

AKAMAI WHITE PAPER

Content Delivery for an Evolving Internet
Choosing the Right CDN for Today & Tomorrow



TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
THE INTERNET OF TODAY AND TOMORROW	1
CONTENT DELIVERY FOR AN EVOLVING INTERNET	2
CDN REQUIREMENT #1: HIGHLY DISTRIBUTED ARCHITECTURE	2
Better Caching Performance	2
Better Dynamic Content performance	3
Better Mobile Cellular Performance	4
CDN REQUIREMENT #2: CUTTING-EDGE PERFORMANCE SERVICES	5
Web and Mobile Experiences	5
Rich Media Experiences	7
CDN REQUIREMENT #3: SOPHISTICATED SECURITY CAPABILITIES	9
Internet-scale DDoS Defenses	9
High-performance WAF with a High-accuracy Rule Set	10
Cloud Security Intelligence	11
Bot Management	11
CDN REQUIREMENT #4: SUPPORT FOR AGILE BUSINESS	12
Fast and Flexible Control	13
Robust Support for Testing and Canary Deployments	13
Full-featured APIs and Reporting	13
Dedicated Expertise and Managed Services	13
WHY AKAMAI	14

Executive Summary

As the Internet continues to evolve at a rapid pace, the choice of a Content Delivery Network (CDN) partner is a critical decision for organizations looking to deliver compelling online experiences to their customers, partners, and employees. While no one can accurately predict what the Internet will look like in five or ten years, partnering with the right CDN provider – one that has a proven track record of staying ahead of the innovation curve — will help organizations successfully harness the Internet’s potential. In this whitepaper, we define the core requirements for such a CDN – a highly distributed architecture, cutting-edge software services, sophisticated security capabilities, and support for agile businesses – and establish why these particular requirements are critical for helping businesses succeed in today’s fast-changing marketplace.

The Internet of Today and Tomorrow

Over the past decade, the Internet has evolved rapidly and tremendously, today connecting over 3.3 billion people¹ through nearly every facet of their lives. Behind the scenes, the content delivery network (CDN) market has had to evolve just as rapidly in order for CDNs to continue their work as fundamental enablers of the Internet – optimizing and securing transactions as well as helping organizations harness the potential of the web as their sites have transformed over time from static repositories to rich, interactive, full-featured applications accessible over a wide variety of devices.

Now, as we look to the end of the decade and beyond, we expect the Internet to change at an even brisker pace – from the devices that access it and the software that runs on it to the fundamental technologies and protocols upon which it is built. Such changes will offer exciting opportunities for agile businesses able to leverage it, but the increasing complexity creates challenges as well. In particular, organizations will need to overcome obstacles inherent in the following key trends:

- **An explosion of devices and network types.** With 13.4 billion connected devices worldwide today – a number expected to triple by 2020² – the Internet must support an increasingly diverse set of interactions, from web and mobile to wearable tech, machine-to-machine, and Internet of Things. Optimizing interactions across a fragmented device marketplace and continually changing contexts is a complex task for the Internet of today – and becoming even more complex for the Internet of tomorrow.
- **Richer and more sophisticated content.** In the last three years alone, web page weight has doubled³ and websites are employing richer and more sophisticated media, stylesheets, JavaScript, and third-party APIs in order to create the engaging experiences users of today expect. If the content delivery technology underlying these advancements does not evolve, these richer, heavier pages will be slower as well. Similarly, the rising availability of last-mile broadband and HD devices continues to raise the bar on video quality, so companies may expect video capacity requirements to grow 5-10X within the next few years.
- **Evolving protocols and formats.** As the Internet continues to grow well beyond its original intentions, some of its fundamental protocols have had to evolve to keep up. Over the years we have seen IPv6, TLS, and DNSSec – among others – introduced to address existing shortcomings, while changes such as HTTP/2, new video and image formats, and evolving streaming protocols are happening now. In each case, the transition can take years if not decades to complete, and in the meantime, uneven support across browsers and devices can make it challenging to deliver optimal and error-free user experiences consistently.
- **Attacks of increasing scale and sophistication.** As online data and transactions increase in value, websites and other online assets are becoming the target of larger, more complex, and more frequent attacks. For example, reflection techniques have enabled DDoS attacks to grow by an order of magnitude, and these DDoS assaults are now often used as diversionary cover for more insidious breaches aiming at data theft or site alteration. Some have estimated that by 2019, cybercrime will cost businesses \$2.1 trillion globally, roughly four times the annual estimated cost today⁴.

- **Rapid pace of change.** The always-on Internet has accelerated the pace of business across nearly every industry, as real-time data feeds, just-in-time services, and the adoption of trends like Infrastructure-as-a-Service (IaaS), DevOps, and Continuous Delivery underscore the push for more frequent site content and functionality updates. To keep pace with the constantly evolving landscape, companies need a site infrastructure that supports rapid iterations and cuts complexity without sacrificing flexibility.

Content Delivery for an Evolving Internet

Given the increasing complexity of the Internet landscape, being able to deliver secure, high-quality interactions to every user, everywhere is more difficult than ever before – so partnering with the right CDN provider is business critical. As a baseline, the right CDN improves the user experience, adapting optimizations to the real-time context and conditions – whether for a rich website to a smartphone over cellular or a 4K stream to UHD displays at broadcast scale. It must also eliminate complexity, secure websites and applications hosted in the cloud, and enable the agile enterprise. Even more importantly, it must always be forward looking, always evolving. It is the CDN for today and also tomorrow, understanding and anticipating the continually changing needs of the Internet and its users and enabling its customers to thrive in that environment.

While there are a number of offerings in the CDN marketplace, significant differences in platform architecture, software capabilities, and support services translate into striking disparities not only in website performance and the end-user experience but in all of the crucial factors just mentioned. To meet the needs of businesses operating in the Internet of today and tomorrow, a CDN must satisfy four key components, working in concert:

- **A highly distributed architecture**, the underlying delivery platform for optimizing performance, reliability, and scale.
- **A cutting-edge set of performance services** able to provide the optimal experience while simplifying the complexities of delivering web and media content.
- **Sophisticated security capabilities** that have the scale, visibility, and expertise to protect against evolving attacks in real time.
- **Support for agile business**, whether the organization requires granular, self-service platform control or high-touch, high-expertise managed services.

We will now look at why each of these components is critical.

CDN Requirement #1: Highly Distributed Architecture

Since the dawn of the CDN market, delivering content to users from nearby servers has always been the key to providing the best possible performance. By being close to the end user – in both geographical and network topological senses – close proximity minimizes latency and avoids congested peering points, Internet routing problems, and other middle-mile bottlenecks. Consequently, having a highly distributed platform has always been the single most important architectural attribute for CDN performance, scale, and reliability.

This holds true now more than ever, as users, devices, and networks become more distributed and content gets more dynamic. Many so-called “next-generation CDN providers” fail to meet the baseline requirement of a highly distributed architecture – instead deploying a centralized CDN architecture with perhaps only 10-30 POPs, or points of presence, to deliver content from. This is largely because it takes a tremendous investment of time, expertise, and capital to deploy a highly distributed platform – requiring the development of relationships with thousands of network providers as well as highly sophisticated software to run the platform efficiently. Unfortunately, centralized architectures are a subpar shortcut: their performance and capabilities simply do not measure up.

Better Caching Performance

A highly distributed CDN architecture is critical to get as close to as many end users as possible. Today, no single network has more than 6% of (non-cellular) Internet access traffic, and the top 30 networks combined add up to only 46%. It takes more than 600 networks to cover 90% of Internet access traffic. This means even the largest centralized CDNs,

with several dozen POPs around the world, are still not within a single network hop of the majority of Internet users. Their “edge servers” actually sit in the centralized backbones of the Internet, not at the Internet’s edge; as a result, delivering content to users often requires going through congested peering points and relying on BGP (Border Gateway Protocol) routing. However, since BGP is not a performance-based protocol, it does not always provide the lowest-latency routes, nor can it respond quickly to outages, errors, or congestion. Physical distance to end users matters as well, since the farther data has to travel, the more latency is introduced. Because of the way TCP is impacted by latency and packet loss, with its slow start, connection setup overhead, and lost-packet retransmission, latency can have an unexpectedly severe effect on performance, particularly for “chatty” web applications and high-quality video. Thus, having a highly distributed platform, along with the ability to accurately map users to nearby servers, is absolutely essential to achieving high levels of performance.

Better Dynamic Content performance

The performance benefits of a highly distributed architecture hold not only for cacheable content that can be delivered directly by the CDN but also for uncacheable content that requires a full round trip back to the origin. In fact, a highly distributed platform is also essential for the acceleration of dynamic content. CDNs can speed server-to-server communications within their platforms using various route and transport protocol enhancements – optimizing TCP parameters, multiplexing connections, or routing around BGP inefficiencies, for example. These optimizations only work within the CDN platform, however, and don’t apply to the data as it travels between the CDN and end user, so having servers close to end users is critical.

The importance of this is revealed when we examine real-world last-mile performance – in contrast to backbone-centric measurements that third-party performance testing platforms often employ. Figure 1 shows North American download times for a dynamic (uncacheable) page served by Akamai compared with that of a competitor having POPs in fewer than 10 North American cities. Akamai saw a modest 6% edge over the competitor when looking only at testing agents deployed within backbone networks. But when broadening the measurements to include agents distributed across many networks – where users are – Akamai has a 63% advantage, reducing page load time from over 7 seconds to fewer than 4.5. Moreover, these results are for North America only – a relatively well-connected region. Internationally, we would typically see an even greater performance differential between a centralized platform and a highly distributed one.

Backbone vs. Last Mile Testing

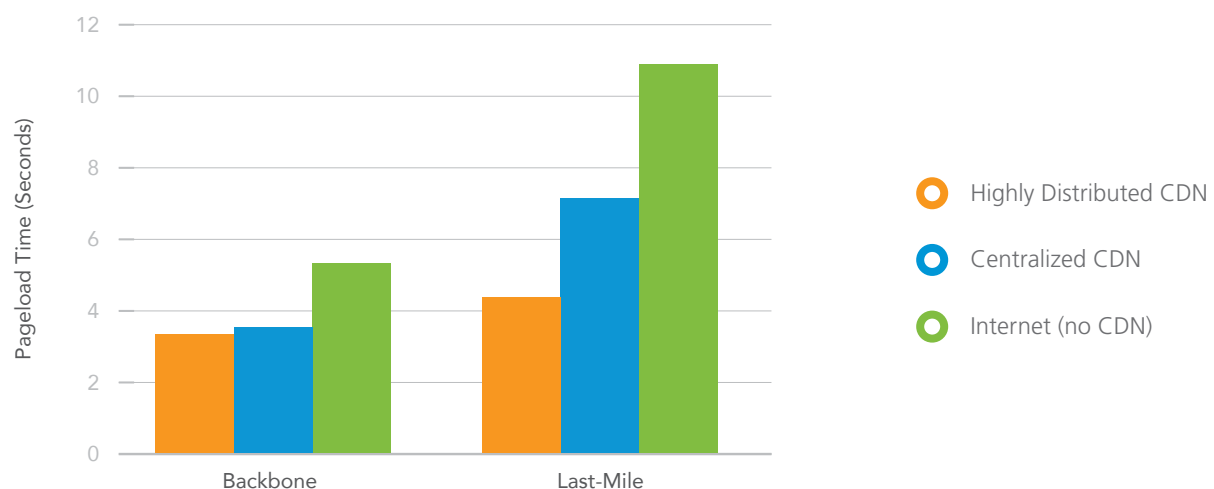


Figure 1: Last-mile testing reveals the real-world performance benefits of a highly distributed architecture compared with a centralized one.

Better Mobile Cellular Performance

Currently, roughly half of all web access on the Akamai network occurs over mobile devices, with about 30% of mobile access occurring over cellular networks. Moreover, mobile network traffic is projected to grow at a 57% CAGR over the next several years, with 70% of the world's population expected to own a smartphone by the end of the decade⁵. As these users increasingly go online over mobile connections, CDN providers will need to extend their platform edges even further.

Achieving good performance for mobile cellular users is particularly challenging due to lower network speeds and higher variability in network congestion rates. Deploying close to the user becomes even more important since high latencies mean high penalties for lost-packet retransmission. The first step is to deploy servers near the mobile gateways and to intelligently map users to the best ones – a nontrivial task because the gateways are not always located in the same city – or even the same state or region – as the users they serve. Beyond the mobile gateway, even better performance can be achieved with CDN servers deployed within the cores of the mobile networks themselves, further reducing latency to the mobile cellular user.

59% Performance Improvement in the Mobile Core

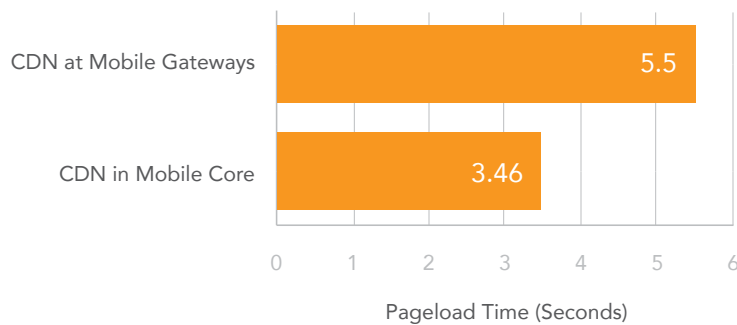


Figure 2: By getting closer to mobile users, CDN servers within the mobile core deliver even faster download times than those at mobile gateways.

Even closer to the user is the radio access network (RAN), and extending CDN reach into the RAN – whether through real-time communications or direct deployments – can offer cutting-edge performance improvements for mobile. For example, the radio scheduler in the RAN is a useful resource for real-time data about available bandwidth – a metric that can vary wildly from one moment to the next, as it is highly sensitive to changes in radio tower connections, signal strength, and interference as well as the number of users sharing a local connection. The rapid fluctuations in available bandwidth make delivering video over cellular particularly challenging, as adaptive bit rate (ABR) technologies often cannot detect and adapt to the fluctuations quickly enough. However, CDNs that have the reach to leverage real-time bandwidth intelligence provided by the local radio scheduler can deliver video streams that are automatically and continually optimized, overcoming fast-changing cellular conditions far more effectively than current ABR technologies can.

CDN Requirement #2: Cutting-edge Performance Services

Just as important as a highly distributed CDN platform are the software services that run on top of it. While many CDN providers support a primary set of performance-enhancing features such as caching, dynamic site acceleration, and adaptive bit rate streaming, the varying levels of intelligence built into these services differentiate their real-world performances. Moreover, as online interactions become increasingly diverse and sophisticated, companies need CDNs with forward-looking capabilities that cut complexity and allow them to leverage advancing web technologies to easily deliver the most engaging Internet experiences possible.

Web and Mobile Experiences

A decade ago, websites were relatively simple and static, and optimizing website performance was primarily about eliminating Internet latency in the “middle-mile”. By efficiently caching content close to end users and intelligently mapping users to the closest servers, CDNs could greatly reduce latency and improve the end-user experience. Today, Internet latency is still hugely important, but the situation is far more complicated, as sites and mobile apps are getting richer, more dynamic, and more complex with increased use of APIs and third-party content calls. Web page sizes have doubled in the last three years alone, due to steady weight increases in images, JavaScript, CSS, and custom fonts⁶. Even worse, Responsive Web Design sites may suffer from “over-downloading” – a site design technique that delivers an optimized experience to both desktop and mobile users and can result in mobile devices requesting unneeded rich media assets meant for larger display devices. Moreover, devices themselves are getting more diverse. In 2012, there were roughly 4,000 different mobile devices in the marketplace; in 2015, there were more than 24,000^{7,8} – creating a hyperfragmented landscape of form factors, browsers, operating systems, and device capabilities to support.

Delivering a speedy and engaging experience to every user, every time in this complex and fast-evolving marketplace requires a CDN with a broad set of intelligent services that work in concert to optimize each end-user experience. These services include advanced caching, dynamic site acceleration, front end optimization, image management, API and mobile app acceleration, and predictive acceleration.

Advanced caching capabilities. While caching is a basic CDN feature, advanced caching capabilities allow a CDN to cache more content – and cache it more efficiently – even as sites become increasingly dynamic. Most CDNs support the ability to set TTLs and ignore or follow various cache control headers, but differentiation comes in the granularity of control over cache rules and cache keys. An advanced CDN will also have powerful cache control engines to support a broad range of cache behaviors through flexible, nested rules with sophisticated pattern matching, and they offer the ability to key off of various request features including cookie values, query string, geo-location, partial URL, HTTP header values, or any combination thereof. This enables caching of many types of content that are typically thought of as dynamic — such as search results, API calls, product category pages, content targeted to different audience segments, and frequently changing content.

For example, by looking for the presence of a “logged in” cookie, a CDN can cache and serve all non-logged-in users one version of a site while fetching personalized content from the origin server for logged-in users. In many cases, a large subset of content may be the same for logged-in and non-logged-in users, and a CDN with advanced cache key mechanisms can serve all of this content from the edge, boosting website performance significantly.

Dynamic site acceleration. Truly uncacheable content requires a combination of dynamic site acceleration techniques including route and transport-layer protocol (TCP) optimizations. Several CDN providers claim to use TCP optimizations, but they are only truly effective with a highly distributed network that sits close to end users, minimizing the distance data travels over unoptimized routes. Moreover, whereas TCP optimizations work primarily by reducing the number of round trips required to render a webpage, another key dynamic site acceleration technology — route optimization — works by actually reducing the latency of each round trip.

Route optimization uses real-time network latency and congestion data to overcome BGP’s inefficiencies. Akamai’s unique route-optimization solution, SureRoute, utilizes Akamai’s highly distributed platform to form an overlay network to the Internet. This allows it to effectively override BGP by sending traffic through faster routes via intermediary servers when the “direct” BGP route is congested or otherwise nonperforming. Route optimization complements TCP optimization – offering a potential multiplier effect on performance improvement when used together – but is particularly critical for chatty Ajax applications and other short, bursty traffic that is highly sensitive to round-trip times. It can deliver significant boosts in performance and reliability for uncacheable traffic, even across short, well-connected geographies, as shown in Figure 3.

Route Optimization: New York to Miami | Aug.1 - Aug.14

Latency (96ms Peak Differential, 290% Peak Improvement)




Figure 3: Using route optimization, Akamai can improve latencies for uncacheable content, even across “well-connected” geographies

Front end optimization. As the matrix of devices, operating systems, browsers, and networks increases, it becomes increasingly challenging to deliver optimized experiences to every user, every time. A capable CDN simplifies this task dramatically for its customers, leveraging up-to-date device characterization along with situational optimization capabilities including front end optimization (FEO) and a broad set of last-mile-focused techniques. Compression, JS/CSS minification, inlining objects, on-demand image loading, domain sharding, asynchronous JS/CSS, and deferral of third-party content calls are among those techniques available. A comprehensive suite of FEO capabilities complements caching and dynamic site acceleration to extract the best possible performance across many diverse web and mobile scenarios.

FEO techniques work by reducing payloads, minimizing connection overhead, preventing “over-downloading”, reducing browser think time, and generally improving perceived performance through smart prioritization of tasks. They are continually evolving, and some CDN providers offer more advanced capabilities that can provide an extra performance edge. For example, to accelerate personalized web pages, Akamai’s unique EdgeStart feature takes advantage of the fact that initial bits of HTML in the page – including the stylesheet and other resources – are likely to be the same across all users, and its edge servers can therefore begin delivering that portion to the user while simultaneously fetching the rest of the page from the origin. This allows the browser to start rendering the page sooner, enhancing end-user response time.

Furthermore, it is critical to partner with a CDN provider that will stay on top of the evolving landscape. For instance, the HTTP protocol is currently transitioning from HTTP/1.1 to HTTP/2 – its first update in over 15 years. HTTP/2 offers a number of performance benefits over its predecessor by supporting multiplexing, pipelining, header compression, and server push. However, this means that certain HTTP/1.1 optimizations and best practices can result in suboptimal performance over HTTP/2. This is particularly relevant to FEO but applies more broadly. As browsers and other clients transition over the next several years, CDNs must be able to help their customers take advantage of HTTP/2’s performance benefits without incurring the costs of a complex transition. More generally, as different browsers work differently and support different capabilities, CDNs need to be able to optimize in real time for each end-user scenario, with HTTP protocol version being just one of many determining factors.

Image management. Today, images make up the bulk of web traffic, representing nearly 65% of web page weight⁹ – and managing these images is becoming a growing headache. In order to optimize for the wide range of browsers, devices, and networks in use today, every image on a website may require dozens to hundreds of derivative images with varying resolutions, aspect ratios, compression qualities, background colors, and even image formats – such as Google’s WebP or Microsoft’s JPEGXR – which allow improved image quality compared with JPEGs of the same size. Because of this, manipulating, storing, and managing images have become resource-intensive and error-prone undertakings. Advanced CDNs can reduce complexity for content providers by offloading this task – enabling content providers to focus on simply creating the original image and leveraging the cloud platform to automatically convert, store, and intelligently deliver images optimized in terms of quality, speed, and form factor for each end user, whether they are using a 4K display, tablet, or watch on high-speed broadband or congested cellular.

API and mobile app acceleration. In recent years, API traffic has grown tremendously on the Internet through both mobile app usage and B2B/M2M data streams. While APIs are generally dynamic in nature, many types of API responses – for example, those returning product descriptions or store locations – are actually cacheable for CDNs with sophisticated caching mechanisms. Non-cacheable APIs can also be accelerated through route and transport-layer optimizations, like those used in dynamic site acceleration, as well as techniques like compression of API response text and intelligent compression of images (for mobile apps). Advanced CDNs can boost API performance in other ways as well such as authenticating requests at the edge, metering or prioritizing API requests during peak traffic, and enabling backward compatibility when rolling out new API versions. Since mobile app traffic consists largely of images and API traffic, CDNs can boost app performance much like they do for websites, through a combination of image caching, image management, and API acceleration.

Predictive acceleration. Looking toward the future, data-driven predictive techniques will be used to push the performance envelope even further for uncacheable content. For example, today, some CDNs offer prefetching of pages to their servers based on preconfigured rules whereby customers identify pages to prefetch following specific requests. With predictive acceleration, the CDN will make these prefetching choices on the fly, leveraging analysis of similar requests to make smart decisions about which pages to prefetch to the edge – or even directly to the device itself, using HTTP/2's new server-push capabilities. By removing the impact of the network wherever possible, predictive prefetching can have a tremendous impact on the end-user experience, particularly over slow cellular connections.

Rich Media Experiences

By 2019, it is estimated that video will account for 80% of all consumer Internet traffic¹⁰ – a staggering statistic that represents both growing audiences and their increasing demand for quality. Within a few years, video bitrate requirements are expected to grow from today's 1 or 2 Mbps to ten times that, fueled by growing numbers of HD mobile device screens and 4K/UHD displays along with rising last-mile broadband speeds. This means publishers may see their video capacity requirements grow ten to one hundred-fold – or more – by the end of the decade, while revenues are likely to rise more slowly. This means publishers need a media delivery partner that will not only help them scale to meet fast-growing audience and quality demands but also scale in a way that supports their near- and long-term business models.

In addition, delivering the experience viewers want is becoming more and more challenging, as complex video workflows, fragmented device and format landscapes, and inefficient delivery protocols hamper organizations' abilities to deliver the high-quality, buffering-free video experiences their audiences demand. The right CDN provider can help companies stay ahead of the curve by offering services focused on quality – from a workflow that preserves quality from pre- to post-delivery and delivery protocols that enable quality at scale to prepositioning capabilities that deliver quality, even when the network can't.

Quality-focused cloud workflow. While many CDN providers focus on the delivery piece of the streaming workflow, achieving the best viewer experience requires diligence from the beginning of the workflow – and maintained throughout. Akamai estimates that roughly 70% of streaming quality issues on its network result from “predelivery” issues that arise during video transcoding, packaging, or storage. These predelivery workflow steps have become enormously complex, as publishers need to package streams to support a wide range of form factors, screen resolutions, and network capabilities – not to mention formats, codecs, and protocols. Secure, reliable storage of the vast collection of resulting files can be a headache as well. By automating and tightly integrating these resource-intensive and error-prone tasks – while fully supporting additional business requirements such as DRM, close-captioning, and ad integration – CDNs can ensure the video delivery chain starts with best possible quality video, and publishers need only worry about producing a single-source video, greatly simplifying their task and speeding time to market. Some CDN providers have also had their cloud workflows audited by organizations such as Farncombe and the MPAA (Motion Picture Association of America), ensuring they meet the stringent security standards required for the delivery of premium video content.


Efficient streaming protocols. The vast majority of video delivery today occurs over TCP-based HTTP, using some type of adaptive bitrate technology. However, as quality expectations continue to push video bit rates higher, the limitations of using TCP-based HTTP for video delivery are becoming apparent. Designed for reliability rather than performance, TCP connections carry significant overhead, particularly when congestion is encountered. With TCP, packet loss requires retransmission and can have a severe impact on performance, resulting in poor-quality streams and buffering delays. In contrast, the lightweight UDP protocol was designed for real-time communications and allows packets to be dropped

in case of congestion so that stream latency is not affected. Evolving technologies now combine UDP with forward error correction techniques to get the best of both worlds, blending TCP's reliability with UDP's speed, even across congested Internet routes.

With an enhanced UDP transport layer, HD and UHD video can be streamed across the Internet without the need to buffer or reduce bitrates when congestion hits. It will be years before the Internet as a whole evolves to support such protocol upgrades, but CDNs that have resilient UDP transport built into both their server platform and client (video player) footprint can help publishers benefit ahead of the curve. Combined with delivery of streams from servers close to each viewer, resilient UDP enables publishers to offer HD+ quality streams to broadcast-size audiences across the globe.

Enhanced UDP improves stream quality for video ingest as well, a critical step in delivering high-quality live video. This makes it possible to produce live HD and 4K streams from any location, no longer requiring a dedicated broadband link all the way to the CDN in order to ensure a full-quality ingest stream.

REBUFFERING



STARTUP TIME



AVERAGE BIT RATE

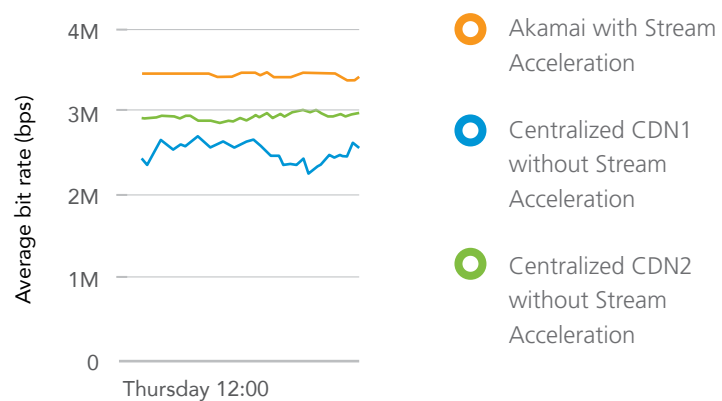


Figure 4: Stream acceleration using enhanced UDP enables much faster start up times, reduced buffering and higher throughput rates for delivering video resulting in higher audience engagement.

Enhanced UDP marks a fundamental change in online streaming, but it is just the beginning. Looking toward the future, capabilities like network-assisted multicast and controlled P2P will help improve quality and scale even further as well as facilitate the economics of broadcast-scale video.

Prepositioned content. For scenarios where last-mile bandwidth is constrained by data plans or simply inadequate, the ability to download video content to the device in advance delivers a viewing experience far better than that supported by the available bandwidth. News clips can be downloaded to a smartphone when it is on Wi-Fi, for example, to be watched later when the user only has 3G cellular coverage. Movies can be automatically downloaded to a set-top box during the day, enabling a household to watch multiple different HD/UHD videos at the same time later that evening.

While the concept is simple, execution is not. iOS places stringent restrictions on the ability to download content in the background, for instance. An intelligent download manager, built into the video player, can navigate OS restrictions while taking into account network availability, storage, battery life, and other requirements to optimize download capabilities. Combining prepositioning capabilities with personalized predictive analytics takes things one step further, helping companies determine the most useful and likely-to-be-watched content for each viewer in order to optimize delivery.

CDN Requirement #3: Sophisticated Security Capabilities

As the volume of high-value data and transactions on the Internet continues to grow, so do the forces of attackers looking to exploit it – and these forces are costing organizations big money. In FY2015, businesses around the world suffered average losses of \$7.7 million due to cybercrime, with U.S. companies seeing the largest losses, averaging \$15 million¹¹. Along with crimes committed by malicious insiders, DDoS and web-based attacks were found to be the most costly.

Unfortunately, the number of DDoS attacks seen across the Akamai network has more than doubled in each of the last two years. Web-based exploits such as SQL injection, cross-site scripting, and local or remote file-inclusion attacks are becoming more common as well, as hackers exploit website vulnerabilities in an attempt to deface, disrupt, or steal from a site. They are also increasingly launched in conjunction with DDoS assaults, using the latter to divert attention while causing more serious damage with the former. In both types of attacks, it is often difficult to distinguish bad traffic from legitimate traffic, and strategies continue to evolve rapidly over time, requiring significant dedicated security resources in order to stay up to date on mitigation strategies.

Given the increasing volatility of the Internet threat landscape, helping to secure websites is a critical CDN requirement. This is a broad topic that spans protection of content, physical security, operational security, compliance (with regulatory requirements such as PCI, ISO, BITS, FISMA, and HIPAA), and even acceptable use policies. While there are many important pieces of the security puzzle, here we focus on cloud-based CDN defense layers that are unique in their ability to protect against the Internet-scale threats of today and tomorrow.

Internet-scale DDoS Defenses

In recent years, the largest DDoS attacks have grown exponentially in size as amplification techniques have allowed hackers to create onslaughts that are hundreds of times larger than before. At the same time, they have become more sophisticated – often coming in multiple waves, using multiple attack vectors, and opportunistically attacking during high-traffic launches and events, when infrastructure is already heavily loaded. They may target the network layer, the application layer, or the DNS infrastructure.

DDoS Size and Frequency as a Function of Time

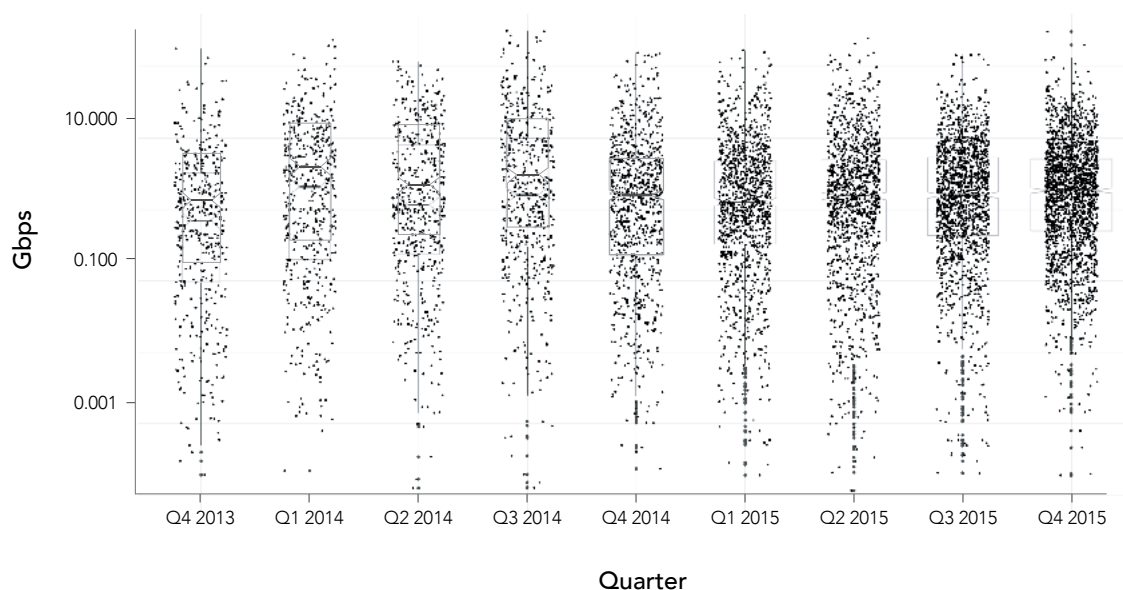


Figure 5: The boxes for each quarter represent the middle 50% of attacks by attack size, while each dot represents an individual attack. The size axis has a logarithmic scale; the upper attacks are many thousands of times larger than the lower ones

With Internet-wide visibility and scale, a highly distributed CDN can provide always-on protection against even the largest of DDoS attacks without affecting the performance for legitimate users. Network-layer assaults like SYN flood attacks are immediately dropped at the CDN edge, while application-layer attacks are mitigated through signature-based filters, geo-blocking, IP blacklists/whitelists, and adaptive rate controls. CDNs can also provide resilient DNS capabilities that accelerate DNS resolutions while protecting against denial-of-service attacks.




Figure 6: DDoS attacks now reach hundreds of Gbps – too large for the vast majority of organizations to handle on their own – but a small fraction of network capacity for a CDN that routinely delivers dozens of Tbps of traffic.

Companies requiring infrastructure protection across all ports and protocols, including UDP-based game server traffic (for example), should look for a CDN provider that can provide on-demand or always-on capabilities through a managed service. Akamai's Prolexic Routed service does just that, offering high-performance distributed scrubbing centers that filter out illegitimate traffic, passing clean traffic through to the origin. Backed by industry-leading time-to-mitigate SLAs, Prolexic Routed has thwarted attacks peaking at over 200 million packets per second, large enough to have taken out Tier 1 routers used by major ISPs.

High-performance WAF with a High-accuracy Rule Set

Another critical layer of defense a CDN should offer is a cloud-based Web Application Firewall (WAF) service, designed to reduce the risk of data breach and protect sites from exploits like SQL injection, cross-site scripting and command injection attacks. By leveraging potentially hundreds of thousands of servers across a highly distributed platform, a CDN can provide inline, scalable WAF protection that can handle even the toughest peak traffic situations, while still delivering a rich and responsive end-user experience.

WAFs rely on dynamic rule sets to distinguish between legitimate and illegitimate traffic. Unfortunately, many WAF implementations fail to provide robust protection due to the lack of an effective and up-to-date rule set. Providing timely updates is a difficult task, given the thousands of continually evolving potential exploits out there.

Traditionally, WAFs have to make a difficult trade-off between false positives and false negatives. Many implementations end up allowing more malicious traffic through in order to minimize the impact on legitimate traffic (i.e., permitting more false negatives to reduce false positives). Moreover, most organizations simply do not have the resources and expertise to keep up with the constant evolution of threat vectors, so their WAF rule set quickly falls out of date and becomes ineffective. Thus, when evaluating WAF solutions, companies need to consider not just scalability and performance but also accuracy and ease of management.

Akamai's Kona Rule Set (KRS) takes a different approach from traditional WAF solutions, using a small number of flexible rules with an anomaly scoring model to improve accuracy and visibility into attacks. Rather than using a separate, more rigidly defined rule for each exploit, Akamai detects exploits based on attributes shared among vulnerabilities, making it more effective not only against known attacks but also new and evolving ones. Its unique scoring model looks at weighted risk scores combined across all of the different rules triggered by a request, generating a more accurate profile of risk. Akamai continuously measures the accuracy of its WAF through automated closed-loop testing using a large set of real-world data, updating scores and weights based on its wide-scale visibility into the legitimate and malicious HTTP requests that run across its network each day. Using this approach, Akamai's rule set achieves far greater accuracy with less noise – significantly reducing false positives and false negatives – providing more robust web application defenses without affecting legitimate traffic.

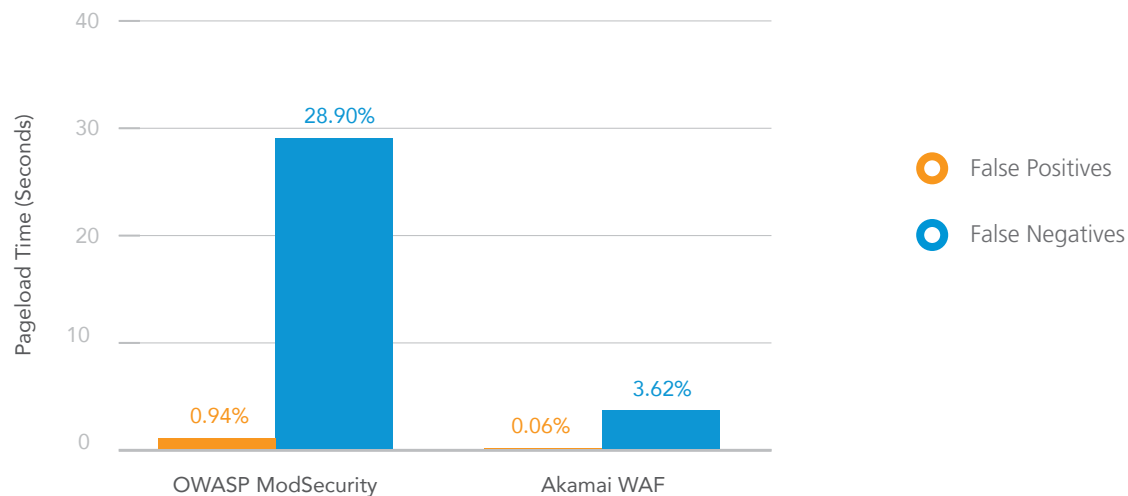


Figure 7: Akamai WAF achieves close to 90% fewer false negatives and nearly 95% fewer false positives than an open source OWASP ModSecurity Core Rule Set through continuous, closed loop testing.

Cloud Security Intelligence

Ultimately, as the online threat landscape continues to evolve and grow, collective data intelligence will become one of the most important weapons in the war on cybercrime. Big data – gathered from across the Internet and analyzed and processed with the right tools – can help identify attack trends, malicious actors, and other important indicators, both in real time and over time. For example, by serving as much as 15 to 30% of the web’s global traffic each day, Akamai has unparalleled insight into Internet traffic patterns around the world. Enriching this data with security statistics, including offline analysis across all its traffic as well as triggers and other security events for its WAF customers, Akamai has created a massive security intelligence platform that incorporates petabytes of historic data with 20 terabytes of attack data added each day. The result is unmatched visibility into Internet threat trends that can be leveraged in many different ways – not only to help mitigate ongoing attacks and improve WAF accuracy as described above, but also to forecast malicious intent and prevent exploits before they occur.

Through its cloud security intelligence platform, Akamai is able to provide real-time client reputation scoring capabilities — for instance, identifying bad IP addresses and assigning them a risk score that predicts the likelihood of their participation in different types of attacks based on their past behavior. To avoid tagging legitimate users as malicious, the risk score calculation looks at many factors, including persistent bad behavior, the severity and magnitude of attacks participated in, and whether or not similar clients are performing attacks. By using the real-time client reputation scores, organizations can decide how they want to handle the request, improving their threat detection and security decision-making.

Bot Management

Bots now play a significant role in the online world, comprising up to 40-60% of traffic for some organizations. Sometimes bots carry out important business tasks, while other times they steal website data, scan for vulnerabilities, perpetrate fraud, or otherwise cause harm. In many cases, whether friend or foe, bot activities also have the unwanted effect of decreasing site performance for human visitors. Unfortunately, effective bot management is far more complex than simply accepting or denying their requests wholesale; organizations need the ability to identify and treat a scraper bot differently than a search-engine bot or an advertising bot differently than an aggregator bot, for example. The ability to accurately categorize the many types of good and bad bots in real time requires significant intelligence capabilities and an in-depth understanding of how Internet bots present themselves. This is another prime application for big-data cloud security intelligence, and CDNs with such capabilities can combine lists of known bots with analysis of bot behavior to help detect and categorize unknown bots. Customers can then leverage this information in real time to apply different policies – such as serving cached content, serving alternative content, sending the request to a different origin, delaying the request, or denying it altogether – to different types of bots as their business strategies dictate.

CDN Requirement #4: Support for Agile Business

As the popularity of trends like Continuous Delivery, DevOps, and IaaS underscore, businesses today need to be agile in order to compete in an era of rapid change and innovation. Whether it's flash sales and daily deals, real-time inventory and pricing changes, or promotional events and product launches, sites are updating features and content ever more frequently — and site infrastructure needs to keep up.

The CDN of today and tomorrow needs to facilitate agility. For some, this means the ability to integrate CDN platform controls and data feeds directly into their DevOps workflow. For others, it means the ability to leverage dedicated CDN expertise to secure and optimize their site so that in-house resources can focus elsewhere. In all cases, it requires a CDN that seamlessly enhances their existing origin and cloud infrastructures while offering the flexibility to offer optimal solutions for a highly diverse set of business use cases – enabling the organization to innovate without bounds.

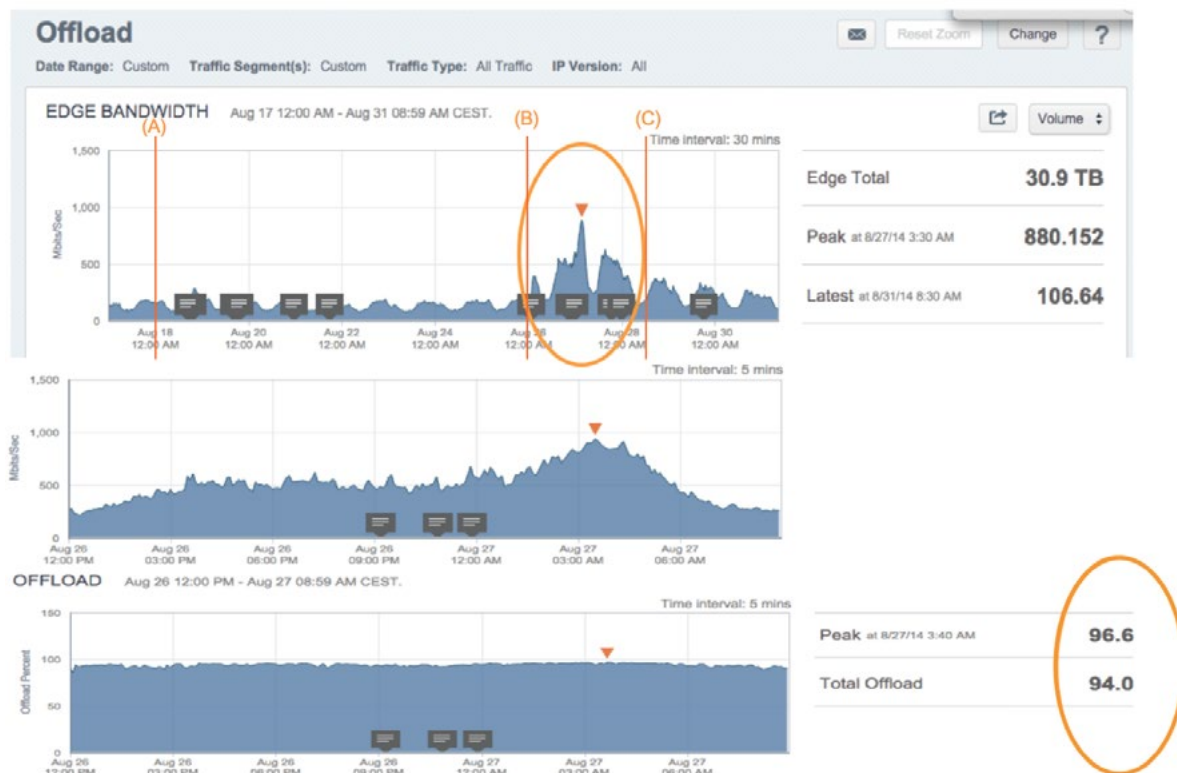


Figure 8: An Akamai customer had to prepare for a major global product launch requiring availability of new product images/information at a precise moment with the expectation of heavy web traffic. (A) The customer and Akamai set up an alternate origin server with new, embargoed content; set up a password and cookie combination for customer technical teams to access the content prior to launch; set passwords' expiration date to the exact time of the product launch and pre-warmed the Akamai platform with the password-protected content to prepare for the switchover. All object TTLs managed via 4-layers of control. Meanwhile, customer's users continued to be served with pre-launch content. (B) All passwords expired precisely at launch and embargoed content became available immediately across the global network; all site visitors received product launch data while Akamai handed peak traffic load, maintaining its protection and offload of the site origin up to 96%. (C) Used fast metadata activation and instant purge for ongoing content changes post-launch.

Fast and Flexible Control

As a baseline, agile business requires an agile CDN platform – one that gives its customers self-serviceable control over its sophisticated capabilities. This is achieved through a combination of features:

- **Advanced cache control**, with the ability to define cache keys and cache control rules at a very granular level, maximizing caching benefits while ensuring fresh content.
- **Fast purge capabilities** that can expunge content across a widely distributed network within a matter of seconds.
- **Flexible content handling rules** that offer granular control over advanced CDN capabilities such as header and cookie handling, performance optimizations, failover behavior, access control, and edge logic.
- **Fast configuration deployment**, enabling cache control and content handling rules to be safely updated across the network within minutes.

While some platforms may have partial capabilities like fast purge, the most powerful CDNs offer all four of the above features, working in concert to deliver the greatest flexibility in meeting different business needs. For example, a website launching a big promotion at a specific time can use a CDN's purge capabilities to remove old content at that time – ensuring fresh, new promotional content within minutes or even seconds of launch. However, a more advanced CDN platform also offers alternatives to better ensure the success of such an important event – such as the ability to easily stage and test the new content on the live network, the ability to prewarm the network for greater origin offload at launch, and the use of a time-based rule to trigger delivery of the new content starting precisely at the desired time.

Robust Support for Testing and Canary Deployments

Just as critical as the ability to deploy changes quickly is the ability to test those changes in a safe and streamlined manner, particularly as organizations move toward Continuous Delivery methodologies and faster, more frequent release cycles. CDNs should facilitate this, not only through staging and test networks but also through safeguards like real-time configuration error checking and support for canary deployments, with the ability to easily roll out (and roll back) new site content in phases, live-testing it with subsets of users before a full-scale rollout takes place.

Full-featured APIs and Reporting

Secure API access to CDN management capabilities allows organizations to further streamline their development process through tight integration of CDN functions like purge, traffic management, failover, and configuration. For example, some of Akamai's customers have integrated purge calls into their content management systems (CMS) so content updates automatically trigger removal of the old content from the Akamai network. Others have broader API integrations where new content is automatically configured within the CMS with the appropriate configuration rules – ranging from caching and cookie handling to mobile device optimization.

CDNs have also become a key source of visibility into real-time usage, performance, and security metrics across an organization's entire infrastructure. To help customers better understand and optimize their online presence, the CDN of today and tomorrow must offer access to their rich, real-time data feeds – both through their own customizable tools and visual interfaces as well as through APIs that enable easy integration into the organization's existing reporting and analytics infrastructure. In addition, CDN providers may provide out-of-the-box plugins to leading third-party performance management solutions.

Dedicated Expertise and Managed Services

Hiring 24/7 staff with the right skillset and expertise to maintain robust site performance and security is costly and time consuming, and many companies simply do not have the resources to do so. However, by partnering with the right CDN provider, organizations can continue focusing on their core business competencies while leveraging the CDN's expert resources to optimize their web applications and online events. Leading CDN providers offer 1) teams with deep experience who have helped to deliver the Internet's biggest events and mitigate its largest attacks, 2) options for dedicated support with response SLAs, and 3) core expertise in several key areas:

Managed delivery services, providing proactive performance analysis to boost conversion rates and reduce abandonment as well as ongoing performance testing, with synthetic and real user monitoring tools to rapidly identify and resolve delivery issues before they impact business.

Managed broadcast services, with 24/7 real-time monitoring across the entire encoding-through-playback workflow, allowing organizations to deliver the highest quality video experience to every audience member. Here, the CDN of today and tomorrow sets itself apart not only through the expertise and experience of its services team but through the sophistication of its monitoring tools. Proactive system component assessments, content integrity and delivery checks, and real-time QoS feedback across the entire viewing audience enable early detection of quality issues along with speedy troubleshooting and resolution.

Managed security services, offering customized WAF rule-set updates, early threat detection, rapid-response attack mitigation, and post-threat recommendations. As today's cyberattacks are more sophisticated than ever before – often shifting strategies midstream or combining multiple attack vectors – there is no substitute for human expertise in combating live threats, and the right CDN will provide access to highly experienced security specialists, staffed 24/7 around the globe with the tactical expertise to minimize an attack's business impact.

Website consulting services, providing an opportunity for companies to leverage CDN proficiency in business-critical areas such as site performance optimization, business continuity risk mitigation, and security vulnerability assessment. Specialists can assist organizations in redesigning or migrating site architecture, preparing for an event or holiday, or expanding internationally – providing the expertise to help minimize time to market and maximize success.

Why Akamai

From the time it pioneered the CDN market nearly two decades ago and throughout the Internet's tremendous advances to date, Akamai has grown and evolved its platform with an ever forward-looking stance. The highly distributed model it championed from the beginning proved to be vital as the web grew more and more dynamic – and this model is even more crucial now with the current rise in mobile traffic. Akamai has continued to expand the intelligent software services running across its distributed platform, building upon core capabilities like advanced caching, dynamic site acceleration, front end optimization, and HD video delivery with innovative services such as image management, API acceleration, and mobile app acceleration, while pushing the leading edge of performance with technologies like enhanced streaming protocols, predictive acceleration, and prepositioned content. Akamai also continues to advance the state-of-the-art in cloud security, not only with its ability to mitigate Internet-scale attacks but also with unique WAF, client reputation, and bot management services – all powered by big data intelligence.

Most importantly, through the power and flexibility of its platform, Akamai helps organizations maintain responsiveness and agility in today's fast-paced environment. Platform control and visibility can be integrated directly into organizations' DevOps workflows through full-featured RESTful APIs, enabling content and configuration updates to be made quickly and easily. Customers also have access to Akamai's industry-leading expertise through a broad array of managed services, enabling them to keep internal resources laser-focused on achieving their core business goals.

In order to help businesses succeed, the CDN of today and tomorrow must constantly push forward and innovate. Today, Akamai continues to work closely with the world's leading brands and most innovative companies, collaborate actively with industry working groups, and invest heavily in research and development – not only to stay on top of the rapidly evolving landscape but to help improve the web and shape the protocols, standards, and best practices in its future. For example, with its engineers highly involved in the IETF HTTP Working Group, Akamai helped define the HTTP/2 protocol and – within seconds after the protocol was ratified – was able to begin serving customer traffic over HTTP/2, seeing immediate performance boosts of up to 68% for mobile cellular users. Akamai's commitment to advancing content delivery for a rapidly changing Internet is proven and long lasting, and by staying ahead of the innovation curve, Akamai helps its customers do the same — harnessing the Internet of today and tomorrow to propel their businesses ever faster forward.

Source

1. <http://www.internetworldstats.com/stats.htm>
2. <http://www.juniperresearch.com/press/press-releases/iot-connected-devices-to-triple-to-38-bn-by-2020>
3. <http://httparchive.org/trends.php?s=All&minlabel=Nov+15+2012&maxlabel=Nov+15+2015#bytesHtml&reqHtml>
4. <http://www.securitymagazine.com/articles/86352-cybercrime-will-cost-businesses-2-trillion-by-2019>
5. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html, <http://www.ericsson.com/res/docs/2015/ericsson-mobility-report-june-2015.pdf>
6. <http://httparchive.org/trends.php?s=All&minlabel=Nov+15+2012&maxlabel=Nov+15+2015#bytesHtml&reqHtml>
7. <https://opensignal.com/reports/fragmentation.php>
8. <http://opensignal.com/reports/2015/08/android-fragmentation/>
9. <http://httparchive.org/trends.php?s=All&minlabel=Nov+15+2015&maxlabel=Jan+1+2016#bytesTotal&reqTotal>
10. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html
11. <http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report/>



As the global leader in Content Delivery Network ([CDN](#)) services, Akamai makes the Internet fast, reliable and secure for its customers. The company's advanced web performance, mobile performance, cloud security and media delivery solutions are revolutionizing how businesses optimize consumer, enterprise and entertainment experiences for any device, anywhere. To learn how Akamai solutions and its team of Internet experts are helping businesses move faster forward, please visit www.akamai.com or blogs.akamai.com, and follow @Akamai on [Twitter](#).

Akamai is headquartered in Cambridge, Massachusetts in the United States with operations in more than 57 offices around the world. Our services and renowned customer care are designed to enable businesses to provide an unparalleled Internet experience for their customers worldwide. Addresses, phone numbers, and contact information for all locations are listed on www.akamai.com/locations.

©2016 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are registered trademarks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice. Published 03/16.