



A Network for Blazing Fast, Immersive Content Experiences

The Internet wasn't originally built to support high-resolution photos, high-definition movies, or crystal clear audio streams. But increasingly, users and businesses rely on the Internet to do just that.

As user expectations for speed, stream quality, and picture quality grow, performance becomes more and more crucial for growing media and entertainment companies.

Here, we will go over why performance matters to your business, the underlying Internet technology that impacts performance, and how you can optimize the performance of your web properties, from improving page load times to offering picture-perfect streaming quality.

PART 1

Why is performance important for media and entertainment platforms?

Users have high expectations for website and app performance. They also expect to be able to interact with a variety of different kinds of media, from podcasts and live streams to HD video and photo galleries. You need to offer these types of content to keep users coming back, but are they slowing down the user experience?

Web performance affects the overall user experience. Websites and applications need to load fast and respond to user actions quickly in order to keep users engaged and increase conversions.

Users have ever-increasing expectations for streaming platforms in particular. Especially for over-the-top (OTT) media providers competing for customers, performance has never been more important.

User engagement

Research shows that users are very willing to abandon applications, websites, and videos that are slow-loading or fail to load altogether:

- 39% of users stop engaging with a website if images take too long to load.¹
- 85% of respondents stop watching video if it takes too long to load.²
- The BBC discovered that 10% more site visitors left for every additional second it took for their webpages to load.³
- 30% of users are likely to stop watching a video if the stream is of poor quality.⁴
- Each 1-second buffering delay for video playback results in a 5.8% percent increase in abandonment rates.⁵

Mobile performance

Because mobile presents unique challenges compared to desktop, it should be considered a separate aspect of performance: a website or app must be specifically designed for mobile to ensure it will perform well on handheld devices.

Mobile devices overtook desktop computers in terms of number of Internet connections in 2016: In October of that year, mobile devices and tablets accounted for 51.3% of Internet usage.⁶ A 2017 survey from venture capitalist firm Kleiner Perkins found that, on average, users spent 3.1 hours per day on mobile and 2.2 hours on desktop.⁷

Consequently, mobile performance is hugely important for businesses:

- 40% of Internet transactions occur on mobile devices.⁸
- Users abandon 53% of mobile sites if they take over 3 seconds to load.⁹

This applies to rich media like video, audio, and images as well. In fact, per research by Verizon, mobile users are the most likely to stop watching if video quality is poor, beating out tablet, desktop, and Internet-connected TV users.⁴

Live streaming expectations

For news sites, social networking sites, TV providers, and others, live streaming is increasing in importance for user engagement. In fact, in a study by Livestream and New York Magazine, 80% of surveyed users said they would rather watch a live stream from a brand than read a blog post.¹⁰ According to the same survey, breaking news is the most often viewed type of live video content for 56% of users, followed by concerts and conferences at 43%.

Expectations are high for live streaming content: 67% of viewers say quality is the most important factor when watching a live stream.

Conversion rates

Embedded video can slow down page load times, frustrating users.¹¹ Pages that load faster convert more users, and apps that respond faster convert more users. Studies have shown that conversion rates go down 7% due to just one additional second of load time.¹² Meanwhile, for Pinterest, a 40% reduction in perceived load times increased sign-ups by 15%.¹³

Search engine optimization and organic traffic

Search engine optimization, or SEO, is the practice of making one's Internet property more visible by improving search results placement, leading to more user clicks and more organic traffic. Site speed is a crucial part of optimizing a website for search. Google has included site speed as a factor in their search rankings since at least 2010.¹⁴ Because of the growing adoption worldwide of mobile devices, Google also began using mobile performance as a ranking factor in 2018.¹⁵

For websites that rely on organic traffic and high placement in search results (such as news websites), following good SEO practices is critical. This entails ensuring that all content loads as quickly as possible.

PART 2

What factors affect performance?



Rich media files are much heavier than the simple webpages of yesteryear. Meanwhile, users are accessing the Internet on a wider variety of devices than ever before, adding further challenges for performance.

For OTT content providers and media and entertainment platforms, maintaining fast performance has become more of a challenge than ever. Let's take an in-depth look at the trends impacting performance today.

The challenges of streaming

Video files are much larger than almost every other file type available over the Internet. Because of this, video is streamed, not downloaded. In streaming, video files reach users a little bit at a time, and apps or browsers on the client side play the video files as they are received. This means that the file is continually sent to the user as they watch it, and an interruption in the stream causes the video to stop playing.

Streaming is subject to the same kinds of delays and performance degradations as other kinds of web content. Because the streamed content is not stored on the client device, hosting location makes a big difference for network latency (see below), as is the case with any type of content accessed over the Internet. If a user in New York is trying to stream from a server in California, the video content will have to cross 3,000 miles in order to reach the user, and the video will have to spend a long time buffering or may not even play at all.

Dropped packets

All data that crosses from one point to another on the Internet is broken up into packets. These data packets travel based on transport protocols.

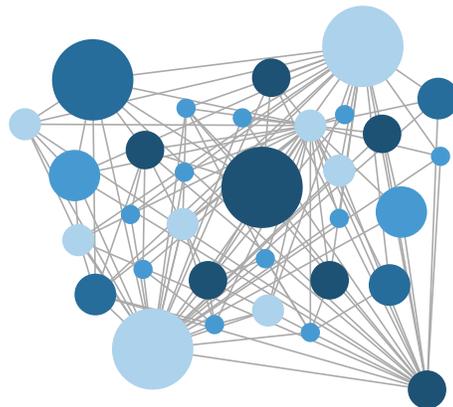
Streaming video and audio files traverse the web using a transport protocol called UDP. In contrast with the TCP protocol, which is used for most content on the Internet, UDP does not establish a connection before sending data packets, nor does it verify that all packets arrive. This makes streaming much more efficient, but if too many packets are dropped in transit, users will experience lower stream quality.

Increased consumer expectations and the usage of mobile

Today's consumers are more connected than ever, and this increases the demand on the backend infrastructure supporting online media. Users connect to apps and websites from around the globe, on all kinds of devices. A 2016 Nielsen survey found that 57% of respondents who made an online purchase in the past six months bought from an overseas retailer.¹⁶

Mobile is the new benchmark for web performance. However, building for mobile presents a new set of challenges. Mobile performance is constrained by network connectivity and availability. Despite the widespread availability of 4G and 5G networks in some countries, 60% of mobile connections worldwide are over 2G.¹⁷ And in some regions, mobile network providers will throttle bandwidth past a certain amount.¹⁸

Despite these challenges, mobile users have high performance standards for their apps: one study by Dimensional Research found that 49% of users expect apps to respond in 2 seconds or less, 55% hold the app responsible for performance issues, and 80% indicated they will only attempt to use a problematic app three times or less.¹⁹



Network issues

Users access Internet properties from all kinds of networks, and network conditions play a huge role in how well your website or app performs.

Bandwidth measures the maximum amount of data that can pass through any given point on a network at once. A user's network may not always provide enough bandwidth for a quality web browsing experience. Limited bandwidth can especially impact the video streaming experience.

Network latency is partially caused by distance. The farther a user is physically from an origin server, the more latency there will be. The speed of light is a hard limit on how fast data can travel, and data will take from a few milliseconds up to nearly a second to travel from the user to the server and back.

The effects of network latency can be somewhat reduced by using a CDN – content delivery network – to cache content closer to users. Major OTT providers such as Netflix make extensive use of CDNs.²⁰

Network congestion occurs when network traffic exceeds bandwidth at a certain point on the network, whether that's within an Internet Exchange Point (IXP), in a data center, or on a LAN router in a home. The resulting network congestion leads to slower Internet speeds for anyone connected to the network. Network congestion can be limited to a certain geographic area that lacks sufficient infrastructure, or it can affect an entire ISP's network.

If multiple people are trying to stream high-definition video at the same time, within the same house for instance, this can cause network congestion. The same principle applies if more users are streaming video within a given region than the infrastructure is equipped to handle. This is another reason why OTT providers make extensive use of CDNs – to limit network chokepoints and network congestion by pushing content as close to users as possible.

Mobile networks are often unreliable, although consumers increasingly rely on them for Internet access. The quality of service over a mobile network depends upon the user's location, the amount of bandwidth offered by their cellular provider, and many other factors. While mobile networks are improving around the world, some parts of the world still struggle with cellular reliability and connectivity.¹⁷

PART 3

What steps can you take to assess and improve performance?

Run a site speed test

How a site or an app performs in a local testing environment is not a good indication of how it will perform for users across a variety of network conditions.

Website speed tests aim to simulate real-world conditions and provide data on how well a website actually performs. The best website speed tests should let you know not just how fast their site or application is, but also what aspects of it are slowing performance.

Speed tests can provide a variety of metrics, including:

- Load time - how long it takes for a web browser to finish downloading and displaying the webpage
- Time to First Byte (TTFB) - how long it takes for the browser to receive the first byte of data from the web server
- Requests - the number of HTTP requests that a browser makes in order to fully load the page

For more speed test performance metrics, see **Appendix: Performance Metrics to Know**.

WebPageTest.org is a well-respected and thorough free testing platform. Google PageSpeed Insights can also help you assess your website.

Additionally, Cloudflare offers a simple testing tool for assessing load time, TTFB, and total requests.

Assess origin server health and load

Monitor server health

Server performance can degrade for a variety of reasons: Server hardware can break down, or server software may become outdated, for instance. The average lifespan of a server is about 5 years.²¹ Servers should be monitored constantly to ensure their health and availability.

Check server load

If origin servers are overloaded, they will perform slowly. Check the memory utilization of your servers.

Are some machines working harder than others? Are some servers using all of their compute power while others aren't? To get the most performance out of your servers and utilize server resources efficiently, it's important to balance workloads across multiple servers.

Offload some content requests via caching

If every user request has to be fulfilled by the origin servers, they may become overloaded. By implementing caching – in the browser, at the network edge (using a CDN), or both – many, if not most, round trips all the way to the origin server can be eliminated.

Identify where site traffic comes from

One of the principal causes of network latency is distance, so the location of users matters a great deal. For instance, if a platform's video content is hosted in New York, U.S., but most of its users are in Sydney, Australia, then most of the users will likely be unable to play the video content at all. The video data will take too long to load.

[Google Analytics](#) is a helpful tool for determining where users are coming from geographically. Once you've identified where site traffic comes from, you can figure out whether or not your web property's infrastructure is set up to serve those locations effectively.

Audit and optimize site images

The user's browser needs to download images before they can be displayed. The larger an image is (in terms of file size, not dimensions), the longer it takes to download. Large images often add to the load time of a webpage unnecessarily, as many devices don't have good enough screen resolution or a large enough screen to make very high resolution images necessary.

Before images can be optimized, you should determine how many images your website has and where they are located by conducting an image audit. Following the audit, as many images as possible should be optimized – meaning, compressed, resized, and converted to a lossy file format such as JPEG. Optimized images will load much more quickly.

Moz.com has step-by-step instructions for crawling all images on your website, identifying which ones need to be optimized, and optimizing them. Screaming Frog's SEO website crawler is helpful for auditing website images.

There are many free image optimizer tools available online. Adobe Photoshop can also compress images and export images in a variety of formats.

Cloudflare Image Resizing, Mirage, and Polish are the best options for companies that already deploy the Cloudflare CDN in order to cache images for faster delivery. Cloudflare Polish can be activated in the Speed tab of the Cloudflare Dashboard.

Check the current performance of your DNS provider

Identify your DNS provider, and then find out if your provider is giving you the best performance possible.

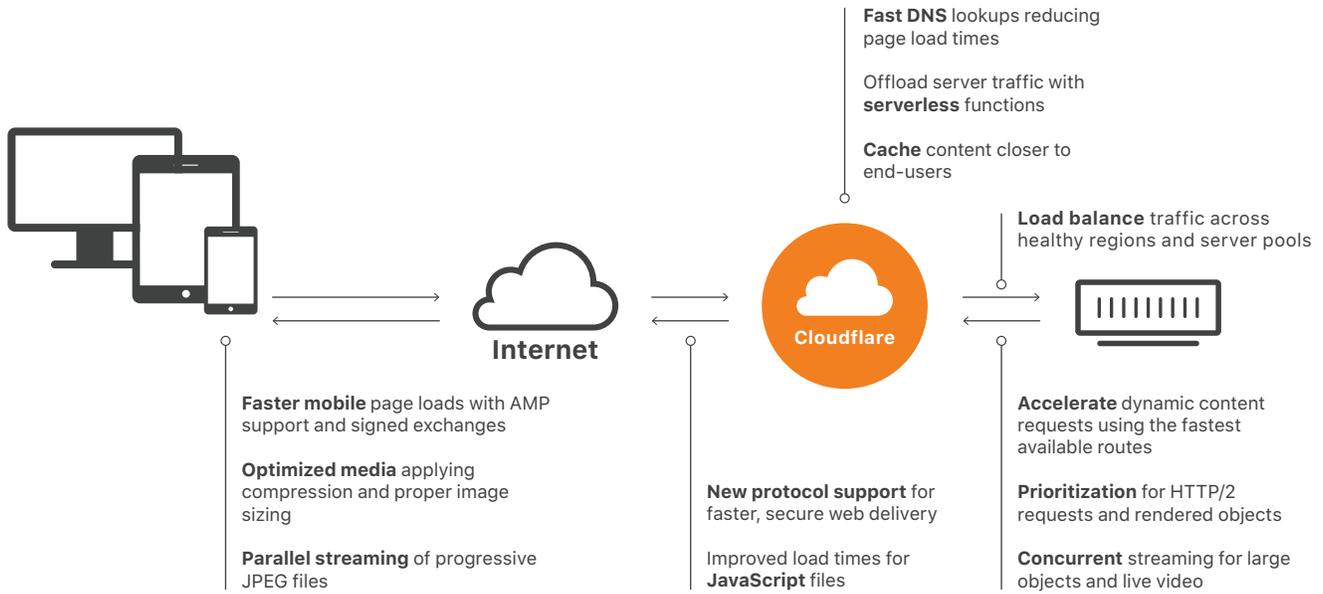
One of the best resources for measuring DNS performance is DNSPerf. DNSPerf regularly tests all authoritative DNS providers and public DNS resolvers. Their results and rankings are available for free at dnsperf.com.

Use a CDN

CDNs are essential for any Internet property that has large media files. A CDN is a network of caching servers that cache and serve content so that requests and responses don't have to travel all the way to and from a web property's origin servers. CDNs move content closer to end users, vastly cutting down on latency. Additionally, by caching content in multiple locations, no one server or network gets overwhelmed with requests for content. Leading OTT content providers such as Netflix²⁰ and Hulu²² make extensive use of CDNs.

PART 4

How Cloudflare solves performance issues for media and entertainment platforms



Faster live and on-demand video

Cloudflare Stream

Cloudflare Stream is an on-demand video streaming platform that combines video storage, encoding, distribution, and a customizable embedded player with per-minute pricing. Videos uploaded to Cloudflare are automatically distributed to the Cloudflare global network, resulting in minimal latency for users. Cloudflare also makes its distribution network available as a standalone offering with Stream Delivery, which caches and delivers video content in conjunction with your existing video encoding and storage infrastructure. For live streaming, Concurrent Streaming Acceleration reduces the latency experienced by simultaneous viewers for a better “spoiler free” experience.

Faster page load times

The Cloudflare CDN spans a global network of data centers that cache static web content closer to users so requests don't need to travel long distances to origin servers, resulting in faster page load times. Cloudflare Argo Smart Routing speeds up load times by delivering dynamic web content over the fastest links available.

Cloudflare offers a number of image optimization features, including Image Resizing, Polish, and Mirage. Image Resizing allows customers to optimize images by resizing, cropping, compressing, or converting them to WebP, a newer image format designed for fast loading. Cloudflare also enables parallel streaming of progressive images to speed up delivery of multiple images on a page.

Better reliability

Cloudflare Load Balancing provides local and global load balancing to reduce latency, either by load balancing traffic across multiple servers or by routing traffic to the closest region. It also includes health checks with fast failover to rapidly route visitors away from failures.

Cloudflare DNS

Cloudflare's DNS is the world's fastest and most reliable authoritative DNS provider.²¹ Cloudflare provides fast and reliable managed DNS as a built-in service on its network. And Cloudflare DDoS protection keeps web properties up and running even in the face of massive DDoS attacks.

Better experience on mobile

Cloudflare offers several products and features specifically for speeding up the user experience on mobile devices. Google AMP allows webpages to load much faster on mobile devices, and Cloudflare supports the use of Signed Exchanges with Google AMP, providing native URL attribution when viewed in the AMP viewer. Rocket Loader from Cloudflare optimizes the prioritization of any assets that need to load before on-page JavaScript can execute, also speeding up page load times.

Serverless computing for more responsive applications

Serverless computing has great potential for creating faster, more responsive apps than ever. Cloudflare Workers allows developers to build serverless applications that run on Cloudflare's network, closer to your users. Applications built with Cloudflare Workers are always available, with low-latency responsiveness.

Conclusion

Today, we rely on the Internet to do something it wasn't designed to do: deliver video, audio, and rich media content, on demand, at a massive scale. However, with technological enhancements and the right product suite, media and entertainment companies can offer high-quality experiences to their users, no matter where they are.

About Cloudflare

Cloudflare, Inc. (www.cloudflare.com / @cloudflare) is on a mission to help build a better Internet. Today the company runs one of the world's largest networks, with nearly 10 percent of the Fortune 1,000 companies and approximately 19 percent of the top 10,000 websites using at least one Cloudflare product. Cloudflare's platform protects and accelerates any Internet application online without adding hardware, installing software, or changing a line of code. Internet properties powered by Cloudflare have all web traffic routed through its intelligent global network, which gets smarter with every request. As a result, they see significant improvement in performance and a decrease in spam and other attacks. Cloudflare was named to Entrepreneur Magazine's Top Company Cultures 2018 list and ranked among the World's Most Innovative Companies by Fast Company in 2019. Headquartered in San Francisco, CA, Cloudflare has offices in Austin, TX, Champaign, IL, New York, NY, San Jose, CA, Washington, D.C., London, Munich, Beijing, Singapore, and Sydney.

Endnotes

1. "The State of Content: Expectations on the Rise." Adobe, <https://blogs.adobe.com/creative/files/2015/12/Adobe-State-of-Content-Report.pdf>. Accessed 24 July 2019.
2. Dahl, Jon. "How do users feel about video streaming quality on their TVs?" Mux, <https://mux.com/blog/how-do-users-feel-about-video-streaming-quality-on-their-tvs/>. Accessed 9 August 2019. h/t <https://www.emarketer.com/content/half-of-u-s-users-stop-viewing-content-slow-to-load>
3. Clark, Matthew. "How the BBC builds websites that scale." CreativeBloq, <https://www.creativebloq.com/features/how-the-bbc-builds-websites-that-scale>. Accessed 22 July 2019.
4. "Quality Matters: Video Streaming Quality Report 2016." Verizon, <https://www.verizondigitalmedia.com/report/quality-matters/>. Accessed 9 August 2019. Gated. Secondary source: "'It's Still Loading?!' Long Load Time Just One Reason People Quit Watching Video." eMarketer, <https://www.emarketer.com/Article/Its-Still-Loading-Long-Load-Time-Just-One-Reason-People-Quit-Watching-Video/1015717>. Accessed 9 August 2019.
5. Krishnan, S. Shunmuga and Ramesh K. Sitaraman. "Video Stream Quality Impacts Viewer Behavior: Inferring Causality Using Quasi-Experimental Designs." IEEE, https://people.cs.umass.edu/~ramesh/Site/HOME_files/imc208-krishnan.pdf. Accessed 9 August 2019.
6. "Mobile and tablet internet usage exceeds desktop for first time worldwide." StatCounter, <http://gs.statcounter.com/press/mobile-and-tablet-internet-usage-exceeds-desktop-for-first-time-worldwide>. Accessed 22 July 2019.
7. Meeker, Mary. "Internet Trends 2017 - Code Conference." Kleiner Perkins, <https://www.kleinerperkins.com/perspectives/internet-trends-report-2017/>. Accessed 24 July 2019. Presentation.
8. "Online mobile transaction statistics." Think with Google, <https://www.thinkwithgoogle.com/data/online-mobile-transaction-statistics/>. Accessed 22 July 2019.
9. An, Daniel. "Find out how you stack up to new industry benchmarks for mobile page speed." Think with Google, <https://www.thinkwithgoogle.com/marketing-resources/data-measurement/mobile-page-speed-new-industry-benchmarks/>. Accessed 24 July 2019. h/t <https://www.marketingdive.com/news/google-53-of-mobile-users-abandon-sites-that-take-over-3-seconds-to-load/426070/>
10. Golum, Caroline. "Live Video Statistics: What Consumers Want [Infographic]." Livestream, <https://livestream.com/blog/live-video-statistics-livestream>. Accessed 9 August 2019.
11. "How to Fix Video That's Slowing Down Your Page Load Time." Jonesen, <https://jonesen.com/blog/how-to-fix-video-thats-slowing-down-your-page-load-time/>. Accessed 9 August 2019.
12. Rodman, Tedd. "Marketing & Web Performance: How Site Speed Impacts Metrics" Yottaa, <https://www.yottaa.com/marketing-web-performance-101-how-site-speed-impacts-your-metrics/>. Accessed 24 July 2019.

- 13.** Meder, Sam et al. "Driving user growth with performance improvements." Pinterest Engineering (Medium), https://medium.com/@Pinterest_Engineering/driving-user-growth-with-performance-improvements-cfc50dafadd7. Accessed 22 July 2019. h/t <https://developers.google.com/web/fundamentals/performance/why-performance-matters/>
- 14.** Singhal, Amit, and Matt Cutts. "Using site speed in web search ranking." Google Webmaster Central Blog, <https://webmasters.googleblog.com/2010/04/using-site-speed-in-web-search-ranking.html>. Accessed 22 July 2019.
- 15.** Wang, Zhiheng, and Doantam Phan. "Using page speed in mobile search ranking." Google Webmaster Central Blog, <https://webmasters.googleblog.com/2018/01/using-page-speed-in-mobile-search.html>. Accessed 22 July 2019.
- 16.** "Global Connected Commerce: Is e-tail therapy the next retail therapy?" Nielsen, <https://www.nielsen.com/bd/en/insights/report/2016/global-connected-commerce/>. Accessed 23 July 2019.
- 17.** Schwarz, Ben. "Beyond the Bubble: Real world performance." Calibre (Medium), <https://building.calibreapp.com/beyond-the-bubble-real-world-performance-9c991dcd5342>. Accessed 24 July 2019.
- 18.** O'Donoghue, Ruadhán. "You've been throttled, but don't stop browsing!" mobiForge, <https://mobiforge.com/news-comment/youve-been-throttled-dont-stop-browsing>. Accessed 24 July 2019.
- 19.** "Failing to Meet Mobile App User Expectations: A Mobile App User Study." Dimensional Research, https://techbeacon.com/sites/default/files/gated_asset/mobile-app-user-survey-failing-meet-user-expectations.pdf. Accessed 24 July 2019. h/t <http://thinkapps.com/blog/post-launch/mobile-app-performance-tips/>
- 20.** Macauley, Tom. "Ten years on: How Netflix completed a historic cloud migration with AWS." Computerworld, <https://www.computerworld.com/article/3427839/ten-years-on--how-netflix-completed-a-historic-cloud-migration-with-aws.html>. Accessed 9 August 2019.
- 21.** "Server FAIL: 3 signs your server is on the brink." Spiceworks, <https://www.spiceworks.com/it-articles/3-signs-server-about-to-fail/>. Accessed 26 July 2019.
- 22.** Adhikari, Vijay Kumar et al. "A Tale of Three CDNs: An Active Measurement Study of Hulu and Its CDNs." IEEE, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.387.4423&rep=rep1&type=pdf>. Accessed 9 August 2019.



1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

© 2019 Cloudflare Inc. All rights reserved.
The Cloudflare logo is a trademark of Cloudflare. All other company and product names may be trademarks of the respective companies with which they are associated.

REV: 190730