

Most websites and applications that we interact with every day are run out of one physical location, but content on the site (like images, text, and video) still needs to travel over wires to the entire world.

For example, if a website's servers are based in New York City, people in Boston will get the content faster than people in San Francisco or Tokyo. The farther away customers are from a company's data center, the slower the website or application loads — creating an inconsistent and frustrating user experience.

Content delivery networks can help solve this problem.

What is a CDN?

A CDN, or content delivery network, delivers content from websites and mobile applications to people more quickly and efficiently, based on their geographic location. A CDN is made up of a network of servers (called "points of presence," or POPs) in locations all over the world. The CDN server closest to a web user is known as the "edge server;" when people request content from a website served through a CDN, they're connected to the closest edge server, ensuring the best internet experience possible.

Websites cache (temporarily store) content on CDNs so that it is delivered from the edge to someone much faster than if it had to be delivered all the way from the origin. When someone is trying to access content from a website or mobile app that has hosted their content on a CDN, then that person's request for content only needs to travel to a nearby POP and back, not all the way to the website origin and back. CDNs also purge (remove and update) content constantly, so that the most up-to-date content is delivered.



If a business isn't using a CDN, their content always has to go back to their origin servers.



With a CDN, content is delivered faster to internet users.



What a CDN can do for businesses

Anybody who has a website or mobile application that's likely to be requested by more than one user at a time can benefit from a CDN, but CDNs are especially useful to large, complex sites with a lot of dynamic content and users spread across the globe.

CDNs offer many specific benefits to different types of businesses and organizations, such as:

Ecommerce: A CDN helps ecommerce sites improve conversion rates and deliver content quickly and efficiently even during times of heavy traffic, like Black Friday and the holidays.

Media / Publishing: Media websites need to deliver timely and up-to-date information, and a CDN can help media companies update headlines and news homepages as stories unfold in real time, and remove data as it becomes outdated.

Finance: CDNs provide banking institutions with a fast, secure, and reliable infrastructure to deliver sensitive data to consumers and analysts.


Government: Large, content-heavy websites can deliver vital information to citizens much more quickly and efficiently by using a CDN.

Mobile apps: A CDN delivers dynamic location-based content for mobile apps, reducing load times and increasing responsiveness across devices.


Technology and SaaS: A CDN helps technology websites serve billions of requests a day to web users without decreasing performance.

Benefits of Using a CDN

Better site performance




Faster load times for web users




Minimizes risk of traffic spikes at point of origin



Decreases infrastructure costs due to traffic offloading (less load on origin)



Quickly scalable during times of heavy traffic



A more secure global network





Glossary of CDN terms

Cache: Websites temporarily store (cache) content on CDNs so that content is delivered faster and more efficiently.

Purge: A purge discards the data that is already cached so new data can be cached in its place. Items that change at the origin need to be removed and updated (purged) from the cache so the freshest content available is delivered by the CDN. Traditional CDNs purge their content every 7 to 15 minutes. Modern CDNs can purge content every 150 milliseconds.

POP: A CDN's points of presence (POPs) on the internet are servers placed at the center of high-density Internet Exchange Points around the world. When someone accesses website content hosted on a CDN, the request for content goes to a POP nearest to that person, rather than all the way back to the origin. For instance, if someone in Tokyo accesses a website hosted on servers in San Francisco, the request for data only needs to go as far as the Tokyo POP.

The edge: The edge is the CDN POP geographically closest to a web user, where the website data is hosted, which sends the content requested to that user.

Latency: Latency is a measure of time delay experienced in a system, the precise definition of which depends on the system and the time being measured. Latency affects the amount of time a person spends waiting for content to be delivered when they visit a website. Ideally, this time delay should be as short as possible, measured in milliseconds. By caching content on CDNs at the edge, websites can minimize latency for their visitors.

Static content: Traditional CDNs can only cache static content. Static content doesn't change, or only changes rarely (ex. images, simple text on a webpage).

Dynamic content: Dynamic content changes often (ex. HTML, AJAX requests, API calls, user-generated content, etc.), and comprises both cacheable and uncacheable content.

Origin: The server where a website's or application's authoritative data is hosted. CDNs pull data from the origin when it is not already in a POP and serve it to users.

Solid-state drive (SSDs): SSDs process requests several times faster than traditional hard drives, so an all-SSD content delivery network can serve content to people even faster.

Configuration: The customer-specific set of rules and logic that manipulate the CDN's caching system.