

Decoupling Bandwidth from Capacity

Burst Buffer and Beyond

DDN[®]
STORAGE

Jean-Thomas Acquaviva, DDN

Step, 2016


2nd SKA CNRS workshop, Paris

Corporate Status: DDN Advanced Technical Center

2

R&D centered on Emerging tech. programs, **Paris, France**

- 25+ R&D engineers



De-correlate bandwidth from capacity

Distributed Virtually Shared Coherent Array of SSDs

3

SSD reshuffles the parameters

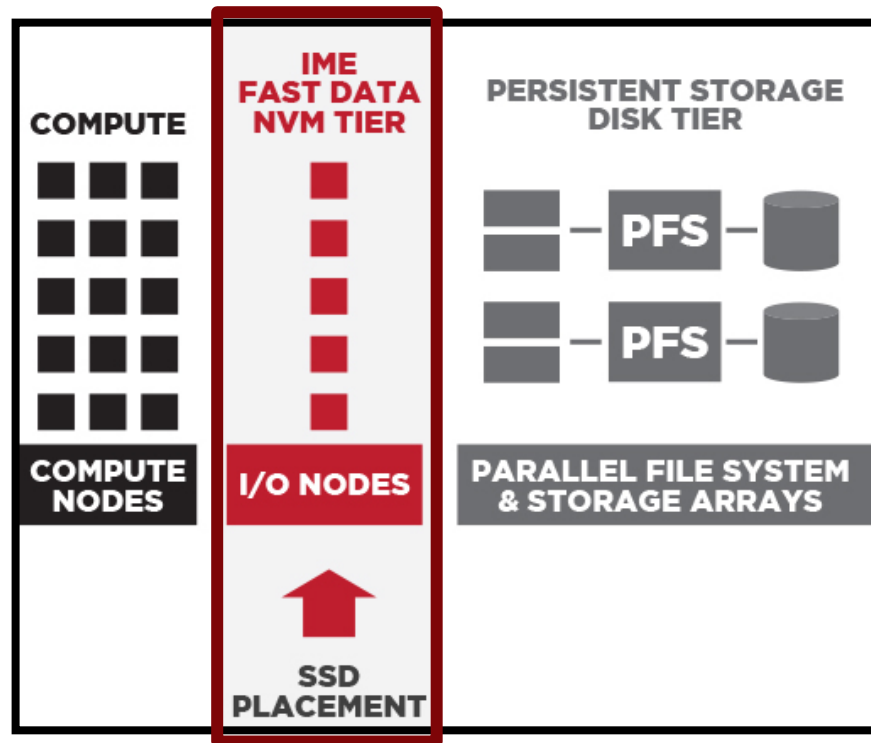
Latency / 40 : 4ms → 0,1 ms

Bandwidth x 3: 150 → 450 MB/s

Capacity / 8 : 8 → 1TB.

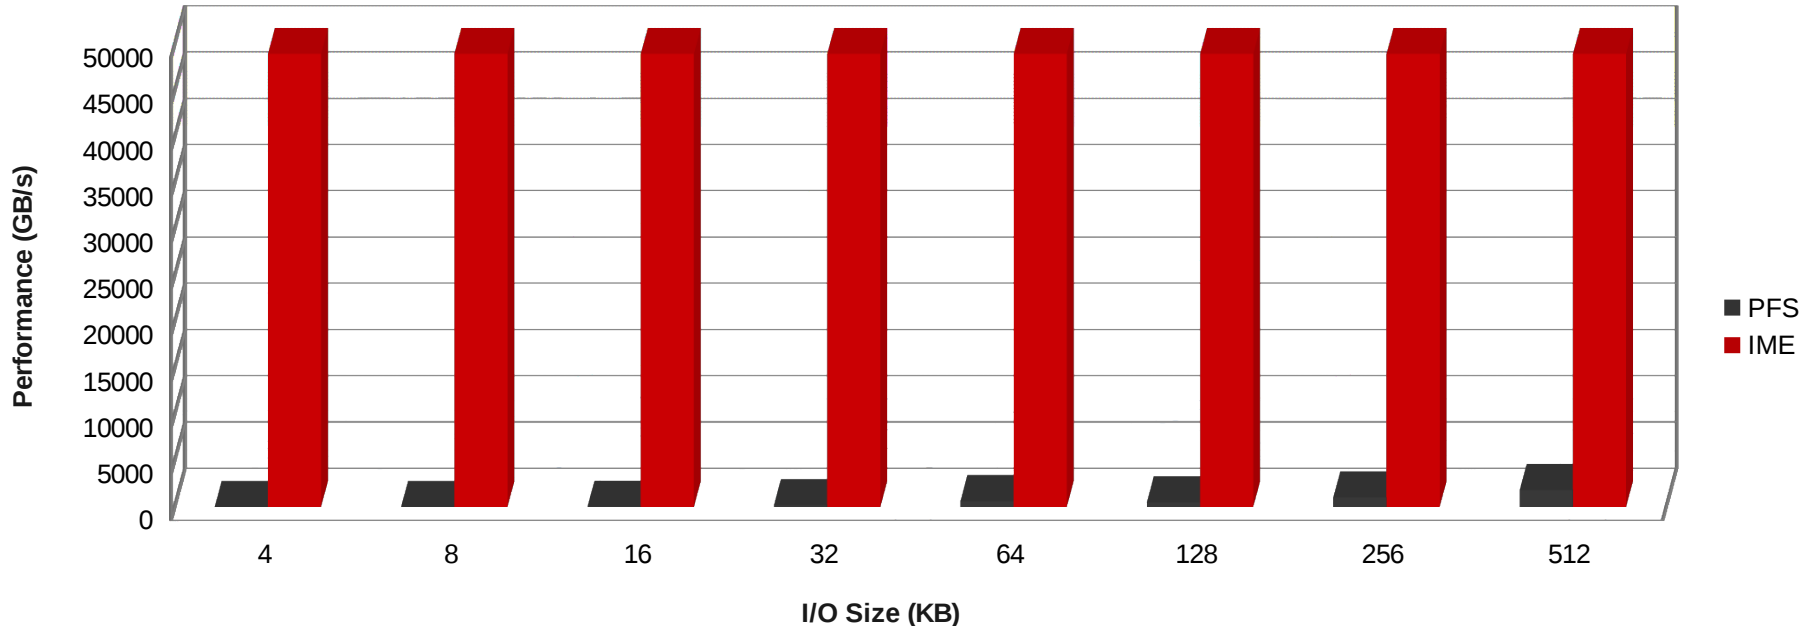
Cost x 10 \$ 0,05/Gbit → \$0.04

What can we do with a costly high bandwidth low latency technology?



Log Structure Removes Artificial Lock Requirements

IOR interleaved access on a single shared file: false sharing impact





John Bent, Garth Gibson, Gary Grider, Ben McClelland, Paul Nowoczynski, James Nunez, Milo Polte, and Meghan Wingate. 2009. *PLFS: a checkpoint filesystem for parallel applications*. In Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis (SC '09).

From Monitoring to Orchestration

5

I/O wait Time for Proxy MPI ranks

48 proxy over 1024 ranks



DDN DIO-pro

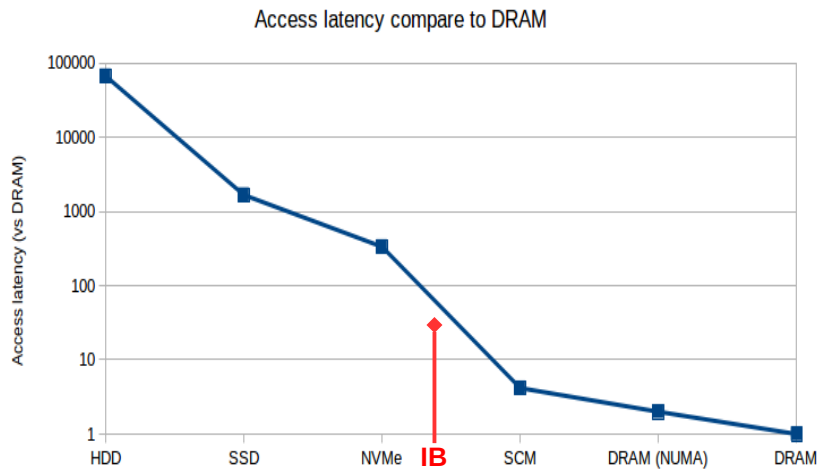
2016

IME maturing from I/O Accelerator to Next gen FS

6

Storage getting closer to the CPU

- ↳ Mechanically same needs will arise



- **Access latency put pressure on the software design**
→ window of opportunity to drastic redesign



Stockage Objet | WOS

Introduction

SKA France | CNRS

Deuxième atelier HPC

Guillaume Mangeot

HPC & Big Data Pre-Sales Engineer, France

Vendredi 9 septembre 2016

Sommaire


- ▶ **Huge Data, Solution actuelle, Limitations**
- ▶ **Solution Objet, comparaison**
- ▶ **Cas concret d'utilisation**

Désolé pour les nombreux anglicismes

Huge Data, solutions actuelles

PFS: Scale-up / Scale-out

Explosion: BigData, DataMining, IOT, VOD, CCTV, Active Archive



Mais :

- Gestion centralisée, un seul FS
- Synchronisation permanente des fichiers
- Concurrence des écrivains

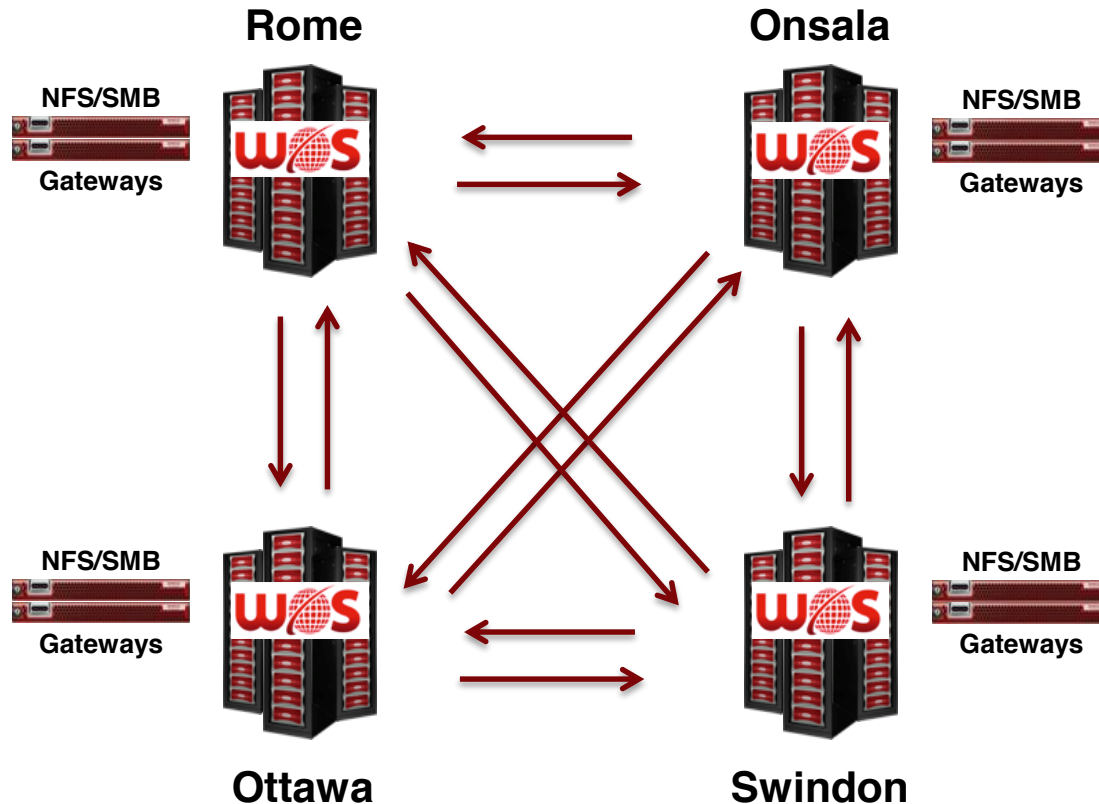
l.u.s.t.r.e.®

GPFS



BeeGFS

Accès Multi-site

Les besoins ont changés

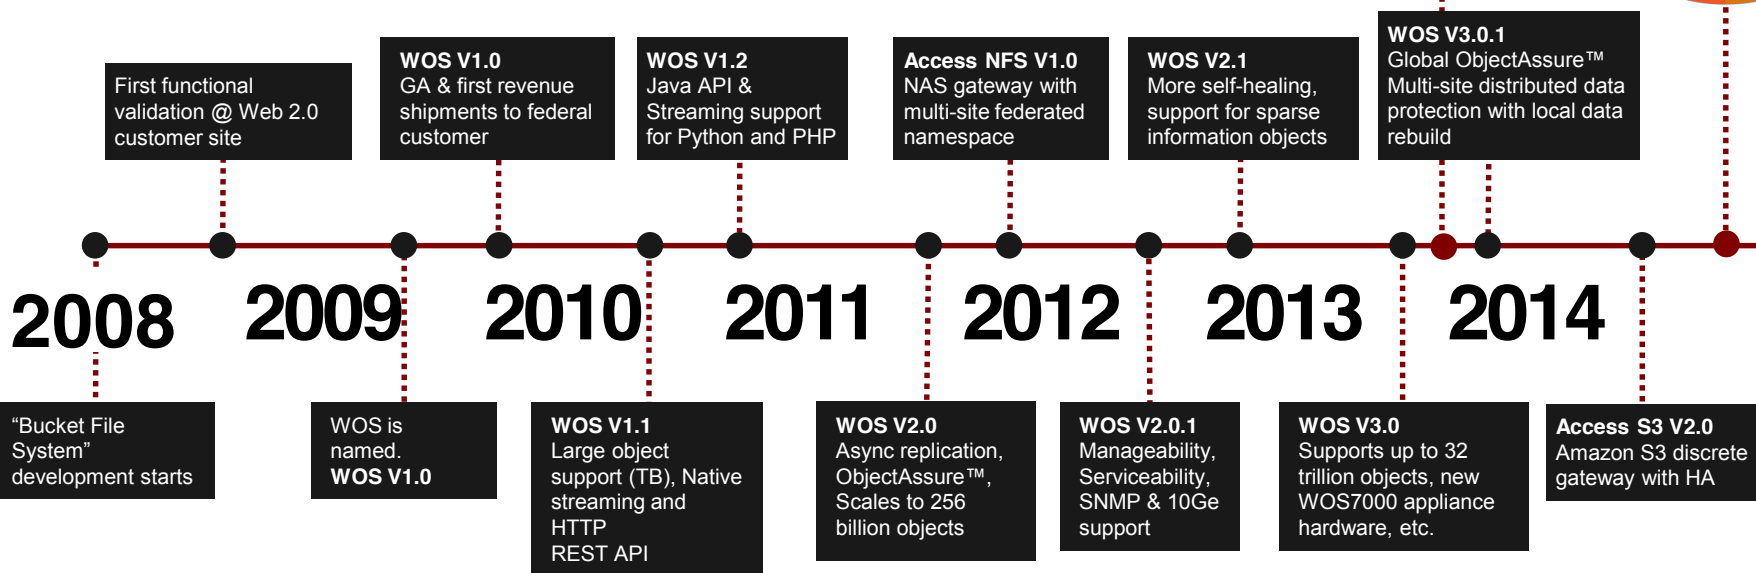


History of Object Storage at DDN



The History of

WOS[®]



Scalabilité de la méthode objet

Réduction de la complexité


- ▶ **Réduction drastique des primitives, POSIX vs.**
 - PUT
 - GET
 - DELETE
- ▶ **Indépendance des objets, figés**
- ▶ **Politiques de stockage, conservation**
- ▶ **Accès multi-site, latency aware**
- ▶ **Gestion décentralisée**
- ▶ **adressage à plat par identifiant**
- ▶ **RAID declustered multi-site, self-healing**
- ▶ **TCO faible (notamment d'administration)**

Object extension, Example Use Cases



Public/Private Cloud Storage

- Storage-as-a-Service enabler
- Unlimited scale, Predictable performance, Global access



Enterprise Collaboration

- File Sync and Share
- Global Workplace productivity, secure and always-on



Active Archives

- Centralized Content Repository
- Storage efficiency, Data durability, No forklift upgrades



Global Content Distribution

- Geo-distributed content
- Low-latency, latency-aware, optimized storage



Cloud Backup

- Always-on, Scalable backup
- Easy recovery

Real cases

WOS extended Architecture

Applications

iRODS


 ctera
 Cloud Storage, Enabled.


 ASG


 commvault
 solving forward®


 owncloud


 ARRIS

API's



HTTP REST



JAVA



C++



Python

WOS Access Gateways



NFS / SMB



S3 / Swift


 WOS
 GRIDScaler Bridge

Tiering



WOS Core

True Object Storage

- ▶ Flat namespace
- ▶ No underlying file system
- ▶ Object-Disk data placement
- ▶ Client-side Global Latency Map

WOS Data Protection

- ▶ WOS Policy engines
- ▶ Replication engine
- ▶ Object Assure™ Erasure Coding

WOS Metadata Management

- ▶ WOS Search

WOS



SMR



Thank You!

Keep in touch with us



sales@ddn.com



2929 Patrick Henry Drive
Santa Clara, CA 95054



[@ddn_limitless](https://twitter.com/ddn_limitless)



1.800.837.2298
1.818.700.4000



[company/datadirect-networks](https://www.linkedin.com/company/datadirect-networks)