



# Varnish Streaming Server

Delivering reliable, high-performance streaming, particularly of live, over-the-top (OTT) and video on demand (VoD) media, is at the heart of the challenge companies face.

---

## HTTP streaming for live, OTT and VoD

As video starts to make up the majority of internet traffic (according to Cisco<sup>1</sup>), streaming challenges become more pressing and specific: achieving speed and consistency, reducing latency, ensuring availability and delivering great user experiences across devices and networks regardless of the kind of video (live or OTT/on-demand).

Whether you're a CDN or a content provider, such as a broadcast network, media outlet, a content aggregator, sports association, or other corporation or organization, Varnish Software gives you the performance and flexibility to meet these challenges.

---

<sup>1</sup> <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.html>

# // HTTP streaming challenge: Deliver any volume of traffic with highest QoE

## Streaming challenges

Before explaining what Varnish can do and how, let's take a look at the practical challenges of HTTP streaming that most companies face.

- Managing the unpredictable: It comes down to **scalability** and using software that scales to any traffic level demand.
- No one wants to wait: **Latency**, or slow content delivery, issues can be a major barrier. Here latency can be two things: time before the user can start the video or the delay between the video being available on the server and via Varnish (a lag in live video).
- Breaking down: A lack of **resilience**, consistency and uptime can lead to streaming problems or total breakdown.
- Adapting to your specific streaming needs: A lack of **flexibility** or adaptability in streaming software solutions can stand in the way of efficient content delivery, as you will want some control over how your streaming works.
- Going blind: "Black box" solutions that take flexibility and **control** out of your hands, remove transparency and possibly tack on extra costs with every feature added.
- Securing yourself and your traffic while also protecting your origin/backend.

## Varnish Streaming Server solutions

Varnish's HTTP-native design is made for solving streaming challenges. Varnish Streaming Server™ gives you solutions that can be used and deployed flexibly to meet your individual streaming needs.

You can implement Varnish Streaming Server to manage the most difficult and persistent streaming challenges.

### 1. LIVE, OTT AND VOD STREAMING VIDEO DISTRIBUTION

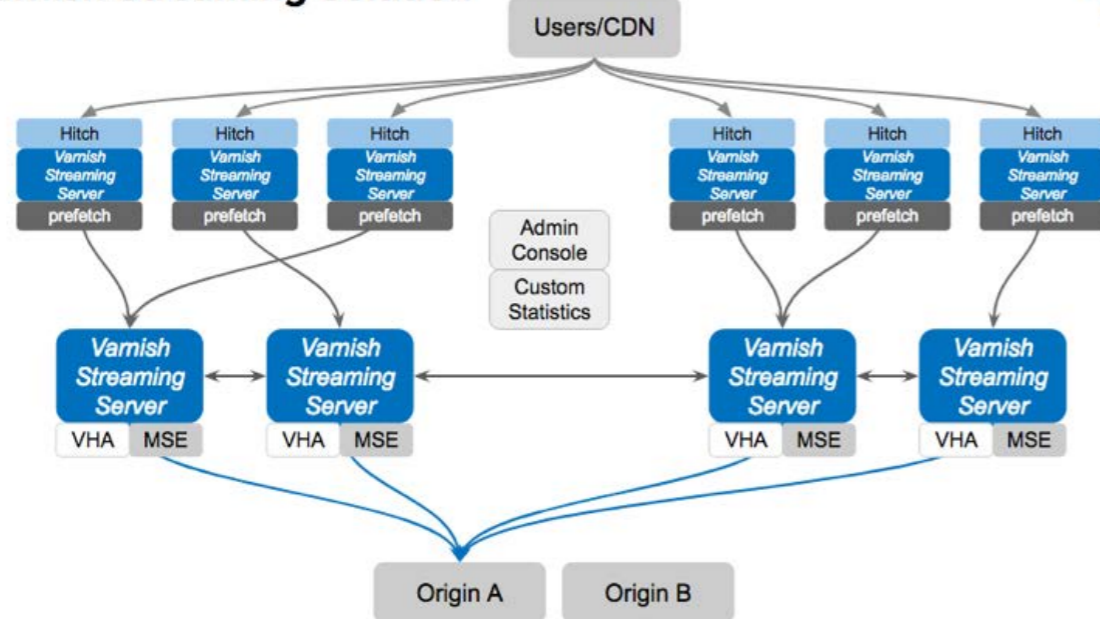
Stream and serve video and efficiently scale out your platform as well as store and serve vast libraries of media content (for high-volume VoD use). Caching live streaming is easier (you still must take care not to kill your origins, of course), thanks to the short lifetime of content, but caching an entire catalogue of VoD, your storage solution needs to scale to tens and hundreds of terabytes, allowing you to cache everything you want, and not just the hottest content. Varnish Massive Storage Engine™ offers:

- Built for up to 100+ terabytes of storage on each node
- Fragmentation-proof allocation algorithm
- Higher cache hit rates thanks to a better eviction policy
- Persistent datastore so you don't lose your entire cache on restart

#### Global live streaming scenario

"Continental" or large-scale video distribution can be deployed using Varnish DIY CDN as multi PoPs (points of presence) distributed across the globe. Each CDN PoP can be located where needed (e.g., UK, US, France, Brazil, Japan, Norway...) to most efficiently deliver video content (live, OTT, VoD) to be as close as possible to the user's video consumption. The CDN PoPs may be deployed using Varnish Cloud instances in those different regions, or on-premise, e.g., reusing your own existing network infrastructure, which is already deployed.

### Varnish Streaming Solution



You only need the Varnish blocks - all the rest are lightweight, modular options that you can add based on your needs, such as Hitch for TLS/SSL termination, VHA high availability modules and MSE for massive storage needs

## // Latency/performance lags are still end-users' number one complaint about streaming media. Prefetching content, keeping your cache warm, is one potential way to boost streaming performance and make the streaming experience smoother

### 2. PREFETCHING: CACHE WARMING

#### Boost performance by anticipating the future: Prefetch to keep your cache warm

Latency/performance lags are still end-users' number one complaint about streaming media. Prefetching content, keeping your cache warm, is one potential way to boost streaming performance and make the streaming experience smoother. Prefetching acts predictively, anticipating what the next logical chunks of content a client will request, letting you prefetch content.

What exactly does this mean? You are loading data into your edge server's cache before it is even requested - making it available and ready to serve immediately when the client requests it. Prefetching content increases the number of cache hits, making it highly useful for video on demand (VoD) and OTT streaming scenarios, pre-warming caches with the desired amount of future chunks as well as providing best Quality of Experience (QoE) for the users.

While the VMOD-http functionality extends beyond prefetching, it is important for streaming. The VMOD lets you execute HTTP requests directly from VCL and supports synchronous and asynchronous operations and connection pooling for higher efficiency. It boils down to the fact that with just six lines of VCL you can warm your cache with the next video segment and keep it ready to continue delivering content.



Prefetch in action: HIT ratios generated by the prefetch feature (left) and video encoded and segmented using MPEG-DASH (right)

### 3. SECURITY & BACKEND PROTECTION

Varnish Streaming Server, while also being about consistent, high-performance content delivery, is underpinned by the concept of protecting your ability to deliver content. This means that a variety of Varnish features ensure and support the smooth storing and serving of content under any conditions.

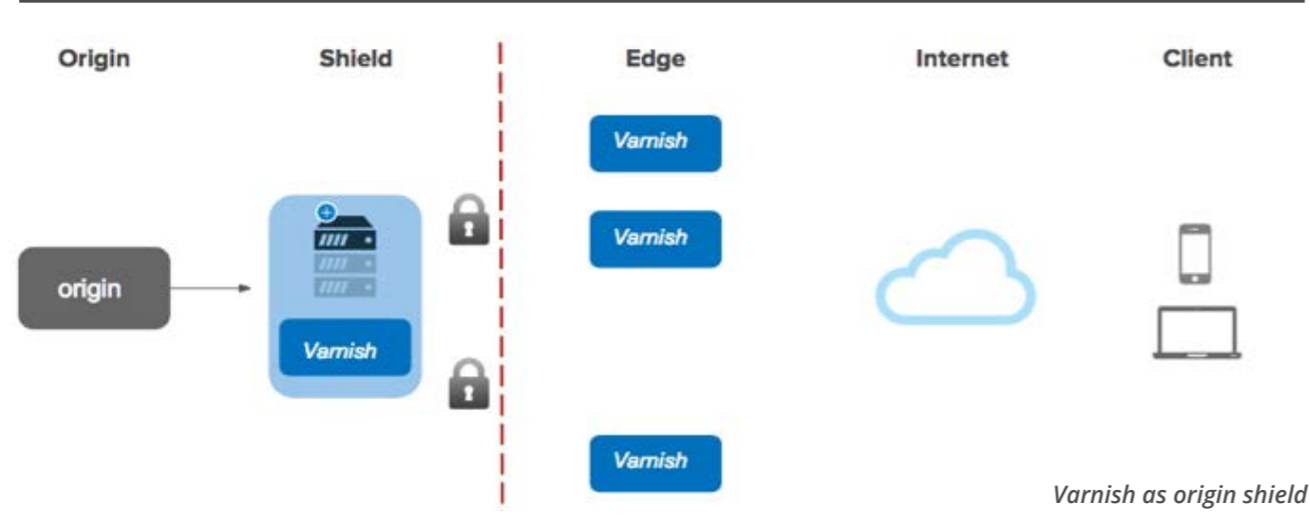
#### 1. Store and serve massive amounts of content

One continuous challenge of VoD streaming is the issue of growing amounts of content - vast libraries and archives of media that must be stored and easily, quickly accessed when requested. Varnish Massive Storage Engine (MSE) provides a significant amount of local cache storage. Caching live streaming is easy, thanks to the short lifetime of content, but caching an entire catalogue of VoD, your storage solution needs to scale to tens and hundreds of terabytes, allowing you to cache everything you want, and not just the hottest content.

#### 2. Origin shield

Protect your backend and ultimately the most valuable asset of all: your content. Your Varnish layer will exist as close as possible to the origin servers and this layer will protect backends and serve as much content as possible to other Varnish nodes. Shield your origin from excess traffic hitting the backend, and achieve solid streaming performance and a reasonable cache-hit rate, even with live streams.

Normally, networks of CDNs turn to the origin to fetch content, and the pressure of an influx of unexpectedly high traffic or an overload drives origin servers into the ground. Varnish acts as a content replication engine at the same time as serving as a protective layer against these traffic floods, acting as a tier of caches - horizontal scaling in front of the traffic. Pressure on the origin is reduced and the reliability and resilience of your streaming service is secured.



#### 3. Control logic and policy: Varnish Configuration Language (VCL)

Flexibility is also key to protecting your content and its delivery, whether it's ensuring that only authorized audiences can access the content or suppressing abuse with rate limiting. In Varnish, this - and much more - is accomplished with Varnish Configuration Language™, which gives you the ability to configure and control the logic governing your content. Because of this flexibility, you can, for example, do things like:

##### • Policy and logic engine

Use VCL as a complex policy and logic engine that enables things like geographically restricted content (geo-blocking), rate limiting or authorization and authentication and more.

### // Everyone is cutting the cord in one way or another.

**With these considerations in mind, companies across industries can benefit from using Varnish as their solution for OTT VoD and live streaming.**

##### • Geo-blocking at country or city level

When content needs to be restricted by geography, Varnish Streaming Server includes a GeoIP VMOD for limiting or restricting access by specific location. While Varnish does not have this functionality built in, it is just one of the many flexible additional modules (VMODs) that can be added thanks to the flexibility of VCL.

##### • Flexible rate limiting and abuse suppression

Designed exactly for these use cases, the Throttle VMOD will act as a guardian, directly in VCL, keeping tabs on requests, letting you refuse too-frequent occurrences.

#### 4. Security

Varnish gives you the option to transmit video with a security layer. One of the challenges here is to transmit a high-performing video and still deliver it with SSL/TLS, as the security layer usually slows down video delivery by taking from your bandwidth, thereby reducing the number of incoming requests. Varnish allows you to deliver secure streaming without undermining its performance.

Most video providers use DRMs to protect files. Varnish Streaming Server is compatible with any DRM in use because Varnish is only a caching layer and does not interfere with encryption logic. We can add power to any video system (which is important as more products aim to offer everything but often neglect the delivery performance aspect of content delivery).

### 4. VARNISH STREAMING SERVER IN PRACTICE

We have helped our customers globally to build advanced, scalable and fast streaming solutions on their own terms through the whole lifecycle of the software: Design, feature development and enhancements, implementation and optimization. Varnish Streaming Server offers all the flexibility and performance to make the streaming experience high-performance, robust and efficient while giving end-users what they want and expect.

**Eurosport:** Eurosport's normal live streaming volume hovers around 20,000 request/second but peaks at 50,000 requests/second during big events like the French Open. They use Varnish Streaming Server to ensure high-performance video streaming and to offload traffic from backend servers.

**SFR:** French telecoms giant needed massive storage to serve its rapidly growing VoD library of content. They also use Varnish to develop their own CDN to save costs and bring in a new source of revenue by offering CDN services onward to their own customers.

**ComHem:** Swedish ISP Com Hem use Varnish to power seamless OTT live and VoD video streaming, manage growing amounts of data and as the foundation of their own private, DIY CDN.

**RTÉ:** Ireland's national radio and television broadcaster uses Varnish Streaming Server in a fairly classic way, employing the Massive Storage Engine and Varnish High Availability, along with the flexibility of VCL.

