Cisco CSS 11000 series switches have the unique ability to dynamically track content requests and identify and replicate hot content to an overflow Web server or cache. Called content replication, this service permits premium service-level agreements (SLAs) based on Web response time guarantees that ensure capacity is available no matter how hard a Web site is hit.

## Content Smart Switch Security

Cisco Web NS includes FlowWall and trade security, comprising four core technologies that enable both site- and system-level security:

- Access control—Cisco provides strong security without compromising site stateful access control; using any combination of source address, destination address, TCP port, or URL allows comprehensive security policies to be built.
- DoS prevention—Cisco prevents common DoS attacks by validating every transaction at initial flow setup time and defending against connection-based DoS attacks and other malicious or abnormal connections.
- Firewall load balancing—For scalable firewall security, Cisco intelligently directs traffic across multiple firewalls, eliminating performance bottlenecks and single points of failure. Firewall load balancing eliminates system downtime that results when a firewall fails or becomes overloaded—breaking Internet connections and disrupting e-commerce purchases or other mission-critical transactions.

Cisco CSS 11000 series switches ensure that all traffic for a given Web flow between a pair of IP addresses, in either direction, will traverse the same firewall. This is accomplished by configuring static IP routes on each switch with the IP address of the adjacent firewall port, the IP address of the remote firewall port, as well as the IP address of the remote Web switch port. In addition, all ports on each side of the load-balanced firewall are utilizing a different IP subnet address. In this configuration, the firewalls cannot be configured to perform NAT—rather all NAT processing is done on the Web Switch on the private network side of the firewall.

Each of the Cisco CSS 11000 series switches secures the paths through the firewall by utilizing health checks to verify that each stage on the paths is responding, including its adjacent firewall port, its remote firewall port, and its remote Web switch port. If any part of the path becomes unavailable, the Web switches reroute all traffic through the surviving paths. Cisco CSS 11000 series switches can be configured to send an alarm or log an event to notify the system administrators of faults. If firewalls are configured to share session-state information and have physical connections to each other, then user sessions will not be interrupted in the case of the failure in any path or firewall.

For additional redundancy, multiple Cisco CSS 11000 series switches can be deployed on both sides of the firewalls being load balanced in an active/passive configuration. This eliminates single points of failure for both firewalls and Web switches. In this configuration, the passive Web switches monitor the health of the active Cisco CSS 11000 series switches, and upon detecting a switch or an uplink failure, the Virtual Router Redundancy Protocol (VRRP) is used to transfer control to the passive switch pair.

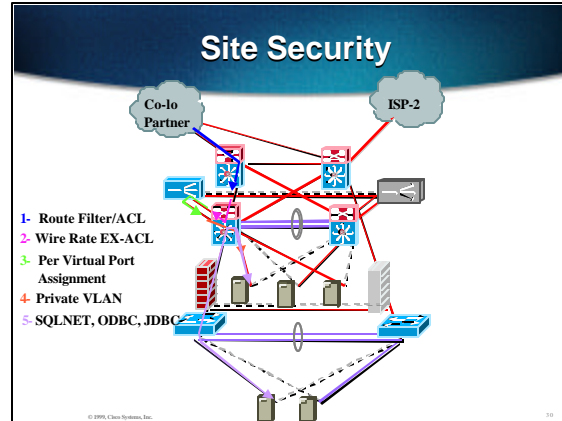
- NAT—Described in RFC 1631, NAT provides translation from public IP addresses and TCP ports to private IP addresses (10.xxx.xxx.xxx) and TCP ports. NAT is important for managing existing allocated IP addresses, a scarce resource, and reducing the need to acquire additional addresses as the network expands. In addition, NAT hides the IP addresses for all devices located behind the Cisco CSS 11000 series switches, eliminating the ability for hackers to attack servers directly by using explicit IP addresses.

## Resource Assurance

Cisco CSS Web switches choose the best currently available resources to fulfill each incoming HTTP request. Cisco CSS 11000 series switches parse the full URL to identify the specific content that is being requested. The switch then determines the best site and server to fulfill the request based on site and server loads, application and content availability, flow duration, and geographic proximity. For local servers, the Web switch monitors server response time for successful connections and creates a normalized load factor,

ensuring the distribution of requests to the best server. In addition, the Web switch periodically performs health checks, ensuring availability for Web servers or back-end servers, and specific applications or content, eliminating error messages and nonresponding requests. In a distributed Web site, Web switches communicate this local knowledge to other switches in the network, so that every Web switch is aware of the load, availability, and proximity of every other Web site. This combination ensures that every request is delivered to the best possible site and server for the content being requested.



## Security with Firewalling and Intrusion Detection



The primary function of an iCommerce network is to conduct financial transactions. Based on this fact, security is of utmost importance regardless of the overall network size. There are varying degrees of network security that apply to an iCommerce implementation depending on comfort levels. Each increasing level of security offers additional granular inspection of IP packets and the connections and transactions to which they belong.

The most basic form of network security is the extended access control list (ACL) within Cisco's IOS software. Using an extended ACL, one can define filters to deny packets matching a 5-tuple of source and destination IP address, TCP/UDP protocol, and source and destination port numbers. Extended ACLs have no knowledge of connection state or packet trends. Extended ACLs are typically used but not as the only form of defense. Using Cisco's IOS firewall feature set (FFS) provides for a moderate level of connection state awareness. In addition, IOS FFS allows for the recognition of certain network attacks including the dreaded TCP-SYN attack.

Adding Cisco's PIX firewall to the design allows for maximum state awareness, control, and accounting available. In addition, using in pairs, the PIX firewall has the ability to perform stateful fail-over should a

 <h3>PIX Firewall Appliance</h3> <ul style="list-style-type: none"> <li>• Stateful, connection-oriented protection</li> <li>• True network address translation</li> <li>• High performance without introducing packet delay</li> <li>• Low cost of ownership - no OS maintenance</li> <li>• Highly reliable - runs from Flash memory; no disk</li> </ul>  <p>© 2000 Cisco Systems, Inc. Cisco.com</p>	<h3>Adaptive Security Algorithm (ASA)</h3> <ul style="list-style-type: none"> <li>• Provides "stateful" connection policy</li> <li>• Connections allowed out—allows return session backflow; incoming connections must be explicitly enabled</li> <li>• Initial TCP sequence number randomized</li> <li>• Tracks source and destination ports + addresses, TCP sequences, and additional TCP flags</li> <li>• Access Control List (ACL) policy support</li> <li>• UDP + TCP session state <ul style="list-style-type: none"> <li>TCP—FIN bit</li> <li>UDP—One minute default timer (except for DNS)</li> </ul> </li> </ul> <p>© 2000 Cisco Systems, Inc. Cisco.com</p>
---	--

single PIX fail. PIX failover provides a mechanism for HW/SW redundancy by allowing two identical PIXes to serve the same functionality in case one failed in an unattended environment. Normally, one PIX is considered the "Active" unit while the other is the "Standby" unit. The Active unit actively performs normal network functions while the Standby unit only monitors, ready to take control should the active unit

fail to perform its functionality.

The two units each have a presence on the network. The Active unit uses the System IP address and the MAC addresses of the Primary unit (the Primary Unit is determined by the unit that has the end of the failover cable marked "Primary" plugged into it). The Standby unit uses the Failover IP address and the MAC addresses of the Secondary unit. If a switchover occurs, the units swap the IP address and MAC addresses they are using so as to replace each other's presence on the network. This action is invisible to the network. The IP to MAC address relationships remain exactly the same so no ARP tables in the network need to time out or be changed.

No other piece of network equipment needs to know about the redundancy or that a switchover occurred. Note that the System IP and the Failover IP addresses must be on the same subnet. There can not be a router between the two units.

## Failover Cable

The failover cable is the only additional HW needed to support PIX failover. It is an essential part of failover. The failover cable is a modified RS-232 serial link cable, the speed setting is 9600 baud. It provides the unit identification, Primary or Secondary (the two units supposed to be identical, without this, identify Syslog message from which unit is a problem), the power status of the other unit, served as communication link for various failover communications between the two units.

## Configuration Replication

The two PIX units must be configured exactly the same (and running the same SW release). This is easily accomplished as configuration replication occurs over the failover cable from the Active unit to the Standby unit in three ways:

- When the Standby unit completes its initial boot-up the Active unit will replicate its entire configuration to the Standby unit.
- As commands are entered on the Active unit they are sent across the Failover Cable to the Standby unit.
- By entering command "write standby" on the Active unit to force the entire configuration in memory to the Standby unit.

The configuration replication only replicates configuration from memory to memory, after replication, a "write mem" is needed to write the configuration into the flash memory. Since the failover cable is used, the replication can take a long time to complete with large configuration. If a switchover happened during the replication, the new Active PIX will have a partial configuration. The unit will reboot itself to recover the configuration from the flash or re-sync by the other unit. When the replication started, "Sync Started", and when complete "Sync Completed" will display on the PIX console. Also, during the replication, the console input is disabled.

## Failover Interface Tests

One important factor to know is how the PIX firewall determines that it is in a failed state. This is critical in designing the collocation site because if the network does not allow a fast PIX firewall failover then the entire site is down for at least a minute.

If a failure is due to a condition other than a loss of power on the other unit, failover will begin a series of tests to determine which unit is failed. This series of tests will begin when hello messages are not heard for two consecutive 15-second intervals. Hello messages are sent over both network interfaces and the failover cable.

The purpose of these tests is to generate network traffic in order to determine which (if either) unit has failed. At the start of each test, each unit clears its received packet count for its interfaces. At the conclusion of each test, each unit looks to see if it has received any traffic. If it has, the interface is considered operational. If one unit receives traffic for a test and the other unit does not, the unit that received no traffic is considered failed. If neither unit has received traffic, then go to the next test.

- Link Up/Down test---This is a test of the NIC card itself. If an interface card is not plugged into an operational network, it is considered failed (for example, the hub or switch is failed, has a failed port, or a cable is unplugged).
- Network Activity test---This is a received network activity test. The unit will count all received packets for up to 5 seconds. If any packets are received at any time during this interval the interface is considered operational and testing stops. If no traffic is received, the ARP test begins.
- ARP test---The ARP test consists of reading the unit's ARP cache for the 10 most recently acquired entries. One at a time the unit sends ARP requests to these machines attempting to stimulate network traffic. After each request the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next machine. If at the end of the list no traffic has been received, the ping test begins.
- Broadcast Ping test---The ping test consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval the interface is considered operational and testing stops. If no traffic is received the testing starts over again with the ARP test.

## Fault Detection

Fault detection in the PIX for a failover scenario is based on the following criterion:


- Failover hello packets are received on each interface. If hello packets are not heard for two consecutive 15 second intervals, the interface will be tested to determine which unit is at fault.
- Cable errors. The cable is wired so that each unit can distinguish between a power failure in the other unit, and an unplugged cable. If the Standby unit detects that the Active unit is powered down (or resets) it will take active control. If the cable is unplugged, a SYSLOG is generated but no switching occurs. An exception to this is at boot-up, at which point an unplugged cable will force the unit active. If both units are powered on without the failover cable installed they will both become active creating a duplicate IP address conflict on your network. The failover cable must be installed for failover to work correctly.
- Failover communication. The two units share information every 15 seconds. If the Standby unit does not hear from the Active unit in two communication attempts (and the cable status is OK) the Standby unit will take over as active.

Finally, the most granular options for network security is a Cisco Secure IDS probe.

### Cisco Secure IDS

**Cisco Secure IDS (NetRanger)**

- Single and Dual Pentium II 400MHz
- 4 GB Hard Disk
- Hardened OS
- Two NICs



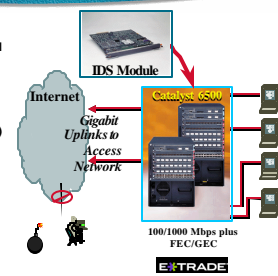
**Command & Control - 10/100BaseT**

**Sniffing/Monitoring - Ethernet, Fast Ethernet, Token Ring and FDDI**

Cisco.com

### Delivering Advanced Security with Integrated Intrusion Detection

- Integrated Intrusion Detection: Initial performance ~100Mbps
- No impact on Performance ("promiscuous" mode)
- Traffic Filtering by vlans (Example: Only inspect HTTP traffic on VLAN2)
- Dynamic Response Actions of "Shunning" & "TCP RSTs"
- Possible to install multiple modules and have each module monitor different vlans
- Full signature set (more than 300+)
  - Updated every other month
  - File download from CCO



100/1000 Mbps plus FEC/GEC

E-TRADE eTV

Cisco.com

Applications are coded by human beings (mostly) and, as such, are subject to numerous errors. These errors can be benign—for example, an error that causes your document to print incorrectly—or malignant—for example, an error that makes the credit card numbers on your database server available via anonymous FTP. It is the malignant problems, as well as other more general security vulnerabilities, that intrusion detection systems (IDSs) aim to detect. Intrusion detection acts like an alarm system in the physical world. When an IDS detects something that it considers an attack, it can either take corrective action itself or notify a management system for actions by the administrator. Some systems are more or less equipped to respond and prevent such an attack. Host-based intrusion detection can work by intercepting OS and application calls on an individual host. It can also operate by after-the-fact analysis of local log files. The former approach allows better attack prevention, while the latter approach dictates a more passive attack response role. Because of the specificity of their role, host-based IDS (HIDS) systems are often better at preventing specific attacks than Network IDS (NIDS), which usually only issue an alert upon discovery of an attack. However, that specificity causes a loss of perspective to the overall network.

This is where NIDS excels. Cisco recommends a combination of the two systems—HIDS on critical hosts and NIDS looking over the whole network—for a complete intrusion detection system.

Once deployed, you must tune an IDS implementation to increase its effectiveness and remove “false-positives.” False-positives are defined as alarms caused by legitimate traffic or activity. False-negatives are attacks that the IDS system fails to see. Once the IDS is tuned, you can configure it more specifically as to its threat mitigation role. As was mentioned above, you should configure HIDS to stop most valid threats at the host level because it is well prepared to determine that certain activity is indeed a threat.

When deciding on mitigation roles for NIDS there are two primary options:

The first option, and potentially the most damaging if improperly deployed, is to “shun” traffic through the addition of access control filters on routers. When a NIDS detects an attack from a particular host over a particular protocol, it can block that host from coming into the network for a predetermined amount of time. While on the surface this might seem like a great aid to a security administrator, in reality it must be very carefully implemented, if at all. The first problem is that of spoofed addresses. If traffic that matches an attack is seen by the NIDS, and that particular alarm triggers a shun response, the NIDS will deploy the access list to the device. However, if the attack that caused the alarm used a spoofed address, the NIDS has now locked out an address that never initiated an attack. If the IP address that the hacker used happens to be the IP address of a major ISP’s outbound HTTP proxy server, a huge number of users could be locked out. This by itself could be an interesting DoS threat in the hands of a creative hacker.

To mitigate the risks of shunning, you should generally use it only on TCP traffic, which is much more difficult to successfully spoof than UDP. Use it only in cases where the threat is real and the chance of the attack being a false positive is very low. However, in the interior of a network, many more options exist. With effectively deployed RFC 2827 filtering, spoofed traffic should be very limited. Also, because customers are not generally on the internal network, you can take a more restrictive stance against internally originated attack attempts. Another reason for this is that internal networks do not often have the same level of stateful filtering that edge connections possess. As such, IDS needs to be more heavily relied upon than in the external environment.

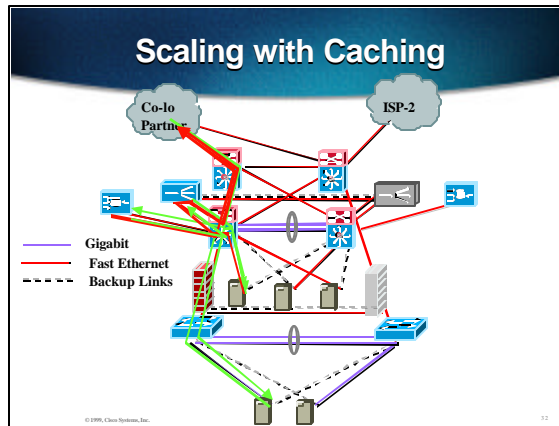
The second option for NIDS threat mitigation is the use of TCP resets. As the name implies, TCP resets operate only on TCP traffic and terminate an active attack by sending TCP reset messages to the attacking and attacked host. Because TCP traffic is more difficult to spoof, you should consider using TCP resets more often than shunning.

From a performance standpoint, NIDS observes packets on the wire. If packets are sent faster than the NIDS can process them, there is no degradation to the network because the NIDS does not sit directly in the flows of data. However, the NIDS will lose effectiveness and packets could be missed causing both false-negatives and false-positives. Be sure to avoid exceeding the capabilities of IDS so that you can get

their benefit. From a routing standpoint, IDS, like many state-aware engines, does not operate properly in an asymmetrically routed environment. Packets sent out from one set of routers and switches and returning through another will cause the IDS systems to see only half of the traffic, causing false-positives and negatives.

Although any combination of the security options can be used, it is recommended that at least extended ACLs and intrusion detection are include in the deployed design.

## Content Caching



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.

However, in a typical iCommerce environment, it is not feasible to locate content caches in close proximity to the user community. Content caching does however have a powerful application within an iCommerce environment. An iCommerce provider can deliver accelerated service 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 web servers. If the content that is being requested is cacheable, the cache engines will fill the request. When the cache cluster fulfills these requests, it off-loads traffic from the web servers thereby minimizing content download latency and increasing web server capacity. Therefore, once a particular piece of cacheable content is requested by a customer, it is cached so that successive requests are not directed repeatedly to a web server. Figure 9 shows the cache application within an iCommerce environment. Within this environment, the cache engine cluster only caches the content that is available on the local web servers. This arrangement is referred to as *Reverse Proxy Caching* function.




Figure 10. Web Cache Cluster Transaction

The key to deploying cache services is to ensure the resultant function will be transparent to the user community. In order to deliver transparent service, Cisco's icommerce architecture presents two techniques to accomplish this. Firstly, Cisco CSS 11000 series switches direct traffic based on "content policies" that can be configured to include, exclude (bypass), or block certain traffic from caches, creating opportunities for new managed services that improve site performance. Cisco CSS 11000 series switches can be configured to perform intelligent cache bypass for transparent caches for noncacheable content requests. When noncacheable (that is, dynamic) URLs are sent to the cache, significant latency is incurred while the cache determines that the requested content is not cacheable, fetches the content from the origin server, and then forwards the response back to the client. In this scenario, the cache essentially becomes a bottleneck, and cache resources are not efficiently used. The intelligent cache bypass capability of Cisco CSS 11000 series switches offloads caches from processing dynamic content requests, thereby improving cache performance by as much as 400 percent.

Cisco has also developed the Web Cache Communication Protocol (WCCP). WCCP is a protocol that runs

between Cisco IOS router and switch platforms and Cisco's 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 10. 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.

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 V2 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

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 Routing / Geographic Load Balancing*

Cisco offers a broad array of Content Routing (CR) technologies. These are also present on a wide-array of platforms.

### **Content Routing Fundamentals**

The goal of Content routing is to make an intelligent decision about where to send a users requests based on load, latency, presence of content or other metrics. The challenge we run into is that there are only two mechanisms or times when we can effectively redirect a user request; DNS Lookup and HTTP Session initiation. Throughout this section, we will use cnn.com as an example.

### **DNS Mode**

When a user is looking up the DNS “A” Record for a particular domain the process involves the user querying their statically configured or DHCP assigned DNS Server. This Client DNS Server is sometimes called a DNS-Proxy or D-Proxy because it queries the DNS hierarchy to learn the IP address for a Fully Qualified Domain Name (FQDN) and then returns the result of this lookup process to the end user. The end-user does not see this process and the authoritative DNS servers do not ever see the end-users IP address.

When we insert a Content Router (CR) into the process we do so by using the DNS Delegation Insertion Mechanism. The CR receives a request from the Client-DNS looking for a particular “A” Record: www.cnn.com It is very important to note that the CR needs to make a forwarding decision as to which site to send the end-user to based on only knowing the clients DNS server IP address. It cannot make an accurate forwarding decision to the end-users IP address in DNS Delegation Insertion. The behavior of the CR after this phase is based on the type of CR deployed and will be covered later in this document. One the end users DNS receives the resolved IP address of www.cnn.com it then passes this back to the end user and the end user’s browser or other application connects to the resolved IP address.

#### *Advantages of DNS Mode-*

- DNS mode is transparent to the end user and does not require ANY special configuration on the client.
- DNS mode works across multiple application types as long as they use a DNS lookup rather than static mapped IP address.
- Most browsers tested cache the DNS response for the duration of the browser session or cache a large number of DNS responses so that this in-effect provides a persistence mechanism to a particular site.

#### *Disadvantages of DNS Mode-*

- Makes a redirection decision to the clients-DNS and not the end-user.
- May be bypassed by some client DNS systems that cache responses regardless of TTL
- May be bypassed by users entering http://cnn.com rather than http://www.cnn.com

### **HTTP Redirect-Mode**

HTTP has built into it a mechanism for redirecting a user to a new location. This is referred to as the HTTP-Redirect or HTTP-302 (the HTTP return code for redirection). The CR is inserted by making it the Virtual IP Address (VIP) that is resolved by the DNS lookup process. HTTP-Redirection occurs after a user initiates a TCP-Session with the CR; goes through the three-way handshaking; and passes the first HTTP-Get to the CR. The CR now has visibility into the actual content being requested and can issue a HTTP-302 redirection back sending a client to the site that is closest and has the exact content being requested.

#### *Advantages of HTTP-Redirection-*

- Visibility into the content being requested
- Redirection made to the end user

*Disadvantages of HTTP-Redirection-*

- Some browsers still have problems with HTTP-Redirects
- Each site must have their own DNS Record thus the site-name must change after a redirection from `http://www.cnn.com` to `http://www-site1.cnn.com` or some other derivative.
- Bookmarking issues arise because of the above where a user can bookmark their browser to a particular site and not the global `http://www.cnn.com` thus bypassing the Content Routing system
- Only works for HTTP traffic

Generally most sites deploy a DNS mode Content Routing system because of the negatives associated with HTTP-Redirect. HTTP-Redirect is a good tool and can be used to alleviate certain problems, especially persistence problems, but DNS is the preferred mechanism for Content Routing today.

## Content Routing Site-Selection Methods

Cisco has a broad array of technologies available to provide the best decision about which site to choose for a particular user request.

### **Distributed Director 7200/ Content Router 7200**

The venerable Distributed Director and its rebranded counterpart the Content Router 7200 have a wide array of metrics available. Most are self-explanatory so we are only going to cover a subset of its functionality. The majority of these metrics use Director Response Protocol to trigger a specific action in a Cisco router infrastructure to gain an understanding of topology, latency, or routing metric and then use this information to make a selection decision.

### **BGP-External Metric**

The BGP-External metric relies on the edge or peering routers at each site to have an accurate “picture” of the Internet learned via BGP updates. It sends a query from the CR to a router at each of the edge sites asking for the “distance” to a particular IP address. Each router at the edge sites receiving this DRP-Query looks in its routing table. If it finds a match it responds back with the number of Autonomous Systems between itself and the destination IP address.

This metric works quite well with a low number of sites that are vastly distributed; i.e. a San Jose, Singapore, and Brussels deployment. It has problems as more than one or two sites are clustered in a particular geography. The reason I state this is because you can get from anywhere in the continental US to about anywhere else in one or two BGP AS-Hops. If the maximum granularity available to make a forwarding decision is “2” then any more than two sites will result in an inaccurate decision being made.

### **DRP-Round Trip Time**

The DRP-RTT metric receives a user request and then forwards a DRP-Query to each of the sites. This query tells the routers there to initiate a Round-Trip Response Time calculation between themselves and the Client-DNS server. This occurs by sending a TCP SYN-ACK on port 53 to the client-DNS. (This SYN-ACK will go through firewalls as well as we can be assured of it being an open port on any client-DNS implementation) The behavior of a TCP stack when it receives an unsolicited SYN-ACK is to issue a RST back to the source of the SYN-ACK. The routers start a timer when they issue the SYN-ACK and calculate the round-trip latency when the RST comes back. This time-value is sent back to the CR who evaluates all of the responses and picks the lowest one and responds back to the client-DNS with the resolved A-record for this name.

This metric involves a lot of steps and can take a while to execute completely and respond back to the client DNS. Because of this it may take two to three seconds to determine site selection. This is not a problem for a forty-five minute MPEG movie; but can be problematic for a six second HTML page load. The reality is users will tolerate a three second delay for a long term flow but when e-companies are striving for a six to eight second page load time a good portion of it can quickly be used up by determining which site to send a user request to before ever starting a TCP session and transferring any data. However, as stated before it is a VERY granular metric and is excellent for long-term traffic flows like streaming files, SNA, Telnet, etc.

### **Content Router 4400**

The CR4400 is a system internally developed and designed to address the requirements of an enterprise or service provider CDN that has between two and twenty locations on the Internet when running in DNS delegation mode. It can scale far beyond this number in edge-insertion mode.

#### *One-Way Race Metric (Boomerang)*

Boomerang is the name of the feature on the CR4400 that we use to determine site selection. There are two components to the CR4400 the Content Router and the Boomerang Agents. The CR is inserted via either insertion mechanism and the agents are deployed at each of the content hosting locations.

The Boomerang process starts by measuring the RTT between the CR and the agents that are deployed at each site. Using this RTT measurement the CR makes an estimation of one-way trip time between itself and the agents. When the CR receives a DNS-lookup it sends this A-record request to each of the agents. It does this by equalizing the distance between itself and the agents, delaying the forwarding of the record request so that they all arrive at approximately the same time.

RTT Site A to CR (140ms) / 2 = 70ms

RTT Site B to CR (60ms) / 2 = 30ms

Difference Between A and B = 40ms

Send Site A request immediately

Send Site B request in 40ms

When each site receives the request they formulate their own response to it that has either their own sites virtual IP address or, if the agent is on a Content Engine, the CE's own IP address as the resolved IP address for that particular FQDN. This response has a destination address of the client-DNS and a source address of the CR.

Since both Site A and Site B are responding at the same time a race-condition has been created. The first response to get to the DNS server is passed through to the end-user as the resolved IP address and the subsequent responses are dropped.

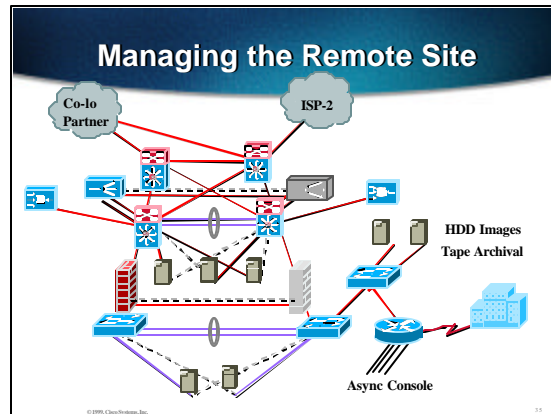
*Advantages of Boomerang Process in DNS Delegation mode-*

- Fast and accurate response to a request
- Never sacrifices a user request to "learn" site proximity
- No delay in deciding which site a client should go to
- Using a One-Way metric works extremely well in third-world countries with satellite Internet access
- Simple to configure and easy to understand.

*Disadvantages of Boomerang Process in DNS Delegation mode-*

- Volume DNS traffic could create additional, unnecessary network load. This is the reason for the recommended deployment not exceeding 20 sites.
- Multiple unsolicited responses may flag DNS server as erroneous.
- Vigorous Anti-Spoofing or RPF filters may impact viability.

## Managing the E-Commerce Site



One of the most common questions asked when the decision to co-locate or host internally comes up is, “How do I manage a remote site?” The biggest fear is that at a co-location site you can’t debug problems, see into the network, use a network analyzer (Sniffer, W&G, etc), or use network management tools. Some of these fears are valid. For instance, in tightening security down tight on all of the equipment you want to shut off Telnet, FTP, SNMP, etc.

There are a couple of solutions though that can allow the remote site to be manageable.

1. On the router that connects the client’s office to the co-location facility install a 32 port Asynchronous module. Use this to connect to all of the networking equipment’s Console port. Secure the Telnet on this router so only IS and other authorized personnel can Telnet to it. Then use CiscoSecure to validate and authorize everyone on a per user basis. This out-of-band console port access will allow all of the equipment to be manageable and problems to be debugged. Additionally, most clients oftentimes locate a modem on one of these ports and turn the modem off and unplug the phone line. If problems arise a phone call to the co-location facility can have their staff plug in the modem and power it on for phone access.
2. Connect a network analyzer to the core switches that the WWW servers connect to. Set the port as a SPAN port. Connect the other switches in the network to the Core switch on different VLANs that are unused by normal traffic flows. Disable the ports they connect into on the core switch. This way if we need to see the traffic on a remote switch we can mirror it’s ports to the interface connecting to the core switch. We can then SPAN the entire VLAN to the network analyzer.
3. Install a third interface into UNIX based “pizza-box” web servers. If you are using an NT based server install a device similar to the Compaq Insight Manager boards. Connect the third interface or Insight Manager modules into a separate LAN switch. Keep a server with a Gigabit Ethernet interface on this same switch. Use it to store Ghost filesystem images in a PC based environment or use Jumpstart in a Solaris/Unix environment. This allows us to do the following-
  - When a new image is completed copy it to the Ghost server
  - Using the Async connections console into the Devices
  - Remove a Web Server from production
  - Open the Insight Manager
  - Partition the hard drive and Format a Web Server
  - Load a new image onto the Web Server from a local Ghost Image Server
  - Bring the Web Server up under a “test” IP address
  - Verify the image loaded properly and the site is functioning

- Bring the web server back into the production pool

This procedure is VERY powerful in that we can now ensure control of every access of a remote server pool.

4. The same method of bringing Network Analyzer traffic, in 2 above, back to the core switch can also be used to distribute NetSonar Active Audit services to the DMZ, WWW LAN, Backend LAN, etc. Simply plug it in a port and leave the port disabled until you need to bring it online. Then move it's port into a VLAN that the servers and routers exist on and run the audit process.

## Connecting to the Corporate Office

Managing the E-commerce site in a co-location environment requires a connection back to the corporate office. T-1 or NxT-1 is still the common method, although T-3 is gaining in popularity as prices fall. This corporate network connection serves two purposes:

- Provide a means for the IS staff to manage the databases and back-end transaction servers in-band.
- Provide connectivity to the out-of-band management network consisting of the asynchronous connection, and the tertiary Ethernet interfaces on the Web servers for software imaging, etc.

Some companies want to use this connection as a means to provide Internet connectivity to the office. This is not strongly recommended but can be facilitated by using a router with 3 Ethernet or Fast Ethernet interfaces and connecting one of them to a separate interface on the outside PIX firewall. The router here should be able to support at least 3 Ethernet interfaces and should have asynchronous port connectivity as well for console connections.

## The Total Cisco Solution – Piecing it Together

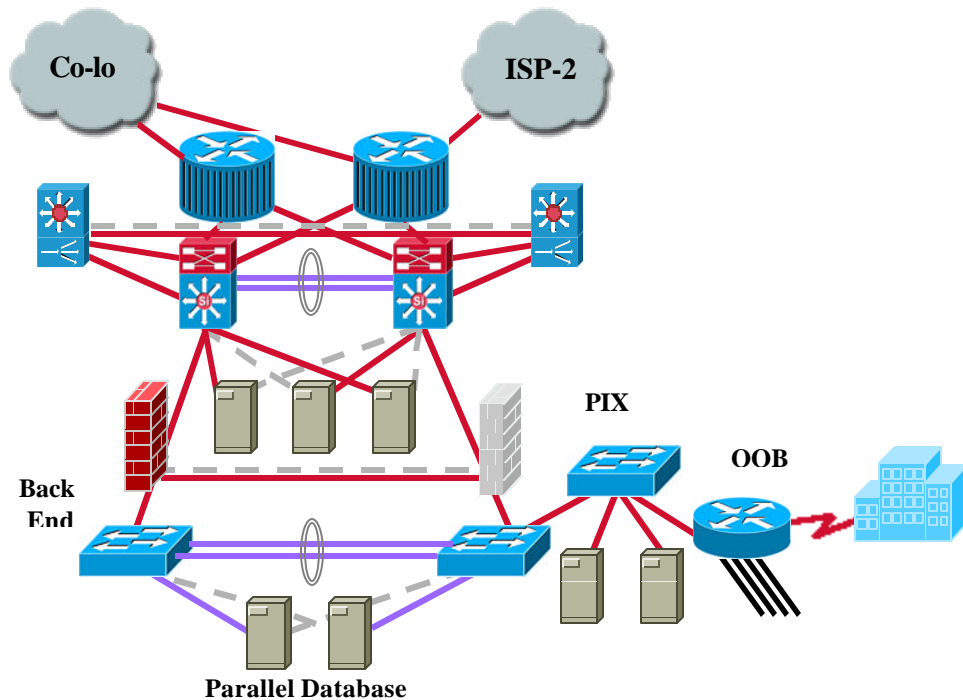
The task of building an iCommerce site seems rather daunting when you look at the complete picture of all the necessary design components. The fundamental purpose of a site is to allow customers to access information and ultimately make purchases across the Internet. To recap, the important factors in building an iCommerce solution are:

- *High Availability*
- *Security*
- *Manageability*
- *Scalability*

Unfortunately, the Internet needs to be treated as a hostile environment, and applications exposed to it must be hardened to survive potential malicious activity. Security is a major consideration when designing the network, but it should not be seen as a hindrance to conducting business on the Internet.

The basis of any iCommerce solution is the application. Understanding the network requirements of the application helps in the design and securing of application for use on the Internet. The iCommerce applications all have the same fundamental building blocks as we have seen previously. There are set of front-end web servers responsible for the delivery of information to the Internet customer. Supporting the front-end servers are a series of back-end database and application servers responsible for administrative duties such as order placement, billing, and customer account management. The traditional method for order fulfillment was a process beginning with a customer placing an order by phone or fax. In the old process, a clerk would then enter that order into a system, the product was built, and shipped to the consumer. The iCommerce solution revolves around the end customer placing the order directly into order-entry systems directly via the Internet. One fundamental problem is how to grant customers access to parts of the internal systems in a secure fashion.

## The Highly Secure Redundant iCommerce Solution



The design illustrated above is the most comprehensive in terms of security and may seem overwhelming at first glance. The web infrastructure is connected to a co-1o and a secondary ISP, so BGP will be running on both the ISP-facing routers for connectivity to the Internet. This arrangement is referred to as being multi-homed. Many times a dot-com will implement this as a phase 2 initiative.

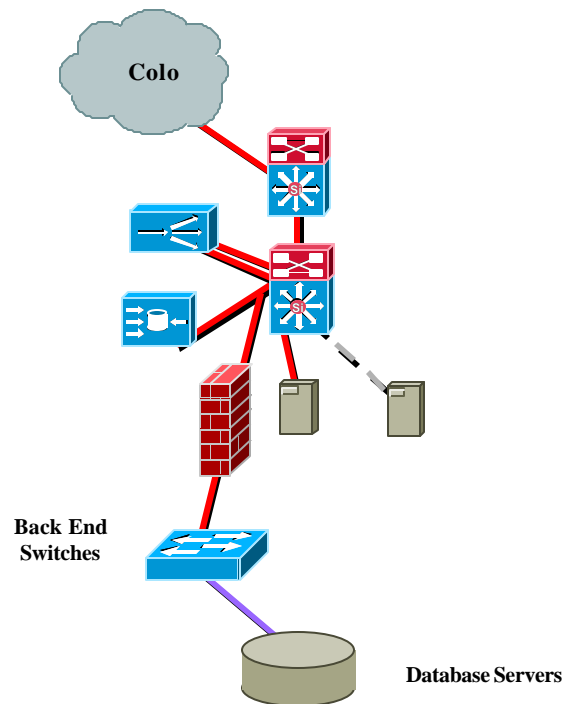
The next module of the network is the Internet-facing web server farm core switches. These switches should be implemented with L3 switching. The load balancers will be attached to these switches along with the web servers. The method for attaching the load balancers to these switches depends upon the mode in which the load balancers operate. There are two web server farm core switches such that if one switch should fail the servers attached to the other switch continue to deliver content. The next module is the web servers themselves. The back-end network is utilized with a separate VLAN of the core switch (in some cases, another set of switches) to access back-end application servers and database servers. The firewalls protecting the back-end database network need only have conduits to pass database traffic.

The Intrusion detection devices are placed on the core switches and the web server farm VLANs. The first question which might spring to mind is, "Why so much Intrusion detection?". The reason for intrusion detection installed on both the outside and inside of the web server farm firewalls is to allow the knowledge of attacks being attempted from the outside of which the firewalls stop. The internal intrusion detection device is to detect if an attack does get through the firewall at which point firewall policy can be changed to prevent future attacks.

The final module of the network is the out-of-band (OOB) terminal server. The terminal server is connected to the consoles of all devices for management purposes. This allows all the networking devices to be configured and monitored out-of-band. One of the major issues in this type of environment is security. For this reason, SNMP is normally not used. However, if SNMP is necessary then devices which utilize it must be carefully tested for security leaks. It is highly desirable that the SNMP connections be conducted on the out-of-band management network. This network, if implemented, should utilize an RFC1918 (private)

address, and be carefully monitored.

### ***The Highly Secure Non-Redundant iCommerce Solution***



Sometimes redundancy needs to be sacrificed due to budgetary reasons, or redundancy is provided through the use of multiple sites which will be described later. The solution illustrated above shows the same solution as before except for the fact that all the redundancy has been removed.

### ***Dissecting the Highly Secure iCommerce Solution***

Having looked at the physical hardware and topology, it is now beneficial to look at the logical topology and exactly how the components are configured and interconnected.

The following table provides a sample configuration for an ISP-facing edge router. In this example, Router configurations provide a first line of security by using ACLs to only allow web, SSL, DNS and BGP traffic to flow through. In addition, the router establishes a BGP peering session with an upstream ISP.

```
!  
version 12.0  
service timestamps debug uptime  
service timestamps log uptime  
no service password-encryption  
!  
hostname sflab-icr1  
!  
boot system flash 7200-1207.bin  
enable password cisco
```

```

!
ip subnet-zero
ip host icr2 10.10.1.2
!
!
interface Loopback0
ip address 10.200.6.1 255.255.255.255
no ip directed-broadcast
!
interface FastEthernet0/0
description # ISP 1 #
ip address 10.205.1.2 255.255.255.0
ip access-group 101 in
no ip redirects
no ip directed-broadcast
full-duplex
!
interface FastEthernet1/0
description Link to ICS1 2/1
ip address 10.200.1.3 255.255.255.0
no ip directed-broadcast
full-duplex
!
interface FastEthernet1/1
description To_Cat6K
ip address 10.200.2.1 255.255.255.0
no ip redirects
no ip directed-broadcast
full-duplex
!
interface ATM2/0
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
router eigrp 200
redistribute bgp 65534 metric 100000 10 255 1 1500
network 10.0.0.0
!
router bgp 65534
network 10.200.1.0 mask 255.255.253.0
network 10.200.1.0 mask 255.255.252.0
neighbor 10.1.100.11 remote-as 65533
neighbor 10.1.100.11 ebgp-multihop 255
neighbor 10.1.100.11 update-source Loopback0
neighbor 10.1.100.11 distribute-list 1 in
neighbor 10.200.6.2 remote-as 65534
neighbor 10.200.6.2 update-source Loopback0
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.205.1.1 100
ip route 10.172.16.0 255.255.255.0 10.173.16.0
no ip http server
!
access-list 1 deny 192.168.0.0 0.0.255.255

```

```

access-list 1 deny 10.200.0.0 0.0.255.255
access-list 101 permit tcp any 10.200.3.0 0.0.0.255 eq www
access-list 101 permit tcp any 10.200.3.0 0.0.0.255 eq 443
access-list 101 permit udp any 10.200.3.0 0.0.0.255 eq domain
access-list 101 permit tcp any any eq bgp
access-list 101 permit icmp any any
access-list 101 deny ip any any
!
line con 0
transport input none
line aux 0
line vty 0 4
password cisco
login
!
end

```

**Figure 18. ISP-Facing Router Configurations**

### ***Internet Connectivity***

The two ISP-facing routers are connected to different service providers (the co-lo and a secondary ISP) and have EBGP connections to the co-lo and ISP2. The two ISP-facing routers also have a link between themselves principally to provide connectivity such that IP packets can take the optimal path to the Internet through either ISP

The most challenging part of building the web server farm is dealing with the BGP peering and IP addressing issues associated with connecting to the Internet. One may not be fortunate enough to have portable address space, but rather have addresses allocated by the ISPs in their CIDR blocks. This may result in the need to do NAT on the BGP-peering routers. If this becomes necessary, the following document provides guidance on configuring your Internet connectivity. The document is entitled “Enabling Enterprise Multihoming with Cisco IOS Network Address Translation (NAT)” and can be found at [http://www.cisco.com/c/prod/collo/collo/cisco/mkt/ios/nat/tech/emios\\_wp.htm](http://www.cisco.com/c/prod/collo/collo/cisco/mkt/ios/nat/tech/emios_wp.htm)

### ***Web Server Switches***

The web server switches, depending upon the load balancing technology being used, will normally need to utilize two different VLANs. In Figure 20, there are two different VLANs, but they represent the same subnet from an IP perspective. The load balancing product being utilized is Cisco’s CSS-11000 series, and it acts as a bridge between the two VLANs. The switches have the following features configured:

- PortFast enabled on all ports
- VLAN trunking disabled on all ports
- Port-Channeling disabled on all ports (EtherChannel)
- VLAN trunk between ports 1/1 on both switches, trunking VLANs 98 and 99

### ***Load Balancing/Content Switching***

In the example, the CSS-11150 is configured with the virtual IP address of 10.0.2.3, and contains a layer 5 content rule. The content rules will make a traffic switching decision based on URL and maintain persistence to the server by examining a cookie that the switch inserts (example: s3). It uses weighted round robin as its load balancing algorithm.

```

!***** CIRCUIT *****
circuit VLAN1

ip address 10.0.0.3 255.0.0.0

!***** SERVICE *****
service server1
ip address 10.0.3.221
keepalive type http
keepalive method get
keepalive uri "/index.html"
keepalive port 80
string s1
active

service server2
ip address 10.0.3.222
keepalive type http
keepalive uri "/seriesb1.gif"
keepalive port 80
string s2
active

service server3
ip address 10.0.3.223
keepalive type http
keepalive method get
keepalive uri "/index.html"
keepalive port 80
string s3
active

service server4
ip address 10.0.3.224
keepalive type http
keepalive uri "/seriesd1.gif"
keepalive port 80
string s4
active

!***** OWNER *****
owner student3

content layer3ipsticky
vip address 10.0.2.3
advanced-balance sticky-srcip
balance weightedrr
add service server1
add service server2
add service server3
add service server4

content layer5stickycookurl
port 80
protocol tcp
vip address 10.0.2.3

```

```
balance weightedrr
advanced-balance cookieurl
add service server1
add service server2
add service server3
add service server4
url "/*"
active
```

**Figure 26. CSS-11150 Configuration**

## **Firewalling**

This is an example of a config for a Cisco Secure PIX firewall located between the web server VLAN and the back-end databases. By default, no traffic is allowed through the firewall. You'll notice that conduits have been configured to only allow database (SQLNET) traffic to pass. In addition, the **fixup protocol** commands let you view, change, enable, or disable the use of a service or protocol through the PIX Firewall. The ports you specify are those that the PIX Firewall listens at for each respective service. The **fixup protocol** commands are always present in the configuration and are enabled by default.

The **fixup protocol** command performs the Adaptive Security Algorithm based on different port numbers other than the defaults.

```
PIX Version 5.0(2)
nameif ethernet0 outside security0
nameif ethernet1 inside security100
nameif ethernet2 pix/intf2 security10
nameif ethernet3 PIX-Stateful security75
enable password 2KFQnbNIdI.2KYOU encrypted
passwd 2KFQnbNIdI.2KYOU encrypted
hostname icfw
fixup protocol ftp 21
fixup protocol http 80
fixup protocol smtp 25
fixup protocol h323 1720
fixup protocol rsh 514
fixup protocol sqlnet 1521
names
pager lines 24
no logging timestamp
no logging standby
no logging console
no logging monitor
no logging buffered
no logging trap
logging facility 20
logging queue 512
interface ethernet0 auto
interface ethernet1 auto
interface ethernet2 auto shutdown
interface ethernet3 auto
mtu outside 1500
mtu inside 1500
mtu pix/intf2 1500
mtu PIX-Stateful 1500
ip address outside 10.200.3.4 255.255.255.0
```

```

ip address inside 10.200.5.1 255.255.255.0
ip address pix/intf2 127.0.0.1 255.255.255.255
ip address PIX-Stateful 10.100.100.1 255.255.255.0
failover
failover timeout 0:00:00
failover ip address outside 10.200.3.4
failover ip address inside 10.200.5.1
failover ip address pix/intf2 0.0.0.0
failover ip address PIX-Stateful 10.100.100.2
failover link inside
arp timeout 14400
static (inside,outside) 10.200.3.4 10.200.3.4 netmask 255.255.255.255 0 0
conduit permit icmp any any
conduit permit tcp host 10.200.3.101 eq sqlnet 10.14.1.128 255.255.255.192
conduit permit tcp host 10.200.3.102 eq sqlnet 10.14.1.128 255.255.255.192
conduit permit tcp host 10.200.3.103 eq sqlnet 10.14.1.128 255.255.255.192
conduit permit tcp host 10.200.3.104 eq sqlnet 10.14.1.128 255.255.255.192
no rip outside passive
no rip outside default
no rip inside passive
no rip inside default
no rip pix/intf2 passive
no rip pix/intf2 default
no rip PIX-Stateful passive
no rip PIX-Stateful default
timeout xlate 5:00:00 conn 5:00:00 half-closed 0:10:00 udp 0:02:00
timeout rpc 0:10:00 h323 0:05:00
timeout uauth 0:05:00 absolute
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
no snmp-server location
no snmp-server contact
snmp-server community public
no snmp-server enable traps
no floodguard enable
telnet timeout 5
terminal width 80
Cryptochecksum:9a73b33f6bc55089e02fba19a64e276c

```

**Figure 27. Cisco Secure PIX Configuration**

### ***Back-end Network***

The back-end network is the network over which content is updated and database access in addition to access to the home office is done. This network is numbered with an RFC1918 address so it can not be directly accessed from the Internet. One of the firewall policies should be not to provide any conduits to this network from outside.

### ***Intrusion Detection***

One of the most important parts of the network is intrusion detection. By knowing from who and from where attacks are being launched on your web site allows you to contact the ISPs involved and have the malicious activity blocked. The other advantage is that the intrusion detection system can respond in real-time to minimize the damage of an attack by applying filters in the ISP-facing routers. The intrusion detection devices in this example are Cisco's IDS sensors. Each Cisco Secure IDS has two interfaces, one which connects to the network being monitored and another connected to the network so the sensor can

communicate with the Cisco Secure IDS director console. The Cisco Secure IDS sensors are connected to the various switches, and the ports to which they are connected are configured to mirror all traffic on the specified VLAN to the Cisco Secure IDS sensors.

### Out of Band Network

The out of band network consists of the terminal server being located inside the firewall. All device console ports are attached to the terminal server allowing out-of-band configuration and management. This allows for the use of TACACS+ or RADIUS authentication on the terminal server to audit who is configuring which devices. If some devices need to be managed through SNMP, a separate ethernet segment should be created upon which the dedicated management station is placed using RFC1918 addresses. In this example all management and configuration is performed via the terminal server.

### Multi-site Redundancy

The final challenge is allow for multiple data centers across the country or around the world. When a user requests a particular application across the Internet, the desire is to connect them to the site that will deliver optimal performance. This is achieved by the use of the DistributedDirector product. When using the application, the user requests to be connected to a particular DNS name or URL. The purpose of DNS is to resolve the symbolic name into a particular destination IP address. DistributedDirector is configured in DNS name server mode, which means it acts just like a DNS caching name server in this example. At each of the Data Centers there are a group of servers for “[www.company1.com](http://www.company1.com)”. Figure 27 illustrates the process utilized by DNS to resolve the host name.

When the user requests ‘www.company1.com’, the workstation sends a request to the DNS server configured to resolve this name to an IP address. The DNS server issues a recursive DNS query to discover the IP address associated with ‘www.company1.com’. The local DNS server returns to the client the IP address associated with ‘www.company1.com’. This is done because the DistributedDirector is configured as the authoritative name server for




Figure 27. Multiple Site Implementation with DistributedDirector

[www.company1.com](http://www.company1.com).

1. Client issues DNS request for ‘www.company1.com’ to the local name server
2. Local name server finds the address for the SOA name server for ‘company1.com’. It tells the local DNS that the authoritative name server for ‘www.company1.com’ is ‘dd-1.company1.com’ (The Distributed Director)
3. Local name server makes the request to ‘dd-1.company1.com’ and it returns back to the local DNS 171.69.121.231 for ‘www.company1.com’
4. The local name server sends back to the client the address resolution for ‘www.company1.com’
5. The client uses the IP address 171.69.121.150 to connect to the application.

The routers in Figure 27 providing Internet connectivity are configured with Director Response Protocol (DRP) agents, so that DistributedDirector has access to the routing tables to determine the optimal server to return to the client. The use of DRP agents allows DistributedDirector to utilize topological proximity metrics based on IGP and BGP information, and also to determine round trip times between clients and servers. To provide redundancy for Distributed Director, multiple name server delegations can be placed in the DNS configuration, and the DNS server will perform round-robin allocation of requests between the configured DistributedDirectors. In the example given above there are two DistributedDirectors utilized 'dd.isp.net' and 'dd-backup-isp.net'. The DistributedDirectors have identical configuration for delegated sub-domains.

### **DNS Configuration**

```
Company1
company1.com. NS      calabogie1.isp.net.
company1.com. SOA    sysadmin.company1.com 1 10800 3600 604800 86400
application NS      dd-1.company1.com.
application NS      dd-2.company1.com.
```

### **Distributed Director Configuration**

```
ip host dd-1.company1.com 10.10.20.231
ip host www.company1.com 171.69.221.150 171.69.111.150
ip domain-name company1.com
ip name-server 171.69.20.254
ip classless
ip dns primary www.company1.com soa dd.isp.net sysadmin@isp.net 21600 900 7776000 86400
drp-e 10 drp-i 50
ip director server 171.69.221.150 drp-association 171.69.221.1
ip director hosts www.company1.com connect 80 interval 300
```

## **Reference**

Enabling Enterprise Multihoming with Cisco IOS Network Address Translation (NAT)  
[http://www.cisco.com/cpropart/salestools/cc/cisco/mkt/ios/nat/tech/emios\\_wp.htm](http://www.cisco.com/cpropart/salestools/cc/cisco/mkt/ios/nat/tech/emios_wp.htm)

Security Design Guide

[http://www.cisco.com/warp/partner/synchronicd/cc/sol/mkt/ent/secur/secur\\_dg.htm](http://www.cisco.com/warp/partner/synchronicd/cc/sol/mkt/ent/secur/secur_dg.htm)

Increasing Security on IP Networks

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs003.htm#xtocid176512>

Cheswick, William & Bellovin, Steven. "Firewalls and Internet Security - Repelling the Wily Hacker", Addison Wesley, 1994. ISBN 0-201-63357-4

Chapman, D. & Zwicky, Elizabeth. "Building Internet Firewalls", O'Reilly & Associates, 1995. ISBN 1-56592-124-0

CSS-11000 Series / DD / PIX / IOS Manuals

RFC1918 Address Allocation for Private Internets