

# Supporting Big Data experiments on Grid'5000

Lucas Nussbaum

lucas.nussbaum@loria.fr

Joint work with Anirvan Basu and David Margery



## Several use cases

- ▶ **Object/Block storage consumer**
  - ◆ Alice would like to deploy a Cloud, and use an object store for VM images, without having to care about the technical details. Best-effort performance is enough.
- ▶ **File-based storage consumer**
  - ◆ Bob would like to test his image-processing application on large datasets, hosted on a POSIX-like filesystem (his application does not support any other data access). Best-effort performance is enough.
- ▶ **Home-made storage framework experimenter**
  - ◆ Carol developed her own storage system and would like to evaluate its performance. Of course, stable performance of the underlying components is critical for reproducibility.
- ▶ **Standard but customized storage framework experimenter**
  - ◆ Dave would like to experiment (measure stuff) with a customized version of Ceph or Hadoop or Spark. Again, stable performance is critical for reproductibility.

# Tentative requirements & current status

- ▶ **Object/Block storage consumer**
  - ◆ Managed (permanent) object/block store
  - ◆ (or dedicated object/block store on reserved nodes)
  
- ▶ **File-based storage consumer**
  - ◆ Managed (permanent) POSIX-like distributed/network filesystem
  
  - ◆ (or dedicated distributed filesystem on reserved nodes)
  
- ▶ **Home-made storage framework experimenter**
  - ◆ Nodes with many disks (SSDs, HDDs)
  - ◆ Managed (permanent) storage to keep data between experiments
  - ◆ Tool for efficient data import (managed storage → reserved nodes)
    - ★ Or disks reserved on nodes, invisible to other users?
  
- ▶ **Standard but customized storage framework experimenter**
  - ◆ (Same as previous, +)
  - ◆ Tools that help with the setup of common Big Data frameworks on reserved nodes

# Tentative requirements & current status

## ▶ Object/Block storage consumer

- ◆ Managed (permanent) object/block store ✓ Two Ceph clusters
- ◆ (or dedicated object/block store on reserved nodes) ✓ Ceph5k (A. Basu) for automated Ceph deployment

## ▶ File-based storage consumer

- ◆ Managed (permanent) POSIX-like distributed/network filesystem ✓ reservation of NFS storage (but no performance guarantees)
- ◆ (or dedicated distributed filesystem on reserved nodes)

## ▶ Home-made storage framework experimenter

- ◆ Nodes with many disks (SSDs, HDDs) ✓ 40 nodes with 4+ disks
- ◆ Managed (permanent) storage to keep data between experiments ✓
- ◆ Tool for efficient data import (managed storage → reserved nodes)
  - ★ Or disks reserved on nodes, invisible to other users?

## ▶ Standard but customized storage framework experimenter

- ◆ (Same as previous, +)
- ◆ Tools that help with the setup of common Big Data frameworks on reserved nodes ✓ Ceph5k

# Tentative requirements & current status

## ▶ Object/Block storage consumer

- ◆ Managed (permanent) object/block store ✓ Two Ceph clusters
- ◆ (or dedicated object/block store on reserved nodes) ✓ Ceph5k (A. Basu) for automated Ceph deployment

## ▶ File-based storage consumer

- ◆ Managed (permanent) POSIX-like distributed/network filesystem ✓ reservation of NFS storage (but no performance guarantees)
- ◆ (or dedicated distributed filesystem on reserved nodes)

## ▶ Home-made storage framework experimenter

- ◆ Nodes with many disks (SSDs, HDDs) ✓ 40 nodes with 4+ disks
- ◆ Managed (permanent) storage to keep data between experiments ✓
- ◆ Tool for efficient data import (managed storage → reserved nodes)
  - ★ Or disks reserved on nodes, invisible to other users?

## ▶ Standard but customized storage framework experimenter

- ◆ (Same as previous, +)
- ◆ Tools that help with the setup of common Big Data frameworks on reserved nodes ✓ Ceph5k