





Leveraging the Cloud for Hybrid Storage, Processing, Delivery, and Archive


David Green, AWS Enterprise Solutions Architect

June, 2017

Hello, world!







Agenda


- Overview
- AWS M&E Segments
- Media Solution Walkthrough
- AWS Cloud Storage Services
- Storage Workload Alignment
- Storage Benchmarking for Media
- Customer Use Cases
- What's New?
- Looking Ahead



Overview

Media Storage & Utilization Trends on AWS

Content Has Gravity (and is getting heavier...)



4K is four times the size of 1080p HD
4K as the “Gold Standard” resolution
4K screens / programming everywhere

4K acquisition standardizing across camera manufacturers?

Activity to create, capture, and display for devices in 8K by 4K content is occurring - with planned implementations in common media systems by the next decade.

Between 2014 and 2020 we expect about a 4.9x increase in the required digital storage capacity used in the entertainment industry and about a 3.7x increase in storage capacity shipped per year (from ~18,050 PB to ~66,291 PB.)



Ultra HD on **amazon**
instant video



(Coughlin Report)

State of the Union (as told by our customers)

Storage Marketplace Maturity


- Asset Management systems
- Pay-as-you-go licensing models
- Integration to cloud-based storage options
- Tape ingest

Distance vs. Content Gravity

- Movement of content
- Global network connectivity
- Running applications remotely
- Content replication/sync across sites

Cost

- Do I always need the durability cloud storage provides?
- Storage price reductions benefits




16 Regions – 42 Availability Zones – 68 Edge Locations


AWS M&E Segments

A Quick Overview of M&E on AWS


AWS Media Segments



AWS Media Segments (w/ service mappings)




AWS M&E Partner Ecosystem



Media Solution Walkthrough


Designing an end to end solution for Content Production & Distribution

A Sample Production Pipeline




Content production and post-production companies are leveraging AWS to accelerate and streamline creative, editing, and rendering workloads with highly scalable cloud computing and storage.

Estimating Storage Capacity is a Challenge




Content Production Systems must deal with multiple Asset Types, Sizes & Access Patterns

Asset Lifecycle Management is Critical



Compute Needs Often Drive Storage Requirements


On-Demand



Pay for compute capacity by the hour with no long-term commitments

For spiky workloads, or to define needs


Reserved



Make a low, one-time payment and receive a significant discount on the hourly charge

For committed utilization

Spot



Bid for unused capacity at a Spot Price - fluctuates based on supply & demand

For time-insensitive or transient workloads


EBS provides flexibility to choose between different models by decoupling storage from compute

Securing the Assets


- *AWS alignment* with the latest **MPAA** cloud based application guidelines
- **VPC private endpoints** for S3 - enables a true private workflow capability
- **Encryption & Key Management** capabilities
- **Glacier Vault** for high-value media/originals
- Fine-grained Identity & Access Management
- **Non-host based** Security Groups & NACLs
- Comprehensive **auditing** via AWS Flow Logs, Config, Inspector & CloudTrail
- Run workloads using **IAM Roles**, not embedded credentials
- Use **AWS KMS** Key Management / **HSM**
- **EBS Encryption** for boot, data and snapshotted volumes




Production Archive



End to End Production Lifecycle




A Typical OTT Service Use Case



AWS Cloud Storage Services

A deeper dive into EBS, EFS and S3, and associated hybrid use cases

EC2 Instance Store




- Local to instance
- Non-persistent data store
- Data not replicated (by default)
- No snapshot/encryption support
- SSD or HDD

Use Cases

- Working set cache (**Rendering**)
- Reverse Proxy Cache (**OTT**)
- Transcoding cache (**Supply Chain, OTT**)

*Applications that require continuity may move temporary files to S3 to enable resuming aborted workflows (**Spot Rendering**)*

EC2 Instance Store vs. EBS Offerings




EBS Volume Types

SSD Based

gp2

General Purpose SSD




Baseline: 3IOPS per GB up to 10,000
Burst: 3000 IOPS (for volumes up to 1TB)
Throughput: 160 MB/s
Latency: Single digit ms
Capacity: 1GB to 16TB

Usage

Bursty workloads (Rendering)
Low latency operations (Streaming)

io1

Provisioned IOPS SSD



Baseline: 100IOPS up to 20,000
Throughput: 320 MB/s
Latency: Single digit ms
Capacity: 4GB to 16TB


Usage

Sustained IOPS
Critical Applications with IOPS
Guarantee (no burst)

HDD Based

st1

Throughput Optimized HDD




Baseline: 40MB/s per TB up to 500MB/s
Burst: 250MB/s per TB up to 500MB/s
Capacity: 500GB to 16TB

Usage

Large Block, high throughput sequential
Transcoding, QC, Packaging
Rendering

sc1

Cold HDD




Baseline: 12MB/s per TB up to 192MB/s
Burst: 80MB/s per TB up to 250MB/s
Capacity: 500GB to 16TB

Usage


Sequential
Logging, Active Archive/Backup

Choosing an EBS Volume Type




Amazon Elastic File System (EFS)

- Designed to support **Petabyte** scale file systems
- **Throughput scales linearly** to storage
- Same latency spec across each AZ
- **Thousands** of concurrent NFS connections
- Works great for **Large I/O sizes**
- Pay for only what you use not what you provision
- Managed with **multi-copy durability**



Data Analytics, Media Processing (Transcoding, rendering, QC, Farms), Distributed Content Management, Web Serving/Publishing, Digital Supply Chain


Amazon S3 Object Store



- **No Filesystem** - everything is GET/PUT/DELETE
- **Requests Per Second (>300)** - avoid hotspots, use front-end buffering
- **Puts vs. Gets** - use Cloudfront for GET heavy workloads
- **Network Throughput** - multi-part uploads, AWS SDK/CLI **Consistency/Commit Latency** - Read after write, factor in REST latency for workload requirements

Active Archive and lifecycle to Cold Archive, Media Processing (Transcoding, rendering, QC, Farms) depending on Application conformity, Mezzanine Storage, VOD Delivery.

Hybrid Storage



Storage Gateway



Amazon
CloudWatch



AWS Identity and Access
Management (IAM)




AWS CloudTrail



AWS Key Management Service
(KMS)

Storage Services (Summarized)



Encryption




Storage Workload Alignment

Usage Patterns for Storage, Aligned with M&E Workloads

Media Workload Mapping




Content Acquisition



Digital Asset Management




Visual Effects & Editing




Playout & Distribution




Publishing



OTT



Media Supply Chain



Analytics

Media Workload Storage-Type Affinity

WORKLOAD	BLOCK	FILE	OBJECT
CONTENT ACQUISITION			
DIGITAL ASSET MANAGEMENT			
VISUAL EFFECTS AND EDITING			
MEDIA SUPPLY CHAIN			
PLAYOUT AND DISTRIBUTION			
OTT			
DIGITAL PUBLISHING			
ANALYTICS			


Storage Benchmarking for Media

A sample Transcode Workflow illustrating Optimization for Cost

Analyzing & Optimizing the Application Layer


1

Known File Access Pattern





2

Parallel Processing Optimization






- Optimizing the application layer's I/O patterns are critical - we can use this to drive down cost
- Analyze the Access Pattern; Optimize the Processing
- Let's benchmark a **common transcoder use case** - 4K, parallel encode @ ~10 bitrates, with HLS packaging
- We will **optimize the transcode** process by only reading the source media into a memory buffer once



Storage Benchmarking for Cost Optimization



GP2 - 1TB EBS Volume @3000 IOPS



ST1 - 1TB EBS Volume



Customer Use Cases

PIX | System & Theory Studios

Customer: Sony DADC

VEN.UE

On-demand cloud-based media supply chain and delivery solution

We have 20 petabytes of content on AWS, the equivalent of more than 800,000 hours of video, available on our platform. We can only move all that content around the world with the scalability we're getting on the AWS Cloud.

~Andy Shenkler, Chief Solutions and Technology Officer

Your content.
Your audience.
Your way.

Problem Statement:

- Provide a performant, secure, and economic media distribution solution
- Challenged by on-prem legacy infrastructure
- Decrease time to market for their customer's finished content

Use of AWS:

- EC2 content processing and SWF, SQS, SNS for media workflow automation
- S3 for storage, Glacier for content archive
- CloudFront CDN for OTT playout

Business Benefits:

- Workflow pipelines can be run in a highly parallelized fashion through AWS elastic scalability.
- Significantly shorten their content delivery SLA with a new AWS enabled target of 1-hr.
- Fully migrating away from on-prem infrastructure.





PIX | SYSTEM

LOS ANGELES | NEW YORK | SAN FRANCISCO | LONDON



24/7 live **global collaboration and asset sharing**. Unrivaled media mining. Industry defining security. And the hardest working support team in show business.

PIX is a trusted partner to the entertainment industry with more than ten years of dedicated service, across **thousands of film and television productions, supporting every major studio**. Our dedicated team of industry experts take pride in helping clients **successfully and securely manage content review and collaboration throughout the creative lifecycle**.





General Flow of Work

- **Users** upload video files ranging from a few MB to multiple GBs.
- **Ingest Servers** encrypt content and store on EBS-based backend filers.
- **Batch Processors** transcode media to multiple formats, add watermarking and perform additional custom workflows.
- **Final Copies** for consumption land on EBS-based filers for serving the content through front-end Web and native OS applications.

Why AWS? **Price.** Ability to handle **large files.** **Easy migration** from on-prem. **High IO with High availability**, in a bursty environment.



Design & Performance

- **Filers** - c4.2xlarge, EBS optimized, Enhanced Networking. Range in size from 48 to 96TB, shared over NFS to clients.
- **Volumes** – 8192GB SC1 EBS for max IO
- **Storage Subsystem** – 2x8192GB Volumes RAID-0 with mdadm software raid, striped with LVM to appear as a single OS volume.
- **Filesystem** – XFS, 1024k block size to optimize for SC1
- **Performance** – NFSv4 provides optimum performance for attached worker nodes. Under heavy R/W IO, filers maintain 1Gbps asymmetrically, with bursts to 1.5Gbps.
- **Peak IO** - *`we have seen above 2.5Gbps in and out of our filer infrastructure!`*



Worldwide Virtual Studio
Original Animation, Visual Effects and Virtual Reality



Top-Rated Virtual Reality
Arcade Game




Collaborated with Barnstorm on
Visual Effects for Season 2



Visual Effects &
Animation for TV show




General Flow of Work



- **Creative work** done on-site with team spread across the globe. Artists work on local workstations.
- **VFX processing** done on AWS utilizing EC2 Spot-based Blender renderfarm.
- Artists **push render jobs** using custom sync tools. These move the content to the rendering service running on AWS.
- **Typical render workload** is ~3TB - consisting of Blender packed assets (in .blend project file) and external assets (meshes, textures, etc.)

Why AWS? *Price. Easy setup, integration and deployment. EC2 Spot History & Spot Pricing.*

Theory Studios | Rendering Pipeline




- **EC2-based Filers** utilizing **ST1 EBS** volumes provide centralized storage.
- **Original Dataset + a Mirrored Dataset** help overcome **500 MB/sec** throughput to multiple machines.
- **Final Render Data** is output to another Instance which syncs frames to the client (customer) site.
- **Challenge** - waiting for **250 Spot Instances** to get data.

What's New?

EBS Elastic Volumes

New! EBS Elastic Volumes




- Dynamically increase capacity, tune performance, and change the type of live EBS volumes.
- Simple; Flexible; Non-disruptive; Automated.
- Modify the configuration of live volumes attached to instances.
- **No downtime, no performance impact.**
- You can automate changes using CloudWatch with Lambda or CloudFormation.
- **No need to plan ahead.** Provision what you need today and change the configuration as business needs change.

Elastic Volumes - Process & Limitations

- Process - Three Steps

1. Issue the modification command
2. Monitor the progress of the modification
3. If increased size, **extend** the volume's file system



- Limits

1. You can modify a volume **once per 6 hours**
2. Supported only for current-generation volumes (**gp2/io1/st1/sc1**).
Not supported for Magnetic/Standard volumes
3. Live changes supported for volumes attached to **current generation instances**

Looking Ahead

November, 2015 - 50TB Snowball

April, 2016 - 80TB Snowball



November, 2016 - 100TB Snowball Edge




November, 2016 - 100PB Snowmobile



I





Thank You!



- Digital Media on Amazon Web Services
<https://aws.amazon.com/digital-media>
- Sony DADC NMS Case Study
<https://aws.amazon.com/solutions/case-studies/sony-dadc>
- EBS Benchmarking
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/benchmark_procedures.html
- FFMPEG multi-output encoding
<https://trac.ffmpeg.org/wiki/Creating%20multiple%20outputs>
- PIX|SYSTEM
<https://www.pixsystem.com>
- Theory Studios
<http://theorystudios.com>