



Websites and mobile applications can't afford to use outdated content delivery networks (CDNs), according to this new 10-page report from Gigaom Research: "Next-generation CDNs: extending customer applications to the edge," by Ernie Regalado.

Traditional content delivery networks were designed to cache static content, and can't handle the dynamic content of today's web and mobile applications. This research report explains how you can get high performance, accelerated content delivery with a next-generation content delivery network.

## Key Takeaways

- Next-generation CDNs are the answer to delivering dynamic content because they can purge dynamic content in real time, making them better suited to the needs of interactive web apps, mobile apps, and APIs.
- Modern CDNs offer huge savings because they don't have to deal with legacy infrastructure and are built for dynamic content acceleration.
- Customers using a modern CDN can improve visitor metrics across all categories, including better conversion and retention, more interactions, and optimized page load times

## How Fastly Can Help

Today's websites need content delivery networks that feature advanced functionality such as instant purging, real-time reporting, and application control at the edge.

Fastly built its CDN infrastructure with all SSD servers using the open source caching software Varnish, which is designed for high-performance content delivery and real-time management.

Fastly's state-of-the-art content delivery network offers immediate visibility into performance analytics, zero-delay Instant Purge in 150 milliseconds across dynamic content, logging and reporting, and API-based tools that give companies even more control over their content.

---

"Anything less than a real-time solution is unacceptable. Dynamic content requires CDN infrastructure that can support millions of users and billions of interactions in a short time and on a global scale."

**Ernie Regalado**, Gigaom Research Analyst

**GIGAOM** RESEARCH

---



**GIGAOM** RESEARCH

# Next-generation CDNs: extending customer applications to the edge

---

Ernie Regalado

November 4, 2014

*This report is underwritten by Fastly.*

## **TABLE OF CONTENTS**

Executive summary .....	3
What are the limitations of traditional CDNs? .....	5
What is a next-gen CDN? .....	8
Next-gen dynamic content caching and traditional CDN dynamic content acceleration .....	8
Dynamic, mobile-aware, interactive web applications.....	11
What does control actually mean? .....	11
Call to action and key takeaways .....	14
About Ernie Regalado .....	15
About Gigaom Research .....	15

## Executive summary

Today's web and mobile applications are wreaking havoc on traditional content delivery network (CDN) infrastructure. Traditional CDNs are unable to cope with the demands of real-time services, as their infrastructures have been built to cache and deliver static content only. Web applications are now dynamic and interactive – and every consumer expects them to function immediately. The traditional CDN, in an effort to overcome its limitations in delivering modern day content, has created new services on top of its legacy software that accelerate content over the middle mile. But building new software on top of old is never ideal.

The next-generation (next-gen) CDN has solved the problem of delivering dynamic content – such as HTML, API, Ajax calls, or anything else that requires application logic – that is generally considered to be uncacheable. By building a modern CDN infrastructure from the ground up, next-gen CDNs can cache rapidly changing content, largely due to features such as instant purging, which removes and replaces stale content in real-time.

Other features that next-gen CDNs offer increase the flexibility and transparency of content delivery and monitoring. These features include instant configuration changes, real-time reporting tools, and real-time logging. Instant configuration changes enable developers to update applications in the cloud while they are running, thus eliminating the need for maintenance downtime. Real-time reporting and logging make it possible to predict traffic spikes, and provide visibility on potential bottlenecks that could impact the CDN. Extended flexibility gives CDN customers complete control over their web application, from development stages through running in production in the cloud. Thus, businesses are no longer constrained to the traditional CDN “black box,” where visibility, functionality, and choices are limited in application and content delivery.

The bottom line is that traditional CDNs are incapable of handling the heavy workloads and demands of modern consumer applications because their architectures are outdated and their software-caching platforms are bloated and unable to keep up with real-time services. Many of these traditional CDNs have failed to re-invent themselves, choosing instead to operate within their comfort zones.

This report examines how strategies for next-gen CDNs will enable online businesses to be more successful by improving web traffic metrics such as conversions, click-through rates, engagement, and retention, while enhancing end-user experience.

**Key Takeaways**

- Today's web applications are complex and resource intensive, and the plethora of content types are severely taxing current traditional CDN infrastructure.
- Traditional CDNs are ill-equipped to deal with the requirements of dynamic content delivery, and don't offer advanced functionality such as instant purging, real-time reporting, real-time logging, and application control at the edge.
- Next-gen CDNs are the answer to delivering rapidly changing content, interactive web applications, mobile applications, and APIs.
- Next-gen CDNs have stepped up their efforts and built platforms that offer customers control over their web applications, giving them the flexibility and visibility they require, matching the features and functionality of leading cloud compute platforms.
- By partnering with the right CDN, a company can enjoy high-performance application and content delivery, advanced reporting and analytics for decision-making, and improved metrics such as conversion rates, retention, click-through, end-user interactions, and optimized page load times.

# What are the limitations of traditional CDNs?

Over the last decade, many online businesses have moved away from websites with static content – product brochures, pictures, videos, and PDFs. Today, businesses are focused on long-tail content, content that is infrequently accessed but has a long internet shelf life. Modern companies want highly dynamic websites that deliver personalized experiences to each visitor, so the first click leads to a more meaningful relationship between brand and consumer. This personalized experience becomes critical with the increasing use of mobile devices and the performance demands they create.

Traditional CDN infrastructure was built to deliver static content, so it cannot properly handle the acceleration of interactive web applications, constantly changing content, or personalized, localized content.

Twitter, Facebook, and similar platforms have changed the way content is delivered and consumed online, as well as how users interact with it. In short, all these interactions are now happening in real-time, which means that anything less than a real-time solution is unacceptable. Dynamic content requires CDN infrastructure that can support millions of users and billions of interactions in a short time and on a global scale.

Typically, CDNs fetch content from the origin server, cache it in various edge locations around the world, and deliver it on request. This process involves caching and delivering content, with the CDN acting as a one-way delivery system. This model starts to break down when end-users begin interacting with applications, posting non-stop user-generated content like comments, pictures, and updates. In this scenario, the one-way delivery system must become a two-way, real-time delivery system that requires different software architecture and underlying hardware infrastructure.

Traditional CDNs face tremendous challenges catering to CDN customers who require real-time services, whether real-time content delivery, real-time updates, instantaneous purging, real-time reporting, or real-time logging. This is especially true for the emerging DevOps community, those engineers who sit between the software developers and operations team within an organization, making sure web applications are scalable, rolled out globally, and can support millions of users without crashing. Their job is ensuring that the hardware, software, network infrastructure, and CDNs can seamlessly work together before and after application deployment.

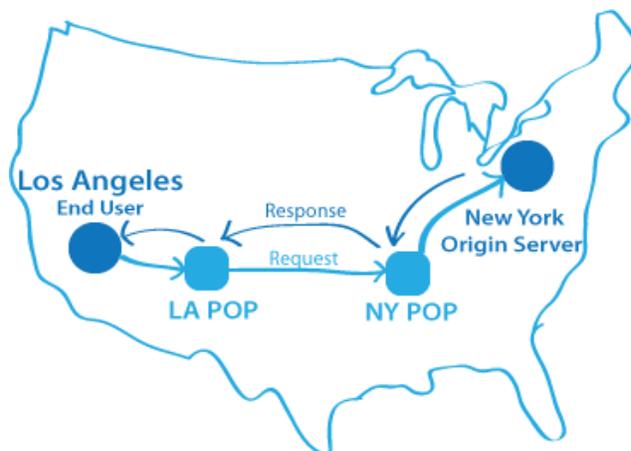
## Middle-mile and TCP/IP optimization

Traditional CDNs handle content through a technique known as TCP/IP middle-mile optimization. Whenever a user requests content, such as base HTML, API, or Ajax, the request must travel long distances to wherever the origin server is located, even if it is in a different country. This technique tweaks the TCP/IP protocol to optimize middle-mile delivery. Although this technique improves dynamic content delivery, it doesn't meet the needs of modern web applications and the rate of content changes. Compounding the problem, traditional middle-mile optimization solutions tend to be very expensive, sometimes costing thousands of dollars per month.

Current CDN middle-mile optimization, often known as dynamic content acceleration, has been around for a decade; Netli, a pioneer of dynamic content acceleration, started doing this in 2000. Unfortunately, although content types have changed over the past 14 years, middle-mile acceleration techniques have not. Manipulating the TCP/IP protocol involves optimizing certain parameters, such as window scaling, which increases the receive window size of a packet, connection pooling that establishes connections and keeps them open, as well as a few others.

The problem with this approach is that packets must still travel long distances, sometimes thousands of miles, at the cost of valuable milliseconds. This amount of time is especially worrying when it comes to Response Times, Time-to-First-Byte (TTFB), and Time to Content Download. Traditional middle-mile acceleration travels very long distances. In the example in the following figure, it travels from an LA-based Point-of-Presence (POP) to a NY-based POP (a distance of 2800 miles), wasting precious milliseconds in the process

### Middle-mile acceleration across a distance of 2800 miles



(Source: Gigaom Research)

## Traditional CDNs at a crossroads

Many traditional CDNs are unable to let go of and innovate beyond their legacy systems because they have invested so much time and energy in their current hardware and software. Even if CDNs upgrade their infrastructures, they must re-architect large parts of their software platforms.

Next-generation CDNs don't have to deal with legacy infrastructure. These companies have emerged recently to address the needs of modern businesses, and have built state-of-the-art CDN infrastructures using the latest hardware, software applications, and best practices available. They've also created new ways to deliver content and offer real-time features that help DevOps break out of that black box where features and functionality are limited, so that businesses can scale their applications to the cloud.

# What is a next-gen CDN?

Long-tail content and user experiences are driving the need for next-gen CDNs that can deliver visibility on how customer content is being consumed. This means visibility into metrics reporting click-throughs from one page to the next – whether, for example, an ad converts to a sale or when a user abandons a page for a performance reason.

The next-gen CDN's mission is to become a part of their customers' DevOps teams and application infrastructure. Giving DevOps teams control over applications allows their companies to respond quickly to customer demands. With access to real-time data, CDN customers can evaluate consumer behavior in real-time, make business decisions in real-time, and deliver up-to-date and relevant content for each user.

Twitter, Facebook, and many leading web destinations have set the bar high, spoiling their visitors with real-time content. Critical CDN components that make real-time experiences possible are instant purge, instant configuration changes, log streaming, high performance API delivery, and access to robust dashboard views of historical and real-time stats.

## Next-gen dynamic content caching vs. traditional CDN dynamic content acceleration

---

In reality, traditional CDNs that introduced middle-mile accelerated solutions convinced businesses that dynamic content could not be cached. These CDNs are now offering add-on solutions that prolong the content-delivery process. However, dynamic content is cacheable with the right CDN. Next-generation CDNs can cache dynamic content because they've eliminated the need for a request to travel thousands of miles cross-country, or across countries.

### Next-gen software caching platform

The foundation of any CDN is the caching platform. Years ago, traditional CDNs deployed caching architectures that were inflexible and required extensive programming to accomplish the smallest caching task. Over time, the caching platforms became even more bloated as new features were piled on.

In contrast, next-gen CDNs deploy the latest caching architectures, taking advantage of advanced programming frameworks, best practices, and the latest server technology. This combination has enabled next-gen CDNs to create high performance global delivery networks that cache static and dynamic

content. Website visitors have continuous real-time content that's never stale. Besides performance, the next-gen CDN gives DevOps teams end-to-end control of applications, so they can push out updates, software modules, or even replace old applications in a second or two, without the need to bring down the service for maintenance. For a popular website with millions of users, this is becoming increasingly important. The figure below illustrates how traditional CDNs and next-gen CDNs deal with content delivery.

### Traditional CDN middle-mile TCP/IP versus next-gen CDN optimization



(Source: Gigaom Research)

### Next-gen CDN infrastructure

As Moore's Law has proven, today's technology world can do more with less – and in no other industry is this more evident than in the CDN industry. What once required multiple racks for compute, storage, and performance now requires one rack. Solid-state drives (SSDs) are the new norm for CDN storage. A traditional CDN is at a disadvantage because it has built its entire network on non-SSDs. Replacing tens of thousands of drives in a global infrastructure is no easy task. And if a traditional CDN decides to replace standard drives with SSDs, then it might as well replace its older servers in production. The problems for traditional CDNs just avalanche from there. If traditional CDNs claim “more is better,” and that they have more hardware located in more places around the world, they are likely running on old infrastructure.

Today's hardware technologies are dramatically better than they were even three years ago, but the rapid change in server and drive technology has presented two challenges for traditional CDNs: they must replace older hardware with newer hardware and they must re-architect part of their software platform to take advantage of higher capacity equipment. This is difficult and extremely time-consuming for a traditional CDN, but not for the next-gen CDN, because its entire infrastructure is built on the latest technology, using 100 percent SSDs, massively scalable SDN switch technology, and dense data centers all over the world at major global peering exchanges. These characteristics enable them to build low-latency hierarchical infrastructure.

# Dynamic, mobile-aware, interactive web applications

Mobile has completely changed the way content is delivered. Websites and web applications must load quickly on smartphones and other mobile devices, not just on desktops. In the mobile world, APIs drive applications, and next-gen CDNs are able to cache APIs, which improves the app's delivery performance, and the end-user experience.

The two key attributes for optimizing APIs are instant cache and instant purge. Companies wanting to improve content-delivery performance for their mobile users should ensure that these two features are available, especially with 5.2 billion smartphones in the market and the number **growing 20 percent yearly**.

Although the industry is making improvements in wireless last-mile content acceleration, it has a long way to go. Architecting websites with mobility in mind is a “must.” Rapidly changing content is here to stay and applications must be architected to support a high level of dynamic interactions. So, CDNs must have infrastructures that support millions and billions of interactions in real-time.

The DevOps engineers responsible for scaling these applications must have as much control as possible to develop robust web applications that can be pushed out to the edge. When control is limited, their ability to create is limited.

## What does control actually mean?

---

- API-based tools
- Instantaneous configuration changes
- Instant purge
- Application logic at the edge

## API-based Tools

CDNs should not only be able to deliver and accelerate a company's API, but they should also give developers access to their *own* API, within the CDN. This is critically important because an API gives DevOps teams more tools that control applications and content. APIs enable them to extend their applications to the edge, integrate within their own application, and perform instant configuration updates, instant purge, and instant caching. A CDN platform lacking strong API-based tools can be a challenge for DevOps teams and limits what they can do with the applications, sacrificing their ability to offer real-time services.

## Instant configuration

Why should DevOps have to wait hours for changes to deploy to their web applications at the edge or work around a maintenance window when it needs to take an application down for updates and upgrades? Next-generation CDNs enable DevOps to push configuration changes to the edge in real-time, where they go into effect immediately. This is largely due to their ability to purge content instantly, in 150 milliseconds.

## Instant purge

Traditional CDNs use unpredictable invalidation frameworks to purge content cached at the edge. In most cases, purge requests can take minutes to hours to propagate throughout their global networks. The more servers in operation, the longer this takes. Nothing is worse for content publishers than having stale content on their websites. They should be able to publish content and purge it instantly, particularly when it is inaccurate or no longer relevant. In one real-world example, a highly popular news site that was using a traditional CDN accidentally published an article. The content publisher panicked when the traditional CDN said it could not purge the content for 12 minutes – which is an eternity for the publisher, who has users requesting up-to-the-minute news.

Older invalidation frameworks, where content is given an expiration time through headers, are inflexible and not real-time. Next-gen CDNs purge old content instantly through a feature called instant purge, making old invalidation frameworks irrelevant. For DevOps, instant purge is a game changer because it gives them flexibility at a granular level to control content at the edge. DevOps can now remove stale content in the cloud within milliseconds. Another benefit is that instant purge makes dynamic content cacheable.

Whereas traditional CDNs built TCP/IP middle-mile acceleration techniques to expedite requests from the browser to the origin server (some located thousands of miles away), instant purge caches dynamic content at a location closest to the website visitor – sometimes only miles away.

## Extending the applications to the edge

DevOps can now extend web applications from a test and development environment to the CDN edge when it is production ready. This means that a business can push application logic to the edge along with content, minimizing latency and improving page load time performance.

Additionally, CDN customers can use Varnish, an open source web accelerator for dynamic content that can improve the performance of their websites, mobile apps, or APIs. CDNs can leverage the power of Varnish Configuration Language (VCL), a robust scripting language, to enable their customers to do more intelligent caching and have parts of the application logic at the edge.

## Real-time services and real-time visibility

Real-time services such as instant purging, instant caching, streaming logs, and real-time visibility are powerful tools that give CDN customers the ability to do real-time monitoring. The next-gen CDN accomplishes this through a real-time visualization library available to the customer. The major benefit of this approach is that customers can create a monitoring system that receives data every second instead of every few hours, the way traditional CDNs do it.

Since traditional CDNs need hours to collect data and deliver it as reporting data, they have been more focused on historical visibility than real-time visibility. The next-gen CDN offers both real-time and historical visibility that provides customers with the most relevant information so they can make informed decisions around content delivery.

## Call to action and key takeaways

The next-gen CDN is the answer for companies requiring high-performance content delivery, purging, and real-time reporting, because they are built from the ground up to handle the demands of dynamic and highly interactive web and mobile applications. CDN infrastructure should be built on the latest software caching platforms with the latest equipment, enabling CDN customers to enjoy real-time services like instant purge, instant configuration changes, real-time reporting and logging, immediate visibility, API-based tools, and application control at the edge.

The next-gen CDN's mission is to become an extension of the customer's application infrastructure, enabling DevOps to roll out the latest applications from the lab to the CDN edge in milliseconds, at the click of a few buttons.

### Key takeaways

- Next-gen CDNs are the answer to delivering interactive web and mobile applications.
- Next-gen CDNs offer CDN customers huge savings by eliminating what traditional CDNs charged for middle-mile acceleration products.
- By partnering with a next-gen CDN, businesses get high performance content delivery, instant purge, instant configuration changes, instant API delivery, application control at the edge, improved visibility, and real-time analytics through API-based tools.
- The result is that CDN customers improve visitor metrics in all categories, including better conversion, retention, and click-through rates, more interactions, and optimized page load times.

## About Ernie Regalado

Ernie Regalado is an Analyst for Gigaom Research and the founder of Bizety.com, one of the most popular CDN blogs globally with a readership that includes executives from all CDNs, 50+ international telecommunication companies and thousands of technology professionals. Ernie is also currently at CenturyLink, where he is a Global Account Director focused on developing new partnerships with leading financial, e-commerce, internet and manufacturing companies. He has 19 years of experience in technology sales, business development, and presales engineering. During this time, he has assisted hundreds of the Fortune 1000, mid-market, internet, e-commerce, media, and entertainment companies in acquiring technology products, services, and solutions.

Prior to CenturyLink, Ernie worked for three Tier 1 Carriers, four CDNS, and a handful of Systems Integrators. He is recognized as a CDN expert, with deep knowledge in video streaming, small file delivery, large file delivery, e-commerce delivery, CDN security, CDN sales, and product marketing.

## About Gigaom Research

Gigaom Research gives you insider access to expert industry insights on emerging markets. Focused on delivering highly relevant and timely research to the people who need it most, our analysis, reports, and original research come from the most respected voices in the industry. Whether you're beginning to learn about a new market or are an industry insider, Gigaom Research addresses the need for relevant, illuminating insights into the industry's most dynamic markets.

Visit us at: [research.gigaom.com](http://research.gigaom.com).

© 2014 Giga Omni Media, Inc. All Rights Reserved.

This publication may be used only as expressly permitted by license from Gigaom and may not be accessed, used, copied, distributed, published, sold, publicly displayed, or otherwise exploited without the express prior written permission of Gigaom. For licensing information, please [contact us](#).