



# **DMF/S3/WOS Testing DMFUG Feb 2015**

Feb, 2015

DDN Australia

## Acknowledgements

- ▶ **Inspired by testing undertaken by SGI in January**
  - Rob Mollard – SGI
  - Susheel Gokhale – SGI
  - Donald Yan - DDN

## DDN | About Us

### 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
- ▶ Worldwide Presence: 20 Countries
- ▶ Installed Base: 1,000+ End Customers; 50+ Countries
- ▶ Go To Market: Global Partners, Resellers, Direct



### World-Renowned & Award-Winning



**Inc.**

**Gartner**

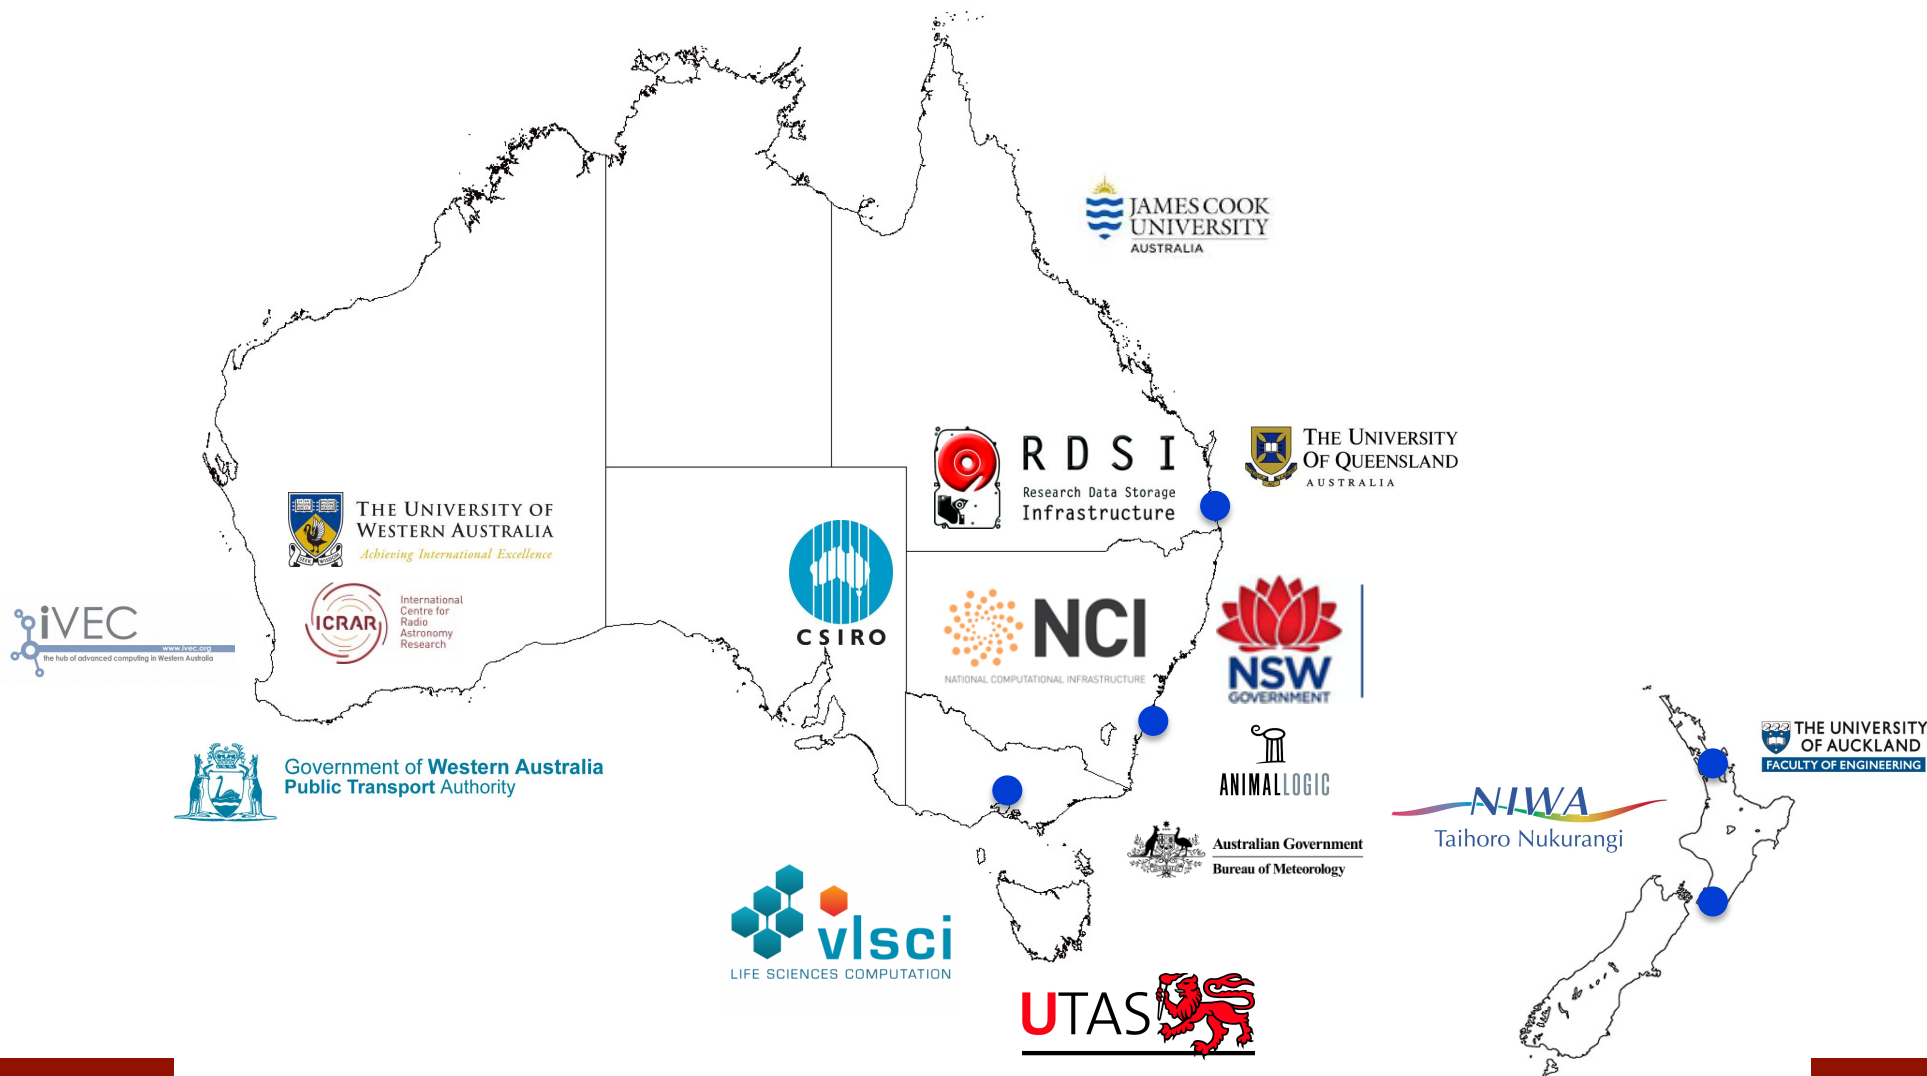
the **451** group

**HPC** WORLD

**STORAGE**

**Federal Computer Week**

# Sample Customers Australia & New Zealand



Functionality test:

Can DMF use DDN's WOS platform with S3 as a storage tier

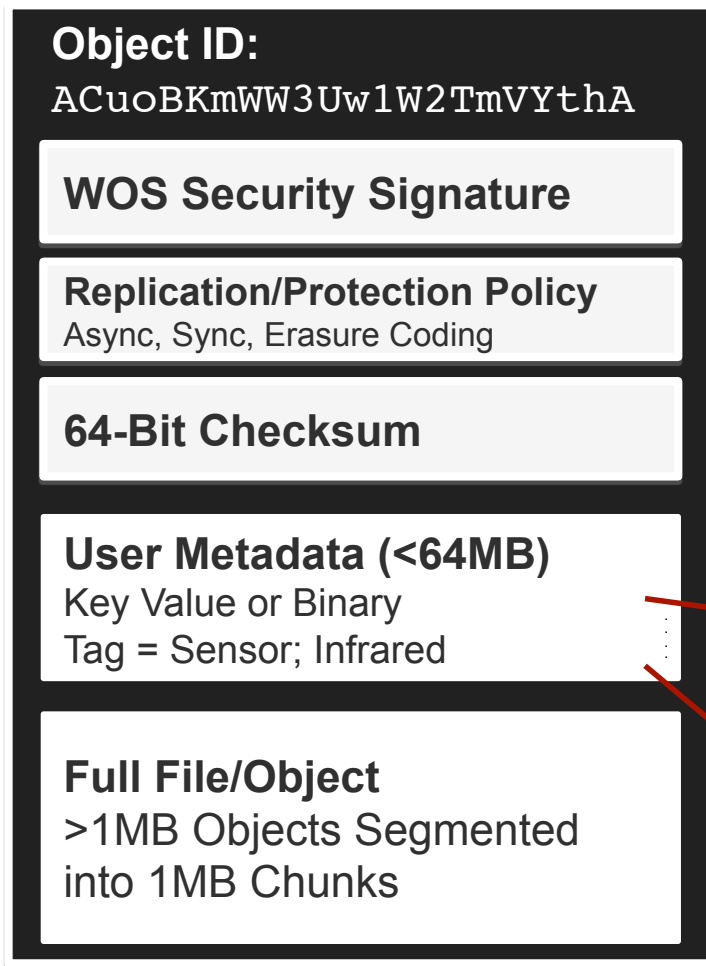
## What is WOS ? (Web Object Scaler)



- ▶ WOS is an object storage platform that enables organizations to build scale-out storage clouds
- ▶ WOS storage nodes can be distributed geographically to build a global storage cloud
- ▶ Data is stored as objects, with an object ID and metadata in a flat namespace
- ▶ A WOS storage cloud is built with pre-installed WOS appliances; intelligent storage containers
- ▶ It is possible to deploy a fully functional storage cloud with just one WOS7000 appliance, and scale as needed

# What is object storage? What is an object?

## WOS Object



- Stores objects not files
- Simple flat structure, no directory hierarchy
- Objects are immutable
- Put, get and delete

Search Example: Return all OIDS with:  
Sensor =infrared AND Latitude < 40 AND Longitude > -50

Key	Value
sensor	infrared
latitude	18.33
longitude	-45.67
date	19 Jan 2012

Indexed fields

# WOS: Efficiency


## High-Density Platform

### ► WOS7000 Platform

- Up To 60 Drives: SAS, SATA, SSD
- 4 x 10GigE, 40GigE, FDR IB
- 2 Object Storage Servers (OSS)
  - Each with 2 CPUs, w/8 cores each, 128GB memory
- Dual port access to drives for high availability

### ► WOS7000 Performance

- Up to 31K Objects per second (Read)
- Up to 2.4GB/s throughput (Read)





# WOS: Scalability

## Exabyte Scalability

Virtually limitless scalability ...

WOS Cluster (max 256 nodes):



### Cluster Maximums

- 256 Nodes
- 7,680 HDDs
- 1T Objects

WOS clusters have 32-bit cluster identifiers which enable customers to build much larger namespaces, today federating 32 clusters into one namespace.

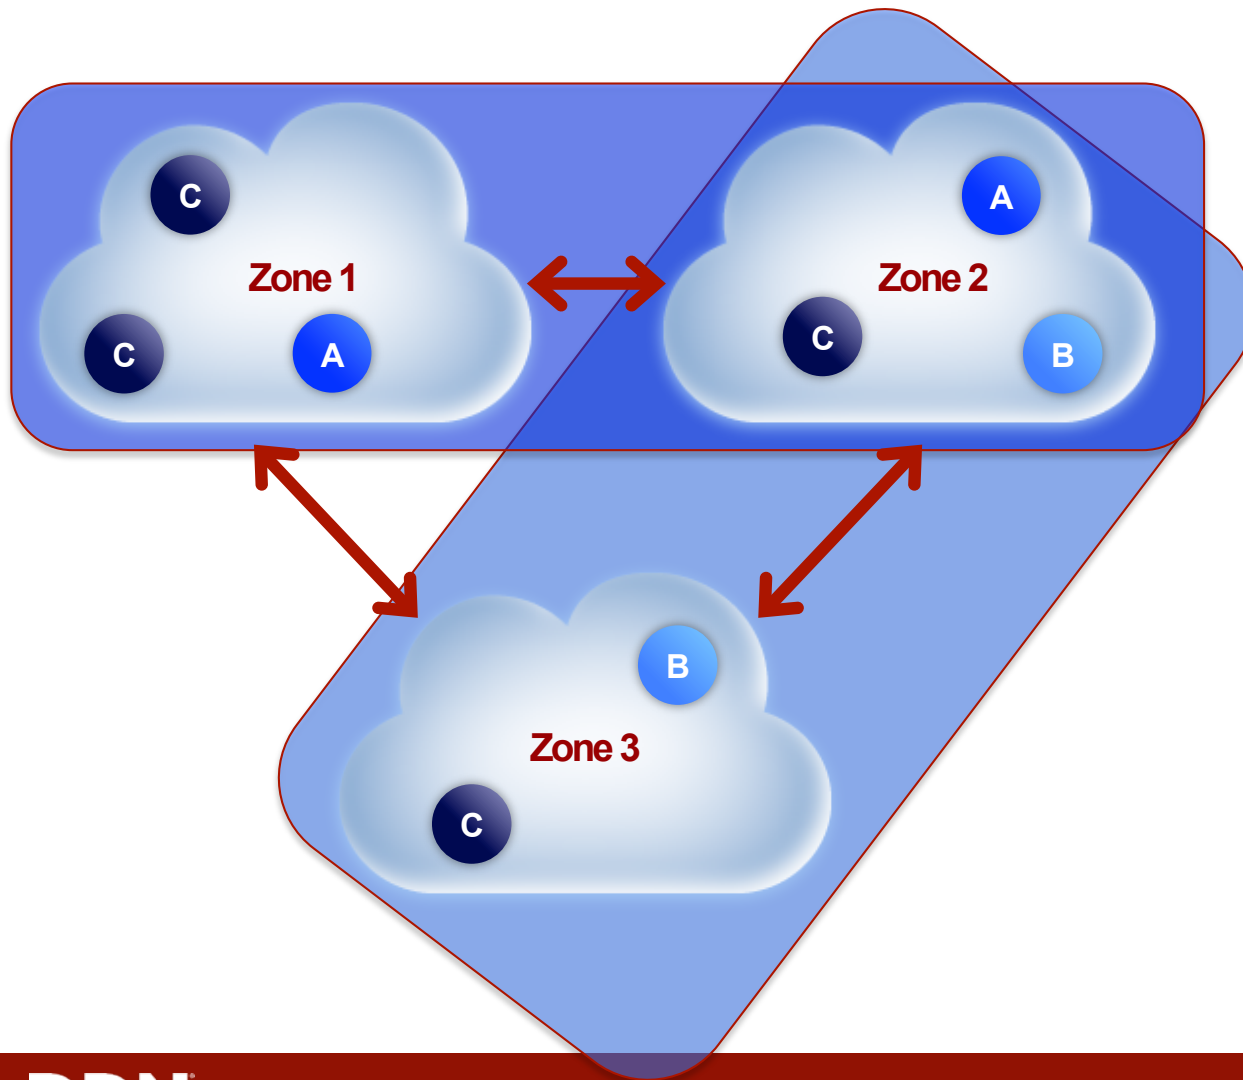
WOS Namespace (max 32 clusters):



### Namespace Maximums:

- 8192 Nodes
- 245,768 HDDs
- 32T Objects

# WOS<sup>®</sup> Zones & Policy Relationship

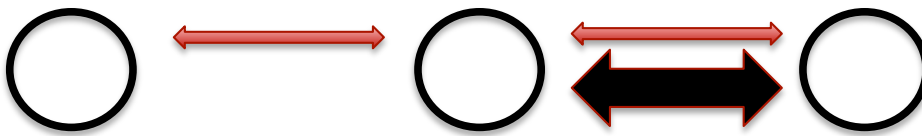


WOS is composed of a number zones containing nodes. A policy refers to the number of replicas an object should have, in each zone and the method that those replica copies should be made. It is a requirement that all nodes in an WOS cluster be reachable from each other.


<b>Policy A</b>	Zone 1 = 1 Zone 2 = 1
<b>Policy B</b>	Zone 2 = 1 Zone 3 = 1
<b>Policy C</b>	Zone 1 = 2 Zone 2 = 1 Zone 3 = 1

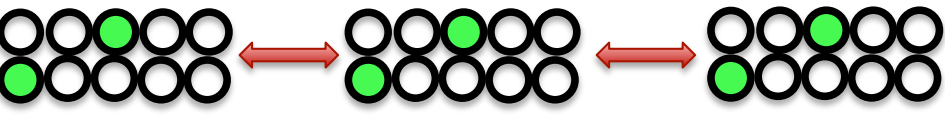
# WOS - Reliability & Policy Relationship

2x+E  OK reliability, Resistant to geographic failure.

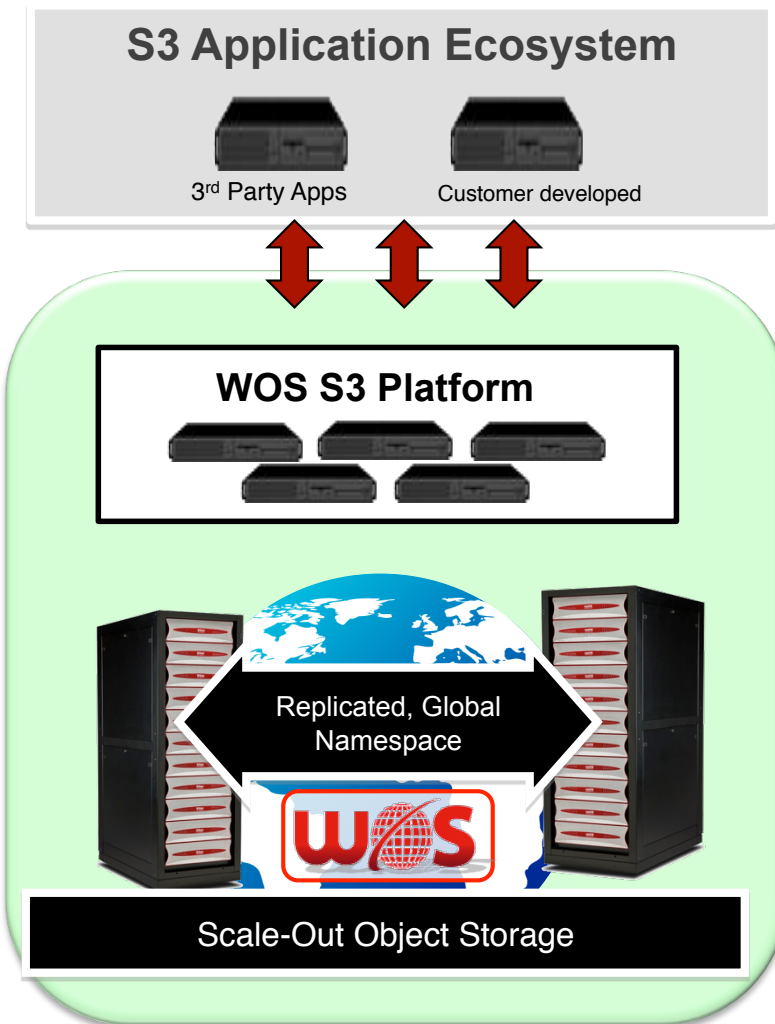
3x+E  Good reliability, Good resistance to geographic failure, large foot print.

OA  
1.25x  Excellent reliability, Not resistant to geographic failure, small foot print.

2xOA  
2.5x  Extremely reliable, Resistant to geographic failure, reasonable foot print.

GOA  
< 1.9  Extremely reliable, Resistant to geographic failure, small foot print, vulnerable to network failure.

# WOS Access S3 Product Introduction



## WOS S3

- ▶ **Scalable & Configurable S3 compatible storage System**
- ▶ **DDN developed & supported**
- ▶ **Embedded or discrete gateway options**
- ▶ **HA & DR protected**

## WOS S3 Components

### ▶ WOS S3 Protocol Nodes

- Namespace & protocol stack
- Highly available, DR protected
- Scalable to hundreds of thousands of users
- Accessed via S3 internal, open source, or partner applications

### ▶ WOS Storage Nodes

- Start small, scale to tens of petabytes
- Replicated, global namespace
- Self healing, always available
- Multi-site for DR & fastest accessibility

# Testing Setup



1U Xeon server

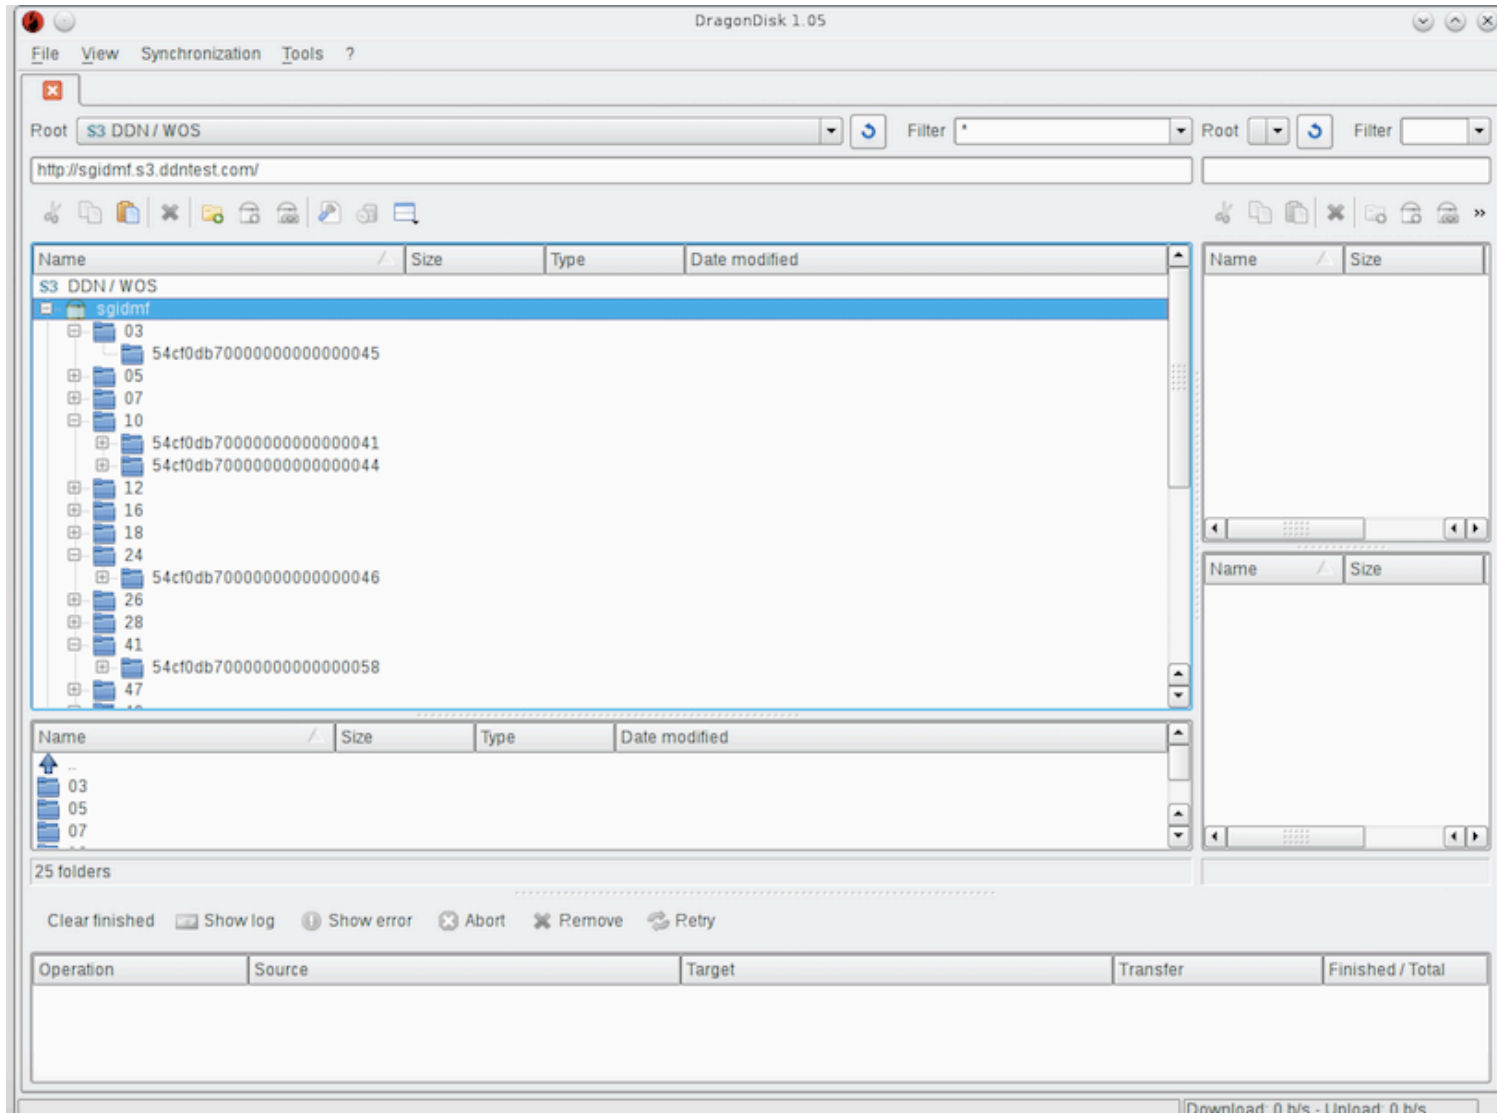
- SLES11SP3
- DMF 6.3



## Simple setup – dmf.conf

	Amazon S3	DDN S3
TYPE	msp	msp
COMMAND	dmcloudmsp	dmcloudmsp
CLOUD_BACKEND	s3	s3
CLOUD_HOST	s3.amazonaws.com	s3.ddntest.com
CLOUD_USER	AKIAJZD4UFQZKCKDSF VQ	AKIA1XF8L30P22H9D56 H
CLOUD_PASSWORD	HOME_DIR/password	HOME_DIR/password
CLOUD_BUCKET	sgidmf	sgidmf
NAME_FORMAT	%p%p/%b	%p%p/%b
CLOUD_REDUNDANCY	standard	standard
MESSAGE_LEVEL	6	6

# “sgidmf” bucket



## Outcome

- ▶ **No functional difference between Amazon S3 and DDN/WOS/S3**

## Conclusions

- ▶ **Private cloud with WOS is a viable option alongside:**
  - Tape, Amazon, local disk.



**DataDirect**<sup>™</sup>  
NETWORKS

**Thank you**

# WOS Access S3 API support

S3 API Call	DDN S3	Cleversafe	RIAK CS	CEPH	EUCALYPTUS
Authentication	Supported	Supported	Supported	Supported	Supported
PUT Objects	Supported	Supported	Supported	Supported	Supported
DELETE Objects	Supported	Supported	Supported	Supported	Supported
GET Objects	Supported	Supported	Supported	Supported	Supported
Create Buckets	Supported	Supported	Supported	Supported	Supported
Delete Buckets	Supported	Supported	Supported	Supported	Supported
Get Bucket Info (Head)	Supported	Supported	Supported	Supported	Supported
List Bucket	Supported	Supported	Supported	Supported	Supported
Multipart Upload	Supported	Not Supported	Partial Support	Partial Support	Not Supported
Object ACLs	Supported	Not Supported	Supported	Supported	Supported
Bucket ACLs	Supported	Supported	Supported	Supported	Supported
POST/Copy Object	Supported	Not Supported	Not Supported	Supported	Supported
Delete Multiple Obj	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Bucket Lifecycle	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Policy (Buckets, Objects)	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Bucket Website	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Bucket Location	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Bucket Notification	Planned	Not Supported	Not Supported	Not Supported	Not Supported
Bucket Object Versions	Planned	Not Supported	Not Supported	Not Supported	Supported
Bucket request payment	Not Supported	Supported	Not Supported	Not Supported	Not Supported