

# IT infrastructure - WP2

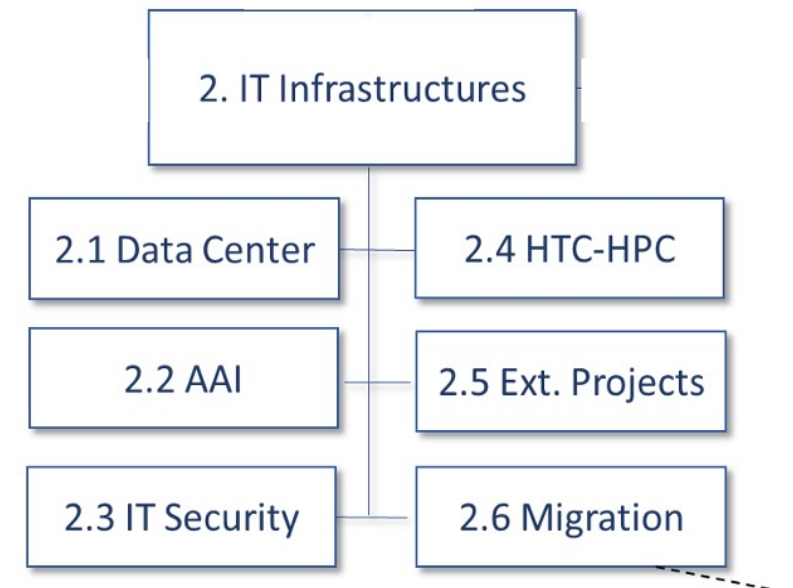
Luca dell'Agnello, Davide Salomoni

---

- Obiettivi e descrizione del WP
- I principali servizi informatici oggetto della ristrutturazione
- Descrizione dei task e sub-task che compongono il WP
- Risorse umane necessarie per ogni task
- Cronoprogramma
- La strategia per la migrazione
- Descrizione dell'architettura del sistema di calcolo e di storage, la topologia della rete
- Modello, interazione e interfaccia con il sistema HPC Leonardo

- Services for scientific computing offered to more than 40 INFN collaborations
  - A “general purpose” HTC farm
  - Data management services with different QoS (e.g., disk and tape)
  - A small HPC cluster with accelerators (for special applications)
  - A cloud infrastructure (INFN cloud and cloud@CNAF) for specific uses (e.g., interactive computing, instantiation of dedicated clusters and ancillary applications)
- Also, a dedicated cloud instance (EPIC, ISO 27001 certified) is available for projects dealing with biomedical and genomic data
  - 2 active projects (Harmony, ACC)
- Furthermore, CNAF is the central ICT service provider for INFN (National Services and Information System)

- Main goal: design and guide the evolution of the services offered by the new CNAF data center taking into account the emerging requirements from WLCG and other communities
- Organized in 5 sub-WPs on the evolution of the infrastructure and services
  - Main one (WP 2.1) focusing on Data Center structure and services
  - WP 2.2 – AAI infrastructure
  - WP 2.3 – IT Security
  - WP 2.4 – Integration of HPC resources into our HTC farm
  - WP 2.5 – External projects
- + a “special” sub-WP dealing with the migration procedure
  - Focusing on relocation of hw and services
  - Also connected to the equivalent sub-WP 3.4

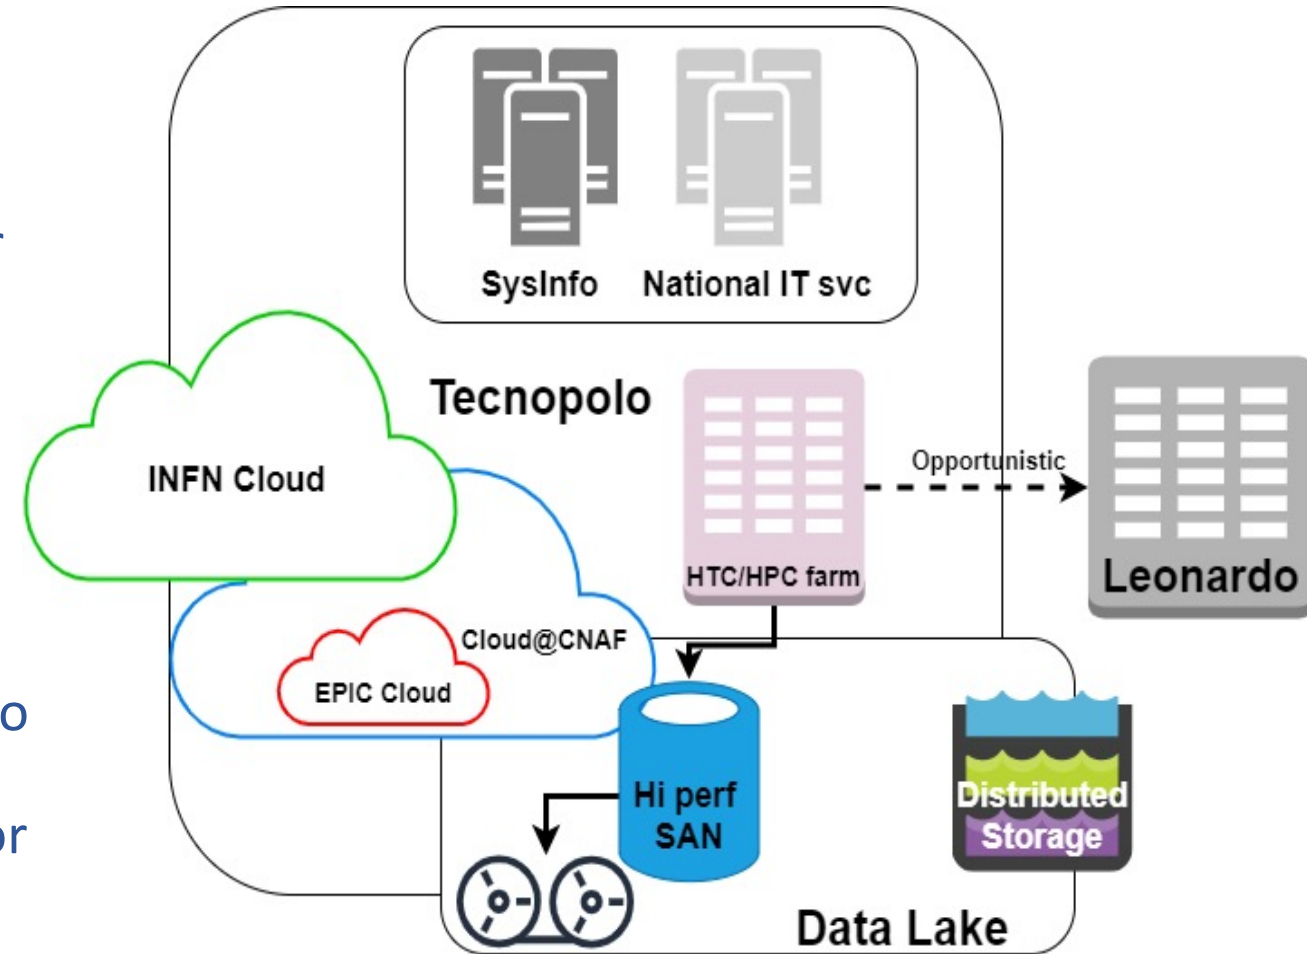


- This is the most complex sub-WP with various tasks and subtasks dealing with the basic services of the data center
  - 2.1.1, 2.1.4, 2.1.5 (Computing) – investigation of the evolution of farm and cloud instances towards a common infrastructure in order to minimize the management effort and to gain flexibility
  - 2.1.2 (DM) – study of the evolution of storage services both at fabric level (i.e., storage models based on open-source alternatives vs. the current one based on GPFS) and infrastructural one (data lake)
  - 2.1.3 (Network) – Evolution of the network both at local level to cope with the increase in dimension of our data center and at geographical level exploiting DCI technology to build data lake
  - 2.1.6 (Operations) – Evolution of auxiliary services (i.e., monitoring, accounting, provisioning, documentation, asset management)
  - 2.1.7, 2.1.8 – Relocation only
  - 2.1.9 (User Support) – Redefinition of the mandate and scope



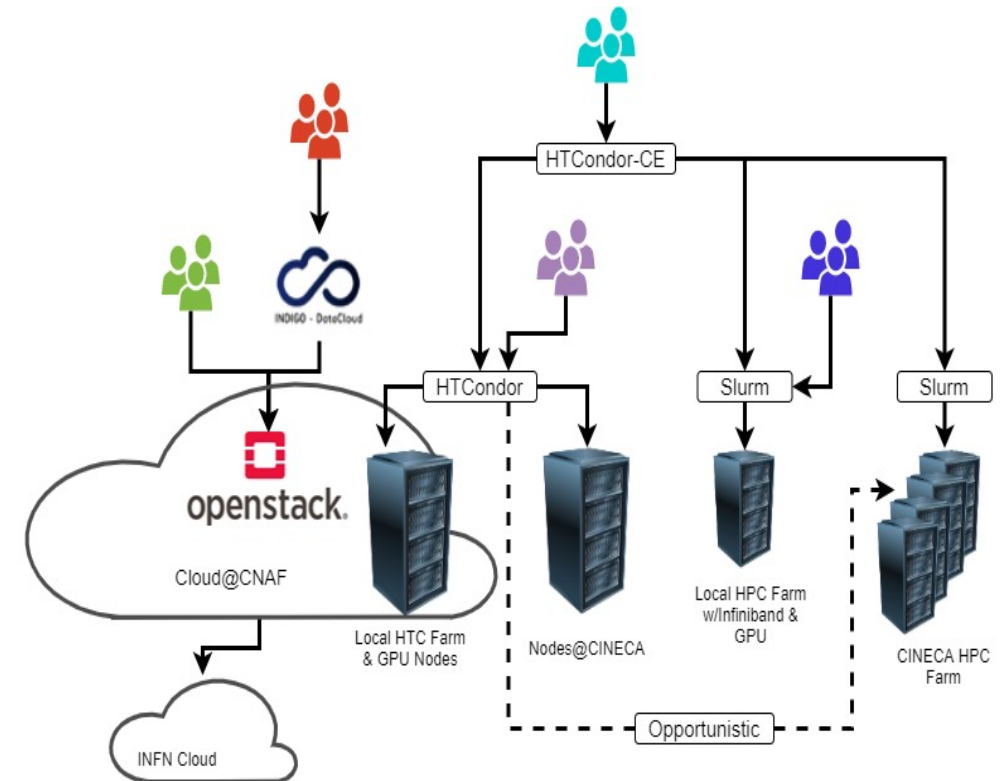
# WP 2.1: the Data Center logical layout

- By 2025:
  - HTC farm: ~1.5 MHS06
  - Disk: ~120 PB
  - Tape: ~200 PB
- Provide two common interfaces to users for computing
  - HTC and HPC farms (including Leonardo)
  - cloud@CNAF (federated with INFN-Cloud)
- Storage resources (disk and tape libraries) serving data to the computing facilities
  - Integrated within INFN and WLCG data lakes
- Strategic role of network connections both internally and with the other INFN centers to exploit cloud and data lake paradigms
- Common fabric layer (i.e., OpenStack with or w/o K8s) for farm and cloud instances
  - Flexibility
  - Less management burden



# Computing (WP2.1.1, WP2.1.4, WP2.1.5)

- HtCondor CE as entry point to both HTC and HPC farms (legacy local access supported)
  - HTC and HPC farms managed by specific LRMS
    - HtCondor (HTC) and Slurm (HPC)
    - Connector HtCondor-Slurm needed
  - Provisioning of GPUs via HtCondor CE
  - Leonardo directly accessed from HtCondor (dedicated resources) or via Slurm (opportunistic, see WP2.4)
    - Transparent to users for dedicated resources
- Cloud@CNAF (federated with INFN cloud)
  - IaaS (based on OpenStack)
    - Virtual clusters for ancillary services
    - User Interfaces for interactive access
    - Specific services for user communities and projects
    - Testbeds for projects
  - PaaS (orchestrated by INDIGO PaaS Orchestrator)
    - Dynamic expansion of the HTC farm
    - Deployment of dedicated clusters
  - Storage back-end based on CEPH/MinIO (see WP 2.1.2)



# WP2.1.1, WP2.1.4, WP2.1.5: farm&cloud

- In order to build a common fabric layer, the HTC farm should become fully virtualized
  - In an extreme vision, the batch system could be a service running on the cloud
  - Various solutions to be evaluated (taking into account performances)
    - WN as docker container on bare metal
    - WN as Kubernetes POD
    - IroniC for bare metal provisioning
  - Kubernetes with or w/o OpenStack
- At higher level evolution of HtCondor (e.g., Token ID based AAI)
- HPC farm will be migrated to Slurm and interconnected at high speed to the storage
- Cloud instances are being upgraded to the latest OpenStack release
  - Definition of the procedure
  - Standardize the release on all instances (EPIC including)

Milestone	Description	Start Date	Delivery Date	FTE Required
M.HTC.1	HTCondor token ID authc/authz	17/5/21	25/6/21	0.8
M.HTC.2	HTCondor-WN as docker container on bare metal	28/6/21	30/11/21	0.5
M.HTC.3	HTCondor-WN as K8s POD	1/12/21	8/2/22	0.8
M.HTC.4	GPU provisioning via HTCondor-CE	10/1/22	4/2/22	0.7
M.HTC.5	HTCondor-CE as gateway to HTC and HPC	4/4/22	27/5/22	0.5
M.HTC.6	Out-of-band node management prototype	1/4/21	30/6/22	0.2
M.HPC.3	SLURM manages the HPC clusters	May 2021	September 2021	1
M.HPC.4	HTCondor-SLURM Connector available and tested	September 2021	January 2022	2
M.HPC.5	Improve connection to the Tier1 Storage system and access GPFS data from the HPC cluster	January 2022	March 2022	0.5
M.Cloud.3	Installation and upgrade of various PaaS orchestrators instances	May 2021	Jun 2021	1
M.Cloud.4	Upgrade of the cloud IaaS level	May 2021	Sep 2021	5



# Farm evolution timeline

Activity	Manpower	2021				2022			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
HTCondor token ID Authc/z	0.8 FTE								
HTCondor-WN as docker container on bare metal: Single whole node + Scalability tests	0.5 FTE								
HTCondor-WN as docker container on bare metal: multiple containers, one per CPU + Scalability tests	0.5 FTE								
HTCondor-WN as docker container on bare metal: one container per job + Scalability tests	0.5 FTE								
HTCondor-WN as K8s POD	0.8 FTE								
GPU provisioning via HTCondor-CE	0.7 FTE								
HTCondor-CE as gateway to HTC and HPC	0.5 FTE						?		
Development of out-of-band node management tool	0.2 FTE								

*Farm HTC: 38 PM*

*Farm HPC: 16 PM*

Activity	Manpower	2021				2022				2023			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Removal of obsolete CPU and storage resources	0.5 (Done)												
Reorganized HPC cluster with a single storage system and new hw included made available	0.5												
SLURM manages the HPC clusters	1												
HTCondor-SLURM Connector available and tested	2												
Improve connection to the Tier1 Storage system and access GPFS data from the HPC cluster	0.5												
Move to Tecnapolo	1												
In production at Tecnapolo -support provided	1 (support)												

# Cloud evolution timeline

Task	Description	2021								2022				2023		
		May-21	Jun-21	Jul-21	Aug-21	Sep-21	Oct-21	Nov-21	Dec-21	Sep-22	Oct-22	Nov-22	Dec-22	Jan-23	Feb-23	Mar-23
M.Cloud.1	Federation of Cloud@CNAF with other cloud infrastructures	█														
	1.1 Federation with EGI Fed Cloud	█														
	1.2 Test of integration with Amazon Web Services		█	█	█	█	█	█	█							
M.Cloud.2	New hardware added to Cloud@CNAF and EPIC Cloud	█														
	2.1 Storage from 2020 tenders	█	█													
	2.2 Net, Compute, Storage equipment (SUPER funding)			█	█	█	█	█								
M.Cloud.3	PaaS orchestrator installation & upgrades	█	█													
M.Cloud.4	IaaS level infrastructure upgrade	█														
	4.1 oVirt upgrade	█	█													
	4.2 Openstack upgrade		█	█	█	█	█									
M.Cloud.5	Cloud Monitoring & Accounting (2.1.4.6, 2.1.4.7, 2.1.4.1.3)															
M.Cloud.6	Validation of new IaaS level services	█														
	6.1 bare-metal provisioning (Ironic)	█	█	█	█	█	█	█								
	6.2 container orchestration engine (Magnum)					█	█	█	█							
M.Cloud.7	Move to Tecnopolo									█	█	█	█	█	█	
M.Cloud.8	Maintenance & Support													█	█	

Cloud: 48 PM

EPIC: 39 PM

# WP 2.1.2 (Data Management)

- Storage resources equally accessible from farm and cloud
- Experimental deployment of CEPH-based storage systems as an alternative to GPFS
  - Larger number of servers, different network infrastructure
  - Evaluation of performances and TCO
- Object Storage (S3 protocol with CEPH and MinIO) provided for cloud
- Evolution of data transfer services
  - Replacing GriFTP with http/WebDAV
  - Already underway for WLCG
- Test and integration of Data Lake solutions
  - Build of INFN DL
  - Join of WLCG one (DOMA, ESCAPE)
  - *Test with Rucio underway for Juno*
- Evolution of tape storage to increase efficiency
  - Dynamic allocation of tape drives aware of multiple libraries
  - Migration to tape based on storage events detection

Milestone	Description	Start Date	Delivery Date	FTE Required
M.DM.1	CEPH cluster in pre-production environment	1/7/2021	31/8/2021	0.6
M.DM.2	Hardware requirements definition for storage tenders	1/10/2021	18/2/2022	1
M.DM.3	Phasing-out GridFTP	1/5/2021	31/12/2022	1
M.DM.4	Evaluation and implementation of Data Lake solutions	1/7/2021	31/12/2023	0.5

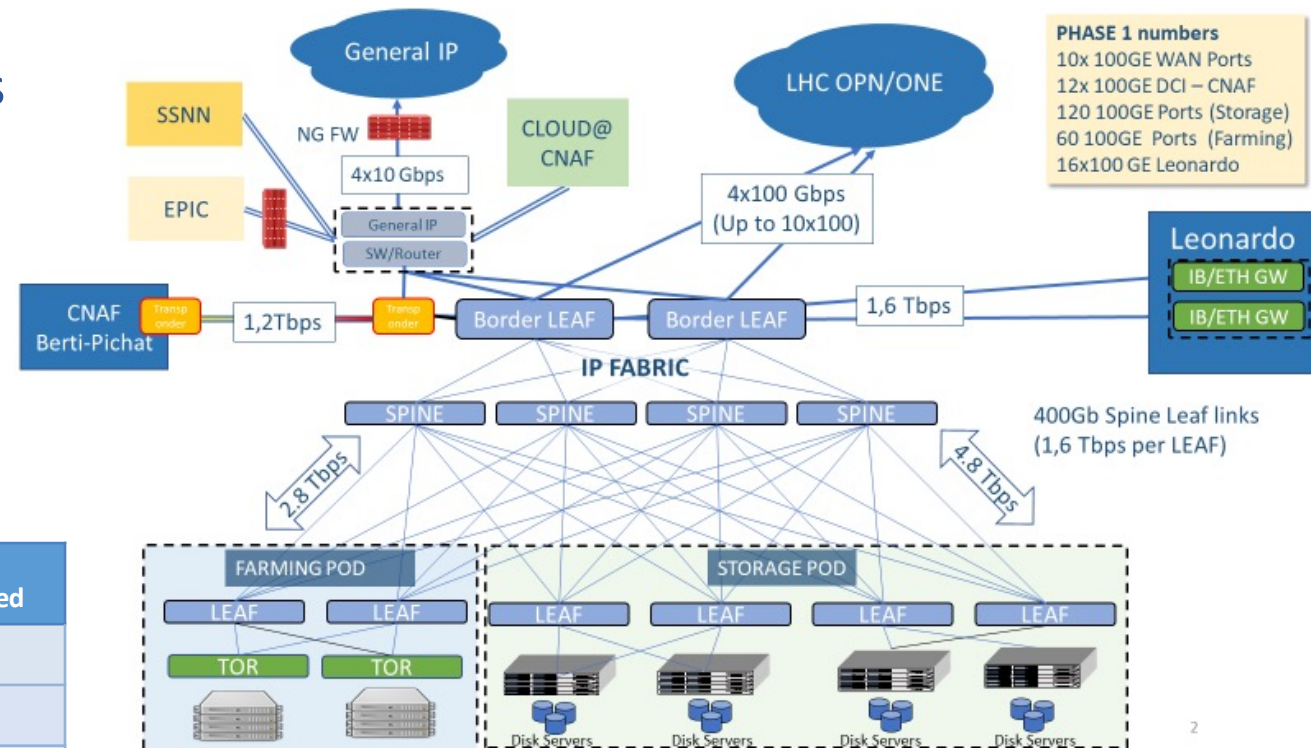
# Data Management: the timeline

		2021				2022				2023			
Task	Activity	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
2.1.2.1	Planning evolution and migration												
2.1.2.2	Initial installation of storage resources												
2.1.2.3	Deployment and evaluation of CEPH												
2.1.2.4	Evaluation and deployment of Data Lake solutions												
2.1.2.5	Software solutions for tape												
2.1.2.6	Evolution of backup service												
2.1.2.7	Evolution of data management services												

DM: 30 PM

# WP 2.1.3 (Network)

- Evaluation of different architectural solutions
  - E.g., spine-leaf vs. central core approach
  - Functionalities and TCO
- Test of DCI with CERN and other INFN centers
  - Interconnection at Tbps scale for Data Lake
- Setup and configuration of DCI with CNAF (crucial for migration)
- Setup and tests of gw IB  $\leftrightarrow$  Eth (Mellanox Skyways) to interconnect to Leonardo



Milestone	Description	Start Date	Delivery Date	FTE Required
M.Net.1	Collection of requirements		30/4/2021	1.5
M.Net.2	Technology tracking and POCs	1/5/2021	31/10/2021	1.5
M.Net.3	Estimation of cost of solutions	1/5/2021	31/10/2021	1.5
M.Net.4	Definition of the technical specifications for cabling and network devices	1/11/2021	31/12/2021	1.5
M.Net.5	Test of DCI with CERN	1/1/2022	31/12/2022	0.5

*Net: 39 PM*

# WP 2.1.6 (Operations) 1/2

- Several auxiliary tools
- Monitoring
  - Upgrade of Sensu and Grafana
  - Evaluation of the feasibility of developing a system similar to SiteMon to test the availability of the services of non-LHC experiments
  - Set-up of a log analysis infrastructure running ELK stack
  - Integration with the BDP (see WP 2.5) to collect monitoring/accounting information, service logs, facility sensors report, etc. and then to analyze them with the help of ML techniques
- Accounting
  - Extend accounting to cloud
- Provisioning
  - Upgrade of Puppet and Foreman
  - Evaluation of Openstack Ironic as lifecycle management for bare metal nodes as an alternative to Foreman
  - Another alternative to be evaluated is the deployment on bare metal with Kubernetes
- Asset Management
  - A new asset management needed
  - openDCIM as candidate solution
- Documentation system
  - Complete the migration from legacy systems (e.g., wiki) to Confluence
  - Evaluation of free alternatives

# WP 2.1.6 (Operations) 2/2

Task	Activity	2021				2022				2023			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
2.1.4.1	Update of monitoring infrastructure		Orange	Orange	Orange	Orange	Orange						
2.1.4.2	Enhancement of access to Grafana portal and segregation		Orange	Orange	Orange	Orange							
2.1.4.3	System similar to WLCG SiteMon test for no-LHC experiments								Orange	Orange	Orange	Orange	Orange
2.1.4.4	Deployment and support for log analysis infrastructure based on BDP			Orange	Orange	Orange	Orange	Orange	Orange	Orange	Orange		
2.1.4.5	Monitoring as a Service for Cloud users			Orange	Orange	Orange	Orange	Orange	Orange	Orange			
2.1.4.6	Accounting data grouped by user and visualization enhancement		Black	Black	Black								
2.1.4.7	Cloud resources accounting		Black	Black	Black	Black	Black	Black					
2.1.4.8	Update of provisioning infrastructure		Green	Green	Green	Green	Green						
2.1.4.9	Harmonization of tools for provisioning and node classification							Green	Green	Green	Green	Green	
2.1.4.10	Evaluation and deployment of asset management tool	Yellow	Yellow	Yellow									
2.1.4.11	Evaluations on documentation tools									Purple	Purple	Purple	Purple

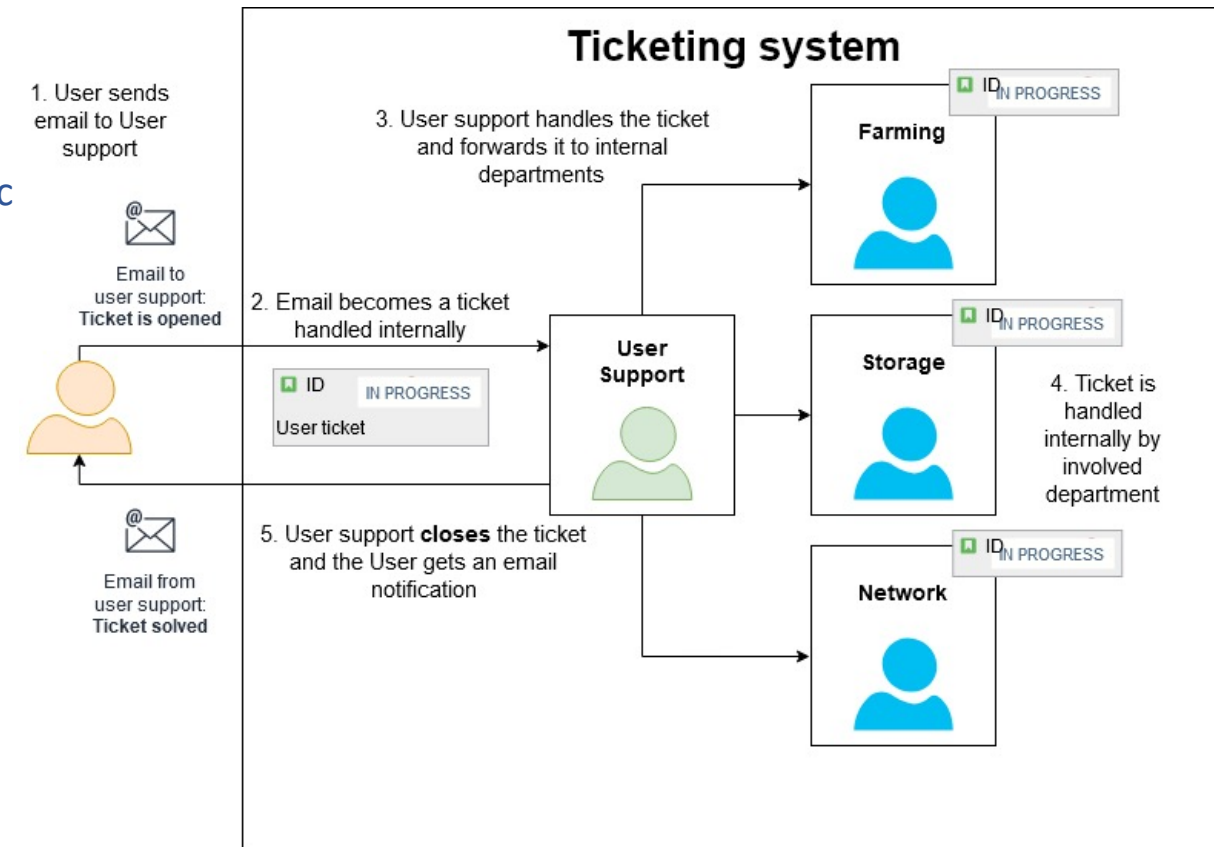
*Operations: 58 PM*

# WP 2.1.9 (User Support)

- Expand the scope to support all CNAF users
  - User Support should represent the only access point for all users, including communities using cloud resources
- Standard procedures
  - Uniform triage and helpdesk procedures for the whole data center and for all services
  - Unique ticketing system (both for users and internal)
  - Unified channel for internal communications (e.g., chat)
  - All the documentation uniformed and provided in an organic way
- Organize tutorials for the users (e.g., Tutorial Days 2020)

Milestone	Description	Start Date	Delivery Date	FTE Required
M.US.1	Mandate update	1/4/2021	31/12/2021	0.1
M.US.2	Uniformed and improved documentation howto	1/4/2021	30/6/2022	0.1
M.US.3	Ticketing system in production	1/4/2021	31/3/2023	0.3

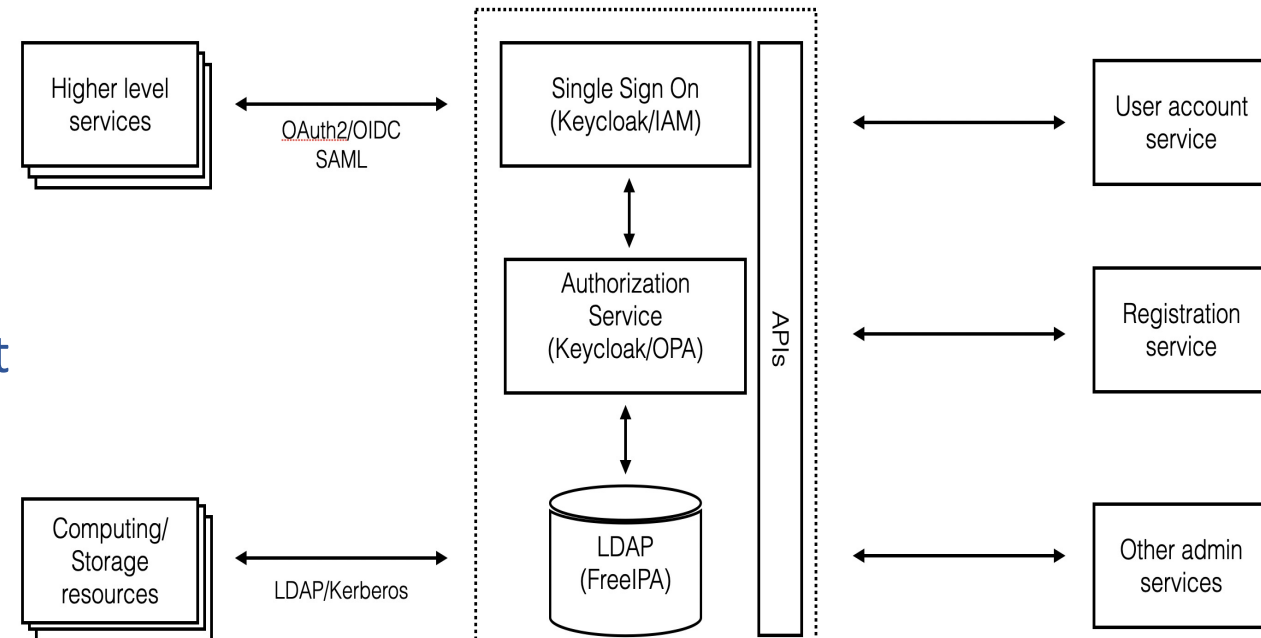
US: 5 PM





# WP 2.2 (AAI)

- Requirements on software solution
  - Open-source software
  - Widely adopted solution
- Deployment of services in HA
- Provide single sign-on on CNAF resources
- Provide flexible account and group management
- Support the delegation of administrative rights in a hierarchical fashion and flexible fine-grained, group-based authorization
- Common operations like resetting password credentials, joining groups or integrating services with the AAI must be well-documented and for CNAF users
- Integrated with IAM
- Investigate possible integration with INFN-IdP
- Adopted solution: FreeIPA, Keycloak

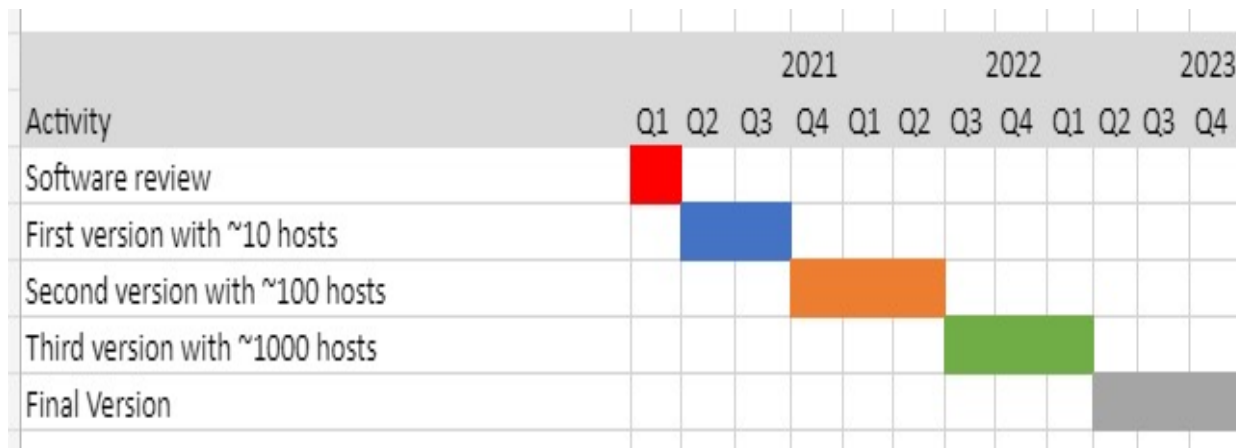


Activity	2021				2022				2023			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Definition of requirements & architecture	█	█										
Deployment core services @ CNAF	█	█	█	█	█	█						
Migration strategy definition & testing	█	█	█	█	█	█	█					
Integration testing		█	█	█	█	█	█					
Design, develop & testing User, Registration & Authorization services	█	█	█	█	█	█	█					
Deployment of AAI services @ Tecnopolo								█	█			
Integration & testing @ Tecnopolo										█	█	

AAI: 72 PM

# WP 2.3 (Security)

- Creation of a Security Operation Centre (SOC) also thanks to activity in the framework of CCR security group
- Set-up of an advanced monitoring system including:
  - A database (updated automatically) to keep track of all CNAF computer (see also Asset Management WP2.1.6) and a list of the installed software
  - Centralized log collection and analysis facility (see also Monitoring WP2.1.6)
  - Active monitor state of all computers to detect anomalous events
  - Periodic scan on all systems (already in place)



*Security: 45 PM*

# WP 2.4 (HPC/HTC integration)

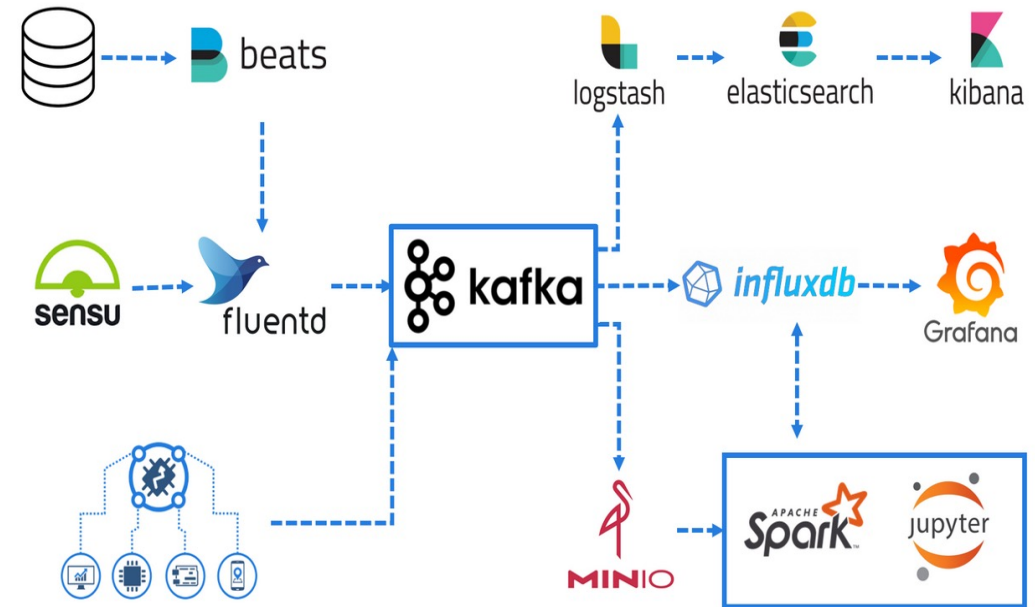
- Opportunistic access to HPC resources
  - WNs completely managed by CINECA, requiring the need to adapt our workflows to their standard setup
- Requirements (from ongoing tests with Marconi A2 and Marconi100)
  - Availability of HTCondor-Slurm interface (see WP2.1.1)
  - Outbound connectivity is one of the key points
    - At least proxying services (e.g., XrootD) will be needed to access storage
    - To investigate if it will be possible to use the Skyway gateways as for dedicated resources
  - Possibility to use virtualization (i.e., Singularity)
  - Extend the WLCG AAI services to CINECA (i.e., adoption of OIDC token-based approach)
  - Possibility to route suitable jobs from HTC at CNAF to Leonardo according to CPUs/GPUs availability

Activity	Configuration	Manpower	2021				2022				2023			
			Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Marconi 100 tests with CPU and GPU workflows;	CNAF@Berti, CINECA@Casalecchio	Playground for the experiments no real commitment	■	■										
Marconi 100 tests with CPU and GPU workflows; preparation of a proper HTCondor CE at CNAF to allow job distribution (local and remote); configuration and deployment of caches for high speed access	CNAF@Berti, CINECA@Casalecchio	0.25 FTE for Condor; 0.2 FTE for caches; 0.25 FTE/exp			■	■								
Net configuration Technopole-BertiPichat	CNAF@Berti, CINECA@Casalecchio	0.2 FTE			■	■	■							
CNAF-Leonardo integration: move or duplication of M100 infrastr. on Leonardo; tests	CNAF@Berti, CINECA@Tecnopolo	0.25 FTE for Condor; 0.2 FTE for caches; 0.25 FTE/exp for integration tests					■	■						
Production Phase	CNAF@Berti, CINECA@Tecnopolo	0.15 FTE for Condor; 0.1 FTE for caches; 0.15 FTE/exp for integration tests						■	■	■	■	■		
Moving infrastructure on storage@Tecnopolo	CNAF@Tecnopolo, CINECA@Tecnopolo	0.5 FTE to integrate storage and network; 0.5 FTE/exp to validate new production env								■	■	■	■	■

*HPC/HTC: 50 PM*

# WP 2.5 (External collaborations)

- Currently 3 main infrastructures for external projects
  - INFN cloud (WP 2.1.4)
  - EPIC cloud (WP 2.1.5)
  - Big Data Platform + IoT Platform
- Big Data Platform relevant also for internal use
  - Log analysis, Security
- CNAF Unified Monitoring Platform
  - It will be deployed in the context of Sub Task 2.1.4 Operations and will be the central monitoring and logging platform.
  - The complete set of logs/metrics will fulfill the requirement of having a centralized log collection and analysis tool expressed by Task 2.3 (Security)



WP2.5: 12 PM

Activity	2021				2022				2023			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Task 1 - Platform Hardening		█	█	█								
Task 2 - Platform Integration with AAI			█	█	█	█						
Task 3 - PBD+IoT Deployment							█	█	█	█		

# Migration strategy (WP 2.6)

- Goal: no (or minimal) down of services
- Foreseen effort: 60 PM over 10-12 months
- Steady increase of pledged resources up to 2025
  - ~1.5 MHS06 of CPU power
    - 500-600 kHS06 From Leonardo
    - 16x40kW racks to be installed for the remaining part
      - Exploring the possibility to use DLC racks (80 kW each)
  - ~100-120 PB-N of disk
    - 23-35 racks (depending on the adopted storage model)
- Installed resources at CNAF at the end of 2022:
  - CPU: ~100-150 kHS06 (~400 kHS06 at CINECA)
  - Disk: ~60 PB-N;
  - 2 tape libraries (~120 PB)
- Only hw resources to be decommissioned after 2023 will be moved
  - We are postponing hw replacement at CNAF to deploy new one at Tecnopolo

- During 2022, tenders will be executed for resources to be installed from the beginning in the new data center

1.5.1	Storage	261 d	Mon 03/01/22	Mon 02/01/23	
1.5.2	Network	261 d	Mon 03/01/22	Mon 02/01/23	
1.5.3	Racks	261 d	Mon 03/01/22	Mon 02/01/23	
1.5.4	CPU	261 d	Mon 03/01/22	Mon 02/01/23	
1.5.5	Hw relocation services	66 d	Mon 03/10/22	Mon 02/01/23	

- After installation and cabling of racks
  - **WP2-M5 23/02/2023**
- Installation of core network devices (router, core switch/spine leafs, DCI)
  - Basic network services (e.g., DNS, DHCP, ...)
- Configuration of DCI between CNAF and CINECA (1.6 Tbps)
  - LAN extension for a smooth migration of services
- Interconnection with Leonardo
  - Tests of Skyway gateways
  - Exploitation of Leonardo as pledged resources
- **Core network devices: WP2-M6 27/02/2023**

Area	# racks	Power (kW)
Farm HTC	16	40
Farm HPC	2	16
Storage	23-35	16
Other Tier1 services	6	16
Cloud General purpose	10	16
EPIC cloud	10	16
National ICT services	5	16
<b>TOTAL</b>	<b>72-84</b>	

- HTC farm mostly located in CINECA premises
  - To be phased out (~400 kHS06)
  - Replaced with new resources (+ Leonardo)
- Small fraction installed at CNAF (~ 150 kHS06) to be relocated
  - Simply switch off and move
- HPC farm (small facility) to be moved/renewed
- These operations can be done at any moment after the DCI is available



# Storage relocation

- Install a buffer of disk to copy the data from CNAF
  - Delayed replacement disk of 2015-2016 tenders (~14 PB)
  - New disk to be installed in 2023 (~7 PB?)
  - Replacement of 2017 and 2018 tenders (~20 PB)
  - **WP2-M7 28/04/2023**
- Move the data from newer systems (green ones) to the buffer and then the hw (one by one)
  - File-systems can be migrated using GPFS tools w/o service interruption
- Move the data from the other systems one at the time (and dismiss the hw)
  - Replacement of orange systems to be installed at Tecnopolo
- Maintenance for the older systems (yellow ones) will be extended until the end of 2023
- **Data transfer completed: WP2-M8 11/09/2023**
- Libraries will be moved one by one after the buffer and the HSM services will have been moved
  - 1 week/library

Production year	End of support	TB-N
2015	March 2021	10.050
2016	May 2022	3.640
2017	June 2023	7.984
2018	Nov. 2023	11.521
2019	June 2025	5.022
2020	Q3 2027	8.700
2020	Q3 2026	2.000
2021	Q2 2027	4.700
2022	2028?	8600

- Migration of National Services (and Sisinfo) **1 week**
  - Since all services are in HA with active replica at LNL, the migration will be done simply switching off CNAF instances and moving them to Tecnopolo
    - Activation of a new direct link Tecnopolo  $\leftarrow \rightarrow$  LNL
- Migration of cloud **2 months**
  - Live migration (40 Gbps bandwidth needed between CNAF and Tecnopolo)
- Migration of EPIC cloud **2 months**
  - Install new hardware at Tecnopolo and move data and services from CNAF
    - Crypted channel will be temporarily established to allow the data transfer

2.6.13	National IT Services relocation	5 d	Mon 31/07/23	Fri 04/08/23	Depends on 2.6.12
2.6.14	Cloud Services relocation	60 d	Mon 20/03/23	Fri 09/06/23	Network and racks availability
2.6.15	EPIC Cloud Relocation	60 d	Mon 20/03/23	Fri 09/06/23	Network and racks availability

# Timeline

	Jan-23	Feb-23	Mar-23	Apr-23	May-23	Jun-23	Jul-23	Aug-23	Sep-23	Oct-23
Tender delivery										
Rack installation& cabling										
Core network installation										
Leonardo validation										
CPU migration										
Dusk buffer available										
Storage migration										
Library migration										
Cloud migration										
EPIC migration										
SSNN/sisinfo migration										