

October 2012

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

ECMWF HPC Workshop: Accelerating Data Management

Massively-Scalable Platforms and Solutions
Engineered for the Big Data and Cloud Era

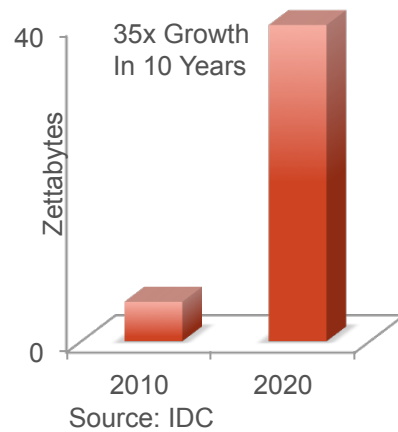
Glenn Wright

Systems Architect, DDN

ddn.com

Data-Driven Paradigm Shift

Massive Data Growth



Storing over 100PB of content¹




33x Sensor Data Growth Over Just Two Years²

Cloud & Web Ubiquity



Information, Today's Global Currency, is Shared & Distributed Globally via Cloud

Big Data Processing



"Information managers must fundamentally rethink their approach to data."

¹ Facebook: 2012 Form S-1 Registration Statement

² Defense Industry Daily: Too Much Information

DDN Big Data Leadership | Validated



“ If any company is well poised to take on the challenges of **EXAScale computing and big data**, it's DDN, since this is its heritage.”



“ Strategic Big Data HW Supplier ”
DDN Ranks #2 in 2012 CIO Survey



June 2012

#1 = IBM | #2 = DDN | #3 = EMC | #4+ = NTAP, FIO, Dell
24% of Respondents Identify **DDN** as Their **Top** Choice for Big Data

DDN | Designed for Big Data & Cloud



Massively Scalable Storage Technology



HyperScale, High Performance Platform

- ▶ DDN's Massively Scalable SFA™:S/W Engine on Enhanced H/W Platforms
- ▶ Over 1TB/s In Only 25 Systems, Millions of IOPS

Cloud Storage & Computing Infrastructure



Peer to Peer Cloud Infrastructure

- ▶ DDN's WOS™ Cloud-Based Data Delivery
- ▶ 55 Billion Objects Per Day, 100+ Locations

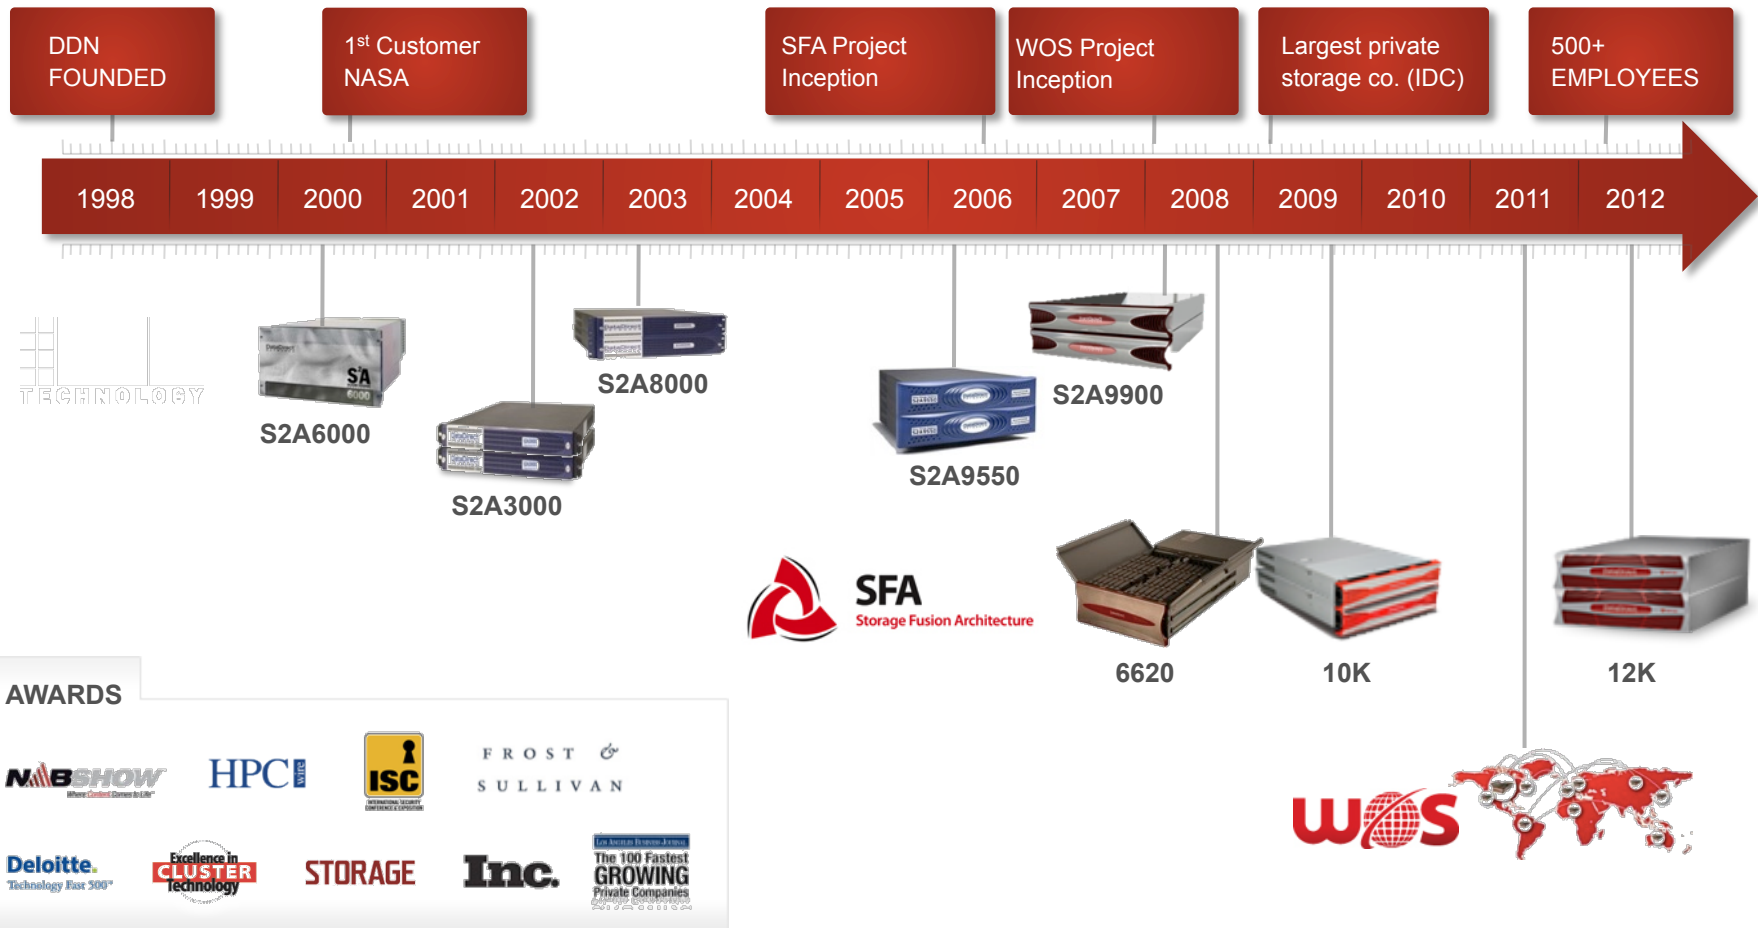
Big Data Processing for Actionable Insight



Big Data Processing System

- ▶ DDN's SFA In-Storage Processing™
- ▶ Nanosecond Latency, 16 Virtual Machines

DDN | A Decade+ Investment In Scale



The Broadest Big Data Portfolio

File Storage




SAN File Storage

xSTREAMScaler™
100s SAN/LAN Clients
HSM Capable
Streaming Optimized




Enterprise NAS

NAS Scaler™
1-16 NAS Servers
Fully-Featured
High Performance



Parallel File Storage

EXAScaler™
10Ks of Clients
1TB/s+



GRIDScaler™
1Ks of Clients
1TB/s+
Scale-Out NAS

Object Storage



WOS®
256 Billion Objects
GeoReplicated
Cloud Foundation
Mobile Cloud Access

Storage Engines

Silicon Storage Architecture

9900



6GB/s in Real-Time
1,200 Drives: 2 Racks

6620



2GB/s, 350K IOPS
120 Drives in 8U

10K




15GB/s, 1.2M IOPS
1,200 Drives: 2 Racks
Embedded Computing

12K



40GB/s/1.7M IOPS
1,680 Drives: 2 Racks
Embedded Computing

Platforms Support Any Drive Technology
to Optimize Data Intense Workflows



SAS

SATA

SSD

DDN | Storage Fusion Architecture (SFA)



Accelerating Big Data and Cloud, Optimizing TCO

Over 1 Million Lines of S/W Code – First Customer Shipped 2008
Designed Specifically for Big Data and Cloud Workloads



DirectMon™: Infrastructure Management

'Scaler File System Family

In-Storage Processing™ Engine & DMA Driver

Storage Fusion Architecture™
[Core Storage S/W Engine]

Low-Latency Connect: FC, IB, Memory

Interrupt-Free Storage Processing

ReACT™ Adaptive Cache Technology

DirectProtect™ Data Integrity Management

Quality of Service Engine

Storage Fusion Fabric™

Parallel State-Machine Design

Maximum Performance, Lowest Latency

Virtualized Processing

Optimized Environment for Big Data Application Hosting

Robust Data Protection

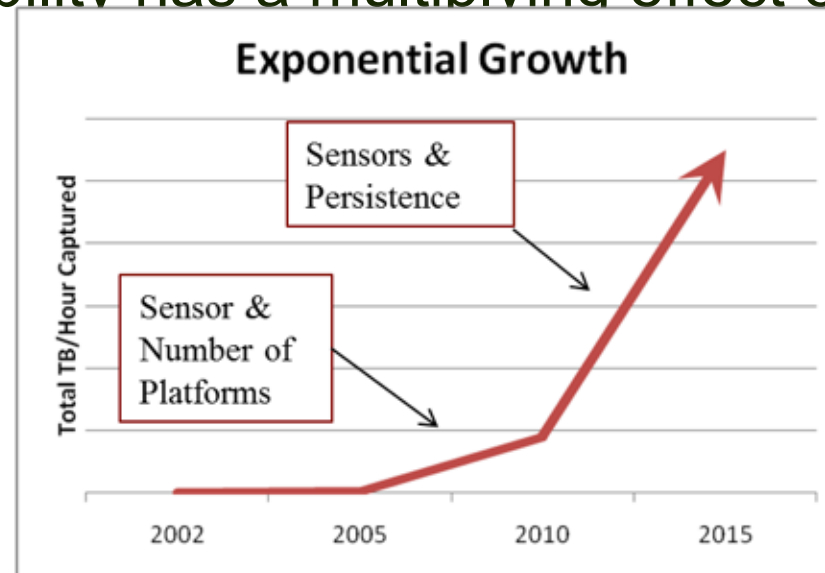
Quality of Service and Performance Without Compromise

Flexible & Massively Scalable

Best-In-Class Scalability and Density

Data Sensors: The Challenge

- Increasing amount of data captured per hour by the next generation sensors.
- Increasing number of multi-sensor, ISR platforms being developed and deployed.
- Increasing linger capability has a multiplying effect on the amount of data



Conventional Storage

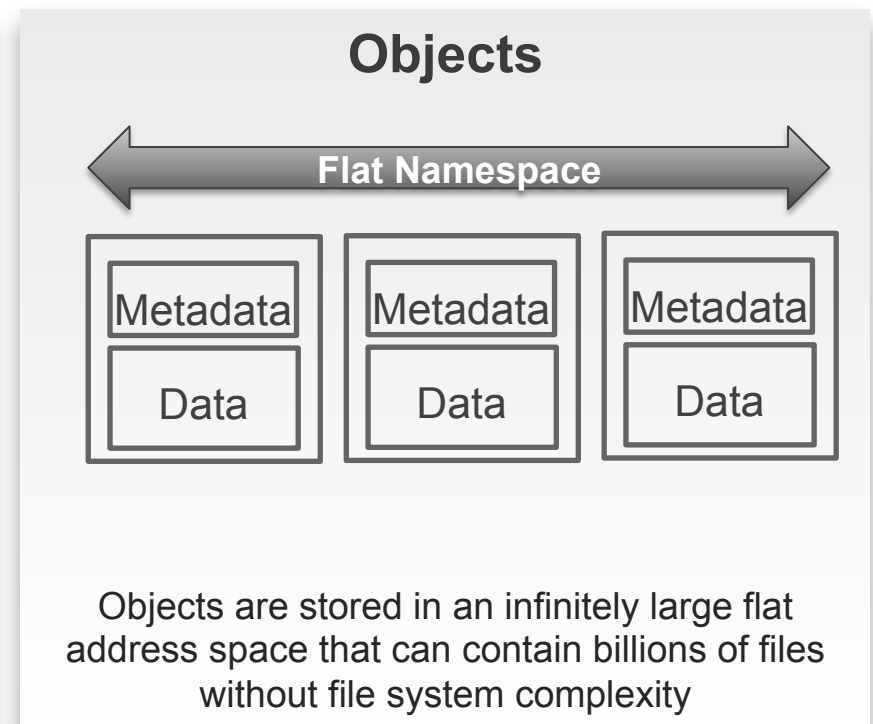
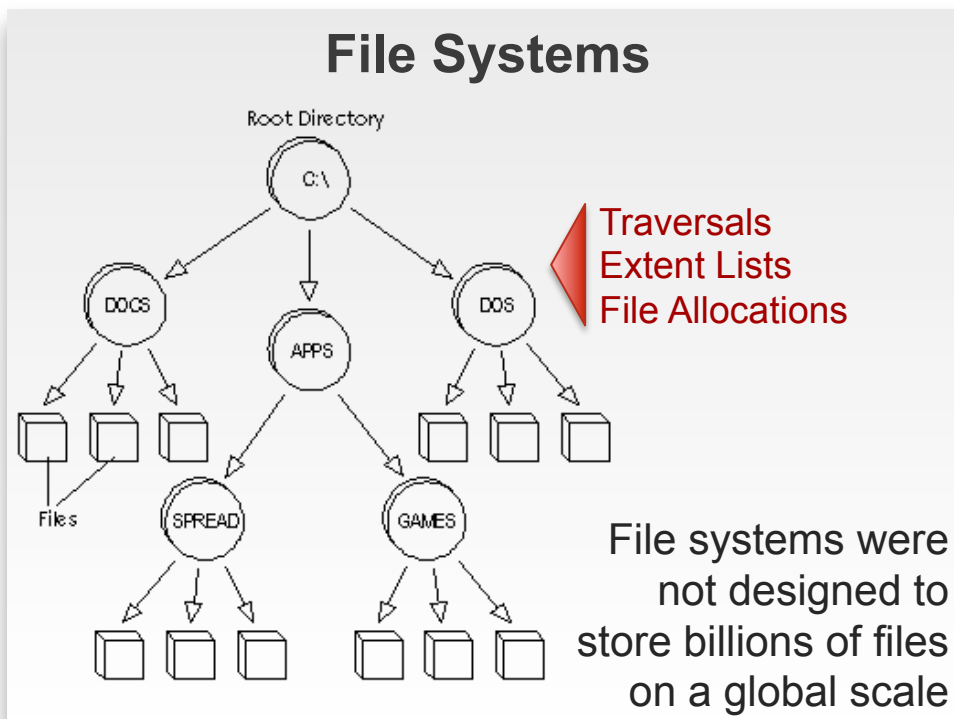
Well understood, but stretched to its limits for unstructured data cloud storage

- Introduced in the 90's
 - With 16TB file system sizes that many still have
- Management intensive
 - LUNs, Volumes, Aggregates,...
 - Heroically management intensive at scale
- Antiquated resiliency techniques that don't scale
 - RAID
 - Cluster failover, "standby" replication, backup
 - File Allocation Tables, Extent Lists
- Inefficient usable space utilization
- Focused on structured data transactions (IOPS)
 - File locking overhead adds cost and complexity
- Prohibitively expensive at cloud scale
 - To purchase, run, manage, protect and expand



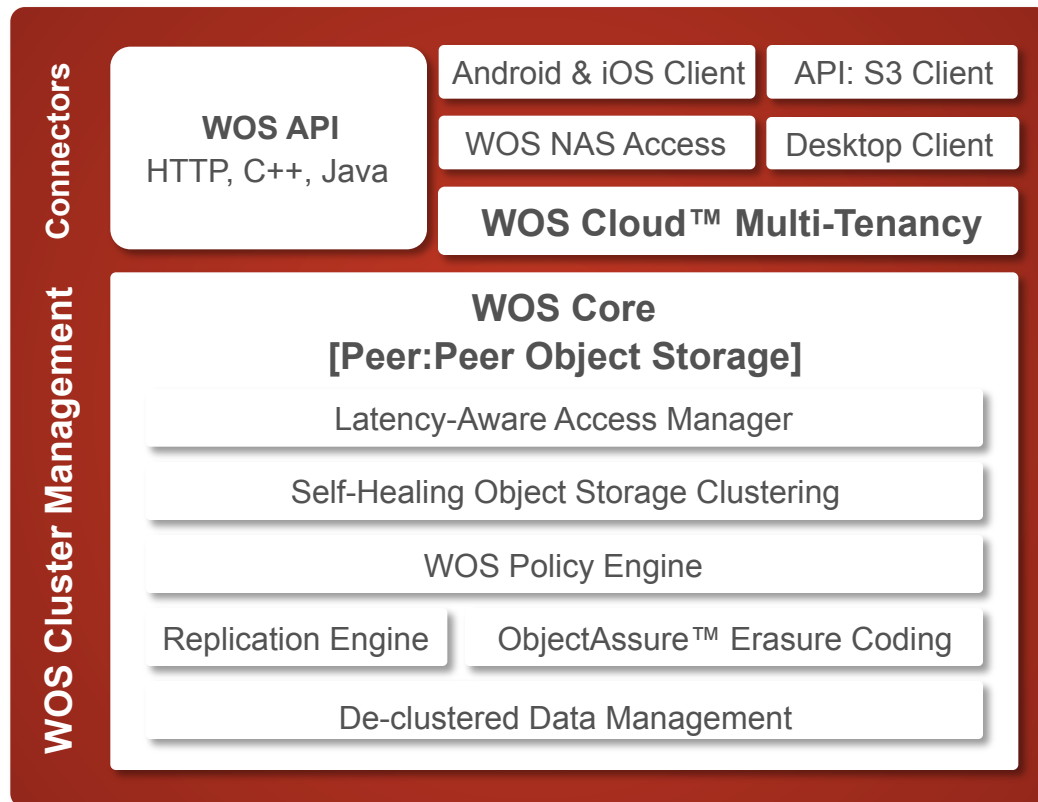
Object Storage | Explained

- Object storage stores data into containers, called objects
- Each object has both data and user defined and system defined metadata (a set of attributes describing the object)



Enabling Real-Time Global Collaboration

Web-Scale, High-Performance Cloud Storage Appliances
99% Efficiency, Petabyte-Class Peer:Peer Technology



Limitless Scalability

Eliminates the limitations of traditional file systems

Store Objects Intelligently

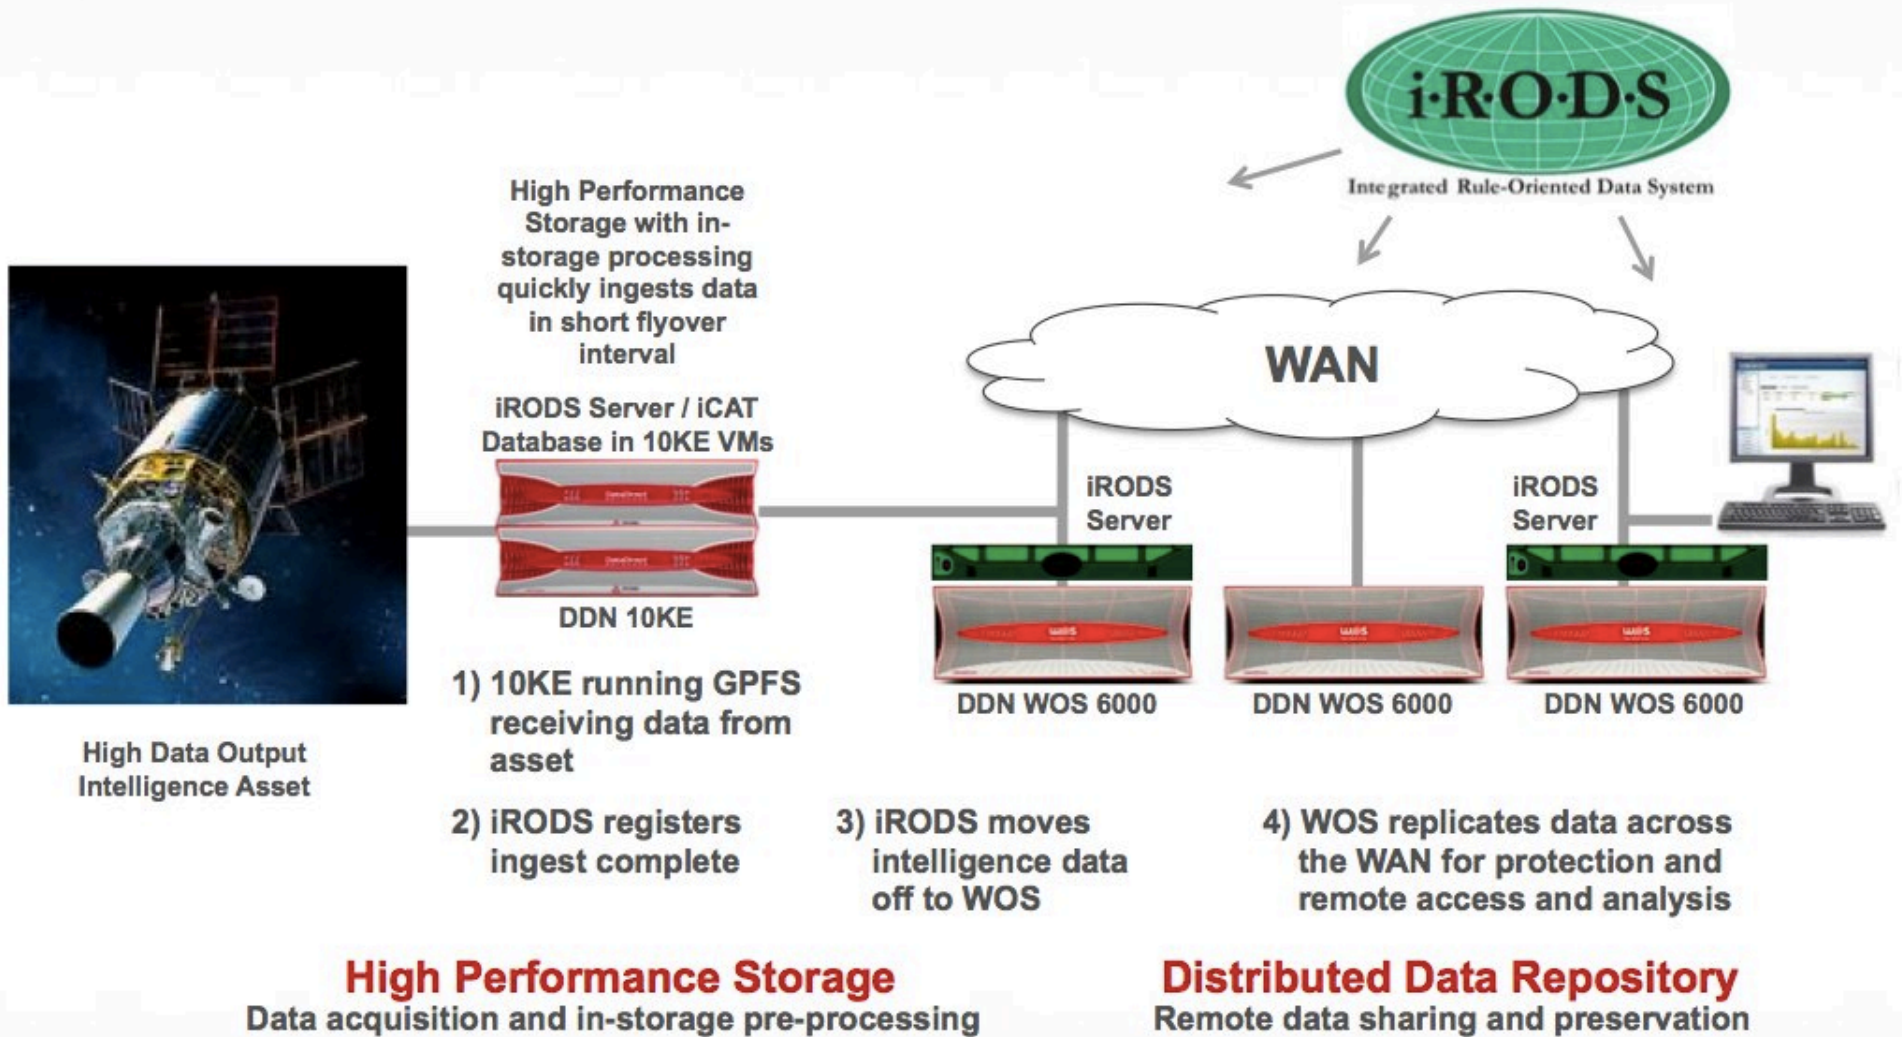
User-Defined Metadata allows customers to understand their data

Global, Peer:Peer

Distribute data across 100s of sites in one namespace

Self-Healing

Intelligent Data Management system recovers from failures rapidly and autonomously



DDN | The “Big” In Big Data



98%
Fortune 500

- ▶ Customers of Online File Service
- ▶ PBs of Cloud Storage



youSENDit™

768
Thousand

- ▶ # of Supercomputer CPUs
- ▶ 35PB Storage Processing System



23
Million

- ▶ Broadcast Viewers
- ▶ 9PB of Real-Time Play-To-Air Systems



If any company is well poised to take on the challenges of exascale computing and big data, it's DDN, since this is its heritage.

451 Group

October 2012

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

Accelerated Data !

Massively-Scalable Platforms and Solutions
Engineered for the Big Data and Cloud Era

Thank You

ddn.com