

Does object storage actually fit into file-based workflows?

Cloud Computing West

Michael King

Director of Product Marketing

Why is Big Data such a Big Problem?



The Economist, February 27th - March 5, 2010

Because we are storing it the wrong way

How much of the data you store never changes?



50%, 60%, 70%, 80%, more?

File & Object Storage

Different, but Perfect Together



**File
Storage**

vs.

**Object
Storage**

- Millions of files
- Locking mechanisms
- Amendable data
- Management complexity
- Not designed for small I/O
- File system hierarchy
- Complex to scale
- TCO increases exponentially
- Great for high bandwidth, transcoding, and editing

- Trillions of objects
- No locking mechanisms
- Immutable data
- Automated, self-healing
- Highly optimized for large & small
- One storage pool, object ID's
- Scales easily, uniformly
- Lowest TCO
- Great for video streaming and content distribution

Object Storage Features

- ▶ No File System, No Hierarchy
- ▶ 99.9% disk utilization
- ▶ Simple - Get, Put, Delete
- ▶ Customizable Metadata
- ▶ Immutable Content
- ▶ Lot of Data Protection Options
- ▶ Easily Geographically Distributed
- ▶ Highly Scalable
- ▶ Easy to Manage
- ▶ Low TCO



Object storage does not allow object/files to be modified. If changes are made new objects are created.

7 Factors for Evaluation


Does object storage fit into your file-based workflow?

- ▶ Production Storage Performance Requirements
- ▶ HSM Abilities
- ▶ Archiving Practices
- ▶ Ability to Use/Monetize Long tail content
- ▶ Data Protection Requirements
- ▶ Distribution Requirements
- ▶ Collaboration & QA Environment



Use Cases

- ▶ Collaboration
- ▶ Tier One Relief
- ▶ Active Archive
- ▶ Tape Cache
- ▶ Online Content Origin
- ▶ Cloud Back-up & DR
- ▶ Application Storage



The Ultimate Solution

- ▶ Optimized, “smaller” production tier
- ▶ Automated, policy based migration to Active Archive Object Storage
- ▶ Very high performance and the lowest TCO

REAL-TIME, HIGH PERFORMANCE





-  Ingest
-  Edit
-  Transcode
-  Render



AUTOMATED
MIGRATION

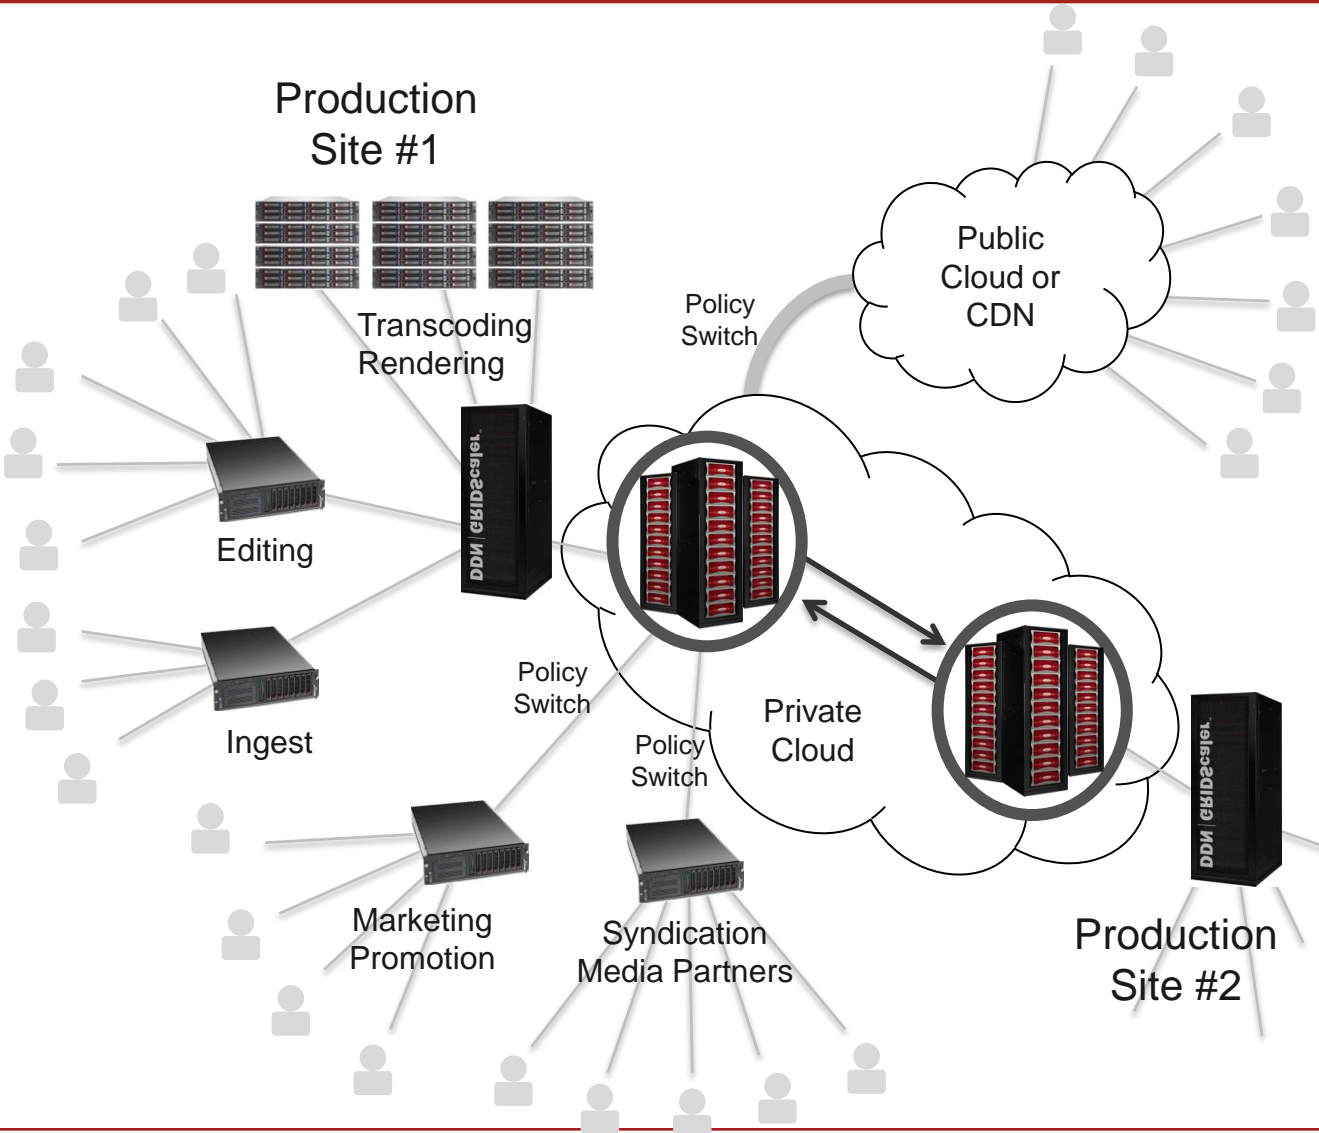


CLOUD & OBJECT PLATFORM

-  Collaborate
-  Distribute
-  Archive
-  Protect



Collaborate, Distribute and Federate




Overview

- ▶ Federation of production namespaces, using object storage as the wide area distribution mechanism
- ▶ Multi-site - ingest, production, distribution, collaboration, & archive
- ▶ Access through production or directly from object storage
- ▶ Availability and protection controlled by policy
- ▶ DR for Free

Post Production Example

- ▶ Large Multi-national Creative Services Group
- ▶ Building an active archive that will allow them to easily collaborate between LA, Vancouver and London
 - Free up expensive Tier-1 Storage
 - Consolidate Tier-2 pools across many or all facilities
 - Increase utilization and efficiencies of geo-dispersed compute & creative personnel resources
 - Get DR “for free”
 - Much easier scaling path over time
 - Improve cost-competitiveness and facility attractiveness



Conclusion

- ▶ Object storage does have a place in file-based workflows
- ▶ Lowest cost, disk based solution
- ▶ Use where content is being stored, retrieved and used but not changing frequently
- ▶ Rethink storage tiering
- ▶ Shift our thinking to incorporate it where it really fits well
 - Active Archive
 - Collaboration
 - Content Distribution



DDN is a Leader in Massively Scalable Platforms and Solutions for Big Data and Cloud Applications

- ▶ **Established:** 1998
- ▶ **Revenue:** \$250M+ – Profitable, Fast Growth
- ▶ **Main Office:** Sunnyvale, California, USA
- ▶ **Employees:** ~550 Worldwide
- ▶ **Worldwide Presence:** 20 Countries
- ▶ **Installed Base:** 1,000+ End Customers; 50+ Countries
- ▶ **Go To Market:** Global Partners, Resellers, Direct



World-Renowned & Award-Winning



World's Fastest & Most Efficient Object Storage Solution

End-to-End Object Store



- DDN's Web Object Scaler (WOS)
- Ground-Up Approach To Resolve Key Scalability, Efficiency, and Cost Challenges
- High Performance, Exascale Solution
- Self-Healing, Highly-Resilient Scale-Out Cloud Storage Software
- True No File System Solution
- Full Enterprise Support
- Lowest TCO of Any Solution



Thank You