

DataDirectTM
N E T W O R K S
I N F O R M A T I O N I N M O T I O N TM

DDN & iRODS

iRODS User Group Meeting
Cambridge, MA
June 18, 2014

David Martin
WOS Product Line Manager

Storage should improve collaboration

... *Not make it harder*

- ▶ Distributed, not centralized
- ▶ Minutes to install, not hours
- ▶ Milliseconds to retrieve data, not seconds
- ▶ Replication built in, not added on
- ▶ Instantaneous recovery from disk failure, not days
- ▶ Built in data integrity, not silent data corruption



DDN is the World's Largest Privately Held Data Storage Company



Key Statistics

Overview:

Providing the data backbone for the world's most extreme, data-intensive environments – enabling organizations across multiple industries to maximize the value of their information globally.

- Established: 1998
- Headquarters: Chatsworth, California USA
- Over 1,000 Customers in 5 Continents

Key Industries:

- High Performance Computing & Life Science
- Cloud & Web Content
- Rich Media
- Intelligence/Federal

Go To Market:

- Global Partners, VARs, Resellers

Industry Validation



World's Largest Privately-Held Storage Co.



Fast500 Technology Company

Inc. 500|5000 High-Growth Company

Best HPC Storage Product



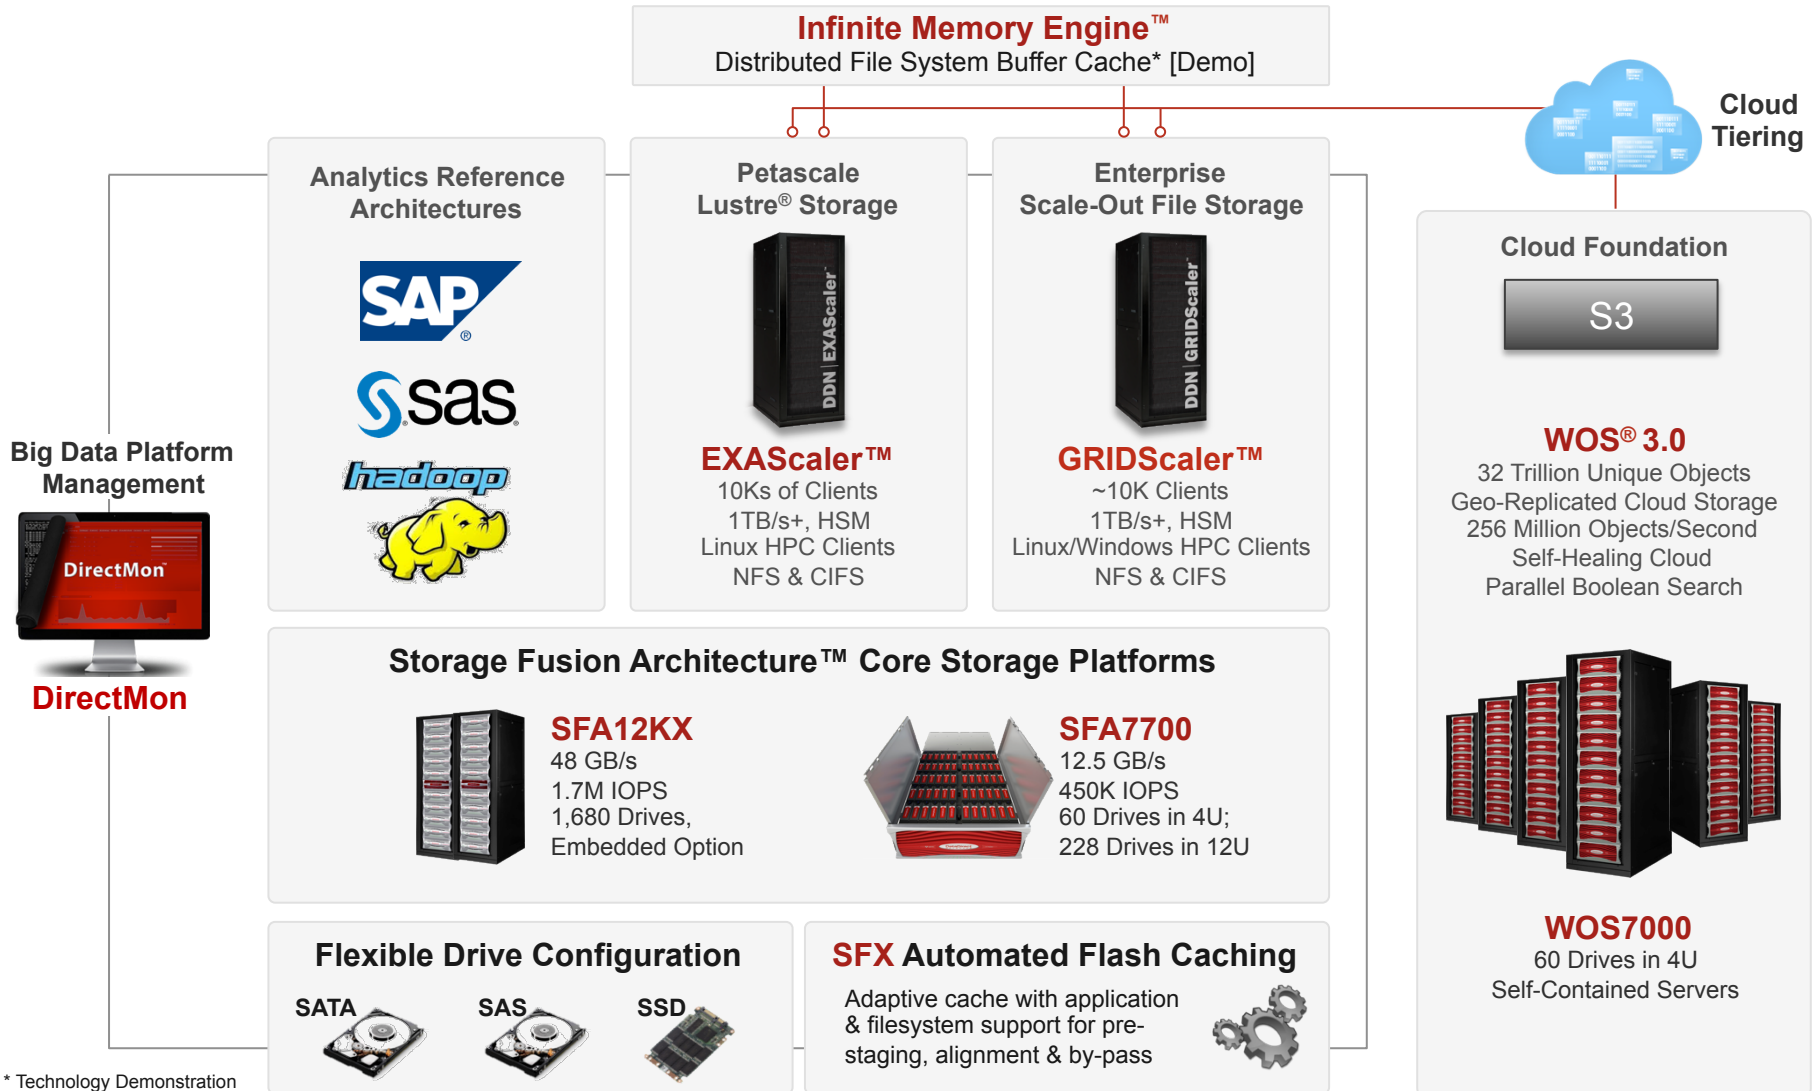
Best Practice for Digital Media

World Class Set of Global Customers

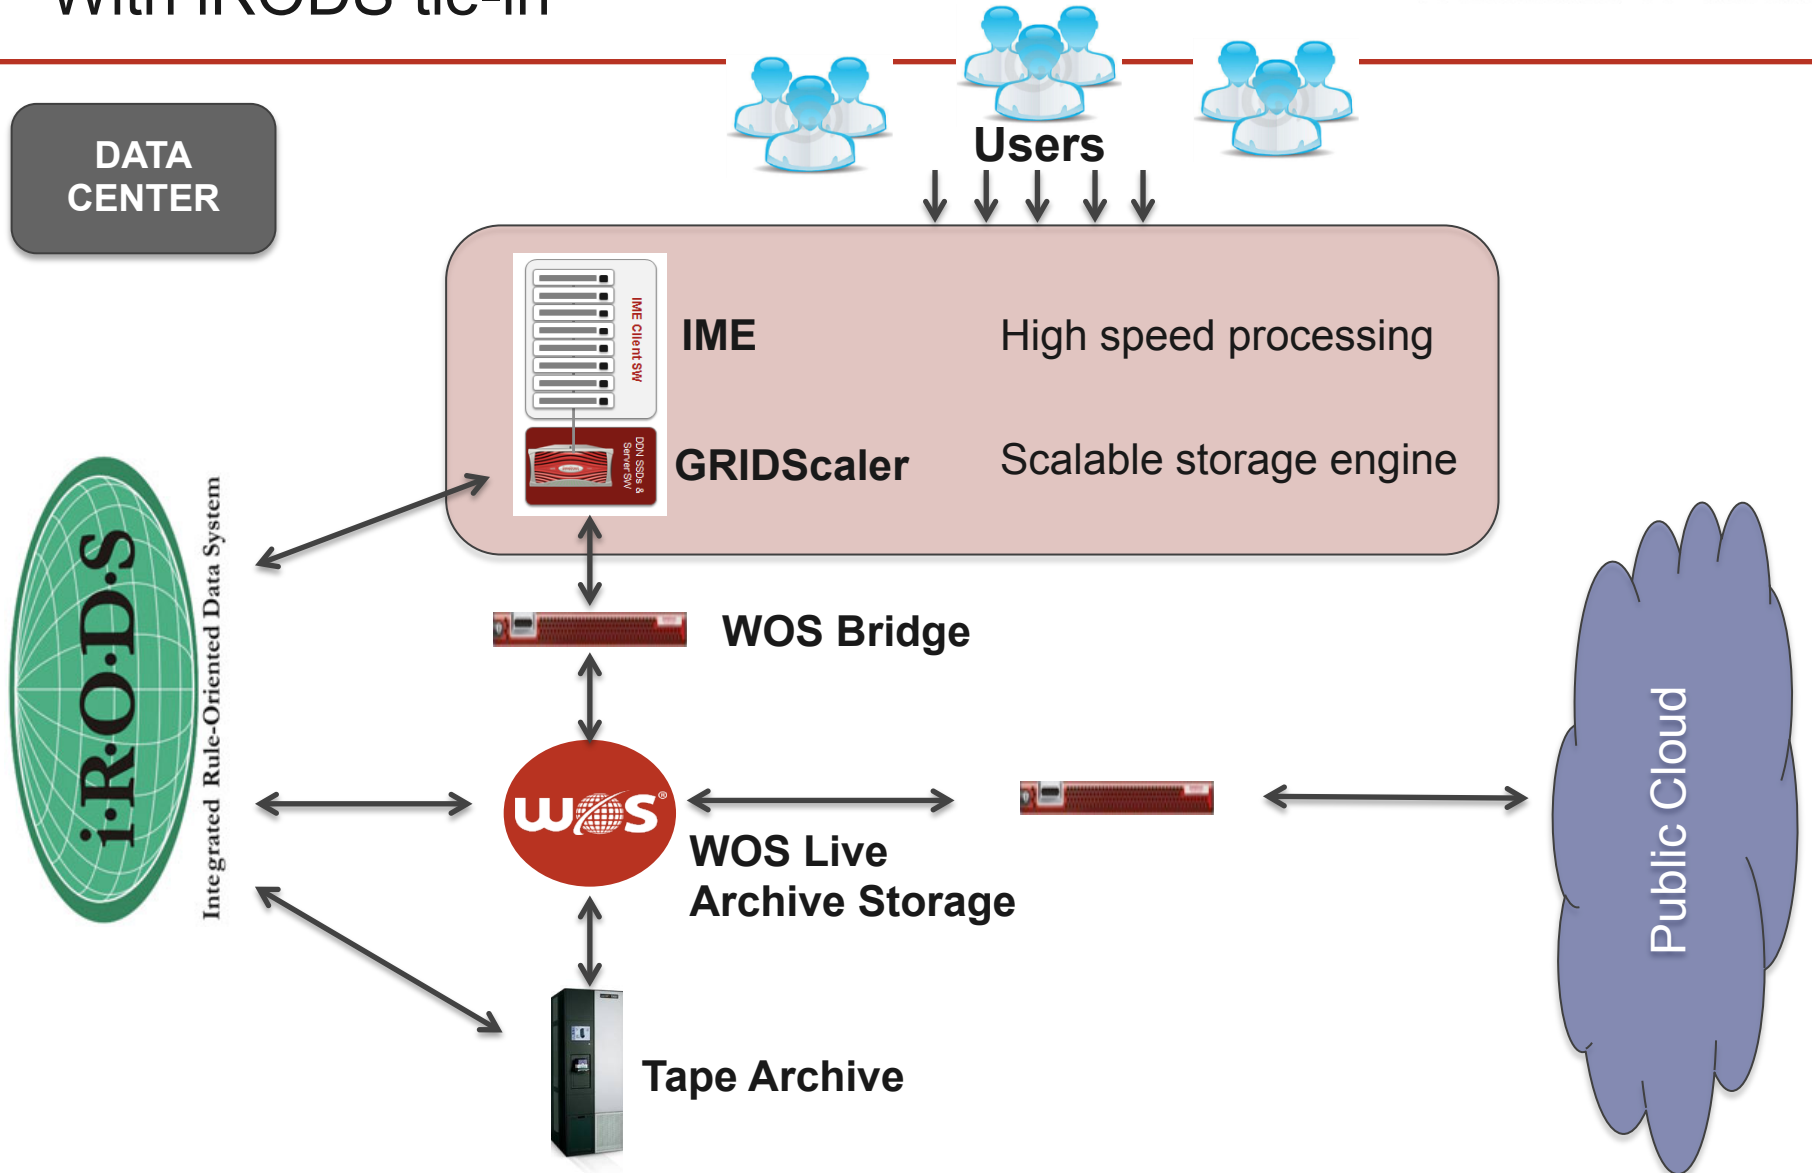


Big Data & Cloud Infrastructure


DDN Announced & Pending Product Portfolio



High Performance Data Center With iRODS tie-in



IRODS Data Grid & WOS Made for Big Data



• Hyperscale • Distributed • Collaborative

• Accessible • Secure • Flexible



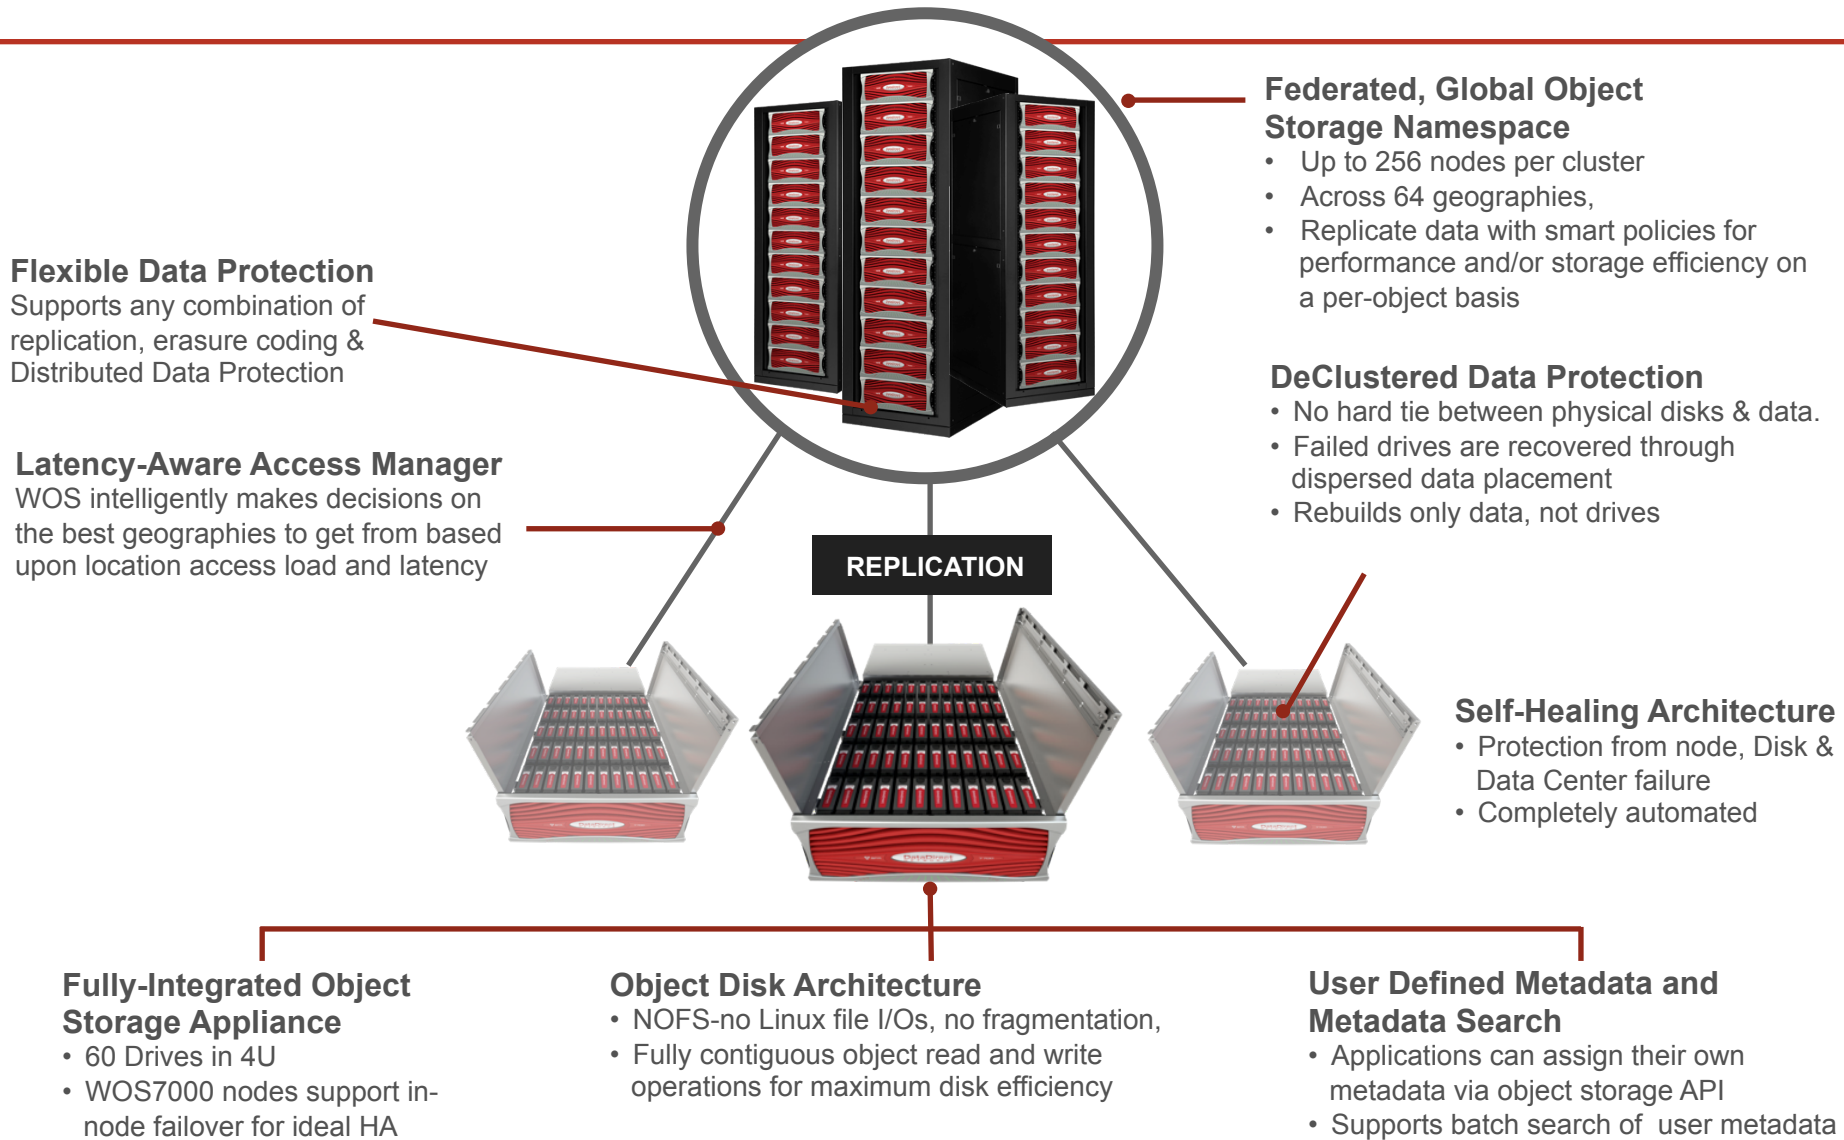
+



= Managed Collaboration

DDN products assure data longevity with low latency over geographies and iRODS adds the important policy management & accessibility layers

WOS Key Feature Breakdown



Why DDN WOS & iRODS are perfect complements

WOS extends iRODS by...

- ▶ WOS makes iRODS more performant & scalable
 - Single iCAT DB entry references multiple WOS copies
 - iRODS retrieves lowest latency WOS object instance

- ▶ WOS enables iRODS users to implement better QOS
 - High speed access to distributed project data
 - Long tail data can be stored on deep archive WOS Nodes

- ▶ WOS dramatically improves iRODS storage efficiency & availability
 - Self healing – at local node SAS bus speeds which increases data durability
 - Highly available – no interruption if disk, node, or site failure
 - WOS implements highly efficient Global Object Assure
 - Lowest TCO in the industry, <1 FTE for multi petabyte multi-site deployments

iRODS extends WOS by...

- Integration with heterogeneous storage environments (multi-vendor disk & tape)
- Additional clients and access methods
- Richer storage rules and migration policies

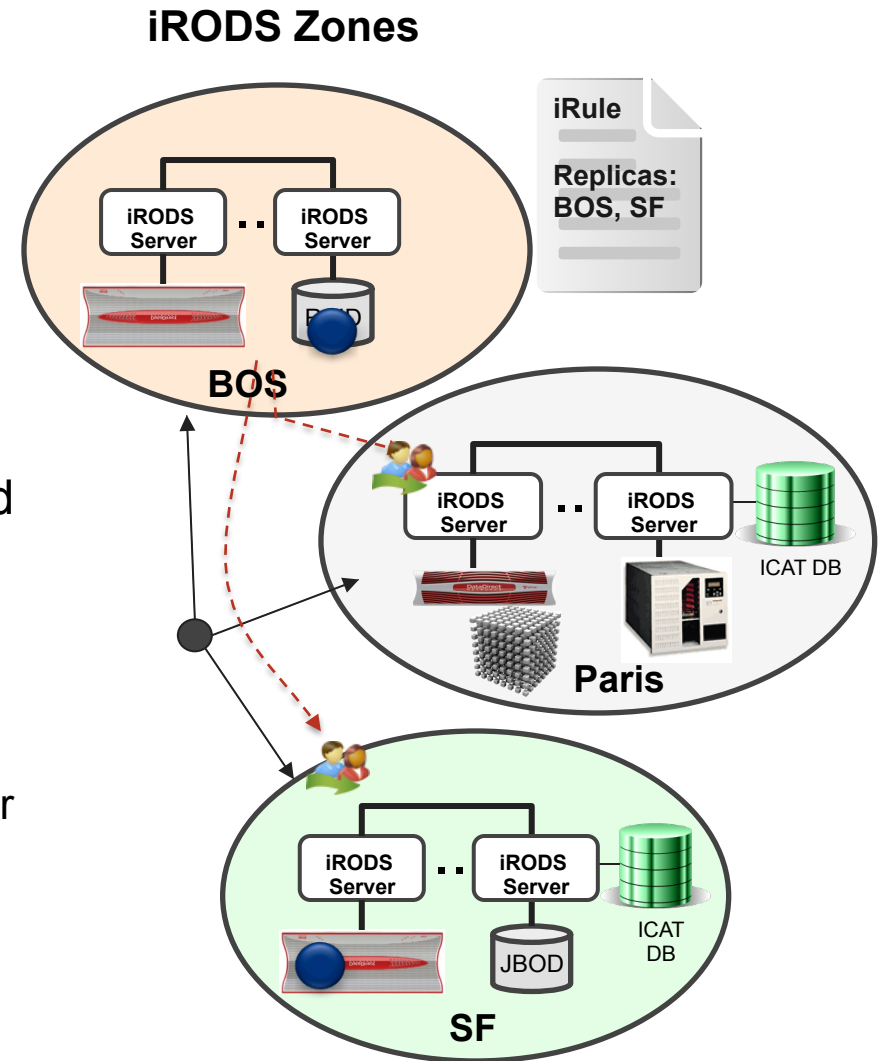
iRODS Standard Operation

iRODS Operation

- iRules determine file store locations
- iCat maintains file location metadata
- All retrieves get primary file location info from iCAT (all retrieves use primary)

Store & Retrieve Operation

- iRule policy determines where file is stored
 - FileA stored in BOS is replicated to SF automatically by iRODS
 - iCat maintains file location metadata
- File Access always reads primary copy
 - FileA read in SF or Paris will go to BOS server (primary)



WOS + iRODS Integration

WOS Location Intelligence

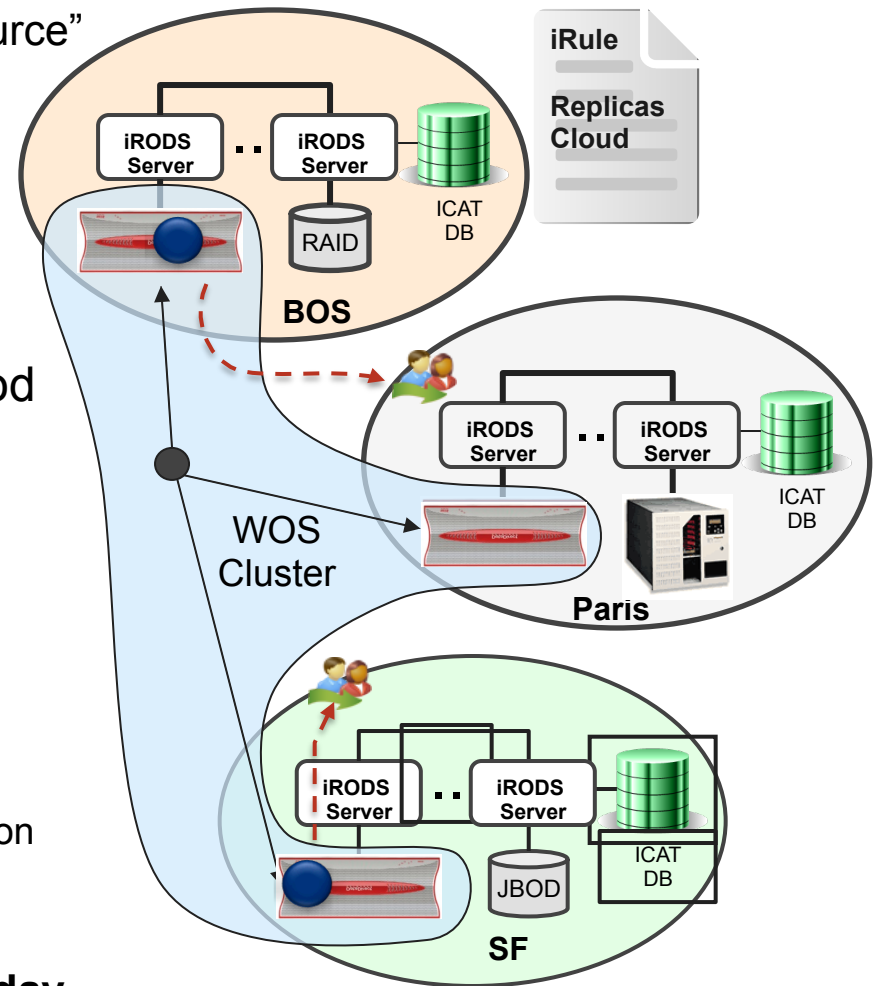
WOS as a iRODS Cloud resource

- WOS Cluster is defined as iRODS “Cloud Resource”
- Multiple WOS replicas stored once in iCAT
- iRODS utilizes WOS data location intelligence

File store / retrieve operation

- iRule policy specifies cloud storage method
 - FileA stored in SF writes to closest node of WOS Cluster
 - WOS replicates file to BOS per WOS storage policy
 - Single entry in iCAT DB
- File read always reads from closest WOS location
 - FileA read in BOS will retrieve from BOS WOS node
 - FileA read from Paris will retrieve from closest location (normally Boston)

iRODS Zones



iRODS optimizations for WOS is available today

Power of WOS Location intelligence

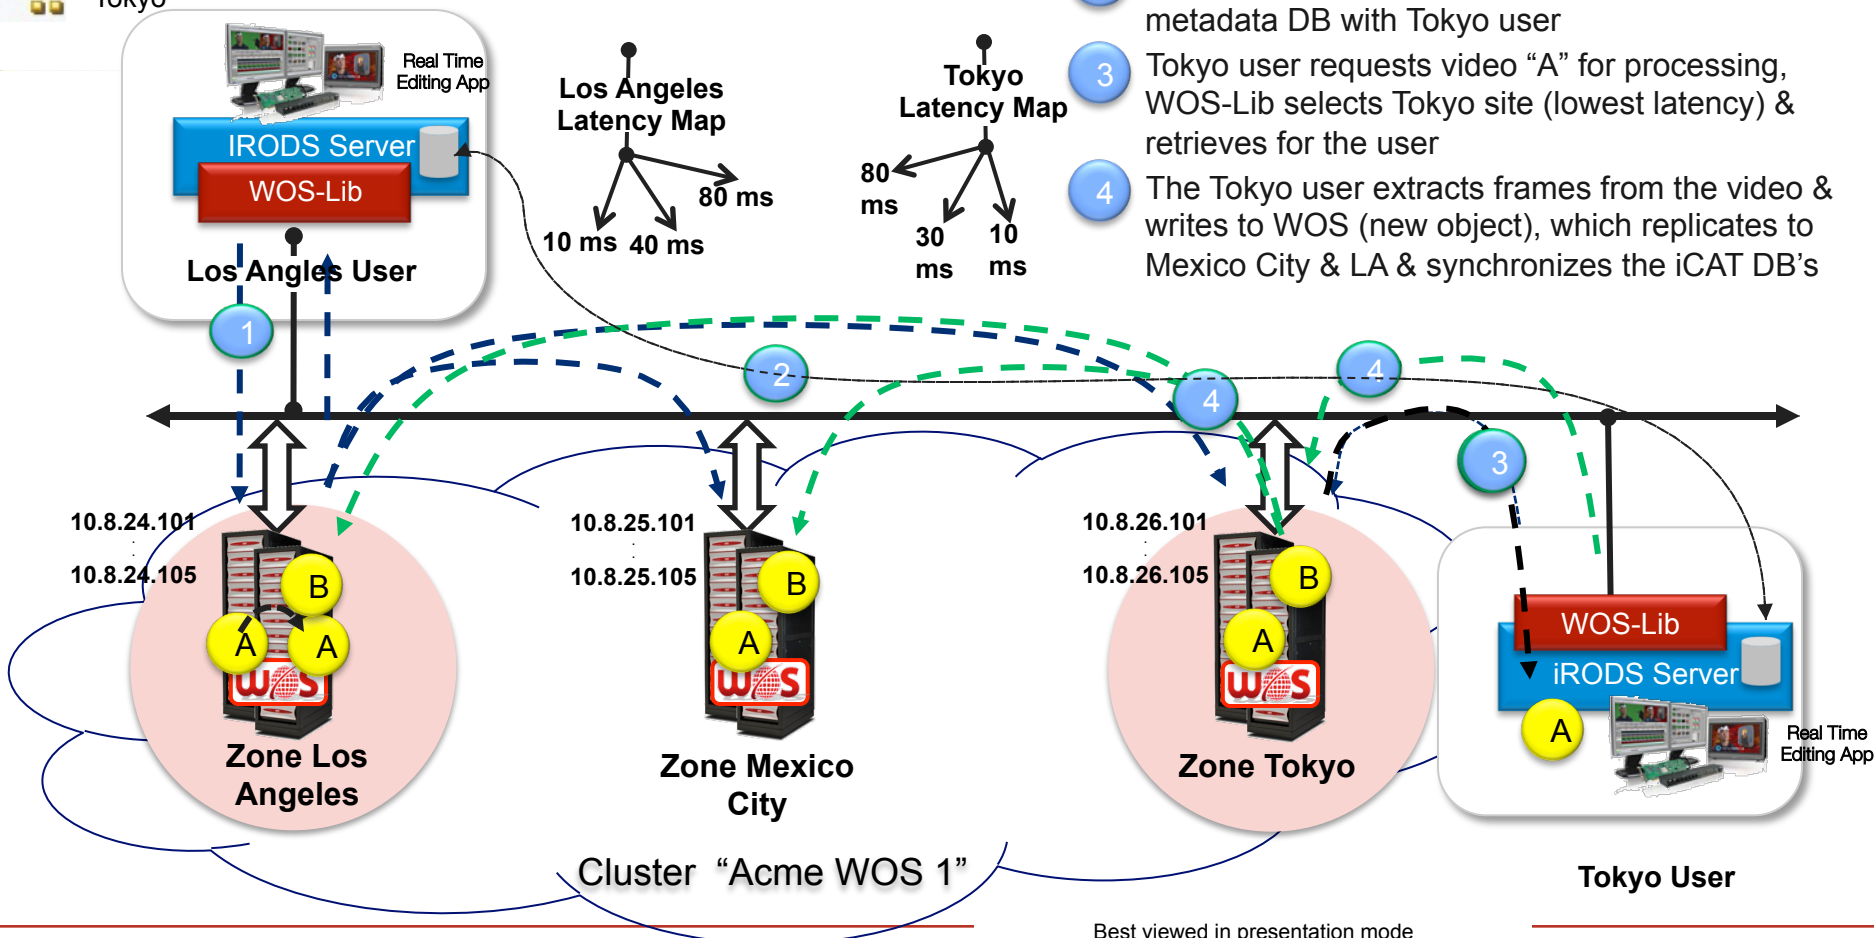
Minimize Latency & Enable True Collaboration

Acme WOS 1

- Los Angeles
- Mexico City
- Tokyo

Data Locality & Collaboration

- 1 LA site user edits video "A", which replicates to Mexico City & Tokyo based on policy
- 2 MP Gateway immediately synchronizes iCat metadata DB with Tokyo user
- 3 Tokyo user requests video "A" for processing, WOS-Lib selects Tokyo site (lowest latency) & retrieves for the user
- 4 The Tokyo user extracts frames from the video & writes to WOS (new object), which replicates to Mexico City & LA & synchronizes the iCAT DB's



Best viewed in presentation mode

WOS Replication


Replication delivers both multi-site collaboration + data & disaster protection

- ▶ Select either Asynchronous or Synchronous replication per policy
 - Definable on a per-object basis
- ▶ Synchronous Replication replicates to remote sites to insure data protection before acknowledging a successful “write” to the application
- ▶ Asynchronous replication to remote WOS nodes increases performance
 - Two copies of data are always written to “local” (lowest latency) node before “PUT” call returns OID to minimize/eliminate risk of data loss
 - Replication occurs in background, once replication successfully completes, extra local copy is removed
 - Especially useful for big files and big data sets




WOS Local & Replicated Object Assure

WOS Local OA (LOA)



WOS Replicated OA (ROA)



Local Object Assure:

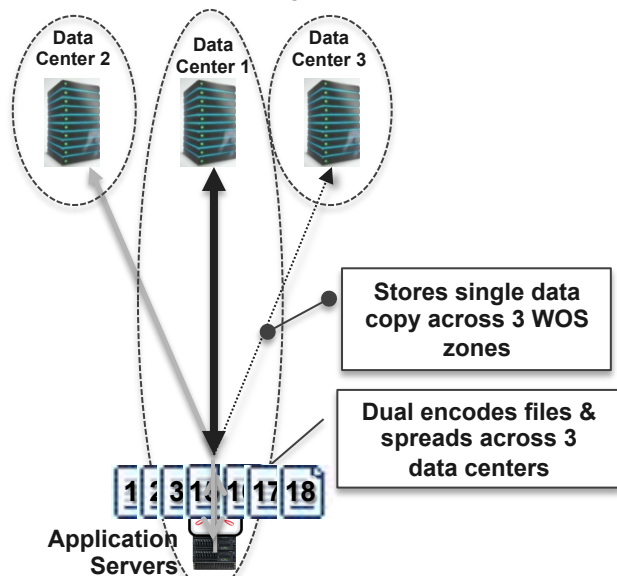
- Erasure-code based declustered data protection
- 8+2 (Data + Parity) for single copy protection
- Each WOS node can withstand up to two concurrent drive failures without loss of data or availability
- Local Data rebuilds

Replicated Object Assure

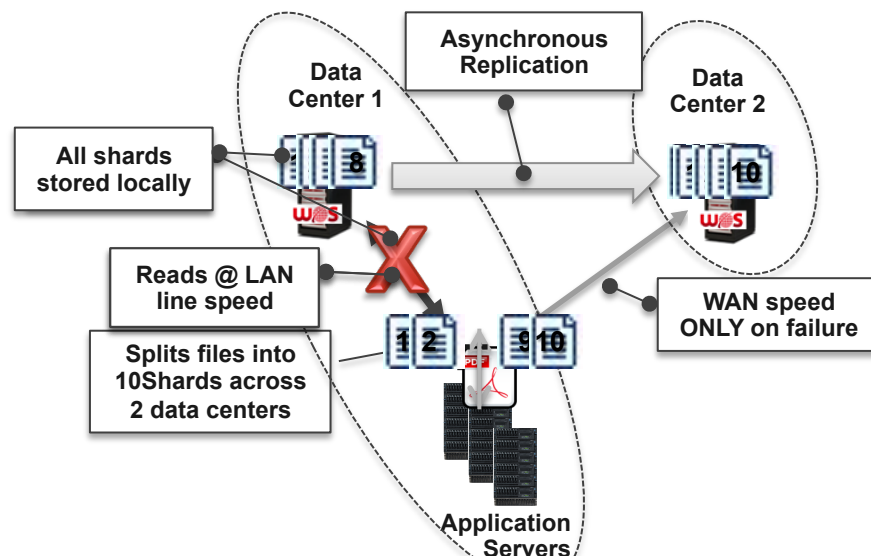
- DR for OA storage methods
- Sync & Async replication supported
- Optimized for large objects
- Two OA replicas provide equal protection as 3 standard replicas

WOS Global Object Assure & Replicated Object Assure

WOS Global Object Assure



WOS OA Replicated



Best viewed in presentation Mode

The Net-Net - Ultimate customer flexibility

Global Object Assure

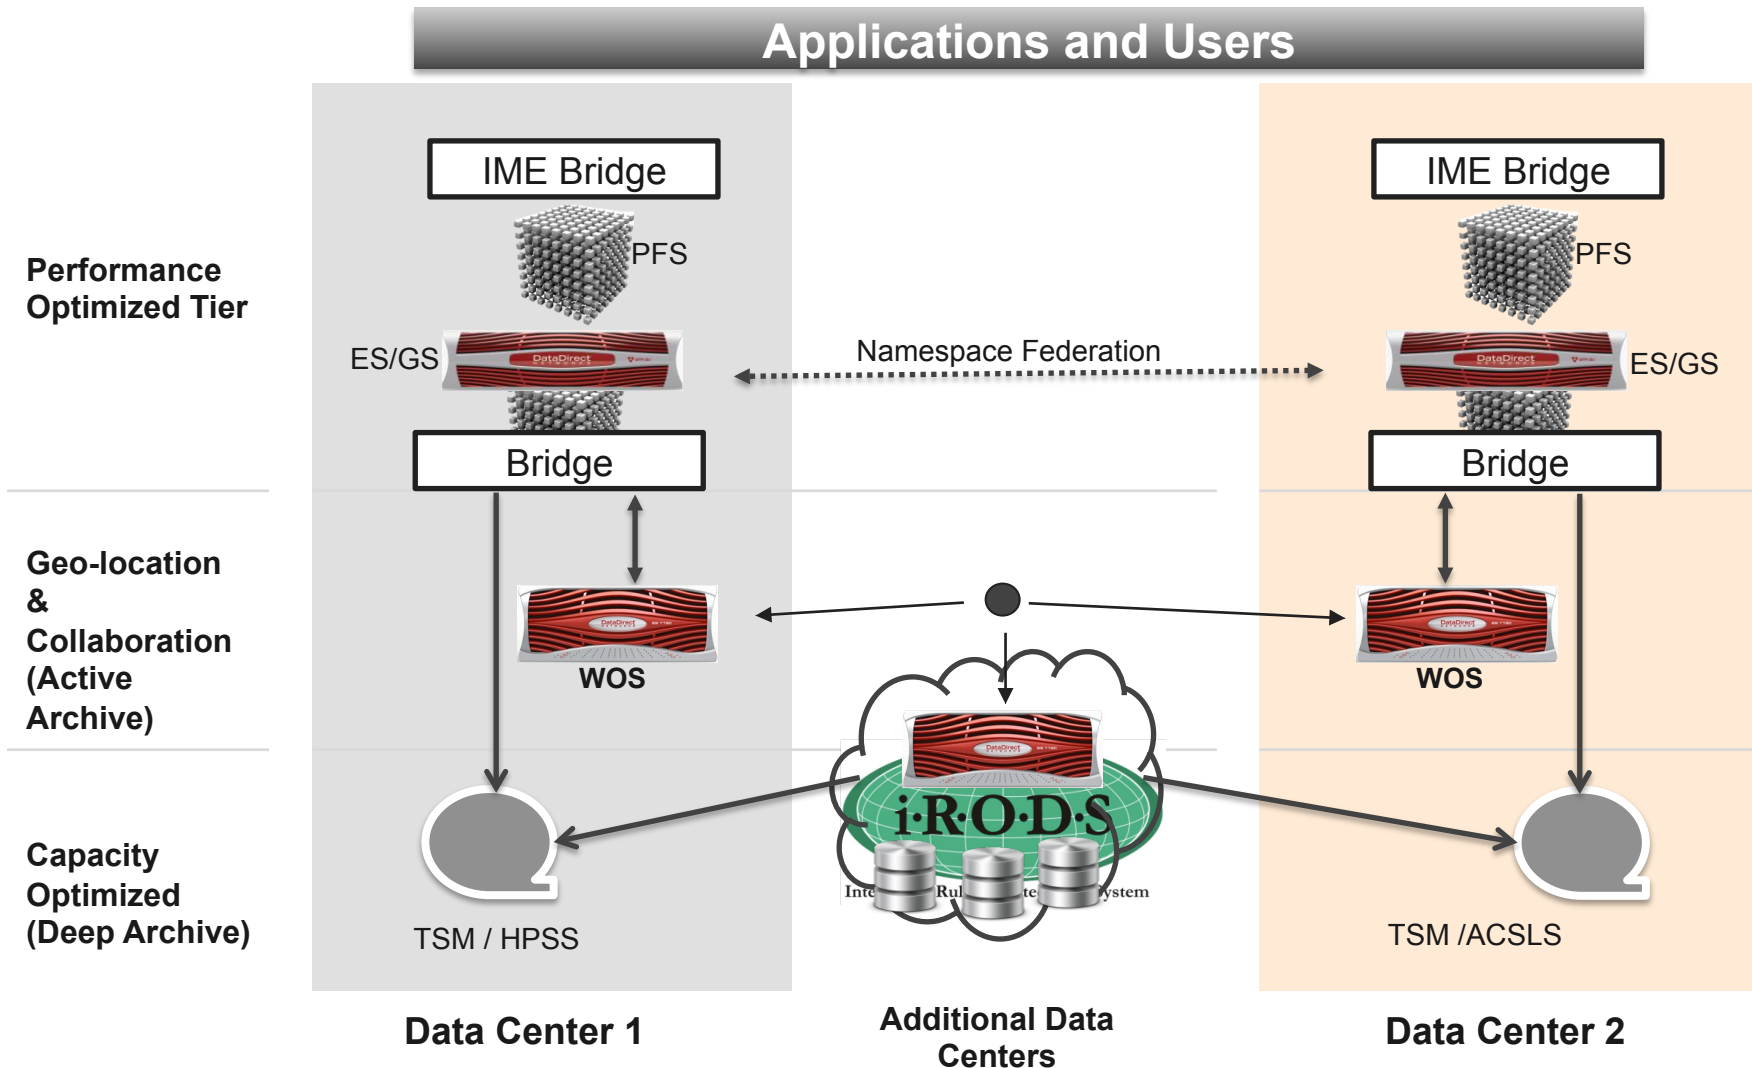
- A minimum of 3 data centers must be deployed to provide disaster protection
- High storage efficiency
- Data rebuilds from drive failures occur @ SAS Bus speeds

Fastest LAN Path = Performance (unless DR scenario)

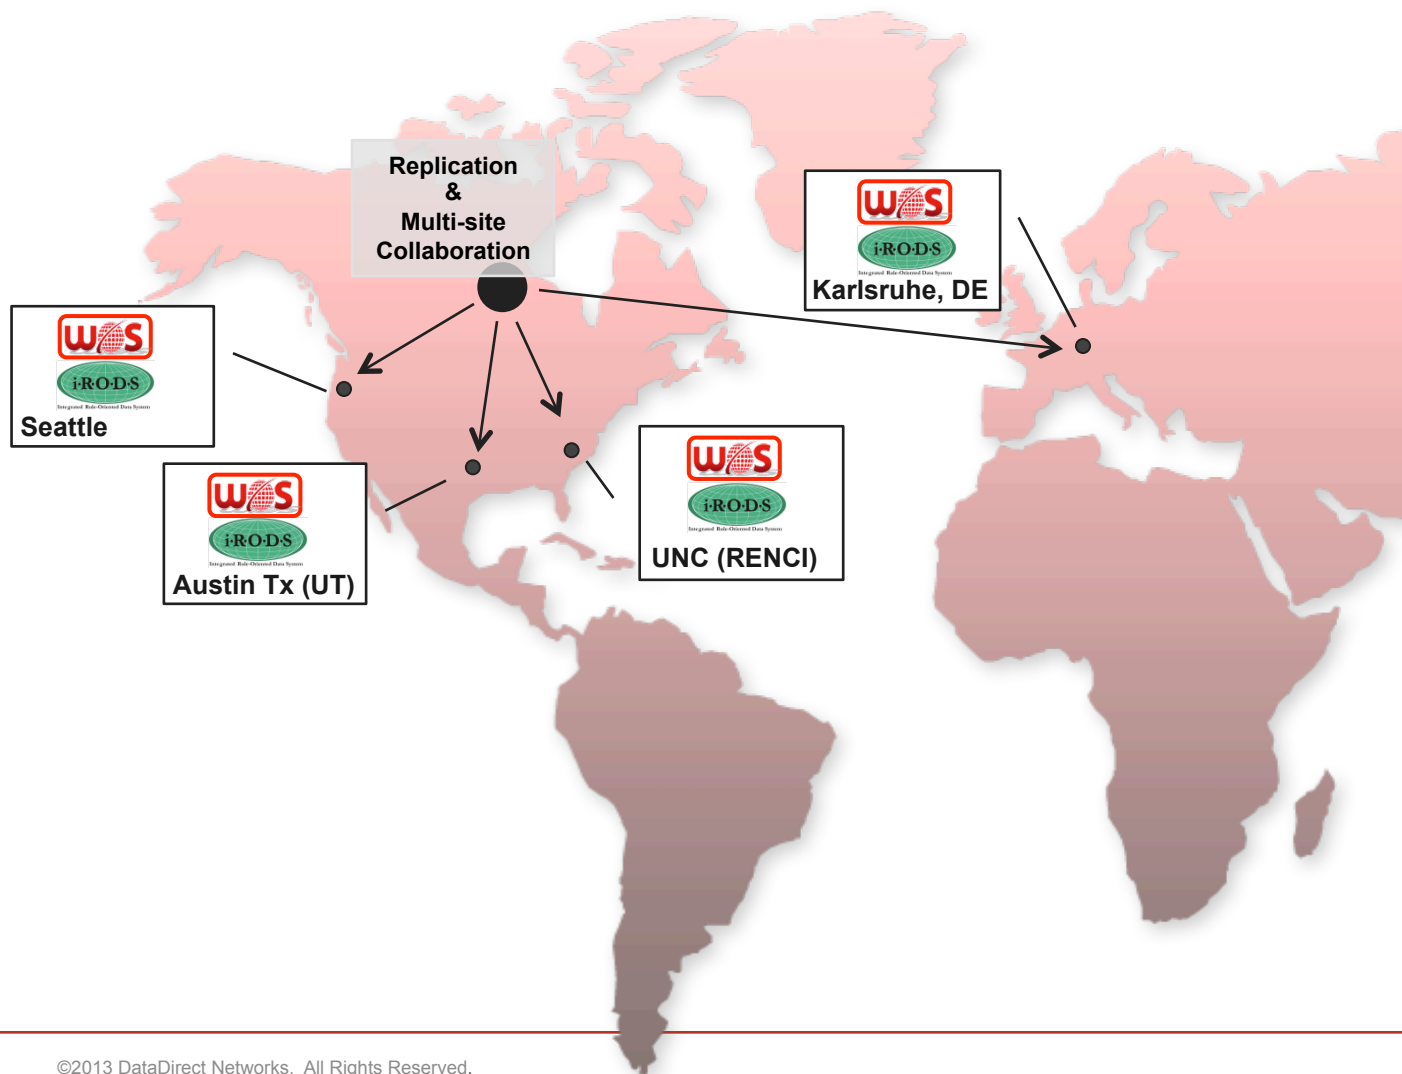
- Two data centers required for full DR
- Store/retrieve data @ full LAN line speeds
- Data replicates to remote asynchronously @ WAN Speeds
- Only retrieve data over slower WAN in DR scenarios

DDN-IRODS Summary

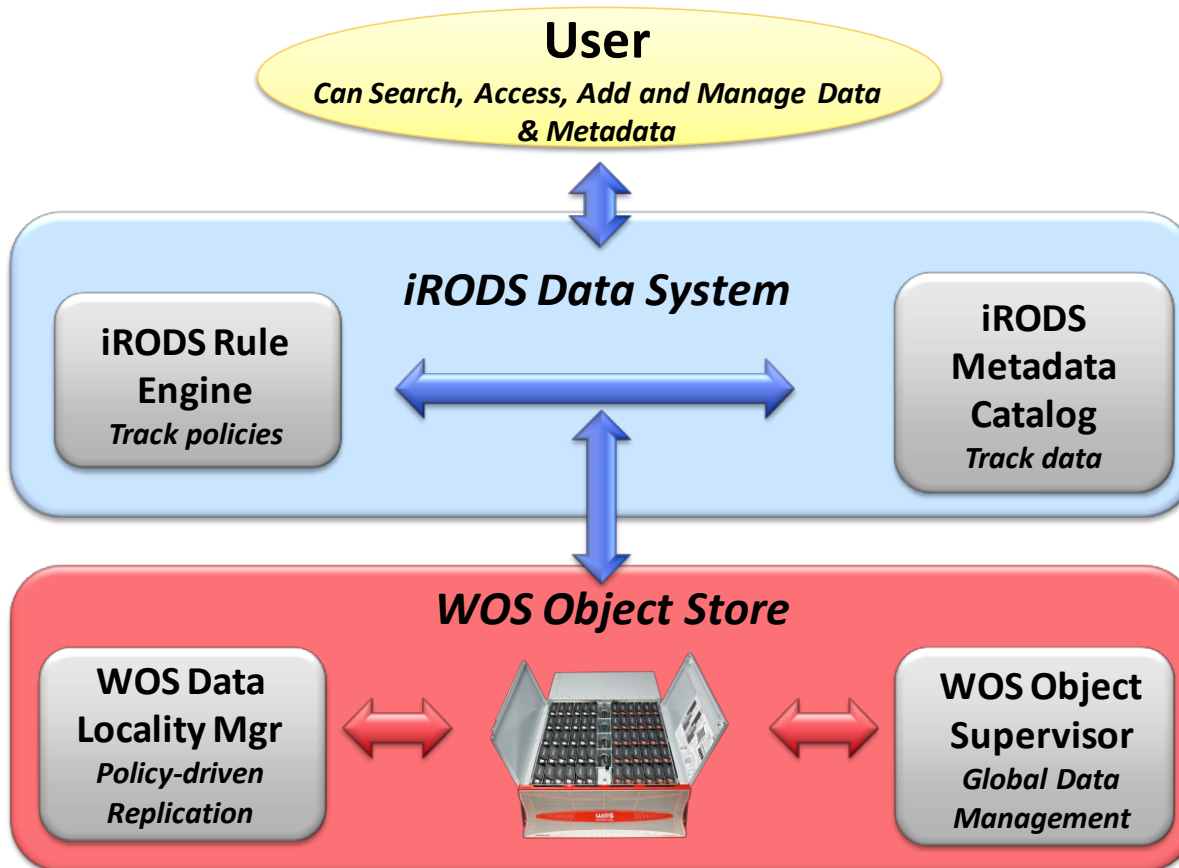
Data Center View



Building a DDN & iRODS storage Grid



Summary WOS & iRODS



WOS – iRODS Integration

- **Petabyte scalability:** Scale out by simply adding storage modules
- **Unrivalled Simplicity:** Management simplicity translates directly to lower cost of ownership
- **Self-Healing:** Zero intervention required for failures, automatically recovers from lost drives
- **Rapid Rebuilds:** Fully recover from lost drives in moments
- **Replication Ready:** Ingest & distribute data globally
- **Disaster Recoverable:** For uninterrupted transactions no matter what type of disaster occurs
- **File Layout:** Capacity and performance optimized
- **Object Metadata:** User-defined metadata makes files smarter