

Deliver High-Definition Internet Media at Peak Demand

Leveraging the Internap® Content Delivery Network (CDN) Architecture

A comprehensive look at Content Delivery Network technologies and architecture, including the different levels of capacity, connectivity, evolution and functionality. This white paper will explore the challenges associated with early, “first-generation” Content Delivery Network providers and how the Internap® Content Delivery Network is positioned to address these challenges as well as best practices for rich media delivery.

An Internap White Paper



Table of Contents

Introduction	3
Challenges of the First-Generation Content Delivery Networks (CDN)	3
The Next-Generation Content Delivery Networks	4
Overcoming Network Strain	4
The Price of Performance	5
Internal Management	5
Leveraging the Internap Architecture	5
Optimized Routing	5
Traffic Management	6
Content Management	7
Stream Performance Characterization	8
Network and Content Security	9
Best Practices for Rich Media Delivery	10
Summary	11
About Internap	11
Internap Services	11



Introduction

Content Delivery Network (CDN) is a term that was coined in the late 1990s to describe a system of gateway or reverse-proxy caches networked together across the Internet to reduce load from the origin site while also attempting to improve performance to the end-user. Content was placed at the “edge” of the network and served from “cached” servers close to the end-user saving the need to retrieve this data across the network for every user request.

Challenges of the First-Generation Content Delivery Networks (CDN)

The core benefit of conventional CDNs is driven by a reactive, user-initiated pull-through of content into a given cache and subsequent cache hits. To drive the pull-through process, all end-user requests for content are redirected to a CDN gateway server. At the gateway, the user request will meet one of two fates.

In the first possibility, the gateway cache has recently seen a request for the same object, cached the prior response, and retained the response in cache (i.e., not aged out based on least-recently-used criteria). If these conditions are met, the gateway cache may still be required to validate the cached content with the origin, but otherwise it will be able to serve the content directly to the user. This defines a cache-hit, where the CDN realizes the full or partial benefit of off-loading traffic from the origin and potentially improving performance.

In the second possibility, the gateway cache will have no knowledge of the requested content and will need to pass the request to the origin website. The gateway cache subsequently mediates the connection between the origin and the user, and will cache some content as resources permit. This defines a cache-miss, where the CDN imposes a performance tax on the transaction.

This process is effectively bottlenecked by whether or not a cache-hit can be generated, and is driven purely by statistics. Large numbers of users hitting a single cache within certain time and/or resource constraints will receive a large benefit. However, as users are distributed across caching servers and across time or in contention of resources, the CDN benefit becomes highly diluted.

Upon inception and through the course of their operating history, first-generation CDNs believed that the closer you place the “edge” server to the end-user, the better performance gain would be achieved. CDNs have pursued a strategy of high-distribution and low concentration: building small numbers of servers into large numbers of networks. Unfortunately, as the number of caching servers increases, the probability of cache hits decreases. Without these critical cache hits, the benefits of using a CDN are lost.

Today’s websites pose additional challenges on the aging first-generation model. These sites are more robust, with richer graphics, dynamic content, and Flash and streaming media. Delivery of these objects is beyond the capabilities of just a few machines within a cache node. Traffic patterns can change much faster than machines can be redeployed and rebalanced. With the first-generation CDN architecture, often too many servers are deployed to networks with less than expected usage and not enough servers deployed to networks with more than expected usage.



As well, many of these servers are located in third-party facilities, creating a significant investment in operational resources such as hands-on labor that must install, reconfigure and troubleshoot remote hardware. The legacy of thousands of aging servers deployed under an outdated model, a complex network management challenge and all the associated costs of maintaining these individual systems are inevitably passed on as higher service pricing or poorer-than-expected performance.

The Next-Generation Content Delivery Networks

With the only route optimized CDN in the industry, the Internap® next-generation CDN leverages a system of performance-boosting technologies and high-density, smart system architecture to ensure the rapid and reliable delivery of content to the end-user.

The Internap content delivery model is comprised of strategically located global content delivery points that are connected to Internap Private Network Access Point (P-NAP®) facilities. The P-NAP facilities provide the intelligent Internet infrastructure to deliver content using the best possible path to the end-user through the use of patented route optimization technology. This unique technology avoids high latency and packet-loss paths across the Internet and enables Internap to provide a high-quality, high-throughput streaming and download experience for its customers. The Internap content delivery points use redundant media server and storage clusters with virtually unlimited capacity to efficiently deliver content to the end-user to anywhere around the globe.

Users experience a faster, more reliable experience to their users that is not constrained by content type or size, Internet performance problems, bandwidth limitations or caching inefficiencies.

Using standards-based methods, Internap customers securely publish their content through the content delivery system to redundant publishing portals available worldwide. Content is immediately cataloged and proactively replicated to the content libraries at all of the major global content delivery P-NAP facilities. This distribution employs various integrity checks and synchronization technologies and uses sophisticated IP acceleration protocols – all done reliably at unsurpassed speeds.

Overcoming Network Strain

Serving extremely popular content can strain all aspects of modern technology, not only at the network layer but also at the system and physical disk layers. It's not uncommon to find an overnight video success story pummeling the spindles of an unprepared storage platform before ever reaching the limits of the network capacity.

To prevent primary storage from being overwhelmed during peak load, there are multiple tiers of caching across all the streaming servers. This is true for all delivery methods including HTTP, Windows® Media Server® and the Flash streaming service. This leverages the net capacity of more than one thousand streaming servers and their disk spindles, as opposed to relying on monolithic direct storage technologies.



The Price of Performance

Leveraging performance-boosting technologies, Internap is able to build smarter, not fatter. With strategically located high-capacity data centers and Points of Presence (POPs), the costs associated with maintaining thousands of servers does not exist and therefore is not passed onto the customer. Internap delivers performance without the bloated price tag.

Internal Management

All of the critical factors that affect the quality of high-quality media delivery are controlled by Internap advanced optimization technologies and not left in the hands of other third-party providers. Many other streaming providers and early first-generation CDNs leave the management of their equipment and connectivity up to a third party. The connectivity, routing decisions and server management are left in the hands of the chosen upstream network provider and/or colocation facility. These other networks may not be well versed in supporting unexpected bursts in traffic that are extremely sensitive to degradation in packet behavior. In contrast, the Internap CDN is hands-on and supported by highly-trained engineers at Internap facilities. Not only does Internap manage its fully-redundant network, but also takes complete advantage of its patented route optimization technologies to provide highly reliable robust delivery of content, applications and communications throughout the Internet.

Leveraging the Internap Architecture

Providing the highest quality, scalable and reliable user experience over the Internet requires the following key technologies to perform in concert:

1. Optimized Routing
2. Traffic Management
3. Content Management
4. Network and Content Security

Optimized Routing

Once a stream or download has begun, it's important to ensure that the data takes the most reliable and quickest path on the Internet. Typically, for a broadband or Internet Service Provider (ISP), there is little that can be done to control how an Internet packet of data reaches a network, since each provider sets up their network according to individual business rules.

The Internap Performance IP™ network service is unique in the industry. Through its proprietary intelligent route-control technology called MIRO (Managed Internet Routing Optimizer™) and a global network of PNAP facilities, Internap continually analyzes the traffic situation on every major Internet backbone, and then selects the optimal path to deliver the stream on the basis of capacity availability, distance and performance. Advanced analytics are used to probe and benchmark all known points on the Internet through each bandwidth carrier and continuously adjust the routes that content uses to reach the end-user. This concept would be similar to attaching “driving directions” to the packets that make up the video/audio stream on its way to the end-user. These driving directions take into consideration such things as the traffic conditions and the speed limits of the various routes along the journey. The control over these paths and routes can be very granular since 99% of the bandwidth used during a transfer is outbound.



The MIRO technology benefits the CDN in a couple of ways. First, by ensuring that whenever possible the customer is directed to a nearby POP, we increase the overall throughput achievable. By the same token, we reduce the swath of Internet that must be traversed by the stream and therefore eliminate potential performance bottlenecks. Second, the constant monitoring of capacity available in each POP allows MIRO to avoid overloading a location. Overloading would result in poor stream delivery performance at the source (the streaming and progressive servers) which translates into stuttering videos and/or pauses and stops in the stream. This across POP load balancing also helps manage capacity better. Finally, MIRO can avoid taking end-users to a POP that has really poor performance for the prefix the end-user is in.

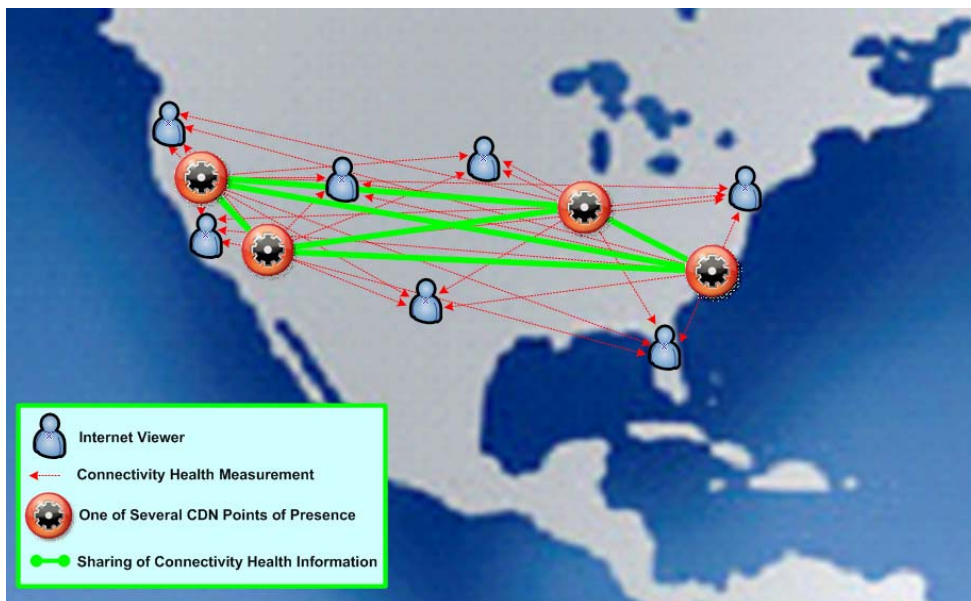


Diagram: All POPs Continually Monitor and Detect Internet Congestion

Traffic Management

Nobody likes to wait, and your viewers/listeners won't either. It's critical that the initial request for content be handled promptly and responded to by facilities, servers and networks that have the capacity on-hand to exceed the requirements needed to serve the stream at the highest level of quality, even in times of peak demand. Several factors are taken into consideration when making this first and most critical decision about directing traffic.

1. Which facilities have round-trip times and latency that are within an acceptable range of the viewer/listener?
2. Which facilities and servers have CPU, memory, disk and network load capacity within an acceptable range that can take on additional requests?
3. Is there a disaster or planned maintenance that should be avoided?



With Internap, the outcome of how the traffic is directed is based on the relationship between the end-user location on the Internet and all of the Internap P-NAP facilities. This data is compiled with these key decision making metrics within a split second. A connection is then immediately established with the preferred POP/P-NAP facility and streaming/downloading begins. A server, network or disk system can quickly become overloaded if balanced incorrectly and it's important to shift traffic to avoid these issues.

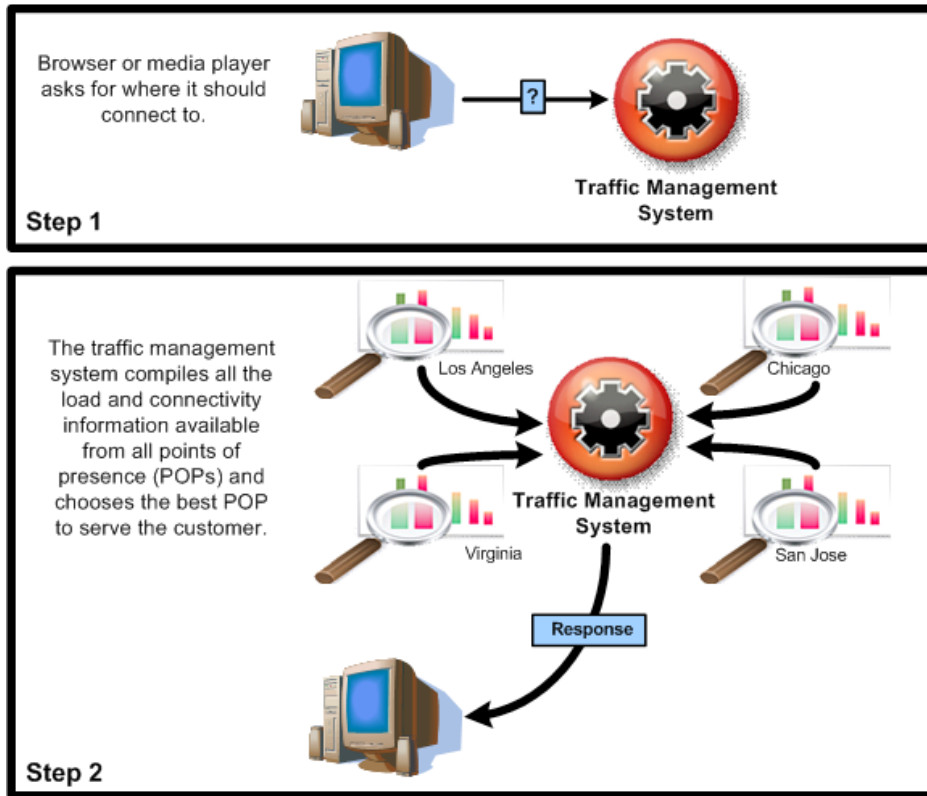


Diagram: The Decision-Making Process Behind Traffic Management

Content Management

After publishing new videos, music or other content, it needs to be available instantly all over the world. The Internap content library is reliably synchronized and integrity checks are continuously run. Having your content securely replicated and distributed throughout the world is a benefit provided by a storage architecture, which is designed with no single points of failure. If for any reason one of the main content library facilities were to encounter an unforeseen disaster, your content would still be available through identical surviving content libraries. Content is pushed quickly at the speed of light using a fiber optic storage network and proprietary IP acceleration technologies.



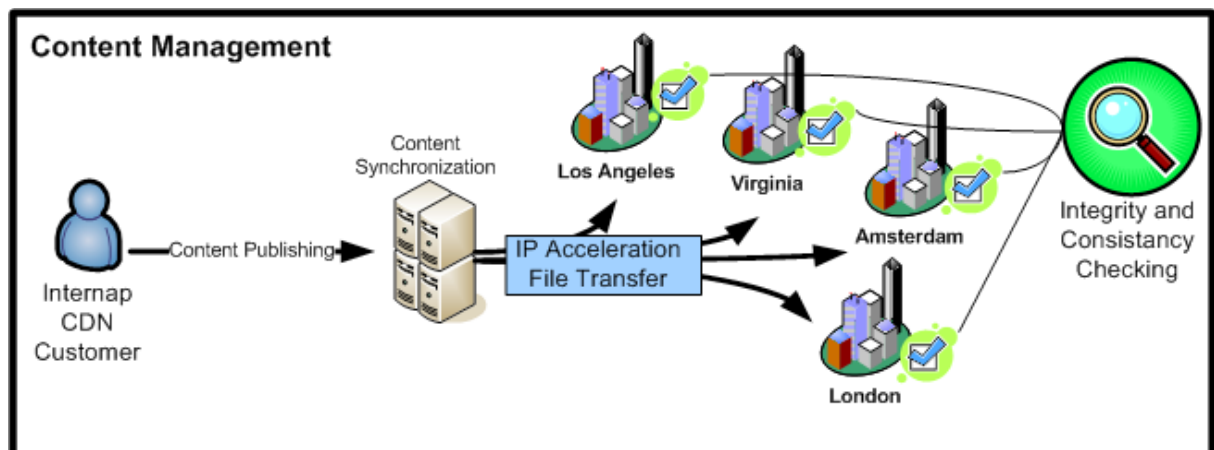


Diagram: How Files are Distributed and Synchronized across the CDN

Content is distributed to and streamed from key locations throughout the Internet where the majority of the Internet's users are well connected. These are highly redundant and powerful points of presence capable of pushing your media out at Gigabit speed.

Stream Performance Characterization

Key to the heart of the Content Delivery Network is a set of advanced measurement and analytical services that keep tabs on the performance of all key aspects of the entire platform. This omnipresent view into the technology is present at all levels of the network's stack including the bits traveling across the wire all the way up to the user's streaming experience.

User-Experience Layer: Stream Buffering, Time to Connect, Automated Stream Sampling by members of the Internap Network Operations Center

Network Layer: IP Packet Latency, Jitter and Drops, Internet Routing Problems, Number of Connections per Server

Physical Layer: Gigabit per Second Throughput, Disk Throughput, CPU Load, Physical Data Circuit Integrity

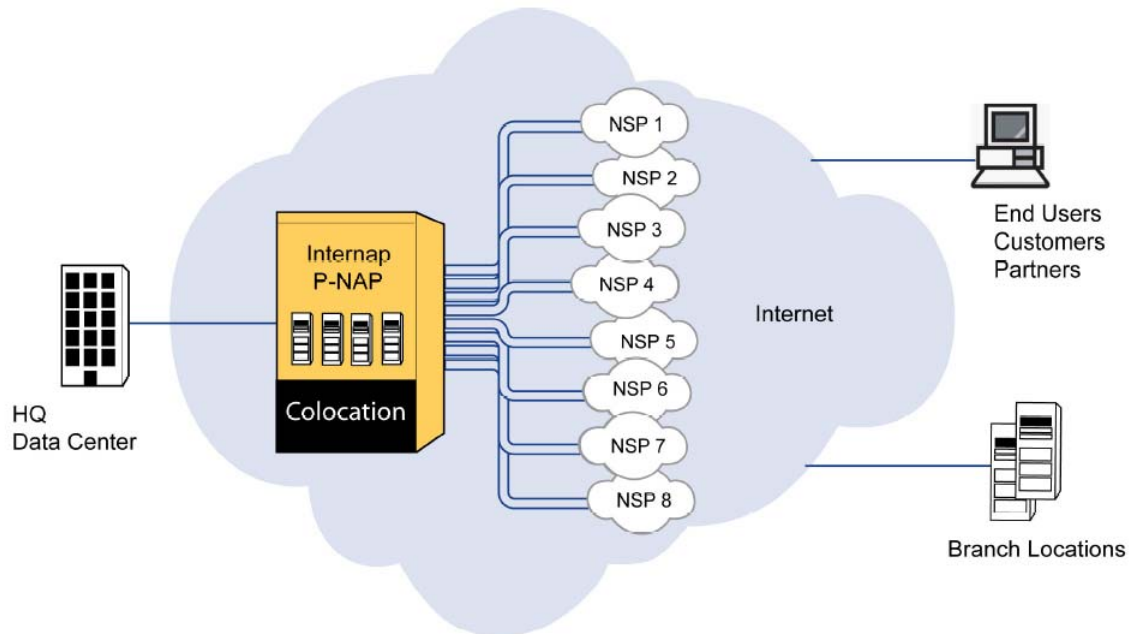
As an unbiased third party, Internap also utilizes Keynote's StreamCheck services to verify its streaming quality. Keynote's StreamQ™ is the de facto industry-standard metric for measuring the streaming media user experience.

StreamQ is measured in letter grades from A+ to F that indicate the amount of time a user had to wait or experience interruptions ("frustration time") versus the time actually spent enjoying the content ("play time").

To determine the StreamQ for a particular check, Keynote adds up the connect time, buffer time and re-buffer time. Then, a two second penalty is added for each re-buffer event. If this final "adjusted frustration time" is less than six seconds, the grade is an A+; 6 to 9 seconds A, 9 to 12 seconds, B+, etc. Internap maintains an A+ StreamQ rating.



Network-Based Route Optimization: Delivered through a series of network access points, featuring direct, high-speed connections to multiple premium carriers, Internap monitors and measures performance and traffic-routing patterns across the path of least resistance.



The Internet is comprised of several of thousands of networks that are connected to one another loosely based on good will and a common interest to do business. There is no governing body or organization that ensures these networks are all integrated correctly or efficiently. Internap continuously probes the Internet searching for these issues and quickly changes its routes to detour around any problem areas.

Network and Content Security

Internap maintains a CDN security program designed to address threats facing our digital media delivery business and yours. Examples of credible scenarios include the following:

- Compromise of the server and network infrastructure underpinning the CDN, leading to a loss of data integrity and system availability;
- Interception or unauthorized access to customers' confidential content in transit; and
- Third-party abuse of customers' accounts and unauthorized distribution of media files.



To address these threats, Internap uses a layered approach to information assurance. Examples of specific controls and capabilities used in our CDN include the following:

- Ingress and egress firewalls at the perimeter;
- Host-based firewalls, intrusion prevention and malware protection;
- Central logging with advanced anomaly detection;
- Proactive detection and management of system and network vulnerabilities;
- 24x7x365 monitoring and computer security incident response;
- Encrypted (SSL) delivery of customer content;
- Authenticated delivery of customer content through token-based authentication;
- Enhanced customer logging and reporting features through Internap MediaConsole®; and
- Delivery of raw media logs to customer for fraud detection.

Best Practices for Rich Media Delivery

1. Guarantee a high-quality user experience – Companies should make sure that the site provides good performance for all visitors, regardless of their location. It must scale to handle high traffic load during peak seasons and events, and it must be available 24x7x365 no matter what the network conditions on the Internet.
2. Match capabilities with business needs – Make sure the partner you select is best suited for the type of media that you are going to deliver. For example, if you are looking to deliver streaming media, you will want to select a partner whose expertise is in streaming media. CDN customers should select a particular flavor of service or service combination that suits their business needs and requirements.
3. Solid Professional Services ability – Carefully assess internal skills that will be required to meet your business needs, both immediate and future. If your project will require custom development, you will benefit from a partner with a solid professional services team.
4. Choose a provider with solid performance measurement tools – Reporting and analytics capabilities are now more critical than ever to evaluate and maintain the effectiveness and performance of your content. Ensure your delivery partner offers monitoring and reporting capabilities that are customizable to fit your specific needs. The key is to have a rock solid understanding of what information will be important to you and your company. Then compare what your needs are to the partners capabilities. If your needs include the ability to track geographic information, or for information to be imported into an existing content management system, make sure these options are available.
5. Evaluate ease and flexibility of service – Choose a partner that provides the most efficient way to manage, deliver and protect your rich media. Look for flexibility in content management and an easy to navigate system. Tools and wizards may be important, but if your needs require an open API or the ability to support third-party integration, make sure the system is flexible enough to handle both.



6. Scalable, reliable, manageable content storage – Although price can be a determining factor, many organizations subscribing to CDN services say issues of reliability eclipse cost. Ensuring your end-users receive the best possible experience will keep them returning to your site. Your media delivery partner should have the ability to scale with your growing business needs as well as feature worldwide connectivity, high bandwidth capacity and maximum route exposure to multiple providers.

Summary

The Internap content delivery system offers a simple-to-use platform for meeting high-performance delivery needs. You can provide scalable high-definition video to your users using a globally distributed technology that adheres to standards and requires no client-side installation. There are no downloads or custom video player complications for viewers to get frustrated with. Internap will continue to focus on bringing the latest and greatest advancements in media delivery technology while providing the highest levels of performance and reliability.

About Internap

Internap is a leading Internet products and services company that provides *The Ultimate Online Experience®* by managing, delivering and distributing applications and content with a 100% uptime reliability service level agreement. With a worldwide platform of data centers, managed Internet services and CDN, Internap frees its customers to innovate, improve service levels and lower the cost of IT operations. Thousands of companies across the globe trust Internap to help them achieve their Internet business goals. For more information, visit <http://www.internap.com>.

Internap Services

Internap delivers a comprehensive set of Internet solutions to unleash your potential and transform your business.

- Internet Services deliver unflinching reliability, unparalleled network performance and proprietary route control technology with 100% uptime guarantees.
- Colocation Services offer a fully-redundant and reliable infrastructure for your business-critical applications.
- CDN Services quickly and securely stream and distribute video, audio, advertising and software to audiences across the globe.
- The Internap MediaConsole provides a single, easy to navigate system featuring Media Asset Management, Digital Rights Management (DRM) and detailed reporting tools.
- Professional Services bring custom development, transcoding solutions and expertise through a team of consultants.
- World-class customer care with a single point of contact for customer service.

