

# KISTI-GSDC Site Report

## HEPiX Spring 2013 Workshop

Gianni M. Ricciardi

KISTI-GSDC

2013 April  
CNAF Bologna (Italy)

# KISTI GSDC Overview

- Global Science Data hub Center: promotion of research collaborations between Korean researchers and national and international institutes around the world
- Part of National Institute of Supercomputing and Networking (former Supercomputing Division of KISTI)
- Providing computing and storage resources for collaborations with HEP experiments (ALICE, KiAF, BELLE, CDF, LIGO, STAR, RENO) and other fields of research (GBRAIN)
- Joining several software development projects
- Running T2 and *T1 candidate* for the ALICE experiment

# Resources

- Totally 4232 cores ( $\sim$ 500 servers) and 3.3PB of disk storage
- Torque/Maui (version shipped with EMI 2) on most of clusters including T1 and T2;
- Condor on clusters for CDF, LIGO and STAR collaborations;
- Mostly SL6 but still several SL5 systems (due to some experiments requirements)
- Configuration management: Puppet
- Monitoring: Nagios, Ganglia, cacti
- In-house developed installation system (evaluating Spacewalk/Satellite)

# T1 Candidate for the ALICE Experiment

- Initial setup in July 2010
- Currently under evaluation
- ALICE collaboration appreciated the progress of KISTI-GSDC as a WLCG Tier-1 site candidate
- A couple of fundamental requirements not met yet: network and 24-hour on-call support
- A road map to achieve the remaining milestones should be provided in May 2013

# T1 Candidate for the ALICE Experiment

- 1512 cores, RAM 3GB/core
- SE: 1PB disk
- Old 100TB SE to be decommissioned soon (data migration in progress)
- Tape installation complete
- Dedicated network link to be upgraded soon

# T1 Candidate - Pledged Resources

	Achieved	Pledged	
	2012	2013	2014
CPU [HS06]	18,800	25,000	31,250
Disk [TB]	1,000	1,000	1,000
Tape [TB]	1,000	1,500	2,000

- Based on WLCG & ALICE Collaboration MoU
- Meeting **ALICE requirements** by 2013
- Providing 2500 cores ( $\sim 31\text{K}$  HS06) and 2PB tape by 2014

# T1 Candidate - Network

- Currently 1Gbps (dedicated) from KISTI to CERN
- 10Gbps required to join LHC OPN (Optical Private Network)
- Dedicated 2Gbps will be deployed by the end of April 2013 (budget secured)
- Plan for 3Gbps upgrade as well as the road map to reach 10Gbps should be presented in May 2013

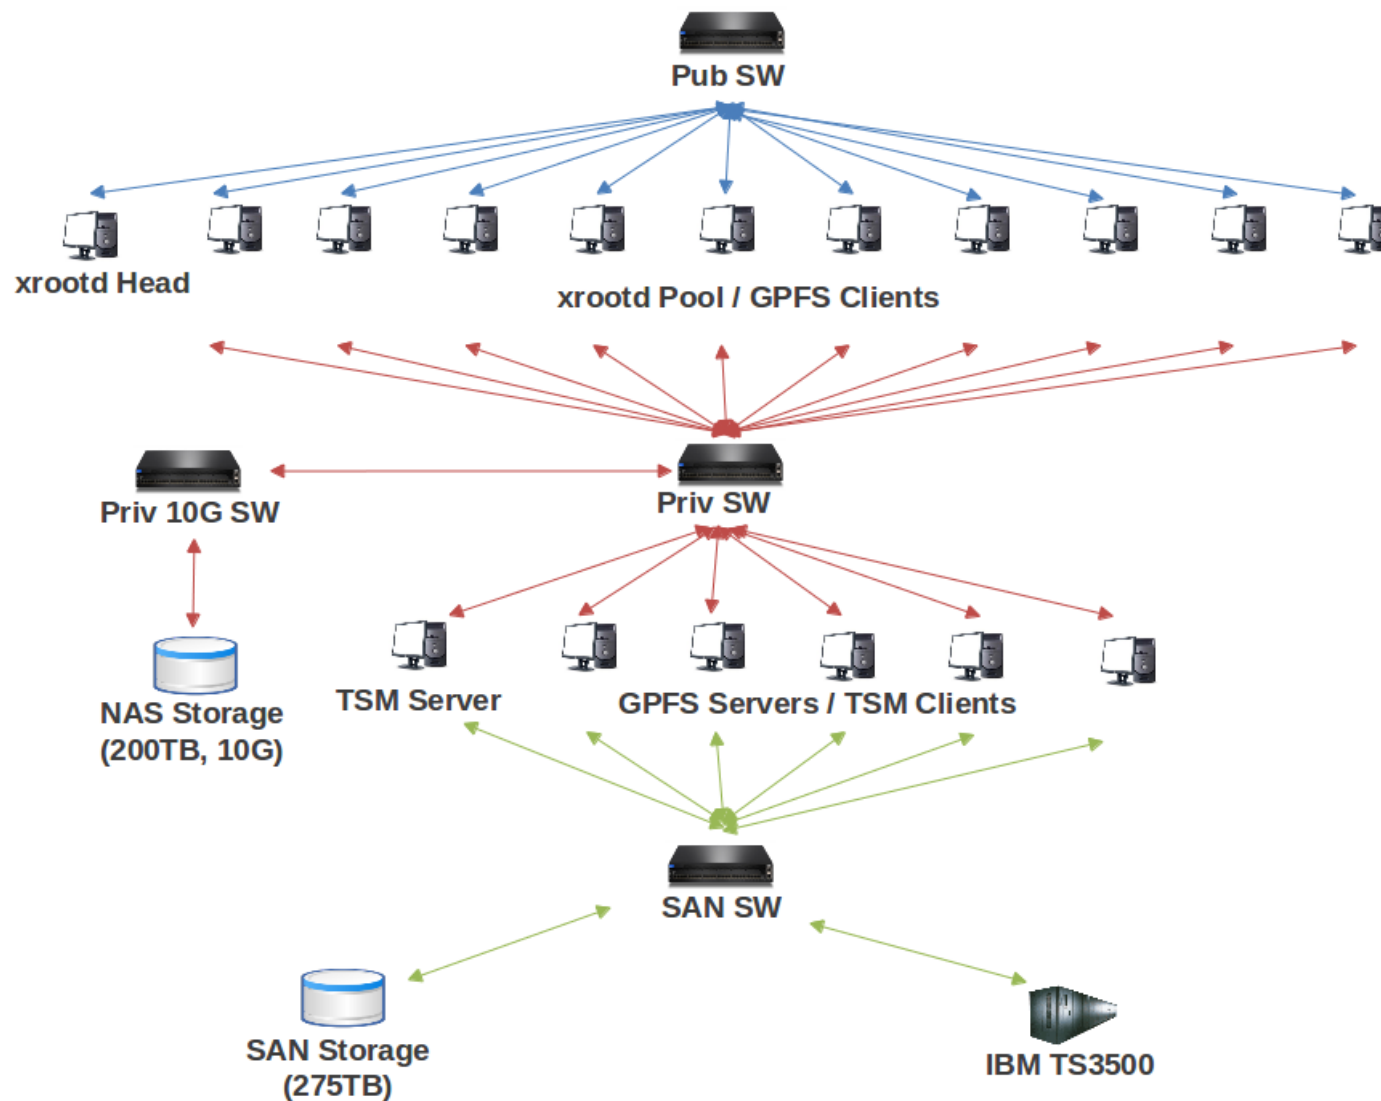


# T1 Candidate - Tape System

- Based on a IBM TS3500
- 8 tape driver (250MB/s each)
- Currently 260 media, 4TB each (~1PB totally)
- Additional media (up to 3PB) to be added this year
- Managed through TSM/GPFS
- Custom FRM hook scripts to interface xrootd to GPFS (migrate/purge)
- 2 specific policies (FRM, GPFS) to migrate files from xrootd to GPFS and from GPFS to tape
- Current testbed:
  - 200TB NAS volume (xrootd pool disk)
  - 275TB SAN volume (GPFS pool disk for cache)
  - Heavy network traffic during migrations to GPFS



# T1 Candidate - Tape System



# T1 Candidate Tape System

- ALICE p-Pb RAW data replication started last March and going on
- Quite useful to test and tune migration policies
- Periodic functional test by ALICE on xrootd for tape passed (tape system fully functional)
- Currently, some issues related to transfer bandwidth from CERN:
  - fairly good use of 1G link but not stable yet
  - some problems with TSM/GPFS configuration: IBM experts are investigating together with the storage team

# T1 Candidate - Tape System

- Future improvements:
  - xrootd pool disk to be deployed on SAN (220T volume)
  - GPFS and xrootd servers installed on the same nodes to mitigate network issues (kind of *single cache* architecture)
  - Any suggestions greatly appreciated 😊

# Questions and (possibly) Answers



Thanks for listening.