WOWZA™ media systems

FINDING THE STREAMING MEDIA INFRASTRUCTURE THAT FITS YOU BEST

ON PREMISES, IN THE CLOUD, OR A HYBRID MODEL

OVERVIEW

With the appetite for anytime-anywhere video and audio consumption growing continually, content providers of all types need to deliver on-demand and live streaming that reaches any device—or risk losing their audiences.

Fortunately, organizations of every size now have access to high-quality, reliable, live and on-demand streaming to any screen, and it's clear they're eager to take advantage of it. Yet all too often they're deterred from the seemingly daunting task of deciding how and where to deploy video-streaming technology: on premises, in the cloud, or using a hybrid model. As video formats and infrastructure requirements evolve and the number of streaming playback devices increases, organizations are left scratching their heads, wondering which workflow will best meet their needs now and into the future.

This is a primer for those of you evaluating the deployment of a live streaming workflow (everything between the encoder and the player) on premises vs. in the cloud vs. a hybrid of the two. We'll highlight key factors to consider and then look at several cloud infrastructure examples to help point you in the right direction.

STREAMING POPULARITY

Is streaming really that popular?

A December 2015 study from Sandvine indicates that 70% of traffic on the Internet is from streaming, and a May 2015 report from Cisco predicts that number will hit 80% by 2019. The numbers reflect the overall growth of video as a part of our daily lives. After all, streaming applies to far more than just expanding traditional broadcast models. Examples include live sports, college lectures, government meetings, real-time web conferencing, Internet radio, church services, weddings, traffic cams, and video game matches—the sky is the limit.

WOWZA™ media systems

SCOPING YOUR NEEDS

Deciding whether a streaming media workflow should use on-premises or cloud-hosted software and/or cloud-based services comes down to a number of who, what, when, where, and how factors. To help discover which approach is best for your organization, start by asking the following questions about your business requirements.

FUNCTIONALITY

What are the key user-experience and business requirements for your streaming?

You might be trying to meet an internal communications goal or delight a global consumer audience, and the scope of features and capabilities between those two goals can be wide-ranging. Evaluate whether you need to secure your content, bill for access, ensure fault tolerance, deliver with low latency, meet certain viewer expectations, track how long each viewer watched, provide closed captions, etc. Each will have cloud vs. on-premises pros and cons to consider.

CONTENT

What kind of content are you delivering?

A seated speaker and limited on-screen text can often be delivered reliably with nonadaptive low-bitrate streams, reducing the computing power and bandwidth you'll need. On the other hand, content with a ton of visual detail (e.g., lots of small text or a surgical procedure) or high motion (e.g., a basketball game) requires higher bitrates to have the same visual quality. Higher bitrates often lead to adaptive bitrate streaming (to dynamically support older viewing devices and fluctuating network connections), which in turn requires your infrastructure to support more processing power and aggregate bandwidth, and this may tip the scales in favor of a cloud infrastructure.

RESOURCES

What people and equipment are you starting with? How much are you willing to invest to build and manage live-streaming workflows?

If you already have staff with hands-on experience building a live-streaming infrastructure, you're off to a good start. Similarly, whether you have cloud-knowledgeable team members is a key factor. Assess if having streaming and cloud infrastructure knowledge in house will be core to your business's value, or if you 're comfortable outsourcing it. Another facet of resources is your existing equipment and infrastructure. If you already have a robust datacenter on

If all the components in the dashed box below are in the cloud, this is a cloud-based deployment. If only some are in the cloud, this is a hybrid deployment.

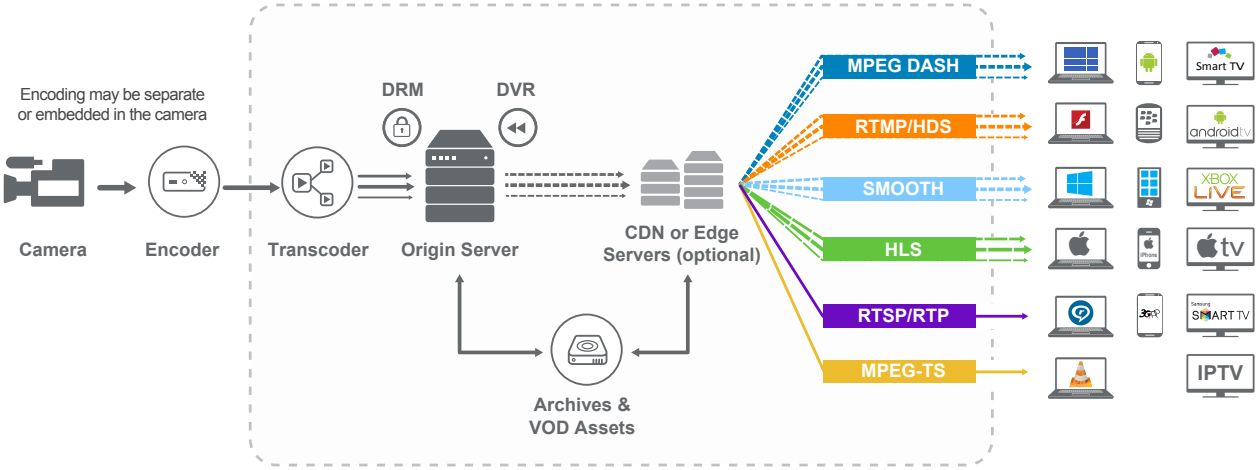

A generic any-screen streaming workflow with processing and delivery fully or partially based in the cloud

premises with fat pipes to the Internet, streaming to a content delivery network (CDN) or directly to a well-defined number of external viewers may be comfortably within your capabilities. Internal live streaming can often be handled by just one or a few servers, depending on content, network topology, multicast network enablement, number of viewers and locations, etc. In fact, companies often jump-start their internal streaming by repurposing old email and database servers.

AUDIENCE

Where and how numerous are your viewers?

Your target audience is arguably the most important consideration in an on-premises vs. cloud decision. If all or most of your viewers are located onsite, an on-premises deployment is the most common approach. Conversely, if most of your viewers are not on your organization's intranet or are geographically dispersed, offloading the processing and delivery to a third-party infrastructure usually makes more sense.

TRAFFIC

How predictable are viewership spikes, and how quickly will you need to scale to consistently deliver the highest-quality video streams?

If your viewing audience is a fairly fixed in size and location, then ensuring sufficient server and network resources may be straightforward. More often, there is some level of audience-size volatility within a live event and from one live event to the next. When using HTTP adaptive bitrate streaming, such unpredictability can often be addressed relatively simply with a CDN. Whether it's an enterprise CDN you've built on premises to service your own intranet users, or a third-party cloud-based CDN that is delivering to global Internet viewers, HTTP caching servers in a CDN typically scale well to accommodate a highly variable viewership.

However, if you're delivering using traditional streaming protocols such as RTMP or RTSP, scaling typically requires deploying more streaming servers. For an audience on the Internet—especially when events are fairly infrequent—it may be more cost-effective to spin up cloud-based servers only when needed.

CONTROL

What level of control and customization do you want over your organization's streaming infrastructure?

An on-premises deployment gives you full control over your workflow. You have both the privilege and the responsibility of managing every level of your technology stack. A cloud deployment cedes some control and management to the service provider. Hybrid deployments fall in between, sometimes providing the best combination of control and third-party management.

SPEED TO MARKET

When do you want to launch live streaming?

Time to market is a critical element of an on-premises vs. cloud evaluation. If you're just starting out with live streaming and still defining your needs, a great way to move forward is to define a minimum feature set that is relatively simple and safe to try, and then learn from each live-streaming experience. We regularly see companies start streaming quickly with a managed cloud service, and then over time they hone their expertise, refine their requirements, and move toward a customized cloud or hybrid model.

COST

What are your budget and preferred financing model?

Your budget will help define whether hiring people or purchasing servers is even an option. The financing model should tell you whether it's preferable to emphasize capital expenditures (capex) or operating expenditures (opex).

WOWZA™
media systems

WORKFLOW OPTIONS

Based on your responses to the preceding considerations, you and your team should now be in a better position to decide which of the following stream-processing and -delivery workflow types is best suited to you.

FULLY ON PREMISES

WOWZA STREAMING ENGINE™

With a fully on-premises solution you have complete control over your streaming infrastructure, from the networking layer up to the viewing application or user experience. You also get lower opex and more predictable workflow costs. However, it may be a costlier choice overall, due to initial and recurring capex for hardware, the need to overprovision for spikes in viewership, and continual operating system and network maintenance.

IN THE CLOUD

WOWZA STREAMING CLOUD™

Cloud-based deployments can scale your capacity up and down in minutes, while also cutting costs for hardware and maintenance. With a cloud deployment, your financing model focuses on opex instead of capex. Organizations using the cloud can rapidly scale to handle fluctuations in viewership and transfer intensive processing jobs to third-party infrastructure. These benefits make the cloud ideal for live-streaming events where viewership quickly fluctuates.

THERE ARE THREE PRIMARY OPTIONS: IaaS, PaaS, and SaaS.

Infrastructure as a Service (IaaS)

In the IaaS model, a self-managed cloud deployment of virtual machines gives you more control of your infrastructure while leaving the core hardware and network management to the provider. With your control comes greater responsibility for maintaining the virtual machines, including operating system updates and security.

Platform as a Service (PaaS)

With PaaS, a cloud-based managed service minimizes the complexities of live streaming by offloading most of the configuration and management to third-party specialists, so you don't have to make big investments in resources and expertise to get started. However, what it offers in scalability it sometimes lacks in control.

As your needs evolve, you may want to work with a cloud-based service that provides advanced features or a REST API for fine-grained control, plus is agile about adding new capabilities and keeping your streaming platform future-proof. Alternatively, at that point you could deploy an IaaS-based solution for the highest level of cloud infrastructure control.

Software as a Service (SaaS)

Finally, with SaaS, a vendor provides the player on a hosted page and abstracts you from all of the underlying processing and delivery components. In this case, you would simply point your viewers to that hosted web page.

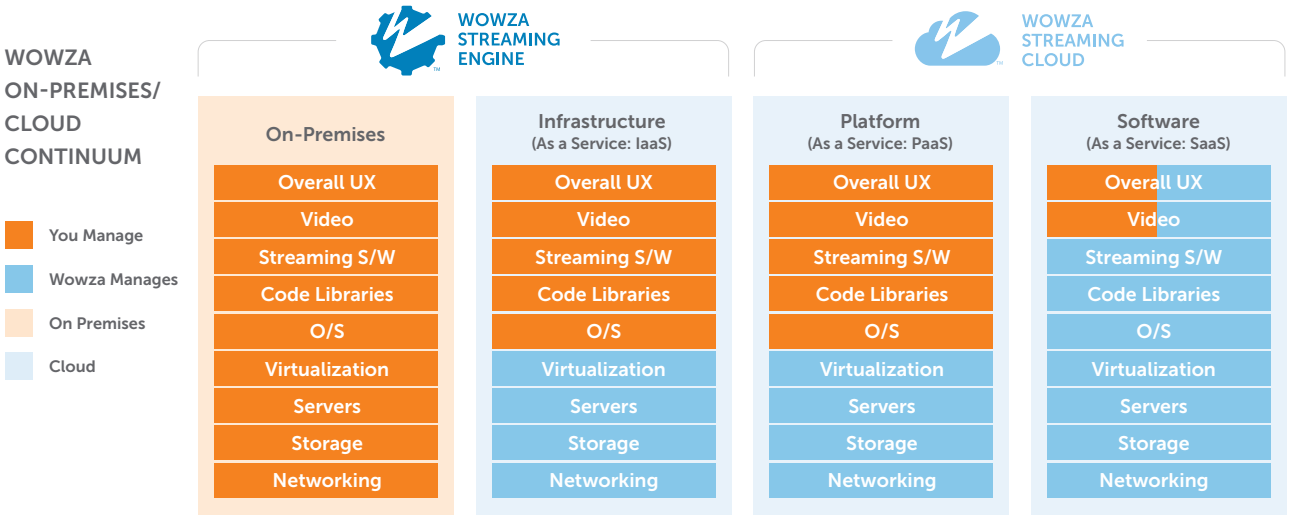

PARTIALLY IN THE CLOUD (HYBRID)

An approach that spans both on-premises and cloud resources is the most flexible in terms of future workflow mobility. Organizations can mix and match components to create the ideal balance of economics and control.

Hybrid deployments come into play to handle all sorts of streaming scenarios, from offloading of peak volume distribution to offsite processing of dozens of long-running live or linear OTT (over-the-top) channels. While a hybrid model can enable organizations to partner with third parties to handle scaling and technology updates, it still requires human resources to locally manage parts of the workflow.

A common hybrid workflow is any scenario in which a live camera feed is encoded on premises and published to a local origin server, which in turn sends the stream to a CDN for broad distribution. Similar to (and sometimes co-located with) other cloud services, CDNs provide globally distributed computing and storage datacenters connected by fat Internet connections. They optimize those resources to move content as quickly as possible to reach the end users who are requesting it, focusing on optimizing delivery routes and methods to achieve low-latency responses.

Let's look at a simple enterprise hybrid (Iaas) deployment example, which may provide additional insights. Traditionally, enterprises (medium-to-large companies and institutions) have had on-premises deployments. This was driven by several factors, including local control, available bandwidth, bandwidth costs, and reduced latency. Today, however, enterprises are increasingly using the cloud to power their streaming workflows. This example could apply to any live event, such as an all-hands meeting.

Locally, an encoder captures live feeds from one or more cameras and creates a single 1080p output stream that it simultaneously records and pushes to a local streaming media server. That server sends the HD stream to a more powerful IaaS cloud instance that is launched an hour ahead of time and configured as a live transcoder plus media server.

The cloud-based transcoder converts the single 1080p stream to a set of streams at different resolutions and bitrates, ranging from 1080p down to 160p, and sends them back down to the onsite media server. The local media server then repackages (requiring little CPU power) the incoming ABR stream set into the adaptive Apple HLS (HTTP Live Streaming) format and, as backup for older devices and browsers, a single traditional Flash RTMP stream at 480p.

Viewers go to an intranet content portal and click on the video player embedded there. Depending on their device playback capabilities, they see either the HLS adaptive stream or the RTMP stream on their screen.

WRAP-UP & NEXT STEPS

Live video streaming is not a one-size-fits-all approach. Almost every scenario has a unique set of requirements. Choosing the right workflow to suit your organization's particular needs significantly affects the benefits you gain.

Determine a few key deciding factors for your scenario, and then—as needed—consider building a spreadsheet that compares likely deployment models. Your spreadsheet could just be a pros and cons list for each model, but since every factor (whether hardware, resources, third-party services, or time to market) has a financial impact, comparing projected bottom-line costs may be more objective and informative. Pick the simplest path forward and start testing. As you experiment, you'll likely find that you rapidly refine your requirements list and with every successive test get closer to the most effective workflow for you.

ABOUT WOWZA MEDIA SYSTEMS, LLC

Wowza Media Systems is the leading software and cloud-based services provider that continuously simplifies the complexities of high-quality video and audio streaming for organizations worldwide. Wowza delivers an award-winning, extensible solution set that provides unparalleled customization, configuration, and control for an unlimited number of streaming use cases. The company's integrated, hybrid model empowers direct end customers, service providers, and an extensive worldwide partner network. Proven, patented Wowza technology enables organizations to expand their reach, more deeply engage with key audiences, and increase their monetization opportunities. Wowza has rapidly built a fan base of more than 19,000 customers, in 170+ countries, across all industries. The company was founded in 2005, is privately held, and is based in Colorado. For more information, please visit www.wowza.com.

To learn more about how others are using Wowza streaming technology, visit wowza.com or contact sales@wowza.com.