



When Data Pools Become Data Lakes: Managing and Storing Massive Quantities of Data

Lance Broell
DDN Storage



Data Growth is Real and Accelerating

➤ Unstructured data is the primary growth driver

- ◆ Digital images and video
- ◆ Electronic documents
- ◆ Sensor data

➤ Impact of unstructured data


- ◆ Increased reliance on file storage pools
- ◆ Growing storage islands
- ◆ Increased storage administration



Growth Drives Complexity and Cost

➤ File storage systems at scale

- ◆ Unpredictable performance
- ◆ Increased complexity and storage administration
- ◆ High Storage TCO



Storage Pools Beyond The Limits

It's Time For A Change!



A New Approach – Object Storage

Data Lakes - Ready For Business



Object Storage Enables the Data Lake

➤ Infinitely scalable storage pool

- ◆ Trillions of objects
- ◆ Exascale capacity
- ◆ 10 minute scale out

➤ Predictable performance

- ◆ Scales linearly with capacity
- ◆ Unaffected by the number or size of objects



Object Storage Next-gen Data Protection

- Unmatched data durability
 - ◆ Replication
 - ◆ Erasure coding
 - ◆ Fast data rebuilds
- Multi-site capability
 - ◆ Collaboration
 - ◆ Disaster Recovery



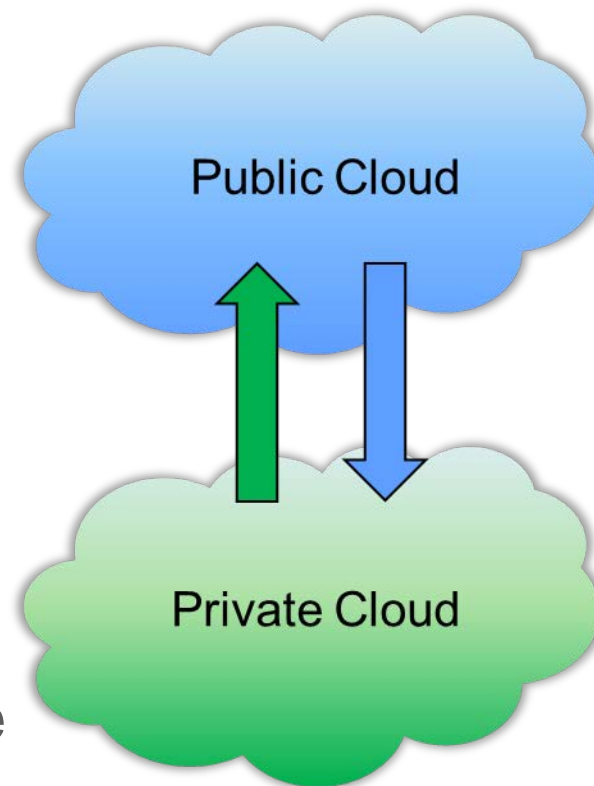
Private or Hybrid Storage Cloud

➤ Private Storage Cloud

- ◆ Scalability of public cloud
- ◆ Behind firewall security
- ◆ LAN performance
- ◆ Lower TCO than public cloud

➤ Hybrid Storage Cloud

- ◆ Private + public cloud
- ◆ Application development, archive



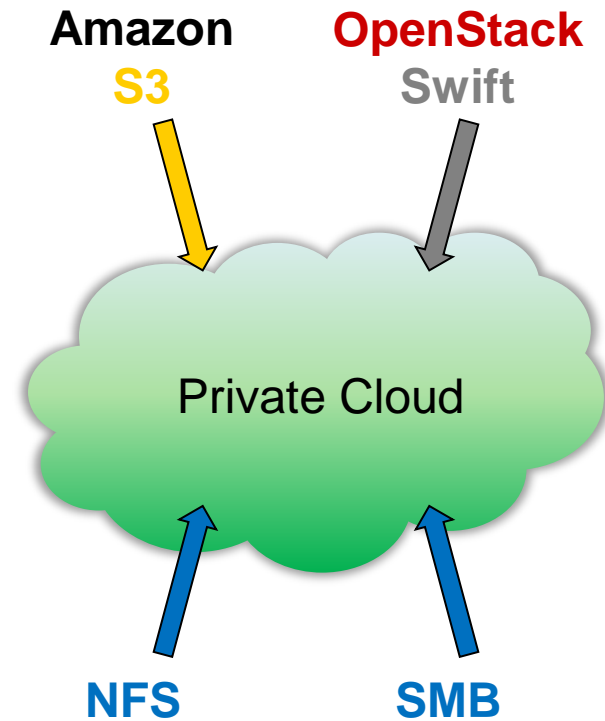
Driving Capacity into the Data Lake

➤ Industry standard interfaces

- ◆ Plug and play simplicity
- ◆ Legacy file protocols
- ◆ Cloud APIs

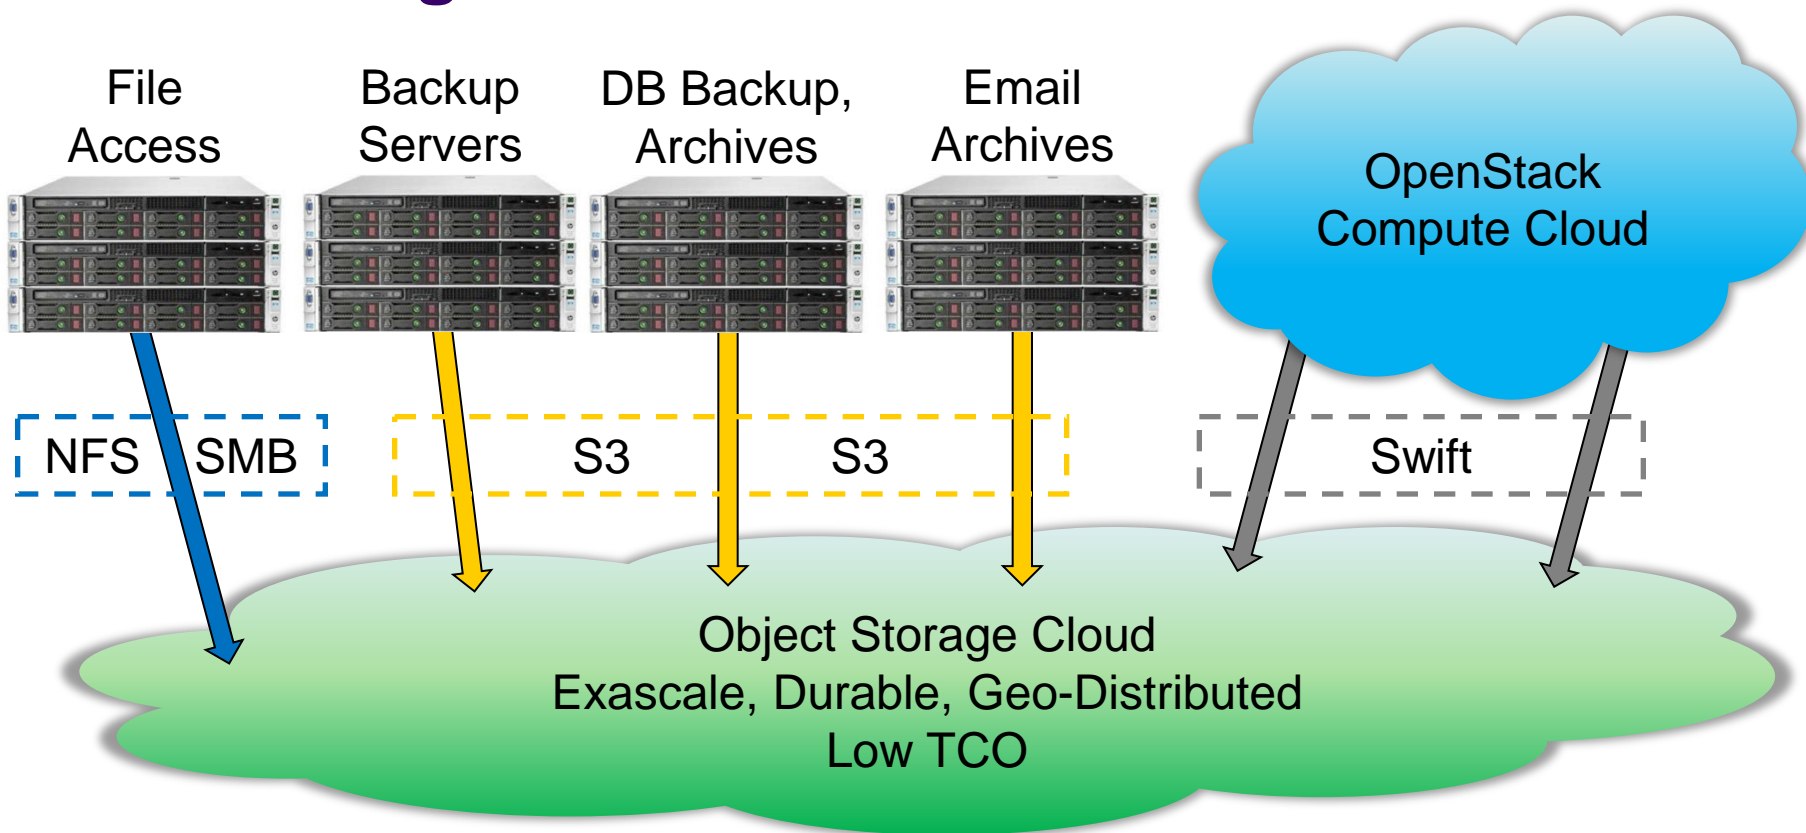
➤ Mix and match

- ◆ Multiple interfaces based on application requirements



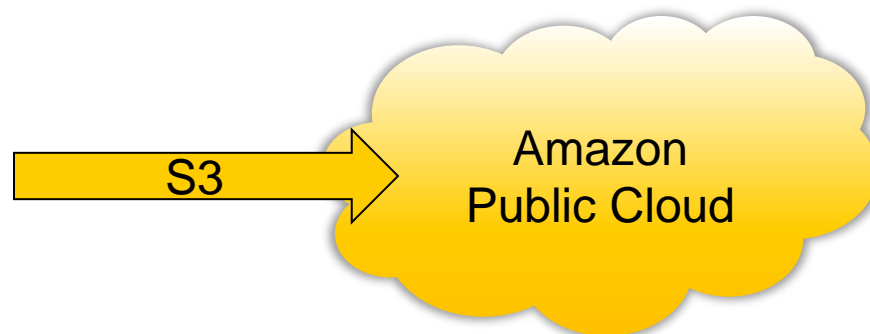
Object Storage – Data Lake Model

Accessing the Data Lake



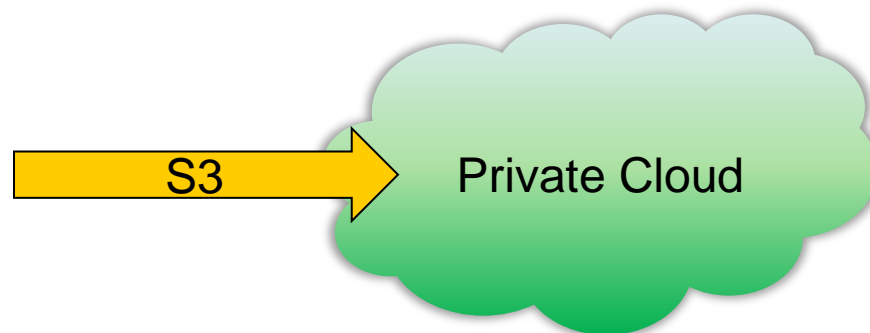
Develop, Test and Launch in the Cloud

- Applications developed in the public cloud using S3 storage



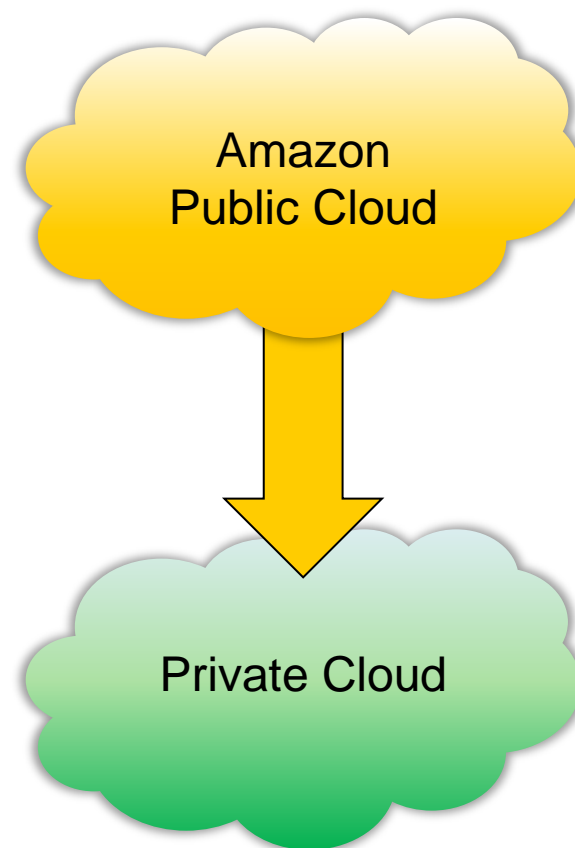
Enables

- Migration to private cloud without code changes



Motivation to Migrate

- Lower cost at scale
 - ◆ 12 to 18 month payback
- Enhanced security
 - ◆ Behind the firewall
- Improved performance
 - ◆ LAN vs WAN



Object Storage Data Lake

- Provides near limitless scale
- Provides superior data durability and mobility
- Lower cost than public cloud with local performance
- Support legacy and emerging interfaces and integrates easily into existing environments
- Reduces complexity and storage administration by eliminating storage silos and data management points
- Delivers the lowest possible TCO

Thank You!

Questions & Answers