WHITE PAPER

The Network Architecture Behind NetAid

NetAid was developed by Cisco Systems and the United Nations Development Program (UNDP) to help end extreme poverty. Confronting the reality that more than a billion people on Earth live on less than US$1 dollar a day, the primary goal of NetAid was to raise awareness and foster enhanced communications over the Internet to contribute to the fight against poverty. Three concert events launched the NetAid Web site, www.netaid.org. The concerts were broadcast over the Internet, television, and radio, on October 9, 1999. According to Diane Merrick of Cisco Systems, who developed the NetAid idea, what distinguished the NetAid concerts from other philanthropic events (such as Live Aid and Hands Across America) was the desire to incorporate Internet technology with a live event to effect social change. The NetAid Web site handled large amounts of traffic during the concert event, and is an ongoing forum where those who have goods and services to offer in the fight against poverty can meet nongovernmental agencies and coordinate relief efforts for poverty-stricken areas around the world.

Cisco underwrote the cost of NetAid—managed the overall program—and led the development and coordination of the technology and marketing.

The Concert Kick-Off

Three NetAid concerts on October 9, 1999 kicked off the online effort to combat extreme poverty. Television's VH-1 led a worldwide broadcast event during the 12-hour concerts in Geneva, London, and New York, along with hundreds of radio stations. Numerous television stations broadcast concert excerpts during the weeks following the live performances. The combined television, radio, and Internet coverage reached an estimated two billion people in 160 countries before and during the concert events. Onstage banners at all three concert sites advertised the www.netaid.org Web site address to the estimated 50 million worldwide viewers, while the world-class performers encouraged people to visit the site throughout their shows. And the only place where viewers could partake of the entire event was online at www.netaid.org.

Web Site Architecture and Internet Ecosystem

This paper describes the elements and relationships of the NetAid Web site architecture and how it enabled the NetAid applications. The architecture successfully delivered cost-effective, multimedia content. Record-setting traffic to the site exceeded any previous online event, including the Victoria's Secret fashion show, the 1998 Men's World Cup, and the 1996 Summer Olympic Games. During the concert events alone, the Web site registered more than 40 million hits with a 99.69 percent successful Web page download rate, and supported a total of 2.4 million video streams with a 99.33-percent successful connection rate (versus an average rate in previous events of 40-percent)—all record numbers. Prior to the concert, about 100 charitable agencies were listed on the site. By the end of the concert event, more than 2000 nongovernmental charitable organizations worldwide had added their information to the registry, and more than US$12 million in donations were made online.

In addition to its unique architecture, the NetAid site provides a great example of the Internet Ecosystem and Internet economy business models in action. The prevailing model of competition in the Internet economy is more like a web of inter-relationships than the hierarchical, command-and-control model of the industrial economy. Unlike the value chain, which rewarded exclusivity, the Internet economy is inclusive and has low barriers to entry. Just like an ecosystem in nature, activity in the Internet economy is self-organizing. The process of natural selection takes place around profit to companies and value to customers, or in the case of NetAid, services to those in need. As the Internet ecosystem evolves both technologically and in population, it will be even easier and likelier for countries, companies, and individuals to participate in the Internet economy.

The NetAid Internet ecosystem consisted of technology, application, system integration, and philanthropic partners. Through these partnerships NetAid leveraged the full power of the Internet. The success of the NetAid architecture and technology demonstrated that the business model of working with key partners, each offering unique expertise to run your business over the Internet, clearly works. Organizations around the world can follow NetAid's example of leveraging an Internet ecosystem to expand their sphere of influence.

Ecosystem Technology Partners

From a technology perspective, the NetAid Web site proved the viability of its architecture for sustaining highly available services on a massive scale. The NetAid Web site sets a new standard for scalability, reliability, security, and high availability. Following its own dogma proposing that the potential for success in the Internet economy can be measured by the quality of a company's ecosystem, the Cisco team entered the NetAid project knowing that for NetAid to truly succeed, they needed to form partnerships with industry's best. Cisco led the development effort, which ultimately included the UNDP, KPMG, Akamai Technologies, Real Broadcast Networks, the University of Oregon, and consulting engineers from Red Hat.

Each technology partner brought unique skills to the team, and together they created a synergy resulting in the largest, most advanced Web site on the Internet. Its architecture offers insights and lessons to those seeking to build next-generation Web identities and e-commerce sites that support the latest multimedia applications.

The Cisco volunteer team of 12 members was drawn from several areas of the company. Martin Kagan was the technology lead for the entire project. Thomas Herbst, Director, Consulting Engineering, and his consulting engineers led the design efforts. Chris Lonvick, Manager of Consulting Engineering, was instrumental in assuring security throughout the site. Barry Greene, Consulting Engineer, helped resolve certain issues with Cisco DistributedDirector technology. Greene and Paul Donner, Consutling Engineer, worked on Border Gateway Protocol (BGP) peering issues. Eliot Lear, Consulting Engineer, Kurt Schmidt, Laboratory Administrator, and J.T. Taylor worked in the NetAid Network Operations Center at Cisco. James Aviani, Manager, Software Development, and his team wrote key Apache Web software plug-ins that enabled scalability features with assistance from Sam Gendler, a consultant to Cisco. Jerry Scharf, consultant, worked to optimize the network using DistributedDirector technology. Thomas Conroy, Network Consultant, was responsible for procuring Cisco equipment used throughout the networks.

Architectural Objectives

The architects of the NetAid Web site had to meet multiple objectives, requiring multiple applications. All applications required the network to support high-availability features, starting with solid Internet connections with backup links and devices incase of primary failure. Other levels of network and application failover protection needed to occur within applications and between servers.

Each user activity had specific application requirements, as follows:

- *Provide information*—visitors could learn about the nature of extreme poverty including statistics, find articles on specific poverty-stricken areas and people, and view other types of information. This activity would need to support heavy use of graphics, including photographs, video clips, and audio clips.
- *Facilitate communication*—visitors and charitable organizations could learn more and respond to specific needs. This activity had to provide links to Web sites of both government agencies and nongovernmental organizations.
- *Collect donations*—on behalf of NetAid-registered organizations worldwide. This section of the site required high-level security to encourage user trust, with ample scalability to support thousands of simultaneous transactions and high availability to ensure continuous service.

- *Enable live video streaming*—front stage and back stage television coverage at the three concert sites would be distributed over the Internet as multibandwidth unicast and multicast video streams. The unicasts required enough bandwidth to sustain thousands of simultaneous unicast streams and robust applications to deliver them. The multicast required CPU processor speed and per-point of presence (POP) bandwidth capable of supporting multiple simultaneous transmissions at several delivery rates to enable users to connect at one of several speeds.

Determining Traffic Load

Along with defining the scope of Web site activities, the Web site design team, led by Herbst and Kagan at Cisco, had to determine how much traffic would come to the site during the concert events. From a planning perspective, the anticipated total number of visitors was merely a starting point. As Kagan put it, the team had to plan capacity to meet the needs of "the worst second," that is, predict the size of that spike and provision a site that wouldn't crumble under the load. But how much traffic would arrive?

To some extent, predicting the worst second was a combination of historical analysis and guesswork to derive a reasonable number. Kagan and his team conducted some research, including results of the 1996 Summer Olympics and 1998 Men's World Cup Web sites. Typically 10-percent of the audience moves from one medium to another, in this case, from the television or radio to the Web. Using this math, of the 500 million predicted viewers 50 million of those could go online at some point during the concert event. However, certain factors could have pushed this number higher. NetAid would be the first event of its kind that would actively drive its viewing audience to the Web site from the onstage banner, to commercial advertisements, to encouragement from the artists themselves. Also, the typical VH-1 viewer in the United States was believed to be far more likely to have a PC and Web access than the average American. A third factor that could cause a spike was the first commercial break on television, when people would take the opportunity to go online. Kagan and the technical team ultimately decided to provision a site capable of sustaining up to 60 million hits per hour, one million hits per minute, or just over 16,000 transactions per second, even if as much as half the network failed.

Web Site Architectural Overview

To ensure success, the NetAid Web design team had to develop a highly scalable, high-availability architecture with multiple failover contingencies to alternate links, sites, and servers. Ultimately, the Web site comprised five networks (figure 1).

- *Web pages*—The primary Web site comprised 60 co-located servers in POPs and load balanced with Cisco DistributedDirector. Content was provided by the UNDP, with pages designed by KPMG. The site was managed overall by Cisco.
- *Web graphics*—A specialist in serving Web content, Akamai Technologies had 1200 servers in place around the world that were available to the NetAid project to accelerate all graphics requests.
- *E-commerce*—To collect donations, a separate, highly secure e-commerce site was built by Cisco and KPMG, and managed by Cisco. (Cisco also hosted the e-commerce site during the concert event.)
- *Unicast video streaming*—The Real Broadcast Network in Seattle, Washington, downloaded satellite television feeds (mixed at the concert site in New York) and distributed streaming video to all 60 POP locations, where users could initiate online sessions.
- *Multicast*—The University of Oregon offered its experience with Cisco IP/TV® multicast systems to serve high-bandwidth multicast streams of the satellite television feed over the Internet2 network, which was accessed and viewed at Internet2 institutions.

Figure 1 Overview of Network(s) for netaid.org Functional Divisions

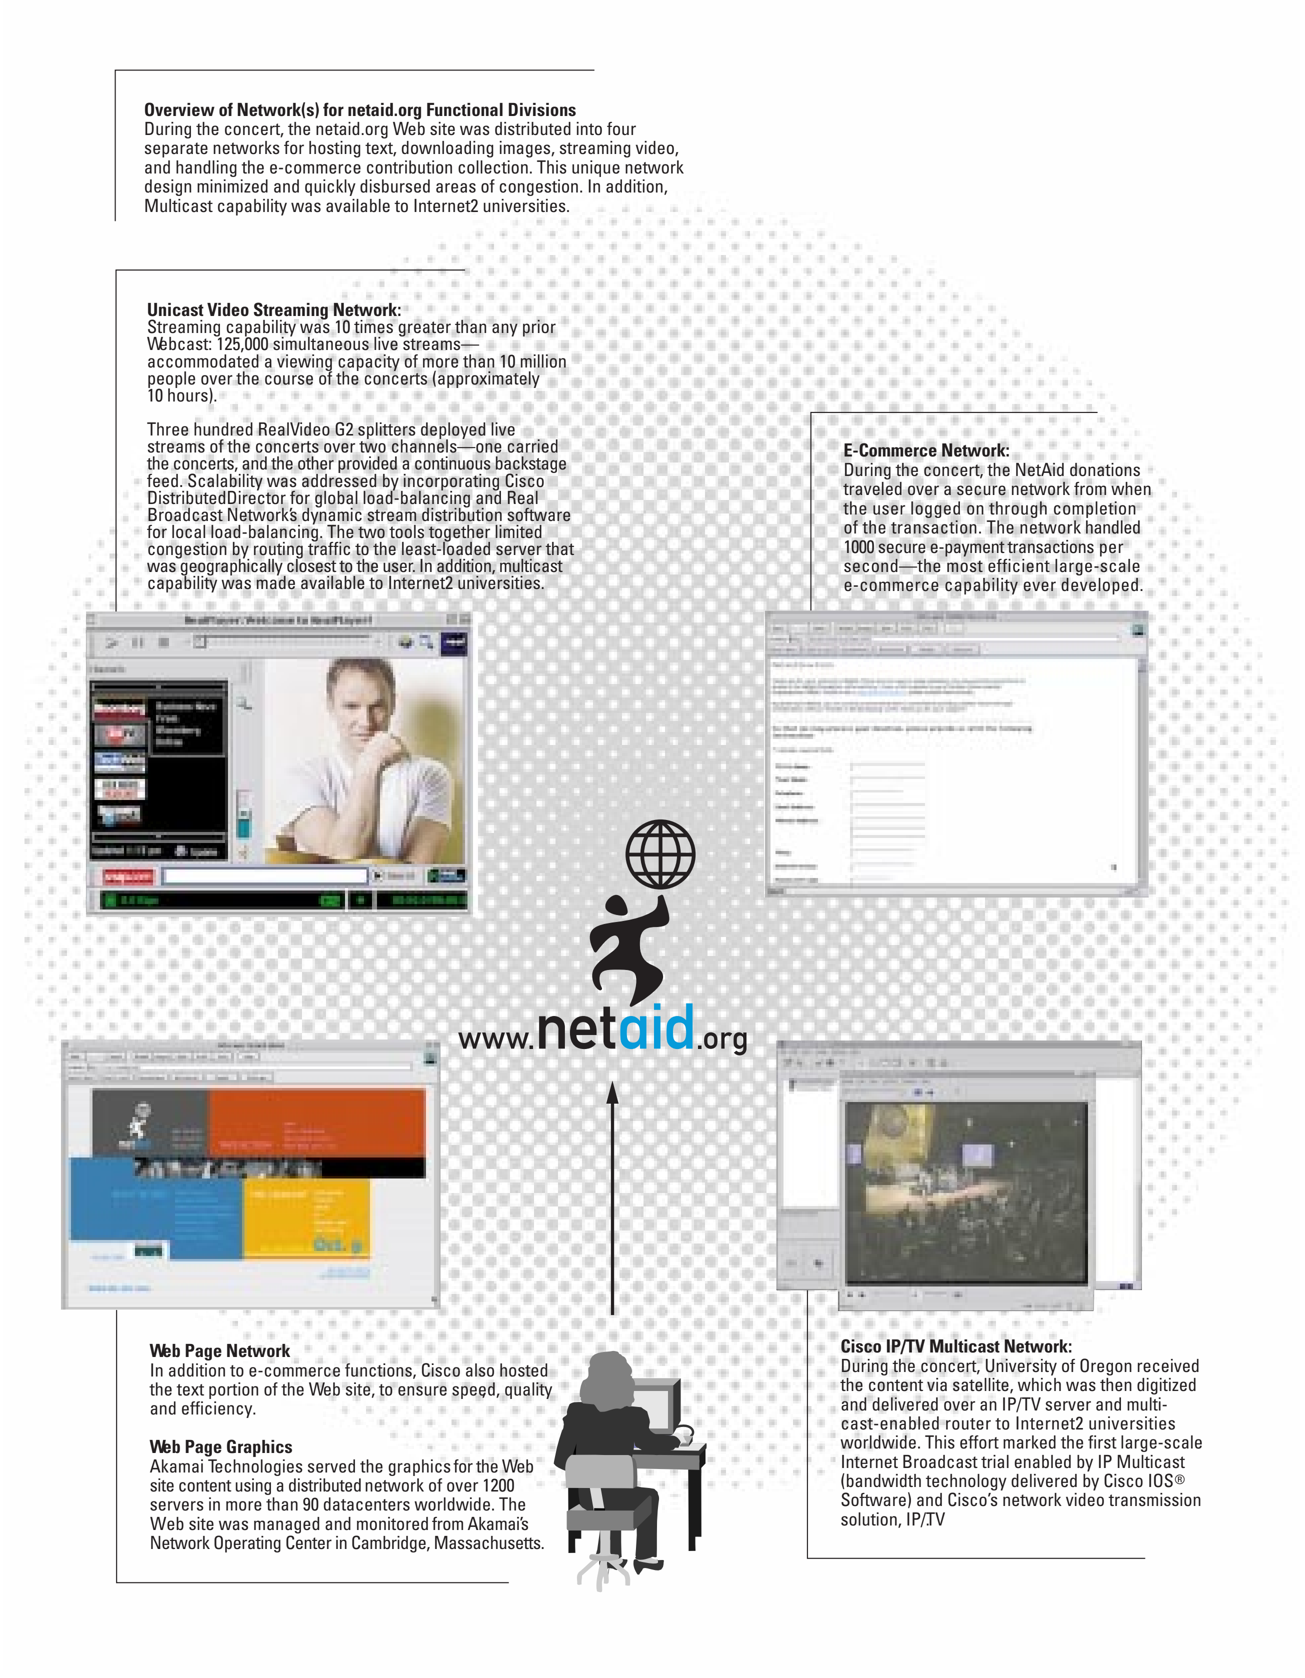

Overview of Network(s) for netaid.org Functional Divisions
During the concert, the netaid.org Web site was distributed into four separate networks for hosting text, downloading images, streaming video, and handling the e-commerce contribution collection. This unique network design minimized and quickly disbursed areas of congestion. In addition, Multicast capability was available to Internet2 universities.

Unicast Video Streaming Network:
Streaming capability was 10 times greater than any prior Webcast: 125,000 simultaneous live streams—accommodated a viewing capacity of more than 10 million people over the course of the concerts (approximately 10 hours).

Three hundred RealVideo G2 splitters deployed live streams of the concerts over two channels—one carried the concerts, and the other provided a continuous backstage feed. Scalability was addressed by incorporating Cisco DistributedDirector for global load-balancing and Real Broadcast Network's dynamic stream distribution software for local load-balancing. The two tools together limited congestion by routing traffic to the least-loaded server that was geographically closest to the user. In addition, multicast capability was made available to Internet2 universities.

E-Commerce Network:
During the concert, the NetAid donations traveled over a secure network from when the user logged on through completion of the transaction. The network handled 1000 secure e-payment transactions per second—the most efficient large-scale e-commerce capability ever developed.

Web Page Network
In addition to e-commerce functions, Cisco also hosted the text portion of the Web site, to ensure speed, quality and efficiency.

Web Page Graphics
Akamai Technologies served the graphics for the Web site content using a distributed network of over 1200 servers in more than 90 datacenters worldwide. The Web site was managed and monitored from Akamai's Network Operating Center in Cambridge, Massachusetts.

Cisco IP/TV Multicast Network:
During the concert, University of Oregon received the content via satellite, which was then digitized and delivered over an IP/TV server and multicast-enabled router to Internet2 universities worldwide. This effort marked the first large-scale Internet Broadcast trial enabled by IP Multicast (bandwidth technology delivered by Cisco IOS® Software) and Cisco's network video transmission solution, IP/TV

Security

Security was a primary concern throughout the design and implementation of every area of the Web site. The NetAid mission would have suffered if any portion of the network failed due to planning oversights or external malicious intent. The Cisco consulting engineering team, headed by Thomas Herbst, began with writing an overall security policy that reflects project goals. Chris Lonvick played a key role by leading the development of the security policy and architecture. They implemented the policy using established products that could assure high availability and scalable performance. The resulting security architecture was based upon previously proven designs that could meet the NetAid requirements. The steps taken to assure security of each of the five network components are detailed in each following section.

Web Pages

KMPG designed the Web pages, and content was provided by the UNDP. These pages were served by 60 Apache Web servers, co-located in 60 POPs worldwide. Each POP had a rack with 12 devices (figure 2). They included:

- One Cisco 7206 router (U.S. installations only)
- Two Cisco Catalyst® 2948 switches
- One Cisco 2511 router
- One Apache Web server running Red Hat Linux operating system software
- Seven RealNetworks G2 servers (for unicast streamed video) running Red Hat Linux operating system software

Figure 2 NetAid POP

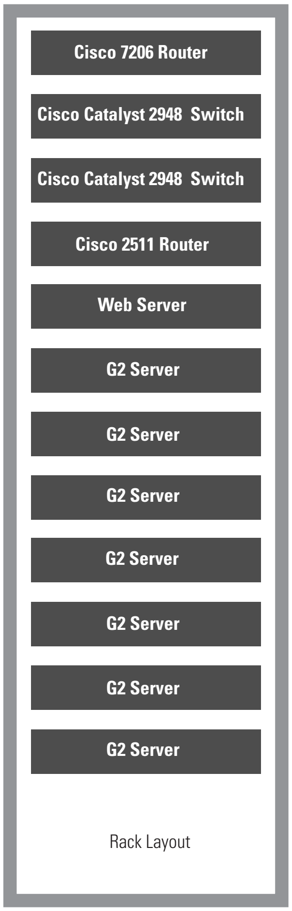

All devices in the POP were interconnected by switched 100-Mbps Ethernet links. All Web servers (serving both pages and graphics) were Intel Pentium-based servers running Red Hat Linux operating system software and Apache Web server software. Pentium-based servers were chosen because of cost efficiencies and post-event reusability and because Cisco wanted to use in-house expertise that was already familiar with the platform. This also proved a good choice because Akamai Technologies loaned more than 400 Pentium servers to the project, making for excellent interoperability.

Red Hat Linux software was selected because of its flexible and open design for tuning performance and security and the ability to recruit consultants from Red Hat and elsewhere to fine-tune the code for the specific applications. In fact the Red Hat consultants, Zack Brown, Ryan Tilder, and Joey Pruitt, were so valuable to the team that they became known as the "Three Musketeers," able to find and solve problems quickly and effectively. Cisco engineers also preferred Red Hat Linux software because of internally available expertise.

Apache Web server software was chosen because of its proven and well-understood characteristics. Reliability was more important than performance, and because Apache software runs on nearly half of the servers on the World Wide Web, it had proven reliability and a widely available pool of experts. Expected performance of Apache-based Web servers serving static hypertext transfer protocol (HTTP) objects from Random Access Memory (RAM) is 500 to 1000 transactions per second. Given estimated traffic load and the requirement that the Web site could accommodate "the worst second" under a 50 percent failure scenario, the site design called for a minimum of 35 servers deployed by October 9, concert day.

Load Balancing

A critical component of the site's scalability is load balancing, which distributed traffic as fairly as possible across the many co-located POPs and again within the POP across the servers. Load balancing helps prevent any single POP site or server to become so overloaded that it would freeze or crash. Global load balancing across the distributed Web server network was done during initial Domain Name Server (DNS) resolution using Cisco DistributedDirector. The domain names netaid.org, netaid.net, and netaid.com are registered with the InterNIC to point to six authoritative name servers. For the World Wide Web (WWW) hostname, these in turn pointed to the six Cisco DistributedDirectors. For high availability, there were three sites in San Jose, California (at Cisco, AboveNet, and Qwest), each with two identically configured Cisco DistributedDirectors, for a total of six devices.

When a Cisco DistributedDirector receives a Internet WWW Domain Name Server (DNS) query, it sends a Director Response Protocol (DRP) message to the distributed network of DRP agents. The DRP agents "bid" on the request and report to the Cisco DistributedDirector on their ability to efficiently service each request. Based on these replies and depending upon the source of the DNS query, Cisco DistributedDirector returns a hostname IP address. The end user resolves the WWW hostname and sends out an HTTP request to that Web server, which in turn sends the requested Web page to the user (figure 3).

Figure 3 Two Approaches to DistributedDirector Deployment

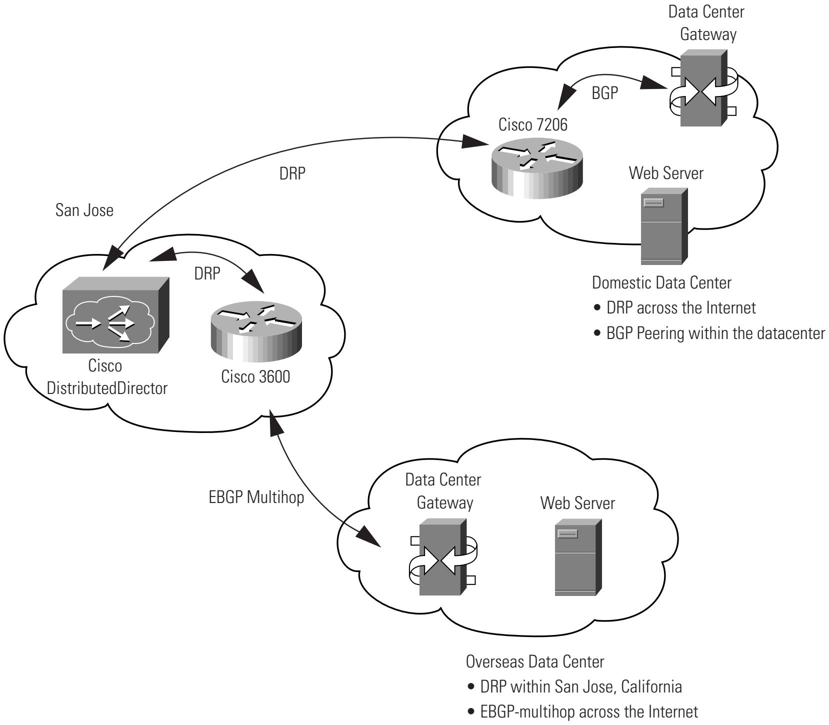
Data Center Gateway
BGP
Cisco 7206
Web Server
DRP
San Jose
DRP
Cisco DistributedDirector
Cisco 3600
Domestic Data Center
• DRP across the Internet
• BGP Peering within the datacenter
EBGP Multihop
Data Center Gateway
Web Server
Overseas Data Center
• DRP within San Jose, California
• EBGP-multihop across the Internet

Each POP with a NetAid Web server also has a Cisco 7200 series or Cisco 3600 series router with an established Border Gateway Protocol (BGP)-peering relationship with the local gateway. The Cisco DistributedDirectors used DRP to communicate with these routers and determine which POP was "closest" to the end user in terms of Internet topology, not necessarily physical geography. Cisco 7206 routers were physically located in each POP for the U.S. racks. International POPs posed a logistical concern, in that routers shipped to certain locations would not clear customs and arrive in time to meet the October 9 deadline. Routers for international POP sites were therefore physically located with the Cisco DistributedDirector clusters in San Jose, but logically connected to their corresponding POP racks via the Internet.

Data Encryption

Another feature of the NetAid Web site invites visitors to register online, either as individuals or organizations. The team believed that the Web server design was probably sufficient to keep data safe. However, that data contained personal information about each donor, so the team determined that additional precautions were required. Personal information requires encryption to protect its confidentiality. U.S. export laws prevented the design team from enabling encryption within Web servers located outside the United States, and each country has its own laws pertaining to the export and import of cryptographic products. The team did not have the time to evaluate each of those laws and obtain the proper permits needed to transfer software to all servers. To solve this issue, an Apache plug-in on each Web server activates a connection with separate domain, join.netaid.org, when a user clicked the JOIN NETAID button. To solve this issue, all of the www.netaid.org servers were programmed to redirect traffic to different Web servers.

Within the United States, the redirection was to one of the 30 "join" servers that were often in the same rack as the other servers. However, outside of the United States, traffic was redirected to one of the 20 "join" servers located at Cisco in San Jose. All server members of the join.netaid.org domain accept user information in cleartext, then use a public encryption algorithm to encrypt data on the server and store it.

Encrypted information stays within each server, providing security without violating the import or export laws of any country. Cisco periodically retrieves that data and delivers encrypted enrollment data to the UNDP, where it is decrypted and processed.

Web Graphics

Akamai was one of the first technology partners to join Cisco in the NetAid project, complementing the existing skill sets with its content hosting expertise. The Cisco team was greatly impressed with the complementary skill sets that existed between the Akamai and Cisco teams. Therefore, the Cisco and Akamai teams collaborated extremely well together.

The NetAid Web site offloaded graphics requests to Akamai servers to reduce traffic load on the primary Web servers by as much as 90-percent.

The Cisco design team was very concerned about their ability to scale the Web site. Graphics on the site could significantly slow performance. Each Web page at the site had nine to 30 graphics. With thousands of potential visitors downloading the same pages at once, there was a strong potential for delay if all content were stored on a single server. Fortunately, the Web offers a workaround for performance. When a browser first requests a page, it contains a plaintext HTML file with text and page layout and references to inline images. Each image is individually requested by the browser and seamlessly laid out on the screen as a single, integrated "page." These inline references can point anywhere.

Inline graphic references on every HTML page in the NetAid Web site were modified during the content-distribution process using the Akamai Launcher script. For example, an HTML tag that looked like:

<img src="http://www.netaid.org/images/netaidlogo.gif">

became:

<img src+"http://a292.g.akamaitech.net/7/292/950/
2db1947852607d/www.netaid.org/images/netaidlogo.gif">

A browser calling up a page at the NetAid Web site automatically requested graphics that were served from the Akamai distributed network of 1200 FreeFlow servers around the world. Each rack in its network contains eight servers and two Cisco Catalyst 2948 switches interconnected by a 100-Mbps Ethernet LAN (figure 4). Load balancing within the Akamai network is handled by their proprietary software.

Figure 4 Web Graphics Network

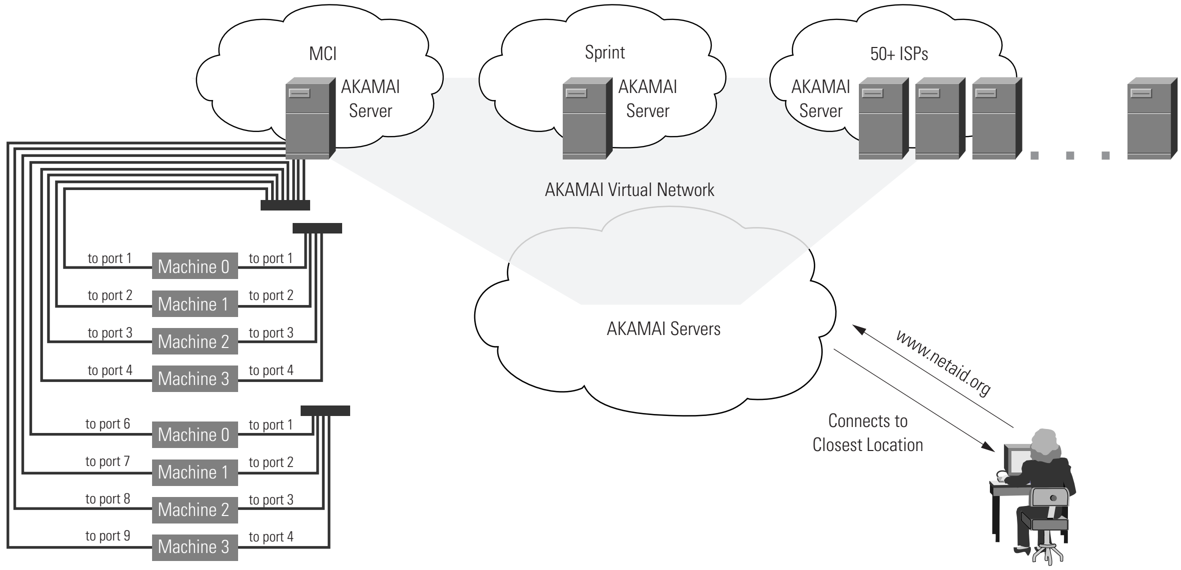

For security, the Cisco Secure Consulting Services team conducted an intensive security scan of the Akamai FreeFlow servers and were satisfied with the results, making no changes. Akamai managed their network the same way they manage their other clients' content, from a 24x7 network operations center (NOC) located in Cambridge, Massachusetts.

E-Commerce

When a visitor to the NetAid Web site at www.netaid.org clicks a button to make a donation, the request is routed to a separate domain, donate.netaid.org. The e-commerce site was designed as a separate domain for scalability and security reasons. The site had to be able to sustain very high traffic and strengthen user confidence with multiple levels of security that protected credit card information. As with the Web pages, anticipating traffic load for the "worst second" helped determine the size of the site. The Cisco design team had a fear that a star such as U2's Bono would tell fans in 160 countries, "On the count of three, I want you all to go to the Web site and donate one dollar." The site was provisioned to support up to 2800 transactions per second, with 28 servers each supporting up to 1000 transactions per second, the performance limit of Rainbow Technologies CryptoSwift secure socket layer (SSL)-accelerator cards in each server.

KPMG brought extensive experience in building transaction-processing systems and took responsibility for designing how credit card transactions were processed. KPMG originally considered a complex architecture that would have required a Java servelet to direct-connect users to the Oracle transaction-processing database. This architecture assumed the need to verify credit card information in real time, before the donor left the site. But unlike many Web sites that must process transactions in real time, the NetAid e-commerce Web site had the luxury of batch processing. Since transactions were all donations, it was not necessary to verify credit card information before a donor left the site. Another factor that simplified the design was that no NetAid merchandise was sold at www.netaid.org. There wasn't any inventory to verify or multiple requests to correlate. Designing a batch-mode architecture reduced the number of lines of code from thousands to a couple hundred.

After reviewing all requirements, the Cisco Consulting Engineering team designed the e-commerce site architecture. To enable scalability, the architecture was split into a front end and a back end. Cisco built the front end, which distributed traffic load and accepted secure credit card donation forms. KPMG built the back-end procedures and Oracle database (figure 5).

Figure 5 E-commerce Network

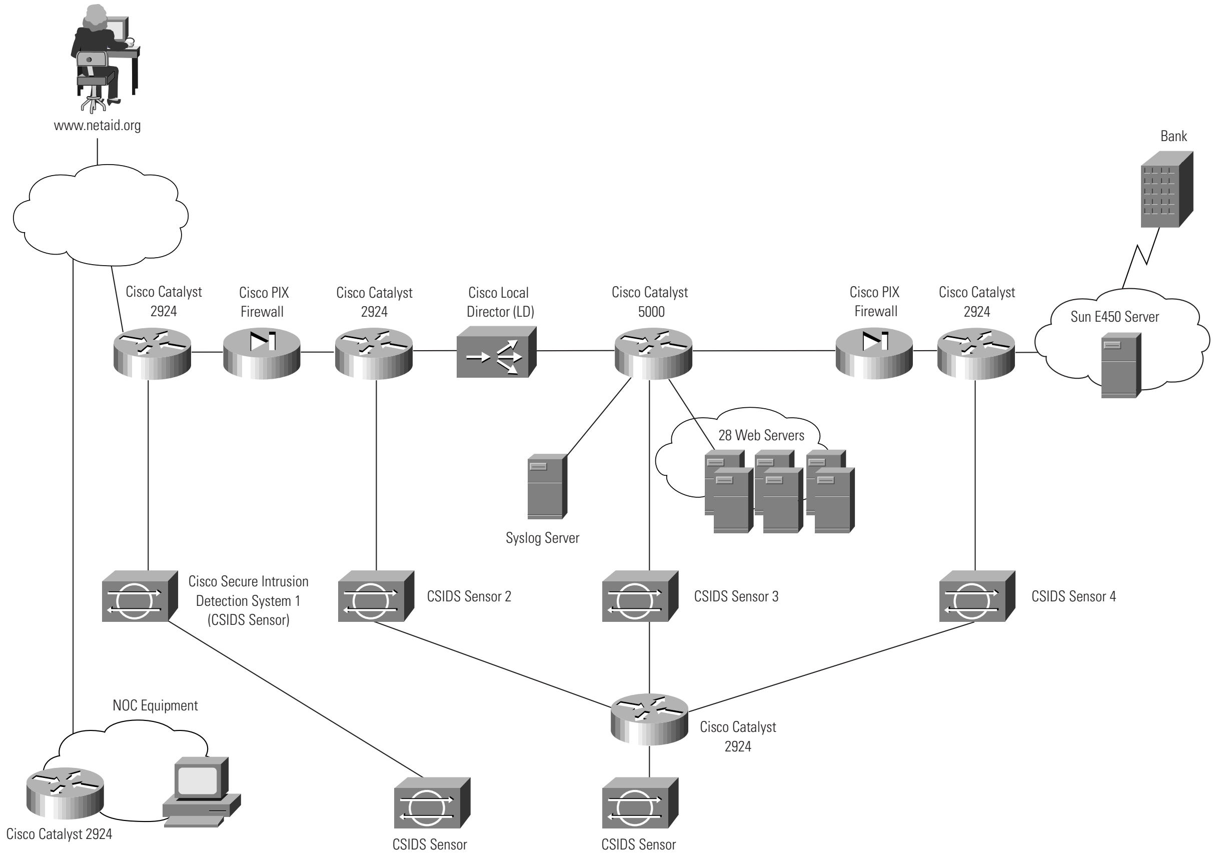

Lonvick and his team chose to be paranoid about the possibility of compromise, and created a robust architecture with several layers of protection. The Cisco PIX™ Firewall separated the trusted e-commerce network from the untrusted, public Internet. Behind the PIX Firewall at the front end were 28 SSL Apache Web servers, each with an SSL-accelerator card from Rainbow Technologies. Like the Web servers, these run Red Hat Linux operating system software. SSL is widely available as the integrated encryption technology in most Web browsers. Each server contained a single Web page that is a donation form with text. Because the single Web page contained no graphics, less bandwidth and CPU was consumed per session to maximize site performance. A Cisco LocalDirector provided load balancing across the server cluster. Donation servers behind the first PIX Firewall were only permitted to accept SSL sessions. Further, SSL servers encrypted donation data and stored it for a short period of time (about 15 minutes). These servers receive and store donations. KPMG loaded an x.509v3 certificate issued by Rivest, Shamir, and Adelman (RSA) encryption onto each donation server to assure donors that they had reached the correct page. The SSL protocol automatically uses this certificate to validate that the site can be verified as the actual donate.netaid.org site.

The "closed" security policy did not allow outbound-initiated sessions; therefore, most security monitoring devices were located inside the PIX Firewall. The monitoring devices were Red Hat Linux devices locked down by removing unneeded services and deleting unneeded code. The Red Hat consultants followed generally accepted practices in their procedures to "harden" these devices. Even though the e-commerce site used a switch in a controlled environment, the network design team set up secure shell protocol (SSH) between all devices for access and file transfers. The Red Hat consultants implemented a TCP-based, syslog protocol for guaranteed delivery of event messages.

A second PIX Firewall segmented the back-end database server from the 28 front-end donation servers. The database server was allowed to periodically initiate sessions through the back-end PIX Firewall to retrieve donation information. Upon verification that donation data was received on the database server, it was deleted from the front-end donation servers. Once a day, transactions were batched and processed via the bank gateway. Verified transaction data was returned to the Oracle server.

Multiple Cisco Secure Intrusion Detection System (CSIDS) sensors watched for undesirable activities in the network. From an operational perspective, the most sensitive area was behind the back-end PIX Firewall where the credit card information is stored. But it was equally important to watch activity on all other segments, including the untrusted network outside of the first PIX Firewall. The sensor outside of the PIX Firewall (essentially on the Internet) was linked to a single CSIDS Director, while the other sensors inside the e-commerce architecture were linked to a separate CSIDS Director. The CSIDS Director outside the first PIX Firewall would signal the first indication of any attacks or probes. The sensors inside the e-commerce architecture would signal any probes or attacks that successfully penetrated the outer PIX Firewall. The sensors were tuned to eliminate false-positive signature matches. This was not difficult inside the e-commerce architecture where operations staff could tightly control both protocols and applications. Under normal conditions, the CSIDS Director displayed green icons, which change color to yellow or red if any anomalies are reported by the sensors.

When most of the e-commerce security was in place, the Cisco Secure Consulting Services team tested this architecture using their Security Posture Assessment (SPA) to review the architecture, then scan and probe the address space trying to find vulnerabilities and any unsafe practices. After the scan, they reported any findings along with recommendations for improvements. Some loose ends that were found in this scan were tied up and the site was opened without any problems. During and after the event, the Cisco Secure Consulting Services team remained on call to assist with security issues. The architecture was locked down several days before the event to provide time for testing and ensure overall stability.

Donation Process

The system handled donations in this way:

- A visitor to www.netaid.org clicks button to make a donation and is redirected to donate.netaid.org.
- The e-commerce cluster establishes an SSL-encrypted session with the visitor, and one of the 28 SSL servers uploads the donation form; the server terminates the SSL session.
- The visitor completes the donation form, then clicks the DONATE button. This initiates a new SSL session with a donation server.
- The SSL protocol encrypts data for transmission. The donation server decrypts it and passes it to the Apache Web server, which processes it.
- An Apache plug-in encrypts data again when writing it to the disk.
- The Oracle server retrieves encrypted data from the donation servers in a round-robin fashion.
- A separate process on the Oracle database decrypts the form, processes the transaction, and adds it to the database. Once there, a process transmits information to the bank.
- Receipts were issued via mail to those who donated US$250 or more.

Cisco LocalDirector Technology

Specific Cisco LocalDirector technology issues were addressed during the design process. First, the original design proposal segmented the SSL servers behind six Cisco LocalDirector clusters. This added bridging and routing complexities and would have required an external "round-robin" method of redirection to spread the load over the six Cisco LocalDirectors. Simplifying this to a single LocalDirector, with a second LocalDirector in failover mode, enabled high availability and eliminated forwarding issues.

A feature called "SSL Sticky" ensures that all traffic from the same user goes to the same server. This is a useful feature in typical e-commerce applications to avoid renegotiation of SSL keys during a transaction. It also allows a typical e-commerce server to maintain continuity of a session during extended time periods while a customer may be making a purchasing decision and filling out a form. Since posting donations does not require a single, continuous SSL session, or even a connection to the SSL server that previously delivered the form to the user, this feature was disabled. Turning off the "cache SSL key" feature on the donation servers increases available memory to streamline overall performance.

Attacks during the Concerts

On the day of the NetAid concert events, the site was hit with a smurf attack. In this type of attack, a machine somewhere on the Internet sends Internet Control Message Protocol (ICMP) Echo Request packets (pings) to a multitude of other hosts using the spoofed source address as the intended target. In this case, one of the DNS servers outside of the first PIX Firewall was the target, and many hosts sent ICMP Echo Replies to it. The DNS server withstood the influx of unwanted packets, but the links to the e-commerce site were saturated with useless traffic. This could have prevented anyone from making donations. The Cisco Information Systems group responded by calling upon all their service providers to place filters against ICMP Echo Replies coming into the Cisco links.

During and after the event, other probes took place. After the event, a smurf variant was launched, where the attacker sent a large amount of session requests to various machines that did not support those services, again with the spoofed source address of the target. Additionally, they sent packets to nonexistent networks that generated "network unreachable" messages. These produced ICMP Unreachable messages that were sent to the target machines, but the amount of traffic was relatively small and ignored by the operations staff.

Satellite Television Feeds

While the television network was not directly part of the NetAid Web site, it delivered content to the two video-streaming portions of the site. Real Broadcast Network and the University of Oregon captured the satellite feed, digitized and encoded it, and delivered unicast and multicast streamed video of all concert events via the Internet.

There were three concert locations in Geneva, London, and New York. All cities had television crews covering onstage performances. Video feeds were broadcast via satellite from Geneva and London to New York, where the producers mixed the feeds into two live shows (one onstage, the other backstage). These were uplinked to the Galaxy 7 satellite for reprocessing into digital video streams on the Internet. Because of the expense of satellite time, there were no redundant feeds.

The front stage performances from London and Geneva were transmitted via satellite to a Don Mischer Productions truck at the New York concert location, where they were produced for the Web. The BBC-hosted backstage interviews from London were carried via satellite to a VH-1 Productions truck, also at the New York location, where they were produced for the Web, along with the VH-1-hosted backstage interview from New York. There wasn't any backstage footage from Geneva because there were no interviews.

The satellite television engineers considered using video compression prior to sending the mixed feeds to the satellite for Web broadcast from the West Coast of the United States. However, this idea was rejected as an unnecessary step that could also have been another point of failure. Another idea that was rejected was to send all five unmixed broadcast feeds (three front stage, two back stage) directly to the Internet so online users would have had more choices. It was decided to uplink two feeds because of cost and management concerns.

Unicast Video Streaming: Real Broadcast Networks

The virtually simultaneous rebroadcasts of the NetAid concerts over the Internet provided two advantages. First, the Internet was the only place where viewers could see the entire 12-hour concert series, because most radio and television stations provided only partial coverage. Second, it gave the video streaming industry an opportunity to earn viewer trust after what became known as "the Victoria's Secret fiasco," where traffic from their online fashion show exceeded the network load limits, limiting its reach and effectiveness. Again the question arose, how many people would watch? Based on initial estimates for overall Web site traffic load, the "worst second" scenario would have 50 million people online with active video streams for the entire event, but this seemed highly improbable. At the low end, MTV reported that they had never served more than 4000 concurrent streams for any televised concert event. Eventually the team decided to provision capacity to support 100,000 concurrent unicast streams across all POP server sites. Actual usage during the concert event totaled 2.4 million streams, making NetAid the most-watched video streaming event on the Internet to date.

Cisco worked closely with Real Broadcast Network to deliver the unicast streams. Real Broadcast Networks has a ready-made network with proven encoding technology that would make content viewable at analog modem speeds with reasonably high quality. The Real Broadcast Networks facility in Seattle has dual satellite downlink systems on separate power sources for full redundancy and high availability. From there, signals passed through a series of encoders and streams sent to all POP sites over the Internet. Users on the Web site could initiate video streaming. For users without RealPlayer software, the Web site also provided a hotlink to download it.

Using proprietary software running on Windows NT, Real Broadcast Networks generated three encodes, each at dual speeds, for a total of six unicast streams (76):

1. Front stage concert audio only (28.8 Kbps and 14.4 Kbps)
2. Front stage concert audio/video (56 Kbps and 28.8 Kbps)
3. Backstage concert audio/video (56 Kbps and 28.8 Kbps)

Figure 6 Unicast Video Streaming “In”

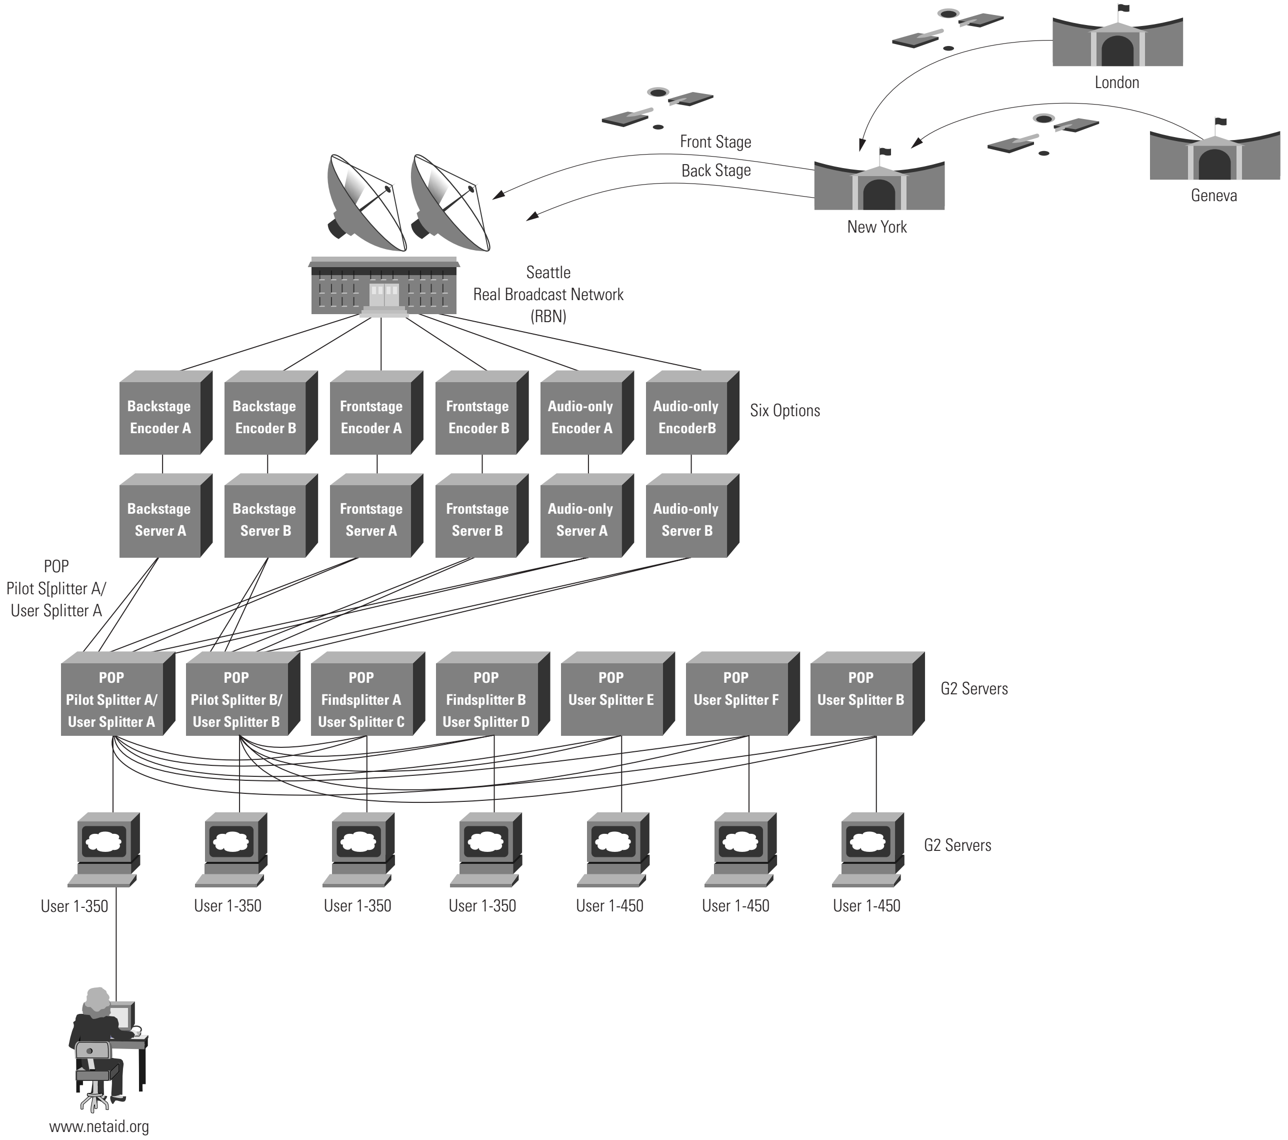

Figure 7 Unicast Streaming Video "Out"

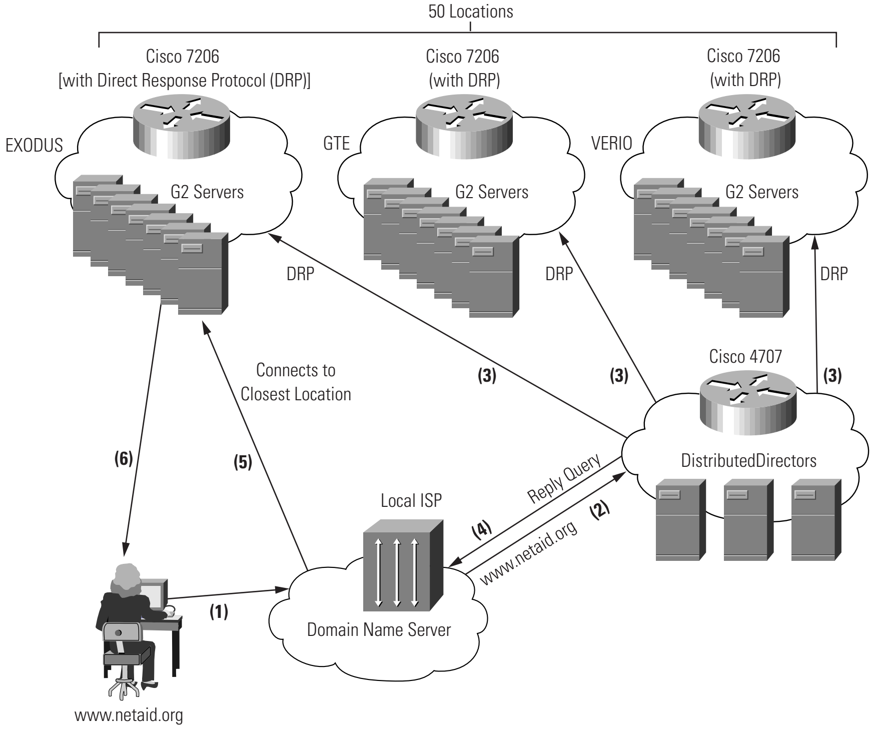

Each POP housed seven G2 servers, all of which ran RedHat Linux operating system software and were fully loaded with the Real G2 software, making any of them capable of serving as a Pilot-splitter, Find-splitter, or User-splitter. The primary configuration contained two Pilot-splitters and two Find-splitters and all seven were configured as User-splitters. Encodes were served by RBN from 12 Red Hat Linux servers in Seattle over the Internet to the redundant pair of Pilot-splitters at each POP location. Pilot-splitters then replicated each stream to the other G2 servers in the rack to conserve bandwidth. Real Broadcast Networks had multiple egresses via several Internet service providers (ISPs) to assure delivery worldwide, which allowed single rack to serve up to 2500 concurrent streams on more that 50 sites. This totaled 125,000 concurrent streams.

Red Hat consultants played an active role in fine-tuning Linux operating system software on the G2 servers and at Real Broadcast Networks to ensure that everything worked smoothly. The Cisco Secure Consulting Services team performed a security test of the Real Broadcast Networks portion of the network and did not find any weaknesses.

By prearrangement, the NetAid video streams were also accessible from other Web sites via deep links. The Web sites that offered front-end links to the NetAid concert video streams were the BBC, MTV, VH-1, and the UNDP. Andrew J. Perez, Director of Operations at RBN, manned the satellite truck in New York, staying in close contact with technicians in Seattle to monitor quality at both ends of the source feed.

A hybrid of Cisco DistributedDirector hardware and Real Networks Find-splitter software provided load balancing across the distributed video network. Visitors to the NetAid Web site clicking a link saw:

Front Stage:	http://play.netaid.org/?url=farm/*/nafront.rm
Back Stage:	http://play.netaid.org/?url=farm/*/naback.rm
Front Stage, Audio only	http://play.netaid.org/?url=farm/*/naaudio.rm

Recalling that the domain name netaid.org was registered with the InterNIC to point to six authoritative name servers for hostname "play," these in turn point back to the six Cisco DistributedDirectors. Cisco DistributedDirector was aware of the two Find-splitter servers in each POP location. Based on a number of criteria ending with random round-robin, Cisco DistributedDirector returns the IP address of a single Find-splitter to the user. The Find-splitter provides local load balancing within each POP location by monitoring the load on each of the other six servers in the rack.

During the events, some encoders and G2 servers did fail, but the high-availability architecture worked. The failover mechanisms kept the NetAid streams active throughout the entire 12-hour event. There were also some predictable interference from sunspots that was compensated for, and a transient failure (less than five seconds) of the fiber backup network from the London concert.

Multicast Video Streaming

The NetAid team teamed up with the University of Oregon located in Eugene, Oregon to deliver multicast video of the NetAid concerts over the Internet2 network. The University of Oregon is a leader in network and Internet research and has established itself as a source of high-quality multicast video. They have been using Cisco IP/TV software since 1996 for University of Oregon-sponsored multicast events. The University of Oregon delivers content over the Internet2 networks using GSR series routers connected to OC-3 lines with a potential audience of millions of students.

Compared to the expense of broadcast television, multicast video is inexpensive and useful for delivering specific content to targeted audiences. Multicast is appealing to both educational institutions and corporations seeking to enable e-learning, company meetings across distributed offices, and other applications.

Joanne R. Hugi, Director of the Computing Center, and Hans Kuhn, Academic User Support Specialist, led the University of Oregon team. The University of Oregon received the NetAid satellite feed into two dishes on campus. Standard video tuners routed signals to the Advanced Network Technology Center (ANTC), where an array of Cisco IP/TV servers captured and encoded live video, which was then streamed over the campus backbone to the Internet2 networks (figure 8).

Figure 8 University of Oregon Cisco IP/TV Multicast Network

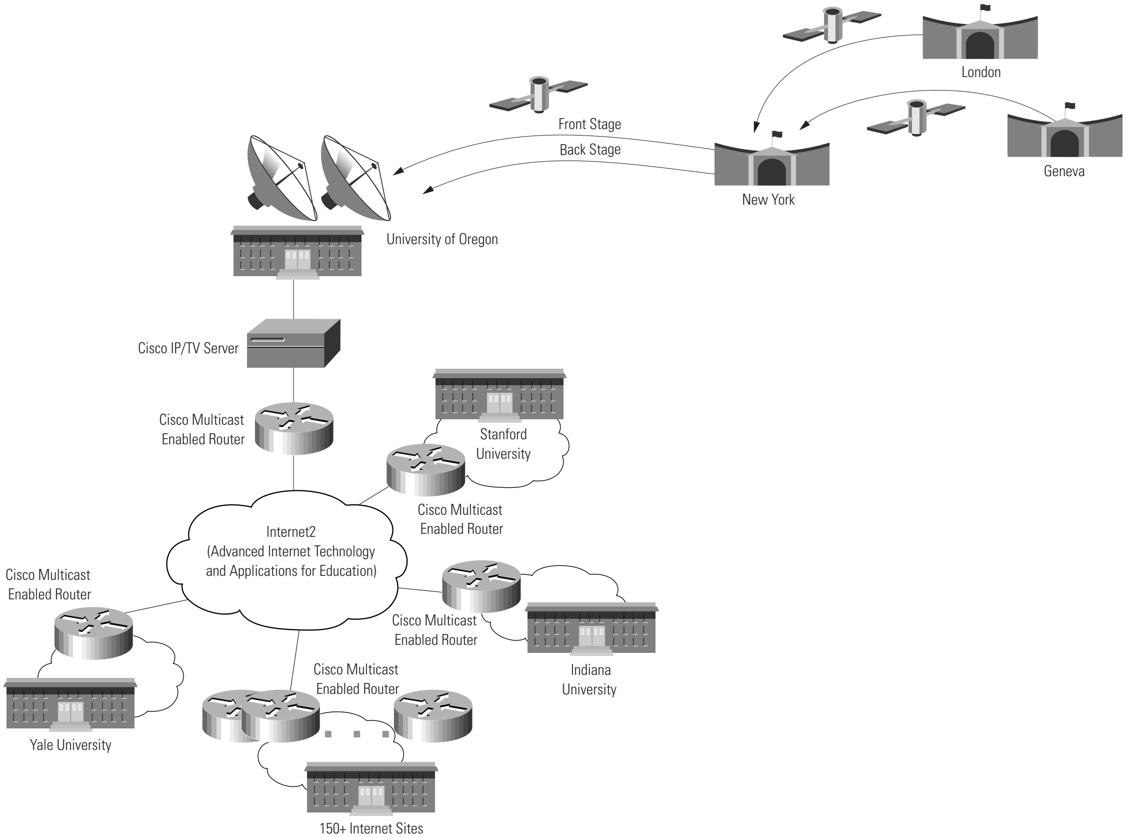
London
Front Stage
Back Stage
Geneva
New York
University of Oregon
Cisco IP/TV Server
Cisco Multicast Enabled Router
Stanford University
Cisco Multicast Enabled Router
Internet2 (Advanced Internet Technology and Applications for Education)
Cisco Multicast Enabled Router
Cisco Multicast Enabled Router
Indiana University
Cisco Multicast Enabled Router
Yale University
150+ Internet Sites

The greatest advantage of multicast technology is that it boosts Internet transmission capabilities by delivering a single stream that is locally replicated near the end user, unlike unicast, which is replicated at the source. So the team at ANTC only needed to provision redundancy for each source.

Cisco technologies for conducting such large-scale, Internet-wide, interdomain multicast are standards-based solutions designed around Protocol-Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and Multicast Border Gateway Protocol (MBGP).

Cisco provided the ANTC team with eight Cisco IP/TV servers. Two of these servers acted as IP/TV content managers, one active and one as backup. Four servers generated the actual multicast streams, and two more servers were on hot standby in case a server failed. The actual content delivered included the front-stage and back-stage satellite feeds from New York, each at two speeds for a total of four streams. The full T1-speed streams used MPEG encoding, while the 325-kbps streams used H.261 encoding. The lower-resolution H.261 stream was transmitted on behalf of users with computer platforms that do not support MPEG streams. Such high-bandwidth multicast technology offers quality comparable to digital video disk (DVD).

The only adjustment that the team had to make to the standard IP/TV configuration was disabling RTP control protocol (RTCP) in all multicast routers. This protocol enables routers to maintain session state for all active multicast clients and so uses processing cycles and memory. The University of Oregon team decided to disable it to avoid any risk of crippling the network in case there were an excessive number of users on a router.

Disabling RTCP made it difficult to measure the number of viewing sites. The University of Oregon team received positive feedback from the University of California at Berkeley, Indiana University, and Georgia Tech, as well as arranged a viewing site on campus in Eugene. There were no reported difficulties with the transmission.

Despite the dual satellite downlinks in Eugene to the ANTC, the University of Oregon also arranged a backup feed via a virtual private networking (VPN) tunnel with Cisco in San Jose, where the NOC team was also downlinking the two feeds. If a feed had been disrupted in Eugene, the ANTC could have switched over to the Cisco feed.

Testing

As components of the NetAid Web site went online, the NetAid design team performed as much testing as time and resources would allow. This key quality-control step helped identify weaknesses and glitches across the architecture and the server operating system software and ensure the strength of the high-availability design. Cisco, Akamai, and Real Broadcast Networks all used in-house, proprietary testing tools to “abuse” or “slam” the network. Cisco primarily tested the Cisco DistributedDirector portion of the network, while Akamai and Real Broadcast Networks tested the server clusters in each POP as they went online. As a result of the testing, a variety of critical software bugs were identified and fixed, and the three Red Hat Linux consultants worked with Akamai and Real Broadcast Networks to identify, patch, and improve operating system performance within servers.

Conclusions

NetAid proved the viability of using high-availability, secure architectures on a massive scale, reaping record numbers for any event-oriented Web site to date. The overall design principles are a model for future, similar endeavors. With a strong emphasis on making NetAid successful through ecosystem partnerships, Cisco was able to bring together a world-class team that built a site to handle crushing traffic loads for a variety of applications. The Internet was the only place where anyone could see the entire event, and the NetAid Web site was the most successful of its kind in terms of serving high-quality unicast and multicast video to millions of viewers worldwide.

The technical success of the NetAid Web site was a result of several factors. First, it was a successful collaboration of talents among the ecosystem partner companies, each of which brought a distinctive strength to the team, yet all of whom shared the NetAid vision for using the Internet to effect social change, in this case to battle extreme poverty. Second, the design team determined to learn from the mistakes of others, not to repeat them. They made every effort to anticipate possible problems and prevent them. Third, the Web site enjoyed the advantages of the best available solutions in networking, content hosting, Web site development, and video streaming, deploying an architecture using robust design and deploying only technologies and products proven to go the distance. Last, the presiding emphasis on a scalable, secure, highly available architecture protected the entire Web site from failure or compromise, so that it was more than able to deliver the record-setting usage numbers generated by the NetAid concert events.

CISCO SYSTEMS

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

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

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

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

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the
Cisco Connection Online Web site at http://www.cisco.com/offices.

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

12/99 BW5728