Cloud Bursting

HPC for Wall Street Conference

April 4, 2016

Scott Jeschonek, Director of Cloud Products at Avere Systems

scottj@averesystems.com

AVERE

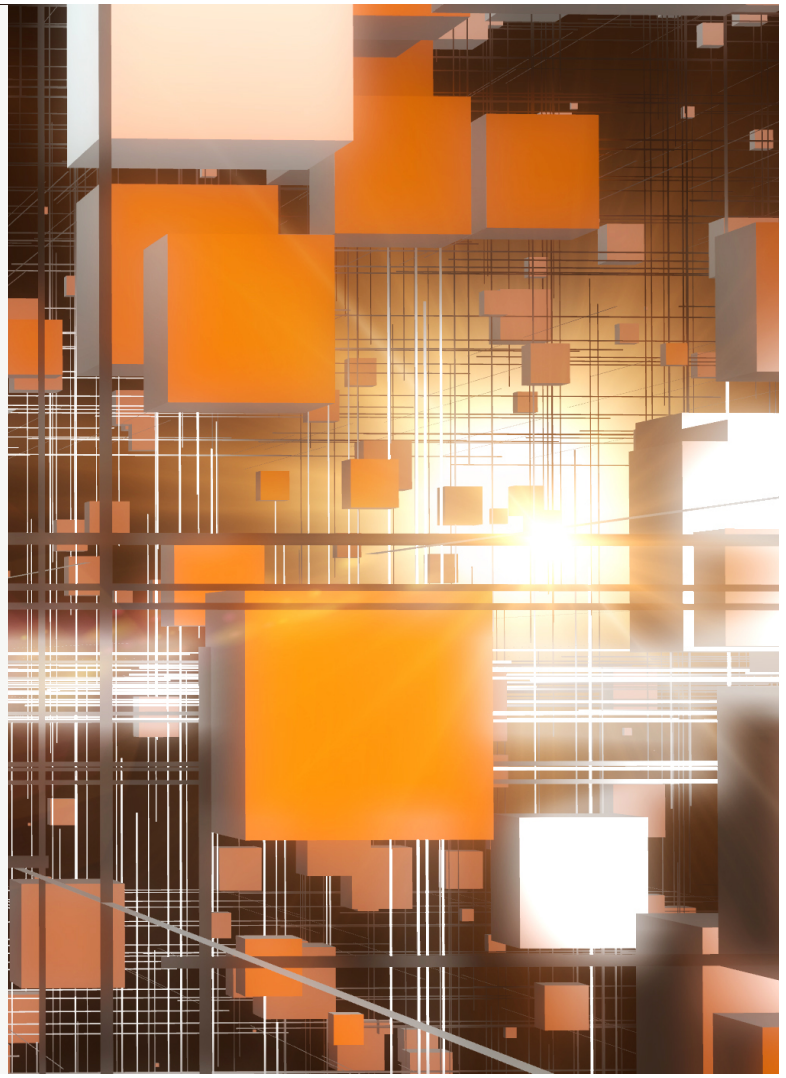

Yesterday, and Today: On Premises Datacenter, Redundant DC

- Multiple remote offices with analysts
- Single large office with analysts
- Own datacenters or leased DC equipment/space
- Site redundancy
- Data local to compute environments

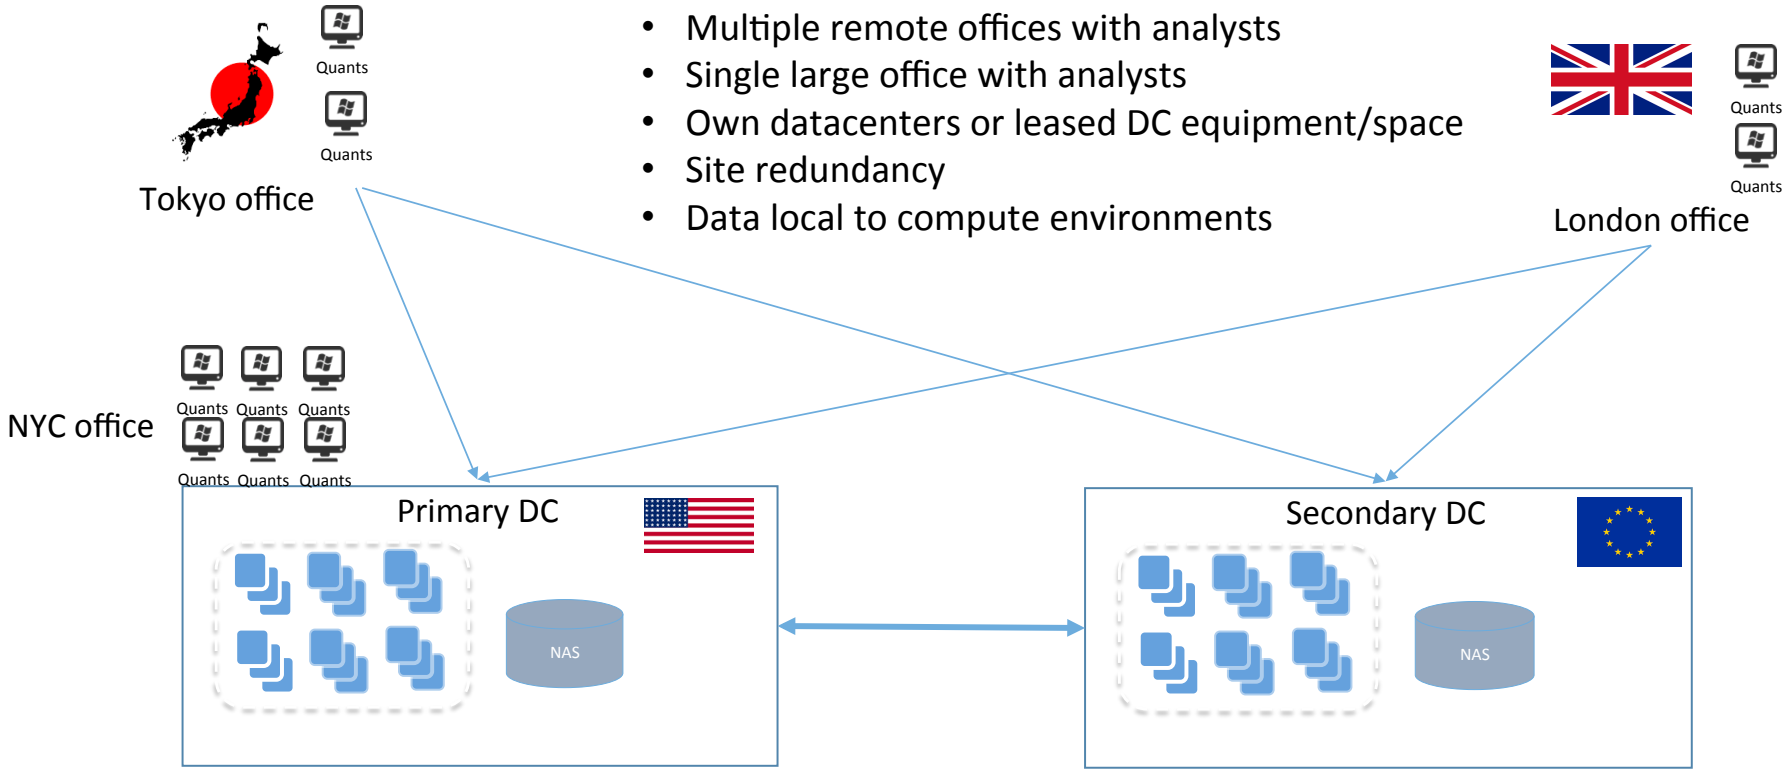

AVERE

Near Future, Hybrid Cloud

- Adoption of one or more cloud providers
  - > 1 hedge on price and SLA
- Mix of on-prem and cloud resources
- Regulatory, proprietary and/or security characteristics will likely keep data in the DC

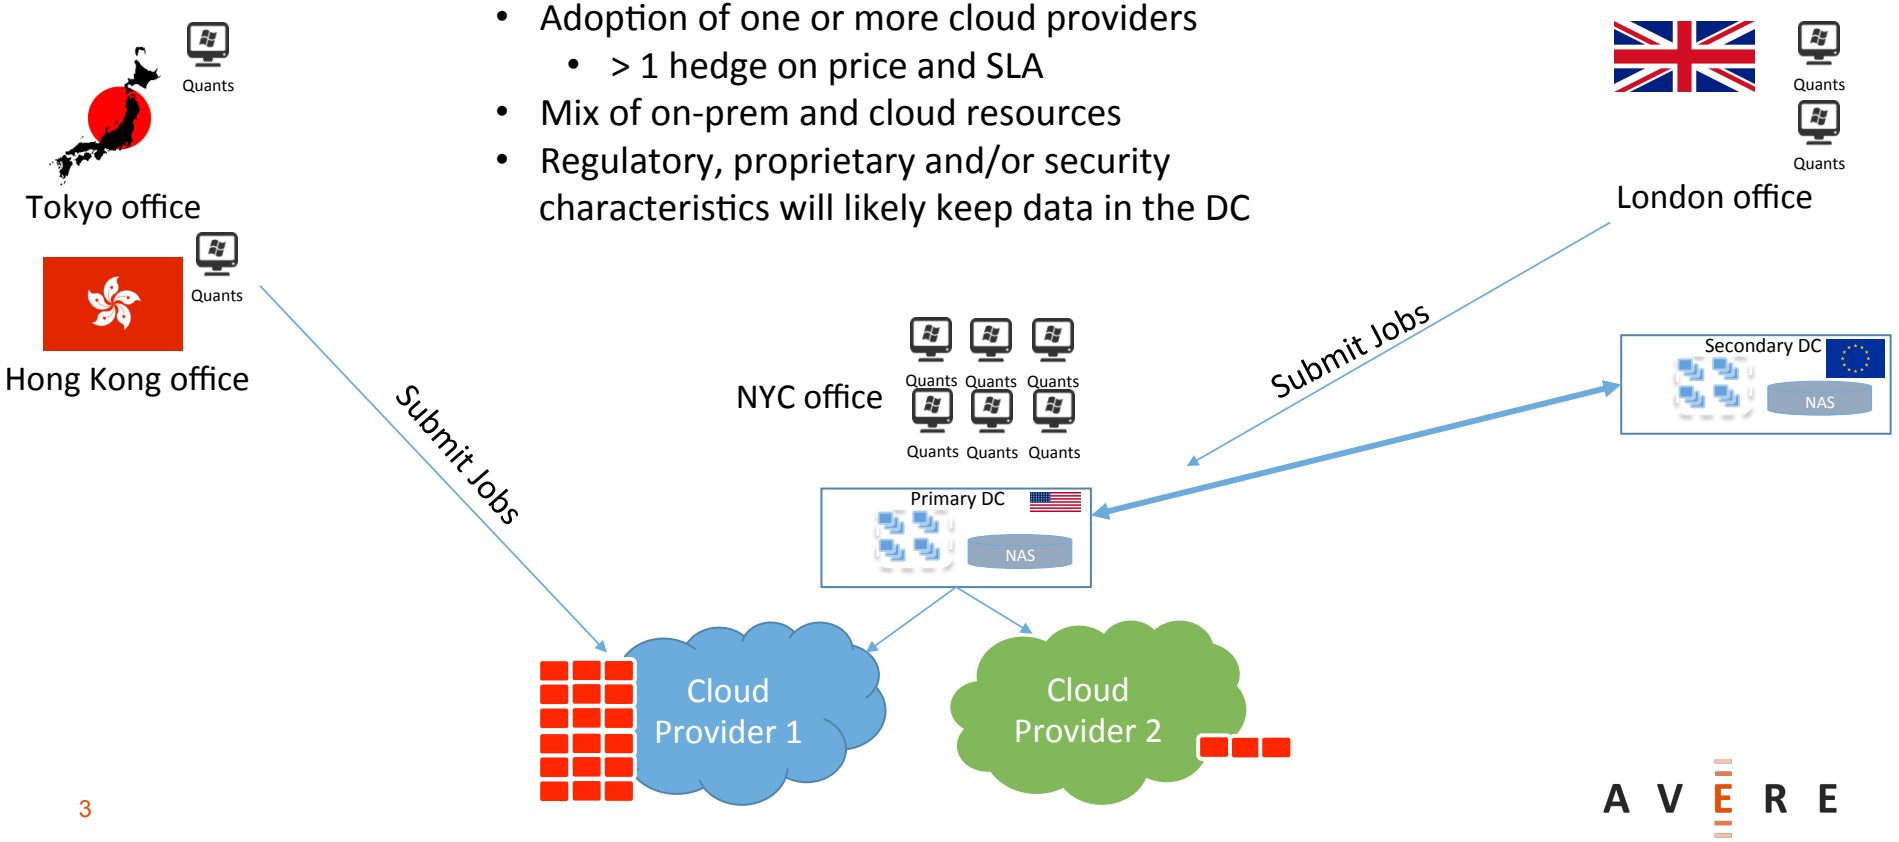

AVERE

Advantages of Someone Else's Cloud

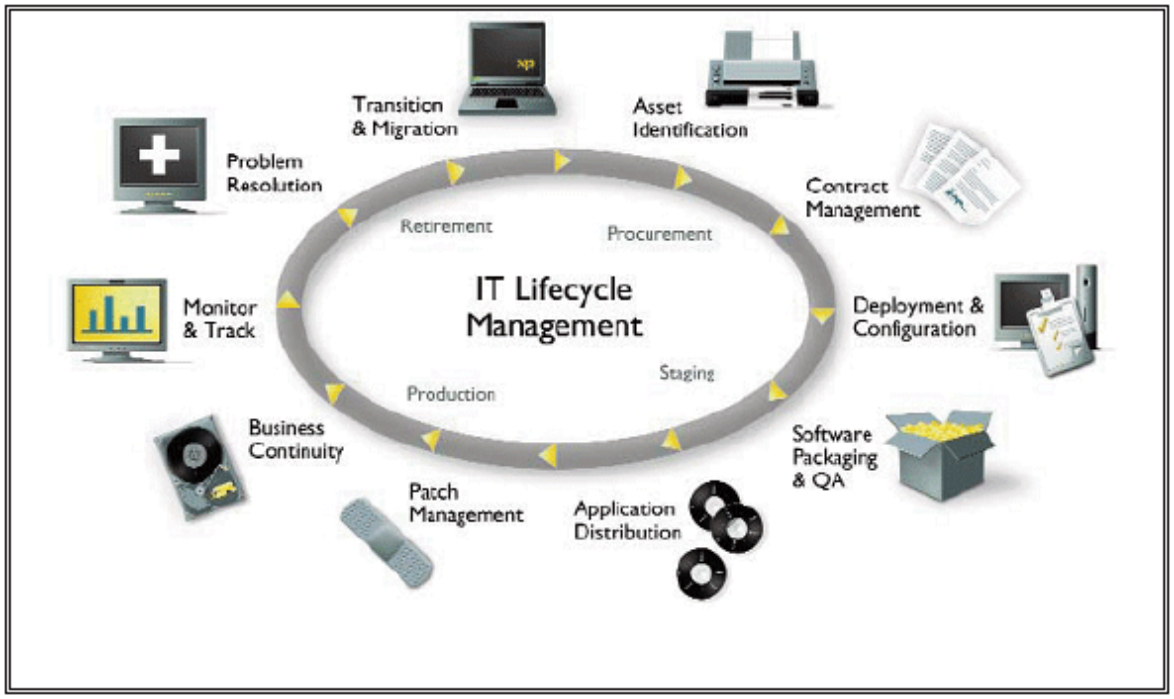

Significantly reduce infrastructure management costs both in money and time

Maintain operational flexibility during scale-out jobs…let the provider deal with scale challenges

AVERE

Analytic Affinity of the Cloud

Run thousands of cores against thousands of (actuarial, financial, genomic, scientific) jobs

Run these cores on-demand

Enable the cores and jobs with a few lines of code or a script, or use your existing CM infrastructure

When finished, execute a few more lines of code or script and tear it all down (if you like)

Relentless ability to repeat this

AVERE

Processing Destined for Cloud Computing

High Core Count Applications:

Computationally intensive: Monte Carlo simulations, Back-testing, Batch

Read-heavy workloads: Reading in large amounts of historical market data

Batch or burst usage model: Running simulations overnight

Evolving to an on-demand pricing/risk option: Near real-time pricing execution

Portfolio rebalancing

AVERE

One Practical Approach to Leveraging Public Cloud

Locate “parity” work in cloud compute:

- Batch
- Back-testing
- Simulation work

Use Cloud Services where applicable (e.g., Query Engines, Machine Learning)

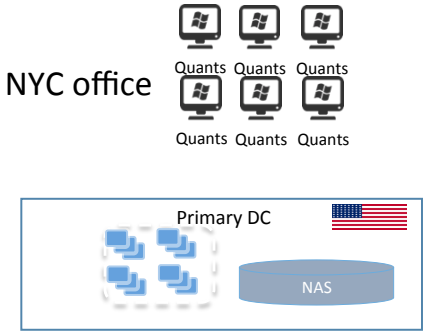

Use local DC in two ways:

1. Until current assets are retired
2. Ongoing R&D / Best-in-Breed technology facility…tech that is ahead of the public cloud lifecycle

AVERE

Implementation Considerations

Latency	Between Compute and Data
Locality of model data	Data sets are large and costly to move
Security / Regulatory	Data may need to be at rest in a specific geography Data encryption Restricted access to data
Cost	20,000 cpu cloud cores, 3 hours a day, annually ~$1.5M

Fundamental Question: How does cloud compute access data?

AVERE

Cloud Compute Data Access Option 1

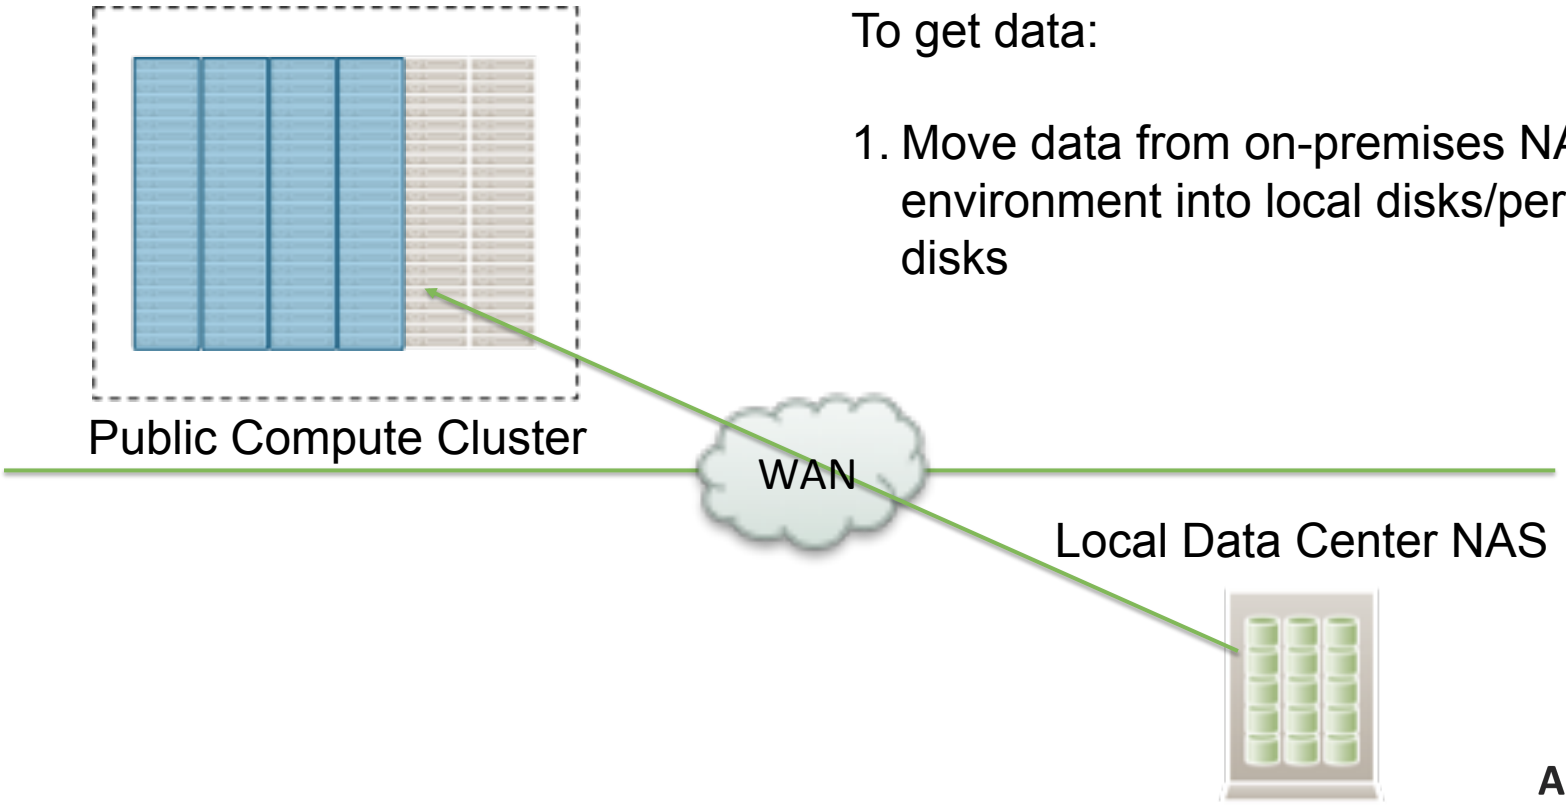

To get data:

1. Move data from on-premises NAS environment into local disks/persistent disks

AVERE

Cloud Compute Data Access Option 2

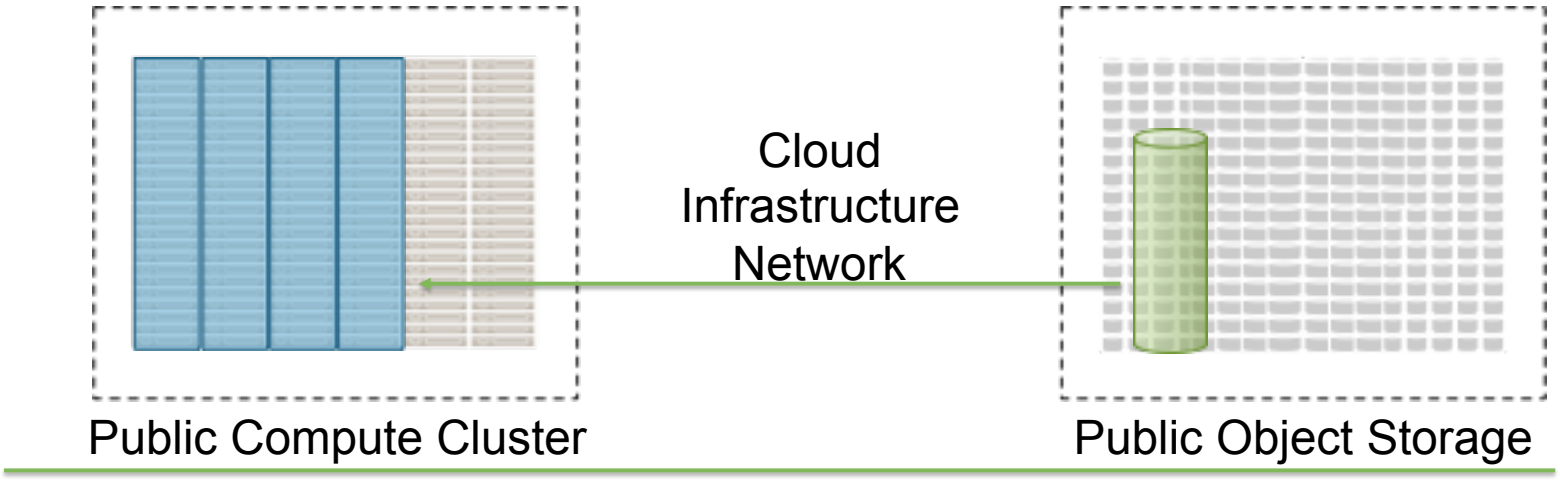

To get data:

1. Ingest Data from Cloud Provider Object Storage

AVERE

Cloud Compute Data Access Option 3

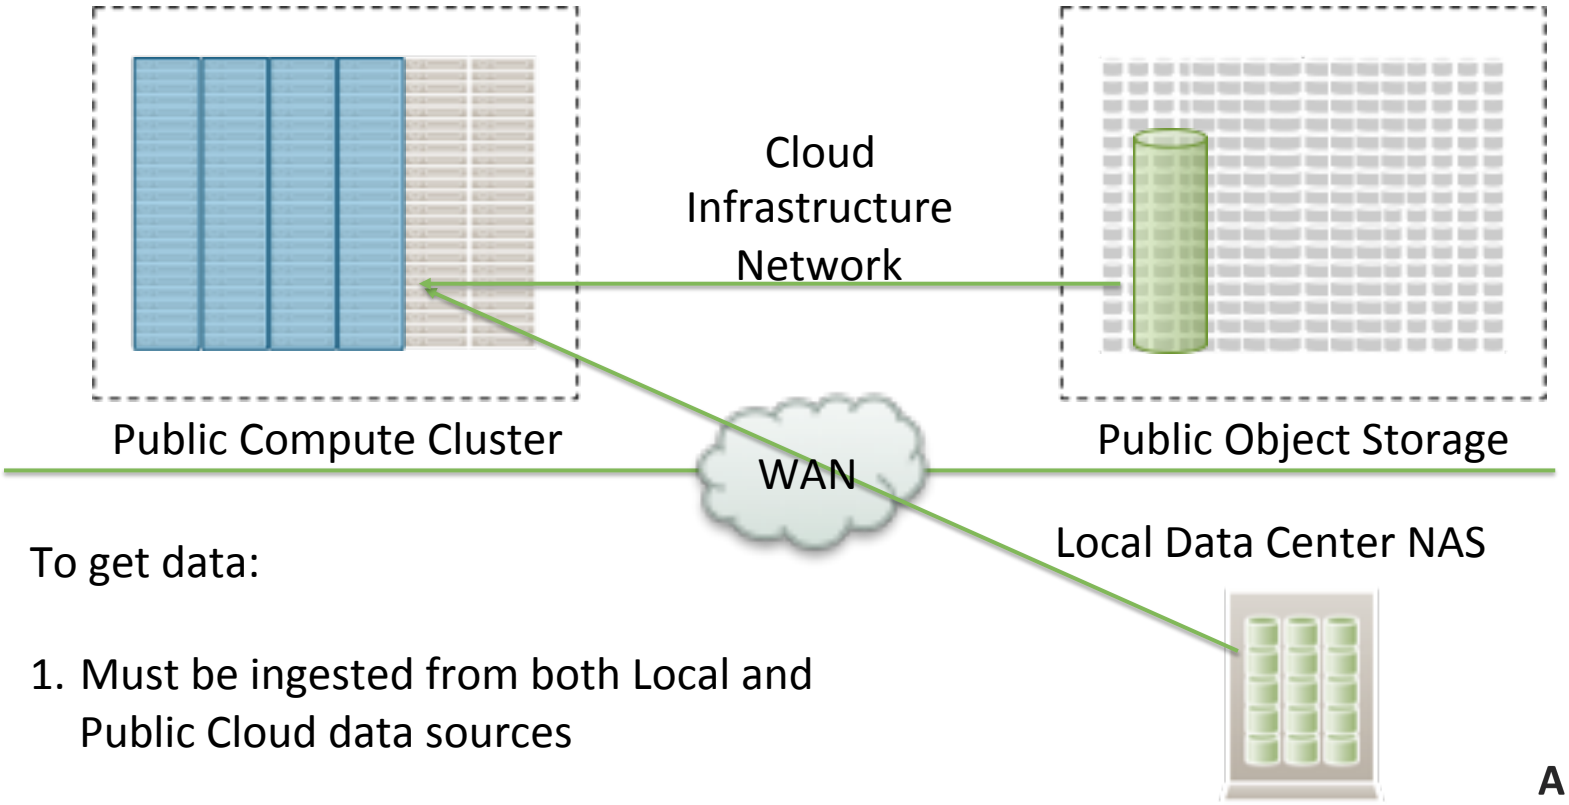

To get data:

1. Must be ingested from both Local and Public Cloud data sources

AVERE

Additional Challenges

- **Ingest Latencies**: Data likely resides in multiple locations across WAN or Cloud Infrastructure connections
  - Increased Latencies for loading will result in longer ramp times for jobs
  - More susceptible to network issues

- **DataNode costs in the cloud** - Disk must be made available for processing. Provisioning latency may grow depending on the size of the local disks

- **Local versus Persistent SSD** — If the cluster is meant to persist, the disks must also persist.
  - Local disks “go away” when the compute node is terminated or even moved in clouds

AVERE

Typical File Access in Hadoop Cluster

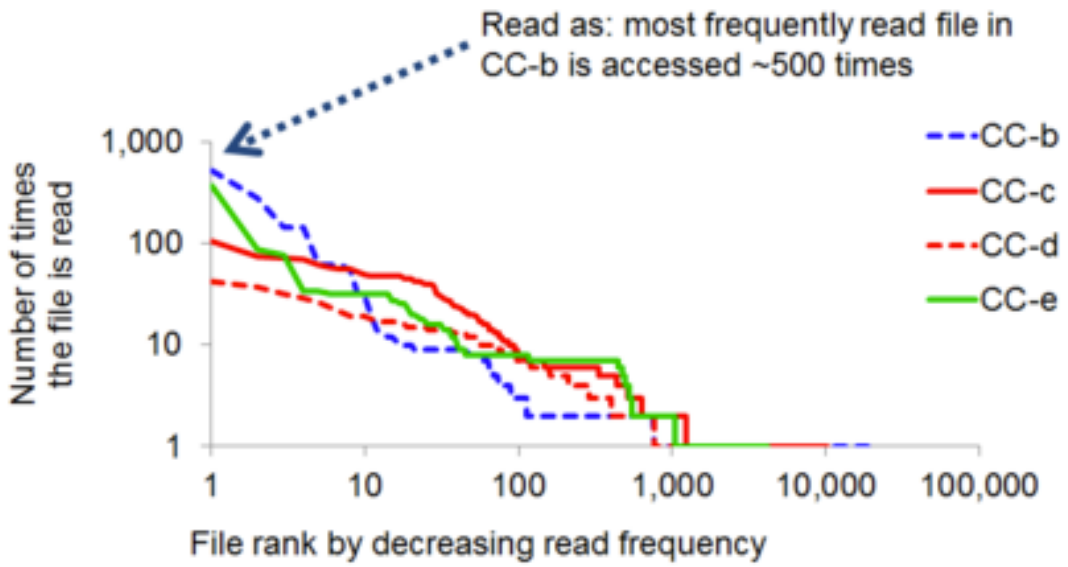

source: http://blog.cloudera.com/blog/2012/09/what-do-real-life-hadoop-workloads-look-like/

Caching files will work for certain types of jobs

Where typical file is accessed By multiple clients

AVERE

Avere Building Blocks

"Avere is uniquely positioned to offer scale across tens of thousands of cloud compute cores while leaving the data where it originates, on premises, with it's global file system and caching capabilities." - Unnamed CTO

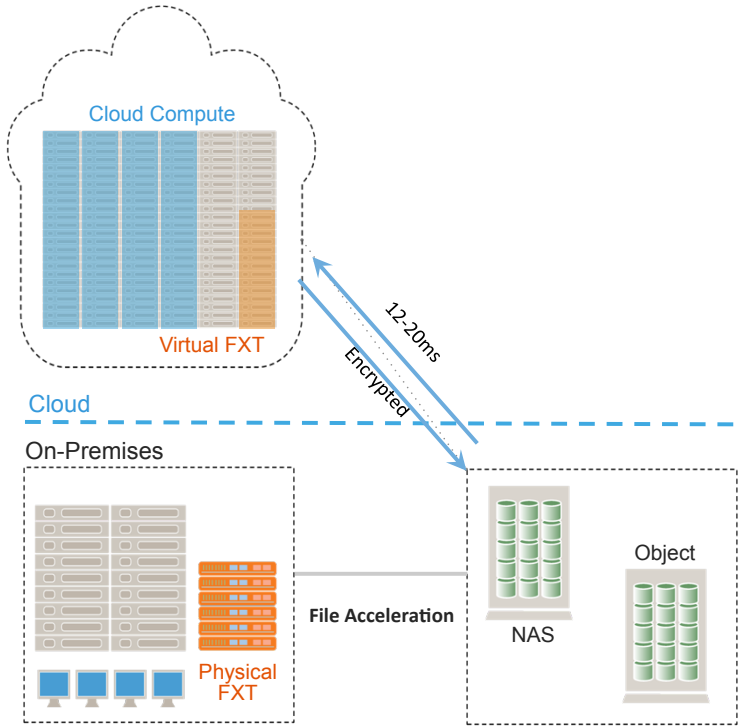

AVERE

Simplify Data Access Across Sources

Global Namespace, FlashMove®, FlashMirror®

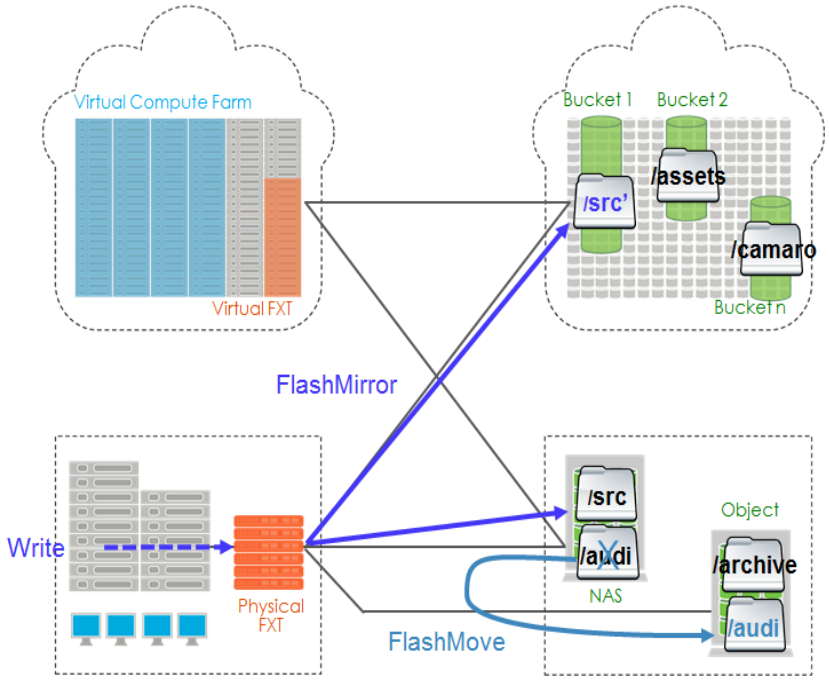

Global namespace (GNS)

- Simplified management
- Single mount point
- Heterogeneous core filers
- Junctions (/sw) for better mgmt

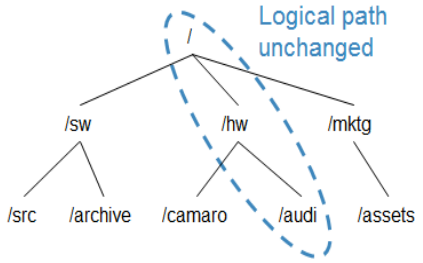

AVERE

Use Case for Avere Hybrid Cloud

Design goal: Run “wonderfully” parallel workloads on cloud compute, securely, economically, and at scale.

Requirements: Model data stays on-premises for daily quant access. Cloud compute needs to scale to double, triple and quadruple on-premises compute resources to drastically reduce computational times. Applications need to be used “as they are”.

Results: Models run 9x faster at ⅕ the cost.

AVERE

“Hybrid Hadoop”

- Leveraging Avere technology and NFS
- Job code access “local” data on the file system
  - Really mapped to NFS mount points on the Avere vFXT cluster
- Avere caching will improve access times for all nodes except the first one or two, as the file is gradually loaded based on demand
- Data is clustered; therefore a failure of one of the Avere cluster nodes does not require a new cold read…
  - HA partner node will have the cached data

AVERE

FS Use Case: Market Risk Analysis with Cloud Bursting

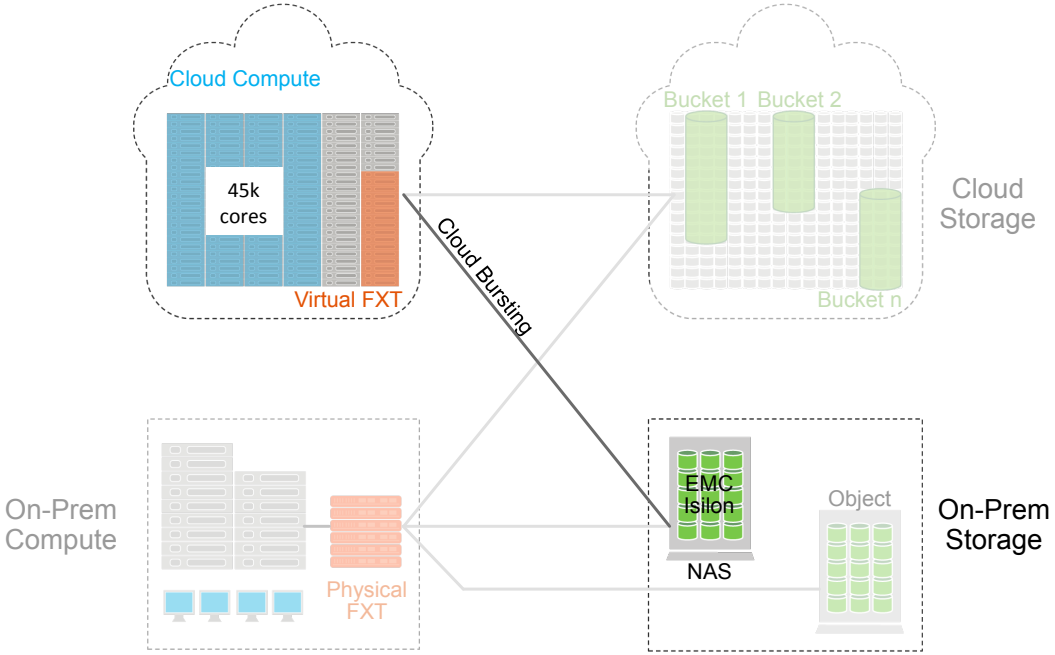

Customer Challenges

- Simulation complexity increasing
- Run 10-100x more simulations
- Finish models in less time
- Reduce $/simulation

Avere Benefits

- vFXT provides scalable file system for cloud
- Scale to more than 45k compute cores
- Auto-caching of data from on-prem storage
- Cloud economics: zero footprint, easy to turn on/off, pay only for what you use

Actual Customer Results

- 45k cores used in Major Cloud Provider

AVERE

Thank you!

www.averesystems.com
sales@averesystems.com
888-88-AVERE

AVÉRE