Speeding and Scaling Web Sites Using Cisco Content-Delivery Technology

Executive Summary

The explosion of the Web has impacted all of us. It has become a part of the daily fabric for consumers and businesses alike, exacerbating the demands on Internet infrastructure. Web traffic has skyrocketed, sites have incorporated ever-richer media, and the trend shows no signs of abating. Slow Web sites have become so familiar that the "WWW" acronym has been commonly parodied as the "World Wide Wait."

At the same time, a company's Web presence has become increasingly mission critical; with ever more sophisticated business-to-consumer and business-to-business applications. Web performance now translates directly into key financial metrics; slow, unresponsive sites have a bottom-line impact.

So what are the underlying causes of slow Web environments? While performance is a complex engineering topic, today's Internet architecture is based on centralized servers delivering files to all points on the Internet. Some basic bottlenecks can surface with this approach:

- Network link congestion due to insufficient capacity on a particular telecommunications link
- Network equipment congestion due to the processing speed of today's networking equipment
- Web server congestion due to the data processing speed of today's Web servers
- Distance delay in the network due to the time associated with data traveling over long distances

Moreover, multiple organizations are typically associated with each step of the connection from a server to a browser. They may be beyond the control of the site owner and primary carrier.

A variety of methods have been employed to address these delays:

- More bandwidth and faster network equipment—This is the simplest "brute-force" approach, but can be expensive and it addresses only bottlenecks that can be controlled.
- Content replication and load balancing across local server farms—This enables Web sites to scale and to offer more reliable service.
- Geographic distribution of servers—This goes a little further in addressing congestion both on the server and in the network.
- Content caching—This addresses the distance delay issue by offering geographical diversity and serving content close to the requester. Caching also saves bandwidth, sending a piece of content only once from the origin server to the cache, and then serving many times from the closer location.

These solutions provide great flexibility in solving many of the problems that plague Web sites and surfers today—but each method does not provide the complete solution alone. Only when you combine these traditional techniques with a new approach called content-delivery networks are all of the problems addressed.

This new approach offers a significant breakthrough in delivering Web content. Content-delivery networks address many of the limitations of previous solutions to deliver high performance by replicating content to the "edge" of the network, thereby minimizing the distance between where content is requested and where it is served. In so doing, content delivery also saves bandwidth. The fundamental principles underlying content delivery include central control of content and the network, efficient distribution of content to the servers at the "edge" of the network, and automatic redirection of content requests to a local edge server.

Cisco offers the industry's only complete content-delivery solution. With these products and the traditional Web scaling technologies of load balancing, higher-speed networks, and caching, network service providers can gain many benefits. These benefits include improving Web-site performance for delivery of on-demand content for both static and streaming media, greatly improving peak load site performance for Web events, lowering bandwidth costs by delivering content from the edge of the network, and enabling new premium Web services.

By enabling new premium Web services, network service providers can enhance their differentiation and create new revenue streams. Some of the new services enabled by the Cisco content-delivery software include on-demand streaming-media services, distributed content vending services, scalable live streaming-media services, and live or on-demand TV-quality media to the enterprise.

This paper examines the current Internet bottleneck problem, the limitations of previous attempts to solve it, how content delivery addresses these problems, and how the Cisco solution solves these problems to speed and scale the Web.

Performance Bottlenecks on the Internet

The Internet has grown from a research tool used primarily by scientists to a global network that links people and businesses together. What used to be a tool of convenience is now a necessity, providing a critical medium for all types of transactions. Today people rely on the Internet to communicate via e-mail, conduct research, trade stocks, conduct online banking, obtain news, play games, and conduct business-to-business/business-to-consumer e-commerce. This global network offers fast, low-cost distribution of information, and it has caused a shift in the balance of knowledge. Small companies can now compete with giant multinational corporations, and consumers have the information they need to be educated shoppers. But this technology is not without limitations or complaints. All through its development, there have been growing pains.

According to Computer Industry Almanac, since 1995 the number of Internet users worldwide has increased nearly eightfold from 39 million to 318 million, and projections show this number doubling to more than 700 million users by 2005. The growth in the number of Web sites has grown from 380,000 in 1996 to 3,400,000 in 1999 (source: estats). Sheer growth in users and Web sites coupled with the increased complexity of sites and associated applications has led to the biggest complaint attributed to the Internet—lack of speed.

According to the Internet Research Group, the time to load the home page of key Internet sites (as measured by Keynote Systems Business 40 Web Sites) has improved consistently over the last four years. But overall page-loading performance is still in the range of 8 seconds—ten times longer than the threshold believed to represent natural human reading/scanning speeds. In a study conducted by GVU Center at Georgia Tech (its tenth WWW User Survey), over 60 percent of the respondents cited speed and slow ads as the major problems with the Web.

Offering unlimited potential as both an entertainment medium and business tool, the Internet is facing barriers that other major technological advances, such as the highway system, faced in their infancies. As we know, these systems flourished after overcoming these obstacles. Unless these constraints are removed from the growth of the Internet, realizing its potential may be delayed or unfulfilled.

Economic Effects of Internet Congestion

Poor Internet experiences such as long download times, slow performance, low-quality video, outdated information, or service interruptions have an economic impact on the Web-site sponsor and host. Studies show that users who have a bad experience with a site likely will not return. The end result is lost opportunity, including lost customers and revenue.

Zona Research's Economic Impact of Unacceptable Web Site Download Speeds tracked "bailout rates" (the percentage of people who did not wait for pages to load but instead went to other pages). Figure 1 reveals some interesting statistics on the bailout rates for Web pages of various sizes. The study noted that over 50 percent of the people attempting to download a page in excess of 70 kB leave before the page is completed. As might be expected, the larger the page size, the greater the bailout rate: Typically a 40-kB page has a 30-percent bailout rate, while a 34-kB page has only a 7-percent bailout rate. What appears to be a mere 6-kB difference—15 percent less information and download time—results in a greater than fourfold reduction in the bailout rate. This data demonstrates users' high sensitivity to download times.

Figure 1 Bailout Rates for Various Web Page Sizes

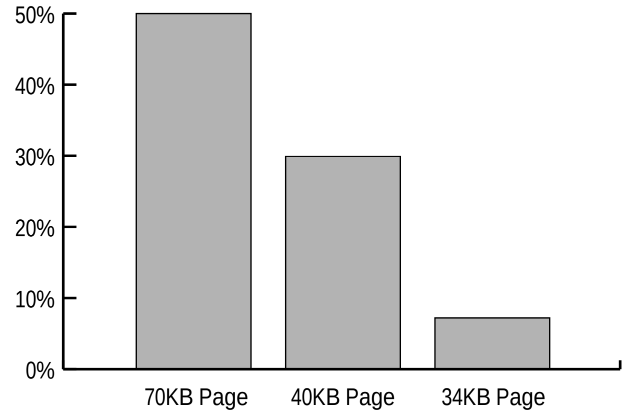

Table 1 The Economic Impact of Web Performance

Site Type	Examples	Impact
Heavily Traveled Sites	Portals, Search, News, Top 100	Lost advertising revenue
e-Commerce	On-line Catalogs and Shops	Lost sales revenue
Streaming Media Sites	WebCasting and Video/Audio on-Demand	Lost event revenue and repeat business

Bottleneck Causes

The Web is slow for several major reasons:

- Network link congestion
- Server congestion
- Network equipment congestion
- Distance delay

Network Link Congestion

Network link congestion is the data transmission delay associated with the capacity (bandwidth) of the telecommunication infrastructure. This is determined by various factors, including the speed of the equipment installed in the network and the capacity of the transmission lines. As the growth in users and applications on the Internet has grown, the telecommunication networks responsible for providing the infrastructure of the Internet have grown in an attempt to keep up with demand. Unfortunately, demand for bandwidth has outstripped telecommunication network capacity, and analysts expect this trend to continue.

Many individual network links may exist between two points in the Internet. A few large carriers operate high-capacity networks, known as the Internet "backbone," which connect the various network service providers. These networks are connected at "peering points." The many Internet service providers (ISP) then connect their networks to one or more backbone providers.

Some technology advances such as terabit routers, dense wave-division multiplexing, and faster transmission equipment have been implemented to help alleviate network delays due to bandwidth. But, it is the network peering points (see Figure 2) that continue to be a major bottleneck.

Figure 2 Network Peering Points

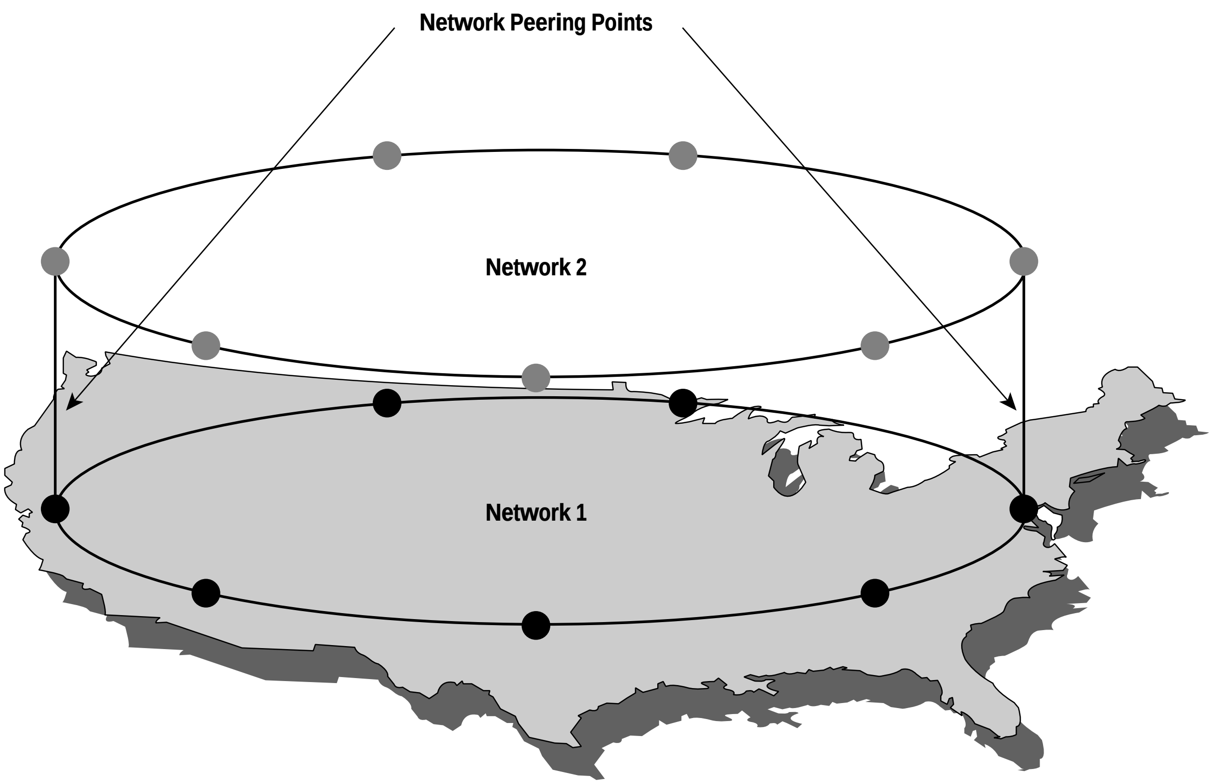

Figure 3 End-to-End Network Speed

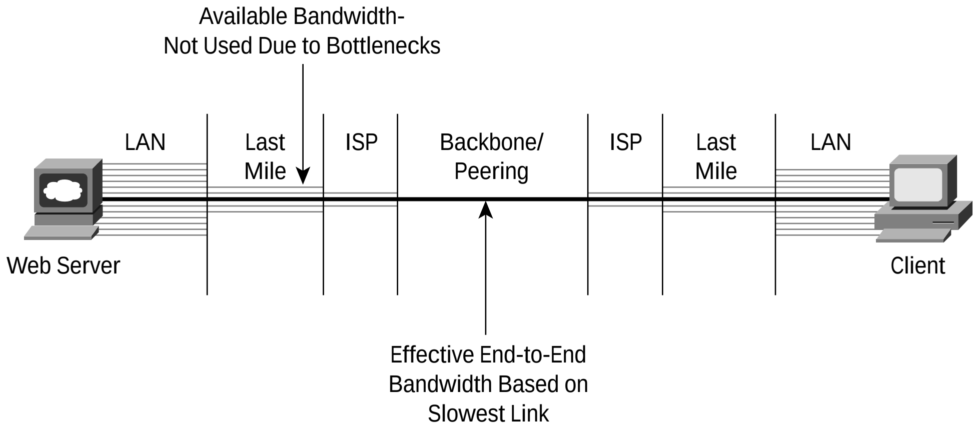

The end-to-end speed of a network is only as fast as its slowest link. Provisioning last-mile connections, even with gigabit connectivity, will do little to solve the speed problem. If the trip of a packet through the core of the network or a peering point can achieve only a 64-kbps performance (equivalent to a dialup line), then additional last-mile bandwidth will not increase performance. Consequently, as more users, richer media, and complex applications are introduced, any added capacity would be consumed.

Network Equipment Congestion

Internet data travels through multiple network devices, including routers and switches that connect the transmission lines between locations. Each network device processes the data before moving it and continuing on to the next item. As network traffic increases, these pieces of equipment must process more and more information packets, which collect in queues waiting to be processed. Even with the tremendous design advances of networking equipment, exponential growth in Internet traffic has led to substantial queuing delays. When overloaded, network equipment discards queued data packets, requiring them to be retransmitted, in turn increasing traffic even more.

Advances in network equipment processing power should continue to follow Moore’s Law, the doubling of processing speeds every 18 months, and enable larger amounts of data to be processed faster. However, if Internet traffic continues to grow faster than the rate at which technology advances, the problem of network equipment congestion will continue.

Server Congestion

Whereas network equipment congestion is associated with throughput speeds of network devices, server congestion is the delay incurred with processing requests for data. When a Web site receives a visitor, the server must process inquiries and transmit the requested data. The larger the traffic flow on a Web site, the more requests must be served and processed by the server. Until the data is served, the requests sit in a queue, resulting in a delay and a slow Web experience.

There are a few approaches to dealing with server congestion. One is to install more powerful servers, with faster processors, multiple processors, and more memory. This will enable the server to process requests faster. However, larger servers cannot address the requirements of a Web site that is expected to serve tens of millions of requests from around the world each day.

Another method is to balance the load over multiple servers. Load balancing is the only way to address high server request rates associated with popular Web sites from a central location. Techniques for distributing a Web site over multiple servers are addressed in the section “Current Web Acceleration Solutions.”

Distance Delay

More bandwidth and better equipment will help resolve some of the limitations of today's Web, but the obstacle of distance is still a problem. There are inherent time limitations in the transmission of data over long distances that cannot be changed without modifying the laws of physics. Fiber-optic cable cannot transmit data faster than the speed of light. No matter how advanced the technology, certain constraints cannot be overcome, even with more powerful equipment or bandwidth.

Serving a simple Web page with text and graphics involves many transactions between the browser and Web server. One transaction occurs for each embedded object plus another for the HTML page. Therefore, the delay begins with the geographic distance the request must travel between the client computer and the Web server hosting the desired content. The further the distance, the longer it takes for that request to be received. Each transaction requires a connection to the Web server and the transmission of the requested data. Without even considering network equipment or network link congestion, the time it takes for signals to travel long distances over fiber-optic or copper cables can constitute up to one-half the time it takes to download a typical Web page.

Assuming that data download is instantaneous (not counting download time for a 35-kB GIF or 3-MB document), the connection setup time just to open a TCP session and the handshake between a browser and Web server can cause an unacceptable delay. For example, assume a home page with 17 embedded objects (for example, GIFs); the associated connection setup time could range from 0.4 seconds (Web server in New York and browser in Denver) to 2.9 seconds (Web server in New York and browser in Melbourne, Australia). These estimates do not even account for the time to download the actual content.

Table 2 Connection Setup Times for Various Distances

Location	Network Miles	Connection Setup Time (seconds)
New York – Denver	1771	0.4
New York – San Francisco	2934	0.7
New York – Singapore	11923	2.7
New York – Melbourne	12940	2.9

Current Web Acceleration Solutions

Current Web acceleration methods take advantage of the technologies in the five categories, as shown in Figure 4:

Figure 4 Methods Available to Increase Speed on the Internet

Three primary approaches incorporate the following methods:

- Local server farms
- Manual distributed site mirroring
- Network caching

These solutions provide great flexibility in solving many of the problems that plague Web sites and surfers today—but each method does not provide the complete solution alone. Each of these methods has advantages and limitations when addressing the issue of speeding Web performance and scaling to address large audiences.

Local Server Farms

Method: Content replication; load balancing
Issue Addressed: Server congestion

One way of reducing server congestion is to add multiple local Web servers to a central site. These servers work together to reduce the time it takes to respond to a request for content. Local server farms spread the content across multiple servers that are installed in a single location.

Advantages:
- This solution enables the combination of multiple servers to act as a single large server to address increased traffic.
- This solution enables central management and maintenance.

Limitations:
- For simple implementations where Domain Name System (DNS) round robin is used and a single Universal Resource Locator (URL) is mapped to multiple Web servers, the status of each server is unknown by all other servers and users could be sent to a server that is not working or is congested, resulting in an even poorer Web experience.
- This solution is susceptible to network connectivity outages, bringing down a Web site.
- This solution addresses only congestion associated with the central site; it cannot address delays associated with network link congestion, network equipment congestion, or distance delay.

Manual Distributed Site Mirroring

Method: Content replication; geographic distribution; load balancing
Issue Addressed: Server congestion; some impact on network link congestion; network equipment congestion; and distance delay

Manual distributed site mirroring consists of distributing Web servers over several geographical locations and copying files manually onto each server. It addresses many of the issues associated with Internet performance in a limited fashion.

Advantages:
- Latency issues associated with network equipment, network congestion, and server congestion can be addressed by deploying multiple servers. This spreads the load over multiple servers in diverse locations and bypasses some of the congested servers.
- Network link congestion and distance delay are addressed with this method, because there are multiple sites with independent Internet connections distributed around the world. However, because techniques for routing requests to appropriate servers are not sophisticated, requests may be sent to a server that is currently congested, or they may not be sent to the closest server.

Limitations:
This solution does address three of the major issues responsible for slow Web sites. In addition, this solution can create other issues for the Web-site owner:
- In many implementations where minimal server selection intelligence is implemented, there is a good chance that the request will go to a server that is not appropriate. Performance, therefore, can be worse than if simply sent to the origin server because of geographical or congestion considerations. The lack of network intelligence may actually result in slower Web performance by sending a request to the farthest geographical location (sending a request from Denver to the Hong Kong server instead of New York) or network performance location.
- Latency associated with other network equipment, such as routers and transmission equipment, are not resolved in this case because traffic may need to transverse congested network links or peering points.
- Managing multiple servers in multiple locations is expensive both in manpower and equipment.
- Because of fixed bandwidth pricing, operating multiple locations could result in overprovisioning of bandwidth and increased cost.
- Keeping content up-to-date and fresh at all sites is a complex and expensive task if not automated. Human error can mean that inappropriate or out-of-date content is served.

Network Caching

Method: Caching

Issue Addressed: Server congestion; some impact on network link congestion; network equipment congestion; and distance delay

Network caches deployed in an ISP's network will deliver faster Web performance and save bandwidth when the requested content is stored in the cache.

Advantages:

- A cache located near the requestor can address bandwidth congestion issues at the core of the Internet by bypassing the network and serving content from the edge—near the user requesting the content.
- A cache can also lower distance delay and eliminate some of the network equipment and server congestion because of the load being distributed among many caches and the data being served from a closer location, allowing for the bypassing of some network segments.

Limitations:

- Cache misses occur when the content requested is not present in the cache. When this happens, the cache must then go back to the origin Web server, pull the content, and send it to the user. This takes longer than simply going initially to the origin Web server. The Internet Research Group estimates a typical network caching solution results in a cache miss of 60 percent.
- When content is served from a cache, the usage data is not automatically reported back to the Web-site owner, and must be retrieved and assembled with other usage data—a scenario that can be cumbersome. If a Web site is charging for ad impressions or a unique number of users, accurate usage tracking is critical. Inability to accurately report impressions can result in lost ad revenue.

Principles of Content Delivery

"Over the next two years, a whole new layer of infrastructure that Forrester calls 'content routing' will emerge. This will optimize content delivery by linking applications with distributed network servers, and distributing requests across the network to maximize user performance."

Forrester Research, Scaling Web Performance, March 1999

Content-delivery solutions speed Web performance and overcome some of the issues associated with current-delivery methods. Content delivery provides for:

- *Control content and network centrally*—All content and network devices are controlled from a central location, thus eliminating the costs associated with remote administration.
  - Content that is delivered by a content-delivery solution has a single URL associated with it, regardless of its server source. A single URL eliminates the need to maintain multiple sets of content or Web sites.
  - All distributed network devices are network appliances that are easy to install and maintain while offering the benefits and functionality of a server. The appliance nature of the product enables the devices to be installed and managed from the central location without remote intervention.
- *Distribute content to the edge of the network*—Content-delivery solutions place appliances throughout the Internet, close to where the content requests originate. This principle mirrors some of the functionality of caches, but with distinct differences. In particular, these appliances work together in a centrally controlled collaborative fashion to ensure overall network performance. Like a cache, content is replicated from the origin server to the cache only once, regardless of the number of times the content is served. However, a CDN provides greater content control. By prepopulating the server, the content will be available for fast delivery to the user, eliminating cache misses and increasing the hit rate. Combining this with on-demand populated caching delivers a totally integrated solution with the benefits of both approaches.

- *Redirect content requests to a local source*—When the URL is requested by a browser, the content-delivery solution intelligently determines a local delivery server based on:
  - Geographical and network location
  - Presence of content
  - Current status of server (both availability and load)
  - Current status of network (traffic load and network errors)

Table 3 Congestion Issues Addressed by Content-Delivery Solutions

Source of Congestion	Content Delivery Solution
Network Link Congestion	Content stored and served locally where bandwidth is more plentiful and less expensive.
Network Equipment Congestion	Content stored and served locally where content can be delivered without encountering many pieces of network equipment where it might be queued.
Server Congestion	If multiple delivery nodes are distributed throughout the Internet, each single delivery node will only serve a portion of the requests for a given Web site. NO single Web server will get overloaded.

Comparing Content Delivery to Current Web Acceleration Solutions

Current Web acceleration solutions (pre-content-routing-era technologies) have been discussed along with their limitations in the previous sections. The chart in Figure 5 compares these solutions with content delivery (content routing in Forrester Research nomenclature).

Table 4 Comparing Content Delivery to Current Web Acceleration Solutions

	Pre-Content-Routing-Era Technologies	Content-Touring-Era Technologies	Impact of Content Routing
Integration	Very poor	Products and services designed to publisher's needs	Process controlled by content owner
Distribution platform	Few Web servers	Many caches dispersed worldwide	Less expensive and easier to manage
Update process	Scripted directory transfers: no guarantee of delivery	Automated, streamed update	Up-to-date content
User redirection	Based on limited server metrics	Factors network capacity, current usage, and content availability	Optimized user performance
User tracking	Inconsistent, incomplete logs must be manually consolidated	Consistent, real-time usage available locally and across the network	Complete view of usage and effectiveness
Scalability	Add more Web servers at each distribution center	Popular content dynamically pre-populated to underutilize content POP's	Improved peak protection

Increasing Effective Bandwidth through Content Delivery

Moving the content serving point from the origin Web server to a content-delivery node that is closer to the content request will cause the content to bypass many network segments. By bypassing the congested segments of the network, the lower-cost, less-congested segments with more available bandwidth are used—resulting in a faster, more robust Web experience. In Figure 5, the content is being delivered from an ISP point of presence (POP) that relies only on the "last mile" and LAN segments of the network. These segments are typically dedicated to the user and have more available bandwidth.

Figure 5 Increasing Effective Bandwidth through Content Delivery

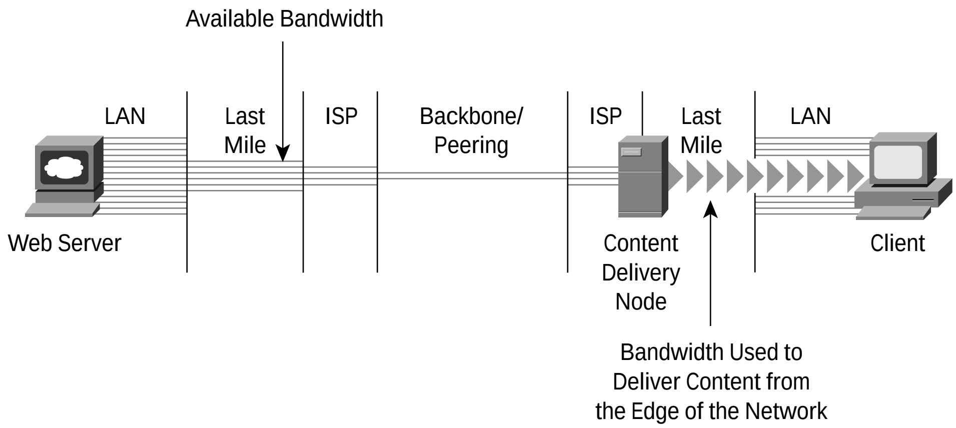

The Elements of Content Delivery

Content delivery addresses the major issues that affect Web speed and utilizes many of the technologies and tools of the current Web acceleration methods. However, a content-delivery system offers a complete, integrated solution that combines these elements, eliminating the limitations of the current delivery methods and the need to integrate multiple technologies from different sources.

Figure 6 The Elements of Content Delivery

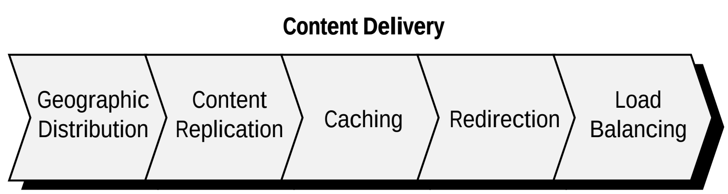

- Geographic distribution—Content-delivery nodes are dispersed geographically around the world on multiple networks.
- Content replication—Content is distributed automatically to multiple content-delivery nodes.
- Caching—Content is stored or cached on the delivery nodes. Content is distributed by a controlled combination of push replication and on-demand pull, based on a set of user-defined policies.
- Redirection—Content requests are processed and redirected to the content-delivery node based on numerous network and delivery node parameters, including proximity (geographic and network), current status, and current load.
- Load balancing—Content requests are spread to multiple content-delivery nodes to ensure scalability, reliability, and fast response.

The Cisco Content-Delivery Networks Solution

Cisco offers the only complete solution for creating content-delivery networks. It resolves the issues of speeding Web performance. Combining a turnkey content-delivery network with legacy acceleration techniques such as data-center load balancing enables a complete end-to-end Web solution and enables service providers to offer new Web services (see the section "Service Opportunities for Content-Delivery Networks" for more details). The Cisco Content-Delivery Network (CDN) solution:

- Enhances Web-site performance for delivery of on-demand content, including static and streaming media
- Significantly improves site performance, including peak load times and Web events
- Lowers bandwidth costs by delivering content from the edge of the network
- Enables premium Web services

The Cisco CDN solution provides the following functionality:

- Serves static content that is pushed to content-delivery nodes located at network edge points instead of the central Web site. Static content includes embedded Web objects such as GIF and PDF files.
- Serves static content that is pulled into the embedded cache. This content also includes GIFs and PDFs.
- Provides a platform to enable on-demand streaming solutions from remote content-delivery nodes instead of central Web sites.
- Provides a platform to enable live streaming media to large worldwide audiences.

The Cisco CDN solution is made up of three components:

- Content Distribution Manager—Network Controller
- Content Router—Request Routing Node
- Content Engine—Content Delivery Node

These components are deployed as shown in Figure 7.

Figure 7 Deployment of CDN Components

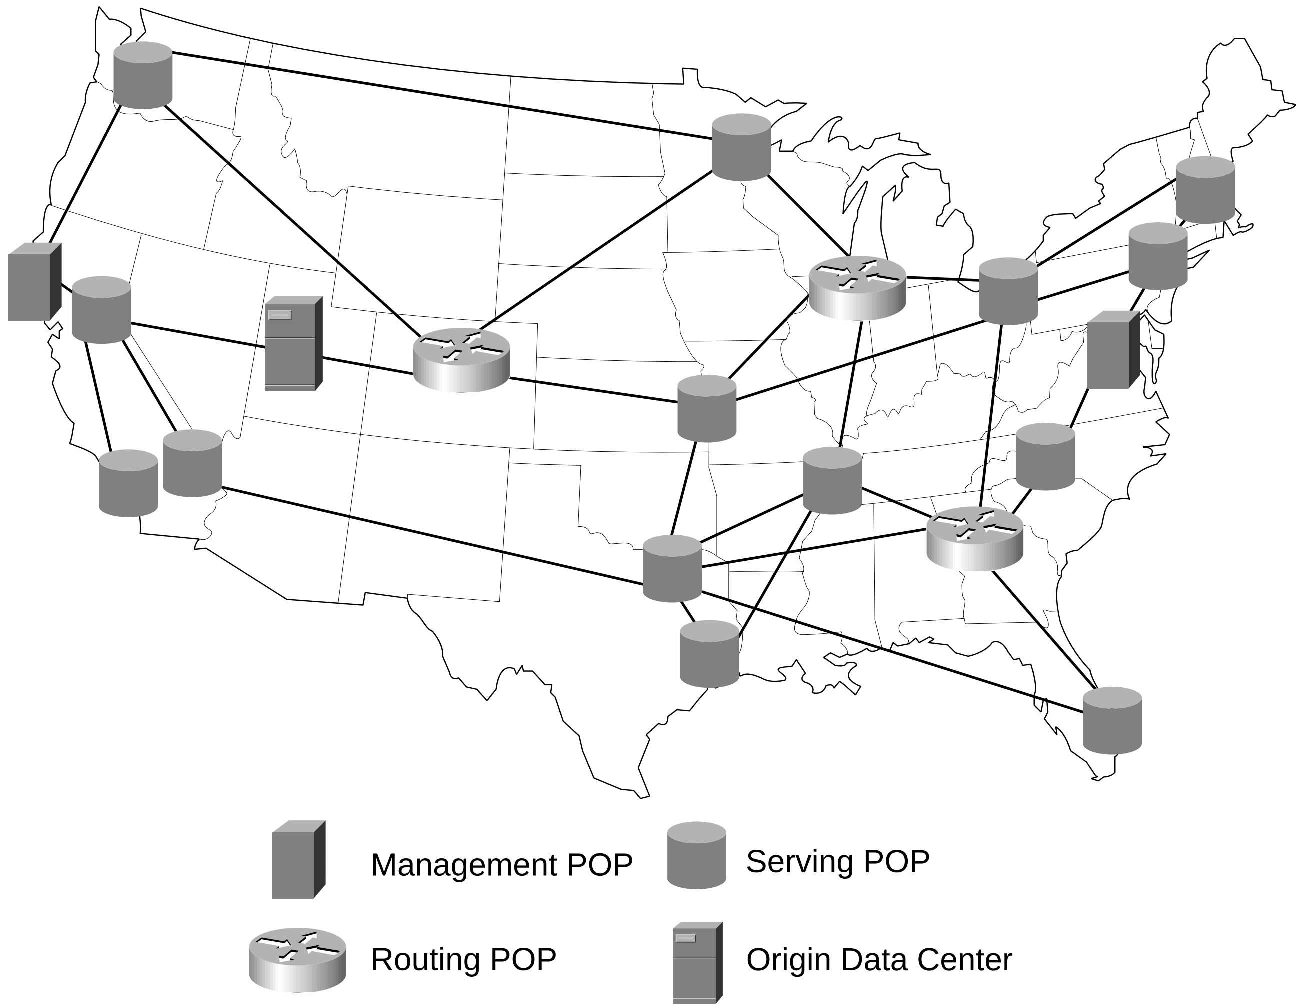

Management POP Deployment

A management POP is where the CDN is managed and is typically a strategically located, highly reliable location within the network, such as the Network Operations Center (NOC). An example of the management POP is shown below in Figure 8:

Figure 8 Management POP

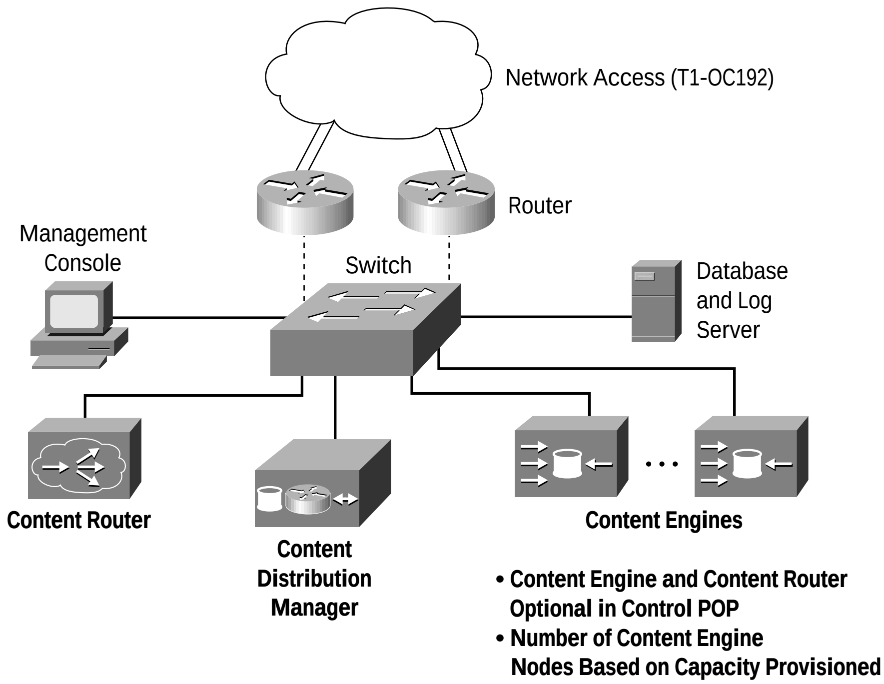

The major component of the management POP is the Content Distribution Manager (CDM), which controls all the policy settings for both content and network devices. This device can be deployed in single or redundant mode, depending on the level of availability necessary. The network will continue to operate when the CDM is unavailable, but the policy changes will not be possible. Also deployed in the management POP is the database and log server used to store network settings as well as usage logs.

Because the management POP is strategically located, a content router (to route user requests) and content engines (to serve user requests) can optionally be deployed.

Routing POP Deployment

The major component of the routing POP is the content router (CR) that routes user requests to content engines (CEs) to deliver content. This device can be deployed in single or redundant mode (either in the same location or distributed throughout the network.) The network will continue to operate when a single or multiple content routers are unavailable, as long as there is a single content router available to service the target domain. Adding multiple routing POPs to the network is as much about performance as reliability. Adding multiple locations enables higher traffic patterns with less delay because the requests are spread across multiple devices in geographically diverse locations.

Because the routing POPs are embedded within the network possibly close to the users, CEs (to serve user requests) can optionally be deployed.

Figure 9 Routing POPs

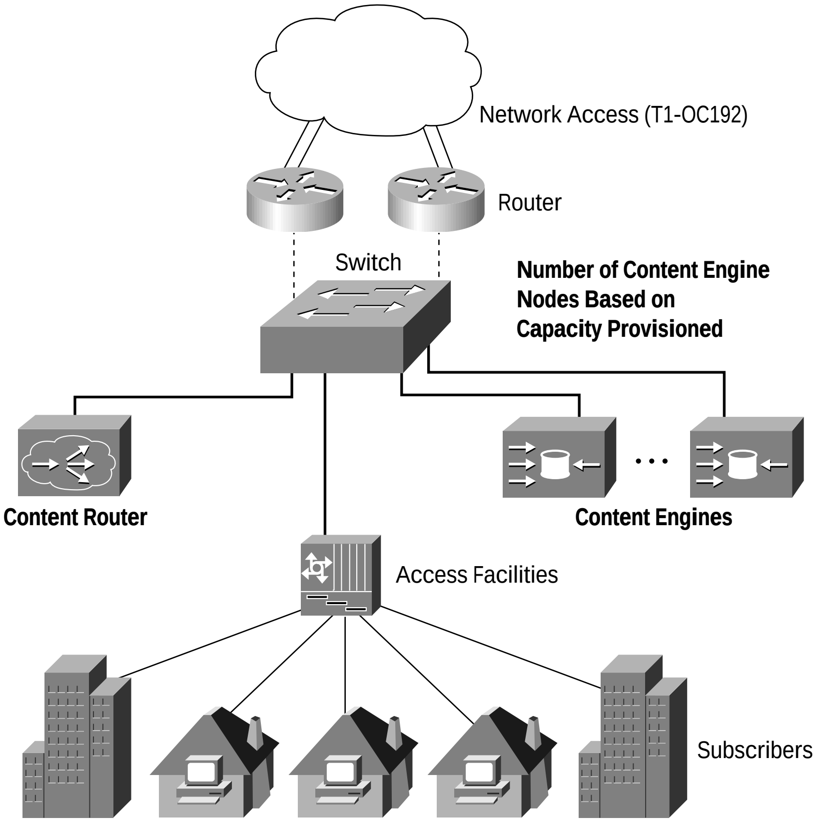

Delivery POP Deployment

The component of the delivery POP is the CE that serves content to users based on redirection from the content routers. Delivery POPs are deployed throughout the network in many locations close to the end users. One or many CEs are installed within the delivery POP. The clusters are provisioned based on the load that each POP is expected to receive. Fault tolerance and load balancing is built into the CDN on both on a global and individual POP basis.

Figure 10 Delivery POP Deployment

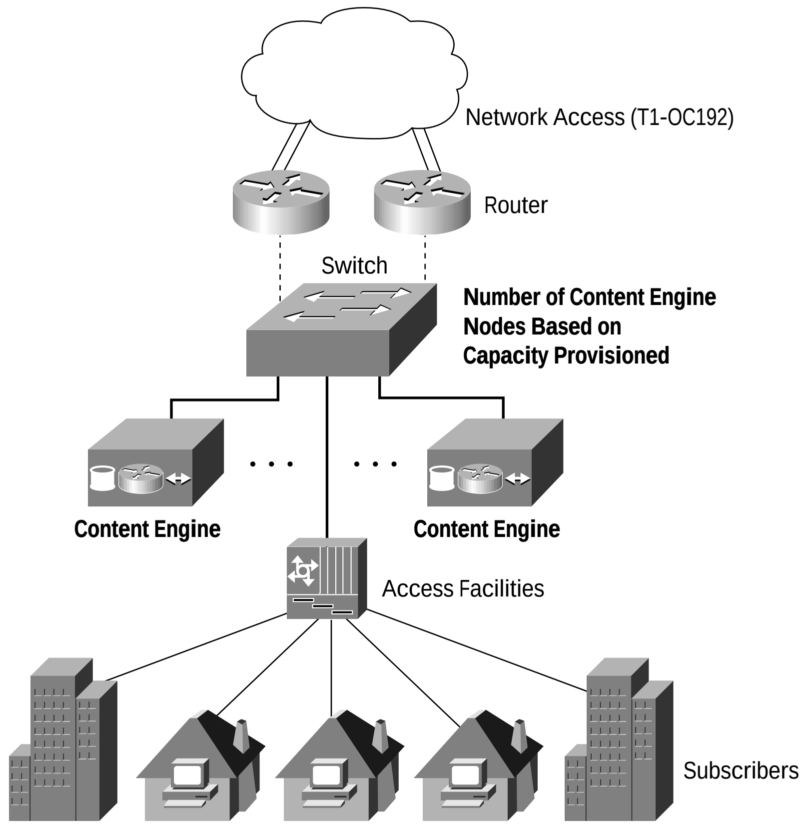

Content Distribution Manager: Network Administration

The content distribution manager is located at the logical central point of the network to control all operations associated with:

- System policies
- Network devices settings
- Content control
- Automatic replication of content
- Interface for live origination
- Billing and event logging

The content distribution manager is deployed either as a single server or in redundant mode; it can either be deployed in a local cluster or distributed geographically.

Other Content Distribution Manager functionality includes:

- Central repository for real-time system monitoring and event logging
- Network management software interface such as Simple Network Management Protocol (SNMP) to integrate easily into most NOC operating software
- Management of content, including registration of Web sites (routed domains) and live streaming to enable acceleration and reliability
- Central management using a Web graphical user interface (GUI) to control configuration, management, and policy settings
- Redundant configuration for fault tolerance
- Publishing tools to enable existing Web sites to easily subscribe to CDN product line without labor-intensive modifications

Cisco Content Router: Request Routing Node

The content router nodes are deployed at strategic locations within the network. Their functionality includes:

- Real-time content request processing using standard DNS by redirecting user requests to an appropriate CE based on geographic location, network location, and network conditions
- Redundant configuration for multinetwork and wide-area fault tolerance and load balancing

Cisco Content Engine: Delivery Nodes

The CEs are located at the edge of the network (such as in ISP POPs) to store and deliver content to users. Other functionality includes:

- Content delivery to end users and other CEs based on the Cisco content routing technology
- Self-organization into a mesh routing hierarchy with other CEs to form the best logical topology based on current network load, proximity, and available bandwidth
- Storage of content replicas
- Endpoint servers for all media types
- Platform for streaming media and application serving

How the Cisco CDN Solution Works

In the basic use case—serving any on-demand content format from GIFs to streaming media—from a remote CE, the solution works in the following manner.

Central Control

- Using Web GUI management, a new customer domain name is registered with the content distribution manager. In addition, the content to be push-published is configured by referencing an easily configured file highlighting the content that should be pushed to the CEs. Content can be either entire Web sites or subdomains.
- DNS is updated to reroute appropriate fully qualified domain names to the content router nodes for the appropriate domain names. The Cisco URL scheme includes either modified URL or delegation of a domain name to the system.
- Domain names or groups of domain names are associated with groups of CEs using distribution networks. A distribution network is a method to map content to the CEs are enabled to serve that content.

Distribute

- Target content is then replicated from the origin Web server to all subscribed and enabled CEs. Replication occurs from origin server to CE and from CE to CE.

Redirect

- When a browser requests content that is under the control of the CDN (either by entering a URL or by having the browser request embedded objects), a DNS request is sent to one of several redundant content router nodes, which then redirects the request to a CE or set of CEs based on the following criteria:
  - Network proximity
  - Presence of requested content stored on the CE
  - Current CE status, including load and availability
  - Current network condition, including load

The CE then serves the requested content to the browser.

The Cisco CDN Features and Benefits

Table 5 Features and Benefits of CDN Product Line

Feature	Benefit
High scalability • 10,000 node networks • Millions of users	• Enables Web sites to address growth while increasing performance
Increases Web performances • Live or On-Demand content	• Provides a better Web experience for both average Web content as well as peak loads associated with live events
Single URL redirection • Central control • Single Web site • DNS and/or HTTP redirection • Content served from optimal server based on network location and current conditions	• Enables lower maintenance cost associated with single centralized Web site • Enables better Web experience by providing content from the best source • Saves bandwidth cost by replicating content once to remote site and serving it multiple times
Format agnostic • Any media type or format	• Provides flexibility to address all Web applications
Complete content caching control • Push population • Pull population	• Ensures that content is available when user requests it
Automatic replication • Bandwidth efficient, spanning tree routing	• Eliminates the management cost associated with keeping servers synchronized • Reduces bandwidth time and cost by using embedded network bandwidth for distribution of content
Virtual distribution networks • Logical grouping of content • Ability to set policy on groups of content of specified domain names	• Lowers content administrative costs • Enables segmentation of content onto EdgeRunner nodes to enable more efficient use of bandwidth and disk space
Application-level multicast • Splits live stream at each network node	• Reduces the amount of bandwidth required to serve live streams while scaling to large audiences
High reliability and fault tolerancy • Dynamically adjusts to network conditions and user load	• Ensures that network resources are always available and that Web site performance is enhanced • Lowers deployment cost
Easy deployment and integration with existing Web sites • Customizable URL deployment	• Enables deployment without modification of origin Web server

SODA: A Key Underlying Technology

SODA stands for Self-Organizing Distributed Architecture; it is the patented technology of the Cisco content-delivery solution. Cisco is unique in providing a single source for technology in the areas of caching, content distribution, load balancing, and redirection. Many other suppliers can provide some of the technology, but only Cisco offers end-to-end turnkey advanced content-delivery solutions.

Self-Organizing Content Engines

Cisco CEs automatically organize themselves into a single cooperating system that can be expanded simply by adding more nodes. When a node is added to the network, it automatically configures itself in the network based on current network topology and performance. Its place in the network is optimized around bandwidth, network congestion, geographical location, as well as its relationship to other nodes.

Each node is part of a sophisticated mesh routing hierarchy that efficiently handles replication of published content. Each node monitors the current state of the network and its peer nodes, to automatically adjust to changing conditions: It automatically understands when nodes have dropped out of the network, whether the network is congested or a node is currently heavily loaded. Changes in conditions can be detected almost instantaneously, and corrections are made to continue operations without service interruptions.

Single-URL Redirection

Cisco provides a minimally intrusive redirection technology that directs user requests to the optimal delivery point in the network. The flexible architecture enables the user to delegate the Web-site domain to the content-delivery system or to do a simple edit to the URL domain name to enable content to be served from the same system.

Entire Web sites can be designated for delivery or just a subset of the content based on a user-defined set of rules (just GIF delivery or all files three levels deep from www.Cisco.com). When the content is flagged for delivery, no changes are required to the existing Web site, and all user requests are routed to the central Web server, which then redirects the request to the optimal CE. This system includes the use of a single Web site that captures all user and impression information while keeping maintenance costs down.

How Content Routing Works

Refer to Figure 10 for redirection steps.

1. Browser requests embedded object with URL of www.sp.com/logo.gif.
2. DNS lookup of www.sp.com is requested from the local DNS server.
3. Local DNS server does not currently have the IP address for the sp.com cached and sends a request to the authoritative DNS server.
4. The authoritative DNS server responds to the DNS lookup of sp.com with IP addresses of multiple content router nodes.
5. Local DNS server sends a request to the content router nodes for the IP address of the www.sp.com server.
6. A content router node returns the IP address of the CE to the local DNS server (where it is cached with a Time To Live [TTL] value).
7. The local DNS server returns to the browser the IP address of the www.sp.com server.
8. The browser requests from the CE the file associated with the URL www.sp.com/logo.gif.
9. The CE receives the request. If the content has been published and is stored on the CE, it is immediately served to the browser. If it is not currently stored on the CE, a pull from the origin www.sp.com Web site is initiated and the content is then served.

On the next request for content associated with www.sp.com, the local DNS server has the IP address for the optimal CE cached, and Steps 3 through 6 are skipped. When the TTL expires for the IP address in the local DNS server, then all steps are repeated.

Service Opportunities for Content-Delivery Networks

The Cisco CDN solution enables service providers to offer higher performance for existing services, and deliver new services that are not feasible with current techniques. Fundamental services upon which applications can be built include:

- High-performance delivery of static Web content
- Reliable, scalable, on-demand streaming media
- Massive live streaming events

High-Performance Delivery of Static Web Content

The Cisco CDN solution can both improve the speed at which static Web sites are delivered and reduce the bandwidth associated with them. Content providers can incorporate richer graphics, documents, software, and other files into their Web environments. At the most basic level, achieving fast site performance is essential for competitive service offerings.

Reliable, Scalable, On-Demand Streaming Media

Streaming media on the Web today is typified by grainy, low-resolution images at best, often jerky and nearly unrecognizable. Moreover, it consumes significant bandwidth, a scenario that can be prohibitively expensive for content providers to absorb into typical business models. With the Cisco CDN product line, service providers can reliably deliver streaming media in a bandwidth-efficient manner, enabling lower costs and an improved viewer experience. With on-site Cisco CEs, service providers can incorporate TV-quality media without impacting their customer's network.

Massive Live Streaming Events

The CDN platform will enable future live streaming services, which are among the most demanding. Scaling them to tens of thousand or millions of users is completely impractical using traditional approaches. The Cisco CDN solution will make these events possible. By splitting live streams and delivering them only from local CEs, the solution offers the scalability, reliability, and bandwidth efficiency essential to their feasibility.

Conclusions

Combining content-delivery networks with traditional Web acceleration techniques results in a complete end-to-end solution to increase Web performance, scalability, and reliability. CDNs offer compelling performance and cost advantages while enabling new and exciting Web services. Content providers and Web viewers will migrate to the highest-performing environments.

The Cisco CDN is the only complete content-delivery solution in the industry, incorporating sophisticated, patented technology that offers network service providers the opportunity to capitalize on this explosive market opportunity.

References

Worldwide Internet Users Estimate and Forecast; Computer Industry Almanac, 1999.

Net Market Size and Growth: Domains, Hosts, and Sites; www.emarketer.com/estats/nmsg_domains.html, 2000

Content Delivery 101: An Introduction to Content Distribution & Delivery; Internet Research Group.

GVU Center's tenth WWW Users Survey; www.gvu.gatech.edu/user_surveys/survey-1998-10/, 1998.

The Economic Impacts of Unacceptable Web Site Download Speeds; Zona Research, 1999.

CISCO SYSTEMS

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

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

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

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the
Cisco.com Web site at www.cisco.com/go/offices.

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

8/00 LW