Stardust.com, Inc.
15575 Los Gatos Blvd., Suite C
Los Gatos, California 95032
Phone: 408-402-0566
Fax: 408-402-0567

www.stardust.com

White Paper - The Ins and Outs of Content Delivery Networks

What they are, why we need them, and how to build them.

Platinum Sponsors

Gold Sponsors

Silver Sponsors

Media Sponsors

Alliance Sponsor

Information Partner

The ins and outs of Content Delivery Networks

What they are, why we need them, and how to build them: A white paper for content providers, CDN service providers, e-business providers, Web hosting companies, ISPs and product vendors

Introduction: The changing nature of Internet content....3

Today's content delivery landscape....4

What is a CDN? 5

Who needs CDNs? 6

The CDN services landscape 14

Industry standards efforts....15

Streaming media and CDNs 16

Conclusion....16

CDN glossary of terms....16

Acknowledgements....20

Introduction: The changing nature of Internet content

The Internet has matured to the point where its value is transcending mere IP connectivity for the support of static Web pages and the exchange of text-based email. Content providers and, publishers and e-businesses are looking to the Web as a high-performance, reliable vehicle for delivering bandwidth-intensive, rich multimedia content such as e-commerce transactions, special events, news and entertainment services.

The efforts to develop dynamic, often interactive multimedia content are shifting the IP networking industry's focus from Layer 3 connectivity issues to the construction of intelligent, Layer 4 - 7 infrastructures that can better support the stringent response-time requirements of these new content types. In particular, the industry emphasis is turning to specially- tuned overlays to the Internet called *content delivery networks or content distribution networks (CDNs).*

Other factors perpetuating CDNs include:

- Speed of development and deployment can be faster at upper network layers
- Need is growing to improve quality of service and quality of experience
- Content providers need to scale to larger audiences who are often consuming the same content

CDNs leverage a strategically-arranged set of distributed caching, load-balancing and Web request redirection systems. They ensure that, based on user proximity and server load, content is served up in the most efficient manner to the user requesting it. This arrangement benefits the end user (and by association, the content provider), as well as network connectivity providers, who reap the operational rewards of streamlined bandwidth consumption.

Generally, content is served from a cache server, also called a surrogate, at the edge of a network, a single hop away from the user. The surrogate is a mirror of the content provider's origin server, located in a CDN service provider's data center. This setup enables CDN service providers to deliver, on behalf of their content provider customers, the best possible experience to end users, who are notoriously intolerant of response-time delays.

Special streaming-media events at a Web site (live sporting events and fashion shows, for example), high-volume e-commerce transactions during holiday seasons and interactive videoconferencing sessions are just a few of the applications driving the construction of CDNs. The special network design requirements of such content are causing increasing numbers of content publishers to consider the expertise and economies of scale offered by CDN service providers.

Keeping these sophisticated content sites strategically dispersed and available to a broad set of users can be a complex and costly undertaking for a content provider, whose primary expertise and business focus is to develop and sell content or products. At a minimum, Web content providers require the following to successfully support their rich multimedia content:

- Near-100 percent server availability and the ability to consistently deliver fast response times to impatient users
- The ability to reach the broadest possible set of customers in a cost-effective, scalable fashion
- Content management tools that allow them to render their content CDN-ready, refresh their content and track user activity at their sites

Today's content delivery landscape

Most large content owners to date have hosted their own content and managed their own Internet connections. However, as use of the Web as a business foundation proliferates and the demands of content distribution increase, many will turn, at least in part, to CDN service providers. The reasons? They can increase the performance and reliability of their content services and lower their total cost of ownership.

Method of content hosting and delivery in enterprises with 500 or more employees

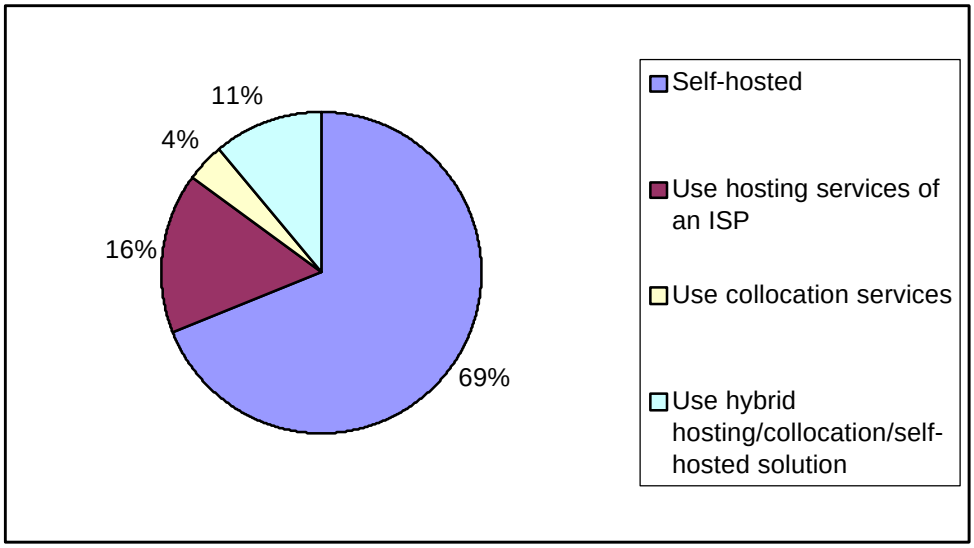

Source: HTRC Group, 2000 Content Delivery Service Study

Figure 1. While most content is currently self-hosted, the scalability and manageability challenges arising with the proliferation of rich multimedia Internet content will eventually drive many content providers into the arms of CDN service providers.

The evolving needs of content providers are driving industry players to develop and standardize new technologies. The trends are also motivating the industry to build and deploy peering and settlement capabilities for the business relationships among CDN service providers necessary to ensure consistent, top-quality service levels. This is a step that has never materialized in the peering environment of Layer 3 Internet access services. In the case of CDNs, though, two industry groups, the Content Alliance and the Content Bridge alliance, are fostering these technical and business standards in concert with the Internet Engineering Task Force (IETF).

What is a CDN?

A CDN is a system, often an overlay network to the Internet that has been built specifically for the high-performance delivery of rich multimedia content. A CDN's raison d'être is to make the Internet a trusted delivery network for mission-critical, content-rich CDN services.

Functional components. The functional components of a CDN that work together to accomplish this goal include the following:

- Redirection services for making sure that a user's Web request is directed to the cache server that is closest and most available
- Distribution services, comprising a distributed set of surrogate servers that cache content and work on behalf of a Web owner's origin server, enabling traffic to bypass congested areas of the Internet. IP Multicast is one technology that might be used as a component of distribution services as a network-efficient vehicle for updating mirrored caches.
- An accounting and billing system that enables the CDN provider to measure, log and bill the content provider based on usage (the amount of bandwidth consumed by users who tap into the content provider's site). In addition, these systems are evolving to enable CDN providers to cross-bill one another for CDN internetworking services.

How do CDNs work? CDNs bring order and quality of service (QoS) to the Internet's IP backbone to eliminate or minimize the infamous "World Wide Wait." Slight network delays or somewhat inconsistent response times go almost unnoticed when the Internet application at hand is text-based email or a static Web page download. In a multimedia world of entertainment services, gaming, live videoconferences and streaming broadcasts, which are sensitive to response-time delays, though, extra steps must be taken to ensure delivery of a quality experience.

CDNs address these stringent response-time demands in large part by minimizing the number of Internet backbones that a Web request and resulting streaming or downloadable content must traverse. One way they do this is by hosting replicas of a developer's content in cache servers called *surrogates,* located within *network edge points of presence (PoPs).* The setup enables CDN service providers, in many cases, to deliver content stored in cache servers just one hop away from the user.

A user request to a content provider's Web site is redirected to a data center of the CDN provider. Typically, the content provider must set up this redirection itself using the tools and encoding method dictated by the CDN service provider (See sections, "Who needs CDNs?"; "Content Providers"; "Editing/redirection of Web links.") The CDN contains a lookup service that steers a content request to the content surrogate that is closest (geographically or shortest travel time) to the user and/or is the least busy.

To determine the content server that is most available to a user at the time of a request for content, CDN service providers make use of load-balancing technology. This technology can be homegrown by the CDN service provider or delivered in the form of router software or a network appliance from a third-party vendor. Load-balancing capabilities

front-end a content server cluster and direct traffic to the least-loaded server, much like a traffic cop steers vehicles around congestion at a street intersection.

Increasingly, CDNs are also making use of *Web switching,* also called *content switching* or *application-layer switching* to further enhance QoS levels. These software capabilities enable the network connectivity elements in a CDN (routers and switches) to examine not only IP address information when determining a best path through the network, but also to take into account the specific response-time requirements of the application or content being requested.

These Layer 4 – 7 switching capabilities can also be delivered in the form of router software from a router/switch vendor or as server software that runs in conjunction with the router or switch.

Meanwhile, the surrogates contain software that extracts logging, usage and billing information that is aggregated up to a central database in the CDN service provider's data center. The CDN service provider uses this information to determine what to charge the content provider for the CDN services rendered. Pricing models are generally usage based, with a fixed fee per megabit per second of usage. The usage information collected and aggregated by the surrogates is also used to deliver content management capabilities, usage trending, and performance information to the content provider customer as a part of the CDN service package.

Who needs CDNs?

There is a value chain of CDN beneficiaries that begins with content publishers, extends to CDN service provider specialists, ISPs, CDN infrastructure component makers, local access telecommunications providers and, ultimately, to content consumers, both in the mass market and in the business environment. In this value chain, generally the CDN service provider "owns" the content provider customer, while the ISP or local access provider "owns" the end customer (the consumer of the content).

Let's take a brief look at each player, why they need CDNs and what their particular interests and requirements are.

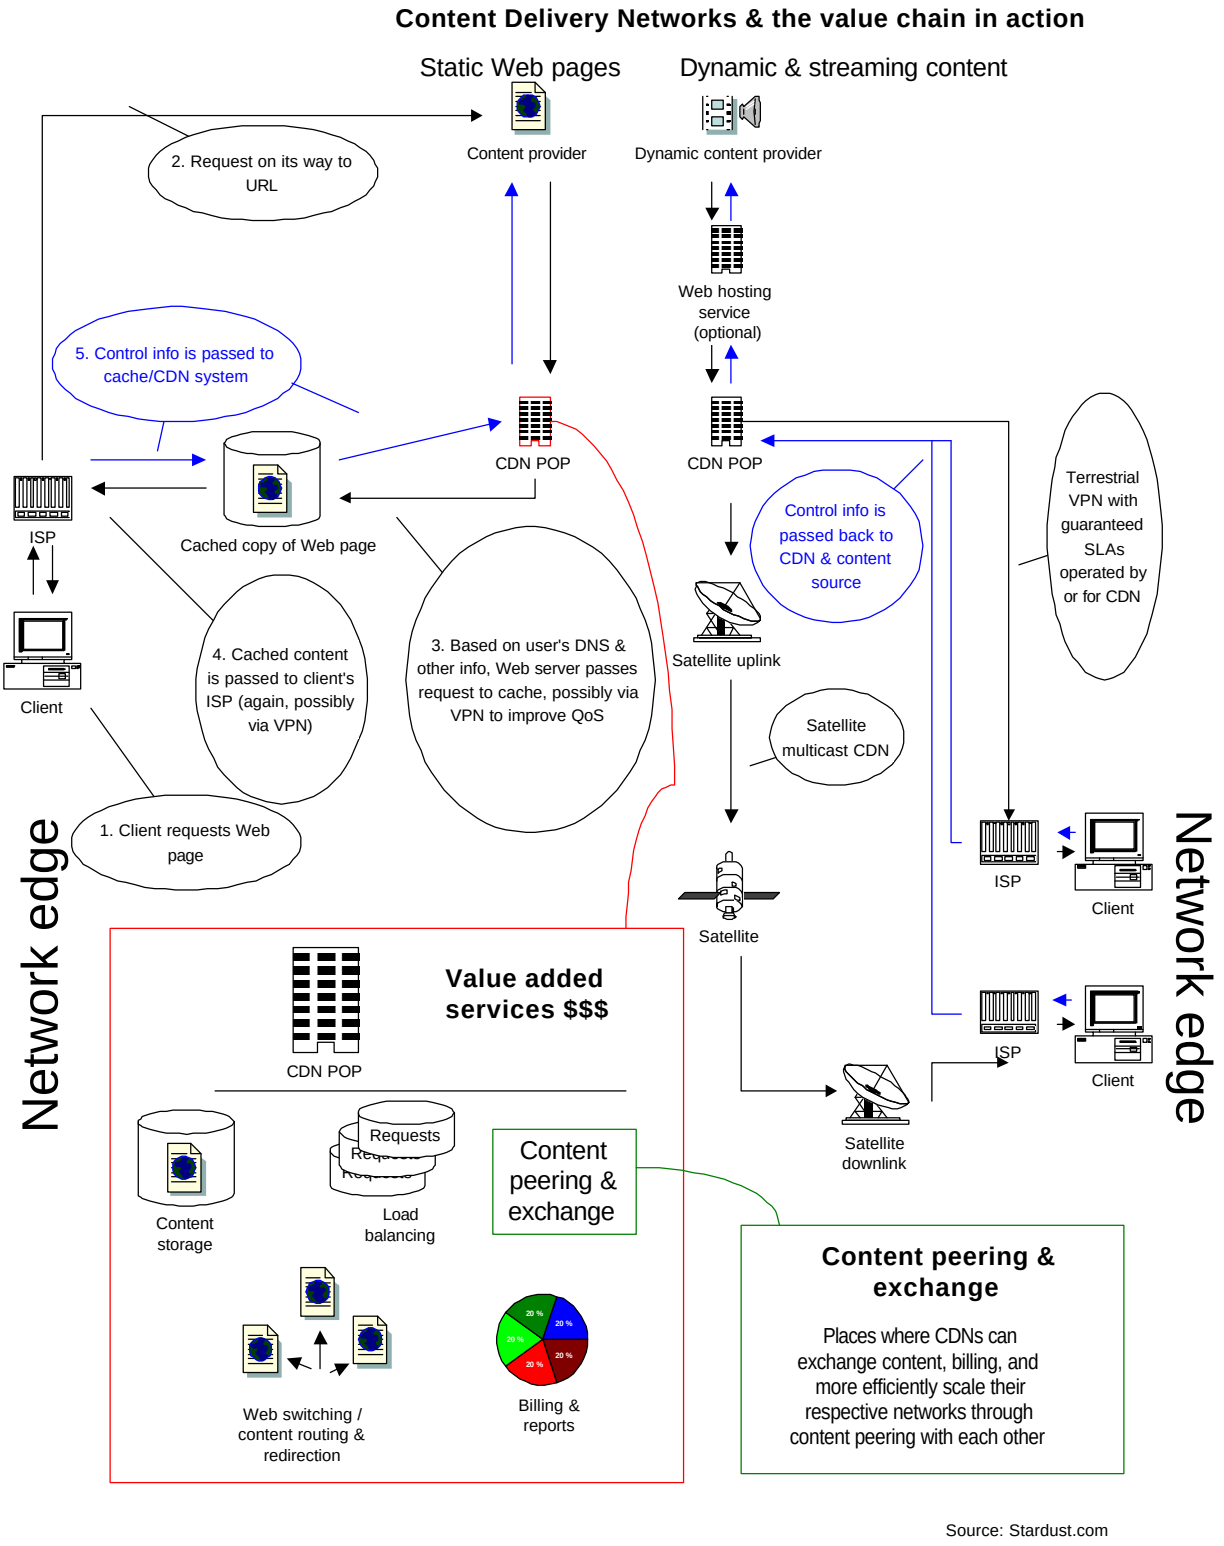

Source: Stardust.com

Content providers

Organizations that build content for the Web are faced with the challenge of delivering content with increasingly dynamic characteristics to customers with consistent, high service levels. As users become increasingly Internet-savvy and demanding, the QoS levels in content delivery can make or break customer attraction and retention. Some examples of organizations that fall in the content provider category are the following:

- Owners of e-commerce sites concerned about response times for browsers and transaction times for buyers.
- Retailers wishing to broadcast a promotional media event from a store
- E-learning developers such as virtual universities or traditional leader-led sales automation training companies that are adding Web-based versions of their offerings
- News organizations wishing to offer Web-based video news coverage
- TV stations, radio stations and entertainment companies leveraging the Web to deliver entertainment services
- Any and all businesses whose Internet-based content is mission-critical

To date, the lion's share of content providers have been hosting and managing their own Internet-connected content sites, sometimes supporting mirrored content servers in two or three locations. The reasons include CDNs being a relatively new and little-known phenomenon and the perception of high costs associated with CDN services.

Today, services from well-known CDN providers such as Akamai Technologies Inc. and Digital Island Inc., for example, list for $1,995 per megabit per second of usage. At a glance, these prices often seem high to content providers, particularly compared with the approximate $600 per megabit per second of usage charged for vanilla Internet hosting and connectivity services from collocation companies such as AboveNet Communications and Exodus Communications Inc. However, as content providers find themselves needing to run multiple data centers to efficiently serve local content around the world, totaling up the hardware resources, network connectivity costs and, most significantly, the human resource requirements to support their 24-by-7 sites makes the build-it-yourself alternative less cost-effective.

Calculations by HTRC Group, a networking research firm in San Andreas, Calif., for example, indicate that as content owners outsource increasing amounts of the content distribution function, their costs decrease while performance actually increases. The reasons parallel those in the typical network outsourcing model: Being able to leverage the economies of scale, skill set and network connectivity breadth of a company in the business of distributing network content for many companies pays dividends to each customer sharing the resources.

What content publishers require from CDNs. Content providers are beginning to demand a combination of hosting and distribution capabilities that are pushing the industry to develop new products and to work together on new technical and business standards. The result has been the formation of two industry groups—the Content Alliance and the

Content Bridge alliance—each following slightly different paths to foster IETF standards for the interoperability of CDNs at the technical and business levels.

- **Multiprovider capabilities.** To date, most publishers who have taken advantage of CDN services have leveraged the services of a single CDN provider. This is in part because the coverage of a single service has been adequate to date for the relatively limited amount of rich multimedia content that has been regularly dispersed across the Internet. However, many content providers are now seeing a need for a broader network reach, but often prefer not to have to strike and manage relationships with multiple CDN service providers. Interoperability among CDNs, whose owners could reimburse one another for reciprocal distribution services, is one way to give content providers a broader reach as well as more flexibility in choosing or changing providers. The interoperability issue is the primary driver behind the formation of the Content Alliance and the Content Bridge alliance.

- **Editing/redirection of Web links.** Content providers making use of CDN services must edit their content to redirect the links in their Web pages to the network of their CDN service provider. Doing so has historically required that they make changes to their naming conventions using the encoding method required by the CDN operator. This is a significant effort for content providers and also tends to lock the provider into the redirection of content to a single CDN—a component that has discouraged some Web site owners from CDN services.

  There are two basic encoding methods used by content providers today to redirect their links to their CDN service provider's network. In one case, CDN service providers require the content provider to refer its Web links in its own Domain Name System (DNS) server to the DNS name of the CDN. A second scenario requires that the content provider encode the name of the CDN service provider into its own DNS—a process called "canonical names" or "C-names." In the first instance, it is easier for content providers to change CDN service providers. In the second, the tradeoff is that generally the content provider will gain richer set of content management and network visibility capabilities, since its server is integrated into the CDN network infrastructure.

  Efforts are under way to make life easier for the content provider. Products are beginning to emerge from companies such as CacheFlow Inc. and Novell Inc. that front-end a content publisher's Web site to perform URL rewriting on the fly. This approach shields the publisher from having to rewrite the Web pages to redirect user requests and precludes the publisher from having to pre-commit to a CDN. For now, these rewriting solutions will work with multiple service providers' CDNs, though they will redirect content to just one.

- **Content management/usage visibility.** The availability of sophisticated, yet easy-to-use content management tools are also a critical success factor for CDN services. Not only do content providers need to be able to redirect their content to CDN sites with minimal effort, they also require access tools that enable them to refresh their content.

For updating distributed surrogates throughout the CDN, most CDN service systems are designed today around a pull model for HTTP objects. The surrogates or cache servers are embedded with an algorithm to detect a mismatch between the surrogate and the origin server. When the mismatch is detected, the surrogates automatically perform an update. In addition, CDN service providers generally supply their content provider customers with a Web-based front end so that the content provider can see what content is being served and purge and update content as necessary.

Content providers also need visibility into usage histories and trends. Those in an advertising-based revenue model, for example, require this information for reporting and promoting their site to advertisers. Content providers also need to see usage trends to determine how to hone their content to make it most attractive to customers. For example, some content providers—or their advertisers—would like to personalize and/or localize content at edge of CDN provider's network to better target users in a given geographical area, much as a local newspaper or Yellow Pages directory would in the print world.

The above capabilities—those that make CDN-enablement more flexible but less invasive on the part of the content provider—will contribute heavily to the acceptance and use of CDN services.

CDN service providers

CDN service providers today include the companies such as Akamai, Digital Island, InterNAP Network Services Corp., epicRealm Inc., Mirror-Image Internet Inc. and Speedera Networks Inc. These companies are basically in the business of bringing management and QoS to what have been largely been best-effort Internet services to date. They are looking to serve the growing demand of their content provider customers to deliver new kinds of content requiring higher service levels to end users. This trend represents a burgeoning opportunity for CDN service providers who can add significant levels of control to the public Internet to build themselves significant revenue streams.

Worldwide CDN Services Market Opportunity			
2000	**2001**	**2002**	**2003**
$97M	$462M	$1.063 B	$2.165 B

Source: HTRC Group, 2000 Content Delivery Service Study

Many CDN service providers have developed their own proprietary technologies for caching, content management and load balancing. Akamai, for example, makes nearly exclusive use of its own technology for optimizing its ability to deliver high service levels to users.

Similarly, InterNAP has built software it calls the ASsimilator that works in concert with the Border Gateway Protocol (BGP) to move traffic directly from the ISP network closest

to the requesting user directly to the ISP backbone to which the content provider customer is connected. To make this work, InterNAP has struck business relationships with the world's largest 11 ISP backbone providers. When a customer sends a request to a Web site, the InterNAP-connected ISP detects it and shuttles the message to an InterNAP data center, which in turn forwards it directly to the content provider's Web site. This setup circumvents ISP-to-ISP network access points (NAPs) for public peering and bypasses much of the congestion of the public Internet.

At this juncture, most CDNs are owned and managed by a single entity. Industry forces, however, are driving a shift to multivendor CDNs to allow content providers and users to benefit by the cumulative reach of multiple providers. CDN service providers that hook into the reach of other CDN service providers, for example, gain a scalable way of adding QoS and connectivity reach without building their own infrastructures to do so. The move to multivendor CDNs will likely result in a new generation of product requirements for CDN service providers.

What CDN service providers require. If standards prevail and CDN service provider players migrate from proprietary home-grown network intelligence to mix-and-match capabilities, some of the software functions they will require from CDN manufacturers include the following:

- **Accounting and billing capabilities** that extend the ability to charge back customers to tracking usage among multiple CDN service providers hosting and delivering content for common customers.
- **Third-party clearinghouse services** for settlement of reciprocal services among CDN service providers. Note: The Content Bridge alliance is currently testing this concept, using Adero Inc. as the third-party clearinghouse in the multivendor delivery of content from America Online (AOL). The Content Bridge was created to prove a multivendor CDN model before defining the technologies to be used in such networks. The group currently has about 10 vendors participating in a real-world multivendor CDN trial. In contrast, the Content Alliance has been working to define technology that would support multiple business models: not only a clearinghouse, but also bilateral private peering.
- **Content signalling technologies** that indicate when content should be invalidated or refreshed and that can be extended across multiple CDNs.
- **E-commerce capabilities** (transaction processing, credit card verification, security) for the delivery of entertainment services such as pay-per-view and gaming

CDN deployment basics and considerations. There are three main architectures in use for deploying CDN services:

- **Facilities-based CDN.** In this scenario, the provider owns its own data centers and can provide network services across a large geographic area to distributed end users. Players in this category might be a large ISP, for example, that has built a CDN overlay to its Internet access and connectivity services.
- **Distributed or multinetwork CDN.** Here, a CDN specialist places its CDN servers in the PoPs of as many facilities-based providers as possible, creating an internetwork

of CDN servers that cross multiple ISP backbones. This model, also called a *meta*-CDN, is currently the one used by Akamai and Speedera, for example.

- **Hybrid CDN.** Companies in this categories category run some of their own facilities and also leverage the infrastructures of other ISPs or CDN service providers. Digital Island, for example, put itself in this category with its acquisition of ISP Sandpiper, Inc.

Network service providers

ISPs and hosting/collocation companies primarily in the business of selling IP connectivity services today are seeking ways to differentiate themselves and add new revenue streams. Meanwhile, ISPs running Internet backbones are already looking heavily at CDN technologies such as load balancing and caching purely for the operational benefits they afford. Smarter, more precise bandwidth management using these technologies is simply good business practice for cost-conscious network service providers looking to keep bandwidth costs down while improving traffic flow for their customers.
As long as they are deploying the technologies anyway, such companies can decide to become full-blown CDN providers in their own right through content peering internetworking arrangements with existing CDNs. Or they can adopt cooperative business models whereby they "plug in" their networks to one or more CDN infrastructures. One business model might be for them to buy "edge" services from CDN providers for delivery of broadband content to their own customers.

In this way, for example, *local access providers* delivering content in the last mile and already using caching and load-balancing capabilities for internal operational benefits could participate in the CDN value chain. Alone, they might not have the network coverage or clout to catch the attention of a CNN or MSNBC, but could get paid for the ultimate delivery of the content from those providers.

Satellite-based network service providers also have a place in the value chain by delivering IP content directly to a local access provider's network edge or even directly to a business or consumer site. For example, satellite company Hughes Network Systems recently announced that it will integrate Inktomi Corp.'s Traffic Server network caches at its worldwide network operations centers (NOCs) to enable the company of provide fast, efficient delivery of IP-based applications directly to businesses and consumers.

Hughes and Inktomi will also cooperate on the development of new satellite-optimized caching and content distribution software that will be embedded within hundreds of thousands of Hughes satellite receivers in homes and businesses worldwide, and will operate as remote extensions to the Traffic Server caches deployed at the NOCs. Finally, Hughes said it will bundle Inktomi Traffic Server caching software with its DirecPC satellite platform to provide an integrated end-to-end caching and content distribution solution.

What network service providers require from CDN service providers and CDN component product makers. ISPs require business relationships with CDN providers to

work in tandem for the expedited content delivery required. In the InterNAP case, for example, direct connections run from InterNAP data centers to major global Internet backbones run by UUNet, Sprint, Cable & Wireless, Genuity, Digex, PSINet, AT&T, Verio and Earthlink These connections are not free public or private peering; rather, InterNAP pays each of these backbones for full transit TCP/IP connectivity. With an economic relationship in place between InterNAP and the major backbone providers, InterNAP is able to reliably provide its customers with differentiated quality of service.

CDN product manufacturers

Because there is a large market for the creation of CDNs, it follows that there is a significant opportunity for the makers of CDN network elements. The customers of the CDN product manufacturers are primarily CDN service providers, ISPs wishing to become CDN service providers, network service providers looking to gain operational efficiencies in their own network operations using CDN technology and enterprises wishing to build corporate CDNs for internal content-rich applications.

Worldwide CDN Equipment/Software Market Opportunity* (in $ millions)			
2000	**2001**	**2002**	**2003**
101	246	436	749

Source: HTRC Group, 2000 Content Deliver Service Study

* CDN-specific products only

CDN product manufacturers need to pay attention to the IETF activity on content distribution internetworking so that they can use industry-standard technologies, where appropriate, for interoperability and off-the-shelf manufacturing economies.

To stay on top of these activities, most CDN product manufacturers belong to one or both of the industry content alliances – the Content Alliance or the Content Bridge alliance. The idea is typical of the standards process: fostering standards will accelerate CDN proliferation, which, in turn, will allow CDN product manufacturers to sell more equipment and software. Both alliances have joined in supporting the IETF as the forum for the development of interoperation among CDNs.

Enterprises

Enterprises are in line for eventually needing the same basic CDN infrastructure capabilities as public content providers. The reason is that the nature of information being distributed by IP technologies within the enterprise is following much the same route that it is across the public Internet. Enterprises, for example, are developing their own e-learning content for internal distribution or may download or stream it directly from the Internet or from a satellite link to one or more sites for local distribution over the corporate IP network. Similarly, for highly distributed enterprise sites, enterprises might wish to tap the economies and reach of the public Internet for videoconferencing and bandwidth-intensive collaboration capabilities.

What do enterprises need? Like a CDN service provider, an enterprise wishing to run its own internal CDN is in the market for similar caching, redirection, load-balancing, and content management tools as those deployed by a CDN service provider.

Consumers

Consumer thirst for tapping entertainment, gaming, live distance-learning, and videoconferencing capabilities using their Internet connections is driving the content providers to build such Web-based capabilities. Generally, the consumer as ultimate customer is "owned" by the consumer's Internet access provider, which may or may not also be a CDN provider.

The CDN services landscape

One reason that most Web site owners have typically hosted their own content is that CDN services are fairly new, and many content providers do not know much about them. In addition, regular use of the Web for the delivery of rich, multimedia services that require significantly higher QoS levels is just beginning to pick up steam. Finally, technical and business standards for CDN service peering and interoperability are still evolving.

CDN industry players include the following types of companies:

- Product vendors that manufacture CDN components, such as caching, load-balancing, and redirection software and network appliances.
- Facilities-based CDN service providers—ISPs that overlay CDN elements onto their Internet backbone services to become CDN providers.
- Multi-network CDN providers—companies that place CDN servers in multiple ISP networks to achieve a broader, cross-network reach.
- Hybrid CDN service providers—those that own some of their own network facilities but also leverage those of an ISP.

Content creation/ personalization tools	Interwoven
	ATG, Broadvision, Documentum, Open Market, Vignette
Streaming	Akamai Technologies Inc., Digital Island Inc., Ibeam, Infolibria, Inktomi Corp.
Content caching solutions	CacheFlow Inc., Cisco Systems Inc., Inktomi Corp., Infolibria, Novell Inc., Persistence
Content distribution services	Akamai Technologies Inc., Digital Island Inc., InterNAP Network Services Corp., Epicrealm Inc., MirrorImage, Speedera
Load balancing	Cisco Systems Inc., F5 Networks Inc.
Traffic monitoring & analysis tools	Keynote Systems, Reasonate, Radview
Content storage/hosting/collocation	AboveNet, Epoch, Exodus
Satellite	Cidera, Edgix, Ibeam, Loral Cyberstar, PanAmSat's Net-36, Orblynx
Data mining	ATG, E-Piphany

Sources: Witsoundview.com, Stardust.com

Industry standards efforts

The CDN standardization efforts address both the technical requirements for multivendor peering as well as standard methods of billing and settlement. As of this writing, the Content Alliance and the Content Bridge alliance have made significant progress with convincing the IETF that a formal IETF working group is needed to define and formalize CDN technical specifications. Such specs have already been submitted to the IETF for consideration by both alliances. These activities could lead to specifications for multivendor CDN product interoperability in about a year's time, according to individuals involved in the efforts.

Technical issues surrounding content delivery were a major focus at the December 2000 IETF meeting in San Diego, Calif. A CDN "Birds of a Feather" (BOF) meeting served as a precursor to the potential formation of four new IETF Working Groups:

- Content Delivery Network Peering (CDNP): Concerned with specifying how CDNs run by different operators can share the information necessary for their CDNs to interoperate across administrative boundaries. Among the specs defined by this working group would be those for how to track usage and exchange billing information across network borders.
- Open Proxy Extension Services (OPES): Concerned with defining standards for how proxy caches would execute code to enable special services, such as redirection to a server storing foreign-language content. OPES standards, among other things, would specify how to encapsulate content and communicate with servers.
- Contextualization of Resolution: Focused on extending DNS naming capabilities to handle more complicated services, such as identifying an IP address of a server storing a foreign-language version of a particular news organization's content.
- Web Replication and Caching: Charged with specifying a common way to replicate, broadly distribute and store copies of content in servers located in many broadly dispersed sites at network edges. Having standards here would give CDN service providers a greater variety of choice among vendors for their CDN components.

At this writing, the IETF's decision whether to officially form these working groups is imminent.

The Content Alliance

The Content Alliance was formed in August 2000, spearheaded by Cisco Systems Inc. The 62-company consortium has taken the approach of first defining and standardizing CDNs technologies that will then be implemented in CDNs. To this end, the Content Alliance formed an internal design team called the Content Peering Working Group. If the IETF sanctions the formation of the working groups described above, the Content Peering Working Group would likely be absorbed into the proposed IETF's formal CDNP. The Content Alliance has a stated goal of supporting a variety of business models.

The Content Bridge alliance

The Content Bridge Alliance, founded by Adero Inc. and Inktomi Corp. in August 2000, is a smaller group (about 10 companies). It is composed of members who have actively participated in building a CDN service with limited commercial availability using Content Bridge-specified to technologies. This group is in the process of proving out the CDN business model by offering multivendor CDN services, including settlement services from Adero, on a small scale before deciding which technologies should be standardized.

Ultimately, the CDN standards efforts of both alliances will merge under the auspices of the IETF.

Streaming media and CDNs

For providers of streaming media content, the content delivery challenge is compounded by incompatible format and bit-rate requirements at the edge, and the varying platforms required by the servers for these streaming media. For instance, Windows Media Player origin servers are based on Windows NT or Windows 2000, while RealPlayer origin servers may be Unix-based. MPEG2 or MPEG4-based origin servers run on every imaginable platform.

Typically, Webcasts have presented a limited selection of bit rates for consumers, that may not be the optimal use of bandwidth at the edge, depending upon the nature of the connection at the edge, on time of day, and on other factors.

This is triggering the growth of storage products and services as integral components of CDNs. The more efficiently such products and services can store and forward multiple formats and bit rates of streaming media, the broader the range of streaming media offerings will be on a given CDN.

For instance, a number of new companies have sprung up to offer “streaming media appliances,” lights-out boxes for collocation facilities at the edge, that offer set-and-forget servers providing streaming media in Windows Media, Real, QuickTime or MPEG formats. Such companies include:

Vividon - http://www.vividon.com
Midstream - http://www.midstream.com
Vingage - http://www.vingage.com
Network Engines - http://www.networkengines.com

Secondly, services have begun to appear that will rationalize the delivery of streams from origin servers to edges capable of providing both a variety of popular media formats and a variety of bit rates.

Software that is enabling such services is being provided by vendors such as AnyStream, whose recently announced Agility Edge allows in-background reencoding of streaming content at the edge to adjust to changing network congestion conditions or consumer

demand for varying formats. This adaptive approach can minimize the work a CDN has to do to deliver the right format at the right bitrate at the right time.

Streaming media CDNs continue to face various challenges, the most notorious being "flash" crowds. According to officials at Navisite, the November 2000 Madonna Webcast involved the construction of the largest-ever FastForward media bridge network ever. But officials note that traffic spikes are wildly different than those observed for on-demand usage of content. Subscribers to live events need to be better educated about IP Multicast, which is best-suited to deal with the spikes of live events; and that unless customers ask for multicast, service providers will continue not to provide it.

Conclusion

The Internet is carrying increasing loads of mission-critical, rich multimedia content. The dynamic and delay-sensitive characteristics of this content require much higher levels of Internet QoS than have traditionally been available. Because best-effort Internet services with no-settlement, network-layer peering are no longer sufficient for the content delivery needs of many Web site owners, the construction of CDNs has gotten underway. CDNs are smart, application-layer network overlays to the public Internet that make use of caching, content replication, load-balancing, accounting, billing, and other content management functions.

CDN services have been in place for about two years from proprietary CDN service providers such as Akamai and Digital Island. The direction that CDN services will be taking in the next few years will likely be toward multivendor CDN interoperability and settlement relationships among providers. The reason is that the linking of CDN infrastructures enables service providers to more quickly deploy services to a broadly dispersed set of users in a scalable, economical fashion.

There are still interoperability philosophies and standards to be worked out, but ultimately, CDNs are vital to the continued success of Web publishers and for empowering business and consumer users in new ways. Over the next couple years, CDNs will likely evolve to render the Internet a “pay-for-performance” environment. In turn, this will impact the business models of ISPs who, by participating in the CDN value chain, will realize new sources of revenue for themselves.

CDN glossary of terms

aggregator – A distributed or multinetwork CDN service provider that places its CDN servers in the PoPs of as many facilities-based providers as possible, creating an internetwork of content servers that cross multiple ISP backbones. Also called a meta-CDN.

application-layer multicast – An emerging form of the one-to-many IP Multicast forwarding model that reduces router configuration and management complexity by setting up forwarding trees at Layer 7. Instead of routers needing to synchronize and process the many versions of the Layer 3 multicast routing protocol currently deployed

across the Internet, they forward traffic based on traditional unicast IP address information. At the application layer, users interested in joining a multicast group send their IP address to a multicast master. The master determines the appropriate router for forwarding to each user, and informs and configures that router to do so. Application-layer multicast is considered simpler to configure and deploy than Layer 3 multicast, but is less network-efficient.

cache, cache engine or cache server – A server in a network service provider's PoP that hosts a mirror image of a Web publisher's content, which is hosted centrally in an origin server. The goal is to have the content replicated in as many places as possible where user congestion is likely to occur

caching proxy – A content server situated near the clients that improves Internet performance problems related to congestion. Caching proxies cache objects based on client demand, so they may not help the distribution load of a given origin server.

CDN – An overlay network to the Internet built specifically for the high-performance delivery of rich multimedia content.

CDN service provider – A company that manages a CDN and sells CDN services to content provider customers. The current value-add of a CDN service provider compared to an ISP is that CDN service providers have deployed sophisticated technologies for bypassing Internet congestion and accelerating the performance of their customer's content delivery.

Content Delivery Network or Content Distribution Network – An architecture of Web-based network elements, arranged for the efficient delivery of digital content. See also "CDN."

content peering – A function by which operators of two different CDNs can share content, maintain consistent content delivery levels across their infrastructures, and bill one another for services rendered.

content provider – A company that develops services and information for delivery over the Internet.

edge services – The delivery of content from a surrogate to an end user across a single last-mile hop. Requires caching at the edge of a service provider's network.

ISP – An IP network service provider that provides access to the public Internet. May or may not also be a CDN service provider.

load balancing - Intelligent functions in IP networks—either bundled into routers or run as separate appliances—that determine which servers are least loaded and balance requests among server clusters accordingly.

origin server – The content server hosted in a centralized data center, such as a data center of a CDN service provider. The origin server is the one that is refreshed by the content provider. The origin server communicated updates to many distributed surrogate servers, often using IP Multicast technology.

PoP – An IP network service provider's central office, which connects an end user, such as a consumer, to the Internet over a last-mile access link.

proxylet – A set of APIs describing a standard set of basic CDN capabilities that can span multiple CDNs. Proxylets, also called "open extensible proxies," are being defined by the IETF to enable content providers to make consistent services available on multiple CDN service providers' platforms.

publisher – In the context of CDNs, a company that develops and posts content for Internet delivery.

QoS – Quality of service. A term used to indicate the preferential or differentiated treatment of certain flows of Internet traffic based on the response-time characteristics of the traffic being delivered and/or the customer's willingness to pay for preferential treatment.

redirector – A tool that enables content providers to redirect requests to their own DNS servers to the DNS server of their CDN service provider. Also, a lookup service that uses metrics such as user proximity and server load to determine which surrogate delivers content to the requesting user.

reverse proxy caching – Use of surrogates or cache servers to extend a publisher's origin point to distributed points of presence (PoPs) that are physically closer to end users.

settlement – In the CDN environment, business relationships combined with technical standards that enable one CDN service provider to charge another for the hosting and/or delivery of content.

surrogate – A cache server in an ISP PoP that hosts content mirroring that of a content provider's origin server.

Web switching – A network capability whereby the packet-forwarding decisions made by network switches and routers take into account the characteristics and requirements of the application (in addition to traditional Layer 3 addressing information.)

Acknowledgements

Stardust.com wishes to express special thanks for assistance in preparing this paper to The HTRC Group, Inktomi Corp., Mirror Image Internet Inc., and Novell, Inc.