



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

GPFS CACHE over Wide Area Network

Lugano, March 15th 2013

Stefano Claudio Gorini, CSCS



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

GPFS 3.5 – New Features

- **Active File Management**
- Snapshot clones
- Snapshot support of filesets
- Independent Filesets
- Improved user and group quotas
- Windows Native Port
- GPFS Native RAID



AFM – Active File Management


- It's a **scalable, high-performance**, file system **caching layer integrated** with the **GPFS** cluster file system.
- It enables you to **create associations between GPFS clusters** and **define the location and flow of file data**, to automate the management of the data.
- This allows you to implement a **single namespace view across sites** around the world.



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Environment



Latency is the main issue in this scenario



Environment

GPFS HOME CLUSTER

main GPFS cluster where the **data** is **safely maintained** and possibly **backed up**.

GPFS REMOTE CACHE

remote site installation with minimal **GPFS clusters** (suggested 3 servers) and a **disk cache**.
Designed to **satisfy the short term needs**.

WIDE AREA NETWORK

link that connects the **two Clusters**:
a **dedicated link** or
a simple **Internet connection**.



AFM-MODE

READ-ONLY

The home data is readable and writeable, but the **remote cached data can't be modified locally.**

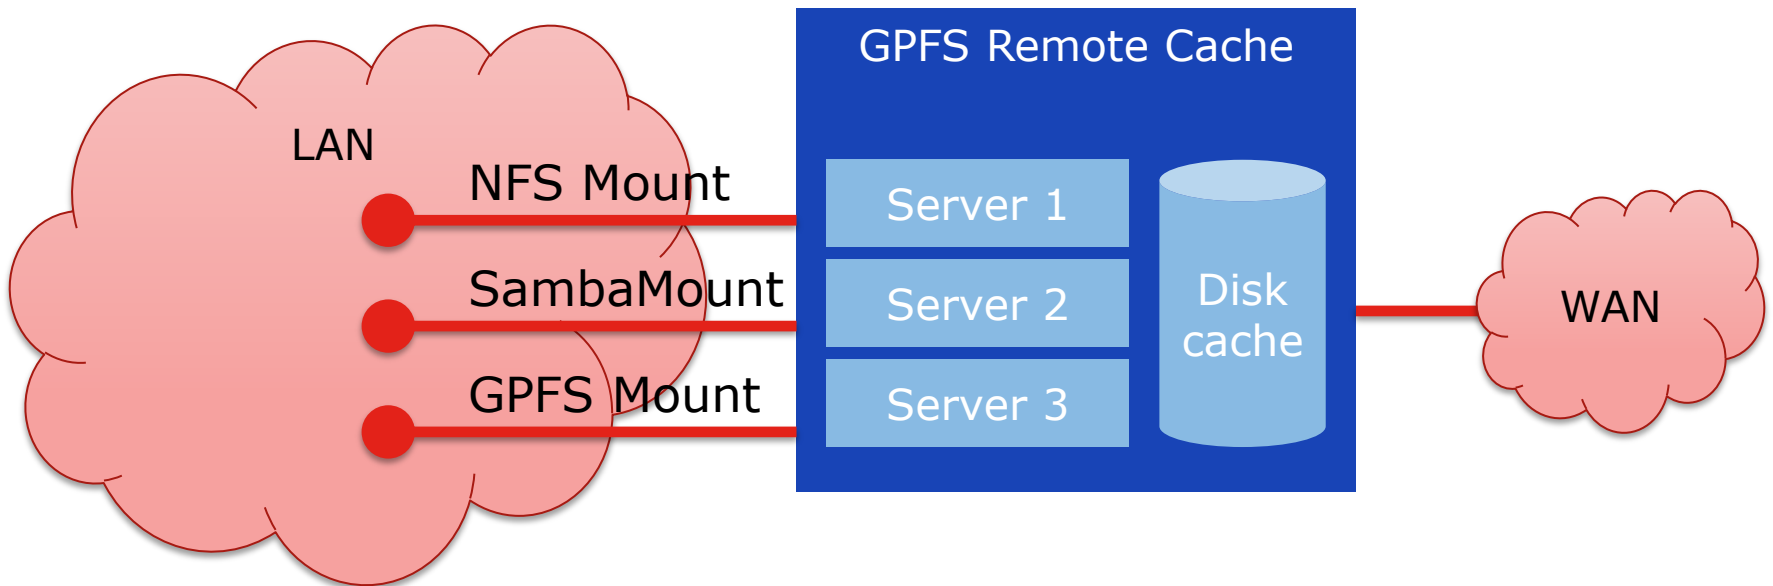
SINGLE-WRITER

Data can only be modified
In the cache.
Home data cannot be modified.

LOCAL-UPDATES

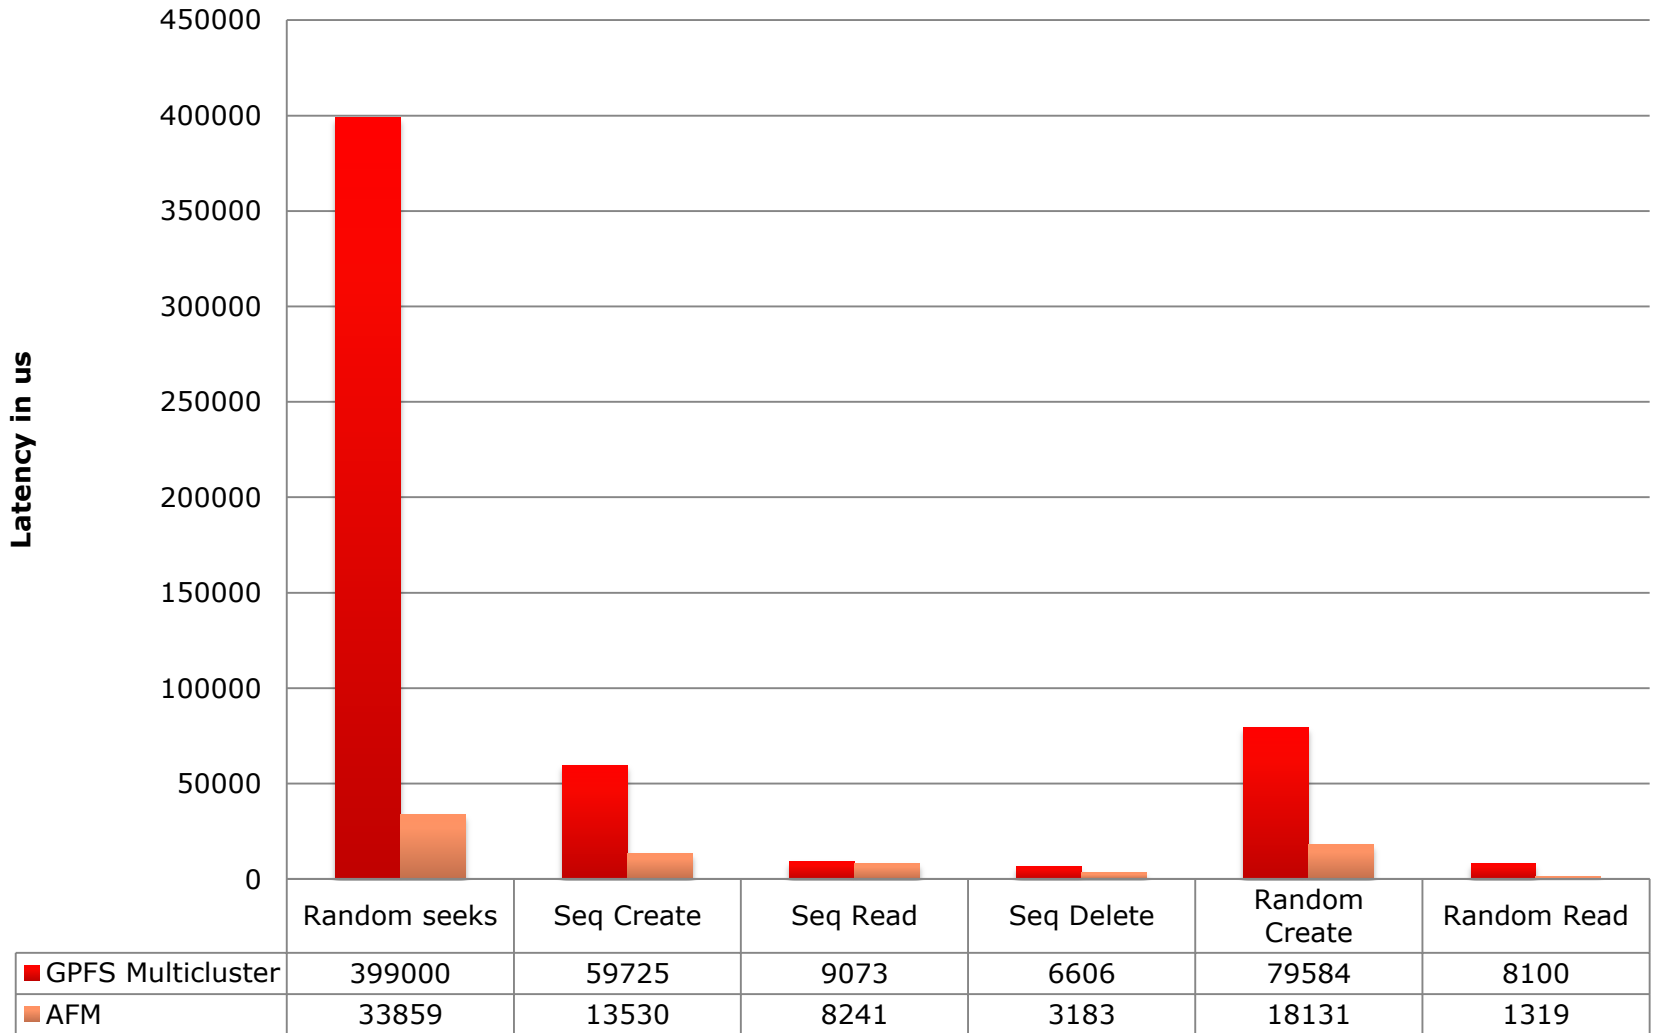
This is like similar to a **snapshot**,
but modified data will
never be synchronized
back, relationship is broken.

Example of GPFS Remote Cache Environment






Experience in Numbers



A different Use Case - **Migrate a generic FS to GPFS**

AFM works on the NFS layer

It doesn't care which FS you have as a source



1. Create the link
2. Fetch all the data
3. Remove the link



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Conclusion

Pros

- Feel at home all around the world
- Reduce the cost
- A new way to migrate data

Cons

- Unique bottle neck

What is missing?

- **Multi-writer mode**



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Thank
you



A special Thanks to:

- Tacchella Davide (CSCS)
- Kalyan Gunda (IBM)