Flawless Video Streaming to Any Device on Any Network

Intelligent Video Caching for the Enterprise

Demand for enterprise video is at an all-time high and continues to grow year over year. In fact, Wainhouse Research's 2019 Enterprise Streaming End User Survey found 83% of respondents reported some usage of streaming business video on at least a weekly basis, up from the 79% usage level reported in the previous year's survey.[1] Further amplified by the global pandemic, this increase is fueling unprecedented demand for live and on-demand video for everything from all-employee meetings and executive messages to immersive training using 360-degree video.

The more you use video, the more it taxes your network. Corporate internet connections and WAN links are usually limited in bandwidth. Streaming video across those connections consumes the throughput, impacting network performance and putting your business-critical applications and operations at risk. You could add more bandwidth—which means more time, resources and money—or you could introduce an enterprise content delivery network (eCDN). An eCDN is a network overlay that optimizes the bandwidth required to stream video behind the firewall.

The more you use video, the more it taxes your network.

Enterprises have long used caching for network optimization, better response times, and improved user performance.

This same concept applied to streaming video achieves similar benefits.

One of the best ways to reduce the amount of video travelling across your corporate network is to move the video source inside those connections and closer to your viewers. The less distance the video has to travel, the less bandwidth it consumes.

Enterprises have long used web caching for network optimization, better response times, and improved user performance. This same concept is applied to streaming video to achieve similar benefits. Ramp OmniCache™ is an intelligent video caching solution that replaces expensive infrastructure upgrades with a lower cost, flexible solution. The lightweight software-only solution runs on your existing infrastructure and is designed specifically for live video and video on demand (VOD) content.

According to Wainhouse Research, more than three quarters (77%) of organizations indicate that they support two or more streaming platforms, with more than one third (37%) having three or more.[2] Because the requirements for streaming platforms vary from one functional area of the business to the other—corporate communications, learning and development, human resources—the need for multiple platforms is here to stay. In contrast, IT's requirements for video delivery are largely the same regardless of the streaming platform or use case. OmniCache is an enterprise software caching solution capable of supporting video from virtually any source. So, if one department is using Microsoft Stream, another is using Brightcove, and yet another is using Kaltura, OmniCache optimizes video delivery for all simultaneously.

Number of Streaming Platforms in Use

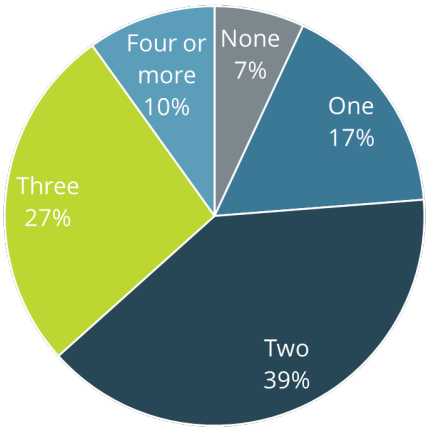

Wainhouse Research, Q3 2019

OmniCache lets you standardize your video network on a common enterprise streaming infrastructure (CESI). In addition to supporting virtually any video source and player technology, it easily expands over time without requiring costly and labor-intensive network upgrades and meets or exceeds enterprise security requirements. Choosing a CESI-compliant enterprise content delivery solution allows IT departments to maintain an efficient network without introducing new constraints.

How Does OmniCache Work?

OmniCache is a robust video proxy cache purpose-built to deliver enterprise video with the lowest possible impact on your network infrastructure and security.

HTTP caching servers designed for storing and redistributing video data are placed in strategic locations around the network. For example, caches can be set up at the edge of the enterprise WAN close to users in regional and branch offices.

When a viewer requests a video, the cache retrieves it from the video source and stores a local copy. When other viewers in the same location request the video, they are served from the local cache. As a result, fewer video streams are pulled across the corporate internet connection and travel the full distance from source to viewer, reducing the overall traffic on the network.

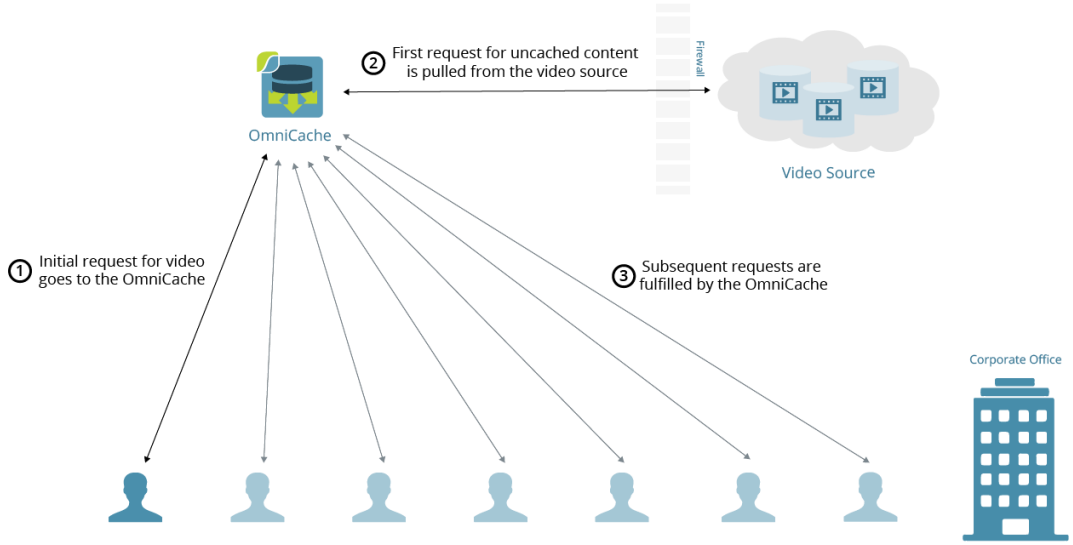

Examples of OmniCache Network Deployments

The way you deploy OmniCache depends on the location of your facilities, the concentration of employees at each location, and your WAN and internet connectivity.

Caches are usually placed on the LAN side of your internet connections for locations directly connected to the internet.

If your company has a single internet connection, or multiple locations are interconnected and share an internet connection, the relative speed of the WAN link becomes an important consideration. Locations with higher concentrations of viewers, or with lower speed WAN links, may benefit from a cache placed behind the WAN connection as well.

Local Cache

The local cache is a simple deployment model used when the network connection for a given location is insufficient to support the expected number of viewers. OmniCache is deployed on site to serve video to the viewers at this location, eliminating the need for each viewer to pull their own stream across the constrained network connection.

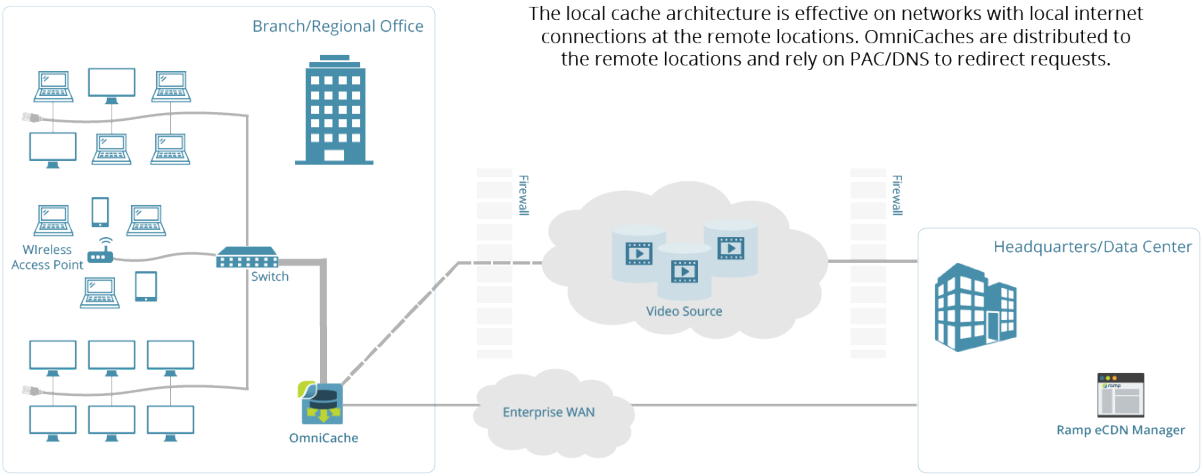

Central Cache

A centralized cache deployment model is used when the corporate network is configured to route all traffic from certain locations through a data center to reach the internet. If the network connections between the data center and the given locations are sufficient to support the expected number of viewers, centralized caches can be set up to serve those locations while drastically minimizing the number of video streams passing through the firewall.

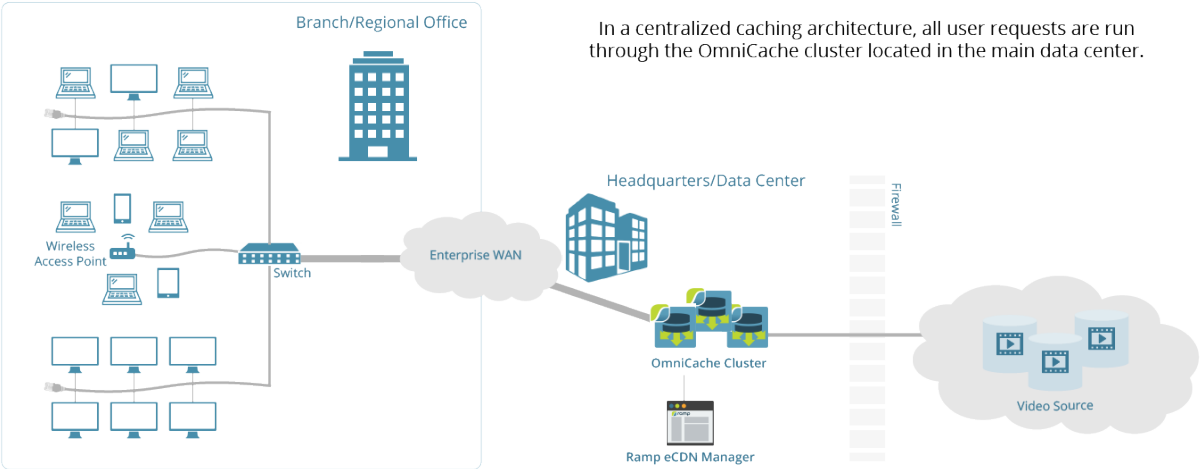

Deployments like this normally require multiple OmniCaches to be located together in a cluster. To ensure traffic is balanced across the OmniCache cluster members, some form of load balancing is used to distribute the traffic.

For small deployments, load balancing can be as simple as a DNS round robin. For larger deployments, a dedicated load balancer is recommended. In both cases, the OmniCaches form a mesh of nodes that actively work together to ensure optimal routing and load balancing.

Central Cache with Routing

Central cache with routing is a deployment model that combines both a centralized cache and local caches.

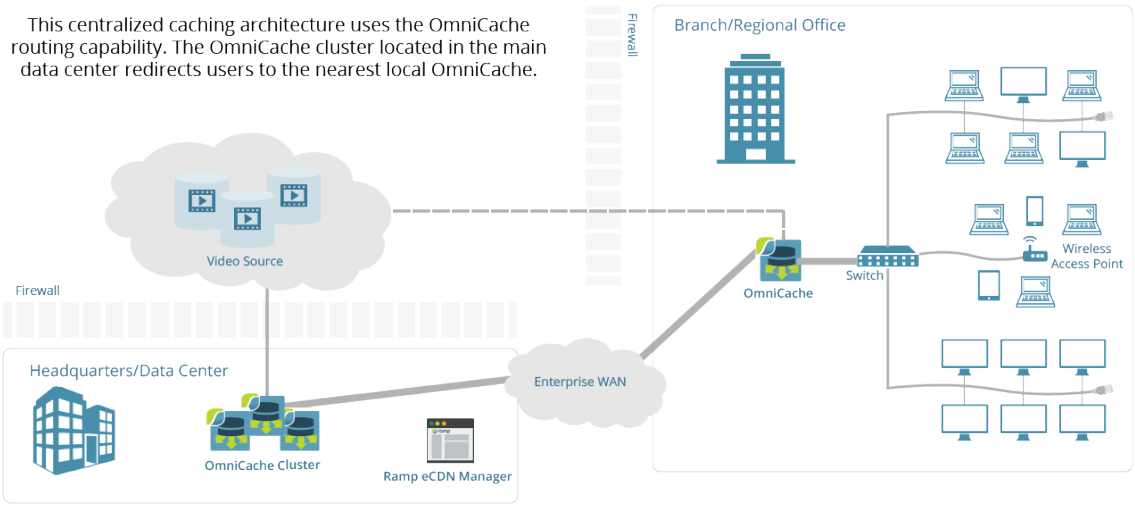

This deployment is recommended when the computing traffic from certain locations is routed through a data center, but the network connection to/from the data center is insufficient to support the expected number of video streams. OmniCaches are deployed at both the data center and at given locations. All video requests are first routed through the central cache in the data center, which then automatically redirects viewers to the nearest local cache to fulfill the request. This routing process ensures users are served by their local cache using their local internet connection.

Directing Viewers to OmniCache

Once OmniCaches are deployed, video traffic needs to be routed to the caches instead of the original video source. This request steering can be done in three ways. The method(s) you use will depend on your video source(s) and your IT preferences.

Integrated Player Request Steering

Many streaming platform providers have partnered with Ramp to integrate their video players with our eCDN. When you use one of these providers, OmniCache operates as a reverse proxy and the video player automatically rewrites URIs to steer video requests to the target cache.

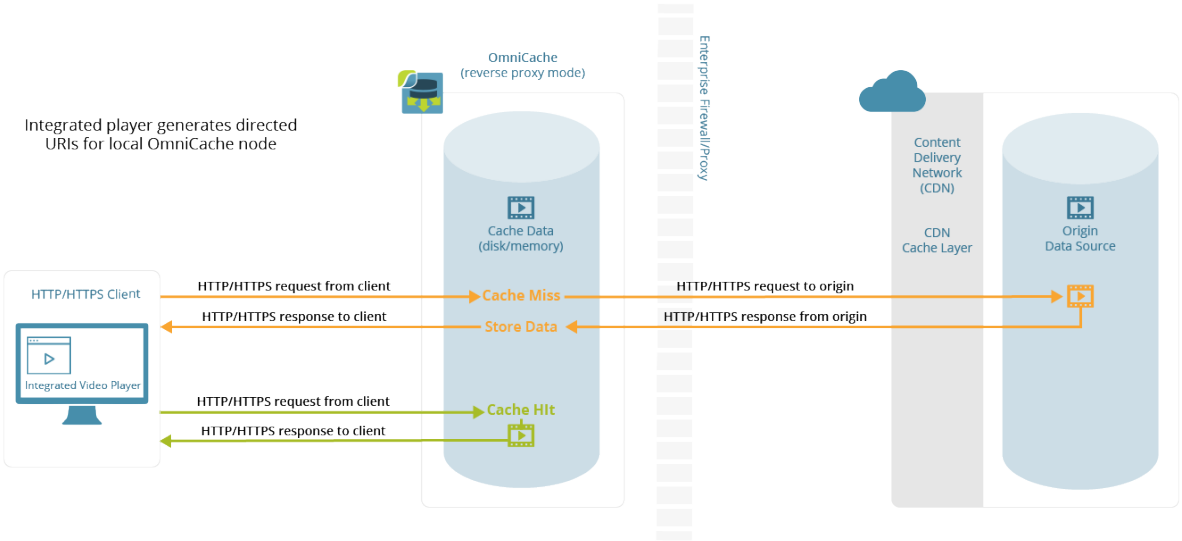

Typical integrated player implementations are tenant-wide, meaning that for a given enterprise, the configuration for accessing OmniCache is fixed to a fully qualified domain name (FQDN). The player queries for the IP address of the OmniCache, allowing your enterprise IT department to control which cache a player uses based on DNS queries—a common practice when network services have local instantiation. Some integrated player implementations are extremely flexible, allowing routing to a specific cache to be determined on an individual video basis.

DNS Request Steering

For video players that are not integrated, you can choose to configure your enterprise DNS to route traffic to OmniCache. With this method, OmniCache runs as a reverse proxy, and you add an enterprise DNS entry for the video source URL to resolve to the OmniCache IP address.

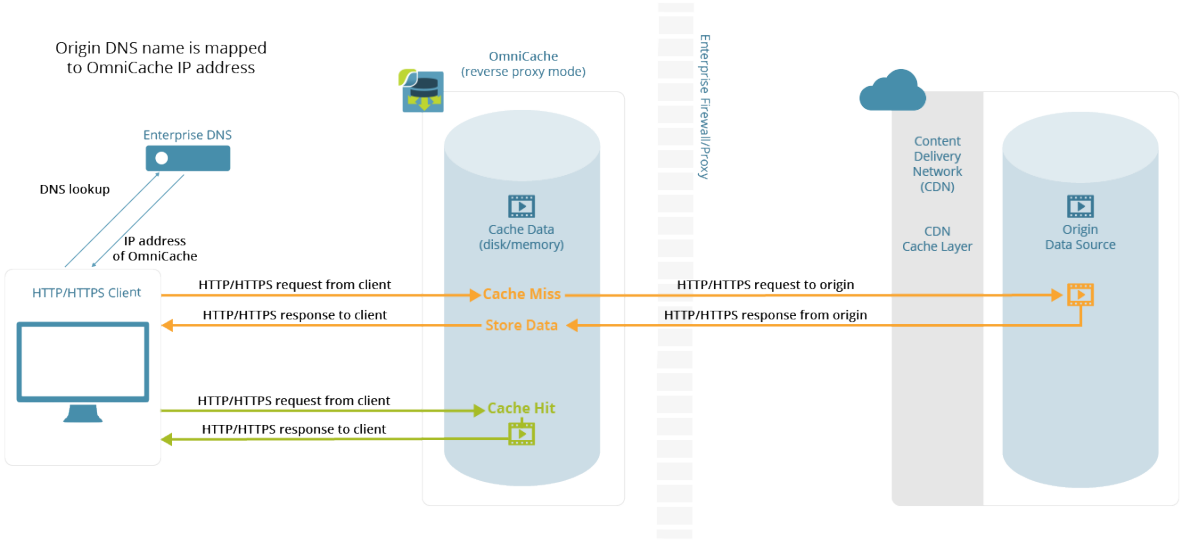

DNS request steering is granular down to the streaming platform level, rather than the individual video. In many cases, this is granular enough.

PAC File Request Steering

Instead of using DNS resolution, you can also choose to use proxy auto-configuration (PAC) files to handle request steering. With this method, OmniCache runs as a forward proxy. PAC files on the viewing devices contain information for the browser to determine which requests to redirect and the routing information for the target OmniCache.

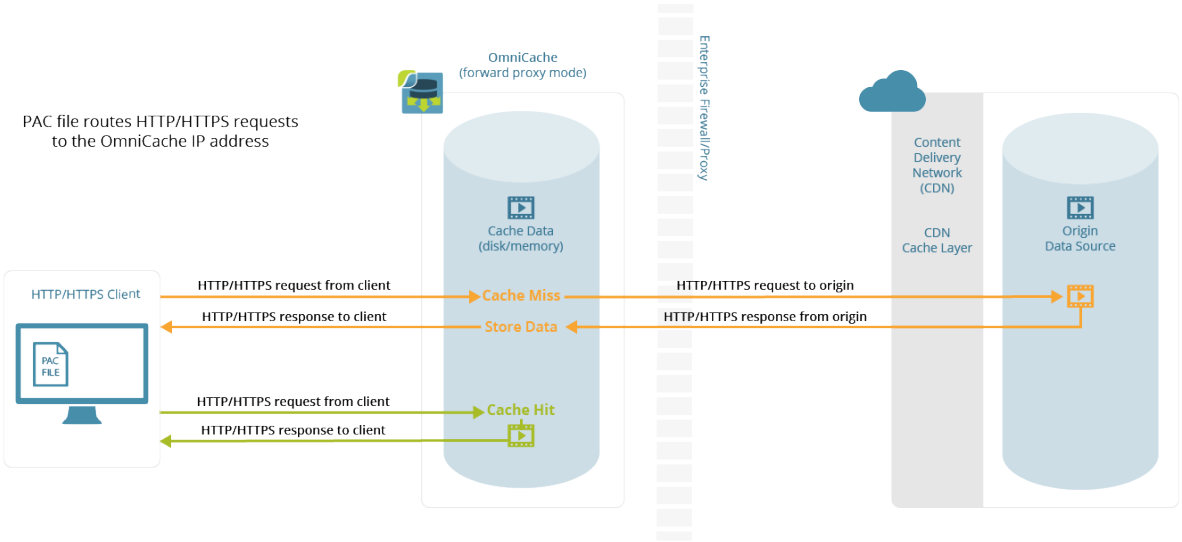

PAC files offer a granular approach to connecting the player to the video source via OmniCache. PAC files work by matching regular expressions, domain names or IP host addresses/subnets with the proxy to be used for each match. Because PAC can be centrally managed, one PAC file can address an entire company.

Benefits of OmniCache: A Purpose-Built Video Cache

OmniCache is a highly flexible, vendor-neutral, enterprise-ready video cache. It offers benefits for more types of streaming video than any other eCDN product or technology. The powerful features of OmniCache maximize viewer experience while minimizing impact on the network.

Vendor Neutral

With more than 75% of enterprises using two or more streaming video platforms and nearly 10% using four or more,[3] a world where one platform can serve all the needs of an enterprise is unlikely. Video streaming and content management platforms are not all created equal. The requirements for a platform that can support global learning and development programs may fall short of meeting the requirements for executive broadcasts, HR communications or quarterly earnings calls.

Because enterprise content delivery networks are network overlays, deploying and managing unique distribution infrastructure for each streaming platform is an IT nightmare—and fortunately unnecessary. Some video platform providers offer video caching as part of their solution, but most of those are proprietary and work only with their respective platforms. OmniCache is deliberately vendor neutral and is designed to support any HTML5 video streaming, as well as legacy solutions such as Microsoft Silverlight (Smooth Streaming) and Adobe Flash's HTTP Dynamic Streaming (HDS) across the corporate network. For enterprises that need to distribute non-HTML5 video as well, OmniCache can generate HTTP Live Streaming (HLS) from encoders supporting RTP or be combined with a media server such as Wowza Streaming Engine™ to ensure complete coverage from a single enterprise delivery solution.

Ramp's solutions prove an eCDN can and should be able to retrieve and redistribute video from any video source. Not only does this simplify the deployment and management of the eCDN, but it provides maximum flexibility for the future as the enterprise video landscape evolves over time.

Widest Variety of Video Formats

Live and VOD

Caching is the only eCDN technology that works as effectively for on-demand video as it does for live streaming. Unlike all-purpose caching solutions, OmniCache is purpose built for storing and distributing large media—designed to minimize latency while maximizing efficiency. A typical OmniCache keeps a fast cache in memory for serving live video (close to real time is desirable) and a slower, disk-based cache for longer term storage and playback of recorded video. In either case, OmniCache can hold an entire video in memory (disk or RAM), giving users the ability to replay, rewind and fast forward. In this way, OmniCache provides DVR capabilities that naturally look and behave just like the origin server.

Because OmniCache supports both live video and video on demand (VOD), it can easily be the only streaming distribution infrastructure you need for your entire enterprise.

Modern and Legacy Video Protocols

Adaptive Bitrate Streaming (ABR) enables highly scalable streaming that adapts to changes in network quality without relying on a feedback mechanism that would burden video servers.

Most video content management platforms initially leveraged Adobe Flash or Microsoft Silverlight as the basic video infrastructure. Both technologies had players that were well supported in browsers and provided ABR capabilities by proprietary protocols, HDS and Smooth Streaming respectively.

With the adoption of HTML5 video players, HLS and Dynamic Adaptive Streaming over HTTP (DASH) have become the preferred streaming methods.

OmniCache supports all these streaming protocols. Companies can deploy a single eCDN infrastructure that will support the migration from an older streaming implementation based on Smooth Streaming or HDS to a modern implementation that uses HLS or DASH.

Progressive Download

Progressive download was one of the early mechanisms for allowing video to start playing before it is fully download. Though HTTP streaming is replacing progressive download as the norm for delivering video flexibly over the internet, progressive download remains common.

OmniCache supports progressive download of MP4 video, ensuring services that don't support HTTP streaming can still benefit from the use of an eCDN.

Player Compatibility and Integration

Unlike other software caches designed to work exclusively with a single proprietary solution, OmniCache can be used with all your platforms to become the only distribution infrastructure you need. It has been tested and verified for compatibility with all the major browser-based video players available from our ecosystem of technology partners, including JW Player, THEOplayer, Bitmovin, Video.js, hls.js, dash.js, Azure Media Player and Flowplayer.

In addition, we have partnered with industry-leading streaming solution providers to integrate OmniCache into their products, making it as easy as possible to optimize your network for video. Our collective goal is to create a seamless experience. In most cases, you simply install Ramp software and select OmniCache in the player's settings. You can review a list of Ramp integrated partners on our website. The full and current list of player compatibilities is available upon request.

Software Solution

Unlike many cache solutions, OmniCache is a software-only solution. Because it is not a physical appliance, you enjoy all the benefits of managing and maintaining software over hardware.

Most enterprises have already invested in virtualization or private cloud deployments. OmniCache can be installed on existing server resources (virtual machines or bare metal) running your company-preferred operating system (Windows or Linux). New OmniCache instances can be turned up as needed, and shut down when not, consistent with any virtualization-/cloud-ready software.

In addition, servers and virtual machines of different sizes and performance specifications can be used to optimize the use of compute resources while relying on the same code instances for all installations.

No Software Client or Plugin

One of the biggest challenges with rolling out any new solution is making sure client software is deployed on all of your user devices. An eCDN is essentially a transparent service, so it can be difficult to detect when a device doesn't have the client installed. Most of the time users can still watch video if they don't have the client properly installed, but the delivery won't use—and reap the benefits of—the eCDN.

Fortunately, OmniCache avoids this issue altogether. The administrative configuration done at the time OmniCaches are deployed automatically and transparently directs your users to the eCDN without the need for any client software to be installed. You get maximum eCDN coverage across your user base without burdening user systems with software that runs constantly in the background.

The lack of a software client also means any device that can play a video from its source will also be able to play it via the cache. Smartphones, tablets and devices with guest access are supported equally well with no user action required.

Centralized Management

Ramp's eCDN management platform is an easy to use, web-based interface for configuring, administering and monitoring your entire OmniCache environment.

OmniCache networks can scale rapidly by creating profiles with settings that are pushed to all OmniCaches. Centralizing the configuration ensures all OmniCache nodes are aware of the full deployment, allowing advanced features such as routing and intelligent sourcing to work seamlessly across the entire network.

Adding and removing OmniCache nodes is easily coordinated through the management platform and often requires nothing more than adding the node securely to the configuration.

All communication between the management platform and the OmniCaches can be secured. Enrolling OmniCache nodes requires specific identity security which ensures rogue nodes are not able to disrupt the network.

Real-Time Monitoring and Analytics

OmniCache gathers and exposes advanced diagnostics and insights giving administrators a complete view of key performance indicators and quality of experience (QoE) information for the eCDN and any video asset. Analytics include aggregated and detailed data on the number of viewers by node, video and location. Additionally, the performance of the individual OmniCache nodes is tracked.

Real-time monitoring gives you a jump on identifying issues, so you can quickly respond before viewers can even report a problem. CPU utilization, memory consumption and cache performance are available to identify and correct any misconfigured, overloaded or faulty nodes.

In addition to the reports available through the management platform, all analytic data can also be made available to external reporting tools so you can perform detailed analysis specific to your environment.

Network Transparency

OmniCache integrates transparently with most corporate network architectures, predominantly a three-tier architecture of core, distribution and access, overlaid with Wi-Fi access points connected to the access layer. This typical architecture is designed around centralized services and servers.

OmniCache software is deployed on virtual machines or servers located in the data center and each user establishes a point to point connection to the cache. The video traffic flow matches the design premise of the entire local area network, ensuring the most efficient use of resources.

This transparent integration into the network is especially evident with Wi-Fi, where all traffic is sent to the Wi-Fi controller before any forwarding decision is made. With OmniCache and the Wi-Fi controller both located in the data center, the video traffic is routed optimally over Wi-Fi without the need for any special configuration or modification.

Optimal Cache Performance

Because OmniCache is built specifically for video, it stores data in a way that maximizes efficiency while minimizing latency. By caching data in larger chunks than all-purpose caches (which are biased toward smaller, text-based file types), OmniCache requires less processing time to store and reassemble these large media files.

In addition, our robust video cache will store data differently depending on whether the video is a live event or pre-recorded. During live events, the members of a large audience are not only accessing the same video at the same time, they are also watching the same segments of video at essentially the same time. Time is of the essence to minimize latency, keep everyone in sync and as close to real time as possible. Therefore, OmniCache uses memory (RAM) for storage so it can retrieve video segments, store them, and serve them back to viewers in a very short amount of time.

With video on demand, viewership is more staggered. OmniCache needs to store multiple segments, if not entire videos, to ensure data is available for subsequent viewer requests. A large cache is the key for getting the best viewing experience with the most bandwidth optimization, so in this case OmniCache uses the disk for storage.

Enterprise Grade Security

Ramp knows security is a top priority. OmniCache deploys 100% on premises, behind your firewall, putting the entire infrastructure and management under the control of your IT department.

Access to videos is protected in multiple ways. First, to access a video, users must pass authentication and authorization checks at the video content management system (VCMS). The VCMS is responsible for determining which content users are authorized to view.

Video traffic is then encrypted while in transit and at rest, so your data is protected at all times. OmniCache uses HTTP(S) encryption to ensure video traffic cannot be snooped. Caches can be configured to support (and reject) specific encryption and signing algorithms, which lets OmniCache stay up to date with modern security standards.

As a cache, OmniCache stores content either in RAM or on HDD/SSD. To protect content while stored, OmniCache implements a local encryption scheme that is fully compliant with NSA-adopted NIST Suite B Cryptography. Keys are under the control of admins and can be expired or replaced as needed. To learn more about Ramp eCDN security, read our paper Securing the Distribution of Enterprise Video Streams.

Virtual Desktop Infrastructure

In many industries, Virtual Desktop Infrastructure (VDI) is commonplace. VDI creates special problems for streaming video due to the resource load placed on the hosts running the virtual desktops themselves.

To address this challenge, Citrix and VMware have implemented browser content redirection. Browser content redirection offloads the rendering of video to the VDI endpoint (i.e., the user's physical desktop) while making it appear as if it is working like any other VDI application.

Because most VDI applications are running on a host machine in a data center, they fetch traffic directly at the server then provide it to the client to render (Server Fetch Client Render). However, when the content is being rendered by the client, it may be more efficient to allow the client to directly fetch the video segments (Client Fetch Client Render). With browser content redirection, both of these implementations are possible.

Since OmniCache is software that is also typically deployed in a data center, it works transparently in these environments. Whether you choose server or client fetch, OmniCache can serve video content equally well.

Virtual Private Networks (VPN)

As the name implies, a VPN creates a private network for employees using the public internet to access corporate resources when they're not in the office. With VPN, users securely access tools, applications and other resources housed behind the corporate firewall. It does this by creating an encrypted tunnel to transfer data between users and your local network.

But with so many business applications hosted in the cloud, the VPN tunnel can become a bottleneck. Data requests come through the VPN and go out to the cloud-based resource over your corporate internet connection. The response data comes back through your firewall and finally out to users over the VPN. This doubling of traffic is called VPN hairpinning.

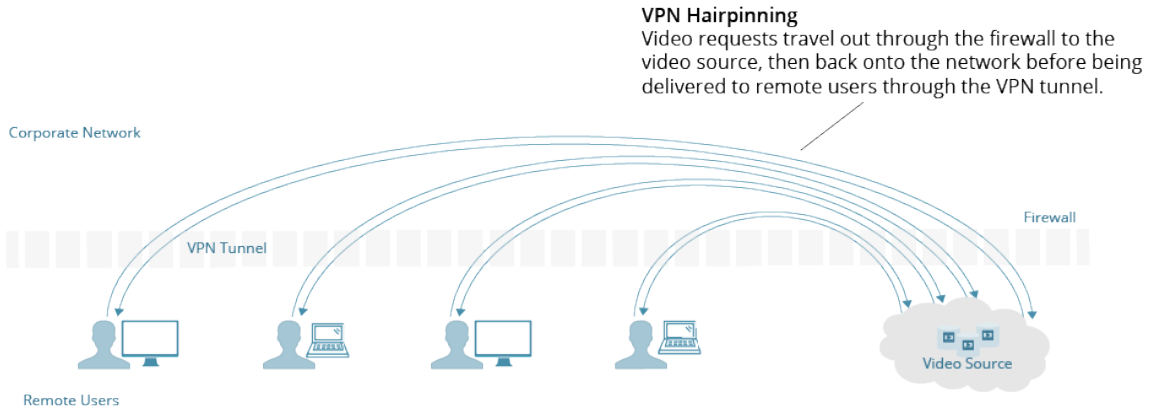

Using OmniCache to serve VPN users, you can eliminate the hairpinning effect. When the first viewer requests a video, the cache retrieves it from the source and stores a local copy. So, only the caches request videos directly from your streaming platform in the cloud (i.e. Microsoft Stream). Then, when another viewer wants to watch a video, they get it from the cache, which resides behind your firewall.

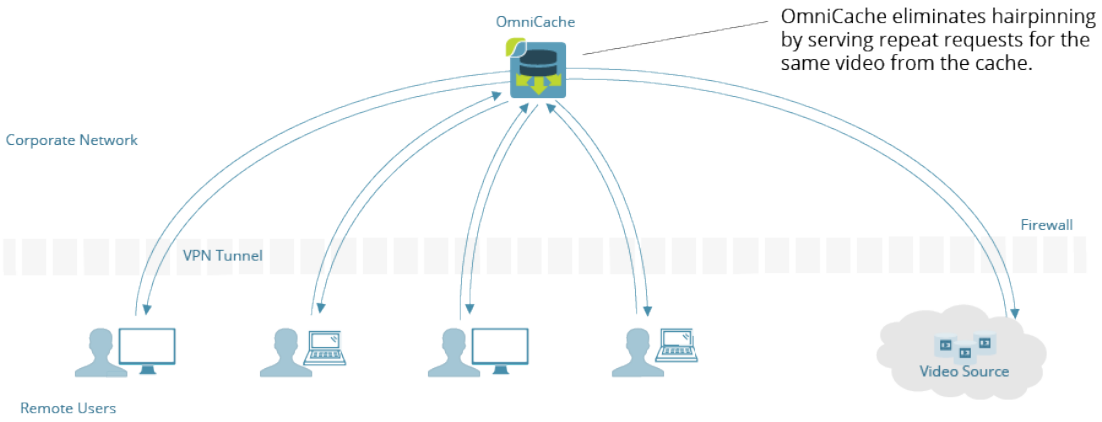

Advanced Bandwidth Optimization

Delivery Profiles

Today's most common video streaming protocols, HTTP Live Streaming (HLS) and Dynamic Adaptive Streaming over HTTP (DASH), are both adaptive streaming technologies designed to serve the best quality of video at any time based on the bandwidth and performance capacity available at the viewing device. To do this, both protocols format their videos as a package of individual video, audio and closed captioning files called tracks. The video tracks differ in quality, resolution and bitrate in order to cater to variations in internet bandwidth. Audio tracks are typically used to support multiple language translations.

When a video is streamed, the player is presented with a menu of video, audio and closed captioning options (a playlist in HLS and a manifest in DASH) as well as the tracks, each of which has been broken up by time increments into a series of smaller files called segments. The player uses the menu to choose the video and audio tracks that best fit its needs. Attributes such as display window size, network speed and network quality all contribute to the selection of the optimal video track. The player can change the version of each track selected on a per segment basis.

Regardless of how many tracks are selected by video players throughout the viewing audience, OmniCache will effortlessly store and redistribute any tracks and segments that have passed through the cache. Therefore, optimization will occur any time a repeat request is encountered.

However, the circumstances that make adaptive streaming desirable on the internet are less applicable on a corporate LAN. On the LAN, enough bandwidth is usually available to safely serve the highest quality video to all the viewing devices. And in fact, by not catering to adaptive streaming, the corporate network can recognize even more bandwidth optimization because only one version of the video track is requested from the video source, stored and redistributed (as opposed to multiple versions).

OmniCache includes a configuration option that allows administrators to select a single video track on behalf of the player. When this delivery profile is invoked, the playlist/manifest is dynamically rewritten to exclude all other track options.

This same feature of OmniCache can be used to control bandwidth consumption at locations with low capacity internet or WAN connections. Instead of selecting the highest quality video for viewers at these locations, the delivery profile can set a maximum limit on the allowed video bitrate that is made available. In this way, admins can ensure video never consumes too much of the available bandwidth at these locations.

Collapsed Forwarding

When OmniCache receives a request for a video segment that is not currently available in the local storage, it forwards that request to the video source. The source returns the video segment to OmniCache, OmniCache stores it locally, then distributes it to the viewer that made the request. This process takes some computing transaction time, even though the amount of time is undetectable by humans.

During a video event like a live stream, it is common for subsequent requests for the same video segment to arrive at the cache before the segment has been retrieved and stored. In a very simple cache implementation, the cache would continue to forward these repeat requests to the video source until the first request is satisfied and the segment is available in storage. These repeat requests erode the level of bandwidth optimization recognized by the eCDN.

OmniCache has a built-in capability called collapsed forwarding that prevents this issue. When collapsed forwarding is enabled, the cache is aware that a request for the video segment is already in motion. It holds all subsequent requests until the segment is retrieved and available to be served locally. Because the video players are anticipating and requesting video segments before they are needed, the slight delay in holding repeat requests does not impact the viewing experience. However, because more requests for segments are served from cache when collapsed forwarding is used, less WAN bandwidth is consumed.

Live Prefetch

Unlike video on demand, where the entire video is complete before any segments are made available, live streaming creates video segments as the video is recorded. Because most viewers of live events are watching at about the same point in the video, their players are requesting the same segment at the same time, leading to a lot of requests for the same segment at once. While collapsed forwarding definitely improves this situation, it would be even better if the first request was already served from cache, thus preventing the need to hold requests and having them time out if the server responds slowly.

By understanding how segments are numbered and when they are likely to be available, the live prefetch feature of OmniCache predicts the next segment and attempts to retrieve it in anticipation of the next request. Live prefetch maximizes the likelihood requested segments will already be available in the cache.

Live prefetch combined with collapsed forwarding gives OmniCache the best opportunity to serve video segments from cache, thus maximizing ROI without any compromise to the user experience.

Proxy Implementations

OmniCache appears as an HTTP and HTTPS proxy to web browsers, ensuring simple and native support for any modern browser. By behaving as a proxy, no client software is needed, and directing traffic to OmniCache is simple.

In corporate networks, proxies are primarily used as a man-in-the-middle system providing some level of access control to the internet and network resources. OmniCache supports both forward- and reverse-proxy modes.

In both cases, the client establishes a connection with OmniCache, and OmniCache establishes a connection with the VCMS, allowing OmniCache to cache the video, audio and closed captioning segments associated with the video.

If SSL/TLS (HTTPS) is used, OmniCache needs an X.509 certificate. To prevent errors, the certificate must be trusted by the browsers, which is typically done by creating a certificate that is signed by the company's certificate authority (CA). The certificate can include an exhaustive list of sites that OmniCache serves, or OmniCache can use a CA certificate and dynamically create site certificates as needed.

This choice is a matter of preference for the company's IT department. The former provides a more deterministic, controlled environment as the certificate only supports previously determined services. The natural behavior of the web browser upon reading an invalid certificate (i.e., when the user accesses a site via OmniCache that isn't one of the previously determined services) is to reject the connection. In contrast, the CA approach ensures all video services are secured. Enforcing access only to trusted sites can be implemented using other rules within the OmniCache.

Forward Proxy

Forward proxy, or standard proxy, is typically used when clients access the internet. Browsers are directed to the forward proxy either using the default browser proxy configuration or using a PAC file.

When deployed as a forward proxy, OmniCache accepts HTTP and HTTPS connections (HTTPS requires a certificate) and forwards the requests for segments to the VCMS.

The browser is aware of the presence of the proxy in this configuration.

Reverse Proxy

A reverse proxy is typically deployed in front of a server farm and acts as a load balancer and SSL termination point.

OmniCache can be deployed as a reverse HTTP or reverse HTTPS proxy (HTTPS requires a certificate). In this case, the client browser is not aware that it is working through a proxy.

Multiple Proxies

While individual proxies are useful, one of the values of OmniCache is the ability to support multiple proxies at the same time. Because different streaming platforms may require different proxy configurations, OmniCache easily supports companies with more than one streaming platform in use.

Predictive Video Delivery

Bandwidth is a precious commodity, especially if the network extends to customer-facing applications like self-service kiosks or point-of-sale systems. Large scale, on-demand video events, such as pre-recorded executive messages or required training, can cause a temporary spike in bandwidth utilization and put those systems at risk. Instead of waiting for the first viewer requests to populate the caches, videos can be pre-positioned on OmniCache during periods of low network activity. Ramp's eCDN management platform lets you schedule when to pre-position videos, monitor the status of videos during pre-positioning, and pre-schedule when to remove video assets from the OmniCaches. This feature better manages network resources and benefits viewers with faster video start times.

Automatic Cache Discovery and Routing

AutoDiscovery is a feature of OmniCache that brings a network of caches into service more quickly than using a manual configuration. OmniCache and the management platform work together to form a complete network of nodes that take responsibility for servicing requests for content.

When an OmniCache is first enabled, it automatically discovers its local IP subnets (the networks it is connected to) and reports that information to the management platform. The management platform uses this information to build a complete routing table of all known local IP subnets. The routing table is used to identify the local OmniCache when a video request is received. Instead of issuing a standard HTTP redirect message, the playlist/manifest is rewritten, which directs the current and all subsequent requests for the video by the given user to the preferred OmniCache. This automatic discovery and rewriting of the playlist/manifest is very powerful when OmniCaches are deployed at numerous branches or remote offices because it automates much of the routing configuration.

However, when OmniCaches are deployed in data centers, the local IP subnets are specific to the data center and don't have users connected to them. The local discovery of IP subnets can be overridden or, more accurately, augmented by manual configuration of IP subnets (i.e., the admin can assign clusters of OmniCaches to server specific groups of users by IP subnet). These manually added IP subnets are added to the routing table so users are always serviced by the right cluster of OmniCaches.

Self-Healing Mesh

A network of OmniCaches can intelligently heal itself when it detects a problem within the video ecosystem. If an OmniCache cluster goes down while video(s) are streaming, the impacted caches are removed from the routing table and the requests fall back to the default OmniCache.

Summary

Companies are embracing the value of video to reach across the enterprise and connect with audiences around the globe. Despite the advantages of video, its impact on network bandwidth can lead to performance problems that impact business operations. Furthermore, a poor experience with low-quality transmissions erodes the intrinsic value of video.

Ramp OmniCache is a standards-based solution that supports live and on-demand video from virtually all streaming protocols, including HLS, DASH, HDS and Smooth Streaming. It brings video closer to viewers to deliver a high-quality and stable video experience to everyone. OmniCache also reduces costs and complexity because it requires no proprietary hardware infrastructure, customized players, or software clients on viewing devices. OmniCache offers the following:

- Lightweight, flexible software that runs and scales on your existing network without proprietary hardware or labor-intensive network upgrades.
- Because OmniCache manages video distribution on-premises at the network layer, it never exposes your network to outside risks.
- Supports any modern streaming video source, any player technology, any device (including mobile), VDI and VPN environments.
- Simultaneously supports all live and on-demand video streaming platforms you use today—and in the future.
- Encrypts your videos while in transit and at rest.
- Content pre-positioning gives you the power to push content to local caches during times of low network activity, which optimizes network resources and avoids bandwidth spikes.
- Centralized management, monitoring and analytics via the Ramp eCDN management platform.
- Easily scales to serve more employees or a new location. The management platform automatically discovers new OmniCache nodes and adds them into the network.
- Routes traffic based on proximity, availability and capacity, and if a problem occurs, it intelligently redirects viewers to the closest usable cache for unsurpassed resiliency.
- OmniCache gathers and exposes advanced diagnostics and insights giving administrators a complete view into key KPIs for specific videos and live events. Analytics include aggregated and detailed data covering OmniCache node performance, bandwidth used, and cache hits and misses.

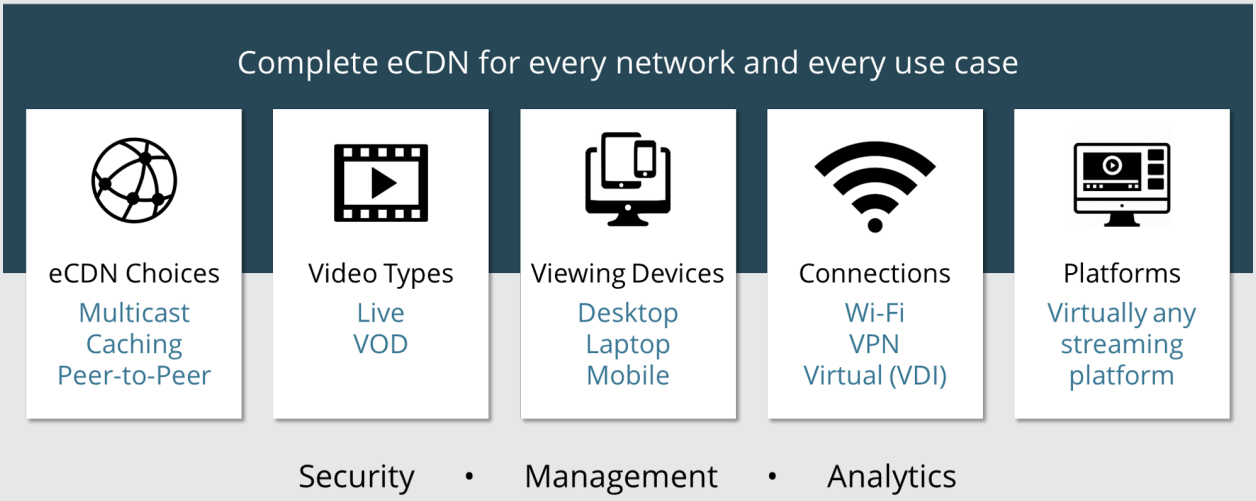

About the Author

Ramp is focused on helping every organization tap into the power of live and on-demand streaming video. Our enterprise content delivery network (eCDN) solutions drastically reduce the bandwidth needed to stream uninterrupted, high-quality video on corporate networks. Using multicasting, video caching, peer-to-peer networking, or any combination, Ramp is the eCDN for all—all enterprises, all networks, all use cases, and all streaming platforms. Ramp works with virtually any modern platform and is tightly integrated with leading streaming video solutions. Our software deploys entirely behind your firewall for maximum security and scales easily as demand for video grows. With centralized management, monitoring and insightful analytics, you get unprecedented visibility into and control over network.
Visit rampecdn.com for more information.

Footnotes

[1] State of the Market: Enterprise Streaming Solutions and Services - 2020
https://insight.wainhouse.com/reportaction/EV-SOTM-2020-Streaming/Toc

[2] Survey Insight: Enterprise Video, ITDM Perspectives - North America Q3 2019
https://insight.wainhouse.com/reportaction/EV-SURV-19Q3-ITDM/Toc

[3] Survey Insight: Enterprise Video, ITDM Perspectives - North America Q3 2019
https://insight.wainhouse.com/reportaction/EV-SURV-19Q3-ITDM/Toc