IDC HPC User Forum
April, 2014

DataDirect™ NETWORKS
INFORMATION IN MOTION™

WOS

Web Object Storage (WOS™)

Distributed Hyperscale Collaborative Storage

Mike Vildibill

VP, Product Management, and
Emerging Technologies Development

Big Data & Cloud Infrastructure

DDN Product Portfolio

DataDirect NETWORKS
INFORMATION IN MOTION

Infinite Memory Engine™
Distributed File System Buffer Cache

Big Data Platform Management

DirectMon

Analytics Reference Architectures

SAP
sas
hadoop

Petascale Lustre® Storage

EXAScaler™
10Ks of Clients
1TB/s+, HSM
Linux HPC Clients
NFS & CIFS

Enterprise Scale-Out File Storage

GRIDScaler™
~10K Clients
1TB/s+, HSM
Linux/Windows HPC Clients
NFS & CIFS

Cloud Tiering

Cloud Foundation

CIFS/NFS & S3/Swift

WOS® 3.0
32 Trillion Unique Objects
Geo-Replicated Cloud Storage
256 Million Objects/Second
Self-Healing Cloud
Parallel Boolean Search

WOS7000
60 Drives in 4U
Self-Contained Servers

Storage Fusion Architecture™ Core Storage Platforms

SFA12KX/E
40 GB/s
1.7M IOPS
1,680 Drives,
Embedded Option

SFA7700
Midrange Block, File & Object Platform:
- 13 GB/s
- 600K IOPS

NextGen SFA

Flexible Drive Configuration

SATA

SAS

SSD

SFX Automated Flash Caching

Adaptive cache with application & filesystem support for pre-staging, alignment & by-pass

©2013 DataDirect Networks. All Rights Reserved.

WOS Core Feature Baseline

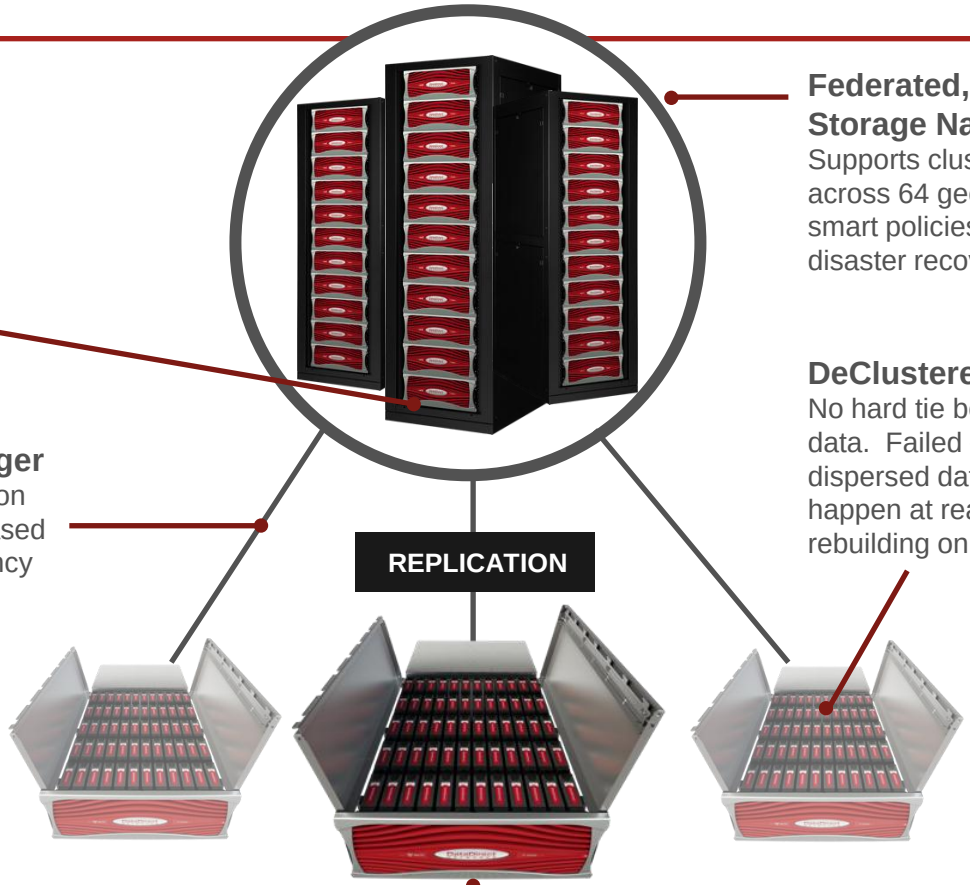

Federated, Global Object Storage Namespace
Supports clustering of up to 256 nodes across 64 geographies, replicate data with smart policies for performance and/or disaster recovery on a per-object basis

Flexible Data Protection
Chose between replication, erasure coding or both

DeClustered Data Protection
No hard tie between physical disks and data. Failed drives are recovered through dispersed data placement – rebuilds happen at read, not write, speed – rebuilding only data.

Latency-Aware Access Manager
WOS intelligently makes decisions on the best geographies to get from based upon location access load and latency

Fully-Integrated Object Storage Appliance
60 Drives in 4U

Object Disk Architecture
Formats drives with custom WOS disk file system, no Linux file I/Os, no fragmentation, fully contiguous object read and write operations for maximum disk efficiency

User Defined Metadata and Metadata Search
Applications can assign their own metadata via object storage API, WOS now also supports simple search of WOS user metadata

©2013 DataDirect Networks. All Rights Reserved.

Distributed Hyperscale Collaborative Storage

Global View, Local Access

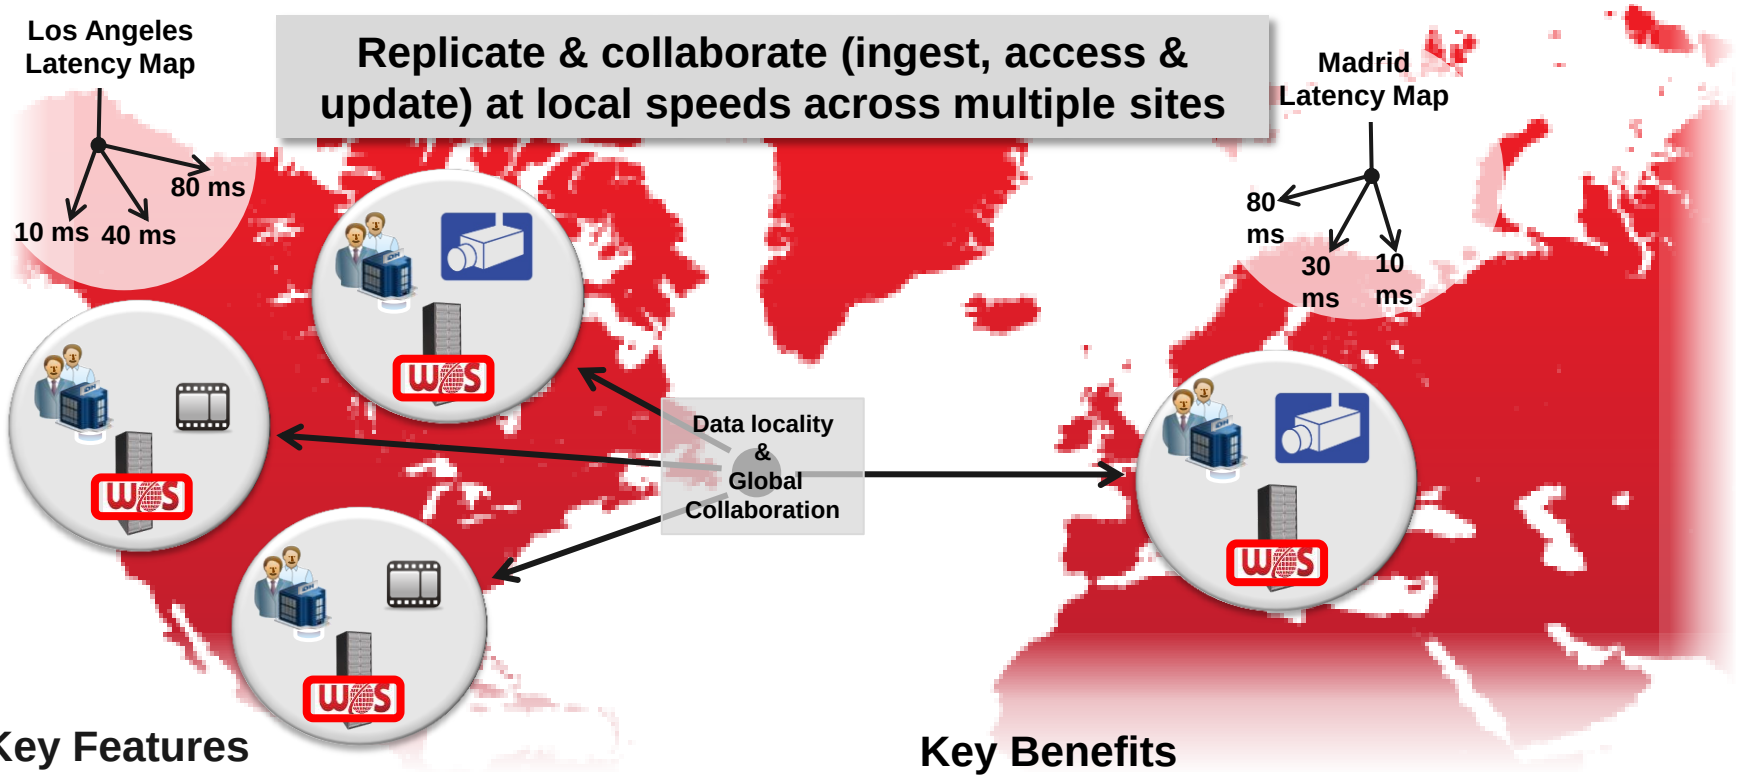

Key Features

- Asynchronous or Synchronous Replication across up to 4 sites
- Geographic, location, & latency intelligence
- NAS data access @ LAN speeds
- Data and DR protected

Key Benefits

- Users can access and update data simultaneously across multiple sites
- Increased performance & optimized access latency
- No risk of data loss

©2013 DataDirect Networks. All Rights Reserved.

WOS Reliability
Complete Choice of Protection Schemes

DataDirect NETWORKS
Information in Motion

3 Copies

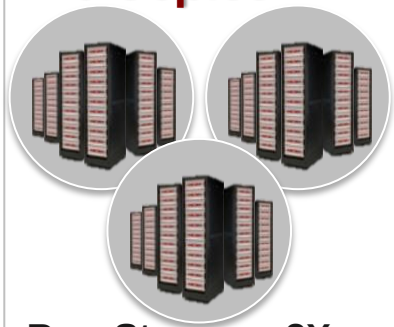

Raw Storage : 3X

- Performance
- Efficiency
- Reliability
- Scalability

Use Case:
- Distributed Collaboration

Local Copy OA

Raw Storage: 1.25X

- Performance
- Efficiency
- Reliability
- Scalability

Use Case:
- Low cost local centralized storage

Replicated OA

Raw Storage: 2.5X

- Performance
- Efficiency
- Reliability
- Scalability

Use Case:
- High throughput, streaming media, collaboration

Global OA

Raw Storage: <1.88X

- Performance
- Efficiency
- Reliability
- Scalability

Use Case:
- Archives, Clouds

©2014 DataDirect Networks. All Rights Reserved.

DataDirect™
NETWORKS
INFORMATION IN MOTION™

WOS Access: NFS, CIFS, S3, SWIFT, …

©2013 DataDirect Networks. All Rights Reserved.

WOS Access NFS / CIFS
Features & Value Proposition

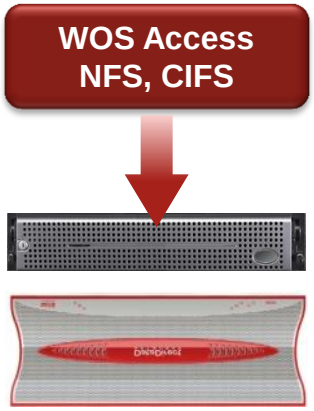

WOS Access
- NFS v3 & V4
- CIFS SMB 1 & 2
- HA & DR Protected
- Namespace sync between gateways
- Access Controls
- Built–in DR
- Simple UI for administration

WOS Access NFS/CIFS Features

- Scalable to multiple gateways across multiple sites
- Access “closest” instance of data via NFS/CIFS
- Synchronized NFS/CIFS namespace across remote sites
- Local read & write cache

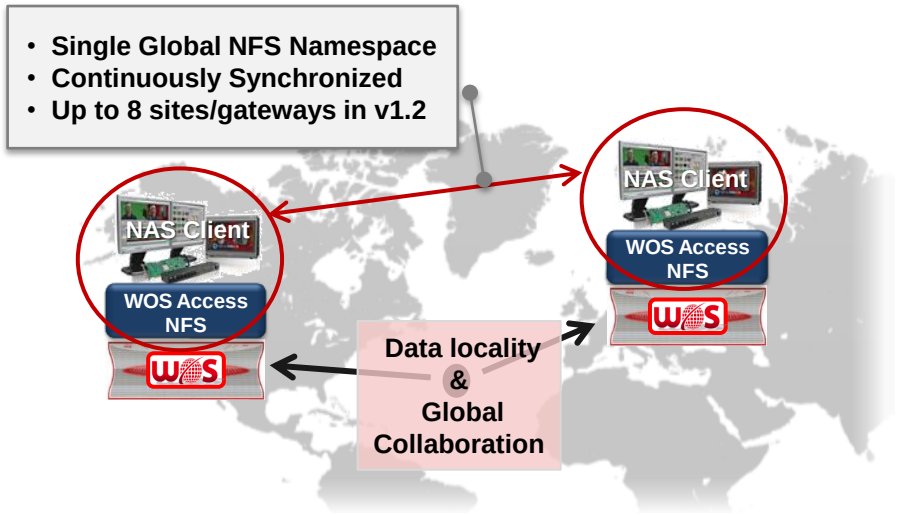

With WOS & WOS Access Gateway, distributed users can collaborate at local speeds across multiple locations using NAS protocols

©2013 DataDirect Networks. All Rights Reserved.

NFS & CIFS Scale Out

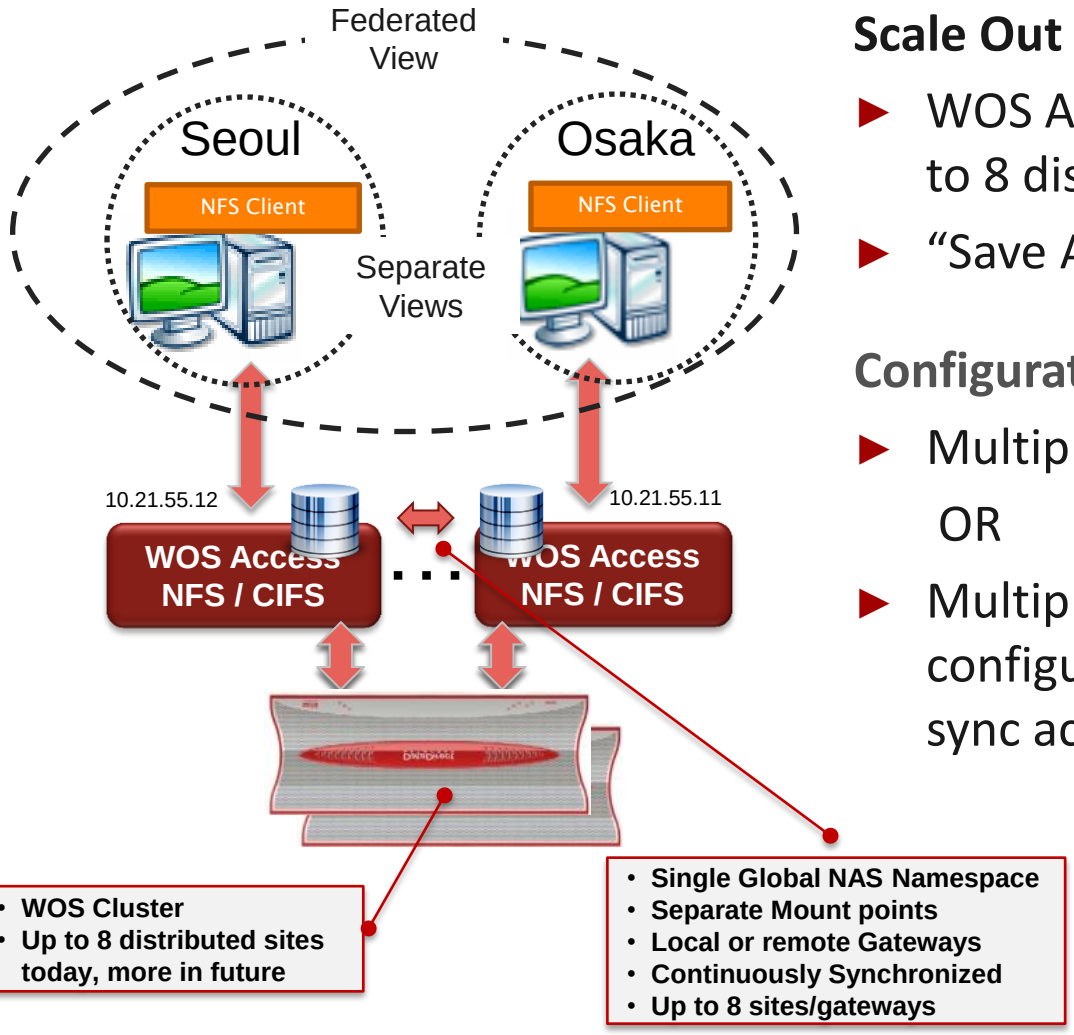

Scale Out + Collaboration

- WOS Access NFS/CIFS can scale out to up to 8 discrete gateways in same cluster
- "Save As" writes across multiple sites

Configuration Options

- Multiple sites can be federated

  OR

- Multiple separate mount points can be configured on single WOS cluster (no DB sync across clusters)

©2013 DataDirect Networks. All Rights Reserved.

WOS Access NFS

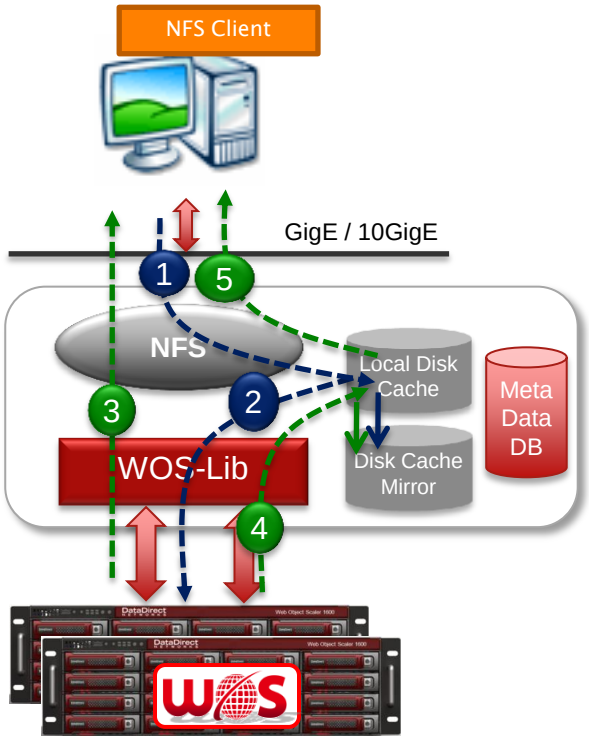

NFS Writes

1. NFS applications write files directly to gateway local disk cache & mirror
2. When File is closed, cache is flushed to WOS via WOS-Lib. Local cache will be managed per user policy (high water marks, LRU, Security, etc)

NFS Direct Read Mode

3. Files are read directly from WOS (typically faster) by default

NFS Cached Read Mode (option)

4. If file is in GW local NFS disk cache, it will be accessed directly from cache
5. If file is not in GW cache, file is read from nearest WOS into local gateway cache and then returned NFS Client

©2013 DataDirect Networks. All Rights Reserved.

WOS Access CIFS

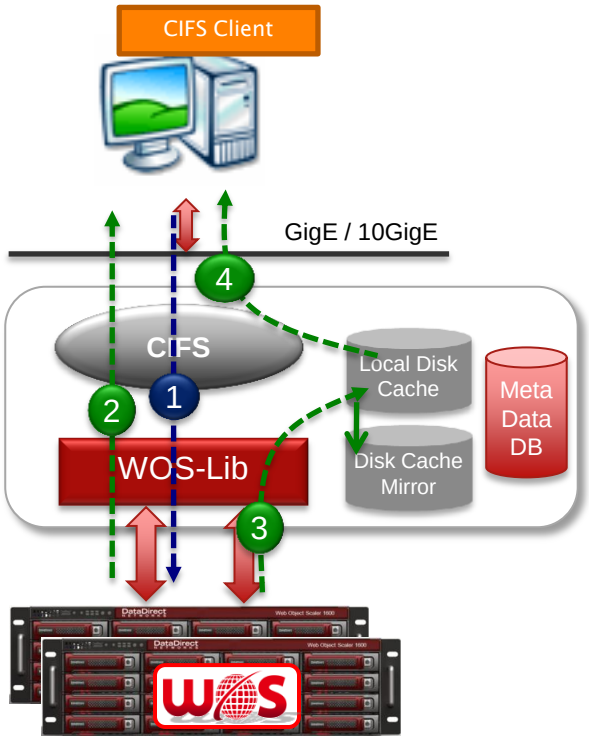

CIFS Writes

1. CIFS applications write files directly to WOS

NFS Direct Read Mode

2. Files are read directly from WOS (typically faster) by default

NFS Cached Read Mode (option)

3. If file is in GW local NFS disk cache, it will be accessed directly from cache
4. If file is not in GW cache, file is read from nearest WOS into local gateway cache and then returned NFS Client

©2013 DataDirect Networks. All Rights Reserved.

WOS Cloud Storage

Interfaces and Ecosystem Partners

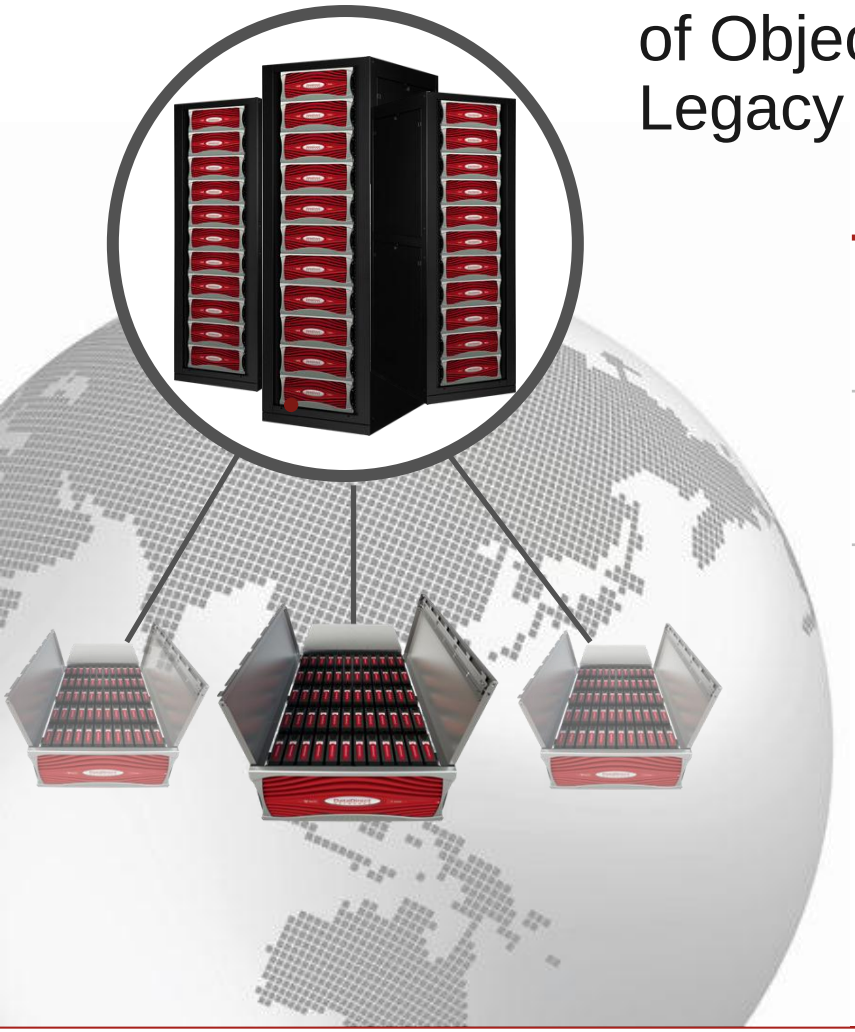

WOS's Ecosystem Enables Broad Adoption of Object Storage Technology By Both Legacy and Next-Generation Applications.

Supported Interfaces

WOS	WOS C++ API WOS Java API WOS Python API WOS PHP API HTTP w/ caching
DataDirect Networks	EXAScaler™ GRIDScaler™ WOS Access NAS (NFS/CIFS)
WOS Developer	iRODS, commvault, ctera, PIXIA CORP, Nirvana, Atempo, ownCloud, Tiger Technology, pogoplug

©2013 DataDirect Networks. All Rights Reserved.

DataDirect™ NETWORKS
INFORMATION IN MOTION™

WOS Bridge: GRIDScaler (GPFS), EXAScaler (Lustre), Ceph, Hadoop, …

©2013 DataDirect Networks. All Rights Reserved.

Enable High Speed Edge Computing with GRIDScaler WOS Bridge

- **Offload**: Offload/archive GRIDScaler (GS) files to WOS to free up space & improve performance
- **Distribute & Federate:** Replicate & federate namespaces across other GS + WOS Access sites for collaboration & disaster protection
- **Collaborate**: WOS Access & GS users at remote sites can review & update files at local LAN speeds & share with GS & WOS Access users at other sites

©2013 DataDirect Networks. All Rights Reserved.

GRID-WOS

Collaborate, Distribute and Federate

DataDirect NETWORKS Engineering Plan of Record

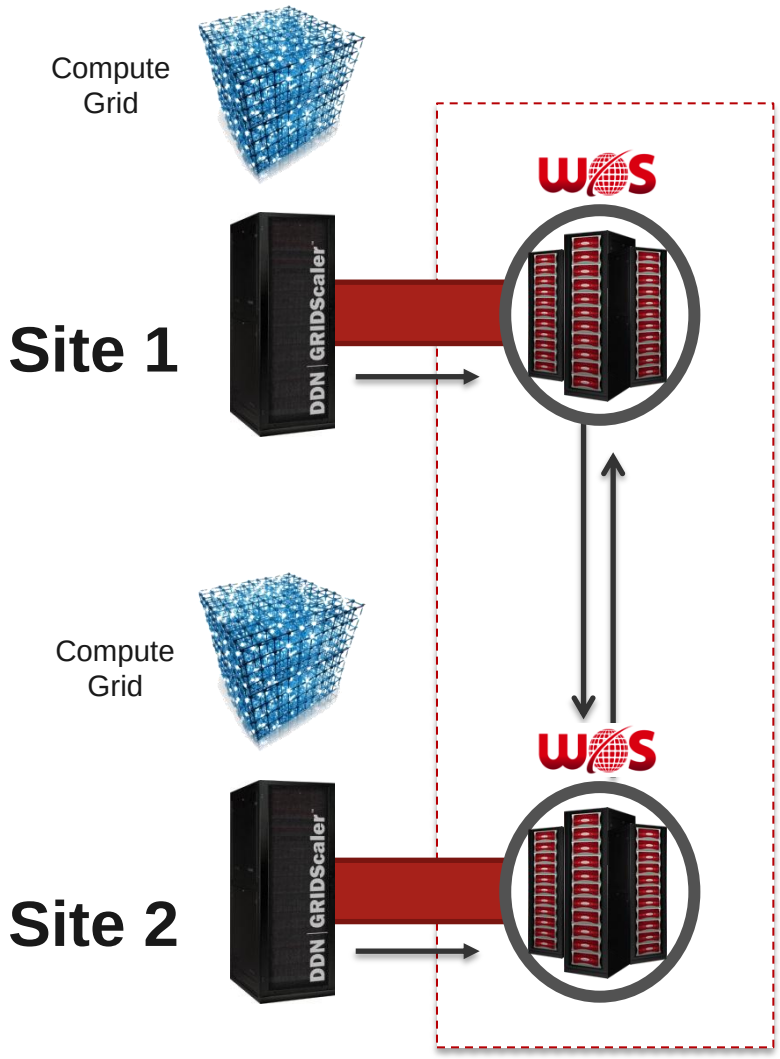

Multiple sites can work on a common data set where data is published worldwide automatically & available immediately for access via any protocol.

Overview

- Federation of GRIDScaler namespaces, using WOS as the wide area distribution mechanism
- Multi-site ingest, distribution & collaboration capacity for up to 8 GRIDScaler systems

Access Mechanisms

- Data is ingested and written into GridScaler & copied to WOS, or written directly to WOS – and then exposed as a NAS mount point
- Data is ingested into WOS via NFS or CIFS and accessible to GridScaler as a NAS mount point

©2013 DataDirect Networks. All Rights Reserved.

GS-WOS Bridge v1.5 Features
GridScaler-GridScaler-WOS Access Federation

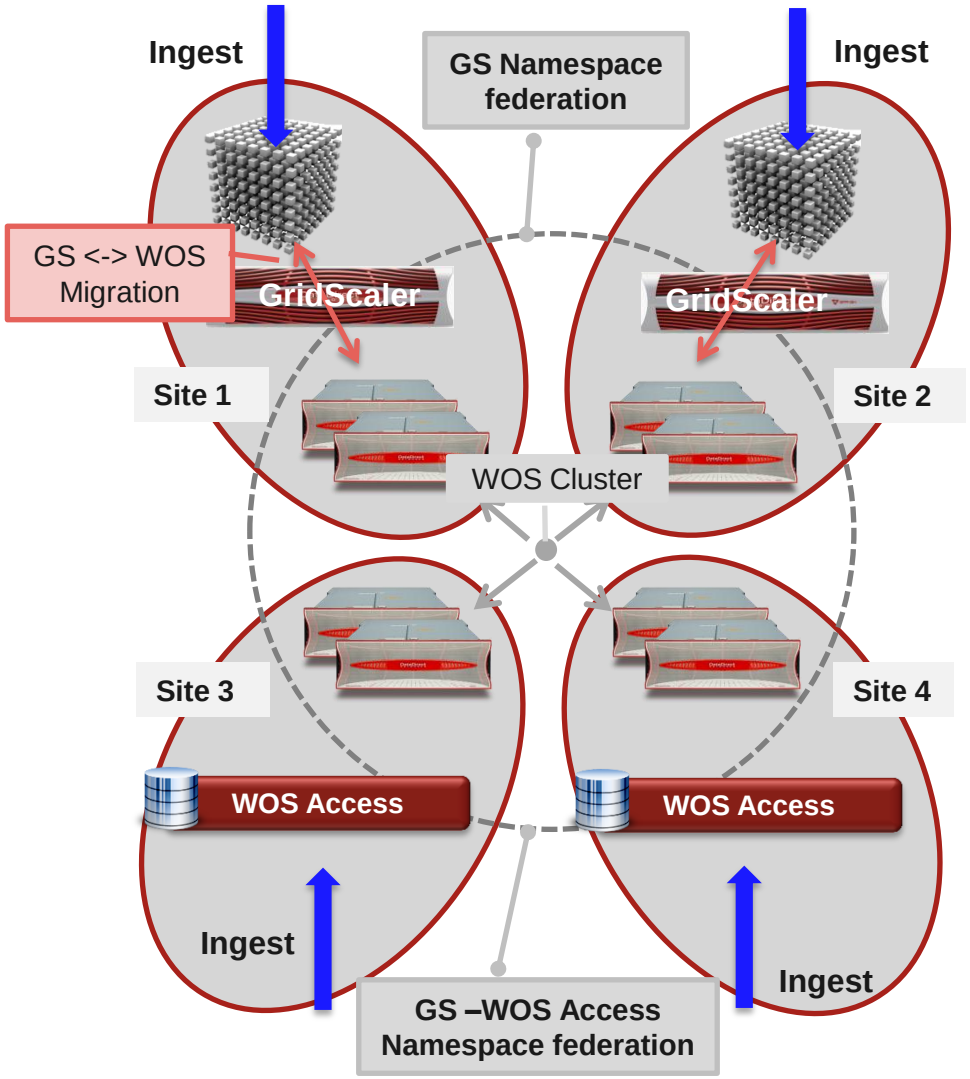

- GS<-> GS <->WOS Access Multi-site ingest & namespace federation
- Automatic migration of files to WOS via GS ILM Policies
- Write access to migrated files (automatic rehydration)
- Migrate files using multiple WOS policies
- E-platforms support qualification

©2013 DataDirect Networks. All Rights Reserved.

DataDirect™
NETWORKS
INFORMATION IN MOTION™

WOS Customer Uses Cases

©2013 DataDirect Networks. All Rights Reserved.

Customer Collaboration
SSERCA: Cross-university Collaboration

DataDirect NETWORKS
Information in Motion

► The Challenge:

- Build a storage cloud to support a state-wide community of dispersed scientists
- Support multi-site collaboration via NAS & sync and share tools

► The Solution:

- 12PB WOS Storage Cloud, distributed over 6 SSERCA sites
- Data sharing across NAS & OwnCloud
- Products: WOS Access, WOS Core & OwnCloud. Parallel filesystems in future

► The Benefits:

- WOS enables SSERCA to consolidate traditional storage layers, including backup and archive, and allows research teams to maximizing productivity and efficiency.
- WOS enables SSERCA's to store research data from multiple sources, of any type and size

©2013 DataDirect Networks. All Rights Reserved.

SSERCA: Cross-university Data Sharing

Rich Collection of Integration & Interfaces

DataDirect NETWORKS
INFORMATION IN MOTION

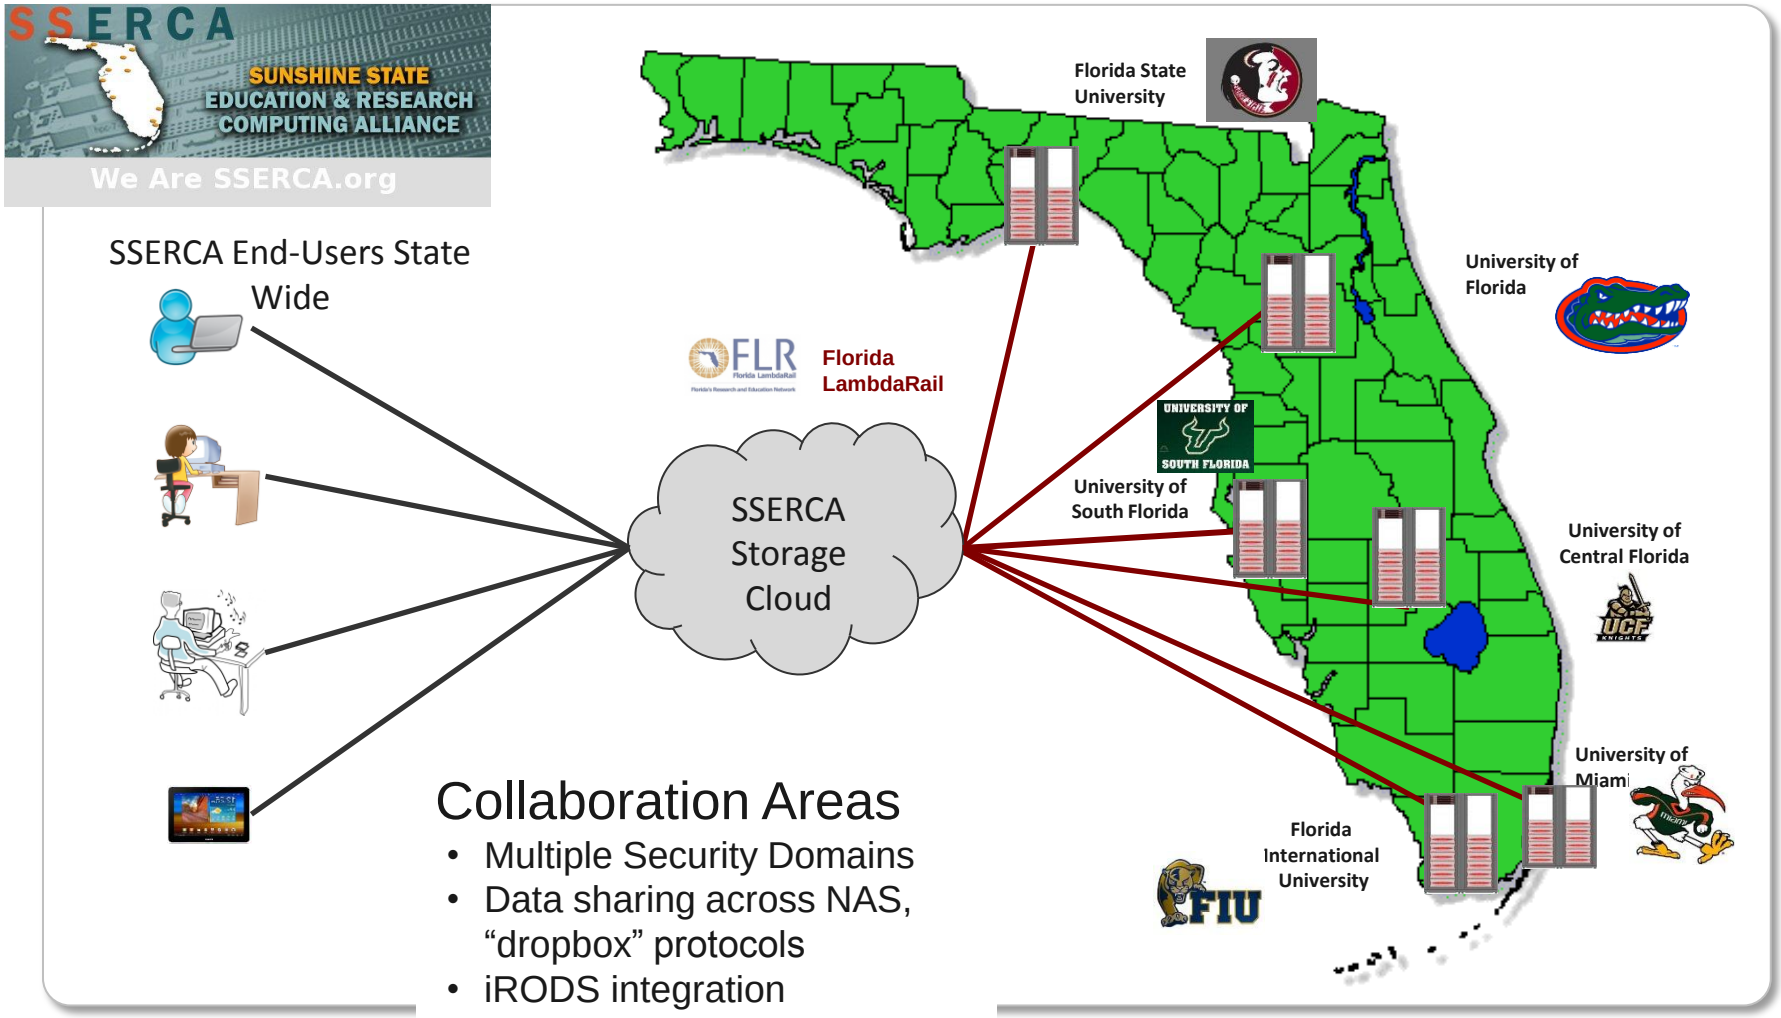

Collaboration Areas

- Multiple Security Domains
- Data sharing across NAS, “dropbox” protocols
- iRODS integration
- ES-WOS Bridge

©2013 DataDirect Networks. All Rights Reserved.

Customer Collaboration
UCL: Accelerate Media Workflows

DataDirect NETWORKS
Information in Motion

► The Challenge:

- Share project-based research data
- Scale in both volume and velocity

UCL

► The Solution:

- GRIDScaler File Storage provides high performance and low latency required for their HPC applications
- WOS provides unlimited scalability and local collaboration

► The Benefits:

- Collaborative platform to accelerate workflows & time-to-discovery
- Highest storage density enables future scalability within their expensive, limited downtown London floor space

©2013 DataDirect Networks. All Rights Reserved.

UCL: Accelerate Media Workflows

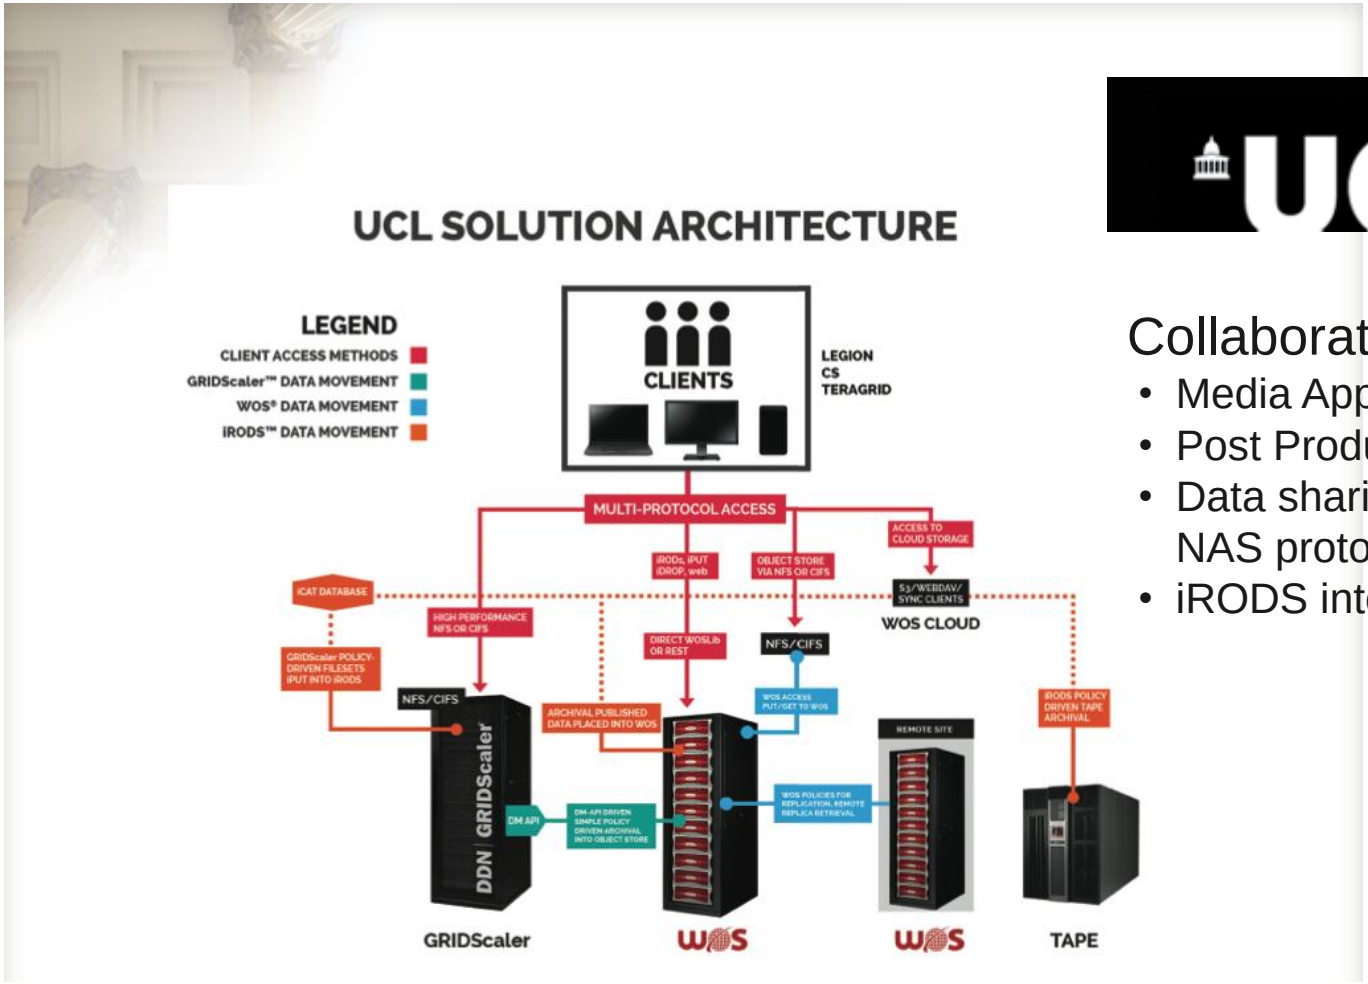

UCL

Collaboration Areas

- Media App Integration
- Post Production via GS
- Data sharing between NAS protocols
- iRODS integration

©2013 DataDirect Networks. All Rights Reserved.