

Deconstructing the CDN: An Executive's Guide to High-Performance Web Delivery

This white paper attempts to simplify the decision tree by outlining the main features and functions of CDNs relative to different types of content. The goal is to help you understand the feature set that is most important to your required functionality so that you can make more informed decisions.

A White Paper

©2010 Internap Network Services Corporation. All Rights Reserved. 09/10



©2010 Internap Network Services Corporation. All Rights Reserved. 09/10



Table of Contents

Introduction	4
What Do You Want to Do?	5
CDN Applications: Where Do You Fit In?	6
Beginning with the Basics	6
Regional Capacity	6
Internet vs. Private-Network CDN	7
Service Level Agreement (SLA)	7
Pricing	7
CDNs for Media Distribution	8
CDNs for Website Acceleration/Caching	9
CDNs for Large File/Software Delivery	10
Supplemental Benefits of CDNs	11
Conclusion	11
About the Author	12



Introduction

Imagine a business where making someone wait one extra second can cause them to lose interest and go elsewhere.

Your first impression might be that these customers have unrealistic, sky-high expectations (and are not worth the anxiety). But the truth is that these are your customers, if you have an online presence and rely on an engaged web audience to further your organization's goals.

Society has progressed so far with connectivity that anything less than instantaneous, on-demand communication and content is considered a problem. To put this in perspective: Web analytics show that for an online retailer, a one-second delay in website performance can reduce customer conversion as much as 7 percent.¹

Whether you're running an online store, an Internet TV station, a place of worship or an online gaming company, the performance of your website has to exceed the expectations of web savvy consumers. Great content isn't great if the page loads in patchwork pieces...or a checkout cart locks up...or a video stream has the jitters.

A CDN (Content Delivery Network) is one of the most effective and efficient ways to create a satisfying user experience. It's also a highly flexible option that addresses a wide range of needs – making it possible for you to, for instance, simulate a broadcast network over the Internet, or cache large files for faster delivery to your intended audience. But not all CDNs are created equal. There are many decisions you need to make and considerations you need to weigh – and if you're not in a technical role at your organization, it's easy to get lost in the terminology.

This white paper attempts to simplify the decision tree by outlining the main features and functions of CDNs relative to different types of content. The goal is to help you understand the feature set that is most important to your required functionality so that you can make more informed decisions.

¹ "The Performance of Web Applications: Customers Are Won or Lost in One Second," Bojan Simic (Aberdeen Group, November 2008).





What Do You Want to Do?

Let's look at some of the ways a CDN might fit into your online strategy. While a CDN makes sense for virtually any type of website, there are three primary business applications they address:

- Media distribution
- Website caching/small file delivery, and
- Large file and software delivery

Across these three primary applications, CDNs also fulfill critical needs for:

- Redundancy
- A globally distributed network, and
- Scaling for peak traffic periods



CDN Applications: Where Do You Fit In?

CDN Use Case	Used For:	Examples
Media Delivery	<ul style="list-style-type: none"> • Live broadcasts over the Internet • Entertainment • Corporate Communications 	<ul style="list-style-type: none"> • Sporting Events • News/Weather • Worship Services • Presidential Inaugurations/ State of the Unions • Television Programs • Movies • Earnings Reports • Executive Podcasts • Training/Continuing Ed • Interviews
Website Acceleration/Caching	<ul style="list-style-type: none"> • Sites that have users dispersed over a wide geographic area; and/or • Content is driven by a database 	<ul style="list-style-type: none"> • Retail Sites • Reservations/Ticketing Sites • Investor/Trading Sites • Online Newspapers and Magazines
Software or Large File delivery	<ul style="list-style-type: none"> • Large File delivery 	<ul style="list-style-type: none"> • Gaming companies

Beginning with the Basics

Chances are that your website falls into one or more of the categories listed (or that's the plan). Regardless of how your site or business is evolving, there are some important things to consider about your CDN, independent of specific features.

Regional Capacity



Where are your users located? What is the expected geographic distribution of your content? If you are doing live broadcasts, where is the nearest ingress point to the network? Some CDNs are regionally “weighted”; for example, one may be strong in the Pacific Rim but weaker in South America. Others are global, but more costly. You can typically find the regional “spread” of a CDN on the CDN’s website. If your organization only needs to communicate with customers and employees in North America, it might not make economic sense to use a global CDN today – but keep future scalability on the radar.

Besides geographic location, you also need to consider the type of functionality your website will offer. If you want to stream live flash in South America, the question isn’t whether the CDN has strong large-object HTTP regional capacity there; what you’ll be concerned with is the streaming capacity. Likewise, if you care about large-object capacity in the Pacific Rim, the CDN’s capacity for Window’s Media in that region will be immaterial. Generally, in today’s world CDN capacity is not fungible in that respect.

Internet vs. Private-Network CDN

There are two basic high-level architectures of a CDN:

- Internet-based CDN: POPs (or Points of Presence) are interconnected via the Internet.
- Private-Network CDN: POPs are interconnected via a private, dedicated network.

Opinions differ within CDN circles about which is the better option. Generally speaking, with a route-optimized CDN, there is little to no difference between these network types. From a performance standpoint, these networks can be made to behave identically. While each has its strengths, a route-optimized Internet CDN offers better uptime because it doesn’t rely on a single network for transit. It’s also likely more economical, since the overhead of a private network makes it difficult to “commoditize” below a certain price threshold. In other words, as transit costs continue to decline, the owners of an Internet-based CDN can take advantage of the reduced costs while not being burdened by the maintenance and upgrade costs of owning their own network. These savings are passed on to the customer.

Service Level Agreement (SLA)

What type (if any) SLA does your content require – and how does that match with the SLA that is provided by the CDNs that you are considering? It’s important to ask if the SLA is proactive, since most are not. A proactive SLA means the provider will credit your account for an outage without you having to “catch” them in one. Many outages go undetected by customers; unfortunately, users will be well aware of these occurrences.

Pricing

What are the elements of pricing? What is fixed and what is based on consumption? Where items are based on consumption, what are the overage charges? Can you buy buckets of bandwidth/storage? Is the contract automatically renewed? Generally speaking costs within the CDN continue to fall for bandwidth and storage. This ‘commoditization’ is advantageous for the customers. What is not so easily commoditized are the value added features that CDN’s provide within the 3 categories of usage.

Now, let’s look at the different use cases of CDNs and the specifics to address with each.



CDNs for Media Distribution

Distribution of media on the Internet has become a very common use case. All major content creators are now delivering their content over the Internet. Generally, there are three types of media delivery:

1. **Live Delivery.** Today, anyone can operate their own “broadcast network.” Live delivery over the Web has become a critical strategy not only for big media but also for smaller “niche” markets. President Obama’s inauguration in January 2009 broke all records for viewership of live web broadcasts. At the same time, thousands of small and mid-sized churches and faith organizations broadcast their services multiple times a week to tens of hundreds of viewers. High school football and Little League baseball are regularly broadcast over the Internet. A camera, a mixer, an encoder and a CDN ingress point are all you need to reach your intended audience.
2. **Video on Demand.** This is how media distribution on the Web is primarily used today. Sites such as YouTube and Hulu have made Internet video a routine part of life, so organizations of all types are following suit and enhancing their online presence with video.
3. **Linear Channels.** This is an interesting but still nascent application of media delivery. A linear channel is a 24/7 channel of content that is not video on demand, but rather scheduled to play at specific times – much like a traditional television channel. This is also referred to as “simulated live” since it is prerecorded but scheduled.

Depending on the form of media distribution you want to do, there are several questions to ask about your potential solution:

- Does this CDN support all three of the above distribution models – or more importantly does it support the one(s) you care about? If live broadcasting is a must-have, find out the number of POPs and redundancy of the network ingress points and how quickly they are provisioned upon a new request. Ask about the service provider’s support in general. Does the CDN have a Network Operations Center (NOC) available 24/7? What is their call response policy? Don’t worry about asking too many questions; if you were in the field at a live event and it went poorly, you would have wished you would had asked.
- Does the CDN support integration with an Online Video Platform (OVP)? An OVP is an application that allows content creators and owners to “onboard” their content onto a CDN for either live or on-demand delivery. The Video Platform often allows users to insert ads, syndicate the content and report on usage. Most people using a CDN for video distribution would require these features. Assuming you are going to need an OVP, which one? Does the CDN that you are working with own the platform, or is this a third party that you have to contract with separately? Often the OVPs are resellers – meaning that you need to have a relationship with the reseller and not the CDN directly. This may or may not fit with how you want to manage things.
- Does the CDN support your intended distribution format, such as:
 - Adobe® Flash®
 - iDevice
 - Microsoft® Silverlight®
 - HTTP Streaming



- Icecast

It's perfectly valid to inquire about their future plans for these products. If you are streaming your radio station over Icecast™, that is a very different model than Silverlight – and a CDN's plan to support those products is likely to be different for each. If your content is encoded into Flash, ask what version server the CDN will be on, and what the provider's plans are for the next release.

There are additional considerations depending on the type of functionality you're after:

High Definition (HD). High-definition video on the Internet is still an evolving concept. If the long-term goal is to broadcast your content in HD, it's important to understand the roadmap of the CDN vendor that you are engaged with. A CDN vendor should have a specific technology answer to the question, "How will your CDN manage HD traffic?"

Digital Rights Management (DRM)/Licensing. If your content is licensed or needs to be protected – such as a pay-per-view event or movie – it's important that the CDN you are considering either has a DRM solution or partners a company that provides one.

Mobile. Does your CDN support mobile delivery? If so, what flavors?

Transcoding. Transcoding transforms a file into the codecs or formats compatible with the user's device – such as H.264, Flash, Silverlight, iDevice and Quicktime. This allows for maximum distribution of your content, because it's not dependent on a specific device or format.

Pay-Per-View. If your content is intended to be viewed only by people that are paying for it, determine if the CDN can provide that feature or integrate with a PPV vendor that does.

Analytics. For live broadcasts, real-time reporting is highly desirable since it allows you to customize your content for the people watching. For instance, a faith based broadcaster can recognize someone signing on to the broadcast from Arkansas and thank that viewer in real-time for joining the service. For VOD, it is valuable to know which are the most frequently visited assets, and be able to drill into them to extract data such as average viewing time and geolocation of viewers or view their DMA results.

CDNs for Website Acceleration/Caching

Website acceleration is a general term for four complimentary technologies that are offered by CDNs to improve the overall performance of websites. These technologies apply to some sites more than others; therefore, you'll need to do some analysis to determine if a particular solution makes sense for your desired functionality.

The four technologies are:

1. **Static Site Caching.** Also known as "reverse proxy caching," this is a prime feature of CDNs. Static site caching allows the CDN to point its distributed network at an "origin"



server and cache the content of that origin in geographically diverse areas through a mechanism called GEO-DNS² (see also additional illustration notes at the end of this document). Though many factors (such as current loads) contribute to this routing process, the main factor is the proximity of information. For instance, in a CDN that deploys servers on both coasts of the U.S., a request coming from California most likely will be served by an edge server located on the West Coast rather than the East Coast.

2. **Dynamic Site Caching.** Dynamic data is data that is time sensitive – such as pricing or account balances. Generally, this relates to information that has a short lifespan and cannot be cached for any length of time.
3. **Web Application Acceleration.** Whether you're running a reservations system or selling products, web application acceleration will improve how this function performs. The overall focus of these methods is to improve the HTTP or HTTPS performance. Strategies may include HTTP protocol acceleration and compression; tuning of the packet sizes; and other techniques. These solutions tend to be very browser-centric, as the defaults in the browser are often the best places to find the opportunity for acceleration.
4. **IP Acceleration.** As opposed to web application acceleration, IP acceleration is focused on improving the performance of Transport Control Protocol (TCP) traffic. Because TCP was designed more for accuracy than speed, some CDN providers have developed (or partnered with companies) for solutions for accelerated IP.

CDNs for Large File/Software Delivery

Call it a progress bar, call it a harmonogram – either way it's a misnomer. Users don't feel a sense of "progress" or "harmony" when a dense download ties up their system. Any organization that distributes large files or software over the Web – enterprise application vendors, antivirus providers and game developers come to mind – need the infrastructure to handle bulky, spiky traffic. A CDN offers an excellent way for the organization to deliver these assets without the capital expense required to build a global network of sufficient bandwidth.

A "best practice" CDN for large file/software delivery will include certain components to ensure secure, speedy transmission. Look for:

- **Secure Content Delivery.** When delivering software/games, the ability to secure the content is mandatory. This may mean delivery over secure socket layer (SSL). If so, you need to determine whether you are using a shared certificate or your own. If the latter, you need to determine that the CDN supports third-party certificates.
- **Authenticated Access.** If your content requires authenticated access, you need to know what methods the CDN provides to enable this. This could be domain-based authentication, user authentication or some other method.
- **Large File Optimization.** A large file (100 MB or bigger) has special needs with regard to delivery. (Note: that it doesn't take much these days to create a 100 MB file – a high-quality

² GEO-DNS is the mechanism by which a CDN is able to locate place from which a request for content is coming from and redirect that request to a Point of Presence that is near the original requestor.



image or graphics-laden PowerPoint might be this size.) If your files are in this range, you need to determine if the CDN you're considering is equipped to distribute files that large. One of the main features you'll want to have is the ability for the user to pause a file during download and then resume at a later time.

- **Download Manager.** This is a client software package that helps users download content. It often increases download performance by accelerating the file transfer. Sometimes this is a "brandable" piece of software that the organization can distribute to its web visitors. It also insures the delivery even if there is a broken connection by picking up where the file transfer left off and completing the download.
- **Download Analytics.** You need real-time reporting and detailed historical reports to determine your ROI and track record in delivery of your content. The CDN should report on all deliveries as well as any failed attempts. It can also send a 'receipt' back when a download is complete in real time so that the transaction is confirmed – especially important for paid transactions where some sort of support is implied.

Supplemental Benefits of CDNs

Whether you're using a CDN to enhance media delivery, web caching or file downloads, there are additional infrastructure benefits to be gained. CDNs address single-point-of-failure, global delivery and scalability concerns.

	Redundancy	Global Delivery	Scale for High-Volume Traffic Periods
Web Caching	Your web presence must be on 24/7	High costs to set up and maintain servers around the world	No need to build out your network only for occasional traffic spikes
Media Delivery	There is no time machine for live events gone awry	Address buffering and latency due to single homed networks	High volume periods
File Download	Mission-critical applications are mission critical for a reason	Stop the time out due to latency	High volumes are usually experienced only for 1-2 days after a release

Conclusion

Ultimately, when your target audience can't view and interact with your online presence, they go elsewhere. Sometimes they go to your competitor. Sometimes they lose interest in whatever they were searching for, and sometimes they'll pick up the phone and use more costly, "high-touch" forms of customer service that impact your bottom line. But rarely do they come back to your website.

With some knowledge about how CDNs work and the features that will facilitate your objectives, you're well-positioned to implement a CDN that will not only boost your online presence, but yield measurable efficiencies – and keep up with increasing expectations.



About the Author

Pete Mastin is Director for CDN Engineering at Internap.

About Internap

Internap is the leading world-class organization providing data center colocation facilities and managed Internet services to Fortune 500 companies worldwide.

Additional Notes

GEO-DNS is the mechanism by which a CDN is able to locate place from which a request for content is coming from and redirect that request to a Point of Presence that is near the original requestor.

Below is an illustration of how Geo DNS works:

