

adaptTo()

APACHE SLING & FRIENDS TECH MEETUP
BERLIN, 28-30 SEPTEMBER 2015

Adobe Managed Services Complicated Cloud Deployments

Adam Pazik / Mike Tilburg – Adobe Systems



Who are Adobe Managed Services? Backup and Restore Key AWS Learnings Complicated Clouds

Adobe Managed Services - AEM



150+

Enterprise grade,
global
AEM hosted
customers



30+

CloudOps and AEM
SysAdmin
expert engineers



Basic

Mid-Market model launched

1500+

AEM instances
running
globally


500tb+

Storage

24/7

Global, follow the
sun support model

WE TOOK THE HOSTAGES,
SECURED THE BUILDING, AND
CUT THE COMMUNICATION
LINES LIKE YOU SAID.



EXCELLENT.

BUT THEN THIS GUY CLIMBED UP
THE VENTILATION DUCTS AND WALKED
ACROSS BROKEN GLASS, KILLING
ANYONE WE SENT TO STOP HIM.



AND HE RESCUED
THE HOSTAGES?

NO, HE IGNORED THEM.
HE JUST RECONNECTED
THE CABLES WE CUT,
MUTTERING SOMETHING
ABOUT "UPTIME".



SHIT, WE'RE
DEALING WITH
A *SYSADMIN*.

Just some of our enterprise customers...



Westfield



Anheuser-Busch

bp



abbvie



CONDÉ NAST



MOTOROLA

Coca-Cola

RBS
The Royal Bank of Scotland



Backup and Restore Methodology

At Adobe Managed Services

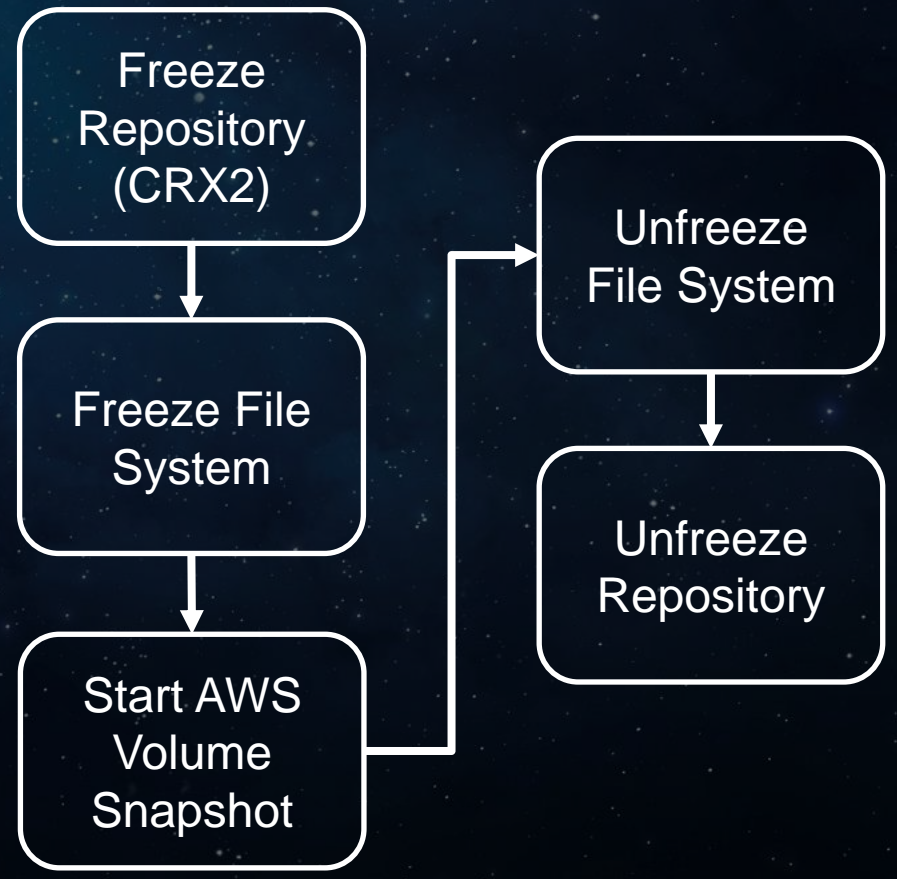
AWS Snapshot Technology

Fast

Incremental

Asynchronous

High Level Backup Workflow



Block Repository Writes (CRX2 Only)

JMX Console

IP:Port/system/console/jmx/com.adobe.granite%3Atype%3DRepository

java.lang.Void	<u>blockRepositoryWrites()</u> Blocks all repository writes
java.lang.Void	<u>unblockRepositoryWrites()</u> Unblocks all repository writes

Using granite-mbeans-cmdline-1.0.jar

Wrapper to call the operation blockRepositoryWrites/unblockRepositoryWrites on com.adobe.granite:type=Repository

Verify integrity of lucene index

Fast

Swap volume from last snapshot

Automated through our cloud controller

EBS Pre-warming is a consideration

Key AWS Learnings

At Adobe Managed Services

AWS Learnings for AEM

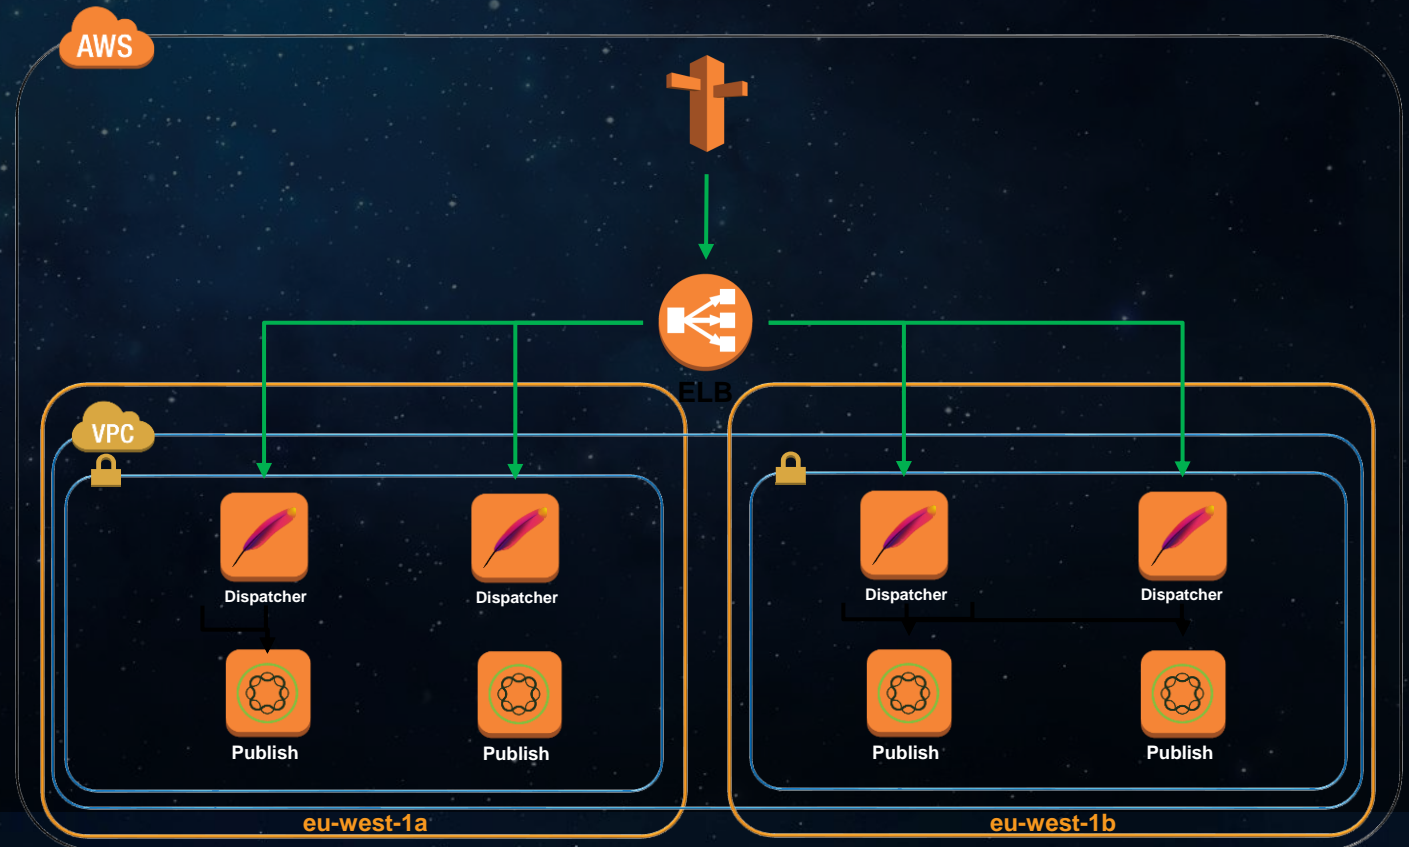
- Use EBS Optimized instances
 - Greater throughput to storage
- Watch out for EBS pre-warming
 - 5-50% degradation of IOPS
- Utilise the fast ephemeral storage where possible
 - Java temp dir, other processing (ImageMagick etc)
- Correctly scale your ELB (load balancer) with AWS
- Use VPC S3 EndPoint when using S3 DataStore

Complicated Cloud #1

Rapid Traffic Spikes

Complex customer one....Rapid Traffic Spikes

- “Must” move to cloud
- Traffic data incomplete from previous provider
- Generally static site
- Very adverse to CDN due to historical issues with old provider



www.client.com

AWS



Route53
DNS



ELB

VPC



Publish



Publish



Publish

10.22.1.0/25

eu-west-1a



Publish



Publish



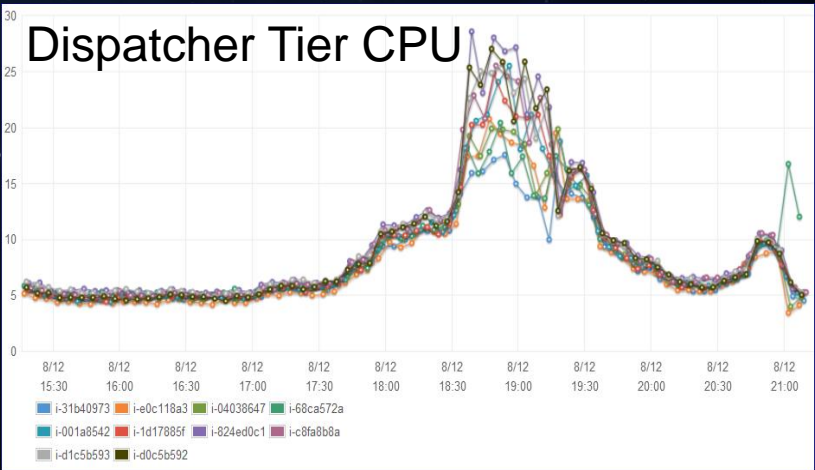


Publish

10.22.1.128/25

eu-west-1b

Publis

Complex customer one...Rapid Traffic Spikes



Complex customer one - Remediation

Analyse traffic patterns (splunk/goaccess)

Dispatcher cache serving a lot of traffic but...

Just too much traffic

High client context calls hitting publish tier

Scale dispatcher tier based on ELB RPM

Globally distributed client base

Cache client context calls

Deploy CloudFront CDN....

...but not without challenges

Complex customer one – CDN Pain Points

- Entire cache must be purged on new code release....

...but entire cache purge is not possible with CloudFront

UPDATE: Full cache purge now supported!

...and individual object invalidation is SLOW

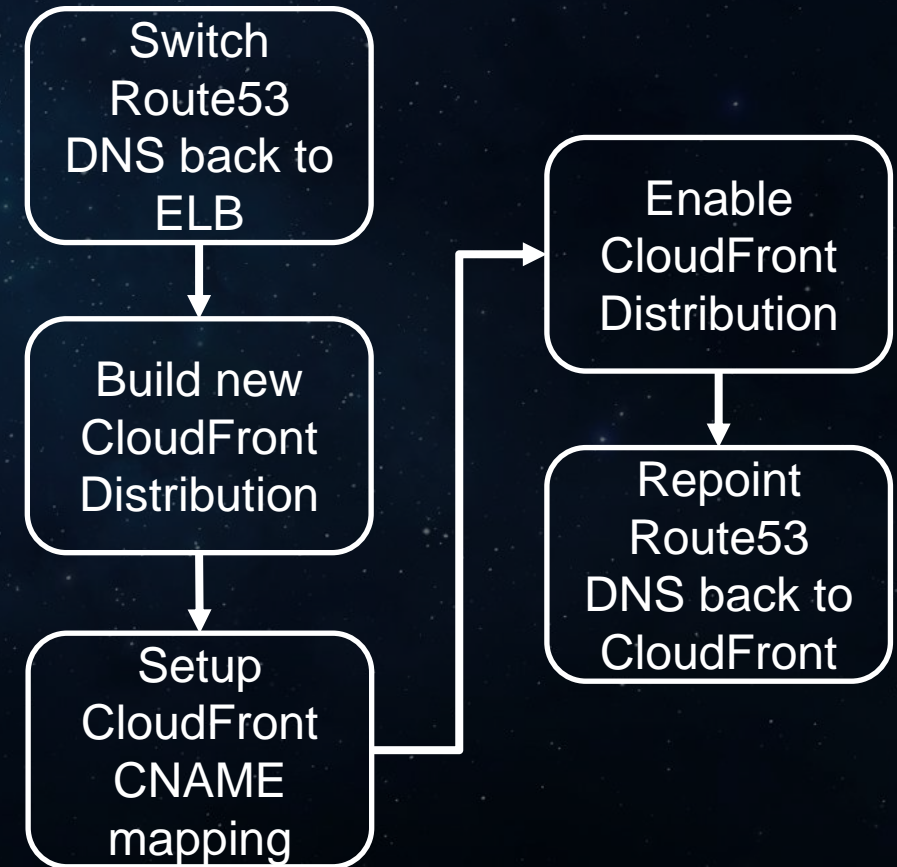
Solution: Rebuild the CloudFront distribution
Fresh cache within **~20 minutes**

- Dynamic content should be served from cache instantly


Solution: Implement URL fingerprinting to ensure unique content is cached

Very low expiry header for html pages

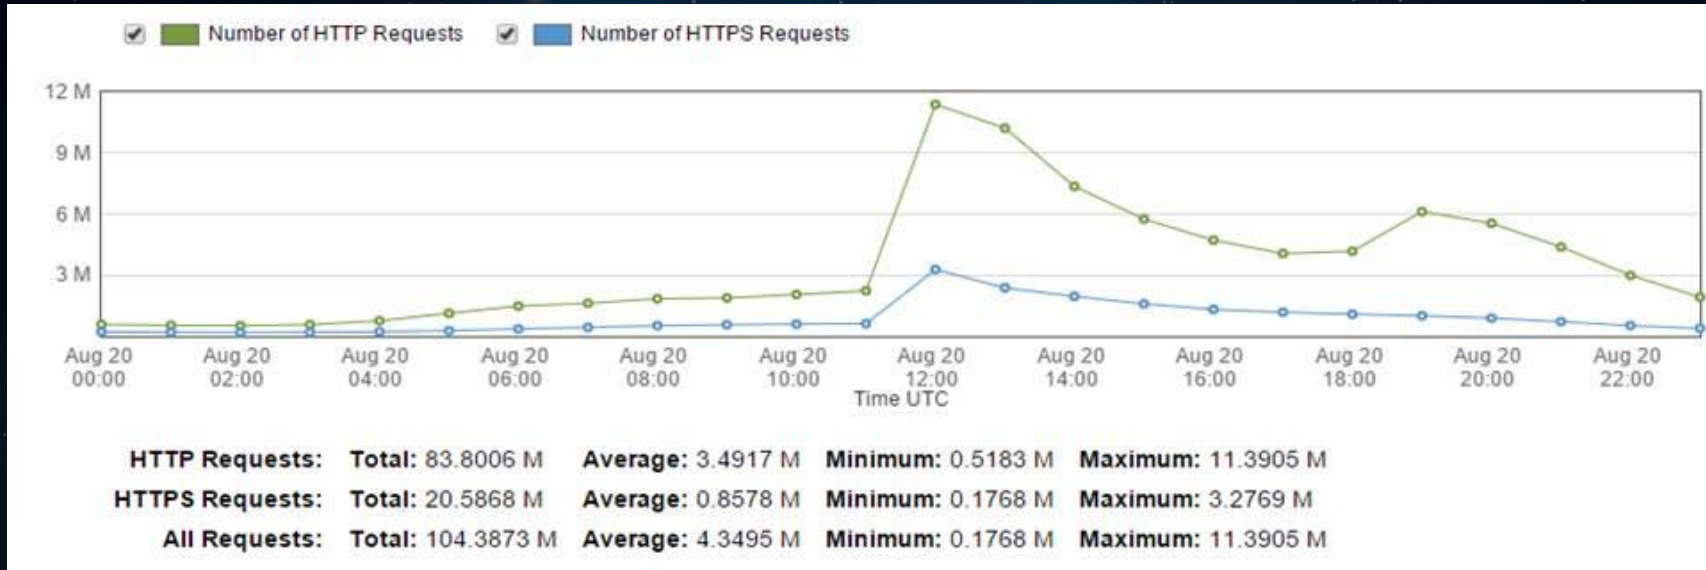
CloudFront rebuild process (aws cli/chef)




Complex customer one – Result



Complex customer one – Recent traffic spike



Zooming in on the spike... ~48k requests per minute at 11:59 (GMT+1) to a peak of ~335k requests per minute at 12:11.



Complicated Cloud #2

Global DAM

Complex customer two...the Global DAM

CONDÉ NAST

- A cloud-based DAM solution for CNI's 85 magazine titles in 11 markets across the Americas, Europe and Asia.
 - Manage assets for publication across multiple channels (print, tablet, web)
 - Facilitate reuse of assets between brands globally through a controlled process
 - Ensure assets are not published without commercial rights having been agreed.
 - Improve searchability by enforcing metadata standards, and linking assets to published pages.

Complex customer two...the Global DAM

CONDÉ NAST

- A cloud-based DAM solution for CNI's 85 magazine titles in 11 markets across the Americas, Europe and Asia.
 - Manage assets for publication across multiple channels (print, tablet, web)
 - Facilitate reuse of assets between brands globally through a controlled process
 - Ensure assets are not published without commercial rights having been agreed.
 - Improve searchability by enforcing metadata standards, and linking assets to published pages.
- In addition to archiving and search, CNI wanted the DAM to support WIP
 - Managing assets from initial upload by photographers, through selection and layout, delivery to external repro houses for retouch, publication and archive.

Complex customer two...the Global DAM - Challenges

SCALABILITY

Each title publishes many thousands of assets each year, the DAM also holds the 'overs' which weren't published.

Assets can be large – several hundreds of MB for a single image

Repository (Data Store) is going to become huge over time.

Complex customer two...the Global DAM - Challenges

SCALABILITY

Each title publishes many thousands of assets each year, the DAM also holds the 'overs' which weren't published.

Assets can be large – several hundreds of MB for a single image

Repository (Data Store) is going to become huge over time.

RESPONSIVENESS

The DAM manages assets throughout the production of the magazine & must be responsive even with large assets.

All markets are both contributing to the DAM and consuming assets from it – single 'global' DAM

....but there is no single AWS location which is appropriate for all CNI markets.

Complex customer two...the Global DAM - The (current) Solution

S3 DATA STORE CONNECTOR

Scalable, cost-effective storage

Enables 'shared' Data Store topologies in AWS – binaryless replication

Accessed over HTTP(s), so very different performance to local, SAN or NFS Data Store



Complex customer two...the Global DAM - The (current) Solution

S3 DATA STORE CONNECTOR

Scalable, cost-effective storage

Enables 'shared' Data Store topologies in AWS – binaryless replication

Accessed over HTTP(s), so very different performance to local, SAN or NFS Data Store



Local file cache

Downloaded binary records are cached as files on local disk in a LRU cache.

Complex customer two...the Global DAM - The (current) Solution

S3 DATA STORE CONNECTOR

Scalable, cost-effective storage

Enables 'shared' Data Store topologies in AWS – binaryless replication

Accessed over HTTP(s), so very different performance to local, SAN or NFS Data Store



Local file cache

Downloaded binary records are cached as files on local disk in a LRU cache.

Async Upload to S3

New binaries are added to local file cache, and asynchronously uploaded to S3 later

Complex customer two...the Global DAM - The (current) Solution

S3 DATA STORE CONNECTOR

Scalable, cost-effective storage

Enables 'shared' Data Store topologies in AWS – binaryless replication

Accessed over HTTP(s), so very different performance to local, SAN or NFS Data Store



Local file cache

Downloaded binary records are cached as files on local disk in a LRU cache.

Async Upload to S3

New binaries are added to local file cache, and asynchronously uploaded to S3 later

Proactive Caching ^{1.4+}

Checking size of record triggers the background download of binary from S3.

Complex customer two...the Global DAM - The (current) Solution

S3 DATA STORE CONNECTOR

Scalable, cost-effective storage

Enables 'shared' Data Store topologies in AWS – binaryless replication

Accessed over HTTP(s), so very different performance to local, SAN or NFS Data Store



Local file cache

Downloaded binary records are cached as files on local disk in a LRU cache.

Async Upload to S3

New binaries are added to local file cache, and asynchronously uploaded to S3 later

Proactive Caching ^{1.4+}

Checking size of record triggers the background download of binary from S3.

Rec Length Cache ^{1.5+}

In-memory cache of record lengths to prevent reliance on local file cache.

The (current) Solution

GLOBAL DAM ARCHITECTURE

A global architecture consisting of separate author clusters in EMEA and APAC to minimize latency between users and the DAM.

Completely transparent to the end user – system feels like a ‘single DAM’

Complex customer two...the Global DAM - The (current) Solution

GLOBAL DAM ARCHITECTURE

A global architecture consisting of separate author clusters in EMEA and APAC to minimize latency between users and the DAM.

Assets are initially uploaded to the market's 'primary DAM'



Complex customer two...the Global DAM - The (current) Solution

GLOBAL DAM ARCHITECTURE

A global architecture consisting of separate author clusters in EMEA and APAC to minimize latency between users and the DAM.

On publication, assets are replicated to 'secondary DAM(s)



Complex customer two...the Global DAM - The (current) Solution

GLOBAL DAM ARCHITECTURE

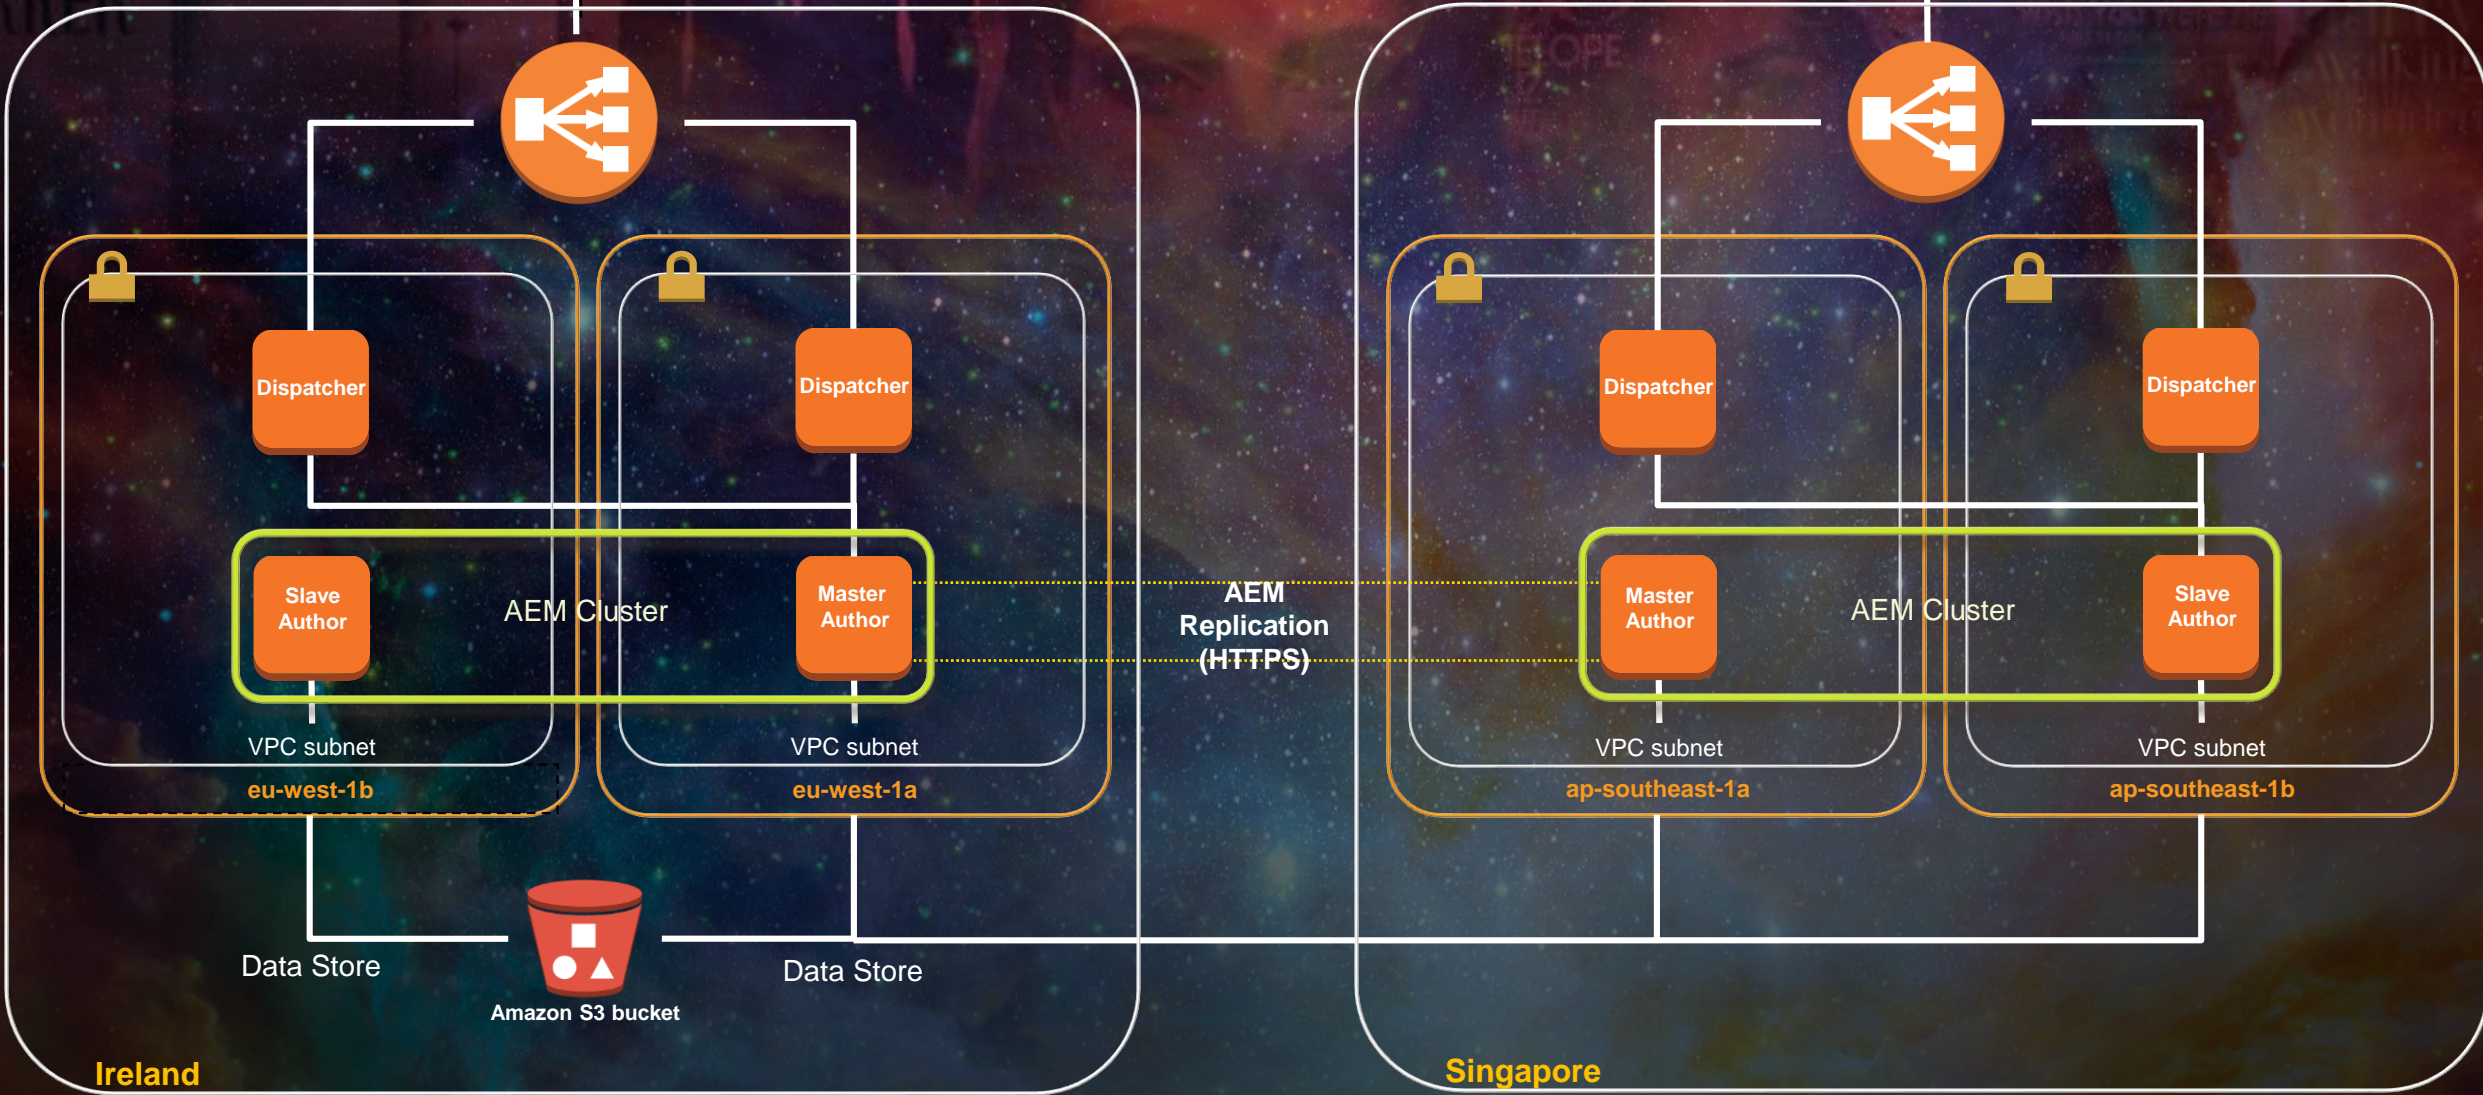
A global architecture consisting of separate author clusters in EMEA and APAC to minimize latency between users and the DAM.

All modification of assets takes place on the 'primary DAM', and asset is re-distributed



Complex customer two...the Global DAM - The (current) Solution

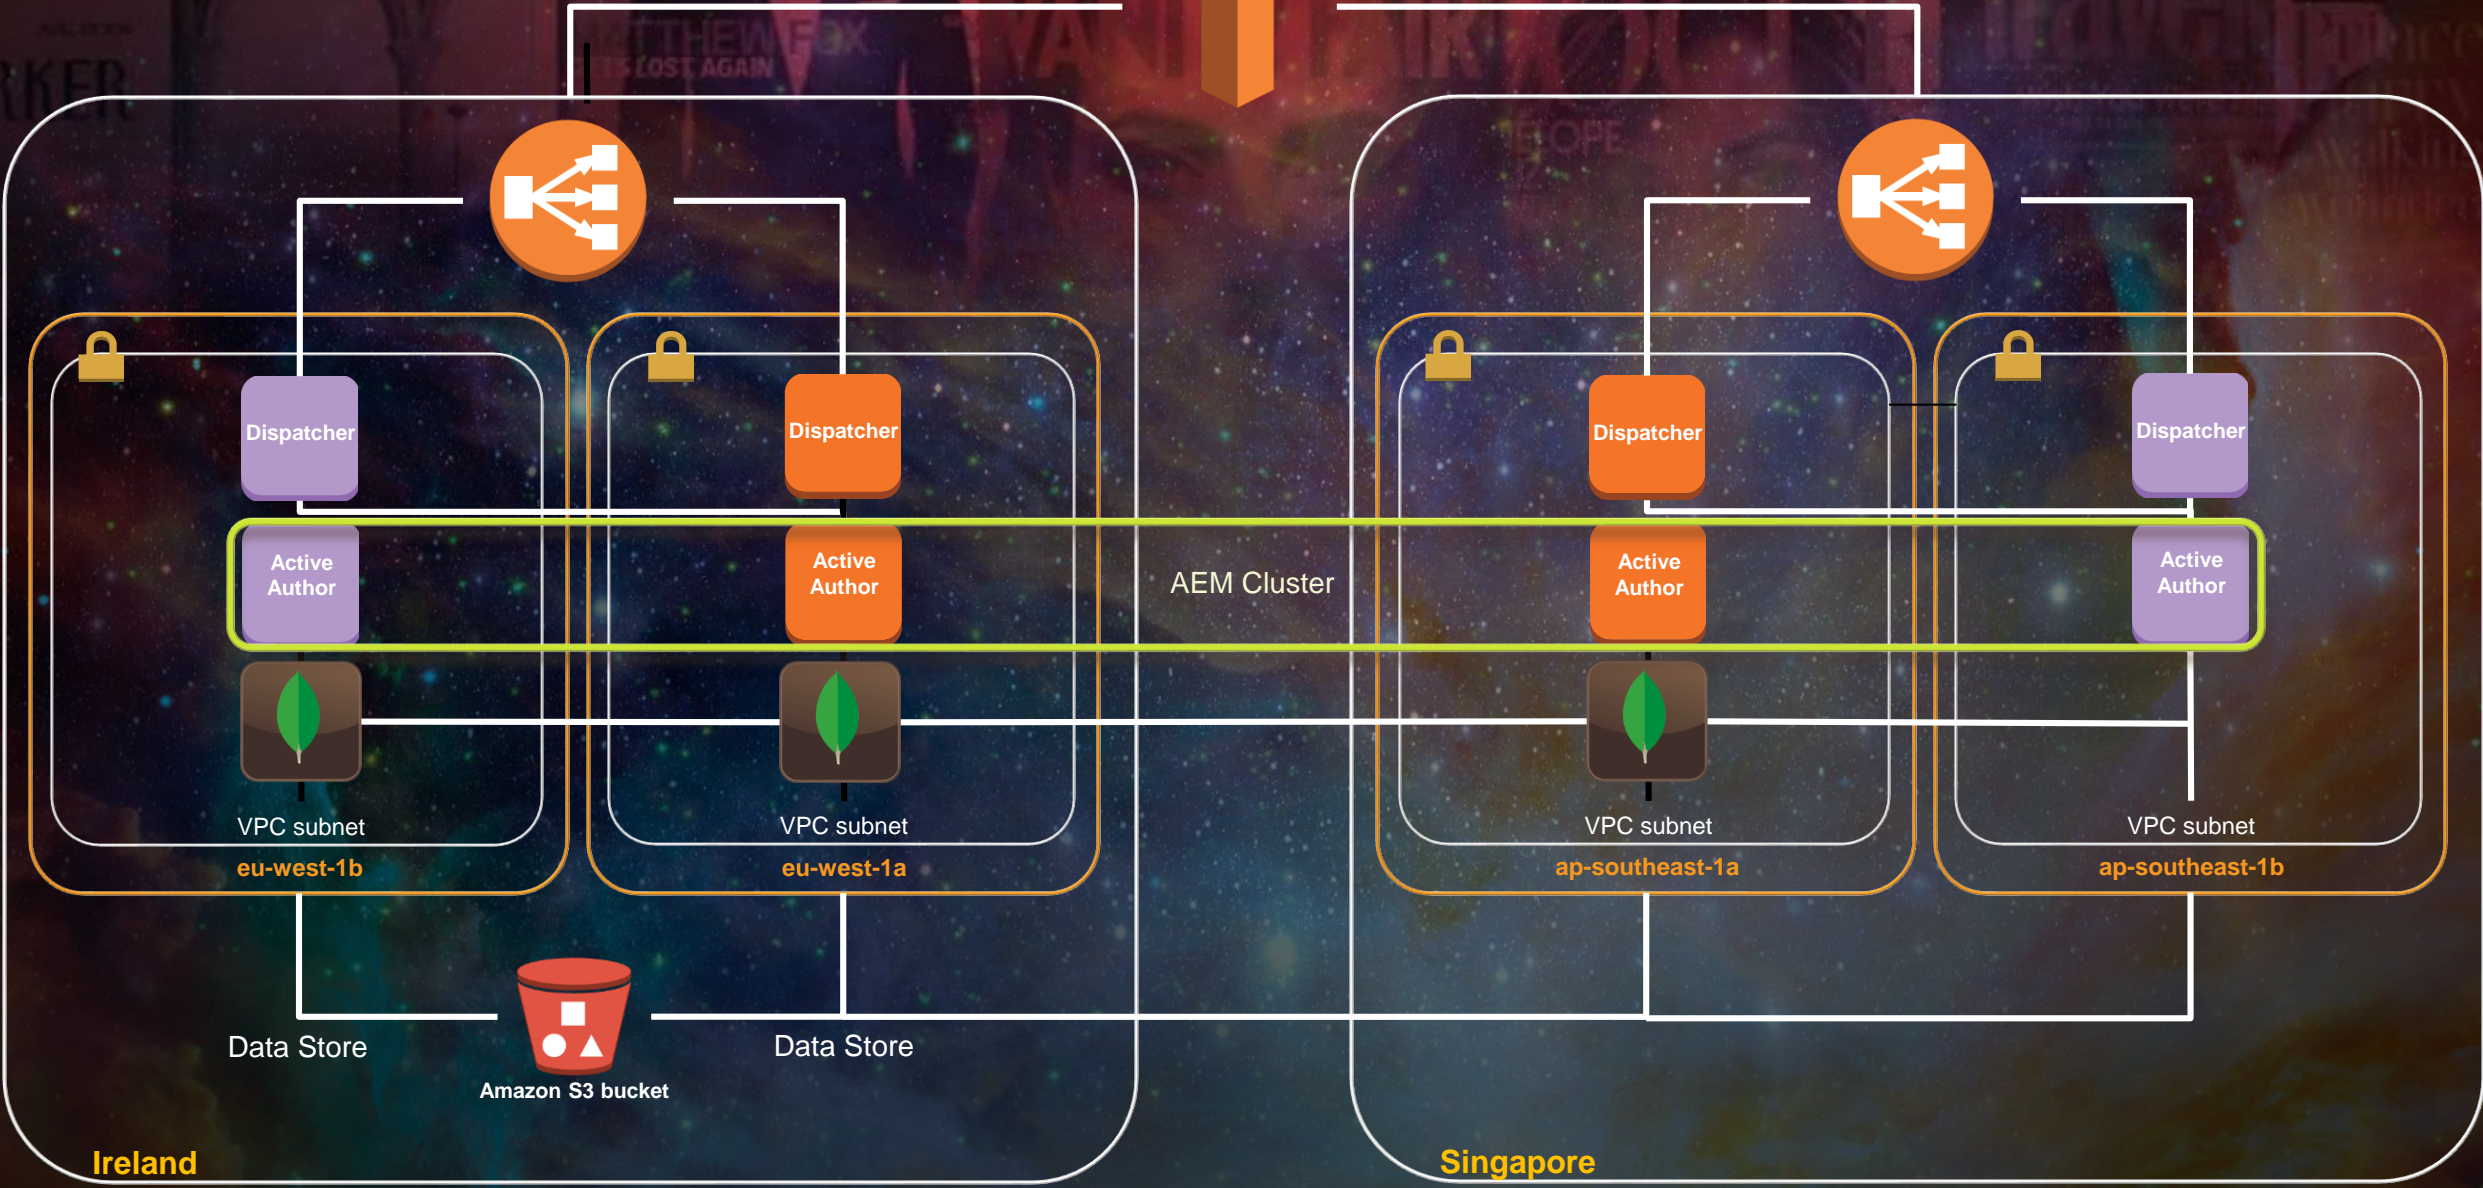
Neustar UltraDNS



Ireland

Singapore

Complex Customer Two - The Future





Who are Adobe Managed Services? Complicated Clouds Backup and Restore Key AWS Learnings

Thank You


For more information contact:

Adam Pazik

EMEA Team Lead, Adobe Managed Services

T +44 7810 636 994

pazik@adobe.com



Adobe

AEM 6.0 learnings

- Keep upgrading Oak
 - 1.0.15 released – **hotfix 6561**
 - Contains important compaction and cold standby fixes
- Run Offline tar compaction at least monthly
 - Biggest impact on authors
 - 2gb segment store (nodes/properties) reduced to ~200mb
 - Requires free space to run
 - Automate
- Automate BlobGC to defrag datastore

Importance of Java temp directory

Upload 400 (1 MB) JPG images to AEM 6.1 DAM					
		Bytes			
	Folder Location	READ	WRITE	TOTAL	% Share
1	/crx-quickstart/segmentstore	637,906,188	1,767,242,540	2,405,148,728	38.0%
2	/crx-quickstart/index	-	50,438,060	50,438,060	0.8%
3	java.io.tmpdir	2,981,566,566	881,919,725	3,863,486,291	61.1%
4	/crx-quickstart/logs	294,912	5,530,452	5,825,364	0.1%
		3,619,767,666	2,705,130,777	6,324,898,443	

400 1mb jpegs uploaded
61% utilization at java temp dir

Always put java temp dir on fast storage!
AWS = ephemeral