



INFN-T1 site report

Andrea Chierici

On behalf of

INFN-T1 staff

Outline

- News on management
- Data management & Storage
- Farming
- Projects and activities



News on management

Management

- Gaetano Maron terminated mandate as CNAF director
- New CNAF director: Luca dell'Agnello
- New INFN-T1 director: Daniele Cesini

Tecnopolo

- The main focus of our management is on Tecnopolo migration
 - TDR delivered to project referees
- 2021-2022 Infrastructure preparation
- Q1 2023 actual migration
 - Exploit of Leonardo (pre-exascale) for part of CPU pledges
- Current DC in production until end of 2023



Data Management and Storage

Storage resources

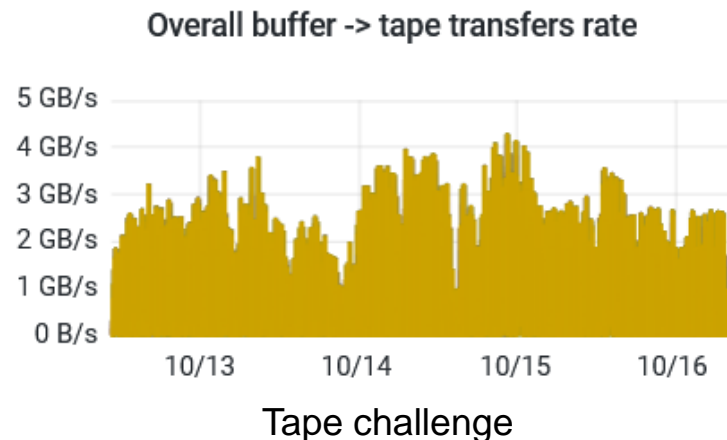
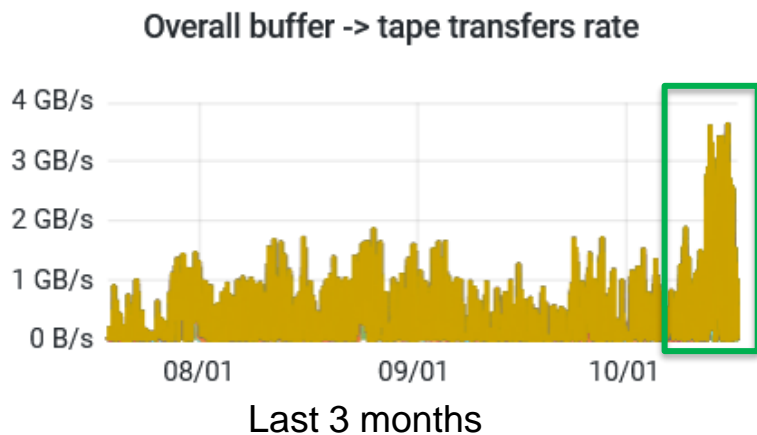

- Disk
 - 41 PB installed – 36 PB used
- Last tender
 - 8.7 PB usable with GPFS back-end (10 TB disks)
 - Installation ongoing
 - 5.2 PB raw usable as CephFS back-end (18 TB disks)
 - Pre-production testing with Alice
- Tape
 - 96 PB installed – 89 PB used
 - 2 libraries
 - Oracle SL8500, almost full
 - IBM TS4500 in production since Feb 2020, 15 PB installed

Storage services

- Migration from gsiftp to https/XRootD, in view of Globus retirement
 - WLCG DOMA TPC working group initiative
 - Token-based AuthN/Z fully supported on our StoRM WebDAV endpoints
 - HTTP in production for:
 - Disk ATLAS (GSI-FTP no more used)
 - Disk CMS and LHCb (GSI-FTP still used)
- SRM+https tape tests
 - ATLAS: Over 100TB recalled, ready for production when FTS servers will be upgraded
 - CMS: first functional test succeeded

Tape challenge

- New GEMSS in production
 - manages recall queues from multiple tape libraries





Farming

Computing resources


- Farm power: approx. 440 KHS06
 - 250 kHS06 at CNAF
 - 180 kHS06 at CINECA
 - 10 kHS06 at Bari-ReCaS
 - Just installed latest tender
 - 42 KHS06 (Epyc 7282)
 - ~~Within summer we plan to add 240 KHS06~~
 - New batch of nodes from CINECA (same configuration as current one)
 - We will reach 2021 pledge and will be able to dismiss very old hardware
- New estimate: before end of the year**

Software updates


- Rebooted the whole farm twice during summer
 - Dynamic upgrade procedure works smoothly
- HTCondor upgrade started last week
 - 8.8 → 9.0.x
 - CE 3.4 → 5.1
 - TokenId auth*n, Grid Scitokens aware

Batch System

- 6xHTC-CE, 1xSN, 1xCM, ~870WN, 39Kcores



hs06 Avg: 426 K pledge Avg: 588 K hs06_mc Avg: 162 K hs06_sc Avg: 265 K
 hs06_pot active Avg: 432 K hs06_pot down Avg: 2.03 K hs06_pot total Avg: 434 K



Grid Running Slots Avg: 36.0 K Local Running Slots Avg: 1.91 K Grid Pending Jobs Avg: 3.77 K
 Local Pending Jobs Avg: 9.63 K Grid Running Jobs Avg: -24 K Local Running Jobs Avg: -1.4 K



Other projects and activities

Common tools

- OpenDCIM is the new asset management tool for INFN-T1
- Plan to insert only recent hardware
- Tool for Tecnopolo

INFN-T1 Inventory
Data Center Device Detail

Asset Tracking

Device ID: 1063
Status: Production
Label: ENC-205-10-32
Serial Number: RK016820
Asset Tag:
Primary IP / Host Name:
Manufacture Date: 1970-01-01
Install Date: 2021-08-16
Warranty Company: E4
Warranty Expiration: 1970-01-01
Last Audit Completed: Audit not yet completed
Departmental Owner: Farming

Escalation Information

Time Period:
Details:

Primary Contact: Chierici, Andrea
Tags:


Custom Attributes

Classification:
Department:

Physical Infrastructure

Cabinet: CNAF-DC / 205-10
Device Class: Supermicro - CSE-217BHQ+-000NBP2
Height: 2
Position: 32
Half Depth:
Back Side:
Number of Data Ports: 0
Nominal Draw (Watts): 0
Weight: 0
Power Connections: 2
Chassis Total Draw: 0
Chassis Total Weight: 0
Device Type: Chassis

Device Images



SNMP Configuration

SNMP Version: 2c
SNMP Read Only Community:
Consecutive SNMP Failures*: 0
*Polling is disabled after three consecutive failures.

Chassis Contents

Number of Slots in Chassis: Front 0 Rear 4

Slot #	Height	Device Name	Device Type
1	1	WN-205-10-32-01-A	Server
2	1	WN-205-10-32-02-A	Server
3	1	WN-205-10-32-03-A	Server
4	1	WN-205-10-32-04-A	Server

Notes:



RHEL agreement

- Signed agreement that allows CNAF to use unlimited RH licenses.
- Products covered by “campus 2”
 - Red Hat Enterprise Linux
 - Red Hat Enterprise Linux Workstation
 - Red Hat Enterprise Linux Desktop
 - Red Hat x86_64 Linux Server
 - Red Hat Linux HPC
 - Red Hat Linux for SAP Applications
 - Red hat Software Collections
 - Red Hat Developer Toolset
 - Red Hat Linux Atomic Host
 - Red Hat Enterprise Virtualization (RHEV)
 - Red Hat OpenStack,
 - Red Hat Gluster Storage
 - Red Hat Insight

Cloud@CNAF

- Hardware:
 - 10 storage nodes for Ceph cluster → total of ~ 4PB raw disk space
 - Compute nodes with GPU: NVIDIA T4 (8), Tesla V100 (8), A100 (2)
 - New compute nodes & switches
- Software
 - OVirt - Upgrade of infrastructure hosting some of the core services
 - GPFS still used as backend for some of Openstack services → planned to be replaced by Ceph
 - Openstack – upgrade ongoing (on CentOS 8 Stream)
 - Ironic - under testing
 - CERN-CNAF collaboration, gaining experience from Openstack Baremetal SIG (IRONIC operators)
- Cloud@CNAF federated with **INFN Cloud & EGI FedCloud**, hosting among many projects:
 - VIRGO K8S testing infrastructure
 - ML-INFN leveraging the above-mentioned GPUs
 - PLANET (**P**ollution **L**ake **A**nalysis for **E**ffective **T**herapy)
 - DODAS
- CNAF is hosting one of the 2 [INFN Cloud](#) backbone sites
 - New hardware (compute & storage), new [services](#) under testing, users support



EPIC Cloud

- Enhanced **P**rivacy and **C**ompliance
- A cloud instance **ISO 27001** certified (Information Security Management System), **ISO 27017** (Security in cloud services) and **ISO 27018** (Managing Personally Identifiable Information in public clouds)
- EPIC Cloud offers an **IaaS Cloud** for high security demanding research communities
- Migrated from a bare-metal to a cloud platform
- In March 2021 we got a new certification, with an extended scope:
 - *“Hosting of physical and virtual systems for biomedical data access and conservation, and management of data analytics applications targeted to biomedical and genomic research. **Delivering cloud services in IaaS mode by applying ISO/IEC 27017:2015 and ISO/IEC 27018:2019 guidelines**”*
- We had to learn new work methodologies
 - Very useful in Tecnapolo transition



EPIC Cloud

- EPIC Cloud is **based on the same technologies of INFN Cloud** (OpenStack, IAM, Onedata), with various enhancements to meet higher security and privacy standards
 - IAM is based on FreeIPA -> provides **2FA**, integration with web services, SSH and VPN
 - Onedata has advanced **auditing functionalities**
 - **Network segregation between OpenStack tenants**
 - Shared responsibility model for IaaS:
 - Users manage data, applications, runtime, middleware and OS
 - CNAF staff manage networking, storage, servers, virtualization
 - Advanced logging and auditing services (**centralized syslog managed applying the *segregation of duties* principle**)
 - Advanced backup services providing client-side encryption
- New services based on MinIO storage and HashiCorp Vault secret manager under development

K8s & Distributed FS

- Mount a Distributed FS and leverage its native APIs in a container can improve flexibility in installing, monitoring and recovering data management and transfer services
- Successful deployment and performance tests exploiting Kubernetes/OpenShift & Docker/Podman container engines with the following FS:
 - IBM Spectrum Scale (GPFS)
 - CephFS
 - Lustre-ZFS
- Successful integration and performance tests
 - CERN EOS services on CephFS/RBD storage backend via K8s/Docker



Questions?