

Make websites  
fly with super-fast  
page loads in  
all conditions

Brochure

# Varnish Web and API Acceleration

# Varnish Web and API Acceleration provides the speed and stability at scale that websites need

## Varnish Web and API Acceleration

All websites, whether they cover news, information, media, e-commerce, or entertainment, have a common challenge: they need to satisfy every visitor with a fast, responsive experience at all times. The internet is a highly competitive place, and users demand immediate access to content. There is little patience for long page load times, with a 90% increase in bounce rate when webpage load time goes from 1 to 5 seconds.<sup>1</sup> Visitors could be anywhere in the world, on any device, but they all expect latency-free experiences. This can include internal web services too; all businesses rely on maintaining fast access to internal systems for communications, finance and data management.

The irony is that popularity results in high traffic and can actually affect web performance and cause downtime, so content providers need tools to handle high demand, scaling to match audience sizes while minimizing latency. When a website is revenue-critical, maintaining web performance is a must, even when serving dynamic or personalized content and pushing real-time updates to time-sensitive web services.

## Speed up page loads for everyone

Varnish Web and API Acceleration provides the speed and stability at scale that websites need to deliver excellent web experiences for their users while protecting backend servers and minimizing operating costs. As a caching layer that sits between origin servers and website visitors, it enables up to 99% backend offload while decreasing page load times significantly, keeping visitors onsite and engaged while reducing pressure on the backend.

It scales easily to match any demand and helps websites stay online and responsive, including when delivering personalized or dynamic content. The comprehensive feature set includes multi-terabyte disk storage and content replication and its flexible, architecture-agnostic design means it syncs easily with the full range of CMSes, e-commerce platforms and CDNs.



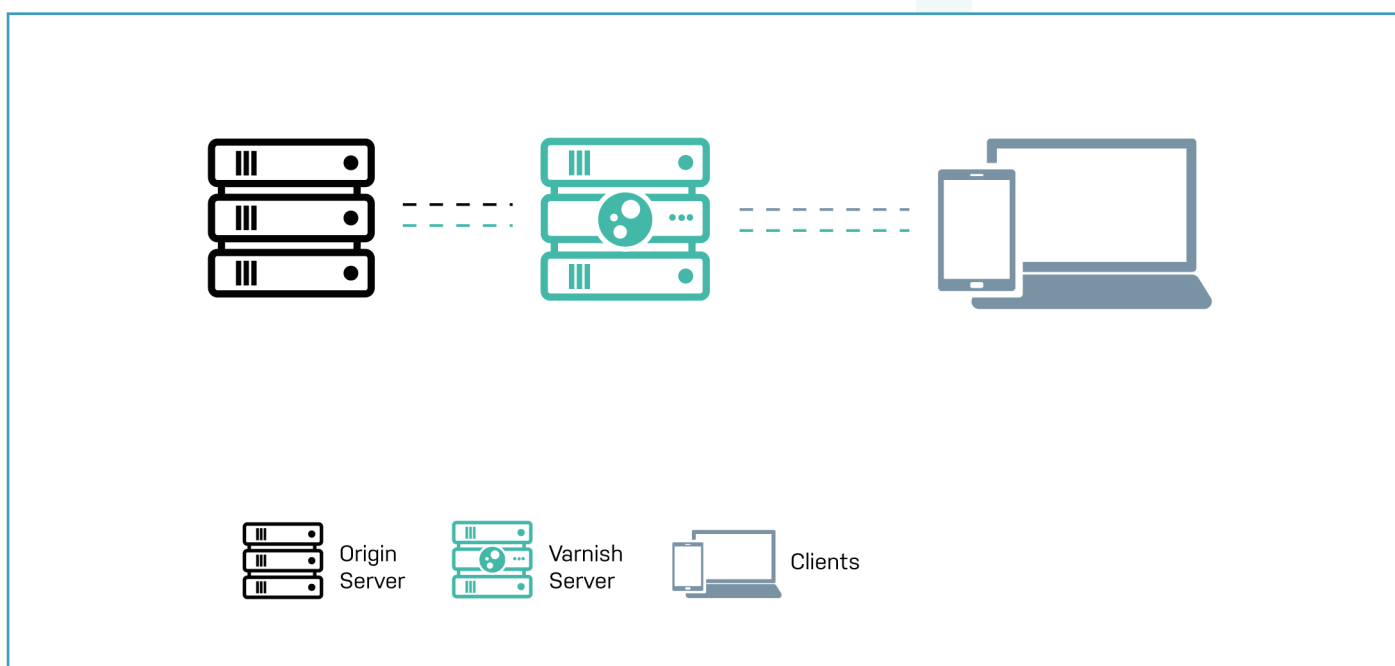
<sup>1</sup> <https://www.thinkwithgoogle.com/marketing-resources/data-measurement/mobile-page-speed-new-industry-benchmarks/>

# How Web and API Acceleration works

Web and API Acceleration functions as a caching layer that sits between a content provider's web servers and their users, helping to deliver anything from breaking news, real-time online shopping updates, and data for internal web services.

It takes a copy of the content passing through it the first time a user requests something from the website or web service. Future requests for the same content can be served by Web and API Acceleration instead of being sent all the way from the backend servers, which is how it reduces server load and decreases page load times significantly, ensuring high-performance HTTP and API content delivery.

Putting Web and API Acceleration in front of backend servers protects the infrastructure from flooding and lets the cache do the heavy lifting, helping web services stay online and satisfyingly responsive even when delivering personalized and complex, dynamic content to large audiences. Web and API Acceleration is incredibly flexible so it also excels as a load balancer, SSL proxy, WAF, and HTTP gateway, for streamlining operations while increasing capacity, performance and resilience.



Web and API Acceleration is flexible and hardware agnostic, ready to be deployed on existing infrastructure, alongside existing tools, on-premise, in-cloud or as a hybrid set up.

# What can Web and API Acceleration do?


Web and API Acceleration can do many things, but some of the reasons our customers rely on it are so they can:

- **Accelerate web content delivery** with a highly configurable reverse proxy that greatly reduces latency.
- **Protect critical infrastructure** by offloading requests for content, handling more users with less bandwidth.
- **Support huge, unpredictable demand** by scaling rapidly to deliver content at speed while protecting origin servers.
- **Manage costs of high-traffic content delivery** by provisioning alongside existing CDN architecture to prevent cost escalation.
- **Simplify and future-proof web architecture** with a platform-agnostic solution that adds robustness and efficiency to web operations.

---

Scale easily to match any demand and help websites stay online and responsive, including when delivering personalized or dynamic content

---



# Web and API Acceleration features

## Performance

Deliver web content up to 1000x faster, with an average server response of 100µs and the potential to hit 150Gbps per server while supporting 100k+ requests a second.

## High Availability

No matter the traffic demand, high availability scales up to ensure capacity and redundancy to avoid costly downtime.

## Origin protection

Increase your cache-hit ratio and reduce your backend server requests by at least 80%, making sure your origin is shielded from heavy traffic and DDoS attacks.

## Security

Secure cache traffic with front and backend TLS, cache encryption and Varnish WAF, plus support for granular authentication and JWTs.

## Edge logic and personalization

VCL adds true flexibility and edge power, for setting granular policies and delivering personalized user experiences, including for mobile and dynamic content.

## Memory Governor

Scale the cache size automatically to make the most of memory capacity without risking saturation

## Edge storage and persistence

Storage multi-terabyte data sets at the edge, for fast content delivery and adding persistence so that even in an outage it can serve content.

## Cache invalidation

Advanced cache invalidation ensures that up-to-date content is ready for visitors while maintaining website speed, regardless of content changes.

## Real-time insights

Get instant access to real-time traffic, performance and caching statistics from everything that travels via HTTP for actionable analysis.

---

Putting Web and API Acceleration in front of backend servers protects the infrastructure from flooding and lets the cache do the heavy lifting

---

# Web and API Acceleration customer success

## CBC

The Canadian national broadcaster boosted their website performance by putting Varnish in front of Apache servers to accelerate object delivery up to 100x and reduce the number of servers required, leading to a 30% reduction in CAPEX and OPEX over five years.

In an outage situation, the whole site can be offline, and we still have a working website, thanks to Varnish technology. The CBC website could not do without it.

*- Massimo Mollica,  
Manager at CBC/Radio-Canada*

We immediately saw how every tweak translated to performance in real time. Nikon is really excited about how fast the site feels now.

*- Cindy Jeffus,  
Sr. Manager of Internet Operations,  
Nikon Inc.*

## Nikon

Web and API Acceleration helped Nikon keep its e-commerce web and mobile apps performing well, with a >50% increase in cache-hit ratio and a significant reduction in infrastructure costs.

## La Nación

La Nación gained increased performance and control, plus seamless integration with pre-existing CMS and backend systems and the ability to scale easily in high-traffic situations.

Varnish implementation was very simple. The platform is very versatile when making changes and adapting it to SSO, web services, etc., unlike in other applications where everything is more proprietary and closed.

*- José Luis Falvo,  
Manager of Infrastructure and Operations  
of Digital Media, La Nación*

In implementing the Varnish solution, we have come to rely on high speed and the functionality of Varnish High Availability - placing an object in cache and being able to deliver it seamlessly and quickly through two different Varnish instances. And finally, having timely granular control over caching and cache invalidation is crucial.

*- Cristian Tuduca,  
ETH IT Services*

## ETH Zurich

The major university deployed Varnish to replace Adobe Dispatcher and gain control over cache logic, unlock higher caching performance and scalability, and achieve a 99.99% SLA across 600+ internal websites.

---

# Web and API Acceleration excels as a load balancer, SSL proxy, WAF, and HTTP gateway, for streamlining operations while increasing capacity, performance and resilience.

---

## Support the way you need it

Many enterprises value Web and API Acceleration's 24x7x365 support provision, with access to short SLAs, technical expertise, long-term software lifecycles and managed updates:

- Day-to-day troubleshooting with a team of Varnish experts on hand to support general usage.
- Each release is stable, tested and hardened, to enable confident enterprise rollout.
- New releases and feature updates are communicated in advance, with clear upgrade pathways.
- Security updates are backported, for long-term protection.

## Professional services from the Varnish experts

Another benefit of Web and API Acceleration is access to professional services with a team of industry-leading caching and content delivery experts, for help planning, sizing and implementing the solution, and fine-tuning once installed. No one goes it alone, with dedicated teams on hand to support usage of Web and API Acceleration every step of the way, no matter the use case or challenge. Web and API Acceleration is flexible enough to grow and adapt in line with business requirements, for future-proofed technology that is expertly tuned for every situation.

## Experiencing Web and API Acceleration

Extended, free trials are available and easy to set up and run, for a no obligation, low-effort way to experience the performance benefits of Web and API Acceleration first hand.

## Learn more

Founded in 2010, Varnish Software has offices in Oslo, Stockholm, London, New York, Los Angeles and Paris. To find out more, please visit [www.varnish-software.com](http://www.varnish-software.com).



Los Angeles - Paris - London  
Stockholm - Singapore - Karlstad  
Dusseldorf - Oslo - Tokyo

