THE FULL MONTY: MOVING BEYOND THE HYBRID CLOUD

M. Yanchyshyn

Amazon Web Services, Inc., USA

ABSTRACT

High-end media projects require geographically diverse talent as well as powerful hardware and software. Fluctuating workloads, unpredictable budgets, and globalized production teams have prompted many companies to adopt hybrid workflows where parts of the production pipeline, such as long-term storage and processing, are moved to the cloud. Many cloud providers have on-demand, pay-as-you-go services that map to hybrid production models and variable workloads, but this solves only part of the problem. Many media companies still buy expensive hardware to handle parts of the pipeline on-premises. This includes synchronizing large media files, using high-performance storage, and providing employee workstations. This paper addresses these barriers to fully embracing the cloud by presenting an end-to-end media production pipeline built entirely in the cloud. Learn how to combine the latest advances in application streaming, high-performance storage, clustered media processing, and high-speed file delivery to form a complete pipeline in the cloud that lowers costs and meets the needs of today's professional media projects.

INTRODUCTION

Many media companies are moving some of their video and visual effects workflows to the cloud, but leaving some components—such as primary data storage and editing workstations—on-premises. Recent developments, including accelerated file transfers; high-quality application streaming for GPU-accelerated software; high-performance, shared storage in the cloud; automation of cloud-based, distributed media file processing jobs; and flexible licensing are removing most barriers to adopting end-to-end media solutions in the cloud. This paper explores techniques for building cloud-based, high-quality media workflows that are designed to lower costs, reduce operational complexity, and increase productivity.

TRANSFERRING DIGITAL MEDIA TO AND FROM THE CLOUD

File sizes continue to grow every year as new technologies and consumer trends emerge. These files often have higher bitrate and video resolution requirements, such as 4K and more recently 8K or FUHD. A five-minute HDTV video can result in a file size of roughly 2.5 GB when packaged as MXF OP1a with MPEG-2 and PCM encoding. This size is modest compared to files of similar lengths that have scenes with complex visual effects. Large file size and the associated time and cost of file transfer present significant challenges for organizations that want to adopt cloud-based media workflow, particularly if their editing workstations, primary storage, and other components are located on-premises.

For example, a visual effects artist may use compositing software on a local workstation, but then want to perform the final render on a remote cluster of multiple virtual machines. Traditionally, the FTP protocol would typically be used to transfer the required assets from the artist’s workstation to the remote render farm; however, given the large file sizes associated with today’s media projects a significant amount of network bandwidth would be required to transfer the assets quickly enough to meet production deadlines. FTP also performs poorly when transferring files over long distances due to the overhead associated with how TCP protocols handle window sizing and packet retransmission. Security is also a concern because FTP transmits passwords in clear text. A number of alternatives to the FTP protocol have emerged in recent years in response to the need for higher throughput plus more reliable and secure data transfers. They employ a mix of techniques to improve performance and security: encryption, parallel data transfer streams, data compression, disabling congestion control, de-duplication, and the use of the UDP protocol.

Multi-Part, Parallelized Uploads

Amazon Simple Storage Service (Amazon S3) is an example of a cloud-based data storage service that facilitates large file transfers with a multi-part upload API (1). Clients and SDKs that support this API automatically open a number of parallel HTTPS connections to upload a single file as a set of unordered parts. Retransmission of failed parts is handled transparently. This results in higher data transfer throughput, quick recovery from any network issues, and the ability to begin an upload before you know the final file size.

TCP Alternatives

Because HTTPS is a TCP-based protocol, it can suffer from some of the same performance issues as FTP when used over long distances or poor network connections. Parallel transfers, properly tuned TCP window settings, removal of congestion control, and other techniques can mitigate this; however, these modifications are beyond the control or comfort level of many users. As a result, protocols built on UDP have grown popular in the media industry for transferring large digital files. There are many solutions available, but for those users who want to move assets to Amazon S3, there are only a handful of commercially supported offerings that integrate with its APIs: Aspera (2), Signiant (3), File Catalyst (4), and Data Expedition (5).

The implementation and associated ecosystem of tools differ greatly among the file transfer solutions offered by these companies, but they share much in common. UDP is used for the bulk of the data transfer because it typically allows for higher throughput due to several factors: UDP is a connectionless protocol and does not degrade over long distances or with high latency, it tolerates packets that are delivered out of order, and it does not require packet receipt acknowledgement or retransmission. Some of these traits taken in isolation can result in higher error rates and less reliable transfers, so the accelerated file delivery solutions mentioned above implement additional features on top of UDP, including error correction, multiplexing, data transfer rate control, compression, and encryption. The net result is that file transfers are considerably faster and more reliable, particularly over long distances.

Test results vary depending on many factors; however, a recent test of UDP-based file transfer protocols moving 650 GB files from virtual machines hosted in Amazon Web Services (AWS) datacenters in Tokyo, Japan to virtual machines hosted in AWS datacenters in Virginia, USA showed a ten-fold speed increase over standard TCP-based file transfer protocols (6). Performance gains are often even higher (7, 8, 9).

Private Connections

Private connectivity to cloud providers is another way to reduce latency, increase available network bandwidth, and remove the burden that large file transfers put on your primary Internet connection. For example, AWS Direct Connect lets you establish a dedicated, private network connection from your datacenter, office, or co-location facility to AWS. AWS Direct Connect links your internal network to an AWS Direct Connect location (10) over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, and the other to an AWS Direct Connect router. If you don't have equipment hosted in the same facility as AWS Direct Connect, you can use a network provider to connect to AWS Direct Connect (11). With this connection in place, you can create virtual interfaces directly to services in the AWS platform such as Amazon EC2 and Amazon S3, bypassing Internet service providers in your network path.

Shipping Physical Storage Devices

Some media projects are housed in facilities with limited bandwidth, are unable to procure a private network connection, or have digital assets that are simply too large or numerous to upload in time to meet production deadlines. AWS Import/Export (12) is one example of a service that addresses this common problem by letting you ship your data to the cloud by means of portable storage devices. AWS Import/Export transfers your data directly onto and off of storage devices using the high-speed internal network of Amazon, bypassing the Internet or a private link. For significantly large datasets, this approach is often faster and more cost effective than upgrading connectivity.

MOVING YOUR WORKSTATIONS TO THE CLOUD

Although the protocols and other suggestions that we’ve discussed so far greatly improve the speed at which media files can be moved to and from the cloud, they still introduce considerable delays and complexity. One solution is to move the workstations and software used by media professionals to the cloud so that they are co-located with the digital assets.

Virtual Desktop Infrastructure (VDI)

Virtual Desktop Infrastructure (VDI) refers to the process of running a user desktop inside a virtual machine that lives on a server in a datacenter (13). Most VDI solutions offer enhancements over standalone Remote Desktop Protocol (RDP) or Virtual Network Computing (VNC), such as session virtualization, user management, enhanced security, and improved performance. Although they may work well for business applications and software without special hardware requirements, running professional media software on VDIs has historically proven more difficult due to a lack of support for OpenGL or DirectX and, in some cases, GPU acceleration. Most professional media software also requires considerable compute and graphics resources.

In response, several VDI vendors have released solutions that run on GPU-enabled hardware and include driver support for OpenGL and DirectX. Popular examples include VMWare Horizon View, Citrix XenDesktop HDX 3D Pro, and Microsoft RDP8/RemoteFX. However, until recently these solutions were not compatible with large-scale cloud providers that offered pay-as-you-go pricing models due to a lack of virtual machines and associated hypervisors with integrated GPU-enabled graphics cards.

GPUs in the Cloud

NVIDIA has worked with a handful of cloud service providers to offer virtualized hardware with GRID graphics cards (15). IBM SoftLayer introduced a new service in late 2012 where virtual machines running on hypervisors with NVIDIA GRID (Kepler) GPUs could be rented on a month-by-month basis. While this was an improvement over large capital expenditures for GPU-enabled hardware, SoftLayer's monthly pricing model is not granular enough for the highly variable workloads of many media professionals. In late November, 2013, AWS addressed this gap by releasing its "g2.2xlarge" virtual machine type with an hourly pay-as-you-go pricing model. In April 2015, a new version was released with four times the capacity. The graphics cards on both g2 virtual machine types feature on-board hardware video encoders, each designed to support up to eight real-time HD video streams (720p@30fps) or up to four real-time full HD video streams (1080p@30fps) (16).

The ability to meet hardware requirements of professional media software and pay for it by the hour solves only part of the problem; it is arguably an even greater challenge to deliver an acceptable user experience when using this software over the Internet, regardless of the remote hardware capabilities. In the case of high-end video or visual effects production, any significant degradation of the quality of the images displayed on the user's screen is unacceptable, as is any noticeable latency in cursor movements, UI response, and keyboard actions. Accurate color and quality audio output are also important. These requirements can be met by making efficient use of the GPU, by choosing the right protocol, and by reducing network latency.

Amazon EC2 g2 virtual machines use a technique called "GPU pass-through" to give each virtual machine direct access to a full GPU card. Direct access with native NVIDIA drivers results in considerably better performance due to full usage of the graphics card's capabilities and less resource usage than software GPU emulation or GPU virtualization via API intercepts.

Choosing an Application Streaming Protocol

Microsoft Remote Desktop (RDP) with RemoteFX (19), Teradici PCoIP (20), and Citrix HDX 3D Pro (18) are three of the most popular remote-access protocols with support for GPU acceleration and OpenGL/DirectX. They are all proprietary, Layer 7 protocols that are based on a mix of TCP and UDP. Like the previously mentioned accelerated file transfer protocols, they employ a variety of techniques to improve performance and enhance security such as encryption, data compression, intelligent error correction, and connection optimization.

Support for operating systems and virtualization platforms differs among protocols. PCoIP, HDX, and RDP8 officially support only Microsoft Windows. Users who want to use Linux can use a combination of VirtualGL (21) and TurboVNC (22); however, at this time these do not support audio output. Microsoft RDP with RemoteFX is also primarily targeted at users of the Hyper-V platform, so there is no documentation or support for leveraging it to deliver 3-D applications running on Amazon EC2. For users who want to take advantage of the hourly pricing model and g2 virtual machines of Amazon EC2, Teradici's PCoIP and Citrix XenDesktop with HDX 3D Pro are well supported options.

Alternatives and Newcomers

In addition to the protocols discussed in the preceding sections, new offerings have emerged in recent years for users who want to stream GPU-enabled applications from virtual machines in the cloud. Amazon AppStream (23) uses a proprietary protocol called STX to deliver a high-quality, GPU-accelerated application stream. It supports high-quality audio and up to 4:4:4 chroma subsampling for accurate color representation. Frame (formerly Mainframe2) takes another approach, sending the screen image as an interactive H.264 video stream. This allows users to access remotely hosted, GPU-accelerated software through their web browsers without any plugins (24). OTOY's Octane Cloud and ORBX.js take a similar approach to Frame, leveraging the proprietary ORBX streaming codec, HTML5, and JavaScript to deliver a high-quality, interactive application stream within a web browser. Octane supports both Windows and Linux application streaming (25, 26).

Improving Network Performance for Application Streaming

In addition to choosing the right protocol to meet your platform and performance needs, it is also important to reduce latency and sustain high network throughput to maintain a high-quality experience when using applications remotely. Cursor movement and user interface (UI) response are especially sensitive to elevated latency, and low or variable network throughput can result in visual artefacts, distortion, and discoloration.

A private connection to your cloud provider such as AWS Direct Connect is one way to reduce latency and increase connection reliability. In addition, WAN optimization can greatly improve application streaming performance. Citrix's CloudBridge is one example of a network acceleration solution that you can deploy as a hardware or software appliance to act as a network acceleration gateway for application streaming. It employs a variety of techniques to improve performance, including adaptive TCP flow control, compression, caching, and de-duplication, as well as protocol optimizations such as minimizing handshakes and payload optimization (27, 28).

SHARED STORAGE IN THE CLOUD

High-performance, shared storage is a key requirement for any media organization that wants to migrate their workstations or do batch processing of media files in the cloud. Although cloud-based object storage is a great choice for parts of many media workflows, a large number of professional media software packages are not designed to interact with the HTTPS APIs used by the most popular cloud object storage services. Cloud-based, compatible alternatives include distributed file systems such as Amazon Elastic File System (Amazon EFS), Azure File Storage, Lustre, and GlusterFS; protocols such as NFS and SMB/CIFS; and cloud-aware filers like NetApp Cloud ONTAP and Avere FXT.

A key differentiator of cloud-based media workloads is the ability to size your infrastructure and associated costs to match the variable nature of the work being done by most media organizations. Many cloud-based storage systems can be resized on demand and, in some cases, removed and quickly restored as needed. Other factors to consider include performance, operational complexity, scalability, and hybrid deployments.

Intel Lustre

Intel Lustre (31) is a popular file system for providing the storage needed for clustered workloads such as transcoding and rendering, or as a shared file system where multiple people and processes access pooled media assets from remote workstations and services. Lustre has grown particularly popular for high-performance computing (HPC). The popular analytics software vendor SAS recently demonstrated how they achieved aggregate storage throughput of 3 GB/s at 50 MB/s per CPU core when using Lustre as a shared storage system for their SAS GRID product (30). This level of performance highlights Lustre's viability as a high-performance file system for media workloads, and puts it on par with many enterprise-grade SAN solutions.

Lustre is available on the AWS Marketplace and can be deployed automatically using infrastructure management tools such as AWS CloudFormation. Cfncluster is another tool that greatly reduces the time it takes to create an Intel Lustre cluster on AWS. Cfncluster handles the deployment of numerous components including virtual networking, security, monitoring, storage, and compute resources needed to power your Lustre cluster (35, 36).

Amazon Elastic File System

Amazon EFS (29) is an elastic file system designed for the cloud. Like Lustre, it can be deployed on demand, and you pay only for what you use. Unlike Lustre, Amazon EFS is a fully managed service, so the complexity of deployment and data storage capacity and performance scaling is greatly reduced. Data stored on Amazon EFS is automatically replicated across multiple datacenters across a wide geographical area, leading to very high data durability. Amazon EFS supports the NFSv4 protocol, so you can mount its file systems on a large number of operating systems and maintain compatibility with most media software. Some of the issues traditionally associated with the NFS protocol have been addressed in Amazon EFS such that failover is seamless, and very large NFS file systems with many concurrent user connections perform consistently and reliably.

Avere Hybrid Cloud NAS

Avere Hybrid Cloud NAS is a good option if you need direct support for the SMB/CIFS protocol in addition to NFS, or if you need a global namespace for your storage that spans on-premises and the cloud. Avere FXT Edge Filers provide high-end NAS functionality with SAN-like performance and can span both on-premises storage appliances and cloud-based object storage systems such as Amazon S3 and Google Cloud Storage (32, 33).

DISTRIBUTED MEDIA PROCESSING AND AUTOMATION

Whether you are running a large transcoding job, rendering a 3-D animation scene, or doing some other form of media processing, a key factor in benefiting from the cloud is efficient resource infrastructure use and job scheduling. Many cloud services such as Amazon EC2 have a pay-as-you-go pricing model. By leveraging frameworks that distribute large jobs across multiple virtual machines, and spinning them up and down only as required, media companies can reduce processing time by using larger clusters and, at the same time, lower costs by paying for fewer total hours of infrastructure usage.

Thinkbox Deadline

ThinkBox makes render farm management tools that were built to address complex visual effects software pipelines. Deadline, its oldest and most popular software component, splits up the frames from your visual effects scene and distributes them for processing across multiple virtual machines in a render farm. Deadline can automate computation tasks by running programs and also features a plugin architecture. The host computers that participate in processing jobs can be in geographically diverse locations, including virtual machines running on cloud services like Amazon EC2. ThinkBox's Virtual Machine Extension (VMX) and Cloud Wizard extend Deadline, adding the ability to automatically provision virtual machines in the cloud with pre-installed visual effects software solutions. VMX can automatically scale the number of virtual machines participating in the processing of jobs in the queue. You can also use it track costs and prevent budget overruns by setting a maximum budget for each job. This auto scaling, provisioning, and budgeting framework allows studios to lower costs by limiting the infrastructure required to process a scene to the number of jobs actually in flight.

Telestream Vantage Cloud

Telestream has made their popular transcoding and workflow coordination software, Vantage, available on a subscription basis for AWS. Vantage Cloud Subscriptions allow users to provision Vantage transcoding environments in just a few minutes and host them within an Amazon Virtual Private Network (VPC). The number of virtual machines required for each transcoding job can be scaled up and down automatically, and jobs can be triggered using standard Vantage APIs. Vantage Cloud can also leverage some of the accelerated file transfer solutions mentioned earlier to facilitate moving files to and from the cloud for transcoding. Specifically, Amazon S3 multi-part HTTPS API and Aspera FASP, in addition to more traditional protocols like FTP and SFTP, can be used to transfer files automatically from "watch folders" during transcoding workflows (34).

Amazon Elastic Transcoder

Amazon Elastic Transcoder takes a different approach to pay-as-you-go transcoding workflows in the cloud. Instead of charging by the hour, customers pay by the number of minutes transcoded multiplied by the number of output files produced. It is also a managed service, so resource scaling is handled automatically on the backend with no input required from users or third-party software. Amazon Elastic Transcoder is integrated with Amazon S3 so it can transfer to and from this cloud storage service automatically during transcoding workflows.

LICENSING

Many cloud platforms have a usage-based provisioning and cost model that includes both the software and hardware. For example, Amazon EC2 allows you to deploy virtual machines running Microsoft Windows or various Linux and BSD distributions and pay by the hour for both the infrastructure and associated operating system license. Telestream Vantage and Amazon Elastic Transcoder are examples of an emerging wave of solutions that offer pay-as-you-go pricing models where licensing is bundled into the offer. On the front-end, Adobe Creative Cloud allows users to pay for popular media software such as Adobe Photoshop and Adobe After Effects on a monthly basis.

Many media software solutions rely on license servers to support multi-user environments. Some license managers, such as Flexera's FlexNet Publisher (formerly FlexLM) that are widely used by companies such as Autodesk and The Foundry, are fully supported for use on cloud services such as Amazon EC2 (39). Other license servers such as V-Ray License Server require a hardware dongle attached to the license server hardware, preventing its use with many cloud platforms. Reprise License Manager (RLM), used by Solid Angle's Arnold rendering software and elsewhere, is another example that is limited to only on-premises deployments and Google Cloud Engine. When users want to use media software that is licensed through these and similar license management solutions on platforms like Amazon EC2, it is possible to host the license software remotely on a supported platform. Software VPN solutions such as OpenVPN (37) or IPSec VPN functionality built into virtual networking services such as an Amazon VPC virtual private gateway (38) can allow users to easily establish secure links to remote locations where the license server is hosted.

CONCLUSION

Only a few years ago, the idea of moving high-performance professional media software to the cloud would have been dismissed. This is now possible with today's advanced application streaming protocols, connectivity options, and GPU-enabled virtual machines. Likewise, cloud-based shared storage that can perform on par with top on-premises solutions is now a possibility thanks to enterprise support from the likes of Intel, NetApp, and Avere. The final pieces—making it easy to coordinate complex media processing and cloud-friendly licensing—are making steady progress. Many other top players in the media industry are taking cloud-based media processing very seriously.

Significant challenges still remain, notably the lack of granular, subscription-based licensing for a handful of popular media software. But the shift to granular, subscription-based models in the cloud is well underway. Ultimately, it is media industry customers who will drive this transformation, as they increasingly demand the benefits of the infinitely scalable, cost-effective and increasingly powerful resources available in the cloud for their end-to-end professional media workflows.

REFERENCES

1. Amazon S3 multi-parti API documentation: https://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.html

2. Aspera On Demand for AWS: http://asperasoft.com/fileadmin/media/Asperasoft.com/Product_Briefs/AWS_AoD_AsperaPB.pdf

3. Signiant Flight for Amazon S3, featured on the AWS Big Data Blog: https://blogs.aws.amazon.com/bigdata/post/Tx2QZODBZN1RD1I/Moving-Big-Data-Into-the-Cloud-using-Signiant-Flight

4. Filecatalyst Direct: http://filecatalyst.com/solutions/filecatalyst-direct/

5. Data Expedition ExpeDat Gateway for Amazon S3, featured on the AWS Big Data Blog: https://blogs.aws.amazon.com/bigdata/post/Tx20YXSQ49507II/Moving-Big-Data-Into-The-Cloud-with-ExpeDat-Gateway-for-Amazon-S3

6. TsunamiUDP file transfer benchmarks: http://blogs.aws.amazon.com/bigdata/post/Tx33R88KHCWEOHT/Moving-Big-Data-into-the-Cloud-with-Tsunami-UDP
7. Signiant file transfer protocol benchmarks: http://www.signiant.com/technology/acceleration/
8. Aspera FASP file transfer protocol benchmarks: http://asperasoft.com/technology/transport/fasp/#faspsolution-464
9. Data Expedition MTP file transfer protocol benchmarks: http://www.dataexpedition.com/MTP/performance.html
10. AWS Direct Connect locations: https://aws.amazon.com/directconnect/details
11. AWS Direct Connect partners: https://aws.amazon.com/directconnect/partners/
12. AWS Import/Export: https://aws.amazon.com/importexport
13. Citrix VDI definition: https://www.citrix.com/glossary/vdi.html
14. The Foundry NUKE hardware requirements: https://www.thefoundry.co.uk/products/nuke/sys-reqs
15. NVIDIA GRID Cloud Computing Services: http://www.nvidia.com/object/gpu-cloud-computing-services.html
16. Amazon EC2 virtual machine types: https://aws.amazon.com/ec2/instance-types
17. NVIDIA GRID GPU pass-through: http://www.nvidia.com/object/dedicated-gpus.html
18. Citrix HDX 3D Pro: http://www.citrix.com/products/xendesktop/features/high-def-experience/hdx-3d-pro.html
19. Microsoft RDP 8.1 and RemoteFX explained: http://blogs.msdn.com/b/rds/archive/2012/11/26/remotefx-features-for-windows-8-and-windows-server-2012.aspx
20. Teradici PCoIP: http://www.teradici.com/pcoip-technology
21. VirtualGL: http://www.virtualgl.org
22. TurboVNC: http://www.turbovnc.org
23. Amazon AppStream: https://aws.amazon.com/appstream/
24. Frame: https://fra.me
25. OTOY ORBX: http://home.otoy.com/stream/orbx/
26. Octane Cloud Workstations on AWS Marketplace: https://aws.amazon.com/marketplace/seller-profile?id=795808b7-f99d-426d-bb03-8aa79ff5b65e
27. Citrix CloudBridge Data Sheet: https://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/cloudbridge-technical-overview.pdf
28. Citrix CloudBridge virtual application on the AWS Marketplace: https://aws.amazon.com/marketplace/seller-profile?id=fb9c6078-b60f-47f6-8622-49d5e1d5aca7

29. Amazon EFS: https://aws.amazon.com/efs
30. Running a High Performance SAS Grid Manager on AWS with Intel Lustre: https://blogs.aws.amazon.com/bigdata/post/Tx1DSZMORS09PQE/Running-a-High-Performance-SAS-Grid-Manager-Cluster-on-AWS-with-Intel-Cloud-Edit
31. Intel Lustre: https://www-ssl.intel.com/content/www/us/en/software/intel-solutions-for-lustre-software.html
32. Avere Hybrid Cloud NAS for AWS: http://www.averesystems.com/amazon-web-services
33. Avere products on the AWS marketplace: https://aws.amazon.com/marketplace/seller-profile/ref=sp_mpg_product_vendor?ie=UTF8&id=4da57cb4-798c-4822-86eb-5e7d908731f6
34. Telestream Vantage Cloud: http://www.telestream.net/pdfs/datasheets/dat-Vantage-Cloud.pdf
35. CfnCluster: https://aws.amazon.com/hpc/cfncluster/
36. Intel Lustre documentation: https://www-ssl.intel.com/content/dam/www/public/us/en/documents/reference-architectures/ICEL-RA.pdf
37. OpenVPN: https://openvpn.net/
38. Amazon VPC virtual private gateway VPN: https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html
39. Flexera licensing for Amazon EC2: http://www.flexerasoftware.com/producer/company/news-center/press-releases/FlexNet-Producer-Suite-Enhancements-Ready-App-Producers-Cloud.html

©2015, Amazon Web Services, Inc. or its affiliates. All rights reserved